It's Okay to Be Square If You're a Flexagon

ETHAN J. BERKOVE
Lafayette College
Easton, PA 18042
berkovee@lafayette.edu

JEFFREY P. DUMONT
35 Forest St.
Claremont, NH 03743

It has been said that a mathematician can be content with only paper and pencil. In fact, there are times when one doesn't even need the pencil. From a simple strip of paper it is possible to make a surprisingly interesting geometric object, a *flexagon*. The flexagon can credit its creation to the difference in size between English-ruled paper and American binders. The father of the flexagon, Arthur Stone, was an English graduate student studying at Princeton University in 1939. To accommodate his smaller binder, Stone removed strips of paper from his notebook sheet. Not being wasteful, he creased these lengths of paper into strips of equilateral triangles, folded them in a certain way, and taped their ends. Stone noticed that it was possible to flex the resulting figure so that different faces were brought into view—and the flexagon was born [**4**]. Stone and his colleagues, Richard Feynman, Bryant Tucker, and John Tukey, spent considerable time cataloging flexagons but never published their work.

Like many geometric objects, flexagons can be appreciated on many levels of mathematical sophistication (the first author remembers folding flexagons in elementary school). With so adaptable a form, it is not surprising that flexagons have been studied from points of view that vary from art to algebra. Our interest in flexagons was sparked by a question posed in a paper by Hilton, Pedersen, and Walser [**9**]. They studied one of the *hexaflexagons*, so-named because the finished model has the shape of a hexagon. They calculated the group of motions for a certain hexaflexagon, then inquired about other members of the hexaflexagon family. We have determined that the trihexaflexagon is exceptional, as it is the *only* member of the hexaflexagon family whose collection of motions forms a group.

Working to generalize this result, we shifted our attention to *tetraflexagons*, which are constructed from strips of squares folded into a 2×2 square final form. We discovered that tetraflexagons are, if anything, more complicated and interesting than their hexagonal cousins. These results convinced us that tetraflexagons, often only mentioned in passing in the literature, deserve to be brought into the limelight. In this paper, we will summarize the results of our investigations. Although most of the material on hexaflexagons is known, the material on tetraflexagons includes new results and open questions. It is our intent to give interested readers enough material to start their own explorations of these fascinating objects.

The hexaflexagon family

It is easy to fold a flexagon, and we highly recommend making one of your own as this experience will be helpful in following the results in this section (and it's fun). Construct a strip of nine equilateral triangles and a tab as in FIGURE 1; this strip is called the *net* of the flexagon. Each triangle in the strip, and in general each polygon in a flexagon net, is called a *leaf* of the flexagon. You may want to label both sides of

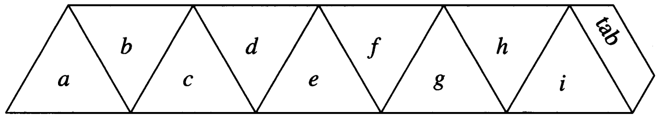

Figure 1 Trihexaflexagon net

each leaf and precrease all edges in both directions. Hold the leaf marked a in your hand, fold leaf c over leaf b, f over e, and i over h. Finish the flexagon by gluing or taping the tab onto leaf a. The final model should look like the flexagon depicted in FIGURE 2. Clockwise from the top, one can read off the leaves f, d, c, a, i, g. We call (f, d, c, a, i, g) a *face* of the flexagon.

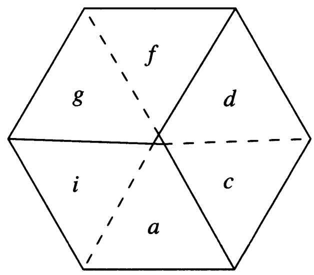

Figure 2 The (f, d, c, a, i, g) face of the trihexaflexagon

To flex your new creation, bring the three corners at the dashed lines down together so they meet. The hexagon will form a Y, at which point it will be possible to open the configuration at the middle. (This is the only possible way to perform a *flex-down* for this flexagon. There is also an inverse operation, a *flex-up*.) The result is a different face (f, e, c, b, i, h) of the flexagon. The flex can be repeated to get a third face (g, e, d, b, a, h), and one more flex returns the flexagon to (g, f, d, c, a, i), the original face rotated clockwise through an angle of $\pi/3$. Since it has three distinct faces, this flexagon is known as the *trihexaflexagon*—it is the simplest member of the hexaflexagon family. The three faces can be seen more easily if they are marked somehow: Wheeler **[15]** shows how to color the net so each flex brings out a new color, and Hilton et al. **[9]** give a way of marking the net so flexes bring out happy and sad pirate faces.

We would like a way to keep track of all the faces of a flexagon while we flex, which we can do using a graph: vertices represent the faces of the flexagon, and an edge joins two vertices if there is a flex that takes one face of the flexagon to the other. On occasion we will use a directed graph, where an arrow points towards the face that is the result of a flex-down. We choose to ignore the orientation of a face in the graph as this has a tendency to make the graph overly complicated. The completed graph is called a *structure diagram*. The cycle in FIGURE 3 is the structure diagram for the trihexaflexagon. It shows the three distinct faces, as well as their relationship via flexing.

The flex-down we described is a *motion* of the flexagon, a transformation that takes one hexagonal face of the flexagon to another hexagonal face. We require that our flexagons have no faces containing loose flaps that can be unfolded or moved so the hexagonal shape is lost. This becomes a significant issue as the number of triangles in the net increases. Indeed, in larger nets it is increasingly likely that a random folding of the net yields a face containing a loose flap, which in turn causes the entire flexagon to fall apart into a Möbius band with multiple twists. Therefore, we only consider flexagons that are folded in such a way that every flex is a motion.

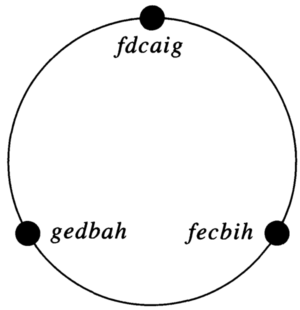

Figure 3 The structure diagram for the trihexaflexagon: vertices denote faces, edges denote flexes between faces

To determine the structure of the set of motions, let f denote a flex-down (hence f^{-1} denotes a flex-up), and denote by τ a flip along the x-axis of the flexagon. Then $f^3 = \tau^2 = id$ and it is easy to confirm that $\tau f \tau = f^{-1}$. The group with this presentation is well known as S_3, the symmetric group on 3 letters, which has 6 elements.

This analysis ignores the fact that f^3 is not strictly the identity flex, but is instead a rotation through $\pi/3$ degrees. The complete set of motions of the trihexaflexagon, including rotations, is analyzed in [**9**]; the only difference from the argument above is that $f^{18} = id$, and the resulting group is D_{18}, the dihedral group with 36 elements.

Motions of the hexahexaflexagon We perform a similar analysis for the member of the hexaflexagon family with six faces. This *hexahexaflexagon* can be constructed from the net of 18 triangular leaves with a tab as in FIGURE 4. To create the flexagon, label all leaves front and back and precrease all edges as before. Fold leaf a under the rest of the strip. Then fold the edge between leaves c and d so that a and d are adjacent. Next fold the edge between e and f so that c and f are adjacent. Continue rolling the strip in this manner until o and r are adjacent. The finished roll should look like the initial strip for the trihexaflexagon with leaf b on the far left. Fold this like the trihexaflexagon, then tape the tab to a to complete the model.

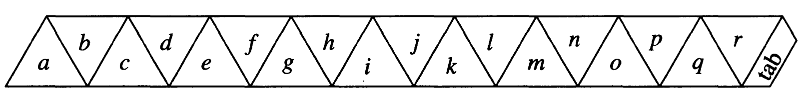

Figure 4 The hexahexaflexagon net

In contrast to the trihexaflexagon, from the initial face of the hexahexaflexagon either alternating set of corners flexes down. We can distinguish the two flexes by looking at what they do to the number of leaves in a triangular segment of the hexagon. Following Oakley and Wisner [**11**], we call the entire triangular segment a *pat*. In our newly folded hexahexaflexagon, pats alternately contain 2 and 4 leaves, or $(2, 4)$ for short. One of the flexes, which we call f, preserves thicknesses so is *pat-preserving*. The other flex, g, is *pat-changing*, from $(2, 4)$ to $(1, 5)$ and vice versa. The pat thicknesses are invariant under τ, a flip along the x-axis.

Starting with the initial face of the hexahexaflexagon, one finds that f^3 rotates the hexagon clockwise through an angle of $\pi/3$, so $f^{18} = id$. In addition, as $\tau f \tau = f^{-1}$, f and τ generate a copy of D_{18} in the collection of hexahexaflexagon motions. On the other hand, g leads to a face where the only possible flex is the pat-preserving f, which leads to a face where the only possible flex is the pat-changing g. The three-flex combination gfg rotates the hexagon clockwise through an angle of $\pi/3$, and g

satisfies $\tau g \tau = g^{-1}$. Therefore, there is at least one other copy of D_{18} in the motions of the hexahexaflexagon.

We turn this information into the structure diagram in FIGURE 5. Solid arcs between faces denote pat-preserving flexes, while dotted arcs denote pat-changing flexes. There is a primary cycle that can be traversed using pat-preserving flexes, and three subsidiary cycles that can be entered at faces where two flexes are possible. From the cycle structure we learn an important fact: the collection of motions of the hexahexaflexagon must have at least two generators, f and g. However, there are faces where only one of f and g can be applied, so it is not always possible to apply f or g twice in a row. In fact, g^3 does not even make sense. We now have an answer to the question posed in **[9]**: the hexahexaflexagon's motions do not form a group!

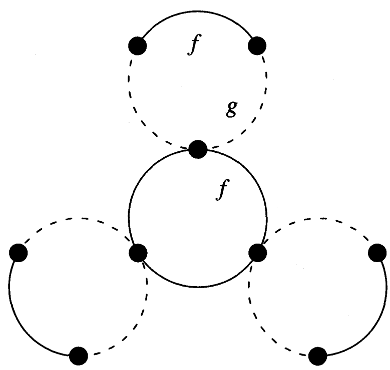

Figure 5 Structure diagrams for the hexahexaflexagon: dotted edges are pat-changing flexes, solid edges are pat-preserving flexes

This conclusion surprised us, as the collection of motions for every other geometric object we know of has a group structure. Furthermore, all hexaflexagons except the trihexaflexagon share this characteristic, as they contain pat-changing flexes. Pat-changing flexes occur at faces with two possible flexes, and the only hexaflexagon structure diagram without intersecting cycles is the trihexaflexagon's. Structure diagrams are necessarily finite, and results in Wheeler **[15]** imply that cycles cannot form a closed link. As a result, in any hexaflexagon there are only so many times the pat-changing flex can be applied. Thus, for every hexaflexagon with a pat-changing flex g there is a k such that g^k is undefined, which implies that the collection of motions for that hexaflexagon cannot form a group.

The astute reader might have noted that the structure diagram in FIGURE 5 seems to have nine faces, not six, as the hexahexaflexagon's name suggests. This discrepancy can be explained by carefully studying the hexahexaflexagon's faces. Upon closer inspection, three identical sets of triangles occur in two separate faces, but in different orders. If a face depends on both the triangles and their order, then indeed the hexahexaflexagon has nine faces. If order is disregarded, however, there are only six faces. The latter is the standard accounting, hence the name (although Oakley and Wisner **[11]** distinguish between the six *physical* faces and the nine *mathematical* faces).

We can construct other members of the hexaflexagon family by increasing the number of triangles in the initial net. For example, starting with a straight strip of 36 leaves, the strip can be rolled, then rerolled to yield the net in FIGURE 1, then folded as in the trihexaflexagon case to yield the dodecahexaflexagon. One can fold other members of the family by starting with nets that are not straight. The article by Wheeler **[15]** and Pook's book **[14]** contain some nice directions for folding the tetrahexa- and pentahexa- cases. There is also a HexaFind program [**3**], which generates all nets for hexaflexagons with any given number of faces.

We have only discussed a bit of what is known about hexaflexagons. In [**11**], Oakley and Wisner introduce the concept of the pat, then use it to count the total number of hexaflexagons with a particular number of faces. Madachy [**10**], O'Reilly [**12**], and Wheeler [**15**] describe connections between hexaflexagons and their structure diagrams. McIntosh [**7**], Madachy [**10**], and Pook [**14**] provide lengthy bibliographies to other work on hexaflexagons. Gilpin [**8**], Hilton et al. [**9**], and Pedersen [**13**] study the motions of the trihexaflexagon. Furthermore, the group of motions of the trihexaflexagon is identified in [**8**] and [**9**].

The distant cousins: tetraflexagons

We were introduced to tetraflexagons in a Martin Gardner *Mathematical Games* column in *Scientific American* [**6**]. We immediately noticed some differences; the direct analog of a strip of triangles, a straight strip of squares, makes for a poor flexagon; when you fold square over square you end up with a roll that does not flex at all (and has a trivial structure diagram). Therefore, we allow the nets to have right-angled turns. These turns occur at what we call *corner squares*, those attached to their neighbors on adjacent edges rather than opposite ones. Because so much is known about hexaflexagons, we felt that the tetraflexagon family would be readily analyzed. Our actual experience mirrored Stone's, as reported by Gardner [**6**]: "Stone and his friends spent considerable time folding and analyzing these four-sided forms, but did not succeed in developing a comprehensive theory that would cover all their discordant variations." However, we have established some notation and proved some results; our investigations have convinced us that tetraflexagons are as fascinating as, and more subtle than, their hexaflexagon relatives.

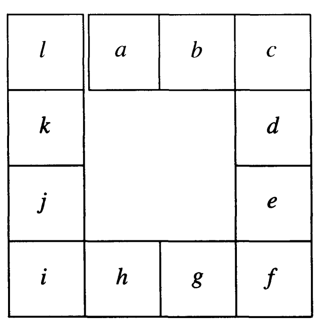

Figure 6 A "4 × 4 ring" net

Here is one way to fold a tetraflexagon using a 4×4 net with a cut adjacent to a corner square. Label a net as in FIGURE 6 and lay it flat. Mark the cut edges of a and l so you can tape them together after you have folded the tetraflexagon.

1. Fold leaf c over leaf d. Leaves a and b will remain to the left of c, but will flip over.
2. Fold c and d over e; the back of leaf c will touch the front of leaf e, and again, a and b are left free.
3. Fold f over g; move the flap with a and b so it points to the right.
4. Fold f and g over h so f and h are touching. Then slip the flap with a and b under j.

5. Fold j over i (without moving a and b), and flip the partially folded tetraflexagon upside down. Position the tetraflexagon so k and l are at the top of the model and a and b are upside down and facing to the left.
6. Fold i and j over k so i and k are touching.
7. Slide the flap with a and b under l, then fold a over l. Finally, tape a and l together on the right, along the edges that were originally cut.

The finished tetraflexagon will have 90-degree rotational symmetry, and should look like FIGURE 8a.

Now that you have a tetraflexagon in front of you, let's introduce some notation (see also FIGURE 7). The tetraflexagon has four pats, which conveniently look like the four quadrants in Cartesian coordinates. Therefore, we will refer to pats by the quadrant they are in: pat I, pat IV, etc. If you look carefully at your tetraflexagon, you'll notice that adjacent pats are connected by a layer of paper, called a *bridge*. The two leaves, one in each quadrant, that make up this layer are *bridge leaves*. Leaves b and c are bridge leaves, as are k and l. Bridges will figure prominently in our analysis.

Define a *pocket* as a strict subset of leaves within a pat attached to the rest of the tetraflexagon on two adjacent sides. Without the strictness condition, an entire pat is a pocket, connected to the rest of the flexagon by the two bridges; this is an extreme case we wish to avoid. On the other hand, a pocket may consist of a single leaf, in which case the pocket is a bridge leaf. As an example, squares f and g form a pocket in the tetraflexagon you just folded and this is the only pocket in its pat. The attached sides of the pocket are *connected*, whereas the other two sides are *free*. When a free side of a pocket lies along the outside edge of the tetraflexagon (like the outside edge of h and i), we call the entire tetraflexagon edge a *foldable edge*. The foldable edge will be folded in half during the flex.

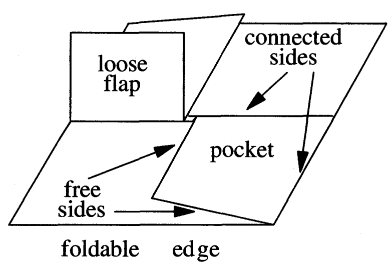

Figure 7 Tetraflexagon notation

There are two conditions that a tetraflexagon must satisfy in order to flex: there must be a pocket with a component bridge leaf, and the bridge in the pocket must be at right angles to the foldable edge. (Since the bridge goes from a pocket to an adjacent pat it cannot cross the pocket's free side.) To perform a flex, orient the tetraflexagon so the foldable edge is forward and the pocket lies on top of the tetraflexagon, as in FIGURE 8a. Fold the tetraflexagon in half perpendicular to the foldable edge so that the pocket remains on the outside. Put your thumbs into the pocket and its kitty-corner companion (FIGURE 8b), then pull outwards in the direction of the arrows. The pocket layers will separate from the rest of the pat, rotating outwards 180 degrees but staying in the same quadrant. The rest of the pat layers will maintain their orientation but move to the top of the adjacent pat. Simultaneously, the layers in the adjacent pats will rotate outwards 180 degrees. When you flatten the tetraflexagon, you will see a new face as in FIGURES 8c and 8d.

We call the flex with the pockets up and the foldable edge forward a *book-down flex*. It is only one of four possibilities, depending on the position of the pocket and foldable edge. We therefore distinguish between four types of flexes: flexes up or down along the y-axis (*book-up* and *book-down*) and flexes up or down along the x-axis (*laptop-up* and *laptop-down*). From FIGURE 8, we see that an up flex is the inverse of a down flex and vice versa. As before, the top layer of leaf of the tetraflexagon, like (a, d, g, j), is called a *face*, and a flex is a *motion* if it takes one tetraflexagon face to another.

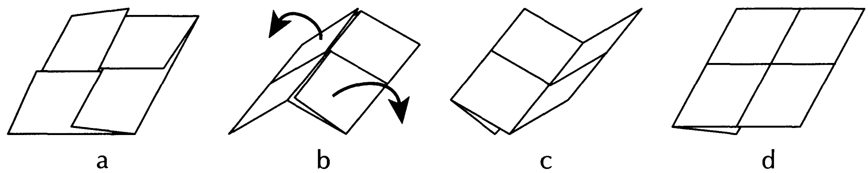

Figure 8 Flexing a tetraflexagon

Before we go further into our exploration, we introduce some assumptions about the tetraflexagons we consider to aid in our analysis.

1. *Net* assumption. All tetraflexagon nets have exactly four corner squares. The two vertical strips are each composed of m leaves, and the two horizontal strips of n leaves. All nets are folded into four-pat (2×2) tetraflexagons.
2. *Winding* assumption. There is only one bridge between adjacent pats of the tetraflexagon. In other words, as the net is folded, the strip of paper enters and exits each quadrant exactly once. This assumption is easily justified; we will see in the Flexing Lemma below that subsequent flexes maintain all four bridges as single layers.
3. *Rigidity* assumption. No tetraflexagon face contains a loose flap. A *loose flap* is a collection of two or more leaves connected to the rest of the tetraflexagon on only one side (see FIGURE 7). Since a loose flap determines one of two possible tetraflexagon faces depending on its position, movement of a loose flap can be considered a motion. The rigidity assumption ensures that the only tetraflexagon motions are flexes.
4. *Symmetry* assumption. Folded tetraflexagons have 180-degree rotational symmetry (in contrast to the hexaflexagons' 120-degree symmetry).

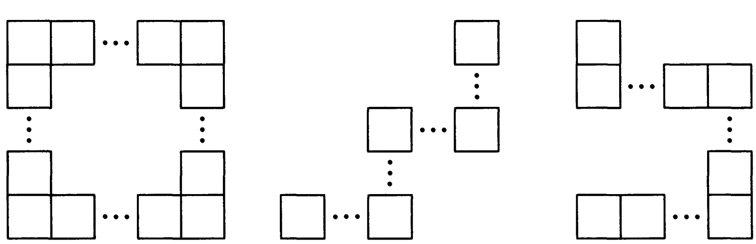

Figure 9 Rings, bolts, and snakes

The net assumption implies that nets are one of three shapes, depending on parity. All nets are shown in FIGURE 9: when m and n are both even, the only possible net is a *ring*; when both m and n are odd, the only possible net is a *lightning bolt*; and when exactly one of m and n are odd, the only possible net is a *snake*. We remark that analogs of the other three assumptions hold for most hexaflexagons.

Basic analysis: folding 101 and structure diagrams with dead ends Flexing for tetraflexagons is a bit more complicated than for hexaflexagons. We carefully describe what happens during a flex in the following

FLEXING LEMMA. The following hold for tetraflexagons under our assumptions:

1. Corner squares remain in their pats during flexes.
2. A pocket that has a component bridge leaf is either a single layer (that is, it is the bridge leaf), or consists of every leaf in the pat except for either the top or bottom one. In the latter case, the leaf not belonging to the pocket is also a bridge leaf.
3. A flex maintains all four bridges as single layers.

Proof. To simplify the arguments, we will assume for this proof that the pocket is always on top of pat IV with the foldable edge forward, as in FIGURE 8a. The other cases will follow by rotation and mirror images. Moreover, by the symmetry assumption we only need focus on pats III and IV.

We start with the first claim. When we perform the flex shown in FIGURE 8, all leaves in pat III remain in pat III, which means that the only corner square that can change pats lies in pat IV. However, if the book-down flex moves this corner square to pat III, then the flex is also forced to move the strip of leaves connecting the corner squares in pats I and IV. This is impossible as the corner square in pat I is fixed by the symmetry assumption.

For the second claim, assume that the pocket with a bridge leaf consists of more than one layer. We claim that the corner square in pat IV must be a leaf of the pocket. By way of contradiction, assume the corner square is below the pocket. Notice that the corner squares in pats I and IV are connected by a vertical strip. As this strip is woven from the corner square in pat IV to the corner square in pat I, it must at some point become the unique bridge to pat I (by the winding assumption). One of the bridge leaves is a leaf of the pocket, so the vertical strip of squares must also cross from the bottom of pat IV to the pocket. Because the vertical strip is woven back and forth, it can only cross to the pocket on the front or back edge of pat IV. The vertical strip cannot cross the gap at the front edge of pat IV because that edge is free by assumption. The vertical strip cannot cross the gap at the back edge of the pocket either, as that would lock the pocket to the rest of pat IV, making a flex impossible. We conclude that the corner square must be a leaf of the pocket.

An analogous argument shows that the vertical strip cannot cross from the corner square to the bottom layer either. By the rigidity assumption, this bottom layer must therefore consist of a single leaf. Now if the bridge leaf to pat III were any layer but the bottom, it would force the left side of the pocket in pat IV to be connected, rather than free. This is impossible as all pockets have two free sides.

For the final claim, the winding assumption implies that, when first folded, the tetraflexagon has single layer bridges between pats. Consider the effect of a book-down flex. This flex flips the bridge leaves that lie in the "pages" of the book upside down, so those bridges remain single layers. For the two bridges that cross the spine of the book, focus on pats III and IV, and assume first that the pocket is a single layer; then the book-down flex leaves a single layer in pat IV, which must be the bridge leaf. Otherwise, the pocket has more than one leaf and the bottom leaf of pat IV is a single layer. In this case, the completed flex moves the single layer from pat IV to pat III, again maintaining the bridge as a single layer. ■

The Flexing Lemma gives us insight about what happens to a tetraflexagon during a single flex. What about the larger picture? What can we say about the set of all the motions of the tetraflexagon we folded? Again, we answer these questions with a structure diagram, but now use the orientation of the graph's edges to distinguish among the four possible types of flexes. Denote a book flex by a horizontal line segment in the structure diagram, a laptop flex by a vertical one, and let arrows point to the faces that are the result of a flex down. With a little flexing and experimentation, we find that the structure diagram for the tetraflexagon we folded is an L shape, with five vertices (corresponding to faces), and four edges, two in each part of the L (corresponding to flexes). On all edges, the arrows point away from the corner of the L. This structure diagram, which appears in the bottom left of FIGURE 11, has a feature we have not seen before: there are vertices incident to a single edge. We call the faces associated to these vertices *dead ends*, because once these faces are reached via a flex, the only motion that can be applied is the flex's inverse.

For the hexaflexagon cases, a flex of finite order appeared in the structure diagram as a cycle. The structure diagram we just constructed has no cycles, hence no elements of finite order. In addition, any flex can be applied at most twice, at which point some other flex must be applied. In other words, there is really no group structure at all in the motions of the tetraflexagon we constructed. An immediate question is whether there is a tetraflexagon whose collection of motions forms a group.

There is, and it can be folded from the 3×3 lightning bolt net shown in FIGURE 10. Start with a copy of this net—you may want to specially mark the shaded sides of a and h as shown in the diagram, as those are the edges that get taped together at the end. Fold leaf b over c, d (and a, b, and c) over e, and g over f. Make sure h lies on top of a, then tape a and h together on the right, along their shaded sides. Your finished tetraflexagon should have 180-degree rotational symmetry, as per the symmetry assumption. From this initial face, the only possible down flex is a laptop-down. If this is followed with book-down, laptop-down, and book-down flexes (again, the only possible down flexes) you get back to the initial face. The resulting structure diagram is a square cycle, corresponding to the symmetry group $\mathbb{Z}/4\mathbb{Z}$, the cyclic group with 4 elements. We believe that this is the *only* tetraflexagon whose collection of motions forms a group (although a rigorous proof eludes us).

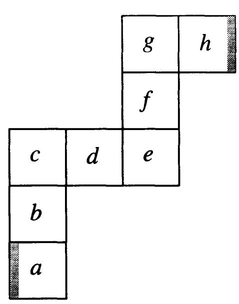

Figure 10 A 3×3 bolt net

You may have noticed that although we call the net in FIGURE 10 a 3×3 lightning bolt, there are only two leaves, g and h, in one row. Actually, that row does contain a third leaf, a, which becomes part of the row after we finish taping. In general, when we refer to a net as m by n, we include the corner squares in the count. For the cases of the bolt and snake nets, this means that either one row or one column will be a leaf short when the net is first constructed.

By reducing larger nets to the 3×3 case, one can construct many tetraflexagons whose structure diagrams contain cycles. For example, starting with the net in FIGURE 6, folding b over c and i over h turns the 4×4 ring net into a 4×3 snake net.

Another pair of folds results in a 3×3 lightning bolt net, which can then be folded as above. More generally, one can turn an $m \times n$ net into either an $m - 1 \times n$ net or an $m \times n - 1$ net by making appropriate folds. Here, appropriate means that the pair of folds results in a net that satisfies the symmetry assumption. Also, a given sequence of edge choices must result in a tetraflexagon that satisfies our four assumptions. Once the net is a 3×3 lightning bolt, it is folded to guarantee at least one cycle in the structure diagram. This process is reminiscent of how a hexaflexagon is folded from a strip containing $9(2^n)$ triangles by rolling it n times to form the trihexaflexagon net with 9 triangles. A difference, though, is that there are many ways of choosing edges in pairs to reduce an $m \times n$ net to the 3×3 net, so a generic tetraflexagon net can yield a large number of tetraflexagons with distinct structure diagrams.

As an example, through exhaustive folding, both figurative and literal, we have identified in FIGURE 11 most (all?) of the possible structure diagrams that result from a 4×4 ring. Some of the diagrams are very basic, but notice that three of them contain cycles, all but one contain dead ends, and two contain both. This variety in structure diagrams from the same net demonstrates how difficult it is to classify even simple tetraflexagons.

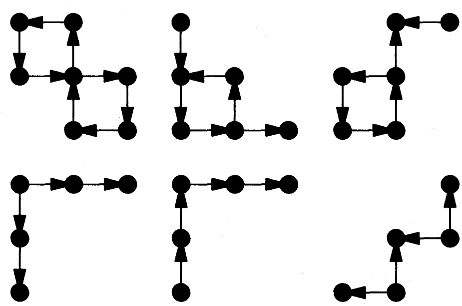

Figure 11 Tetraflexagon structure diagrams folded from a 4×4 ring

Despite their "discordant variations," the tetraflexagons we investigated have some features in common. In a number of the examples we folded, we came across faces where all four flexes were possible. We call such faces *crossroads*. For the reader who wants to see a crossroad in action, start with the net in FIGURE 6, labeled front and back with each letter. Fold b under c, c over d, e (and a–d) under f, g (and the rest of the net) over f, i over h, j over i, and l over k. Lift a over l, then tape a and l together on the right (outside) edge. The final tetraflexagon is shown in FIGURE 12d. (What is the resulting structure diagram?) Crossroads are among the most interesting features in a structure diagram, and we wondered how many crossroads a structure diagram could contain. We noticed that after we flexed a crossroad, we never saw another crossroad, which led us to prove the

CROSSROAD THEOREM. Crossroad faces are never adjacent in tetraflexagons built from ring, lightning bolt, or snake nets.

Proof. Let's analyze the shape of a tetraflexagon at a crossroad face. In order to perform all four types of flexes, every edge of the tetraflexagon must be a foldable edge and there must be two pockets per edge to allow both types of flexes. It is easy to show that the pockets cannot be on the same side of the foldable edge, as in FIGURE 12a. If the pockets were on the same side, the remaining leaves of the two pats would be the bridge between the pats, by the second claim of the Flexing Lemma. By the symmetry assumption, the same would be true for the other two pats of the tetraflexagon, and the entire tetraflexagon would simply fall apart.

We conclude that one pat along the foldable edge must have a pocket on the top of the pat to allow a flex down, while the adjacent pat along the edge must have a pocket on the bottom of the pat to allow a flex up. There are essentially two cases to consider, shown in FIGURES 12b and 12c (FIGURE 12d is FIGURE 12c's mirror image). We can show that FIGURE 12b is impossible by a careful analysis of pat IV. Since one of the two pockets in pat IV contains all the layers but one, assume, without loss of generality, that it is the pocket associated to the foldable edge on the right. Then the bridge to pat III crosses the edge marked by the arrow, and the component bridge leaf in pat IV is one of the leaves in the pocket by the second claim of the flexing lemma. However, this implies that the edge marked by the arrow is a connected edge, which makes a book-down flex impossible. Therefore, the only possible crossroad faces have configurations like FIGUREs 12c and its mirror image 12d.

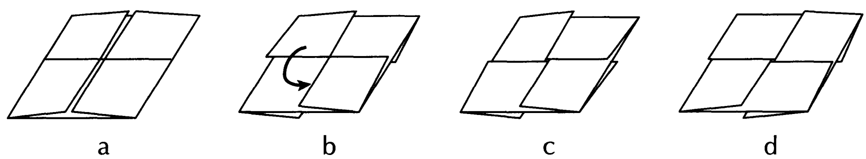

Figure 12 Two impossible configurations and two crossroads

We next determine when a flex from a crossroad yields another crossroad. By symmetry arguments, it is sufficient to consider FIGURE 12c. We claim that is impossible to perform a book-down flex from FIGURE 12c and get the crossroad in FIGURE 12d. This time, consider pat III. Note that by the second claim of the Flexing Lemma one of the two pockets in pat III must consist of every layer in the pat except one; assume (by symmetry again) that it is the bottom pocket. After a book-down flex is applied, the original layers in pat III are flipped upside down, and are covered by a single layer from pat IV. This single layer must be part of the new bridge to pat IV. Therefore, it is not possible for the right edge of pat III to be free, as in FIGURE 12d.

Thus, if two crossroads are connected by a flex, they must be identical. We rule out this possibility too. Consider the enlargement of pat IV of FIGURE 12c, shown in FIGURE 13. In order for a book-down flex to yield the same configuration, pat IV must contain a pocket as part of the bridge. But the bridge must be a single layer, by the second claim of the Flexing Lemma. ∎

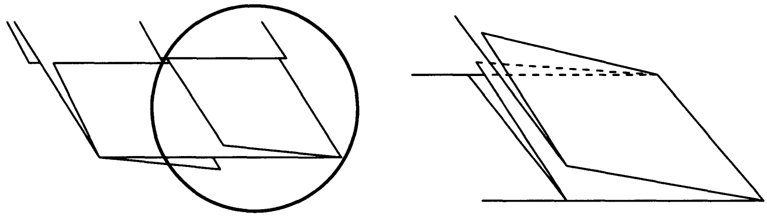

Figure 13 Closeup of pat IV

The Crossroad Theorem tells us that cycles cannot be too dense in a structure diagram. For example, we have a corollary to the theorem, named after the symbol that appears on all Purina pet foods.

PURINA COROLLARY. Under our assumptions, there is no "Purina Tetraflexagon" with the structure diagram shown in FIGURE 14.

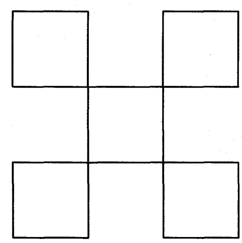

Figure 14 Purina structure diagram

Surgery: adding new parts to an old flexagon So far, we have a very general idea of what can and cannot be part of a structure diagram. On the other hand, it would be desirable to add a given component to a structure diagram, allowing us to tailor-make tetraflexagons with interesting properties. There is a technique we learned from Harold McIntosh's online notes **[7]** that allows us to do exactly that at an appropriate spot in a structure diagram. We call this technique *surgery*. One performs surgery on a tetraflexagon by symmetrically grafting two strips of squares to the model to add the new feature.

An instance of surgery is shown in FIGURE 15. Start with a tetraflexagon at a face where one of the pats is a single leaf, that is, a corner square. Place this tetraflexagon so the single leaves are in pats II and IV with the foldable edge forward, as in the top left picture in FIGURE 15. By considering mirror images, if necessary, we may assume that the pocket in pat III is on the bottom of the pat. Tape a 2×1 strip to the right edge at pat IV and make a cut in the tetraflexagon along the positive x-axis as in the top right picture. Fold the strip on top of pat IV, then fold the strip up and in half as in the bottom left picture. Finally, tape the top edge of the strip (in pat IV) to the cut edge in pat I as in the bottom right picture. Perform the same procedure symmetrically

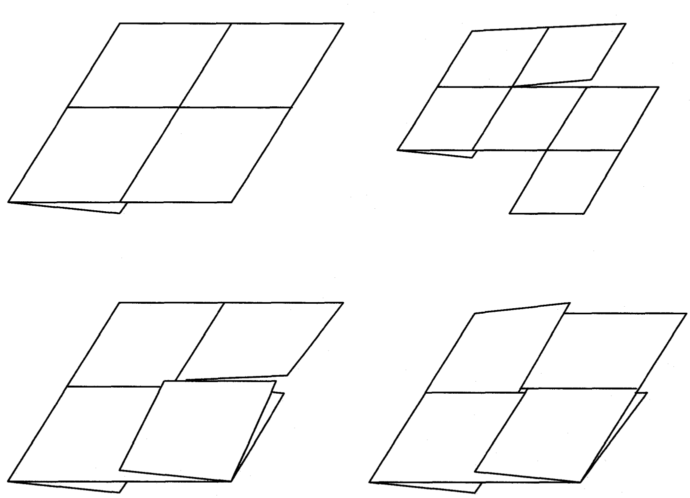

Figure 15 How to add a cycle to a tetraflexagon

on pat II; when you are done, the front and back foldable edges will have pockets on both top and bottom. The resulting tetraflexagon's structure diagram will be the same as before, but with a cycle tacked on at the appropriate face.

Following this technique, it is straightforward, at least in theory, to build a tetraflexagon whose structure diagram consists of an arbitrarily long row of four-cycles connected one to the other at opposite corners. (In light of the Crossroad Theorem, in some sense this is as crowded as collections of cycles can become.) In practice, the tetraflexagons quickly become too thick to flex easily as the size of the net increases. We challenge the reader to use surgery to construct a tetraflexagon with the structure diagram shown in FIGURE 16.

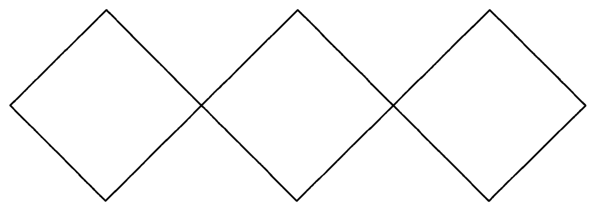

Figure 16 A challenging structure diagram

One can also use surgery to add dead ends to a structure diagram. Start with a tetraflexagon in the same position as before and make the same cut along the positive x-axis, but this time add a $1 \times n$ strip horizontally to the right edge of pat IV. Roll the strip counterclockwise until it lies on top of pat IV, then tape the top edge of the middle square of the roll to the cut edge of pat I. Repeat this symmetrically on pat II. This procedure will add n horizontal segments to the structure diagram at the appropriate place.

To finish this section we mention that besides appearing in Gardner articles **[5]** and **[6]**, "square" flexagons are the topic of a short note by Chapman **[2]** where he uses primary and secondary colors to distinguish tetraflexagon faces. This note also contains directions for constructing tetraflexagons whose structure diagrams are a cycle and two linked cycles. For more information on surgery and tetraflexagons, the reader should consult McIntosh's notes **[7]** or Pook's book **[14]**.

Where do we go from here?

At this point, we know something about tetraflexagons—how they flex, how and where cycles can occur, and why some structure diagrams are impossible—but there are many questions we haven't answered. What are the possible tetraflexagon structure diagrams? How many tetraflexagons, up to rotation and mirror image, can be made from an $m \times n$ net with four corners? We can get a rough upper bound by recalling that the corner squares in a tetraflexagon stay fixed in their quadrant. Flex a tetraflexagon so that one pat consists of a single leaf and its adjacent pats contain $n + m - 1$ leaves. With the exception of the bottom (or top) leaf, these might be in any order, so there are at most $(n + m - 2)!$ possible tetraflexagons that can be folded from a given $m \times n$ net. Many of these tetraflexagons will not satisfy our four assumptions; is there a better upper bound? Finally, there is the question that provided our initial motivation to study flexagons: Are there any other tetraflexagons besides the one built from the net in FIGURE 10 whose collection of motions forms a group?

We developed most of our tetraflexagon results before looking through McIntosh's material [7], and you can imagine our surprise when we saw directions for the Purina Tetraflexagon in his notes! McIntosh introduced this creature in a discussion on surgery, and we were no less surprised when we constructed it and confirmed that it worked. The apparent discrepancy between beautiful theory and ugly counterexample was explained once we took a scissors to the flexagon and opened it up; in the net, every one of the 24 leaves was a corner square! Clearly, in order to do surgery while maintaining the net assumption, one must be careful that the transplanted leaves do not add any more corner squares.

More importantly, this example shows that there are many interesting tetraflexagons that do not satisfy the net assumption. In fact, perhaps our favorite flexagon violates both the net and symmetry assumptions, as it is folded from regular pentagons. Of course, this flexagon cannot be folded flat, but we were impressed that such an object could exist at all. This flexagon, and many others, are described in McIntosh's notes [7] as well as Pook's book [**14**]. We highly recommend these references, as well as an impressive report by Conrad and Hartline [**1**] (also found in [**7**]) for the reader interested in deepening his or her background about flexagons and their kin.

We have found flexagons to be interesting mathematical objects at many levels. Since they are easily folded, they are a good addition to an introductory mathematics class, giving students an opportunity to look for patterns and explore a topic at their own pace. In a more advanced setting, such as a course in geometry or algebra, flexagons can be used to introduce notions of symmetry and transformation. They are also an excellent topic of study for a senior thesis, as most of the flexagon materials are written at an elementary level and are appropriate for a student's introduction to the reading of mathematical articles. And, of course, there are many open questions that are easily asked but difficult to answer. Despite their long history, flexagons still give an exciting twist to an otherwise boring strip of paper, and are well worth a little study.

Acknowledgments. We would like to thank Harold McIntosh for informative correspondence, and Liz McMahon, Kyra Berkove, and the anonymous referees, whose comments greatly improved the exposition of this paper.

REFERENCES

1. Anthony Conrad and Daniel Hartline, Flexagons, *RIAS Technical Report 62-11*, Baltimore, MD, 1962.
2. P. B. Chapman, Square flexagons, *Math. Gaz.*, **45**, (1961), 192–194.
3. Antonio De Queioroz, The HexaFind Program, http://www.coe.ufrj.br/~acmq/hexaflexagons.
4. Martin Gardner, *Hexaflexagons and Other Mathematical Diversions*, University of Chicago Press, Chicago, 1988.
5. ———, *The 2nd Scientific American Book of Mathematical Puzzles and Diversions*, Simon and Schuster, New York, 1961.
6. ———, Mathematical games, tetraflexagons and tetraflexigation, *Sci. Amer.*, May 1958, 122–126.
7. Harold McIntosh, Flexagons, http://delta.cs.cinvestav.mx/~mcintosh/oldweb/pflexagon.html.
8. Michael Gilpin, Symmetries of the trihexaflexagon, this MAGAZINE, **49**:4 (1976), 189–192.
9. Peter Hilton, Jean Pedersen, and Hans Walser, The faces of the trihexaflexagon, this MAGAZINE, **70**:4 (1997), 243–251.
10. Joseph Madachy, *Mathematics on Vacation*, Charles Scribner, New York, 1966.
11. C. O. Oakley and R. J. Wisner, Flexagons, *Amer. Math. Monthly*, **64**, (1957), 143–154.
12. Thomas O'Reilly, Classifying and counting hexaflexagrams, *J. Recreational Math.*, **8**:3, (1975), 182–187.
13. Jean Pedersen, Sneaking up on a group, *The Two Year College Math. J.*, **3**: 2, (1972), 9–12.
14. Les Pook, *Flexagons Inside Out*, Cambridge University Press, Cambridge, 2003.
15. Roger Wheeler, The flexagon family, *Math. Gaz.*, **42**, (1958), 1–6.