

Some Toroidal (and Non-toroidal) Rearrangement Puzzles

Barry A. Cipra
bcipra@rconnect.com

In 1998 I came across an etching by American artist Sol LeWitt titled “Straight Lines in Four Directions and All Their Possible Combinations,” in an exhibit catalog I found in a used bookstore on a trip to Norman, Oklahoma. (I was there for a conference on tornado forecasting.) The picture, redrawn below, is relatively straightforward: Each of the sixteen squares in a 4x4 array either does or doesn't have a horizontal, vertical, up diagonal or down diagonal line segment drawn in it.




Figure 1: The Sol LeWitt puzzle: 16 squares with all combinations of lines in horizontal, vertical, up-diagonal and down-diagonal directions.

LeWitt arranges the four singlets across the top row, the six doublets across the next row and a half, the four triplets after that, then the square with line segments in all four directions, and finally the “empty” square with nothing at all in it. (Actually, LeWitt’s version does not include the empty square; in the book where I first saw the etching, the title appears there.) LeWitt’s arrangement also draws horizontal lines first, then verticals, then up diagonals, then down diagonals, e.g., the first three doublets are h-v, h-u, and h-d, followed by v-u and v-d, and ending with u-d. I call this kind of systematic approach to laying things out a “Sol LeWitt” arrangement.

My eye (more precisely, my brain...) noticed that some lines continue from one square to another, but rarely all the way across; only a few diagonals make it all the way from one outer edge of the array to another outer edge. This made me wonder: Could the sixteen squares be rearranged, without rotating any of them, into some different 4x4 array so that all lines *do* continue all the way from outer edge to outer edge?

It turns out they can. Not only that, but the solutions have a rather remarkable property: If you move the top row of squares to the bottom, or the left column to the right, you still have a solution. That is, the solutions are all “toroidal” - there's a wrap-around effect, as if the squares were drawn on a donut.

It isn't surprising - it's obvious, in fact - that the horizontal and vertical lines in a solution behave toroidally, but it is a surprise that the diagonals do as well; there's nothing in the question itself that requires it. I eventually wrote this up for an earlier Gathering for Gardner; that paper

was published in **Puzzler's Tribute: A Feast for the Mind** (pp. 387-393).

I've tried since to come up with other, similar puzzles whose solutions all have the same toroidal property. So far my efforts have all failed. My first effort appeared in the aforementioned **Puzzler's Tribute** paper:




Figure 2: A “Circle” LeWitt puzzle, with quarter circles drawn or not drawn, centered at the four corners of each square.

The design criterion here is that each square either has or doesn't have a quarter circle centered at each of its four corners, and the problem is to rearrange the given 4x4 array of squares, sans rotations, so that each arc continues from square to square. There are 32 quarter circles in all, so a toroidal solution will have 8 complete circles (some of which appear as pairs of semicircles on opposite sides of the array). This puzzle does have toroidal solutions, but it also has solutions that are non-toroidal, so in that sense it's a failure.

I later tried a variant I called the "Sine" LeWitt problem: Along each edge of each square, either do or don't draw a half period of a sine wave, and then try to rearrange the squares, again without rotations, so that you get sets of sine curves running from left to right and top to bottom. (Note, there's not much difference, visually, between a half sine wave and a quarter circle, especially when they're hand drawn. This opens the possibility for an alternative puzzle in which the goal is to arrange the squares so that each quarter circle is part of a complete circle.)




Figure 3: The "Sine" LeWitt puzzle.

But like the Circle LeWitt puzzle, Sine LeWitt has both toroidal and non-toroidal solutions, so it's another failure.

Before I go on, a word about rotations: I self-imposed the non-rotation rule mainly to keep the sixteen squares all different. If, for example, you rotate the Sol LeWitt square with a single

up diagonal by a quarter turn, it becomes a duplicate of the down-diagonal square. (Some squares, of course, don't change if you rotate them by a quarter turn, and all squares are invariant under half turns.) I usually label the squares in my designs in a way that subtly discourages rotations. But if you want to rotate pieces, go right ahead. Just know, it's a somewhat different problem then. In particular, if you allow rotations, the original Sol LeWitt problem has additional solutions that are **not** toroidal.

Recently, in 2019, I decided to turn the whole problem on its head, and designed a set of sixteen different squares for which no matter how you arrange them, you get continuity from square to square, with toroidality understood to occur at the outer edges:




Figure 4: A toroidal looping puzzle, with labels to discourage rotation of squares, in “Sol LeWitt” order, with single-crossing squares first, then two-crossing squares, etc.. (Figure courtesy of Donna Dietz – see <http://www.donnadietz.com/cipra/CipraPuzzle.html> for a playable version of the puzzle.)

Each square has four arcs in it, with each arc connecting two “thridpoints” (an invented term for midpoints that divide an interval into thirds) of adjacent edges; the key rule that limits the number of different patterns to 16 is that the two arcs emanating from the thridpoints on each edge must connect to thridpoints on *opposite* edges. The four-bit label in each square specifies whether the two arcs emanating from the thridpoints of the left, top, right, and bottom sides of the square, in that order, do or do not cross, with “1” if they do and “0” if they don’t. One of the labels’ roles is to discourage rotation, but you are, as before, welcome to refuse to be discouraged, and rotate to your heart’s content.

A note about “thridpoints”: It’s a purely aesthetic choice to divide each side of the squares into thirds; any two points on each pair of sides will do; the important thing is that arcs continue from one square to the next. Indeed, one way to discourage rotations would be to choose “thridpoints” asymmetrically, so that continuations would be disrupted if any of the squares were rotated. It’s also an aesthetic choice to use quarter circles and quarter ellipses for the arcs; the essential property is continuity, not smoothness. An interesting question to ponder is whether aesthetic choices enhance the process of mathematical discovery or restrict it – or both!

It occurred to me later to incorporate a consistent rule for passing one arc over another wherever there's a crossing, which makes the set of squares, if you "fatten" the arcs so they look like stretches of string, look like this:




Figure 5: Toroidal looping puzzle "fattened" into over- and under-passes, arranged in "binary" fashion, with labels ordered from 0000=0 to 1111=15. (Photo courtesy of Pete Benson at CherryArborDesign.com - the puzzle can be purchased there.)

One pleasant surprise here is that, no matter how you rearrange the squares - and even if you allow rotations - the sequence of

over- and under-crossings always alternates. Experts in knot theory undoubtedly see this as obvious; the rest of us can be content to scratch our heads or work out an ad-hoc proof.

I debuted this puzzle at the 2019 MOVES conference at the National Museum of Mathematics in New York, without specifying the particular puzzle I had in mind for the pieces. You might notice I haven't done so here either (yet). I did so in part to see what ideas others would come up with for what could be done with the pattern. I invite readers to pause at this point and think for themselves of something interesting to do with the pieces. (At MOVES I did not even hint that the pattern should be cut into separate squares; some people came up with the idea of cutting, but along the *arcs*, like a jigsaw puzzle.)

One person at MOVES (I'm sorry, I don't remember who it was) observed that the pieces looked like "Tsuru" tiles, named after a popular board game of relatively recent vintage. Tsuru tiles also connect the "thridpoints" on the four sides of a sqare, and some of them are identical with the tiles in my puzzle, but others are not. I'm not sure what rule (if any) governs the set of Tsuru tiles; as mentioned above, I chose a rule that produces exactly 16 different patterns.

So here's the challenge I had in mind when I invented the puzzle: Can you rearrange the tiles so that there is exactly one loop that runs through all the arcs of all the squares?

Since each tile has four arcs, there are 64 arcs in all. It's convenient to talk about the "length" of a loop as the number of separate arcs it consists of. If you patiently count them, you will find that the "binary" arrangement in Figure 5 has four

loops, each of length 16. The “Sol LeWitt” arrangement in Figure 4 has a pair of loops of length 4 that are fairly easy to spot; the rest of its arcs belong to two loops, each of length 28.

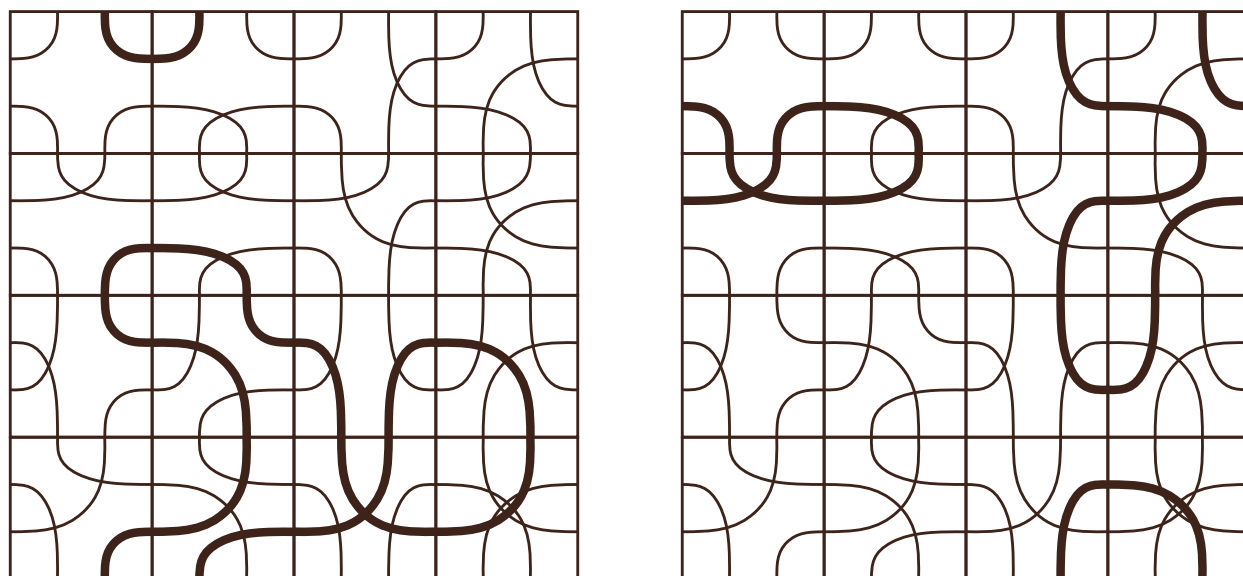


Figure 6. Two loops, each of “length” 16, in the “binary” arrangement from Figure 5. (Note, the arcs here are poorly drawn quarter circles and ellipses, as evident from a careful look at the bottom rightmost tile.)

Notice that all those loop lengths are multiples of 4. It’s not hard to see that loop lengths must be even; a two-color checkerboard proof does the trick: Each loop passes back and forth between black and white squares. To show the number of arcs is a multiple of 4, use a four-coloring of 2x2 patches, say rows of alternating Red/Blue alternating with rows of alternating Green/Yellow (hence columns of alternating Red/Green alternating with columns of alternating Blue/Yellow). If, in following a loop, you pass from Red to Blue, you’ll next pass from Blue to Yellow no matter which way you turn (up or down), then from Yellow to Green, then from Green back to Red, after

which you'll wind up repeating the color sequence again and again.

Peter Winkler took an interest in the puzzle at MOVES; by the end of the afternoon he had a proof that a single, 64-arc loop is impossible. A key observation was that all attempts finding a one-loop arrangement invariably left an *even* number of loops. What Peter finally proved was that, no matter what set of tiles you use (i.e., let each square be any tile, even if you repeat some tile patterns multiple times and not use others at all), the parity of the number of loops is equal to the parity of the number of 1's in the tiles' labels. Donna Dietz, who was also at the MOVES conference, wrote up Peter's proof, along with other observations the three of us made, in a paper posted on the ArXiv: <https://arxiv.org/abs/1908.05718> . She also posted a playable version of the puzzle on her website, as noted in Figure 4. (Clicking on any two squares there interchanges them, so you can move pieces wherever you want.)

Peter's proof works for any even-by-even array of my tile patterns; it doesn't work if one of the dimensions is odd.

Jim Propp, another MOVES attendee, had an interesting suggestion at the meeting: Instead of toroidal connections, whenever an arc came to the outer edge of the 4x4 square array, connect it to the arc in the nearest neighboring square, with three-quarter-circle connections at the four corners. For the initial "binary" array, you get this:




Figure 7: Jim Propp's non-toroidal suggestion for the looping puzzle problem.

When Jim showed me this, I suggested connecting adjacent thridpoints **within** each edge around the perimeter instead:




Figure 8: My alternative to Jim Propp's non-toroidal suggestion for the looping puzzle problem.

To our considerable surprise, this *does* consist of one single loop! It's still unclear, to me at least, if there's anything behind this beyond mere happenstance.

Later, in some email correspondence, when he saw the over- and under-passing version of the looping puzzle, Jim complained that the 16 tiles were no longer a complete set of possible patterns: There should really be a separate tile for each assignment of which arc goes under the other when two arcs cross. This would lead to a 9x9 puzzle with a total of 81 different tiles. That's a bit big for my taste, but I urge anyone undaunted by the size to see if there's anything of interest in it.

In response to Jim's complaint, I designed a 4x4 "Toroidal Trellis" problem:

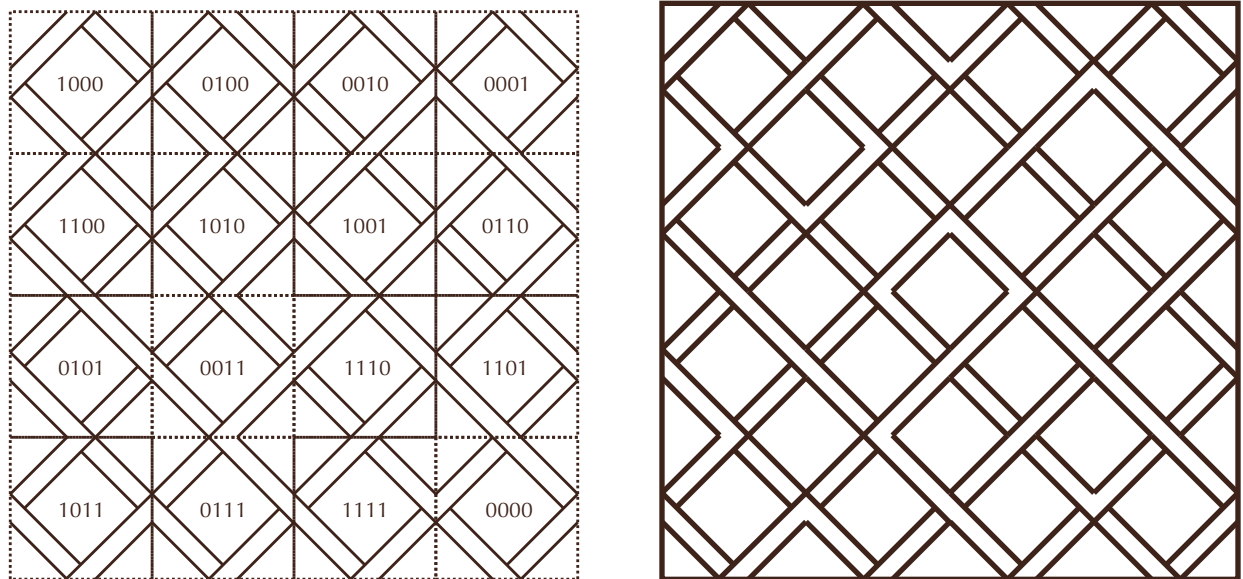


Figure 9: A 4x4 "Trellis" puzzle, with squares in a "Sol LeWitt" arrangement, shown with labels and guidelines (left) and as pure trellis (right).

The idea is to think of each tile as containing four thin slats of wood, running diagonally by quarter turns, with one slat lying over the other where they meet at the tile's edge. The binary numbers indicate whether the slat "entering" an edge (in a clockwise direction) lies over or under the slat "exiting" the edge, starting at the tile's topmost edge. One can now picture the pattern as a trellis, by joining the "upper slats" that meet from the two sides of each edge and likewise for the "lower" slats. Since there are 64 slats altogether, one can again ask if there's an arrangement of the tiles so that the trellis, again with toroidal connections at the outer edges of the 4x4 array, consists of a single loop. I again don't know the answer.

Alternative to toroidal identifications, one can imagine the pattern in a 4x4 arrangement as a ribbon that reflects with a crease when it hits an outer edge of the array. Amazingly, the "Sol LeWitt" arrangement in Figure 9 above *is* a single loop! So are the 4x4 "binary" arrangement and a "magic square" arrangement:

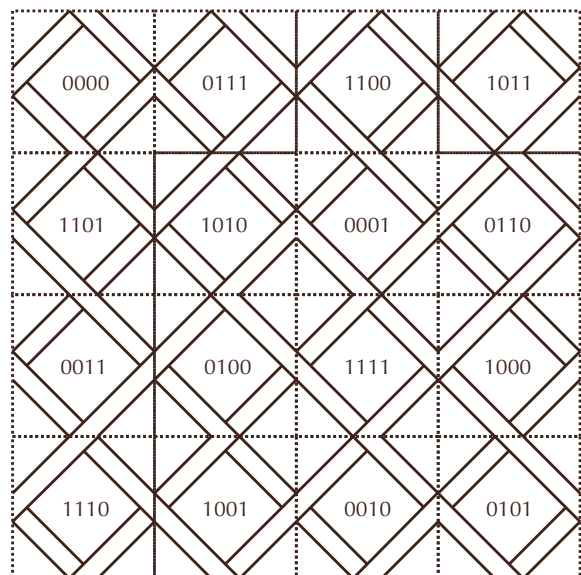
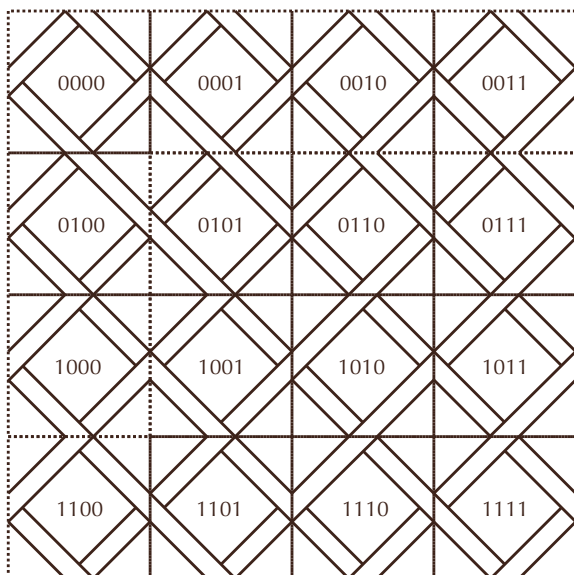


Figure 10: The Trellis puzzle in its “binary” arrangement (left) and a “magic square arrangement (right). Viewed as (non-toroidal) ribbons creased at the outer edges of the array, each is an example of a single loop.

Certainly not *every* arrangement of the nontoroidal “ribbon” trellis consists of a single loop, since it’s easy to arrange the tiles so as to produce short loops of length 4. But I suspect that a large number of arrangements do give a single loop. My only evidence of this, however, is the fact that the first time I tried a “random” arrangement, it turned out to be single-looped. Since there are $16! = 20,922,789,888,000$ different non-toroidal ways to lay out the 16 tiles, I’d be surprised indeed if I just got lucky. It might be worth someone’s time counting the exact number of single-loop arrangements. (It might be worthwhile doing the same for Jim Propp's and/or my non-toroidal versions of the looping puzzle problem.)

More recently, in 2021, I got to wondering if I could reduce the size of the looping puzzle from 4x4 to 3x3. That is, could I come up with a set of *nine* patterns that exhaust all the possibilities for some design criterion? (It also occurred to me to see if I could reduce things yet further to a 2x2 version of a puzzle. The ultimate, of course, would be to come up with a challenging 1x1 puzzle!) The problem is that 9 doesn’t easily relate to 4. But what finally occurred to me is that among the 24 permutations of four objects, exactly 9 are *derangements*, i.e., permutations that have no fixed elements. So this suggested two pairs of possibilities:




Figure 11: Two labellings of the thridpoints of a square (left) and two labellings of the corners and midpoints (right) that lend themselves to a “deranged” looping puzzle. The idea is to connect each “a” point to a “b” point with a *different* number.

The “thridpoint” and corner-midpoint labellings in Figure 11 produce these two sets of 9 different tiles:

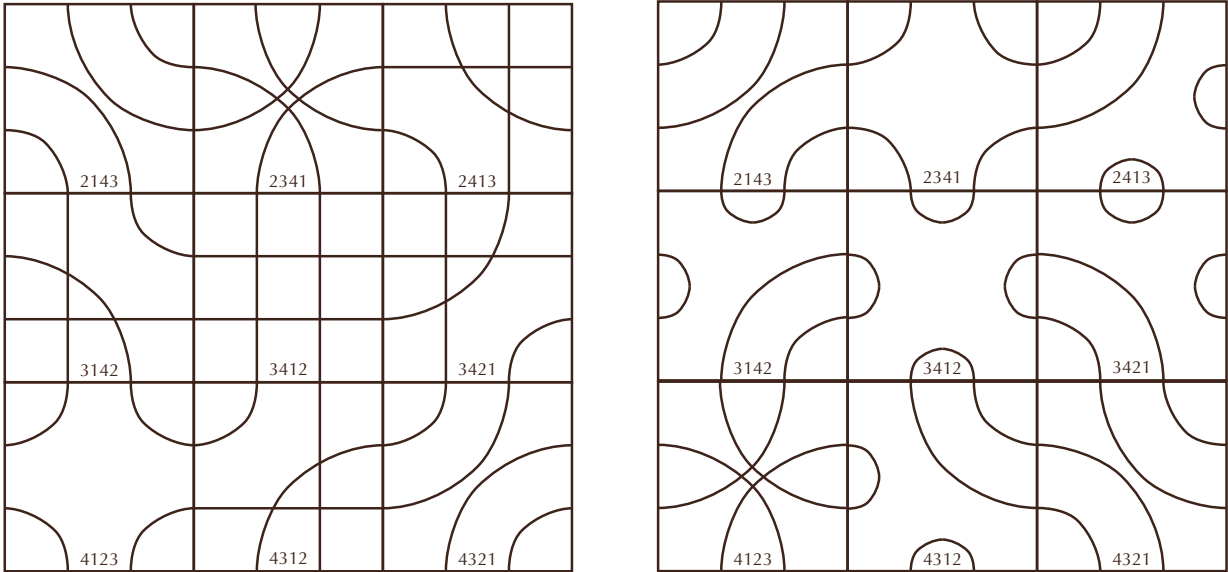


Figure 12: Two toroidal “derangement” puzzles based on connecting thridpoints as labeled in Figure 11 (left). Can either of these be rearranged so as to have a single toroidal loop?

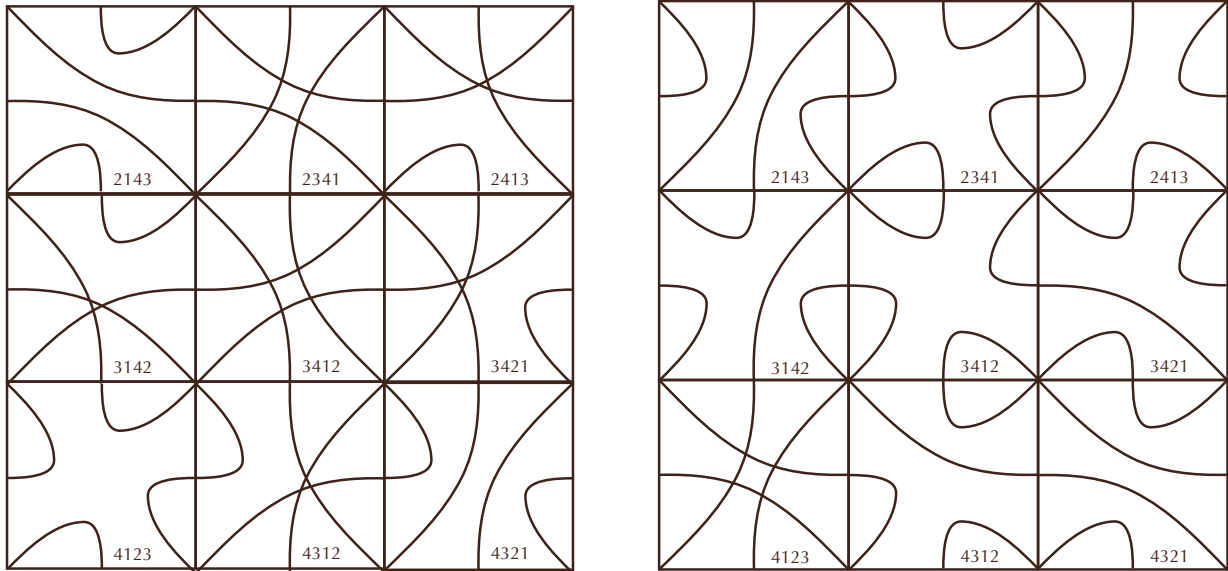


Figure 13: Two toroidal “derangement” puzzles based on connecting corners to midpoints as labeled in Figure 11 (right). Can either of these be rearranged into a single toroidal loop? (Note, the convention at corners is to continue from one square into the diagonally adjacent square.)

These final four puzzles are small enough that a brute-force (computer) search could easily resolve them: There are, after all, only $8! = 40,320$ toroidally different arrangements. (The 4×4 puzzles have $15! = 1,307,674,368,000$ toroidally different arrangements, which is big enough to call for some clever pruning.) I myself have not spent any time, either brute-force or cleverly, looking systematically for arrangements that give a single loop, but I have noticed that for two of the puzzles, randomly rearranging the squares often produces a single-loop solution, whereas for the other two I’ve yet to find a single-loop arrangement. I leave it to the reader to guess (and then check) which two are which – and, ideally, to figure out why.