

**A Pedestrian Approach to
A Method of Conway,
or,
A Tale of Two Cities**

James Propp
Massachusetts Institute of Technology

May 14, 1997

Tilings in two dimensions have long fascinated both professional and recreational mathematicians. The appeal is easy to understand: the objects being studied are concrete, and a person can develop useful intuitions about them through hands-on experimentation that verges on play. One special charm of the subject is that questions about tiling often take the form “Can you tile a ... with ...?”, so that an affirmative answer can be embodied in a single picture. The solver of such a problem may have spent hundreds of hours devising the picture, but a reader can verify its validity in a matter of minutes. (See [2] for an example of this.)

In contrast, a negative answer to a tiling problem — that is, the assertion that a certain kind of tiling does not exist — may not be so simple to verify. If the region to be tiled is finite and there are only finitely many allowed tiles, then it is possible in principle to determine the status of that tiling problem (possible versus impossible) by brute-force, examining all possibilities; if all possibilities can be eliminated then the desired tiling does not exist. But checking the validity of such a proof takes nearly as long as constructing the proof in the first place. What’s more, the brute-force approach does not allow one to make the jump from isolated cases to infinite families of tiling problems.

Fortunately, there are techniques for expeditiously proving that some tiling problems cannot be solved. One of the cleverest of these is a method invented by John Conway thirty-five years ago and developed further by Jeff Lagarias [1]. Although the Conway-Lagarias article is cast in the language of combinatorial group theory, the method often produces simple geometric criteria for non-tileability that can be applied to specific problems without knowledge of the algebraic machinery that gave rise to them.

Conway's theory often permits one to construct simple "certificates of non-tileability" for a region. For instance, we will see that Figure 7 is in a certain sense a visual proof of the unsolvability of a particular tiling problem. I like to think of it as a "proof-mandala": if one prepares one's understanding in the proper way (by absorbing some preparatory theorems about the set of allowed tiles) and contemplates the picture in the right frame of mind (by noting certain facts about the untiled portion of the region), then one can achieve enlightenment (perceive the futility of striving to tile the region). Putting it less fancifully, Figure 7 is the visual culmination of an argument; it clinches a proof of non-tileability once Conway's theory has been used to formulate a specific concrete criterion.

My purpose here is to give you a glimpse of Conway's ideas in as accessible a way as possible, avoiding combinatorial group theory entirely. The price that you will pay for my pedestrian approach is that some of the ideas in the proof of the main theorem will seem to come out of nowhere. I hope that the novelty and power of the method will intrigue you sufficiently that you feel impelled to read the original articles and learn about the algebraic framework that the proof fits into. In addition to Conway and Lagarias' article [1], there is also a paper of Thurston [8] that gives the method a slightly different slant. I will close the article by briefly discussing the more advanced viewpoints that these two articles take.

Tiles, tilings and tileability

The tiles we will consider are of the kind known as **polyominoes**. Every polyomino can be formed by taking a union of cells in a square grid, and the number of cells determines whether the polyomino is classified as a **monomino**, **domino**, **tromino**, **tetromino**, or what have you. (For a precise definition of polyominoes, see [5] and [7].) The five tetrominoes are shown in Figure 1. These shapes have received recent notoriety through the video game Tetris.

Figure 1. The tetrominoes.

When a large polyomino (a union of many cells) can be written as a union of small polyominoes that are disjoint except along their boundaries, we say that the large polyomino has been **tiled** by the small ones. It is easy to show that each of the small polyominoes must be a union of some subset of the cells that jointly constitute the large polyomino.

The book [7] discusses many of the tiling properties of tetrominoes, and I will not attempt a general overview here. Instead I am going to focus on one particular sort of tetromino, the so-called **skew tetromino**. This is the tetromino standing at the very left of our family portrait. The skew tetromino cannot be rotated so as to yield its own mirror-image without leaving the plane. Later, one of the skew tetromino's siblings will elbow its way into the action, namely, the **square tetromino**, standing off to the right — but that is getting ahead of our story.

Figure 2. The skew tetromino.

I am going to give criteria for recognizing when a region can be tiled by skew tetrominoes, using all four of the orientations shown in Figure 2. These criteria are necessary conditions for tileability, so that when one of

them fails, you can conclude that the region in question cannot be tiled. While such criteria can never give you assurance that a tiling problem has an affirmative answer, they can often save you from wasting time looking for a tiling where none exists.

One sort of region that will defeat any would-be tiler who has only skew tetrominoes on hand is a rectangle. For consider the skew tetromino that covers the upper-left cell of the rectangle; without loss of generality, we may suppose it is placed as shown in Figure 3. The placement of this skew tetromino forces the placement of another skew tetromino below it, and so on, leading to a configuration in the lower-left corner that cannot be extended to a tiling. This is a “quasi-local” obstruction to tiling, in that we only need to look at part of the region to be tiled to deduce untileability.

Figure 3. Failed tiling of a square.

Such quasi-local arguments for non-tileability fail us when we move on to the tiling problems that are the main subject of this article. Define an **Aztec diamond of order n** as a region consisting of $2n(n+1)$ unit cells, arranged in centered rows of lengths $2, 4, 6, \dots, 2n-2, 2n, 2n, 2n-2, \dots, 6, 4, 2$. Aztec diamonds were introduced in [3], where it was shown that the Aztec diamond of order n has exactly $2^{n(n+1)/2}$ tilings by dominoes. Figure 4 shows the Aztec diamond of order 5. This region also appears in problem 81 of [5], which asks whether the region can be tiled by the twelve pentominoes in such a way that each pentomino gets used exactly once; this problem (which also appears as problem 5.22 of [7]) was solved in the negative by Andy Liu [6] ten years ago.

Figure 4. Aztec diamond of order 5.

The difficulty with tiling an Aztec diamond is global: Figure 5 shows that we can tile a sizable portion of the Aztec diamond of order 17, leaving

uncovered only a 6-by-6 block that is far away from the boundary. On the other hand, since (as is easily shown) we can tile the whole plane with skew tetrominoes, we can start tiling our Aztec diamond from the center and work outward until only a narrow fringe along the boundary remains untiled. Thus, it is not difficult to tile any particular portion of the Aztec diamond — what is hard is tiling the whole thing.

Figure 5. Failed tiling of an Aztec diamond.

A well-known technique for proving that particular tiling problems cannot be solved is the use of **coloring arguments**. For instance, to show that the Aztec diamond of order 5 cannot be tiled by skew tetrominoes, impose the coloring shown in Figure 6, with two-by-two blocks of black cells alternating with two-by-two blocks of white cells (along with a few left-over white cells). A skew tetromino of any orientation placed anywhere in the region must cover three black cells and one white cell or vice versa. In particular, it must contain an odd number of black cells. Since any tiling of the entire region by skew tetrominoes must use exactly 15 tiles, and each tile contains an odd number of black cells, the entire region must contain an odd number of black cells. Since the number of black cells in the region is even, we have reached a contradiction, and no such tiling exists.

Figure 6. Colored Aztec diamond.

Suitably generalized, the coloring-scheme of Figure 6 tells us that the Aztec diamond of order n cannot be tiled by skew tetrominoes if $n(n+1)/2$ is odd, which happens whenever n is 1 or 2 more than a multiple of 4. But what about values of n for which $n(n+1)/2$ is even? A bit of experimentation should convince you that no tiling exists when $n = 3$ or 4, so one might suppose that a different coloring argument might be devised to

handle such cases. However, it can be shown (see [1]) that there cannot be a coloring argument (at least in the simplest sense of the term) that proves the untileability of the Aztec diamond of order n for even *one* value of n for which $n(n+1)/2$ is even.

This is where Conway's approach comes to our aid. Following his lead, we will soon see that for *no* positive integer n can the Aztec diamond of order n be tiled by skew tetrominoes. In fact, we will formulate a criterion (the main theorem) that will let us look at Figure 7 and, after a moment's inspection, announce that we are satisfied that the Aztec diamond of order 7 cannot be tiled by skew tetrominoes.

Figure 7. The proof-mandala.

Shadowing paths

To initiate you into the mysteries of such proof-mandalas, I am going to take you on an odd sort of imaginary excursion: you and I will in a sense be walking together, but we will be doing our walking miles apart, in two different cities.

For want of better names, I will call these cities Hoboken and Manhattan, with no slander intended toward either great city. Any resemblance between either of the imaginary places that I will describe and the real places whose names they bear is mostly coincidental.

In Hoboken, all the streets and avenues are one-way, with streets running west to east and avenues running south to north. If you like, you can imagine a two-way ring road around the city that will somewhat alleviate the inconvenience resulting from the city's eccentric traffic system. (See Figure 8.)

Figure 8. Hoboken.

The city is naturally associated with a portion of an infinite square grid;

for example, one can interchangeably speak of “vertices” (in the grid) or “intersections” (in the city). To keep track of the traffic pattern, let’s mark the edges of the grid with arrows that indicate the direction of traffic (solid arrows for the streets and unfilled arrows for the avenues). Note that each vertex has an ingoing solid arrow, an outgoing solid arrow, an ingoing unfilled arrow, and an outgoing unfilled arrow, and that every intersection looks just like every other.

As a pedestrian in Hoboken, I am free to go with or against vehicular traffic; at each corner I can go in any of the four possible directions. In this way I can trace a path in the grid whose edges connect adjacent vertices. Given my starting point, my journey is uniquely specified by the set of decisions I make as to whether to travel on a street or an avenue and whether to travel with traffic or against it. Conversely, any sequence of such decisions corresponds to an actual path through the city.

Figure 9. Manhattan.

Hoboken is the city in which I am going to do my walking; you, however, are going to be walking in Manhattan. This city also has one-way streets and avenues, but the direction of traffic along streets (and along avenues) alternates, instead of being in a single fixed direction throughout the city. (See Figure 9.) As before, each vertex has an ingoing solid arrow and an outgoing solid arrow and an ingoing unfilled arrow and an outgoing unfilled arrow (the first two horizontal and the last two vertical). There are four different sorts of intersections in Manhattan, but they are all related to one another by symmetries (reflections and rotations, to be precise).

Say a path in Manhattan **shadows** a path in Hoboken if they are described by the same instructions, in the following sense. As I take my walk in Hoboken, you and I will talk by cell-phone and I will tell you what I am doing: *not* whether I am going north, south, east, or west, and *not* whether

I am turning left or right, but merely whether I am going on a street or an avenue, and whether I am going with traffic or against it. If you and I walk at the same pace and you imitate my path as described by me over the cell-phone (which can be done in one and only one way, given your starting point), you are tracing out the **shadow** of my path.

Figure 10. Shadowing a path.

For example, consider the path in the top half of Figure 10 that starts at the marked point and traverses the boundary of the skew tetromino in the counterclockwise sense, ending where it started. If I traverse this path in Hoboken, then the description of my path is $S^+S^+A^+S^+A^+S^-S^-A^-S^-A^-$, where S and A denote street and avenue respectively and $+$ and $-$ determine whether I am going with or against traffic. If I describe my path to you in this fashion by cell-phone, and you imitate it in Manhattan, you will travel along the path in the bottom half of Figure 10, traveling clockwise around the cell on the right and counterclockwise around the cell on the left.

Let's call a walk **closed** if its terminal point coincides with its initial point. The shadow of a closed walk in Hoboken need not be closed in Manhattan; a walk that encircles a single block in Hoboken is an example of this. However, we do have the following crucial fact:

Claim 1: If a region in Hoboken can be tiled by skew tetrominoes, then any closed walk in Hoboken that makes a complete circuit of the boundary of that region is shadowed by a closed path in Manhattan.

Proof: We use induction on the area of the region, necessarily a multiple of 4. If the tileable region is just a single skew tetromino then we are in the situation of Figure 10. To see that this figure is the only one we need to examine for the base case of our induction, note first that if you had chosen

a different starting point for your walk in Manhattan, your walk would still have been closed, since Manhattan has symmetries (translations, rotations, and reflections) carrying any vertex to any other and respecting the pattern of arrows. Neither would the situation be affected if I had chosen a different starting point for my walk along the boundary of the skew tetromino, for that would be tantamount to having you start your walk at a different point in Manhattan. Finally, we have dealt with only one of the four orientations a skew tetromino can have, but the symmetries of Hoboken and Manhattan imply that what works for one orientation must work for the other three.

Now suppose we've proved the claim for all tileable regions with area less than $4n$. Consider a region R with area exactly $4n$ that comes equipped with a particular tiling. To prove that the boundary of R (a closed curve in Hoboken) is shadowed by a closed curve in Manhattan, I am going to modify my itinerary a bit, as suggested by Figure 11. Specifically, I am going to travel from a starting point p on the boundary of R to another point q on the boundary of R ; then travel from q to p in the interior of R , traveling *only* on the boundaries of tiles in my tiling of R ; then return from p to q by the same interior route; and finally travel in the same direction as before (clockwise or counterclockwise) from q to p along the boundary of R , completing my tour. The polygonal arc pq divides R into two regions, which we will call A and B .

Figure 11. The induction step.

Since A and B each have area less than $4n$, and since each is tileable by skew tetrominoes, the boundary of each of them is shadowed by a closed path in the shadow-grid. Thus, the modified itinerary that I described above (from p to q to p to q to p) is shadowed by a closed path. However, the second and third legs of this journey (from q to p to q in the interior of R) are “inverses” of each other. It follows that as your walk in Manhattan

shadows my (modified) walk in Hoboken, the second and third legs of your journey will also be inverses of each other, and the two together will leave you right where you were at the end of the first leg, when I was first at q . Hence, excising this pointless detour from my journey (and from yours), we see that the shadow of a path encircling R in Hoboken is also a closed path in Manhattan. This verifies the induction and completes the proof. \square

Signed area and the main theorem

We can strengthen Claim 1 using the notion of the signed area “enclosed” by a shadow path. To define this, we must first define the winding number of a finite closed curve around a point not on the curve. If we draw a ray emanating from the point, the winding number is simply the number of times that the curve crosses from one side of the ray to the other in the positive direction (counterclockwise) minus the number of times that the curve crosses from one side of the ray to the other in the negative direction (clockwise). If the closed curve consists of grid-edges, then the winding number of the curve around a point is the same for any other point in the same cell. We call this the winding number of the curve around the cell. For all but finitely many cells, the winding number of the curve around the cell is zero. Thus, we can speak of the sum of the winding numbers of the curve around all the cells; this is what I mean by the signed area enclosed by a curve. The reader should check that when the closed curve is simple, the signed area as I have defined it is equal to plus or minus the area enclosed by the curve as defined in the usual sense (positive if the curve winds counterclockwise, negative if the curve winds clockwise).

If the region R in Hoboken has a boundary whose shadow in Manhattan is a closed curve, we define the (signed) **shadow-area** of the region R as the signed area enclosed by the shadow of its boundary. We now show that a region that can be tiled by skew tetrominoes must have shadow-area equal

to zero.

Claim 2: If a region in Hoboken can be tiled by skew tetrominoes, then any closed walk in Hoboken that makes a complete circuit of the boundary of that region is shadowed by a closed path in Manhattan that encloses signed area 0.

Proof: Again we use induction. The same pictures that worked before still work; we just have to examine them in a different frame of mind. In Figure 10, we need to notice that the shadow-curve has winding number $+1$ around one cell, -1 around another cell, and 0 around every other cell, giving it a total signed area of 0. As for the induction in Figure 11, we need to observe that for every place where the shadow of the internal path from q to p crosses a ray, the shadow of the return path from p to q crosses the same ray in the opposite direction. Thus the shadows of the second and third legs of your trip in Manhattan make canceling contributions to the winding number around any particular cell, and thus make canceling contributions to the signed area of the complete path. Excising the detour, we obtain the desired induction. \square

You (the reader, not the walker) are now in a position where you can convince yourself of the untileability of many specific regions, including as it happens Aztec diamonds of any order. Specifically, take the boundary of your region in the Hoboken grid and shadow it in the Manhattan grid; if the resulting curve is not closed, or if it is closed but encloses non-zero signed area, then your region cannot be tiled by skew tetrominoes.

Satori

If you want to convince someone else that the region cannot be tiled, using the method as we have discussed it so far, then that person has to do

essentially everything that you did. That is, he or she must carefully shadow the entire boundary and (if the shadow-path closes) work out the signed area of what may be a rather crazy self-intersecting closed curve. This is certainly a better way for you to convince people that a tiling problem is unsolvable than forcing them to read through a three-inch stack of coffee-stained sheets of graph paper in which all possibilities are tried and eliminated. But there is a still better way, in which a region R can be marked up in such a manner that, in the case where the boundary of R is shadowed by a closed curve, the proof-checker can see at a glance how much signed area is enclosed by the shadow-curve.

Recall that a square tetromino is a 2-by-2 square, as shown on the far right in Figure 1. The shadow of the boundary of a square tetromino in Hoboken is the boundary of a square tetromino in Manhattan, but the orientation may switch, giving it signed area $+4$ or -4 . We need to nail down the sign exactly. Let's choose a particular corner in Hoboken and call it the Hoboken origin; we will say that another corner in Hoboken is **even** if it can be reached from the origin in an even number of steps (lengths of a city block) and **odd** if it can be reached from the origin in an odd number of steps. Meanwhile, let's find an intersection in Manhattan of "Hoboken type" (i.e., where the street goes east and the avenue goes north), and let that be the Manhattan origin. We define evenness and oddness of Manhattan intersections in an analogous way. If we assume that you and I begin our walks at vertices of the same parity, then it follows that you and I will be at vertices of the same parity as one another forever afterwards, since every block we walk changes the parity of your location and mine from odd to even or from even to odd.

Call a square tetromino in either grid **even** if its corners are at even vertices and **odd** otherwise. Consider the boundary of such a tetromino in Hoboken, traversed in the counterclockwise direction. The shadow of this path is also the boundary of a square tetromino, and with a little doodling you can check that the shadow path in Manhattan encircles signed area $+4$

or -4 according to whether the original square tetromino was even or odd.

We can now state our

Main theorem: Suppose a simply-connected region in the plane can be tiled by a mixture of skew tetrominoes and square tetrominoes. Then the number of even square tetrominoes minus the number of odd square tetrominoes does not depend on what tiling one chooses; i.e., it is an *invariant*. In particular, if the difference is non-zero for one such tiling, then the region cannot be tiled by skew tetrominoes alone.

Proof: The difference in question, when multiplied by 4, is just the signed area enclosed by the shadow of the boundary of the region to be tiled, because each tetromino is shadowed by a path enclosing signed area $+4$, -4 , or 0 according to whether it is an even square tetromino, an odd square tetromino, or a skew tetromino. The justification is the same as in the proof of Claim 2, namely, the additivity of signed area. \square

Now look back at Figure 7. The Aztec diamond of order 7 has been decomposed into a number of skew tetrominoes along with four square tetrominoes. Because all four square tetrominoes are even, the invariant has value $+4$. It follows that the region cannot be tiled by skew tetrominoes alone. In this sense, Figure 7 can be a proof-mandala for the impossibility of tiling the Aztec diamond of order 7 by skew tetrominoes, once the mind has absorbed the main theorem.

The mandala has even more to teach the receptive spirit. Notice that the Aztec diamond of order 7 has an Aztec diamond of order 6 sitting inside it, fringed above by skew tetrominoes and a single square tetromino. This order-6 diamond in turn contains of an order-5 diamond fringed above by skew tetrominoes. And so on. The mandala shows a clear iterative pattern for reducing an Aztec diamond of order $2k$ to an Aztec diamond of order

$2k - 1$ plus some skew tetrominoes, and for reducing an Aztec diamond of order $2k + 1$ to an Aztec diamond of order $2k$ plus some skew tetrominoes and a single square tetromino, in such a way that the square tetrominoes all have the same parity. Thus, for general n , any tiling of the Aztec diamond of order n by skew tetrominoes and square tetrominoes must have an excess of exactly $\lfloor \frac{n+1}{2} \rfloor$ square tetrominoes of one particular parity. In particular, for $n \geq 1$, there can be no tiling of the Aztec diamond of order n by skew tetrominoes alone.

It would be interesting to know of a different way to prove the main theorem. One possible approach would be to mimic Donald West's proof [9] of Conway's triangle-tiling theorems, and show that every tiling of a simply-connected plane region by skew tetrominoes and square tetrominoes can be obtained from every other such tiling by means of a small repertoire of "local moves," each of which preserves the difference between the number of even square tetrominoes and odd square tetrominoes. Readers might experiment with such local moves and see if they can come up with a demonstrably complete set. My own guess is that a complete finite set of local moves does exist.

Extensions

As was mentioned earlier, the shadow-path method can be used to show that many particular regions R , and not just Aztec diamonds, cannot be tiled by skew tetrominoes. In fact, if one chooses R "at random," then it is likely that the shadow of its boundary will not be a closed path, or if it is closed, that it will not enclose signed area 0. Even if it does enclose signed area 0, a slight refinement of our approach may permit us to rule out the existence of a tiling. For, observe that the shadow of the boundary of a skew tetromino winds clockwise around one cell in the shadow-grid and counter-clockwise around another cell *of the same color*, relative to the coloring shown in Figure 12. Thus, we can talk about (signed) A-area, B-area, C-area, and

D-area, whose sum will be the signed area enclosed by a curve. In order for R to be tileable by skew tetrominoes, the shadow of the boundary of R must enclose A-area, B-area, C-area, and D-area all equal to 0.

Figure 12. Coloring the grid-squares.

This strengthened version of the main theorem goes a long way toward closing the gap between demonstrably tileable regions and demonstrably untileable regions. However, the gap is not altogether shut. For instance, the region of area 8 shown in Figure 13 is not tileable by skew tetrominoes, despite the fact that the main theorem (even in its strengthened form) does not tell us this. It would be valuable to know of an efficient algorithm that would close the gap completely, by deciding whether a given region is or is not tileable by skew tetrominoes.

Figure 13. An untileable region.

It is also interesting to change the game and weaken the notion of “tileability” so that the preceding necessary condition becomes sufficient as well. We define a **tile homotopy** of a path in the grid as a process of perturbing the path by “pulling it through tiles”. More precisely, an elementary homotopy between two closed grid-paths replaces a part of the path (call it P) joining two vertices p, q by another grid-path P' joining the same two vertices, such that P and P' together form the boundary of a tile. Figure 14 shows a series of elementary tile homotopies between the path shown in Figure 13 and the trivial loop; first the boundary is pulled outward by adding a skew tetromino on the outside and pulling the boundary through the new tile; then the path is pulled inward, using a tiling of the enlarged region by skew tetrominoes. It can be shown using combinatorial group theory that a closed path in Hoboken is tile-homotopic to the trivial loop if and only if its shadow

in Manhattan is closed and encircles A-area, B-area, C-area, and D-area all equal to 0.

Figure 14. Tile homotopy.

The shadow-path method can be used to prove that in any partial tiling of the Aztec diamond of order n , the diameter of the untiled portion is bounded below by a constant times \sqrt{n} . For, one can find a small rectangle that covers the untiled portion and approximate the boundary of this rectangle by a closed loop L that travels only along the boundaries of tiles but is as direct as possible subject to that constraint. Since L is tile-homotopic to the boundary of the Aztec diamond, whose shadow encloses signed area roughly $2n$, the shadow of L must have length at least $c\sqrt{n}$ for some constant c . Hence L itself must have length at least $c\sqrt{n}$, implying that the untiled portion of the Aztec diamond has large diameter. It is hard to imagine a proof of such a result by means of the methods that predated Conway's work.

On the other hand, Aaron Meyerowitz has shown that there can be no analogous bound on the area of the untiled portion, by pointing out that for any $n \geq 1$ it is possible to tile the Aztec diamond of order n by $\frac{n(n+1)}{2} - 1$ skew tetrominoes, a single L -shaped tromino, and a single monomino. Specifically, we can tile the Aztec diamond of order n minus a 2 -by- $2\lfloor \frac{n+1}{2} \rfloor$ rectangle that butts up against a corner of the diamond, as shown in Figure 15. This rectangle can then be tiled by skew tetrominoes leaving only a tromino and a monomino unaccounted for. Note, however, that the tromino and the monomino are far apart, as the preceding lower bound on the diameter of the untiled portion requires.

Figure 15. Variant proof-mandala.

One disadvantage of the elementary approach to Conway's invariants that

has been adopted here is that the reader may be left feeling convinced but mystified: How might anyone dream up the traffic patterns of Hoboken and Manhattan that proved so useful here? And how can the arguments used here for skew tetromino tilings be generalized to handle other tiling problems? The answer lies in the notion of the tile homotopy group; the interested reader should consult [1]. A major idea in the Conway-Lagarias paper is to regard graphs like Figures 8 and 9 as Cayley graphs of groups, as I have implicitly done in the way I labeled and oriented the edges. One satisfying feature of the tile homotopy viewpoint is that coloring arguments of the sort considered earlier turn out to be a special case of Conway’s method. Specifically, coloring arguments are associated with abelian homomorphic images of the tile homotopy group.

Thurston’s follow-up paper [8] presents a more geometrical way of looking at tile homotopy that suggests that Cayley graphs of groups are not at the heart of the method. Given a collection of subsets of the plane (to be viewed as the set of all allowed locations of tiles), Thurston invites us to create a topological space by taking the disjoint union of all those subsets (imagined if you like as floating above the plane) and identifying two points on the boundaries of two such tile-regions if they lie above the same point in the plane. Assuming that the tiles are all simply-connected, it is easy to see that the boundary of any tileable region corresponds to a path in Thurston’s space that is homotopic to the trivial loop. Thus we obtain a necessary condition for tileability from homotopy considerations, though actually exploiting this connection may be difficult in practice without recourse to Cayley-graph tricks. It turns out that the idea of boundary invariants can be extended beyond the realm of tilings, using Thurston’s more geometrical approach; details appear in [4].

I will end with an accessible puzzle that has a positive solution. The mandala teaches us that when $n = 2k^2 - 1$, the Aztec diamond of order n can be tiled by skew tetrominoes and k^2 square tetrominoes (all having the same

parity). Figure 5 shows us that in the case $k = 3$, we can arrange things so that these k^2 square tetrominoes are all at the center of the diamond, forming a $2k$ -by- $2k$ square. Can you find a way to do this for all k ?

The research was supported in part by NSA grant MD904-92-H-3060. I thank Art Duval, Jeff Lagarias, Aaron Meyerowitz, and Bruce Smith for their comments on early drafts of this article. Thanks are also due to Marjorie Senechal for inviting me to attend the Regional Geometry Institute at Smith College in July of 1993, and to the National Science Foundation for providing funding for my participation; this paper benefited enormously from the opportunity to present Conway's method to a diverse audience. No funds from either the Hoboken or Manhattan Chamber of Commerce were accepted during the writing of this article.

References

- [1] J. H. Conway and J. C. Lagarias, Tilings with polyominoes and combinatorial group theory, *Journal of Combinatorial Theory A* **53** (1990), 183–208.
- [2] K. A. Dahlke, The Y-hexomino has order 92, *Journal of Combinatorial Theory A* **51** (1989), 125–126.
- [3] N. Elkies, G. Kuperberg, M. Larsen and J. Propp, Alternating sign matrices and domino tilings, *Journal of Algebraic Combinatorics* **1** (1992), 111–132, 219–234.
- [4] W. Geller and J. Propp, The projective fundamental group of a \mathbf{Z}^2 -shift, *Ergodic Theory and Dynamical Systems* **15** (1995), 1091–1118.
- [5] S. Golomb, *Polyominoes: Puzzles, Patterns, Problems, and Packings* (2nd edition). Princeton University Press, 1994.
- [6] A. Liu, Golomb's twelve pentomino problems, in *The Lighter Side of Mathematics*, Proceedings of the Eugène Strens Memorial Conference on Recreational Mathematics and its History, R. K. Guy and R. E. Woodrow, editors. Mathematical Association of America, 1994, pp. 206–211.

- [7] G. Martin, *Polyominoes: A Guide to Puzzles and Problems in Tiling*. Mathematical Association of America, 1991.
- [8] W. Thurston, Conway's tiling groups, *American Mathematical Monthly* **97** (1990), 757–773.
- [9] D. West, An elementary proof of two triangle-tiling theorems of Conway and Lagarias, preprint.