Aperiodic Tiling

Andrew Glassner

Microsoft Research

People love patterns. We find recurring patterns everywhere we look—in the structures of rocks, the personalities of our friends, and the cycle of seasons in the natural world around us. Theme and variation have been a staple of creative invention since time immemorial, spanning every form of creative endeavor. I suggest that theme and variation represent a balance between the extremes of regularity and sheer randomness.

Imagine a million grains of sand, placed one next to the other in a straight line. Dullsville. Now imagine those grains of sand scattered at random over a rocky surface. That's just as dull as before. But take those grains and stack them up, let the wind play over them and shape them into dunes, and you've got something beautiful and interesting. Balancing between repetition and randomness can lead to patterns that draw us in and keep our interest.

One of the most interesting ways of assembling small units is along one of the lattices that make up crystals. In this column I'll live entirely in a 2D world, so our "crystals" will be nothing but collections of polygons in the plane. It's well known that there are only three regular ploygons that can tile the plane, as shown in Figure 1. Here the verb "tile" means to cover the infinite plane with a set of polygons so that no gaps or overlaps exist among the polygons. Each polygon is called a "tile," and the composite pattern is called a "tiling."

The tilings in Figure 1 may be theoretically interesting, but because they're so perfectly repetitive, they look boring. Usually when you use one of these patterns to cover a wall or floor, you decorate the tiles with colors or shapes to make the whole thing a little more interesting—again balancing the tiling's regularity with the variation in the tiles themselves. I did this in Figure 1 by using two or three colors, but it didn't help much. You can generate more interest by allowing more than one regular polygon in the game, resulting in the semi-regular tilings in Figure 2. This looks better, but regularity still dominates the patterns.

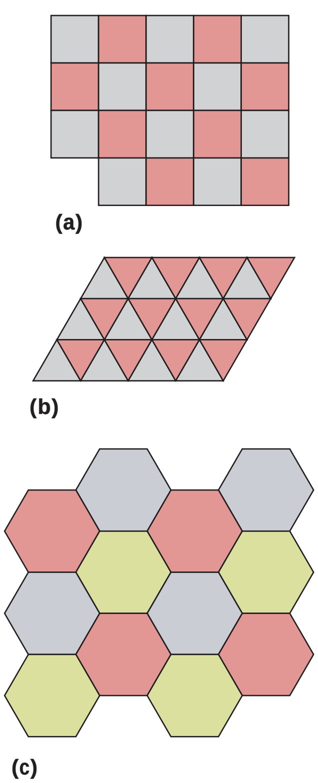

1 Three regular tilings made of many copies of a single regular polygon.

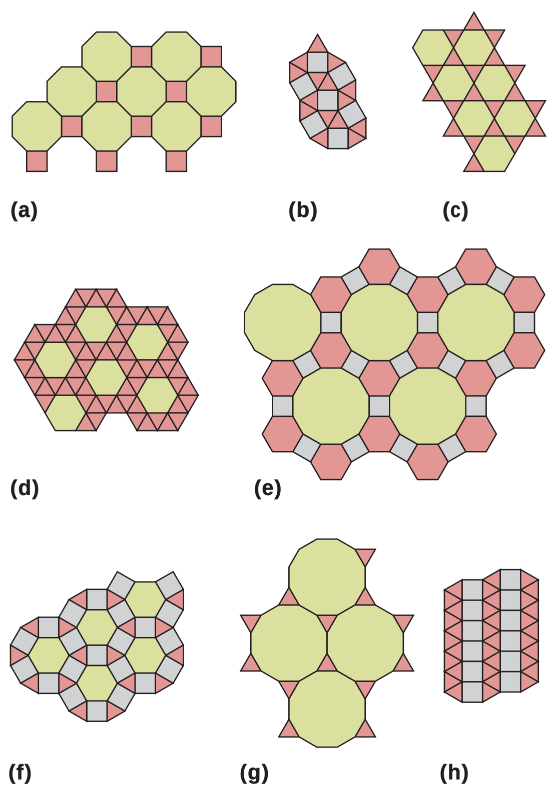

2 Eleven semi-regular tilings made of one or more regular polygons. Every vertex contains the same kinds of polygons meeting in the same order.

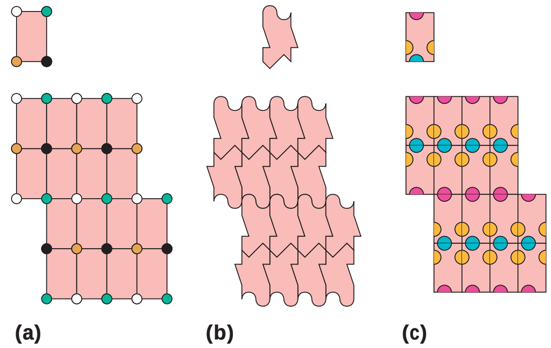

3 **Matching rules. (a) Vertex rules. Corresponding colors must overlap. (b) Edge rules. The deformed edges must link together. (c) Face rules. The decoration on the faces must be continuous.**

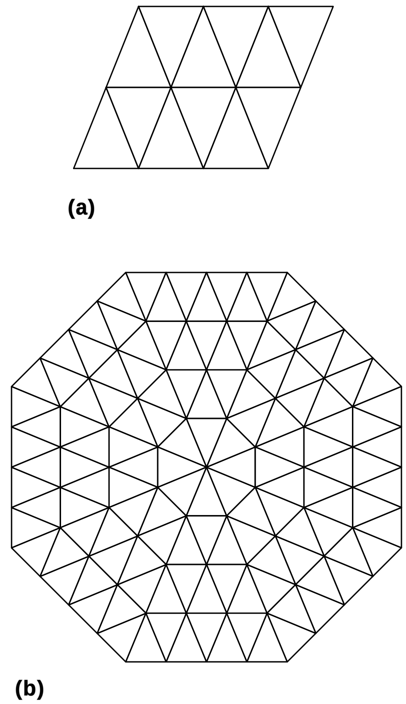

4 **(a) Isosceles triangles can tile the plane periodically. (b) The same triangles can create a radial tiling, which is not periodic.**

We can control our tiling patterns by adopting *matching rules*. These describe the permitted ways in which tiles can be placed next to each other. Figure 3 shows a few examples for assembling such a tiling from a single piece. All the examples in this figure share a matching rule that says tiles may only be placed so that the shared edge is of the same length (that is, we can't push a short edge up against a long one). I also distinguished the two short ends. In Figure 3a, I used a matching rule that says adjacent vertices must have the same color. Figure 3b enforces the same result by modifying the common edge into a mountain, so the tiles can only interlock in the desired way. In Figure 3c the rule stipulates that colored bands must be continuous across tile edges. Whether described as vertex, edge, or face rules, we get the same pattern.

The patterns in Figures 1, 2, and 3 are *periodic*. Essentially this means you find a group of tiles that you can pick up and use as a rubber stamp. In other words, the pattern is created by a single patch—called the *fundamental cell*—that repeats by translation to cover the plane. The distance from one copy to the next is called the *period* of the pattern. The translational part of the definition is important. Figure 4a shows a periodic tiling built from an isosceles triangle. Figure 4b shows a rotationally symmetric pattern built from the same triangle, but this one isn't a periodic tiling, even though there's obviously a great deal of internal structure.

You can create tilings in surprising ways. Figure 5 shows a spiral tiling (which is also nonperiodic), inspired by a construction by Heinz Voderberg. Although I used colors for decoration, this pattern is made up of many copies of only one tile. I found that the easiest way to construct a cell for this sort of tiling is to start with an acute isoceles triangle, as shown in Figure 6. The tiling works if the acute angle goes into 360 degrees an even number of times (I used 16 for this figure, so the acute angle is 360/16 = 22.5 degrees). Let's see how this works. First, create a parallelogram from the triangle. Draw a curve, and then rotate that curve by 180 degrees to get it to match up with the opposite corner. Now take another copy of the triangle and move its acute point to the lower right corner of the parallelogram. Make a copy of the curve and rotate it around the lower right corner of the parallelogram by the acute angle. Now join the sides of the

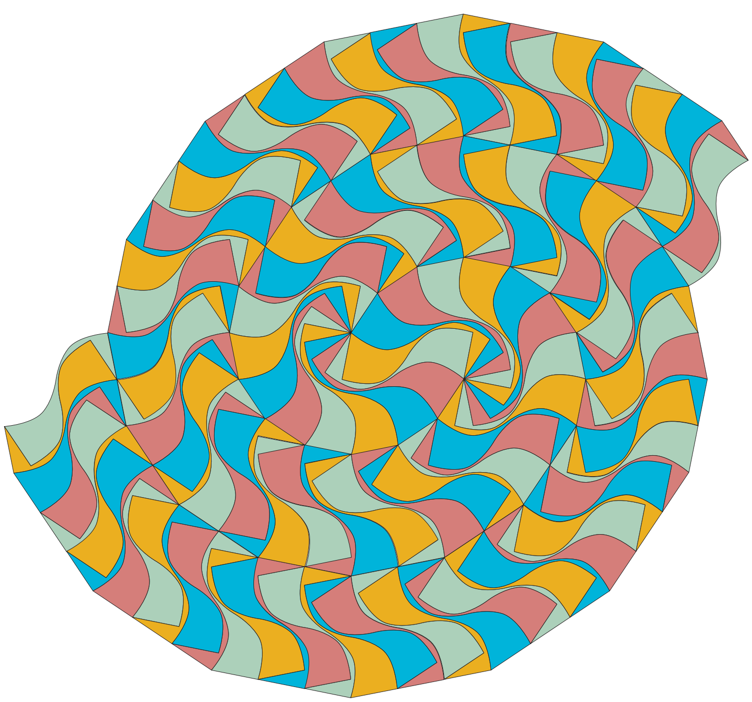

5 **A spiral tiling made of a single tile. The tiling can be extended to cover the entire plane.**

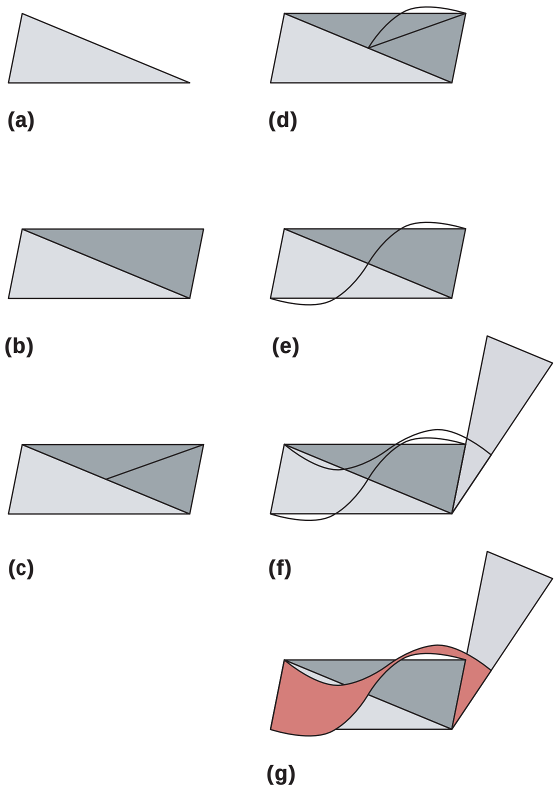

6 **Creating a tile for Figure 5. (a) Start with an isosceles triangle with an angle of 360/*n* degrees (here *n* = 16). (b) Make a parallelogram. (c) Find the midpoint. (d) Draw a curve from the upper right corner to the midpoint. (e) Rotate the curve 180 degrees. (f) Rotate around the lower right corner by an angle equal to the acute angle of the triangle. (g) Join the edges.**

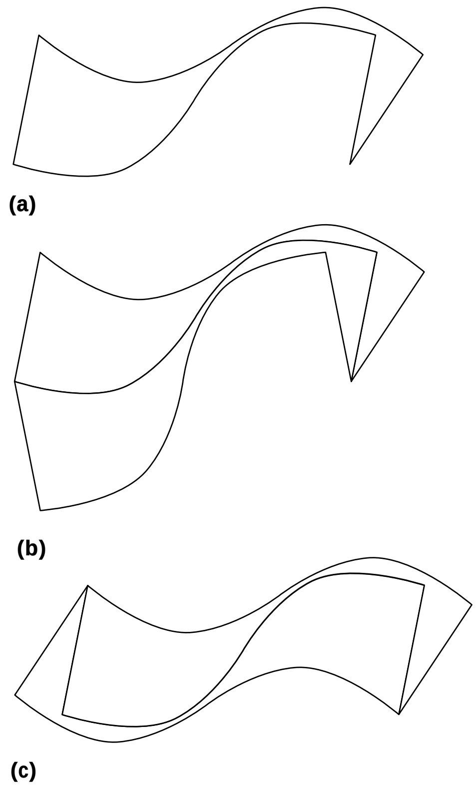

7 **(a) The tile made in Figure 6. (b) This tile fits into itself by rotation. (c) The tile also fits into itself by a vertical reflection, making a block that can tile the plane periodically.**

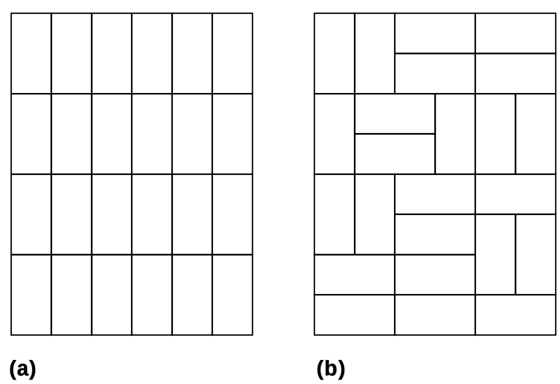

8 **(a) The rectangle can tile the plane periodically. (b) The rectangle can also tile nonperiodically.**

curves, and voilá. The only tricky part is ensuring the curves don't intersect each other.

Figure 7 shows that a tile produced this way has two important properties: it fits into itself both by a small rotation equal to the acute angle of the original triangle and by a full 180-degree rotation. Notice that the spiral tiling in Figure 5 uses both properties: the two centers use slightly rotated copies of the tile, and they join each other with a half-rotated pair. From there it's pretty straightforward to add tiles and follow the spiral outwards. Notice, though, the pattern changes slightly each time it wraps around 180 degrees.

Typically in any periodic tiling an infinite number of fundamental cells exist, but often there's only one smallest such cell. For example, in Figure 1b the smallest fundamental cell is a single square, but you could use two squares adjacent horizontally or vertically, or a 2-by-2 square of squares, and so on.

It's easy to build nonperiodic tilings—simply don't create periodic patterns. Figure 8a shows a periodic tiling of rectangles, and 8b shows the same rectangles in a nonperiodic tiling. Theoretically we could extend the pattern of Figure 8b, always mixing things up, so that the final result is not periodic.

Can you create tiles that only tile the plane nonperiodically? The rectangles of Figure 8 don't fill the bill, since they can be coerced into a periodic pattern. If there's a set of one or more tiles that fit together to create exclusively nonperiodic tilings, that set of tiles—and the resulting pattern—is called *aperiodic*.

To my eye, aperiodic tilings made of a few distinct tiles fit into that desirable class of patterns that balance regularity (because of the recurring instances of just a few tile shapes) and randomness (because the pattern never repeats). This column is the first of two that talk about aperiodic tiles, how to make them, and what you might do with them. The second column will focus on the details of a specific family of tiles.

From 20,000 to 2

The question of whether any aperiodic sets of tiles existed at all went unanswered until recently. In 1961 Hao Wang conjectured that there were no aperiodic tiles. That is, any set of tiles that could tile the plane aperiod-

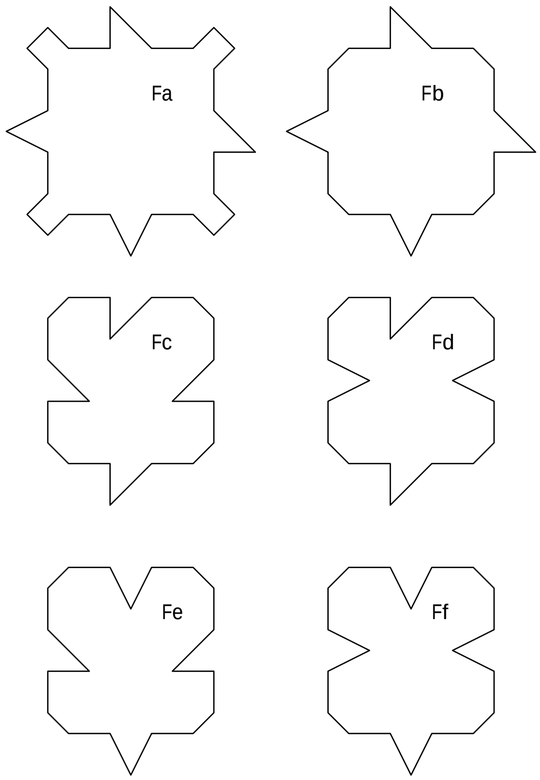

9 **The Robinson tiles. Pieces can only go together so that the tabs fit the slots and the corners are covered.**

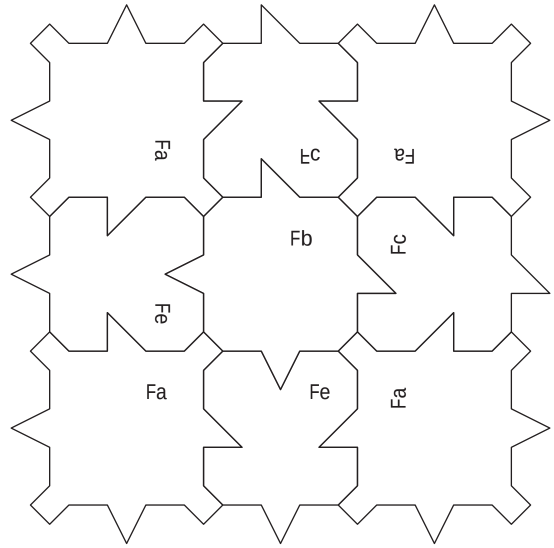

10 **Building a 3-by-3 Robinson block.**

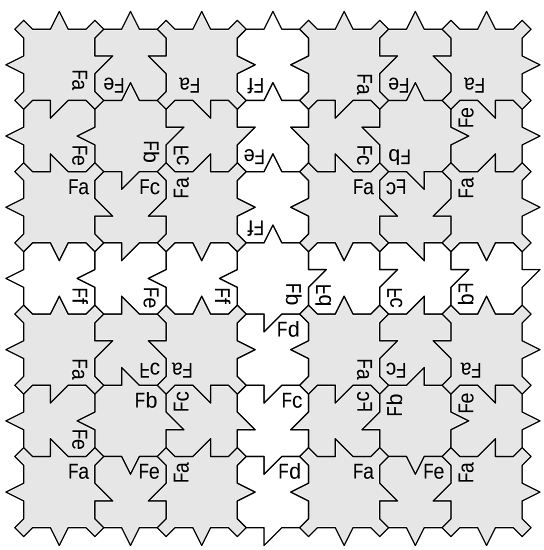

11 **A 3-by-3 Robinson block (shaded grey) occupies each corner of this 7-by-7 pattern. The glue tiles make a large plus sign.**

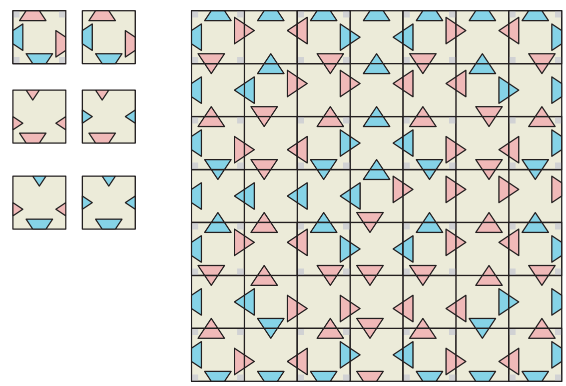

12 **A decoration of Figure 11. The top of row tiles are Fa and Fb, the second row Fc and Fd, and the third row is Fe and Ff.**

Further Reading

The fundamental reference on tiling is *Tiling and Patterns* by Grünbaum and Shephard, published by W.H. Freeman (1987, New York). It's a beautiful book, in addition to being the standard reference work on the subject. Martin Gardner has written several columns on tiling. Two are reproduced as chapters in his book *Penrose Tilings to Trapdoor Ciphers*, also published by W.H. Freeman (1989, New York). There are also some very interesting discussions in *Connections* by Jay Kappraff, published by McGraw-Hill (1991, New York).

ically could also do so periodically. In 1964 Robert Berger disproved this conjecture by inventing a set of more than 20,000 dominoes that formed an aperiodic set.

Once aperiodic tilings were proven to exist, many people set about to find smaller sets. Berger later found a set of 104 tiles, and Donald Knuth found a set of 92 tiles. Robert Ammann and Raphael Robinson independently found several different aperiodic sets of six tiles. Roger Penrose holds the current record for the smallest set—his sets contain only two tiles.

It's interesting that the Ammann tiles and Robinson tiles are both based on squares. This is great for computer graphics, where we deal with square grids in everything from sampling grids to texture patterns. So I'll kick things off with a look at the square-based Robinson tiles.

Robinson tiles

Figure 9 shows one set of Robinson's tiles. Following Grünbaum and Shephard, I used edge modifications to show how the tiles lock together (see the Further Read-

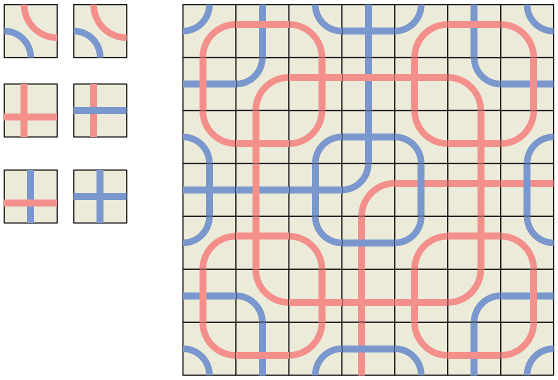

13 **A decoration of Figure 11.**

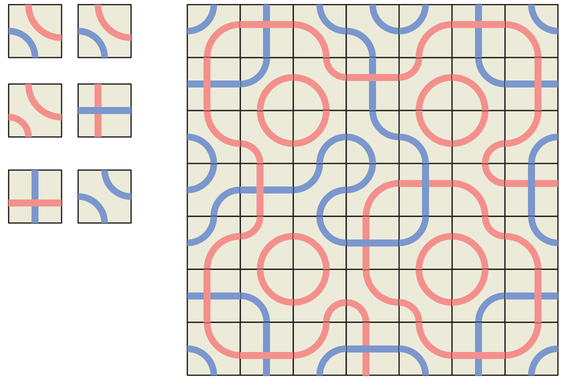

14 **A decoration of Figure 11, and a variation of Figure 13.**

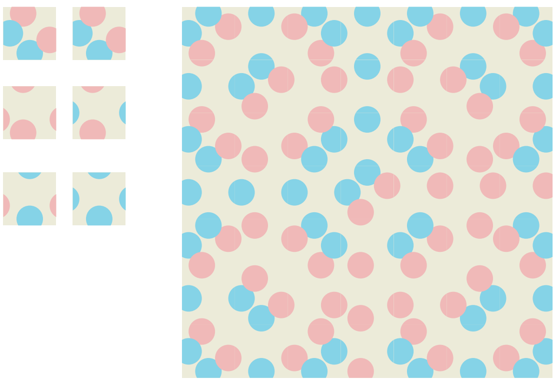

15 **Another decoration of Figure 11.**

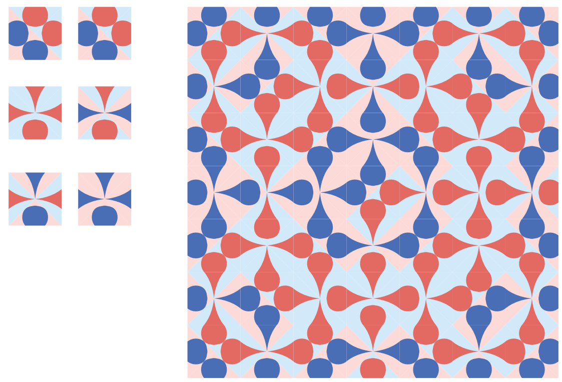

16 **A different decoration of Figure 11.**

ing sidebar for more information). I also included labels a through f on each tile, as well as the letter F to help show the tile's orientation.

Notice that tile Fa has a little tab in each corner, while the other tiles don't. This means that exactly one Fa tile must appear at each vertex of the tiling. To assemble the pattern, just put the pieces together so they fit. You can rotate and flip these tiles as needed.

Starting with a single Fa tile, you can build the 3-by-3 block in Figure 10. Except for some choices of orientation, this is the only 3-by-3 block you can make with these tiles.

To grow the pattern, you can run a new row of tiles above it and to its right. Then rotate the 3-by-3 block and affix copies of it to the right and top of the 4-by-4, and then place a fourth rotated copy into the upper right corner. Figure 11 shows the resulting 7-by-7 grid. You can then run a new row of tiles above and to the right of the 7-by-7 grid to grow a 15-by-15 block, and so on.

Part of the fun of working with aperiodic tilings is decorating the tiles so that you get interesting-looking visual patterns. You can approach these decorations in two ways. One is to come up with designs that enforce the matching rules, so you can dispense with the edge modifications. In this case, you'd simply draw designs on six square tiles, and then build the pattern to make the designs continuous across tile edges. I found this difficult to achieve with the Robinson tiles because I couldn't find a decoration that would force one and only one copy of the Fa type tile at each vertex. The other approach is to decorate the tiles that have edge modifications and assemble the pattern according to the edge rules. In this case the decoration does not force the rules, but simply comes along for the ride. I found this to be much easier for this set of tiles.

Figures 12 through 16 show three different decorations for the Robinson tiles. In all cases, the 7-by-7 block is exactly the one in Figure 11, except I used the indicated patterns on the tiles. It's surprising how much variation you can get out of the completed pattern just by changing the decoration a little bit.

Figure 17 shows an interesting variation on the basic

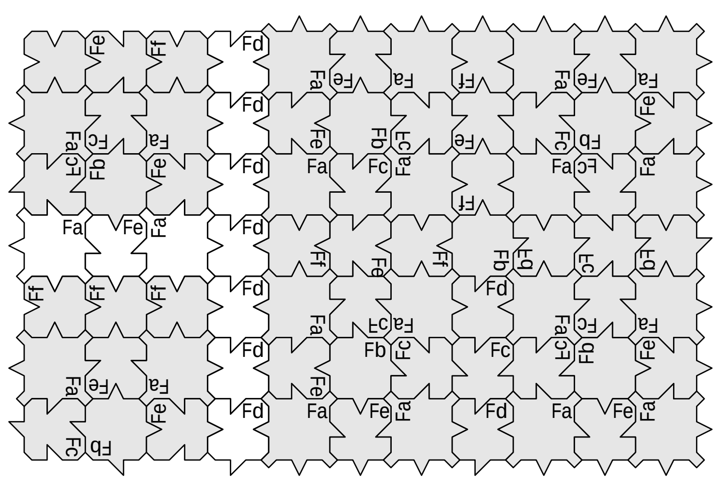

17 **Introducing a fault into a Robinson tiling. The fault lines are in white. The gray tiles are sections of the 7-by-7 block in Figure 11.**

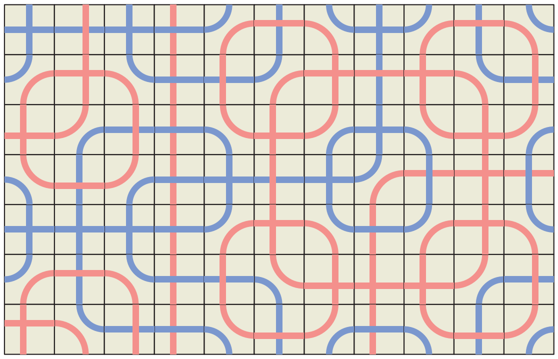

18 A decoration of Figure 17, using the tiles of Figure 13.

Robinson tiling pattern. Here I grayed out the pattern of Figure 11 and placed to its left two "fault lines." These are strings of identical tiles that can knit together two larger patterns. Notice that this also gives us the chance to break the regular pattern of the tiles with corners—each time we hop over one of these infinite lines, we can displace the corner tiles by one unit. Figure 18 shows a decorated example of this fault-ridden tiling using the decoration of Figure 13.

This fault-tiling technique might be useful when tiling a polygonal model. You can pave each polygon with Robinson tiles, and then use the fault lines on the edges to stitch individual polygons together.

Now that we know how to build patterns with these tiles, we can reasonably ask if they're actually aperiodic. Remember that the decoration can be misleading, since it doesn't necessarily force the conditions of tiling. A simple example of the problem is a decoration consisting of a single circle in the center of each tile. Even if the matching rules are followed exactly, the resulting pattern would be nothing but a perfectly regular grid of circles. So we need to look at the geometry of the tiles rather than their decorations.

The insight into the periodicity of the pattern rests on the construction of the basic 3-by-3 block (we'll ignore fault lines in this discussion). Rather that get bogged down in the nitty-gritty, I'll give the general idea here. Basically the trick is based on extending the construction we saw in Figures 10 and 11. Schematically, it's shown in Figure 19. I used my favorite orientable symbol, the letter F, to represent one of the 3-by-3 blocks. Figure 19 shows the construction of the 7-by-7 grid from four copies of the 3-by-3 grid plus the row and column of glue tiles to hold them together. Clearly there's no single rectangular unit that will generate this whole pattern by translation.

To see the big picture, consider it this way: The construction builds up blocks of ever-greater size, built from sub-blocks. None of the sub-blocks will work, since they get rotated by the step that glues them together. So however large a block you choose as your unit tile to stamp out the pattern, there's going to be a rotated copy, or some glue tiles, that get in the way. I hope I've suggested to your intuition that the Robinson tiles are aperiodic, but I certainly haven't proven anything. You can find a complete proof in Grünbaum and Shephard.

The Robinson pattern isn't periodic, but instead it's described by the technical term *almost periodic*. Again ignoring fault lines, you'll notice that each 3-by-3 block repeats with period 8, and the 7-by-7 blocks repeat with period 16. Basically each block that is 2^{n-1} on a side repeats with a period of 2^{n+1} tiles. So you can easily find huge repeating patterns, even though the whole pattern never repeats. Furthermore, this implies another property called *local isomorphism*, which states that any patch of tiles (without fault lines) will repeat infinitely often in the pattern. These traits seem to be shared by many aperiodic sets of tiles.

Ammann tiles

In 1977 Ammann developed a closely related set of tiles. Figure 20 shows the tiles themselves: Once again we have six tiles based on squares, with one tile marked as having special corners that fill in the gaps left by the other tiles. While the Robinson tiles had only two kinds of notches (one symmetrical and one asymmetrical), the Ammann tiles use three different matching conditions on the edges.

Following Grünbaum and Shepherd, I again marked these tiles with edge deformations that force them to link up correctly. Here we have a symmetrical point, an asymmetrical wedge, and a blunted asymmetrical wedge. Each tab may only go into its correspondingly shaped slot, though tiles may be rotated and flipped over as needed.

We can build patterns from these blocks using the same strategy as for the Robinson tiles. Figure 21 shows a 3-by-3 block of Ammann tiles, and Figure 22 shows a 7-by-7 block. These tiles can also be decorated, as shown in Figures 23 through 26.

Making tiles

There's nothing like actually playing with tiles to get a feeling of how they go together. While writing this article I spent a lot of time assembling sets of real tiles. I experimented with new kinds of decorations and really got a visceral feel for how they link up. I encourage you to make your own tiles and play with them. You can

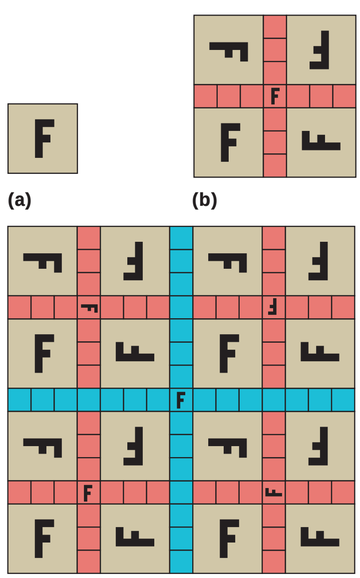

(a) (b) (c)

19 (a) A single 3-by-3 Robinson block. (b) A 7-by-7 block is formed by gluing together three 3-by-3 blocks. Notice that they rotate when assembled. (c) A 15-by-15 block made by rotating and gluing together 7-by-7 blocks.

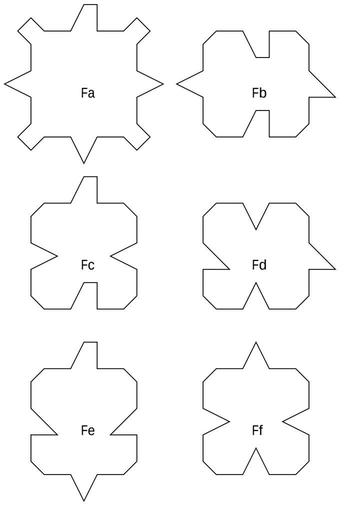

20 **The six Ammann tiles. Note the three edge modifications.**

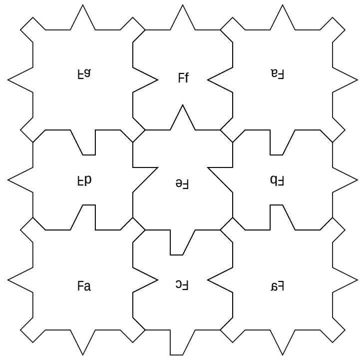

21 **Assembling a 3-by-3 Ammann block.**

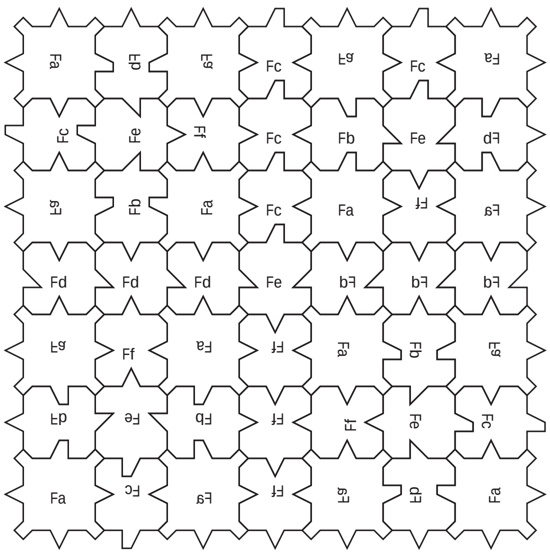

22 **Assembling a 7-by-7 Ammann block from 3-by-3 blocks.**

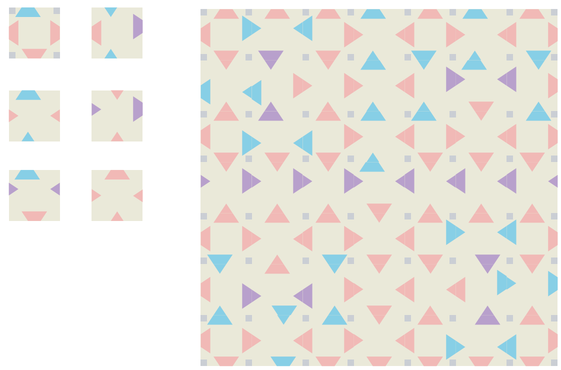

23 **A decoration of Figure 22.**

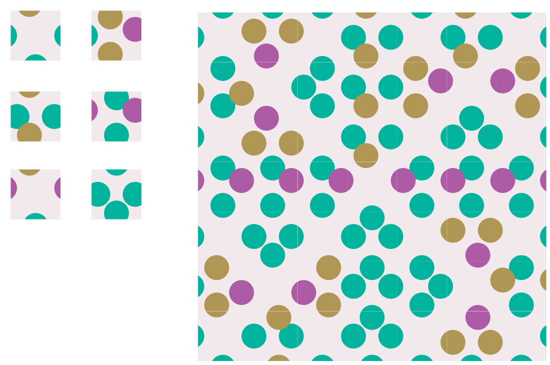

24 **Another decoration of Figure 22.**

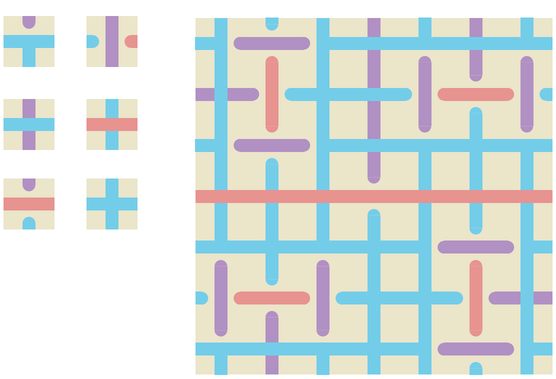

25 **A different decoration of Figure 22.**

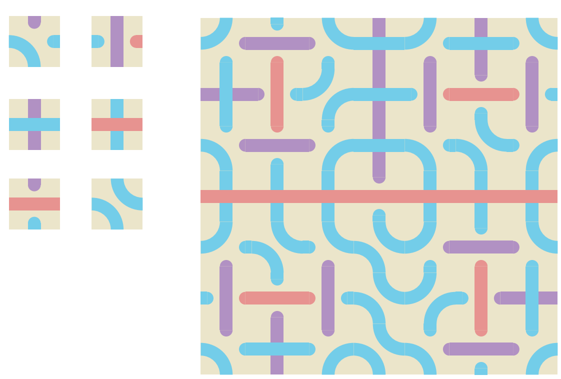

26 **A decoration of Figure 22, and a modification of Figure 25.**

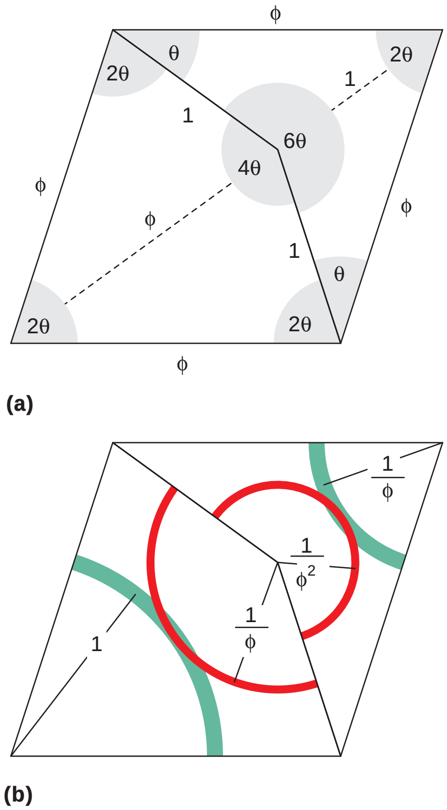

27 **(a) The geometry of the Penrose kites and darts. The Golden Ratio ϕ is equal to about 1.618. The angle θ is 36 degrees. (b) Face decoration for the tiles. Tiles can only be placed so that the arcs are continuous.**

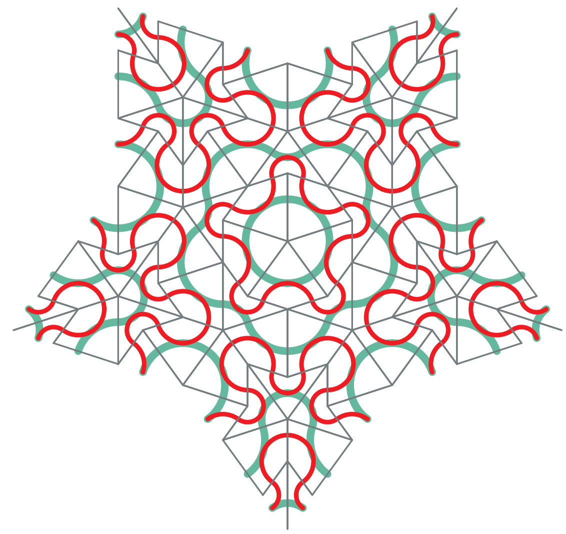

28 **An example of a pattern created with Penrose kites and darts.**

find PostScript templates for all the tiles I talk about in this article on my Web site.

I recommend that you print out the front-view page, glue it to a sheet of cardboard, and then cut out the pieces. Although gluing the tiles to cardboard may seem unnecessary, consider that tiny little pieces of paper are surprisingly aerodynamic—they have a tendency to curl, and they're much lighter than you might think. If you forego the cardboard step, you'll change your mind the first time you sneeze on your carefully assembled pattern of 85 tiles, or someone opens the door quickly and the wind creates a tiny snowstorm of paper where your tiling used to be.

Because these are two-sided tiles, you'll want to decorate the backs as well. I recommend cutting out the back-view pieces from paper and then gluing them to the pieces already cut out of cardboard—it's much easier than trying to align the two paper sheets on opposite sides of an opaque card before cutting.

Penrose tiles

I can't conclude this column without at least mentioning the Penrose "kites and darts," which are perhaps the most famous aperiodic tiles.

Figure 27 shows the basic geometry behind these tiles. The starting shape is a rhombus with a length of ϕ, the Golden Ratio, which is $(1+\sqrt{5})/2 \approx 1.618$. The rhombus is cut by two lines, each of length 1, which meet on the diagonal through the acute vertices. All the angles are a multiple of the same basic angle $\theta = 36$ degrees. The matching rules are given by two arcs on each tile, here drawn in red and green. The center of each arc is shown along with its radius. Note that the arcs overlap in the center. I have found it reads well if I make one thinner than the other, and place the thinner one on top. This decoration—and the names for the tiles—are due to John Conway. The kite is the larger tile, and the dart is the smaller, pointed one. An implicit matching rule is the obvious one that only sides of equal length may sit next to one another. Of course, the colored bands must also join up continuously. Notice that this rule means that you can't reassemble the original rhomb from which the tiles are built, since the bands don't line up.

There is a lot to say about the Penrose tiles, and they will occupy our attention in the next installment of this column. Until then, you may want to play around with them on your own. A few notes may help you save time and sanity. First, these tiles are symmetrical, so you only need to decorate one side. Second, you'll need more kites than darts—in fact, the ratio of kites to darts in a complete tiling will be the Golden Ratio. Third, growing a Penrose tiling is tough. It's surprisingly easy to create a configuration where you just can't add any more tiles and still obey the matching rules. That's okay. You just need to back up and try again. If you can assemble 100 tiles by hand, you probably qualify as a master of the form. We'll see some construction recipes in my next column. Figure 28 shows a piece of a Penrose tiling to help you get started.

Next time we'll look at Penrose tiles in much more detail. In the meantime, you might want to look at the many kinds of theme and variation surrounding us all the time. The leaves on a tree, the waves crashing into shore, and the songs of a bird are all similar to each other but different, balanced on the axis between repetitive and random. There is much beauty to be found there. ■

Readers may contact Glassner at Microsoft Research, e-mail glassner@microsoft.com.