# Andrew Glassner's Notebook

http://www.research.microsoft.com/research/graphics/glassner

## Penrose Tiling

**Andrew Glassner**

*Microsoft Research*

Theme and variation are part of nature. Birdwatchers identify the species of a bird by its distinctive markings, even though the specific colors and shapes vary from one bird to the next. Every thunderclap sounds roughly like thunder, but each one is different. And of course snowflakes are beloved for the hexagonal symmetry they all share, as well as the delicate patterns unique to each one.

In the May/June 1998 issue of *IEEE CG&A* I discussed the topic of *aperiodic tiling* for the plane. This technique helps us create patterns with lots of theme and variation, like the leaves on a tree. I'll continue discussing the subject, so let's first briefly summarize the main ideas from last time.

The basic approach is to take a bunch of 2D shapes and impose rules on how they can connect, like the pieces of a jigsaw puzzle. Suppose that you have an infinite supply of these shapes, or *tiles*, and you cover the plane with them, out to infinity in all directions. You might be able to find a region of the pattern that you could pick up and use as a rubber stamp, and by stamping it out an infinite number of times (without rotating or scaling it), fill the plane with the identical design that you started with. So in essence you've reduced the original set of pieces and their interlocking rules to just many translated copies of a single, larger piece. If you can do this, then the overall pattern—and the set of tiles used to make it— are *periodic*.

If you can't find such a single, big piece that replicates the pattern by translation, the pattern is called *nonperiodic.* Many tiles can create both periodic and nonperiodic patterns, depending on how they're laid down.

If you happen to have a set of tiles that can *only* make nonperiodic patterns, and you can prove that's the case, then you have an *aperiodic* set of tiles.

The quest for aperiodic tiles has gone on for several decades. The first set discovered involved 26,000 tiles. In the last issue I discussed two far smaller sets of aperiodic tiles, named after their creators. These Robinson and Ammann tiles are based on squares, which make them very convenient for computer graphics applications like texture mapping and sampling. Here I'll look at perhaps the most famous set of aperiodic tiles, discovered by Roger Penrose. Actually, Penrose found two distinct sets of aperiodic tiles, one called "kites and darts" and the other called "rhombs." Each set is made of only two tiles. Let's look at the kites and darts first.

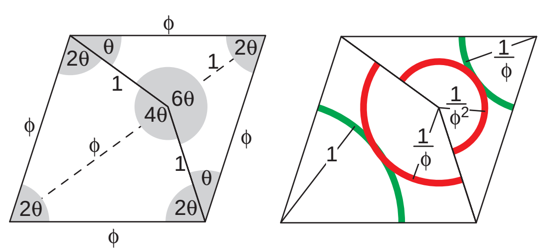


(a) (b)

**1 The geometry of the Penrose kite and dart (a), and a decoration that forces the matching rules (b) . $\phi=(1+\sqrt{5})/2$. $\theta$=36 degrees.**

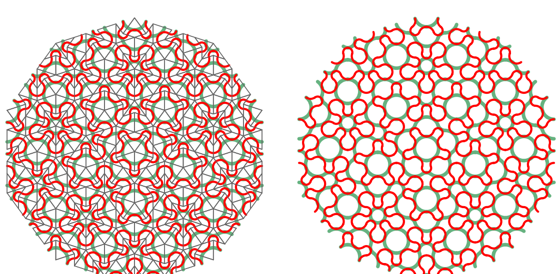

(a) (b)

**2 A pattern created by Penrose kites and darts. (a) With kites and darts outlined. (b) Just the decoration.**

### Aperiodicity

Figure 1 shows Penrose's kite and dart. The "kite" (the larger tile) and the "dart" (the smaller, pointed one) may be assembled only by placing them so that edges of similar length are adjacent and the colored bands are continuous. (The names kite and dart, and this decoration of the tiles, are due to John Conway.) Figure 2 shows an example of a pattern created by these tiles.

As you can see, these tiles are indeed aperiodic. A full proof would take more room than I have, but a general outline of the approach conveys most of the key ideas. I'll use a construction technique that's based on substitution rules. Like the rules of a formal grammar in computer science, or an L system used in botanical simulations, I'll take each tile and replace it, in position, with another set of tiles.

Let's simplify the problem for a moment and look just at 1D patterns.

Suppose that you had an infinite string, and you wanted to fill it up with an aperiodic pattern of white and black beads. I'll use two production rules to do the job. First, every white bead will be replaced with one white and two black beads in that order. I write this W→WBB, as shown in Figure 3a. Similarly, we'll replace each black bead with the rule B→BWW, as in Figure 3b.

Let's start with a single white bead and then apply these rules to start filling up the string. Figure 3c shows the first few steps. Each time you apply the rules, you take the original string of beads and apply all the rules at once. So the first generation is simply WBB. Now the new W goes to WBB, and each of the two Bs independently go to BWW, giving us the new string WBBBWWBWW. Let's run the process forever, so that we have (theoretically) filled up the infinite string with beads.

If the pattern generated this way is periodic, then we could find some chunk of beads—perhaps short, perhaps enormously long—that repeats. We could then take one of these chunks and glue it together into a super bead, and then match the original pattern by simply filling another string with an infinite number of copies of this single super bead. If we can't find such a super bead, the pattern is either nonperiodic or aperiodic. You might be able to convince yourself that in this example we would never be able to find such a super bead.

We follow a similar process in 2D for Penrose tiles. The technique is typically called *deflation* since the new tiles are smaller (or deflated) versions of their predecessors. Figure 4 shows the deflation rules for the kite and the dart. Notice that two of the newly created darts in the dart rule fall off the side of the tile. The matching rules come to the rescue here—you can prove that each half-tile is exactly completed by each of its legal neighbors when that neighbor tile is deflated. If you run the process back the other way and reduce the number of tiles in a tiling, it's called *inflation*, since each tile gets bigger.

Figure 5 shows a few deflation steps applied to a small starting pattern of Penrose tiles.

You may notice that some configurations of tiles seem to appear several times in these tilings. This is because you can only assemble tiles around a vertex in a limited number of ways. For example, take the vertex at the tip of the dart and enumerate all the ways that other kites and darts can be assembled around that vertex while obeying the matching rules. Then repeat the process for the other vertices on the dart, and then for the kite.

If you run through this process and eliminate duplicates, you'll find that only seven different kinds of clusters can be formed. These clusters taken together are called the *atlas* of the tiling. Figure 6 shows the atlas for the Penrose kites and darts.

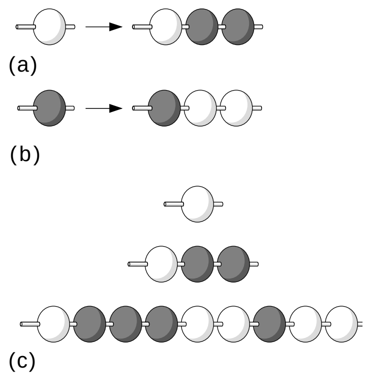


**3** **Creating an aperiodic bead sequence. (a) The rule W→WBB, (b) the rule B→BWW, and (c) three steps starting with a white bead.**

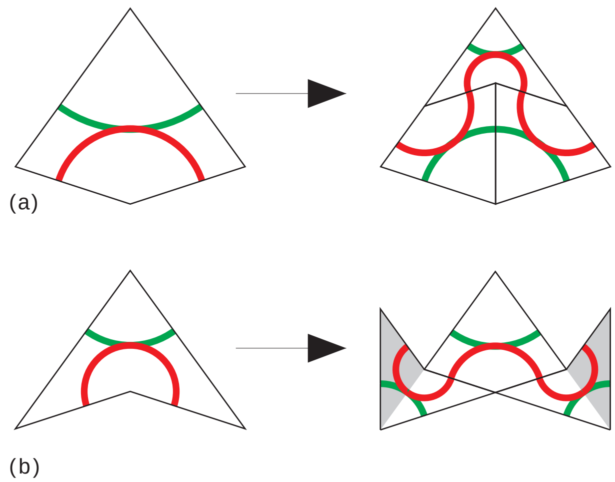


**4** **(a) Deflating a kite turns it into two kites and a dart. (b) Deflating a dart turns it into a kite and two half-darts.**

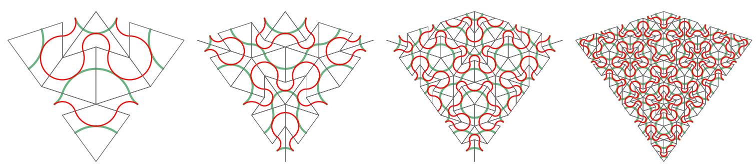

**5** **A Penrose "Ace" deflated four times.**

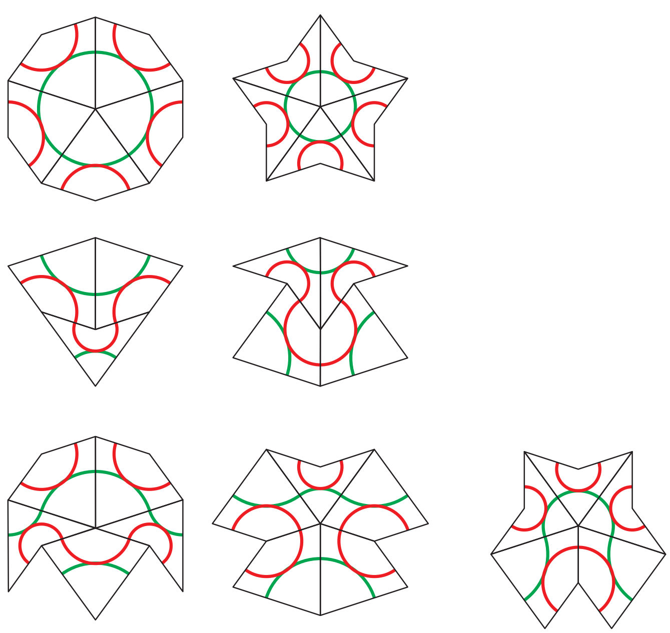

**6** **The atlas for kites and darts. Sun and Star (top); Ace and Deuce (middle); and Jack, Queen, and King (bottom).**

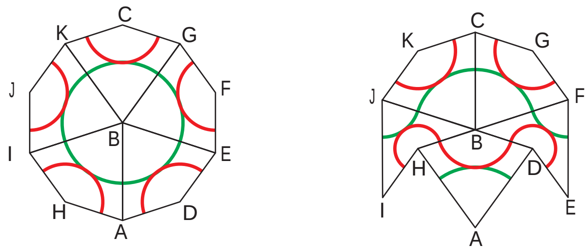


**7 Analytic expressions for vertices in the Sun and Jack. In these expressions, $k^2 = 1 + \phi^2 - 2\phi\cos(2\pi/5)$.**

Sun:

| | |
|---|---|
| A | (0,0) |
| B | $(0, \phi)$ |
| C | $(0, \phi + k)$ |
| D | $(\cos(\pi/10), \sin(\pi/10))$ |
| E | $\phi(\sin(2\pi/5), 1 - \cos(2\pi/5))$ |
| F | $k(\sin(2\pi/5), \phi + \cos(2\pi/5))$ |
| G | $(\cos(\pi/10), \phi + k - \sin(\pi/10))$ |
| H | $(-\cos(\pi/10), \sin(\pi/10))$ |
| I | $\phi(-\sin(2\pi/5), 1 - \cos(2\pi/5))$ |
| J | $k(-\sin(2\pi/5), \phi + \cos(2\pi/5))$ |
| K | $(-\cos(\pi/10), \phi + k - \sin(\pi/10))$ |

Jack:

| | |
|---|---|
| A | (0,0) |
| B | $(0, k)$ |
| C | $(0, \phi + k)$ |
| D | $\phi(\cos(3\pi/10), \sin(3\pi/10))$ |
| E | $k(\cos(\pi/10), \sin(\pi/10))$ |
| F | $(\phi\cos(\pi/10), k + \phi\sin(\pi/10))$ |
| G | $(\cos(\pi/10), \phi + k - \sin(\pi/10))$ |
| H | $\phi(-\cos(3\pi/10), \sin(3\pi/10))$ |
| I | $k(-\cos(\pi/10), \sin(\pi/10))$ |
| J | $(-\phi\cos(\pi/10), k + \phi\sin(\pi/10))$ |
| K | $(-\cos(\pi/10), \phi + k - \sin(\pi/10))$ |

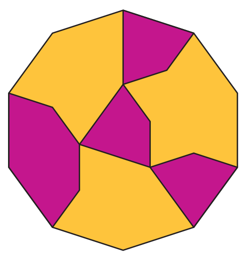

**8 The Gummelt decagon.**

Each of these clusters repeats an infinite number of times throughout any Penrose tiling and can serve as a good starting point for a tiling process.

You can derive analytic expressions for each point in the atlas. For example, Figure 7 gives the locations of the points in the Sun and Jack. You may find it fun to derive the vertex locations for the other patterns—I give them all on my Web site.

## Growing a tiling

Now that you've seen something about Penrose tilings, how do you generate them?

It's very difficult to "grow" a covering of the plane using kites and darts. Suppose that you start with just a single tile, and place new ones that are always in accordance with the matching rules. After a while, you'll probably find yourself in a state where you can't add any new tiles. What went wrong in this process?

Nothing actually went wrong, it just didn't go right enough. As an analogy, suppose you were writing a limerick, and you got this far:

There once was a penguin named Bryce
Who lived in a house made of ice.
He painted it orange,

Now what? There's nothing technically wrong with this limerick (the third line is kind of clunky, but let's overlook that). The major problem here is that you just can't rhyme with orange—you're stuck and it's nobody's fault. You just have to back up a line or two and try again.

It's the same thing with building Penrose tilings. You can follow all the rules and find yourself unable to add any more tiles. You need to take some tiles away and try again, in a trial-and-error process. Obviously if we want to create and use Penrose tilings regularly we need a more reliable approach for generating the patterns.

We saw one approach above: start with a single tile and deflate it over and over, producing an ever denser (or ever larger) tiling. Any finite region can be covered with a finite number of deflations, and the infinite plane can be covered if we deflate forever.

This is fine mathematically. A problem occurs though when we try to use this technique to explain the physical phenomenon of *quasicrystals*.

## Quasicrystals

In 1984, Dan Shechtman and his colleagues at the National Institute of Standards melted together samples of aluminum and manganese. They then quenched, or quickly cooled, this molten metal by squirting it at a rapidly spinning wheel, causing the temperature of the metal to drop by about a million degrees kelvin per second. When they examined the structure of the resulting material using electron diffraction, the pattern looked like one typical of crystalline structures, except that it had a five-fold rotational symmetry.

This was very, very strange. All the laws of standard crystallography disallow five-fold symmetry. Just try to put five regular pentagons together at a single point—you can't do it unless you let them flop over on one another. This alloy seemed to be a crystal by most measures, but it had this bizarre, and disallowed, internal symmetry. The discoverers called this perplexing structure a *quasicrystal*. Since then, many other compounds have been discovered that fit into the quasicrystal category, including materials with five-, eight-, ten-, and twelve-fold symmetries.

As described by Grünbaum and Shephard (see the "Further Reading" sidebar), in 1984 Levine and Steinhardt wrote one of the first papers that tried to explain the internal structure of quasicrystals. They showed that a 3D generalization of the Penrose tiles matched the diffraction results of the real material, and thus seemed to be a plausible model for the quasicrystal's internal structure.

This explanation worked mathematically, but in practice it posed three problems. First, the traditional way to think of crystal formation is that it begins with lots of copies of a single atomic configuration, but the Penrose model requires two building blocks. Second, the Penrose matching rules are straightforward on paper, but it's hard to imagine how simple atomic structures would cooperate to pull them off. Third, as we've seen, it's hard to grow a Penrose tiling without getting into a situation where you can't grow any more; this would make it very difficult for large crystals to form. So the Penrose techniques we've seen so far don't look like a good explanation for quasicrystals.

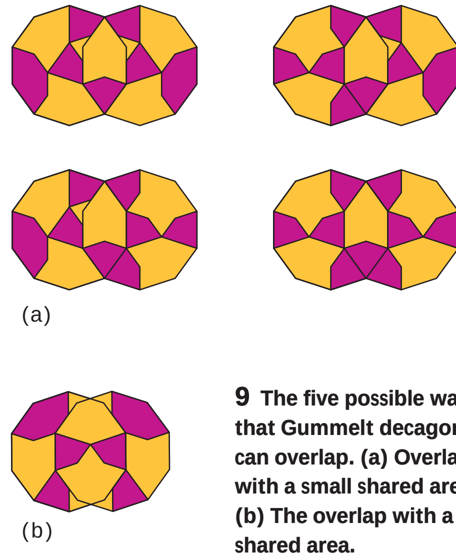


**9** **The five possible ways that Gummelt decagons can overlap. (a) Overlaps with a small shared area. (b) The overlap with a large shared area.**

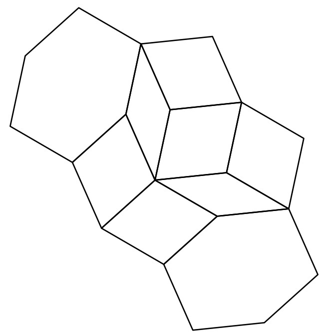

**10** **The Jeong and Steinhardt cluster.**

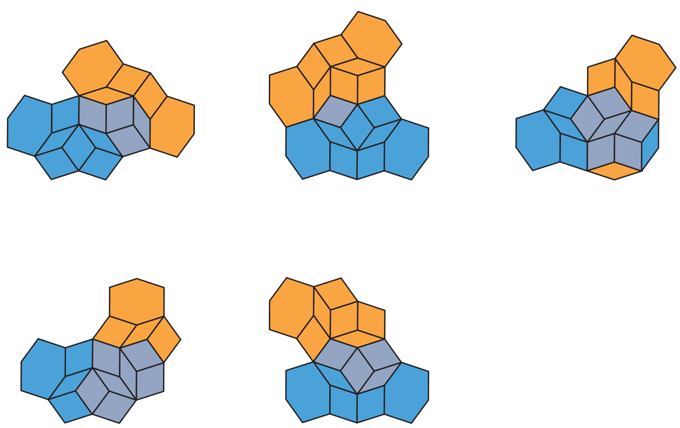

**11** **Five ways that the cluster of Figure 10 can overlap with itself.**

A number of alternative theories have been put forth to explain quasicrystal structures, but Penrose tiles have always been in the running, despite their problems. The challenge to those who believed in the Penrose explanation was to simplify the matching rules and deal with the puzzle of two basic units, rather than one. And closer to the topics in this column, such rules would also help us build new patterns of arbitrary size without getting stuck.

A breakthrough in this quest came in 1996 in a paper by Petra Gummelt (see the "Further Reading" sidebar). She showed that a particular decagon, shown in Figure 8, could be used to grow a Penrose tiling by overlapping. The overlap rules were only that overlap actually occurred (that is, each decagon must overlap with at least one other) and that colors had to match in the overlapped region. Figure 9 shows the five decagon overlap rules. You can break down any Penrose tiling into a collection of these overlapping decagons.

This work showed that a single structure, albeit a complex one, could be used to create the aperiodic Penrose tilings in the plane. This addressed the problem of requiring multiple building blocks. But the matching rules were still a problem. How do you imagine random atomic clusters, floating around in a soup, obeying complicated overlap assembly rules?

On a technical note, observe that the Gummelt decagons do not strictly tile the plane because they overlap; a true tiling creates no overlaps and leaves no holes. Gummelt suggested that this be called a *covering* rather than a tiling.

In their 1997 paper Jeong and Steinhardt discussed Gummelt's decagon and simplified her original proof that it generates a valid Penrose tiling (see the "Further Reading" sidebar). They then present a second technique, also based on overlapping clusters of Penrose rhombs. Figure 10 shows one of these clusters. Jeong and Steinhardt proved through a complicated argument that if many copies of this cluster overlap (as in Figure 11) so that the complete tiling has the highest possible density, the result is a Penrose tiling. They point out that the maximal possible overlap corresponds to the minimum possible energy.

This approach addresses the two chief difficulties with using the Penrose model to explain 2D quasicrystal formation and structure: we're now down to a single structure, and we have eliminated the matching rules in favor of a more physically plausible juggling about for minimal energy. Whether this cluster argument can be extended to 3D—and whether or not it accurately predicts new quasicrystals yet to be found—only time will tell.

Materials scientists and engineers are excited by quasicrystals and developments such as these that seek to explain them. Such understanding can lead to new

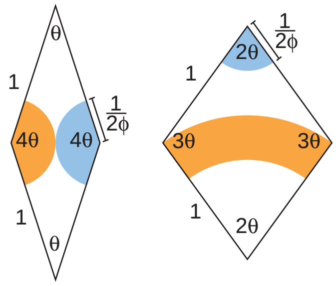


**12** **The geometry of the Penrose rhombs.**

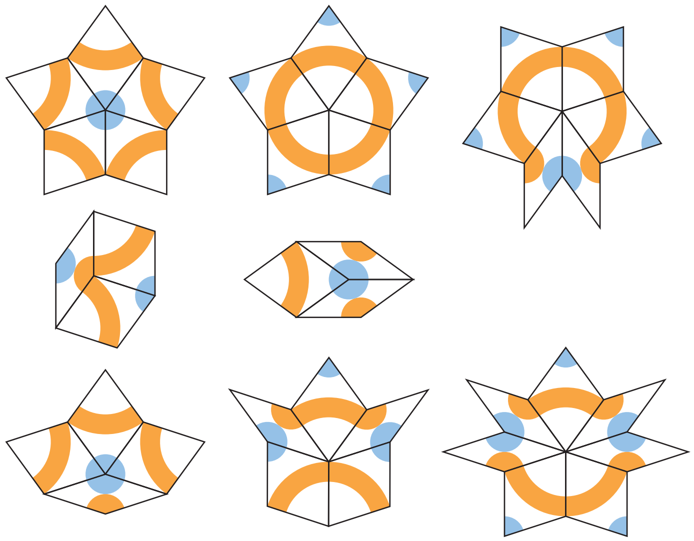

**13** **The atlas of Penrose rhombs. Dotstar, Ringostar, and Medal (top); Fat Cube and Thin Cube (middle); and Bud, Leaves, and Flower (bottom).**

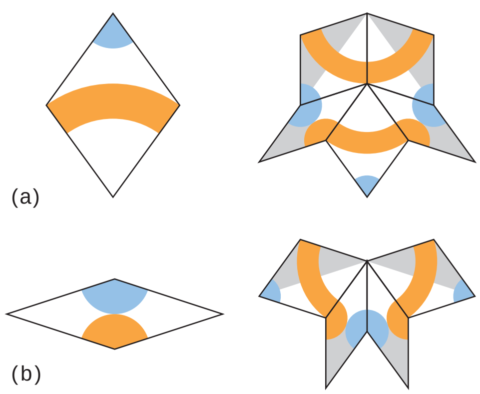


**14** **How to deflate a thin and thick rhomb.**

materials, which can exhibit new properties that have commercial and technical value.

## Penrose rhombs

In addition to the kites and darts, Penrose also came up with another pair of aperiodic tiles, shown in Figure 12. These are usually called the *Penrose rhombs*, and the two tiles are simply named *thin* and *thick*. The figure also provides the matching rules—as before, the color bands must be continuous. The edge-length conditions are easier here, since all sides have the same length. Figure 13 shows the atlas for the rhombs. There's no set of shorthand names for these that I know of, so I've suggested a set of names for these shapes in the caption.

You can deflate rhombs directly, just like kites and darts. Figure 14 gives the recipe. It's interesting to note that you can convert a kites-and-darts tiling into a rhomb tiling, and vice-versa. Figure 15 shows the basic idea. Figure 16 shows the rhomb equivalent of Figures 2 and 5. There's clearly a lot of common ground between the two kinds of tiles and their decorations. Figure 16c shows what they look like when both versions are drawn together.

## Implementation

How would you write a program to create Penrose tilings? In this section I'll concentrate on kites and darts. You can use Figure 15 to adapt these ideas to the rhombs.

The simplest method is to apply deflation rules to a starting tile or set of tiles. The easiest way to deflate a pattern is to first turn it into triangles. Each kite and dart can be split into two congruent triangles along the dashed line in Figure 1. Each of the dart triangles contains an obtuse angle (of $4\theta = 144$ degrees), while the kite triangles don't, so they're typically referred to as *obtuse* and *acute* triangles, respectively.

The deflation process is now simply a matter of running through the list of triangles and replacing each one with the appropriate set of new triangles as shown in Figure 17. The new triangles can be reassembled into kites and darts if desired, or they can immediately go through the substitution process again.

If you don't draw the black lines around each kite and dart, it doesn't matter how you store the triangles. But if you do want to draw the composite tiles, make sure that when you draw each triangle, you don't draw a line down the edge that's internal to the kite or dart to which the triangle belongs. You can see the effect of these suppressed lines around the perimeter of Figure 5. You can use some conventional scheme to identify that edge, such as always making it the first one in the list.

The data structure for this process can be very simple. Just maintain a list of triangles, each containing its three points and a flag indicating whether it's obtuse or acute. Then run through the list, building up a new list. When you're done, go through the list again, drawing the appropriate decoration on each triangle.

To convert to rhombs, use Figure 14 to build a new list of thin and thick rhomb triangles from the obtuse and acute kite and dart triangles.

## Uses

Now that you know how to make Penrose tilings, what good are they to us in graphics? The most obvious answer is to use them to create endless aperiodic texture on surfaces. Simply draw some interesting face decoration that obeys the matching rules, pick an initial tile and orientation, and then deflate until you cover the surface with tiles. You can use this to create a nice big texture on a flat surface, such as the side of a building or a floor.

If you're careful, you can carry this construction across polygons so you don't get a seam where two polygons adjoin. But how could you apply this technique to a sphere—or simpler yet, a cube? Of course, formally you can't do this at all. If you followed the pattern from the front face, say, around to the right face, then in back, then to the left, and then back again to the front face, you've completed a loop: the pattern is periodic in the distance covered by the lengths of the four faces. And since Penrose tilings are by definition aperiodic, you can't use them to cover a cube, a sphere, or even trickier surfaces like donuts or Möbius strips.

There may be modifications of the tilings that will match up around the surfaces of these objects. I haven't been able to find any such instances, but that doesn't mean they don't exist. Other uses of these tilings include generating nonuniform sampling patterns for stochastic sampling, such as that used in distribution ray tracing.

Another use of Penrose tilings is to create image maps that control models. For example, suppose you want to lay out a city that grew up rather than being planned out in advance. You could decorate your tiles with building bases, and create a tiling of the plane. Then use the result as a blueprint upon which you erect office towers, homes, apartment buildings, and so on. They will have a visible large-order structure, but they won't be on a regular, boring grid. You could use such blueprints to specify geometry, materials, the density of flowing lava, or whatever else you'd like to create with some large-scale but nonrepeating structure.

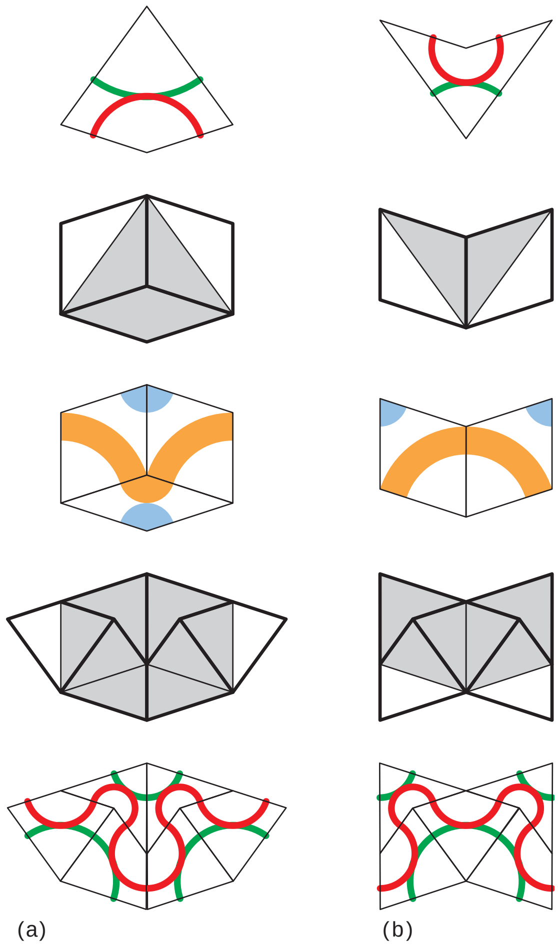


**15 How to convert rhombs into kites and darts, and vice versa.**

## Decoration

Inventing new face decorations for the Penrose tiles is fun, but each of the tiles must be symmetrical across the edges marked with a dashed line in Figure 1. I haven't seen this restriction discussed anywhere, but it's not too hard to find the reason. Consider the Sun pattern in Figure 6. Let's assume that the long edges are different; we'll call them types 1 and 2. We'll see that assuming we have these two distinct edges leads to a contradiction.

The proof is easy. Pick any radial spoke of the Sun and label it 1. Then label its clockwise neighbor 2, then the next clockwise neighbor 1 again, and so on, around the Sun. You'll find that you have the sequence 12121, which means that the last 1 is up against the first 1, which violates our assumption that the kite had two distinct edges. You can prove the same thing for the long sides of the dart using the Star pattern.

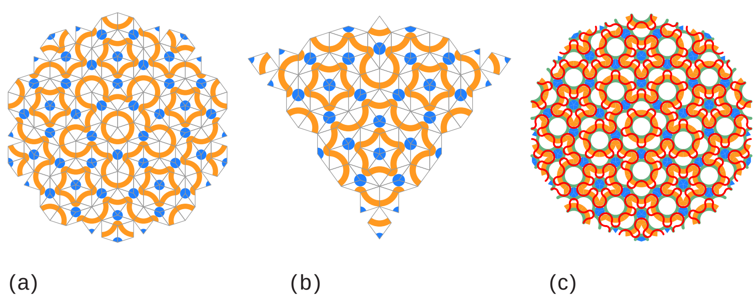


**16 (a) The rhomb version of Figure 2 and (b) Figure 5, and (c) two decorations for the deflated Sun pattern overlaid.**

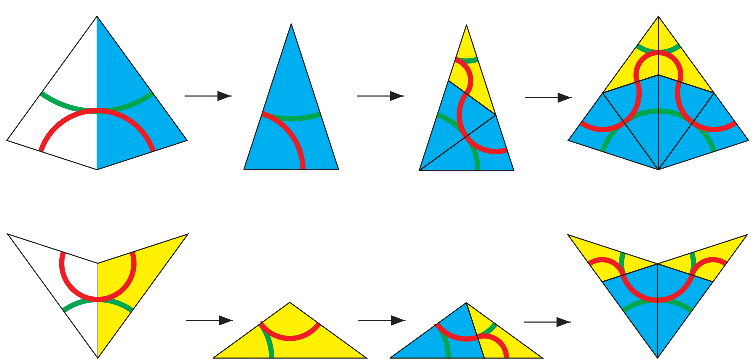

**17 Using triangles to deflate the kites and darts. I color-coded the triangles for clarity: blue for acute triangles and yellow for obtuse triangles.**

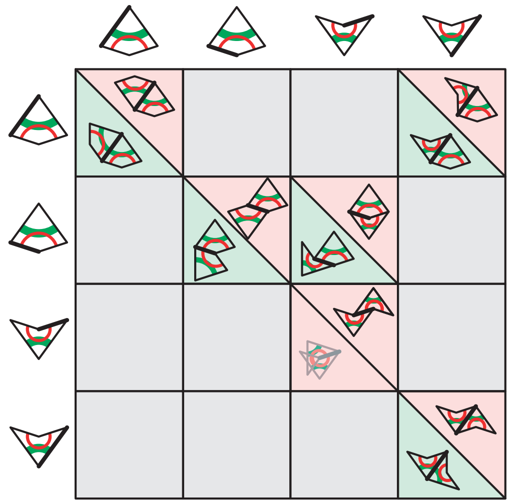

**18** **Dart and kite configurations. Illegal pairs are surrounded with red, legal pairs in green. The bold segment at the edge of each row and column shows the segment being matched. The lower left corner of the grid is symmetrical to the upper right, so I left it empty. The four gray squares in the upper right indicate illegal pairs because the edges are not the same lengths.**

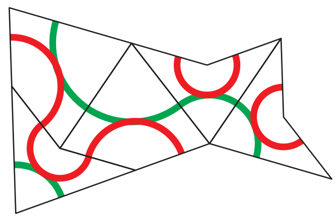

**19** **A single cluster that tests all of the legal kite-dart edge combinations.**

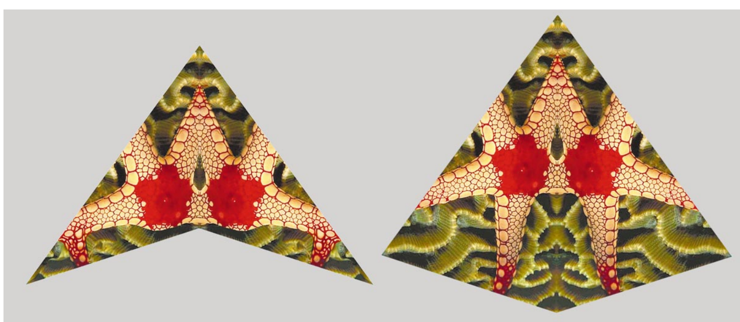

**20** **A symmetrical kite and dart decoration.**

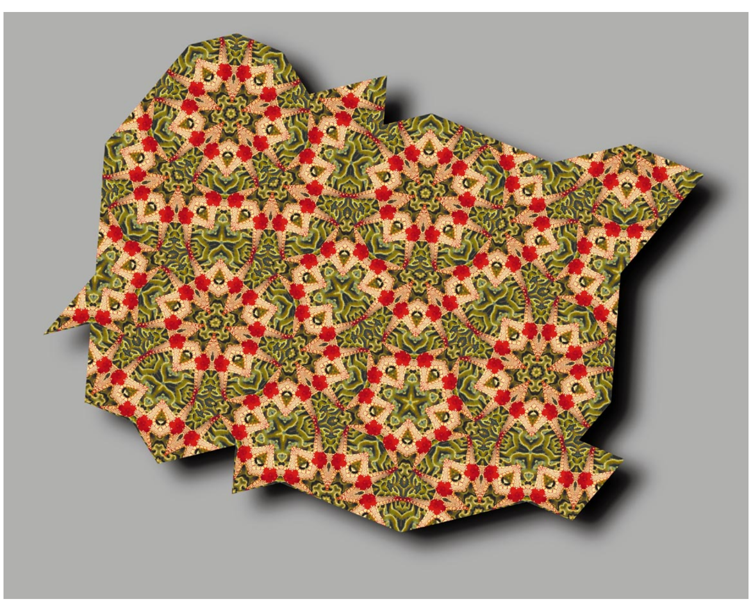

**21** **A Penrose tiling made from the tiles of Figure 20.**

The restriction on the short sides follows the same process. Surround a Sun with darts and assume that the short sides are different. By simply alternating the two side labels you can see that we again reach a contradiction.

Of course, there's no reason not to create tile patterns that are asymmetric on the inside, as long as the edges match up. Or for fun, you could have a variety of different internal designs (as long as the edges are all the same) and mix them up. You can choose the decoration of each tile randomly, or according to a procedural scheme. If you're going for a complicated texture, this is an easy way to make the resulting pattern more complicated. Remember that introducing complexity in the picture moves you away from pattern and towards randomness, so you'll want to approach this process with caution.

To create allowable and attractive decorations, I found it useful to make a chart of the legal and illegal configurations. Figure 18 shows all of the pairings of kite and dart edges, split into legal and illegal pairs. I still like to cook up designs using paper and colored pencils. I found it useful to make a little cluster, as in Figure 19, that shows all the legal configurations in one unit so that I can get a general idea of how the decorations look.

Figure 20 shows a pair of tiles that I made from a photo of a starfish—note that these are symmetrical across their central axis. In Figure 21 I lovingly hand-assembled these into a Penrose tiling. To show the value of asymmetrical designs, Figure 22 shows the tiles after I redrew them to make them asymmetrical, and Figure 23 shows the resulting pattern.

Of course, if you want a really random look, you can generate some kind of random pattern on the tiles using fractal noise, diffusion-reaction patterns, or any other pattern generator you like, and then enforce the boundary conditions programmatically.

## Wang tiles

Although it's a bit off the subject of aperiodicity, there's another important tiling topic that I can't resist mentioning briefly. The subject is the close connection between tiling and computing. Recall that a *Turing machine* is a conceptual machine made up of an infinite piece of tape, a read-write head, a state marker, and a collection of transition rules. When running, the machine looks at the current value on the tape (typically either a 0 or 1), and then looks up its current state. The machine then rummages through its rules until it finds one that matches these inputs. This rule tells the machine what number to write back onto the tape, what state to consider itself to be in next, and whether the read-write head should be moved one step left, right, or not at all.

Although you'd never want to write a big program on a Turing machine, it's not too hard to write little programs. The remarkable thing is that Turing machines are powerful enough, in theory, to execute any computable algorithm. They might be slow and clunky, but anything that can be done by the biggest computer imaginable can also be done by the humble Turing machine.

In 1975, Wang showed that Turing machines could be simulated by tilings (see Grünbaum and Shephard for details). Rather than go through the theory, I'll show an example of the idea. I've chosen a simple operation inspired by something that we do all the time in computer graphics: *z*-buffering. An essential *z*-buffer step is finding the minimum of two numbers. To make things simpler, I'll pose the problem this way: find the smaller of two positive integers, each larger than 2. Figure 24 shows a set of tiles to accomplish this goal and the computation. For convenience, I also numbered the columns of the tiling. The origin of the number line is fixed at 0 by the tile marked α. Two other tiles, marked β, mark the two numbers we want to compare; in this example they're placed at 5 and 11. Their minimum appears to the left of the origin at a tile marked δ, located at –5. This tiling puts δ at –min($\beta_1$,$\beta_2$). How does this work?

As always, tiles can only abut if their edges match. Here I marked the edges with numbers, so the numbers have to be the same across each edge. Edges that are unmarked are implicitly numbered -1, so most of the infinite plane is filled with tiles that are -1 on all four edges. For convenience, the plane is only half-infinite; the upper boundary is marked with edges 0. The tiles in gray are assumed to match along the marked edges and have labels of –1 on the others. I've shown the indices of the number line above the grid to make it easier to see what's happening. α is located at 0, and our two input βs are at 5 and 11.

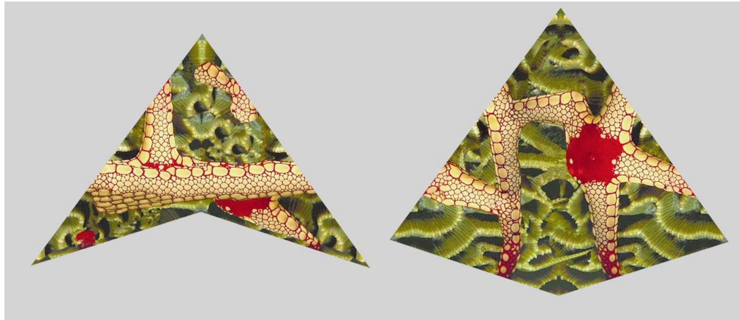

**22** **An asymmetrical modification of the tiles of Figure 18.**

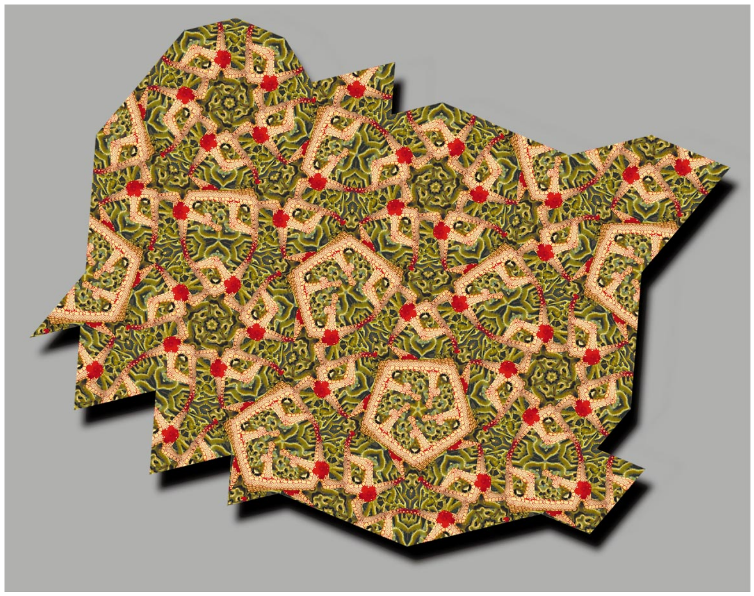

**23** **A Penrose tiling made from the tiles of Figure 20.**

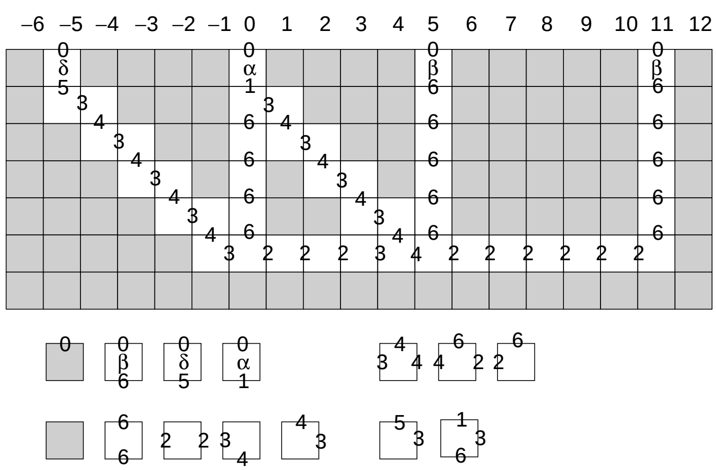


**24** **A Wang tiling that finds the minimum of two numbers. All unmarked edges are assumed to have the value –1. The tile marked α is the origin. The tile marked δ is to the left of α by the same distance as the closer of the two tiles marked β: δ = –min ($\beta_1$,$\beta_2$).**

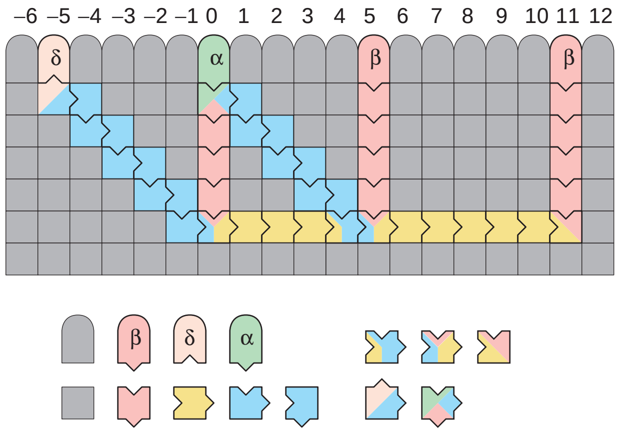


**25** **A color-and-shape version of Figure 24.**

Figure 25 shows another representation of the identical tiling. I'll refer to this figure from here on, since I believe the color coding and interlocking shapes make the discussion clearer. The little bumps at the top of the tiling capture the top row; imagine that they fit up against an infinite row of half-circle concavities. Rather than add gray to every colored tile and make things more complicated, I adopted the convention that any straight edge is implicitly colored gray.

Let's start with the β tiles. Each one connects to the upper border with a half-circle and generates a downward-heading column of red tiles. These columns will continue downward forever unless something interferes. Of course, I've set things up to create interference!

Now look at the α tile, marked in green. It sends down a tile that forks off into a downward-heading red column and a blue diagonal that heads down and to the right. The blue diagonal continues until it reaches the nearer of the red columns created by the β tiles. A couple of multicolored "glue" tiles mediate the collision. At this point of intersection, two new rows are created, marked in yellow. One row moves to the right, where it eventually reaches the outermost red column and stops it (using a multicolored glue tile). The other row moves to the left, until it encounters the red column sent down from the origin. That collision, again handled by a multicolored tile, signals a new blue diagonal up and to the left. Eventually this gets to within one square of the top border, where the pink tiles take over. The δ tile then connects the diagonal to the top of the grid.

Because the two blue diagonals are the same length, and both start directly underneath the origin, the distance from δ to α is the same as the distance from α to the nearer β, and we're finished. Note that although I spoke of columns "moving" in one direction or another, that was just to help analyze how the pattern came together. This computation is represented by a single, static tiling.

Simply by creating the right set of 14 tiles, we've computed the minimum of two positive integers. Wang showed that this process could be carried out in general, so that any Turing-machine program could be converted into a set of tiles and connection rules. Now I'm not suggesting that you run out and sell your workstation for some tiles and a big checkerboard, but it's worth knowing that theoretically any program and its inputs, computation, and results can be represented by one of these tiling patterns—though admittedly anything beyond a toy problem would be enormous.

You might see a similarity in spirit between this form of computation and cellular automata. Notice though that these tiles are static—the act of simply creating a stable and consistent tiling pattern is the computation, unlike automata, which themselves are little processors that run local programs.

I really like the idea of using a computational Wang tiling to decorate a kitchen or bathroom floor. It would be quite an inside joke, of course, but it would provide an occasional reminder of the fact that the humble geometrical mosaic can handle any computational task, from predicting weather and understanding speech to rendering 3D images.

## An aperiodic tile

I'll end this issue's column with an open question: Is there a single aperiodic tile? That is, is there a single shape that can cover the plane with no gaps or overlaps, and creates a pattern that cannot be also created by translation of a subset of the pattern? Nobody knows. The trend from 26,000 to 2 seems to suggest that the step to 1 might be possible.

I've fantasized that this would be a terrific one-page doctoral thesis. Simply titled "An Aperiodic Tile," it would contain a one-line abstract, a picture of the tile, a picture of the inflation rule(s), and a single reference to Grünbaum and Shephard. I hope somebody writes it! ■

## Further Reading

The fundamental reference on tiling of all sorts is *Tiling and Patterns* by Grünbaum and Shephard (W.H. Freeman, New York, 1987). A very good book on the mathematics of quasicrystals and aperiodic tilings is *Quasicrystals and Geometry* by Marjorie Senechal (Cambridge University Press, Cambridge, England, 1995). A popular account of Penrose tiles appears in Martin Gardner's book *Penrose Tilings to Trapdoor Ciphers* (W.H. Freeman, New York, 1989). Gardner discusses many fascinating properties of these patterns that I haven't had the space to cover here. In the text, I referred to "Penrose Tilings as Coverings of Congruent Decagons" by P. Gummelt (*Geometriae Dedicata*, Vol. 62, 1996, pp.1-17) and "Constructing Penrose-like Tilings from a Single Proto-tile and the Implications for Quasicrystals" by Hyeong-Chai Jeong and Paul J. Steinhardt, (*Physical Review B*, Vol. 55, No. 6, 1997, pp. 3520-3532). There are some great online references. A good place to start is the Quasitiler page at the Geometry Center, located at http://www.geom.umn.edu/apps/quasitiler/. A quasicrystal bibliography is available at http://gene.wins.uva.nl/~kerres/quasicrystals.html.