In order to comment on a picture and ask a question about it, the original reason for writing this, I need to indicate how it was constructed. To do that I want to call upon some visual imagination, as I believe the picture calls upon visual imagination. I write this wondering how much it calls upon it. The picture is the result of my wondering how to display effectively a colouring of the surface of a cube, a colouring by weaving the surface, but how the surface is coloured isn't the main difficulty. The difficulty is that one can see at most three surfaces of a cube at one time, and so one is limited in two dimensions in seeing how the colourings of the different surfaces are related. Those relations are specifically what I wanted to display. My first attempt at doing this was to have photographed both the cube itself and a fake that would be the image of the cube in a mirror set up behind it. That allows all six surfaces to be seen and, with some effort, to imagine how the surfaces behind are related to the surfaces in front. One can also see how the strips weaving the surfaces pass across edges of the faces whether they are in the same or different pictures because the hexagonal boundaries match up point to point.
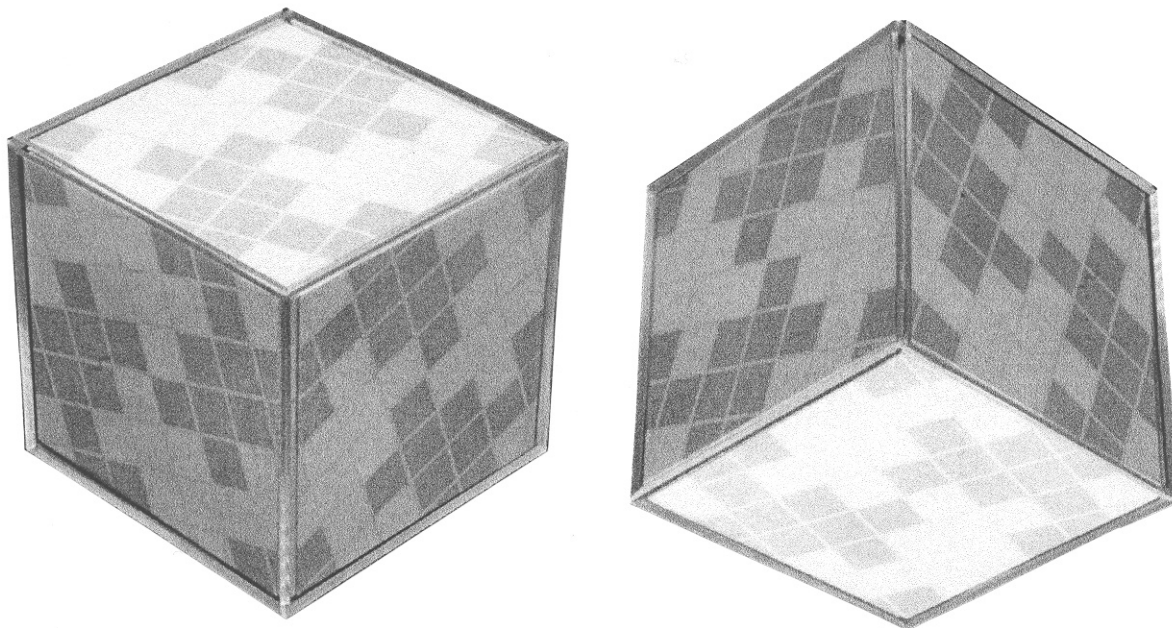


After I had these pictures I realized that one could use many copies of the cube and if one did not use perspective, so that the cube appeared in two dimensions as a regular hexagon, make a big picture that would display the whole cube many times over, perhaps being interesting in its own way. I want to describe how the cubes are arranged, and to do that I'll use a die, since everyone knows what dice look like. For the description, I'm going to write of black and white dice, but the end result will just be a picture of a lot of white dice. On a table, let a row of white dice be arranged three uppermost in a straight line edge-to-edge so that viewed from one side one die has one and two showing and the dice next to it to both right and left have the opposite faces, five and six, showing. Next to the pair with five and six showing in both directions a pair with one and two, and so on alternating indefinitely. Behind the row of white dice, fill the half-plane with black dice. It may from here on be worth looking at the sketch of this arrangement. On top of the front row of black dice, arrange a row of white dice four uppermost in the same way as the first row except that the dice are now the other way up. (This is the higher of the two widest rows of dice in the diagram.)
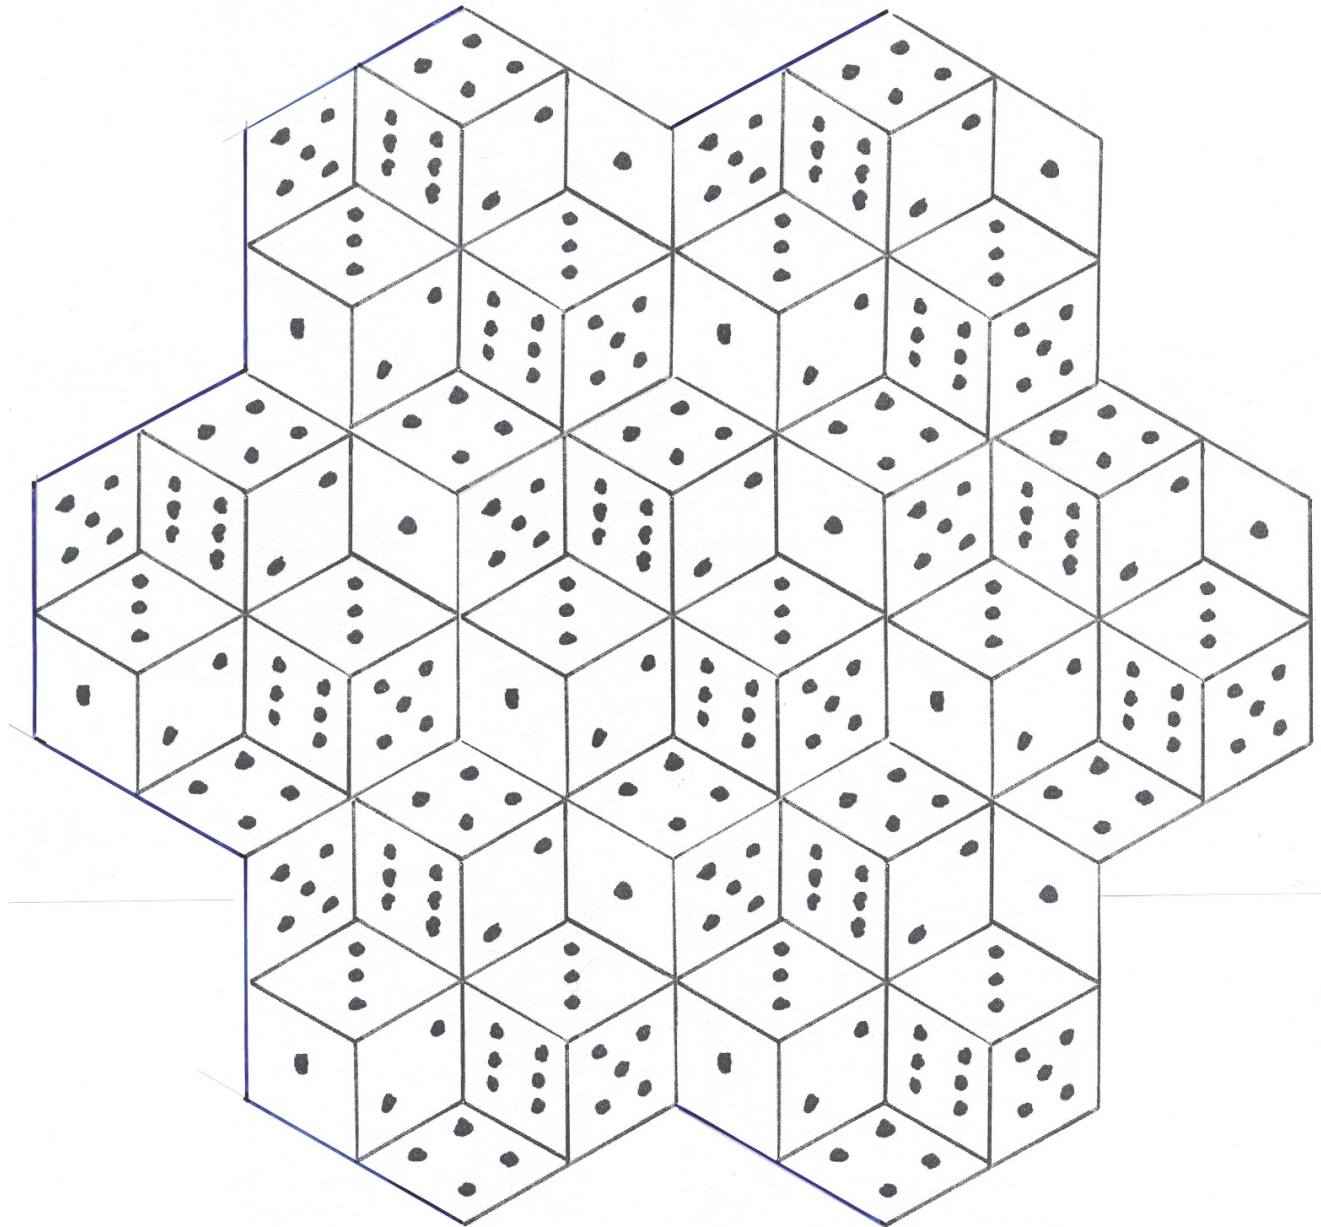
Above the two and six faces on different dice below, put the six and two of one upper die. Showing on the pair of dice surrounding it will be their one and five, and so on alternating as in the row below. Fill in behind with black dice. Then make the next layer of white dice on the front black row be the same as the bottom row but displaced. That is, above and behind the die showing one and two put one showing six and five, and of course above those showing six and five, ones showing one and two. And so on upward. Now that the pattern is established, eliminate the table and make the ramp of dice be infinite downward as well as upward to form an infinite puckered surface that I have tried to illustrate with the diagram. This diagram has the feature that, whenever your attention crosses a division between the faces of two different dice, what you see is what you would see if you looked instead beyond the same boundary on the same die, just as you do when you cross a boundary on the same die. When the boundary is convex, you stay on the same die; when the boundary is concave, you are changing dice, but what you see is as though you had not done so. Just as the horizontal row of faces wraps around and around the die, so do the rows at sixty degrees to the horizontal. One gets a highly redundant picture of all of the faces of a die.

A couple of comments on the diagram. If one looks at it at sixty degrees to the horizontal in either direction, one sees a similar but different stacking of the dice. It takes a moment to adjust what one sees from one view to another! Another visualization that is possible in any of the three orientations is to perceive from the under side a ramp of dice looming over one. These are not symmetries. The diagram has some symmetry despite the fact that a die, unlike an unmarked cube, has no symmetry. If one rotates the diagram about its centre through one hundred and eighty degrees, then it is unchanged. As to its boundary, this is a feature of how the diagram was drawn, but as to the implicitly infinite arrangement of the cubes it is a common feature of the arrangement rotated about any of the vertices where three dice meet (that is, half of the vertices). The other vertices are where three faces of a single die meet. So the arrangement of dice has half-turn symmetries. It also has translational symmetry, since the centres of half-turns can be translated one to another provided they are surrounded by the same die faces. The centres fall into two classes, those with threes to right and left and those with fours to left and right. This is the wallpaper group called *p2*. I have put a parallelogram with translation vectors as sides and centres of half-turns on a second copy of the diagram. These parallelograms tessellate the plane with lattice units; lattice units giving distinct tessellations have one corner at each of the top and right mid-sides and also at the centre of the one drawn. This much symmetry is a result of the way the dice are stacked and says nothing about the dice. More symmetrical cubes would produce more symmetric configurations.

A word on the faces of the dice. The faces have both less and one, in a sense, more symmetry than the squares that bound them. A square has axes of mirror symmetry both vertical and horizontal and diagonal and the consequent centre of half-turn where both pairs intersect. Because all four axes are there, that centre is a centre of quarter-turn symmetry where they intersect. The six-face of a die has vertical and horizontal axes of reflection and the consequent half-turn. The two-face and the three-face have diagonal mirrors and the consequent half-turn. The faces with one, four, and five pips have both pairs of axes and so have the full symmetry of the face. What I have said about the faces is mostly true of just the motifs of pips, but the single pip is the exception, since it has the full rotational symmetry of a circle, much more symmetry than its face. Something I wonder about is the contribution of the symmetry of faces and motifs to one's appreciation of the symmetry of the whole. There seems to me to be little contribution on the part of the faces in this configuration. Does one think of the dice as having the symmetry that they have in spite of the fact that they are arranged in an unmoving way and cannot move independently of the others? And what contribution is made by the symmetries of the rhombic face diagrams, which are different from the square faces themselves: none can have greater symmetry than the rhomb itself, and the six-face has less. But the others all have the full symmetry of the rhomb itself. Let's look at a similar diagram with the pips replaced by six colours similarly arranged.
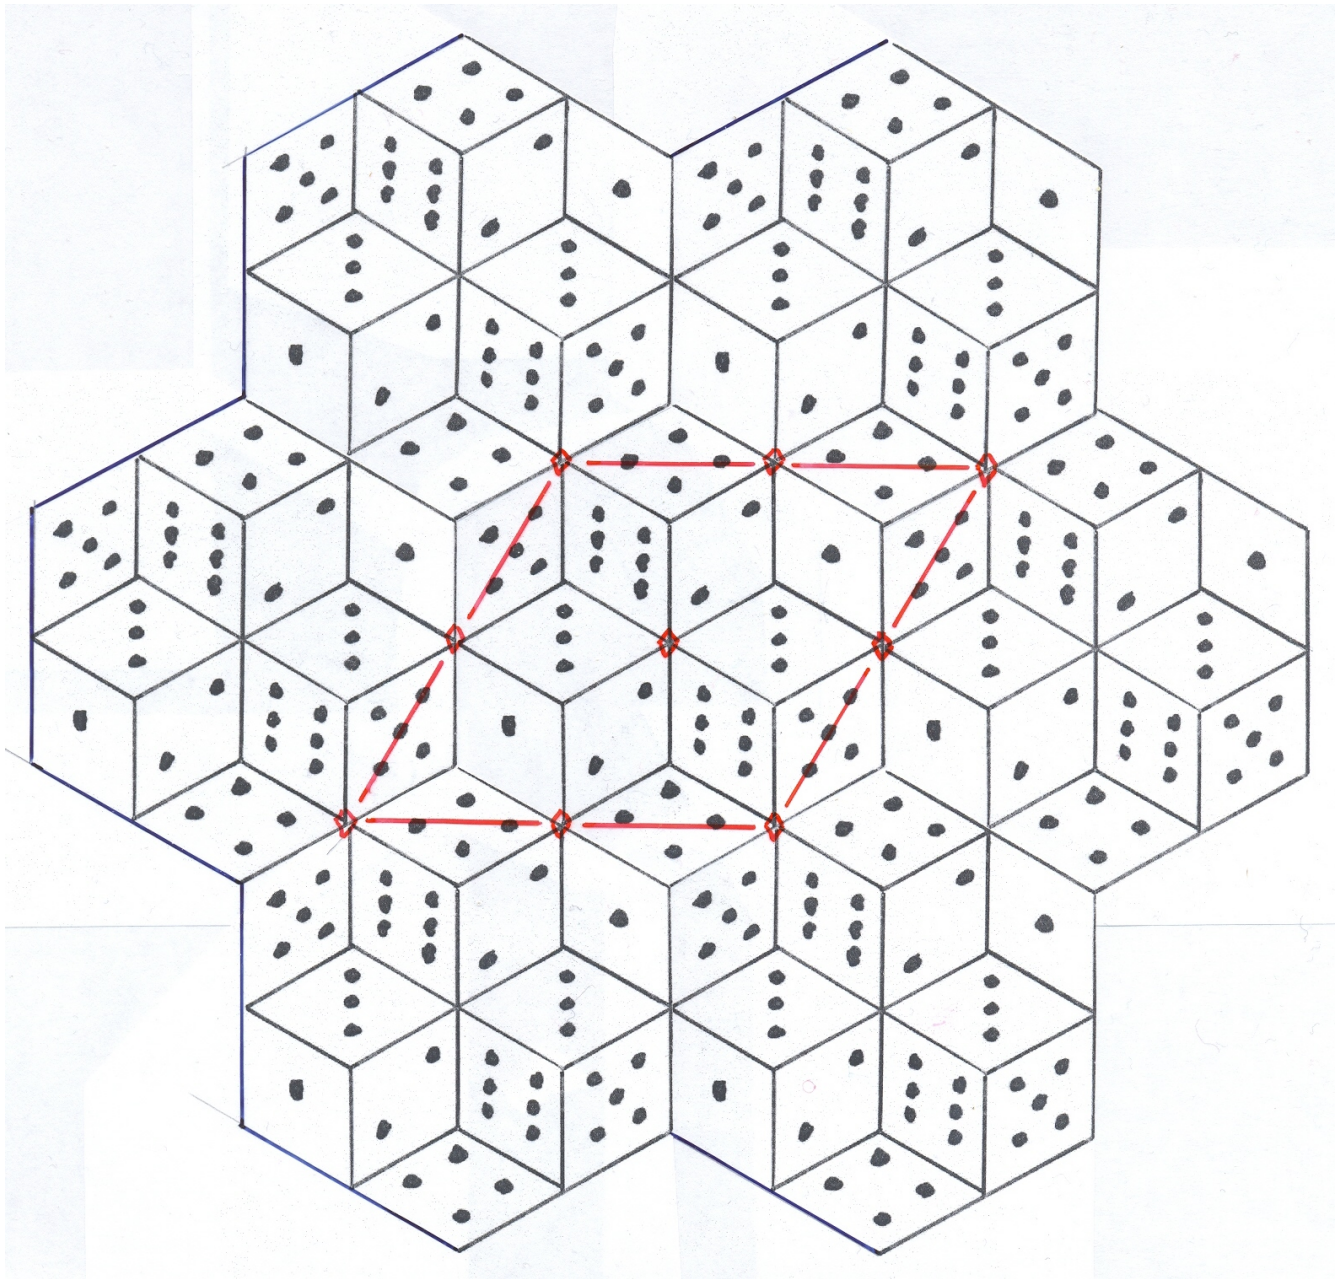
This diagram introduces what is called colour symmetry, as distinguished from simple symmetry. The simple sort has to do with what you can (in prinicple) do to something and have it look the same. Our approximate bodily bilateral symmetry has to do with reflecting the surface of our bodies in a mirror down the middle, certainly something possible to consider only in principle. Colour symmetry has to do with operations that have something look the same except for colour, and the change of colour, if any, has to be coherent so that *everything* that is red has to be turned, say, blue, and everything blue has to be turned something else in that case. If nothing is turned blue, then it can stay blue. The half-turn symmetry of the dice is preserved when the pips become colours with no colour change. But now there are a variety of other transformations that the array of rhombs can undergo that change colour. One is the 120-degree rotation about the top front corners of the cubes. If the cube has colours yellow, green, and brown visible, then those three colours are permuted and also blue, red, and orange. On the other hand, if the cube has orange, green, and red visible, then they are permuted and so are yellow, blue, and
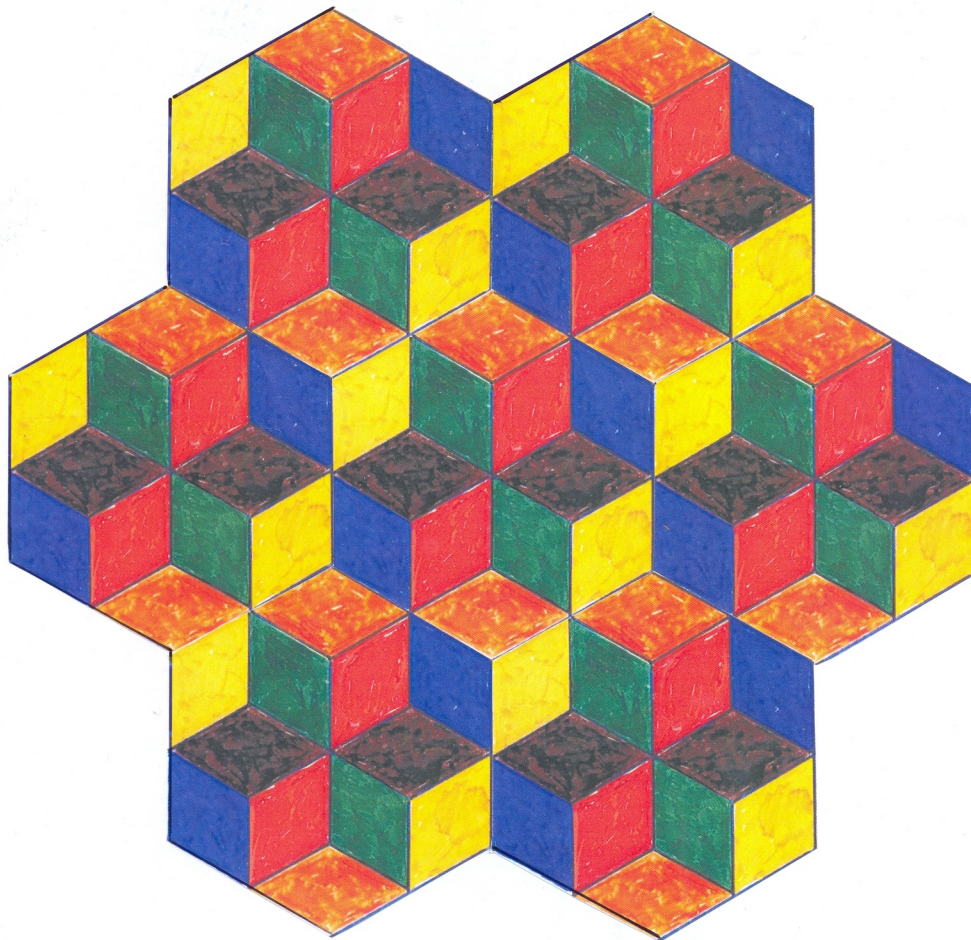
brown, and so on for the other two such kinds of corner (red, blue, and brown and yellow, blue, and orange). There are also 60-degree rotations about the centres formerly of half-turns where three cubes come together and centres of half-turns in the centre of each rhomb. The latter rotation leaves the colour of the rhomb unchanged while interchanging some other colours. For example, a half-turn about the centre of a yellow face leaves yellow and red faces yellow and red but interchanges those that are blue and green and those that are orange and brown. And I am not mentioning axes of reflection and glide-reflection. This group is the wallpaper group *p6m,* which is the most complicated wallpaper group. It seems natural to think of this as a sort of symmetry, although it does not seem reasonable to say quite equivalently that the same transformations leave the configuration of dice faces invariant except for a permutation of numbers of pips. We have a more natural response to changes of colour than numbers of pips.

Another example of colour symmetry, albeit with only two colours, is the pair of woven cubes in the first diagram. The pattern on those cubes is not left invariant by all of the 24 rotations of the solid cube, although the weaving is. While the 120-degree rotations about a diagonal from a vertex through the cube to the opposite vertex does leave the pattern unchanged, the half-turns around axes from mid-edge to mid-edge and the quarter-turns about the centres of faces reverse the two colours. Those rotations of a cube are important in what follows now that the introduction is out of the way.
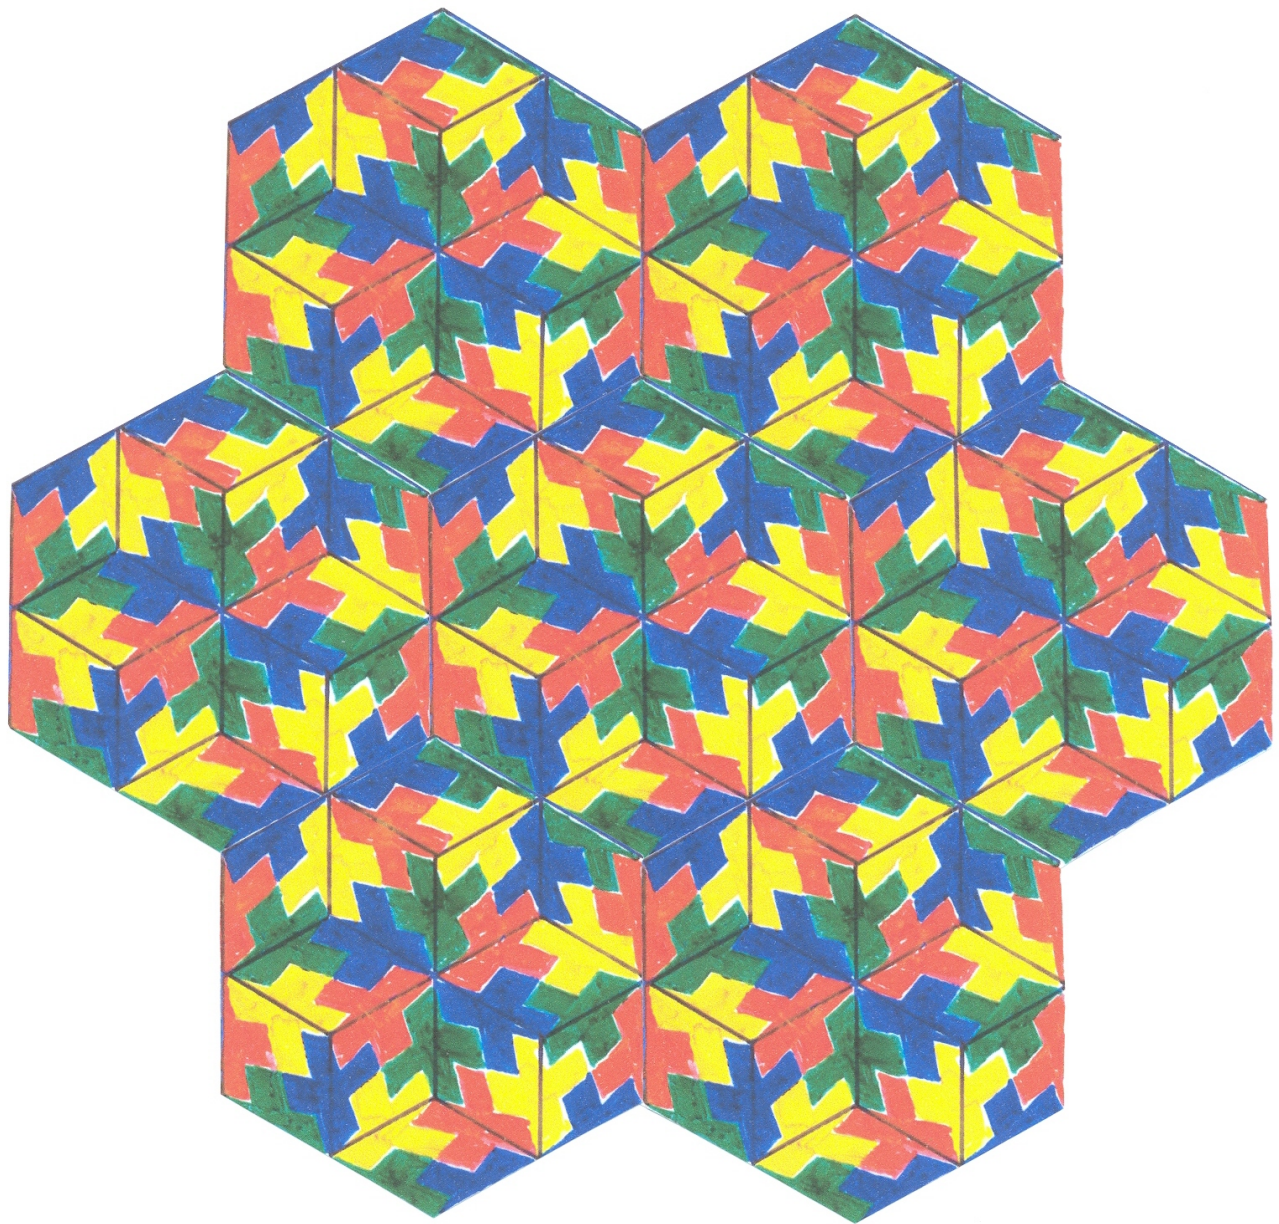
The cubes in this diagram are also woven but with four colours. In the middle of each rhomb boundary a strand is crossing from one face to the adjacent face pretty much straight on, whereas another runs pretty much along that boundary. There is, on each such boundary, one strand doing each thing. Since they are the same colour in this pattern, you can't tell by looking at the area showing where one crosses over the other which one is visible by the colouring of that area. So you need to be told which strand is over which strand. The strip just crossing the rhomb boundary crosses over the strip containing the rhomb boundary, which is drawn in to make it clearer in two dimensions that the diagram is one of cubes and not just a strange picture. Let me say something about the cubes. Each face is different from each other face; there are six permutations of four colours taken four at a time and six faces of the cube to accommodate them all. The quarter-turns, half-turns, and three-quarter-turns about the centres of faces permute the colours, as do the 120-degree rotations about vertices. The half-turn about an axis from the centre of one cross to the centre of the cross on the opposite edge leaves those two crosses invariant but also the other two crosses of those two colours invariant in colour while reversing the colours of the crosses of the other two colours. E.g., consider the axis running from a red

cross on the front edge of a cube to the green cross pictured on the front of the cube beside it. The green crosses adjacent to the red cross are interchanged, as are the red crosses adjacent to the green cross, while the blue crosses are interchanged with the yellow crosses. One of the amusing aspects of these cubes is where the strips go. Each one is a loop that makes two appearances on each face, once in each direction. And on a cube there is only one strip of each colour. These facts can be checked on the diagram; that was what motivated constructing it in the first place.

There is a lot of symmetry to the cubes pictured here, but I am mainly interested in drawing attention to the picture. It has only the strict symmetry that is forced upon it by the way the cubes are arranged, exactly as the dice were arranged. There are half turns where three cubes meet. But the colour symmetry contains the same rotations as the coloured cubes, 60-degree rotations where three cubes meet, 120-degree rotations about the front corners of cubes, and half-turns in the centres of rhombs. It lacks completely the mirror symmetries of the plainly coloured cubes, having the symmetry group *p6*. I
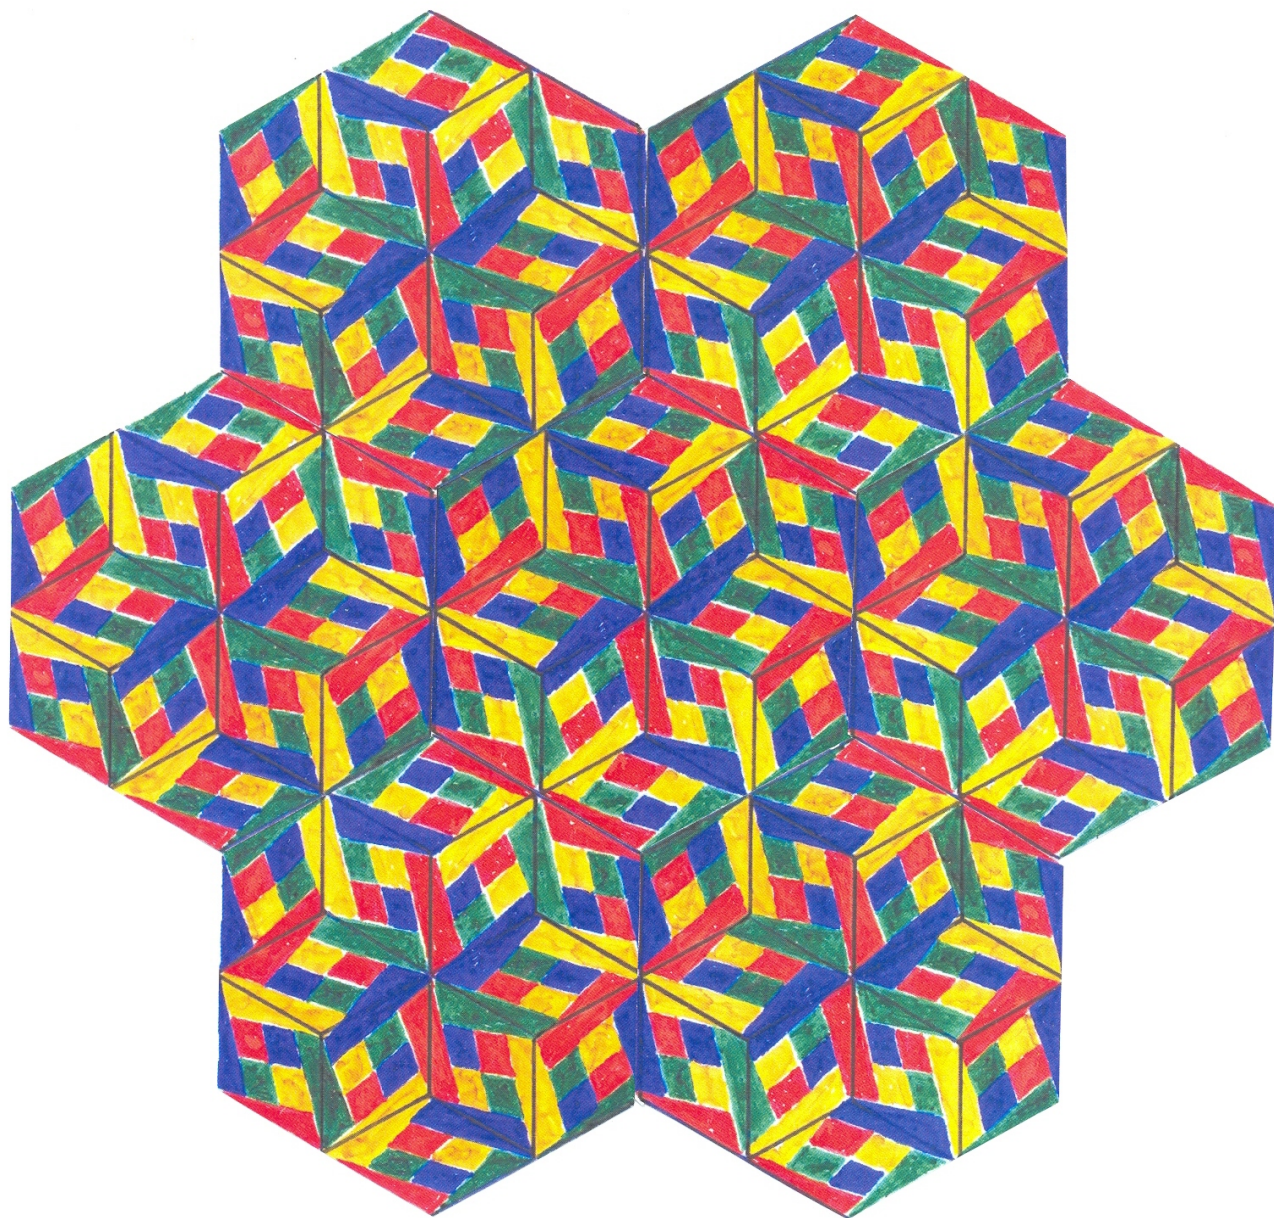


find symmetries attractive, but I find this diagram interesting too because it seems to have two

symmetries that it does not have. It appears to have symmetries of the cubes that do not hold of the two-dimensional diagram. More weakly it seems to have the quarter-turn symmetry of the cross motif, but more strongly for me it seems to have quarter-turn symmetry about the centres of the cube faces. One might say that one *notices* the symmetry of the motif, but the way the crosses fit together gives me a *feeling* of rotational symmetry. The fact that the quarter-turns squared *are* symmetries helps. I wonder whether others have this feeling, which could be explained by our tending to imagine the cubes, each of which does have that symmetry.
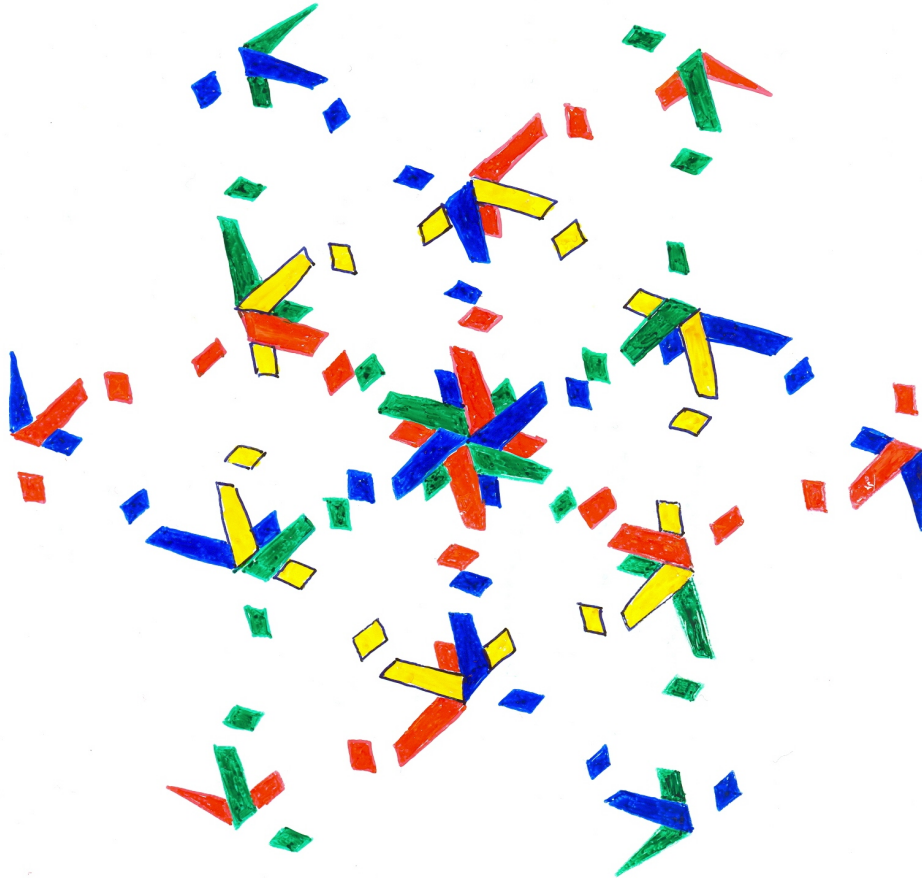


Being curious about this question and in particular how much the crosses contribute to my feeling, I have constructed a diagram of the other similar woven cube available, which does not have the crosses. In this diagram, when a colour crosses itself the strip along the edge of the cube passes over the strip at right angles to it; that is, the latter strip does not pop up to fill that cell but leaves it to the strip filling the adjacent cells along the edge. This cube – because it is a single cube pictured over and over – has the same symmetries as the other but its motif is the disconnected stick with a dot on each side separated from it as though the cross had been pulled apart (division sign instead of plus sign). That eliminates the cross-symmetry illusion if it was there before. But I do not find that it eliminates the imaginary quarter-turn of the rhombs. I wonder about others. This is the question to which I alluded at the start.

The tracing of strips can refer (as above) to their behaviour on a single cube, but it can also refer to the picture itself. It is easier, I find, to trace them on the latter diagram. As I said earlier, each strip on a cube is a loop, but each strip in the picture is also a loop. For example, consider the pair of three yellow cells in a vertical row adjacent to each other corner to corner at about eight o'clock in the diagram. The strip of the upper right of them can be traced through 22 rhombs in 12 cubes before returning to where it began (a total of 24 rhombs). The path is symmetric about the centre of the diagram. Most such

circuits partly escape the diagram. The green, red, and blue strips, touching each other in monochrome pairs at the centre and crossing one another both there and elsewhere and crossing the yellow strip twice, are pictured separately as far as they go in the diagram. That picture does give some idea how loops of different colours interact, but it gives no idea at all how the many loops of the same colour interact. To show that I have done another picture of just yellow loops in the original diagram. I have coloured them with different colours to make clearer how they are arranged in two touching triples around and linked to a central loop and linked to each other. These are traced and extended from the initial picture for the sake of distinctness.
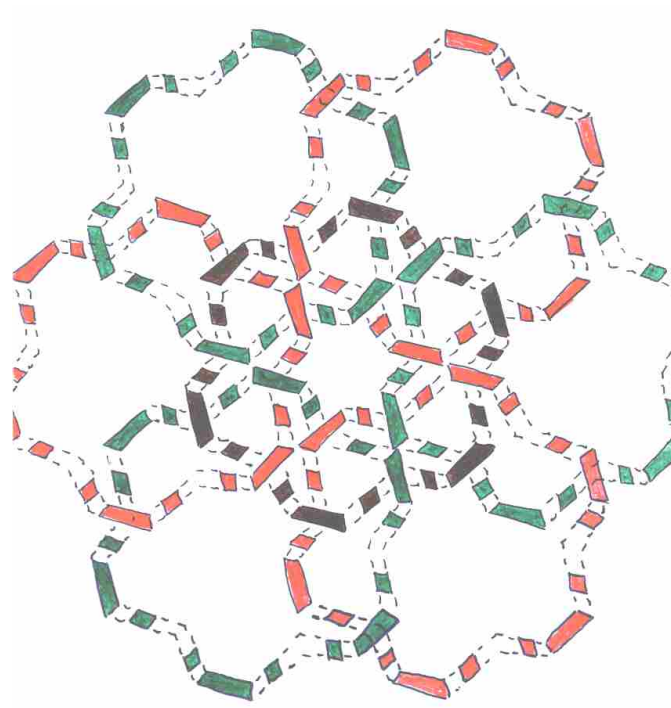


Since you have come this far, I'll show you one more picture, a coloured tiling of the plane by birds. They are not very realistic birds partly on account of my being unable to draw well but mainly on account of the severe constraints placed on their shape by their being more or less plus signs. Being more or less plus signs, they make a tiling with the symmetries of one of the woven-cube diagrams. What about the illusion that they are subject to quarter-turns at wing-tips? While the rest of the picture would no doubt have made him cringe, I think that Escher would have enjoyed the illusion of such an impossible transformation.

Exercise. (a) Repeat with fish. (b) Animate a transition from fish to birds.