
Contents

List of Figures	ix
Preface	xi
Introduction	1
Chapter 1. Ergodic Theory	17
1. The idea behind ergodic theory	17
2. Some mathematical structure	20
3. Substitution tilings	25
4. Finite type tilings	33
Chapter 2. Physics (for Mathematicians)	55
1. Diffraction of light waves	55
2. Statistical mechanics	58
Chapter 3. Order	63
1. Spectrum and order	63

Chapter 4. Symmetry	69
1. Substitution and rotational symmetry	69
2. Spectrum and X-rays	81
3. 3-dimensional models	86
Chapter 5. Conclusion	99
Appendix I. Geometry	103
Appendix II. Algebra	105
Appendix III. Analysis	109
References	113
List of Symbols	117
Subject Index	119

Preface

“In the world of human thought generally, and in physical science particularly, the most important and most fruitful concepts are those to which it is impossible to attach a well-defined meaning” – an intriguing idea from the physicist H.A. Kramers [Kra]. Even within mathematics, where the approach to knowledge is somewhat different from the physical sciences, a version of this thesis would appeal to those engaged in research: in prosaic terms – anything well-understood is less promising as a research topic than something not-well-understood.

This book is motivated by something not-well-understood, a class of structures exemplified by the “Penrose tilings”, or, more specifically, the “kite & dart tilings” (Fig. 1, on p. 2). These structures are not-well-understood on a grand scale, having had significant impact in physics and mathematics, and originating from work done in the 1960’s in philosophy! A few decades is a short time in mathematics, so it is reasonable that such a fertile subject is not yet well-understood. Yet it is the subject of this book.

Such tilings differ in a variety of ways from any studied before, and it may be many years before we are confident of how they fit into the body of mathematics, and the most useful ways to view

them. Our search to understand them will draw us into many parts of mathematics – including ergodic theory, functional analysis, group representations and ring theory, as well as parts of statistical physics and crystallography. Such breadth invites an unusual format for this mathematics text; rather than present a full introduction to some corner of mathematics, in this book we try to display the value (and joy!) of starting from a mathematically amorphous problem and combining ideas from diverse sources to produce new and significant mathematics – mathematics unforeseen from the motivating problem.

The background assumed of the reader is that commonly offered an undergraduate mathematics major in the US, together with the curiosity to delve into new subjects, to readjust to a variety of viewpoints. The book is self-contained; subjects such as ergodic theory and statistical mechanics are introduced, *ab initio*, but only to the extent needed to absorb the desired insight.

I hope I have imparted in the text the excitement I have enjoyed in the journey through the diverse subject matter. Part of the pleasure has been learning from friends and colleagues. My formal training was in physics, and it would be impossible to name all to whom I am indebted for that; as for mathematics, it is a great pleasure to acknowledge: Persi Diaconis, Charles Fefferman, Richard Kadison, Raphael Robinson, Hao Wang, and especially John Conway and Mark Kac, for insight and inspiration you cannot find in books.

It is also a great pleasure to thank Marjorie Senechal and Jeffrey Lagarias for a great number of useful comments and encouragement; this book would certainly not have appeared without their help. And finally, special thanks are due Stuart Levy for creating the program “subst” for the display of substitution tilings; in particular I used it to make the figure of the quaquaversal tiling used in this book.

Most of this book was written in 1997; for further results see the papers listed in <http://www.ma.utexas.edu/users/radin/>.

Introduction

The “kite & dart tiling”, pictured in Fig. 1, has been widely publicized in the last decade or two. *Why?* There are lots of reasons actually, and in this book we will concentrate specifically on the mathematics which they have inspired, scattered in totally unexpected directions.

The story began in the early 1960’s, with the philosopher Hao Wang modeling certain problems in logic [Wan]. It slowly evolved into geometry, in large part from influential work of Raphael Robinson [Rob] and the kite & dart tiling of Roger Penrose [Gar]. We will pick up the story there, and follow it through the twists and turns it has undergone, to the new mathematics that is emerging. It has been a highly interdisciplinary journey, and though we will strongly emphasize the mathematics (chiefly geometry and modern analysis), we will not flinch from analyzing those ideas in physics and crystallography which will help us understand the mathematics. Personally, I find that a good part of the fun of the subject.

Now patterns like Fig. 1 are pretty, but at least once in a while it is useful to face the essence of our endeavor in its raw form. The feature of Fig. 1 that we first emphasize is the *large number* of polygons in it. (There are infinitely many in the full tiling of course.) Analogous

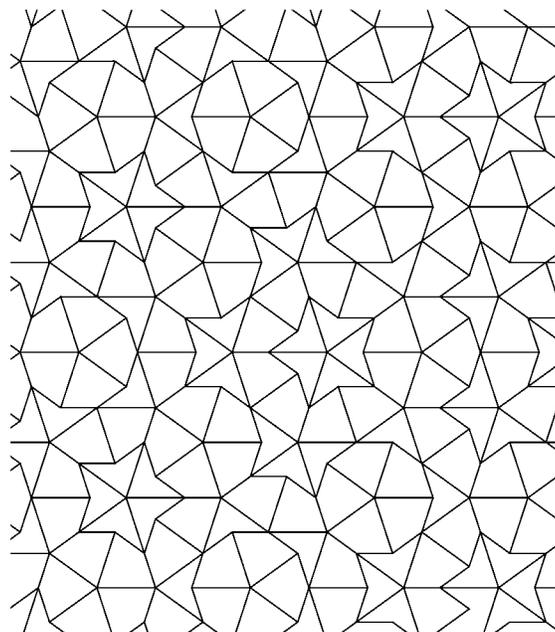


Figure 1. A Penrose “kite & dart” tiling

physical patterns are: a quartz rock made of many atoms; a snail made of many cells; or a beach made of many grains of sand. In general, we will be analyzing “global” structures made out of many small components.

We concentrate not on the external shape of the global structure, such as the facets of a quartz rock, but on the pattern made at a much smaller scale by the small components. For a snail this is quite complicated: the cells gather together into intermediate-size structures which we call organs. For beaches the small scale structure is rather “random”. But rocks are neither as random as beaches, nor as exotic as snails; the atoms in rocks form patterns of intermediate complexity, called crystals.

We know roughly why atoms in rocks form crystals while sand grains on the beach or cells in a snail do not; what we say is that *there are different laws governing the production of these structures* (laws studied in physics, biology etc.), and these different laws naturally lead to different results. But this is all a bit vague, and when one examines this explanation carefully there are serious but very interesting difficulties with it. This will be our subject – why certain kinds of laws or rules seem to produce very special global structures, such as quartz crystals. (Snails are a much harder problem, and beaches are too simple; neither will be mentioned seriously again!) In the next few paragraphs we must get more specific about this idea of rules which produce structures.

This book is about mathematics, not physics, so it will be useful to have in mind models of this structural phenomenon other than rocks, with all their irrelevant details. The general model we will use is the jigsaw puzzle, one made with *very* many pieces with bumpy edges, pieces which we will call “tiles”. If one imagines the tiles to have various colors painted on them it is easy to see how a tile contributes to a global structure or pattern. We will ignore any such colors on the tiles, and only concentrate on their shapes – as if we turned the jigsaw puzzle pieces over; we think of the global pattern as “consisting of” these special shapes fitted together. The bumpy edges of the tiles play an important role in determining how the tiles are allowed to fit together to form the global pattern. (We will also consider 3-dimensional versions of the more traditional 2-dimensional puzzles.)

But rocks will still be useful to us; we will use features of rocks to guide us in our mathematical analysis of patterns. For instance, the tiles in a jigsaw puzzle could represent *any* structure if there were no restrictions. If you wanted certain star-shaped pieces to lie in particular places in a puzzle, you could just chop up the intervening space into other pieces to accommodate the stars. Now one of the reasons the atoms in rocks do not appear in arbitrary (local) structures is

that there are only a small number of possible components, the 92 different naturally occurring kinds of atoms. So we ask: if you were a manufacturer of jigsaw puzzles and were limited to using, say, 92 different tiles of your choice (but could make as many copies of each shape as you wanted), what kind of giant patterns could you produce? Remember that for us the pattern has nothing to do with colored pictures on the tiles, but merely with the manner in which the various tiles fit together. With this constraint of only 92 different shapes to use it is no longer clear that you could make *any* (local) pattern of tiles, and the question is: what kinds of patterns could you make? This is where our discussion about rules of production has led us; we assume we have some small number of different kinds of elementary building blocks, and ask what kinds of patterns can be made out of them given the restriction that the pieces have to fit together like a jigsaw puzzle. Later we will discuss why the restriction that the pieces fit together should be thought of as a law of production.

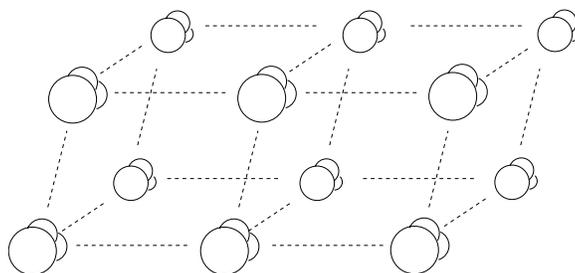


Figure 2. 12 unit cells of a periodic configuration

Now let's refine what we mean by "kinds of patterns". Again we look at crystals for inspiration. The distinguishing special feature of crystals is that they each have a unit cell from which the global pattern is obtained by translation, as in Fig. 2. (The dashes in the figure show the translations separating the 12 unit cells, each unit cell consisting of three spheres.) This repetition is an example of an

“order” property, very different from the randomness of the positions of the sand grains of a beach. Also, periodic structures such as crystals are limited in the symmetries they may have; for instance, a crystal can only have axes with 2, 3, 4 and 6-fold rotational symmetry [HiC; p. 84]. So in studying the “kinds” of structures, we will concentrate on the order and, especially, the *symmetry* properties they may possess.

It is time to examine some examples. First consider the periodic jigsaw using the following K^2 different shapes, where $K \geq 2$ is fixed but arbitrary. They are all basically unit squares, but with certain bumps coming out of some edges and dents going into some edges (any bump fitting into any dent), following the formula: the tile labeled (i, j) (where $1 \leq i, j \leq K$) has i dents on its top edge, $(i + 1)$ bumps on its bottom edge, j dents on its left edge and $(j + 1)$ bumps on its right edge. The exceptions are: $(1, j)$ has K dents on its top edge; $(i, 1)$ has K dents on its left edge; (K, j) has 1 bump on its bottom edge; and (i, K) has K bumps on its right edge. Fig. 3 shows tile $(3, 5)$, assuming $5 < K$. It’s not hard to show that the only way to put these together into a big jigsaw puzzle makes a “ $K \times K$ unit cell” (several are shown in Fig. 4 for $K = 6$), and builds a periodic pattern from copies of the cell. Notice that although the specification of the unit cell is not unique (see an alternative choice in Fig. 5), in a sense the full pattern is unique: there is really only one way to build a global structure from these tiles, up to an overall rigid motion of the plane. (See Appendix I for a review of congruences of Euclidean 2- and 3-dimensional space.)

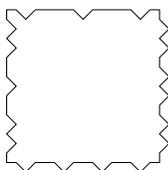


Figure 3. The tile $(3,5)$

1,1	1,2	1,3	1,4	1,5	1,6	1,1	1,2	1,3	1,4	1,5	1,6
2,1	2,2	2,3	2,4	2,5	2,6	2,1	2,2	2,3	2,4	2,5	2,6
3,1	3,2	3,3	3,4	3,5	3,6	3,1	3,2	3,3	3,4	3,5	3,6
4,1	4,2	4,3	4,4	4,5	4,6	4,1	4,2	4,3	4,4	4,5	4,6
5,1	5,2	5,3	5,4	5,5	5,6	5,1	5,2	5,3	5,4	5,5	5,6
6,1	6,2	6,3	6,4	6,5	6,6	6,1	6,2	6,3	6,4	6,5	6,6
1,1	1,2	1,3	1,4	1,5	1,6	1,1	1,2	1,3	1,4	1,5	1,6
2,1	2,2	2,3	2,4	2,5	2,6	2,1	2,2	2,3	2,4	2,5	2,6
3,1	3,2	3,3	3,4	3,5	3,6	3,1	3,2	3,3	3,4	3,5	3,6
4,1	4,2	4,3	4,4	4,5	4,6	4,1	4,2	4,3	4,4	4,5	4,6
5,1	5,2	5,3	5,4	5,5	5,6	5,1	5,2	5,3	5,4	5,5	5,6
6,1	6,2	6,3	6,4	6,5	6,6	6,1	6,2	6,3	6,4	6,5	6,6
1,1	1,2	1,3	1,4	1,5	1,6	1,1	1,2	1,3	1,4	1,5	1,6
2,1	2,2	2,3	2,4	2,5	2,6	2,1	2,2	2,3	2,4	2,5	2,6

Figure 4. Part of a periodic tiling, with unit cells outlined

Next consider some tilings, due to Jarkko Kari and Karel Culik, made from copies of the shapes in Fig. 6. As is shown in [Cul], one can build arbitrarily large collections of these tiles, but none with a unit

1,1	1,2	1,3	1,4	1,5	1,6	1,1	1,2	1,3	1,4	1,5	1,6
2,1	2,2	2,3	2,4	2,5	2,6	2,1	2,2	2,3	2,4	2,5	2,6
3,1	3,2	3,3	3,4	3,5	3,6	3,1	3,2	3,3	3,4	3,5	3,6
4,1	4,2	4,3	4,4	4,5	4,6	4,1	4,2	4,3	4,4	4,5	4,6
5,1	5,2	5,3	5,4	5,5	5,6	5,1	5,2	5,3	5,4	5,5	5,6
6,1	6,2	6,3	6,4	6,5	6,6	6,1	6,2	6,3	6,4	6,5	6,6
1,1	1,2	1,3	1,4	1,5	1,6	1,1	1,2	1,3	1,4	1,5	1,6
2,1	2,2	2,3	2,4	2,5	2,6	2,1	2,2	2,3	2,4	2,5	2,6
3,1	3,2	3,3	3,4	3,5	3,6	3,1	3,2	3,3	3,4	3,5	3,6
4,1	4,2	4,3	4,4	4,5	4,6	4,1	4,2	4,3	4,4	4,5	4,6
5,1	5,2	5,3	5,4	5,5	5,6	5,1	5,2	5,3	5,4	5,5	5,6
6,1	6,2	6,3	6,4	6,5	6,6	6,1	6,2	6,3	6,4	6,5	6,6
1,1	1,2	1,3	1,4	1,5	1,6	1,1	1,2	1,3	1,4	1,5	1,6
2,1	2,2	2,3	2,4	2,5	2,6	2,1	2,2	2,3	2,4	2,5	2,6

Figure 5. Part of a periodic tiling, with different unit cells outlined

cell from which a tiling could be constructed by repeated translation as above.

Instead of talking about arbitrarily large collections of tiles, we will take the plunge and discuss from now on infinite collections which

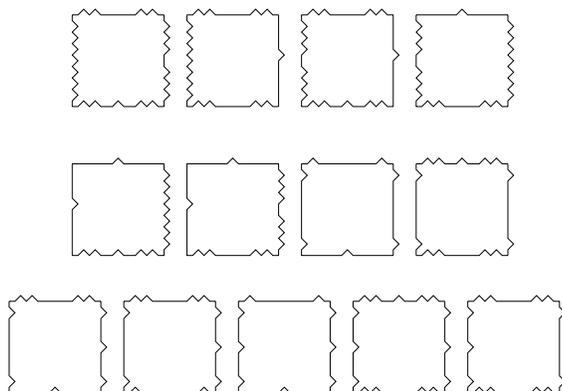


Figure 6. The Kari-Culik tiles

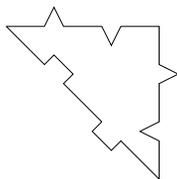


Figure 7. The “random tile”

fill the whole Euclidean 2- or 3-dimensional space. We will also use the word “tiling” in place of jigsaw puzzle for collections of tiles; by a tiling of a space we just mean a collection of tiles which completely covers the space, and such that for each pair of tiles in the tiling the interiors have empty intersection. As we shall see, there is not really much difference in dealing with infinite rather than arbitrarily large collections of tiles, and it will simplify some discussions. In particular, from the tiles of Fig. 6 one can make uncountably many tilings of the plane no two of which are congruent, but as noted above one cannot make a periodic one, that is, one made up of repeated translations of a unit cell as in Fig. 4.

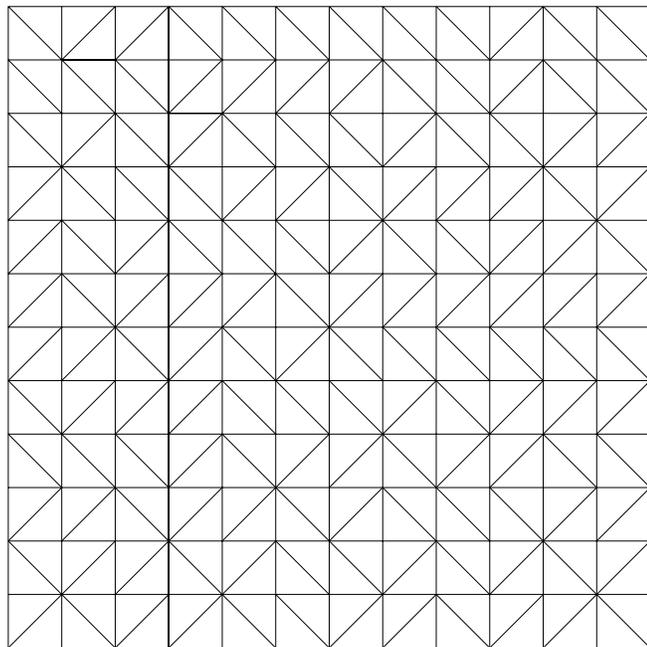


Figure 8. Part of a “random tiling”

We next consider the “random” tile in Fig. 7. The only tilings one can make out of copies of this tile are vaguely like a checkerboard, with the tiles pairing up along their hypotenuses to make the squares of what we will call a “random checkerboard” as in Fig. 8. Note that in pairing up, each square of the checkerboard is filled with a pair of tiles in one of two possible orientations, and that these two possible orientations (think of them as “red” and “black”) are independent in different squares. We could thus think of the possible tilings with these tiles as *any* tilings made with (aligned) red and black squares – not just those alternating in color as in a real checkerboard. This means we *can* tile in very complicated ways; but in contrast with

the previous example, we can *also* tile periodically, that is, in a very simple way, for instance the usual red-black checkerboard.

We want to contrast one feature of the tiles of Fig. 4 and of Fig. 7. From the first set we found that there was essentially only one tiling that could be made; any two tilings were in fact congruent. From the second set we found we could make a very wide variety of tilings, which had little to do with one another. We noted earlier rules or laws that produce global structures from components, and this is an appropriate place to expand on that.

It would be convenient if our rules took a set of tiles such as those in Fig. 6 and produced a specific tiling. However we are trying to understand how structures such as crystals are made, and the rules of nature are not of this simple type. As we shall see in Chapter 2, given a specification of particles such as iron atoms (or better yet, iron nuclei and electrons), the physical rules or laws that govern the production of bulk iron at low temperature do not actually pick out a specific particle configuration such as a particular crystal; the laws (called statistical mechanics) actually specify a large *collection of particle configurations*. So too the rules we will deal with will associate with a given set of tiles such as Fig. 6 not one specific tiling, but a (large) *collection of tilings*.

For instance, with a set \mathcal{A} of tiles such as Fig. 6 we can associate the set $X_{\mathcal{A}}$ of *all possible* tilings that one could make using congruent copies of those tiles. We think of this as a “rule” for \mathcal{A} . (We will consider other types of rules for \mathcal{A} later, which associate special subsets of $X_{\mathcal{A}}$, but unless otherwise indicated, given a set \mathcal{A} of tiles “the” rule for \mathcal{A} associates the set $X_{\mathcal{A}}$ of *all possible* tilings).

Now if one is given a tiling x and wants to determine whether or not it “follows from the rule associated with \mathcal{A} ”, that is, whether or not x is in $X_{\mathcal{A}}$, one must check three things: that the tiles making up x are congruent to elements of \mathcal{A} ; that in x they leave no uncovered gaps in the plane; and that they never overlap in x . Notice that this process could be carried out by examining x in any fixed disk D of

diameter large enough to properly contain any of the tiles of \mathcal{A} , and moving this viewing window throughout x ; it can be seen whether or not x belongs to $X_{\mathcal{A}}$ by such a *local* examination, where the word “local” emphasizes that at no time do we need to examine a portion of x larger than the finite size of our fixed disk D . In summary: a rule associates with some finite set \mathcal{A} of tiles not one tiling but some set X of tilings, and the essence of the rule is what one must do to see if a candidate tiling x satisfies the rule, that is, belongs to X . It is in this way that we have classified rules as local or not.

Getting back to the contrasted feature of the set \mathcal{A}_4 of tiles in Fig. 4 and \mathcal{A}_7 of the tile in Fig. 7, the difference we noted is that all the tilings in $X_{\mathcal{A}_4}$ are very similar to one another (in fact any two are congruent), whereas there are tilings in $X_{\mathcal{A}_7}$ which are very dissimilar from one another. The set \mathcal{A}_6 of tiles shown in Fig. 6 is of a different sort. We said that $X_{\mathcal{A}_6}$ contains tilings which are not congruent to one another. But these tilings are still very similar in a sense slightly weaker than congruence: every finite substructure of any one tiling in $X_{\mathcal{A}_6}$ has congruent copies in any other tiling in $X_{\mathcal{A}_6}$, with the consequence that one cannot tell the difference between the tilings in $X_{\mathcal{A}_6}$ by inspecting only finite portions of them. Now the point is, we consider the cases of $X_{\mathcal{A}_4}$ and $X_{\mathcal{A}_6}$ as satisfactory, but not that of $X_{\mathcal{A}_7}$. (This is partly motivated by the statistical mechanics of solids. After developing some ideas in Chapter 2 we will think of $X_{\mathcal{A}_7}$ as resulting from an “accidental symmetry” which allows both orientations of the pairs of tiles making up squares.) With this prejudice firmly in place, we will concentrate on sets \mathcal{A} of tiles without such accidental symmetries, tiles which only produce sets $X_{\mathcal{A}}$ of tilings in which each pair, while not necessarily congruent, is locally indistinguishable in a sense made precise later. One consequence is that for the sets \mathcal{A} of tiles which we will consider, the tilings in $X_{\mathcal{A}}$ are either all simple (for instance periodic), or all complicated.

Consider next the tiles in Fig. 9 and the following “substitution rule” for making tilings from them. For each tile T we have a way

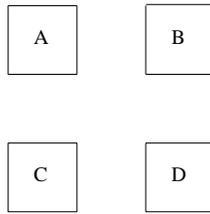


Figure 9. The Morse tiles

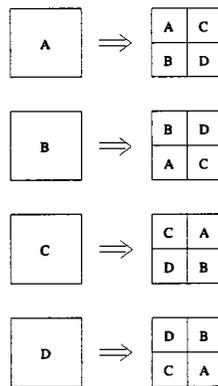


Figure 10. The Morse substitution system

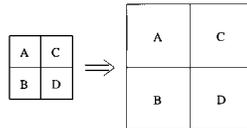


Figure 11. Expansion

(Fig. 10) to associate with it a collection of tiles at a smaller scale (a factor $1/\gamma < 1$ times those in Fig. 9, with $\gamma = 2$ in this case), in a

D	B	B	D	B	D	D	B	B	D	D	B	D	B	B	D
C	A	A	C	A	C	C	A	A	C	C	A	C	A	A	C
C	A	A	C	A	C	C	A	A	C	C	A	C	A	A	C
D	B	B	D	B	D	D	B	B	D	D	B	D	B	B	D
C	A	A	C	A	C	C	A	A	C	C	A	C	A	A	C
D	B	B	D	B	D	D	B	B	D	D	B	D	B	B	D
D	B	B	D	B	D	D	B	B	D	D	B	D	B	B	D
C	A	A	C	A	C	C	A	A	C	C	A	C	A	A	C
C	A	A	C	A	C	C	A	A	C	C	A	C	A	A	C
D	B	B	D	B	D	D	B	B	D	D	B	D	B	B	D
D	B	B	D	B	D	D	B	B	D	D	B	D	B	B	D
C	A	A	C	A	C	C	A	A	C	C	A	C	A	A	C
D	B	B	D	B	D	D	B	B	D	D	B	D	B	B	D
C	A	A	C	A	C	C	A	A	C	C	A	C	A	A	C
D	B	B	D	B	D	D	B	B	D	D	B	D	B	B	D
C	A	A	C	A	C	C	A	A	C	C	A	C	A	A	C

Figure 12. A Morse tiling

fixed relation in space to T . What we do then is start with some tile somewhere in the plane, “break it up” as is done to tile A in Fig. 10, then expand the small tiles by the factor $\gamma > 1$ as in Fig. 11. We repeat this for each of the tiles we now have, again and again; see Fig. 12. After infinitely many iterations we have a tiling of the plane, at least if we are careful in choosing each time the place about which to expand. It turns out that different sequences of choices of points about which we expand can lead to tilings which are noncongruent; so a “substitution rule” associates a large number of special tilings with a set \mathcal{A} of tiles.

This is a useful way of producing tilings; one can produce interesting tilings by this substitution method. But in a way it seems

too easy; we expect that a more complicated mechanism should be required to govern the construction of the complicated structures of interest to us. One hint to an understanding of this situation is that the substitution rule is not as “local” as was the one used for the tiles of Figs. 4, 6 or 7, because to check if a tiling is correctly made by this substitution process one must examine arbitrarily large regions of the tiling to see if they are put together correctly. We need to discuss this difference, because it goes to the heart of our subject. After all, we said we thought the atoms in quartz make different patterns from the cells of a snail because they follow different rules, so we need to be fussy about what kind of rules we use in making our tilings.

One of the big lessons in physics in this century was the choice made, following Einstein, to follow the field approach to physics of Maxwell’s theory of electromagnetism rather than the action-at-a-distance approach of Newton’s theory of gravitation [EiI]. In electromagnetism the influences at a given point (and time) are determined by the immediate environment of that point, while in classical gravitation the motion of a planet at one location is influenced (instantaneously) by the position of the sun (and planets) far away. Postponing the details until Chapter 2, we just note here that the laws governing the atomic structure of solids are also of the local variety, and we use this fact to “prefer” such rules in our tiling problems. As we noted before, the rule which, with any given set of tiles, associates *all* tilings that one can make with those tiles, is local, while a rule which associates with a set of tiles only those tilings made by a special technique may or may not be local; for instance a substitution rule is not local. However, substitution rules will still be useful – for instance they will lead to new ideas about the notion of symmetry.

Fundamentally, symmetry means the invariance of something when that something is acted upon in some way. This is so general, and so useful, it is hard to imagine altering it. Indeed, when we say there are new ideas about symmetry, what we are referring to is much more specific. By far the most useful symmetries of patterns

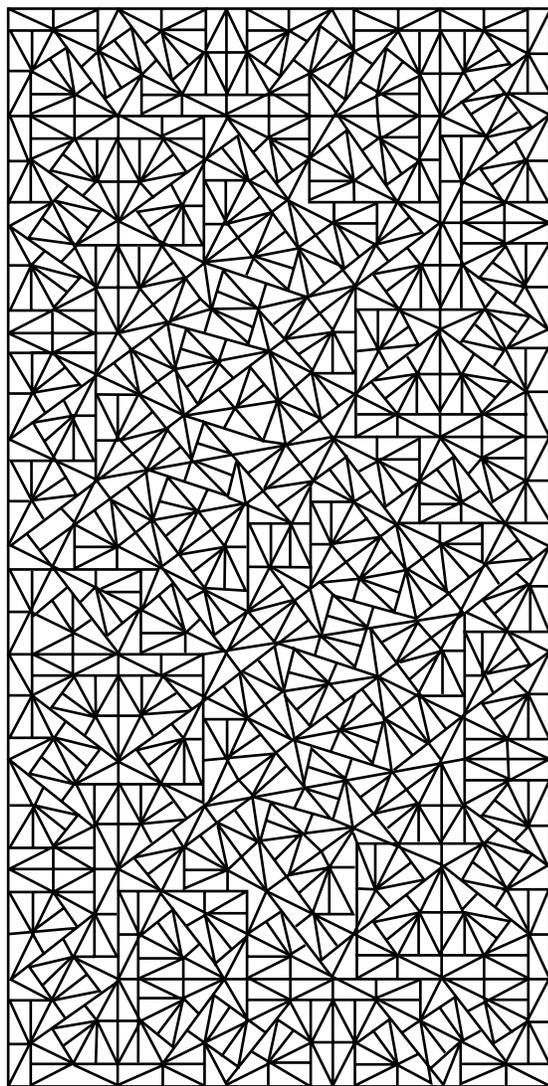


Figure 13. A pinwheel tiling

in space correspond to invariance of the pattern under very special transformations: rigid motions of the underlying space. The patterns we will be dealing with, such as the kite & dart (Fig. 1) or pinwheel (Fig. 13) tilings, are usually not invariant under any nontrivial rigid motion. But the notion of symmetry for such a tiling can be altered by focusing not on the tiling itself but on the relationships between its component tiles.

For each tiling we can construct a set of “frequencies” of its finite parts. That is, for each finite collection p of tiles in the tiling x , count the number of times p appears – in the same orientation – in a ball of volume N , and divide by N , defining, in the limit $N \rightarrow \infty$, a frequency $\nu(p)$. (We will deal with the existence of these limits.) One can ask whether such frequencies are invariant when the tiling x is moved by a rigid motion, such as a rotation. The interesting thing about, say, a kite & dart tiling, is that all its frequencies are invariant under a rotation by $2\pi/10$ even though no kite & dart tiling is itself invariant under a rotation by $2\pi/10$.

So the new notion of symmetry is still geometric, in the sense of corresponding to rigid motions; what is new is a shift in the quantity that is invariant – instead of the tiling itself, we focus on the frequencies of its finite parts.

In a nutshell our book will be about how global patterns are produced, that is, the kinds of production rules they have, and the ways in which these global patterns can exhibit order and symmetry. The production rules we consider do not necessarily produce single or unique structures, even up to congruence. But as discussed earlier they almost do this; any two structures produced by such a rule must be indistinguishable by “local” inspection, looking at any finite portion of the structure.