OXFORD MASTERCLASSES IN GEOMETRY 2014.
Part 2: Lectures on Penrose Tilings,
Prof. Alexander F. Ritter.

Comments and corrections are welcome: ritter@maths.ox.ac.uk

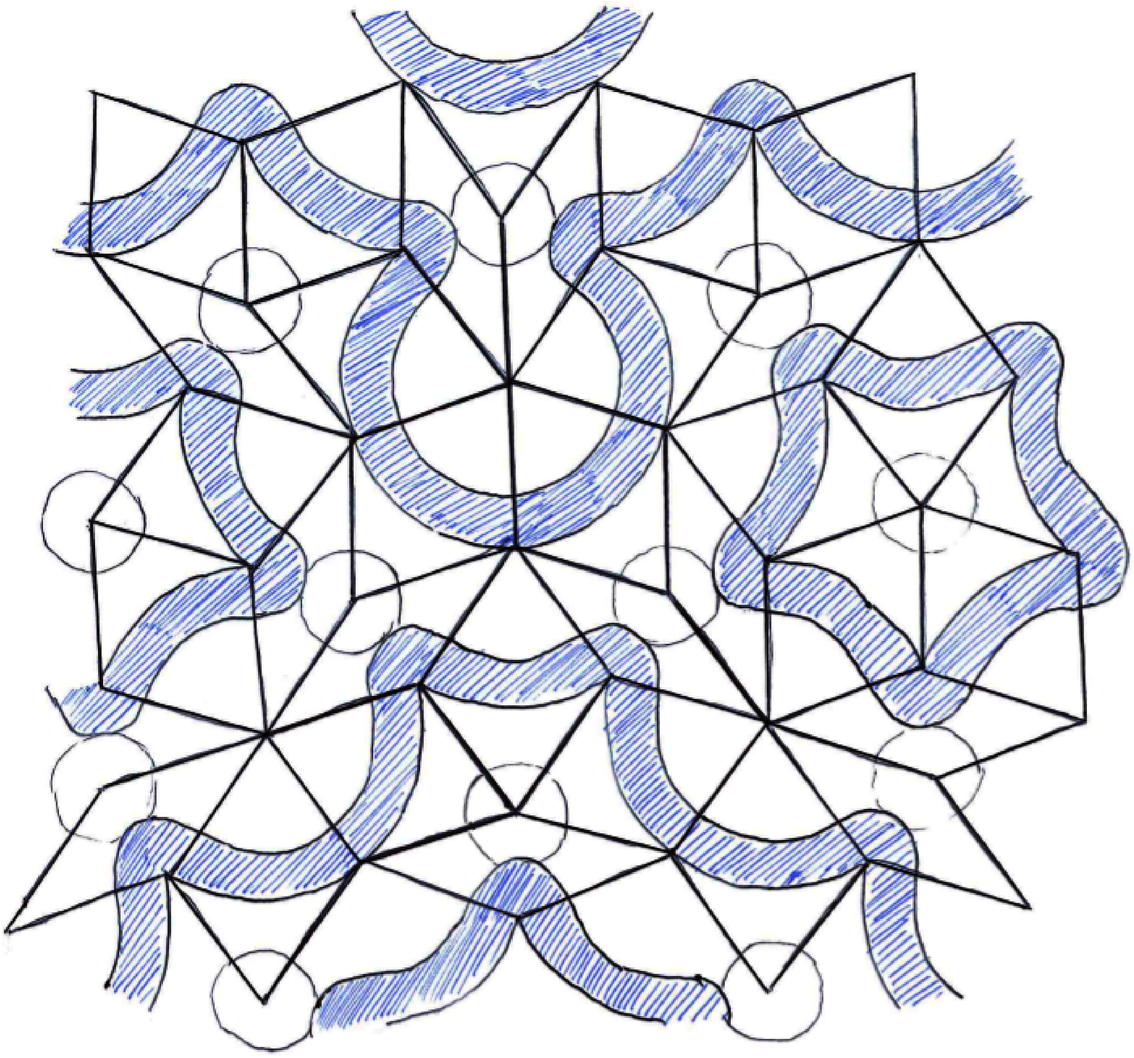

Date: This version of the notes was created on September 8, 2014.

Contents

1. What is a tiling?

1.1. Examples.

A *tiling* (or *tessellation*) of the plane by polygons is a covering of the plane by polygons, so that every point of the plane lies in some polygon, and the polygons do not overlap except possibly along their boundaries (on edges or vertices).

Example.

(1) *The tiling of the plane by unit squares:*

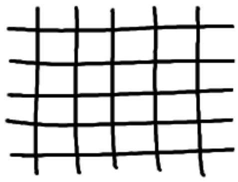

(2) *By decomposing the squares into triangles, the above turns into a tiling by isosceles triangles:*

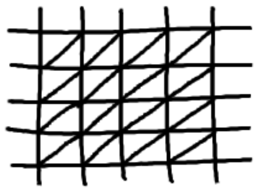

In the Exercises you show that the plane can be tiled by a regular n-sided polygon (with fixed side length) if and only if $n = 3, 4, 6$ (so regular triangles, squares or hexagons).

1.2. More complicated tilings.

Usually, we care about tilings using tiles from a certain finite collection of polygons. In the above example: by a particular square, or a particular triangle. One can also allow more general shapes than polygons, such as any subset of the plane that is homeomorphic to a disc.[1] For example, the following tiling is obtained from the tiling by squares after deforming the tiles, using two types of tiles.

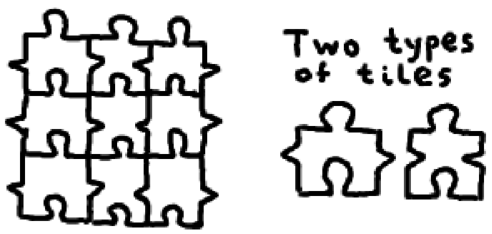

1.3. Markings and matching data.

In the above example, since drawing these squiggly edges is tiring, it is often more convenient to put *markings* on the edges, which tell you how the tiles must fit together.

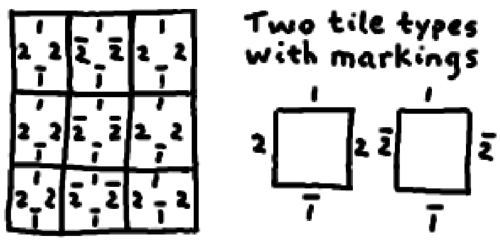

[1] So any region in the plane bounded by a closed curve with no self-intersections.

Above, we use labels 1, 2 which are required to always be glued onto $\overline{1}, \overline{2}$ respectively. You can also use arrows to prescribe in which direction you allow them to be glued.

Markings can also come in the form of artistic decorations on top of the tile types, and the matching condition is that the decorations fit together nicely. For example, the Penrose tiling using Penrose rhombi with decorations:

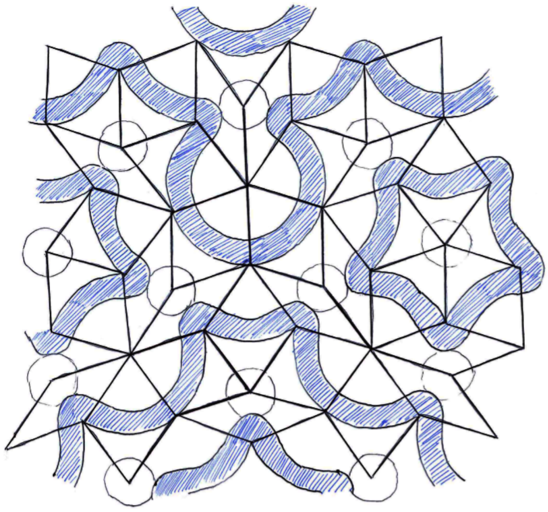

is made using the following two tile types with decorations:

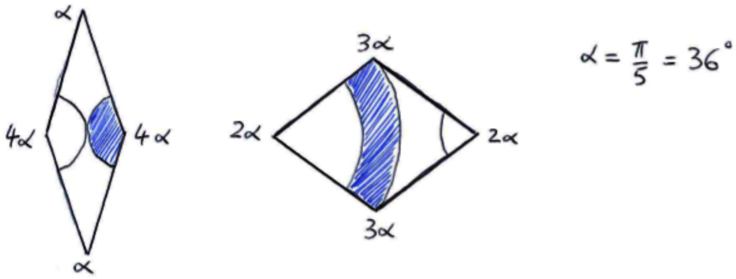

1.4. **Bad tiles.**

Regarding the general definition of tiling. One needs some care in saying exactly what a tiling is, because we want to avoid bad tiles (e.g. a tile made up of disconnected pieces, or a tile having holes, or tiles that have parts which become infinitesimally thin), we want to avoid tiles that are too large (e.g. unbounded tiles, like infinite strips), and nasty things can happen if we allow infinitely many tile types (e.g. you usually do not want there to be infinitely many tiles covering a finite region, for example this can happen if you use the collection of tile types given by squares of any side length).

Exercise 1. *Consider the infinite collection of tile types consisting of discs of any positive radius. Can you tile the whole plane using only copies of tiles taken from this collection?*

1.5. **Periodic tiling, non-periodic tiling, and aperiodic tile sets.**

A tiling is called *periodic* if the tiling has two translation[2] symmetries in two non-parallel directions.[3]

Example. Our favourite unit square tiling, pictured above, has translation symmetries $x \mapsto x+n, y \mapsto y+m$ *for* $n, m \in \mathbb{Z}$. *So, for example,* $x \mapsto x+1, y \mapsto y$

[2]A *translation* is the addition of constants to the coordinates: $x \mapsto x+a, y \mapsto y+b$ for some $a, b \in \mathbb{R}$. We assume that a, b cannot both be zero.

[3]Meaning, the translation constants (a, b), (a', b') are not proportional: $(a, b) \neq (ka', kb')$ for all $k \in \mathbb{R}$.

and $x \mapsto x, y \mapsto y+1$ are two non-parallel translations. So this tiling is, of course, periodic.

A tiling is called *non-periodic* if it has no translational symmetry.[4]

Example. Take the usual tiling by unit squares, divide all squares along one of the diagonals, except for one square, which you divide along the opposite diagonal. This gives a non-periodic tiling:

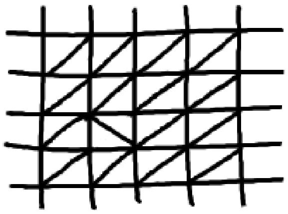

A set F of tiles is called *aperiodic* if every tiling of the plane using copies of tiles from F is always non-periodic. That is, no matter how you tile the plane with them, the tiling will never have any translational symmetry.

We will only be studying the simplest aperiodic tile sets, discovered by Penrose in 1974. However, historically, the first aperiodic set of tiles was discovered in 1966 by Robert Berger (20,426 tiles, later reduced to 104 tiles), and in 1971 Raphael M. Robinson found an aperiodic set of 6 tiles. Other aperiodic tile sets were discovered in the late 1970s by Robert Ammann.

1.6. **Penrose tiles.**

The Penrose Kite K and Dart D are defined by:

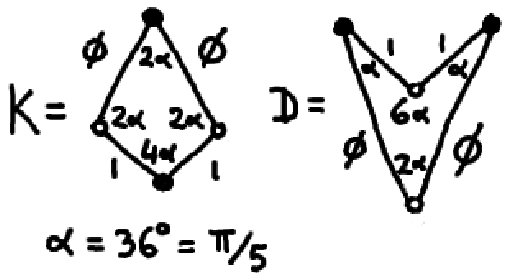

where ϕ is the *golden ratio* $\phi = \frac{1}{2}(1+\sqrt{5})$ (roughly 1.618). Notice K, D have markings (the vertices are coloured white or black).

Another collection of two tile types are the Penrose Rhombi (recall a rhombus is an equilateral parallelogram):

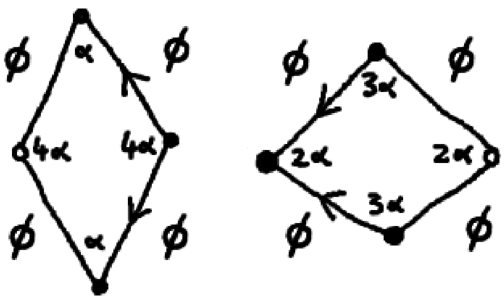

Notice these rhombi have markings (the vertices are coloured white or black, and the edges with equal vertex colours are marked with arrows pointing towards the smaller angle). Although here you could have picked any side-length (provided both rhombi have the same

[4]Notice: "not periodic" does not necessarily imply "non-periodic". There could be one translational symmetry.

side-length), it is convenient to use ϕ, since then there is a natural way to pass from a tiling by K, D to a tiling by such rhombi, and vice-versa.

A third collection of six Penrose tiles is made up of pentagons and pentacles:

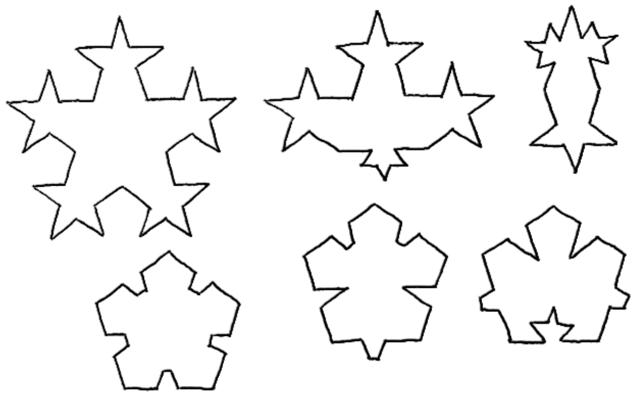

As usual, rather than these squiggly tiles, it is easier to work with tiles with markings:

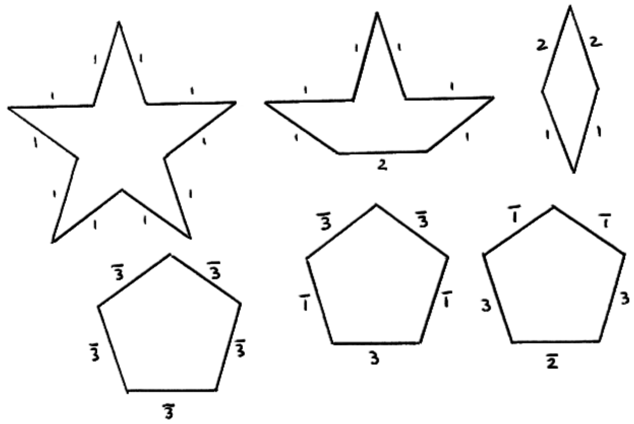

In the Exercises you will show that you can naturally pass from a tiling by pentagons/pentacles to a tiling by K, D.

1.7. **References on Penrose tilings.** There is very little published on Penrose tilings beyond survery papers which do not contain any proofs. The nicest surveys I know of are:

(1) Roger Penrose, *Pentaplexity*, Eureka 39, 1979.[5]
(2) Martin Gardner, *Penrose Tiles to Trapdoor Ciphers*, CUP, 1997.[6]

The only reference that I know of, which contains proofs, is the following book, but it seems to be out of print:

(3) Branko Grünbaum and Geoffrey C. Shephard, *Tilings and Patterns*, W.H. Freeman & Company, 1986. (Careful: the reprint paperback is allegedly abridged, missing chapters 8-12, of which 10 was the chapter on Penrose tilings)

2. The Extension Theorem

2.1. **How do you know that you can tile the whole plane?**

Some tilings, like our favourite tiling by squares above, have such symmetry that it is obvious that you can tile the entire plane even though you have only shown, with a drawing, how to build a tiling of a finite region. However, for tilings involving more complicated tiles,

[5]Reprint of *Pentaplexity: A class of non-periodic tilings of the plane* published in The Mathematical Intelligencer 2.

[6]The first chapter, on Penrose tilings, is a reprint of *Extraordinary non-periodic tiling that enriches the theory of tiles*, Scientific American 236, 110–121, January 1997.

even though you may be able to tile a very large region, if you don't see a repeating pattern then how do you know that you can tile the whole plane?

If you play with Penrose tiles, you will in fact notice that, very often indeed, when you try to build a larger and larger region you eventually get stuck because no other tile fits. This is because the tile that you choose to place in a certain spot may affect your chances of extending the tiling later on in some completely different spot.

More dramatically, Penrose tilings have a *non-locality property*, in the following sense. You have built a large patch of your tiling so far, and it turns out that on opposite sides of this large patch you have to make a choice of which tiles to place: one for each of the two opposite sides. It turns out, in some situations, that the tile that you place on one of the two opposite sides determines uniquely the tile that you must place on the opposite side – otherwise you will get stuck later on in a completely different spot! This suggests that locally, there is no recipe (algorithm) to decide which tile you must place, because you must also keep track of the global geometry of your tiling (namely, the other tile you picked on the opposite side, affects your chances of extending the tiling).

Exercise 2. *Can you write an algorithm for a computer which, given a finite tiling patch (by Penrose tiles), can tell you in finite time whether this patch can be indefinitely extended in some way to give a Penrose tiling of the entire plane?*[7]

So with the exception of some special Penrose tilings (the cartwheel tiling, the sun, the star), it seems to be rather difficult to build a Penrose tiling in practice. The extension theorem allows one to prove the existence of Penrose tilings, without building them explicitly.

2.2. What the extension theorem says, intuitively.

To avoid confusions, we will call *partial tiling* the tiling of a connected subset of the plane (rather than the whole plane).

The extension theorem says that if you use a finite set F of tile types, and if you can find a sequence of partial tilings $S_1, S_2, \ldots$ covering larger and larger discs $B_{r_1}(0), B_{r_2}(0), \ldots$ centred at zero (so the radii $r_n \to \infty$), then there exists a tiling of the whole plane.

Examples. *If the partial tilings* $S_1 \subset S_2 \subset S_3 \subset \cdots$ *are contained inside each other concentrically,*[8] *then obviously you get a tiling of the whole plane:* S_1 *will never be changed after step 1,* S_2 *will never be changed after step 2, and so on: so we are building the tiling inductively.*

The extension theorem does not require the S_i to be contained inside each other concentrically. In fact, the S_i may be completely unrelated to each other.

Examples.

(1) *Consider the following partial tilings by the unit square:*

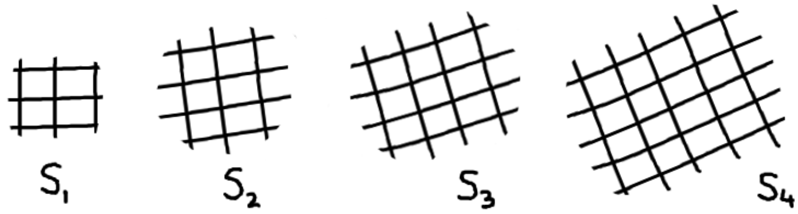

[7]Later, we will show that any finite region of a Penrose tiling has an identical copy in any other Penrose tiling. So you could ask the computer to check the finite patch against a known tiling T of the plane: so if you find a copy of the patch inside T, then you know that you can extend the patch using T! But this checking will not stop in finite time if there is no copy of the patch inside T.

[8]If these partial tilings are contained inside each other, but not concentrically, then it is not clear whether you can find a tiling of the whole plane. For example, the balls $B_1(1), B_2(2), B_3(3), \ldots$ are contained inside each other but their union is $\cup B_n(n) = \{(x, y) \in \mathbb{R}^2 : x > 0\}$, only half of the plane.

These are patches of our usual tiling by unit squares except tilted by increasing angles $\alpha_1, \alpha_2, \ldots$ which converge to 45 degrees. So none of the partial tilings $S_1, S_2, \ldots$ are subsets of each other, and the limit tiling (the usual unit square tiling rotated by 45 degrees) does not match up with any of the S_n.

(2) *Suppose instead the angles $\alpha_1, \alpha_2, \ldots$ are random angles:*

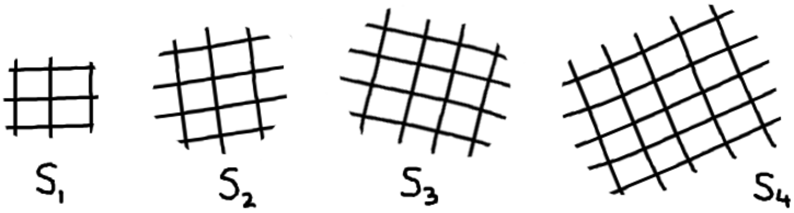

The sequence $S_1, S_2, \ldots$ will not converge. But there is some subsequence, $S_{i_1}, S_{i_2}, \ldots$, which will converge: just pick a subsequence of the angles $\alpha_{i_1}, \alpha_{i_2}, \ldots$ which converges (modulo $2\pi\mathbb{Z}$).

2.3. **The grasshopper trick.**

We will later refer to the following Exercise you had:

Exercise 3. *Consider a closed square S in $\mathbb{C}$. Suppose a grasshopper jumps around on the square. Its positions are $z_1, z_2, z_3, \ldots$.*
Chop up the square into 4 equal squares. Use the pigeonhole principle to show that for one of these four smaller squares, say S', the grasshopper is in S' infinitely many times (at positions, say, $z_{i_1}, z_{i_2}, \ldots$ for some $i_1 < i_2 < \cdots$).
Use this chopping idea inductively to show that, some subsequence of positions $z_{j_1}, z_{j_2}, \ldots$ of the grasshopper converges to some point z in the square S. In other words, the grasshopper jumps infinitely many times arbitrarily close to z.
Show in general, that for any closed bounded region $S \subset \mathbb{R}^n$, and any sequence of points $z_1, z_2, \ldots$ in S, there is a subsequence $z_{j_1}, z_{j_2}, \ldots$ which converges to some point z in S (z need not be unique).
Cultural remark. A (metric) space is called *compact* if any sequence of points has a subsequence which converges to a point in the space.

2.4. **The selection theorem.**

Here, and later, when we talk about convergence of tiles, we are using the Hausdorff distance. I will briefly recall the definition below, although you will notice later that we never really use this detailed definition. It is enough if you are happy with your own intuition of what it means for two copies of the same tile to be close or not.

Recall that the Hausdorff distance $d(S_1, S_2)$ between two closed bounded subsets $S_1, S_2 \subset \mathbb{R}^2$ of the plane is defined by[9]

$$d(S_1, S_2) = \min \left\{ \delta \geq 0 : \begin{array}{c} \text{every point } s_1 \in S_1 \text{ has } d(s_1, s_2) \leq \delta \text{ for some } s_2 \in S_2, \\ \text{and every point } s_2 \in S_2 \text{ has } d(s_1, s_2) \leq \delta \text{ for some } s_1 \in S_1 \end{array} \right\}$$

where the distance $d(s_1, s_2)$ between points s_1, s_2 refers to the Euclidean distance.

We say that tiles $S_1, S_2, \ldots$ converge to a tile S if $d(S_n, S) \to 0$ as $n \to \infty$. We abbreviate this by writing $S_1, S_2, \ldots \to S$.

Theorem 4 (Selection Theorem). *Suppose we are given:*

(1) *a point $p \in \mathbb{R}^2$;*
(2) *a tile T (any closed and bounded subset $T \subset \mathbb{R}^2$);*

[9]An equivalent definition is $d(S_1, S_2) = \max\{\max_{s_1 \in S_1} \min_{s_2 \in S_2} d(s_1, s_2), \max_{s_2 \in S_2} \min_{s_1 \in S_1} d(s_1, s_2)\}$.

(3) *a sequence of tiles* $S_1, S_2, \ldots \subset \mathbb{R}^2$ *in the plane each containing* p*, and each of which is a copy of* T *(so all the* S_j *are "the same" up to translation and rotation).*

Then some subsequence $S_{j_1}, S_{j_2}, \ldots$ *converges to a tile containing* p *which is a copy of* T*.*

Proof. First choose, once and for all, an identification $S_n \cong T$ for each $n = 1, 2, \ldots$. Via this identification, the point $p \in S_n$ gives a point $q_n \in T$.

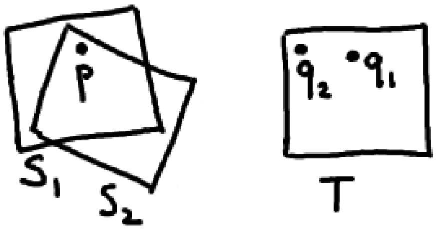

Now use the "grasshopper trick": you can pick a convergent subsequence $q_{i_1}, q_{i_2}, \ldots \to q$ inside T. Notice that if you place the tile T inside $\mathbb{R}^2$ so that $q \in T$ lies on top of $0 \in \mathbb{R}^2$, then your only freedom to place T is to rotate it about q. Thus each S_{i_n} differs from T by a small translation by $q - q_{i_n}$ and by some rotation angle α_{i_n}. Again, by the "grasshopper trick", you can pick a subsequence of the rotation angles $\alpha_{j_1}, \alpha_{j_2}, \ldots$ which converges (modulo $2\pi\mathbb{Z}$). Call α the limit. But now, observe this means that the tiles $S_{j_1}, S_{j_2}, \ldots$ converge to the tile S which is a copy of T rotated by α with $q \in T$ placed onto the point $p \in \mathbb{R}^2$. □

Corollary 5 (Selection Theorem)**.**

Suppose we are given a finite collection F *of tile types (each of which is a closed and bounded subset of* $\mathbb{R}^2$*). For example* $F = \{Kite, Dart\}$*.*

If $S_1, S_2, \ldots \subset \mathbb{R}^2$ *is any sequence of tiles in the plane, each of which is a copy of some tile from* F*, and each containing the point* $p \in \mathbb{R}^2$*, then there is a subsequence* $S_{j_1}, S_{j_2}, \ldots$ *converging to a tile* S *containing* p *which is a copy of some tile* $T \in F$*.*

Proof. Since there are only finitely many tiles in F, there must be some tile $T \in F$ which appears, as a copy, infinitely often in the sequence $S_1, S_2, \ldots$. So pass to a subsequence $S_{k_1}, S_{k_2}, \ldots$ all of which are a copy of T. Now apply the previous Theorem. □

Exercise 6. *Show that these results can fail if you remove some hypothesis. For example, if you do not assume that all tiles contain* p*, or if you do not assume that the tiles are bounded, or if you assume that* F *is an infinite collection of tiles.*

2.5. **The extension theorem.**

Theorem 7 (Extension Theorem)**.**

Suppose we use a finite collection F *of tile types each of which is a bounded closed subset of* $\mathbb{R}^2$ *which is homeomorphic to a disc.*[10] *If for any arbitrarily large disc you can find a partial tiling covering the disc, then there exists a tiling of the entire plane.*[11]

Proof.

Step 1. Pick $\delta =$ any positive radius such that each tile of F contains some disc of radius δ. Now let $m = \frac{\delta}{100}$. Consider the lattice $\mathbb{Z} \cdot m \times \mathbb{Z} \cdot m \subset \mathbb{R}^2$ of points whose coordinates are integer multiples of m. Notice that, by the choice of δ, any copy of a tile from F placed inside $\mathbb{R}^2$ must contain at least one lattice point in the interior.[12]

[10] Equivalently: any region in the plane bounded by a closed curve with no self-intersections. We need this condition to avoid bad tiles – can you see why?

[11] As usual, when using a collection F of tile types, the tiles we place in $\mathbb{R}^2$ are each obtained from a tile $T \subset \mathbb{R}^2$ of F by translation and rotation.

[12] Just ask yourself: since the tile contains a disc of radius δ, look at where in $\mathbb{R}^2$ the centre of that disc is, and ask yourself whether this disc contains points of the lattice.

Step 2. The lattice is a countable collection of points, meaning we can number the points $x_1, x_2, x_3, \ldots$. For example, we can do this explicitly as in the picture, by spiralling outwards starting from $x_1 = (0, 0)$:

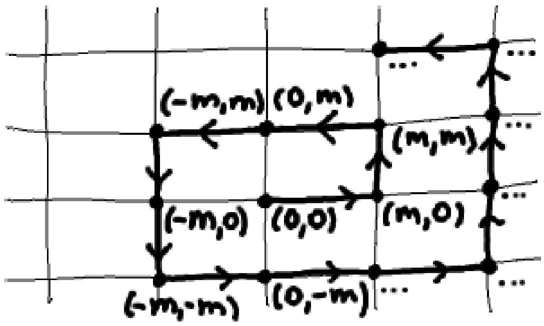

Step 3. By the assumption of the Theorem, we know that there exists some partial tiling T_r which covers any given disc $B_r(0)$.

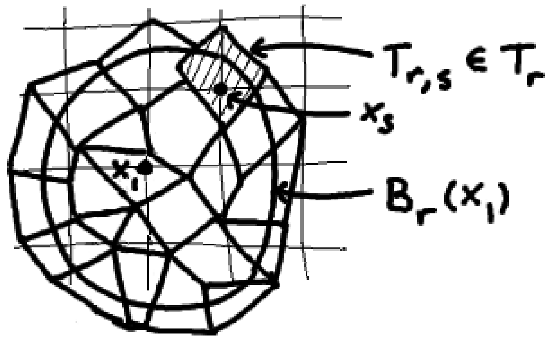

From this partial tiling T_r, there is some tile $T_{r,s}$ which contains the lattice point x_s. There may be more than one such tile, when x_s lies on the boundary of several tiles, but for the proof it does not matter which tile $T_{r,s}$ you pick.

Step 4. For any choice of radii $r_1, r_2, r_3, \ldots \to \infty$, you therefore obtain a sequence of tiles

$$T_{r_1,s},\ T_{r_2,s},\ T_{r_3,s},\ \ldots$$

which contain the point x_s.

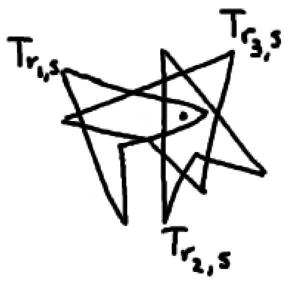

By the Selection Theorem (Corollary 5), there is a subsequence[13]

$$T_{r'_1,s},\ T_{r'_2,s},\ T_{r'_3,s},\ \ldots \to S$$

converging to a tile S containing x_s which is a copy of some tile from F.

[13]So, for some subset $\{r'_1, r'_2, r'_3, \ldots\} \subset \{r_1, r_2, r_3, \ldots\}$.

Step 5. Run Step 4 for the radii $1, 2, 3, \ldots$, for $s = 1$ (so for the tiles containing x_1) to obtain a subsequence $r'_1, r'_2, r'_3, \ldots$ of $1, 2, 3, \ldots$ so that:

$$T_{r'_1,1},\, T_{r'_2,1},\, T_{r'_3,1},\, \ldots \to S_1$$

where S_1 is a copy of a tile in F with $x_1 \in S_1$.

Now run Step 4 for the radii $r'_1, r'_2, r'_3, \ldots$, for $s = 2$ (so for the tiles containing x_2) to obtain a subsequence $r''_1, r''_2, r''_3, \ldots$ of $r'_1, r'_2, r'_3, \ldots$ so that:

$$T_{r''_1,2},\, T_{r''_2,2},\, T_{r''_3,2},\, \ldots \to S_2$$

where S_2 is a copy of a tile in F with $x_2 \in S_2$.

Now run Step 4 for the radii $r''_1, r''_2, r''_3, \ldots$, for $s = 3$ (so for the tiles containing x_3) to obtain:

$$T_{r'''_1,3},\, T_{r'''_2,3},\, T_{r'''_3,3},\, \ldots \to S_3$$

where S_3 is a copy of a tile in F with $x_3 \in S_3$.

Keep going inductively. In general, if we abbreviate $r_j^{(n)}$ to mean that we put n dashes $'$ on r_j, once we have found $r_1^{(n)}, r_2^{(n)}, \ldots$, we run Step 4 for the radii $r_1^{(n)}, r_2^{(n)}, r_3^{(n)}, \ldots$, for $s = n+1$ (so for the tiles containing x_{n+1}) to obtain a subsequence:

$$T_{r_1^{(n+1)},n+1},\, T_{r_2^{(n+1)},n+1},\, T_{r_3^{(n+1)},n+1},\, \ldots \to S_{n+1}$$

where S_{n+1} is a copy of a tile in F with $x_{n+1} \in S_{n+1}$.

Step 6. *Useful remark. If a sequence $z_1, z_2, \ldots \to z$ converges, then any subsequence also converges to the same limit: $z_{i_1}, z_{i_2}, \ldots \to z$.*

This is simply because if *all* points z_n approach z, then of course also *some* points z_{i_n} approach z (if $d(z_n, z) \to 0$ as $n \to \infty$, then also $d(z_{i_n}, z) \to 0$ as $n \to \infty$).

By construction S_i is the limit

$$S_i = \lim_{n\to\infty} T_{r_n^{(i)},i}$$

Therefore it is also the limit of any subsequence of $T_{r_n^{(i)},i}$, such as the subsequence $T_{r_n^{(i+1)},i}$, or $T_{r_n^{(i+2)},i}$, or generally $T_{r_n^{(k)},i}$ for $k \geq i$. So:

$$S_i = \lim_{n\to\infty} T_{r_n^{(k)},i} \qquad \text{(for any } k \geq i\text{)}.$$

Step 7. We claim that $S_1, S_2, S_3, \ldots$ defines a tiling of the entire plane!

First, to clarify, these tiles $S_j \subset \mathbb{R}^2$ need not be distinct: some of them may be repeated in the list. To show that these really tile the entire plane, we need to check two things:

(1) two tiles S_i, S_j only overlap along the boundary (unless they are equal $S_i = S_j$);
(2) any point $p \in \mathbb{R}^2$ lies in some tile S_i (possibly several tiles).

Let's prove (1) first. Suppose S_i, S_j overlap on their interiors. Then also the two tiles $T_{r_n^{(k)},i}, T_{r_n^{(k)},j}$ which converge to S_i, S_j must overlap on their interiors for large n (here we use that the tiles are homeomorphic to discs, and we use Step 6 taking $k \geq \max\{i, j\}$). But those two tiles belong to the partial tiling $T_{r_n^{(k)}}$, so they cannot overlap unless they equal. So $T_{r_n^{(k)},i} = T_{r_n^{(k)},j}$. So taking the limit as $n \to \infty$, we deduce that $S_i = S_j$. This proves (1).

Now we prove (2). Suppose $p \in \mathbb{R}^2$ does not belong to any of the S_i, by contradiction.

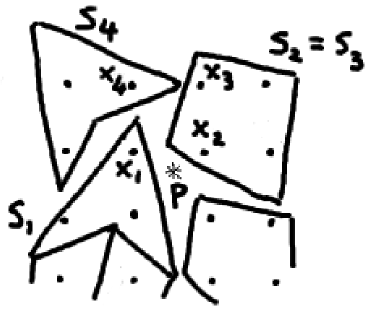

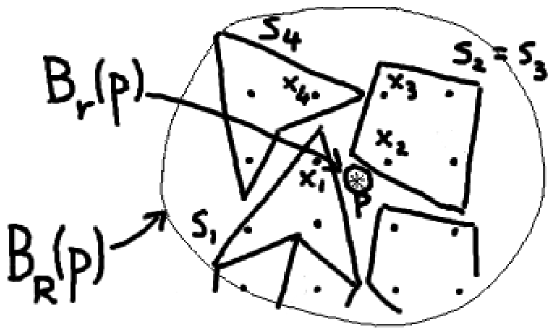

Let's consider a large disc $B_R(p)$ around p, so large that any copy of a tile from F which touches p must lie entirely inside this ball.[14] There are only finitely[15] many tiles S_i which lie inside or intersect this disc. The union of these finitely many tiles S_i is a closed set, so their complement is open. Therefore, since p lies in this open set, there is also a small ball $B_r(p)$ around p lying inside this open set. So all S_i are at least a distance r away from p. Therefore[16] the tiles $T_{r_n^{(k)},i}$ which converge to S_i are at least $\frac{r}{2}$ away from p for large n. For large n, the partial tiling $T_{r_n^{(k)}}$ will cover $B_R(p)$, so some tile X of $T_{r_n^{(k)}}$ must contain p. But X must contain a lattice point in the interior, say x_i. But then $X = T_{r_n^{(k)},i}$ (there can be only one tile of the partial tiling which contains x_i in the interior, since tiles must not overlap). So $T_{r_n^{(k)},i}$ contains p, contradicting that it should be $\frac{r}{2}$ away from p. □

3. Penrose Kites and Darts

3.1. **The Kite and the Dart.**

The Kite K and the Dart D are defined as follows:

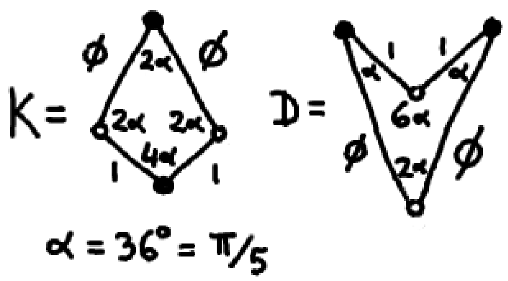

Here ϕ, roughly 1.618, is the *golden ratio* defined by

$$\phi = \frac{1+\sqrt{5}}{2}.$$

We emphasize that part of the data of the Kite and Dart is the colouring of the vertices – black or white. These markings impose a *matching condition*: when you tile with K, D you are required to place the tiles so that if two tiles touch at a point which is a black vertex of one tile, then that point must also be a black vertex of the other tile (similarly, white vertices must match when placing tiles).

[14]Notice, here we use again that there are finitely many tiles in F, and that each tile is bounded.

[15]This uses the assumption that the tiles are bounded, and the fact that each tile contains a disc of radius δ and thus occupies at least area $\pi\delta^2$.

[16]More figuratively: you are driving too fast on a road, and there is a lamppost on the left of the road, and a lamppost on the right of the road, and you crash into the lamppost on the left. Once crashed, you are far away from the lamppost on the right. Then just before the crash you were also far away from the lamppost on the right! In this metaphor, p is the lamppost on the right, S_i is the crashed car in the left lamppost, and $T_{r_n^{(k)},i}$ is the position of the car just before the crash.

Without these markings, the tiles K, D would sometimes[17] give rise to periodic tilings, which we do not want (we want the tile set $F = \{K, D\}$ to be aperiodic, which we will prove later).

We mentioned in Section 1.3 how markings and matching conditions work with tilings. One can draw two honest tiles without markings by replacing the edges above with squiggly lines so that the tiles can only fit together in the way that we want (i.e. as they would when we have straight edges and vertex colour markings).

Exercise 8. *Try drawing the two tiles with squiggly edges so that the matching conditions are automatically imposed.*

3.2. **The golden ratio.**

The golden ratio $\phi = \frac{1}{2}(1+\sqrt{5})$ appears in many situations in geometry, for example in the construction of the regular unit pentagon:

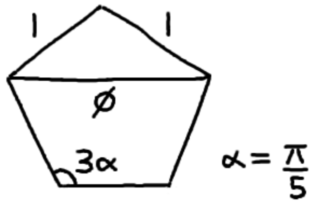

Notice that ϕ is a solution of the equation

$$\phi^2 - \phi - 1 = 0.$$

From this equation, we obtain the useful equalities: $1+\phi = \phi^2$ and $1 + 2\phi = \phi^3$ (since $\phi^3 = \phi\phi^2 = \phi(1+\phi) = \phi + \phi^2 = \phi + (1+\phi) = 1 + 2\phi$).

Exercise 9 (Fibonacci numbers)**.** *The numbers* $1, 1, 2, 3, 5, 8, 13, 21, \ldots$ *are the* Fibonacci numbers*, where each number is obtained by adding the previous two numbers. Call these numbers* F_n*. The defining equation is* $F_{n+2} = F_{n+1} + F_n$ *with initial conditions* $F_1 = F_2 = 1$*. Show inductively that from the equation defining* ϕ *you also obtain the equations*

$$F_n + F_{n+1} \cdot \phi = \phi^{n+1}.$$

The golden ratio has many beautiful properties, for example:

$$\phi = 1 + \cfrac{1}{1 + \frac{1}{1+\frac{1}{1+\cdots}}}$$

Exercise 10. *Show that the continued fraction on the right in the above equality really does*[18] *converge to* ϕ*.*

Look at the fractions you obtain when approximating the above continued fraction:

$$\begin{aligned}
1 &= \tfrac{1}{1} \\
1 + \tfrac{1}{1} &= \tfrac{2}{1} \\
1 + \tfrac{1}{1+1} &= \tfrac{3}{2} \\
1 + \tfrac{1}{1+\frac{1}{1+1}} &= \tfrac{5}{3} \\
1 + \tfrac{1}{1+\frac{1}{1+\frac{1}{1+1}}} &= \tfrac{8}{5}
\end{aligned}$$

[17]Insert a K into the D, so that the 4α and 6α add up to a full angle 2π. This is a parallelogram. In the Exercises, you will prove that any parallelogram can tile the plane periodically. Notice, on the other hand, that the markings do not allow you to insert a K into D in that way!

[18]*Hint. Suppose it does converge to something, say* x*. Then can you find an equation that* x *must satisfy given the self-similarity of the continued fraction expression?*

Exercise 11. *Do you see a pattern above? Show by induction on n that the n-th fraction you obtain above (approximating the continued fraction of ϕ) equals $\frac{F_{n+1}}{F_n}$. Deduce that*

$$\phi = \lim_{n\to\infty} \frac{F_{n+1}}{F_n}.$$

4. Composition, Decomposition, and Inflation

4.1. Subdivision trick: from K, D to A, B.

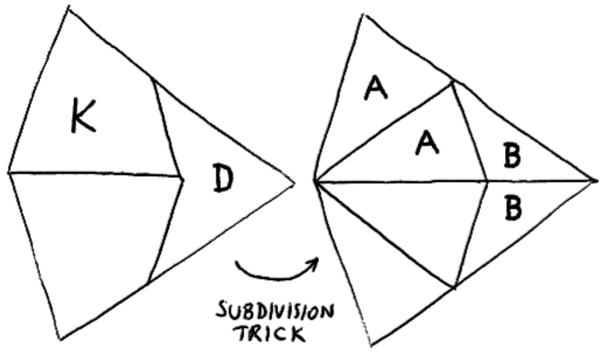

The tiles K, D can be subdivided along the axis of symmetry to give isosceles triangles A, B:

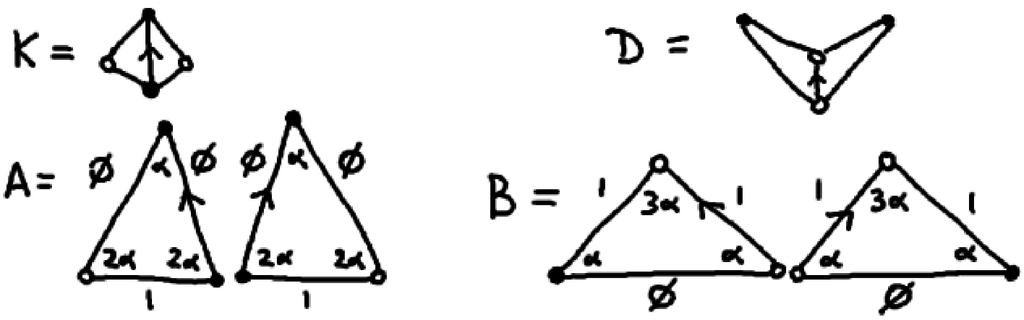

Although we write two letters A and B, there are actually four tiles: two copies of A and two copies of B. The two copies are different because they have different vertex markings (you would need a reflection to obtain one from the other). So we really should talk about tile types A_1, A_2, B_1, B_2, distinguishing the left-half and the right-half of the subdivision of each of K, D, but for sake of brevity we will just say A, B.

It is clear that there is a unique way to go from K, D to As and Bs by subdividing as shown. However, we also want there to be a unique way to get from a tiling by A, B tiles to a tiling by K, D. The problem is, for example, that two As may be glued along the edge having black vertex colours in the opposite way as that which arises from subdiving K.

For this reason, we put an additional marking on the A, B: we put an arrow pointing to the peak of the isosceles triangle on the edge which has equal vertex colours. In any tiling by A, B we then require the additional matching condition that the arrow directions must match when two tiles share an edge with arrows.

With the arrows and vertex colours it is now obvious that there is a unique way to go from tilings by A, B to tilings by K, D, and vice-versa.

4.2. Composition trick 1: from A, B to $A, \phi B'$.

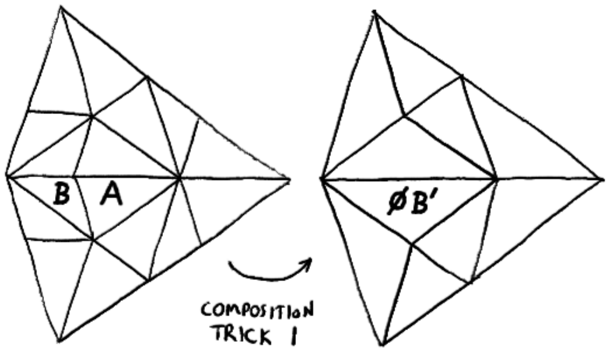

Now, notice that if you have a tile B, then there must always be a tile A attached to B along the shortest edge (of length 1) as follows:[19]

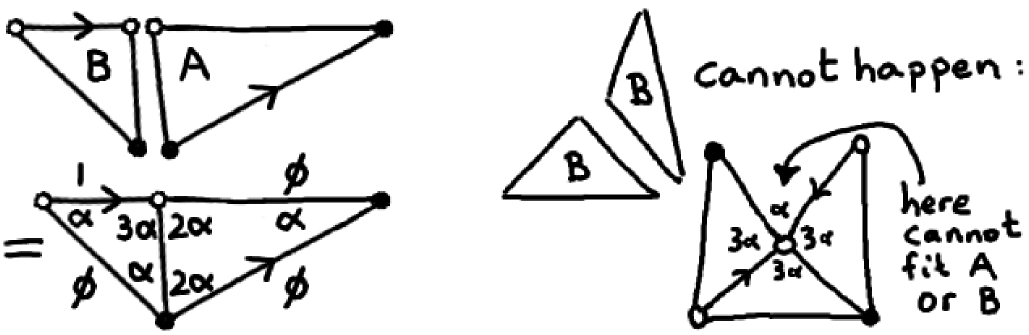

Of course, such an A is not glued onto another B, since there is a unique shortest edge in A. Thus, every such glued $B + A$ can be replaced by the resulting tile, which we call $\phi B'$:

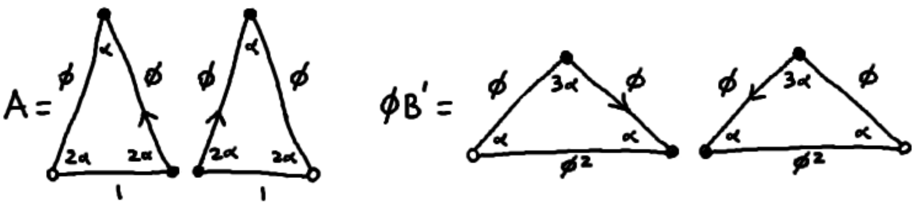

We call it $\phi B'$ because it is obtained from B by rescaling by ϕ and by switching colours and reversing arrows (so the arrow now points away from the peak of the isosceles triangle, along the edge with equal vertex colours). We will later continue to use the dash symbol $'$ to mean: switch the colours white$\leftrightarrow$black and reverse the directions of arrows.

Applying this composition rule, all Bs will have disappeared and we will only be left with $A, \phi B'$ tiles. So we can uniquely pass from a tiling by A, B to a tiling by $A, \phi B'$ tiles.

Vice-versa, a tiling by $A, \phi B'$ determines uniquely a tiling by A, B tiles: you just need to subdivide $\phi B'$ as in the above picture to replace each $\phi B'$ with two glued tiles A, B. Again, notice there is a *unique* way to make this subdivision.

4.3. Composition trick 2: from $A, \phi B'$ to $\phi A', \phi B'$.

[19]Notice, on the right of the picture, that if you try to attach two copies of B along the short edge, then you get stuck: either the vertex colours fail, or the angles around a vertex add up to more than 2π.

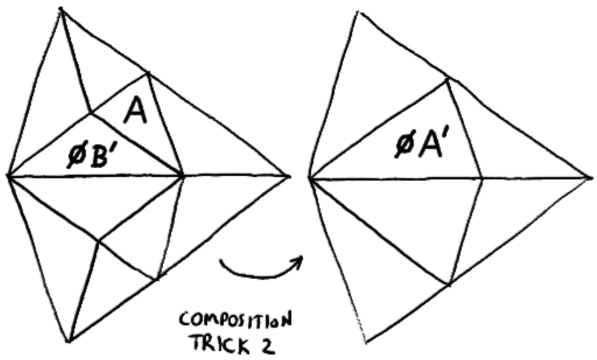

Now, notice that if you have a tile A, then there must always be a tile $\phi B'$ attached to A along the edge which has an arrow (of length ϕ) as follows:[20]

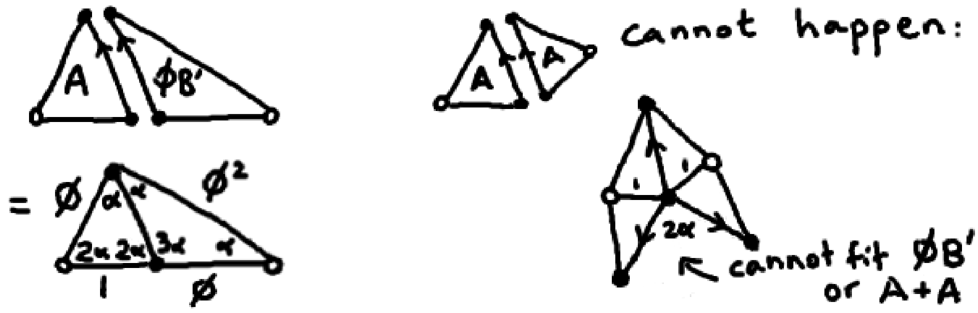

Thus, this composition gets rid of all As, by replacing them with $\phi A'$ (so A rescaled by ϕ with reversed colours/arrows). So from any tiling by $A, \phi B'$ we obtain uniquely a tiling by $\phi A', \phi B'$ tiles.

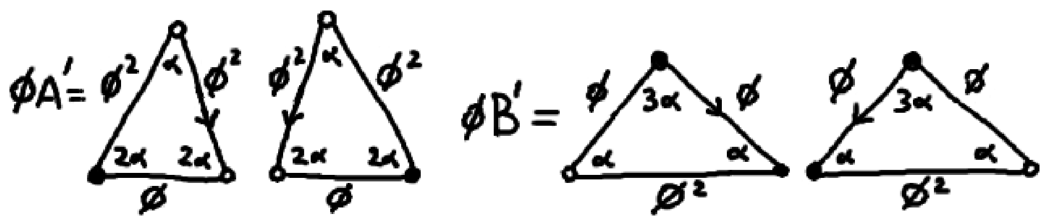

Vice-versa, a tiling by $\phi A', \phi B'$ determines uniquely a tiling by $A, \phi B'$ tiles: simply subdivide $\phi A'$ in the unique way to give rise to an A and a $\phi B'$ tile.

4.4. Composition.

Applying the subdivision trick and composition tricks 1, 2 changes tile types as follows:

$$K, D \implies A, B \implies A, \phi B' \implies \phi A', \phi B'.$$

We now switch colours/arrows, and then apply the reverse of the subdivision trick (erase the edges with arrows):

$$\phi A', \phi B' \implies \phi A, \phi B \implies \phi K, \phi D.$$

Notice that we end up with larger tiles: the Kites and Darts are rescaled by ϕ.

Thus: from a tiling by K, D we have uniquely obtained a tiling by $\phi K, \phi D$. I will call this *Composition.*

[20]Notice, on the right of the picture, that if you try to attach two copies of A along the edge with the arrow, then eventually you get stuck. Since only the A tiles have an edge of length 1, you would in fact need four A tiles attached, but then you get stuck: the vertex colours or the arrows fail, or the sum of the angles around a vertex add up to more than 2π.

Remark 12. Double-composition *consists of: subdivision, composition tricks 1, 2, 1, 2, and finally reverse subdivision:*

$$K, D \implies A, B \implies A, \phi B' \implies \phi A', \phi B' \implies \phi A', \phi^2 B \implies \phi^2 A, \phi^2 B \implies \phi^2 K, \phi^2 D,$$

where we used that $A'' = A$, $B'' = B$ (reversing twice does nothing). This is often convenient, as it avoids the unnatural step of reversing colours/arrows.

4.5. **Decomposition.**

Since each individual step in the *Composition* procedure can be uniquely reversed, we can also carry out all steps in reverse. In reverse, we pass from a tiling by K, D to a tiling by $\frac{1}{\phi}K, \frac{1}{\phi}D$ tiles. I will call *Decomposition* the reverse of *Composition*.

Example. Below are the steps of the Decomposition *applied to an Ace (two Kites and a Dart):*

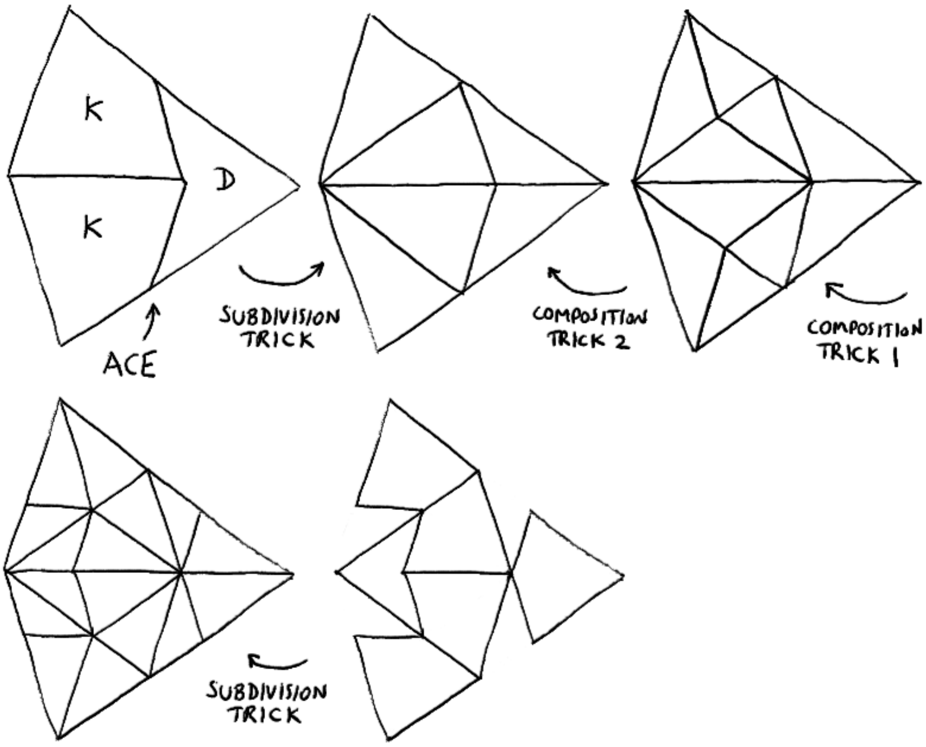

The last move above is to erase the edges with arrows. Near the boundary, there are half-tiles, and it is customary when working with partial tilings to complete those tiles. In the above case, we obtain Batman (The Dark Knight Falls):

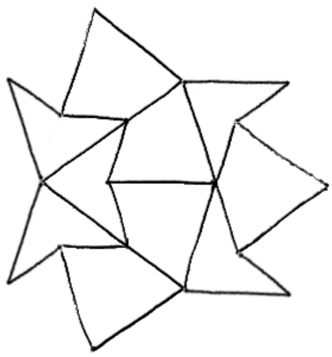

Notice that both *Composition* and *Decomposition*, in practice, just means taking a black marker and drawing a tiling on top of your given tiling. In the *Composition* process, you will draw bigger tiles (rescaled by ϕ, which is approximately 1.618). In the *Decomposition* process, you will draw smaller tiles (rescaled by $1/\phi$, approximately 0.618). In both cases, if your original tiling was only a partial tiling, covering some region R, then the decomposed/composed partial tilings will also cover the same region R.

4.6. **Inflation.**

If we apply *Decomposition* and then rescale by ϕ, we obtain a partial tiling by tiles of the same size as originally, but covering a larger region (namely ϕ times larger). I will call *Inflation* this procedure: *Decompose* and then rescale by ϕ (equivalently: first rescale by ϕ then *Decompose*). Inflating a partial tiling by K, D covering the ball $B_r(0)$ will produce a partial tiling by K, D of the ball $B_{\phi r}(0)$. Summarizing:

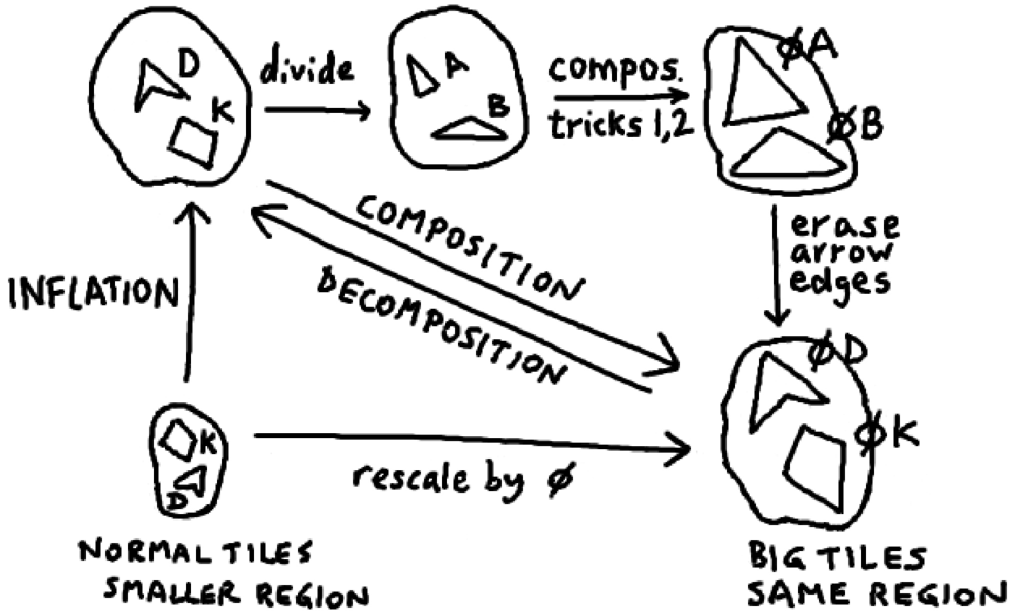

Example. In the Decomposition of the previous example, rescaling the final Batman gives:

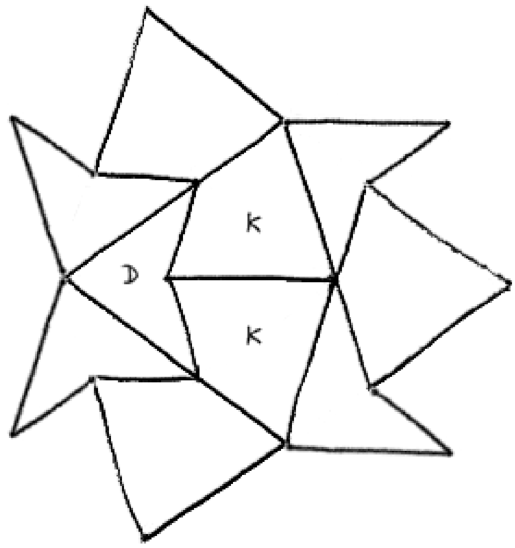

Notice that Batman has an Ace up his sleeve (but the Ace is flipped).
The tiles K, D at the end are the same size as originally, but the region has expanded by a factor of ϕ (roughly 1.618).

5. Existence of a Penrose tiling

5.1. The existence of a tiling of the entire plane by K, D.

So far, we never proved that the plane can be tiled by K, D. We will use the extension theorem. Start with any partial tiling by K, D covering some small disc $B_r(0)$ around 0. By applying *Inflation* n times, you obtain a partial tiling by K, D covering the disc $B_{\phi^n r}(0)$. Thus, we can obtain partial tilings covering arbitrarily large discs. Notice that, as far as we know, none of these partial tilings have anything to do with each other! But the extension theorem doesn't care: it will produce for us, from a subsequence of these inflated partial tilings, a limit tiling of the entire plane![21]

Philosophically speaking, we have not actually in practice found any tiling of the entire plane by Kites and Darts. We only know that one exists.

In one of the Exercises, you will be asked to build a tiling by K, D which have a pentagonal[22] symmetry about the origin, and to show that in fact there are exactly two such tilings (called the *sun tiling* and the *star tiling*). So you will know how to build two tilings in practice.

Exercise 13 (Hard)**.** *How many possible tilings of the plane by K, D are there? More precisely, you do not want to distinguish two such tilings if one can be obtained from the other by translation and rotation. So, the question is: can you classify all tilings by K, D up to congruency?*

The above exercise is hard because, as we will see later, *any* finite region, no matter how large, of a given tiling by K, D must also arise (in fact infinitely often!) in any other tiling by K, D. So far, for all we know, the extension theorem may always give the same tiling up to congruency. Also so far, for all we know, maybe not all tilings by K, D arise by repeated inflation via the extension theorem.

6. Penrose Rhombi

From a tiling by $A, \phi B'$ tiles (obtained from a K, D tiling after the subdivision trick and composition trick 1), we can obtain a tiling by two types of rhombi:

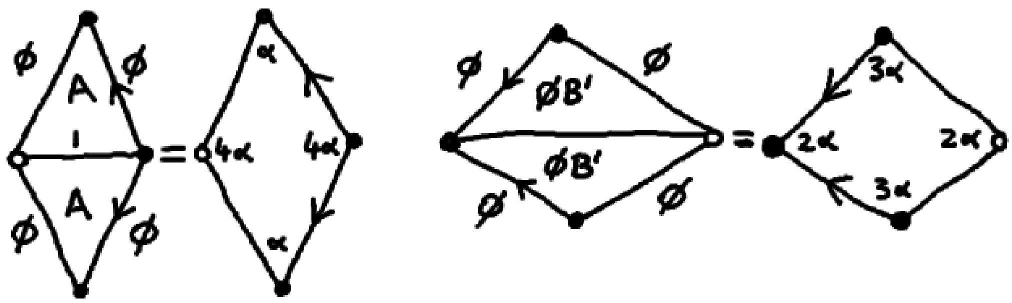

Simply glue the A tiles along the shortest edges that appear (of length 1), and glue the $\phi B'$ tiles along the longest edges that apper (of length ϕ^2).

As usual, there are markings on the tiles (vertex colours, and arrows on the two edges which have equal vertex colours). We remark that despite the markings, there really are just two rhombi, not four.[23]

[21] In fact, so far we don't even know whether a large partial tiling obtained by *Inflation* can actually be continued to give a tiling of the whole plane.

[22] i.e. rotating the tiling about 0 by integer multiples of the angle $2\alpha = 2\pi/5$ will give you back the exact same tiling.

[23] For example, there are two ways to glue two copies of A, depending on whether you put the white vertex on the left or on the right in the above picture. But the resulting two rhombi are related by a rotation by π, so you only obtain one marked rhombus from gluing two copies of A. Similarly, when gluing two $\phi B'$.

There is a unique way to glue As or $\phi B'$s to obtain rhombi. Conversely, there is a unique way to subdivide the two rhombi to obtain As and $\phi B'$s: simply subdivide the rhombi along the diagonals of length 1 and ϕ^2 respectively. Thus there is a unique way to pass from tilings by K, D to tilings by Penrose Rhombi, and vice-versa.

7. Proof that the Penrose tiles are aperiodic

Recall from Section 1.5 that a set F of tiles is called *aperiodic* if every tiling of the plane using copies of tiles from F is always non-periodic (that is, it does not have any translational symmetry).

Theorem 14. *The Penrose tiles K, D are aperiodic.*
Remark. the same holds for the tiles A, B, and the Penrose rhombi, as you can always pass from these tilings uniquely to a tiling by K, D and vice-versa.

Proof. Suppose not, by contradiction. (A mathematician's finest weapon!)

Then you know there exists a tiling of the plane by K, D which has a translational symmetry.

So consider a large patch of the tiling. When you move this first patch by the translation symmetry, then you get a second patch. The two patches must agree on the overlap, since that translation is a symmetry.

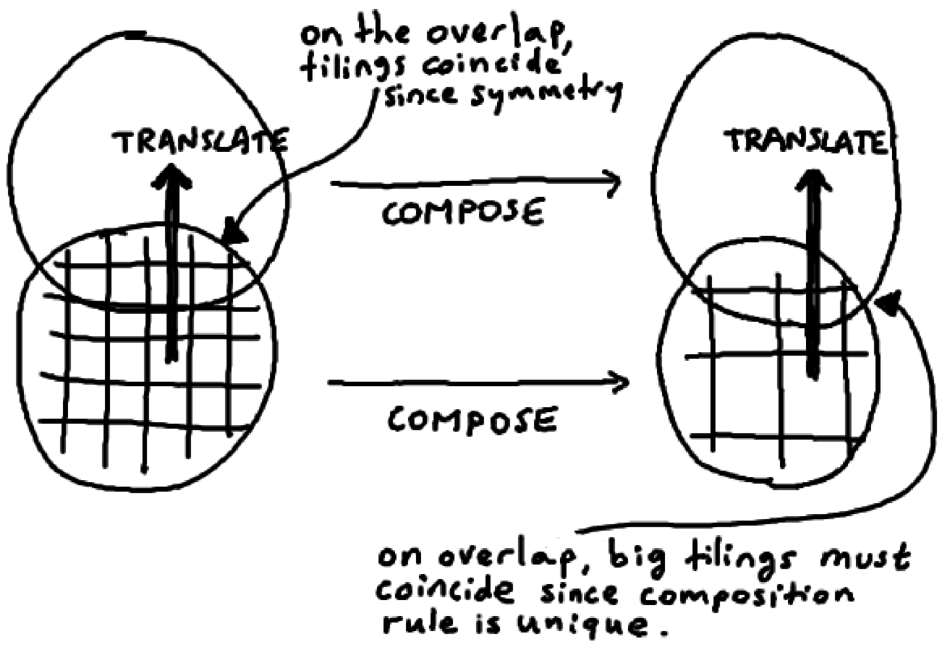

Now, apply *Composition* to the first patch (recall, in practice, this just means that you draw some thick dark lines on top of your existing patch, which gives a tiling over the same region but using bigger tiles).

Next, apply *Composition* to the second patch.

The two patches you obtain after *Composition*, must also agree on the overlap. Why? Because the *Composition* involves *unique* rules, so since the original two patches agree on the overlap, also the composed patches must agree on the overlap (you have unique rules on how to compose the tiles in the original overlap of patches).

Upshot: the two patches obtained after *Composition* respect the same translational symmetry.[24]

[24]Using the *same* constants of translation, we do *not* rescale the constants of translation!

Therefore, *Composition* applied to the entire tiling of the plane by K, D gives a new tiling, involving tiles which are larger (the new tiles are $\phi^2 K, \phi^2 D$ instead of K, D), but nevertheless we still have the same translational symmetry, say $z \mapsto z + c$.

Thus, applying *Composition* many times, say N times, we obtain a tiling by $\phi^{2N} K, \phi^{2N} D$ with the same translational symmetry $z \mapsto z + c$. But, for large N, this is absurd:

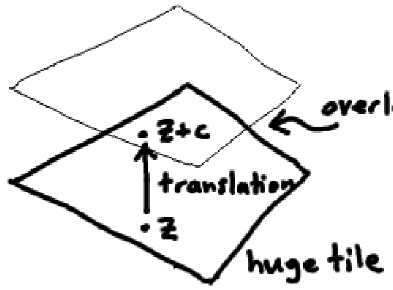

For N very large compared to c (for example, $N > 1000 \cdot c$ works), the translation will move a given tile $\phi^{2N} K$ to a new tile which overlaps with the given tile: but this contradicts that the translation is a symmetry.[25] (and similarly for a given tile $\phi^{2N} D$). Contradiction. □

So why does the above argument not work for tilings by unit squares (with vertices at the lattice points $\mathbb{Z} \times \mathbb{Z}$)? The issue is that the rule for *Compositions* is not unique in this case. For example, suppose we want our *Composition* rule to double the side lengths of squares. Then there are four ways of *Composing*, depending on whether you want the new lattice of composed squares to be $2\mathbb{Z} \times 2\mathbb{Z}$, $(2\mathbb{Z}+1) \times 2\mathbb{Z}$, $2\mathbb{Z} \times (2\mathbb{Z}+1)$ or $(2\mathbb{Z}+1) \times (2\mathbb{Z}+1)$. In the above proof, the translation symmetry $z \mapsto z + (0, 1)$ of the original tiling by unit squares will not give a translation symmetry of the composed tiling. Instead, the translation will send one choice of *Composition* of squares to another choice of *Composition* of squares, and in the case of $z \mapsto z + (0, 1)$ these are two different choices:

8. The cartwheel tiling

8.1. **The cartwheels C_n and the cartwheel tiling C.**

Take an Ace as in the Example of Section 4.5, and apply inflation several times. This yields a sequence of partial tilings: C_0 (Ace), C_1 (Batman), C_2, ...

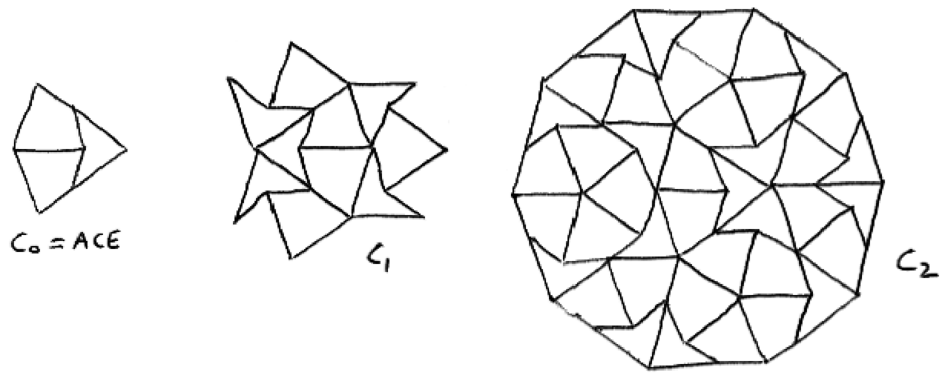

[25] a symmetry sends a tile to a tile, and in a tiling the tiles are not allowed to overlap.

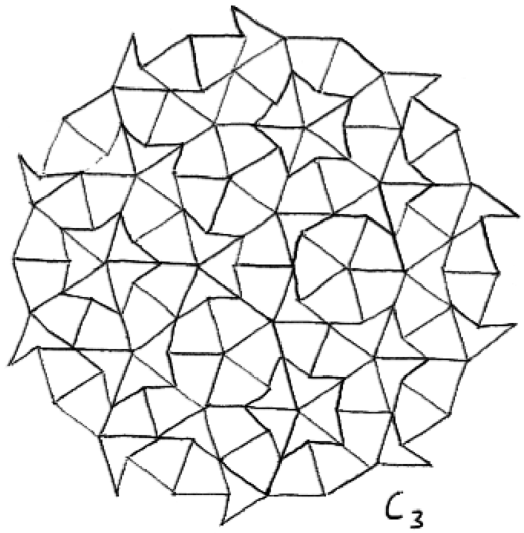

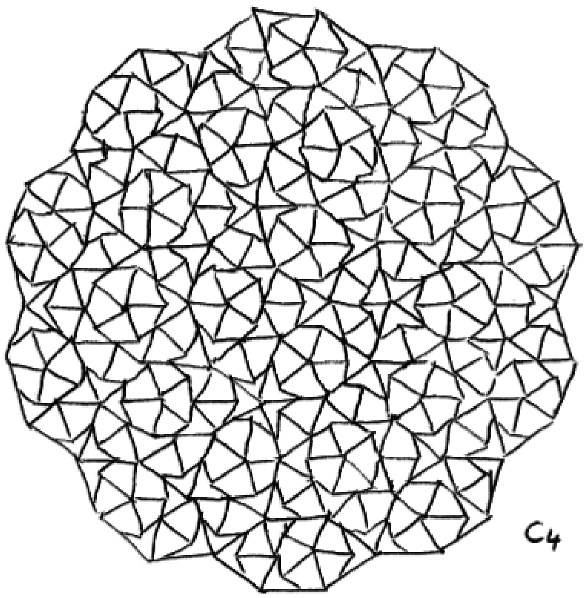

In the Exercises, you will prove that $C_0 \subset C_2 \subset C_4 \subset \cdots$ are contained inside each other concentrically. The odd $C_1, C_3, \ldots$ also appear inside $C_2, C_4, \ldots$ respectively, but they are flipped (reflected). In the picture below, on the left you can see the Bat-Signal C_2 reflected in the Gotham Skyline C_3 (the picture also highlights Batman C_1 and his flipped Ace C_0).

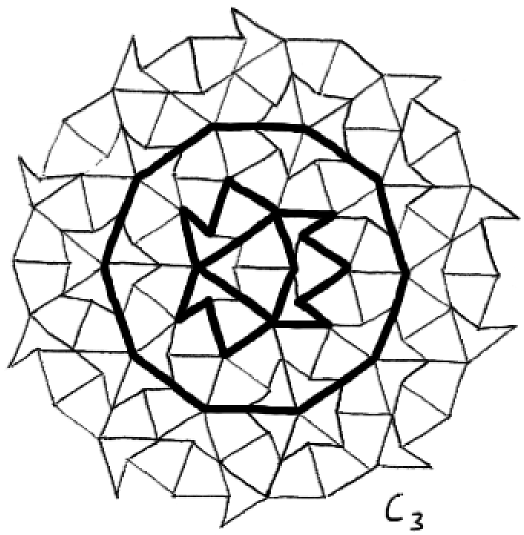

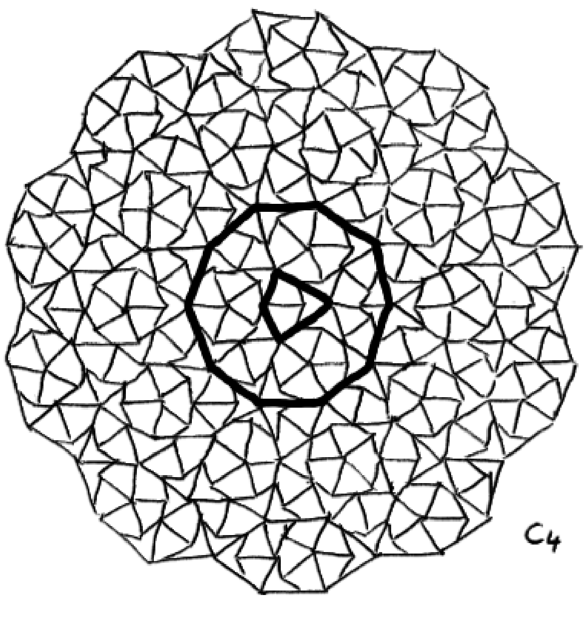

In the picture on the right, you can see $C_0 \subset C_2 \subset C_4$ concentrically.

The limit

$$C = \lim_{n \to \infty} C_{2n}$$

is called the *carthweel tiling* (the limit makes sense since in general each C_{2n} is contained concentrically in the next C_{2n+2}, as you will prove in the Exercises).

Remark. The Extension Theorem applied to the sequence $C_0, C_1, C_2, \ldots$ would have spat out one of two types of convergent subsequences: those which eventually will only involve even cartwheels which converge to C, or those which eventually will only involve odd cartwheels which converge to a reflection of C.

8.2. **Every tile sits inside arbitrarily large cartwheels.**

Theorem 15. *Given any tiling by K, D, every tile lies inside a copy of the cartwheel C_{2n} for each n.*

Proof. **Step 1. Given any tiling by K, D, every point $p \in \mathbb{R}^2$ lies inside an Ace.**

Proof. First suppose $p \in D$. Look at the picture of D: inside the slot where the 6α angle is, you cannot put a copy/copies of D since the vertices would not match. You also cannot put only one copy of K, with the angle 4α inserted next to the 6α of D, because the vertex colours would not match. You can only insert two copies of D, which will form an Ace, as required.

Next suppose $p \in K$. If you put a copy of D along the short edge (of length 1) then again you are in the above situation: the D will give rise to an Ace. If you do not put a copy of D, then we can assume that on the two short edges of K we put two more copies of K. But three copies of K cannot glued in this way along the short edges: it would create an angle of $3 \times 4\alpha > 2\pi$.

Step 2. A clever Composition-Decomposition trick.

Given a tiling T by K, D. Call T^{2n} the tiling obtained from T by *Composing* $2n$ times (this new tiling uses the huge tiles $\phi^{2n}K, \phi^{2n}D$).

Now $p \in T^{2n}$. Therefore, by Step 1 applied to the tiling T^{2n}, the point p lies inside some huge Ace of T^{2n}. Now apply *Decomposition* $2n$ times. Then

$$p \in (\text{decompose the huge Ace } 2n \text{ times}) = C_{2n} \subset T.$$

(Here we used that *Composition/Decomposition* are unique, and by definition the repeated *Decompositions* of an ace give cartwheels using rescalings of the tiles K, D) □

Remark 16. *The above theorem does not imply that every tiling by K, D is a carthwheel tiling! Even though each tile lies inside an arbitrarily large cartwheel C_{2n}, these copies of C_{2n} need not be concentric around the given tile. In the above proof, you do not know where p lies inside the huge Ace, so you do not know where p lies inside the copy of C_{2n}. There is no guarantee that p will always be at the centre of the huge Aces as n varies.*

9. Any finite region of a Penrose tiling is repeated infinitely often

9.1. **Vertex neighbourhoods.**

Given a vertex, there are 7 possible ways of arranging K, D tiles around this vertex. We call these the 7 *vertex neighbourhoods*:

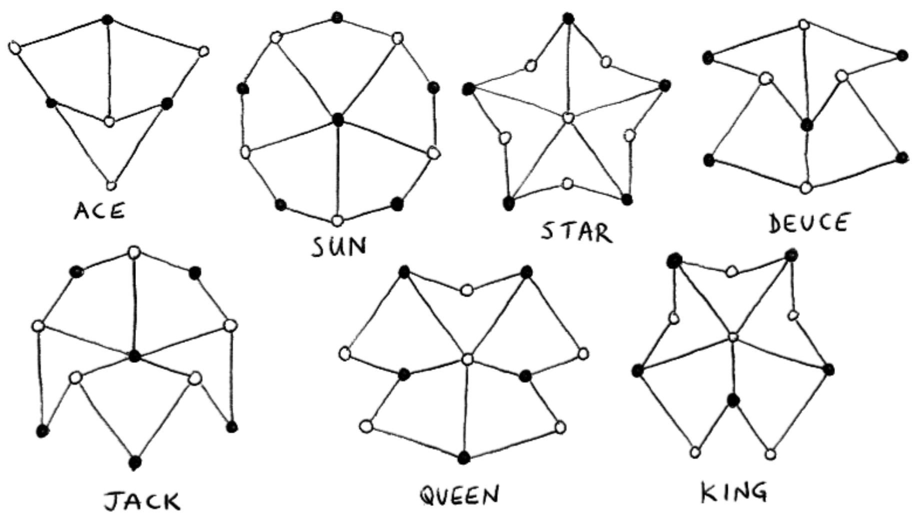

Theorem 17. *In every tiling by K, D, each of the 7 vertex neighbourhoods occurs infinitely often.*

Proof. By the previous Theorem, each tile sits inside a copy of C_4. Now look at the picture of C_4: you will find copies of all 7 vertex neighbourhoods inside C_4. □

9.2. **Any finite region of a Penrose tiling occurs infinitely often.**

Theorem 18. *Given any finite region R of a tiling by K, D. There are infinitely many congruent (translated and rotated) copies of R in the tiling. In fact, there are infinitely many congruent copies of R inside* any other tiling *by K, D as well.*

Proof. Let T be the given tiling. Let T^n be the tiling obtained from T by applying *Composition* n times. For large n, this involves a tiling by huge tiles $\phi^n K, \phi^n D$. If R is in the interior of one of these huge tiles, let's say a huge Dart X, then simply consider the infinitely many huge Darts that are in the tiling T^n (these huge Darts are copies of X but they may be translated and rotated).

By *n-Decomposition* we will mean: apply *Decomposition* n times. Observe that if you apply *n-Decomposition* to T^n you get back T, and the *n-Decomposition* of all those huge Darts will give rise to infinitely many copies of the *n-Decomposition* of X (since *Decomposition* involves *unique* rules). But by construction, the *n-Decomposition* of X contains R (since X arose from *n-Composition* on a region containing R).

If the region R is not contained in the interior of a huge tile of T^n, then R must either intersect some edge or some vertex of a huge tile. Then simply consider a huge vertex neighbourhood X which contains R and run the same argument as above (using the previous Theorem that the tiling T^n must contain infinitely many copies of this huge vertex neighbourhood X).

We have found infinitely many copies of R inside the given tiling T, but we can run the same argument with any other tiling Q by K, D: just consider the *n-Composition* Q^n, and consider the infinitely many copies of the huge vertex neighbourhood X inside Q^n. □

Remark 19. *In the Exercises you will improve the above result, by showing that in fact within a distance of ϕ^5 (just under 12) times the diameter of your region you can find another copy of the region (the factor ϕ^5 is not optimal, but it is fairly easy to obtain).*

Remark 20. *The theorem says that you can find infinitely many* congruent *copies of R, but in fact you can find infinitely many* translated *copies of R. Indeed, all the copies of the tiles D in a tiling by K, D are related to each other by rotation by finitely many possible angles: the integer multiples $0, \alpha, 2\alpha, \ldots, 9\alpha$ of α. Similarly all Ks are related by rotations by those angles, and all vertex neighbourhoods are related by those angles. The rest of the argument now follows by the next exercise.*

Exercise 21. *Check that in a large enough cartwheel C_{2n} you can find all vertex neighbourhoods rotated by all those angles. If you prefer, you could also use the* sun tiling *(or the* star tiling*) to find those neighbourhoods, and then apply the above theorem.*
Example. The centre of the sun tiling *has ten overlapping Aces meeting in the centre, thus all ten of the above angles occur for the Ace vertex neighbourhood.*

10. Average number of K, D per unit area of the plane

Consider a tiling by K, D. Recall $B_r(0)$ is the disc centred at 0 of radius r. Then, we will now prove that,

$$\frac{\#\text{K inside } B_r(0)}{\#\text{D inside } B_r(0)} \to \phi$$

as $r \to \infty$. So the average number of Kites over Darts on larger and larger discs will approximate the golden number (!). There is nothing special about the ball $B_r(0)$: you could also use larger and larger squares, or larger and larger hexagons.

In general, you can take any sequence $R_1, R_2, \ldots$ of larger and larger convex regions. "Larger and larger" means we require that, given any large ball $B_r(0)$, there is some N so that all R_n contain $B_r(0)$ for $n \geq N$.

Exercise 22. *In the process of inflation, count how many ϕK get produced from each K and from each D. Similarly count how many ϕD get produced. Deduce that, starting with any partial tiling by K, D,*

$$\lim_{n\to\infty} \frac{\#\phi^n K \text{ after the n-th Inflation}}{\#\phi^n D \text{ after the n-th Inflation}} = \phi.$$

Thus, for tilings by K, D arising by Inflations+Extension Theorem, if you take $R_1, R_2, \ldots$ to be the (possibly non-convex) regions obtained by repeated Inflations, then ϕ is the limit of the ratios (number of K inside R_n)/(number of D inside R_n) as $n \to \infty$.

In the Exercises, you will show that, in general, you cannot drop the assumption about convexity: if you were allowed to choose non-convex regions $R_1, R_2, \ldots$, then you can in fact make the ratios (#K inside R_n)/(#D inside R_n) converge to any number you like.[26]

Theorem 23. *Take any sequence $R_1, R_2, \ldots$ of larger and larger convex regions of the plane. Then for any tiling of the plane by K, D the ratio of numbers of K, D tiles converge:*

$$\lim_{n\to\infty} \frac{\#K \text{ inside } R_n}{\#D \text{ inside } R_n} = \phi.$$

Proof. (There are probably shorter and easier proofs of this – try to find one!)

By the subdivision trick, we can pass from a tiling by K, D to a tiling by A, B. Each K gives rise to two A, and each D gives rise to two B. So:[27]

$$\frac{\#\text{A inside } R_n}{\#\text{B inside } R_n} = \frac{2\#\text{K inside } R_n}{2\#\text{D inside } R_n} = \frac{\#\text{K inside } R_n}{\#\text{D inside } R_n}.$$

So it remains to show that the ratio of the number of As over the number of Bs converges to ϕ. Recall the two composition tricks:

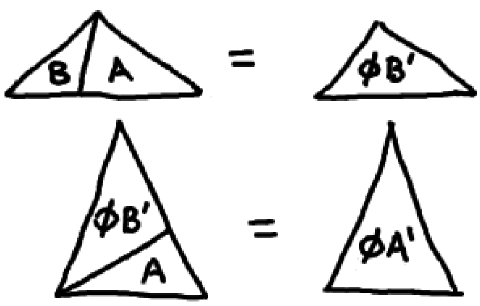

Ignoring markings (since we only care about counting tiles), these compositions can be summarised by the equations:

$$\begin{aligned} \phi B &= A + B \\ \phi A &= A + \phi B \\ &= A + (A + B). \end{aligned}$$

[26] The idea is similar to the Exercise, mentioned in my other Masterclasses Lecture notes, where you rearrange a certain series so that the infinite sum converges to any number you like.

[27] The first equality is not quite right, but it is correct in the limit $n \to \infty$. There may be a few K, D intersecting the boundary ∂R_n which subdivide to give copies of A, B which are inside R_n.

Let's see what happens when we keep rescaling by ϕ the tiles[28]

$$\begin{array}{rclrcl}
\phi A &=& A+(A+B) & \phi B &=& A+B \\
\phi^2 A &=& 2A+3(A+B) & \phi^2 B &=& A+2(A+B) \\
&=& 5A+3B & &=& 3A+2B \\
\phi^3 A &=& 5A+8(A+B) & \phi^3 B &=& 3A+5(A+B) \\
&=& 13A+8B & B &=& 8A+5B
\end{array}$$

Can you see a pattern?

The coefficients in the equation are the Fibonacci numbers $1,1,2,3,5,8,13,\ldots$, from Section 3.2.

Exercise 24. *Show by induction on n that*

$$\begin{array}{rcl}
\phi^n A &=& F_{2n+1}A + F_{2n}B \\
\phi^n B &=& F_{2n}A + F_{2n-1}B.
\end{array}$$

Now consider what happens when you apply *Composition* n times to the given tiling T by K, D. You obtain a tiling T_n by $\phi^n K, \phi^n D$.

Easy Case: suppose a region E is exactly equal to a union of some tiles of T_n (we don't care whether E is convex or not).

Let's say that E consists of a copies of $\phi^n A$ and b copies of $\phi^n B$. Now apply the above formulas for $\phi^n A$ and $\phi^n B$:

$$E \text{ contains } aF_{2n+1} + bF_{2n} \text{ copies of } A$$
$$E \text{ contains } aF_{2n} + bF_{2n-1} \text{ copies of } B.$$

Therefore:

$$\begin{array}{rcl}
\dfrac{\#\text{A inside } E}{\#\text{B inside } E} &=& \dfrac{aF_{2n+1}+bF_{2n}}{aF_{2n}+bF_{2n-1}} \\
&=& \dfrac{a\frac{F_{2n+1}}{F_{2n}}+b}{a+b\frac{F_{2n-1}}{F_{2n}}} \\
&\simeq& \dfrac{a\phi+b}{a+b\frac{1}{\phi}} = \dfrac{a\phi+b}{a\phi+b}\cdot\phi = \phi.
\end{array}$$

Explanation of the above: we divided top and bottom by F_{2n} because we wanted to obtain fractions $\frac{F_{2n+1}}{F_{2n}}$ and $\frac{F_{2n}}{F_{2n-1}}$ whose limits we know are ϕ, by Exercise 11. So for large n those fractions are approximately equal to ϕ: $\frac{F_{2n+1}}{F_{2n}} \simeq \phi$ and $\frac{F_{2n}}{F_{2n-1}} \simeq \phi$.

Easy Case: suppose we have a sequence of larger and larger regions $E_1, E_2, \ldots$, such that, for each n, E_n is exactly equal to some union of tiles from T_n.

Then we can take the limit of the above approximations (see the next exercise):

$$\lim_{n\to\infty} \frac{\#\text{A inside } E_n}{\#\text{B inside } E_n} = \phi.$$

Exercise 25. *There is a small gap in the above argument: the values a and b also depend on n, so we cannot treat them as constants when we take the limit $n\to\infty$ (that is: $a=a_n$, $b=b_n$ depend on n and will typically also grow to infinity). You need to check[29] that the above approximation nevertheless converges to ϕ.*

[28]To find $\phi^{n+1}A$ in terms of A, B you cannot simply replace ϕ^{n+1} by $\phi^{n+1} = F_n + F_{n+1}\cdot\phi$, the formula we found in Section 3.2. That's because the equalities "$\phi B = A + B$" aren't really meaningful as numbers, it is just an abbreviation for how tiles are built out of other tiles.

[29]*Hints. Write $\frac{F_{2n}}{F_{2n-1}} = \phi + s_n$, where you think of s_n as a small error (indeed $s_n \to 0$ as $n\to\infty$). Plug this formula into $\frac{a_n\frac{F_{2n+1}}{F_{2n}}+b_n}{a_n+b_n\frac{F_{2n-1}}{F_{2n}}}$ to obtain a fraction of the form $\frac{a_n\phi+b_n+error}{a_n\phi+b_n+error}\cdot(\phi+error)$, where you need to check that the errors indeed do not matter in the limit.*

Hard Case: suppose we have larger and larger convex regions $R_1, R_2, \ldots$ (possibly unrelated to T_n).

One of the Exercises asks you to finish the proof in this case. Below are some hints – Spoiler alert: stop reading if you do not want any hints!

Inside the convex region R_n, find the largest possible region E_n which is exactly equal to a union of some tiles from T_n. These huge tiles are $\phi^n A, \phi^n B$ so two vertices within a tile are at most ϕ^{n+1} distance apart. Let's call P_n the perimeter of the region R_n. Then the gap $G_n = R_n \setminus E_n$ between the two regions $E_n \subset R_n$ has width at most ϕ^{n+1} (the distance between adjacent vertices), therefore

$$\text{Area}(G_n) \leq P_n \cdot \phi^{n+1}.$$

For large n, this is very small compared to Area(E_n) (here we use the fact that the regions $R_1, R_2, \ldots$ are becoming "larger and larger", meaning: for sufficiently large n the regions R_n contain the arbitrarily large disc $B_r(0)$).

Finally, writing $A(X), B(X)$ for the number of A or B tiles in region X, and using that $A(R_n) = A(E_n) + A(G_n)$, and similarly for B, we obtain

$$\frac{A(R_n)}{B(R_n)} = \frac{A(E_n)}{B(E_n)} \cdot \frac{1 + \frac{A(G_n)}{A(E_n)}}{1 + \frac{B(G_n)}{B(E_n)}}$$

By the easy case, $A(E_n)/B(E_n) \to \phi$, so it remains to check that

$$\frac{A(G_n)}{A(E_n)} \to 0 \qquad \frac{B(G_n)}{B(E_n)} \to 0.$$

But this follows immediately from the fact that the area of G_n is tiny compared to E_n for large n.

Exercise 26. *Fill in the details for the very last step. It needs a little care since, you need to check that it is not possible for the large E_n to have very few tiles of a certain type, say A, because otherwise, in principle, G_n may have many tiles of type A making the fraction above large, even though the area of G_n is small compared to E_n. To prove that, you need to look at the possible vertex neighbourhoods (you need to expand the "sun" and the "star" vertex neighbourhoods, which contain only one tile type, to see that the other type of tile must arise nearby)* □

11. Algorithm to decide whether a patch of tiles can be extended

Spoiler alert: this Section answers most of Exercise 2.

Theorem 27. *A computer can decide in finite time whether a partial tiling by K, D can be extended to a tiling of the plane.*

Proof. Suppose that the partial tiling R can be extended to a tiling of the whole plane. Recall in the proof of Theorem 18 that one looks for a huge vertex neighbourhood X in T^n (the tiling obtained after n *Compositions*) which contains the region R. One can write down an explicit formula for an n that works in terms of the diameter of the region R.

Now the cartwheel C_4 contains all vertex neighbourhoods, and applying n *Decompositions* we obtain C_{4+n} (rescaled by ϕ^{-n}). Therefore, a copy of R must sit inside C_{4+n}.

Thus, given any partial tiling R, it is enough for the computer to build the cartwheel C_{4+n} (with n determined by the diameter of R) and check if there is a copy of R (this only involves a finite number of calculations). If there isn't, then R cannot be extended to a tiling of the plane. □

12. There are many different Penrose tilings of the plane

Spoiler alert: this Section answers most of Exercise 13.

Theorem 28. *There are uncountably many different tilings by K, D. Indeed, there is a bijection between $\mathbb{R}$ and the possible different tilings (they have the same cardinality).*

Sketch proof. Given a tiling by K, D, apply the subdivision trick to obtain a tiling by A, B. More precisely, this is now a tiling by four types of tiles: A_1, A_2, B_1, B_2. Indeed recall that, for example, subdividing K gives two halves, the left-half is A_1 and the right-half is A_2: due to the markings these are actually different tiles, differing by a reflection.

Exercise 29. *Show that composition trick* 1 *($A, B \Rightarrow A, \phi B'$) is of the following two types:*

$$B_2 + A_1 = \phi B_2' \quad \text{or} \quad B_1 + A_2 = \phi B_1'$$

(the first type is drawn in the picture in Section 4.2).
Show that composition trick 2 *($A, \phi B' \Rightarrow \phi A', \phi B'$) is of the following two types:*

$$A_1 + \phi B_1' = \phi A_1' \quad \text{or} \quad A_2 + \phi B_2' = \phi A_2'.$$

(the first type is drawn in the picture in Section 4.3).

Given any point $p \in \mathbb{R}^2$ you can now specify "coordinates" which tell you in which tile p is located in the tiling, as follows.

Starting with the given tiling $T = T_0$, apply *Composition* many times to obtain: $T_1, T_2, T_3, \ldots$ After n *Compositions* we obtain T_n, which is a tiling by $\phi^n A_1, \phi^n A_2, \phi^n B_1, \phi^n B_2$.

One[30] of these tiles contains p, let's say this is a tile of type $\phi^n T_{n,p}$ where $T_{n,p} \in \{A_1, A_2, B_1, B_2\}$. Thus, to p, we have associated a sequence involving the symbols A_1, A_2, B_1, B_2:

$$T_{0,p},\ T_{1,p},\ T_{2,p},\ T_{3,p},\ \ldots.$$

These are called the coordinates for p, and the $T_{n,p}$ are called the entries.

Example. In the following picture, after 4 Compositions the point p lies in the huge tile $\phi^4 A_2$, so $T_{4,p} = A_2$. The first five entries of the coordinates for p are: A_1, B_2, B_2, A_2, A_2. Try to check this with a pencil (in the next example, we will explain this further).

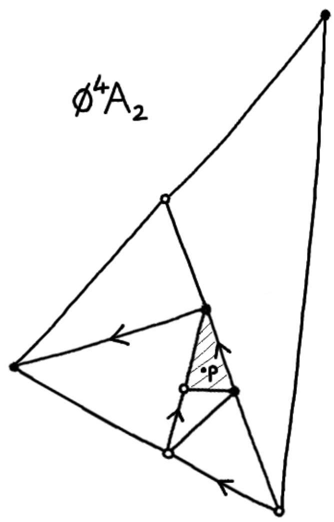

[30] There may be several choices when p lies on an edge or on a vertex of a tile of T, but this is not a serious issue. We could simply stipulate that p should be in the interior of some tile of T.

Notice that, given $T_{n,p}$, if we apply *Decomposition* n times (that is, the reverse of the *Compositions* above), then we are back to a patch of tiles R_n inside T: so we have obtained a neighbourhood of tiles surrounding p.

Exercise 30. *Check that these patches $R_1 \subset R_2 \subset \ldots$ are roughly*[31] *concentric around p, and that the R_n converge to the tiling T as $n \to \infty$.*

Thus, the "coordinates" of a point p tell you how to build T in practice, with centre p.

Example. For example, if the coordinates for p are: $A_1, B_2, B_2, A_2, A_2, \ldots$, then the first three steps of the reconstruction are below, and the fourth step gives the above picture of $\phi^4 A_2$. Notice that in the first two steps only one of the two composition tricks was used, whereas in the third step we used both (of combined type: $B_2 + A_1 + A_2 = \phi A_2'$).

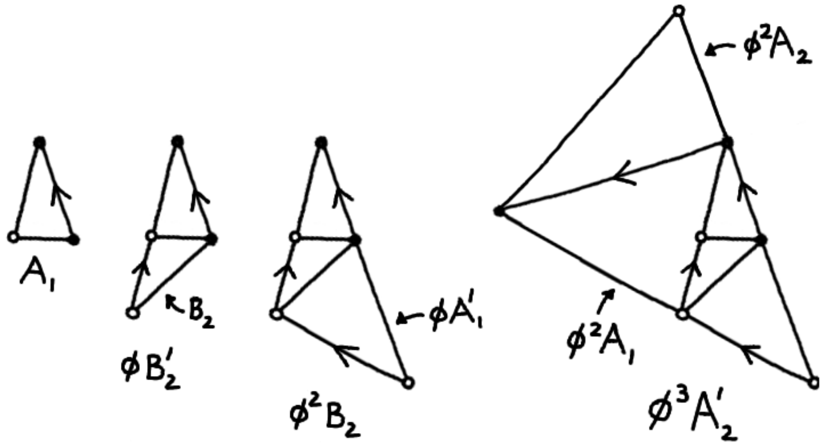

Now consider two points $p, q \in \mathbb{R}^2$. The first few entries of their coordinates may be different, but eventually the entries will be the same[32] because the tiles become so huge that both p, q will be in the same huge tile $\phi^n T_{n,p} = \phi^n T_{n,q}$ for all large enough n.

Thus, the entries of the coordinates of any two points can only differ in finitely many places. We can define an equivalence relation on the collection of all sequences involving A_1, A_2, B_1, B_2: we say that two sequences are equivalent if they only differ in finitely many places. Thus, the coordinates of points all belong to a unique equivalence class, and once you know the equivalence class you can reconstruct the tiling in practice, as we saw above.

More precisely, the equivalence class reconstructs the tiling up to translation and rotation (the choice of the point p fixes the translation, and of course the coordinates we get do not

[31] This is not entirely correct, there are exceptions caused by a reflection symmetry, see the next footnote.

[32] This is not entirely correct, there are exceptions caused by a reflection symmetry:

Suppose that in the above example, the sequence of moves repeats itself after reaching the first A_2. So for example from $\phi^3 A_2'$ instead of $\phi^4 A_2$ (obtained by composition trick 2) we instead get $\phi^4 B_1$ (via composition trick 1), just like the first move when we passed from A_1 to $\phi B_2'$ (composition trick 1). Then the sequence we get is $A_1, B_2, B_2, A_2, B_1, B_1, A_1, B_2, B_2, A_2, \ldots$. The problem is that the shaded tile (the original tile containing p) would always be touching the edge with the arrow of the A-triangles that we build. So we would actually be constructing a tiling of the half-plane (consider the limit of the nested A-triangles). If this is to become a tiling, there must be a copy of the original A-tile reflected on the other side of the line that defines the half-plane. Now you can check that a point in that tile, A_2, must have the same coordinates except with all indices switched (so $A_2, B_1, B_1, A_1, B_2, \ldots$) otherwise the two half-plane tilings will not match along the line separating them. So the two half-plane tilings are reflections of each other.

Thus, in this case, two points p, q on different sides of the separating line will have sequences which do not eventually agree: they only eventually agree if we switch all indices of the sequence for q.

In any case, such sequences with repetitions will not play a serious role in determining the cardinality since they are countable (just like $\mathbb{Q}$ is the countable subset of decimal numbers in $\mathbb{R}$ which eventually have a block of decimal numbers that keeps repeating).

change if we rotate the tiling, whereas a reflection will be detected by the coordinates since that switches all $A_1 \leftrightarrow A_2$, $B_1 \leftrightarrow B_2$).

The collection of all possible sequences involving four symbols (such as A_1, A_2, B_1, B_2) has the same cardinality as $\mathbb{R}$ as follows. It is enough to check that it has the same cardinality as the interval $[0, 1]$ (observe that $[0, 1]$ has the same cardinality as $\mathbb{R}$ since $\mathbb{R}$ is a countable union of intervals of length 1). Now write decimal numbers in $[0, 1]$ in base 4 and replace $0, 1, 2, 3$ by the symbols A_1, A_2, B_1, B_2 to get the corresponding sequence. This is not quite correct, since some decimal numbers in base 4 are identified (such as $0.3333\ldots = 1.0000\ldots$): but there are only countably many such, so the cardinality is not affected by this mistake. Finally, the collection of equivalence classes also has cardinality $\mathbb{R}$ since each equivalence class contains only countably many sequences.[33]

So we are done? Not quite: we have not shown that each equivalence class arises from some tiling by K, D. In fact, this is not true. The following exercise asks you to study which equivalence classes arise from tilings.

Exercise 31. *Not all sequences involving A_1, A_2, B_1, B_2 arise from tilings. Can you describe which ones do?*[34] *Check that this subset of the equivalence classes still has cardinality $\mathbb{R}$.* □

Remark 32. *Notice that the above proof also gives you a very explicit way to build lots of Penrose tilings in practice, thus answering much of the discussion in Section 2.1.*

13. Acknowledgements

Many thanks to the CMI-PROMYS scholars who attended the 2014 Oxford Masterclasses in Wadham College, and who contributed to shape these notes with their many insightful comments and questions.

[33]pick a sequence, then all other sequences in the class are given by changing finitely many entries: that's countably many choices.

[34]*Hint. the four types of composition tricks mentioned in Exercise 29 impose constraints on the sequences, since the sequences must be consistent with* Composition. *I think you should get that: after a B-letter the next letter must have the same index (for B_1: either B_1, B_1 or B_1, A_1 is allowed), whereas an A-letter can be followed by anything except a B-letter of the same index (so not A_1, B_1 and not A_2, B_2).*