Theorem 7.1 *There exists a protoset of polygons that admits a valid tiling but does not admit a valid tiling with a non-trivial symmetry.*

□

The smallest aperiodic Wang protoset contains 11 tiles, but with geometric tiles a single tile is enough to force non-periodicity! A polygonal prototile (the *hat*) was recently reported that is alone aperiodic: there exist monohedral tilings of $\mathbb{R}^2$ using the hat but none of these tilings has a translational symmetry. We discuss this tile in Section 7.4 below.

An older aperiodic protoset consists of two polygons. In 1974 R.Penrose presented this famous aperiodic pair called *kite* and *dart*:

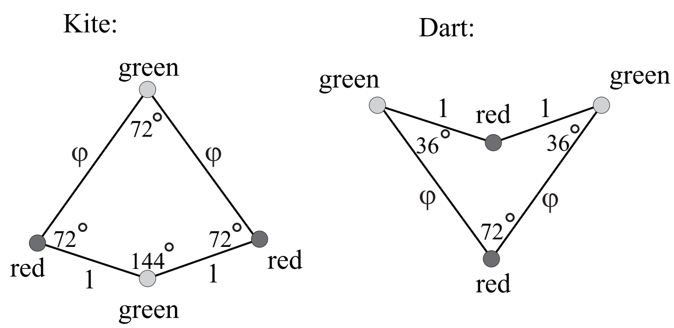

The Penrose tiles are obtained by cutting in two a rhombus that has a 72° angle. The resulting quadrilaterals have edges of length 1 and $\varphi = (1+\sqrt{5})/2 = 1.618\ldots$, the golden ratio. The vertices are colored red and green, and in valid tilings equal edges must be placed together and also the colors at the vertices must match. (This condition prevents one from gluing the kite and dart back together to form the rhombus.) These matching rules can be easily enforced using geometric constraints only by, say, using bumps and dents as follows:

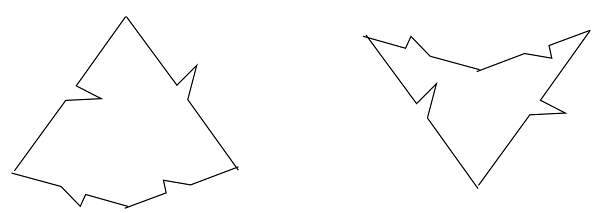

Here is a part of a tiling using kites and darts. (For clarity, the bumps and the dents are not shown):

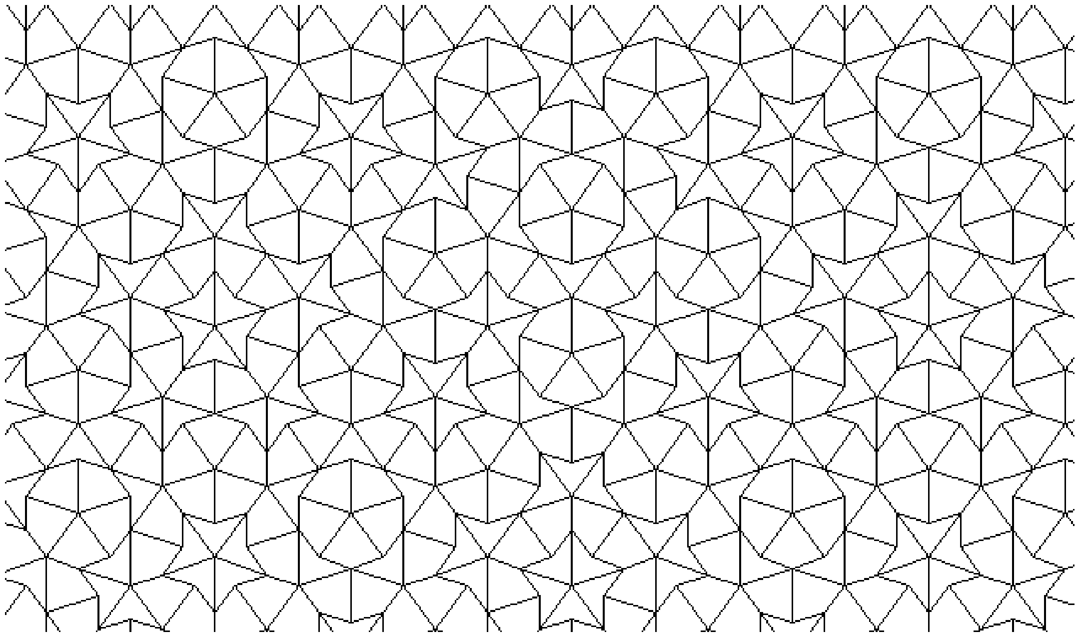

The following result will be proved in the homework problems:

Theorem 7.2 *Penrose kite and dart are an aperiodic pair of prototiles. They do admit valid tilings with a 5-fold rotational symmetry.*

□

7.1 Substitutions

Substitutions are a popular method to construct hierarchical, non-periodic tilings. As a simple example, consider the *chair substitution* where an L-tromino is cut into four smaller, similar shapes:

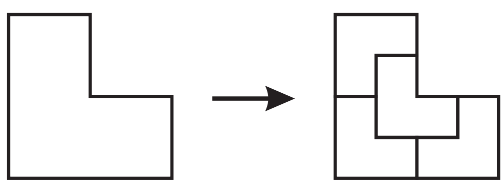

Starting from a single tile, we repeat the following operations:

(i) Replace each tile by four smaller copies as above,

(ii) Magnify the obtained pattern by factor two horizontally and vertically.

Here are the first two iterations (the meaning of the grey tile will be explained later):

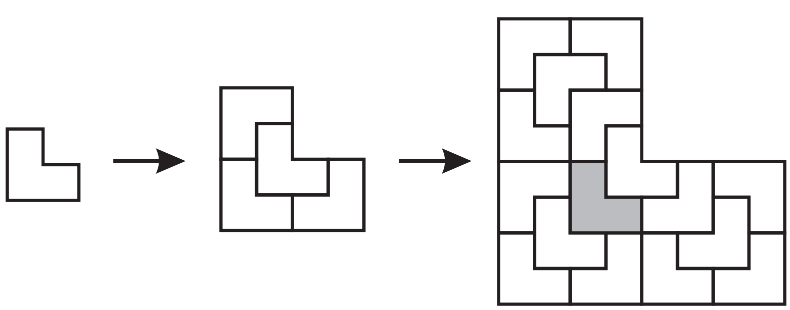

The first and the second steps produce “supertiles” that consist of four and sixteen copies of the tile. The k'th supertile is a pattern of 4^k tiles. It is clear that in this way we obtain ever larger areas covered by the tiles. To obtain a tiling of the infinite plane, we suitably position the obtained supertiles on the plane so that the next pattern always expands the previous one, and take the limit of the process. Note that the k'th supertile consists of four copies of the $(k-1)$'st supertile, so we can position it (in four different ways) on the plane so that the $(k-1)$'st supertile is its subpattern. Moreover, we can do this positioning so that the patterns grow in all directions, so that each point of the plane gets eventually covered by a tile.

For example, in the illustration above, we can align the gray tile of the second supertile over the initial tile, and repeat this positioning after all even rounds. Because the grey tile is not on the boundary, it is guaranteed to get surrounded by more and more tiles on all sides. The following illustration shows the position of the grey tile in the fourth supertile.

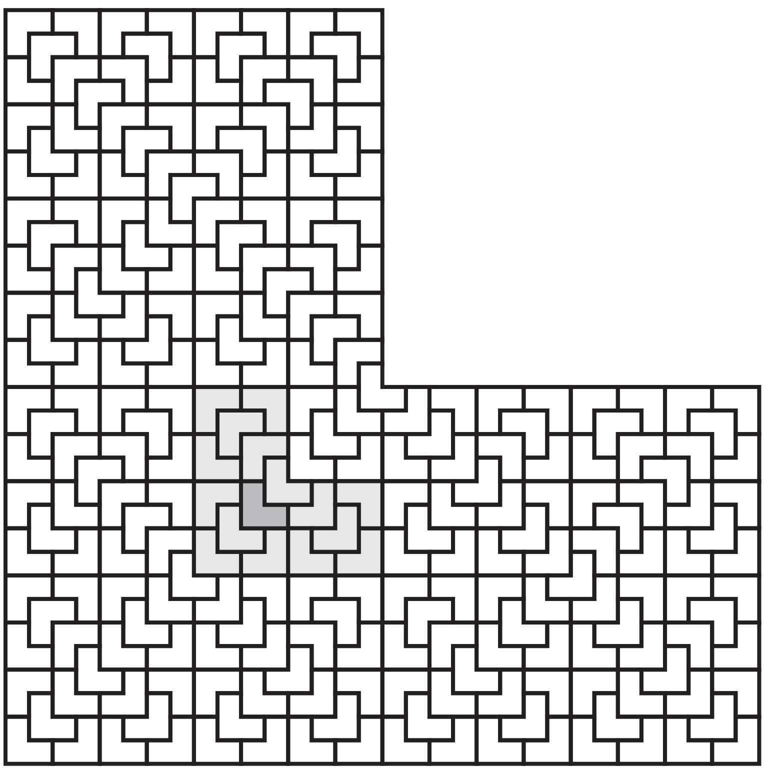

In the limit we obtain a tiling t of the infinite plane. It is clear that in this tiling each tile belongs to a supertile, which in turn belongs to a second level supertile, and so on.

To prove that the obtained tiling t is not periodic, we make the important observation that two supertiles cannot overlap: each tile in the tiling is part of a *unique* supertile. To see this, consider the four parts of a supertile:

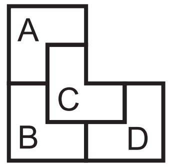

We observe that neighbors A and B meet at the end of the L, that is, they are adjacent to each other at an edge denoted by $*$ here:

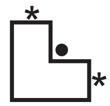

Also neighbors B and D meet each other at their $*$-ends. In contrast, the $*$-ends of tile C are not adjacent to $*$-ends of other tiles. So we can recognize the center tiles C of all supertiles simply by the property that neither $*$-end is adjacent to a $*$-end of another tile. In the other cases A, B and D, the neighbor at the inner corner (indicated by $\bullet$ above) is the center tile C of the supertile, and hence the supertile containing the tile is unique also in this case.

We conclude that tiling t has a unique partitioning into a tiling by the supertiles. The same reasoning then applies to the next levels, so that tiling t can be partitioned in a unique way into a tiling by k'th level supertiles, for every k.

Suppose now that tiling t would have a period $\vec{n} \neq (0,0)$. Applying translation $\tau_{\vec{n}}$ to the supertiling produces a supertiling. But since t remains the same under $\tau_{\vec{n}}$, and since the coresponding supertiling is unique, the supertiling must have period $\vec{n}$. The same reasoning applies to supertilings of level k, for every k. But for sufficiently large k, the k'th level supertile overlaps its translate by $\vec{n}$, so $\vec{n}$ cannot be a period of the k'th level supertiling, a contradiction.

We have proved that the tiling obtained by iterating the chair substitution is non-periodic. It was essential in the proof that the partitioning of the tiling into supertiles is unique. Note that the L-tromino is of course not aperiodic: it trivially tiles a 2×3 rectangle, which yields a periodic tiling of the plane. However, there are general methods to decorate tiles in such substitutions so that non-periodicity is forced. This increases the total number of tiles as several variants of the tiles are needed with different decorations. We skip the general method, but discuss in detail Amman's aperiodic tile set that is also based on the substitution concept.

7.2 Amman's aperiodic tile set

The following pair of tiles, due to R.Amman in 1977, forms an aperiodic tile set.

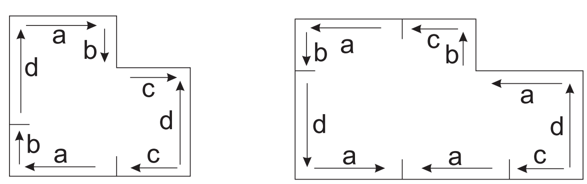

The tiles may be rotated and flipped in any orientation. The labeled arrows along the edges give a matching rule: each arrow must fit against an arrow with the same label and oriented in the same direction. The lengths of the arrows are arbitrary (yet positive), but all arrows with the same label have also the same length. (A remark: this matching rule can not be implemented geometrically with bumps and dents. It can be implemented geometrically, though, by adding a third key tile, so that the construction provides an aperiodic set of three geometric tiles.)

The following illustration shows how an A-tile and a B-tile fit together into a *super-A*, and how an A-tile and two B-tiles form a *super-B*:

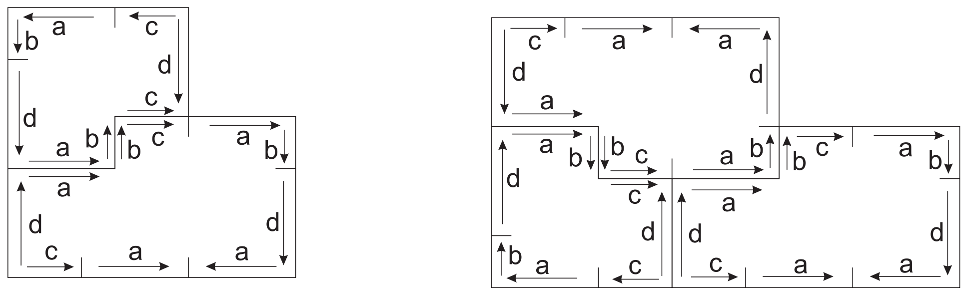

Any tiling by the resulting supertiles

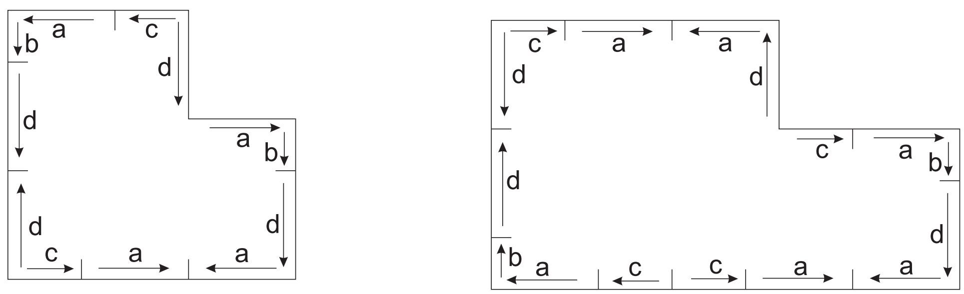

can hence be broken into a tiling by the original tiles. Let us redecorate the supertiles with arrows labeled a', b', c' and d', where the new arrows represent combinations of old arrows as follows:

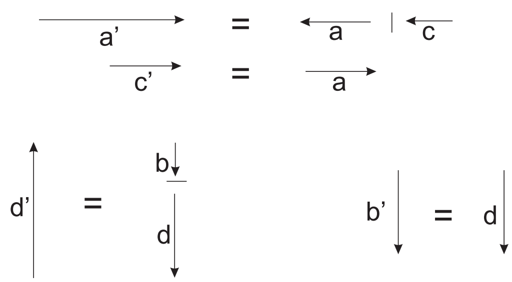

The redecorated supertiles

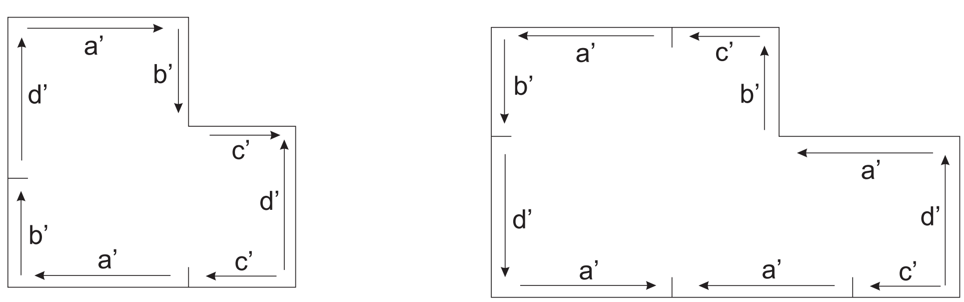

are called *expanded* A and *expanded* B, respectively. Of course, in any valid tiling by the expanded tiles, the tiles can be replaced by the corresponding supertiles, and the tiling remains valid. But also the converse holds: in any tiling by the supertiles, replacing the supertiles by the coresponding expanded tiles yields a valid tiling. To see this, note that in the supertiles the c-arrow is always immediately followed by an a-arrow in the same direction. Hence two neighbors with a common c-segment also share the following a-segment. So replacing every ca -segment by a new a'-arrow leaves the tiling valid. Analogously, every

b-arrow is immediately followed by a d-arrow, so replacing bd -segment by a new d'-arrow keeps the tiling valid. But the obtained tiles are (up to renaming the arrows) exactly the expanded A- and B-tiles. We conclude that the supertiles and the expanded tiles admit exactly the same tilings.

Next we observe that the expanded tiles are isomorphic to the original tiles, where the arrows with labels a', b', c' and d' correspond to the arrows a, b, c and d, respectively. (However, the ratios of the arrow lengths may change, so the shapes of the expanded tiles are not necessarily similar to the original tiles. Similarity in shapes is obtained if the length of arrow d is φ times the length of b, and the length of a is φ times the length of c, where φ is the golden ratio, i.e., the positive number satisfying $\varphi^2 = \varphi + 1$. But similarity in shapes is not necessary in the reasoning that follows.)

We can now build supertiles of level two by simply combining the expanded tiles the same way we combined the original tiles to build the first level supertiles. Iterating the reasoning allows us to build supertiles and expanded tiles of levels two, three, four and so on. These provide tilings of larger and larger regions of the plane by the original A- and B-tiles. As with the chair substitution, we can take the limit, which yields a valid, hierarchical tiling of the infinite plane. The following illustration shows the first levels of the process:

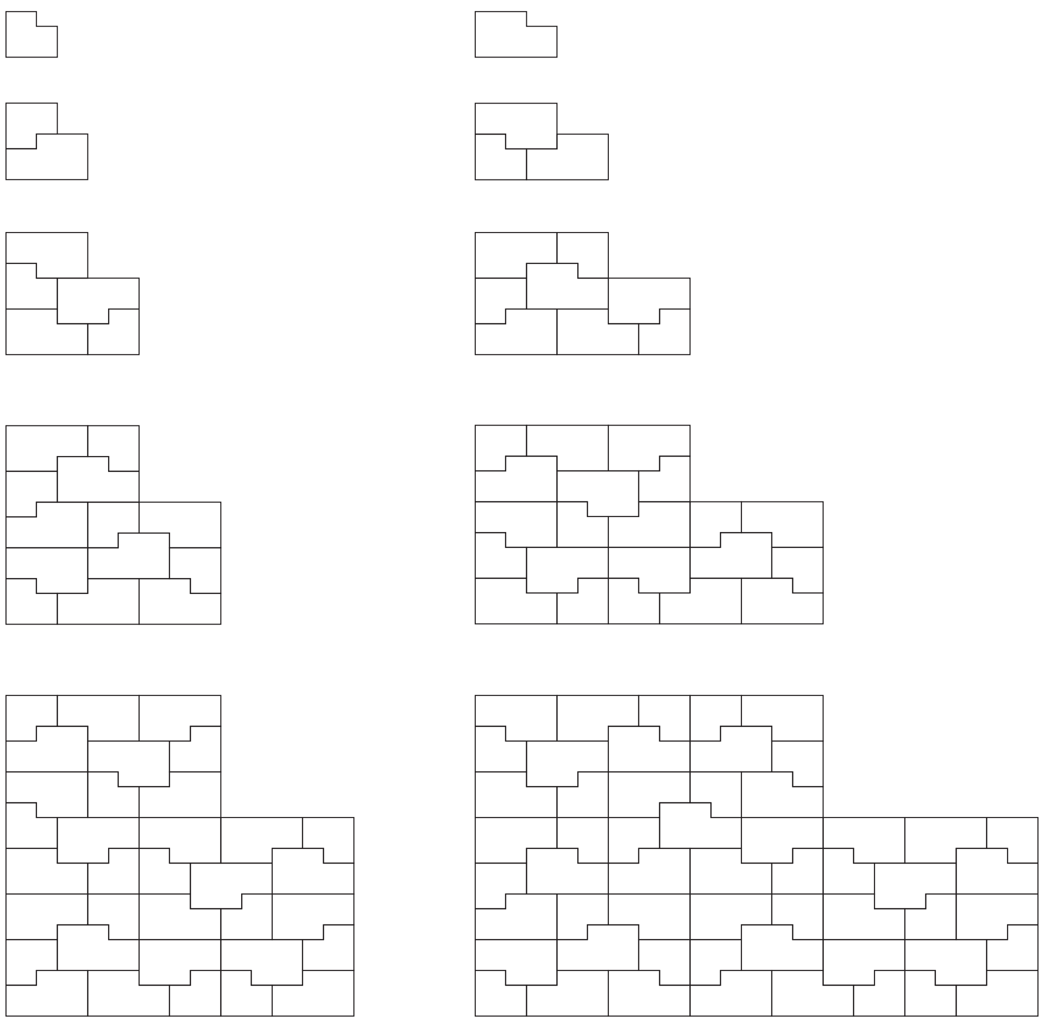

Consider now an arbitrary tiling of the plane by A and B. First we prove that every tile belongs to a supertile. Consider an A-tile and its neighbor at the inner corner *.

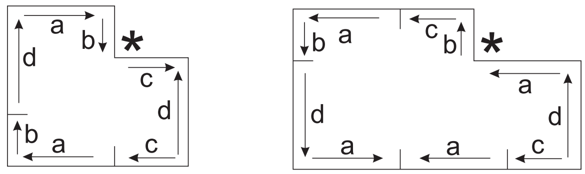

A simple case analysis (based, for example, on the b-arrows) shows that the only tile to match is the B-tile, oriented to form the super-A.

Consider then a B-tile and its inner corner *. An A-tile fits in the corner in two different ways: one provides super-A, as seen above; the other (shown below) cannot be expanded into a tiling of the plane, because no tile fits in the corner •:

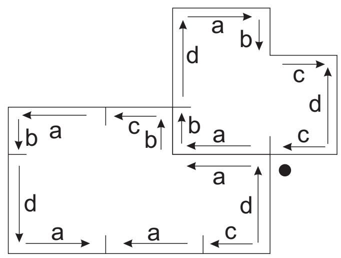

On the other hand, there is only one possible way to fit a B-tile in the *-corner of the B-tile:

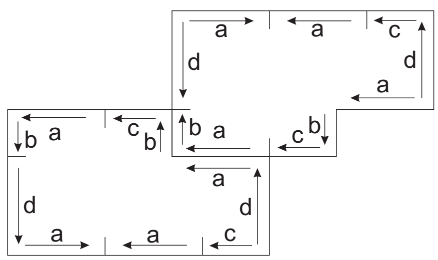

The only tile to fit in the *-corner of the second B-tile is the A-tile, so that the three tiles together form super-B. We conclude that every tile of a valid tiling is part of a super-A or a super-B.

In fact, the tiles of any tiling by A- and B-tiles can be grouped into non-overlapping supertiles and, moreover, such grouping is unique. To see this, find first all B-tiles whose *-neighbor is also a B-tile. These necessarily are part of a super-B, as discussed above. The super-B tiles do not overlap. All remaining tiles must be part of super-A tiles, so the unique grouping is concluded.

We are ready to make the following conclusion:

Theorem 7.3 *The A- and B-tiles form an aperiodic pair of tiles.*

Proof. By iterating the substitution to form supertiles, while properly aligning the obtained supertiles, we obtain a sequence of growing, correctly tiled patterns. In the limit, a valid tiling is obtained.

Let us show that no periodic tiling is possible. Suppose the contrary: A valid tiling t exists that is invariant under the translation by $\vec{n} \neq \vec{0}$. The tiles in t can be partitioned in a unique way into supertiles. If this tiling t_s by supertiles is translated by vector $\vec{n}$, a tiling t'_s by the supertiles is obtained. However, when the supertiles in t'_s are broken into their A- and B-pieces, the obtained tiling is the $\vec{n}$-translation of t, hence it is equal to t. But the supertiling obtained from t is unique, so $t'_s = t_s$, and tiling t_s is invariant under the translation by $\vec{n}$.

When the supertiles in t_s are replaced by the corresponding extended tiles, a valid tiling t_e by the extended tiles is obtained that has period $\vec{n}$. By repeating this argument on t_e in place of t, and iterating the reasoning, we see that there are valid tilings by extended tiles of all levels that are $\vec{n}$-periodic. This is not possible since an extended tile of a sufficiently high level overlaps with its translation by $\vec{n}$.

We conclude that the A- and B-tiles do not admit a periodic tiling, and hence they are an aperiodic pair of tiles. □

7.3 The extension and the periodicity theorems

To prove that repeating a substitution leads, in the limit, to a tiling of the plane was easy in our examples. We used the fact that the pattern obtained at iteration $k+1$ contains the level k pattern as its subpattern. This means that the sequence of obtained patterns have a well-defined limit. Moreover, in our examples the patterns grow on all sides, so that the limit covers the whole plane and is thus a valid tiling of the plane.

But it can be shown that even if the obtained patterns do not contain previous ones as subpatterns, a valid tiling exists as long as arbitrarily large disks can be covered. This result (which we present without a proof) extends Corollary 4.3 from Wang tiles to geometric tiles. It is important to note that the theorem only considers finite protosets of topological disks:

Theorem 7.4 (Extension theorem) *A finite protoset $\mathcal{P}$ of topological disks admits a tiling if and only if, for every $r > 0$, a disk of radius r can be covered by copies of the prototiles. (That is, there is a collection of tiles, all congruent to elements of $\mathcal{P}$, such that (i) the interiors of the tiles are pairwise disjoint, and (ii) a disk of radius r is included in the union of the tiles.)*

□

Theorem 7.4 above generalizes Corollary 4.3 from Wang tiles to geometric tiles. The other important basic property of Wang tiles states that a Wang tile set that admits a one-way periodic tiling admits also a two-way periodic tiling (Theorem 4.1). Also this theorem can be generalized to finite sets of polygonal prototiles and edge-to-edge tilings:

Theorem 7.5 (Periodicity theorem) *Let $\mathcal{P}$ be a finite set of polygons. Assume that there exists an edge-to-edge tiling by the protoset $\mathcal{P}$ that is one-way periodic (=invariant under some translation). Then there also exists an edge-to-edge tiling by $\mathcal{P}$ that is two-way periodic (=invariant under translations by two linearly independent vectors).*

□

7.4 Hat: an aperiodic monotile

The polygonal tile *hat* was recently proved to be an aperiodic monotile, *i.e.*, there exist monohedral tilings of the plane using the hat, but none of these tilings are periodic. The tile is a union of eight *kites* of a grid formed by overlapping the regular tilings by equilateral triangles and regular hexagons: