Exploring Aperiodic Tilings with Inflation

by
Forrest Trepte

A Thesis
Submitted in partial fulfillment of
the requirements for the Degree of Bachelor of Arts with Honors
in Computer Science

WILLIAMS COLLEGE
Williamstown, Massachusetts
June 2, 1996

Abstract

This work focuses on realities of implementing an abstract datatype that traverses infinite aperiodic tilings of the plane and three-space. The approach is to develop a generalized mechanism for specifying finite sets of prototiles that tile space nonperiodically. When the tiling exhibits an invertable inflation property, tiles may be addressed in a way that supports arbitrary, possibly oriented, wanderings through space. Techniques are presented that can be used to discover lines of symmetry in a tiling and to construct an atlas of its local configurations. Also developed are a number of theoretical results about this class of tilings, including a concise statement of a local isomorphism theorem. This pretty result suggests that any finite portion of a specific tiling by prototiles appears infinitely often in all other tilings by the same set of tiles.

Our analysis is inspired by the index sequences that are commonly used to uniquely specify tilings. An index sequence identifies a tiling about a point by locating the point within an infinte hierarchy of self-similar tilings known as inflations. Our work builds an addressing system with index sequences of finite rather than infinite length and uses addresses to represent the dual graph of a patch of tiles. An n-digit address identifies a tile within n inflations of some prototile. We prove that any tiling represented in our form can tile the plane or space and must do so nonperiodically. An algorithm is constructed to find the addresses neighboring any given tile. Appending digits to an address extends the currently represented region, placing it within a larger patch to simulate the exploration of an infinite expanse of tiles. The information maintained grows as needed and the number of digits stored is logarithmic in the number of tiles in the patch represented. Our system could be used to investigate very large quasiperiodic tilings and the principles behind the data structures illuminate many beautiful properties of the tilings themselves.

Acknowledgements

I would like to thank my advisor, Duane Bailey. This project began as an independent study that stretched into winter study, but it was destined to become a thesis. Working with Duane was too much of a high to give up. His amazing abilities of spatial perception, love of writing, and contagious fascination with the rather odd topic of quasicrystals made this work possible. Duane is both a hacker and a scientist, but most importantly an artist. Thank you Duane.

I am also greatful to Bill Lenhart for his careful reading and commentary on my writing and to my fellow thesis students A.J., Leaf, and Chris for their assistance and late-night companionship.

Contents

List of Figures

Chapter 1

Background

In 1966, Berger discovered a set of 20,426 shapes that will only tile the plane aperiodically. By 1974, R. Penrose had discovered an aperiodic set of just two shapes. A finite collection of shapes that tile the plane is called a set of *prototiles* It is currently an open question whether a single prototile exists that forces an aperiodic tiling. Penrose tilings can be decorated with Ammann lines: line segments placed on the prototiles connect to form a grid of parallel lines passing through the tiling. Two lengths of intervals separate successive parallel lines: long and short. The pattern of long and short intervals formed is quasiperiodic and is a pattern known as the Fibonacci sequence.

1.1 Quasicrystals

Today, there are a number of techniques for generating quasiperiodic tilings and lattices with any symmetry in two, three, or more dimensions[Ste86]. Janot, Dubois, and Boisseiu[JDdB89] present an overview of matching rules, inflation, projection, and generalized dual techniques.

One method follows the technique used by Penrose[Pen84]. Many sets of prototiles can be assembled to produce an aperiodic tiling of the plane. It is more challenging, however, to find a set of prototiles such so that every tiling of the plane composed of these prototiles is aperiodic. Penrose developed *matching rules* for his set of two rhombuses that limit the ways that tiles are allowed to neighbor other tiles, and this prevents Penrose's tiles from forming any periodic tilings. If local rules are used as a mechanism for growing nonperiodic structures, backtracking is required since some placements allowed by the local rules result in patches that are impossible to extend to form a tiling. However, any tiling formed according to local rules is guaranteed to be aperiodic. Local matching rules have attracted considerable attention. Gähler, Baake, and Schlottmann[GBS94] study the relationship between the unmarked shapes of tiles and matching rules that can be used to constrain their placement so that only nonperiodic tilings can be formed. Klitzing and Baake[KB94] investigate a stronger form of local rules for tilings with eight and twelve-fold rotational symmetry. Sometimes it is the case that a patch of tiles can be constructed so that at all times at least one placement is *forced*. That is, there is always some spot at which no other tile can be placed and thus such a tiling can be grown by using matching rules without backtracking. Conway discovered that this is even the case if certain defects are placed at the center of a tiling by Penrose's kites and darts.[SG91] Other approaches use growth rules that are local except that tiles that are forced are always added first. Steinhardt, DiVincenze, and Socolar[OSDS88] grow defect-free tilings by Penrose rhombuses by adding a thick rhombus in a particular way whenever there are no forced placements. Later, Socolar[Soc91] uses a similar technique to analyze the stages of growth of a tiling by Penrose rhombs and presents a physical model that uses random probabilities to account for the condition when no vertices are forced.

A Penrose tiling also obeys a self-similarity property called *inflation*. The tiles of any Penrose tiling can be attached together to form a tiling made up of larger, equivalent pieces. This process can be repeated to form tilings by arbitrarily large pieces. Inflation and matching rules are used to prove that the Penrose rhombuses tile the plane and can do so only nonperiodically. Other types of tilings can be proven aperiodic by use of similar methods. Our approach for representing tilings relies exclusively on this approach for generating tilings. We define inflation precisely in Chapter 4.

Another approach generates quasicrystal structures by projection from a crystal structure in a higher dimension. A hyperplane of a slope with irrational coefficients is passed through a regular grid of lattice points. All lattice points within a certain distance of the hyperplane are orthogonally projected onto the hyperplane. Because each axis of the hyperplane is irrationally sloped with respect to the grid, no projected points along any axis fall in a repeating pattern. The points projected on the hyperplane can be used to construct a corresponding tiling of the plane that is quasicrystalline.

The generalized dual method constructs quasicrystal structures dual to intersecting grids.[dB81] A *grid* is formed by a set of parallel curves. A *N-grid* is a set of N grids such that every pair of curves from different grids intersect at exactly one point. An N-grid can be mapped by a dual transformation onto a tiling of the plane such that every intersection of grid lines becomes a tile and every space between grid lines becomes a vertex of the tiling. This approach can also be used to construct tilings of space or higher dimensions by taking the dual of intersecting surfaces. Durand[Dur94] developed software that generates graphic images of non-periodic tilings by projection from five-dimensional space using an approach given in deBruijn[dB81].

The generalized dual method can produce a wider class of tilings than the projection method and the Penrose approach produces a narrower class of tilings than either of the other approaches.[SS86] In general, only the Penrose approach produces aperiodic prototile shapes—shapes that cannot tile the plane periodically. Typically, *complex matching rules* are required to force the prototiles generated by the other methods to tile the plane aperiodically. In these circumstances, determining whether the placement of a tile is legal requires the examination of a larger area around the tile than its immediate neighbors. Complex matching rules cannot be used to deform the edges of the prototiles to equivalent shapes that do not require the use of matching rules. The projection and generalized dual methods only produce tilings by rhombuses while the inflation method can be used to produce tilings with a variety of shapes.

Shechtman, Blech, Gratias, and Cahn in 1964 melted an alloy of aluminum and manganese and then rapidly cooled it. Electron diffraction of the substance that solidified revealed that it had icosahedral symmetry. Previously, pure solids were classified as either crystals or glasses. Glasses are formed as random packings of atoms that do not produce a structured diffraction pattern. The atoms in a crystal are arranged in parallel planes that reflect incoming particles in such a way that reflected waves add to produce a diffraction pattern. However, a central tenet of crystallography is that only twofold, threefold, fourfold, and sixfold axes of rotational symmetry are possible for a crystal structure while the aluminum alloy had fivefold rotational symmetry. The glass, random-tiling, and Penrose models were advanced to explain these results.[Ste86] Later investigators discovered other alloys and improved techniques for growing crystal-like structures with diffraction patterns that indicate translational order, but with symmetry types forbidden for crystals. These substances are called *quasicrystals*. Current evidence suggests that these substances are formed from an arrangement of atoms that corresponds to aperiodic tilings. The precise arrangement of atoms in these structures and the physical process by which they are formed is as yet unknown.

1.2 Recommended Reading

Stephens and Goldman[SG91], and Steinhardt[Ste86] provide an overview of physics results relating to quasicrystals. A pair of articles by Steinhardt with Levine and

Socolar[LS86][SS86] discuss in detail methods for generating one, two, and three-dimensional quasicrystals and predicting their diffraction patterns. Penrose[Pen84] provides an account of the insight and techniques that lead to the discovery of his first aperiodic set consisting of six prototiles. Gardner's article[Gar77] describes the discovery and properties of the Penrose kites and darts. Chapter 10 of Grünbaum and Shephards' book[GS87] is required reading for anyone interested in aperiodic tilings. Chapter 10 of this book provides an extensive compilation of the properties of various aperiodic tilings and the rest of this book surveys the entire tiling field. Lunnon and Pleasants[LP87] generalize Penrose's method for constructing aperiodic tiles with inflation and matching rules and construct higher-dimensional aperiodic tilings from one-dimensional quasiperiodic sequences. They also discuss in depth generalized notions of properties we are very interested in: irreducible protosets, local isomorphism, and vertex configurations.

A great variety of aperiodic tilings are described in the literature. Grünbaum and Shephard introduce the Wang and Robinson tiles based on square shapes, three sets of Penrose tiles and the equivalent Robinson triangles, and Ammann's tiles whose sides can be continuously varied in length according to certain formulas. The inflation properties of these tilings are discussed and all the sets of prototiles can have their edges deformed so as to force an aperiodic tiling without matching rules. Baake, Ben-Abraham, Klitzing, Kramer, and Schlottman[BBAK+94] catalog all possible arrangements of tiles around a vertex for several tilings with five-fold, tenfold, twelvefold, and three-dimensional icosahedral symmetry. Steinhardt and Socolar[SS86] construct a three-dimensional analog of the Penrose tiles and Danzer[Dan89] gives a related set of four tetrahedra made from planes related to a regular icosahedron. We are most interested in the Robinson triangles and Danzer's tetrahedra since they use a small number of shapes and their inflations are simple and neat.

Chapter 2

Example: Penrose and the Triangles Inside

A *tiling* is a covering and a packing; tiles are placed to fill the plane without overlapping. Figure 2.1 is a tiling constructed from three shapes: a regular triangle, square, and octagon. We study tilings formed from a finite (usually small) number of shapes. A *periodic* tiling is identical to a translation of itself. A periodic tiling translated a certain distance in some direction is indistinguishable from the original. Two such translations are indicated by arrows in Figure 2.1 and all combinations of these work as well. Consequently, the tiling can be formed by the repeated placement of a single rhombus as shown in Figure 2.2.

Two important sets of shapes are Penrose's kites and darts, and thin and thick rhombuses shown in Figure 2.3. Any protoset containing a rhombus admits periodic tilings, but Penrose marked his tiles with decorations known as *matching rules* that constrain tile placements. The arrows and dots on the corners and edges of a Penrose tile must be identical to those of its neighbor and Figure 2.4 is a tiling with kites and darts placed in this manner. The matching rules are used as a shorthand for tiles with more complicated shapes. Deformed versions of Penrose's kite and dart can be formed that can only be placed according to the matching rules. Since no periodic tilings can be formed by Penrose's tiles in accordance with the matching rules, it is said that the kites and darts, and rhombuses are *aperiodic*. In Figure 2.4 only a finite section of the tiling is visible, so it is impossible to register that it is not periodic, but clearly there is an irregularity here, as compared with Figure 2.1. A proof of the aperiodicity of Robinson's triangles is found in Section 2.2.2.

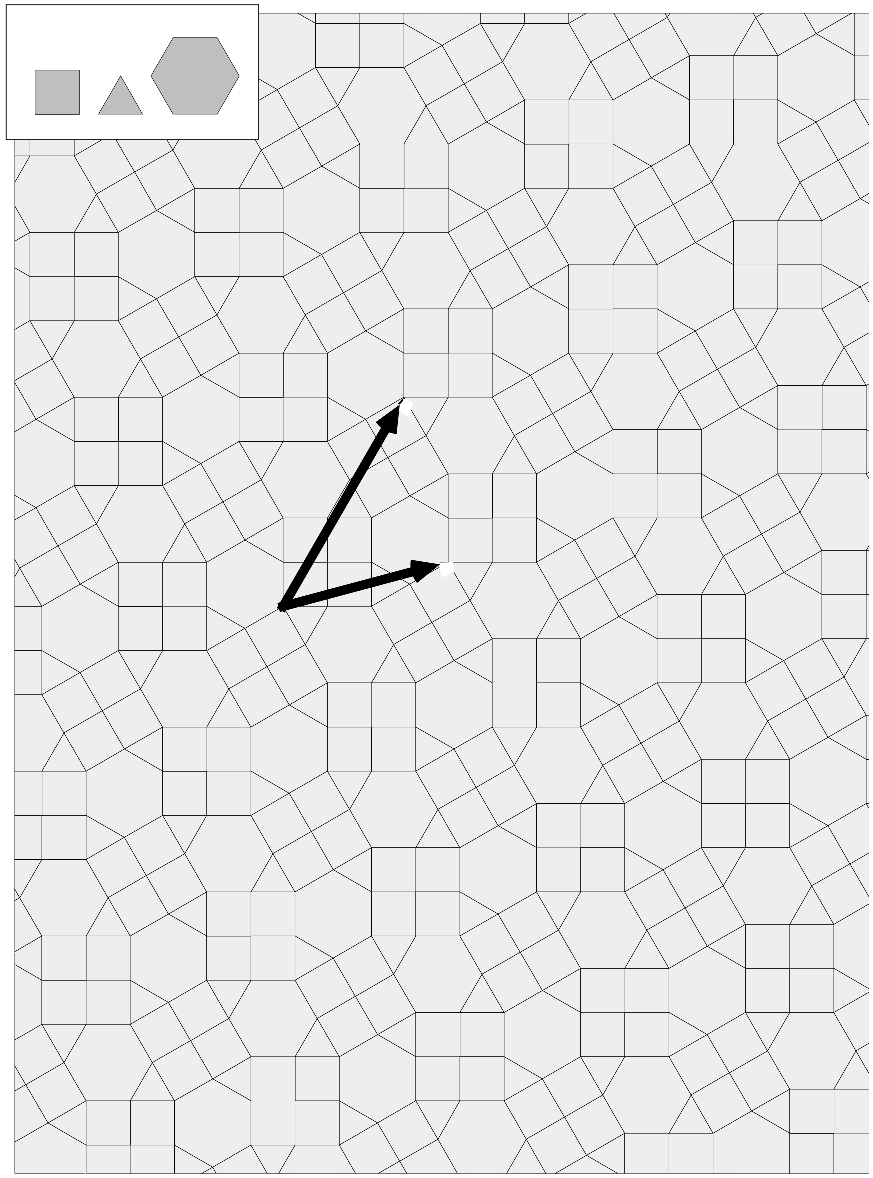

Figure 2.1: A periodic tiling. The upper-left corner shows the 3 shapes that are placed without overlapping and fill the plane. The tiling can be translated according to either of the two arrows without effect. Tilings with such translations are *periodic*.

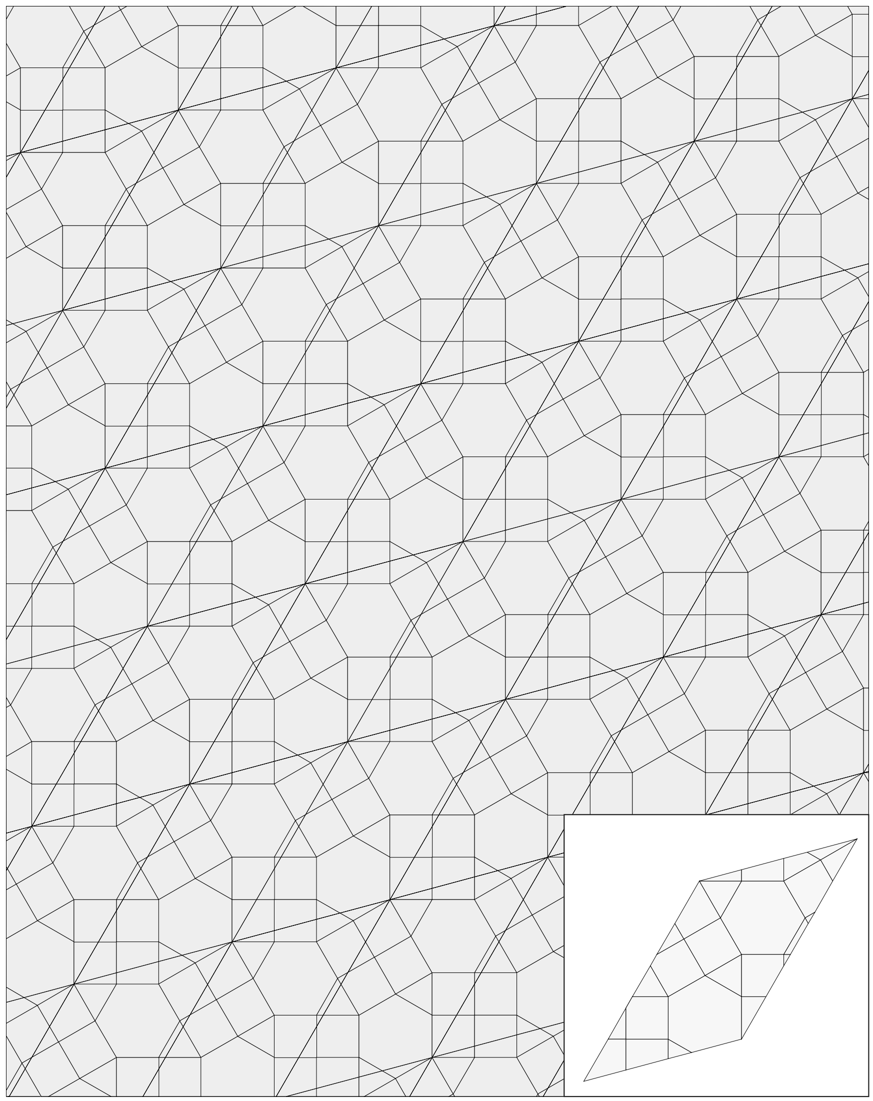

Figure 2.2: A periodic tiling as a grid of period parallelograms. Any combination of the two translations of Figure 2.1 maps the tiling to itself, implying the grid shown above. Any periodic tiling can be formed by regular placement of a period parallelogram.

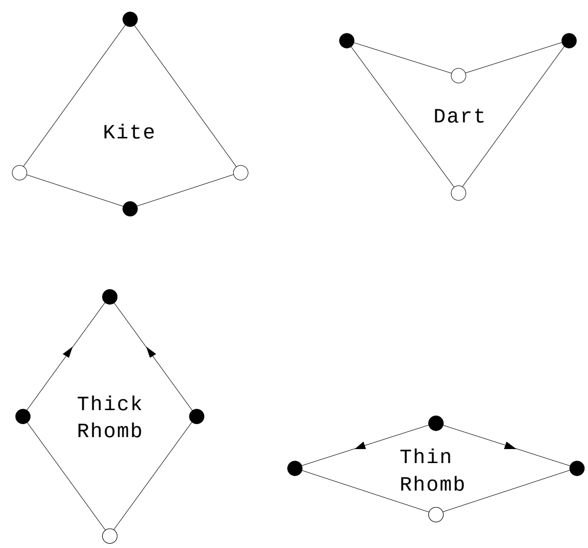

Figure 2.3: Shapes of Penrose's kites, darts, and rhombuses. Penrose discovered two related aperiodic sets of shapes: the kites and darts (above) and the rhombuses (below). Arrows and the black and white dots must match for adjacent tiles.

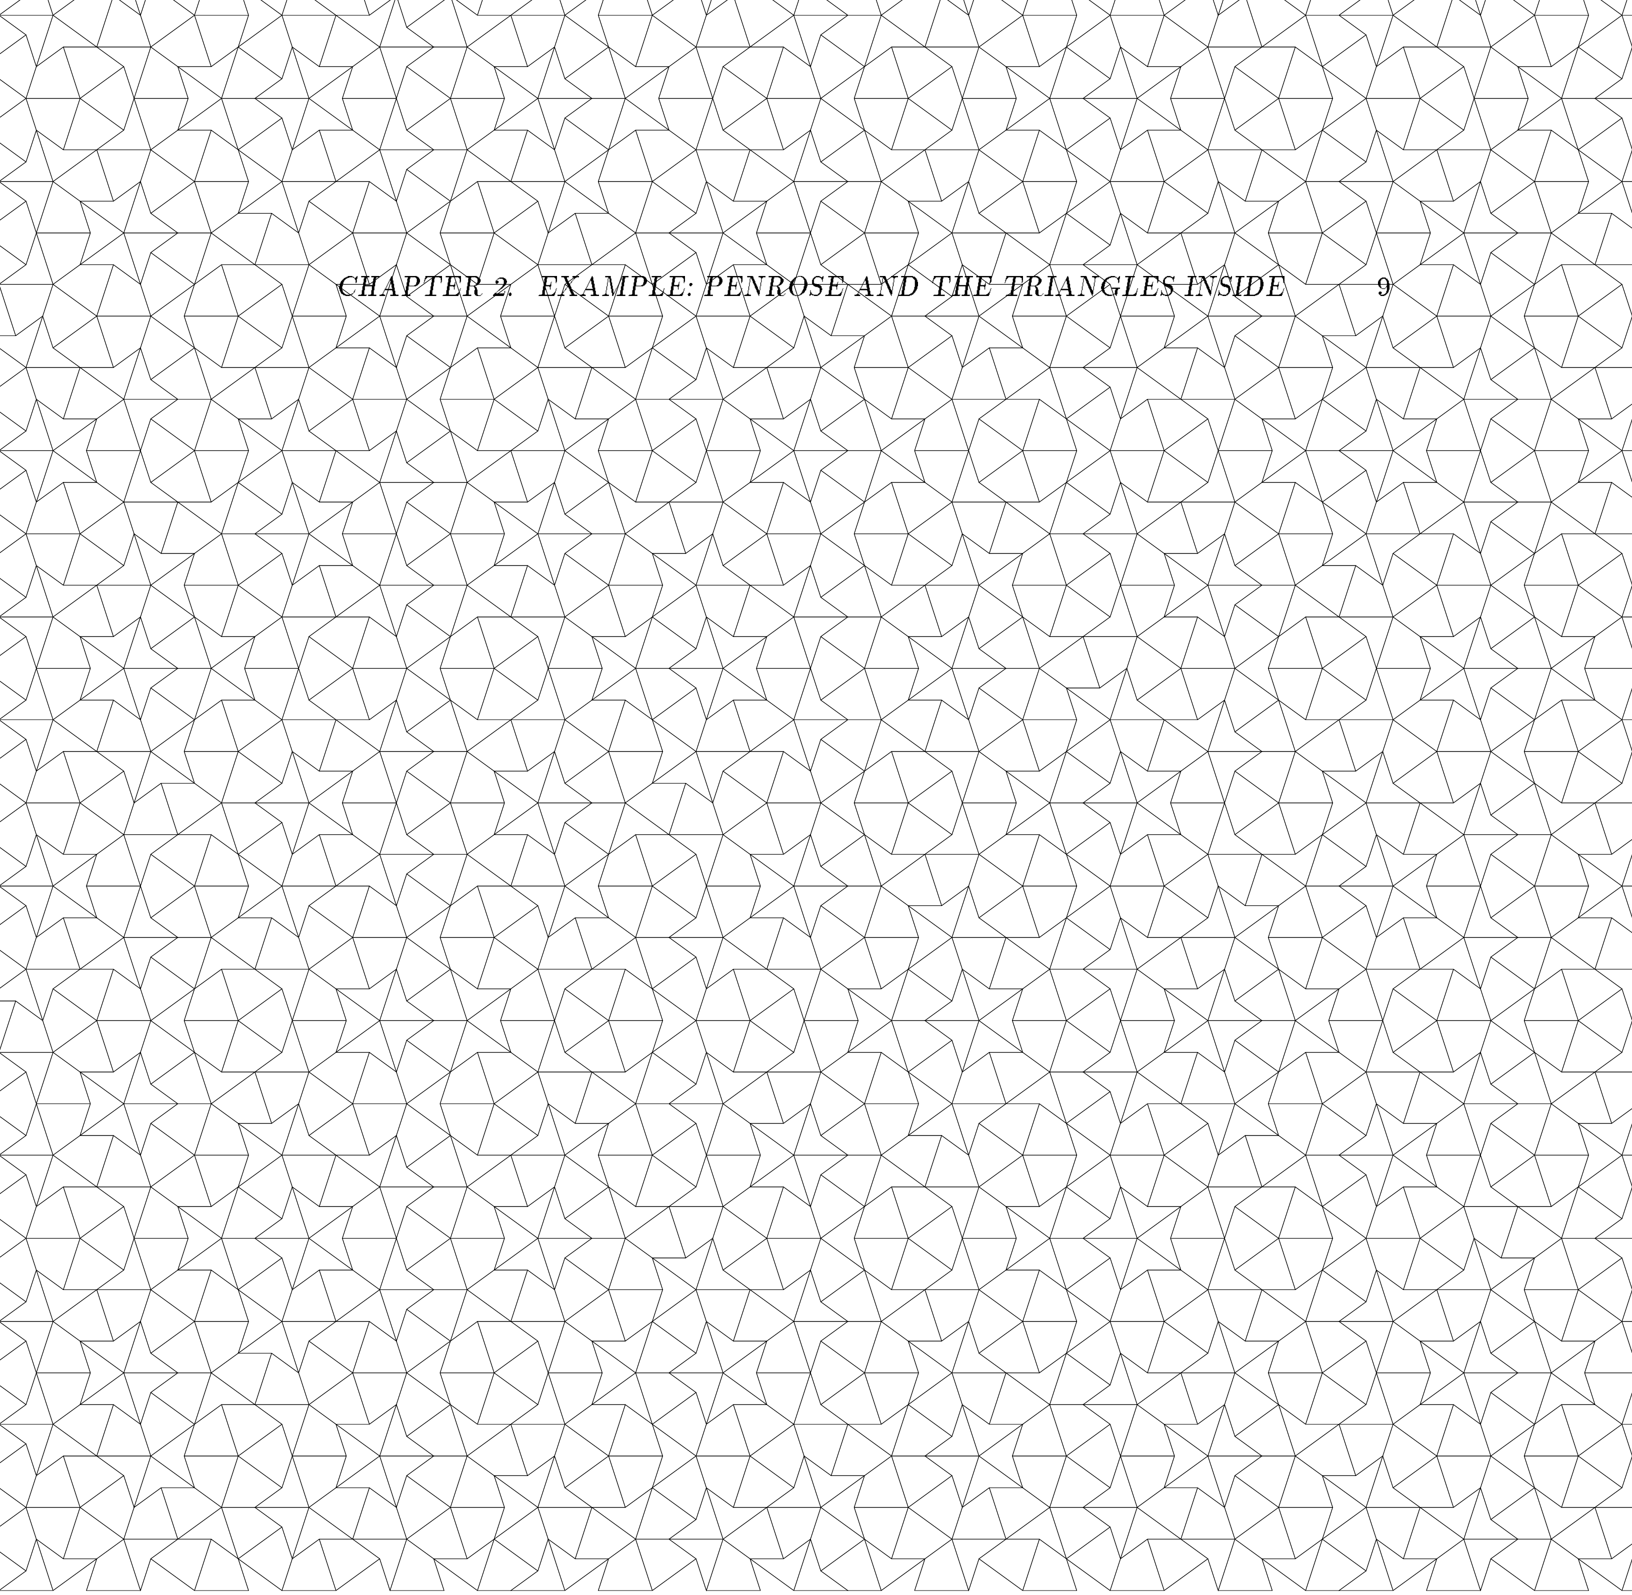

Figure 2.4: An aperiodic tiling by Penrose's kites and darts. Kites and darts are placed according to the matching rules to fill the plane. This tiling is highly structured, but not in a simple, regular way. Since there are symmetries by translation, this tiling is *aperiodic*.

2.1 Meet Robinson's Triangles

Robinson's triangles are formed by cutting Penrose's tiles in half to form the two different triangular shapes seen in Figure 2.5. The two isosceles triangles are used (with different relative sizes) to form two sets: the A-tiles and B-tiles. Again, they are decorated with dots and arrows, and more complicated undecorated tiles can be formed that match equivalently, and this is shown for the A-tiles in Figure 2.6. Each tile may also be turned over and placed. Another way of thinking of this is that there is a right-handed and left-handed version of each tile, making four A-tiles and four B-tiles. Left-handed tiles are shown in grey and labeled with primes (') in Figure 2.5. The matching rules force tiles to pair so that large A-tiles are kites, small A-tiles form darts, large B-tiles form thick rhombs, and small B-tiles form thin rhombs. We assign a symbol to each type of tile by using the letter A/a for A-tiles, B/b for B-tiles, capitalizing based on size, and adding a prime (') for left-handed tiles. Thus the symbol $\mathbf{b}'$ indicates a small B-tile. Many properties of Robinson's triangles are related to $\tau = \frac{1+\sqrt{5}}{2} = 1.618034\ldots$, the golden ratio.

The A-tiles have dimensions and angles as seen in Figure 2.7. **A** has a base of length 1 and legs of length τ and the smaller tile has a base of length τ and legs of length 1. The area of the large tile is τ times greater.

Figure Figure 2.8 gives the dimensions and angles of the B-tiles. **A** and **b** are congruent. **B** and **a** are similar triangles where the sides of **B** are τ times as long and its area is τ^2 times greater.

2.2 Inflation

In any A-tiling or B-tiling, some of the tiles can be joined in pairs, resulting in a larger tiling of the other type. Each joined pair forms a tile equivalent to a large tile of the new type. Some large tiles will not be joined with small tiles and these tiles are already equivalent to small tiles of the new type. The process of attaching together tiles to form a tiling by a set of larger shapes is known as *composition*.

A-Tiles Build Larger B-Tiles It has already been shown above that **A** and **b** are identical. The **A** and **a** can be grouped together to form a tile equivalent to the **B**. To accomplish this, **A** is turned over and its base is attached to **a**. The tile formed by grouping **A** and **a** has an extra arrow and a white dot at its base that can be removed without effect since this edge will only match with itself in either case.

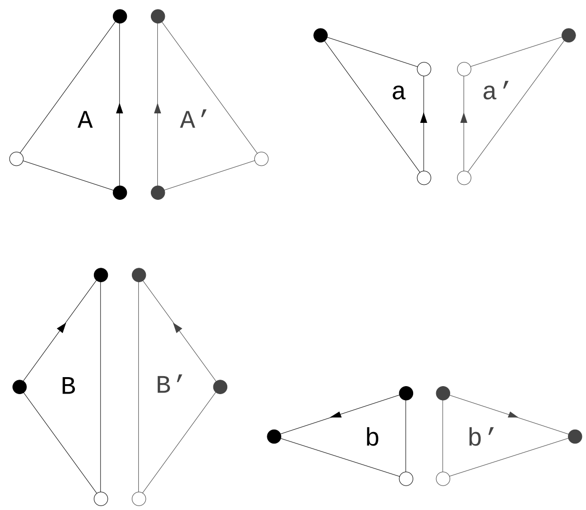

Figure 2.5: Cutting Penrose's tiles in half forms Robinson's tiles. Cutting apart kites and darts forms A-tiles (**A** and **a**). Cutting apart rhombuses forms B-tiles (**B** and **b**). There is a right-hand and left-hand version of each tile and left-handed tiles are indicated with a prime ('). In any Penrose tiling, any tile can be replaced with a right-handed and left-handed Robinson tile of the appropriate type. Due to the matching decorations, in any Robinson tiling the tiles must pair in couples whose union is a Penrose tile.

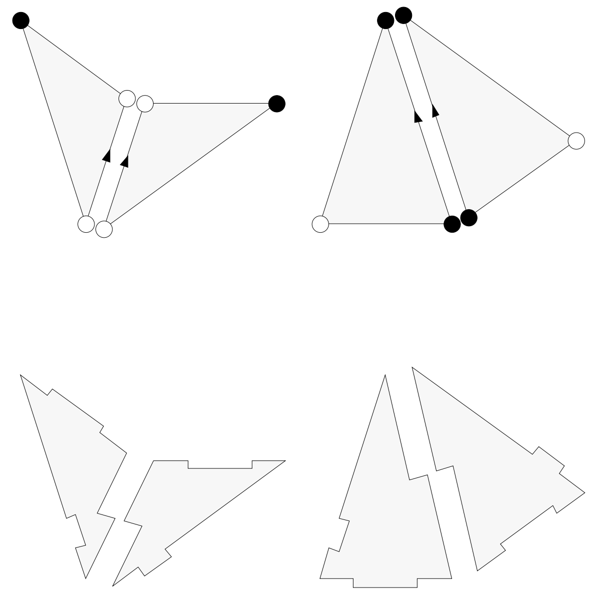

Figure 2.6: Deforming Robinson's tiles eliminates the need for matching rules. The matching decorations for the A-tiles (above) and equivalent undecorated tiles (below) tile in exactly the same way.

B-Tiles Build Larger A-Tiles B-tiles can be attached together to form A-tiles scaled by τ. The **B** and **a** (scaled) are identical except that the color of their dots is reversed. Joining the **B** and **b** along their sides with two black dots results in a tile equivalent to the **A** (scaled). The color of its dots must also be reversed. Reversing the dots of the **A** and **a** is allowable since this has no overall effect on the matching properties of the A-tiles formed. The side with an extra dot is equivalent to a side with an arrow instead, since this side is required to match only with itself and either decoration forces it to match in the same orientation.

These compositions can also be reversed. If every large tile of an A-tiling or B-tiling is divided into two pieces and alterations are made to the dots and arrows as described above then a tiling of the other type results. Repeating this operation will form tilings by smaller and smaller shapes that alternate between A-tiles and B-tiles. If the tiling $\mathcal{T}$ is divided to form $\mathcal{T}'$ with smaller pieces of the opposite

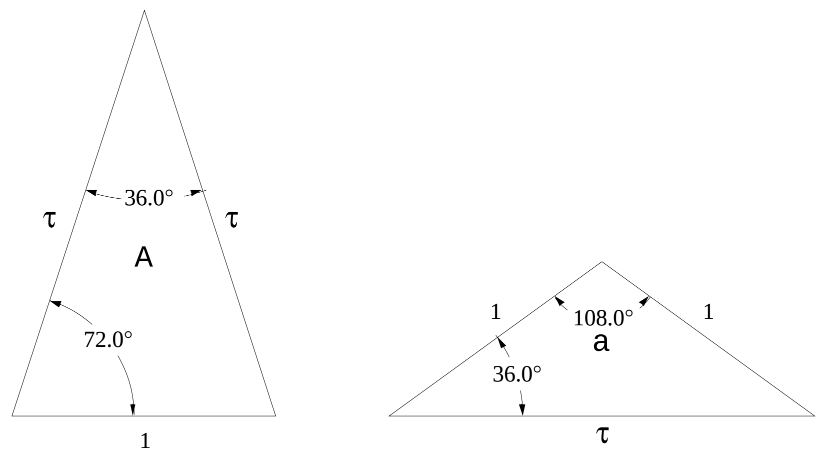

Figure 2.7: Dimensions and angles of the A-tiles.

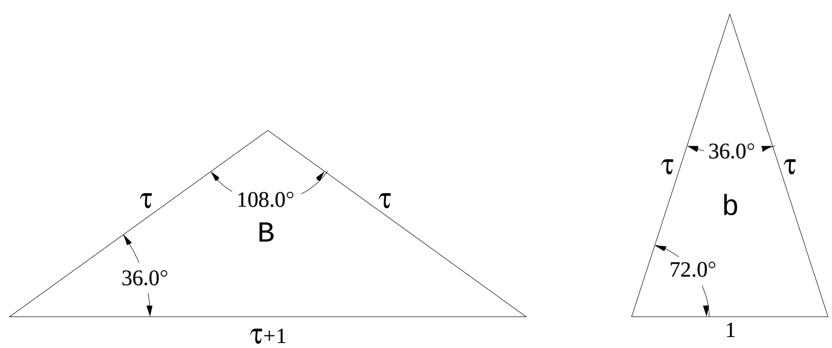

Figure 2.8: Dimensions and angles of the B-tiles.

type then we say $\mathcal{T}'$ is an *inflation* of $\mathcal{T}$ and $\mathcal{T}$ is a *deflation* of $\mathcal{T}'$.[1] Inflation and deflation can be iterated to form tilings by arbitrarily large or small A-tiles or B-tiles. The area of the large or small tiles at successive levels of inflation is smaller by a factor of τ.

Deflation can be used to prove that Robinson's triangles tile the plane. Begin by placing single **A** tile. Inflate the tile, decomposing it into two smaller tiles of the opposite type. Continue inflating the tile, forming increasing numbers of smaller and smaller tiles. After n inflations, a patch of tiles results with tiles that are $\frac{1}{\tau^n}$ times as large as the originals. If this patch is scaled by τ^n, a patch of tiles results with shapes of the original size and an area τ^n times as great as the seed tile. Increasingly large values of n can be used to produce increasingly large patches of B-tiles or A-tiles for odd or even n. Since this method can be used to form arbitrarily large patches made up of arbitrarily many original unit tiles, it follows that the plane can be tiled with either A or B-tiles.[2]

2.2.1 Relative Frequency of Tiles

The **B** tile is formed from an **A** and an **a** while the **b** is formed from a single **A**. If even larger A-tiles are then formed from these B-tiles then the large **A** will contain three original unit A-tiles and the large **a** will contain two. At each successive level of inflation, the new large tile will contain a number of unit tiles equal to the sum of the unit tiles contained in both the large and the small tiles from the previous level of inflation. The new small tile will contain a number of unit tiles equal to the number of unit tiles in the large tile from the previous level of inflation. If L_n and S_n are the number of unit tiles contained in single large and small tiles at n levels of inflation then:

$$\begin{aligned}
S_n &= L_{n-1} \\
L_n &= L_{n-1} + S_{n-1} \\
S_{n-1} &= L_{n-2} \\
L_n &= L_{n-1} + L_{n-2}
\end{aligned}$$

Since at the first level of inflation both tiles are unit tiles, $L_1 = S_1 = 1$. L_n is the n-th Fibbonacci number (fib_n) where $fib_1 = 1$, $fib_2 = 2$, $fib_3 = 3$, $fib_4 = 5$, $fib_5 = 8$,

[1] Inflation and deflation are properly defined in combination with a scaling to produce tiles of the original size. We neglect this detail to simplify later discussion and because illustrations of inflation and deflation are easier to visualize when smaller tiles lie within each other.

[2] This is proven by use of the Extension Theorem as stated in Grunbaum.[GS87] "Let L be any finite set of prototiles, each of which is a closed topological disk. If L tiles over arbitrarily large circular disks D, then L admits a tiling of the plane."

and $fib_n = fib_{n-1} + fib_{n-2}$. S_n is fib_{n-1}. In a complete tiling of the plane, the ratio of large to small tiles is $\lim_{n\to\infty} \frac{fib_n}{fib_{n-1}}$. This ratio is equal to τ.[GKP94]

2.2.2 Unique Inflation

Above, we have defined inflation and deflation as opposites. If this is to make sense, we need to show that there is only one way to inflate or deflate any particular tiling. We call this property *unique inflation*. Inflation is defined above as a specific way of replacing each tile with a patch of smaller tiles of the opposite type. We need to show that all tilings can also be composed to form larger tilings of the opposite type by attaching pieces in this way, and that only one such composition is possible.

The leg of an **a** that has no arrow can only match the base of an **A** and thus in an A-tiling, every **a** must border an **A** in this way.[3] This connection is the only way in which a large B-shaped tile can be formed by grouping an **a** with adjacent tiles. The tiles of an A-tiling must be composed to form a B-tiling by connecting every **a** to an **A** in this way. Each pair becomes a **B** and the remaining **A** tiles become **b** tiles. The new B-tiles will all match up properly because the A-tiles they are formed from matched. Given an A-tiling, there is exactly one way to connect tiles in order to form a B-tiling.

Similarly, there is only one way to form an A-tiling from a B-tiling. The side of a **b** with an arrow must border a **B** so that these tiles compose an **A** and this is the only way that an **A** can be formed.[4] Connecting each **b** to the appropriate **B** forms the **A** tiles of the next level of inflation while the leftover **B** tiles become **a** tiles.

Now we can prove that Robinson's tiles are aperiodic. A tiling is periodic if some nontrivial translation exists that maps the tiling to itself. A set of tiles is aperiodic if no periodic tiling can be constructed from the tiles.

Assume that there exists a periodic tiling $\mathcal{T}$ of A-tiles or B-tiles. Then there is a translation T of distance d that maps $\mathcal{T}$ onto itself. Deflate the tiling by a sufficient number of levels, forming the tiling $\mathcal{T}^*$ so that the shortest side of any tile in $\mathcal{T}^*$ is longer than d. Now pick a point v_1 at a corner of some tile of $\mathcal{T}^*$. Translate v_1, obtaining $v_2 = T(v_1)$. This point does not lie on a vertex of the inflated tiling. Since the tiling is periodic, the entire uninflated tiling surrounding v_1 is identical to that surrounding v_2. Inflating $\mathcal{T}^*$ eventually forms $\mathcal{T}$ and the translation of $\mathcal{T}^*$, $T(\mathcal{T}^*)$

[3] If this edge were to match with another **a** a flaw in the tiling would arise because the other leg can only match with an **a** and these three tiles leave a 36 degree wedge that cannot be filled.

[4] Every **b** tile is adjacent to another at its base, forming a rhombus. If a **b** was adjacent to another **b** along the side with an arrow then two such rhombs would be adjacent in such a way that a wedge with three black dots would remain. The arrows on the wedge point outward and no tiles can fill this space.

also inflates to $\mathcal{T}$. Thus, another different inflation of the tiling can be constructed with the same structure as the inflated tiling surrounding point v_1 placed instead around point v_2. However, this means that there are two possible deflations of $\mathcal{T}$ and this contradicts unique inflation. Therefore, no such periodic tiling could exist and A-tiles and B-tiles are aperiodic.

Another interesting property of Robinson's triangles is *local isomorphism*. Any finite patch of tiles in an A-tiling is congruent to infinitely many locations in every A-tiling. This property also holds for B-tiles. Local isomorphism is related to the concept of vertex configurations. A *vertex configuration* is a patch of all the tiles that meet around a single corner. If a single tile is inflated a sufficient number of times then it will contain all possible vertex configurations. In Theorem 3 we show that this means that all tilings are locally isomorphic.

2.3 Index Sequences

An *index sequence* is used to classify a tiling as the limit of an infinite sequence of inflations. Our exposition follows very closely that of Grünbaum and Shephard [GS87] and all results in this section are stated there as well. Lunnon [LP87] applies index sequences to a more general class of tilings.

In Appendix A, Figure A.1–Figure A.10 show a series of deflations of a tiling. One of the tiles in the first panel is specially marked and this mark is place at the same position in each panel of the series. The index sequence of the marked tile is the type of tile (large or small) containing it at each level of inflation. We write 0 for large and 1 for small and alternating digits represent A-tiles and B-tiles. Thus, 0001010101...is the address of the marked tile in the first tiling.

No index sequence contains two consecutive 1's because small tiles are always found within large tiles when a tiling is deflated. Since there is never more than one way that a tile of a particular type can be contained within a tile of another type, different tiles have the same address only if they are located within an infinite series of different, nested, increasingly large tiles of the same types. Informally, the limit of such a series of tiles is an infinitely large tile composed from unit tiles the size of the indexed tile. More than one infinitely large tile can occur in a tiling if the nested tiles grow as a half-plane or fan from a point as shown in Figure 2.9. A series of tiles that grows this in this way can only be joined to tile the plane if tiles of the same type and opposite handedness are adjacent. These configurations result in tilings with a line of reflection or fivefold rotational symmetry, while if one infinite tile fills the plane the resulting tiling has no symmetries. We work out the circumstances in which these cases apply below.

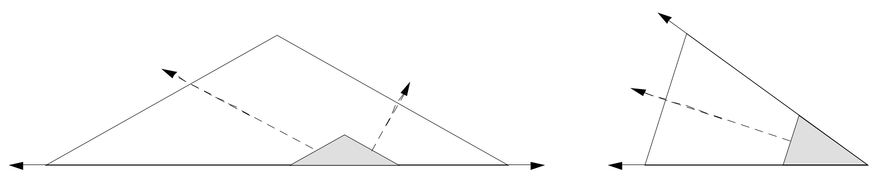

Figure 2.9: Infinite subregions of the plane generated by deflation. The shaded triangle lies inside a larger, similar tile of the same orientation that contains it in a deflated tiling. When the grey tile lies on one edge of the deflated tile (left), deflations fill a half-plane. When the grey tile lies on two edges of the deflated tile (right), deflations fill a fan bounded by two rays.

In addition to specifying the location of 1, 2, or 5 tiles (depending on the symmetry group) in a tiling, an index sequence completely determines the tiling itself. Any two tiles in the same tiling must be contained within the same tile or the same tile in different symmetric regions after a sufficient number of inflations.[5] Thus, two index sequences differ in a finite number of bits if and only if they identify tiles in the same tiling. This allows a simple proof that there are uncountably many different tilings.[6] We will see later that the symmetry type of a tiling can be identified from its index sequence.

2.4 Addresses

We are interested in representing finite regions of tiles called *patches*. We could do this by growing the dual graph of a patch directly, as described in Section 5.2, but this graph could easily become prohibitively large. There is only one way that one type of tile can be contained in any other, so an index sequence can be viewed as a series of instructions for forming increasingly large tiles containing the indexed tile. The first digit of an index sequence specifies the type of a single tile and each later digit nests the patch described by the previous digits within a larger patch.

We examine index sequences of finite length that we call *addresses*. An n-digit

[5] We discuss symmetry in Section 2.7.

[6] Fleshing out this proof requires careful examination of what it means for tilings to be different and how this actually occurs for tilings as the limit of index sequences. Since index sequences are infinite sequences of 0 and 1 (without two consecutive 1's), they correspond to the reals.

address represents the location of a tile within a patch containing fib_n or fib_{n-1} tiles if its last digit is 0 or 1 respectively. Therefore, the length of an address is logarithmic base τ in the number of tiles. Since no addresses contain consecutive 1's, valuing their digits according to the Fibonacci numbers uniquely maps addresses of length n to integers between 0 and $fib_{n+1} - 1$. We will often write addresses as strings of tile types rather than 0's and 1's and by convention we use **A** or **a** as the first digit, representing only A-tilings. Thus 01001001 is written **AbABaBAb**.

Addresses carry information about the patch a tile is within as well as its location. Say A is the address of a tile in the patch P. Removing the first digit of A results in the address of a tile in the deflation of P that contains A. Similarly adding a digit to the front of an address results in the address of a tile in the inflation of P that is contained by A. The choice to add a 1 or 0 places the resulting address inside the large or small tile inflated from the large tile addressed by A. If A addresses a small tile, there is no choice and both addressed tiles are congruent.

The patch P is congruent to some number of inflations, n, of one of the A-tiles. Adding a bit to the end of P represents a patch congruent to $n + 1$ inflations of one of the A-tiles. The new tile is located within a patch congruent to P that lies inside this inflation. If the last digit of A is 0 then adding a 1 to the end of A represents a small tile at this next level of inflation, which is identical to P. If a 0 is added then the patch represented by A is located within a large tile at the next level of inflation. This large tile is formed from a large and small tile from the level of inflation of A, and P is congruent to one of these. We call adding a bit to the end of an address *extension*. Small tiles can only be extended in one way, but large tiles can be extended in two different ways. Following a large tile by a small one does not make the patch larger, but the next extension after this will. Extension produces a patch at the next level of inflation that contains the current patch and whenever a 0 is added this results in a larger patch. A 0 is added with at least every other extension. If an address ends in 0 then there are two different ways that it can be extended and choosing one of them determines the structure of an inflated patch surrounding the current patch. Extension is used to simulate an infinite tiling by extending finite patches whenever necessary.

2.5 Adjacency

In order to tour tilings using addresses, we need a technique for finding which addresses are adjacent. The sides of each tile are labeled left (l), right (r), and middle (m). We call the base of a triangle the middle side and if the base is downward then the left and right legs are on the left and right. An opposite scheme is used if the tile was turned over before being placed. In other words, the left and right sides of a left-handed tile are the mirror-image of a right-handed tile. We want to know which addresses are adjacent and the sides by which they are adjacent.

If a move leaves a tile t with address $A = (a_1, a_2, \ldots, a_n)$ by side s, we want to determine the address A' of the tile t' that the move enters. We do this by examining the consequences of the move for successive digits of A. Figure 2.10 shows that a move may result in a series of crossings of the sides of tiles at several levels of deflation. A move that passes across an edge of the n-deflated tiling may result in a change to the value of digit $n+1$ of the address. Whenever an n-deflated tile is crossed, i-deflated tiles are also crossed for $i < n$. We first examine a_1 and a_2 and assume that the first digit of the address represents a type of A-tile. a_1 gives the type of t (**A** or **a**) and a_2 gives the type of the larger, once-deflated tile that contains it (**B** or **b**). There is only one way that any particular A-tile can be contained in any particular B-tile and we can use this relationship to determine the effect on a_2 of the move leaving digit a_1 by its side s. The move can be of two types. An *internal* move leaves a_1 and crosses to another tile within a_2. An *external* move leaves a_1 and also leaves a_2 by one of its sides e_2. If we encounter an external move then we repeat the process. The once-deflated tile represented by a_2 lies inside the tile represented by a_3 in some particular way and leaving a_2 by side e_2 is either an internal move or is an external move that leaves a_3 by some side s_3. This process is continued until an internal move is reached. Ignoring symmetric tilings for now, both the origin and destination tiles (t and t') of the move must be contained inside the same tile at some level of inflation and the two addresses are the same from this digit on.[7]

Figure 2.11 can be used to determine the result of any external or internal move. A table of the ways in which any tile a can be related to the deflated tile $a*$ that contains it appears in Figure 2.12. An entry of the form $as \rightarrow a^*s^*$ is an external rule. Leaving a by s leaves a^* by s^*. An internal rule has the form $as \rightarrow a's' \subset a^*$ and indicates that leaving a by s remains within a^* and crosses into the tile a', entering by its side s'.

Using the table, we can follow a series of external moves. We will discover that leaving the tile addressed by A along side s implies a series of external moves, and eventually the digits a_l and a_{l+1} will fit an internal move of the form $a_ls_l \rightarrow a'_ls'_l \subset a_{l+1}$. This means that leaving A by s crosses the side s_l of the $(l-1)$-deflated tile represented by a_l and enters the $(l-1)$-deflated tile represented by a'_l through its side s'_l. If $A' = (a'_1, a'_2, \ldots, a'_n)$ is the destination address of the move then $a_i = a'_i$

[7]Two adjacent tiles across a line of symmetry are in different tiles at every level of inflation and this process breaks down.

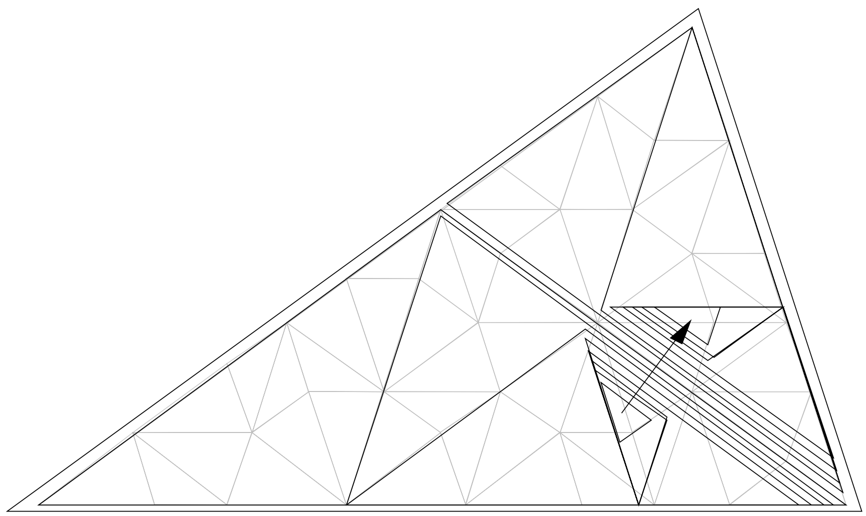

ABaBaBABA -> ABaBaBAbA

Figure 2.10: A move between two tiles that results in a series of crossings at lower levels of inflation. The move exits tile **ABaBaBABA**and enters **ABaBaBAbA**. This aerial view raises the higher levels of inflation containing each tile by slanting them away from each other. Since the adjacent tiles are in different sections of the **A** tile containing the entire patch, this crossing requires traveling down and up a deep valley.

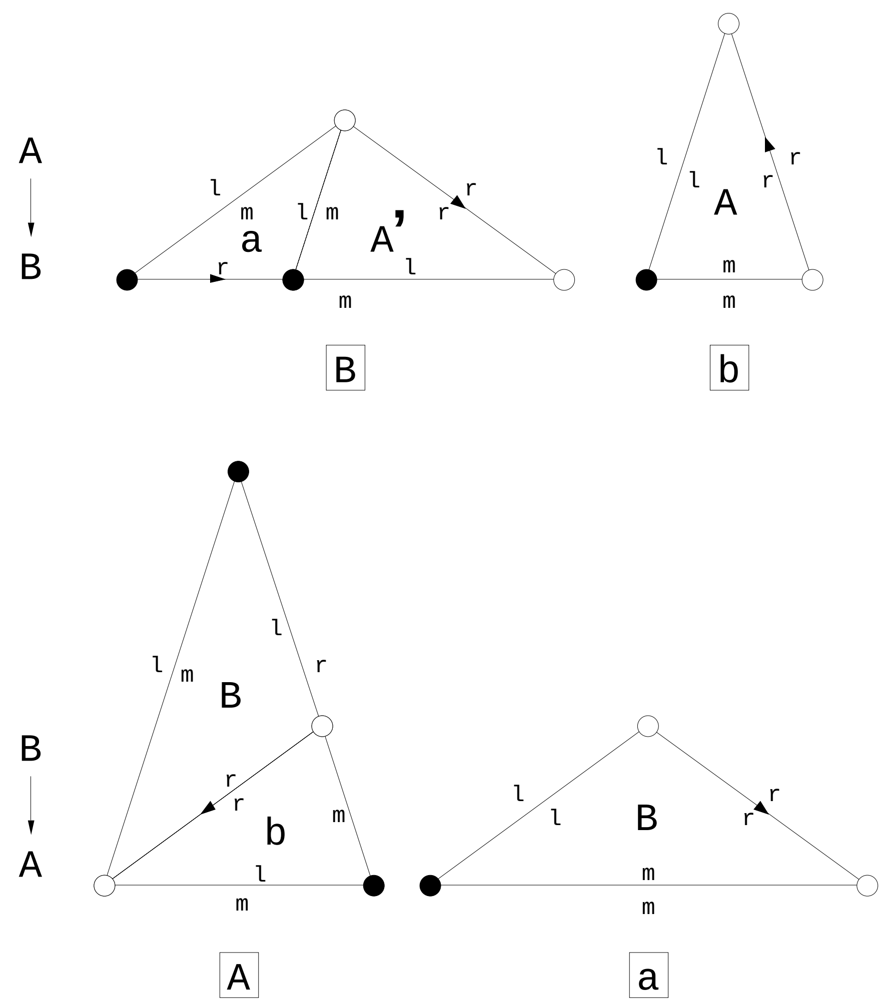

Figure 2.11: Relation of a tile's side to the side of the inflation patch containing it. Left (l), middle (m), and right (r) sides of the tiles are labeled. Each of the four patches shown is the way that the type of tile labeled inside the square decomposes in order to form the next level of inflation.

a	m	$\rightarrow$	**B**	l		
a	l	$\rightarrow$	**A**	m	$\subset$	**B**
a	r	$\rightarrow$	**B**	m		
A	m	$\rightarrow$	**b**	m		
A	l	$\rightarrow$	**b**	l		
A	r	$\rightarrow$	**b**	r		
A	m	$\rightarrow$	**a**	l	$\subset$	**B**
A	l	$\rightarrow$	**B**	m		
A	r	$\rightarrow$	**B**	r		
b	m	$\rightarrow$	**A**	r		
b	l	$\rightarrow$	**A**	m		
b	r	$\rightarrow$	**B**	r	$\subset$	**A**
B	m	$\rightarrow$	**a**	m		
B	l	$\rightarrow$	**a**	l		
B	r	$\rightarrow$	**a**	r		
B	m	$\rightarrow$	**A**	l		
B	l	$\rightarrow$	**A**	r		
B	r	$\rightarrow$	**b**	r	$\subset$	**A**

Figure 2.12: External and internal rules for Robinson's Triangles. An external rule $as \rightarrow a^* s^*$ indicates that leaving a by s leaves a^* by s^*. An internal rule $as \rightarrow a's' \subset a^*$ indicates that leaving a by s remains within a^* and crosses into the tile a', entering by its side s'. Note that **Ar** and **Bm** each appear on the right-hand side of two external rules.

for $i > l$.

Now we work backwards, filling in the digits a'_i for $1 \leq i < l$. Finding a rule that matches $a'_i s'_i \rightarrow a'_{i+1} s'_{i+1}$ identifies the new value of digit i of A'. In this manner, we can fill in the remaining digits of A' from last to first. However, there is a special case to consider. All types of tiles and sides appear as the right-hand side of some external rule, but two of these (**A**r and **B**m) appear on the right-hand side of more than one rule and we need to show how to choose which rule to apply. Figure 2.13 shows the alignment of tiles for which this is possible. Whenever this situation occurs the digit (**A**r or **B**m) represents entry into a large tile along a long side, and two smaller tiles at the next level of inflation have edges on that long side. It happens that according to the matching rules, sides **A**r and **B**m can only occur next to their reflections—the same type tile with opposite handedness. This means that whenever one of these sides is encountered while working backwards then the remaining digits of A' are identical to those of A in the same positions.

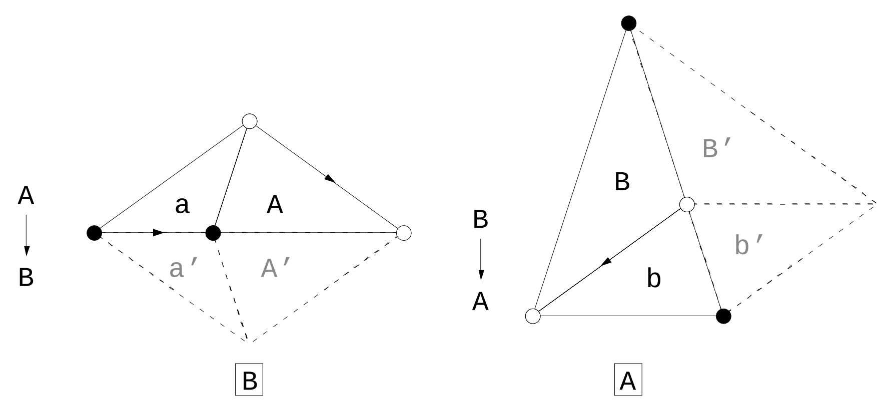

Figure 2.13: Symmetrical matching disambiguates non-unique backwards moves. The rules **bm** $\rightarrow$ **Ar** and **Bl** $\rightarrow$ **Ar** and the rules **ar** $\rightarrow$ **Bm** and **Al** $\rightarrow$ **Bm** describe sides of inflation patches that contain two smaller tiles. Since these sides only match with the opposite-handed tiles, bits of the addresses preceding one of these rules are identical.

If the move crosses outside of the patch represented by A then we must extend this patch until A does contain the destination. In practice, we do this while following external rules. If we determine that a_n is left by side e_n, but A only has n digits then we extend A by adding another digit and proceed as normal. Then if we request the adjacency of A, the result, A' will be a longer address because the algorithm used extensions of A to compute A'.

This algorithm can be used to compute the adjacencies of any address and requires time proportional to the length of the resulting address. It can be used to explore an unlimited expanse of Robinson triangles. Additional bits are added as exploration proceeds increasingly far from the origin and each additional digit extends an address to encompass a patch containing τ times as many tiles.

2.6 Orientation and Handedness

The above algorithm can be extended to maintain information about the handedness and rotational orientation of tiles that are visited. A tile can have one of ten different orientations. We name these orientations with the compass-like directions N, NE, EN, ES, SE, S, SW, WS, WN, and NW. The orientation of a tile is the direction in which the wedge of its two legs points.

Every pair of adjacent tiles has differing handedness (one of the tiles was turned over when placed) and the orientation of any tile can be determined from the orientation of one of its neighbors and the sides by which the two tiles are adjacent. Figure 2.14 illustrates these relationships and Figure 2.15 gives a set of rules for determining the handedness and orientation of a tile from one of its neighbors and the sides they share. Thus, if the origin of an exploration is assigned an address, handedness, and orientation then our data structure can produce the address, handedness, and orientation of the other tiles encountered as exploration proceeds as a series of moves to adjacent tiles.

Using the address of a tile within an inflated tile with known orientation and handedness, we can deduce the orientation of the unit tile relative to the inflated tile. Figure 2.11 can be used to deduce the orientation of tiles formed by inflation from others. These rules are given in Figure 2.16 and they relate the orientation of tiles at different levels of inflation.

2.7 Symmetry

So far, we have left an important detail unaddressed. If certain index sequences are used to extend addresses then leaving some addresses by certain sides results in infinite series of external moves. The situation in which this occurs corresponds to the crossing of a line of symmetry in the tiling. If two tiles are adjacent across a line of symmetry then no single deflated tile will contain both of them. This problem also blights our approach for representing a generalized class of tilings and is discussed at length in Section 6.1.

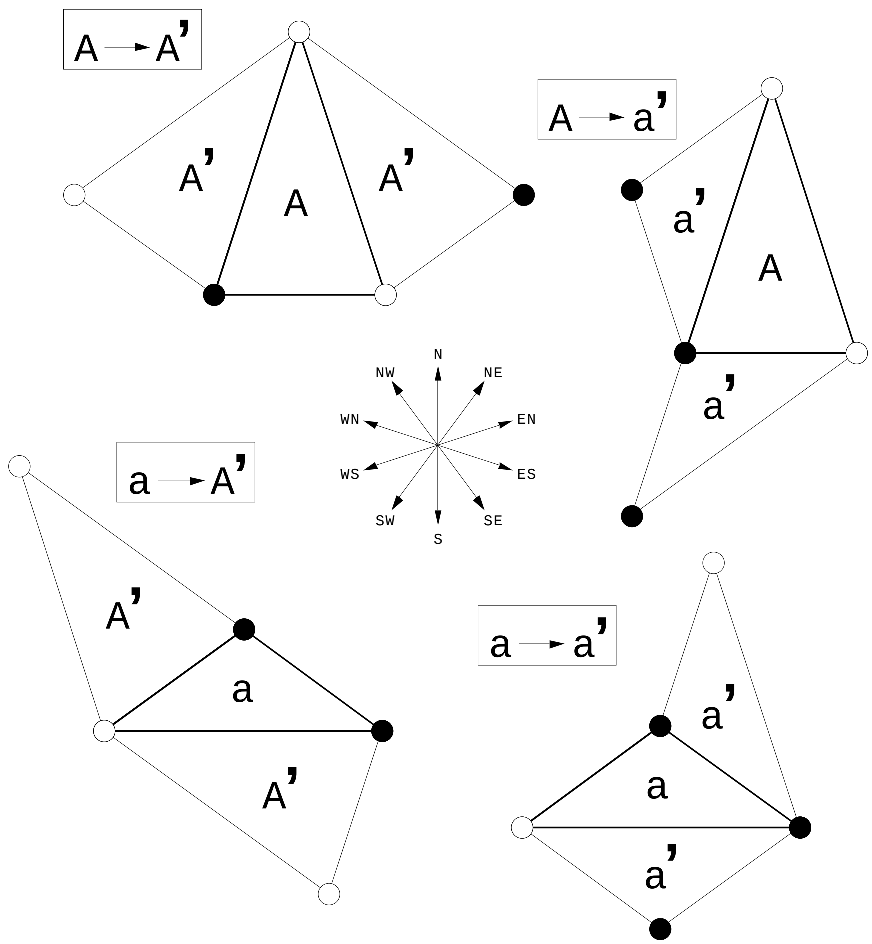

Figure 2.14: Potential neighbors of an A-tile. **A** and **a** are shown neighboring other tiles in every possible way. Note that the only options are **A**l and **a**m and that adjacent tiles have opposite handedness.

A	r	N	$\rightarrow$	**A**$'$	r	NW
A	l	N	$\rightarrow$	**A**$'$	l	NE
A	l	N	$\rightarrow$	**a**$'$	m	WN
A	m	N	$\rightarrow$	**a**$'$	l	NW
a	l	N	$\rightarrow$	**A**$'$	m	NW
a	m	N	$\rightarrow$	**A**$'$	l	WN
a	m	N	$\rightarrow$	**a**$'$	m	S
a	r	N	$\rightarrow$	**a**$'$	r	WS

Figure 2.15: Orientation rules for adjacent tiles. The table contains a rule of the form $aso \rightarrow a's'o'$ if a tile of type a and orientation o can neighbor a tile of type a' and orientation o' so that they are adjacent along sides s and s'. Primes (') after the types of tiles indicate that they left-handed. Rules are only given for right-handed tiles with the orientation N. Any other orientation or handedness can be handled by reversing the handedness of both the tiles of a rule or rotating them.

B	ES	$\rightarrow$	**A**	N
b	WS	$\rightarrow$	**A**	N
B	N	$\rightarrow$	**a**	N
A$'$	ES	$\rightarrow$	**B**	N
a	SE	$\rightarrow$	**B**	N
A	N	$\rightarrow$	**b**	N

Figure 2.16: Orientation rules for inflations of tiles. The table contains a rule of the form $ao \rightarrow a'o'$ if a tile of type a and orientation o is a tile from the inflation of a tile of type a' and orientation o'. Primes indicate left-handed tiles. The rules are only given for large right-handed tiles with the orientation N but are analogous to rules with different handedness and other orientations.

To examine sequences of external rules, we form a directed graph called the *edge extension roadmap*. Nodes of the graph are pairs of tile types and sides. For each external rule $as \rightarrow a^*s^*$, an edge is placed from node as to node a^*s^*, forming the structure shown in Figure 2.17. All infinite sequences of external moves correspond to a path through the edge extension roadmap and all paths through the roadmap generate (using the tile types of each node) an index sequence for a tiling with at least one line of symmetry. All infinite paths through the roadmap pass through the node Bm infinitely many times. Neglecting the first few digits before Bm is reached, all paths can be formed by repeated concatenation of one of the three cycles from Bm, as shown in Figure 2.18. These three cycles extend patches as shown in Figure 2.19. A legal address (no consecutive 1's) is in the regular language $(\mathbf{A}(\mathbf{bA})^*\mathbf{B}|\mathbf{aB})^*$. An address that contains a line of symmetry is formed by a legal prefix followed by one of the three loops and thus has the form $(\mathbf{A}(\mathbf{bA})^*\mathbf{B}|\mathbf{aB})^*(\mathbf{A}|\mathbf{a})(\mathbf{BA}|\mathbf{BaBABa}|\mathbf{BAbAbABa})^\infty$. Three special tilings—the cartwheel, the sun, and the star—have addresses of $(\mathbf{AB})^\infty$, $(\mathbf{ABaB})^\infty$, and $(\mathbf{aBAB})^\infty$. We can see that the cartwheel is a repetition of the first cycle, and the sun and star alternate between the first and second cycles. The index sequence of a mirror-symmetric tiling is made up of loops that extend patches along a half-plane and five-fold rotationally-symmetric index sequences have patterns that extend patches in infinite fans as seen in Figure 2.9.

Figure 2.20 shows a **B** tile inflated 9 times and two **B** tiles with the same orientation and handedness as the resulting patch appear at two of its corners. No such tiles appear at the previous levels of inflation. The addresses of the two tiles are the repeating digits of the sun and star index sequences. Whenever **B** occurs in an address but is not followed by the address of one of these tiles then the address corresponds to some position other than the corners. If the sequence disagrees with the three roadmap cycles in infinitely many places this corresponds to selection of a tile within **B** and the tiling grows to fill the plane and contains no lines of symmetry. Figure 2.19 shows each of the three cycles from **B**b in the edge inflation roadmap by finding tiles corresponding to the addresses of within inflations of a **B** tile. If after some prefix, the index sequence is formed from the three roadmap cycles then this corresponds to selection of tiles along the middle edge of **B** and the tiling grows as a half plane, forming a line of mirror symmetry. The sun and star are formed from different phases of the alternating pattern of cycle 1 and cycle 2 and are the only tilings with five-fold symmetry. The sun and star have 10 lines of mirror symmetry as well.

We can incorporate one of these index sequences into our addressing scheme as follows. Adjacent tiles across lines of symmetry have the same address. Since any address identifies a tile in each region of symmetry, we store with the address a

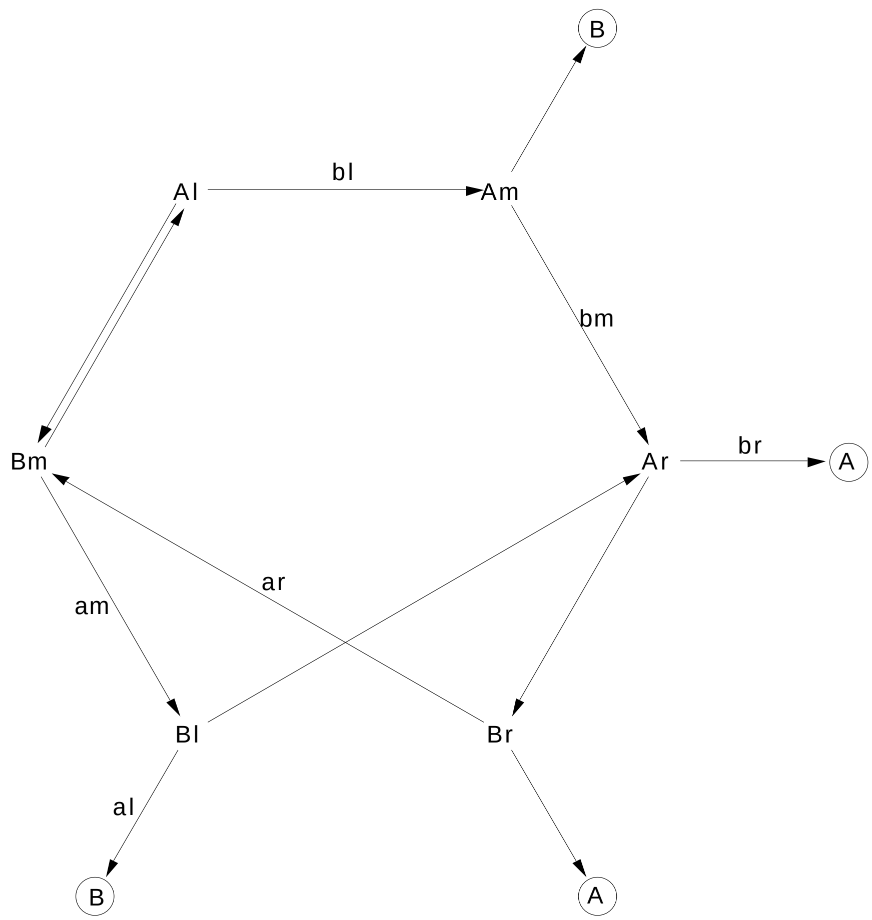

Figure 2.17: Roadmap for Robinson's tiles. Nodes without circles are the sides of tiles and transitions between them are external rules. Any tiling that contains a line of symmetry has an index sequence that is an infinite tour within this graph. Circled nodes and transitions to them are internal rules that exit such a pattern.

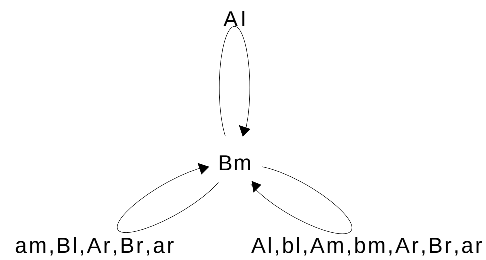

Figure 2.18: Symmetry cycles. Only three cycles are possible from the state **B**m in Figure 2.17.

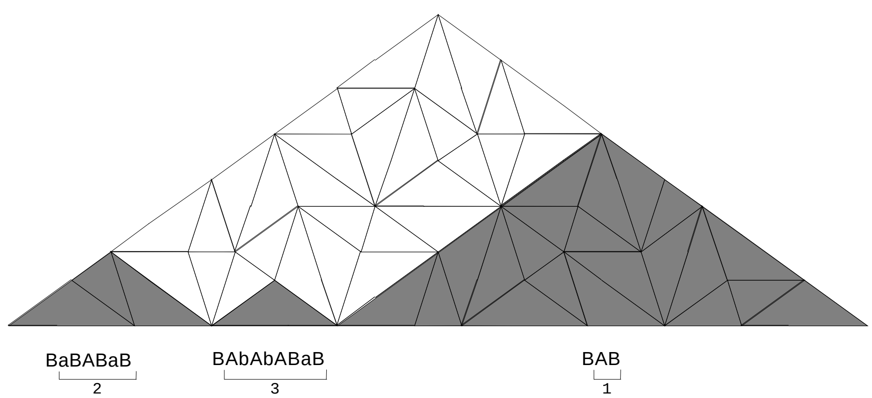

Figure 2.19: Location of each symmetry cycle within an inflated patch. Each cycle in Figure 2.18 descends into one of the grey sections inside a **B** tile. Inflations that follow these cycles will grow a half-plane along the **B**m side or a fan along one of its corners.

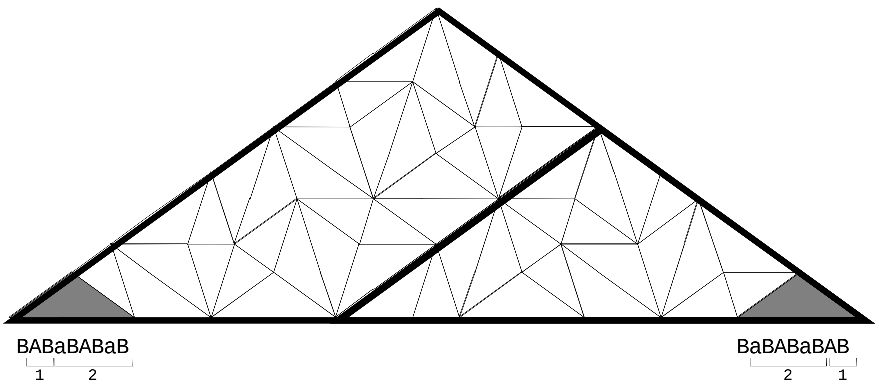

Figure 2.20: Address patterns with five-fold symmetry. When a **B** tile is inflated 9 times, two **B** tiles with the same orientation and handedness as the entire patch appear at the two shaded corners. These addresses are formed by alternations of the first 2 symmetry cycles. Since any other addresses following a **B** tile in this way correspond to tiles that are not at these corners, alternation of these two cycles is the only way to form tilings with five-fold rotational symmetry.

special digit that is the location of the tile in one particular region of the tiling. This is a number between one and five for five-fold rotationally symmetric tilings such as the sun and star and is binary for mirror-symmetric tilings such as the cartwheel. It is not possible to tell when a move crosses a line of symmetry unless a pattern used to extend addresses is known. If this pattern is known, moves that cross lines of symmetry can be identified in the roadmap. If the pattern is a repeated cycle, then moves that cross a line of symmetry can be identified with the parser shown in Figure 2.21. This analysis can be performed in advance of actual use of particular addresses. When the adjacency algorithm encounters one of these moves, it immediately terminates and returns the same address it was given. When this occurs, the special digit is changed to indicate a new region of the tiling. In five-fold symmetric tilings, we examine the pattern of digits in the infinite address and determine the handedness of the **B** tile for the **B**m that is reached at the end of each occurrence of the second roadmap cycle. As can be seen in Figure 2.20, if this tile is right-handed then the move is clockwise around the region of symmetry.

```
 1       2          3
BA | BaBABa | BAbAbABa
12   123456   12345678
```

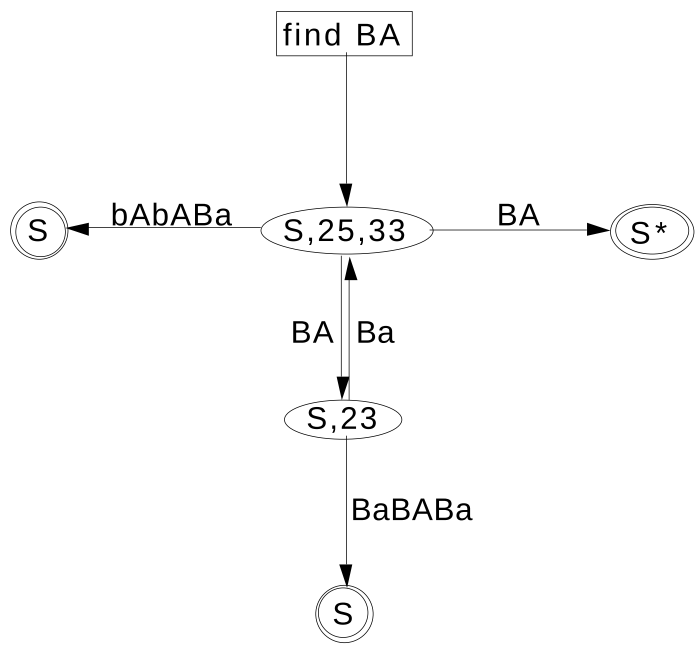

Figure 2.21: A parser to identify the beginning of loops of symmetry. Positions within the three loops are represented with pairs of numbers. The first number is the loop (1–3) and the second is the position of one of the digits of the loop. States indicate the set of such locations at which the string could possibly be located. State S indicates a position after finishing one of the loops and before starting a new one. A state with a single S is the goal. The state S' indicates that the previous state *was* the start and we can move one step back to find that position. Input to the parser is an infinite string of digits from an index sequence that is formed by a repeated cycle. We begin at the top by finding **B** followed by **A**. We feed the parser digits until a final (circled) state is reached or until digits do not match a transition. If no transition is found for some digits then the tiling formed is not symmetric.

Chapter 3

Three Dimensions: Danzer's ABCK Tetrahedra

The concept of a tiling can be extended to three dimensions. Tiles are volumes that match together face-to-face rather than edge-to-edge. A tiling is a placement of these volumes that fills space with no overlapping tiles. Now we adapt our data structure to a set of three-dimensional tiles discovered by Danzer[Dan89]. We must account for two new complexities. An addresses involves more than two different shapes, and second, analyzing backwards moves is not as neat as with Robinson's triangles.

Danzer says that he discovered these tiles by examining the mirror planes of a fixed regular icosahedron. Tiles can have four different shapes, all of which are tetrahedra with edge-lengths and dihedral angles given by Figure 3.1.

In Figure 3.2, we provide fold-up cutouts for each tile. The tiles are named A, B, C, and, of course, K. Mirror images are also allowed, and we denote these as A', B', C', and, K'. The tiles have one matching restriction. Any face participating in a dihedral angle of 90 degrees must match with its mirror image. The A, B, and C tiles each have one such edge, while the K tile has two. Thus valid tilings contain A, B, and C tiles in groups of four, while K tiles form octahedra.

3.1 Inflation

Danzer's tiles can be assembled to form a larger version of each tetrahedron. This occurs directly, rather than cycling through different protosets as with the A and B-tiles. Assembling groups of tiles in this way yields A, B, C, and K tiles scaled by τ and with the same matching rules. Figure 3.3 suggests the structure of these

	1–2	2–3	3–1	2–4	1–4	3–4
A	$36°$, a	$60°$, τb	$72°$, τa	$108°$, a	$90°$, 1	$60°$, b
B	$36°$, a	$36°$, τa	$60°$, τb	$120°$, b	$108°$, $\frac{a}{\tau}$	$90°$, 1
C	$36°$, $\frac{a}{\tau}$	$60°$, τb	$90°$, τ	$120°$, b	$72°$, a	$36°$, a
K	$36°$, a	$60°$, b	$72°$, $\frac{a}{\tau}$	$90°$, $\frac{\tau}{2}$	$60°$, $\frac{1}{2}$	$60°$, $\frac{1}{2\tau}$

Figure 3.1: Edge lengths and dihedral angles for Danzer's tetrahedra. Vertices are numbered from 1 to 4 and an edge is given as the pair of vertices at its endpoints. Edges are listed along the top of the table and the four types of tetrahedron are listed along the left side. Entries in the table indicate the dihedral angle and length of one of the edges of one of the tetrahedron types. Three constants are used for the edge lengths: τ is the golden ratio, a is $\frac{\sqrt{(10+2\sqrt{5})}}{4}$, and b is $\frac{\sqrt{3}}{2}$. This table is from Danzer[Dan89] with some notations changed and an erroneous angle corrected.

assemblages. We call each group of tiles an *inflation patch* and they are shown in Appendix B, Figure B.1–Figure B.4. In order to inflate a tiling by Danzer's tetrahedra, each tile is replaced with the corresponding inflation patch and the result is scaled by τ, yielding a new tiling of the original size. Seven levels of inflation of the K tile are presented in Figure B.5–Figure B.9. It is interesting to note that a K-tile appears at the corner shown at every level of inflation of a K-tile.

3.2 Addresses

We form addresses using the inflation patches. The first digit of an address is just the type of the addressed tile—A, B, C, or K. The second digit is the location of this tile within one of the inflation patches. The A tile can only be found inside a larger C-tile, but all other tiles occur within patches at the next level of inflation in multiple ways. We assign unique numbers to tiles of the same type in the same inflation patch according to the labels of nodes in Figure 3.3 to ensure that locations are uniquely specified. The address $(K2B)$ represents a K-tile inside a B-tile, specifically the second K-tile inside the B-tile, or $K2$ in the inflation patch for B in Figure 3.3. We call the patch within which an address locates a tile its *extent*. Thus the extent of a two-digit address is one of the inflation patches. We consider the letters of an address its digits. Numbers between digits indicate the unique placement of digits in relationship to succeeding digits.

The extent of a three-digit address is one inflation of one of the inflation patches. The inflation of an inflation patch is constructed by scaling the inflation patch and

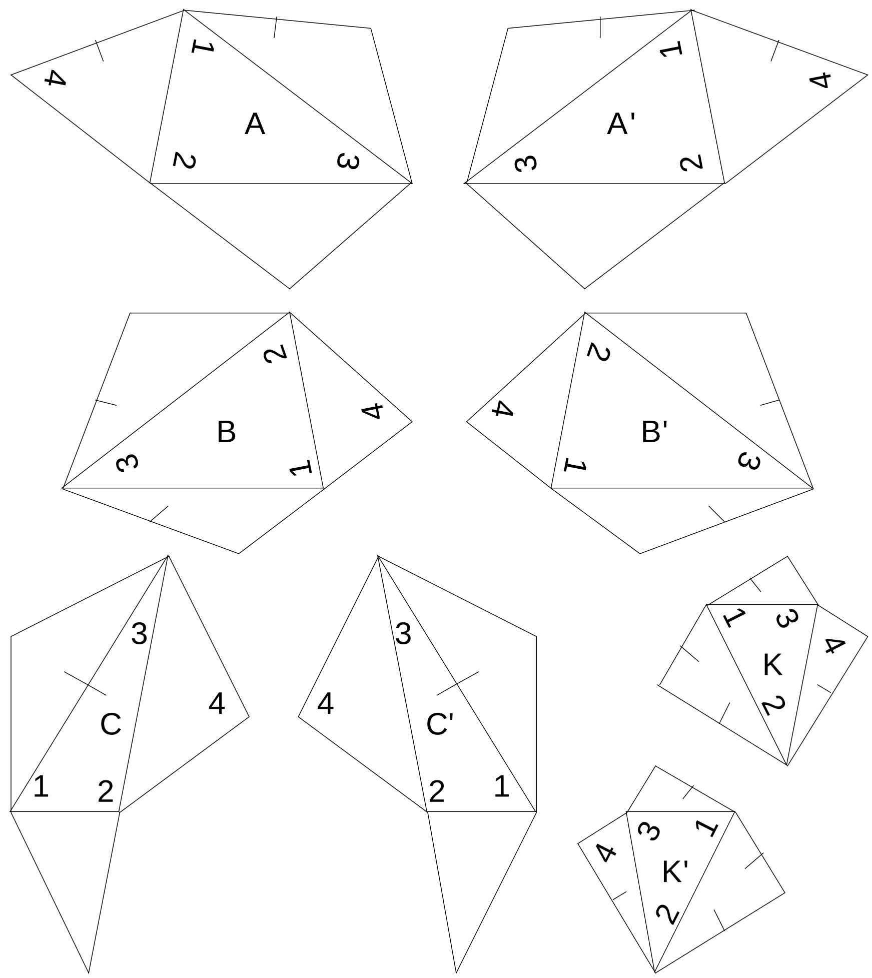

Figure 3.2: Folding templates for the ABCK tiles. Edges on faces that must match with the same type of tile with opposite handedness are marked with a slash. Vertices are labeled with their numbers (1–4). Left-handed tiles are built by folding the tiles backwards or printing a mirror image of this page.

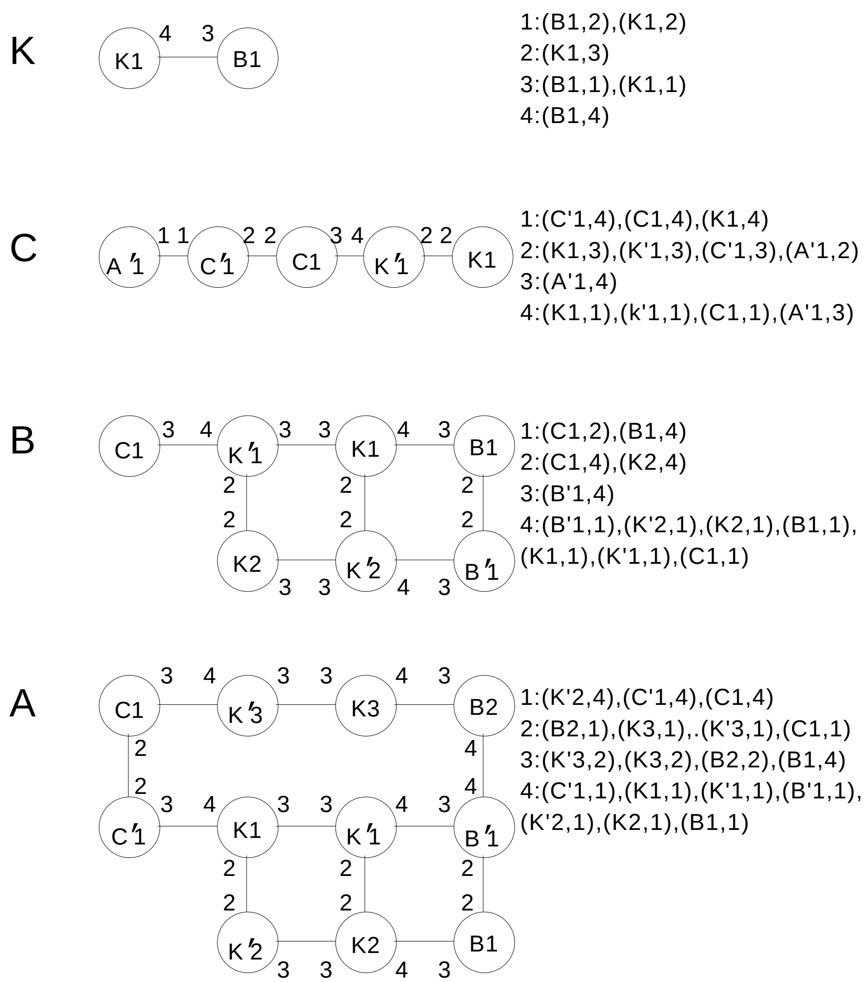

Figure 3.3: Structure of inflation patches for Danzer's tetrahedra. Nodes are the tiles which make up the inflation patch for each type of tetrahedron. Labels indicate the types of these tiles and primes (') indicate that a tile is left-handed. Numbers were assigned arbitrarily to ensure ensure that each node's name is unique within its inflation patch. An edge between two nodes indicates that the two tiles are adjacent by the faces opposite the numbered vertices labeling the edge. To the right of each inflation patch, the individual faces that make up each face of the patch are listed. If a tile is adjacent to no other tiles then it appears in one of these lists.

replacing each of these scaled tiles with the inflation patch of its type. The third digit of an address identifies one of the scaled tiles that was replaced. Then the first two digits identify the location of a single tile within the inflation patch selected by the third digit.

Each digit of the address works in a similar fashion. The n-th level of inflation for a tetrahedron of type p is formed by expanding the unit inflation patch for p by a factor of $\tau^{(n-2)}$ and then replacing each of these expanded tiles with the appropriate inflated patch from level $n-1$. Note that we let the first level of inflation be a single tile of type p. The extent of an n-digit address is the n-th inflation of the tetrahedron with the type of the n-th digit. Digit number $n-1$ selects a tile from the unit inflation patch that was scaled to form the level of inflation corresponding to the extent of the address and previous digits select a single tile from the patch that replaces the selected tile.

3.3 Adjacency

External and internal rules can be formed for the Danzer tiles in a manner similar to the rules for Robinson's tiles. Every edge between nodes of the graph in Figure 3.3 is an internal rule and sides that are part of a face of an inflation patch are external rules. There are a total of 100 rules. In Figure 3.4 we show the eight rules from the inflation patch of the K-tile.

As for Robinson's triangles, a move from a tile implies a series of external moves that eventually reach an internal move. Once an internal move is reached, all the digits after a certain point are known. The trick is to determine the previous digits of the destination address. For the Robinson triangles, only two sides appear on the right-hand side of multiple external inflation rules, but these only occur if the tile being entered is a mirror-image of the current area. For Danzer's tetrahedra, there are many multiple rules. However, the only matching that is not an exact mirror image occurs because face 4 of K sometimes matches with face 3 of B, sometimes with face 3 of C', and sometimes with itself (K'). This face of the K and B tiles consists of side 4 of a single B tile while the face of the C tile consists of side 4 of an A tile. Thus if a move crosses from a K to a C' tile then the previous digit of the origin address must have been a B and the previous digit of the destination address must have been an A. The move that occurs thus changes an address of the form (...xB1K...) to an address of the form (...yA1C'...). Here x and y represent the unique location of previous digits within the large tiles of type B and A. (For example, if the previous digit is a K, it could be located within an A or B tile in 6 or 4 ways.) Ellipses (...) indicate some prefix and suffix of the two addresses.

K1	1	$\rightarrow$	K	3		
K1	2	$\rightarrow$	K	1		
K1	3	$\rightarrow$	K	2		
K1	4	$\rightarrow$	B1	3	$\subset$	K
B1	1	$\rightarrow$	K	3		
B1	2	$\rightarrow$	K	1		
B1	3	$\rightarrow$	K1	4	$\subset$	K
B1	4	$\rightarrow$	K	4		

Figure 3.4: External and internal rules from the K tile's inflation patch. An external rule of the form $Xn\ s \rightarrow Y\ s'$ indicates that leaving the tile Xn inside the inflation patch for Y by its side s also leaves the inflation patch by its side s'. An internal tile of the form $Xn\ s \rightarrow Ym\ s' \subset Z$ indicates that leaving the tile Xn inside the inflation patch for Z enters another tile, Ym, within Z by its side s'. Here n and m are always 1, but for other inflation patches they are needed to differentiate tiles of the same type.

Now we note that A and B tiles can only match together in such a way that tiles of the same type but with opposite handedness lie across from each other. Thus, the previous digits that form the head of each address are identical. The way in which side 4 of a B tile matches with the fourth side of an A tile can be used to determine y from x. The same reasoning is also true in reverse for a move from a digit C' to a digit K, and all rules also hold when the handedness of every digit is reversed. Thus a move from (...yA'1C...) leads to (...xB'1K'...) where x can be deduced from y.

Therefore, a data structure can be formed for Danzer's tiles that works in the same manner as the Robinson data structure described in the previous chapter. We have not analyzed the symmetry types that are possible for Danzer's tiles because this analysis is significantly more complicated. We do, however discuss the problem of symmetry in general in section Section 6.1 and the techniques presented there are applicable to Danzer's tiles. In forming a data structure for addressing and calculating adjacencies in inflation patches of Danzer's tiles we only rely on particular features of the tiles in order to determine how to calculate previous digits once an internal rule is reached. This stage of the adjacency algorithm can be solved more generally by constructing a table that is used to look up previous digits. The calculation required to produce such a table is similar in nature to our deduction above that Bx and Ay precede K1 and C'1 and that y can be deduced from x. In

Section 5.3, we show how to generate a data structure for exploring tilings generated by an arbitrary set of inflation patches and this technique could be used to more precisely define a data structure for Danzer's tetrahedra.

Chapter 4

Tilings in Space

Thus far, we have been operating with the intuitive notion that a tiling is the placement of tiles to fill space with no overlaps, but it is important to define this concept more precisely before we make more general statements about inflation and aperiodicity[1]. We will restrict our definitions to tilings with certain well-behaved properties and the standard form we will use to specify aperiodic tilings we wish to explore further restricts the class of allowed tilings even further.

4.1 Protosets and Patches

Our definitions will apply to tilings and patches in three dimensions with the understanding that analogous definitions hold in two or more than three dimensions. Where we say "space", we mean "plane" in two dimensions, a section of a surface in three dimensions is an arc in two, and volume in three dimensions is area in two.

A *prototile* is a closed ball of points . Any prototile has a *boundary*, which is a subset of the prototile and contains all the other points. A family of prototiles, or *protoset*, is a set of prototiles where the prototiles have some minimum and maximum size. This means that every prototile must contain a sphere with radius $r > 0$ and must be contained by a sphere with radius r'. A *tile* is a set of points congruent, without reflection, to some prototile. This congruent prototile is called the tile's prototile type, or *prototype*. A *patch* is a (possibly infinite) set of tiles for

[1] Grünbaum and Shephard[GS87] stress the importance of carefully defining tiling terminology and describe confusion that has resulted from subtly inconsistent definitions in the literature. Our definitions are consistent with those of Grünbaum and Shephard with some restrictions and alterations to suit our application. Most notably, we extend their definitions to three dimensions and we place more emphasis on the concept of patches, considering tilings a special kind of infinite patch.

which the intersection of any two tiles has zero volume and the union of all tiles is a ball. Here, we have chosen to draw a distinction between the set of shapes that may be placed to form patches (prototiles), and the actual instantiation of prototiles in some patch (tiles). We will not normally speak of tiles in any other context than their appearance in some particular patch. A protoset *admits* a patch if every tile of the patch has a prototile type in the protoset. We choose to ignore differences between patches due to rotation or translation. Two patches are considered *the same patch* if one can be mapped to the other by rotation and translation.

An *edge* of a patch is a region of intersection between exactly two tiles[2]. Edges have positive area. Regions of intersection between three or more tiles have zero area. Edges are open sets with zero area, no two of which intersect. Tiles that share an edge are said to be *adjacent*. If one or more edges of a tile lie on the boundary of the patch then that tile is called an *exterior* tile, and a tile adjacent to other tiles along each of its edges is *interior*.

Consider a patch consisting of two cubes. These cubes could be joined together in an infinite variety of ways by sliding faces by each other in small increments. We wish to constrain these adjacencies. A patch is *edge-to-edge* if each of the prototiles is decorated such that:

1. The decoration consists of a finite number of nonintersecting open discs, called *facets*.

2. The facets form a cover for the boundary of the prototile.

3. Given any two tiles in the patch, placing prototile decorations in the same positions at which prototiles were placed to instantiate the two tiles results in the placement of a facet from each tile consisting of exactly the same points as the edge shared by the adjacent tiles.

Any set of prototiles may be altered by adding bumps to its faces so that it will only admit edge-to-edge patches. Hereafter, we will only consider prototiles with defined facets and whenever we say a protoset admits a patch, we mean that the protoset admits an edge-to-edge patch according to this fixed decoration of facets. If we need to emphasize that non-edge-to-edge patches are prevented by the shape of a protoset's facets, we say that protoset *forces edge-to-edge placement*.

In fact, the geometry of the decoration of a prototile completely determines its shape, since the union of its facets covers the tile's boundary. Note that a given set

[2]The term "edge" may be misleading because in three dimension, edges have area. The name "edge" is taken from the two-dimensional case, in which edges are arcs.

of prototiles may still admit uncountably many different patches, but that no two will have all the same adjacencies.

The facets of Robinson's triangles and Danzer's tetrahedra are marked in special ways, and only similarly-marked facets can be placed together. Many similar systems of marking, known as matching rules, are often used in tiling literature to constrain the ways in which tiles are allowed to fit together. Usually, as is true for Robinson's and Danzer's prototiles, these markings are used as a shorthand and the facets can be deformed in such a way that unmarked, deformed tiles match in exactly the same ways as the marked tiles. If matching rules are used, we require that they can be represented by a table telling whether any pair of facets from the protoset is permitted to match or not. We will not examine matching rules that involve the examination of surrounding area beyond two adjacent tiles to determine if they match. It is important to note that local matching rules are not generally sufficient to grow infinite tilings without backtracking, but they are necessary.

4.2 Tilings and Aperiodicity

If a union of the tiles of a patch fills space, the patch is called a *tiling*. Whenever a tile is instantiated in a tiling, every facet of its prototile type is aligned with an edge of the tiling, and thus all tiles are interior. A consequence of our requirement that the prototiles of a protoset have minimum size is that tilings have countably many tiles.

A tiling is *periodic* if there is a translation by a positive distance that maps the tiling to itself. This definition of periodic is nonstandard. Traditionally, periodic tilings in two dimensions must have at least two nontrivial translational symmetries in nonparallel directions. If this is the case, then the entire tiling can be formed by repeated placement of a single decorated parallelogram in a single orientation, as can be seen in Figure 2.2. Three-dimensional tilings with three nonparallel translations can be formed from a period parallelepiped. Now consider a three-dimensional tiling that has only one direction of translational symmetry. This entire tiling can be formed by repeated placement of a single planar strip perpendicular to the direction of translational symmetry that has a width that is a multiple of the distance of translation. Such a tiling cannot be formed by repeated placement of any finite pattern. However, we are chiefly interested in tilings that are *nonperiodic*. Two-dimensional tilings taken from slicing planes perpendicular to the direction of translational symmetry are not periodic, but we do not consider such a structure sufficiently "irregular" to be called nonperiodic.

Nonperiodic tilings can be constructed by placing tiles radially or by placement of several intersecting grids spaced quasiperiodically, but it is difficult to find a finite family of shapes that will only tile aperiodically. A protoset is *aperiodic* if it admits at least one tiling, but admits no periodic tilings. Penrose's rhombs and the kites and darts are aperiodic protosets consisting of only two tiles. It is currently and open question whether there is an aperiodic protoset consisting of a single tile.

4.3 Inflation

Inflation is a technique for cutting up a patch in order to form another patch consisting of smaller tiles with shapes equivalent to those of the original tiles. Depending on the ways in which these cuts are made in relation to the original tiles and whether the cutting process is invertible, we can use the concept of inflation to prove that protosets are aperiodic and to form data structures that represent the patches they admit. Inflation can iterated to form patches by increasing numbers of successively smaller shapes, as can be seen by viewing Figure A.1–Figure A.10 in reverse order. Inflation is best understood visually and it will be helpful to spend some time contemplating these pictures before we define the process more carefully below. To do so, we first define what it means for tiles to have equivalent shapes and define operations for cutting apart and reassembling tiles.

Earlier, we mentioned that prototiles could be deformed so that they can only be placed edge-to-edge or can only be placed adjacent according to certain matching rules. A set of prototiles can also be deformed with no effect on the ways in which the prototiles can be placed to form patches. We consider two protosets *surrogates* if they have basically the same shapes[3] and equivalent matching rules. In two dimensions, we think of forming transparent sheets with an overlaid destination prototile shape and position for each source prototile. Two protosets are surrogates if a single set of transparencies can be constructed so that given any patch by one protoset, each of its tiles can be replaced by a destination tile and the result can be scaled to form a patch admitted by the other protoset. Specifically, two protosets P and P' are surrogates if and only if a metric equivalence H deforms the prototiles and facets of P such that given any patch G admitted by P, if every tile t in G were instead instantiated as $H(t)$ a patch would result that is a uniform scaling of some patch admitted by P'. We will refer to the magnitude of this scaling as

[3]Grünbaum and Shephard[GS87] have a concept that is similar to that of surrogate protosets. They say tiles have the same "basic shape" if one can be obtained from the other by altering its facets. Gähler, Baake, and Schlottmann[GBS94] define another similar concept as a relationship between tilings. They say that two tilings are *mutually locally derivable* if (after rotation and scaling) any ball in one tiling is uniquely determined by a ball in the other tiling with a radius that is a constant amount larger.

the *expansion factor* from G' to G. H^{-1} can be used to reverse this process and equivalent patches admitted by P and P' can be mode to correspond by a bijection.

A patch G' is a *composition* of the patch G if every tile of G' is a union of tiles of G. Similarly, a protoset P' is a composition of the protoset P if all prototiles of P' are each congruent to patches admitted by P. Note that if the protoset P' is a composition of P then every tiling admitted by P' is a composition of some tiling admitted by P and if a single tiling G' is a composition of G then the protoset that admits G' is a composition of the protoset that admits G. The term *decomposition* will refer to the reverse of composition. If the patch G' is a composition of G then G is a decomposition of G', and similarly for protosets. A composition is formed by attaching together smaller tiles to form larger tiles while a decomposition is formed by dividing larger tiles into smaller tiles.[4]

Theorem 1 *If the protosets P and P' are surrogates, P' is a composition of P, P forces edge-to-edge placements, and some prototile $p' \in P'$ is composed from a patch containing an interior tile of prototype p such that p and p' are corresponding surrogate prototiles, then P tiles the plane.*

Proof: We will prove this statement by constructing a series of increasingly large patches, each of which lies inside its successor. We will see that each of these patches can be extended in all directions, hence the patch at the limit of this process is a tiling. The first patch in this series, G_p consists of a single tile of prototype p. For each patch G in the series, the next patch, $\mathrm{INF}(G)$, is obtained by the following process:

1. Find G', the surrogate patch corresponding to G.

[4] Decomposition is not always defined in this way. An alternative definition of inflation[GS87][LS86] concerns a decoration of prototiles that is used to obtain inflated tilings. Decorations consist of edges that become smaller tiles when each tile is replaced with its decoration and the original edges are removed. This allows the creation of new tiles that overlap two or more tiles from the previous level of inflation. We do not use this definition because it is less useful for describing finite patches, it makes proof of aperiodic properties more difficult, it may be more complicated to invert, and the data structures we eventually use to represent tilings cannot represent this type of inflation. The Penrose kites and darts, and rhombuses are classes of tilings that are most simply described with inflation decorations. However, Penrose used a definition according to inflation to analyze the first aperiodic set he discovered and his tilings by two prototiles can be analyzed in terms of composition as well. A potential area for future investigation would be the study of whether inflation by this type of decomposition decoration is more or less general than the method of inflation we use here and whether it is possible to convert a representation in one form to the other.

2. Find $\text{INF}(G)$, admitted by P, by dividing each tile of G' in the same way that its prototype is composed from prototiles of P.

We write $\text{INF}(\text{INF}^n(G_p))$ as $\text{INF}^{n+1}(G_p)$, and our series of extending patches is:

$$G_p = \text{INF}^0(G_p), \text{INF}(G_p) = \text{INF}^1(G_p), \text{INF}^2(G_p), \text{INF}^3(G_p), \ldots$$

Now we need to demonstrate that for all $n \geq 0$, $\text{INF}^n(G_p)$ is contained inside $\text{INF}^{n+1}(G_p)$. By contained inside, we mean that a subset of tiles of $\text{INF}^{n+1}(G_p)$ is the same as $\text{INF}^n(G_p)$ and the tiles in this subset are all interior. This will demonstrate that the patch $\text{INF}^{n+1}(G_p)$ can be extended in all directions. This property is true when n is 0 because $\text{INF}^1(G_p)$ is the patch that p' is composed from and the theorem states that this patch contains an interior tile of prototype p. Because G_p consists of a single tile of prototype p, $\text{INF}^0(G_p)$ is contained inside $\text{INF}^1(G_p)$. To show the property holds for $n = 1$, we note that $\text{INF}^1(G_p)$ was formed by dividing the single tile of $\text{INF}^0(G_p)$ according to the patch p' was composed from. Since $\text{INF}^1(G_p)$ contains a tile of prototype p, the formation of $\text{INF}^2(G_p)$ will also divide this tile according to the patch p' was composed from. Since this tile is internal, $\text{INF}^1(G_p)$ is contained inside $\text{INF}^2(G_p)$. This reasoning holds for all $n \geq 0$. Since $\text{INF}^0(G_p)$ is contained by $\text{INF}^1(G_p)$ it is also true that $\text{INF}^n(\text{INF}^0(G_p))$ is contained by $\text{INF}^n(\text{INF}^1(G_p))$ because the generation of $\text{INF}^n(\text{INF}^1(G_p))$ involves n decompositions of the patch $\text{INF}^0(G_p)$, which lies inside $\text{INF}^1(G_p)$ and this will form a patch congruent to $\text{INF}^n(\text{INF}^0(G_p))$ contained inside $\text{INF}^n(\text{INF}^1(G_p))$.

$$\lim_{n\to\infty} \text{INF}^n(G_p) \text{ is a tiling admitted by } P. \ \Box$$

Observations on Theorem 1:

An *inflation specification* is a protoset such as P' above, satisfying the conditions of the theorem, and includes the correspondence between tiles that compose each prototile of P' and the surrogate prototypes of these tiles. An inflation specification completely determines the above construction. We will make extensive use of INF as above. We say $\text{INF}(G)$ is the *inflation* of G and we call $\text{INF}^n(G)$ the *n-th level of inflation* of G. We will generalize the concept of inflation to refer to any patch admitted by the protoset P. $\text{INF}(G)$ is taken by decomposing the surrogate patch corresponding to G so that every tile is divided in the same way that its prototype was composed. Note that $\lim_{n\to\infty} \text{INF}^n(G_p)$ is invariant under inflation.[5] If a protoset admits a tiling then every surrogate of that protoset admits a tiling as well.

[5] Lunnon uses a very similar construction for a tiling that is invariant under inflation.[LP87]. In fact, this is a unique fixed point. $\text{INF}(G)$ results in a tiling of the same size as G. We scale

We also define INF for any protoset P surrogate to an inflation specification P'. If G is admitted by P and G' is the corresponding patch admitted by P' then $\text{INF}(G)$ is the patch admitted by P that corresponds to $\text{INF}(G')$.

In the theorem, we require that P force edge-to-edge placements so that whenever a patch admitted by P' is decomposed we are guaranteed that the resulting patch admitted by P is also edge-to-edge. Any surrogate for P admits a tiling corresponding to $\lim_{n\to\infty} \text{INF}^n(G_p)$ so we can use Theorem 1 to prove that a protoset that does not force edge-to-edge placement tiles the plane provided we can find an inflation specification composed from the protoset in which tiles instantiated from the inflation specification can only be placed adjacently if the tiles they are composed from match up edge-to-edge along the edge of the adjacencies.

Any inflation specification can be used to construct infinitely many equivalent inflation specifications with each prototile formed by increasingly large compositions. If the inflation specification is as in the theorem, then for any $n \geq 1$, $P^n = \{\text{INF}^n(G_p) \mid p \in P, \text{ patch } G_p \text{ contains a single tile of prototype } p\}$ is also an inflation specification. We say an inflation specification is *minimal* if it cannot be constructed in this manner from a protoset with fewer prototiles. Any prototiles of P that do not occur in the composition patches of prototiles of P' are never instantiated in $\lim_{n\to\infty} \text{INF}^n(G_p)$. If a subset of the prototiles of P' that includes p' is formed from patches that contain no tiles that have a surrogate prototype that is outside the subset then no prototiles outside the subset will ever appear in the tiling generated by Theorem 1 and these prototiles can be discarded. If no such subset exists then the inflation specification is *irreducible*.[6]

Our theorem requires that some prototile of P' is composed from a patch containing a tile of its surrogate prototype. The existence of this tile guarantees that the series of patches generated will always contain a tile of that type and also that the number of tiles at each inflation level is greater than the number of tiles for the previous level. The internal location of this tile guarantees that the patch at each level of inflation extends the previous level in all directions. However, minimal and primitive inflation specifications often do not contain such an internal tile. In fact, neither Robinson's triangles nor Danzer's tetrahedra have inflation patches with any

---

each inflation by the expansion factor of the surrogate protosets, producing a patch with smaller tiles that can be translated to fit directly above the previous level of inflation. In the limit, this produces a patch of infinitely small tiles packed into the shape of a single original tile (distorted by many surrogations). It can be shown that if inflation is defined in this manner then, using the Hausdorff metric, the inflation of any patch is closer to $\lim_{n\to\infty} \text{INF}^n(G_p)$ than to the uninflated patch. Thus, $\lim_{n\to\infty} \text{INF}^n(G_p)$ is the unique fixed point of INF. In fractal geometry, this is the Contraction Mapping Theorem.[Bar93]

[6] The term irreducible comes from Lunnon [LP87] and he develops these concepts more fully. We draw on his analysis in the next paragraph.

internal tiles at all. However, it is possible to construct a larger protoset that does by inflating the patches that the prototiles are composed from, as described previously. If no such inflation of the prototiles contains an internal tile then inflation can only generate a strip of tiles rather than a tiling. Note that this does not mean that protoset cannot tile the plane, but rather that the inflation specification cannot be used to create a tiling. If an inflation of one of the prototiles eventually contains an internal tile, then it is possible to show that a later inflation of the prototile will contain a tile of that prototile type. Higher levels of inflation will contain inflations of the internal tile. If one of these does not contain a tile of the desired prototile type then the inflation specification is not primitive. The prototile that does not appear in these inflations can be discarded and the construction repeated with fewer prototiles. If inflation can generate a tiling then an inflated protoset containing a prototile of the desired form will eventually be found. If the inflated protoset admits a tiling, then the original, surrogate protoset admits a tiling as well.

In our proof of Theorem 1, we showed that if $\text{INF}^0(G_p)$ is contained inside $\text{INF}^1(G_p)$ then $\text{INF}^n(\text{INF}^0(G_p))$ is contained inside $\text{INF}^n(\text{INF}^1(G_p))$. It will be useful to further describe the relationship between tilings at corresponding levels of inflation. Consider an arbitrary patch, G, and color two of its tiles blue and green. As $\text{INF}^n(G)$ is constructed, always decompose tiles of one of these colors into tiles of the same color. Then $\text{INF}^n(G)$ will contain distinct blue and green regions. The subset of tiles having either of the colors forms a patch. If two adjacent tiles of G were chosen, then tint the edge between them pink. Whenever a patch is inflated, if a tile has a pink facet, then all facets of the tile's decomposition that lie along the pink facet are tinted pink as well. Then, every inflation of G will have a blue and green region that meet at a pink disc made up of the edges between blue and green tiles.

Consider the patches G and $G' = \text{INF}(G)$. A finite number of inflations of a patch that is not a tiling will never generate a tiling. Thus, if G is a tiling then G' is a tiling and if G' is a tiling then so is G. If INF is invertible over the range of patches that are tilings then we say that the inflation specification that generates INF possesses the property of *unique inflation.*

Theorem 2 *If an inflation specification possesses unique inflation then the inflation specification admits only aperiodic tilings.*

Proof: If an inflation specification is unique then the inflation of any two different tilings is different. Every iteration of INF decomposes a tiling and reexpands the result by the expansion factor of the inflation specification so that tiles of the same size and shape are obtained. If a translation of distance d maps a tiling to itself, then a translation of distance dm maps an inflation of the tiling to itself where m is the expansion factor of the inflation specification. If the inflation specification is unique then INF on tilings is invertible and the converse is also true. The inverse of inflation, *deflation*, can be iterated, and translational symmetry is preserved with distances diminishing according to powers of the inverse of m.

Assume that G, admitted by a unique inflation specification P, is a periodic tiling with a translation of distance $d > 0$ that maps G to itself. Every prototile of P contains a sphere of radius r. $\text{INF}^n(G)$ is taken to itself by a translation of distance $\frac{d}{m^n}$. However, the tiles of $\text{INF}^n(G)$ still contain balls of radius r. This is a contradiction because for sufficiently large n, $r > \frac{d}{m^n}$. No such translation can exist and therefore, P is aperiodic. □

A *vertex configuration* is a minimum patch of tiles that meet at a single point and completely surround that point. If all vertex configurations occur in a tiling inflated from a single tile then any local configuration of any tiling admitted by the protoset can be found in any other tiling.

Theorem 3 (Local Isomorphism Theorem) *If some inflation of a tile,* $\text{INF}^n(p)$ *structurally contains every vertex configuration then given any two tilings* G *and* G'*, every finite subpatch of* $\mathcal{T}$ *is structurally equivalent to infinitely many subpatches of* $\mathcal{T}'$*.*

Proof: Every tiling contains an infinite number of tiles of prototype pP. Any tiling can deflated n times and since its deflation contains an infinite number of occurrences of t, the original tiling contains an infinite number of patches congruent to n inflations of t. Therefore, any tiling contains an infinite number of occurrences of every vertex configuration.

Given G, a subpatch of $\mathcal{T}$, we deflate $\mathcal{T}$ until G is contained within a single vertex configuration of $\text{INF}^{-k}(\mathcal{T})$ for some k. The length of the smallest edge in $\text{INF}^{-k}(\mathcal{T})$ grows with k and when this edge is larger than the radius of a circle containing G then G is within a single vertex configuration v.

Now, given another tiling $\mathcal{T}'$, we can find G' at infinitely many places within $\mathcal{T}'$. We know that the deflation $\text{INF}^{-k}(\mathcal{T}')$ contains an infinite number of occurrences of v. Since $\mathcal{T}'$ is the k-th inflation of this deflated tiling, it contains infinitely many occurrences of the k-th inflation of v and each of these contains G. Therefore, any finite patch in any tiling can be found in infinitely many places in any tiling. □

A consequence of the Local Isomorphism Theorem is that every finite region G of any tiling admitted by the protoset can be found inside the patch produced by an inflation of some prototile, $\text{INF}^l(p)$.

Chapter 5

Generalization: Uniquely Inflatable Tilings

In Chapter 2 we described the construction of data structures for exploring a particular class of tilings: Robinson's triangles. Similar techniques can be used to form data structures for exploring tilings formed from any prototiles that obey unique inflation rules by composition of tiles. Figure 5.1 demonstrates the relationship between a tiling by sets of points in space and its *dual graph*, a graph with a vertex for each tile, and an edge between vertices for each pair of adjacent tiles. A tiling's dual graph represents its structure—interconnections between adjacent tiles—while discarding geometric information such as the shapes and positions of tiles.[1]

Given a description of a family of prototiles and an inflation process, dual graphs can be formed that represent the structure of the tilings admitted by the prototiles. An inflation specification, as defined in Chapter 4 contains all the information that is required in order to generate such a data structure.[2] A tiling is explored by constructing the dual graph of increasingly large patches as more distant tiles are visited. Locations within a graph can be held as addresses of logarithmic length in terms of the number of tiles in the currently constructed region. Because of the way in which we form addresses, we will never actually store the dual graphs of the patches we explore: an address, alone, holds the information we need to investigate the dual graph of the patch it represents. This pretty result allows us to

[1] Maintaining this information would, of course, make the structure richer, at the possible expense of significantly increasing storage. We discuss in Section 6.3 a method for including the orientation of tiles.

[2] Actually, in the case of symmetric tilings, additional analysis is required to recognize the ways in which symmetries can occur, as is discussed in Section 6.1. Otherwise, explorations that cross a tiling's line of symmetry would cause the examination of infinitely-long addresses.

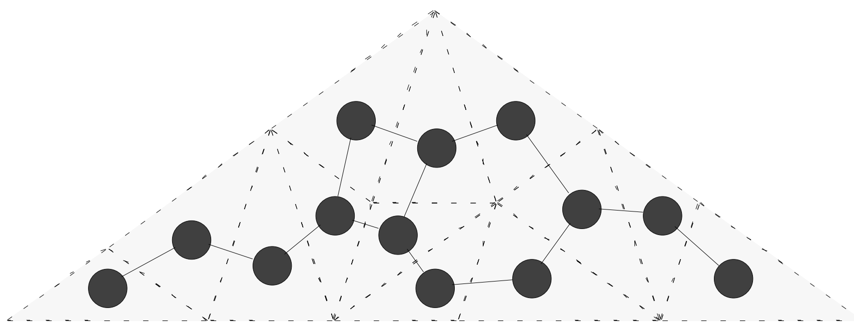

Figure 5.1: A patch of Robinson triangles superimposed with its dual graph. Tiles and edge adjacencies of the patch correspond to vertices and edges of the dual graph. Thus, for example, land-locked tiles have dual vertices of degree 3.

explore absolutely precise representations of very large tilings using a small amount of storage space.

5.1 Specifying a Protoset and Inflation Rules

In Chapter 4, we defined prototiles, edges, tiles, and patches as sets of points in space. Below, we define these concepts for tilings represented as dual graphs. These discrete definitions are the representations we will use to form data structures. If we need to clarify that we are talking about a patch as set of points in space we will refer to a *spatial* patch and a patch represented as a dual graph will be referred to as a *dual* patch. Section 5.1.6 describes a procedure for converting spatial representation into dual representations.

A *protoset specification* defines a family of prototiles and their facets and an *inflation specification* defines a set of patches equivalent to each of the prototiles at one level of inflation. The protoset specification and inflation specification contain sufficient information to generate a data structure for exploring the dual graphs of tilings admitted by the prototiles.

5.1.1 Protoset Specification

Tilings are made up of tiles selected from a *protoset*, a finite set of *prototiles* $P = \{p_i \mid 1 \leq i \leq n\}$. The boundary of each prototile $p \in P$ is made up of a finite set of facets, $\mathcal{F}(p) = \{p^i \mid 1 \leq i \leq n\}$. A *prototile specification* consists of a protoset P and the facets of each prototile, $\mathcal{F}(p)$.

5.1.2 Patches

Tiles are the instances of prototiles that actually appear in a tiling. We say that the *prototile type* of a tile is the prototile from which it was created. A tile's edges are instances of the facets of its prototile type.

The structure of a region within a tiling is represented by a patch. A *patch* is a triple of the form $(T, \mathcal{P}, \mathcal{A})$. $T = \{t_i \mid 1 \leq i \leq n\}$ is a set of all tiles in the patch. $\mathcal{P} : T \Rightarrow P$ is a total function that maps every tile to its prototile type. The function $\mathcal{A}$ relates neighboring tiles in the patch and is defined more precisely momentarily.

An *edge* in a patch has the form (t, e) where t is a tile and e is one of the facets of its prototile type. Given a patch $(T, \mathcal{P}, \mathcal{A})$, $\mathcal{F}(T, \mathcal{P}) = \{(t, e) \mid t \in T, e \in \mathcal{F}(\mathcal{P}(t))$ denotes the set of all edges in the patch.

If two edges cover each other spatially they are said to be *coincident* and the two tiles they lie on are *adjacent*.[3] Every edge is coincident with at most one other edge. The automorphism $\mathcal{A} : \mathcal{F}(T, \mathcal{P}) \Rightarrow \mathcal{F}(T, \mathcal{P})$ determines the pairing of coincident edges. In particular, $\mathcal{A} : (t, e) \mapsto (t', e')$ if and only if (t, e) and (t', e') are coincident. Note that $\mathcal{A}$ is not defined on edges that are adjacent to no tile and that $\mathcal{A} = \mathcal{A}^{-1}$ on the range of $\mathcal{A}$. Edges in the range of $\mathcal{A}$ are called *interior* edges and edges not in the range are *exterior*. The set of exterior edges of a patch G is called its *boundary* and is denoted $\mathcal{B}(G)$. See Figure 5.2 for an example of a spatial patch and its representation as a dual patch.

The patch $(T, \mathcal{P}, \mathcal{A})$ is considered a *subpatch* of $(T^*, \mathcal{P}^*, \mathcal{A}^*)$ if $T \subseteq T^*$, $\mathcal{P} = \mathcal{P}^*$ over the domain T, and $\mathcal{A} = \mathcal{A}^*$ over the domain $\mathcal{F}(T, \mathcal{P})$. If $(T, \mathcal{P}, \mathcal{A})$ and $(T', \mathcal{P}', \mathcal{A}')$ are subpatches of $(T^*, \mathcal{P}^*, \mathcal{A}^*)$ then $(T, \mathcal{P}, \mathcal{A})$ and $(T', \mathcal{P}', \mathcal{A}')$ are *distinct* if $T \cap T' = \emptyset$.

Given two patches $G = (T, \mathcal{P}, \mathcal{A})$ and $G^* = (T^*, \mathcal{P}^*, \mathcal{A}^*)$, we say that G is *structurally contained* by G^* if there exists a function $\mathcal{H} : T \Rightarrow T^*$ such that:

1. $\mathcal{H}$ is an injective function mapping tiles of G to tiles of G^*
2. for all $t \in T$, $\mathcal{P}(t) = \mathcal{P}^*(\mathcal{H}(t))$
3. for all $(t, e) \in \mathcal{F}(T, \mathcal{P})$, if $\mathcal{A} : (t, e) \mapsto (t', e')$ then $\mathcal{A}^* : (\mathcal{H}(t), e) \mapsto (\mathcal{H}(t'), e')$

[3] This definition of edge is not strictly analogous to the definition in Chapter 4 in that it distinguishes two sides of each edge. In Chapter 4, there was one edge for each pair of adjacent tiles while here, if t is adjacent to t' then t has an edge to t' and t' has an edge to t.

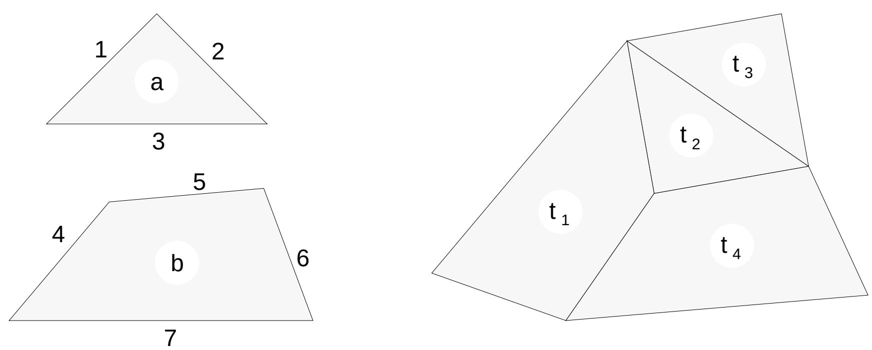

Figure 5.2: Notation for a patch. A set of prototiles (left) and a patch formed from them (right). If the patch is $G = (T, \mathcal{P}, \mathcal{A})$ then:

$$\begin{aligned}
T &= \{t_1, t_2, t_3, t_4\} \\
\mathcal{P} &= \{(t_1, b), (t_2, a), (t_3, a), (t_4, b)\} \\
\mathcal{A} &= \{((t_1, 4), (t_2, 2)), ((t_1, 5), (t_4, 4)), \\
&\quad ((t_2, 1), (t_4, 5)), ((t_2, 2), (t_1, 4)), ((t_2, 3), (t_3, 3)), \\
&\quad ((t_3, 3), (t_2, 3)), ((t_4, 4), (t_1, 5)), ((t_4, 5), (t_2, 1))\} \\
\mathcal{B}(G) &= \{(t_1, 7), (t_3, 1), (t_3, 2), (t_4, 6), (t_4, 7), (t_1, 6)\}
\end{aligned}$$

If G is structurally contained by G^* and G^* is structurally contained by G then $\mathcal{H}$ is bijective and we say that G and G^* are *structurally equivalent*. Figure 5.3 shows that Danzer's B-tile is structurally contained by the A-tile.

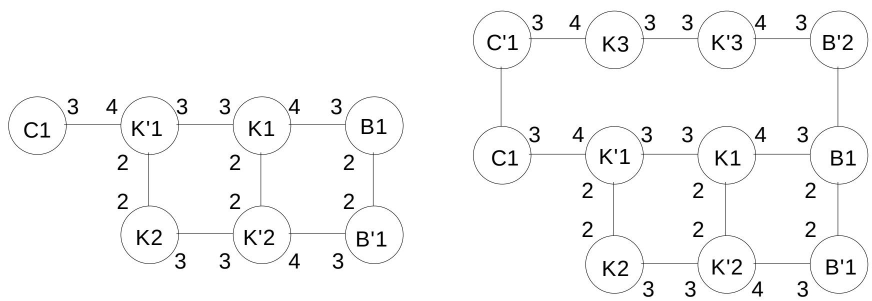

Figure 5.3: The patch on the left is structurally contained by the patch on the right. These are the inflation patches of Danzer's B-tile and A-tile (left-handed). The function $\mathcal{H}$ as in the definition of structural equivalence would map tiles in the patch on the left to tiles of the same name in the patch on the right. The fact that the inflation patch of the B-tile is contained in the inflation patch of the A-tile and that the inflation patch of the A-tile contains no A-tiles leads us to believe that there might be a nicer way to partition the Danzer tiles that would involve fewer pieces.

5.1.3 Patch-Sides

It is often necessary to describe the spatial arrangement of the boundary of a patch by breaking the edges that lie on it it into sections and establishing an ordering on the edges within those sections. A *patch-side* is an ordered list of exterior edges of the form $((t_1, e_1), (t_2, e_2), \ldots, (t_n, e_n))$. A patch-side enumerates the edges of a contiguous section of the boundary. In 3-space, a union of closures of the edges of a patch-side is a disc. The set of patch-sides of a patch G is denoted $\mathcal{S}(G)$ and partitions $\mathcal{B}(G)$.

We say that the patch-side S is *spatially contained* by S' if there is a rotation and translation X and a mapping $\mathcal{M} : S \Rightarrow S'$ such that for every edge $e \in S$, the set of points contained by $X(e)$ is congruent to the points of $\mathcal{M}(e)$. If S is spatially contained by S' and S' is spatially contained by S then S and S' are *spatially congruent*. We require that a convention be established so that all pairs of spatially congruent patch-sides $S = (e_i \mid i \leq 1 \leq n)$, $S' = (e'_i \mid i \leq 1 \leq n)$ are formed such that $\mathcal{M}(e_i) = (e'_i)$ for all $i \leq 1 \leq n$. Intuitively, this simply means that all patch-sides that have edges of the same shape and in the same arrangement must have their edges listed in the same order. In two dimensions, listing edges in clockwise order satisfies this requirement. In three dimensions, each class of spatially congruent patch-sides must be identified and assigned some fixed ordering.

If two patch-sides from two different patches are spatially congruent and the patches can be slid together so that their patch-sides coincide while the patches lie opposite each other across the patch-sides, then pairs of individual edges from the boundaries of the two patches will be juxtaposed. Our ordering convention on patch-sides allows us to find these coincident edges because the edges listed in one patch-side coincide spatially with the edges of the opposing patch-side listed in reverse order, as can be seen in Figure 5.4 and Figure 5.5. Given patch-sides $S = ((t_1, e_1), (t_2, e_2), \ldots, (t_n, e_n))$ and $S' = ((t'_n, e'_n), (t'_{n-1}, e'_{n-1}), \ldots, (t'_1, e'_1))$, the total function $\mathcal{M}_{SS'} : S \Rightarrow S'$, which is one-to-one and onto, pairs edges in this way. Specifically, $\mathcal{M}_{SS'} : (t_i, e_i) \mapsto (t'_i, e'_i)$. This function will be used to join patches together along spatially-congruent patch-sides.

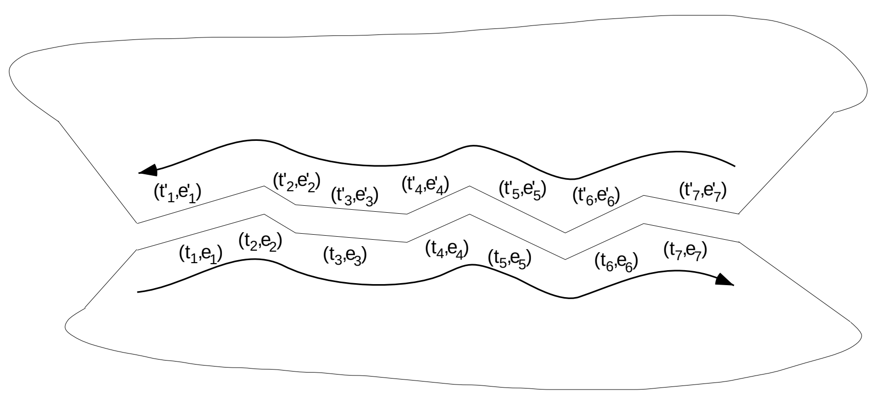

Figure 5.4: Joined two-dimensional patch-sides. When two patch-sides are placed together, their edges match in pairs. If $S = ((t_1, e_1), (t_2, e_2), \ldots, (t_n, e_n))$ and $S' = ((t'_n, e'_n), (t'_{n-1}, e'_{n-1}), \ldots, (t'_1, e'_1))$ then $\mathcal{M}_{SS'} : (t_i, e_i) \mapsto (t'_i, e'_i)$ is the pairing of matching edges. In two dimensions, listing edges of patchsides in clockwise order ensures that coincident edges are always paired by $\mathcal{M}$ appropriately.

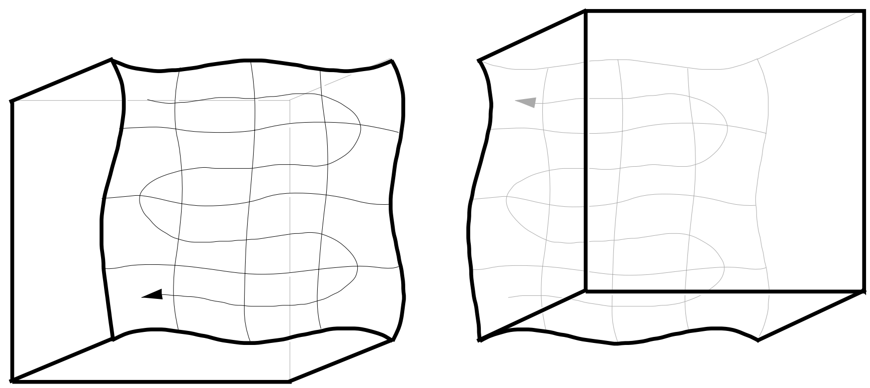

Figure 5.5: Joined three-dimensional patch-sides. Two patches that can be joined together along spatially-congruent patch-sides. The edges of one of these patch-sides would be paired with the edges of the other in reverse order. Since the patch-sides have the same shape, but different sides face outward from the patches, the patch-sides are formed by enumerating the edges in opposite order, as shown by the arrows.

5.1.4 Composition

Composition is a method of forming a patch by attaching together smaller patches along their patch-sides. If $\mathbf{G} = \{G_i \mid i \leq 1 \leq n\}$ is a set of patches and $\mathbf{S} = \{(S_i, S_i') \mid i \leq 1 \leq m\}$ is a set of pairs of patch-sides then $C_{\mathbf{GS}} = (\mathcal{T}_{\mathbf{GS}}, \mathcal{P}_{\mathbf{GS}}, \mathcal{A}_{\mathbf{GS}})$ denotes the *composition* of $\mathbf{G}$ along $\mathbf{S}$. Figure 5.6 presents a composition of patches spatially and, analogously, as dual graphs. Each pair in $\mathbf{S}$ must have patch-sides that are from different patches in $\mathbf{G}$ and that can be attached. Thus, for each pair (S, S') in $\mathbf{S}$, $|S| = |S'|$ and $S \in \mathcal{S}(G), S' \in \mathcal{S}(G')$ for some $G, G' \in \mathbf{G}, G \neq G'$. Here we have used $|S| = |S'|$ to indicate that S and S' have the same length. Spatially, in order for this composition to make sense, all the pairs of patch-sides of S must be spatially congruent.

The composed patch $C_{\mathbf{GS}}$ is formed as a union of the patches in $\mathbf{G}$ with additional adjacencies between tiles that are connected by patch-sides in $\mathbf{S}$, as follows:

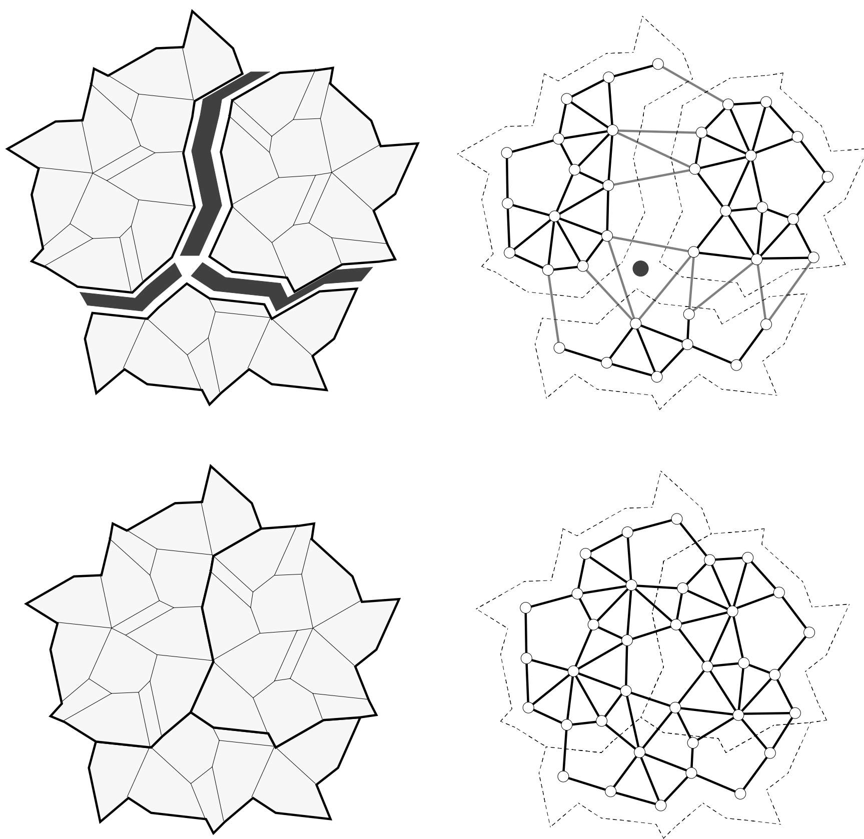

Figure 5.6: Composition of a set of patches. Patches are composed spatially (left) and as dual graphs (right). Separate patches (above) are attached below. Dark zigzags (upper-left) are placed between the three pairs of patch-sides that will be attached together. Grey edges (upper-right) indicate new adjacencies formed by the composition.

$C_{\mathbf{GS}} = (T_{\mathbf{GS}}, \mathcal{P}_{\mathbf{GS}}, \mathcal{A}_{\mathbf{GS}})$ where:

$$\text{for all } G \in \mathbf{G},\ \text{let } G = (T_G, \mathcal{P}_G, \mathcal{A}_G)$$

$$\begin{aligned} T_{\mathbf{GS}} &= \bigcup_{G \in \mathbf{G}} T_G \\ \mathcal{P}_{\mathbf{GS}} &= \bigcup_{G \in \mathbf{G}} \mathcal{P}_G \\ \mathcal{A}_{\mathbf{GS}} &= \bigcup_{G \in \mathbf{G}} \mathcal{A}_G + \bigcup_{(S,S') \in \mathbf{S}} \mathcal{M}_{SS'} \end{aligned}$$

All tiles in the patches of $\mathbf{G}$ are contained in $C_{\mathbf{GS}}$ and have the same prototile type. Two tiles of $C_{\mathbf{GS}}$ are adjacent if and only if they were adjacent in some patch of $\mathbf{G}$ or they lie opposite each other along joined patch-sides of $\mathbf{S}$.

5.1.5 Inflation

In Section 4.3, we defined the concept of an inflation specification, a set of prototiles composed from surrogate prototiles and showed that if the inflation specification is unique, inflation can be used to generate aperiodic tilings. This is the technique our data structures will use to generate the dual graphs of the tilings we wish to model.

Each prototile p is associated with an *inflation patch* that is written I_p. This patch describes the structure that will replace tiles of prototype p when they are decomposed in order to inflate a tiling. The boundary of an inflated prototile is divided into sections corresponding to the prototile's facets and these sections are used to form the patch-sides of the inflation patches. For a prototile p, $\mathcal{S}(I_p) = \{I_{(p,e)} \mid e \in \mathcal{F}(p)\}$, where $I_{(p,e)}$ denotes the patch-side along facet e of the inflation patch for p. The inflation patches I_p and their patch-sides $I_{(p,e)}$ make up the *inflation specification*.

Once an inflation specification is defined, we can construct an inflation of any patch G by composing inflation patches associated with individual prototiles. For every tile of G, we substitute an inflation patch of that tile's prototile type and the inflation patches are joined together along the adjacencies of G. Since the patch-sides of inflation patches match together in exactly the same way as the prototiles themselves, the patch-sides joined by these adjacencies will be spatially congruent. Because tiles of the patch we are inflating may have the same prototile type, it is important to form the composition from structurally equivalent duplicates of the inflation patches. Given a patch $G = (T, \mathcal{P}, \mathcal{A})$, its inflation, $\text{INF}(G)$ is constructed as follows:

for all $t \in T$

let $p = \mathcal{P}(t)$

define G_t, a patch structurally equivalent to I_p where $\mathcal{H}$ is the function that takes tiles in I_p to corresponding tiles in G_t.

for all $s \in \mathcal{F}(p)$

define $G_{(t,s)} = ((\mathcal{H}(t), e) \mid (t, e) \in I_{(p,s)})$

if $\mathbf{G} = \{G_t \mid t \in T\}$ and $\mathbf{S} = \{(G_{(t,e)}, G_{(t',e')}) \mid \mathcal{A}((t, e)) = (t', e')\}$ then $IG = C_{\mathbf{GS}}$

Levels of inflation are written with increasing powers: $G = \text{INF}^0(G)$, $\text{INF}(G) = \text{INF}^1(G)$, $\text{INF}^2(G)$, $\text{INF}^3(G)$, $\ldots$. It is sometimes convenient to use the concept of inflating a single prototile. When we write $\text{INF}(p)$, this is a shorthand for $\text{INF}(Gp)$ where Gp is a patch that contains a single prototile of prototype p. $\text{INF}(p)$ is structurally equivalent to I_p.

5.1.6 Converting Spatial Representations

Since dual graphs contain no positional information, translational symmetry has no meaning in this context and thus it is difficult to define aperiodicity. However, in Section 4.3 we showed that an inflation specification generates aperiodic tilings if it possesses unique inflation. Here, we show how to convert such a spatial inflation specification into an inflation specification for dual patches. It is important to work in this direction, because a spatial inflation specification does not necessarily exist that is equivalent to an arbitrary inflation specification for a dual patch. Only if a set of prototiles tiles the plane, can a dual patch be said to represent their structure. If inflation specification for dual graphs is a representation of a spatial inflation specification with unique inflation, then the inflation specification for dual graphs models the structure of aperiodic tilings.

Given a tiling and a region $\mathcal{R}$ within it, it is straightforward to construct a patch $G = (T, \mathcal{P}, \mathcal{A})$ that represents the dual graph of $\mathcal{R}$. An element $t \in T$ is created for every tile in $\mathcal{R}$ and $\mathcal{P}(t)$ is defined according to the shape of this tile. Coincident edges of $\mathcal{R}$ result in adjacencies defined by the function $\mathcal{A}$. If $\mathcal{R}$ is an entire tiling or an unbounded region then T will contain a countably infinite number of tiles. However, we will primarily be concerned with finite patches consistent with one or more infinite tilings.

Patch-sides can also be converted in this manner. For each edge on a patch-side, an element of the dual graph's patch-edge is formed by taking the prototype of the adjacent tile and the facet that corresponds to the edge. It is important to order these edges in such a way that whenever patches are composed, the appropriate tiles of their patch-edges are joined together. Whenever a patch-side is enumerated, we first check to see if a spatially congruent patch-side has been enumerated previously. If this is the case and the same side of these patch-sides face outward from their patches, then the edges of the new patch-side are enumerated in the same order as the other. If opposite sides face outward then the edges are enumerated in reverse. Figure 5.5 shows two patch-sides that would be enumerated in opposite order. In two dimensions, it is sufficient to list all edges in clockwise order. Whenever patch-sides S and S' can be juxtaposed, their edges will be enumerated in opposite order and the mapping $\mathcal{M}_{SS'}$ will correctly join adjacent tiles.

Spatially, an inflation specification is a protoset P' with prototiles composed of tiles from another protoset P in such a way that P and P' are surrogates. Each prototile p of P corresponds to a patch that forms a prototile of P'. I_p is formed by converting this spatial patch to a dual patch. All the I_p and $I_{(p,e)}$ of an inflation specification equivalent to the spatial inflation specification are constructed in this way. Tilings generated by this inflation specification will model tilings admitted by the spatial inflation specification.

5.2 Exploring Inflated Patches

The construction of successive levels of inflation of a patch by composition of the inflation patches I_p can be conducted algorithmically according to a process that follows naturally from the notations we have developed. A protoset specification, inflation specification, the operations $\mathcal{M}$ and C, and the construction of $\text{INF}(G)$ from G can represented in a straightforward manner.

Given a patch G we can form arbitrarily large patches $\text{INF}^n(G)$ inflated from that patch. Because we will be examining tilings with unique inflation (see Section 4.3), every finite region $\mathcal{R}$ of any tiling formed from the protoset can be found inside the patch produced by some level of inflation l of some vertex configuration v. If Local Isomorphism also holds for the protoset (Theorem 3 holds if some inflation of a tile contains all configurations) then a single tile $\text{INF}^0(p)$ will, after some number of inflations m, contain every vertex configuration, and then $\text{INF}^{m+l}(p)$ will contain $\mathcal{R}$. Here we will restrict our attention to the growth of inflated patches from a single prototile, although vertex configurations or arbitrary patches can be inflated in the same manner.

5.2.1 Efficiency

Now we need to consider the efficiency of constructing such patches. Say that we have $\text{INF}^l(p)$ and wish to construct $\text{INF}^{l+1}(p)$, which will contain $\mathbf{k}$ tiles and that none of our prototiles contains more than $\mathbf{f}$ facets. $\text{INF}^{l+1}(p)$ will have $O(\mathbf{kf})$ edges and so can be represented using $O(\mathbf{kf})$ storage space. Recall from Section 5.1.4 and Section 5.1.5 that $\text{INF}^{l+1}(p)$ is formed as $C_{\mathbf{GS}}$ where $\mathbf{G}$ is a set of patches to be joined along the patch-sides $\mathbf{S}$. $\text{INF}^{l+1}(p)$ is formed by the following sequence of steps:

1. For every tile in $\text{INF}^l(p)$ of prototype p, create a structurally equivalent copy of the inflation patch I_p and add it to $\mathbf{G}$.
2. For every adjacency in $\text{INF}^l(p)$, create a pair (S, S') and add it to $\mathbf{S}$.
3. Form a union of the elements (T,P,A) of each of the patches in $\mathbf{G}$.
4. Add to this union the adjacencies $\mathcal{M}_{SS'}$ for each (S, S') in $\mathbf{S}$.

Each step can be completed in time proportional to the size of the result, $\text{INF}^{l+1}(p)$. All the patches of $\mathbf{G}$ are structurally contained by $\text{INF}^{l+1}(p)$. Thus copying them (step 1) and unioning them (step 3) require time proportional to the size of $\text{INF}^{l+1}(p)$. Step 2 adds an element to $\mathbf{S}$ for each adjacency of $\text{INF}^l(p)$ and there are fewer adjacencies in $\text{INF}^l(p)$ than in $\text{INF}^{l+1}(p)$. In step 4, for some (S, S'), $\mathcal{M}_{SS'}$ is formed by traversing the edges of S forwards while traversing the edges of S' backwards. This will produce the adjacencies added by the connection of S and S' in time proportional to the number of these adjacencies and step 4 creates only a subset of the adjacencies of $\text{INF}^{l+1}(p)$. Therefore, $\text{INF}^{l+1}(p)$ can be formed in $O(\mathbf{kf})$ total time.

Note that no reference is made to the size of the inflation patches I_p in this result. A larger or smaller inflation specification does not affect the order of operation when the result of the construction contains a fixed number of tiles. At successive levels of inflation, the number of tiles is multiplied by a factor proportional to the number of tiles in the inflation specification, but this is accounted for by the inclusion of the number of tiles, $\mathbf{k}$, itself, in the order of operation. Because $\mathbf{k}$ grows exponentially in terms of the level of inflation, l, the number of tiles in $\text{INF}^l(p)$ is some fraction of the number of tiles in $\text{INF}^{l+1}(p)$. In order to construct $\text{INF}^l(p)$ from $\text{INF}^0(p)$, the previous levels of inflation, $\text{INF}^i(p)$ for $2 \le i \le (l-1)$, must first be constructed. However, because $\mathbf{f}$ is fixed for this construction and the number of tiles in these patches grows exponentially, the entire sequence of constructions still requires time proportional to the number of tiles in the patch $\text{INF}^{l+1}(p)$, $\mathbf{k}$. Therefore, given a prototile specification and an inflation specification, the inflated patch $\text{INF}^l(p)$ can be formed in $O(\mathbf{kf})$ time where $\mathbf{k}$ is the number of tiles in $\text{INF}^{l+1}(p)$ and $\mathbf{f}$ is the number of facets of the prototile with the most facets. We normally regard $\mathbf{f}$ as fixed by some particular protoset and consider the construction of $\text{INF}^l(p)$ as $O(\mathbf{k})$.

5.2.2 Limitations

The method of exploring aperiodic tilings just described is useful for applications that visit most or all vertices of a fixed inflated patch. In this case, $O(\mathbf{k})$ time is unavoidable and if information is stored with vertices that are visited then $O(\mathbf{k})$ space is required as well. However, this method is inconvenient if a user wishes to explore only a small section of an inflation patch or if the area to explore is determined while the exploration is being performed. For example, a physicist might want to simulate the electrical properties of an infinitely large quasicrystalline tiling in which each of the tiles is initially assigned a neutral charge. Assigning charge to some tile would then set off a process that would modify the charge of neighboring tiles and these tiles would, in turn, affect their neighbors. Eventually these effects would be sufficiently dissipated as to be considered irrelevant, but the distance at which this happens and the direction the charge travels would not be known in advance.

The above approach for representing aperiodic tilings has two important shortcomings. First, it is not expandable—the region to be explored must be determined in advance. A user of this system might be forced to request the construction of a patch large enough to encompass everywhere their exploration might conceivable proceed, but the actual exploration path might only involve a small portion of this patch. Second, this representation fails to take advantage of self-symmetries of the inflation process. Consider the l-th level of inflation of some prototile, $\text{INF}^l(p)$. If l is suitably large, then this patch will contain many tiles of any particular prototile type. Say that $\text{INF}^l(p)$ contains d prototiles of its own type p. When the next level of inflation $\text{INF}^{l+1}(p)$ is constructed, every tile of type p is replaced with the same inflation patch, I_p. This results in the repeated insertion of a large amount of duplicate information into the data structure. At level $l+2$, tiles of the same type in $\text{INF}^l(p)$ will have been expanded twice to form $\text{INF}^{l+2}(p)$, which will contain d duplicates of the patch $\text{INF}^2(p)$ and these copies will, in turn, contain d^2 copies of the patch I_p.

These problems can both be addressed by the observation that every patch $\text{INF}^l(p)$ is structurally contained by some patch $\text{INF}^{l+1}(p')$ at the next level of inflation. This property allows us to extend the size of a patch, constructing additional tiles beyond its boundary by finding a patch at the next level of inflation that structurally contain the current patch. Because any given patch could lie in an infinite number of tilings, there will often be several patches at the next level of inflation that could contain the current patch and the user will have the option of choosing which of these candidates the current patch actually lies inside. If a user exploring the patch $\text{INF}^{l}(p)$ discovers that the investigation needs to proceed beyond the patch's boundary then our data structure can construct a patch $\text{INF}^{l+1}(p')$ and pretend that the user had, in fact, really been exploring a subsection of this larger patch all along. This property will allow us to construct a method of addressing the tiles within a patch so that representing the structure of a patch is dramatically more efficient in terms of storage space. In Figure 5.7, larger and larger patches must be grown in order to encompass a user's path in a tiling by Robinson's triangles. In the following section, we prove our observation, concerning the self-similarity of inflation.

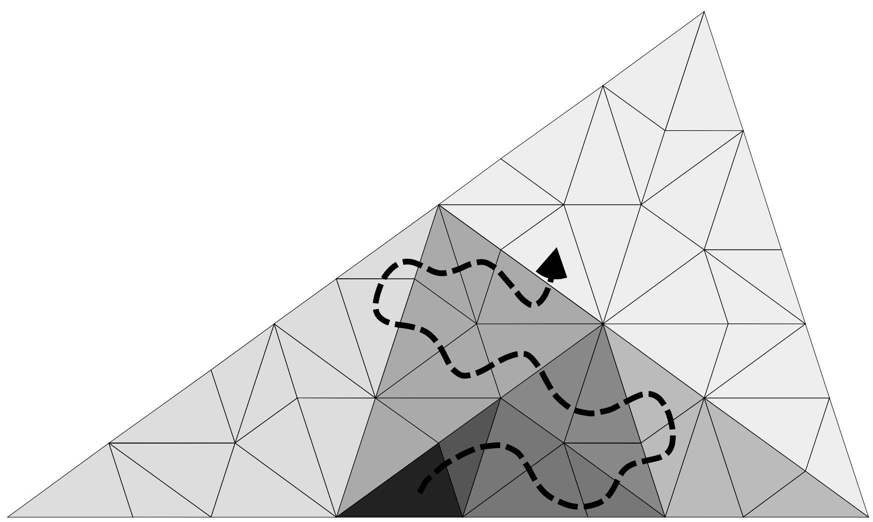

Figure 5.7: Growth of larger patches as a tiling by Robinson's triangles is explored. Initially, our data structure models only the single tile (the darkest triangle) from which the exploration begins. Whenever the user leaves the currently represented area, a larger patch must be constructed. Lighter grey shades indicate the areas added by an exploration path indicated by the arrow.

5.2.3 Self-Similarity of Inflation

The inflation specification consists of inflation patches that are inserted in place of each tile in a patch to form the patch's inflation. When we examine the series of inflations $\text{INF}^0(p)$, $\text{INF}^1(p)$, ... for all prototiles p, the inflation specification also defines the macro structure of these prototile inflations. $\text{INF}^1(p)$ is the inflation patch for p, $\text{INF}^2(p)$ is formed by inserting inflation patches in place of the tiles of $\text{INF}^1(p)$, and $\text{INF}^3(p)$ is formed by performing this substitution yet again. Each tile of $\text{INF}^1(p)$ has been substituted for twice in the formation of $\text{INF}^3(p)$. Another way of describing this inflation is that $\text{INF}^3(p)$ is formed by replacing every tile in $\text{INF}^1(p)$ of prototype p with the patch $\text{INF}^2(p)$. An alternative way of understanding the inflation process is that the patch $\text{INF}^l(G)$ is formed by substituting the patch $\text{INF}^{l-1}(p)$ for every tile of type p in G. This is not quite precise because we have not defined patch-sides for the $\text{INF}^n(p)$, but it is easy to imagine how this would be accomplished. Now it can be seen that any prototile inflation $\text{INF}^l(p')$ is structurally contained by the larger prototile inflation $\text{INF}^{l+1}(p)$ if and only if a tile of type p' lies in the inflation patch for p. If there are several such tiles, then $\text{INF}^l(p')$ can be located within $\text{INF}^{l+1}(p)$ in several ways.

Theorem 4 *If the inflation patch of a prototile p, $I_p = (T, \mathcal{P}, \mathcal{A})$, and $t_1, t_2, \ldots, t_n \in T$ are tiles of the same prototype $\mathcal{P}(t)$ then $\text{INF}^l(\mathcal{P}(t))$ is structurally equivalent to n distinct subpatches of $\text{INF}^{l+1}(p)$.*

To prove this, we first need to show that inflation preserves the structural equivalence of patches.

Lemma 4.1 *If G is structurally equivalent to n distinct subpatches of G^* then $\text{INF}(G)$ is structurally equivalent to n distinct subpatches of $\text{INF}(G^*)$*

We show this by examining the construction of the inflation of a patch from Section 5.1.5. Inflation of the patch G occurs by forming patches G_t for each tile t in G, where G_t is a copy of the inflation patch for the prototype of t. $\text{INF}(G)$ is formed by sewing together these G_t along patch-sides corresponding to adjacent edges of G.

If two patches G and G' are structurally equivalent then a function $\mathcal{H}$ pairs tiles of G with equivalent tiles of G'. Equivalent tiles in each patch have the same prototypes and are adjacent to equivalent tiles. The inflation of a patch is determined by the prototypes and adjacencies of its tiles, so the construction of $\text{INF}(G)$ is the same as $\text{INF}(G')$. Thus if G and G' are structurally equivalent then $\text{INF}(G)$ and $\text{INF}(G')$ are structurally equivalent as well.

Given G and G', distinct subpatches of G^*, $\text{INF}(G^*)$ can be divided into distinct subpatches structurally equivalent to $\text{INF}(G)$ and $\text{INF}(G')$. These subpatches are unions of the G_t from the construction of $\text{INF}(G^*)$ for all t in G or G' respectively. Now we can find the distinct subpatches promised in the lemma:

> Say G is structurally equivalent to $G_1, G_2, \ldots, G_n$, subpatches of G^*.
>
> For all $i \leq 1 \leq n$:
>
> > Let $G_i^* = (T, \mathcal{P}, \mathcal{A})$ where $T = \bigcup_{t \in G_i} G_t$
>
> $G_1^*, G_2^*, \ldots, G_n^*$ are distinct.
>
> $\text{INF}(G)$ is structurally equivalent to $G_1^*, G_2^*, \ldots, G_n^*$. □

The lemma carries us almost to Theorem 4. The theorem is trivially true for $l = 0$ since we are given that the tiles $t_1, t_2, \ldots, t_n$ are in $I_p = \text{INF}^1(p)$. Thus, $\text{INF}^0(\mathcal{P}(t_i))$ is structurally equivalent to n distinct subpatches of $\text{INF}^1(p)$ and the lemma can be iterated to show that $\text{INF}^l(\mathcal{P}(t_i))$ is structurally equivalent to n distinct subpatches of $\text{INF}^{l+1}(p)$ for all $l > 0$. □

5.3 The Address Adjacency Algorithm

An *address* is a series of digits that identifies a single tile within some patch $\text{INF}^l(p)$ where l is the number of digits in the address. Digits are chosen from tiles in the inflation patches. Given such an address, we call $\text{INF}^l(p)$ its *extent*. A single-digit address A_1 is a tile within the inflation patch of some prototile p_1. If a more significant digit is added to A_1 to form the address A_2, this two-digit address locates a tile within the second level of inflation of some prototile p_2. However, we form addresses in such a way that the extent of A_1 is structurally contained by the extent of A_2. If I_{p_2} contains n tiles of type p_1 then Theorem 4 says that n different sections of the extent of A_2 are structurally equivalent to A_1 and the second digit of A_2 is used to identify one of these sections. Each additional more significant digit of an address describes the location of the patch addressed by the previous digits within a larger patch at the next level of inflation. Given an address A_l of length l, the tile's location within the patch $\text{INF}^l(p)$ is found by taking the digits from most to least significant. The most significant digit of the address identifies a subpatch of $\text{INF}^l(p)$ within which the addressed tile is located. This patch is structurally equivalent to $\text{INF}^{l-1}(p')$ where p' is the type of the address made from the first $l-1$ digits of A_l. This smaller address can, in turn, be used to narrow the location of the tile addressed by A_l to an even smaller subpatch of $\text{INF}^l(p)$ that is structurally equivalent to $\text{INF}^{l-2}(p')$. When this process is continued for all l digits, the subpatch is narrowed to a single tile.

It is notationally convenient to define a disconnected patch $I = (T_I, \mathcal{P}_I, \mathcal{A}_I)$, which contains all of the inflation patches of the form I_p. If for any $p \in P$, $I_p = (T_p, \mathcal{P}_p, \mathcal{A}_p)$ then let $I = (\bigcup_{p \in P} T_p, \bigcup_{p \in P} \mathcal{P}_p, \bigcup_{p \in P} \mathcal{A}_p)$. Then an *address* is a list of the form $(a_i \mid a_i \in T_I, 1 \leq i \leq l)$ such that for all $1 \leq i \leq (l-1)$, if $a_i \in T_p$ then $\mathcal{P}_I(a_{i+1}) = p$. We will often write the number of digits, l, of an address A as $|A|$. If the addresses A and A' have the same last digit and $|A| = |A'|$ then the addresses have the same extent and we say the two addresses are *consistent*.

Each digit of an address is the position of a tile in one of the inflation decorations I_p inside T_I. The only limitation on the digits of an address is that the type of decorated prototile that the digit is inside, p, must be the same as the prototile type of the next digit. This requirement guarantees that if a digit is added to an address the extent of the new address structurally contains the extent of the shorter address. $|A|$ is logarithmic in terms of the number of tiles in the extent of A, which allows us to store locations within very large patches using a relatively small number of bits. If $|A|$ is countably infinite then we call A an *index sequence*. An index sequence is, in a sense, the location of a tile within an infinitely large inflated patch. An index sequence completely determines a tiling that surrounds a single tile.

5.3.1 Adjacency

It is not remarkable that we can address patches of tiles using a logarithmic number of digits. However, our addressing scheme follows the self-symmetrical structure of the inflations of tiles. This is important because we would like to be able to determine which tiles are adjacent by examining their addresses. Tiles that are near each other are likely to have the same trailing digits. For example, two consistent addresses that differ in only their first two digits locate tiles that are both within a section of their extent that is structurally equivalent to some $\text{INF}^2(p)$. Spatially, tiles that differ only in their first n digits can be contained by a ball of radius r^n where r is the radius of a ball that contains the largest inflation patch. However, not all nearby tiles have similar addresses. Two adjacent tiles might differ in every digit but the last.

In order to use addresses to representing patches of tiles, a method must be found for computing adjacencies. The first step towards finding such a method is to analyze the ways in which two adjacent tiles can be related at two successive levels of inflation. Throughout this discussion, we will use subscripts to refer to elements of the inflation patches where $I_p = (T_p, \mathcal{P}_p, \mathcal{A}_p)$ for all $p \in P$.

Deflation Adjacency Rules We first form two mappings X and N called the *external* and *internal edge deflation rules*. Consider the one-digit address $A = (a)$ that is the location of a tile a in an inflation decoration I_p. We are interested in the location of one of the tile's sides $e \in \mathcal{F}(\mathcal{P}_I(a))$. If the edge (a, e) is external then it lies on some patch-side $I_{(p,e')}$ and $X : (a, e) \mapsto (e')$. If the edge is internal then it is coincident with some other edge (a', e') and $N : (a, e) \mapsto (a', e')$. Figure 5.8 shows tiles related by internal and external rules. X and N are formed with respect to any edge in I as follows:

$$X = \{((a, e), (e')) \mid a \in T_p, (a, e) \in I_{(p,e')}\}$$
$$N = \{((a, e), (a', e')) \mid a, a' \in T_p, \mathcal{A}_p : (a, e) \mapsto (a', e')\}$$

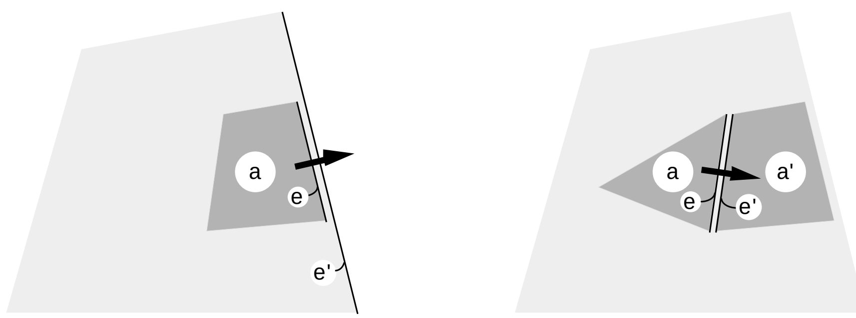

Figure 5.8: An external (left) and internal (right) edge deflation rule. Moves leaving the tile a by its side e are shown with arrows. An external move also leaves the inflation patch by its patch-side e', while an internal move enters another tile a' by its side e'.

Inflation Adjacency Rules The *backwards edge inflation rules*, B, involve a second inflation decoration $I_{p'}$ joined to I_p along some patch-side. Because of the way in which we will be using B later, we let $a^* \in T_I$ be any tile with prototile type p' and $e^* \in \mathcal{F}(p')$ be a facet where $I_{(p',e^*)}$ is the edge of $I_{p'}$ along which the joining occurs. If (a, e), an edge in I_p is coincident with the edge $(a', e') \in I_{(p',e^*)}$ then $B : (a, e, a^*, e^*) \mapsto (a', e')$. For a picture of tiles related by a backwards rule, see Figure 5.9. The mapping is defined over all (a, e, a^*, e^*) as follows:

$$B = \{((a, e, a^*, e^*), (a', e')) \mid a \in T_p, S = I_{(p, X(a,e))}, \\ S' = I_{(\mathcal{P}_I(a^*), e^*)}, \mathcal{M}_{SS'} : (a, e) \mapsto (a', e')\}$$

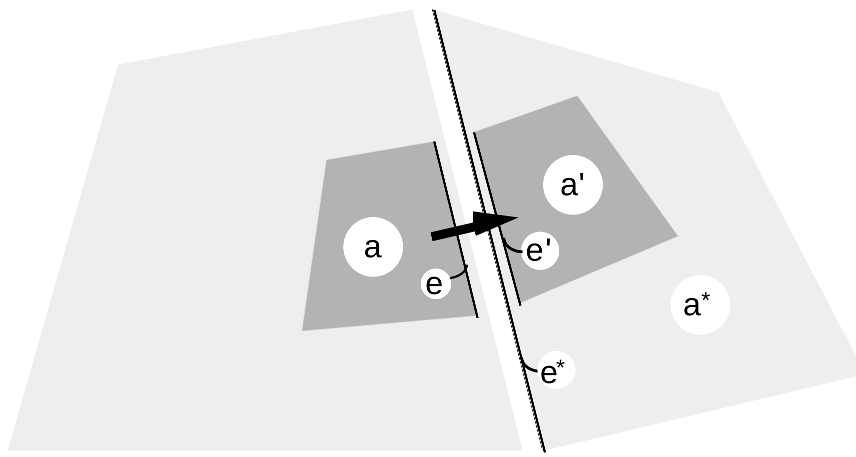

Figure 5.9: A backwards edge inflation rule. When the inflation patch containing tile a borders the side $e*$ of an inflation patch $a*$, side e of tile a is coincident with side e' of tile a'.

The Adjacency Algorithm The address of a tile gives its location within increasingly large patches of the form $\text{INF}^l(p)$. If a tile is located on the boundary of the first n of these patches then exiting the tile by the side that lies on the boundary of $\text{INF}^n(p)$ also exits the first n inflation patches. This move does not leave the patch $\text{INF}^{n+1}(p)$ and all digits after the first n would be unchanged. The move would, in turn, enter n new, successively smaller subpatches of $\text{INF}^{n+1}(p)$ that determine the address of the move's destination.

Given the address A of a tile t and one of its edges e_1, the address $Asym'$ of the tile adjacent to t along e_1 can be found by an algorithm that uses external rules to scan from left to right across digits of A until an internal rule is encountered. Then, the algorithm scans back from right to left while using backwards rules to fill in the digits of $Asym'$ that differ from $Asym$. Each step of this process involves a fixed-size table lookup and the algorithm is $O(|A'|)$ in the worst case. We use a left-facing arrow ($\leftarrow$) to indicate assignment to variables.

$A' \leftarrow A$
$d \leftarrow 1$
while $(e_{d+1}) \leftarrow X(a_d, e_d)$
    $d \leftarrow d + 1$
$(a_d^{'}, e_d^{'}) \leftarrow N(a_d, e_d)$
while $d > 1$
    $(a_{d-1}^{'}, e_{d-1}^{'}) \leftarrow B(a_{d-1}, e_{d-1}, a_d^{'}, e_d^{'})$
    $d \leftarrow d - 1$

We have assumed above that d never becomes larger than $|A|$. In fact, this will not be true whenever a move leaves the extent of A. In practice, whenever a_d is accessed but has not been defined then the algorithm asks the user to provide an additional digit that it will use to extend A. This digit defines the structure of the area beyond the extent of A. In this way, the algorithm can be used to travel along a path of adjacent tiles in a tiling corresponding to an arbitrary index sequence.[4] Because a move may extend an address, the algorithm's efficiency is $O(|A'|)$ rather than $O(|A|)$. The system only stores enough digits to represent a current inflated patch that contains every previously visited tile and whenever a tour ventures out of this region the user must choose a new patch at a higher level of inflation to locate the current patch within.

[4]Note that not all tilings admitted by the protoset can be generated by inflation of a single tile according to an index sequence. If a tiling contains a line of symmetry then moves that cross this line result in infinite extension of the address and nontermination of the algorithm. Also, if local isomorphism does not hold for the protoset then some tilings cannot be formed by inflation of a single tile. Section 6.1 refines the algorithm for addressing tiles within inflations of a vertex configuration rather than a single tile in order to eliminate these problems.

Chapter 6

Care and Feeding: Additional Refinements

This section presents problems, extensions, and other details we fell it is important to mention concerning our exploration data structure. In order to use our technique, a user must first provide a protoset specification, an inflation specification, and a method for generating additional digits when an address must be extended.[1] Once this information is provided, if the user chooses the address of some tile and one of its sides then a data structure can find the address and side of the adjacent tile. However, if the user has chosen the edge of a tile that lies on a line of symmetry then the algorithm presented in Section 5.3.1 will not terminate. Recognizing these situations require additional analysis that we describe in Section 6.1. Users may also want more information about tiles than their adjacencies. Physicists and mathematicians are often interested in classifying the local configurations of tiles that can occur in a tiling and we discuss this concept in Section 6.2. Section 6.3 provides a method for determining the orientation of a tile. These concepts suggest potential avenues for future research.

[1]A user could provide the index sequence by giving a repeating pattern that forms the infinite address of the tiling to be explored. Alternatively, whenever the result of a user's request for an adjacency results in a longer address the system could ask the user to provide additional digits to use for extending the address. Note that these extensions must be appropriate. According to the definition of an address, only certain digits may follow others.

6.1 Symmetry and Infinite Crossings

In Figure 6.1 an inflation patch contains an external tile of its own prototype. Furthermore, both the tile and the inflation patch have the same orientation. If we attempt to determine the adjacency of an edge of the tile that lies on the boundary of the inflation patch using the adjacency algorithm of section Section 5.3.1, the algorithm will encounter in infinite sequence of external moves. Whenever the edge of a tile lies on the boundary of the inflation patch extending the tile at every level of inflation, we say the edge is an *infinite crossing*.

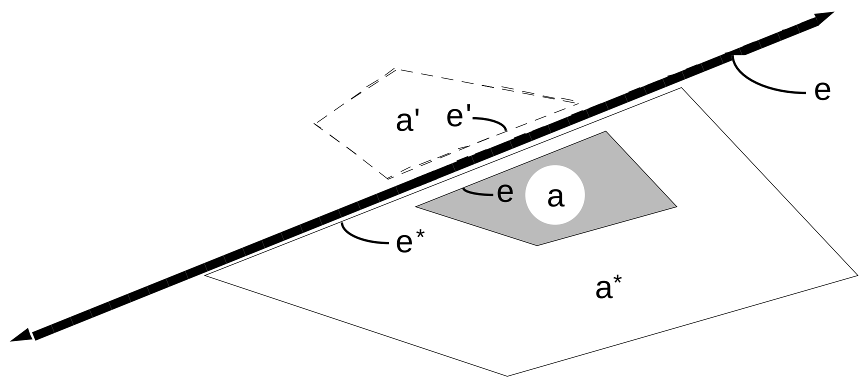

Figure 6.1: An infinite crossing results in nontermination of the adjacency algorithm. (a) The inflation patch $a*$ contains an external tile a of the same type and orientation. If addresses are extended in such a way that this relationship occurs at every level of inflation then no inflation patch contains both a and its neighbor a'. The edge e of a that lies on the patch-side e^* of a^* is an infinite crossing and the adjacency algorithm will never reach an external rule for this edge. e lies on a patch-side of every inflation of a, or the ray e^∞.

6.1.1 Seams

If a tiling contains a line of symmetry then the edges along this line are all infinite crossings.[2] Symmetrical tilings cannot be formed by inflation of a single tile and

[2]This is true provided no prototiles are perfectly symmetrical. If a tile crosses a line of symmetry then it must be symmetrical along this line. Furthermore, all inflations of the tile would be symmetrical as well. Such a protoset can be modified by cutting in half all prototiles that are

thus it is no surprise that our data structure has difficulty exploring them. If an edge is an infinite crossing, we know that an infinitely large section of edges are also infinite crossings. In Figure 6.1, a is identified as an infinite crossing because it is of the same type as the inflation patches. However, every edge along e^∞ is also an infinite crossing since after some number of inflations the inflation patches of any tile along e^∞ are identical to inflation patches of a. We call such a ray of infinite crossings a *seam*. In three dimensions, a seam is topologically equivalent to a planar section with infinite area. Seams are infinitely long in one or both directions and are not necessarily straight.

In two dimensions, extending tiles may generate no seams, one linear seam, or two ray-like seams, as shown in Figure 6.2. Extensions in three dimensions may have no seams, one planar seam, two seams that are half-planes, or at least 3 fans bounded by two rays emanating from a point. Figure 6.3 shows inflations that would generate seams of these types and the infinitely extending tiles that are the limit of these inflations. Since the tilings we consider can be deflated and the topology of seams is preserved under inflation, we can see that all seams must meet at least a point. Figure 6.4 is a hypothetical two dimensional tiling with seams that do not meet at a point, and it can be seen that this structure is inconsistent with inflation. Thankfully, all tilings can be generated by inflation of a single tile, two adjacent tiles, an edge configuration, or a vertex configuration, as shown in Figure 6.5. Any of these shapes can be represented as a patch and we call such patches *configurations*.

6.1.2 Addresses

It is not difficult to add additional information to our addressing system to allow the representation of tilings containing seams. An additional initial digit is added to addresses that identifies one section of a configuration. Instead of a tile in one of the inflation patches, the first digit is now a tile in a patch that is the configuration we wish to explore. The first digit of a tile's address narrows its location in one of these sections and later digits pinpoint its location extended from that section. This corresponds to inflation of an external tile rather that an internal tile in the manner of Theorem 1. The limit of such an inflation is not space, but a group of these inflations can be placed together at a point in the manner of a vertex configuration in order to fill space. The first digit of an address selects the location of a tile in one of these large regions of the tiling and later bits identify the tile's position within that region in the same way as an ordinary address. The prototype of the configuration section identified by the first digit must be the same as the

symmetrical in this way, forming left and right-handed prototiles as Robinson's triangles are for Penrose's kites and darts and Danzer tiles are for the octahedra they form.

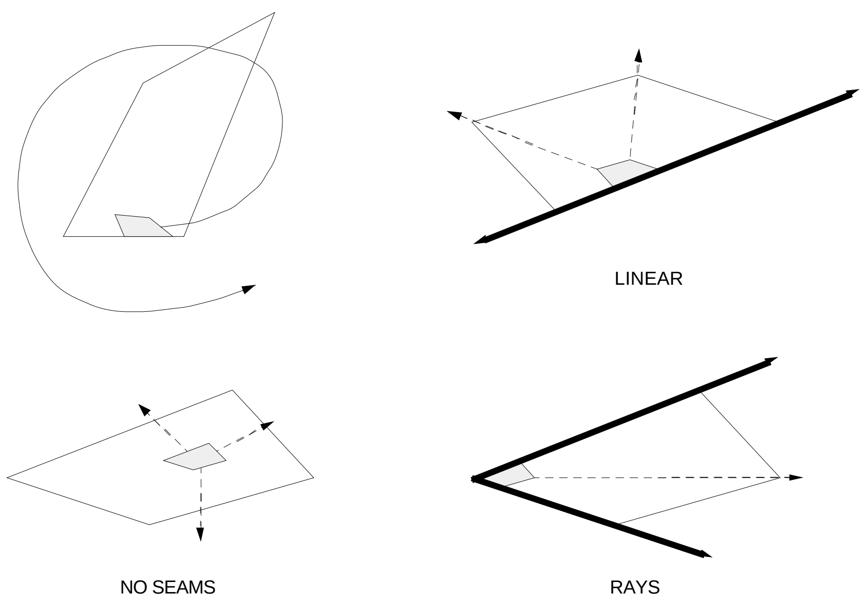

Figure 6.2: Classification of seams in two dimensions. Every tile is located within some other tile at the previous level of inflation. The pattern of tile within tile can grow in the three ways shown above. If an edge of the small shaded shape lies in the middle of the same side of the larger shape then a half-plane is grown with a linear seam at the boundary. If the small shaded shape is at a corner than a fan bounded by two ray seams is formed. Otherwise, continued extension tiles the plane and no seams are created.

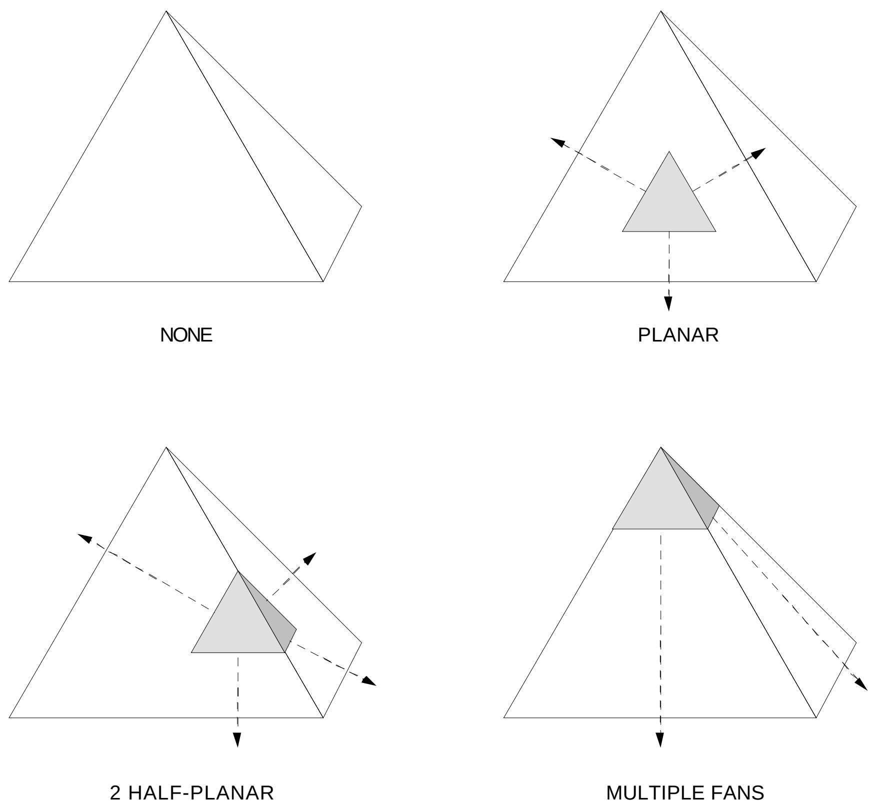

Figure 6.3: Classification of seams in three dimensions. In each section, a small shaded pyramid extends to the larger pyramid. Arrows indicate the growth of seams. Note that the shaded pyramid for the upper-left cannot be seen because it is obscured within the larger pyramid. Also, in the lower right the third fan seam is obscured.

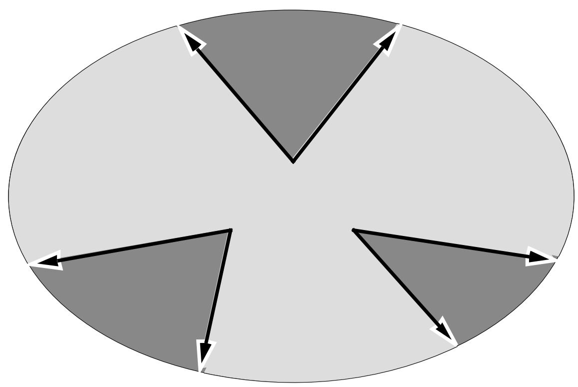

Figure 6.4: All seams must meet at a point. 3 fans are bounded by the 6 seams shown in this hypothetical tiling, dividing the plane into four sections. However, the triangle in the middle section is awkward. It remains a fixed size at all levels of inflation. Under inflation, tiles contain balls with increasingly large radii and one of these will contain one of the three points from which the seams emanate. Such a tile is inconsistent with the definition of a seam so this situation is impossible.

prototype of the second digit and the system for extending addresses must extend them in a manner consistent with this configuration. Not all possible extensions of tiles grow in such a way that they can fit with other extended tiles around a vertex configuration. Whenever sections of the configuration are extended they must be extended in such a way that the infinite crossings making up each of the seams (initially the internal edges of the configuration) map to individual planar sections of the extended inflation patches. The extent of an address of length n is $n-1$ levels of inflation of the configuration being explored.

A tile's first digit will change only when an infinite crossing is encountered. Leaving aside for a moment the problem of recognizing this situation, we can extend the adjacency algorithm (see Section 5.3.1) for addresses with configurations. We assume that the function $\mathcal{E} : (a_1, a_2, \ldots, a_n) \mapsto a_{n+1}$ is used to extend any address by one digit. For addresses of at least two digits, extension occurs in the same way as when addressing tiles within tilings without seams as in Section 5.3.1. An extension of a tile in the configuration is one of the inflation patches that contains a tile of its prototype. Many such extensions may be possible, but the extension function will be chosen so that these extensions are consistent with the seams of the tiling. The configuration tile must be an exterior tile of the inflation patch and the

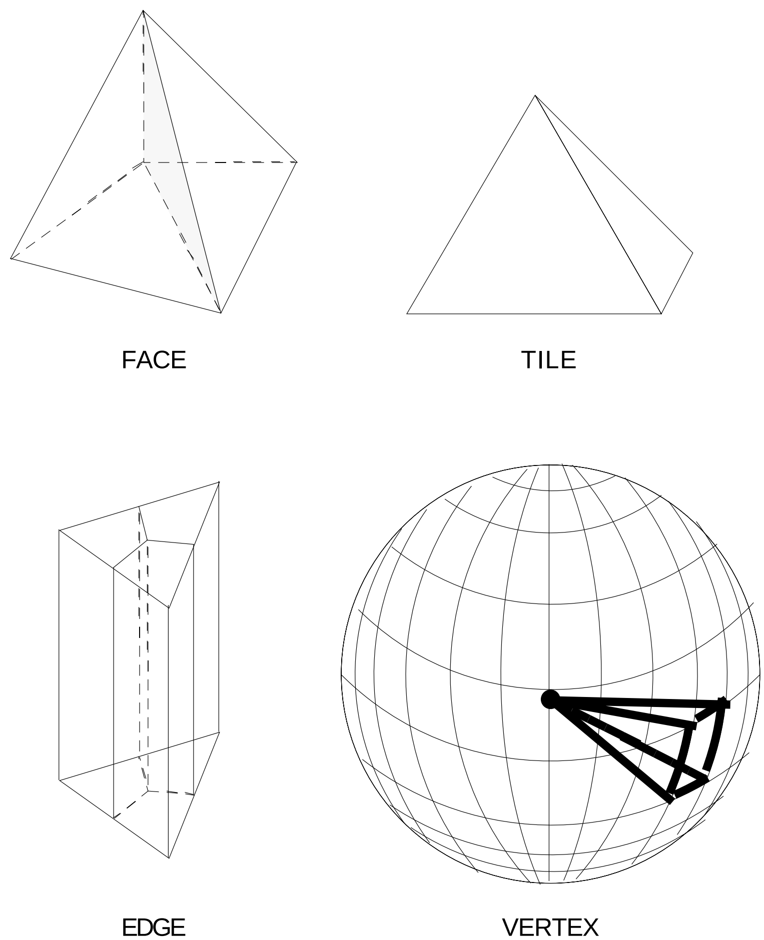

Figure 6.5: Kinds of configuration. A point in a tiling can be contained by tiles in four ways. The point can lie within a tile, on a face between two tiles, at an edge between multiple tiles, or at the corner of many tiles. We call an arrangement of tiles in this manner a configuration. Only one tile of the vertex configuration is shown. Imagine that such shapes radiate from the center of the globe to every section of its surface. Usually we do not include the trivial case of a single tile when we speak of a configuration.

extension patches for all the tiles of the configuration must fit together so that the configuration tiles within them are arranged in the same way as the configuration, as seen in figure Figure 6.6. $\mathcal{E}$ must be defined so that extending a single digit address (a_1) yields the location of the configuration tile (a_1) within its extension patch. This property holds for addresses with two or more digits as well. An address with n digits represents the location of a tile within a patch of the form $In-1p$ and the first digit represents a location within the configuration. The n-th extension of the single digit addresses for each tile of the configuration can be fit together around a point and the configuration will lie around the point in this formation. As can be seen in Figure 6.6, the tiling formed by an infinite series of extensions will have a seam for each adjacency in the configuration.

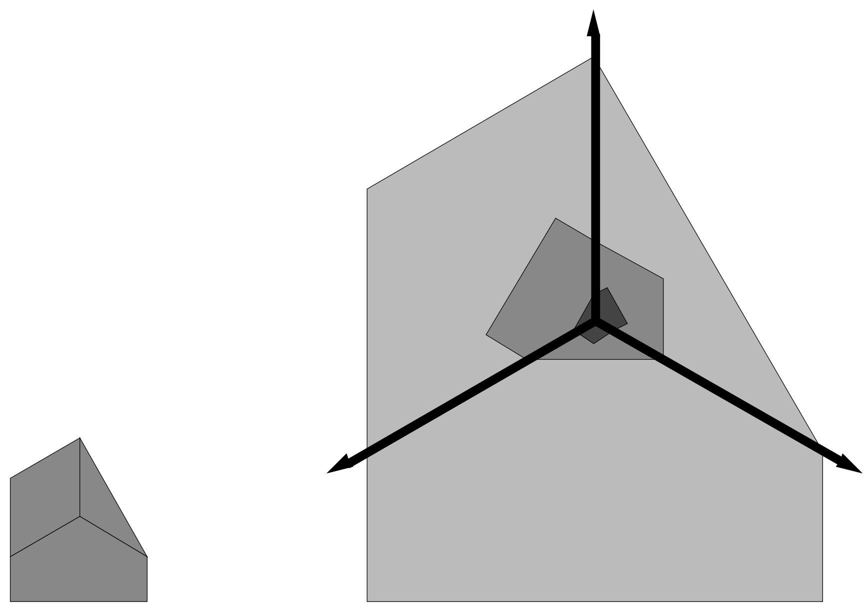

Figure 6.6: Consistent extension of the sections of a vertex configuration. Each section must be extended so that it remains edge-to-edge with its neighbors. In general, sections with the same angle (or more complicated congruent shapes) can be swapped at various levels of inflation.

6.1.3 Adjacency

Determining adjacencies is now more intricate than in Section 5.3.1. We first introduce a set $\mathcal{C} = \{(A_i, e_i, \mathcal{E}_i) \mid 1 \leq i \leq n\}$. Here $A_i = (a_1, a_2, \ldots, a_n)$ is the digits of an address representing a tile t in an inflation patch with an extent of the form $\text{INF}^{n-1}(p)$. e_i is an edge of t, and $\mathcal{E}_i$ is the function used to extend addresses. $\mathcal{C}$ is defined over all addresses, edges, and consistent extension functions. Clearly this infinitely large set cannot be algorithmically represented, but we will postulate its existence in order refine the adjacency algorithm for use with configurations. If A is an address, $G = (T, \mathcal{P}, \mathcal{A})$ is its configuration, and e_2 [3] is the side of the tile whose neighbor we wish to find, then $A' = (a_1, a_2, \ldots, a_{n'})$ is this neighbor along its side $e_2^{'}$.

if not $(A, e_1, \mathcal{E}) \in \mathcal{C}$ then
  let $(a_2^{'}, a_3^{'}, \ldots, a_{n'}^{'})$ and $e_2^{'}$ be the adjacency of
    $(a_2, a_3, \ldots, a_n)$ and e_2 computed by the standard
    adjacency algorithm of Section 5.3.1

otherwise:

  if $a_2^{*} = \mathcal{E}((a_1))$
    $e_3 = X(a_2, e_2)$
    $(a_2^{*}, e_1) \in I_{(\mathcal{P}_{I(a_3)}, e_3)}$
  $\mathcal{A}((a_1, e_1)) = (a_1^{'}, e_1^{'})$
  $d \leftarrow 2$
  while $d \leq |A|$
    $a_d^{'} \leftarrow \mathcal{E}(a_1^{'}, a_2^{'}, \ldots, a_{d-1}^{'})$
    $e_d^{'} \leftarrow X(a_{d-1}^{'}, e_{d-1}^{'})$
    $e_d \leftarrow X(a_{d-1}, e_{d-1})$
    $d \leftarrow d + 1$
  while $d > 2$
    $(a_{d-1}^{'}, e_{d-1}^{'}) \leftarrow B(a_{d-1}, e_{d-1}, a_d^{'}, e_d^{'})$
    $d \leftarrow d - 1$

[3] We call the side e_2 instead of e_1 because we no longer move from the first digit except when an infinite crossing is encountered and in this case, we must calculate the side crossed for the first digit using the algorithm.

If A, e_2 is an infinite crossing, then this algorithm works in three steps. First, the algorithm must identify the effect of the infinite crossing as a move between regions of the configuration—from (a_1, e_1) to $(a_1^{'}, e_1^{'})$. a_2^*, the extension of (a_1) is the location of a_1 within the same inflation patch that a_2 is located within, and this relationship is used to find $(a_1^{'}, e_1^{'})$. Second, we find the extension of $a_1^{'}$ that matches with $a_{|A|}$ and the adjacent sides $e_{|A|}$ and $e_{|A|}^{'}$. Third, we can use the same process of backwards inflation as in Section 5.3.1 to fill in the correct digits of $|A'|$. Here, we stop at $d > 2$ rather than $d > 1$ because the digit a_1 was already calculated specially in the first step. Figure 6.7 illustrates these three steps spatially. The efficiency of this algorithm is linear in terms of the length of A' and in the case of infinite crossings, $|A| = |A'|$. The trick, however, is to find the set $\mathcal{C}$.

6.1.4 Identifying Symmetry

Now we wish to categorize the ways in which addresses can be extended to produce infinite crossings. We will do this by forming a directed graph called an *edge inflation roadmap*. Vertices of this graph are all facets of all prototiles, $\mathcal{F}(P)$. For each external rule $((a, e), (e^{'}))$, we place an arc from e to $e^{'}$, and the arc is labeled with a.

Any cycle in this graph is the address of an infinite crossing. Say $(p_1, e_1) \overset{a_1}{\rightarrow} (p_2, e_2) \overset{a_2}{\rightarrow} \ldots \overset{a_{n-1}}{\rightarrow} (p_n, e_n) \overset{a_n}{\rightarrow} (p_1, e_1)$ is a cycle. Then if the address a_1 is extended along the series $\{\overline{a}_i \mid \overline{a}_i = a_{i \bmod n}, i \geq 2\}$, a move from a_1 by side e_1 will correspond to an infinite series of external rules $((\overline{a}_i, e_{i \bmod n}), (\overline{a}_{i+1}))$ that form a seam.

In fact, any address with a suffix that is an infinite path through the graph produces an infinite crossing and some paths would never cycle. In general, it is impossible to determine whether an address and side are an infinite crossing without encoding an analysis of the extension function in order to analyze whether an infinite series of extensions is a path in the graph. If such a series contains infinitely many internal rules then the series is not an infinite crossing. This problem is the determination of whether an infinite string is in the language accepted by a finite automaton.

If the extension function contains a repeating cycle then the problem can be viewed as parsing for a cycle in the graph. If there are n nodes in the edge inflation roadmap, we make n copies of the roadmap. In each graph, a different node is chosen as a start and stop state. Each graph is a deterministic finite state machine that accepts cycles beginning at its start state and each state has one outgoing transition for every address digit that is an external rule from that state. A nondeterministic finite automaton can then be created with a start state that has an empty transition to the start state of each copy of the graph. This automaton parses the repeating

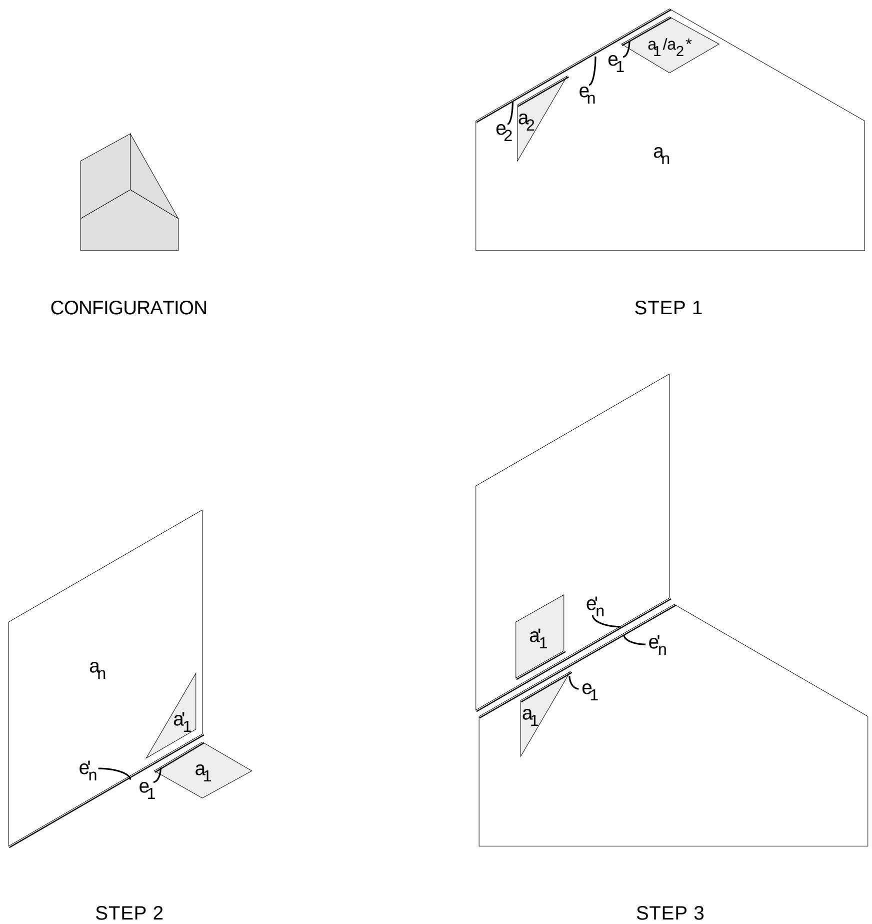

Figure 6.7: A move to a different region of symmetry. The algorithm takes place in three steps. Moving from a_2 by e_2 exits the whole section a_n by e_n. In step 1, we figure out the side of a_1 that also leaves by e_n. In step 2, we find the adjacent region in the symmetry configuration and set the first digit of the new address ($a_1^{'}$). Then we extend it to find the side $e_n^{'}$ that matches with e_n. Then in step 3 we use backwards rules to fill in the rest of A'.

unit of an address and accepts it if the address corresponds to an infinite crossing. If the automaton is converted to a deterministic finite state machine then it will have a state for every way of chosing a node in each copy of the roadmap, or n^n states where n is the number of facets in the protoset. Many of these combinations may be unreachable, but an explosion of states makes this approach practical only for small protosets.

The edge inflation roadmap could be used to automatically categorize the arrangements of seams that can occur in tilings formed from a protoset specification and inflation specification. An infinite path through the roadmap is an index sequence for an infinite, but incomplete section of space. Tilings formed from multiple pieces formed in this way have a seam for each adjacency in the configuration describing the way in which they are connected.

6.2 Configurations

A configuration is a group of tiles that share a point and completely surround it. The configurations of a protoset are closely related to local isomorphism and to the seams that can occur. An automated technique for generating all the possible vertex configurations for a set of prototiles would eliminate the painstaking analysis that is required to calculate combinations of tiles by hand. It is possible to exhaustively catalog all possible configurations of a set of prototiles given their geometry and matching rules by trying all possible ways of attaching their edges together. Each tile can attach to others in only a finite number of ways and depending on how sharply angled the tiles are, only a certain number of tiles can be fit around a point. Baake et al.[BBAK$^+$94] outline such a technique that they used to generate all possible vertex configurations for three-dimensional tilings with icosahedral symmetry made from an oblate and prolate rhombohedra. Trying all possible ways to match tiles together will, in general, produce some degenerate vertex configurations that are admitted by the matching rules but cannot be extended to tile the plane. We will describe a technique that uses—you guessed it—an inflation specification to discover all configurations that actually occur in tilings generated by inflation. We will search for all types of configurations (vertex, edge, and adjacency) as depicted by Figure 6.3. Edge and adjacency configurations can be thought of as partial vertex configurations—tiles that share more than a single point.

We call the set of all configurations an *atlas*. We begin by assuming that local isomorphism holds for the protoset in question. Thus we can say that all configurations occur within some level of inflation of a single tile. We also assume that a method can be developed for finding the configurations within a patch. This information cannot be recovered from a dual graph alone because this structure does not reveal how tiles are related to each other except that they are adjacent by certain sides. Information could be maintained that would identify the vertices that coincide when two tiles are adjacent.[4] The algorithm could then identify all subpatches that share a vertex. Unless such a vertex lies on the boundary of the patch, then the maximal subpatch of tiles surrounding it is a vertex configuration. If two subpatches are structurally equivalent then they are the same configuration.

In the first step of our technique we examine all the inflation patches and add all configurations within them to the atlas. If the inflation patches are so small that they contain no vertex configurations then we inflate them until they do. In Theorem 1 we saw that inflation of one tile produced a tiling of the plane. Because of local isomorphism, we know that every configuration exists within some inflation of that tile. By examining enough inflations of the tile, we can find all these configurations.

When looking for configurations inside inflations of a patch G, it is important to realize that new configurations cannot occur within tiles of the patch, but rather only along the edges where patch-sides are attached together. We see that we only have to check the new adjacencies formed by an inflation for new configurations. Furthermore, there is no need to keep track of an entire patch for multiple levels of inflation. Any configuration that has not yet been discovered must lurk inside an inflation of one of the configurations that was just discovered on the current step. If a vertex configuration has been expanded previously then all vertex configurations within it have already been found. Our algorithm proceeds to search all inflations of the initial patches by expanding the sections in which tiles are attached together in new ways. At each step the algorithm inflates every configuration discovered in the previous step and looks for new configurations along the patch-sides of the inflation patches that are placed together to form this inflation. Since there are a finite number of configurations, eventually no new configurations will be formed. When this occurs the atlas is complete.

An atlas can be used to form the most restrictive possible set of matching rules for a set of prototiles. Every adjacency in every vertex configuration must be allowed by the matching rules and if all these adjacencies are allowed then any tiling admitted by the prototiles can be formed.

The configurations that are generated can also be used to analyze symmetries that are possible with the protoset. Any of the configurations can be inflated. If we observe a point that is surrounded by a configuration and inflate the configuration

[4] There are a finite number of such vertices. Spatially, edges are open discs of points shared by exactly two tiles. Open arcs lie on the closure of two or more edges, and vertices are single points that lie on the closure of more than two arcs. A tile that had an infinite number of vertices would have an infinite number of edges.

then a new configuration will be formed around the point. The tiles surrounding the point will become more and more pointy at increasing levels of inflation. We could form a graph to indicate which configurations subdivide to which others. Since there is only one way to inflate a tiling, there is one arc out of each node of the graph. These levels of inflation will eventually cycle among some set of configurations. Every sequence of inflations of configurations, when reversed, is a system for extending patches of tiles from each section of the configuration in a consistent manner to form a tiling containing seams for every adjacency of the configuration.

6.3 Orientation

To this point, we have been concerned only with capturing the dual graphs of tilings that we explore, but it is possible to use addresses to include information about the directions in which tiles are oriented as well. This would allow a spatial reconstruction of any local region and would give a sense of direction to the exploration. For example, a user may wish to travel in whatever direction is closest to north.

Each tile is oriented as follows. Place all prototiles is space arbitrarily. These are the reference prototile orientations. Actual orientations of tiles are given as rotations in three dimensions from the tiles. If t is of type p, R is such a rotation, and $R(p)$ is a translation and scaling of t then we say that t has orientation R.

Each inflation patch is oriented in the same way as the reference position of its prototype and we calculate the orientation of each tile of an inflation patch. For any t in some inflation patch I_p, let R_t be the orientation of t within its reference patch. Then given any address $(a_1, a_2, \ldots, a_n)$ the addressed tile's orientation is $R_{a_1} \times R_{a_2} \times \ldots \times R_{a_n}$ where $\times$ indicates composition of the rotations.

Taking all the R_{a_i} from the inflation patches, if the closure of these rotations is finite, then this is a finite set of possible orientations and we can represent orientations discretely. Careful positioning of the reference positions will sometimes reduce the number of different possible orientations. Ideally, there is a finite number of orientations and we can compute orientations relative to adjacent tiles as we did for Robinson's tiles (see Figure 2.14). However, it may be that tiles are oriented relative to their inflations at an irrational angle and that a continuous range of orientations must be used.

We can incorporate more spatial information into our data structure by generalizing our transformations to include translation and scale as well. The scaling must always be the uniform expansion factor of the tiles (this was τ for Robinson's triangles). Then it is possible to pinpoint the precise location and shape of any tile by composing the transformations associated with its digits and applying this transformation to the reference prototile for the first digit of the address. In this context, tilings can be seen as the finite application of an iterated function system. This addressing system could be used to form ray-traced images of tilings without building any models larger than the reference prototiles. Hart and DeFanti[HD91] describe such a system for rendering fractals generated by iterated function systems. Computing the intensity of a pixel requires time proportional to the number of levels of inflation of the tile being drawn. Rather than constructing and transforming each tile in the model, the inverse transformations are applied to the rays themselves as they intersect a tile at each level of the inheritance hierarchy. This method would be well-suited for generating images of aperiodic tilings because in this case rays only must to traverse a structure of finite depth, normals are well-defined, and no tiles intersect.

Chapter 7

Conclusion

The bulk of our work has concerned the adaptation of index sequences for use as an addressing system. The inspiration for this approach comes from a discussion of index sequences for Robinson's triangles by Grunbaum and Shephard[GS87].[1] In developing our data structure, we faced four major problems:

1. **properties:** In Chapter 4 we developed careful spatial definitions for our tiling terminology and laid a foundation for the correspondence between tilings as dual graphs and as sets of points. Inflation, aperiodicity, and local isomorphism are well-studied topics in the field of tilings. We investigated these properties using the same model of inflation as is used by our data structure. This allowed us to determine the properties of a protoset that can be used to prove that it tiles the plane aperiodically.

2. **specification:** We developed a notation for representing the dual graph of a tiling and inflation rules for growing tilings. These notations allowed us to precisely define the construction of inflated tilings and the address of a tile within an inflated patch. These principles can be translated almost directly to a mechanical implementation of our data structure.

3. **adjacency:** The adjacency algorithm enabled us to use the space of addresses as if were a graph structure. Each move from tile to tile in this graph requires time proportional to the length of the addresses and addresses have logarithmic length in the number of tiles of the represented patch. An adjacency is computed by examining relationships between consecutive pairs of digits that represent pairs of tiles at consecutive levels of inflation.

[1] Grunbaum and Shephard report that Lunnon studied the use of index sequences for generating images of tilings, but Lunnon never published these results.

4. **symmetry:** Unless seams are taken into account, our data structure can only explore one region of a symmetrical tiling and the adjacency algorithm fails to terminate if exploration crosses a seam. We were able to recognize and handle this case for the Robinson triangles. Our approach for the general case was more difficult and required the study of roadmaps and configurations. Some aspects of this theory would benefit from more investigation and formalization.

We believe that our data structure is an efficient mechanism for exploring very large spaces of aperiodic tilings admitted by inflation. Except for addresses themselves, no storage of the structure of a tiling needs to be maintained beyond the fixed tables of rules that are generated in advance of the exploration of any particular class of tiling. Our system uses addresses with lengths that grow in the same order of magnitude as the minimum number of bits required to identify tiles in a patch of the size that they represent. Moving to adjacent tiles requires time proportional to the number of bits in an address. In neither space nor time is the constant in front of these orders of magnitude particularly large. Three features of our data structure make it a particularly desirable mechanism for touring aperiodic tilings:

1. **general:** Any tilings generated by unique inflation by composition of tiles can be represented using our system. Furthermore, any region within any tiling of a given class can be explored.[2]

2. **discrete:** Our representation is not an approximation, but a construction that precisely models regions within a tiling. Methods dependent on projection are susceptible to rounding errors that can result in flaws in the tiling.[3]

3. **expandable:** During the course of exploration, addresses grow longer as needed in order to encompass travel further from the center of exploration. As far as the user is concerned, the system represents an infinite tiling.

7.1 Future Work

We briefly mention here what we believe are interesting questions that future research might pursue.

[2]Note that symmetric tilings require extra effort to represent properly.

[3]Bell et al[BDHJ83] have studied the use of various regular tilings in a hierarchical addressing system for processing image data. They find it desirable to avoid floating point calculations. QuasiTiler[Dur94] sometimes produces tilings with holes because a point very near a border is incorrectly placed outside it.

Theoretical issues:

- What are the relationships among inflation, unique inflation, aperiodicity, and local isomorphism?
- What is the minimum information required to represent a section of a tiling?

Efficiency analysis:

- How many digits must be used to structurally contain any ball of tiles of a certain radius?
- How many digits are required to represent a random connected patch of tiles?
- On average, how many digits change per move?

Data structure maintenance:

- Formalize algorithms to generate vertex configurations and symmetry types for an arbitrary inflation specification.
- Implement a system to create a data structure for exploring tilings generated by any given inflation specification.
- Attempt to optimize other operations for traveling within a tiling. For example, find the shortest path between two points (in number of moves). The shortest path does not necessarily correspond to the geometric direction of the destination.
- Generate data structures for a wide variety of protosets and assess the suitability of the data structure for representing structures of interest to physicists.
- **EXPLORE!**

Appendix A

An index sequence formed in inflations of Robinson's triangles

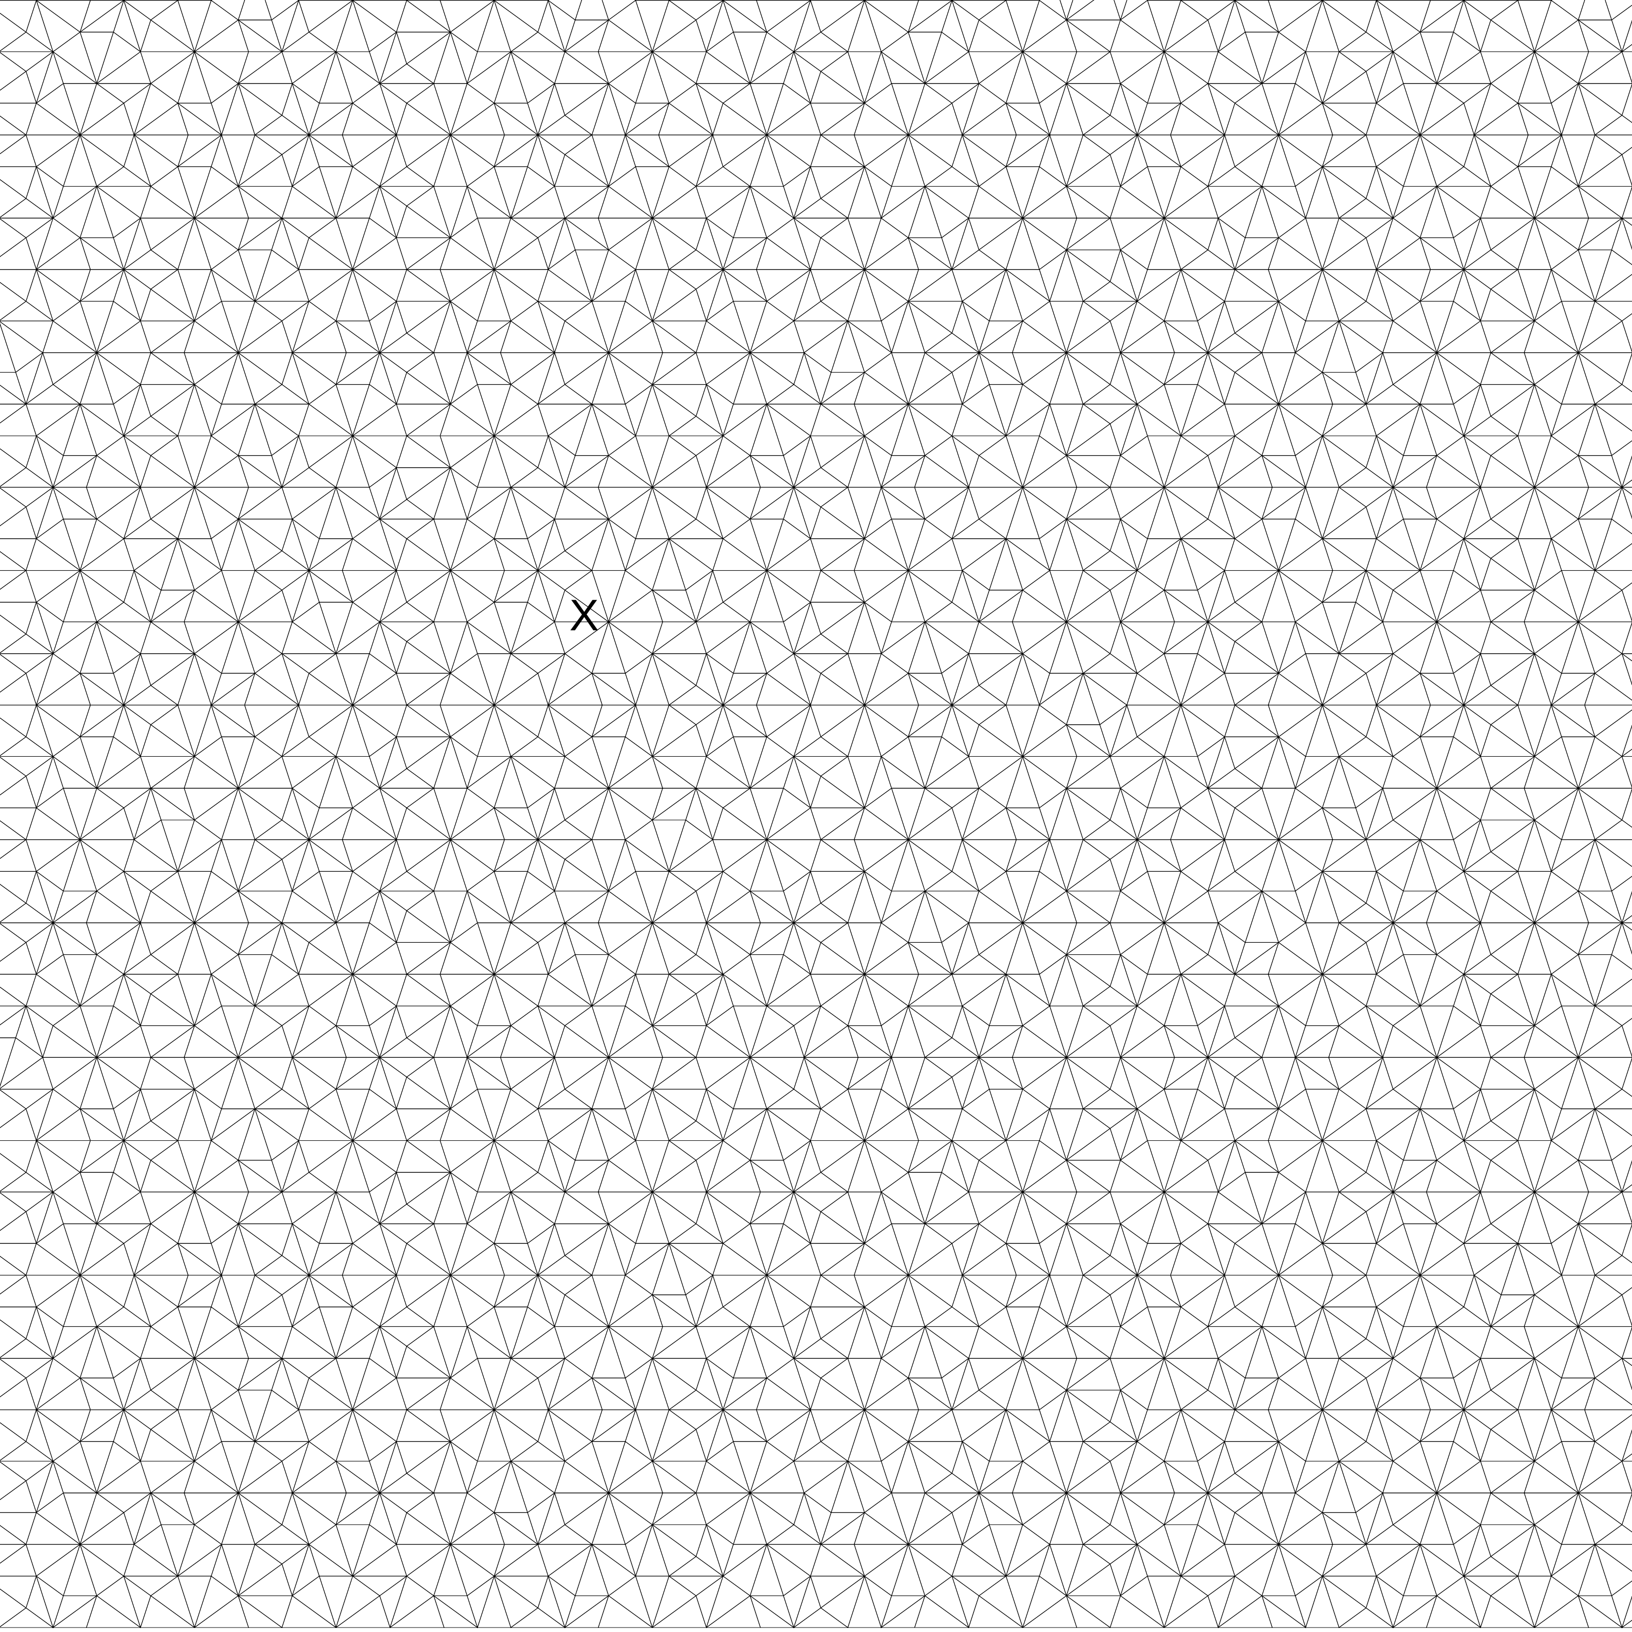

Figure A.1: Formation of an index sequence. The 1st digit is “0”.

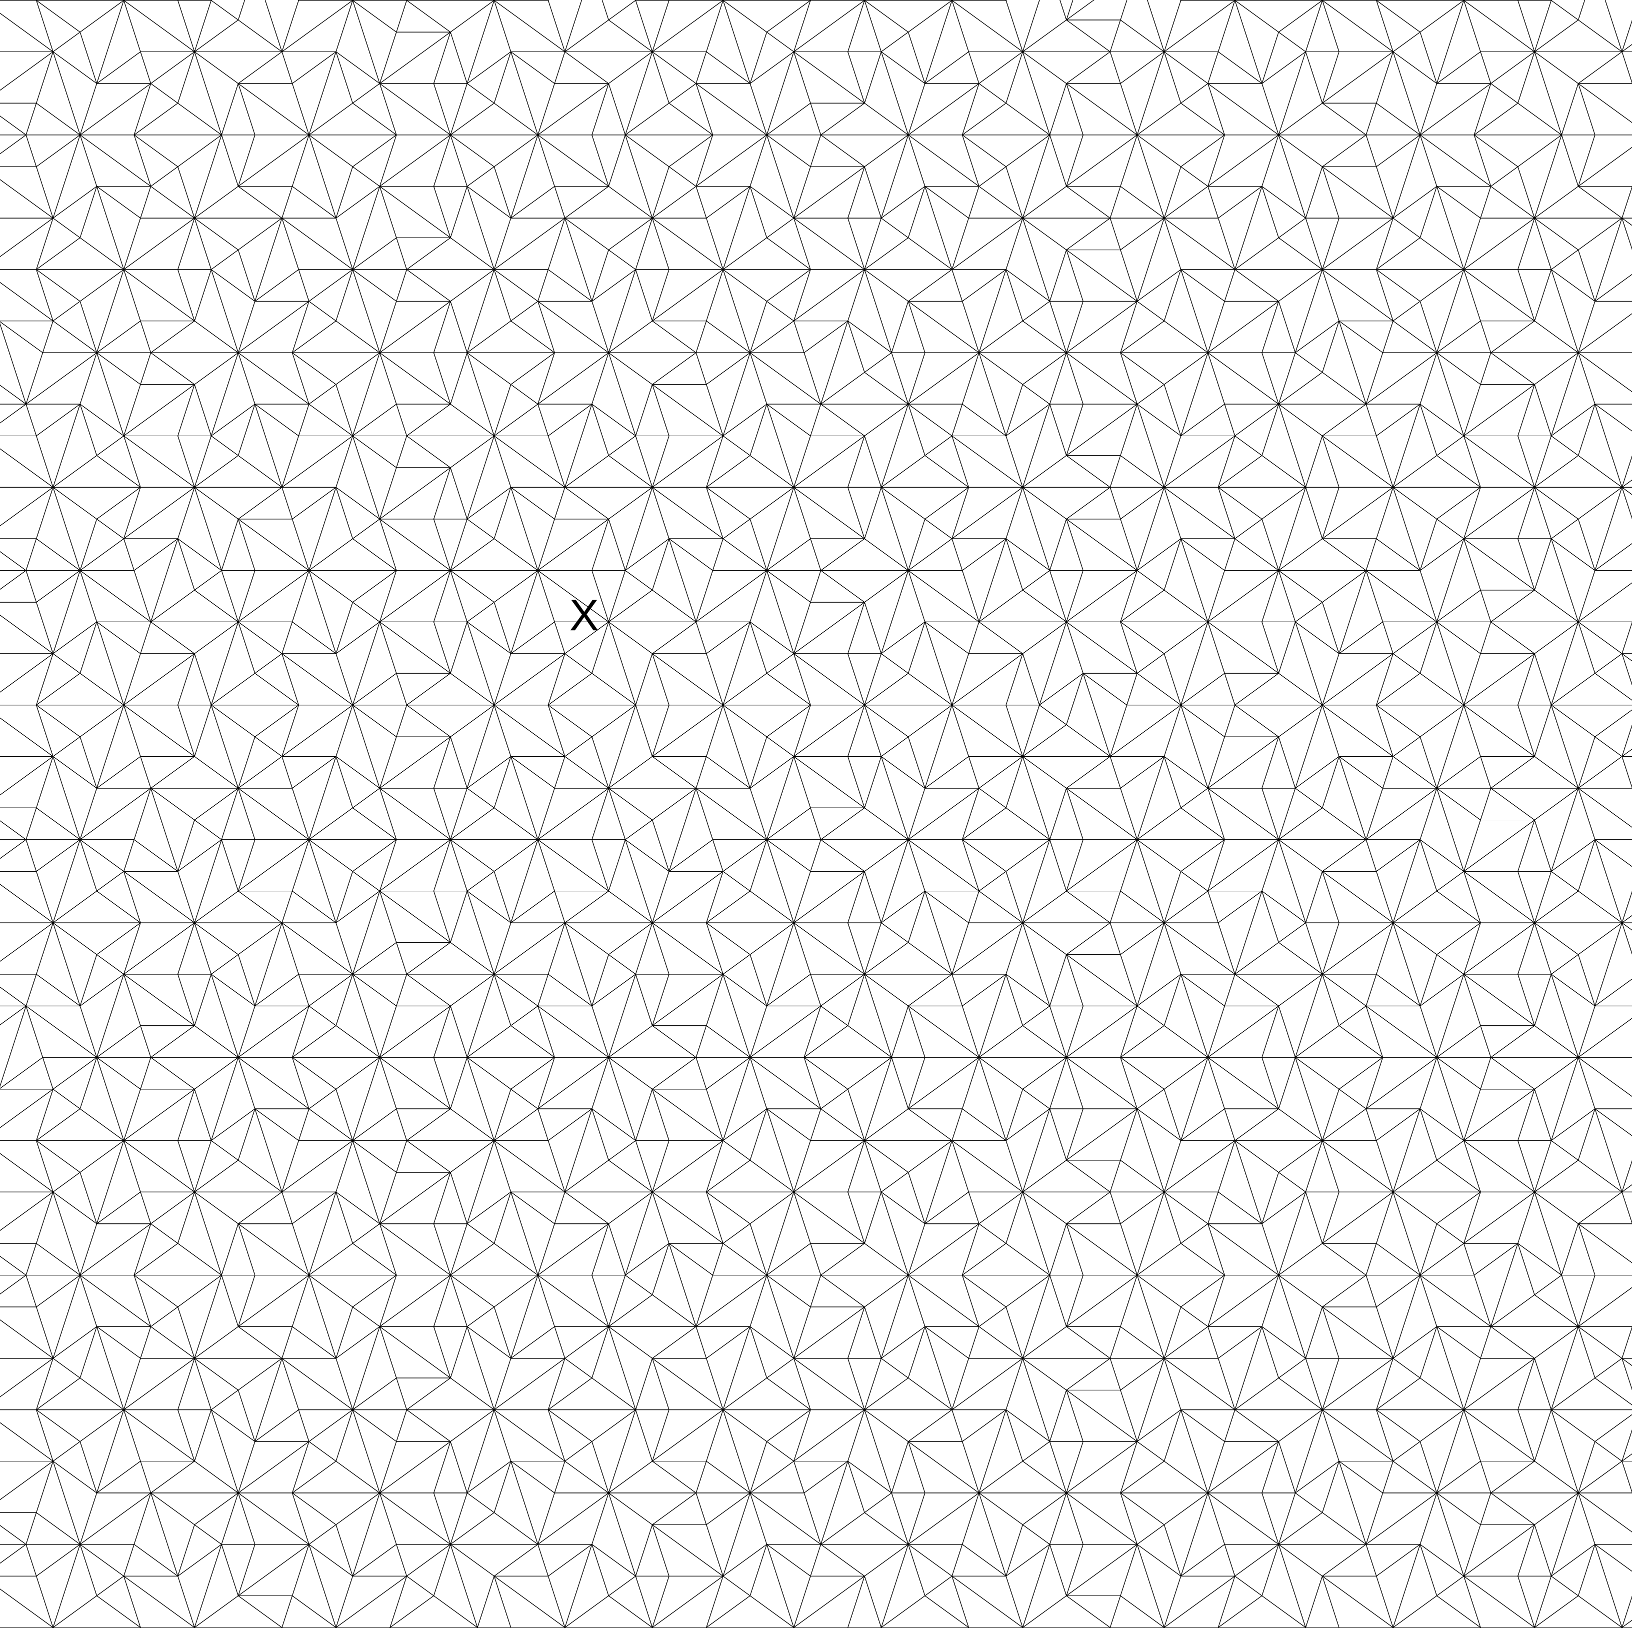

Figure A.2: Formation of an index sequence. The 2nd digit is “0”.

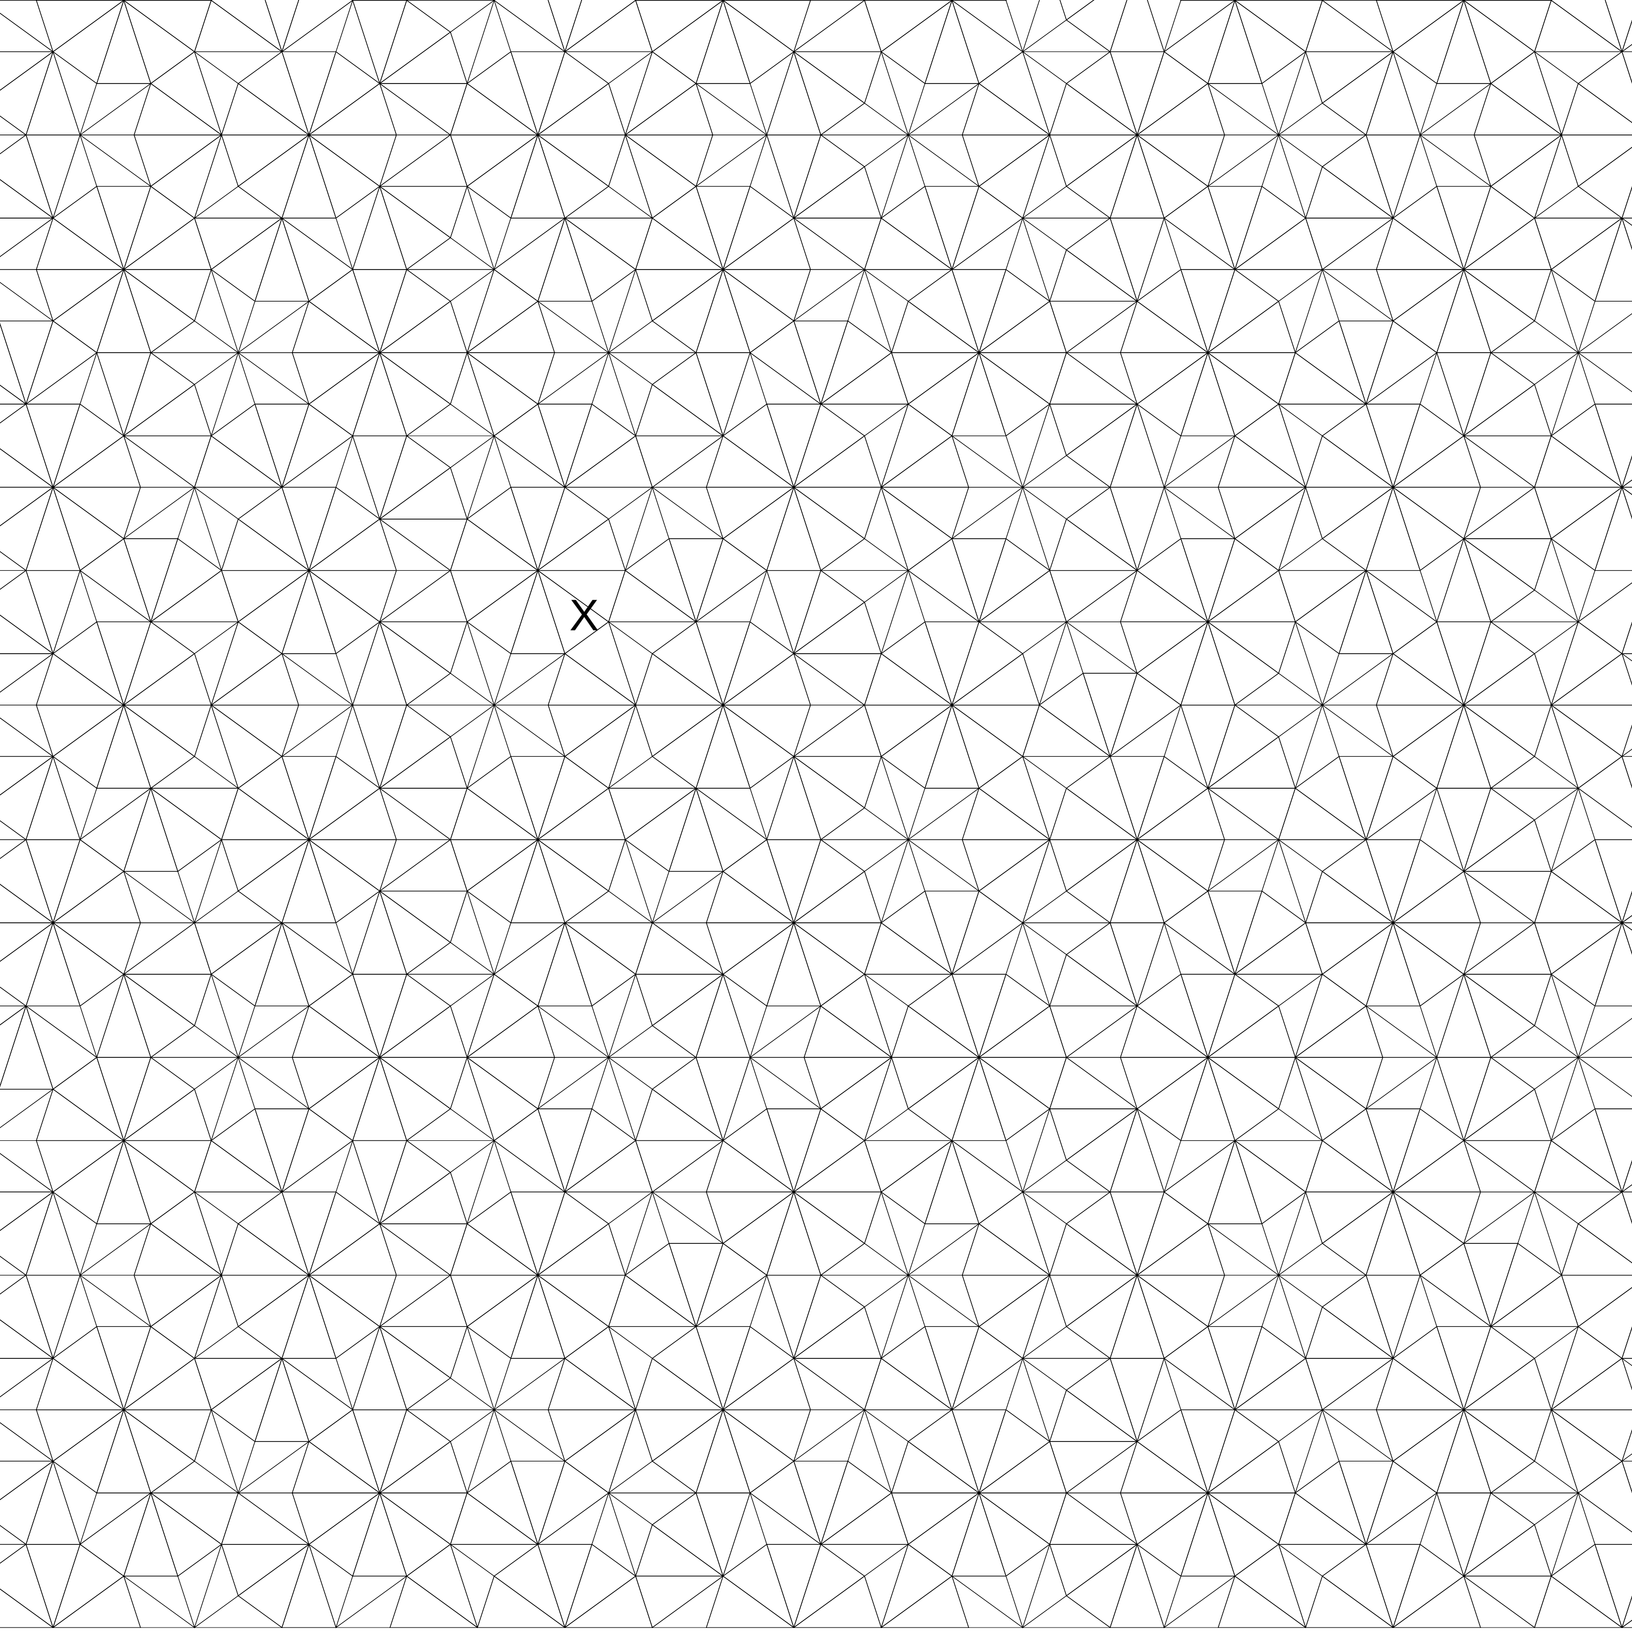

Figure A.3: Formation of an index sequence. The 3rd digit is “0”.

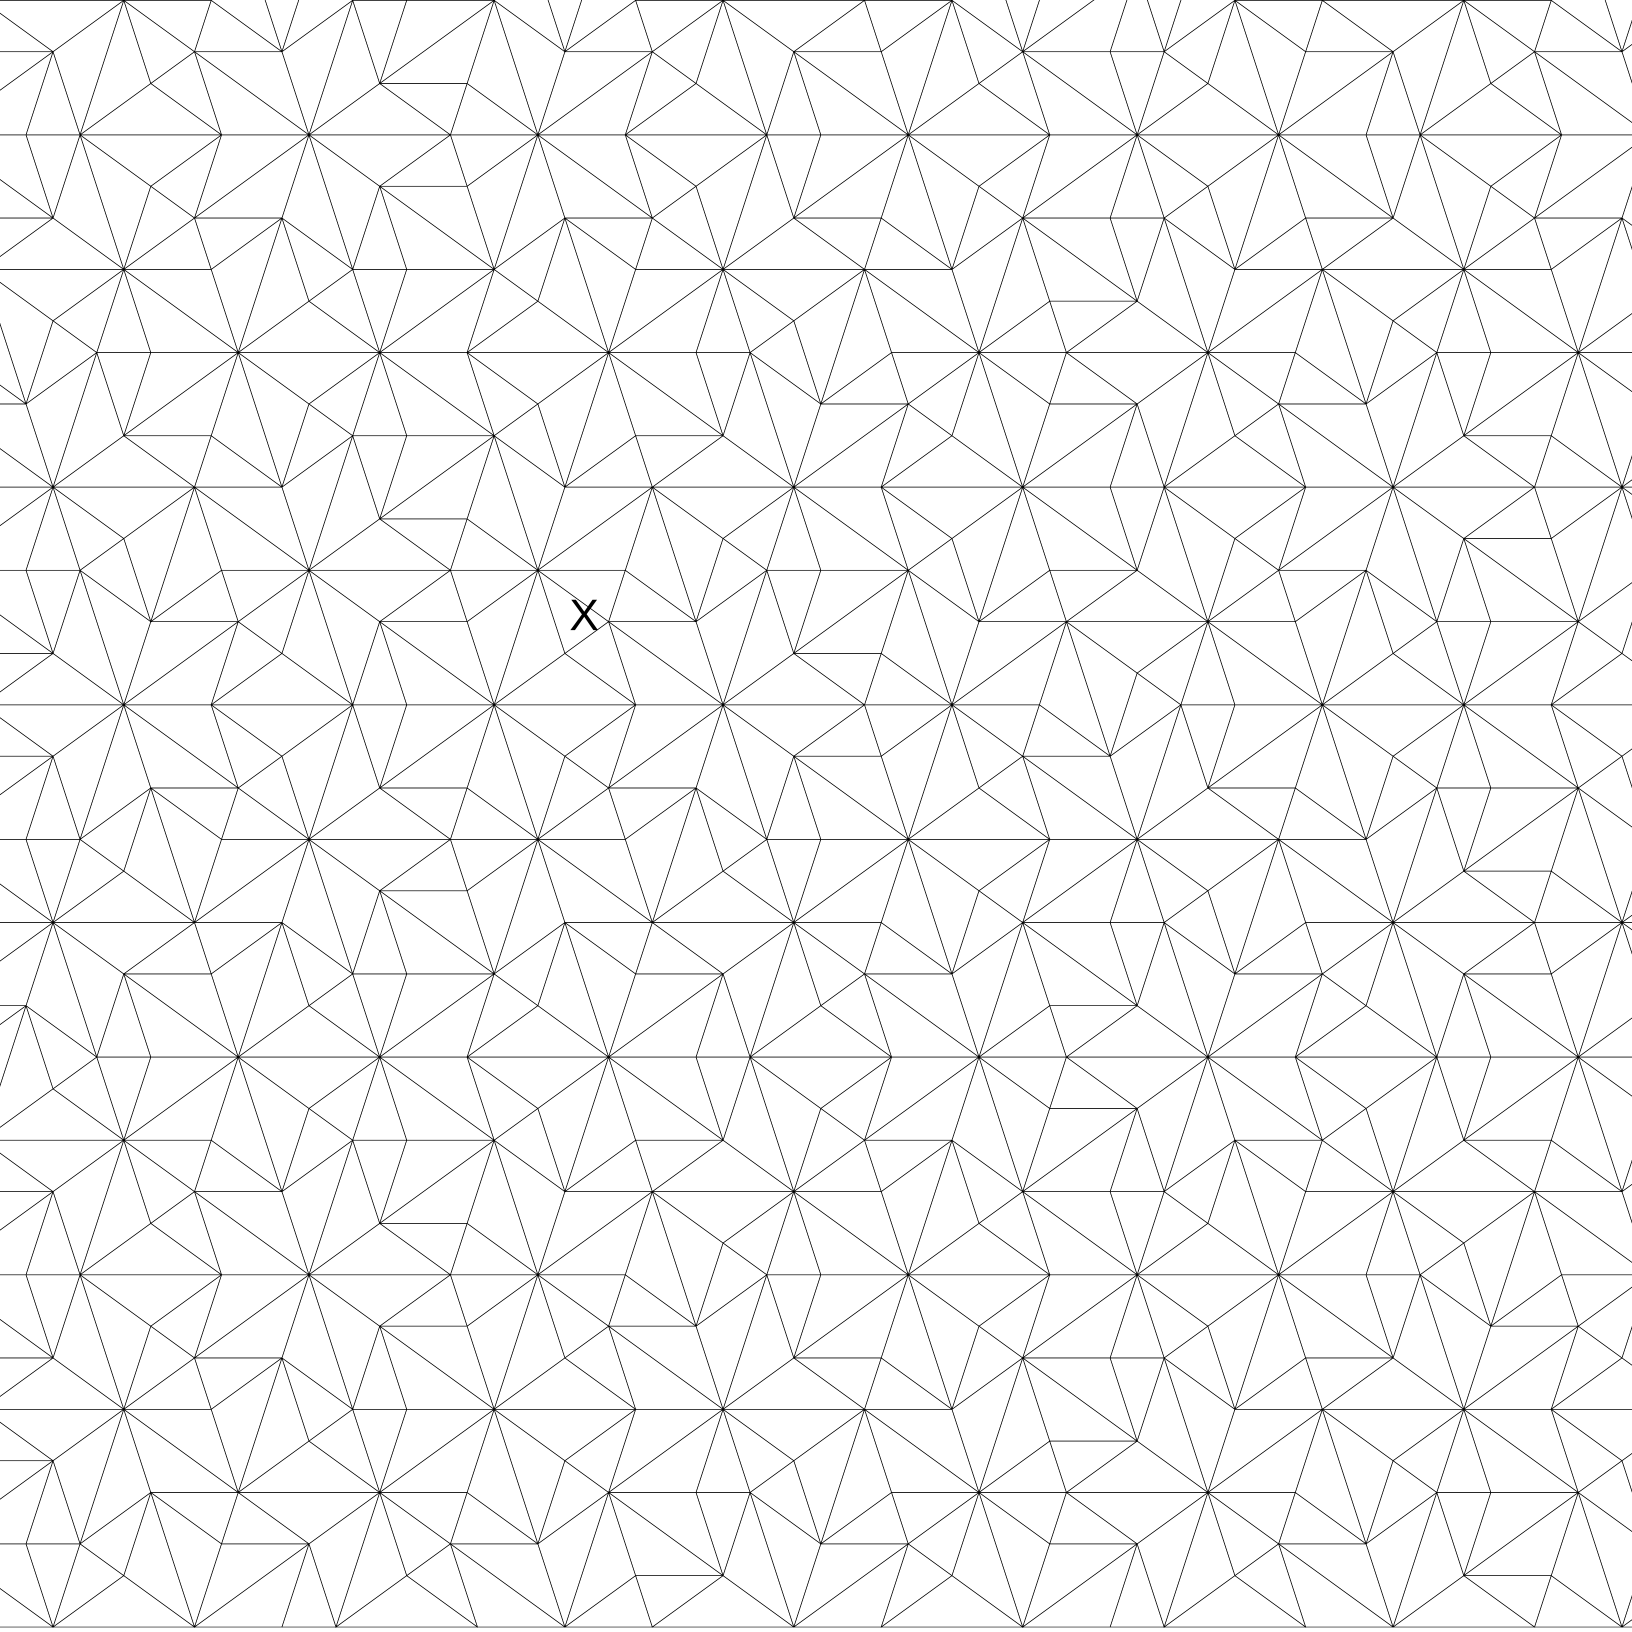

Figure A.4: Formation of an index sequence. The 4th digit is “1”.

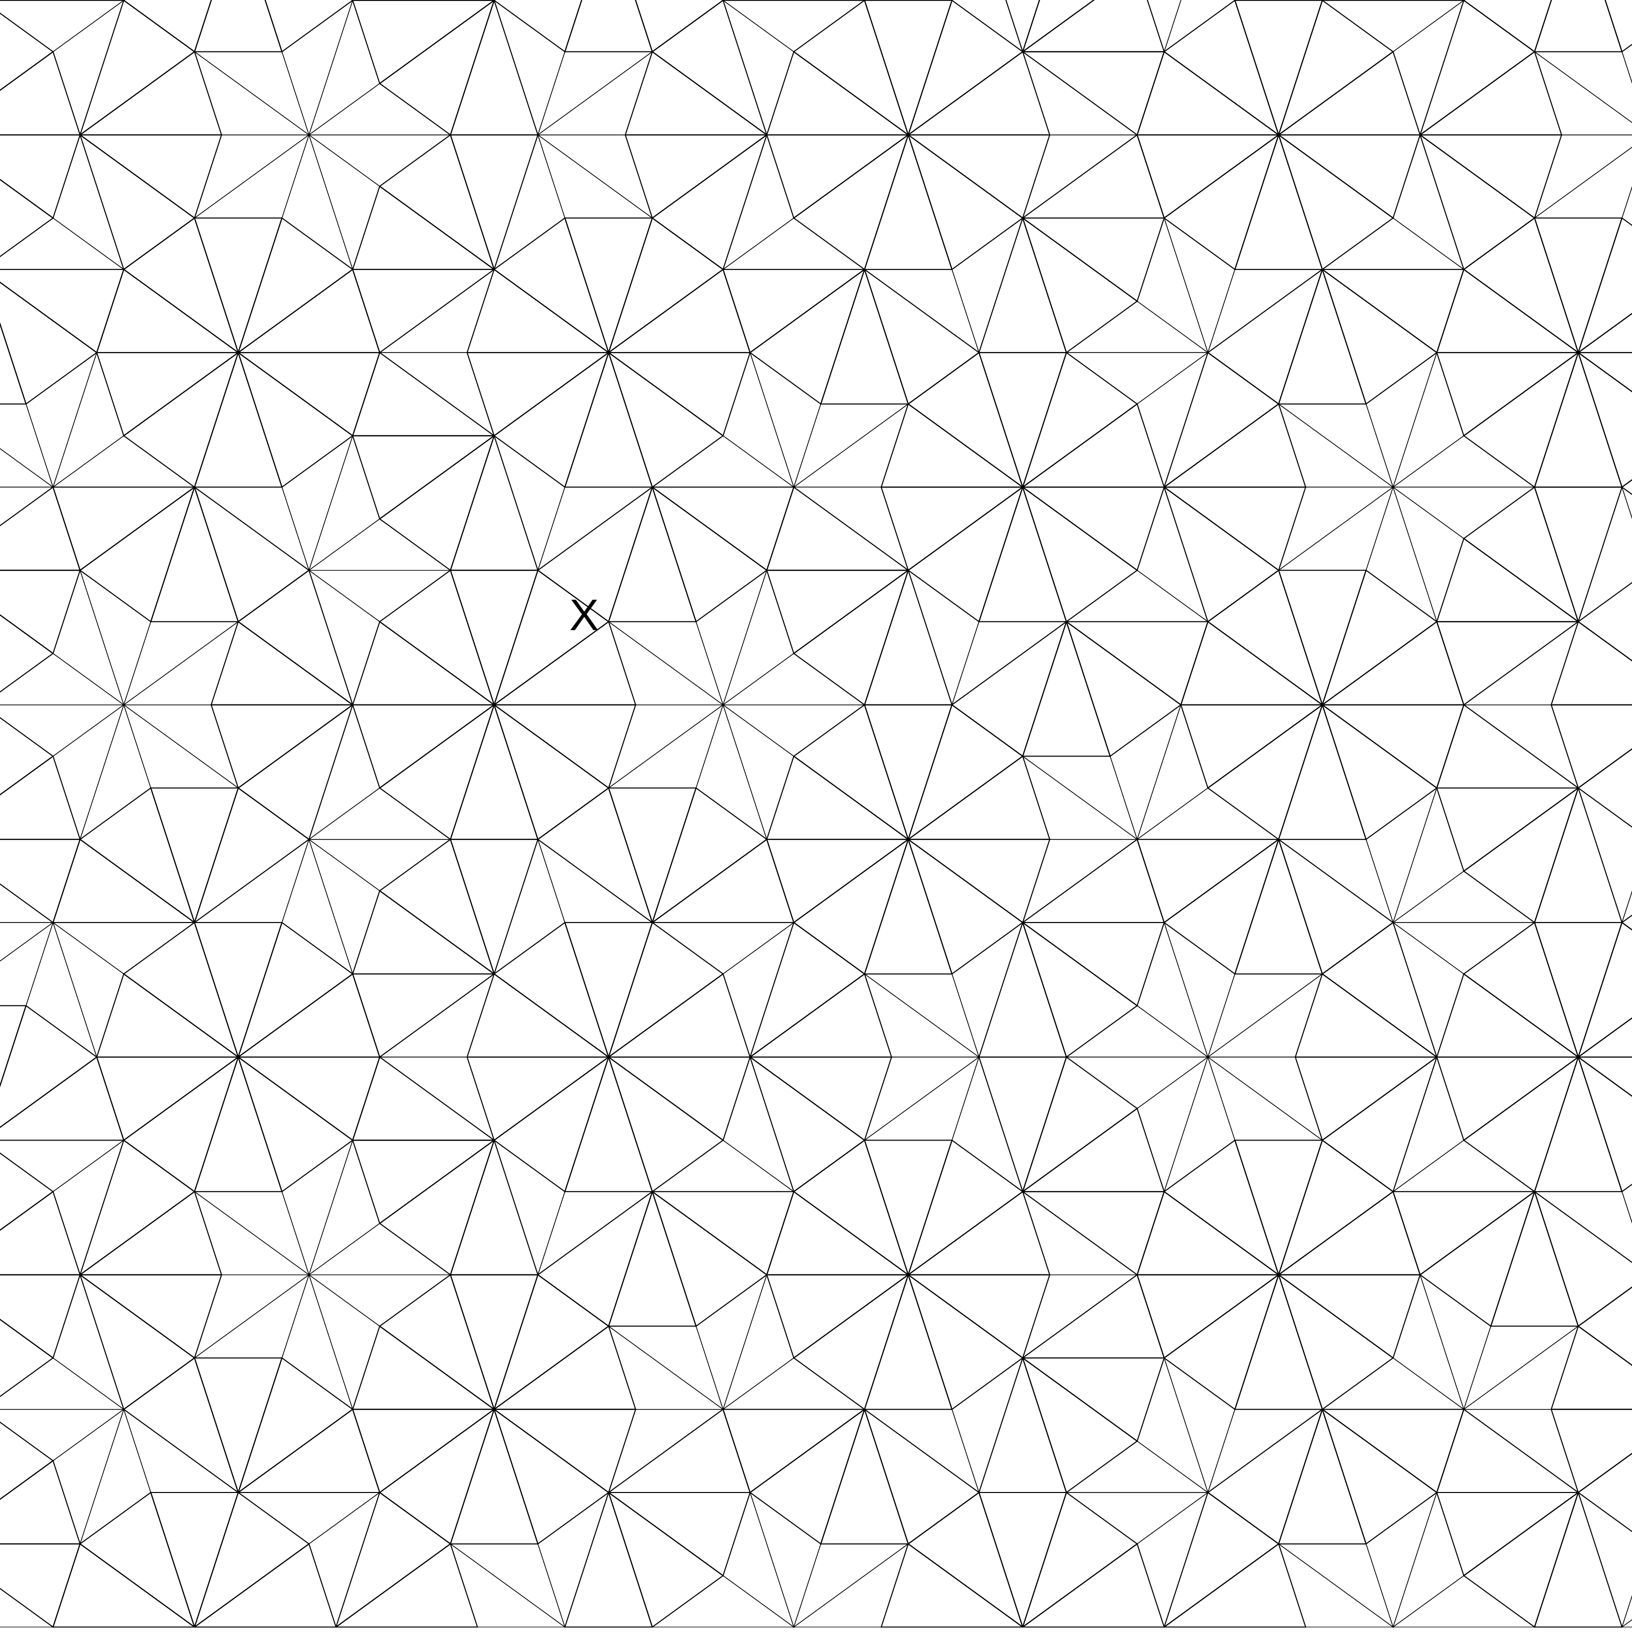

Figure A.5: Formation of an index sequence. The 5th digit is “0”.

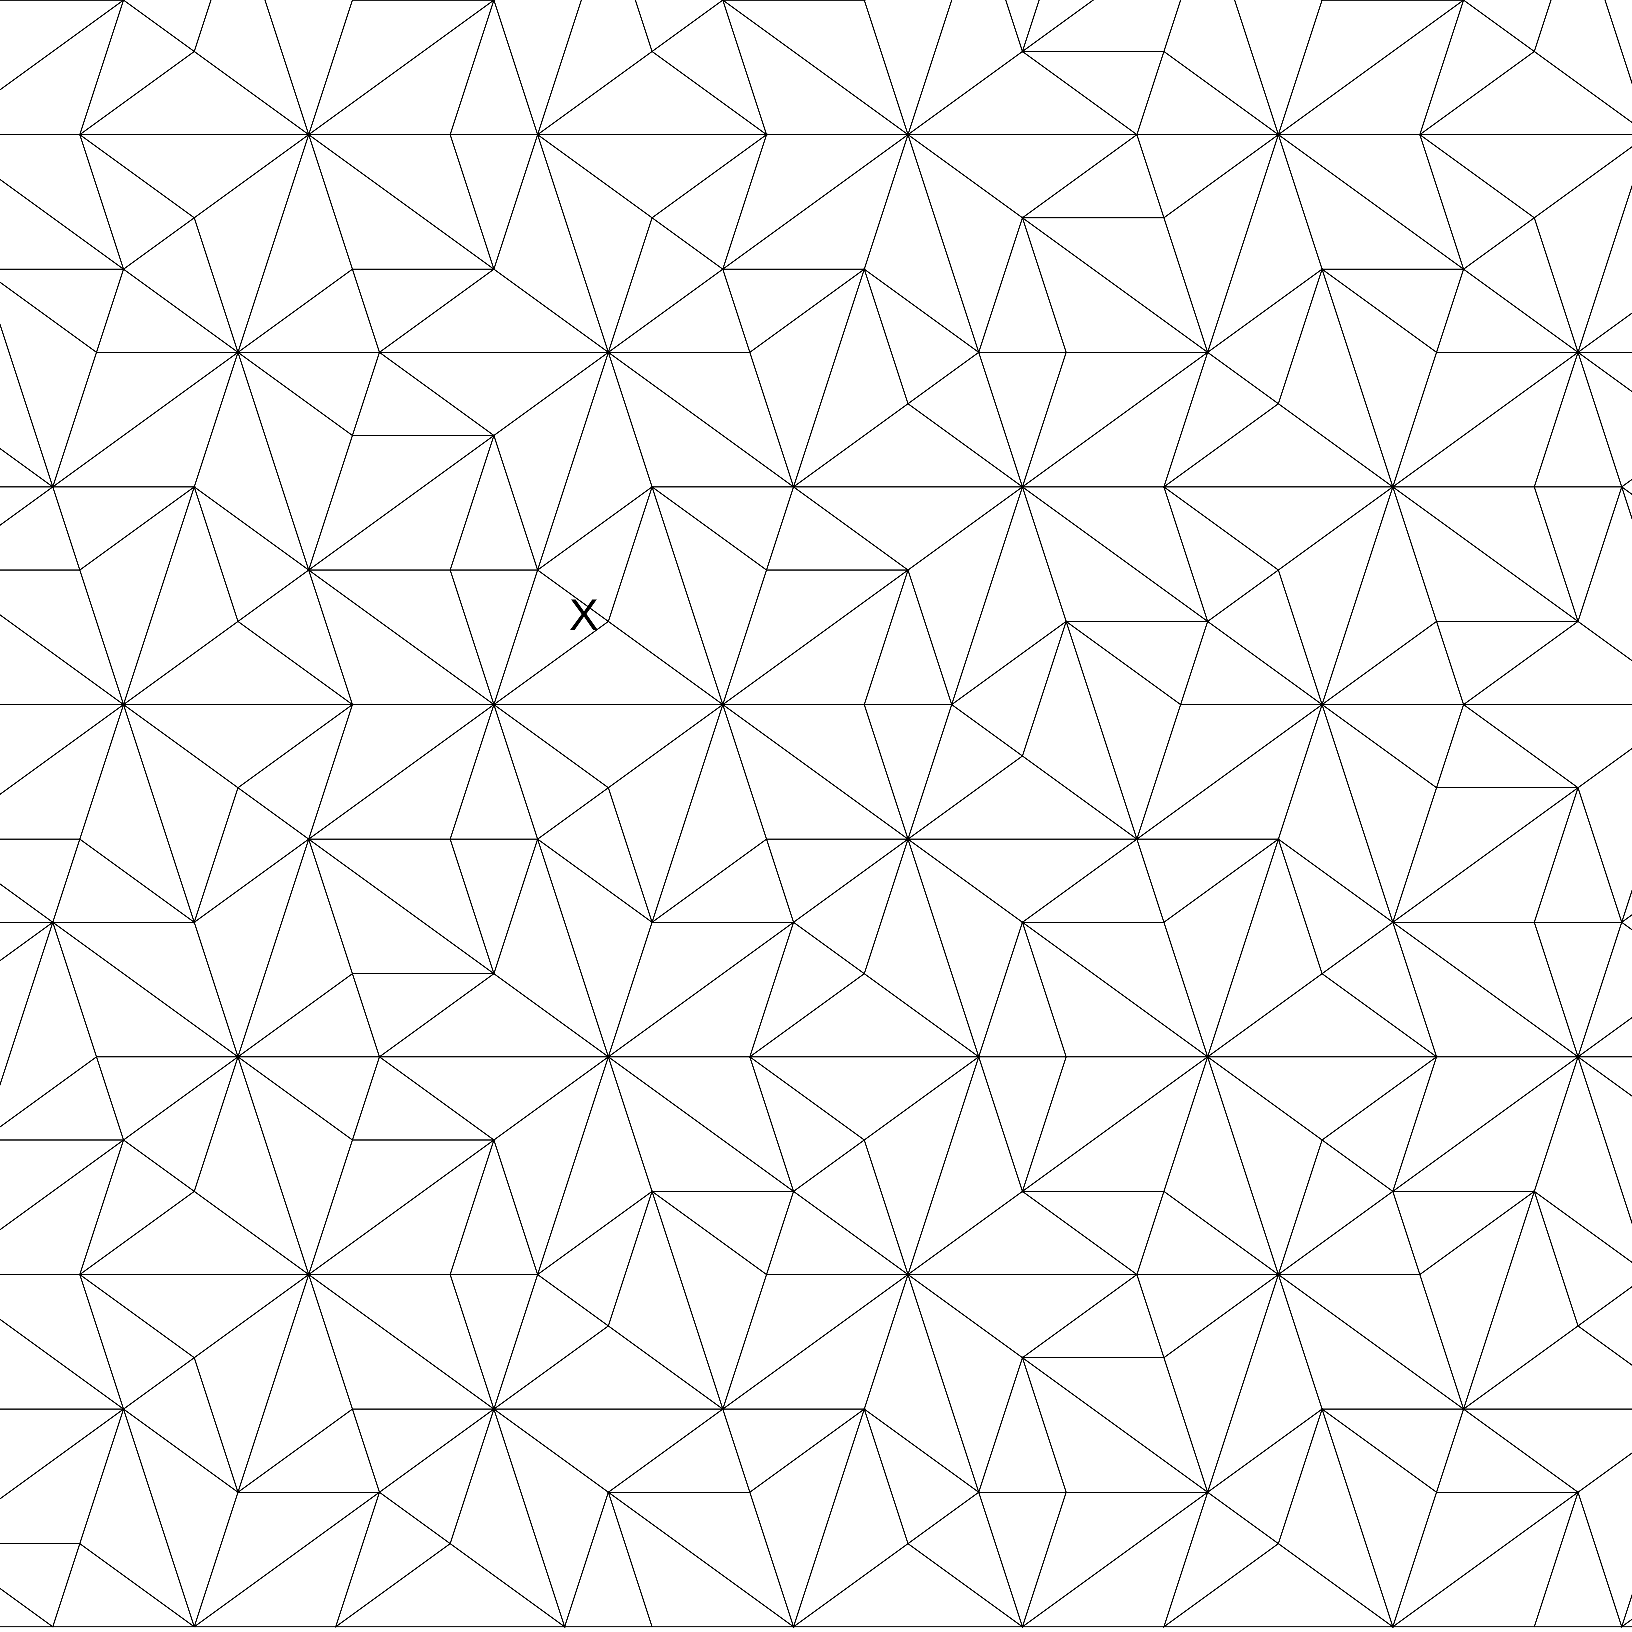

Figure A.6: Formation of an index sequence. The 6th digit is “1”.

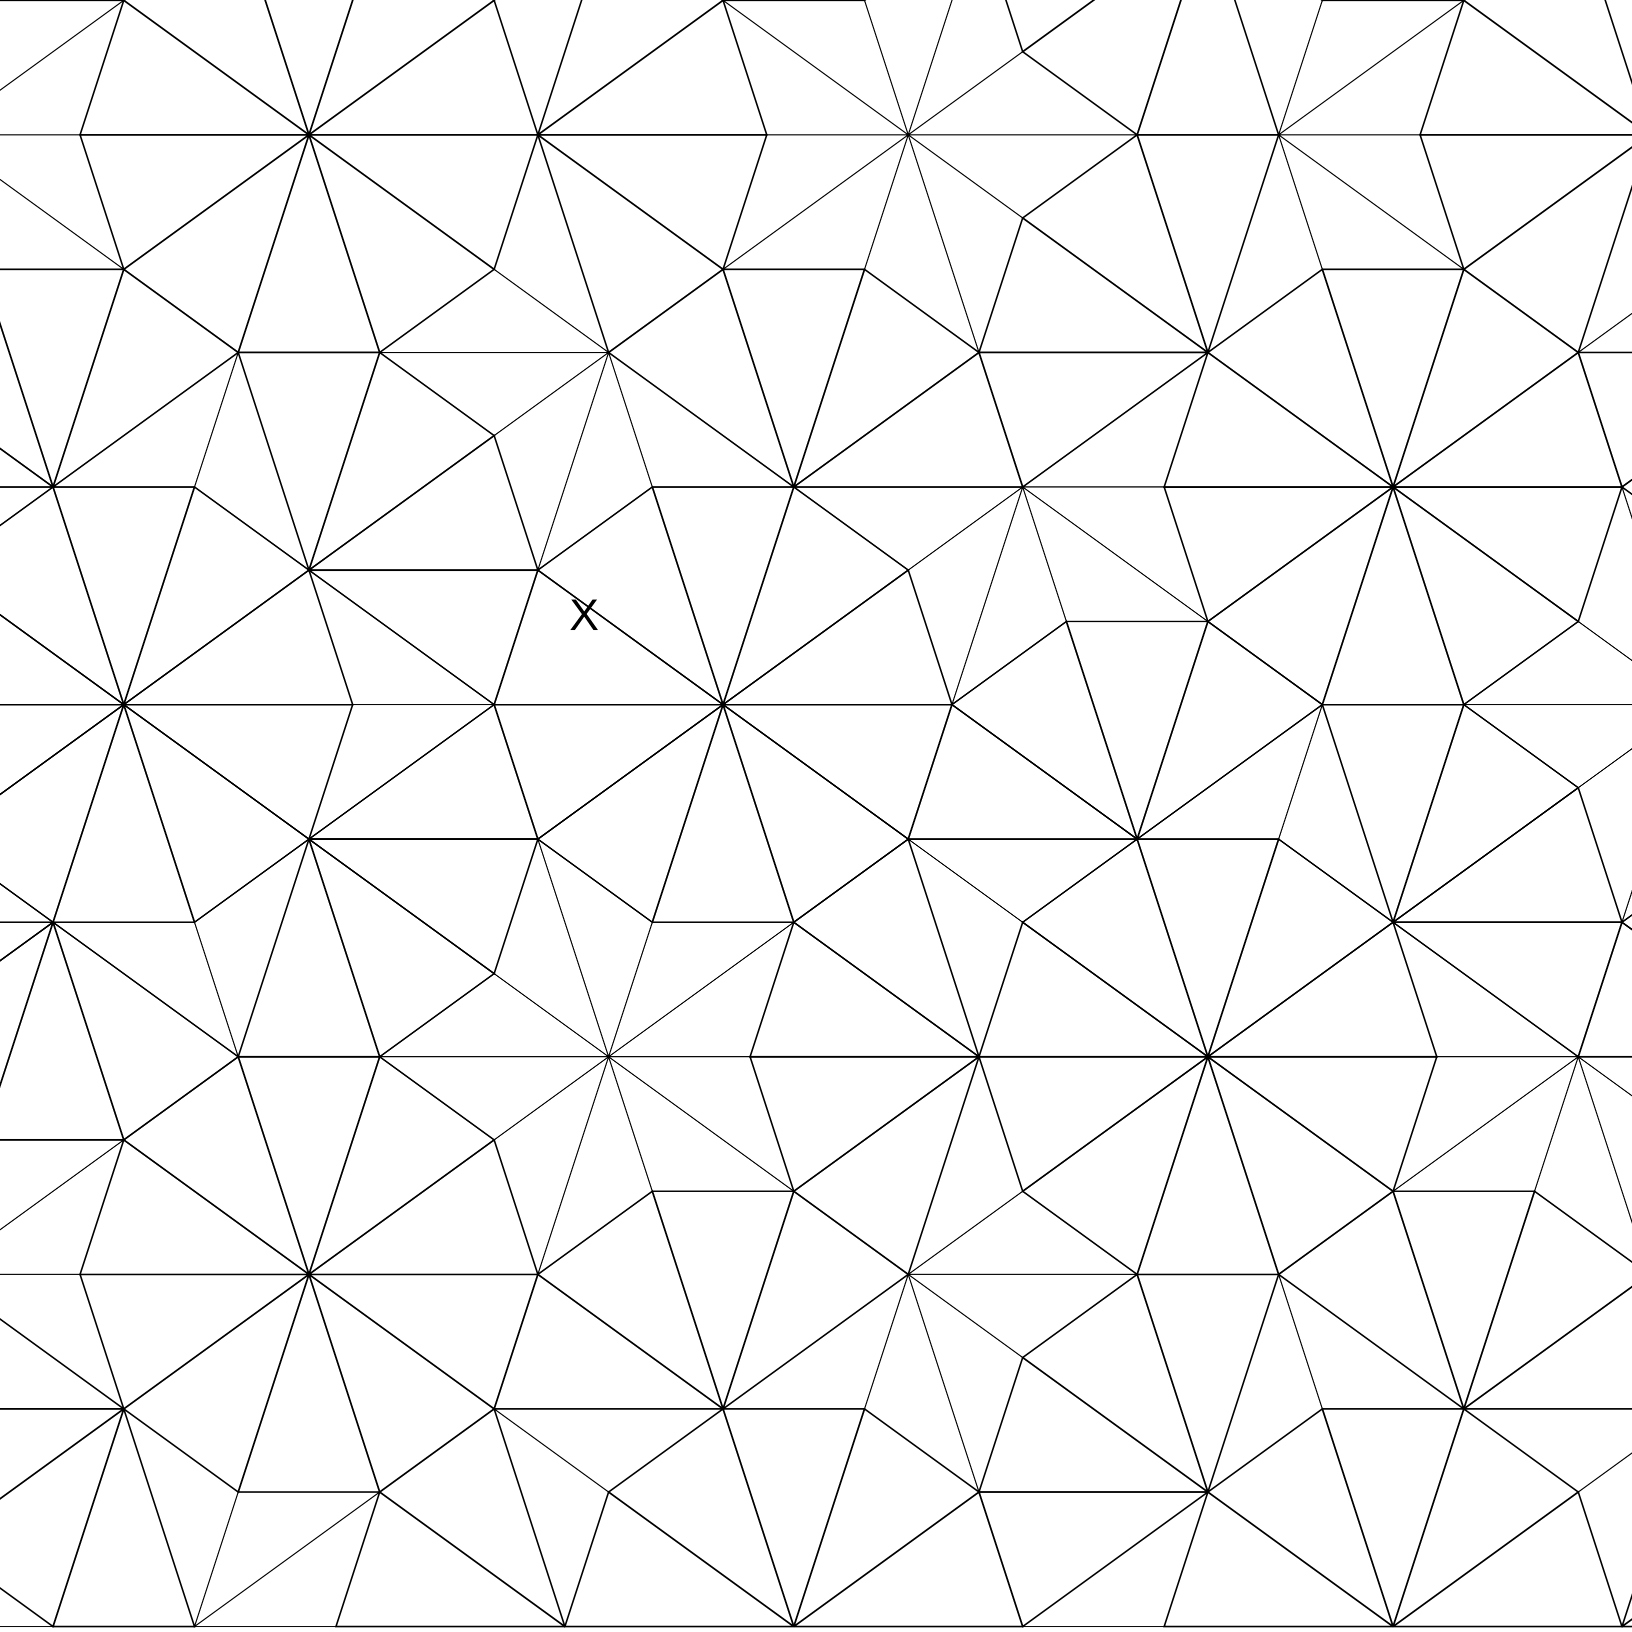

Figure A.7: Formation of an index sequence. The 7th digit is “0”.

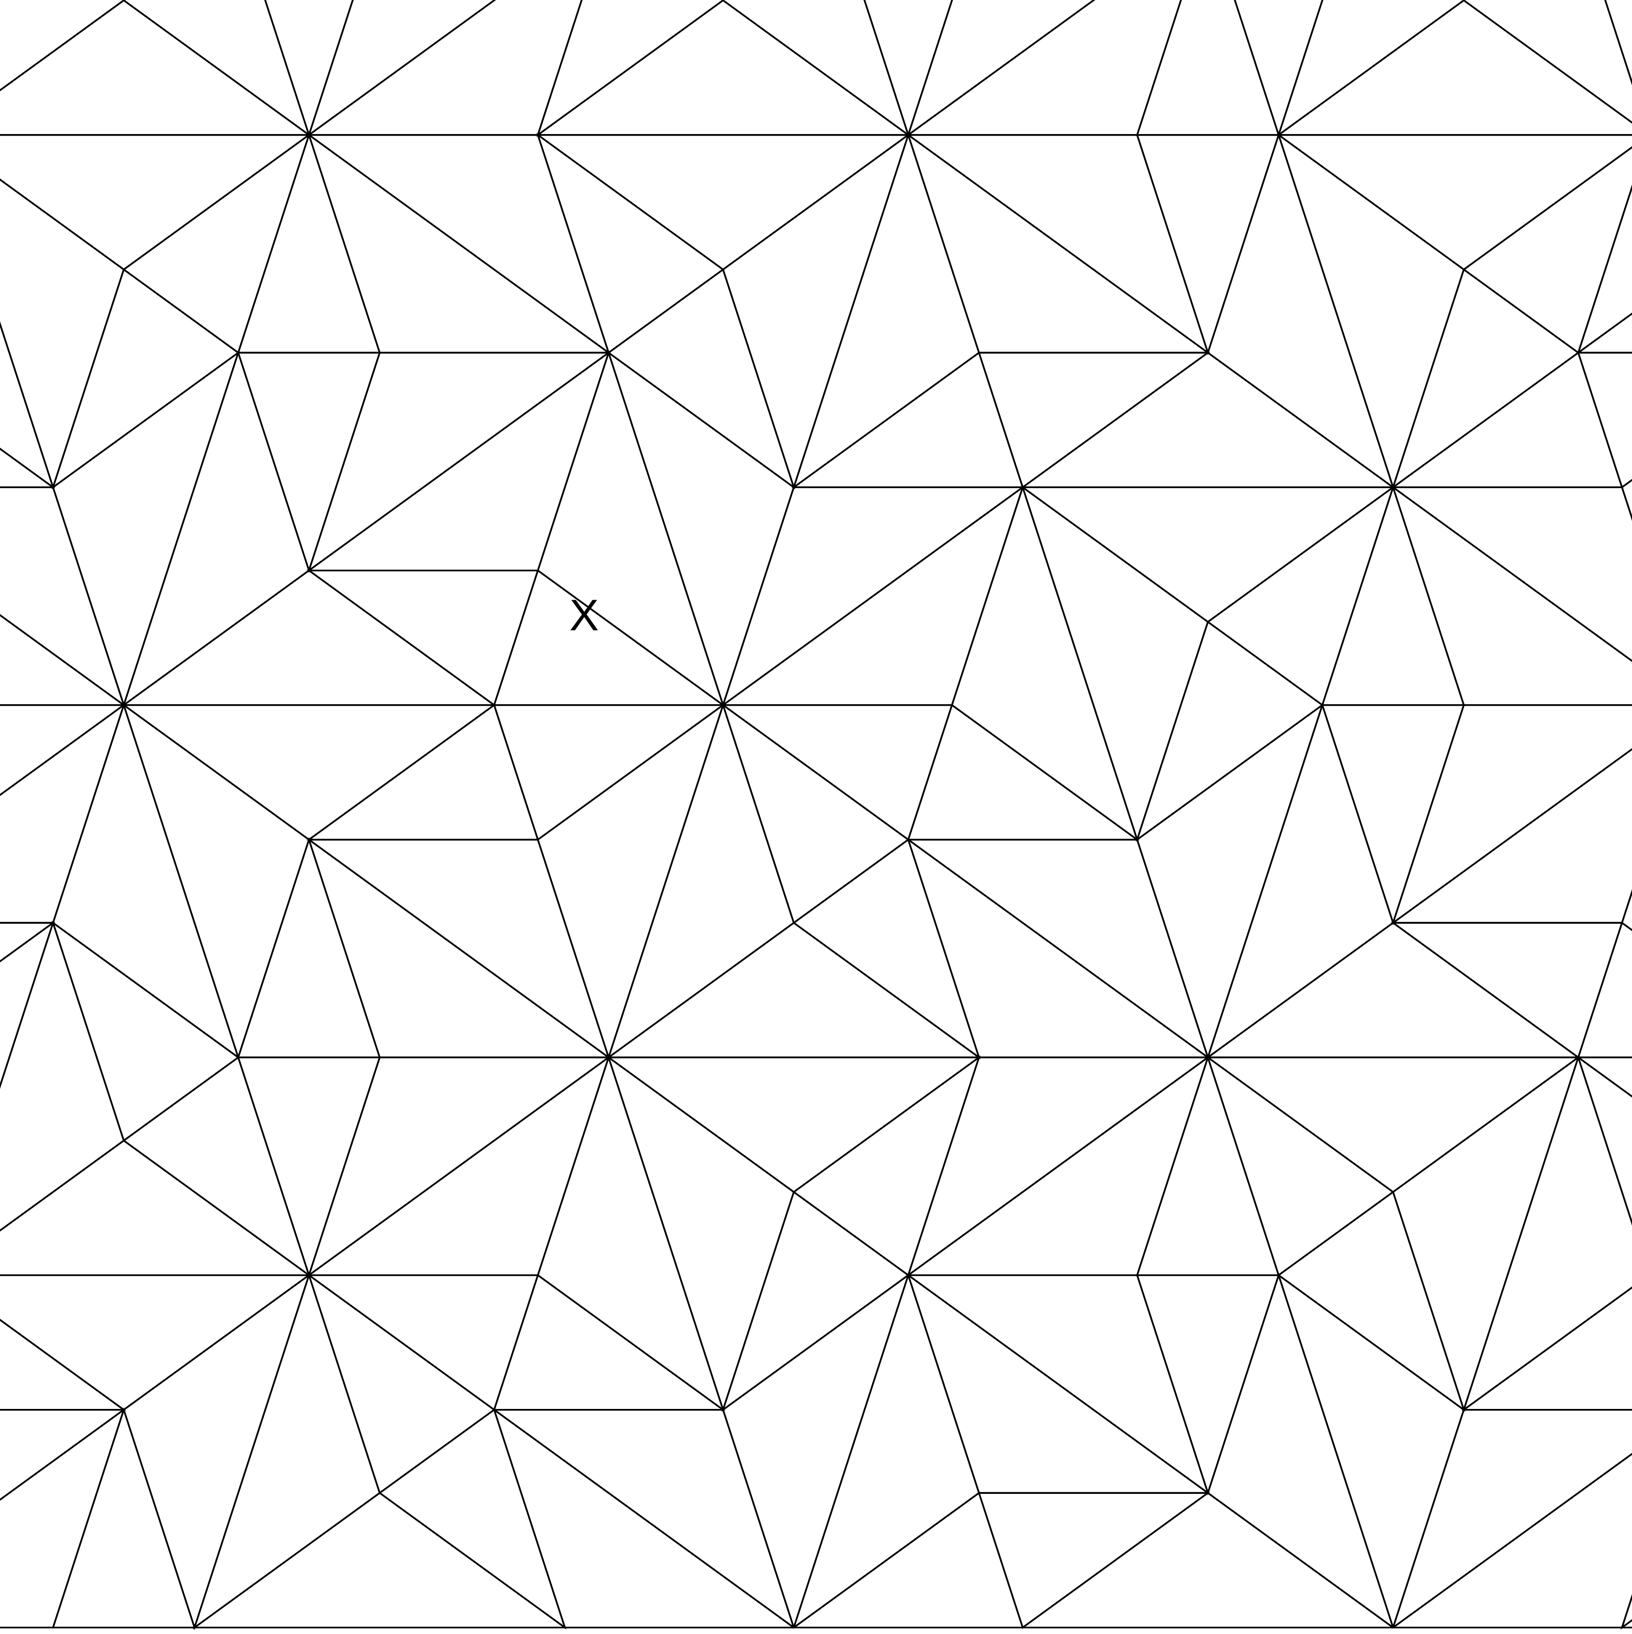

Figure A.8: Formation of an index sequence. The 8th digit is “1”.

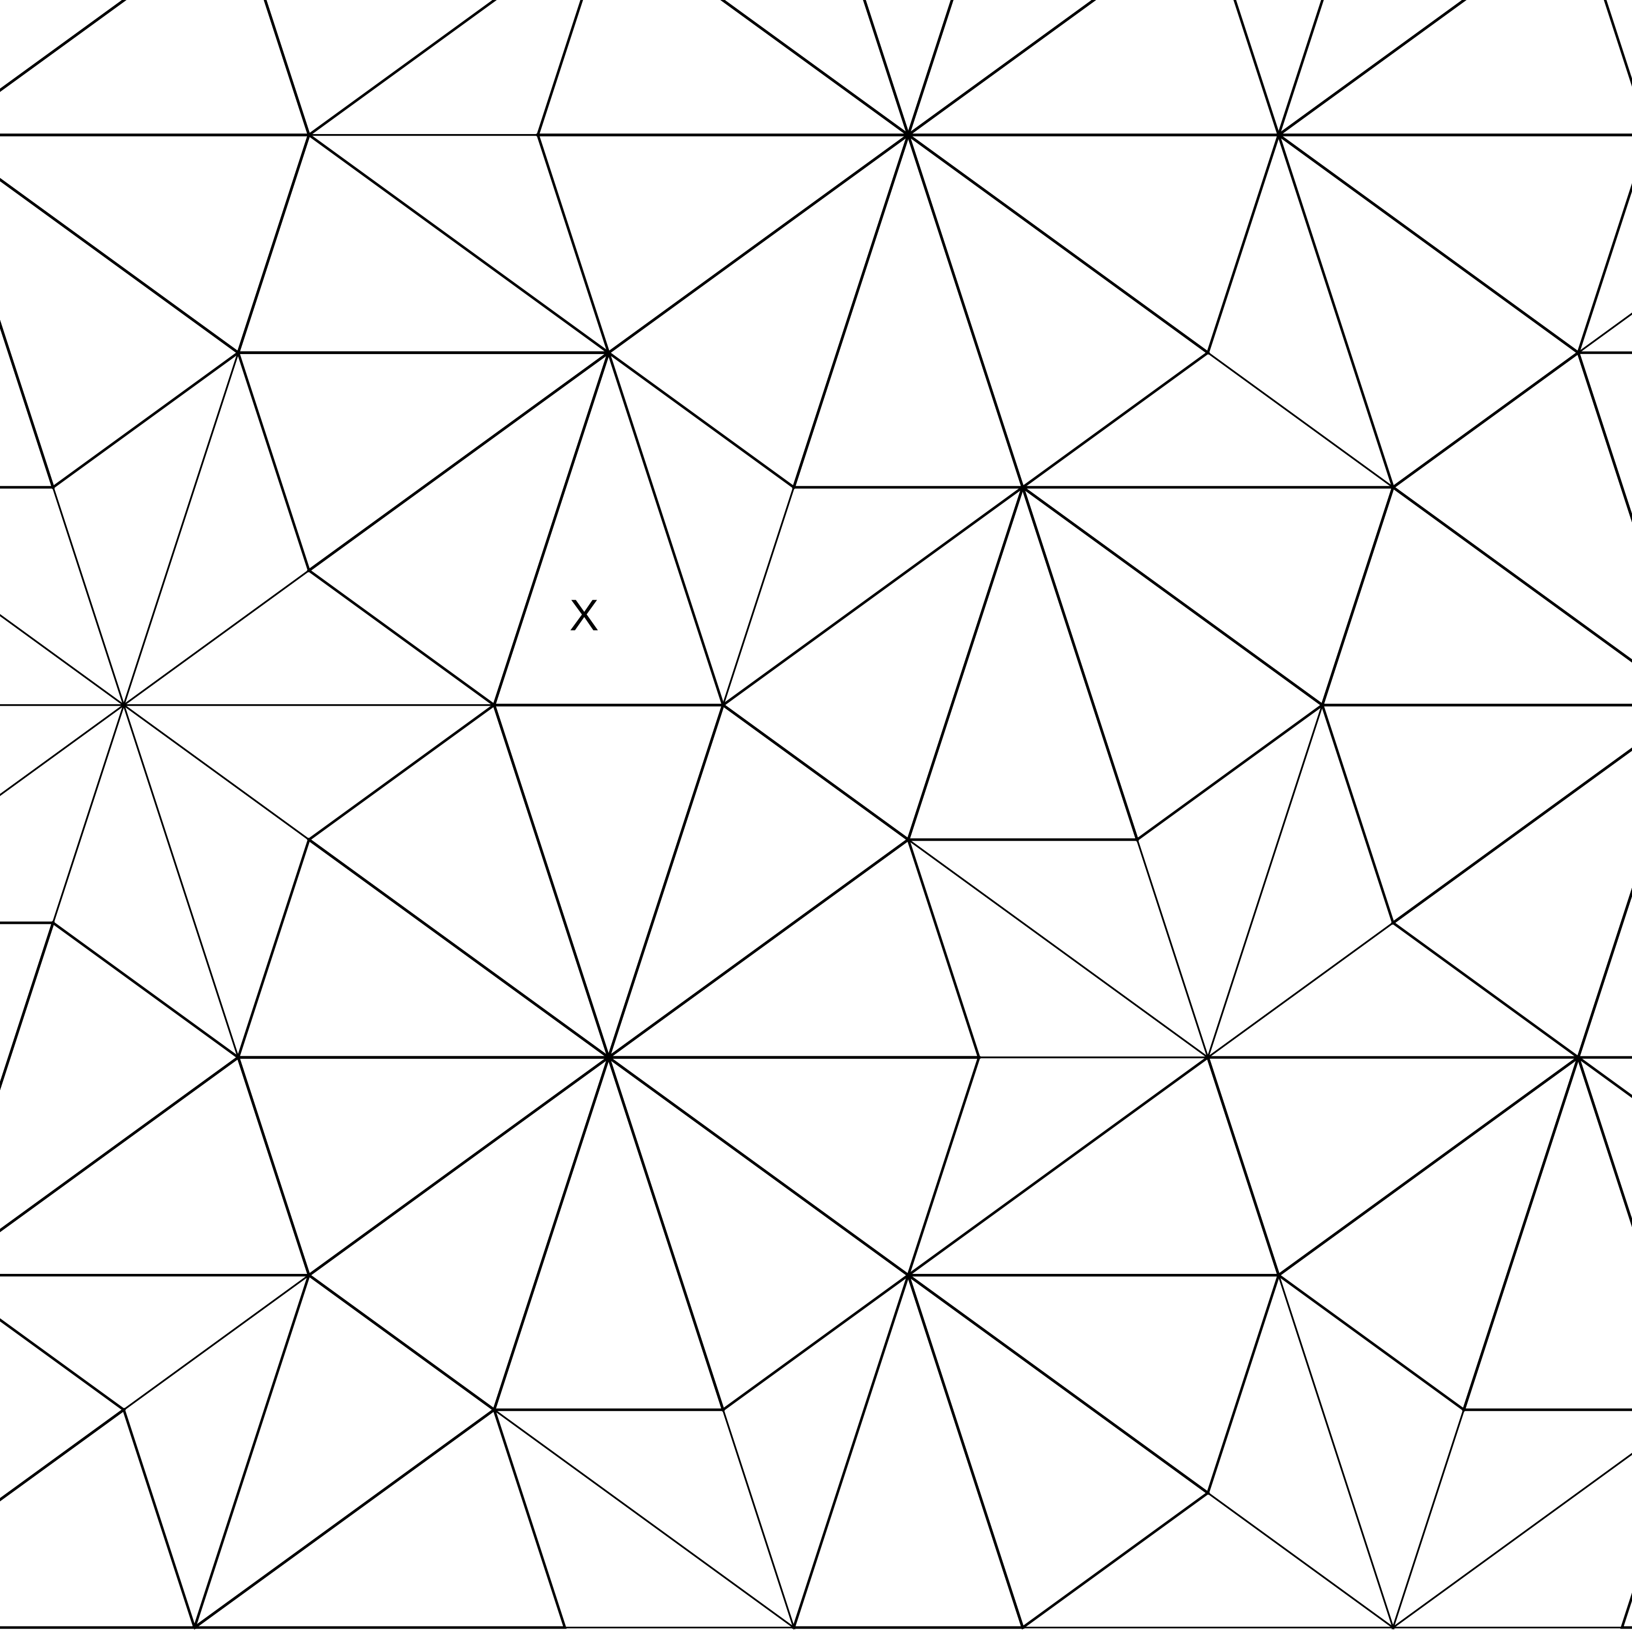

Figure A.9: Formation of an index sequence. The 9th digit is “0”.

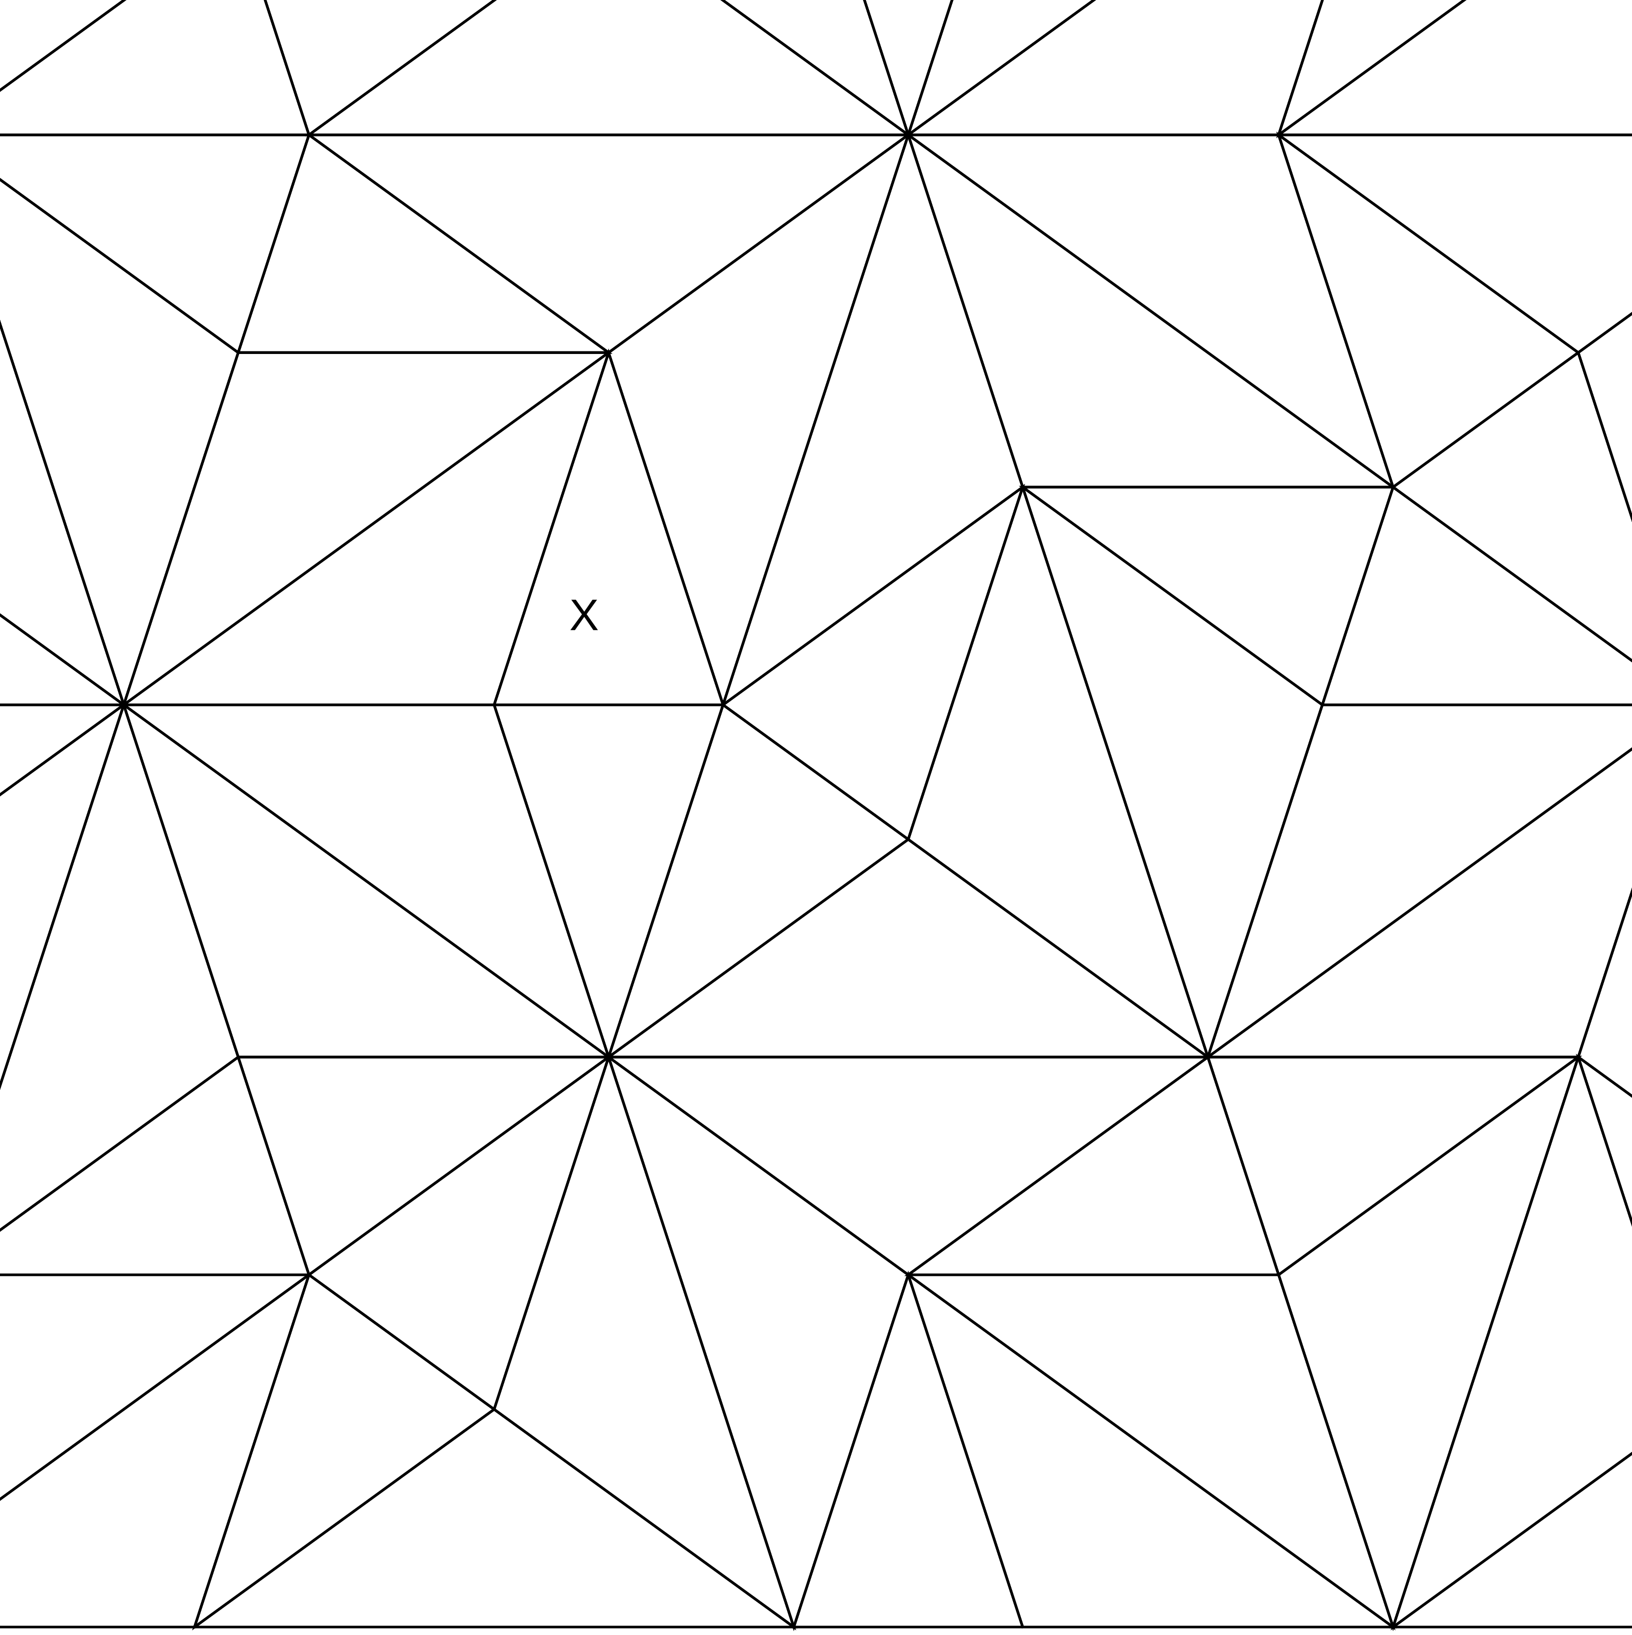

Figure A.10: The 10th digit is “1”. The index sequence is 0001010101...

Appendix B

Inflation Patches and Inflation of Danzer's Tetrahedra

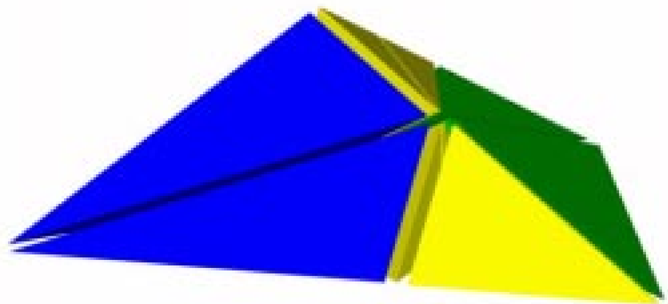

Figure B.1: Inflation patch of an A-tile.

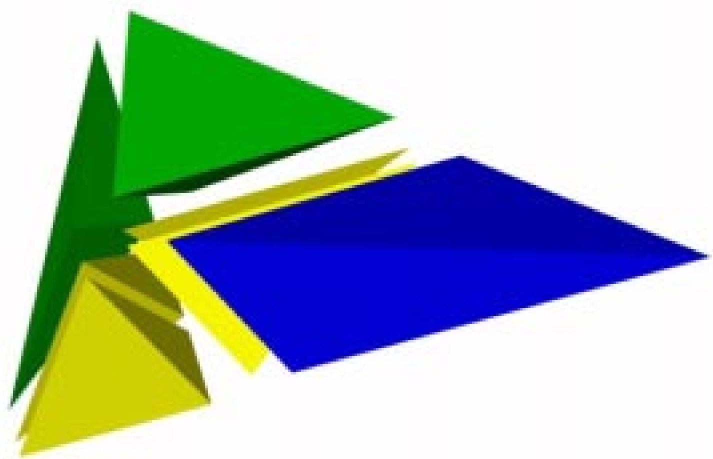

Figure B.2: Inflation patch of a B-tile.

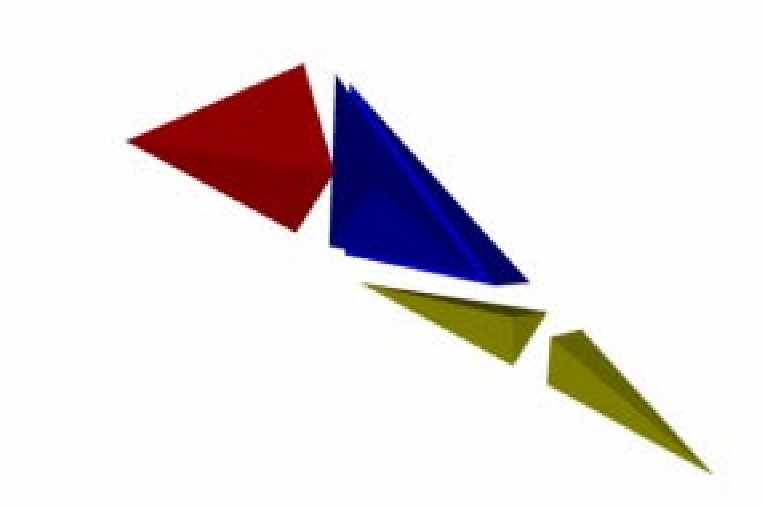

Figure B.3: Inflation patch of a C-tile.

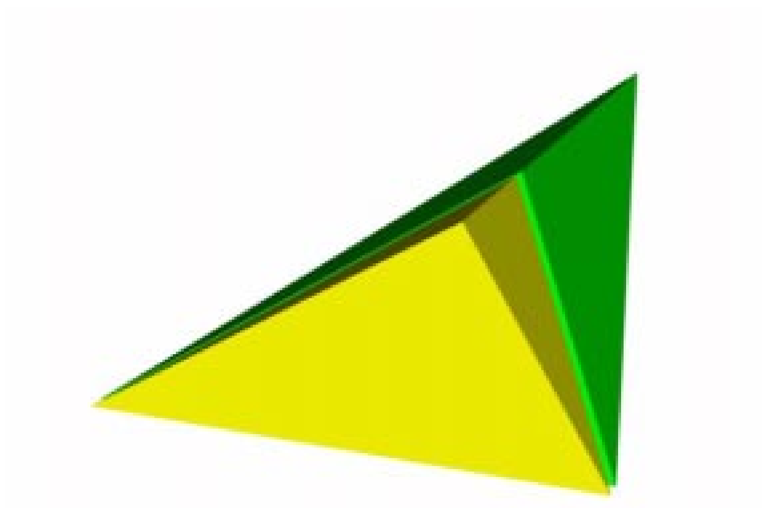

Figure B.4: Inflation patch of a K-tile.

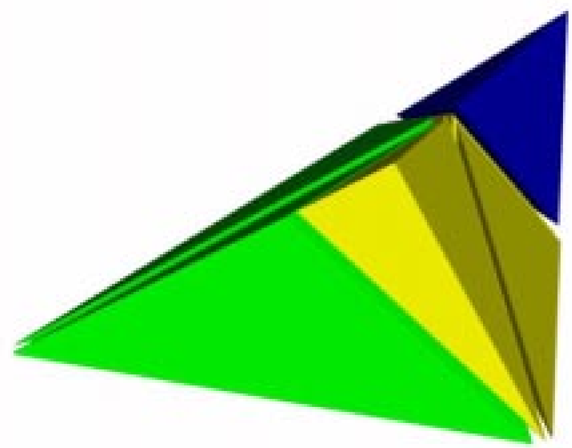

Figure B.5: A K-tile at inflation level 3. (Level 1 is the tile itself and level 2 is its inflation patch.)

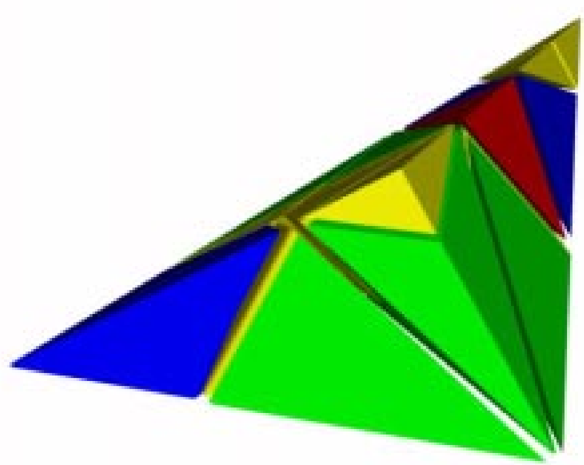

Figure B.6: A K-tile at inflation level 4.

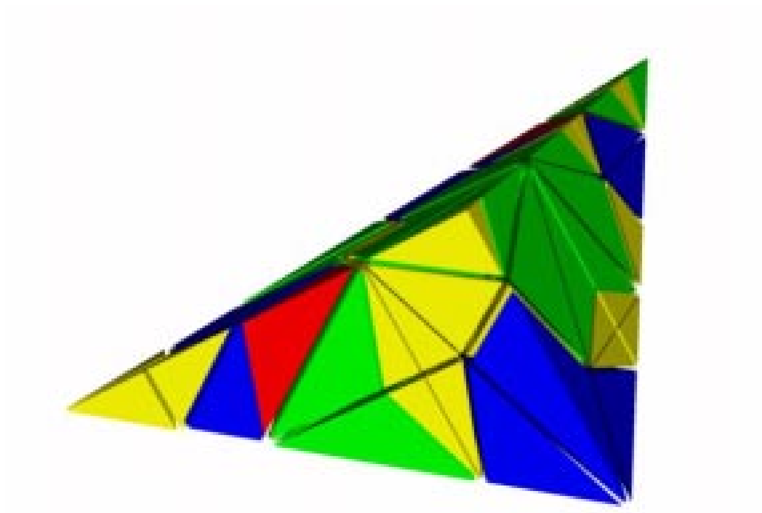

Figure B.7: A K-tile at inflation level 5.

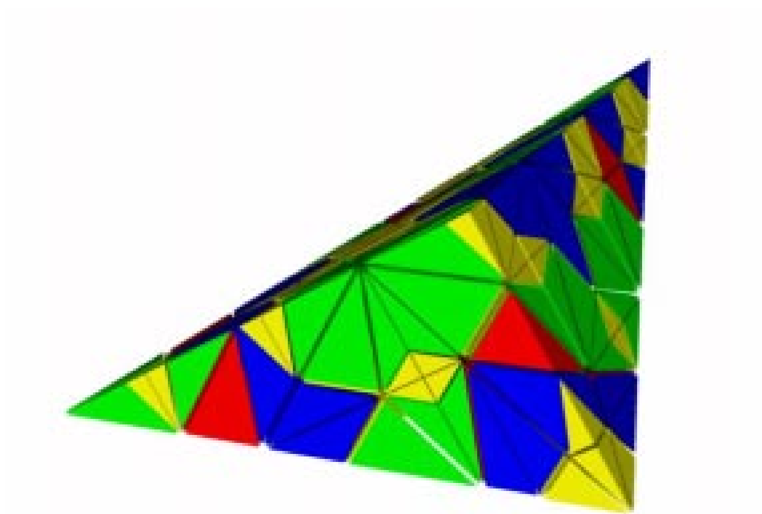

Figure B.8: A K-tile at inflation level 6.

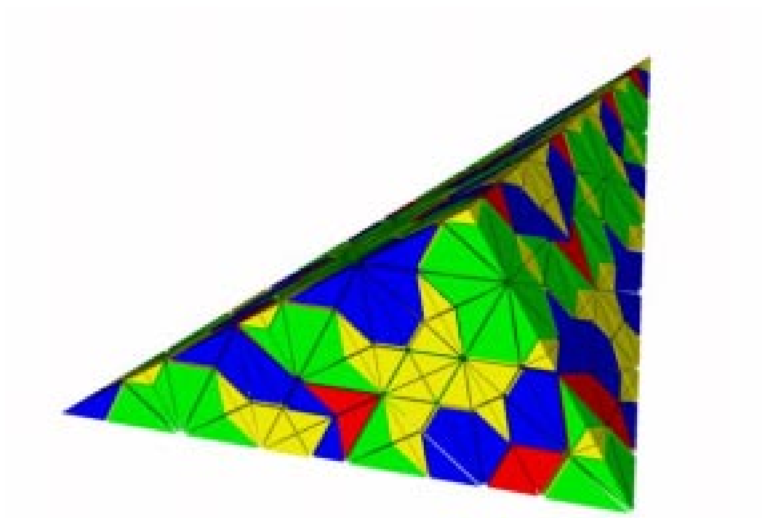

Figure B.9: A K-tile at inflation level 7.

Bibliography

[Bar93] Michael F. Barnsley. *Fractals Everywhere.* Academic Press Professional, Boston, second edition, 1993.

[BBAK+94] M. Baake, S. I. Ben-Abraham, R. Klitzing, P. Kramer, and M. Schlottman. Classification of local configurations in quasicrystals. *Acta Crystallographica*, A50:553–566, 1994.

[BDHJ83] S. B. M. Bell, B. M. Diaz, F. C. Holroyd, and M. J. Jackson. Spatially referenced methods of processing raster and vector data. *Image and Vision Computing*, 4:211–220, 1983.

[Dan89] L. Danzer. Three-dimensional analogs of the planar penrose tilings and quasicrystals. *Discrete Mathematics*, 76:1–7, 1989.

[dB81] N. G. de Bruijn. Algebraic theory of Penrose's non-periodic tilings. *Nederl. Akad. Wetensch. Proc.*, 84:39–66, 1981.

[Dur94] Eugenio Durand. Quasitiler 3.0. http://www.geom.umn.edu/apps/quasitiler/about.html, 1994.

[Gar77] Martin Gardner. Extraordinary nonperiodic tiling that enriches the theory of tiles. *Scientific American*, pages 110–121, January 1977.

[GBS94] Franz Gähler, Michael Baake, and Martin Schlottmann. Binary tiling quasicrystals and matching rules. *Physical Review B*, 50(17), November 1994.

[GKP94] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete mathematics: a foundation for computer science.* Addison-Wesley, Reading, second edition, 1994.

[GS87] Branko Grünbaum and G. C. Shephard. *Tilings and Patterns*, chapter 10. W. H. Freeman and Company, New York, 1987.

[HD91] John C. Hart and Thomas A. DeFanti. Efficient antialiased rendering of 3-d linear fractals. *Computer Graphics*, 25:91–100, 1991.

[JDdB89] Christian Janot, Jean-Marie Dubois, and Marc de Boissieu. Quasiperiodic structures: Another type of long-range order for condensed matter. *American Journal of Physics*, 57(11):972–987, November 1989.

[KB94] Richard Klitzing and Michael Baake. Representation of certain self-similar quasiperiodic tilings with perfect matching rules by discrete point sets. *Journal de Physique I*, 4:893–904, 1994.

[LP87] Waldo F. Lunnon and P. A. B. Pleasants. Quasicrystallographic tilings. *Journal de Math. Pures et Appl.*, 66(3):217–263, 1987.

[LS86] Dov Levine and Paul J. Steinhardt. Quasicrystals. i. definition and structure. *Physical Review B*, 34(2):596–616, July 1986.

[OSDS88] George Y. Onoda, Paul J. Steinhardt, David P. DiVincenzo, and Joshua E. S. Socolar. Growing perfect quasicrystals. *Physical Review Letters*, 60(25):2653–2656, June 1988.

[Pen84] R. Penrose. *Pentaplexity: A class of non-periodic tilings of the plane*, chapter 5, pages 55–65. Pitman, Boston, 1984.

[SG91] Peter W. Stephens and Alan I. Goldman. The structure of quasicrystals. *Scientific American*, 268:44–53, April 1991.

[Soc91] Joshua E.S. Socolar. Growth rules for quasicrystals. In D. P. DiVincenzo and P. J. Steinhardt, editors, *Quasicrystals: The State of the Art*, volume 11 of *Directions in Condensed Matter Physics*. World Scientific, Singapore, 1991.

[SS86] Joshua E. S. Socolar and Paul J. Steinhardt. Quasicrystals. ii. unit-cell configurations. *Physical Review B*, 34(2):617–647, July 1986.

[Ste86] Paul Joseph Steinhard. Quasicrystals. *American Scientist*, 74:586–597, November–December 1986.