Exploring the Danzer Tiling

by
Benjamin Crawford Chaffin

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
May 25, 1998

Abstract

In 1989, Danzer presented a set of four tetrahedra which, along with their mirror images, form aperiodic tilings of space when assembled according to certain matching rules. Though not the first known aperiodic set in three dimensions, the Danzer tetrahedra are attractive for a number of reasons, including several similarities to the two-dimensional Penrose tiles and Robinson triangles.

We present an expanded version of Danzer's paper[Dan89], giving explanations and proofs and removing some of the notation to provide a more readable and comprehensible exposition of these tiles and their basic properties.

The tetrahedra in a Danzer tiling always meet face-to-face and vertex-to-vertex. Thus the next fundamental unit of the tiling beyond the tiles themselves is the vertex configuration, a ball of tiles clustered about and sharing a central vertex. Danzer's paper states that there are exactly 27 vertex configurations. We give a practical algorithm to compute a complete "atlas" of vertex configurations, and show that there are, in fact, 174.

Every patch in a Danzer tiling is duplicated an infinite number of times in every tiling, and the distance to the nearest identical copy is within a bounded multiple γ of the radius of the patch. Using a computational approach, we show how to place a bound on γ.

Acknowledgements

I would like to thank the many people who so frequently and willingly took an interest in my work and helped me with the problem of the day. They are Ed Burger, Bill Lenhart, Jay Sachs, Colin Adams, Peter Nicholas, Derek Sasaki-Scanlon, and Matt Chen, as well as many others.

I thank my reader, Jay Sachs, whose scathing pen and pointed questions were always on the mark.

I am also grateful to my fellow thesis students, Brendan Burns, Linden Minnick, and Jon Burstein, for their good humor and for being an endless (and welcome) source of distraction.

Professor Dick de Veaux of the Williams College Department of Mathematics has been kind enough to allow me to use his Silicon Graphics workstations. Without their powers of visualization I would have been hopelessly lost.

Finally, I thank my advisor, Duane Bailey. His blinding intellect, fascination with our topic, and contagious humor have kept me going through many late nights and difficult problems. This thesis could not have been started or finished without him. Thank you Duane.

Contents

List of Figures

List of Tables

Chapter 1

Background

1.1 Introduction: Penrose Tiling

A *tiling* is a covering of the plane with a (possibly infinite) set of tiles, each of which is a subset of the plane. The shapes of the tiles are known as *prototiles*, and the tiling itself is made up of copies of the prototiles. A tile must be closed; we will consider only the case where the tiles are also bounded and connected, and where the set of prototiles (the *protoset*) is finite. A subset of the tiles in a tiling is called a *patch*. We require that a patch be homeomorphic to a disk. That is, it must be connected, there must be no two regions connected by a single point, and it must contain no holes.

An example of a tiling composed of three prototiles is given in Figure 1.1. Sliding the entire tiling along one of the two arrows shown generates a tiling which is indistinguishable from the original. Tilings with this property have *translational symmetry* and are said to be *periodic*. This seems like a natural property of any tiling. If we cover large areas with only a finite number of shapes, then, intuitively, we must place them in a repeating pattern.

Contrary to intuition, however, there exist finite protosets which are *aperiodic*. That is, they produce only tilings with no translational symmetry. No matter how we translate the tiling, it never looks exactly the same as before the translation. In 1966, Berger discovered a set of 20,426 tiles that would only tile the plane aperiodically. By 1974, Penrose had reduced that number to two. Penrose's thin and thick rhombs, when fitted together according to certain *matching rules*, produce only tilings without translational symmetry. Though we usually find it more convenient to think of the matching rules as restrictions on the way tiles may be fitted together, it is almost always possible to enforce the matching rules by making notches and knobs on the sides of the tiles like pieces of a jigsaw puzzle. The Penrose tiles (and

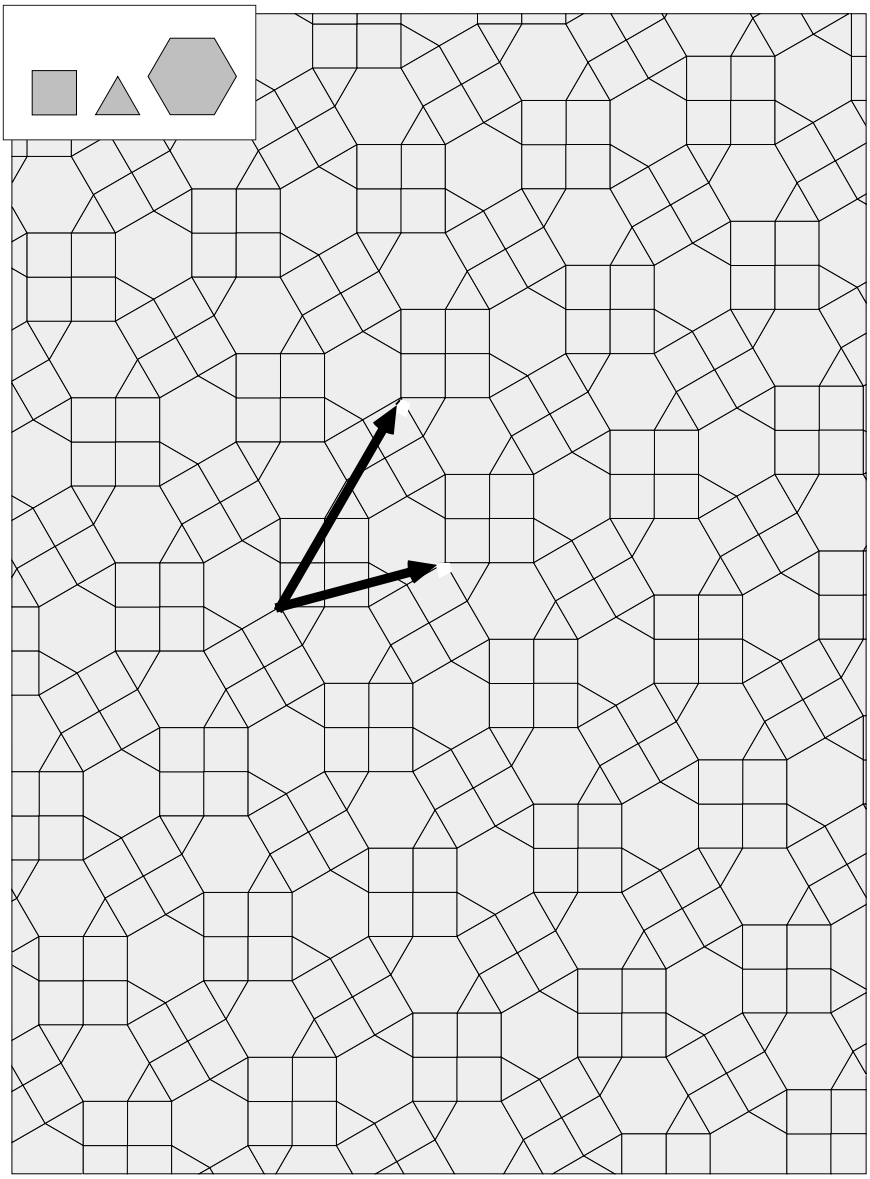

Figure 1.1: A periodic tiling. The upper left corner shows the three prototiles used to construct the tiling. The tiling can be translated according to either of the two arrows without effect. Tilings with such translations are *periodic*.[Tre96]

their close cousins, the Conway kites and darts and the Robinson triangles) have a significant body of work relating to them. We mention only a few interesting results here:

These tilings obey inflation and deflation rules. Kites and darts, for example, can be subdivided into smaller versions of themselves or combined into larger ones. This can be used to prove that perfect infinite tilings exist and that they are aperiodic.

The Penrose tiles can be decorated with Ammann bars, line segments drawn on the tiles. All tiles of a particular shape are decorated in the same way. When the tiles are fitted in a tiling, the line segments join across tiles to form infinite lines. There are five orientations of infinite Ammann lines, each set parallel to one side of a regular pentagon. The distance between two adjacent parallel lines is always one of two values, which we will call *long* and *short*. The ratio of the long to the short interval is the golden ratio, $\tau = (1 + \sqrt{5})/2 \approx 1.618\ldots$, and the sequence of long and short intervals itself forms a particular type of aperiodic sequence called a *Fibonacci sequence*, also known as a *musical sequence*.

The placement of a few tiles restricts the way in which a tiling can be constructed around them. For example, given a patch, it can be shown that although infinitely many tilings can be grown outward from this patch, all extensions must contain a particular tile in a particular position. This tile is said to be *forced* by the patch. The locations of tiles are related to the intersections of Ammann lines. In fact, it can be shown that the Ammann lines determine the tiling rather than the other way around. The Ammann lines determined by a patch can force an infinite number of other Ammann lines. A patch can force tiles at arbitrarily large distances. The collection of all tiles forced by a given patch is called the *empire* of the patch. These non-local effects present a difficulty when trying to "grow" Penrose tilings by adding tiles to the edge of a patch. There are instances when it is impossible to tell from purely local information that a given tile is forced. Stated a different way, a locally legal tile placement can force an inconsistency at a distance, thus producing a defect in the tiling. Recent work by Minnick[Min98] has shown that Ammann lines forced by a patch are not entirely determined by the Ammann lines that pass through it. In some cases there are several possible arrangements of tiles, all of which determine a particular Ammann line.

The fact that a Penrose tiling does not have translational symmetry does not imply that all parts of the tiling look somehow "different." In fact, any patch, no matter how large, is guaranteed to appear an infinite number of times in *every* tiling, and the distance from one instance to its nearest copy is guaranteed to be less than some fixed multiple of its radius. This underlies a proof that the equivalence of two tilings is not decidable.

These tilings exhibit long-range fivefold rotational symmetry. It is not possible to construct a periodic tiling of the plane with fivefold symmetry.

1.2 Quasicrystals

Until 1984, solids were grossly categorized into two classes, glasses and crystals. The atoms in a glass are packed randomly, while those in a crystal are arranged in a regular lattice. The parallel planes of similarly oriented atoms in a crystal reflect incoming particles in such a way that the reflected waves form distinct interference patterns that can be used to deduce the structure of the crystal. A central tenet of traditional crystallography is that the only possible rotational symmetries of a crystal plane are two-fold, three-fold, four-fold, and six-fold, reflecting the fact that triangles, squares, and hexagons are the only regular polygons that can tile the plane with no gaps.

In 1984, Schectman, Blech, Gratias, and Cahn rapidly solidified a molten alloy of aluminum and manganese. Electron diffraction of the resulting solid showed that it had fivefold rotational symmetry, previously considered impossible. It was hypothesized that the arrangement of atoms in this *quasicrystal* resembled a three-dimensional aperiodic tiling with icosahedral symmetry. The work on aperiodic tilings, previously published as recreational mathematics, suddenly had an application in materials science, and since that time articles on tiling have appeared in journals of both mathematics and physical chemistry. Later researchers have refined production methods and have produced quasicrystals with 8-fold, 9-fold, 12-fold, and other previously "forbidden" rotational symmetries. There is some opposition to the conjecture that quasicrystals resemble aperiodic tilings. It remains unclear how atoms attaching themselves to the outer edge of a growing quasicrystal could obey the non-local requirements of forced tiles.

1.3 Recommended Reading

Grünbaum and Shephard's book[GS87] is a broad and readable text on tiling. Chapter 10 is essential reading for anyone interested in aperiodic tiling.

Penrose's "Pentaplexity"[Pen84] offers an interesting account of the development of the Penrose tiles and gives two simple proofs that the tilings they produce are aperiodic.

In a pair of papers, Levine and Steinhardt[LS86] and Socolar and Steinhardt[SS86] give an extensive overview of the structure and properties of quasicrystals. Saitoh, Tsuda, and Tanaka[STT97] give a chemistry-oriented description quasicrystal structure.

Considerable attention has been focused on methods for growing tilings according to purely local matching rules, in order to demonstrate a mechanism by which quasicrystals might form. Jeong and Steinhardt[JS94] present a model of quasicrys-

tal formation that relies on overlapping low-energy clusters. Gummelt[Gum96] gives a single decagonal prototile that can be used to construct Penrose-equivalent tilings by allowing tiles to overlap according to matching rules. Jeong and Steinhardt[JS97] give a simpler proof that the coverings produced by Gummelt's decagons are equivalent to Penrose tilings and show that maximizing the density of the tiles is equivalent to Gummelt's overlapping rules. Steinhardt[Ste96] summarizes these results. Onoda, Steinhardt, DiVicenzo, and Socolar[OSDS88] developed a method of growing perfect Penrose tilings using local matching rules by adding a thick rhombus in a particular way whenever no placements are forced. Socolar[Soc91] gives a detailed account of purely local matching rules for Penrose tilings that result in structures with long-range quasicystalline order. Baake, Ben-Abraham, Klitzing, Kramer, and Schlottman[BBAK$^+$94] catalog all possible arrangements of tiles around a vertex (including those which do not admit a tiling) for two-dimensional and three-dimensional tilings. This paper will be of particular interest to us in Chapter 5 when we do the same for the Danzer tiling.

Steinhardt and Socolar have constructed a three-dimensional analog of the Penrose tiling. Danzer[Dan89] gives a set of four tetrahedra extracted from the symmetry planes of a regular icosahedron that produce an aperiodic solid when assembled according to matching rules. The tetrahedra obey inflation and deflation rules. We are particularly interested in the Danzer tiling because it has simple and elegant inflation rules and consists of only four tiles and their mirror images.

Chapter 2

A Basis for Tiling in Three Dimensions

The German mathematician Ludwig Danzer published a paper in 1989 in which he outlined a model for aperiodic tilings in three dimensions[Dan89]. The paper is in two parts. The first gives a "semi-axiomatic" approach, stating several requirements which, if met by a set of tiles, will ensure that they form only aperiodic tilings with certain desirable properties. The second part gives an aperiodic set of four tetrahedra which meet all of the requirements and gives some of the properties of the tiling they produce. Danzer's paper is quite short and uses a great deal of mathematical notation. All of his theorems are stated without proof, including the several parts of the main theorem concerning the properties of his tiling.

In this chapter we survey the first section of the paper, which concerns some rules an aperiodic set must follow. We follow Danzer's presentation fairly closely, providing explanations of difficult concepts and supplying proofs of the theorems.

2.1 Definitions

A *prototile* is a closed ball of points. A *tile* is a copy of some prototile, a set of points congruent without reflection to it. In general, we use "prototile" and "type of tile" to mean the same thing, and use "tile" only when referring to a particular instance of some prototile. A *protoset* is a set of prototiles each with some maximum and minimum size; that is, each tile must contain a sphere of radius $r_{\min}$ and be contained within a sphere of radius $r_{\max}$. A *patch* is a (possibly infinite) set of tiles such that no two tiles overlap and their union has no holes. That is, the pairwise intersection of the tiles in the patch is empty and the patch is homeomorphic to a ball. If a patch is finite, then it has a diameter, which is defined to be the maximum distance

between any two points in the patch. The radius of a patch is half the diameter. These definitions correspond to the intuitive definition of three-dimensional tiles and tilings: a tile is a single piece of space with no strange features like holes or sections connected only by a point or a line, and a patch is a collection of tiles fitted snugly together.

A tiling is *infinite* if covers all of space. A tiling $\mathcal{A}$ is *aperiodic* if it is infinite and has no translational symmetry, *i.e.* there is no translation $\vec{T}$ such that $\vec{T}(\mathcal{A}) = \mathcal{A}$. A protoset $\mathcal{S}$ is an *aperiodic set* if all tilings that can be created using the tiles of $\mathcal{S}$ are aperiodic. We will generally use 'tiling' to mean an infinite aperiodic tiling unless otherwise specified.

2.2 Basic conditions for a family of prototiles

Let $\mathcal{F}$ be a protoset $S_1, S_2, \ldots, S_m$ of polyhedra. These will be the models for the tiles we use to construct a tiling. To ensure that all tilings that can be formed from these tiles are aperiodic, we may choose to impose matching rules to restrict the ways in which tiles may be placed adjacent to each other. We define an equivalence relation $\sim$ on the set of all the faces of all the prototiles that will encode the matching conditions; two faces may be placed next to each other only if they are equivalent according to $\sim$. We require that $\sim$ be at least as strong as Euclidean congruence. That is, two matching faces must always be congruent, but $\sim$ may impose some additional restrictions. This ensures that the tiling is always face-to-face. An $(\mathcal{F}, \sim)$-*tiling* is a patch $\mathcal{P}$ composed of tiles each of which is a copy of one of the prototiles in $\mathcal{F}$, matched according to $\sim$.

Danzer places five requirements on $\mathcal{F}$ and $\mathcal{P}$. We have already stated the first of these, which is that the tiling must be face-to-face. Danzer gives this requirement the label **(A)**. The other conditions are given in the following four sections. The first two requirements guarantee the existence of infinite $(\mathcal{F}, \sim)$-tilings and that all such tilings are aperiodic. The remaining two requirements ensure some desirable properties of all $(\mathcal{F}, \sim)$-tilings. As we state each condition we will elaborate on its meaning and examine some of the implications it holds for the structure of a tiling.

2.2.1 Inflation

(Infl) Each tile must be uniquely inflatable. That is, we must be able to construct each tile out of smaller copies of some of the prototiles of $\mathcal{F}$, and there must be only one way in which this can be done. All the smaller tiles must be scaled by the constant factor η^{-1}, where $\eta > 1$. If two faces are equivalent under $\sim$ (that is, they may meet in a tiling), then the way in which they are

dissected by an inflation must be congruent. This maintains the face-to-face requirement when adjacent tiles are inflated.

In general, inflating a patch means inflating every tile in it. Inflating all the tiles in an $(\mathcal{F}, \sim)$-tiling $\mathcal{P}$ produces an $(\eta^{-1}\mathcal{F}, \sim)$-tiling $\mathcal{Q}$ which occupies the same space. We call $\mathcal{Q}$ the *inflation* of $\mathcal{P}$, and denote it $\text{infl}(\mathcal{P})$. If we then scale all the tiles in $\mathcal{Q}$ by η, so that they are the same size as the original tiles of $\mathcal{P}$, we obtain $\eta\mathcal{Q}$, called the *expansion* of $\mathcal{Q}$. We can iterate the inflation-expansion process. When we inflate a tile twice, we first inflate it once, and then inflate every tile in the resulting patch. We write the k^{th} iteration of the inflation as $\text{infl}^k(\mathcal{P})$.

If we iteratively inflate and expand a patch, we can generate arbitrarily large tilings. This means that we can generate an infinite tiling by inflating a tile an infinite number of times, expanding by η each time.

We now consider Robinson's triangles as a two-dimensional example of inflation. There are two sets of four Robinson triangles, known as the **A** tiles and the **B** tiles. In each set there are two shapes and their mirror images. We use capital and lower case letters to denote the shape of a tile, and add a prime ($'$) if it is mirrored (or *left-handed*). So the four **A** tiles are A, A', a, and a', and similarly for the **B** tiles. Each set is aperiodic, so we can build a nonperiodic tiling using exclusively **A** or **B** tiles. We require two sets of tiles because when **A** tiles are inflated, they produce **B** tiles, and vice versa. Note that we may choose to perform double inflations at each step to maintain a single set of tiles.

Figure 2.1 shows the two sets of tiles. The dots and arrows are markings which indicate the matching conditions — tiles may only be placed so adjacent markings are aligned. Figure 2.2 shows the inflation of the right-handed Robinson tiles. **B** tiles can be made of smaller **A** tiles, and **A** tiles can be made of smaller **B** tiles. In this case, the scaling factor η is $\tau = \frac{1+\sqrt{5}}{2} = 1.618034\ldots$, the golden ratio. The inflations of the left-handed tiles are the same with the handedness of all the components reversed.

2.2.2 Deflation

(Defl) Every infinite tiling must have a unique *deflation*, written $\text{defl}(\mathcal{P})$, which inverts the inflation mechanism defined above by combining patches into single larger tiles. For every infinite $(\mathcal{F}, \sim)$-tiling $\mathcal{P}$ there is exactly one $(\eta\mathcal{F}, \sim)$-tiling $\mathcal{Q}$ such that $\text{infl}(\mathcal{Q}) = \mathcal{P}$.

The deflation process can also be iterated. We denote the k^{th} deflation of a patch $\mathcal{P}$ by $\text{defl}^k(\mathcal{P})$. Note that only infinite tilings are guaranteed to be deflatable. Not all finite patches can be deflated. For instance, if we take a tile, inflate it to get a

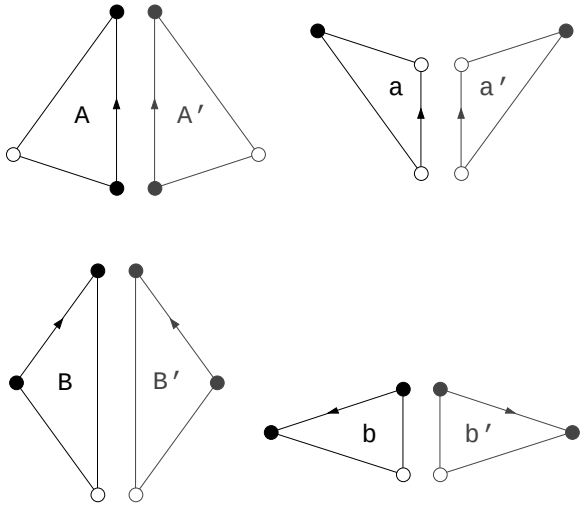

Figure 2.1: The Robinson triangles. The A tiles (top) and B tiles (bottom) form two aperiodic sets. There is a right-handed and left-handed version of each tile; left-handed tiles are indicated with a prime ($'$) and drawn in grey.

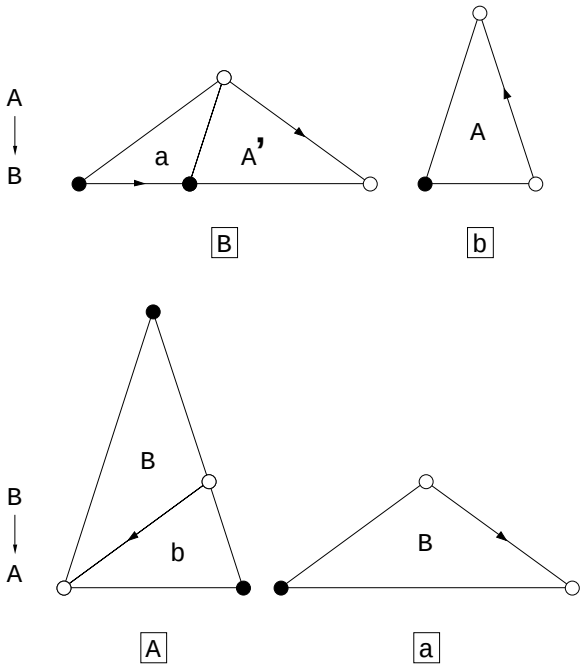

Figure 2.2: Inflating the right-handed Robinson triangles. Each of the four patches shows the inflation of the tile in the boxed label below it.

patch of several tiles, and then delete one of them, we have a patch which cannot be deflated. However, any patch which is part of an infinite tiling is always contained within a finite patch which has a unique deflation, as we will show in the proof of Theorem 2.2.

The fact that inflations and deflations are unique and can be iterated allows us to prove the following Lemma:

Lemma 2.1 *All infinite $(\mathcal{F}, \sim)$-tilings are aperiodic.*

Proof Suppose that we have an infinite $(\mathcal{F}, \sim)$-tiling $\mathcal{A}$ which is periodic. Then (by the definition of periodicity) there is some nontrivial translation $\vec{T}$ such that a translated copy of $\mathcal{A}$ overlaps $\mathcal{A}$ exactly, *i.e.* $\vec{T}(\mathcal{A}) = \mathcal{A}$.

Consider a vertex v_1 of some tile in $\mathcal{A}$ and its image under $\vec{T}$, which we will call v_2. Iteratively deflate $\mathcal{A}$ until the shortest side of any tile is longer than the length of $\vec{T}$, obtaining the tiling $\text{defl}^k(\mathcal{A})$, and suppose that we have chosen v_1 so that it is still a vertex in the deflated tiling. (We can always choose such a v_1 because vertices are never removed by inflation. Thus any vertex in $\text{defl}^k(\mathcal{A})$ is also a vertex in $\mathcal{A}$.) Clearly v_2 cannot lie on a vertex of $\text{defl}^k(\mathcal{A})$, since the distance from v_1 to v_2 is less than the distance from v_1 to the nearest vertex.

Now suppose that we translate $\mathcal{A}$ by $\vec{T}$ before performing the same process as above. Recall that $\vec{T}$ is defined such that $\vec{T}(\mathcal{A})= \mathcal{A}$, so this translation should have no effect on the deflation. As before, we deflate the tiling, obtaining $\text{defl}^k(\vec{T}(\mathcal{A}))$. The deflation produces a vertex at the new location of v_1, which is v_2. But we showed above that when we deflate the untranslated tiling $\mathcal{A}$, there is no vertex at v_2. Thus the tilings $\mathcal{A}$ and $\vec{T}(\mathcal{A})$ are identical, yet they produce two distinct deflations. This violates the requirement that deflation be unique. Therefore there is no such $\vec{T}$, and all infinite $(\mathcal{F}, \sim)$-tilings must be aperiodic. □

See Figure 2.3 for a demonstration of the proof below using a Robinson tiling.

Any periodic tiling obeys inflation and deflation rules; consider a square grid, where groups of 4 squares can be combined to form a larger square. The difference is that the deflation is not unique. A single square could be part of more than one cluster of four, as illustrated in Figure 2.4.

2.2.3 Minimality

A set of tiles which meets the requirements above can be inflated and deflated uniquely and is guaranteed always to produce an aperiodic tiling. There are some other characteristics of tilings which are useful. To ensure that all tilings have some other desirable properties, Danzer imposes additional restrictions on the set

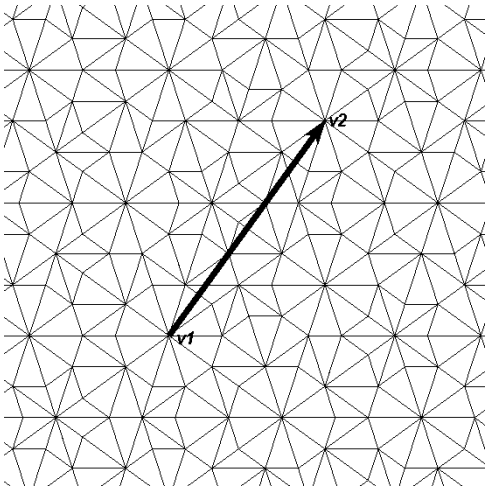

(a) A portion of a Robinson tiling $\mathcal{A}$, showing a proposed translation $\vec{T}$.

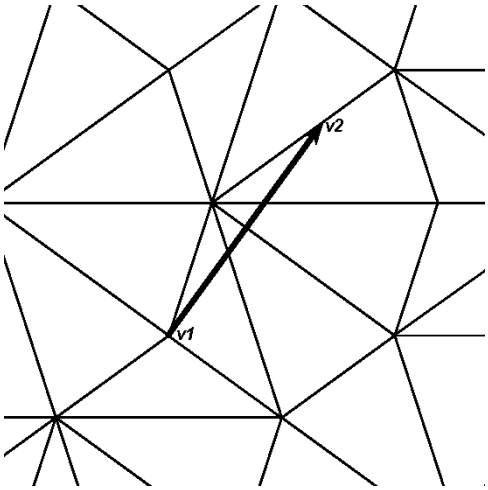

(b) After k deflations, v_1 is still a vertex but v_2 is not.

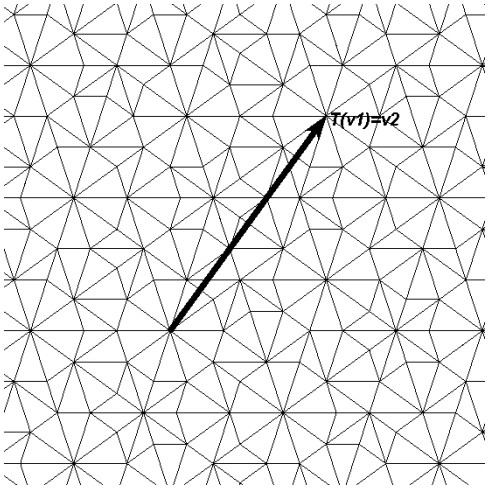

(c) The tiling $\vec{T}(\mathcal{A})$. Vertex v_1 has been translated to the former location of v_2.

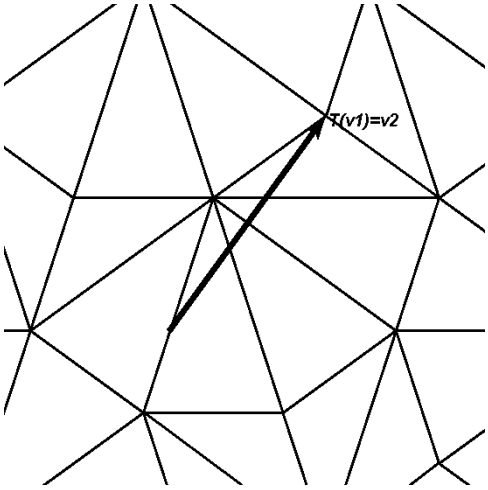

(d) A deflation of $\vec{T}(\mathcal{A})$. There is now a vertex at $\vec{T}(v_1) = v_2$.

Figure 2.3: Unique inflation and deflation guarantee aperiodicity. If $\mathcal{A}$ is periodic, then $\vec{T}(\mathcal{A}) = \mathcal{A}$. There is a vertex at v_2 in $\text{defl}^k(\vec{T}(\mathcal{A}))$, but not in $\text{defl}^k(\mathcal{A})$. Thus two identical periodic tilings can produce two different deflations.

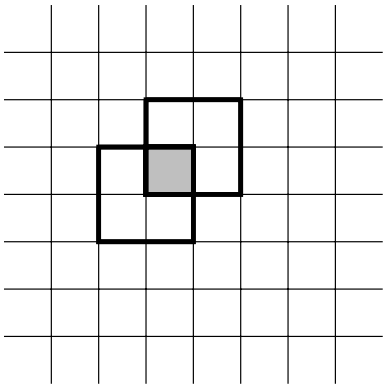

Figure 2.4: Deflation of a periodic tiling is not unique. The shaded tile could be part of either of the larger outlined squares.

of prototiles. Before we look at the first of these restrictions, we will explore an interesting way of denoting inflation and deflation. This will be useful in proving some of the consequences of the next requirement.

We can write the components of the inflation of each tile in an *inflation matrix* $\mathcal{M}$ where $\mathcal{M}_{i,j}$ is the number of copies of prototile i used in the inflation of j. It is clear that we can also represent the second inflation of every tile as a matrix, which we will call $\mathcal{M}^2$. The relationship between $\mathcal{M}$ and $\mathcal{M}^2$ is most easily determined by looking at an example. Consider a system of two tiles with the following inflation matrix:

$\mathcal{M}$	X	Y
X	0	1
Y	2	3

.

This table shows that one X tile inflates to 2 Y tiles, and a Y tile inflates to an X and 3 Y's. Inflating an X tile twice produces two Y's in the first inflation, and two X's and 6 Y's in the second inflation. Performing the same process on a Y tile, we see that after two inflations a Y becomes 3 X's and 11 Y's. So the second inflation matrix looks like this:

$\mathcal{M}^2$	X	Y
X	2	3
Y	6	11

.

To see the relationship between $\mathcal{M}$ and $\mathcal{M}^2$, let's take a closer look at how we derived $\mathcal{M}^2$ and consider the number of Y tiles used in the second inflation of an X tile. We use two Y's for each X in the first inflation and three for each Y. We

multiply the number of X's in the first inflation by the number of Y's each one will produce, and then add the product of the Y's in the first inflation with the number of Y's *they* will produce. In this case, the formula is $(0)(2) + (2)(3) = 6$. The key is that we can also write this as the product of row Y and column X of $\mathcal{M}$: $(2,3) \cdot (0,2)^T = 6$. This formula is true for each location in the matrix, and so $\mathcal{M}^2$ is in fact exactly what its name suggests, the matrix $\mathcal{M}$ squared:

$$\begin{pmatrix} 0 & 1 \\ 2 & 3 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 2 & 3 \end{pmatrix} = \begin{pmatrix} 2 & 3 \\ 6 & 11 \end{pmatrix}.$$

In general, $\mathcal{M}^n$ is the n^{th} inflation matrix. Each column of $\mathcal{M}^n$ represents the number of each type of tile used in the n^{th} inflation of a particular prototile. Now that we can characterize the inflation mechanism as a matrix and any level of inflation as a power of that matrix, we are ready to state the third requirement we would like our tilings to obey:

(Min) The set of prototiles must be minimal; that is, there must be no subset of the prototiles that obeys the same inflation and deflation mechanism.

Minimality gives rise to several interesting properties of an $(\mathcal{F}, \sim)$-tiling. First, note that an alternate statement of the requirement is that if we begin to generate an infinite tiling by repeatedly inflating a single prototile, we must at some point generate a patch containing all types of tiles. If we could inflate forever without ever using a particular prototile, then $\mathcal{F}$ would not be minimal. Thus if we inflate any prototile from a minimal set a sufficient number of times, we get a patch containing all types of tiles. We define J to be the maximum number of times we need to inflate any prototile to get such a patch. Note that J must be strictly less than the number of prototiles: each inflation must introduce at least one new type of tile until all are present. If we inflated once without obtaining a new prototile, then we would never produce a new one after that. Let us examine what this implies in terms of the inflation matrix.

Recall that a column in $\mathcal{M}$, which we will label C, represents the number of each type of tile generated by the inflation of prototile C. Obviously no column in $\mathcal{M}$ can be all 0's, since that would mean that tile C inflates to nothing. If we have a minimal set of tiles, it is also impossible for any row of $\mathcal{M}$ to be all zeros, since that would mean that some tile was never produced by an inflation, and the set of tiles would not be minimal. Minimality may impose more restrictions than these on the form of the inflation matrix $\mathcal{M}$, but these are all we will need in the coming discussion.

Now consider the implications for higher powers of $\mathcal{M}$. If the k^{th} inflation of tile C contains all types of tiles, then column C of $\mathcal{M}^k$ has no entries which are

zero. So when we say that inflating any tile a sufficient number of times produces a patch containing all types of tiles, we mean that for each column of $\mathcal{M}$ there is some power of $\mathcal{M}$ in which all the elements of that column are positive.

In fact, once a column in $\mathcal{M}^k$ has no zero entries, there can never be any zero entries in any higher power. Consider what happens to a uniformly positive column C when we multiply by $\mathcal{M}$. Each entry in $\mathcal{M}C$ is the product of C with some row of $\mathcal{M}$. Because C has no zeros (and because there obviously can be no negative entries in $\mathcal{M}$), the only way for an entry of $\mathcal{M}C$ to be zero is for a row of $\mathcal{M}$ to be all zeros, and since we assumed a minimal set this cannot be the case. Thus if we have a patch containing all types of tiles (a column with no zeros), no type of tile will disappear completely when we inflate the patch. In other words, if a patch $\mathcal{P}$ contains all types of tiles, then $\text{infl}(\mathcal{P})$ also contains all types of tiles.

So far we know that **(Min)** requires that every column in $\mathcal{M}^k$ have no zeros for some $k \leq J$ and that inflating a patch containing all prototiles produces another patch which also contains all prototiles. Combining these observations, we see that in fact inflating any prototile exactly J times gives a patch containing all types of tiles. We may have gotten such a patch before the J^{th} inflation, but we have shown that we have not lost any prototile by continuing to the J^{th} inflation. Any subsequent inflations will also not cause any prototile to disappear completely from the patch. In terms of the inflation matrix, $\mathcal{M}^J$ has no zero entries in any column, and all higher powers of $\mathcal{M}$ also have no zero entries. This conclusion is important enough to be restated more formally in the following theorem:

Theorem 2.1 *There exists some natural number J such that for every $j \geq J$ and every prototile S_n, $\text{infl}^j(S_n)$ contains at least one instance of every prototile. Thus if $\mathcal{M}$ is the inflation matrix, then $\mathcal{M}^j$ contains no zero entries for every $j \geq J$.*

Corollary 2.1 *The distance between a tile and its nearest identical copy is bounded.*

Proof Inflating any prototile $2J$ times produces a patch containing more than one instance of each prototile. Every tile T in an infinite tiling lies within a patch $\mathcal{I}$ that is the $2J^{\text{th}}$ inflation of some prototile. Since $\mathcal{I}$ contains more than one instance of all prototiles, the distance to the nearest copy of T can be no more than the diameter of $\mathcal{I}$. □

Corollary 2.2 *Every prototile occurs with a non-zero relative frequency in an infinite tiling. That is, the ratio of the number of instances of one prototile to the total number of tiles in the tiling is non-zero.*

Proof An inflation increases the number of tiles in a patch by some factor β, which can be no larger than the maximum number of tiles in the inflation of any prototile. Thus if a patch $\mathcal{P}$ contains n tiles, then $\text{infl}^J(\mathcal{P})$ contains no more than $\beta^J n$ tiles. Every tile in $\mathcal{P}$ produces at least one instance of each prototile after J inflations. Thus there must be at least n instances of every prototile in $\text{infl}^J(\mathcal{P})$, and so each prototile must represent at least $1/\beta^J$ of the tiles in the patch. This is true regardless of the starting patch, and thus we conclude that the relative frequency of each prototile is non-zero in arbitrarily large patches. □

2.2.4 Local Isomorphism

By requiring the tiling to be face-to-face, we also ensure that the tiling is vertex-to-vertex. A vertex of one tile can never intersect another tile in the middle of one of its faces or along one of its edges, only at one of its vertices. Thus any interior vertex of a patch is shared by one corner of each of a number of tiles. A *vertex configuration* is an arrangement of tiles around a central vertex so that their union subtends the full solid angle around the center. That is, all points within any sufficiently small ball around the central vertex are covered by some tile. If all the tiles of a vertex configuration C share the vertex v, then we say that C is *centered* at v (though v may not be the center of mass of the vertex configuration).

The final requirement we impose on an $(\mathcal{F}, \sim)$-tiling is as follows:

(EVC) Every vertex configuration which occurs in an infinite tiling must occur in the patch generated by inflating some prototile a sufficient number of times. We will call the maximum number of inflations needed to produce any vertex configuration H. (So there is some configuration which never occurs before the H^{th} inflation of any prototile.)

This final requirement implies that a $(\mathcal{F}, \sim)$-tiling is *locally isomorphic*, as expressed in the following theorem:

Theorem 2.2 *Any patch in an infinite tiling must always be duplicated within some fixed multiple γ of its radius. γ does not depend on the radius of the patch.*

This is a rather startling statement, and it is not obvious that it is true or that it is a consequence of **(EVC)**. We are saying that if we pick any patch of a tiling, then there will always be an *exact copy* of that patch nearby. An obvious corollary is that any patch must be duplicated an infinite number of times throughout the tiling. And yet we know the tiling is not self-similar under translation! We first prove the following Lemma:

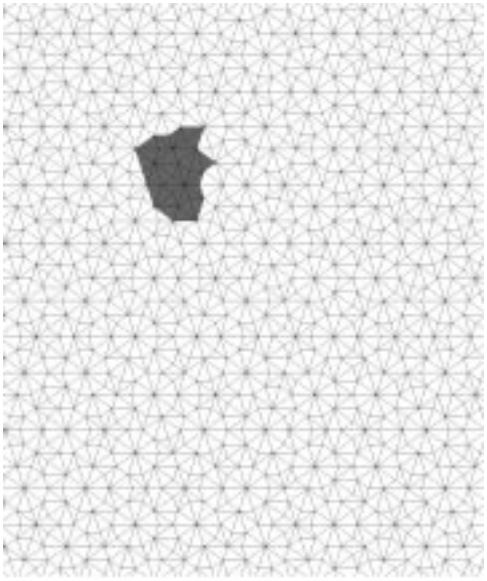

The shaded region is the patch $\mathcal{P}$, shown in a portion of an infinite Robinson tiling.

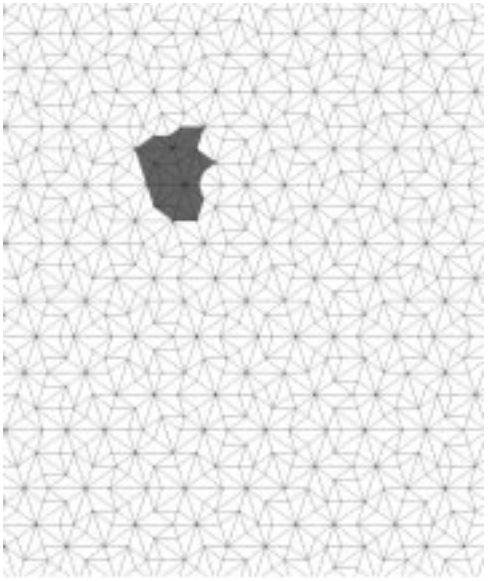

The same region $\mathcal{P}$ after the tiling has been deflated once.

Figure 2.5: Containing a patch in a vertex configuration, steps 1 & 2.

Lemma 2.2 *The image of any finite patch $\mathcal{P}$ in an infinite tiling $\mathcal{T}$ will be entirely contained in some vertex configuration when the tiling is deflated a sufficient number of times. The deflation scales the tiles of $\mathcal{T}$ by a factor proportional to the diameter of $\mathcal{P}$.*

A demonstration of this Lemma is given in Figures 2.5 – 2.7. We use a tiling by the Robinson triangles for this example.

Proof Consider a patch $\mathcal{P}$ with diameter d, contained within an infinite tiling $\mathcal{T}$. The intuition behind the following proof is that if we deflate $\mathcal{T}$ enough, the tiles will be so large that the image of $\mathcal{P}$ cannot possibly span a vertex configuration. It seems clear that if $\mathcal{P}$ is miniscule compared to tile tiles of $\mathcal{T}$, then it must lie within a vertex configuration. We wish to show that the tiles need only be expanded by a factor proportional to the diameter of the patch, however, so we will need a more detailed analysis.

We now consider the image of the patch $\mathcal{P}$ in some deflation $\text{defl}^k(\mathcal{T})$ of $\mathcal{T}$. Recall that a deflation increases the dimensions of the tiles by a factor of η. Thus the tiles in $\text{defl}^k(\mathcal{T})$ have been expanded by a factor of η^k with respect to those in $\mathcal{T}$. We claim that η^k is proportional to the diameter of $\mathcal{P}$.

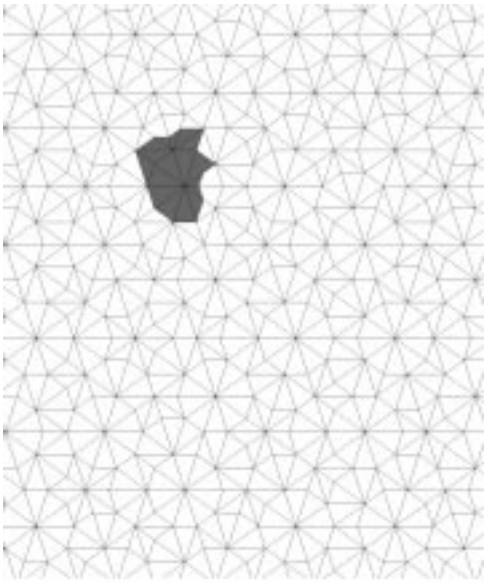

$\mathcal{P}$ after two deflations.

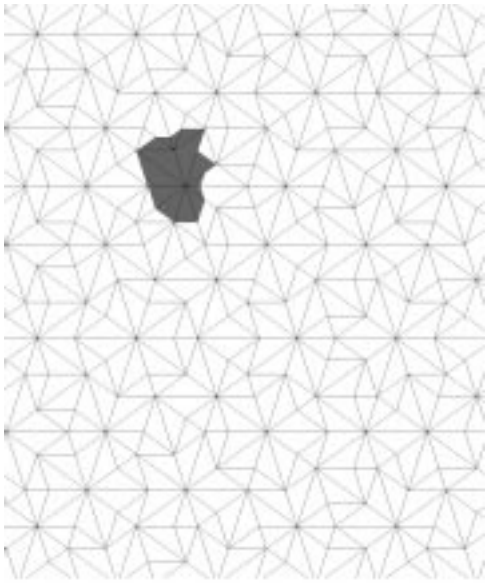

$\mathcal{P}$ after three deflations.

Figure 2.6: Containing a patch in a vertex configuration, steps 3 & 4.

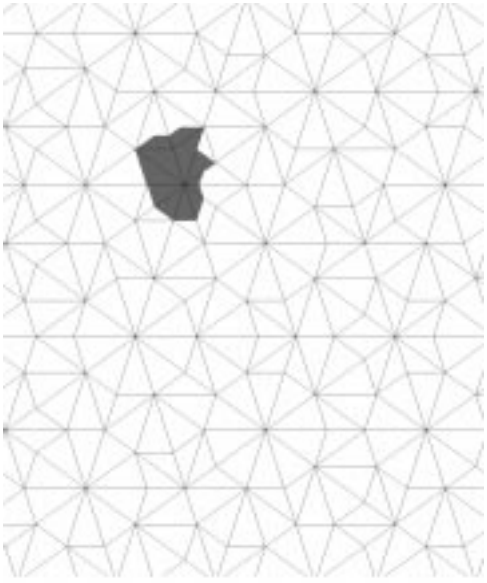

$\mathcal{P}$ after four deflations.

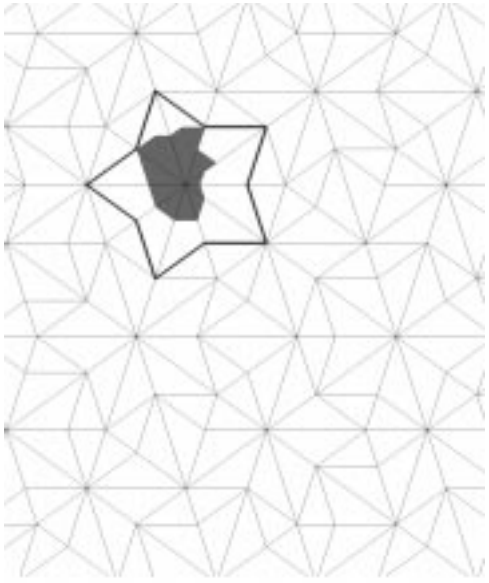

After five deflations, $\mathcal{P}$ is contained within the outlined vertex configuration.

Figure 2.7: Containing a patch in a vertex configuration, steps 5 & 6.

Let $\mathcal{I}$ be the set of tiles intersected by $\mathcal{P}$. If $\mathcal{I}$ consists of only one or two tiles, then clearly $\mathcal{P}$ lies within a vertex configuration. Thus we consider the case when $\mathcal{I}$ contains at least three tiles. Because it intersects more than one tile, $\mathcal{P}$ must cross an edge E between tiles. If $\mathcal{P}$ is contained within any vertex configuration, then it must be one of the configurations centered at an endpoint of E. Thus if $\mathcal{P}$ is not contained within a vertex configuration, then $\mathcal{I}$ must include some tile which does not have a vertex at either of the endpoints of E.

We say that edges which share an endpoint with E are *adjacent* to E. If $\mathcal{P}$ intersects a tile which does not touch an endpoint of E, then it must cross an edge which is not adjacent to E. Thus if the diameter of $\mathcal{P}$ is less than the shortest distance from any point on E to a non-adjacent edge, then $\mathcal{P}$ cannot possibly intersect a tile which does not have a vertex at an endpoint of E, and so $\mathcal{P}$ must lie within a vertex configuration.

Given a vertex configuration, we can compute the smallest distance from any point on an edge E to the nearest point on a non-adjacent edge. We will call this value the *span* of the vertex configuration. The *minimum vertex configuration span* of a tiling, denoted $V_{\min}$, is the smallest span of any vertex configuration of the tiling. $V_{\min}$ is a non-zero constant dependent on the shapes of the tiles. Since a deflation increases all dimensions of the tiles by a factor of η, $V_{\min}$ for $\text{defl}^k(\mathcal{T})$ is larger by a factor of η^k than $V_{\min}$ for $\mathcal{T}$. Thus given a tiling $\mathcal{T}$ with minimum vertex configuration span $V_{\min}$ and containing a patch $\mathcal{P}$, there exists a k such that $\eta^k V_{\min}$ is greater than the diameter of $\mathcal{P}$ by no more than a factor of η. Hence the image of $\mathcal{P}$ lies within a single vertex configuration in $\text{defl}^k(\mathcal{T})$, and the deflation expands $\mathcal{T}$ by a factor proportional to the diameter of $\mathcal{P}$. □

We make two notes concerning this Lemma. The first is that the granularity of the deflation implies that η^k may be greater than the diameter d of the patch by a factor of nearly η. Secondly, the span of a vertex configuration is easily computed if all the tiles are triangles (or higher-dimensional equivalents, such as tetrahedra). In this case, the span is equal to the shortest edge which has an endpoint at the central vertex of the configuration. Thus $V_{\min}$ for a set of triangular tiles is simply the length of the shortest edge of any tile.

Proof of Theorem 2.2 Armed with the Lemma, the proof of local isomorphism is fairly straightforward. The Lemma tells us that if we choose a patch $\mathcal{P}$ within tiling $\mathcal{T}$, then we can be sure that $\mathcal{P}$ is contained within a single vertex configuration C in the k^{th} deflation of $\mathcal{T}$ for some $k \geq 0$.

If we reinflate $\mathcal{T}$, of course, then the reinflation of C will contain $\mathcal{P}$. In fact, any vertex configuration identical to C will produce a patch identical to $\mathcal{P}$ when it is inflated k times. So if we show that there is guaranteed to be another copy of C

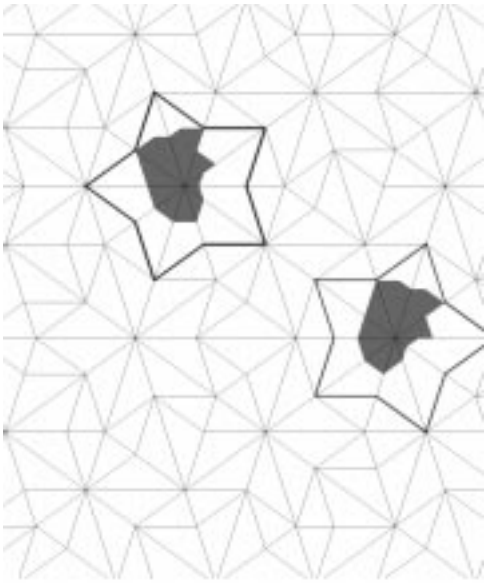

A nearby identical vertex configuration and the corresponding image of $\mathcal{P}$.

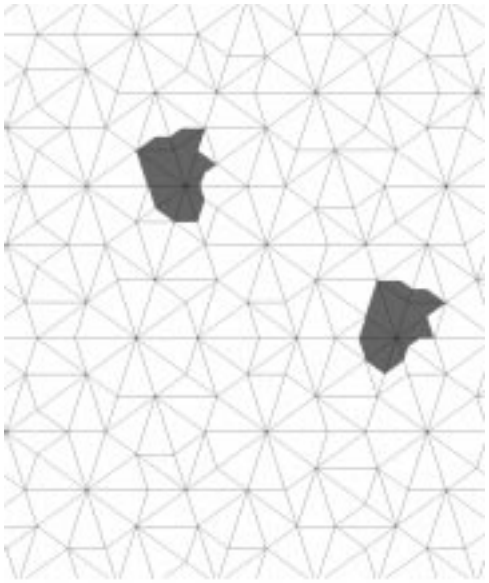

$\mathcal{P}$ and its copy after one inflation.

Figure 2.8: Finding a duplicate patch, steps 1 & 2.

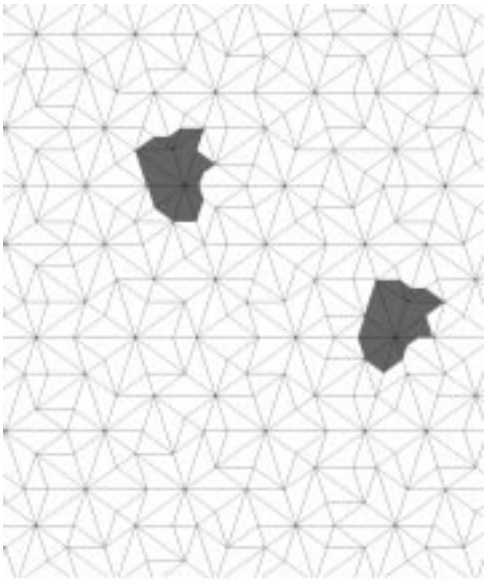

$\mathcal{P}$ and its copy after two inflations.

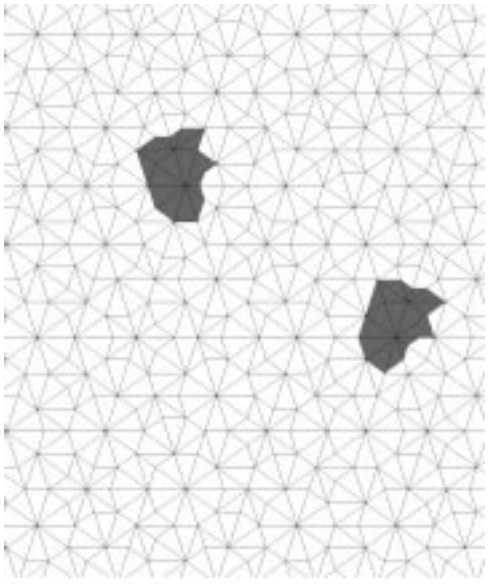

$\mathcal{P}$ and its copy after three inflations.

Figure 2.9: Finding a duplicate patch, steps 3 & 4.

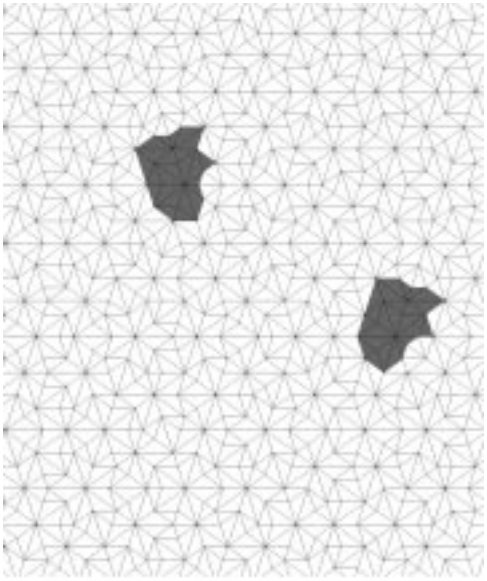

$\mathcal{P}$ and its copy after four inflations.

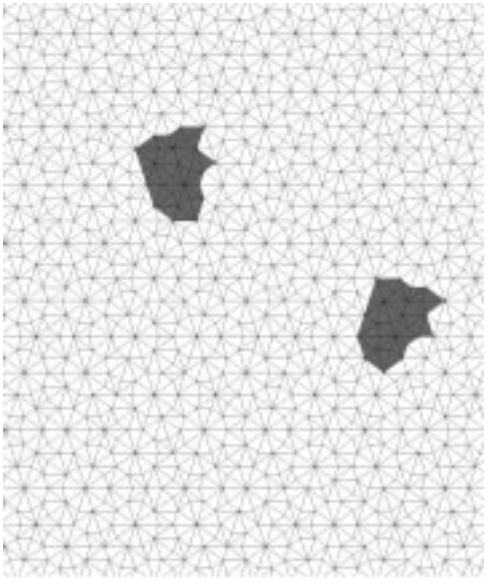

After five inflations, the original tiling is restored, and an exact copy of $\mathcal{P}$ located.

Figure 2.10: Finding a duplicate patch, steps 5 & 6.

nearby, then we know that reinflating $\mathcal{T}$ will produce another copy of $\mathcal{P}$.

Recall that in the discussion of requirement **(Min)** we showed that every prototile occurs an infinite number of times in an infinite tiling. **(EVC)** requires that every vertex configuration occur within H inflations of some tile. Together these facts imply that every vertex configuration occurs an infinite number of times in an infinite tiling. Thus there will always be another copy of C which, when inflated k times, will produce an identical copy of $\mathcal{P}$.

Here is a brief restatement of the method we can use to find a duplicate patch: Begin by picking the patch $\mathcal{P}$. Deflate the tiling k times until $\mathcal{P}$ lies entirely within a single vertex configuration C. This is shown in Figures 2.5 – 2.7. We know by **(EVC)** that there is a prototile S_n which, when inflated $h \leq H$ times, produces a copy of C. Deflate the tiling another h times, search for the nearest copy of S_n (which by Corollary 2.1 is a bounded distance away), and reinflate h times. We have found the nearest copy of C to the particular C that contains $\mathcal{P}$. Reinflating k times will reproduce the original tiling, and the reinflation of the copy of C will produce an exact copy of $\mathcal{P}$. The illustrations skip the process of finding the copy of C, but Figures 2.8 – 2.10 show the reinflation of the tiling and the location of the duplicate patch.

Hence every patch is duplicated elsewhere within the tiling. By the Lemma, the size of C depends of the size of the patch. Once $\mathcal{P}$ was contained within C, we deflated no more than H times, moved a bounded distance, and reinflated. Thus the distance between the two copies of C is no more than η^H times the maximum distance between identical tiles, a constant. When we reinflate k times, we will increase this distance by a factor proportional to the size of the patch. Therefore we have shown that any patch in an infinite tiling must be duplicated within the tiling, and that the maximum distance to the copy is proportional to the radius of the patch. □

Corollary 2.3 *Any patch found within an infinite tiling must occur an infinite number of times in* every *infinite* $(\mathcal{F}, \sim)$*-tiling.*

Proof In our construction, we chose the nearest copy of the prototile S_n to find the nearest copy of C and thus the nearest copy of $\mathcal{P}$, but clearly any instance of S_n would have done as well. Since we know from Section 2.2.3 that every prototile occurs an infinite number of times in every infinite $(\mathcal{F}, \sim)$-tiling, through the process above we can locate an infinite number of duplicates of $\mathcal{P}$. □

Corollary 2.4 *There is no algorithm which can decide in a finite number of steps whether two tilings (given as infinite sets of tiles) are equivalent.*

Proof Any algorithm for comparing tilings must compare a finite number of tiles in a finite number of steps. Thus the maximum distance between any two tiles considered by the algorithm is finite, and they all lie within some patch. By the previous Corollary we know that this patch is duplicated in every $(\mathcal{F}, \sim)$-tiling. Therefore equivalence of tilings is not decidable in finitely many steps. Of course, this applies only to tilings given as infinite sets of tiles. If we were given tilings as generating rules, for instance, equivalence might be easily checked. □

Chapter 3

Danzer's Tetrahedra

This chapter covers the second half of Danzer's paper. We give a family $\mathcal{F}_1$ of four prototiles and a matching relation $\sim_1$ developed by Danzer which meet the above requirements, and discuss some properties of the tiling they produce. As in the previous chapter, we augment the material from the original paper with explanations and proofs.

The first section defines the shapes of the tiles, some notation we will need in referring to them, and the matching relation $\sim_1$. Section 3.2 discusses the adjacencies of tiles allowed under $\sim_1$, knowledge which will prove useful in later sections.

Danzer presents a five-part main theorem concerning the properties of $(\mathcal{F}_1, \sim_1)$-tilings. This chapter will cover the first three parts of this main theorem; the remaining two are more conveniently discussed in the next chapter, since they are related to the atlas presented there. Section 3.3 proves the first part of Danzer's theorem, that the system $(\mathcal{F}_1, \sim_1)$ obeys the four requirements stated in the previous chapter. The final two sections of this chapter prove the second and third parts of Danzer's main theorem.

3.1 Construction

Table 3.1 presents the dimensions of the four tiles. For each edge, we give the dihedral angle between the faces that share that edge and its length. The symbols used in the lengths are $a = (\sqrt{10 + 2\sqrt{5}})/4$, $b = \sqrt{3}/2$, and $\tau = (1 + \sqrt{5})/2$, the golden ratio. Figure 3.1 gives cut-out patterns for the tiles. We also give the solid angle in square degrees subtended by each vertex of each tile in Table 3.2.

The four tiles given above we call (arbitrarily) *right-handed* to distinguish them from their mirror images, which are also necessary to construct a tiling. We will denote the right-handed tiles with capital A, B, C, and K, and the left-handed tiles

Tile	1–2	2–3	3–1	2–4	1–4	3–4
A	$36°, a$	$60°, \tau b$	$72°, \tau a$	$108°, a$	$90°, 1$	$60°, b$
B	$36°, a$	$36°, \tau a$	$60°, \tau b$	$120°, b$	$108°, a/\tau$	$90°, 1$
C	$36°, a/\tau$	$60°, \tau b$	$90°, \tau$	$120°, b$	$72°, a$	$36°, a$
K	$36°, a$	$60°, b$	$72°, a/\tau$	$90°, \tau/2$	$90°, 1/2$	$90°, 1/2\tau$

Table 3.1: Dihedral angles and side lengths for the four Danzer tetrahedra.

	Vertex			
Tile	1	2	3	4
A	18	24	12	78
B	24	12	6	138
C	18	36	6	48
K	18	6	42	90

Table 3.2: The solid angle in square degrees of each vertex of the four Danzer tetrahedra.

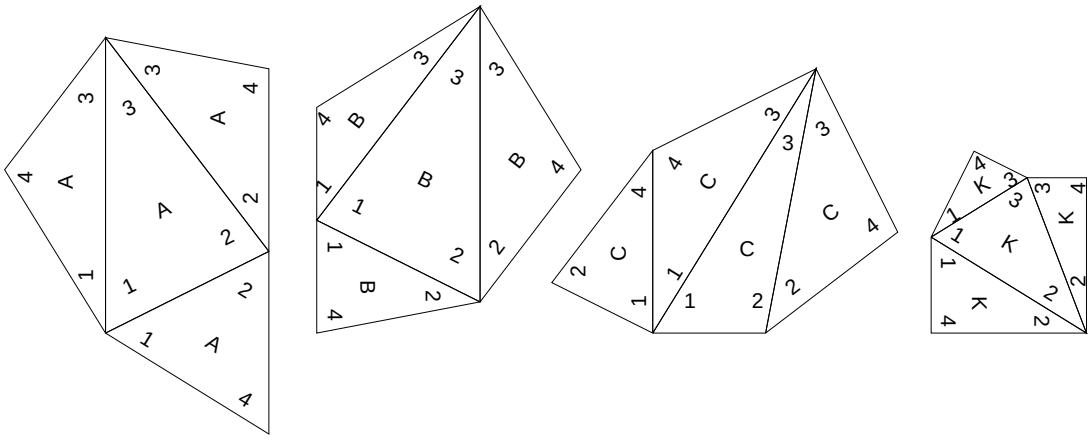

Figure 3.1: Templates for the four Danzer tiles. Left-handed versions can be made by simply folding the tiles so that the writing faces inward.

as lower-case a, b, c, and k. We will also need to refer to specific vertices, faces, and edges on each tile. We use Danzer's enumeration of the vertices and edges, shown in Table 3.1. Each face is given the number of the vertex opposite it. We use subscripts to refer to vertex numbers and normal-size numbers for faces. For example, B_2 refers to vertex 2 of a (right-handed) B tile, and $B2$ is the face opposite it. Edges are written as the name of the tile subscripted by the numbers of the two vertices at the ends of the edge. For instance, C_{24} has a dihedral angle of 120°. Finally, we will use superscripts to denote a level of inflation, so A^0 is simply an A tile, and A^1 is the 11-tile patch obtained by inflating an A tile once.

Now that we have a set of tiles, we can give an equivalence relation $\sim_1$ on them, which is defined by two rules. First, any face may match its mirror image. Second, faces of two different tiles (tiles which are not mirror images of each other) can match only if they are congruent and neither of them has an edge with dihedral angle 90° (a *red* edge). The A, B, and C tiles each have one red edge; the K tile has three. The red edges are A_{14}, B_{34}, C_{13}, K_{14}, K_{24}, and K_{34}. As we investigate the Danzer tiling in the coming sections, we will see how these matching rules enforce an aperiodic tiling.

3.2 Face Adjacencies

The matching relation $\sim_1$ defines 21 legal pairings of the faces of the tiles in $\mathcal{F}_1$. Since every face is allowed to be placed adjacent to its mirror image, sixteen of these are pairings between corresponding faces of opposite-handed tiles. The remaining five are the intertile pairings, which define the only ways in which different prototiles can meet.

We can deduce these pairings simply by inspection of the tiles. Since no face with a red edge is allowed to meet anything but its mirror image, the only faces at which intertile pairings could possibly occur are those without a red edge. These faces are $A1$, $A4$, $B3$, $B4$, $C1$, $C3$, and $K4$, and the equivalent faces of the left-handed tiles. The intertile adjacencies are all pairs of these faces which are congruent: $A1$-$C1$, $A4$-$B4$, $B3$-$C3$, $B3$-$K4$, and $c3$-$K4$. These five pairings, along with the sixteen mirror pairings, are all of the face adjacencies allowed under $\sim_1$. Of course the handedness of both tiles can always be reversed, so $A1$-$C1$ implies that $a1$ and $c1$ also match, and $c3$-$K4$ implies $C3$-$k4$.

In fact, not all of these face adjacencies can occur in an infinite tiling. $A1$-$a1$, $B3$-$b3$, and $B3$-$C3$ cannot appear in an infinite tiling because tiles cannot be placed around them without violating the matching rules. The easiest way to see this is to make models of the tiles from the patterns in Figure 3.1 and try to build around these pairings, but we give brief explanations of why each of them is impossible here.

	A	B	C	K	a	b	c	k
$\tau^{-1}A$	0	0	0	0	0	0	1	0
$\tau^{-1}B$	2	1	0	1	1	1	0	0
$\tau^{-1}C$	1	1	1	0	1	0	1	0
$\tau^{-1}K$	3	2	1	1	3	2	1	0
$\tau^{-1}a$	0	0	1	0	0	0	0	0
$\tau^{-1}b$	1	1	0	0	2	1	0	1
$\tau^{-1}c$	1	0	1	0	1	1	1	0
$\tau^{-1}k$	3	2	1	0	3	2	1	1

Table 3.3: The inflation matrix for the Danzer tiles.

The $A1$-$a1$ pairing leaves an $A3$ face and its mirror image sharing an edge (the A_{24} edge). Since $A3$ can only match $a3$, two more A tiles must be placed here, but they would overlap. The argument is similar for the $B3$-$C3$ pairing. If we construct the $B3$-$C3$ pairing and place face 1 of a b tile against $B1$ (the only neighbor allowed for that face), then $b3$ is partially adjacent to $C1$, which violates the face-to-face requirement. Finally, $B3$-$b3$ presents the same problem on two sides. Both the $B1$ and $B2$ faces must match their inverses, and they cannot do so without causing tiles to overlap.

3.3 Meeting the Requirements

This section concerns the first part of Danzer's main theorem, which states that the system $(\mathcal{F}_1, \sim_1)$ obeys the four requirements **(Infl)**, **(Defl)**, **(Min)**, and **(EVC)**. We will prove that the system satisfies the first three of these. We will leave the proof of **(EVC)** until the next chapter, where we discuss vertex configurations in depth.

3.3.1 $\mathcal{F}_1$ satisfies (Infl) and (Defl) with $\eta = \tau$

We can show by construction that this system obeys the inflation and deflation requirements **(Infl)** and **(Defl)**. The inflation matrix given in Table 3.3 gives the components of the inflation of each tile. For instance, an A tile can be constructed from two (right-handed) B tiles scaled by τ^{-1} and one (left-handed) copy of $\tau^{-1}b$, one copy of $\tau^{-1}C$ and one $\tau^{-1}c$, three copies of $\tau^{-1}K$, and three $\tau^{-1}k$. The construction of an a tile is identical with the handedness of all the components reversed. Danzer presents the inflation matrix for right-handed tiles only, and adds the number of left- and right-handed tiles together. In many cases it is convenient not to distinguish between a tile and its mirror image, but in this case the difference is

	Vertex			
Tile	1	2	3	4
A	B_3, b_3, B_3	B_2, K_2, k_2	C_3, c_3	C_2, k_3
B	b_2, K_2, k_2	B_3, b_3	C_3	$B_1, b_1, C_1, 2 \cdot K_1, 2 \cdot k_1$
C	a_1	a_2, C_3, c_3	K_2	a_3, c_2
K	B_2, K_2	B_3	B_1, K_1	K_4

Table 3.4: The effect of an inflation on the vertices of each tile. The subscript is the vertex number, a coefficient is the number of copies. Thus $2 \cdot K_1$ means two instances of vertex 1 of a K tile.

significant. Inflation is not a symmetric operation; for example, A^1 contains two B tiles and one b tile. For this reason we have expanded the inflation matrix to differentiate between left- and right-handed components.

It is also useful to know the results of inflation on the vertex of a tile. For example, when an A tile is inflated, vertex 2 of a C tile and vertex 3 of a k tile meet at the point that was vertex 4 of the original A tile. Table 3.4 shows the effect of an inflation on each vertex of each of the right-handed tiles. The vertex inflations of the left-handed tiles are the same with all the components reversed.

Figures 3.2 - 3.5 show the inflations of the right-handed tiles.

3.3.2 $\mathcal{F}_1$ satisfies (Min) with $J = 4$

As discussed in Section 2.2.3, we can show that $\mathcal{F}_1$ is minimal if some power of its inflation matrix has no zero entries. Table 3.5 shows the first four powers of the inflation matrix $\mathcal{M}$.

Because Danzer does not distinguish between left- and right-handed tiles (he adds the top and bottom halves of the inflation matrices), he considers $\mathcal{M}^3$ to have no zero entries, and so he states that $J = 3$. The third inflation of a K tile, however, contains an A tile but no a tile. There are no zero entries in $\mathcal{M}^4$, and so any tile, when inflated four times, will produce a patch containing all *eight* prototiles. Thus we define $J = 4$ for the Danzer tiling. By the results from Section 2.2.3, we know that this implies that all prototiles are used in any infinite tiling, that they appear in roughly equal numbers, and that the distance between tiles of the same type is bounded.

By raising the inflation matrix $\mathcal{M}$ to large powers, we can obtain an estimate of the relative frequency of each type of tile in an infinite tiling (see Corollary 2.2). We can also compute the *density* of each prototile, the portion of the total volume of the tiling occupied by tiles of that type. Figure 3.6 shows for each prototile its volume, the fraction of the tiles in an infinite tiling which are of its type, and its

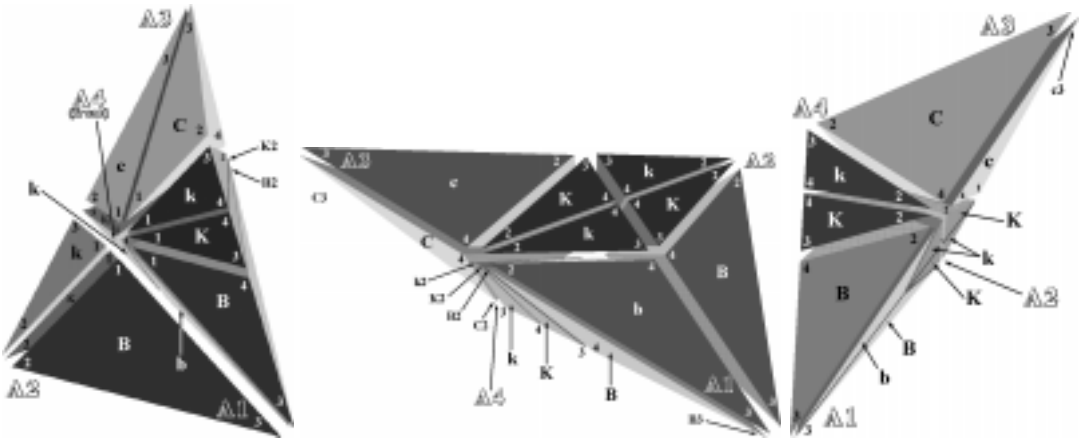

Figure 3.2: Three views of the inflation of an A tile. Outlined labels indicate the vertices of the original tile.

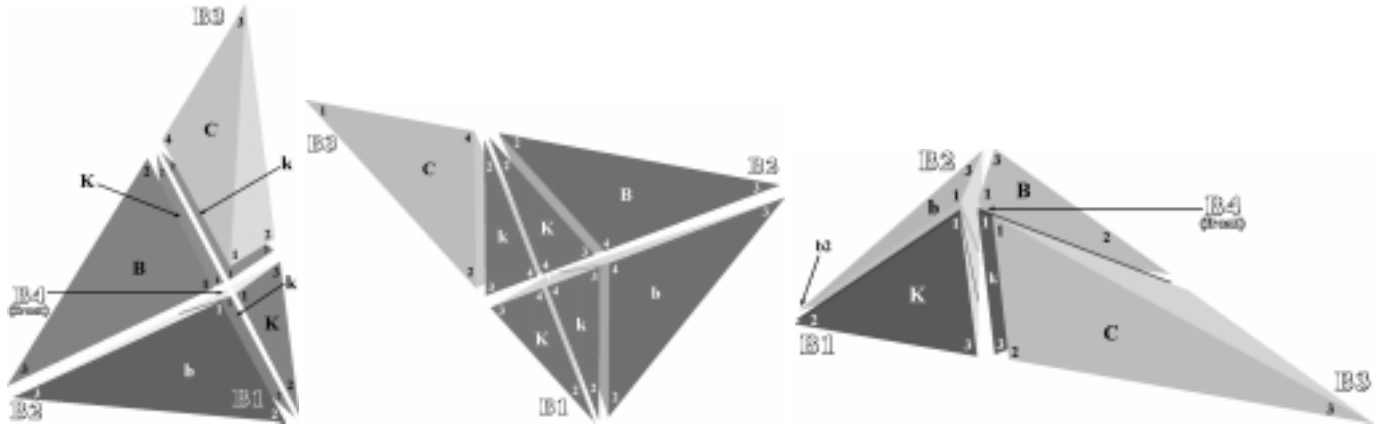

Figure 3.3: Three views of the inflation of a B tile.

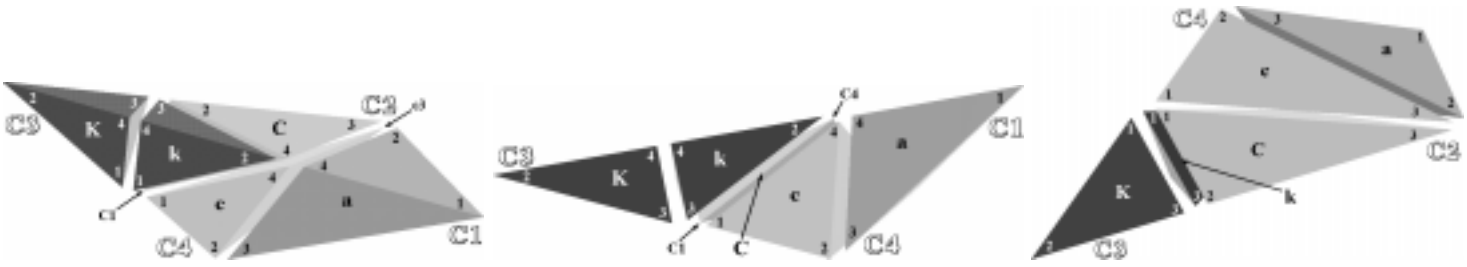

Figure 3.4: Three views of the inflation of a C tile.

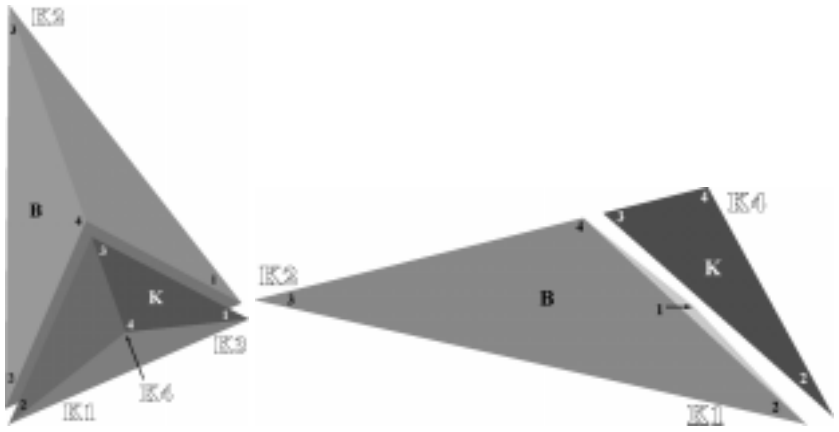

Figure 3.5: Two views of the inflation of a K tile.

$\mathcal{M}$	A	B	C	K	a	b	c	k
$\tau^{-1}A$	0	0	0	0	0	0	1	0
$\tau^{-1}B$	2	1	0	1	1	1	0	0
$\tau^{-1}C$	1	1	1	0	1	0	1	0
$\tau^{-1}K$	3	2	1	1	3	2	1	0
$\tau^{-1}a$	0	0	1	0	0	0	0	0
$\tau^{-1}b$	1	1	0	0	2	1	0	1
$\tau^{-1}c$	1	0	1	0	1	1	1	0
$\tau^{-1}k$	3	2	1	0	3	2	1	1

$\mathcal{M}^2$	A	B	C	K	a	b	c	k
$\tau^{-1}A$	1	0	1	0	1	1	1	0
$\tau^{-1}B$	6	4	2	2	6	4	3	1
$\tau^{-1}C$	4	2	3	1	3	2	3	0
$\tau^{-1}K$	11	7	6	3	11	7	6	2
$\tau^{-1}a$	1	1	1	0	1	0	1	0
$\tau^{-1}b$	6	4	3	1	6	4	2	2
$\tau^{-1}c$	3	2	3	0	4	2	3	1
$\tau^{-1}k$	11	7	6	2	11	7	6	3

$\mathcal{M}^3$	A	B	C	K	a	b	c	k
$\tau^{-1}A$	3	2	3	0	4	2	3	1
$\tau^{-1}B$	26	16	14	6	26	17	14	5
$\tau^{-1}C$	15	9	10	3	15	9	11	2
$\tau^{-1}K$	48	30	28	10	48	30	28	9
$\tau^{-1}a$	4	2	3	1	3	2	3	0
$\tau^{-1}b$	26	17	14	5	26	16	14	6
$\tau^{-1}c$	15	9	11	2	15	9	10	3
$\tau^{-1}k$	48	30	28	9	48	30	28	10

$\mathcal{M}^4$	A	B	C	K	a	b	c	k
$\tau^{-1}A$	15	9	11	2	15	9	10	3
$\tau^{-1}B$	110	69	65	22	111	69	65	22
$\tau^{-1}C$	63	38	41	12	63	39	41	11
$\tau^{-1}K$	203	126	123	40	203	126	123	39
$\tau^{-1}a$	15	9	10	3	15	9	11	2
$\tau^{-1}b$	111	69	65	22	110	69	65	22
$\tau^{-1}c$	63	39	41	11	63	38	41	12
$\tau^{-1}k$	203	126	123	39	203	126	123	40

Table 3.5: Inflation matrices for the first four levels of inflation. The fourth level contains no zero entries, so $\mathcal{F}_1$ is minimal.

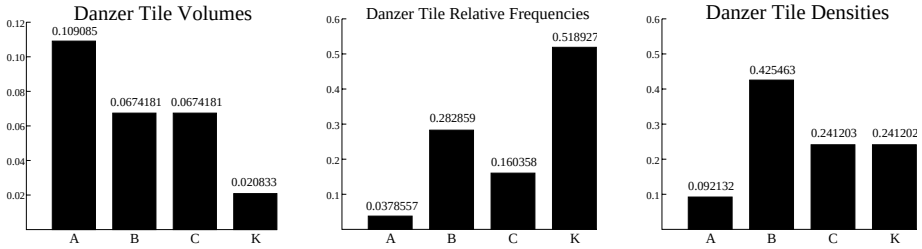

Figure 3.6: Approximate volumes, frequencies, and densities of the Danzer tiles.

Figure 3.7: Three views of the A-octahedron, with the tiles individually shrunk by 30% for clarity. The central edge is A_{14}.

density. These values were measured in A^{11}, and are accurate to ± 0.000001. We will see in the next section that in an infinite tiling each prototile and its mirror image occur in equal numbers, and so their densities and frequencies are the same. Thus we present these values for left- and right-handed tiles combined.

3.4 Octahedra

Theorem 3.1 *In every $(\mathcal{F}_1, \sim_1)$-tiling tiles of the same type cluster to form octahedra. The K-octahedron is formed of eight K tiles (4 of each handedness); the A-, B-, and C-octahedra consist of 4 tiles each.*

Proof We can demonstrate the existence of the octahedra using the face adjacencies discussed in Section 3.2. For example, consider an A tile. The A_{14} edge is a red edge (it has a dihedral angle of 90°), and so the two faces which share it, $A2$ and $A3$, may only meet their mirror images. Thus every A tile must be adjacent to two a tiles, one on each of these faces. The $a2$ face of one of these a tiles is exposed,

Color	Members
Red	A_1, A_4, B_2, C_4, K_2
Green	A_3, B_3, B_4, C_2, K_3
Blue	A_2, B_1, C_1, C_3, K_1
Yellow	K_4

Table 3.6: The four vertex classes. Only vertices of the same class can meet in a tiling. The colors were chosen to correspond to the colors of the tiles in the pictures shown in the appendices.

as is the $a3$ face of the other. These must be matched by an A tile, completing the octahedron shown in Figure 3.7. Every A tile must be a part of such an octahedron.

The argument is similar for the B- and C-octahedra. The K tile, unlike the other three tiles, tile has three red edges, which force an 8-tile octahedron. The K-octahedron is a simple, convex diamond shape with three orthogonal planes of symmetry. The other three octahedra have two orthogonal planes of symmetry. The A- and B-octahedra are concave, and the C-octahedron is convex. The K-octahedron is the only one of the four to contain an interior vertex, and thus it forms a vertex configuration. Vertex configurations will be discussed in more detail in the next chapter.

Corollary 3.1 *The number of left- and right-handed versions of each prototile is equal in an infinite tiling.*

All four octahedra contain equal numbers of left- and right-handed tiles and every tile in an infinite tiling must be part of an octahedron. Thus in an infinite tiling, there are an exactly equal number of instances of each prototile and its mirror image.

3.5 Vertex Classes

Theorem 3.2 *The vertices of the prototiles fall into four classes (or colors) such that at every vertex in an $(\mathcal{F}_1, \sim_1)$-tiling only vertices of the same color can meet. (See Table 3.6.)*

Proof We can extract the vertex classes from the face adjacencies given in Section 3.2. We begin constructing a class with one vertex v, then add all the vertices that can legally be placed adjacent to v, and so on to produce the transitive closure

of the class. If there are any vertices not already classified, we pick one arbitrarily as the starting point for the next class.

Let us start with B_3, and arbitrarily assign it the color green. Of course we can pair any face of the B tile with its mirror image, so b_3 is the same color as its right-handed counterpart. The color of corresponding opposite-handed vertices is always the same. If we place an A tile adjacent to a B tile so that $A4$ matches $B4$, then vertex A_3 meets vertex B_3. Thus A_3 is also green. The $A1 - C1$ pairing places C_2 next to A_3, so C_2 is green. $C3 - k4$ pairs k_3 and C_2, so k_3 is green. Finally, the $B3 - K3$ pairing adds B_4 to the list. We have accounted for all possible pairings of the faces which contain a green vertex, and so the list is complete.

We will spare the reader the derivation of the remaining two classes, but the process is the same. The vertex K_4 is at the center of the K-octahedron, and only ever meets itself. Table 3.6 lists the members of each of the four classes.

Chapter 4

An Atlas of the Danzer Tiling

In the previous chapter we discussed the Danzer tiles and some of the small-scale structure of a Danzer tiling, such as face and vertex adjacencies and the way in which the tiles always form octahedra. In this chapter we begin to investigate the next level of structure, the vertex configuration. A complete knowledge of the legal vertex configurations is useful for several reasons. In the Penrose tiling, the vertex configurations are important in understanding the ways in which tiles are forced, and we believe that the same is true of the Danzer tiling. In addition, if an aperiodic tiling is used as a model of a quasicrystal, then the vertices represent locations of atoms and the vertex configurations are the possible arrangements of neighboring atoms. This chapter begins the construction of an atlas of the vertex configurations of the Danzer tiling. We will show that a Danzer tiling meets the final requirement **(EVC)**, and discuss the final two pieces of Danzer's main theorem:

(Icos) There are exactly three infinite $(\mathcal{F}_1, \sim_1)$-tilings with full icosahedral symmetry, and they permute cyclically under inflation and subsequent expansion.

(Vert) There are exactly 27 vertex configurations, of which 22 occur in an infinite tiling. Except for the vertex at the center of the K-octahedron, every configuration becomes one of the three with full icosahedral group after at most six inflations.

We will begin by showing that we can find all the vertex configurations that occur in an infinite tiling by inspecting a finite patch. We will show that there are exactly 22 such configurations, confirming Danzer's result. After discussing some of the characteristics of these 22 configurations, we will be able show part **(Icos)** above and that the Danzer tiling satisfies the local isomorphism requirement **(EVC)**. As we go, we will dispute or refine several of Danzer's results. Danzer claims that the

constants associated with **(EVC)** have values $H = 6$ and $\gamma = 40$ (see Section 2.2.4 for definitions). We will show that $H = 7$, and we will examine in detail the distances between identical vertex configurations and arbitrary patches, enabling us to place a tighter bound on the value of γ. Part **(Vert)** above states that all vertex configurations inflate to one of the three "balls" in no more than six steps. This is true, but we will show that in fact no more than four inflations are required.

The remaining vertex configurations mentioned in part **(Vert)**, which are locally legal but cannot appear in an infinite tiling, will be discussed in the next chapter.

4.1 Finding the Global Configurations

The first step is to find the vertex configurations which can occur in an infinite tiling, which we will refer to as the *global* vertex configurations. We will call the complete listing of global vertex configurations the *global atlas*, or just the atlas.

We will assume for the moment that the Danzer tiles satisfy **(EVC)**, which requires that every global vertex configuration be produced by no more than H inflations of some prototile. This assumption will be justified in Section 4.4. Given this fact, we can find all the global vertex configurations simply by inspecting the patches produced by inflating each prototile H times. In fact, since **(Min)** guarantees that inflating any prototile J times produces a patch containing all prototiles, we need only consider the inflations of a single tile. Thus if we inflate any tile no more than $H + J$ times, pausing after each inflation to inspect the patch and record any new configurations, we will obtain a complete atlas of the global configurations.

The only remaining problem is to determine the value of H: how do we know when our atlas is complete? This question is addressed by the following Theorem:

Theorem 4.1 *If an inflation generates no new vertex configurations, then no subsequent inflation will ever produce a new vertex configuration.*

Proof Consider the three localities in which a new vertex configuration can occur as the result of the inflation of a finite patch $\mathcal{P}$: at a vertex, along an edge, or on a face of some tile in $\mathcal{P}$. Suppose $\mathcal{P}$ contains a vertex configuration C centered at vertex v. We first note that v must also be a vertex in $\text{infl}(\mathcal{P})$. The inflation may have changed the tiles which share v, however, producing a configuration C' centered at v. In this case we say that configuration C inflates to configuration C'.

This first case essentially permutes the configurations surrounding existing vertices. Inflation also produces new vertices. The prototile inflations presented in Section 3.3.1 show that the first inflation of a prototile never contains an interior vertex. Thus inflating $\mathcal{P}$ will produce new vertices only along the edges and on

the faces of tiles in $\mathcal{P}$. Of course, no vertex configuration can be produced on the boundary of the patch, so every edge on which a vertex configuration is produced will be shared by several tiles, and every such face will be adjacent to another face. We call an interior edge of a patch an *axis*, and a set of tiles all of which share two vertices and the edge between them an *axis configuration*. We give vertex configurations, axis configurations, and face pairings the collective label *interfaces*. All vertex configurations in infl($\mathcal{P}$) must lie on one of the interfaces of $\mathcal{P}$.

Consider the set $\mathcal{I}$ of distinct interfaces in $\mathcal{P}$. Inflating $\mathcal{P}$ produces a new set $\mathcal{I}'$ of interfaces. By the argument above, every vertex configuration in $\mathcal{I}'$ is the result of inflating some interface in $\mathcal{I}$. If $\mathcal{I}' = \mathcal{I}$ then the inflation produced no new interfaces of any type. Thus the set $\mathcal{I}$ of interfaces is closed, and no subsequent inflation will produce any new interfaces.

Note that any two tiles which share a vertex are, by definition, members of the same vertex configuration. Thus for each axis configuration X in $\mathcal{P}$, there is a vertex configuration C (also in $\mathcal{P}$) which contains X. Similarly, every face pairing in the patch occurs as part of a vertex configuration. Thus a description of all the vertex configurations of $\mathcal{P}$ includes all the axis configurations and face pairings in $\mathcal{P}$. Hence when computing the set $\mathcal{I}$ we need not consider all interfaces explicitly. It is sufficient to consider only the vertex configurations of $\mathcal{P}$. If an inflation produces no new vertex configurations, then it produces no new interfaces of any type, and thus no subsequent inflation will produce any new vertex configurations. □

4.2 The 22 Global Configurations

In the previous section we showed that we can obtain a complete list of global vertex configurations by inflating a tile repeatedly, inspecting the resulting patch for new configurations after each inflation and halting when some inflation produces no new configurations. In this section we present the global vertex configurations of the Danzer tiling generated by this process. Table 4.2 lists the basic characteristics of each of the global configurations, and pictures can be found in Appendix A. The designation is the "formal" name of the configuration, consisting of the number of tiles in the configuration and possibly a letter to ensure uniqueness. The name is just for convenience in referring to some of the more visually distinctive configurations. The Composition columns give the number of each type of tile used in the configuration. We give the total number of tiles of each type without regard to handedness since every global configuration contains an equal number of left- and right-handed tiles. For instance, the Bishop (36) contains six B tiles and six b tiles.

The Tile/Level column gives the patch in which the configuration was found. Because the A tile has the most complex inflation, in most cases the earliest oc-

		Composition					Location		
Designation	**Name**	**A**	**B**	**C**	**K**	**Tile/Level**	**X**	**Y**	**Z**
8a	K-octahedron				8	A^2	0.425325	0.375123	0.154508
8b	Sailboat		4		4	A^2	0.425325	0.48738	0.118034
16	Pawn	4		8	4	A^4	0.224514	0.224514	0.072949
18				6	12	A^4	0.32492	0.286568	0.118034
28		8	4	12	4	A^5	0.224514	0.224514	0.072949
36	Bishop		12	4	20	A^3	0.425325	0.48738	0.118034
38	Claw		6	12	20	A^4	0.363271	0.224514	0.072949
40a	Diamond		10	10	20	A^3	0.525731	0.688191	0.118034
40b				20	20	A^4	0.48738	0.750245	0.072949
42	Obelisk	6	12	12	12	A^5	0.32492	0.286568	0.118034
60a		10	40	10		A^4	0.525731	0.688191	0.118034
60b		20	20		20	A^5	0.48738	0.750245	0.072949
64	Taco	12	4	44	4	A^6	0.224514	0.224514	0.072949
70		12	26		32	A^5	0.363271	0.224514	0.072949
80		4	32		44	A^4	0.425325	0.48738	0.118034
90a	Half-discus	10		80		A^5	0.525731	0.688191	0.118034
90b			30		60	A^6	0.32492	0.286568	0.118034
100	Princess		20		80	A^7	0.224514	0.224514	0.072949
110	Cone		10		100	A^6	0.525731	0.688191	0.118034
120a	B-ball		120			A^5	0.425325	0.48738	0.118034
120b	C-ball			120		A^6	0.425325	0.48738	0.118034
120c	K-ball				120	A^7	0.425325	0.48738	0.118034

Table 4.1: Basic characteristics of the 22 global vertex configurations.

currence of a configuration is in some inflation of an A tile. The exceptions are configurations 18, 36, 64, and 100, each of which can be found one level earlier in the inflation of a C tile. For simplicity we always use an A tile as our starting point. Finally, the three rightmost Location columns give the coordinates of the central vertex of an instance of each configuration within the patch indicated in the Tile/Level column. The orientation of the original tile is as follows: vertex 1 of the original A tile is placed at the origin, vertex 2 is on the positive x-axis, and vertex 3 is in the x-y plane. The dimensions of the starting tile are as given in Table 3.1, and no scaling was performed between inflations. When several instances of a new configuration appear simultaneously, one has been chosen arbitrarily.

Configuration 120c, the K-ball, does not appear until the 7^{th} inflation of an A tile, and not until the 8^{th} inflation of any of the other three tiles. All other configurations can be found in less than 7 inflations of some tile, and in 8 or less inflations of any tile. Thus H, the maximum number of inflations absolutely necessary to produce any vertex configuration, is 7 for this set of tiles. This contradicts Danzer's statement in the first part of his main theorem that this set has $H = 6$.

All the global vertex configurations have at least one plane of mirror symmetry. Thus none of them are chiral; that is, each global configuration is congruent to its mirror image. There are three 120-tile balls, 120a, 120b, and 120c, composed exclusively of B tiles, C tiles, and K tiles respectively. These three configurations have all the symmetries of the icosahedral group. (An icosahedron has 6 axes of fivefold rotational symmetry passing through opposite vertices, 10 axes of threefold rotational symmetry passing through the centers of opposite faces, and 15 axes of twofold rotation bisecting opposite edges.) Infinitely inflating one of these balls produces an infinite tiling with full icosahedral symmetry. The three infinite tilings with full icosahedral symmetry mentioned in part **(Icos)** of Danzer's theorem are the three tilings generated by inflating each of the balls infinitely. The unique deflation property shows that any tiling with such symmetry must be equivalent to one of these three.

4.3 Inflating a Vertex Configuration

Inflating a patch may create new internal vertices, but, because the inflation of each tile is self-contained, no vertices ever disappear as a result of an inflation. That is, any point v which is a vertex in $\mathcal{P}$ will be a vertex in $\text{infl}^k(\mathcal{P})$, $k \geq 0$. The number and types of tiles which share v, however, will change under inflation. Figure 4.1 shows a graph of the effect of inflating the 22 global vertex configurations. The three balls (120a, 120b, and 120c) cycle under inflation. All other vertex configurations inflate to the B-ball after at most four steps, except for 8a (the center of the K-

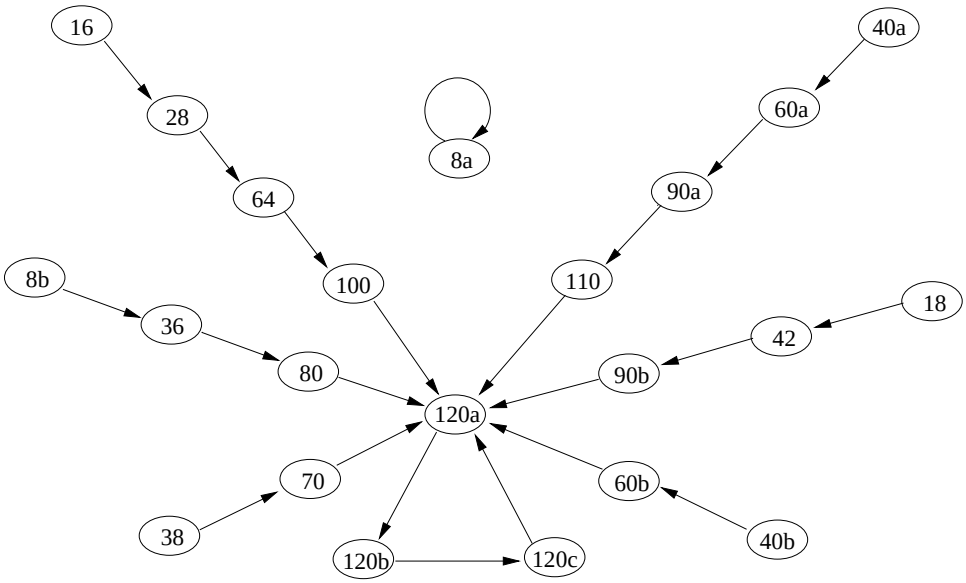

Figure 4.1: The effect of inflation on the 22 global vertex configurations.

octahedron), which is unchanged by inflation. Note that part **(Icos)** of Danzer's main theorem says that all configurations but 8a inflate to one of the three balls after at most six inflations, but in fact four inflations are sufficent and the entry point into the three-ball cycle is always the B-ball.

4.4 $\mathcal{F}_1$ satisfies (EVC)

In Section 4.1 we assumed that the Danzer tiles satisfy **(EVC)**, the requirement that all global vertex configurations be produced by inflating some prototile no more than H times. In this section we will describe a constructive method by which we can show that **(EVC)** is met. The construction is the topic of the next chapter.

The closed-set argument of Section 4.1 shows that there are infinite tilings which contain only the 22 configurations listed above. This does not guarantee that all infinite tilings contain only these configurations, however. One can imagine several ways in which the inflation graph shown in Figure 4.1 could be incomplete. Perhaps there is a single-vertex cycle like 8a, which, though not generated in any inflation of a single tile, would remain fixed under inflation if constructed. There could even be an "alternate universe," a completely separate set of vertex configurations which are closed under inflation.

To show that the Danzer tiles satisfy **(EVC)**, we must show that inflating any

initial patch $\mathcal{P}$ a finite number of times produces a patch containing only the 22 global configurations above. As discussed in Section 4.1, it is sufficient to consider the case of inflating vertex configurations. All vertex configurations produced by inflating $\mathcal{P}$ wil occur on an interface of $\mathcal{P}$, and the vertex configurations of $\mathcal{P}$ completely describe the other interfaces of $\mathcal{P}$.

Thus we wish to show that inflating *any* legal vertex configuration a finite number of times produces a patch containing only the 22 global configurations. The next chapter presents an algorithm for constructing all legal vertex configurations, and shows that all of them inflate to one of the three balls within four steps. Thus all infinite tilings consist of only the 22 global configurations listed above, all of which are found within seven inflations of a prototile, and so the Danzer tiles satisfy **(EVC)**.

4.5 Distance Between Identical Patches

In Section 2.2.4 we also showed the rather surprising consequence of **(EVC)**, local isomorphism. Any patch in an infinite Danzer tiling must be duplicated within a fixed multiple γ of the radius of the patch. We wish to determine an upper bound for γ. Using a computational approach, we construct a bound on γ and give a method for tightening the bound. Throughout this section, when we discuss the distance between identical patches we refer to the maximum distance from any instance of the patch to its *nearest* identical copy.

Recall from Section 2.2.4 that any patch is contained in a single vertex configuration after a sufficient number of deflations. Thus the distance between identical patches is related to the distance between identical vertex configurations and the orientation of the patch within the containing configuration. Our first step will be to determine the maximum distance between vertex configurations of the same type.

Distance Between Identical Vertex Configurations Rather than derive the distance between vertex configurations by studying the inflation mechanism, we will show that there are a finite number of relative orientations of two identical configurations, and that we can construct and measure all of them. The local isomorphism property, expressed in Theorem 2.2, tells us that there is an upper bound on the distance from any vertex configuration to its nearest identical copy. We can view any pair of same-type configurations as demarcating the extent of a patch, which is contained in a larger vertex configuration at some level of deflation. Thus every pair of identical vertex configurations is contained in a patch generated by inflating a single vertex configuration.

These facts imply that we can construct a patch containing two maximally separated vertex configurations by inflating a single vertex configuration a sufficient number of times. Since inflating any tile eight times produces all vertex configurations, inflating a vertex configuration C eight times produces all the ways in which all configurations can be positioned within the inflation of C. A configuration found in C^9 can also be found in the eighth inflation of some configuration in C^1. Thus we can construct all possible relative orientations of vertex configurations by inflating each of the 22 global configurations eight times.

Our algorithm to find the maximum distance between identical vertex configurations proceeds as follows. We inflate a vertex configuration eight times and locate all vertex configurations within the resulting patch. For each of the 22 configuration types, we compile a list of the locations of all its occurrences. We then apply an all-points-nearest-neighbor algorithm to each list. The distance from one vertex configuration to another is measured between their central vertices.

We would like to take the maximum of the resulting distances, but first we must account for boundary conditions. If a point is close to the outer boundary of the patch then it is possible that, were the patch part of an infinite tiling, the point would have a neighbor outside the patch closer than any neighbor inside the patch. This boundary effect could produce artificially large nearest-neighbor distances from isolated instances of a configuration in a corner of the patch. To correct for this effect, we discard any point which is closer to the boundary of the patch than it is to its nearest neighbor within the patch. We lose no information by discarding these points, since any point near the boundary of one vertex configuration is near the center of another, and so will be included in the processing of some other configuration.

We compute the maximum over the nearest-neighbor distance for the points not discarded in the boundary check. We repeat the above algorithm for each of the global configurations. At this point we have 22^2 values: for each type of configuration we have 22 maxima, one from each of the 22 patches we examined. Finally, for each configuration C we take the max over its associated distances to obtain the maximum distance between instances of C in the eighth inflation of any configuration, which is the maximum in any tiling by the argument above. Table 4.2 gives thes results of this algorithm. For each of the global configurations we give its radius and the maximum distance to an identical copy, measured in radii.

These measurements are approximations, but in the case of the K-ball the distance is exactly τ^5. The radius of the K-ball is exactly a, so the maximum distance between K-balls is $\tau^5/a \approx 11.66$ radii, and this is the maximum for any vertex configuration.

Configuration	Radius	Span	Distance
8a	0.809017	$1/2\tau$	1.000
8b	0.809017	$1/2\tau$	0.764
16	0.938424	$\tau/2$	1.066
18	1.13364	$1/2\tau$	0.882
28	1.30902	$1/2\tau$	1.236
36	1.30902	$1/2$	0.764
38	1.48473	$1/2$	0.673
40a	1.24495	$\tau/2$	0.803
40b	1.30902	$1/2$	4.000
42	1.48473	$1/2$	1.090
60a	1.39272	a/τ	1.880
60b	1.30902	$\tau/2$	4.000
64	1.61804	$1/2$	1.618
70	1.30902	$\tau/2$	1.236
80	1.30902	$\tau/2$	1.236
90a	1.61804	a	2.618
90b	1.30902	$\tau/2$	2.000
100	1.30902	$\tau/2$	3.236
110	1.24495	$\tau/2$	3.403
120a	1.53884	1	2.753
120b	1.61804	a	4.236
120c	0.951057	$\tau/2$	11.66

Table 4.2: Maximum distance from each vertex configuration to its nearest copy, in radii, and the span of each configuration. $\tau = (1+\sqrt{5})/2$, $a = (\sqrt{10+2\sqrt{5}})/4$.

Embedding a Patch in a Vertex Configuration Now that we know the distance between identical global vertex configurations, we are ready to consider the distance between identical arbitrary patches. As explained in the proof of Theorem 2.2, the nearest copy of a patch can be located by deflating the tiling until the patch is contained within a single vertex configuration C, locating the nearest copy of C, and reinflating. Because inflation is unique, both copies of C will produce copies of the original patch when inflated.

Lemma 2.2 shows that if the diameter of a finite patch $\mathcal{P}$ contained in an infinite tiling $\mathcal{T}$ is less than the minimum vertex configuration span of $\mathcal{T}$, then $\mathcal{P}$ must lie within a single vertex configuration of $\mathcal{T}$. To produce a bound on the distance from an arbitrary patch $\mathcal{P}$ of a Danzer tiling to its nearest identical copy, we consider three factors. First, we must deflate the tiling until we are certain that $\mathcal{P}$ is contained within a single vertex configuration C. Second, we must determine the worst-case distance to the nearest copy of C. Finally, we must account for the fact that the two instances of C may not be oriented in the same direction. To simplify our computation, we consider deflation to shrink $\mathcal{P}$ rather than expanding the tiles of $\mathcal{T}$. That is, after a deflation we increase our distance scale by a factor of τ so that the dimensions of the tiles in $\mathcal{T}$ remain unchanged and the diameter of $\mathcal{P}$ is scaled by $1/\tau$. The construction of a bound on γ proceeds as follows:

As noted in Section 2.2.4, the span of a vertex configuration composed of tetrahedral tiles is simply the shortest edge with an endpoint at the central vertex of the configuration, and thus $V_{\min}$ for the Danzer tiles is the length of the shortest edge of any tile. The shortest edge is K_{34}, which has a length of $1/2\tau \approx 0.309$. Thus we must deflate $\mathcal{T}$ until, under the metric described above, d is smaller than $1/2\tau$. We also note in Section 2.2.4 that the granularity of the deflation contributes another factor of τ. If after $k-1$ deflations d is very slightly larger than $V_{\min}$, then after k deflations d will be smaller than $V_{\min}$ by a factor of close to (but strictly less than) τ. Thus deflating $\mathcal{P}$ until $1/2\tau^2 < d \leq 1/2\tau$ ensures that $\mathcal{P}$ is contained within a single vertex configuration C.

Next we determine the worst-case distance to the nearest copy C' of C. From Table 4.2 we see that le largest such distance is $\tau^5/a \approx 11.66$. Finally, we consider the possibility that C and C' are not oriented in the same fashion. In the worst case, one is rotated 180 degrees with respect to the other, so that they face in exactly opposite directions. Then the distance between corresponding points in the two configurations is increased by the length of the longest edge with an endpoint at the center of the configuration. The longest edge of any tile is C_{13}, which has a length of τ.

Thus if $\mathcal{P}$ lies entirely within C, the maximum distance to a corresponding region in a copy of C is $\tau^5/a + 2\tau$. We wish to measure this distance in units of d. From the argument above we know that $d > 1/2\tau^2$. Thus we conclude that the maximum

distance from $\mathcal{P}$ to its nearest copy is $(\tau^5/a + 2\tau)2\tau^2 \approx 78.002$ diameters, or twice that many radii.

We have shown that it is possible to place a bound on γ. There is significant room for improvement, however. The three worst cases we mentioned above (the smallest vertex span, the largest distance to the nearest copy of C, and the longest edge) do not occur simultaneously. Thus to improve our estimate we compute a bound on γ for each vertex configuration, taking into account its particular span, nearest-neighbor distance, and longest edge. The maximum of these values is our bound on γ.

Danzer states that $\gamma = 40$ for this tiling. It remains unclear how he arrived at this value. The method described above produces a value of about 46. Though our intuition suggests that Danzer is correct, we are unable to verify his claim. Possible methods for further tightening the bound on γ are discussed in Section 6.1.

Chapter 5

Extending the Atlas

In the last chapter we discovered that there are 22 vertex configurations which occur in an infinite tiling. Examining their profiles (see Appendix A), it seems that we should be able to build more than these 22 without breaking any of the matching rules. For instance, configuration 90a (Figure A.16) looks like a C-ball with roughly the lower third replaced by a disk of A tiles. There is no obvious reason that we could not perform the same replacement on the top half of the ball. Though we cannot find such a configuration in a tiling, one can be constructed without violating the matching rules.

The fact that we can build one vertex configuration which does not appear in an infinite tiling suggests that there may be more (for instance, look at configurations 90b, 100, and 110 — the two cones of the Princess could be directly opposite one another). A knowledge of such configurations could be quite useful in applying tiling theory to quasicrystals. One of the major arguments against the aperiodic model of quasicrystals is what mathematicians call forced tiles and chemists call non-local effects: one small group of tiles can force the positions of other tiles at arbitrarily large distances. It is unclear how atoms attaching to the surface of a growing crystal could "know" about forcing effects from deep inside the crystal. Vertex configurations which are locally legal but do not appear in an infinite tiling could represent the possible defects in a crystal that would occur if an atom were to attach based only on its immediate surroundings. A complete list of legal vertex configurations will also allow us to prove that the 22 global configurations discussed in the previous chapter are in fact the only vertex configurations which can appear in an infinite Danzer tiling.

This chapter will extend the atlas of the previous chapter to include all locally legal vertex configurations by developing an algorithm to construct them exhaustively. The first two sections abstract the problem, allowing us to solve it on a discrete-

valued, two-dimensional surface using a simple representation. We will then show how one can enforce the matching rules in our system, and discuss issues of efficiency. Finally, we present some rather surprising results: our algorithm found 174 locally legal configurations (including the 22 global ones). Part **(Vert)** of Danzer's main theorem suggests that there are a total of 27 vertex configurations.

Our goal in the next several sections is to devise an algorithm that will construct all possible vertex configurations. We will use a depth-first searching algorithm, based on an approach presented by Baake *et al.*[BBAK$^+$94], placing tiles around the central vertex until we have a complete configuration or we reach a point where it is impossible to place a tile without violating the matching rules. A note on terminology: in the coming sections, we will often talk about tiles and their vertices. In this context, we do not use *vertex* to mean a zero-dimensional point but rather the part of the tile contained in a ball of some small but non-zero radius centered at that point. Thus a vertex has a shape and angles associated with it, and is subject to the matching rules. For example, when we place B_3 vertices next to each other, we mean that we are placing B tiles together (according to the matching rules) such that all their 3-labeled vertices coincide.

5.1 The Importance of Being Discrete

The first issue in designing our algorithm is the representation of tiles. There are several difficulties associated with representing arbitrary locations of tiles in three dimensions. One is the issue of accuracy and rounding error. Our task would be much simpler if we could somehow limit each tile to a finite (and hopefully small) number of positions. We could then represent locations and adjacencies explicitly rather than relying on difficult computations using real-valued coordinates.

The key observation is that there is a "common divisor" of every vertex. We concluded from the examination of the effects of inflation on vertex configurations that every configuration, after at most four inflations, inflates to the B-ball (see Figure 4.1). This means that some number of B_3 vertices can be packed together to exactly cover any vertex of any other tile. (This can also be seen by tracing the inflation of a vertex through several steps using the inflations in Table 3.4). Thus we can think of the B-ball as a fixed template upon which we can construct any vertex configuration. Every tile must be placed so that one of its vertices is in the center of the B-ball and so that it overlaps completely some number of the tiles in the B-ball. Because every tile can be placed in this manner, we lose no generality by "justifying" the locations of tiles so that they line up with the B tiles. Any vertex configuration not justified to the B-ball can be rotated so that it is. Figure 5.1 shows how an A tile with A_1 at the center overlaps exactly three tiles on the B-ball.

Next we wish to know exactly how many different justified positions a tile can occupy. Recall from Section 4.2 that the B-ball has full icosahedral symmetry. The icosahedral group contains 60 transformations which map the icosahedron to itself. These transformations will be discussed in more detail in Section 5.3, but for now we note that a tile with a particular vertex at the center of the configuration can occupy 60 distinct positions that are justified with respect to a fixed B-ball.

5.2 Down to Two Dimensions

This last statement implies that every tetrahedron can be in 240 positions relative to the B-ball. As an example, an A tile can occupy 60 positions when A_1 is at the center of the ball, but of course the A_2 vertex has a different shape and covers a different region of the B-ball when placed at the center. Thus each tile can occupy 60 positions for each of its four vertices.

This motivates a conceptual simplification of the problem. Rather than thinking of rather complex three-dimensional tiles covering parts of the B-ball, consider the projection of the B-ball onto a sphere. That is, place a sphere around the B-ball so that the sphere and the B-ball share the same center, and project the tiles outward from the center onto the sphere. Each tiles casts a triangular "shadow", and we obtain a tiling of the surface of the sphere, comprised of 120 congruent spherical triangles. If we flatten each triangle (connecting its corners with a plane rather than a patch of a sphere), we obtain a 120-faceted polyhedron which we will call $\mathcal{B}$ (Figure 5.2). Figure 5.3 is a cut-out pattern for an icosahedron with the projection of the facets of $\mathcal{B}$ drawn on it.

When we place, for example, an A tile with vertex A_1 at the center, we can project it outward onto $\mathcal{B}$, and its projection will completely cover some subset of the facets of $\mathcal{B}$. If instead we place A_2 at the center, we get a projection that covers a different set of facets. Rather than view these sets as different aspects of an underlying three-dimensional tile, we can view them purely as separate sets of facets, to be placed as independent two-dimensional tiles. This gives us four projections of each of the eight tetrahedra, one from each vertex, for a total of 32 two-dimensional tiles (or sets of facets) on the surface of $\mathcal{B}$.

Each of these 32 tiles can occur in 60 positions, corresponding to the 60 icosahedral symmetries. Finally, note that just as we can view different projections of a 3-D tile as different 2-D tiles, we can view different positions of a 2-D tile as separate tiles. Thus we now have $32 \cdot 60 = 1,920$ flat tiles, each of which covers an immovable set of facets on $\mathcal{B}$.

These conceptual modifications simplify the problem of building a vertex configuration considerably. We no longer need to be concerned about which vertex to

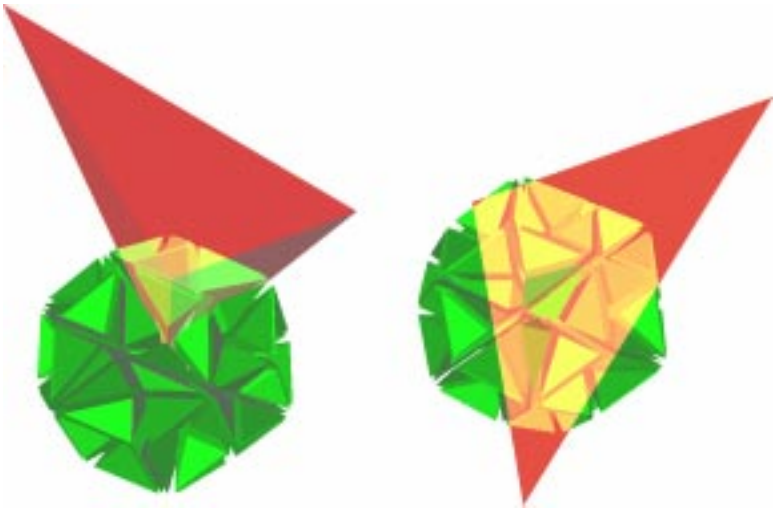

Figure 5.1: A partially transparent A tile superimposed on the B-ball. Vertex 1 of the A tile is at the center of the ball. The tiles of the B-ball have been separated for clarity; there would normally be no gaps. *Left:* The three edges of the A tile intersect the B-ball at vertices. *Right:* The three darker tiles in the center are exactly covered by the A tile.

Figure 5.2: The 120-faceted ball $\mathcal{B}$. Adjacent "plates" of 10 facets have been given different colors for contrast. There are 12 such plates, each centered at one of the vertices of an icosahedron.

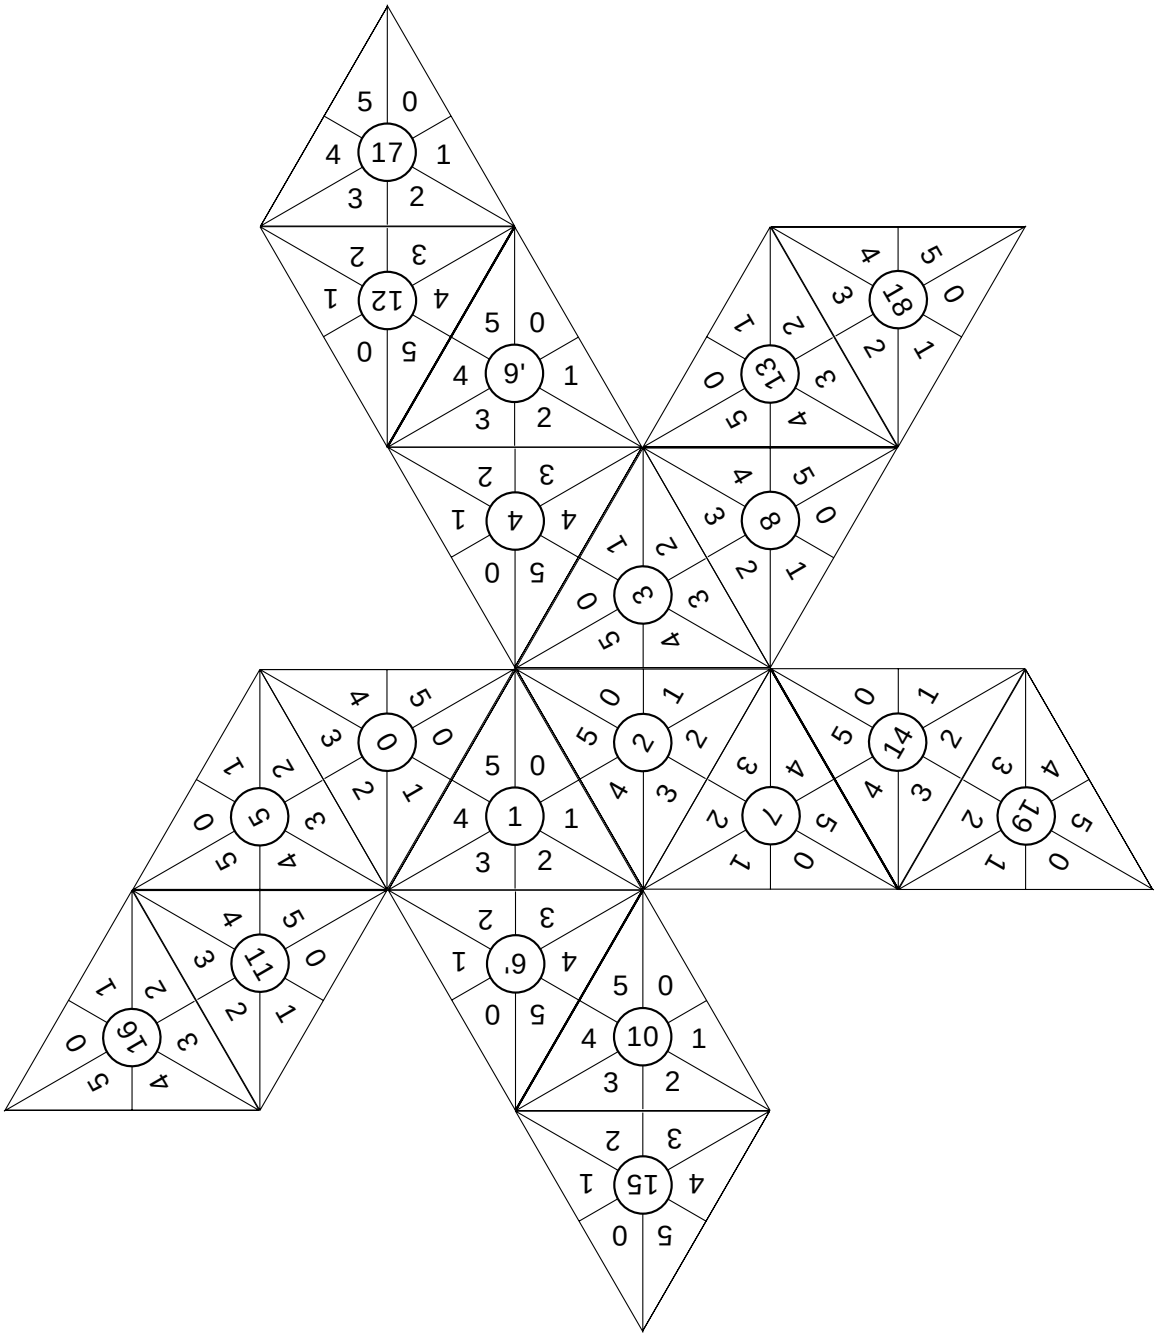

Figure 5.3: A cut-out pattern for an icosahedron. The faces of the icosahedron have been marked with the projections of the facets of $\mathcal{B}$.

place in the center or how to rotate a tile. We must simply find all subsets of our 1,920 tiles which cover $\mathcal{B}$ completely without overlapping. There remain some significant difficulties to be addressed. We discuss the enforcement of matching rules and the tractability of such a large packing problem in Section 5.4. First, however, we will consider the representation of these tiles.

5.3 Representation

Viewing our tiles as sets of facets on the 120-faceted ball $\mathcal{B}$ leads to a convenient representation. A tile T is stored as an array T_0 to T_{119} of boolean values. If T_i is true, then facet i is covered by the tile T. We can check whether two tiles intersect with a bitwise AND, and whether a collection of tiles covers $\mathcal{B}$ with bitwise OR. We also need to know which of the 62 vertices of $\mathcal{B}$ are touched by the border of the tile. This vertex list can be stored in the same manner as the facet list. To avoid confusion between the 62 vertices of $\mathcal{B}$ and the three vertices of a tile on the surface of $\mathcal{B}$, we will refer to the vertices of each tile as *corners* for the rest of this section.

The list of facets and vertices covered by each of the 32 basic tiles is presented to the algorithm explicitly, and then the 60 positions of each tile are computed by applying each of the 60 icosahedral transformations. A transformation is simply a mapping from $\mathcal{B} \mapsto \mathcal{B}$, stored as an array A in which each element A_i contains the location of facet (or vertex) i under the transformation. For example, if $A_6 = 50$, then when $\mathcal{B}$ is rotated according to this transformation, the facet that is currently in location 6 will be in location 50.

Fortunately, we do not have to construct all of these transformation arrays manually, for the following reason. An icosahedron has six axes of fivefold rotational symmetry passing through opposite vertices (*white* axes), 10 axes of threefold rotational symmetry passing through the centers of opposite faces (*green* axes), and 15 axes of twofold rotation bisecting opposite edges (*red* axes). The three basic transformations which map the icosahedron to itself are rotation by 72° about a white axis, rotation by 120° about a green axis, and rotation by 180° about a red axis. All other icosahedral symmetries are equivalent to the composition of a number of these three elemental transformations. Thus we choose an arbitrary numbering for the facets (the identity transformation), explicitly compute facet and vertex mappings for the three basic transformations, and then construct the remaining 56 by composition.

Thus our final representation of each tile is a 120-element array denoting the covered facets of $\mathcal{B}$ and a 62-element array denoting the vertices of $\mathcal{B}$ which lie along the boundary of the tile. There are 1,920 such pairs of arrays, each pair corresponding to one tetrahedron with a particular vertex placed at the center of

the configuration, occupying a particular location with respect to the B-ball.

We maintain two global arrays which represent the facets and vertices of $\mathcal{B}$, in which store keep the union of all tiles we have placed so far. As we will see in the next section, this allows us to easily verify whether a tile about to be placed will conflict with any portion of the partially assembled vertex configuration.

5.4 Enforcing the Matching Rules

Our algorithm for constructing a vertex configuration will be based on a recursive depth-first search. We are given some portion of the configuration which has already been constructed in the global arrays representing $\mathcal{B}$. Since every outer edge of this patch must be matched eventually, we pick one edge arbitrarily to be the site of the next placement. Then for each tile that can be placed adjacent to the chosen edge without a conflict, we place the tile, recurse, and then remove the tile. After each placement we check whether all the facets of $\mathcal{B}$ have been covered; if so, we have a complete vertex configuration. To determine whether the configuration we have constructed is actually new, we compare it under all 60 transformations to each configuration we have already found. If it is not congruent to any of them, then we add it to the list of configurations found.

The key phrase in the paragraph above is "without a conflict." We must ensure that the three-dimensional vertex configuration represented by the tiling of the surface of $\mathcal{B}$ is face-to-face and obeys the matching rules for the Danzer tiles. We must also reduce the branching factor at each level of recursion if the algorithm is to be computationally tractable.

We can accomplish several of these goals using the vertex classes (or colors) described in Section 3.5. Recall that for each of the 1,920 fixed tiles we have an array denoting which of the vertices of $\mathcal{B}$ are on the boundary of the tile. Rather than using simple boolean values for these marks, we assign each vertex one of six colors. A color of black indicates that the vertex is not on the boundary of the tile, and grey indicates that a vertex is on an edge of the tile but is not a corner. The remaining four colors (red, green, blue, and yellow) are reserved for the vertices of $\mathcal{B}$ which coincide with corners of the tile and indicate the class of the vertex of the corresponding tetrahedron.

When considering a tile for placement, we require that the non-black vertex colors of the tile match the colors of the corresponding vertices of $\mathcal{B}$, with the understanding that an unassigned vertex on $\mathcal{B}$ will match anything placed upon it. The use of a separate color for vertices along the edge of a tile ensures that we will only place tiles edge-to-edge, since the color of a corner will never match the color of a point in the middle of an edge.

It should be noted that these restrictions alone are sufficient to make the problem tractable. The branching factor at each level, while still significant, is sufficiently limited by choosing only tiles that lie along a particular edge of the existing patch and that have corners of the appropriate color. This version of the algorithm completed in about 14 hours on an SGI Indigo2.

Unfortunately, while vertex color has the feature that the two-dimensional tiles will be placed edge-to-edge, it is not sufficient to completely enforce the matching rules. The projection of a tetrahedron onto $\mathcal{B}$ (the shape of one of the 32 basic flat tiles) is determined only by the angles at the central vertex, not by the shapes of its faces or the lengths of its edges. There are a number of faces of the tetrahedra which have the same projection onto $\mathcal{B}$ and the same vertex colors, but which are not congruent. This algorithm produced almost 500 vertex configurations, the majority of which were illegal.

To prevent some of the illegal pairings allowed by the algorithm, we also carry a length value with each vertex, which represents the point's distance from the center of the configuration, and we require that these match in the same manner as the vertex colors. This guarantees that the two adjacent three-dimensional faces represented by adjacent edges on the surface of $\mathcal{B}$ will be congruent. This further restriction reduced the running time to roughly 8 hours and the output to 268 configurations, but still allowed some illegal pairings. For example, the $C2$ face is isosceles and the two vertices of the base, C_1 and C_3, are the same color (blue). Thus two right-handed C tiles pointing in opposite directions can be placed together to form a five-sided pyramid whose base is a parallelogram, a pairing which is not allowed under the matching rules for the Danzer tiles but which does not violate the color or length restrictions above.

The final step in enforcing the matching rules makes use of the fact that the Danzer tiles always form octahedra, as described in Section 3.4. Essentially, we are taking into account the most local forcing effects. During the placement of a tile, if we know that there is only one tile that can be placed adjacent to one of its edges, then we immediately place that tile rather than waiting until a later level of recursion. An octahedron formed from four tetrahedra sharing a central edge has two vertices at which all four tetrahedra meet (the endpoints of the central axis) and four vertices at which two tetrahedra meet. Thus depending on which two-dimensional tile we are placing (that is, which vertex of the tetrahedron it represents is at the center of the configuration), we can always place either two or four tiles at a time in this manner. Rather than implementing this as separate tile placements, we create a set of "meta-tiles," each of which is the union of two or four of the 32 basic two-dimensional tiles. We then seek to cover $\mathcal{B}$ with these meta-tiles. This modification does not discard any generality; every meta-tile is a pair or quadruple of tiles that must *always* be placed together. We simply avoid determining this fact

anew each time a tile is placed.

The combination of vertex color, length matching, and meta-tiles is sufficient to completely enforce the matching conditions (the algorithm produced no illegal vertex configurations). The meta-tiles also greatly reduce both the branching factor and the depth of the recursion, reducing the running time of our implementation from over 8 hours to 2 minutes and 40 seconds. The results of the execution of this algorithm are presented in the next section.

5.5 Results

The results of running the algorithm described above were rather surprising. We found 174 vertex configurations, including the 22 global configurations. This is in contrast to part **(Vert)** of Danzer's main theorem, which claims that there are exactly 27 configurations.

The 152 new configurations fall into three major groups. Group 1 has 21 members which are similar to global configurations 60b, 70, 80, 90b, 100, and 110 and the new configuration 100n. They are formed by replacing parts of a K-ball with cones of B tiles and clusters of four A tiles. Group 2 has 121 members, all based on a single replacement. Six A tiles can be placed so that they meet on faces 1 and 2 and share the A_{34} edge, forming a 6-tile "divot" which can be substituted for 12 B tiles on the B-ball in any of twenty different locations. Up to eight such replacements can be made simultaneously. The remaining ten vertex configurations are largely distinctive. Appendix B contains pictures of the distinctive new configurations and a representative from each of the other two groups, and Table 5.1 gives some of their characteristics. Figure 5.4 shows the effects of inflation on all vertex configurations.

The new configurations have a number of interesting characteristics, which we list below in no particular order:

- All of the 22 global configurations have at least one plane of mirror symmetry. In contrast, configurations 40na and 40nb are completely asymmetrical, and are mirror images of each other.

- Configuration 44n is the only configuration to use two different vertices of the same type of tile at its center. Of the 20 C tiles it contains, eight have C_1 at the center and twelve have C_3 at the center. Thus this is the only configuration (besides the K-octahedron) in which all the members of a vertex class are present (see Section 3.5).

- 48n and 54n are variations on the globally occurring 60a, produced by the same replacement which gives rise to the Group 2 configurations.

		Composition			
Designation	**Name**	**A**	**B**	**C**	**K**
36n	Shuriken		12	24	
38n			6	12	20
40na				20	20
40nb				20	20
44n	Snowplow	6	6	20	12
48n		22	16	10	
50n	Torch			10	40
54n		16	28	10	
60n	Discus	20		40	
100n	Prince		20		80
Group 1		4-24	10-26		20-84
Group 2		6-48	24-108		

Table 5.1: Basic characteristics of the ten distinctive new configurations and ranges for the members of the other two groups.

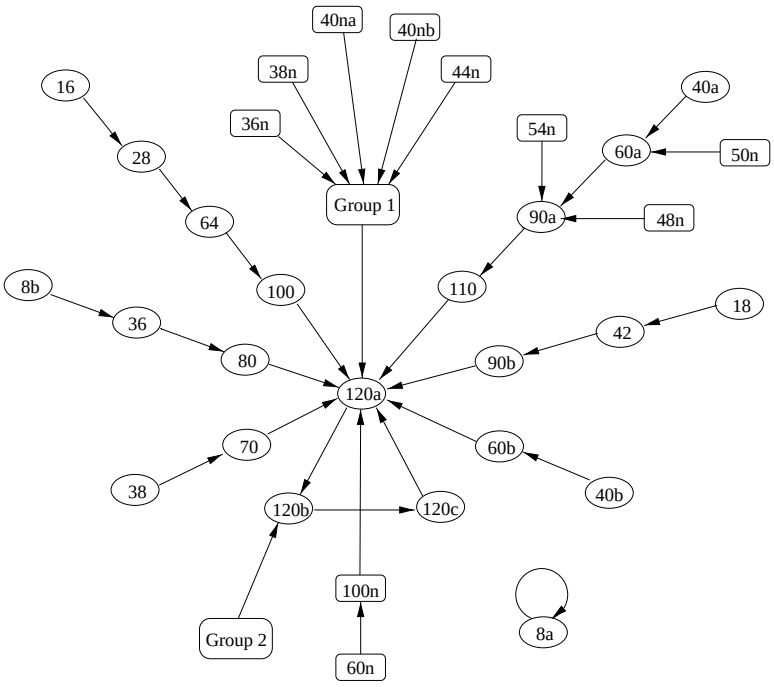

Figure 5.4: The effect of inflation on all vertex configurations. The new configurations are drawn as rounded rectangles.

- 60n is a natural extension of 90a, replacing both sides of the C-ball with medallions of ten A tiles.
- 100n is a variation on configuration 100 (the Princess), and is part of the same class of modifications to the K-ball as 60b, 70, 80, 90, 100, 110, and all the members of Group 1.
- 36n, 38n, and 44n contain the face pairing $B3 - C3$, which does not appear in a global tiling because it cannot be extended (see Section 3.2). The divot of 6 A tiles used in 48n, 54n, and all the members of Group 2 contains the $A1 - a1$ pairing, which also cannot be extended.
- All the new configurations except those in Group 2 become the B-ball after no more than four inflations, as do the global configurations. The Group 2 configurations all become the C-ball after one inflation. They are the only configurations which inflate to a ball other than the B-ball.

Chapter 6

Conclusion

Our work has focused primarily on the vertex configurations of the Danzer tiling. We have produced an atlas of all legal vertex configurations, and identified those which occur in an infinite tiling. We have provided an elucidation of the paper in which Danzer first presented his tiles, and have produced proofs for the majority of his conclusions. We have given a method of computing the local isomorphism constant γ. Finally, we have extended and refined a library of tools for computationally investigating and visualizing Danzer tilings.

The atlas of vertex configurations is useful for a number of reasons. We may wish to use the Danzer tiling as a model for the atomic structure of a quasicrystal. In this case, the vertex configurations represent the immediate neighborhoods around atoms. The non-global configurations could correspond to defects in the crystal, an important concern given the questions surrounding the incremental formation of quasicrystals.

From a more theoretical point of view, the atlas provides insight into the structure of the tiling as a whole. In the study of the Penrose tiling, vertex configurations are critical to the understanding of the locations of the Ammann bars and the tiles forced by a patch. We hypothesize that there are Ammann planes in the Danzer tiling. If they exist, then a understanding of the global vertex configurations will certainly be necessary in determining their locations and the manner in which they force tiles.

We invested a significant amount of work into proving the theorems and conclusions presented in Danzer's paper. A clearer explication of the construction and properties of these tiles is a valuable resource for future research. The proofs themselves often turned out to be surprisingly complex, and led to results that were not stated in Danzer's paper. Properties of a tiling such as the fixed density of tiles and local isomorphism are subtle, and the process of proving them facilitated a deeper

understanding of the tiling. The method of bounding γ, for instance, grew out of the general proof of local isomorphism. In the course of tracing Danzer's steps, we have refined or corrected several of his results:

- Danzer states that $H = 6$. Configuration 120c, the K-ball, is not found until the 7th inflation of any prototile, and thus the Danzer tiling has $H = 7$.
- Though Danzer states that every vertex configuration except the K-octahedron inflates to one of the three 120-tile balls within six inflations, only four inflations are necessary.
- There are 174 legal vertex configurations. Danzer states that there are 27.

We have laid the groundwork for a tighter bound on the value of γ, a constant which is critical to the structure of the tiling. The constructive method of proof is particularly interesting because it might be applicable to other tilings which obey inflation and deflation rules.

Finally, we have built tools which allow us to construct, visualize, and study Danzer tilings computationally. These tools have been invaluable in our work, and they provide important and useful methods of both confirming hypotheses about the tiling and actually proving results, as in the constructive proof of the completeness of the global atlas.

6.1 Future work

There remain many unanswered questions relating to the Danzer tiling. The most tractable of these, we believe, is improving the bound on γ. In Section 4.5 we accounted for different orientations of nearest-neighbor vertex configurations by adding twice the length of the longest edge to the distance between their centers. It seems likely that this term is much larger than necessary. We suspect that in the worst case the distance between the centers of two identical patches will be the distance between their containing configurations plus twice the radius of the patch. This leads us to the following conjecture:

Conjecture: $\gamma < 28$ *for the Danzer tiling.*

This value for γ is based on a worst-case distance between vertex configurations which we suspect may be due to a rounding error in the distance computation. Thus we conjecture, with slightly less confidence than previously, the following:

Conjecture: $\gamma \leq 18$ *for the Danzer tiling.*

In the case of the first conjecture, the computed value for γ was rounded up to 28. In the second conjecture, however, 18 is a precise value. It is possible that in fact $\gamma = 18$ is a tight bound. That is, it may be possible to locate a patch whose nearest duplicate is arbitrarily close to 18 radii away. Proof of this claim would require a much deeper and more subtle analysis of the problem than we have given, however.

Besides more tightly bounding the distance between identical patches, there are a number of other interesting questions which open avenues for future work:

- Can the constructive method of determining a bound on γ be applied to other tilings? In particular, the theoretical bound on γ for the Penrose tiling is much larger than the observed bound. Can we tighten the bound for the Penrose tiling?
- It could be useful to build a computational tool for measuring the distance from an arbitrary patch to its nearest copy. Is the observed bound on γ in the Danzer tiling less than the provable bound, as in the Penrose tiling?
- Are there Ammann planes in the Danzer tiling, and if so, where do they lie? What are the empires of the vertex configurations?
- Can we extend the algorithm presented in Chapter 5 to produce second-order vertex configurations (that is, patches containing all tiles adjacent to a vertex configuration)?
- Corollary 2.2 shows that the relative frequency of each prototile is non-zero. This implies that in the inflation matrix $\mathcal{M}$, the ratio of an element to the sum of its column is non-zero for all powers of $\mathcal{M}$. This prompts the following conjecture:

  Conjecture: *Let $\mathcal{M}$ be an $n \times n$ matrix containing only positive entries. Then*

  $$\lim_{k \to \infty} \frac{\mathcal{M}_{i,j}^k}{\sum_{c=1}^n \mathcal{M}_{c,j}^k}$$

  exists, is non-zero, and is the same for all $1 \leq i, j \leq n$.

  That is, the ratio of an element to the sum of its column converges to a non-zero limit when any $\mathcal{M}$ containing no zeros is raised to arbitrarily large powers, and that limit is the same in each column.

  A related question is whether any square matrix $\mathcal{M}$ such that some power of $\mathcal{M}$ contains no zero entries could be the inflation matrix for some tiling.

Appendix A

Global Vertex Configurations

This appendix presents pictures of the 22 global vertex configurations. The pictures were rendered on a Silicon Graphics Indigo2 using powerflip, a demo shipped with the operating system. A tiles are drawn in red, B tiles in green, C tiles in blue, and K tiles in yellow. In these pictures, no distinction is made between left- and right-handed tiles. Each tile has been shrunk by 30% toward its center of mass to allow us to see inside the configuration. While each individual image is an accurate rendering, scales are not the same between figures or between images within a figure. Figure A.23 shows all the configurations drawn to scale.

The caption for each figure gives how many of each prototile are used in the configuration, and the number of the central vertex. Every configuration has an equal number of left- and right-handed tiles, so for brevity we combine them. For example, $4 \cdot K_3$ means that the configuration contains two K tiles and two k tiles, and that for all four of those tiles vertex 3 is in the middle of the configuration.

Although we show several views for each configuration, it is always difficult to infer three-dimensional structure from flat pictures. The data files containing the locations of the tetrahedra are available in powerflip, POV, and raw formats at http://www.cs.williams.edu/tiling/ for those readers interested in viewing these objects more dynamically.

Figure A.1: Configuration 8a, the K-octahedron. Composition: $8 \cdot K_4$.

Figure A.2: Configuration 8b, the Sailboat. Composition: $4 \cdot B_4$, $4 \cdot K_3$.

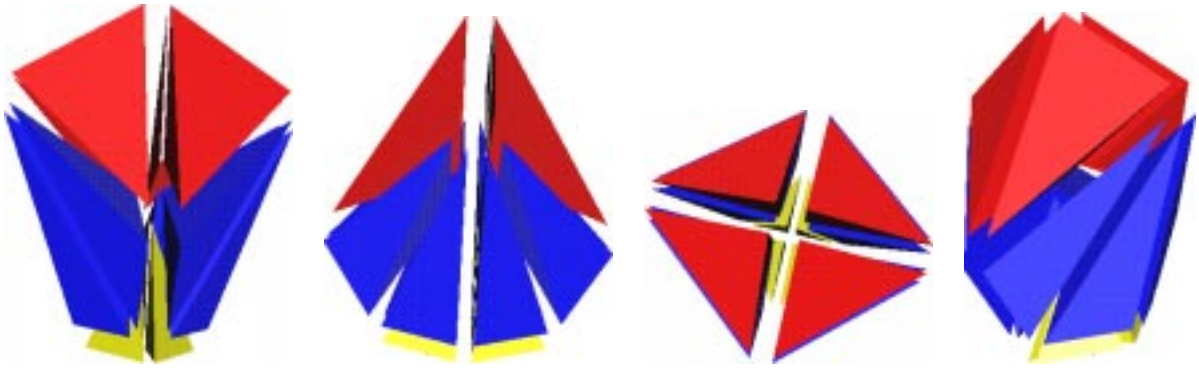

Figure A.3: Configuration 16, the Pawn. Composition: $4 \cdot A_4$, $8 \cdot C_4$, $4 \cdot K_2$.

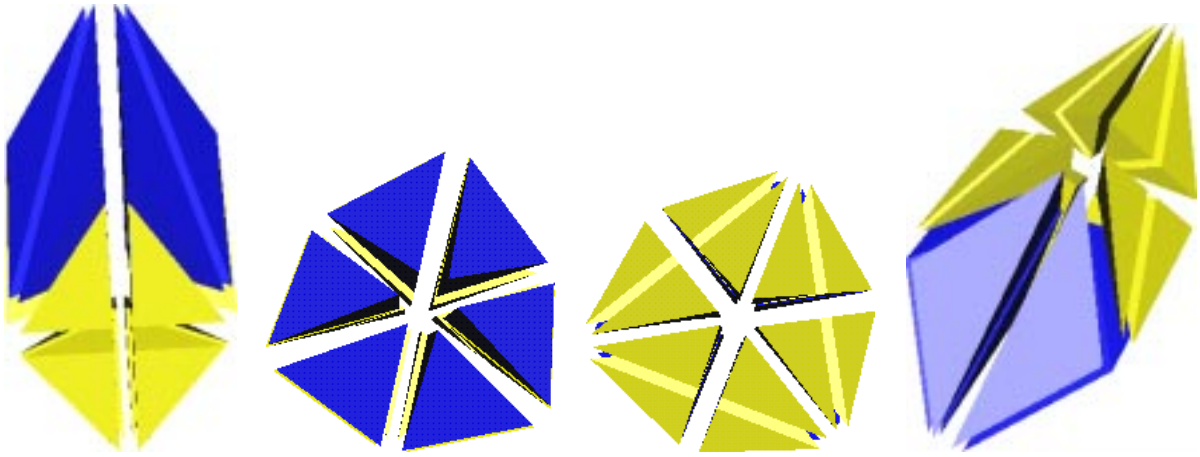

Figure A.4: Configuration 18. Composition: $6 \cdot C_2$, $12 \cdot K_3$.

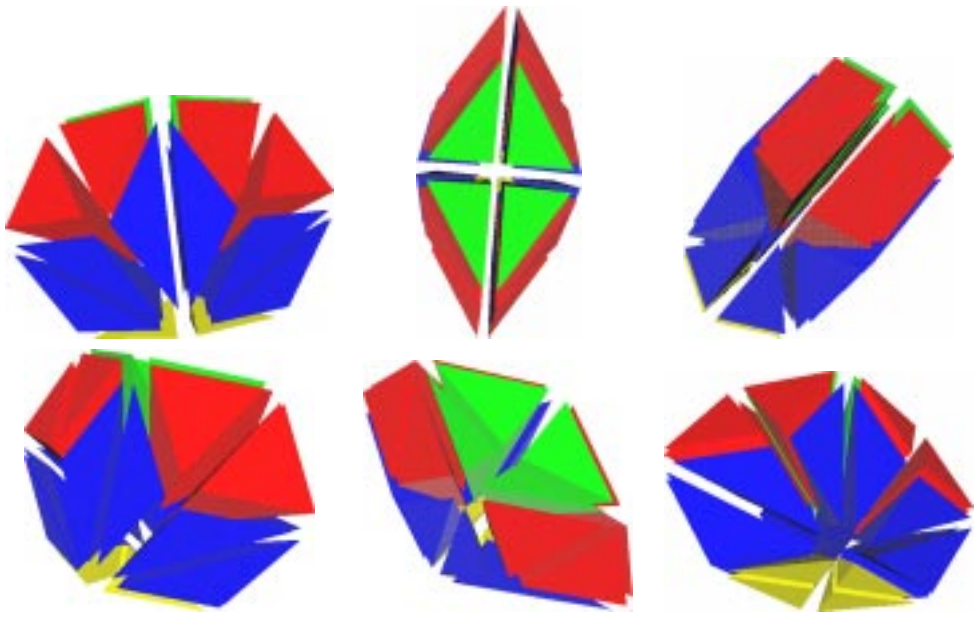

Figure A.5: Configuration 28. Composition: $8 \cdot A_3$, $4 \cdot B_3$, $12 \cdot C_2$, $4 \cdot K_3$.

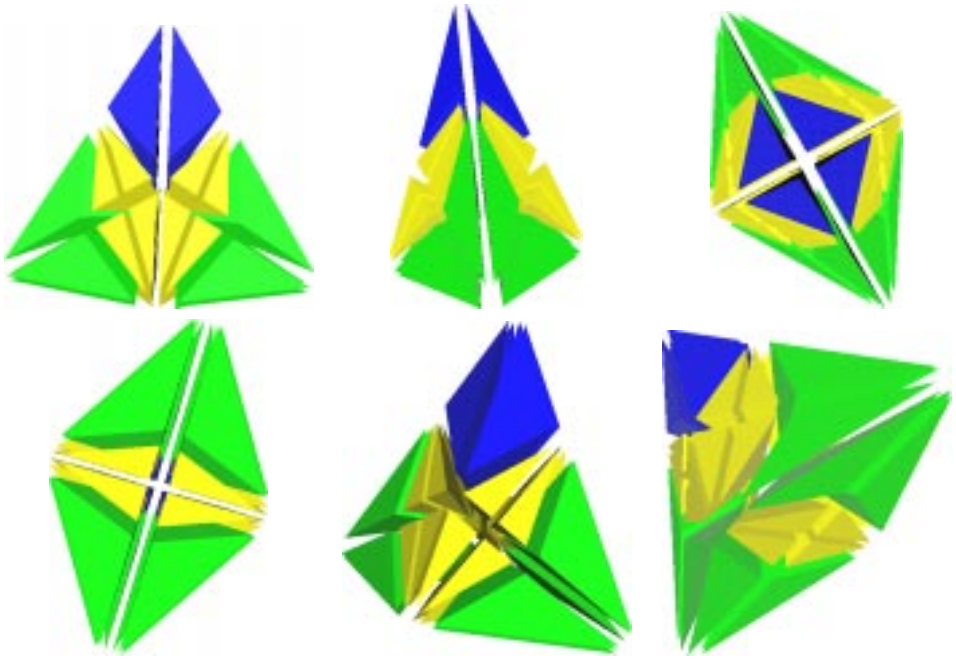

Figure A.6: Configuration 36, the Bishop. Composition: $12 \cdot B_1$, $4 \cdot C_1$, $20 \cdot K_1$.

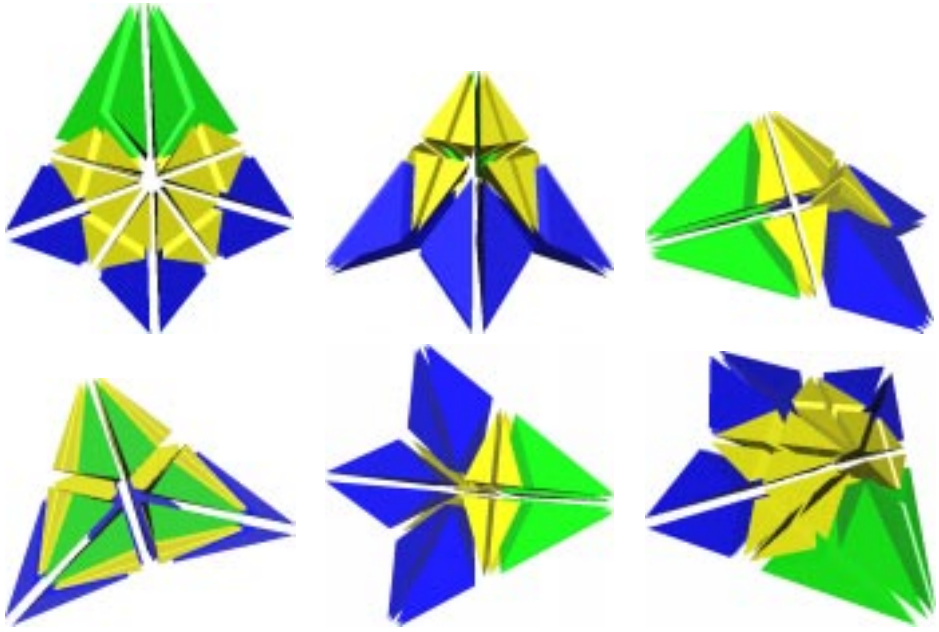

Figure A.7: Configuration 38, the Claw. Composition: $6 \cdot B_1$, $12 \cdot C_1$, $20 \cdot K_1$.

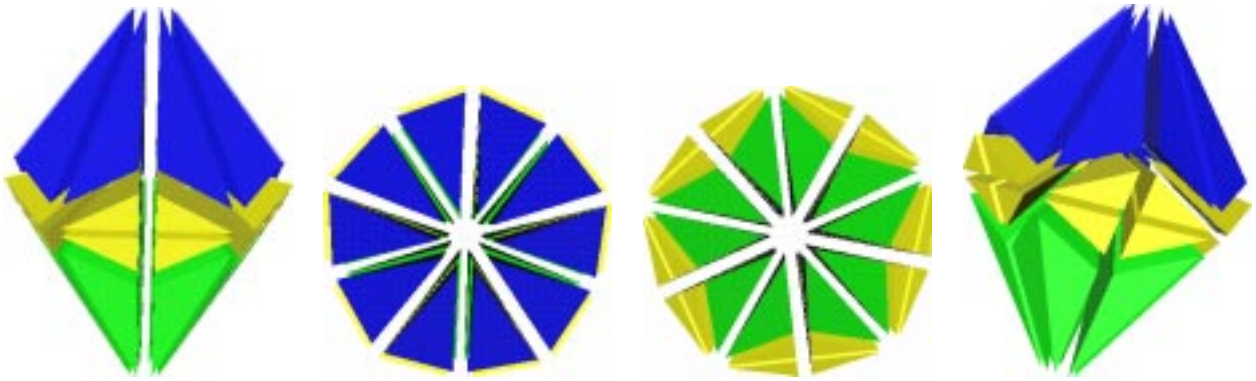

Figure A.8: Configuration 40a, the Diamond. Composition: $10 \cdot B_2$, $10 \cdot C_4$, $20 \cdot K_2$.

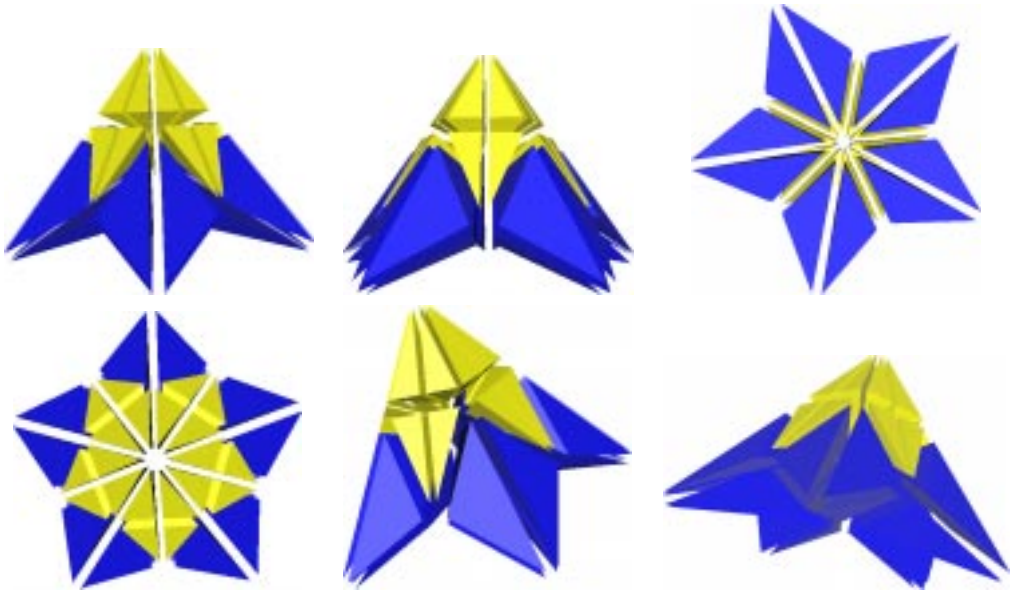

Figure A.9: Configuration 40b. Composition: $20 \cdot C_1$, $20 \cdot K_1$.

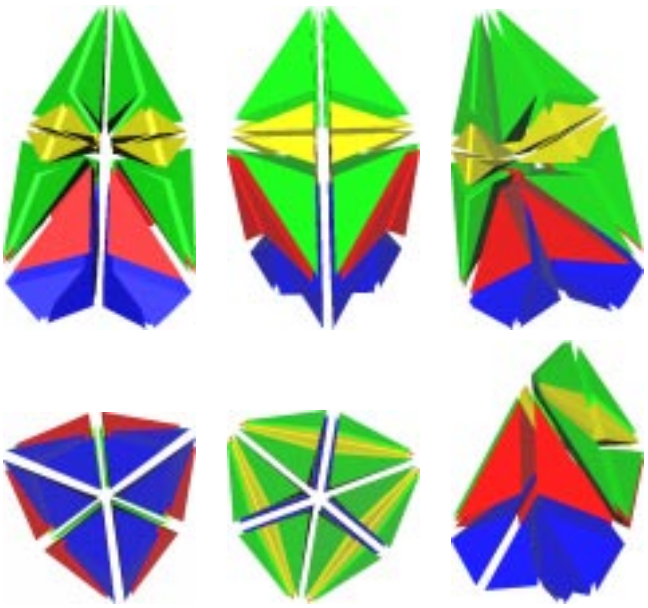

Figure A.10: Configuration 42, the Obelisk. Composition: $6 \cdot A_2$, $12 \cdot B_1$, $12 \cdot C_3$, $12 \cdot K_1$.

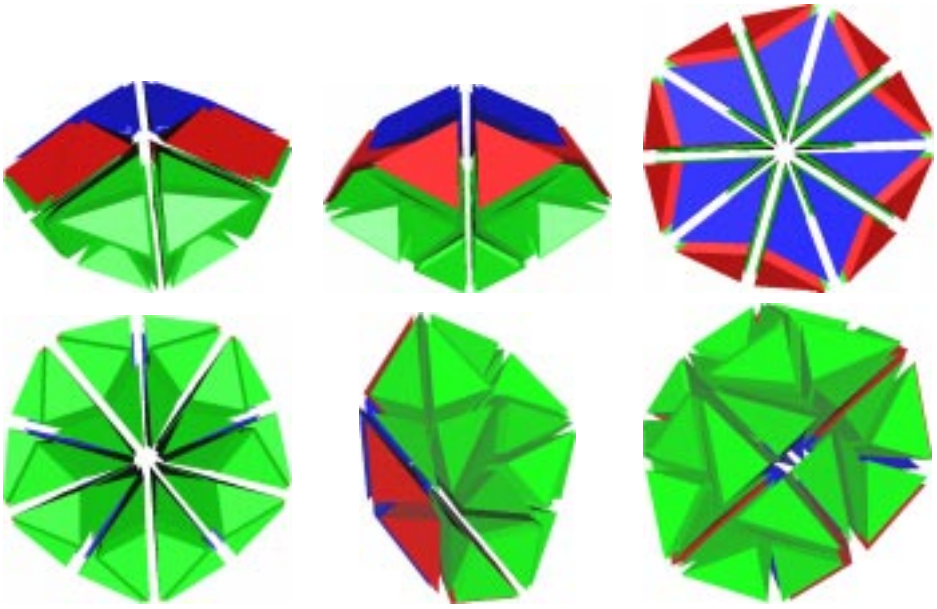

Figure A.11: Configuration 60a. Composition: $10 \cdot A_3$, $40 \cdot B_3$, $10 \cdot C_2$.

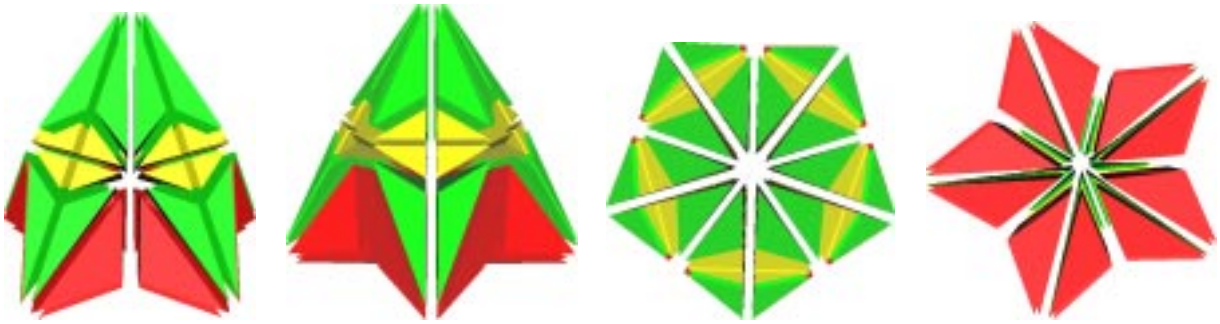

Figure A.12: Configuration 60b. Composition: $20 \cdot A_1$, $20 \cdot B_2$, $20 \cdot K_2$.

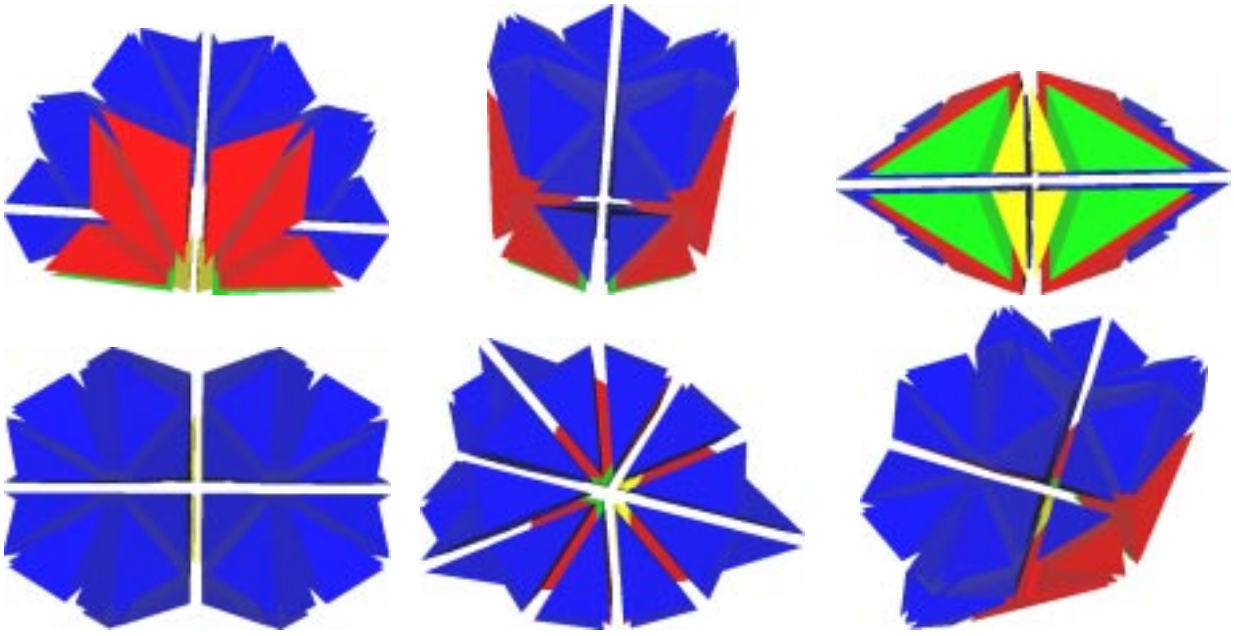

Figure A.13: Configuration 64, the Taco. Composition: $12 \cdot A_2$, $4 \cdot B_1$, $44 \cdot C_3$, $4 \cdot K_1$.

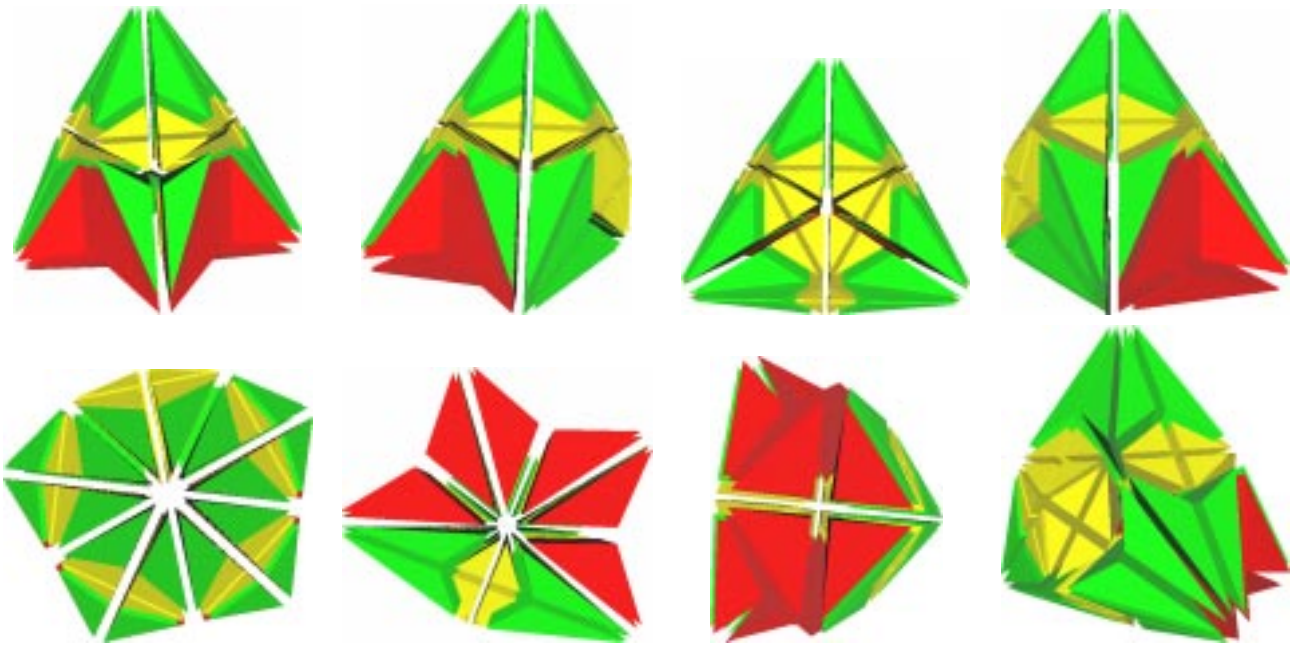

Figure A.14: Configuration 70. Composition: $12 \cdot A_1$, $26 \cdot B_2$, $32 \cdot K_2$.

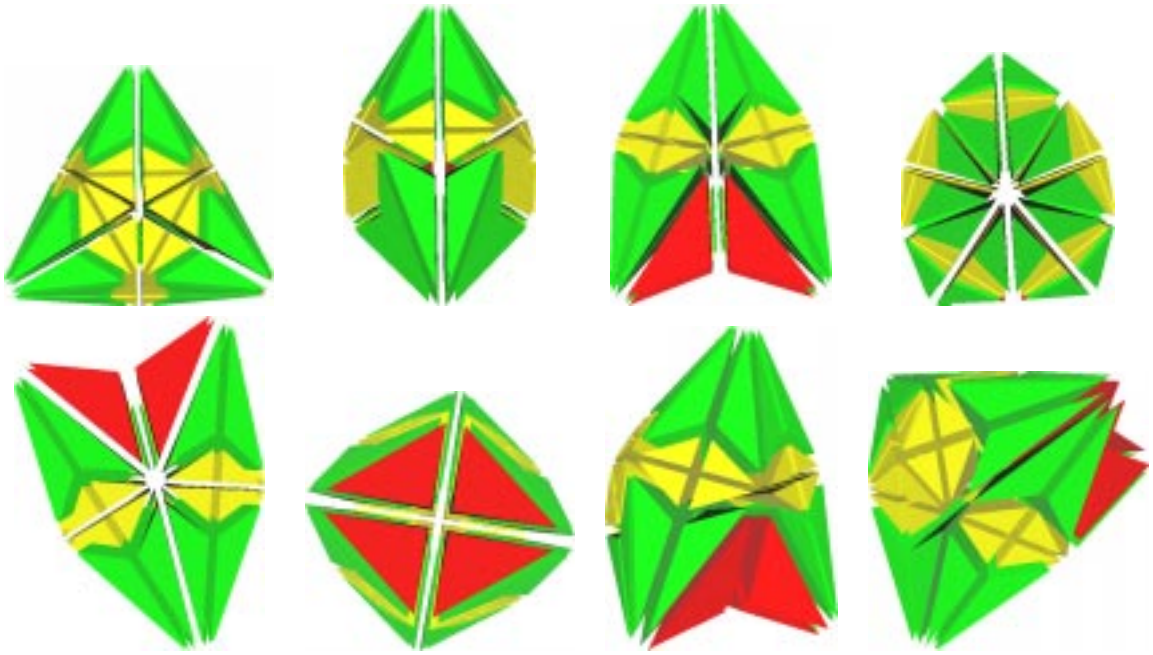

Figure A.15: Configuration 80. Composition: $4 \cdot A_1$, $32 \cdot B_2$, $44 \cdot K_2$.

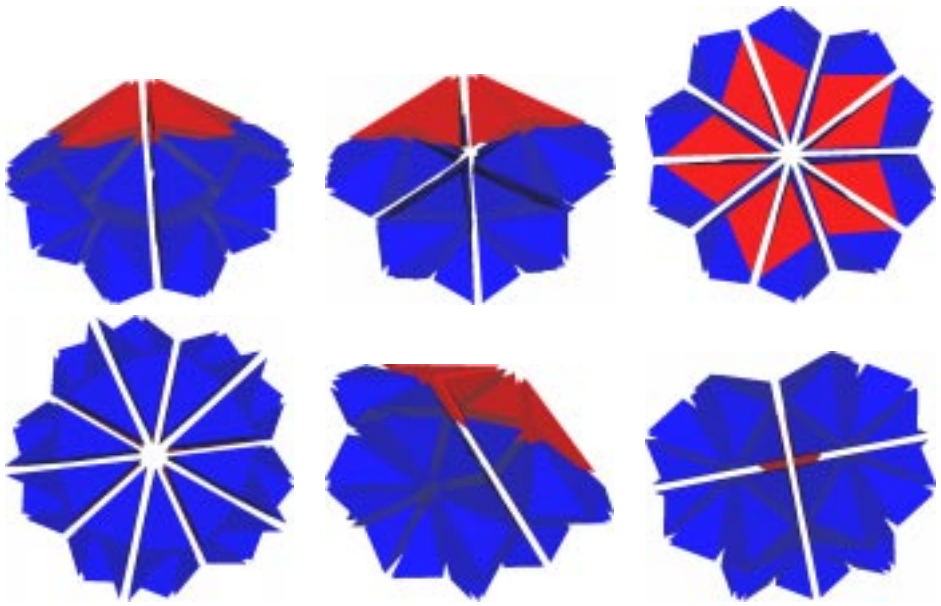

Figure A.16: Configuration 90a, the Half-discus. Composition: $10 \cdot A_2$, $80 \cdot C_3$.

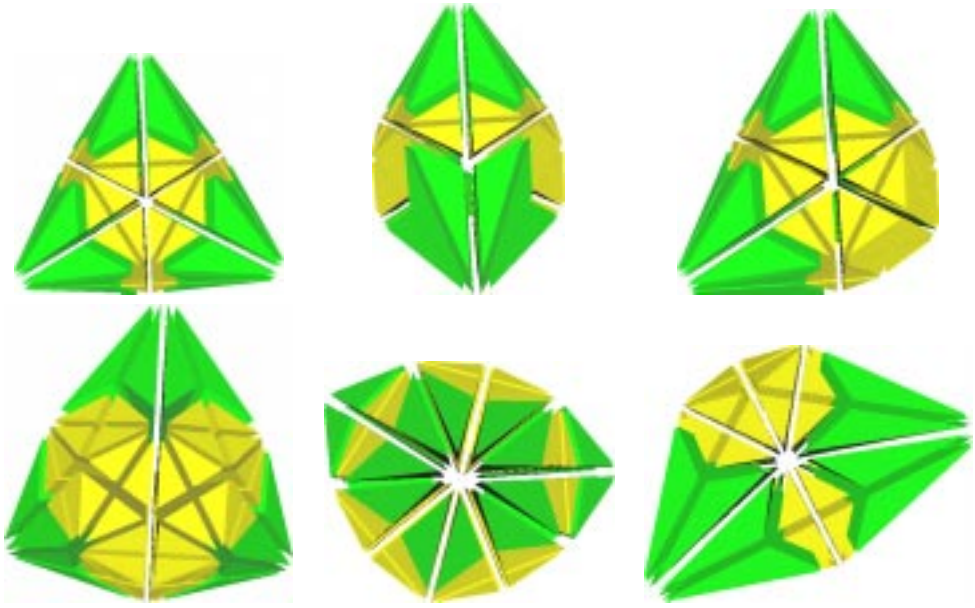

Figure A.17: Configuration 90b. Composition: $30 \cdot B_2$, $50 \cdot K_2$.

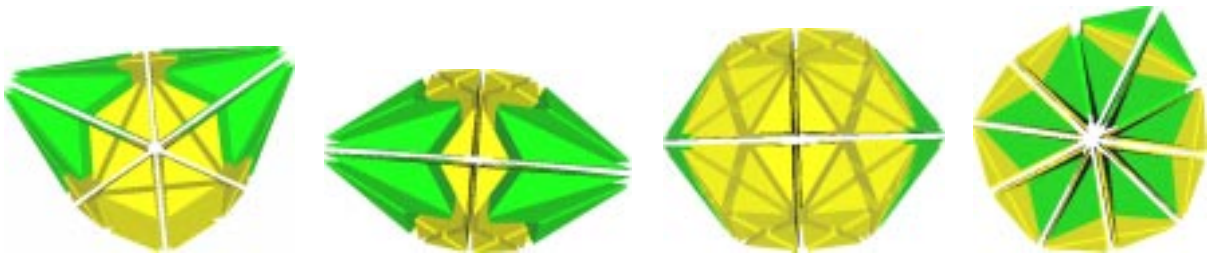

Figure A.18: Configuration 100, the Princess. Composition: $20 \cdot B_2$, $80 \cdot K_2$.

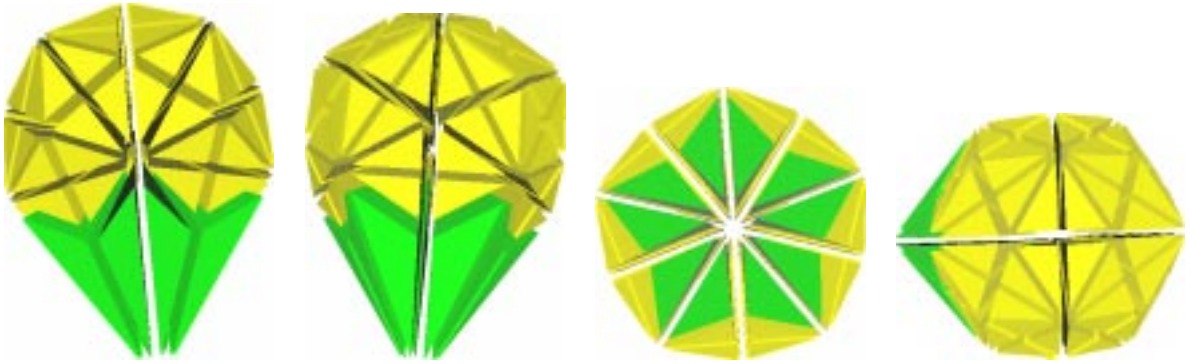

Figure A.19: Configuration 110, the Cone. Composition: $10 \cdot B_2$, $100 \cdot K_2$.

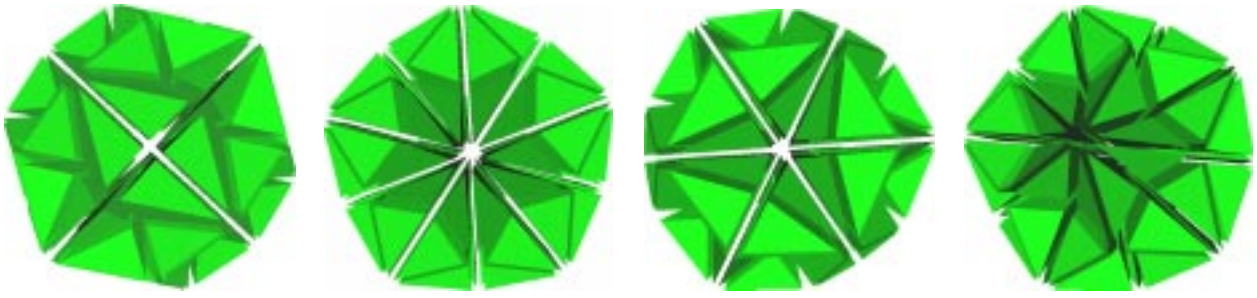

Figure A.20: Configuration 120a, the B-ball. Composition: $120 \cdot B_3$.

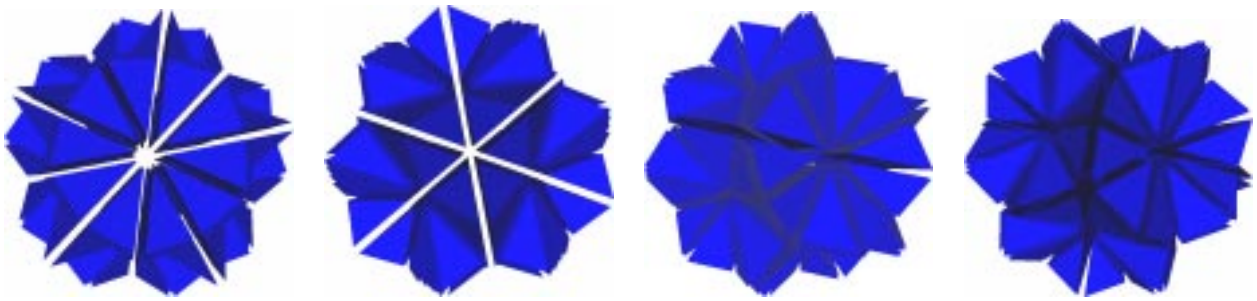

Figure A.21: Configuration 120b, the C-ball. Composition: $120 \cdot C_3$.

Figure A.22: Configuration 120c, the K-ball. Composition: $120 \cdot K_2$.

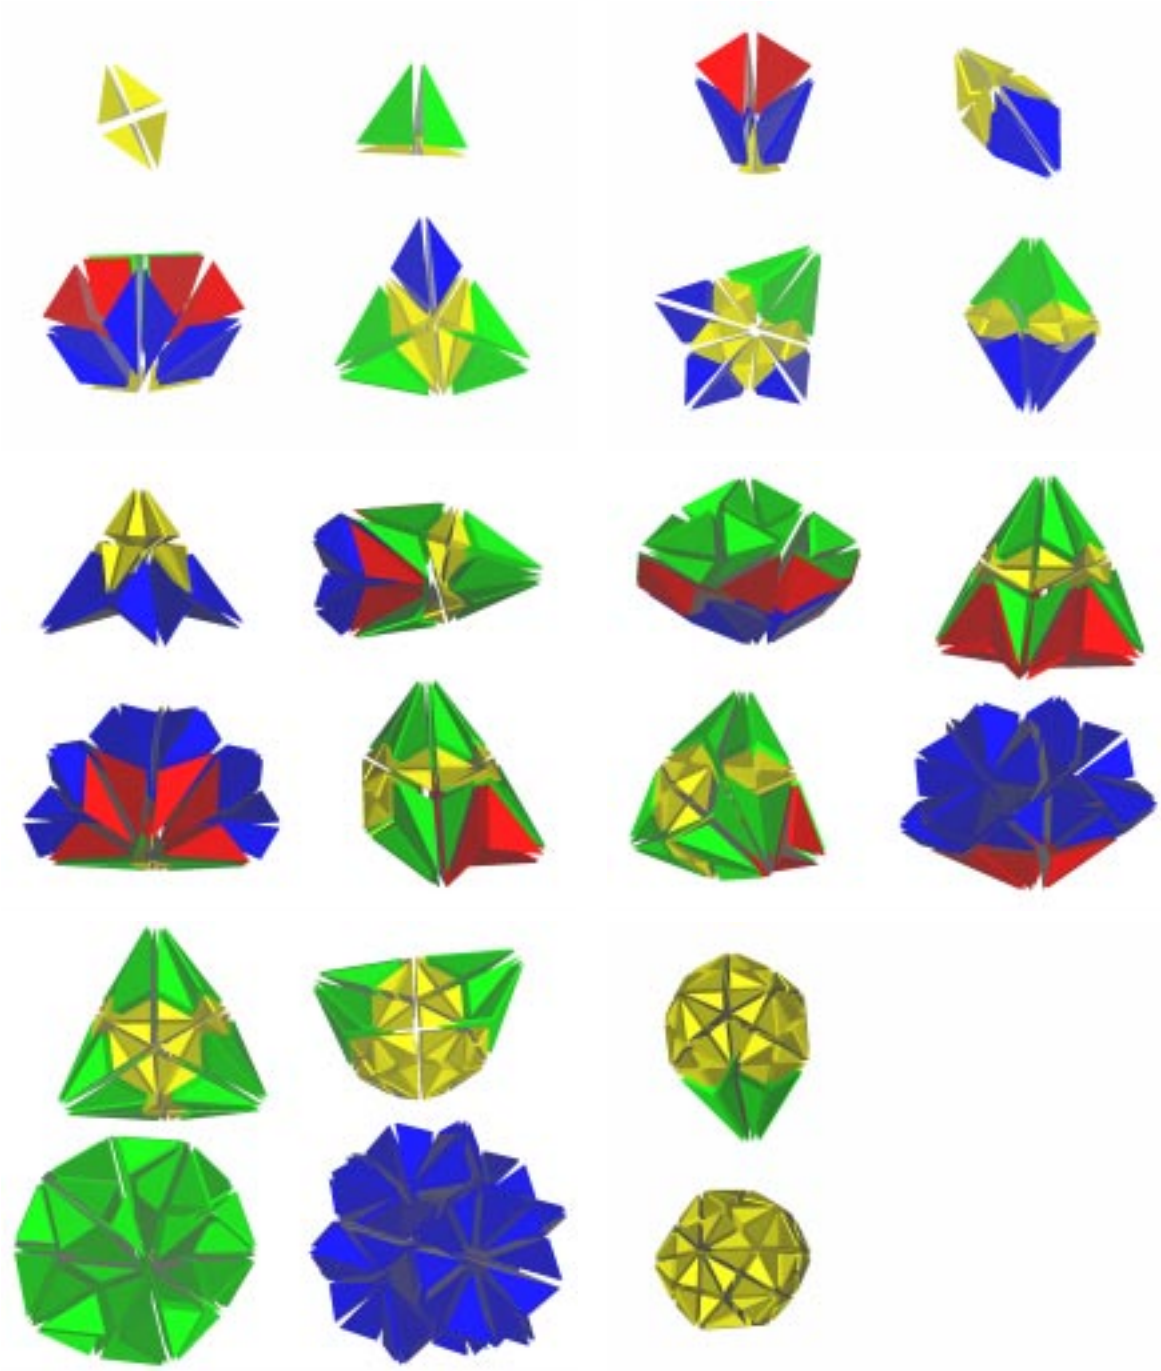

Figure A.23: All 22 global configurations drawn to scale.

Appendix B

New Vertex Configurations

This appendix contains pictures of the non-global vertex configurations constructed by the algorithm described in Chapter 5. There are 152 artificial configurations, which fall into three major groups. Two of the groups, containing 21 and 121 members, consist of many similar vertex configurations resulting from permutations of a few local replacements. We give pictures of one representative from each of these groups. The remaining 10 configurations are distinctive, and we give pictures of all of them.

The pictures were rendered on a Silicon Graphics Indigo2 using powerflip, a demo shipped with the operating system. A tiles are drawn in red, B tiles in green, C tiles in blue, and K tiles in yellow. In these pictures, left-handed tiles are shaded darker than their right-handed counterparts. Each tile has been shrunk by 30% toward its center of mass to allow us to see inside the configuration. While each individual image is an accurate rendering, scales are not the same between figures or between images within a figure. Figure B.13 shows all the configurations drawn to scale.

The caption for each figure gives how many of each prototile are used in the configuration, and the number of the central vertex. Every configuration has an equal number of left- and right-handed tiles, so for brevity we combine them. For example, $4 \cdot K_3$ means that the configuration contains two K tiles and two k tiles, and that for all four of those tiles vertex 3 is in the middle of the configuration.

The data files used to generate these pictures are available in powerflip, POV, and raw formats at http://www.cs.williams.edu/tiling/ for those readers interested in viewing these objects more dynamically.

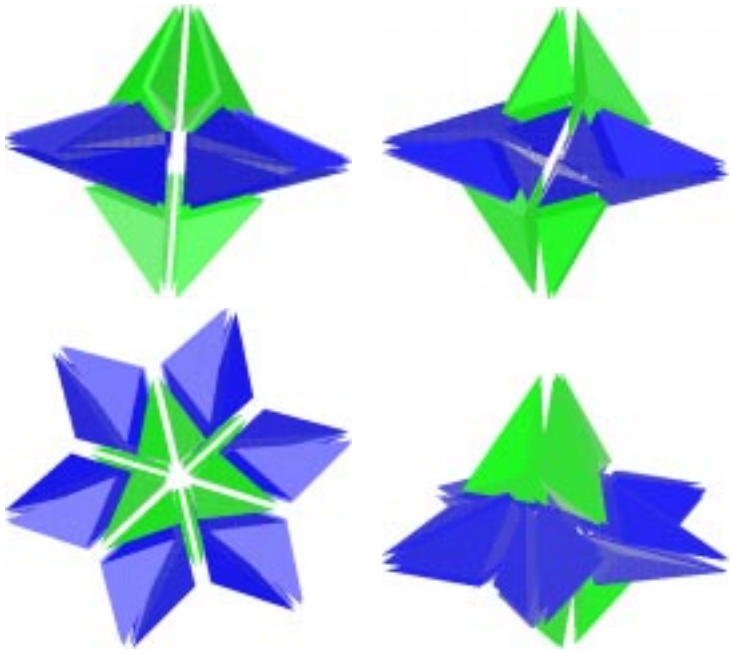

Figure B.1: Configuration 36n, the Shuriken. Composition: $12 \cdot B_1$, $20 \cdot C_1$. Contains the $B3 - C3$ face pairing, and so cannot be extended (see Section 3.2).

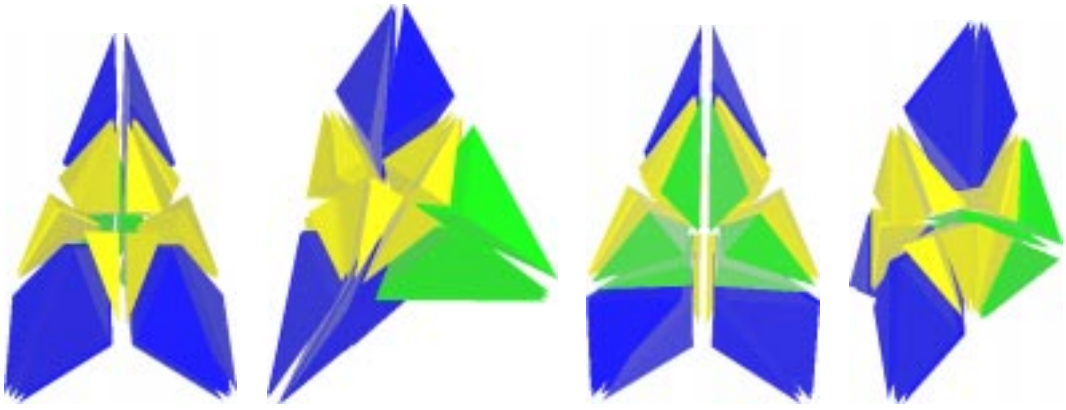

Figure B.2: Configuration 38n. Composition: $6 \cdot B_1$, $12 \cdot C_1$, $20 \cdot K_1$. Contains the $B3 - C3$ face pairing, and so cannot be extended.

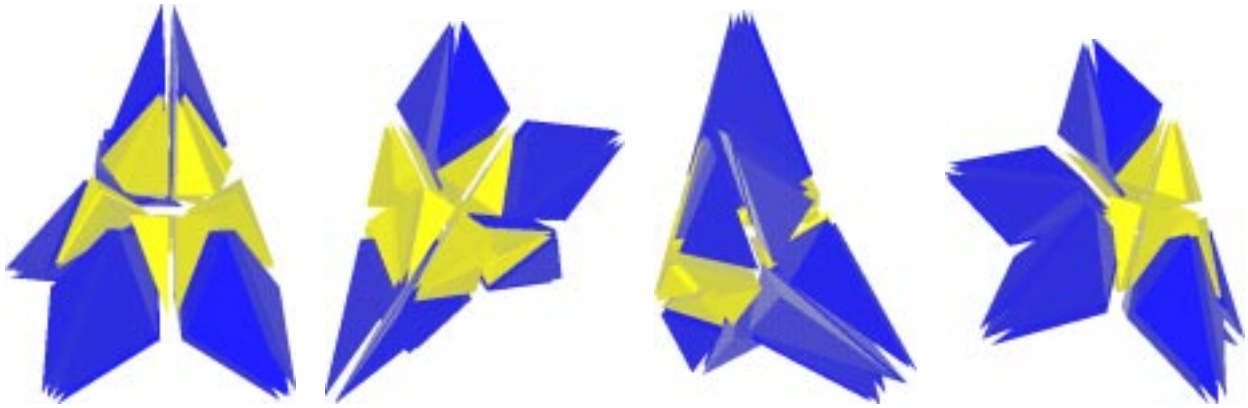

Figure B.3: Configuration 40na. Composition: $20 \cdot C_1$, $20 \cdot K_1$. Has no plane of symmetry; mirror image of 40nb.

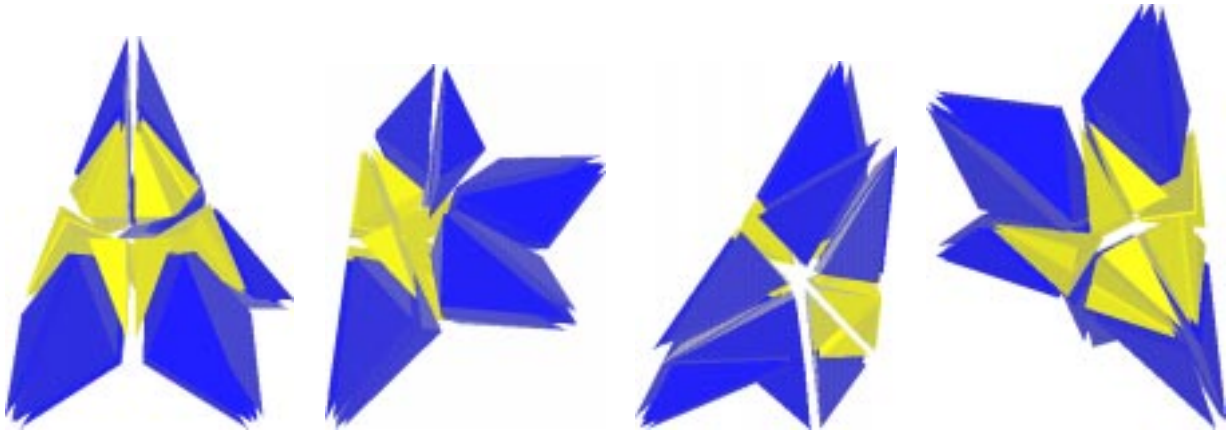

Figure B.4: Configuration 40nb. Composition: $20 \cdot C_1$, $20 \cdot K_1$. Has no plane of symmetry; mirror image of 40na.

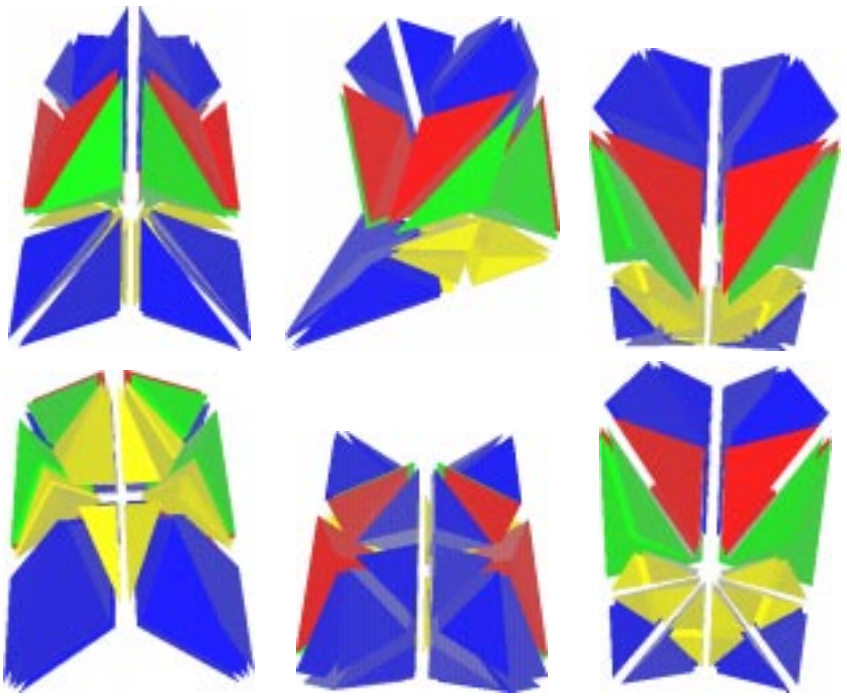

Figure B.5: Configuration 44n, the Snowplow. Composition: $6 \cdot A_2$, $6 \cdot B_1$, $8 \cdot C_1$, $12 \cdot C_3$, $12 \cdot K_1$. The only configuration to use two different vertices of one type of tile (C_1 and C_3), and thus the only configuration besides 8a to contain all the members of its vertex class. Contains face pairing $B3 - C3$, so cannot be extended.

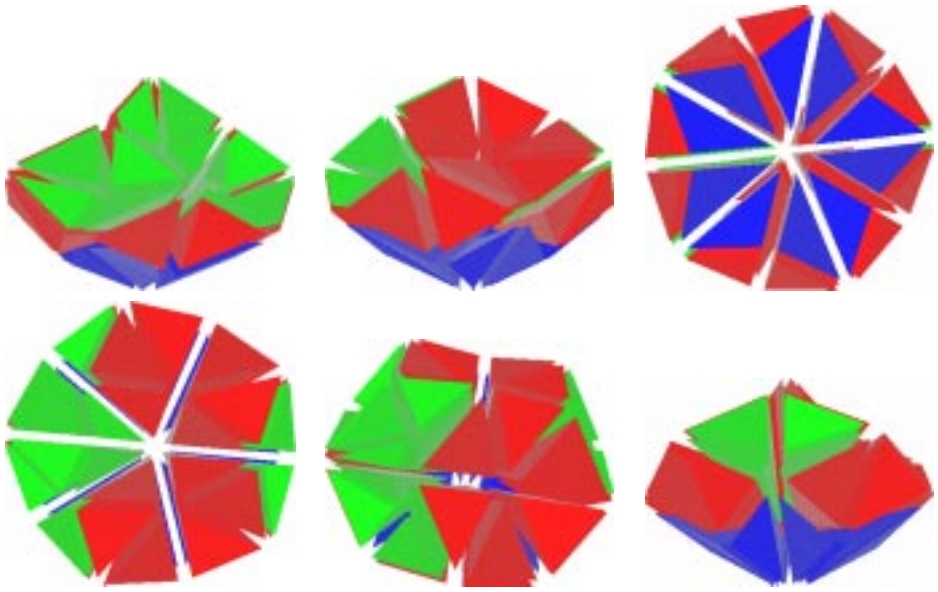

Figure B.6: Configuration 48n. Composition: $22 \cdot A_3$, $16 \cdot B_3$, $10 \cdot C_2$. A modification of 60a with two of the same replacements used to form Group 2. Contains the $A1{-}a1$ face pairing, so cannot be extended.

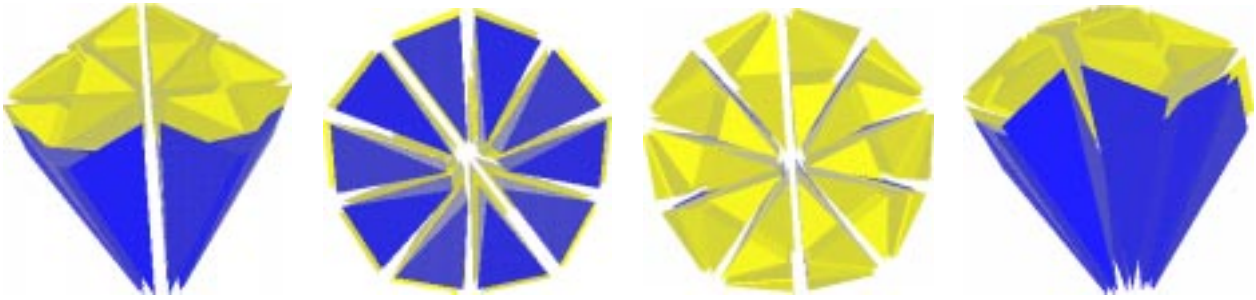

Figure B.7: Configuration 50n, the Torch. Composition: $10 \cdot C_4$, $40 \cdot K_2$.

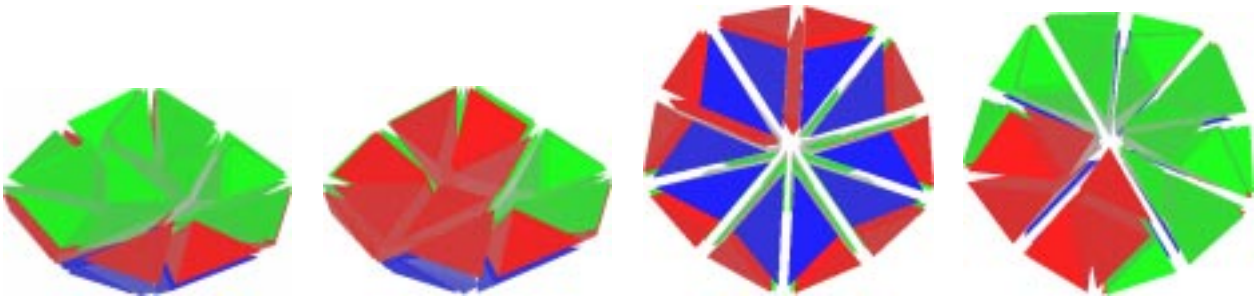

Figure B.8: Configuration 54n. Composition: $16 \cdot A_3$, $28 \cdot B_3$, $10 \cdot C_2$. A modification of 60a with 12 B tiles replaced by one divot of 6 A tiles. Contains the face pairing $A1 - a1$, so cannot be extended.

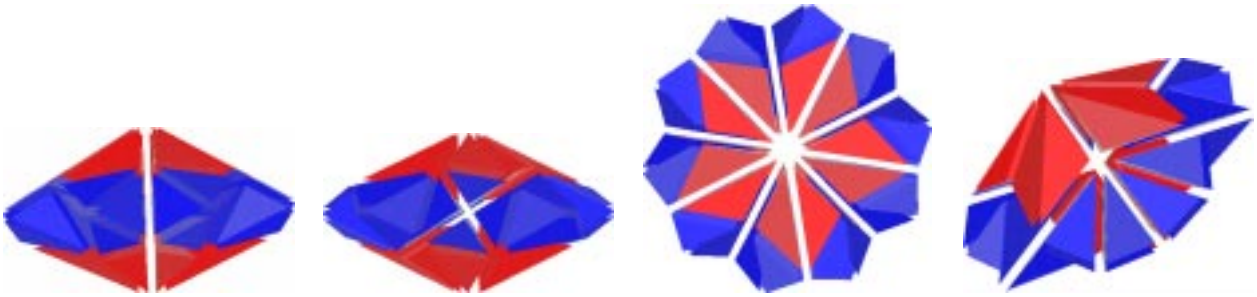

Figure B.9: Configuration 60n, the Discus. Composition: $20 \cdot A_2$, $40 \cdot C_3$. A natural extension of configuration 90a.

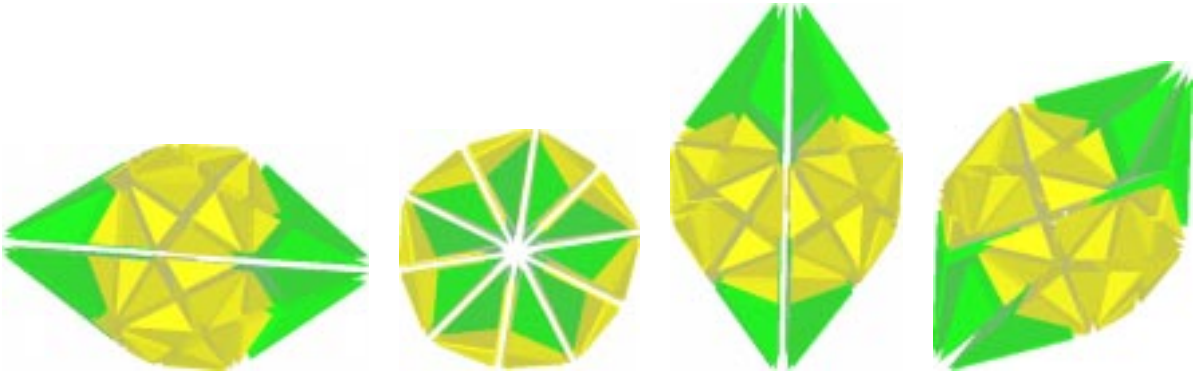

Figure B.10: Configuration 100n, the Prince. Composition: $20 \cdot B_2$, $80 \cdot K_2$. A modification of the Princess, configuration 100.

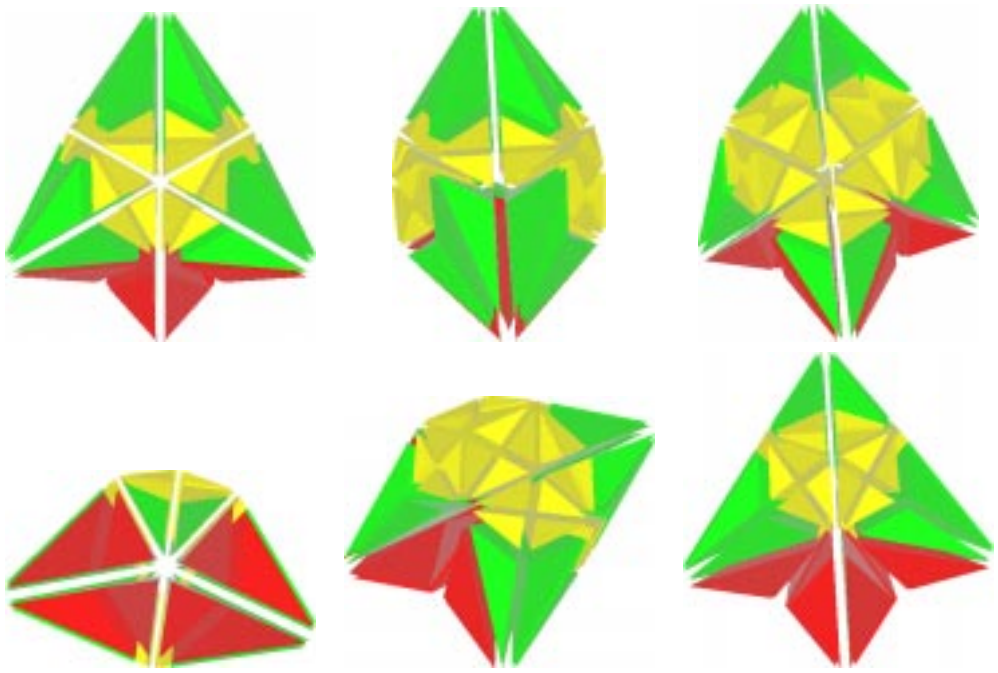

Figure B.11: Group 1 representative, one of 21. See global configurations 60b, 70, 80, 90b, 100, and 110 for other examples of this class of configurations.

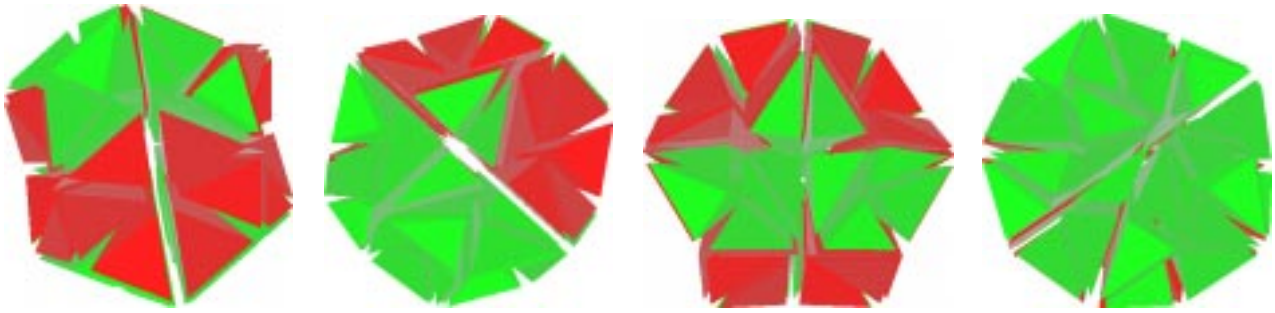

Figure B.12: Group 2 representative, one of 121. Formed by replacing groups of 12 B tiles in the B-ball with divots of 6 A tiles. The divot cannot be extended, so no Group 2 configuration can be embedded in a tiling.

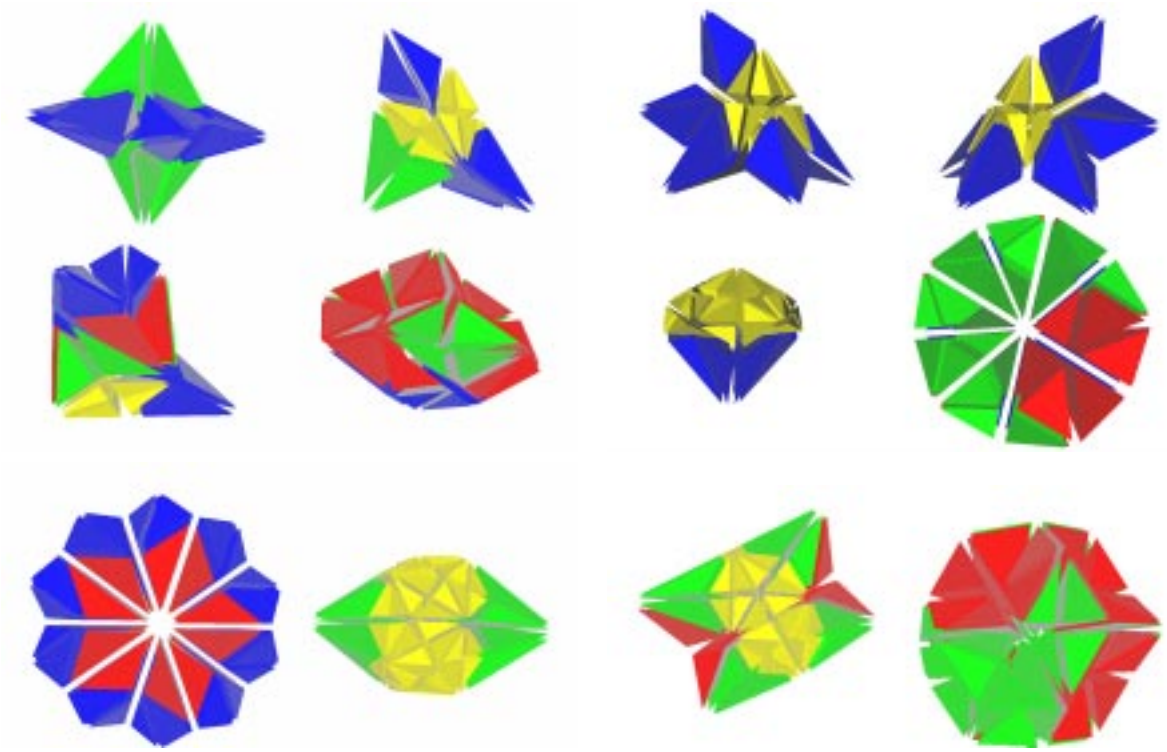

Figure B.13: The 10 distinctive new configurations and a representative from each of the other groups drawn to scale. Left to right, top to bottom, they are: 36n, 38n, 40na, 40nb, 44n, 48n, 50n, 54n, 60n, 100n, Group 1, Group 2.
The scale is the same as in Figure A.23. The group representatives are not the same as those in Figures B.11 and B.12.

Bibliography

[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.

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

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

[Gum96] Petra Gummelt. Penrose tilings as coverings of congruent decagons. *Geometriae Dedicata*, 62(1):1–17, August 1996.

[JS94] Hyeong-Chai Jeong and Paul J. Steinhardt. Cluster approach for quasicrystals. *Physical Review Letters*, 73(14):1943–1947, October 1994.

[JS97] Hyeong-Chai Jeong and Paul Steinhardt. Constructing penrose-like tilings from a single prototile and the implications for quasicrystals. *Physical Review B*, 55(6):3520–3532, February 1997.

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

[Min98] Linden Minnick. Generalized forcing in aperiodic tilings. Technical report, Williams College Department of Computer Science, 1998.

[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.

[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.

[Ste96] Paul Steinhardt. New perspectives on forbidden symmetries, quasicrystals, and penrose tilings. *Proceedings of The National Academy of Sciences Of The United States of America*, 93:14267–14270, December 1996.

[STT97] K. Saitoh, K. Tsuda, and M. Tanaka. Structural models for decagonal quasicrystals with pentagonal atom-cluster columns. *Philosophical Magazine A*, 76(1):135–150, 1997.

[Tre96] Forrest Trepte. Exploring aperiodic tilings with inflation. Technical report, Williams College Department of Computer Science, 1996.