

The Four Color Theorem meets Shapes of Polyhedra

Richard Evan Schwartz *

April 8, 2026

Abstract

We consider solutions to the 4-color problem on sphere triangulations with degree sequence $6, \dots, 6, 4, 4, 4, 4, 4, 4$. We sort these solutions into combinatorial types and show that each generic type is parametrized by the set of integer points inside a 4-dimensional rational polyhedral convex cone that is equipped with an integral quadratic form. We relate this structure to the octahedral stratum of Thurston's moduli space of flat cone structures on the sphere.

1 Introduction

1.1 Context

Here are two formulations of the famous Four Color Theorem.

1. You can color the vertices of any sphere triangulation with 4 colors in such a way that no two adjacent vertices get the same color.
2. You can color the triangles of any sphere triangulation black and white, so that the numbers of black and white triangles incident to each vertex are congruent to each other mod 3.

* Supported by N.S.F. Grant DMS-2505281

Property (1) defines a piecewise affine map ϕ from the sphere to the regular tetrahedron whose vertices have been given the same 4 colors. Color a triangle in the sphere triangulation white if and only if ϕ is orientation preserving on that triangle. This gives Property (2). To go from Property (2) to Property (1) you define the map ϕ using the orientation data and then pull back the vertex coloring of the tetrahedron. The mod 3 condition guarantees that ϕ is well defined.

See [W] for a lot of general information and history about the Four Color Theorem.

Certain families of triangulations closely resemble moduli spaces, and Property (2) above gives a way to study solutions to the 4 color problem on infinite collections of triangulations simultaneously. This paper will explore this perspective in an interesting special case, that of triangulations having degree sequence $6, \dots, 6, 4, 4, 4, 4, 4, 4$. In some sense, this paper is a sequel to my paper [S1], which explores Property 2 relative to triangulations having degree sequence $6, \dots, 6, 3, 3, 3$. However, the aims of this paper are different, and one can read this paper independently from [S1].

A sphere triangulation has *non-positive curvature* if there are at most 6 triangles incident to each vertex. If you build such a triangulation out of equilateral triangles you produce a metric space Σ called a *flat cone sphere*. All but finitely many points of Σ are locally isometric to the plane, and the remaining points, corresponding to vertices of degree less than 6, are locally isometric to Euclidean cones. In the case we are interested in here, Σ has 6 cone points, all having cone-angle $4\pi/3$. That is, Σ is intrinsically an equiangular octahedron.

In his paper *Shapes of Polyhedra* [T], Bill Thurston organized the triangulations of non-negative curvature along these lines. Thurston fixes the list of k cone points and corresponding cone angles, and then considers flat cone spheres with these cone points. The result is a moduli space \mathcal{M} of flat cone spheres having prescribed (and, for ease of discussion, labeled) cone points. The space \mathcal{M} is locally modeled on $\mathbf{C}^{1,k-3}$, a complex vector space with a Hermitian form of signature $(1, k-3)$. Moreover, \mathcal{M} has a system of coordinate charts in which the points with coordinates in the Eisenstein lattice correspond to triangulations. When one considers these flat cone spheres modulo complex scaling, which amounts to projectivizing the moduli space, the resulting space is a complex hyperbolic manifold of complex dimension $k-3$.

From now on we set

$$\mathcal{M} = \mathcal{M}(4, 4, 4, 4, 4, 4). \tag{1}$$

The space \mathcal{M} is locally modeled on $\mathbf{C}^{1,3}$. The projectivized space $P\mathcal{M}$ is an open complex hyperbolic 3-manifold whose metric completion is the quotient of \mathbf{CH}^3 by a certain lattice known as a *Mostow-Deligne lattice*.

1.2 Nice Colorings

We say that a *nice polygon* is any convex polygon whose interior angles all have the form $k\pi/3$ for $k \in \{1, 2\}$. A nice polygon has at most 6 sides. There are 5 kinds: triangles, parallelograms, trapezoids, pentagons and hexagons. We call a nice polygon *unit triangulable* if it is tiled by equilateral triangles having unit integer side lengths. For instance, the regular hexagon, chosen to have unit side lengths, is unit triangulable.

We say that a *nice coloring* of a flat cone sphere Σ in \mathcal{M} is a partition of Σ into nice polygons, alternately colored black and white, such that

1. There are 4 polygons around each regular vertex.
2. There are 2 polygons around each cone vertex.

Here the *regular vertices* are the vertices of the partition that are contained in disk neighborhoods. We call the nice coloring *unit triangulable* if the polygons in the partition are all unit triangulable. In this case, the coloring gives a solution of the 4 coloring problem for the corresponding sphere triangulation.

Figure 1.1 shows an example of a nice coloring. In this example, the edges of the polygon glue together to make the flat cone sphere. Each boundary edge is paired isometrically to another boundary edge in a direction-reversing way. For instance, the rightmost two edges are glued together to create one of the cone points. Figure 2.1 shows how the gluing works in this example.




Figure 1.1: A nice coloring of a flat cone sphere.

Figure 1.2 shows the example from Figure 1.1 and three other examples having the same combinatorics and gluing pattern.




Figure 1.2: Four combinatorially identical examples

The examples shown in Figure 1.2 come from a 4-parameter family of examples, naturally parametrized by a 4-dimensional convex cone $\mathcal{C} \subset \mathbf{R}^E$, where E is the number of edges. The cone \mathcal{C} is a rational polyhedral convex cone. The 4-dimensional linear subspace extending \mathcal{C} intersects \mathbf{Z}^E in a co-compact lattice. The points of $\mathcal{C} \cap \mathbf{Z}^E$ give rise to unit triangulable nice colorings, all having the same combinatorial structure. In this way, the same combinatorial pattern describes an infinite family of solutions to the 4 color problem. The projectivized space $P\mathcal{C}$ is a 3-dimensional convex polytope, and it has a dense set of points corresponding to triangulable nice colorings.

Remark: The precise local conditions we gave for a nice coloring do not cover all possible colorings which correspond to solutions of the 4-color problem. The other kinds of colorings by nice polygons (with more general kinds of local pictures) also come in families, but they are lower dimensional. This is why we say that the 4-color solutions corresponding to nice colorings are the generic kind. For the sake of simplicity, we do not treat the other kinds of 4-coloring solutions even though we could use similar methods.

1.3 Main Results

Our first result generalizes the discussion surrounding the example from Figure 1.2.

Theorem 1.1 *Let C denote any nice coloring of a flat cone octahedron. Let E denote the number of edges in the coloring. The set of nice colorings combinatorially equivalent to C is parametrized by a rational polyhedral convex cone $\mathcal{C} \subset \mathbf{R}^E$. The subset of triangulable nice colorings combinatorially equivalent to C is parametrized by $\mathcal{C} \cap \mathbf{Z}^E$.*

In §3, we will give a concrete recipe for constructing \mathcal{C} given the combinatorics of C . The construction can be implemented by computer, and in fact my computer program implements it.

The nice colorings seem to be quite abundant, and some are quite beautiful. In §2 I will show a lot of examples. The two most interesting examples I found are each indexed by the set of finite integer sequences a_1, \dots, a_k where $a_i \geq 2$ for all $i = 1, \dots, k$, and $k \geq 1$ can be as large as we like. I have not tried to prove rigorously that these families exist, but the numerical evidence is overwhelming. The existence of these families would show that the number of combinatorially distinct nice colorings with $2n$ polygons grows at least exponentially in n .

We call \mathcal{C} a *cone of colorings*. We set $\mathcal{C}_{\mathbf{Z}} = \mathcal{C} \cap \mathbf{Z}^E$. Our next result relates \mathcal{C} to Thurston's moduli space \mathcal{M} . Just as Thurston has a natural Hermitian form on his moduli spaces, we will have a natural quadratic form. Our quadratic form is very reminiscent of the one in [BG].

Theorem 1.2 *Let \mathcal{C} be a cone of colorings. There is an integral quadratic form Q on \mathcal{C} and a locally affine map $A : \mathcal{C} \rightarrow \mathcal{M}$. For each $V \in \mathcal{C}_{\mathbf{Z}}$, the expression $Q(V, V)$ computes 3 times the number of triangles in the triangulation associated to V . The map A is locally injective if and only if Q is a non-degenerate form. If Q has signature $(1, 3)$ then the projectivized map PA from PC to PM is locally injective.*

The quadratic form Q turns out to be the pullback of the real part of Thurston's Hermitian [T] form when it is suitably normalized. We will give a direct formula for Q in §4.5. The formula is completely combinatorial; it does not depend on the geometry of any particular nice coloring associated to \mathcal{C} . Theorem 1.2 would be more powerful with the following conjecture.

Conjecture 1.3 *The quadratic form Q always has signature $(1, 3)$. In particular, Q is always non-degenerate.*

The conjecture has some nice payoffs. First, if Q has signature $(1, 3)$ then Q imparts a real hyperbolic structure to PC . Second, the conjecture combines with Theorem 1.2 to show that the polyhedron PC always maps in a locally injective way into PM . Conjecture 1.3 lines up with the situation in [BG] and [T], but I don't think that the conjecture follows from anything in those papers. In any given case I can compute Q in a basis for the subspace spanned by \mathcal{C} . In all the many cases I have checked, the signature is $(1, 3)$.

Here is a related conjecture.

Conjecture 1.4 *The map $PA : PC \rightarrow PM$ is (globally) injective.*

This paper is still somewhat preliminary. In any case, there is plenty more to say. Here are some additional topics I might add, either to this paper or to a sequel.

- Explicit calculations of the quadratic form on some infinite families of examples. In some of the families, and with respect to a suitably chosen family of bases, the sequence of quadratic forms converges to a limiting form, and the limit (also) has signature $(1, 3)$.
- A discussion of the other combinatorial kinds of colorings of a flat cone octahedron by nice polygons.
- A discussion of how the various convex cones from Theorem 1.1 fit together, and relatedly how the nice coloring degenerates when we move to a face of the corresponding convex cone.

Also, I still need to make the software for this paper public. Right now, it lacks a lot of documentation. I hope to fix this within a few weeks.

1.4 Overview

This paper is organized as follows. In the beginning of §2 I will define the basic combinatorial objects of study, which is essentially the multigraph dual to the nice coloring. Following this, in §2.2, I will exhibit the infinite family of nice colorings. I essentially give a proof, using a grafting idea, but I stop short of giving all the details. After this, I will present a gallery of other examples of nice colorings. In §3 I will prove Theorem 1.1. In §4 I will prove Theorem 1.2.

2 A Gallery of Examples

2.1 The Dual Multigraph

In this section we explain how to encode our nice colorings as a multigraph. For us, an *octahedron* is always a flat cone sphere with 6 cone points, all having cone angle $4\pi/3$. Also, our *graphs* are allowed to have multiple edges between vertices. Sometimes I will say *graph* and sometimes *multigraph*.

Let P be an octahedron with a nice coloring. We form a *dual multigraph* G as follows. There is one vertex of G for each polygon in the triangulation. We join two vertices by one edge for each of the edges they have in common. We draw this graph on P by choosing the vertices to be the barycenters of the faces and by drawing each (blue) edge of the graph so that it only crosses the corresponding edge of the coloring of P . Figure 2.1 shows an example. Here, the parts of the blue arcs which stray outside the black and white polygon would disappear when the polygon is glued up to make the octahedron.




Figure 2.1: The dual multigraph

This construction gives G the structure of an embedded planar multigraph. The regions of G bounded by bigons are in bijection with the cone points of P . So, there are 6 bigon regions in total. The remaining regions are quadrilaterals, and they are naturally in bijection with the ordinary vertices of P . Figure 2.1 shows the example that corresponds to Figure 1.1.

Each quadrilateral face has a distinguished edge e , the one corresponding to the two acute angles around the vertex corresponding to the face. We add a red edge to G just inside the face and essentially parallel to e . Figure 1.1 shows these edges drawn in red. When we add in all these red edges we get what we call the *enhanced multigraph*, and we denote it by \widehat{G} .

Lemma 2.1 *Each vertex in \widehat{G} has 6 edges incident to it.*

Proof: Each vertex v in the multigraph G , corresponding to a nice k -gon P , has k blue edges incident to it, one per side of P . We want to see that there are $6 - k$ additional red edges incident to v . We call a vertex w of P *acute* if its interior angle at w is acute.




Figure 2.2: The dual multigraph

We split each blue edge e incident to v into two infinitesimally close *half-edges*. Figure 2.2 shows this schematically. Thus, there are $2k$ half-edges emanating from v and they are cyclically ordered. Each half edge e' defines a unique vertex w of P . One can connect e' to w by an arc that crosses no other blue edges of G . We call two half-edges *partners* if they are associated to the same vertex w of P . We call these partners *acute* if the interior angle at w is acute. Note that there are exactly $6 - k$ acute partners.

Given any pair (e'_1, e'_2) of acute partners, let w be the associated vertex. There are three other nice polygons incident to w , and exactly one of them has w as an acute vertex. Thus, exactly one of our half edges in the set $\{e'_1, e'_2\}$ has a parallel red edge. This gives us one red edge for each pair of acute partners. Hence there are $6 - k$ red edges incident to v . ♠

2.2 A Simple Infinite Family

Figure 2.3 shows the beginning of a family of cell divisions of the sphere into quadrilaterals. Our convention is that the outside of the big square is also part of the cell division.




Figure 2.3: A family of cell divisions of the sphere

Thus, these cell divisions respectively have 2, 3, 4, 5 faces. Likewise they have 4, 5, 6, 7 vertices. We are interested in the ones which have an even number of vertices. These are bipartite. Based on the even members of this family, we create red/blue multigraphs $\widehat{G}_6, \widehat{G}_8, \widehat{G}_{10}, \dots$ as shown in Figure 2.4.




Figure 2.4: The graphs $\widehat{G}_6, \widehat{G}_8, \widehat{G}_{10}$

The continuation of the pattern may not be entirely clear, so we say more about this. There is a spiral path in G_{2k} which starts at the bottom left vertex and moves around until it reaches the last placed vertex. We have highlighted this path in yellow in each case. We put double red edges along this path, starting with the bottom edge and then continuing in a way that skips over every other edge. We then add 6 more blue edges to make all the degrees 6.

Figure 2.5 shows two nice colorings associated to \widehat{G}_{16} . For the one on the left, we have indicated the gluing pattern both by letters and by some red zig-zags. The one on the right has the same gluing pattern. We included both

pictures to indicate somewhat how the shapes change with the parameter in the cone \mathcal{C} guaranteed by Theorem 1.1.

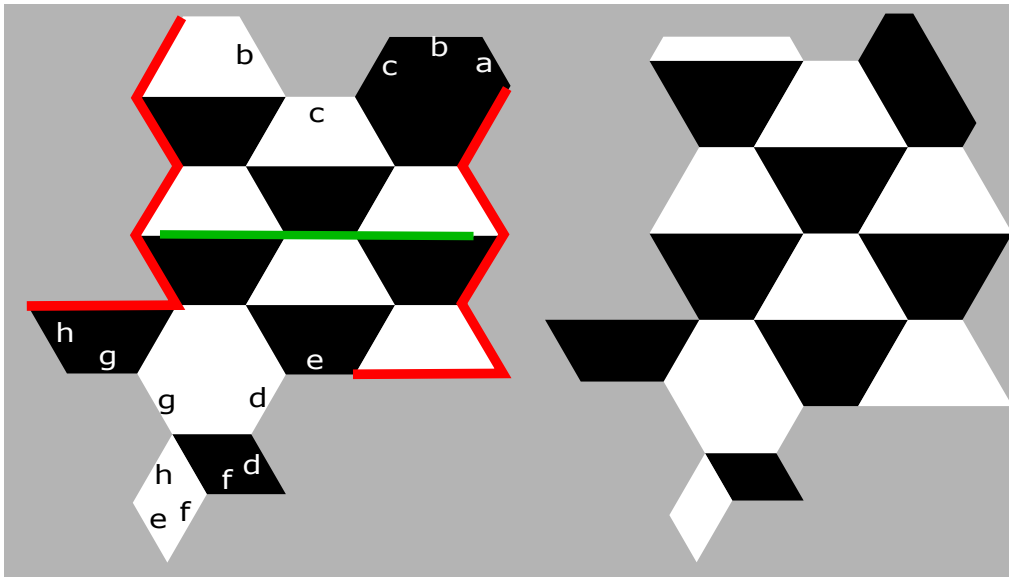


Figure 2.5: A nice coloring associated to \widehat{G}_{16} .

Notice the red zigzags on either side of the figure on the left. These are glued together by a translation. This creates a flat cylinder inside the octahedron. We can lengthen this cylinder by grafting in a 2×3 layer of trapezoids. We slit open the coloring along the green segment, then make the graft using trapezoids which have precisely the same shape as the ones above and below the green slit. We then modify the gluings so that longer red zigzags are identified by a translation. This produces a nice coloring with a longer cylinder. This bigger nice coloring corresponds to \widehat{G}_{22} . This gives us a provably infinite family of nice colorings.

The rest of this chapter presents a number of examples without proof. In each case, we will show one nice coloring and its associated dual multigraph. We include this material to illustrate the richness of the construction.

2.3 A Checkerboard Family

Figure 2.6 shows one example from an infinite family. The gluing around the boundary is very much like the folding of a wallet or a taco. There is an orientation-reversing isometry of the boundary which fixes the two vertices

marked by the two dots, and then the rest of the gluing is implemented by this isometry. A grafting operation like the one presented in the previous section would construct infinitely many examples from this one.




Figure 2.6: A checkerboard themed example

Figure 2.7 shows the multigraph on which this example is based. The graph is made by gluing two identically decorated disks together with an orientation reversing isometry, in such a way that the arrowed edges on the upper left of each side are identified. A similar pattern like this exists for any grid with an even number of squares, though I did not explicitly check that all such graphs actually correspond to nice colorings.




Figure 2.7: The dual multigraph associated to Figure 2.6.

2.4 A Lamination Themed Family

Figure 2.8 shows two examples from what is presumably an infinite family. Unlike the previous two families, I don't see an easy way to prove that all the members in the family exist. These examples (to me) give a hint of geodesic laminations.




Figure 2.8: Nice colorings which hint at geodesic laminations

Figure 2.9 shows the dual multigraph associated to the left side of Figure 2.8. In this picture we glue together two squares in the pattern shown to

make the sphere. The lettering indicates the gluing pattern. (I made the identifications more explicit in this example so that the reader could trace out the paths made by the red edges.) For what it is worth, I have used several shades of red to indicate the 3 connected components of the red subgraph.




Figure 2.9: The multigraph corresponding to the left side of Figure 2.8

2.5 Two Staircase Families

I will illustrate the two big families with examples. The examples are based on the sequence 2, 3, 2, 4. These numbers control the successive lengths of a path of squares which always moves upwards and to the right. Figure 2.10 shows the dual graph for the corresponding member of the first family.




Figure 2.10: Member 2324 of the first staircase family

Figure 2.11 shows the corresponding nice coloring. One thing I should say about my program is that it does not necessarily lay out the most efficient gluing diagram. After computing the shapes, according to the method

described in the next chapter, the program lays the shapes down in a kind of random order, snapping them on to a growing union of shapes.




Figure 2.11: The coloring associated to Figure 2.10.

I am not showing the explicit folding pattern, but you can deduce it from the lengths of the edges on the boundary.

figure 2,12 shows the graph for the same member of the second family.




Figure 2,12: Member 2324 of the second staircase family

Figure 2.13 shows the corresponding nice coloring.




Figure 2.13: The coloring associated to Figure 2.12.

Again, I am not showing the explicit folding pattern.

Remark: In almost all the nice colorings I have presented, the dual graph is made by gluing together nearly identical disks. There are plenty of examples which do not have this feature, but I have not looked for infinite families based on two very different halves. The “two halves” approach is partly an artifact of my program. I have a feature where I make one side, then duplicate it, then let the program find the correct gluing of the disks to make a sphere.

3 Realizing Nice Colorings

3.1 Axioms for the Multigraph

Now we want to reverse the construction in §2.1. That is, we want to start with an enhanced multigraph and try to recover a partition of an octahedron into nice polygons. All our graphs will be planar multigraphs which divide the sphere into quadrilaterals and bigons. Our multigraphs will have red and blue edges. We say that the *blue multigraph* is the graph made entirely from the blue edges. We call a red edge and a blue edge *parallel* if they share the same two vertices. Here are the conditions we need:

1. The blue multigraph should divide the sphere into 6 bigons and some finite number of quadrilaterals.
2. Each quadrilateral face of the blue multigraph should contain one red edge that is parallel to one of the blue edges of the face.
3. The total number of edges incident to any vertex is 6.

We call a multigraph satisfying these conditions a *plausible multigraph*.

We call the plausible multigraph *nice* if it actually comes from a nice coloring of an octahedron. All the multigraphs shown in the previous chapter are nice. Given a plausible multigraph, we can extract the shapes that would be involved in a nice coloring of the sphere if such a thing existed. Each vertex v of blue degree k in the plausible multigraph corresponds to a nice polygon with k sides. The acute vertices correspond to faces in the (full) graph which are incident to red edges.




Figure 3.4: Getting the shape of the nice polygon from the graph

Figure 3.4 shows some local pictures of the multigraph superimposed over the corresponding shapes. This is not quite an exhaustive list of possibilities.

If \widehat{G} is a nice multigraph, we can assign a positive variable to each edge e of \widehat{G} , namely the length of the polygon edge that e crosses. If \widehat{G} is merely a plausible multigraph, we want to find out if there exists a positive assignment of variables to the edges which corresponds to a nice coloring of the sphere. To do this, we solve a system of linear equations and then see if the solution set intersects the positive quadrant. We now explain how this is done.

3.2 The Shape System

Let $\omega = \exp(\pi i/3)$. We can view a nice hexagon as a 6-tuple of positive numbers $\ell = (\ell_0, \dots, \ell_5)$ subject to the following constraint.

$$\ell \cdot (1, \omega, \omega^2, \omega^3, \omega^4, \omega^5) = 0. \quad (2)$$

Taking suitable linear combinations of the real and imaginary parts we get the equivalent relations:

$$\ell \cdot v_1 = \ell \cdot v_2 = 0, \quad v_1 = (1, 1, 0, -1, -1, 0), \quad v_2 = (0, 1, 1, 0, -1, -1) \quad (3)$$

The other nice polygons are limiting cases of nice hexagons, where we set some of the variables equal to 0. For instance, if we have a trapezoid whose long side is parallel to $\omega^0 = 1$ then we set $\ell_1 = \ell_5 = 0$ and then change the indices so that they are consecutive. This gives $\ell = (\ell_0, \ell_1, \ell_2, \ell_3)$ and

$$\ell \cdot v_1 = \ell \cdot v_2 = 0, \quad v_1 = (1, 0, -1, 0), \quad v_2 = (0, 1, 0, -1)$$

Lemma 3.1 *Suppose P is a nice polygon having all integer sides. Then P has a triangulation by equilateral triangles with integer side lengths.*

Proof: When P is a triangle the result is obvious. Say that a *basic piece* is a nice 4-gon with integer side lengths. If P is a basic piece our proof goes by induction on the sum of the side lengths of P . In this case we can chop off an integer equilateral triangle T so that $P' = P - \text{interior}(T)$ is a basic piece with smaller total side length. The left side of Figure 3.5 shows this.




Figure 3.5: An inductive proof

When P is a pentagon, we can chop off a basic piece and leave either an integer equilateral triangle or a basic piece. The idea here is to use a basic piece which shares the shortest side of P . When P is a hexagon we can chop off a basic piece and leave one of the pieces already considered, thereby proceeding inductively. The right hand side of Figure 3.5 shows how we reduce a nice hexagon to a trapezoid with two chops. ♠

Let \widehat{G} be a plausible multigraph, as above. To (try to) find the nice colorings of the sphere associated to \widehat{G} we set up a system of equations where we assign one real variable to each blue edge of \widehat{G} . The role played by the red edges is that they help determine the constraints.

We have E_b variables, where E_b is the number of blue edges of \widehat{G} . Each vertex of \widehat{G} gives two constraints on the edges, as discussed in the previous section. We represent the constraint as a vector of length E_b which is only nonzero in the positions corresponding to the edges incident to the vertex. We have $2V$ constraints, where V is the number of vertices of \widehat{G} . We call this system of equations the *shape system*. We call the shape system *nice* if it has a positive solution.

Now let's compare the number E_b of variables with the number $2V$ of constraints. We have

$$2F_b - 2E_b + 2V = 4, \quad 2F_b = E_b + 6, \quad \implies \quad E_b - 2V = 2. \quad (4)$$

The first equation is Euler's formula. The last equation, what we want, is a consequence of the first two equations. Here is how to derive the middle equation. Set $F'_b = F_b - 6$ and $E'_b = E_b - 6$. These respectively are the number of faces and edges we get when we collapse each bigon to a single edge. That is, we identify the edges bounding each bigon. In the collapsed graph we have $2F'_b = E'_b$, because each face is a quadrilateral and each edge is incident to 2 faces. Rewriting this in terms of F_b and E_b we get the middle equation in Equation 4.

3.3 The Solution Space

Equation 4 tells us that we have two more variables than equations. So, if we have one positive solution we have (at least) a 2-dimensional space of convex solutions. It turns out that the constraint vectors are not independent. In the next result we work with the complex notation, as in Equation 2. The

next result, stated in the complex notation, also says that there are 2 real relations amongst the constraint vectors.

Lemma 3.2 *The complex constraint vectors are linearly dependent.*

Proof: The key to proving this result is to find a good basis for the constraint vectors. We will give the proof under the assumption that there exists a positive solution to the nice system. The proof works in general, but in general it is harder to interpret the proof geometrically.

Here is the key idea: Let Σ be the flat cone sphere which comes from a positive solution of the shape system. There is a well defined *folding map* $f : \Sigma \rightarrow \mathbf{C}$. The map is such that f is an orientation preserving isometry on each white piece and an orientation reversing isometry on each black piece.

Once we stipulate f maps some particular edge into the real line, f gives us a consistent way to label every edge in the dual graph \widehat{G} by 6th roots of unity. We orient the edges of the nice polygons in Σ so that they go clockwise around the white faces. Once we do this, we assign the 6th root of unity $L(e)$ that is parallel to the image of the oriented edge e under f . We give this same label to the edge e' of \widehat{G} which crosses e .

Now the constraint function for any given vertex v in \widehat{G} is

$$\pm \sum_{j=1}^k L(e_j) \ell_j. \tag{5}$$

Here k is the number of blue edges incident to v and e_j is the j th such edge and $L(e_j)$ is the unit complex label of e_j . We choose (+) if v is a white vertex (corresponding to a white face) and (−) if v is a black vertex.

Every edge label appears with a (+) in one constraint vector and with a (−) in another constraint vector and otherwise not at all. Hence the sum of all the constraint vectors is 0. This is our relation. ♠

Remark: Even without a positive solution we could work out the labels of the edges of Σ . This is a purely algebraic construction, and the geometry only serves to interpret it.

The next result says that the relation from Lemma 3.2 is the only relation, up to scaling.

Lemma 3.3 *The shape system has real rank $E_b - 4$.*

Proof: To prove this we just need to see that the only relations amongst the constraint vectors, considered as complex vectors, are scalar multiples of the one from Lemma 3.2. Let R_1 be the relation from Lemma 3.2. Call the complex constraint vectors V_1, \dots, V_n . (Here $n = E_b - 2$.) The relation R_1 is $\sum V_i = 0$.

Suppose for the sake of contradiction we have another relation R'_2 that is not a scalar multiple of R_1 . Note that R_1 involves all the constraint vectors. By subtracting real multiples of R_1 from R'_2 we can find a new relation R_2 which does not involve all the constraint vectors. Call two constraint vectors V_i and V_j *adjacent* if they are associated to adjacent vertices in the dual graph \widehat{G} . After re-indexing if necessary, we can find adjacent constraint vectors V_1 and V_2 such that V_2 is involved in R_2 and V_1 is not.

Re-writing the relations R_1 and R_2 we have

$$V_1 = - \sum_{k \neq 1} V_k, \quad V_2 = \sum_{k \neq 1, 2} c_k V_k.$$

Here the constants c_k are not important to us; some could vanish. Plugging the second equation into the first we see that both V_1 and V_2 are linear combinations of the remaining constraint vectors. The other constraint vectors do not involve the edge e incident to the vertices corresponding to V_1 and V_2 . Thus, if the remaining constraint vectors are satisfied we can assign any variable we like to e and guarantee that V_1 and V_2 are also satisfied. But this is absurd. If we just vary the value of e , at least one of V_1 or V_2 changes its value. This contradiction finishes the proof. ♠

3.4 Putting it Together

Now we assemble the ingredients and prove Theorem 1.1. Suppose that we have some nice coloring of the sphere. We let \widehat{G} be the plausible multigraph that encodes the nice coloring. The shape system has E_b variables and, by Lemma 3.3, has rank $E_b - 4$. Hence the space of solutions to the shape system is 4-dimensional. Let \mathcal{S} denote the solution space.

The space \mathcal{S} is a 4-dimensional linear subspace of \mathbf{R}^{E_b} . Also, \mathcal{S} is defined by integer equations. So \mathcal{S} is a rational subspace. We are interested in the

subset where all variables are positive. This is the intersection of \mathcal{S} with the positive orthant, and this intersection is a rational convex polyhedral cone. This is our cone \mathcal{C} from Theorem 1.1.

Thanks to Lemma 3.1, each point of $\mathcal{C} \cap \mathbf{Z}^{E_b}$ corresponds to a triangulable nice coloring. Again, since \mathcal{C} is a rational cone, the linear subspace \mathcal{S} intersects \mathbf{Z}^{E_b} in a lattice.

This completes the proof of Theorem 1.1.

4 The Map and the Quadratic Form

In this chapter we prove Theorem 1.2, which gives a locally affine map from \mathcal{C} to \mathcal{M} and relates it to a quadratic form on \mathcal{C} .

4.1 Thurston Coordinates: A First Pass

We will first give an elementary account of Thurston's coordinates on the space \mathcal{M} . After this elementary account we will give a second account that is more conducive to the construction of our linear map. We give the first account mostly for expository purposes.

From Complex Numbers to Octahedra: As always, an *octahedron* for us is a flat cone sphere with 6 cone points having cone angle $4\pi/3$. First we describe how to build an octahedron given complex numbers z_0, z_1, z_2, z_3 . Once z_0 is chosen, the other three complex numbers need to be chosen with some care, but an open set of choices works. Let \mathbf{Eis} denote the usual Eisenstein lattice. This is the ring of integer combinations of 1 and $\omega = \exp(\pi i/3)$. The number z_0 specifies a lattice of points in the plane, namely $\Lambda = z_0 \mathbf{Eis}$. What we are doing here is scaling up the Eisenstein lattice by z_0 . We let Γ be the group generated by the order 3 rotations about the elements of Λ . The quotient \mathcal{C}/Γ is a flat cone sphere with 3 cone points with cone angle $2\pi/3$. This is a doubled equilateral triangle.

Let Γ_λ denote the order 6 group of rotations which fix $\lambda \in \Lambda$. The stabilizer of λ in Γ has index 2 in Γ_λ . Let $\lambda_1, \lambda_2, \lambda_3$ respectively be the points $0, z_0, \omega z_0$. Notice that the three orbits $\Gamma \lambda_k$ for $k = 1, 2, 3$ are pairwise disjoint, and their union is all of Λ .

Around λ_k we delete a hexagon h_k with the following description: One of the vertices of h_k is $\lambda_k + z_k$ and the others are the orbit of this one vertex under Γ_{λ_k} . We need to choose z_1, z_2, z_3 so that the three hexagon orbits Γh_k for $k = 1, 2, 3$ are pairwise disjoint, and indeed the grand union of all the hexagons under the action of Γ consists of disjoint hexagons. The sphere is

$$\Sigma(z_0, z_1, z_2, z_3) = (\mathcal{C} - \Gamma h_1 - \Gamma h_2 - \Gamma h_3)/\Gamma. \quad (6)$$

Why does this work? As mentioned above, the simpler quotient \mathcal{C}/Γ is a flat cone sphere with 3 cone points with cone angle $2\pi/3$. When we cut out each of the quotients $\Gamma h_k/\Gamma$ we are cutting off the cone point with cone angle $2\pi/3$ and replacing it by 2 cone points with cone angle $4\pi/3$.

Reversing the Process: Starting with an octahedron $\Sigma \in \mathcal{M}$ we can group the 6 cone points in pairs and connect the two cone points in each pair by an arc. We can do the pairing and connecting so as to minimize the total length of the three arcs. Once we do this, the arcs will be embedded. (If the arcs cross, we can surger them locally and find a new collection that is shorter.)

Let α_k for $k = 1, 2, 3$ be the three arcs. Since we have minimized the total length, each of these arcs is a geodesic segment. If we develop a small tubular neighborhood of α_k into \mathcal{C} we get a region of the form $U_k - h_k$ where U_k is some open topological disk and h_k is a hexagon with 6-fold symmetry. Figure 4.1 shows what we mean.




Figure 4.1: Developing the complement of an arc into \mathcal{C}

Let \bar{h}_k be the quotient of h_k by the order 3 (and index 2) subgroup of rotational symmetries of h_k . We can isometrically glue \bar{h}_k to Σ for $k = 1, 2, 3$. The result is a flat cone sphere Σ' with 3 cone points all having cone angle $2\pi/3$. The flat cone sphere Σ' (after a choice) determines a lattice $z_0\mathbf{Eis}$ and then the lifts of the glued in copies of \bar{h}_k determine the hexagons h_k for $k = 1, 2, 3$ and the complex numbers z_1, z_2, z_3 .

There is one subtle point in this discussion. We say *after a choice* because the geometry of Σ' only determines z_0 up to multiplication by a unit complex number. However, once we make some choice, the collection of sufficiently nearby 4-tuples (z_0, z_1, z_2, z_3) will parametrize a neighborhood of Σ in \mathcal{M} . In this moduli space, the spheres $\Sigma(z_0, z_1, z_2, z_3)$ and $\Sigma(uz_0, uz_1, uz_2, uz_3)$ are considered distinct even though, when u is a unit complex number, the underlying spheres are isometric.

These coordinates just give a local coordinate system, for several reasons. First, we had to make a choice of how to pair up the cone points based on the geometry of Σ . Second, if we replace z_k by ωz_k we get a different parametrization of the same sphere. So, we can only change the parameters a little bit if we want bijective parametrization of a neighborhood of Σ in \mathcal{M} .

4.2 Thurston Coordinates: A Second Pass

Again following [T], we take a more sophisticated approach to Thurston's coordinates on \mathcal{M} . The new system we get is equivalent to the old one *via* a complex linear change of coordinates.

Let Σ be a flat cone sphere. Let c denote the set of 6-cone points of Σ and let $\Sigma' = \Sigma - c$. Let $\tilde{\Sigma}'$ denote the universal cover of Σ' . Let $\pi : \tilde{\Sigma}' \rightarrow \Sigma'$ denote the universal covering map. The space $\tilde{\Sigma}'$ is not metrically complete. We can have a Cauchy sequence $\{\tilde{p}_n\}$ of points in $\tilde{\Sigma}'$ such that $\{p_n\}$ converges to a cone point. Here we have set $p_n = \pi(\tilde{p}_n)$. The limit of $\{\tilde{p}_n\}$ will not exist in $\tilde{\Sigma}'$.

We let $\tilde{\Sigma}$ denote the metric completion of $\tilde{\Sigma}'$. The space $\tilde{\Sigma}$ is known as the *orbifold universal cover* of Σ . We get $\tilde{\Sigma}$ from $\tilde{\Sigma}'$ by adding in points which serve as the limits of all the Cauchy sequences considered in the previous paragraph. The difference $\tilde{\Sigma} - \tilde{\Sigma}'$ is a discrete countable set of points which π maps to c . We call these points the *special points*.

The following is a lemma from [T].

Lemma 4.1 *$\tilde{\Sigma}$ has an equivariant triangulation whose edges are geodesic segments having special points as endpoints.*

Proof: Here is a sketch. Each special point p defines a \widehat{V} oronoi cell K_p consisting of the closure of all the points in $\tilde{\Sigma}$ which are closer to p than to any other special point. About each point in the intersection of at least 3 Voronoi cells is a metric disk that contains a finite number of points in its boundary. The convex hull of these points is a convex polygon whose vertices are special points. These polygons are called *Delaunay cells*. Typically the Delaunay cells are triangles; in case they are not, they may be further triangulated into triangles by the addition of some of the diagonals. This further triangulation may be done in an equivariant way. ♠

Since the triangulation of $\tilde{\Sigma}$ is equivariant, it induces a triangulation of the original flat cone sphere Σ . We can choose a collection of 4 edges e_1, e_2, e_3, e_4 of the triangulation of Σ which make a tree, and then lift this tree to $\tilde{\Sigma}$. To be artistic about it, we can choose the tree to look like a quadrupod, with the central vertex being equal to v . Then we can take the lifted tree to have \tilde{v} as the corresponding lifted vertex.

Let \tilde{w}_k be the other endpoint of the lift \tilde{e}_k of e_k . The 4 complex numbers $d(\tilde{w}_k)$ for $k = 1, 2, 3, 4$ serve as local coordinates for the space \mathcal{M} .

4.3 The Affine Map

Now we describe a real affine map from \mathcal{C} into \mathbf{C}^4 and then we interpret this map as a locally affine map from \mathcal{C} into \mathcal{M} .

If we have a nice coloring of Σ , we can lift this to get a nice coloring of $\tilde{\Sigma}$. To normalize our constructions, we fix a triple (v, e, f) where f is a polygon of the nice coloring, and e is an edge of f , and v is a vertex of f incident to e . We also suppose that v is a cone point. Now we choose a lift $(\tilde{v}, \tilde{e}, \tilde{f})$ of (v, e, f) to $\tilde{\Sigma}$. We call this lift our *favorite flag*. We have a unique locally isometric map $d : \tilde{\Sigma} \rightarrow \mathcal{C}$ such that

- $d(\tilde{v}) = 0$.
- $d(e) \subset \mathbf{R}$ and $d(e - v) \subset \mathbf{R}^+$.
- $d(f - e) \subset \mathbf{R} \times \mathbf{R}_+$.

Figure 4.2 shows our normalization relative to the coordinate axes in \mathcal{C} .




Figure 4.2: Normalizing the developing map.

We pick a point in \mathcal{C} . This defines a nice coloring for us. We then build $\tilde{\Sigma}$ based on this coloring, normalize as above, and then take the point in \mathcal{M} with coordinates $d(w_k)$ for $k = 1, 2, 3, 4$. This is our map A .

The map from \mathcal{C} to \mathbf{C}^4 is real linear. To see this, note that there is some finite chain of nice polygons P_1, \dots, P_m which cover the geodesic segment connecting our favorite vertex \tilde{v} to the vertex \tilde{w}_k . We can express $d(w_k)$ as the sum of certain complex numbers describing certain of the edges of $d(P_j)$ for $j = 1, \dots, m$. When we change the edge labels for our coloring, this sum changes in a linear way.

One thing we note is that the map $A(\mathcal{S}) \subset \mathbf{C}^4$ to \mathcal{M} need not be injective. Since we are working in a local coordinate chart, all we can say is that $A : \mathcal{C} \rightarrow \mathcal{M}$ is locally affine and locally injective.

4.4 The Thurston Hermitian Form

Thurston introduces a Hermitian form on the space \mathcal{M} . In the local coordinates described in §4.1, the form is given by

$$\langle Z, W \rangle = 6z_0\bar{w}_0 - 6z_1\bar{w}_1 - 6z_2\bar{w}_2 - 6z_3\bar{w}_3. \quad (7)$$

Here $Z = (z_0, z_1, z_2, z_3)$ and $W = (w_0, w_1, w_2, w_3)$. The diagonal part of this form $Z \rightarrow \langle Z, Z \rangle$ computes a multiple of the area of Σ . The multiple is such that the form counts 3 times the number of triangles when the tiling is by equilateral triangles having side length 1.

In general, the changes of coordinates between Thurston's various coordinate systems are always complex linear. The Hermitian form has a local expression in each coordinate system. The secret to lining up all the different local expressions is that their diagonal parts all compute the area, and a Hermitian form on a complex vector space is determined by its diagonal part. Equation 7 reveals that the form has signature $(1, 3)$.

Let me describe what the Hermitian form looks like in the coordinate system defined in §4.2. This alternate perspective connects with the way we define the quadratic form on \mathcal{C} in the next section. We have triangulated the universal cover $\tilde{\Sigma}$. Each triangle in the triangulation belongs to an orbit. We choose a finite collection of triangles, one per triangle orbit. We then develop these triangles in the plane. The locations of the points will be complex linear functions of $d(e_k)$.

Suppose we have two such triangulations, corresponding to different points of \mathcal{M} . Let τ and τ' be two corresponding triangles. We choose a vertex v of τ and the corresponding vertex v' of τ' . We let z_1 and z_2 be the complex numbers which describe the difference between the other vertices of τ and v . Likewise define z'_1 and z'_2 . For ease of notation we set $w_k = z'_k$. We then consider the expression

$$\frac{1}{4i} \left(z_1\bar{w}_2 - z_2\bar{w}_1 \right) \quad (8)$$

When $\tau = \tau'$ this expression computes the area of τ . In these coordinates, the Hermitian form is the sum of these expressions, appropriately scaled to match Equation 7.

4.5 The Quadratic Form

Suppose we have a nice hexagon with consecutive side lengths $\ell = (\ell_1, \dots, \ell_6)$. Then,

$$Q(\ell, \ell) = 2 \sum_i \ell_i \ell_{i+1} + \sum_i \ell_i \ell_{i+2} \quad (9)$$

computes $8\sqrt{3}$ times the area. The indices are taken cyclically. You can derive this by the following procedure:

- Triangulate the hexagon.
- Compute the areas of the triangles using Equation 8
- Symmetrize the resulting expression by averaging over all cyclic permutations of the variables.

We get similar formulas for other nice polygons by setting various of the variables equal to 0.

With the understanding that we are writing just one formula for all combinatorial types of nice polygon, we introduce the real quadratic form.

$$Q(\ell, m) = 2 \sum_i (\ell_i m_{i+1} + \ell_{i+1} m_i) + \sum_i (\ell_i m_{i+2} + m_i \ell_{i+2}) \quad (10)$$

We will only apply this form for nice polygons having the same combinatorics. To explain the normalization, we note that when this formula is applied to the pair $\ell = m = (1, 1, 1, 1, 1, 1)$ it returns 18, which is 3 times the number of unit equilateral triangles tiling the corresponding unit hexagon. The reason for the normalization is that we wanted to get an integral quadratic form.

Supposing that our nice coloring has n nice polygons, we define the following quadratic form on \mathcal{C} :

$$Q = \sum_{i=1}^n Q_i, \quad (11)$$

where Q_i is the quadratic form associated to the i th nice polygon. This is our quadratic form on \mathcal{C} . This normalization computes 3 times the number of triangles in Σ when Σ is triangulable by equilateral triangles of side length 1.

4.6 Non-Degeneracy and Injectivity

We have our real linear map $A : \mathcal{C} \rightarrow \mathcal{C}^{1,3}$. This map extends to a real linear map $A : \mathcal{S} \rightarrow L \subset \mathcal{C}^{1,3}$. Here \mathcal{S} is the 4-dimensional subspace spanned by \mathcal{C} and L is a 4-dimensional real linear subspace. At least on an open neighborhood of the good coloring on which this whole apparatus is based, the map from L into \mathcal{M} is injective.

Let us define Q' to be 6 times the real part of Thurston's Hermitian form, restricted to L . The diagonal part of Q' computes $(8\sqrt{3})$ times the area, at least when restricted to subsets of L corresponding to points in \mathcal{M} . This is the same factor that comes up for Q .

Lemma 4.2 *Q' is non-degenerate on L .*

Proof: If Q' is degenerate then there is some nonzero vector $V \in L$ such that $\langle V, W \rangle$ is pure imaginary for all $W \in L$. But this is not true for $W = V$. Hence Q' is non-degenerate. ♠

Let A^* be the pullback operator for quadratic forms. The two forms $A^*(Q')$ and Q have the same diagonal parts: on suitable open sets, both compute area. But a quadratic form on a real vector space is determined by its diagonal part. Hence $A^*(Q') = Q$.

Lemma 4.3 *If A is locally injective then Q is non-degenerate.*

Proof: If A is locally injective then A is an isomorphism from \mathcal{S} to L . But then Q is non-degenerate if and only if Q' is. Since Q' is non-degenerate, so is Q . ♠

Lemma 4.4 *If Q is non-degenerate then A is locally injective.*

Proof: Suppose for the sake of contradiction that $A : \mathcal{S} \rightarrow L$ has a kernel. Then for any nonzero vector $W \in \mathcal{S}$ in the kernel of A we have $A^*(Q')(W, W) = 0$. This contradicts the fact that $A^*(Q') = Q$ and that Q is non-degenerate. ♠

Lemma 4.5 *Suppose that Q has signature $(1, 3)$. Then the projectivized map PA is locally injective on PC .*

Proof: If PA is not locally injective, then there are two points $V, W \in L$ such that $W = \lambda V$ where $\lambda \in \mathbf{C}$ is non-real. But then L contains the \mathbf{C} -span of V . We have

$$Q'(V, V) > 0, \quad Q'(iV, iV) > 0, \quad Q'(V, iV) = 0.$$

Then the restriction of Q' to the \mathbf{C} -span of V has signature $(2, 0)$. This means that Q' cannot have signature $(1, 3)$ on L . ♠

5 References

[BG] C. Bavard and E. Ghys, *Polygones du plan et polyèdres hyperboliques*, *Geom. Dedicata* 43(2), (1992) pp 207–224

[T] W. Thurston, *Shapes of Polyhedra and triangulations of the sphere*, *Geometry & Topology Monographs*, Vol. 1: The Epstein birthday shrift, pp 511-549 (1998) See also arXiv:math/9801088.

[W] D. B. West, *Introduction to Graph Theory, 2nd Edition*, Prentice Hall (2001)