

Continued Fractions and the 4-Color Theorem

Richard Evan Schwartz *

December 28, 2023

Abstract

We study the geometry of some proper 4-colorings of the vertices of sphere triangulations with degree sequence $6, \dots, 6, 2, 2, 2$. Such triangulations are the simplest examples which have non-negative combinatorial curvature. The examples we construct, which are roughly extremal in some sense, are based on a novel geometric interpretation of continued fractions. We will also present a conjectural sharp “isoperimetric inequality” for colorings of this kind of triangulation.

1 Introduction

1.1 Background

The Four Color Theorem, first proved (with the assistance of a computer) by Wolfgang Haken and Kenneth Appel in 1976, is one of the most famous results in mathematics. See [W] for a thorough discussion. Here is one formulation. If you have any triangulation of the 2-sphere, it is possible to color the vertices using 4 colors such that no two adjacent vertices have the same coloring. This is called a *proper vertex 4-coloring*.

Certainly one can properly 4-color the vertices of a tetrahedron. A proper vertex 4-coloring of a triangulation \mathcal{Z} (with the same colors) canonically defines a simplicial map f from the sphere to the tetrahedron: Just map each vertex of \mathcal{Z} to the like-colored vertex of the tetrahedron and then extend linearly to the faces. The map f in turn defines a 2-coloring of the faces of

*Supported by N.S.F. Grant DMS-2102802

\mathcal{Z} . One colors a face of \mathcal{Z} black if f is orientation preserving on that face, and otherwise white.

The associated face 2-coloring has the property that around each vertex the number of black faces is congruent mod 3 to the number of white faces. This derives from the property that 3 triangles of the tetrahedron meet around each vertex. We call a face 2-coloring with this property a *good coloring*. Conversely, a good coloring for \mathcal{Z} defines a simplicial map to the tetrahedron and thus a proper 4-coloring of the vertices of \mathcal{Z} . So, an equivalent formulation of the 4-color theorem is that every triangulation of the sphere has a good coloring.

So far, the Four Color Theorem only has computer-assisted proofs. Perhaps one can get insight into the result by looking at examples of good colorings. The good coloring version has a geometric feel to it, and so perhaps some geometric insight might help. The purpose of this paper is to look at the geometry of these good colorings in some special cases.

A *triangulation of non-negative combinatorial curvature* is one in which the maximum degree is 6. All the vertices have degree 6 except for a list v_1, \dots, v_k which have degrees $2 \leq d_1, \dots, d_k < 6$. Euler's Formula gives the condition on the degrees:

$$\sum_{i=1}^k (6 - d_i) = 12. \tag{1}$$

In particular $k \leq 12$. The quantity

$$\frac{\pi}{3} \times (6 - d_i)$$

is the *combinatorial curvature* at v_i . Equation 1 translates into a discrete version of the Gauss-Bonnet theorem, which says that the total combinatorial curvature is 4π .

The triangulations of non-negative combinatorial curvature form an attractive family to study. In [T], William Thurston organized these triangulations into moduli spaces. To give some idea of how this works, a triangulation of non-negative combinatorial curvature defines a flat cone structure on the sphere with non-negative curvature: we just make all the triangles unit equilateral triangles. The set of all triangulations with the same list d_1, \dots, d_k of defects includes in the moduli space of flat cone structures on spheres with appropriately prescribed singularities. So, even though the triangulations

don't exactly vary continuously, one can think of them as special points inside moduli spaces consisting of structures which do vary continuously. Also, if the triangulations involve many triangles and the defects are well spread out, one can imagine that the defects almost vary continuously.

Given the nice structure of the totality of such triangulations, it seems like an interesting idea to study the space of good colorings as a kind of partially defined bundle over these moduli spaces. Perhaps the structure of such colorings is related somehow to the placement of the defects. As the defects vary around, perhaps the good colorings vary in a nice way to some extent. I imagine that the total picture, seen all at once, would be spectacularly beautiful. All this is very speculative. In spite of making a lot of computer experiments over the years – every time I teach the graph theory class at Brown I play with this project – I don't have much to report. In this very modest paper I will consider the simplest cases. The cases I have in mind are where $k = 3$ and $d_1 = d_2 = d_3 = 2$.

1.2 The Continued Fraction Colorings

The $6, \dots, 6, 2, 2, 2$ triangulations are indexed by the nonzero Eisenstein integers. An *Eisenstein integer* is a number of the form

$$a + b\alpha, \quad a, b \in \mathbf{Z}, \quad \alpha = \frac{1 + i\sqrt{3}}{2}.$$

(It is more common to use $\omega = \alpha^2$ in place of α , but α is more convenient for us.) To see the connection, let \mathcal{E} denote the ring of Eisenstein integers. The points of \mathcal{E} are naturally the vertices of an equilateral triangulation \mathcal{T} of \mathbf{C} . Given some nonzero $\beta \in \mathcal{E}$, we let $\beta\mathcal{T}$ be the bigger equilateral triangulation obtained by multiplying the whole picture by β . We let G_β denote the group of symmetries generated by order 3 rotations in the vertices of $\beta\mathcal{T}$. The quotient

$$\mathcal{T}(\beta) = \mathcal{T}/G_\beta$$

is the desired triangulation.

Since all the vertices of $\mathcal{T}(\beta)$ have even degree, $\mathcal{T}(\beta)$ always has a good coloring. One just colors the triangles alternately black and white in a checkerboard pattern. Indeed, this face coloring corresponds to a proper 3-coloring of the vertices.

Figure 1 shows a different good coloring for $\beta = 2 + 3\alpha$. We call this coloring $\mathcal{C}(2+3\alpha)$. To get the triangulation of the sphere, cut out the specially

outlined central rhombus, fold it up like a taco, and glue the edges together in pairs. Figure 1 is really showing part of the orbifold universal cover of the coloring. Figures 7 and 8 below show more elaborate examples.

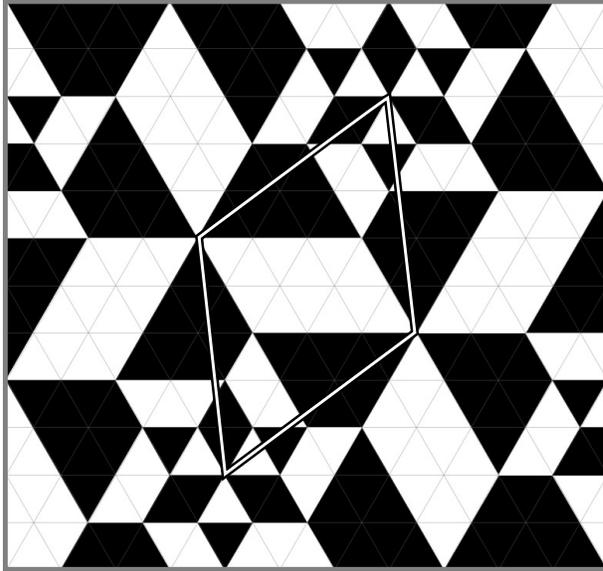


Figure 1: The good coloring $\mathcal{C}(2 + 3\alpha)$ of $\mathcal{T}(2 + 3\alpha)$.

The coloring in Figure 1 is (at least experimentally) extremal in a certain sense. Define the *fold count* of a good coloring to be the number of edges which form black-white interfaces. For the alternating coloring, the fold count is $3/2$ times the number of faces. For $\mathcal{T}(2 + 3\alpha)$ this comes to 57. In contrast, $\mathcal{C}(2 + 3\alpha)$ has fold count 23. It seems that $\mathcal{C}(2 + 3\alpha)$ minimizes the fold count amongst all good colorings of $\mathcal{T}(2 + 3\alpha)$.

The number of good colorings of $\mathcal{T}(\beta)$ grows exponentially with $|\beta|$, for an easy reason. In any good coloring, if we can find a vertex of degree 6 where the colors alternate, we can switch the colors and get another good coloring. In terms of the original vertex coloring, the neighbors of such a vertex v are colored using just 2 colors, and so we have an option to switch the color of v to the other available color. Starting with the alternating coloring, we can take a large family of non-adjacent vertices and make these swaps according to any binary sequence we like. On the other hand, these examples seem rather similar to the alternating coloring. Their fold count is linear in the number of the faces.

The colorings with smaller fold counts, and in particular with minimal fold counts, seem to be much more rigid and interesting. To use an analogy

from statistical mechanics, the colorings with large fold count are sort of like a fluid or a gas, and the colorings with small fold count are more like solid crystals. The example in Figure 1 is part of an infinite sequence of examples. We call these examples $\mathcal{C}(\beta)$, where β ranges over the primitive Eisenstein integers. (An Eisenstein integer $\beta = a + b\alpha$ *primitive* if it is not an integer multiple of another one. Equivalently, a and b are relatively prime.) We will see that $\mathcal{C}(\beta)$ is a special good coloring of $\mathcal{T}(\beta)$ which geometrically implements the continued fraction expansion of a/b . We call these colorings *continued fraction colorings*.

1.3 Properties of the Continued Fraction Colorings

The main result of this paper is that these continued fraction colorings exist, but I will prove some additional results about them. One interesting property is that these colorings have the same number of black and white triangles. Indeed, any good coloring of a triangulation having all even degrees has the same number of black and white triangles. See §3.1 for the quick proof which I learned from Kasra Rafi.

Our next result concerns the asymptotics of the fold count for the continued fraction colorings. Let $\{a_n/b_n\} \in (0, 1)$ be a sequence of rationals. Let

$$\mathcal{T}_n = \mathcal{T}(a_n + b_n\alpha), \quad \mathcal{C}_n = \mathcal{C}(a_n + b_n\alpha).$$

Let f_n denote the fold count of \mathcal{C}_n and let F_n denote the number of faces in \mathcal{T}_n .

Theorem 1.1 *The following is true about the continued fraction colorings.*

1. *If $\{a_n/b_n\}$ converges to an irrational limit then $\lim_{n \rightarrow \infty} f_n/F_n = 0$.*
2. *If $\{a_n/b_n\}$ is the sequence of continued fraction approximants of a quadratic irrational, then $\lim_{n \rightarrow \infty} f_n^2/F_n$ exists and is finite.*

Statement 2 of Theorem 1.1 motivates the following definitions.

Definition: Given a coloring \mathcal{C} we define

$$\eta(\mathcal{C}) = \frac{f^2}{F}, \tag{2}$$

where f is the fold count for \mathcal{C} and F is the number of triangles in the triangulation which \mathcal{C} colors. We call $\eta(\mathcal{C})$ the *Eisenstein Isoperimetric Ratio* of \mathcal{C} .

Definition: Given a quadratic irrational $\eta \in (0, 1)$ let

$$\eta(\zeta) = \lim_{n \rightarrow \infty} \eta(p_n + q_n \alpha),$$

where $\{p_n/q_n\}$ is the sequence of continued fraction approximants of ζ . Theorem 1.1 guarantees that this limit exists. We call $\eta(\zeta)$ the *Eisenstein Isoperimetric Ratio* of ζ .

In §3.3 we will show, among other calculations, that

$$\eta(\phi^{-1}) = \phi^6, \quad \phi = \frac{\sqrt{5} + 1}{2}. \quad (3)$$

Here ϕ is the golden ratio. Our proof of Statement 2 of Theorem 1.1 will show more generally that $\eta(\zeta) \in \mathbf{Q}(\zeta)$, the quadratic field containing ζ . See §3.3 for some examples.

For comparison we prove the following easy result.

Theorem 1.2 *Let \mathcal{C} be any good coloring of any triangulation with degree sequence $6, \dots, 6, 2, 2, 2$. Then $\eta(\mathcal{C}) \geq 3$.*

Theorem 1.2 indicates that some of the continued fraction colorings roughly minimize the fold count. Here is a strong conjecture along these lines.

Conjecture 1.3 *Suppose $\{\mathcal{G}_n\}$ is any infinite sequence of distinct colorings of primitive sphere triangulations with degree sequence $6, \dots, 6, 2, 2, 2$. Then $\liminf_{n \rightarrow \infty} \eta(\mathcal{G}_n) \geq \phi^6$.*

Here *primitive* means that the triangulation has the form $\mathcal{T}(\beta)$ for a primitive Eisenstein integer β . Equation 3 says that this conjecture is sharp. Since $\phi^6 = 17.944\dots < 18$, Theorem 1.2 says that the conjecture is true up to a factor of 6.

Sometimes there are good colorings which have lower fold count than the corresponding continued fraction coloring. See Figure 9 in §3 for an example. Here is one last conjecture about the general situation.

Conjecture 1.4 *There is a constant Ω with the following property. For any triangulation \mathcal{T} with degree sequence $6, \dots, 6, 2, 2, 2$ there is a good coloring \mathcal{C} of \mathcal{T} with $\eta(\mathcal{C}) < \Omega$.*

1.4 Organization

In §2, after a discussion of the slow Gauss map and its connection to continued fractions, I will construct the continued fraction colorings. The building blocks are what I call *capped flowers*, and these in turn are made in layers from cyclically arranged patterns of trapezoids which I call *trapezoid necklaces*. (Look again at Figure 1.) I will explain how the set of trapezoid necklaces is naturally the vertex set of the infinite rooted binary tree (modified to have an extra vertex at the bottom). Taking a path in this tree defines the capped flower. In §3 I will prove Theorem 1.1, establish Equation 3, and give some evidence for Conjecture 1.3. In §4 I will prove Theorem 1.2.

1.5 Acknowledgements

I'd like to thank Ethan Bove, Peter Doyle, Jeremy Kahn, Rick Kenyon, Curtis McMullen, Kasra Rafi, and Peter Smillie for various conversations (sometimes going back some years) on topics related to the material here.

2 The Main Construction

2.1 The Slow Gauss Map

Given $r \in (0, 1)$ we define $\gamma(r)$ to be whichever of the two numbers

$$\frac{1}{r} - 1 \quad \text{or} \quad \frac{1}{\frac{1}{r} - 1} \quad (4)$$

lies in $(0, 1]$. When $r > 1/2$ the first choice works and when $r < 1/2$ the second choice works. When $r = 1/2$ both choices give $\gamma(1/2) = 1$. The map $\gamma : (0, 1) \rightarrow (0, 1]$ is called the *slow Gauss map*.

Let \mathcal{S} denote the set of rationals in $(0, 1]$. On $\mathcal{S} - \{1\}$ we have:

$$\gamma\left(\frac{a}{a+b}\right) = \gamma\left(\frac{b}{a+b}\right) = \frac{a}{b}. \quad (5)$$

We can think of \mathcal{S} as the infinite rooted binary tree (modified to have an extra bottom vertex). We make this tree by joining each member of $\mathcal{S} - \{1\}$ to its image under γ . Figure 2 shows the beginning of this tree.

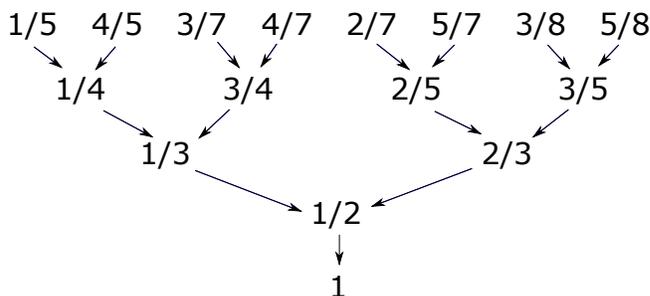


Figure 2: The beginning of the tree of rationals.

The slow Gauss map is connected to continued fractions. The traditional Gauss map is

$$\gamma^*(p/q) = (q/p) - \text{floor}(q/p). \quad (6)$$

For each p/q there is some *comparison exponent* k such that

$$\gamma^*(p/q) = \gamma^k(p/q).$$

In other words, the (suitably) iterated slow Gauss map has the same action as the traditional Gauss map; it just works more slowly. The continued fraction expansion of p/q is obtained by recording the sequence of comparison exponents we see as we iteratively apply γ and γ^* to p/q .

2.2 Trapezoid Necklaces

An *isosceles trapezoid* is a quadrilateral with two parallel sides such that the other two sides are non-parallel but have the same length. We call the longer parallel side the *top*, the shorter parallel side the *bottom*, and the other two sides the *diagonal sides*. We allow the degenerate case of an isosceles triangle. In this case the bottom has length 0. An *Eisenstein trapezoid* is an isosceles trapezoid whose edges lie in the 1-skeleton of \mathcal{T} , the planar equilateral triangulation whose vertices are the Eisenstein integers. Figure 1 above and Figure 3 below feature some Eisenstein trapezoids.

Up to symmetries of \mathcal{T} , an Eisenstein trapezoid X is characterized by the pair (a, b) where a is the length of a diagonal side of X and b is the length of the top. We call X *primitive* if a, b are relatively prime. When X is primitive, we define the *aspect ratio* to be a/b . The aspect ratio determines the primitive Eisenstein trapezoid up to symmetries of \mathcal{T} . Thus, modulo symmetry, the primitive Eisenstein trapezoids are naturally in bijection with the set \mathcal{S} of rationals considered above. The Eisenstein trapezoids in Figure 1 are all primitive, and their aspect ratios are variously $1/1$ and $1/2$ and $2/3$. The Eisenstein trapezoids in Figure 3 have aspect ratio $3/5$.

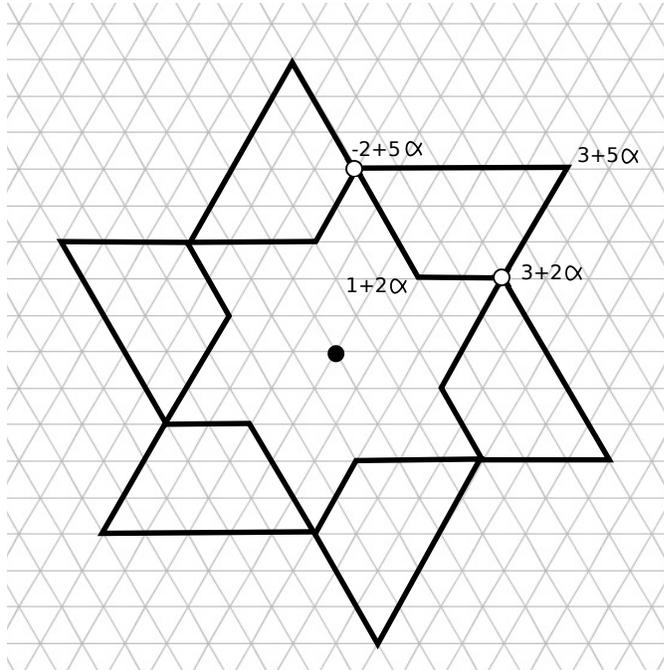


Figure 3: A trapezoid necklace of aspect ratio $3/5$.

Figure 3 illustrates what we mean by a *trapezoid necklace*. This is a union of 6 primitive Eisenstein trapezoids X_1, \dots, X_6 which has the following properties:

- The trapezoids have pairwise disjoint interiors.
- $X_i \cap X_{i+1}$ is a single point, a common vertex, for all i .
- An order 6 rotation ρ of \mathcal{T} has the action $\rho(X_i) = X_{i+1}$ for all i .

In this description the indices are taken mod 6. We define the *center* of the necklace to be the fixed point of ρ . When the center is 0, the map ρ (or perhaps its inverse) is multiplication by α . We define the aspect ratio of the necklace to be the common aspect ratio of the 6 individual trapezoids.

Up to symmetry of \mathcal{T} , there exists a unique Eisenstein necklace of aspect ratio $a/b \in \mathcal{S}$. If we normalize the picture so that 0 is the center, then one of the trapezoids X_1 has vertices

$$a + b\alpha, \quad (a - b) + b\alpha, \quad (2a - b) + (b - a)\alpha, \quad a + (b - a)\alpha.$$

The intersection of X_1 and $X_2 = \rho(X_1)$ is the point $(a - b) + b\alpha$ because

$$\alpha \times (a + (b - a)\alpha) = (a - b) + b\alpha.$$

This little calculation uses the fact that $\alpha^2 = \alpha - 1$.

2.3 Empty Trapezoid Flowers

Each trapezoid necklace X defines a smaller trapezoid necklace $Y = \gamma(X)$ having the same center. The defining property is that the top side of each trapezoid Y_i in Y is a side of a trapezoid X_j of X , and one of the diagonal sides of Y_i is a side of one of the trapezoids of X adjacent to X_j . This awkward definition is very much like a written description of how to drink a glass of water. A demonstration says a thousand words. Figure 4 shows the trapezoid necklaces

$$X \rightarrow \gamma(X) \rightarrow \gamma^2(X) \rightarrow \gamma^3(X)$$

alternately colored black and white. Here X is as in Figure 3. The respective aspect ratios are given by $3/5 \rightarrow 2/3 \rightarrow 1/2 \rightarrow 1/1$. We call the union of these necklaces the *empty 3/5-flower*.

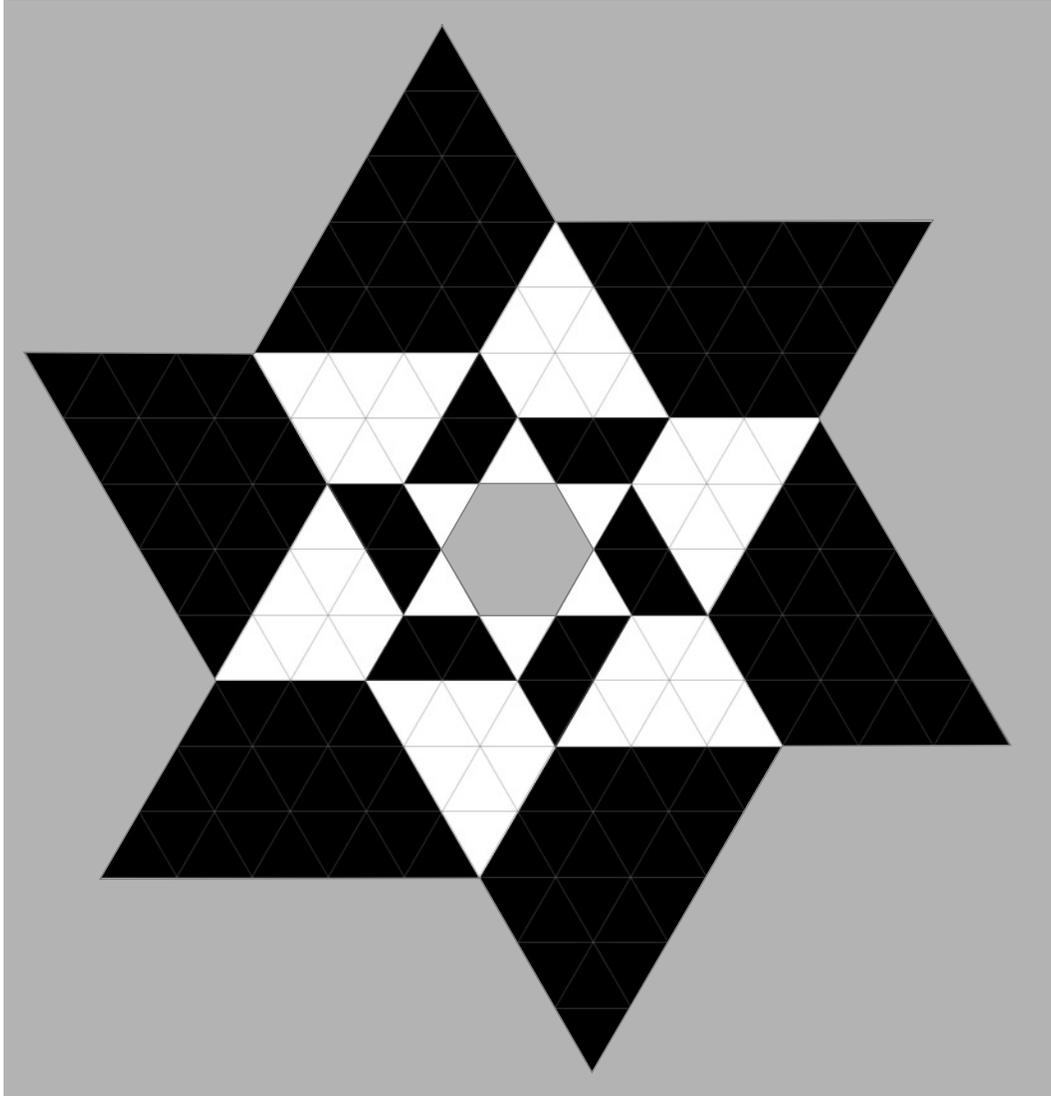


Figure 4: The empty $3/5$ -flower.

As is suggested by our notation, the action of γ here mirrors the action of the slow Gauss map γ defined above. That is, if r is the aspect ratio of X then $\gamma(r)$ is the aspect ratio of $\gamma(X)$.

We discussed above how the slow Gauss map is related to continued fractions. Here we continue the discussion. As we now illustrate, our construction also precisely implements the continued fraction expansion.

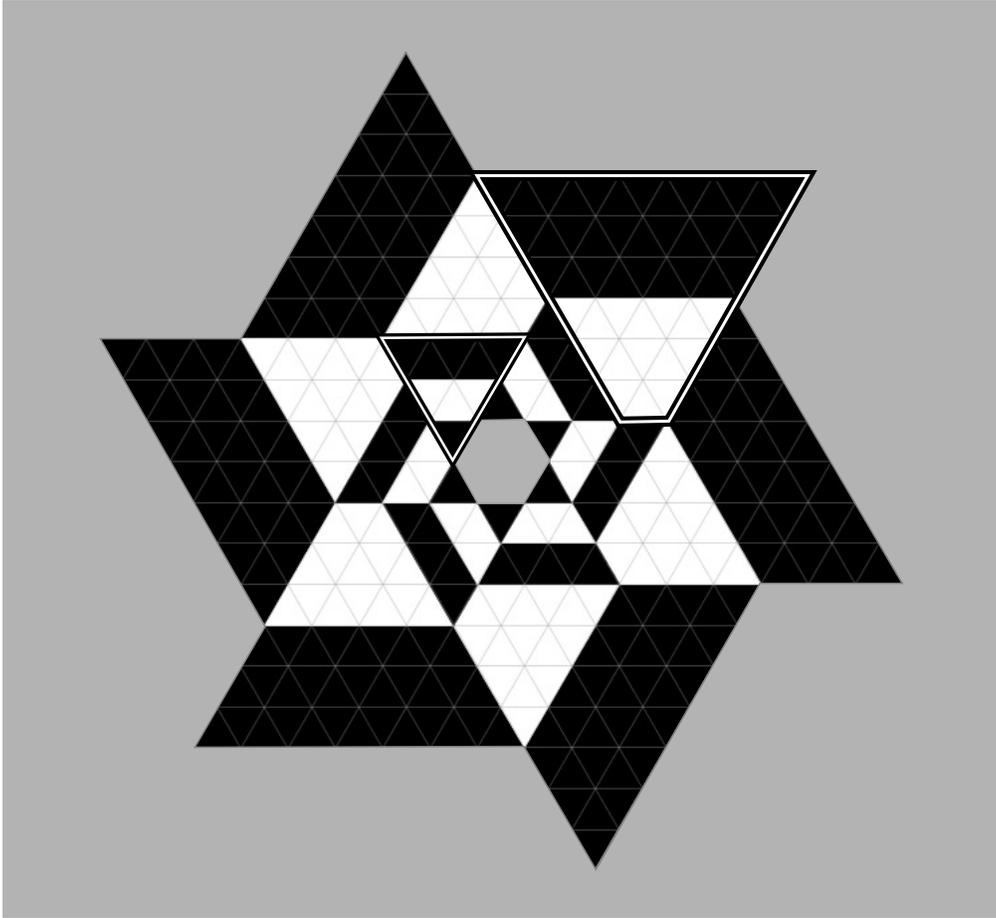


Figure 5: The empty $3/7$ -flower.

Figure 5 shows the empty $3/7$ -flower, corresponding to the tree path $3/7 \rightarrow 3/4 \rightarrow 1/3 \rightarrow 1/2 \rightarrow 1/1$. The specially-outlined trapezoids (and their rotated images) are the maximal trapezoids in the flower. They have respectively 2 and 3 “stripes”. For comparison, $3/7$ has continued fraction $0 : 2 : 3$. That is

$$\frac{3}{7} = 0 + \frac{1}{2 + \frac{1}{3}}.$$

In general, the empty p/q -flower starts with a p/q -necklace and then fills in the full γ orbit, alternately coloring the necklaces black and white. The innermost necklace always has aspect ratio $1/1$. Up to symmetries of \mathcal{T} , the empty p/q -flower is unique. One can read off the continued fraction expansion of p/q by counting the stripes of the maximal trapezoids.

2.4 Filling and Capping the Flowers

We fill an empty flower by coloring the remaining 6 triangles black and white in an alternating pattern. Figure 6 shows this for the $3/5$ -flower. Up to symmetries of \mathcal{T} (and swapping the colors) there is a unique p/q -flower. The empty flowers have 6-fold rotational symmetry but the (filled) flowers have 3-fold rotational symmetry. We break a symmetry to define the filling.

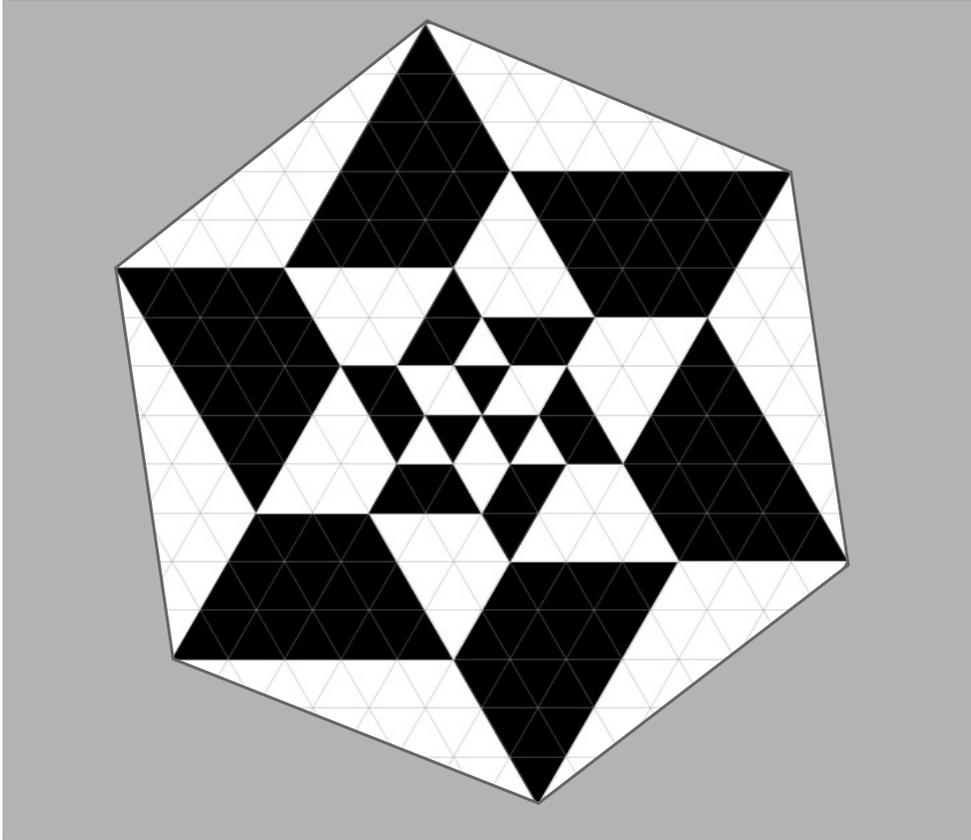


Figure 6: The filled and capped $3/5$ -flower

Figure 6 also shows what we mean by *capping* a flower. We take the convex hull of the flower and color the complementary triangles the color opposite the color of the outer necklace. The capped flowers, which are regular hexagons with Eisenstein integer vertices, are the building blocks of our colorings. When two translation-equivalent capped flowers meet along a common boundary edge, the triangular regions merge to become an *Eisenstein parallelogram* – i.e., one whose boundary lies in the 1-skeleton of \mathcal{T} .

2.5 Defining the Colorings

Consider the hexagonal tiling of the plane by translates of the capped a/b -flower. Because the capped flowers have 3-fold rotational symmetry, the resulting planar coloring is invariant under the group G_β generated by order 3 reflections in the vertices and centers of the hexagons. The quotient of this planar coloring by G_β is $\mathcal{C}(\beta)$. By construction, $\mathcal{C}(\beta)$ is good.

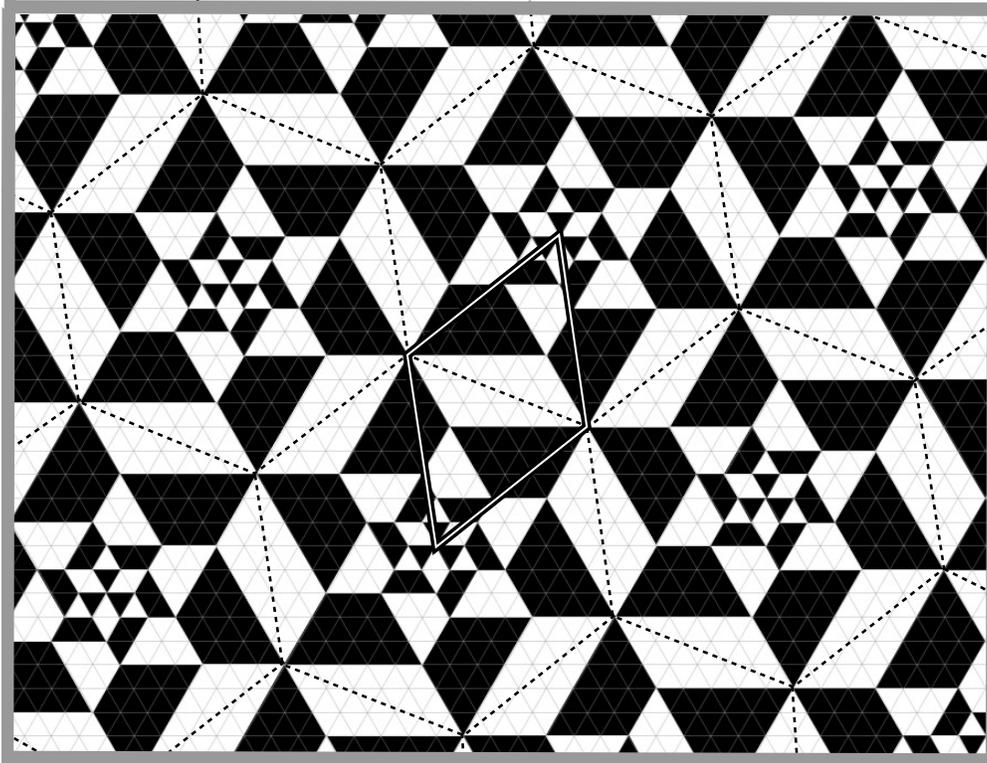


Figure 7: The universal cover of $\mathcal{C}(3 + 5\alpha)$.

Figure 7 shows the construction for $\beta = 3 + 5\alpha$. The region bounded by the big central rhombus is a fundamental domain for the action of G_β . The colorings exhibit a lot of variety. Figure 8 below shows $\mathcal{C}(a + 13\alpha)$ for $a = 1, 3, 5, 7, 9, 11$. It is worth noting that for parameters like $1/13$ the fold count is linear in the number of triangles.

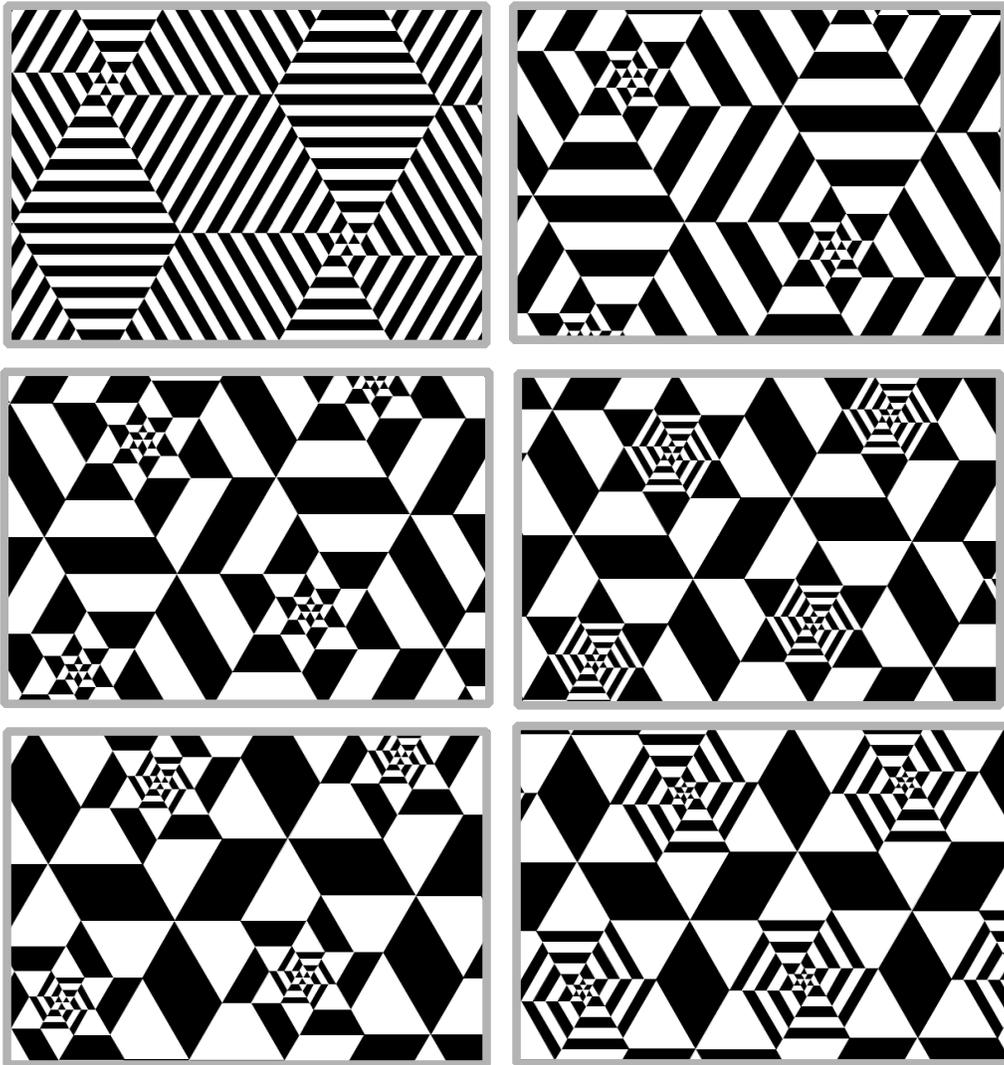


Figure 8: The orbifold covers of $\mathcal{C}(a + 13a)$ for $a = 1, 3, 5, 7, 9, 11..$

3 Properties of the Colorings

3.1 Zero Degree

Let T denote the regular tetrahedron. Call a triangulation of the sphere *even* if it has all even degrees. In this section we prove that any good coloring of an even triangulation Σ has the same total number of black and white triangles. This is equivalent to the statement that the associated map $f : \Sigma \rightarrow T$ has topological degree 0. I learned this proof from Kasra Rafi.

We equip \mathbf{R}^2 with the unit equilateral triangulation. Let $\pi : \mathbf{R}^2 \rightarrow T$ be the branched covering which maps triangles to triangles in an isometric way. A good triangulation has the property that, around each vertex, the number of black triangles is congruent mod 3 to the number of white triangles. In the case of an even triangulation, a good coloring has the stronger property that around each vertex the number of black triangles is congruent mod 6 to the number of white triangles. But this means that the map $f : \Sigma \rightarrow T$ lifts to a map $\tilde{f} : \Sigma \rightarrow \mathbf{R}^2$ such that $f = \pi \circ \tilde{f}$. The map \tilde{f} is homotopic to a point and hence so is f . But then f has topological degree 0.

3.2 Asymptotic Properties

In this section we will prove Theorem 1.1. Recall that \mathcal{S} is the set of rational numbers in $(0, 1]$. Now suppose that $\{p_n/q_n\}$ is an infinite sequence of elements of \mathcal{S} having an irrational limit ζ .

Statement 1: Let $G(p_n/q_n)$ denote the sequence of rationals we get by starting with p_n/q_n , iteratively applying the slow Gauss map, then listing the terms in reverse order. For instance,

$$G(3/7) = \frac{1}{1}, \frac{1}{2}, \frac{1}{3}, \frac{3}{4}, \frac{3}{7}.$$

Lemma 3.1 *For any D there are constants M, N such that if $n > N$ then all but the first M terms of $G(p_n/q_n)$ have $p_n > D$.*

Proof: Our sequences stabilize in the following sense. For any given M , we can take N sufficiently large so that the first M terms of $G(p_n/q_n)$ are the same as long as $n > N$. What makes it work is that the thinner sequence of continued fraction approximants of p_n/q_n converges to the sequence of

continued fraction approximants of ζ , and we can derive the G -sequences from these thinner sequences.

We also note a second property of our sequences: The numerators of the terms in $G(p_n/q_n)$ are monotone increasing as we move along the sequence. When we combine the increasing numerator property with the stabilization property, we get the result of this lemma. ♠

In terms of the flowers, Lemma 3.1 says that all but the first M inner trapezoid necklaces Q have the property that

$$\frac{p(Q)}{A(Q)} < \frac{100}{D}.$$

Here $p(Q)$ denotes the perimeter of Q and $A(Q)$ denotes the number of equilateral triangles comprising Q . We picked an unrealistically large constant of 100 here to avoid having to think about the fine points of trapezoids.

Our analysis shows that, within the n th flower, the average ratio of the perimeter of a trapezoid to the area of the trapezoid tends to 0 as $n \rightarrow \infty$. This immediately implies that the ratio f_n/F_n converges to 0.

Statement 2: Recall that the quadratic irrational limit $\zeta = \lim p_n/q_n$ has an eventually periodic continued fraction expansion and (equivalently) is pre-periodic under the action of the slow Gauss map.

We think of $\mathcal{T}(p_n + q_n\alpha)$ as a coloring of a doubled equilateral triangle. We then scale the metric by a factor of q_n^{-1} . The corresponding sequence of doubled equilateral triangles converges (say, in the Gromov-Hausdorff topology) to another doubled equilateral triangle, namely \mathbf{C}/G_β where $\beta = \zeta + \alpha$. Again, G_β is the group generated by order 3 rotations in points of $\beta\mathcal{E}$.

The rescaled colorings also converge. The limiting coloring is made from “infinite flowers” and parallelograms. The infinite flowers are limits of the rescaled flowers associated to the Eisenstein integers $a_n + b_n\alpha$. Scaling appropriately, we arrange that one of the outer trapezoids in the infinite flower F has aspect ratio ζ , top right vertex β and horizontal top and bottom. The remaining trapezoids have aspect ratios which we find by applying the slow Gauss map iteratively to ζ . These necklaces converge to the origin.

Let $F[k]$ denote the infinite flower obtained by trimming off the outermost k trapezoid necklaces of F . Given the preperiodicity of the slow Gauss map on ζ , there are integers m, n and some $\lambda \in (0, 1)$ such that

$$F[m + kn] = \lambda^k F[m], \quad k = 1, 2, 3, \dots \quad (7)$$

In other words, if we strip off the outer m layers of F then the resulting infinite flower is self-similar with “period” n and similarity factor λ .

The area of a unit equilateral triangle is $\sqrt{3}/4$ and the side length is 1. Let us redefine the Eisenstein isoperimetric ratio to be the quantity

$$\frac{\text{perimeter}^2}{\text{area}} \times \frac{4}{\sqrt{3}}. \quad (8)$$

This new definition coincides with the old definition for colorings based on unit equilateral triangulations. Also, the new definition is scale invariant. Thus, to prove Statement 2 of Theorem 1.1 we just need to show that this quantity is finite and well-defined for our limiting coloring. This suffices because the invariant then varies continuously.

We first compute the area. This is given by

$$\text{area} = C_1 + C_2 \sum_{k=1}^{\infty} \lambda^{2k} = C_1 + \frac{C_2}{1 - \lambda^2}.$$

Here C_1 is a constant that depends on the m outer layers of the infinite flower and on the triangular regions defining the cap of the flower F . The constant C_2 is determined by the outer n layers of $F[m]$.

Now we compute the perimeter. This is given by

$$\text{perimeter} = C_3 + C_4 \sum_{k=1}^{\infty} \lambda^k = C_3 + \frac{C_4}{1 - \lambda}$$

The constants C_3, C_4 have a similar dependence as C_1, C_2 . So, the quantity in Equation 8 exists and is finite. This completes the proof of Statement 2.

Remarks:

(1) The quantities $C_1\sqrt{3}$ and $C_2\sqrt{3}$ both belong to $\mathbf{Q}(\zeta)$. Likewise C_3, C_4, λ also all belong to $\mathbf{Q}(\zeta)$. Therefore, the limiting value $\eta(\zeta)$ lies in $\mathbf{Q}(\zeta)$. I will give some calculations below.

(2) Our insistence that $\{p_n/q_n\}$ be the sequence of continued fraction approximants of ζ is more restrictive than need be. The limiting argument works as long as $\{p_n/q_n\}$ limits to ζ and the corresponding continued fraction expansions are uniformly bounded. Thus, the Eisenstein isoperimetric ratio is well-defined for any irrational in $(0, 1)$ with bounded continued fraction approximation.

3.3 Some Calculations

Let us first establish Equation 3. Rather than take the general approach from the last section we work more concretely with the sequence $\{a_n/a_{n+1}\}$ of continued fraction approximants to ϕ^{-1} . Here

$$a_1, a_2, a_3, a_4, a_5, \dots 1, 1, 2, 3, 5, \dots$$

is the sequence of Fibonacci numbers. We will use the well-known asymptotic formula

$$a_n \sim \frac{\phi^n}{\sqrt{5}}. \quad (9)$$

For $n \geq 2$ let f_n denote the fold count of $\mathcal{C}(a_n/a_{n+1})$ and let F_n denote the number of triangles. From our program we compute the following table:

n	a_n	a_{n+1}	f_n	F_n	$\phi^{-6} f_n^2 / F_n$
2	1	2	13	14	.672718...
3	2	3	23	38	.775795...
4	3	5	39	98	.864923...
5	5	8	65	258	.912601...
6	8	13	107	674	.946633...

After a little trial and error we find the following formulas

$$f_n = -1 + 2 \sum_{k=1}^{n+2} a_k, \quad F_n = 2 + 4 \sum_{k=1}^n a_k a_{k+1}. \quad (10)$$

These formulas give the same answers as in the table above for $n = 2, 3, 4, 5, 6$. An easy inductive proof, which we omit, establishes them. Using the approximation in Equation 9, and the familiar formula for the partial sums of a geometric series, and the fact that $1 + \phi = \phi^2$, we find that

$$f_n^2 \approx \frac{4}{5} \times \phi^{2n+8}, \quad F_n \approx \frac{4}{5} \times \phi^{2n+2}.$$

Here \approx means equal up to a uniformly bounded error. Dividing the one equation by the other gives Equation 3.

Now we describe some additional (nonrigorous) calculations we did in Mathematica [Wo]. Starting with a quadratic irrational ζ we do the following.

1. We take a close continued fraction approximation p/q to ζ . For our calculations we used the command **Rationalize** $[\zeta, \text{Power}[10, -1000]]$. This gives us a continued fraction approximant which agrees with ζ up to 1000 decimal places. (We needed roughly this much accuracy in a few cases.)
2. We compute the continued fraction expansion of $\eta(p + q\alpha)$.
3. We extract what appears to be the start of a preperiodic continued fraction expansion for a quadratic irrational and we then reconstruct this quadratic irrational from the (implied) preperiodic continued fraction expansion. (In all cases, the fractional part of the expression was strictly periodic.)
4. After finding what appears to be the period of the continued fraction expansion, we reconstruct the quadratic irrational which has this continued fraction expansion. We guess that this is probably $\eta(\zeta)$.

Here we show some examples. Let $\zeta_n = \sqrt{n} - \text{floor}(\sqrt{n})$. We have

$$\begin{aligned}
\eta(\zeta_2) &= \frac{75+53\sqrt{2}}{7} & \eta(\zeta_3) &= \frac{132+72\sqrt{3}}{13} \\
\eta(\zeta_5) &= \frac{321+137\sqrt{5}}{19} & \eta(\zeta_6) &= \frac{27+9\sqrt{6}}{2} \\
\eta(\zeta_7) &= \frac{3100+856\sqrt{7}}{259} & \eta(\zeta_8) &= \frac{1569+370\sqrt{8}}{98}
\end{aligned} \tag{11}$$

One could likely adapt the above proof of Statement 2 of Theorem 1.1 to rigorously justify these calculations.

3.4 Evidence for the Isoperimetric Inequality

Most of our evidence for Conjecture 1.3 comes from calculations within the family of continued fraction colorings. Given primitive Eisenstein integers

$$\beta = a + b\alpha, \quad \beta' = a' + b'\alpha,$$

we write $\beta \preceq \beta'$ if $b \leq b'$. It seems that

$$\eta(\mathcal{C}(\beta)) < \eta(\mathcal{C}(\beta')) \tag{12}$$

whenever $\beta \preceq \beta'$ and β is a Fibonacci Eisenstein integer and $\beta' \neq \beta$.

I checked Equation 12 when β is any of the first 10 examples and for all $|\beta'| < 500$. I think that this is very strong evidence. The proof of Equation 12 in general should be purely a matter of number theory. I haven't yet looked for a proof. Equation 12 would establish the Eisenstein Isoperimetric Inequality for the continued fraction colorings. These, of course, are only a tiny fraction of the good colorings.

The only other evidence I have is the fact that $\mathcal{C}(1 + 2\alpha)$ and $\mathcal{C}(2 + 3\alpha)$ seem to minimize the fold counts for good colorings of $\mathcal{T}(1+2\alpha)$ and $\mathcal{T}(2+3\alpha)$ respectively. My program is not good enough to convincingly check this for larger Fibonacci triangulations. Combining this scant evidence with the very good evidence for Equation 12, I make the following conjecture.

Conjecture 3.2 *Suppose that β is a Fibonacci Eisenstein integer and that $\beta' \in (0, 1)$ is a different Eisenstein integer such that $\beta \preceq \beta'$. Let $\mathcal{C} = \mathcal{C}(\beta)$. Let \mathcal{C}' be an arbitrary good coloring of $\mathcal{T}(\beta')$. Then we have $\eta(\mathcal{C}) < \eta(\mathcal{C}')$.*

This conjecture would combine with Equation 3 to prove Conjecture 1.3.

It is worth pointing out that the continued fraction colorings are far from the best in many cases. Figure 9 compares the continued fraction $\mathcal{C}(1 + 5\alpha)$ (left) with another coloring of $\mathcal{T}(1 + 5\alpha)$ (right) which seems to be the one which minimizes the fold count.

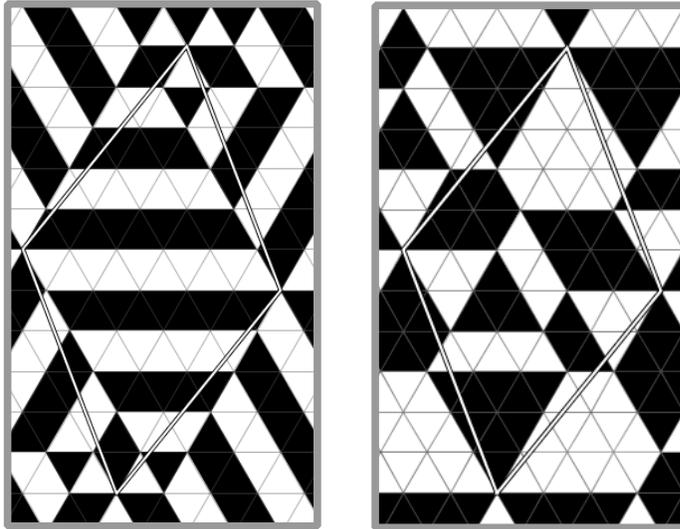


Figure 9: Two good colorings of $\mathcal{T}(1 + 5\alpha)$.

4 A Bound on the Isoperimetric Ratio

In this chapter we prove Theorem 1.2. We first establish an isoperimetric inequality for a special kind of polygon and then we give the main argument.

4.1 Special Hexagons

Say that a *special hexagon* is a convex hexagon whose interior angles are all $2\pi/3$. The regular hexagon is an example, and all other special hexagons are obtained from the regular one by pushing the sides in and out parallel to themselves. Our first result is a version of the isoperimetric inequality for special polygons. This result is undoubtedly well-known.

Lemma 4.1 *The length and area of a special hexagon satisfy the inequality*

$$\frac{\text{length}^2}{\text{area}} \geq 8\sqrt{3}. \quad (13)$$

Proof: It is an elementary exercise, and it is also discussed at length in [T], that the area of such a hexagon is a quadratic function of the side lengths ℓ_1, \dots, ℓ_6 . By symmetry, the desired function is symmetric in the arguments. Therefore, we have

$$\text{area} = a(\ell_1 + \dots + \ell_6)^2 + b(\ell_1^2 + \dots + \ell_6^2).$$

Considering the unit regular hexagon and the unit equilateral triangle, which is a limiting case, we get the relations

$$36a + 6b = \frac{3\sqrt{3}}{2}, \quad 9a + 3b = \frac{\sqrt{3}}{4}.$$

Solving these equations leads to the formula

$$\text{area} = \frac{1}{6\sqrt{3}}(\ell_1 + \dots + \ell_6)^2 - \frac{1}{4\sqrt{3}}(\ell_1^2 + \dots + \ell_6^2). \quad (14)$$

If we hold the perimeter fixed, we maximize the area by minimizing the sum of the squares of the lengths. This happens when all lengths are equal, and the ratio of interest is scale-invariant. Plugging in $\ell_1 = \dots = \ell_6 = 1$ and calculating, we get the advertised inequality. ♠

4.2 The Main Argument

We define an *Eisenstein polygon* to be a convex polygon whose sides are contained in the 1-skeleton of the unit equilateral triangulation. Up to scaling, Eisenstein polygons are all limits of the special hexagons considered in the last section. Let $f(P)$ denote the perimeter of an Eisenstein polygon and let $F(P)$ denote the number of triangles it contains. As an immediate consequence of Lemma 4.1, we have

$$\frac{f^2(P)}{F(P)} \geq 6. \quad (15)$$

Let \mathcal{C} be a good coloring of some triangulation \mathcal{T} with degree sequence $6, \dots, 6, 2, 2, 2$. Let f denote the fold count and let F be the number of faces. Our goal is to show that $\eta(\mathcal{C}) = f^2/F \geq 3$. Note that \mathcal{C} partitions \mathcal{T} into a finite number of monochrome regions with connected interior.

Lemma 4.2 *Each monochrome region is isometric to an Eisenstein polygon.*

Proof: Let P be such a region. Without loss of generality assume that P is colored white. To make the argument clearer, we shave off the outer ϵ of P , so that its boundary does not self-intersect. Let P_ϵ be this slightly smaller polygon. Since both colors occur at the triangles touching each degree 2 vertex, we see that P_ϵ does not contain any degree 2 vertices. Hence P_ϵ is locally Euclidean.

Consider ∂P_ϵ . Each vertex v_ϵ of ∂P_ϵ is within ϵ of a unique vertex v of P . Because P comes from a good coloring, v has at most 3 white triangles around it. Hence P_ϵ is locally convex near each boundary component.

Consider a lift \tilde{P}_ϵ of P_ϵ to \mathbf{C} , the orbifold universal cover of our flat cone sphere. This lift is also locally convex near its boundary components. If \tilde{P}_ϵ is unbounded then the local convexity near the boundary forces it to be an infinite strip. But then P_ϵ is an annulus with geodesic boundary. This is impossible in our flat cone sphere. Hence \tilde{P}_ϵ is bounded. The local convexity forces \tilde{P}_ϵ to have just one boundary component. If the map $\tilde{P}_\epsilon \rightarrow P_\epsilon$ is not a global isometry then P_ϵ contains some degree 2 vertices, a contradiction. We conclude that \tilde{P}_ϵ is a convex polygon and the map $\tilde{P}_\epsilon \rightarrow P_\epsilon$ is an isometry. Hence P_ϵ is a convex polygon.

Letting $\epsilon \rightarrow 0$ we see that P is also a convex polygon. But then P is also an Eisenstein polygon. ♠

Let P_1, \dots, P_k be the white connected regions. By Lemma 4.2, each P_j is an Eisenstein polygon. Let $f_j = f(P_j)$ and $F_j = F(P_j)$. We have

$$\min_j \frac{f_j^2}{F_j} \geq 6, \quad f = f_1 + \dots + f_k, \quad \frac{F}{2} = F_1 + \dots + F_k. \quad (16)$$

The middle equality comes from the fact that each black-white interface in particular lies in the boundary of some P_j , and that ∂P_i and ∂P_j are disjoint except perhaps for vertex intersections when $i \neq j$. Hence

$$\frac{2f^2}{F} = \frac{(f_1 + \dots + f_k)^2}{F/2} \geq \frac{f_1^2 + \dots + f_k^2}{F/2} = \frac{f_1^2}{F_1} \oplus \dots \oplus \frac{f_k^2}{F_k} \geq 6. \quad (17)$$

Here \oplus denotes the Farey sum: We add the numerators and we add the denominators. As is well known, the Farey sum of some fractions lies between the minimum and the maximum of the fractions.

5 References

[**T**], W. P. Thurston, *Shapes of polyhedra and triangulations of the sphere*, Geometry and Topology monographs 1: pp 511–549 (1998)

[**W**], D. B. West, *Introduction to Graph Theory, 2nd Ed.*, Prentice-Hall (2000)

[**Wo**] S. Wolfram, *Mathematica* (2020) wolfram.com/mathematica.