

Symplectic Tiling Billiards, Planar Linkages, and Hyperbolic Geometry

Richard Evan Schwartz *

July 27, 2023

Abstract

The purpose of this paper is to unite two games, symplectic billiards and tiling billiards. The new game is called symplectic tiling billiards. I will prove a result about periodic orbits of symplectic tiling billiards in a very special case and then show how this result is related to planar linkages and hyperbolic geometry.

1 Introduction

Billiards, of course, needs no introduction. However, it has two exotic cousins which are less well known, *symplectic billiards* and *tiling billiards*. In this paper I will unite these two topics. I call the new game *symplectic tiling billiards*. Perhaps anyone who knows about both symplectic billiards and tiling billiards could stop reading now and define symplectic tiling billiards for themselves just based on the name.

For ease of exposition I will stick to the polygonal cases of all these topics. Symplectic billiards is perhaps best played on a pair of polygons, A and B , as shown in Figure 1. Starting with a pair $(a_1, b_2) \in \partial A \times \partial B$ one produces a pair $(a_3, b_4) \in \partial A \times \partial B$ using the rule below.

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




Figure 1: Symplectic Billiards Defined

In words, the line connecting a_1 to a_3 is parallel to the side of B containing b_2 and the line connecting b_2 to b_4 is parallel to the side of A containing a_3 . One then iterates and considers the dynamics. I first learned about symplectic billiards from Peter Albers and Serge Tabachnikov. We later wrote a paper [ABSST] about the subject, proving a few foundational results. There is also upcoming work of Fabian Lander and Jannik Westermann on symplectic polygonal billiards. The two-table perspective appears in their work.

Tiling billiards is a variant of billiards played on the edges of a planar tiling. Figure 2 shows the rule.




Figure 2: Tiling Billiards

The rule is essentially the same as for billiards, except that the trajectory refracts through the edges rather than bouncing off them. I first learned about tiling billiards from Serge Tabachnikov. Now there is a growing literature on the subject. One recent paper is [BDFI]. A later version of this paper will have a more complete bibliography.

In §2 I will define *symplectic tiling billiards* and make a few general remarks about it. I will also show the results of a few easy experiments.

In §3 I will consider a special case of this game, in which the planar tilings involved consist of a finite number of infinite sectors bounded by rays emanating from the origin. I will prove a result about periodic orbits in this situation.

The periodic orbits result from §3 turns out to be useful for giving a clean correspondence between convex planar mechanical linkages and convex polygons with fixed angles. This correspondence is akin to the one given by Kapovich and Millson [KM], but it is more direct. The Kapovich-Millson correspondence involves the Riemann mapping theorem while the one here is elementary.

One can use our correspondence (or the one given by Kapovich-Millson) to give a hyperbolic structure to the moduli space of planar linkages. The hyperbolic structure here comes from Thurston's approach to the moduli spaces of polyhedra given in [T]. In other words, the correspondence simply allows one to import the Thurston hyperbolic structure to the linkage case. I will explain all this in §4.

I was inspired to think about this topic while listening to a great talk given by Juergen Richter-Gebert [R-G] about the special case of equilateral pentagons. Richter-Gebert has a different way to give a hyperbolic structure in this case, which perhaps is more canonical. In any case, our method works in general and is readily computable.

This paper is a meal that I threw together based on ideas that were dropped on my plate while I dined in Heidelberg and Marseille this summer. I would like to thank Peter Albers, Diana Davis, Aaron Fenyes, Fabian Lander, Juergen Richter-Gebert, Sergei Tabachnikov, and Jannik Westermann for stimulating conversations about this material. I probably had the key idea while in free-fall riding the Hurricane Loop waterslide at Miramar water park in Weinheim. I would also like to thank the University of Heidelberg for their continued support, in the form of a Mercator Fellowship, and CIRM (Luminy) for their support. I'd also like to thank the National Science Foundation for their continued support.

2 Symplectic Tiling Billiards

2.1 Basic Definition

For us, a *tiling* is a subdivision of the plane into convex polygonal regions. These polygonal regions are allowed to be unbounded. For instance, an acute sector bounded by two rays emanating from the same point would count as a polygonal region for us.

Given a tiling A , a *particle* on A is a point contained on the interior of an edge of A , together with a choice of a direction pointing into one of the two regions of A adjacent to e . One could encode the direction by a vector transverse to e .

We say that two tilings A and B are *transverse* if no edge of A is parallel to an edge of B . Suppose that (A, B) are transverse and (a_1, b_2) are a pair of particles, with a_1 being a particle of A and b_2 being a particle of B .

We define a_3 as follows. The line connecting a_1, a_3 is parallel to the edge of B containing b_2 . The direction at a_3 goes in the same direction as the direction at a_1 . That is, one and the same vector along the line $\overline{a_1 a_3}$ would serve as a transverse vector. We define b_4 in the same way, swapping the roles of A and B . Figure 3 shows the construction.




Figure 3: Symplectic Billiards Defined

2.2 Remarks on the Definition

Here are some comments about the basic definition.

Unbounded Tiles: One subtle point about this definition is that the point a_3 or the point b_4 might be defined in case A or B has unbounded tiles. What happens here is that the relevant ray simply heads off to infinity without intersecting an edge of the tiling. We allow this, and indeed it might present an interesting case to study, but the squeamish reader could avoid this problem by only considering tilings with bounded tiles.

Lack of Transversality: One might also want to consider the case when A and B are not transverse. For instance, one might like to play this game on a single tiling, setting $A = B$. As in symplectic billiards, one requires that the particles a_1 and b_2 are not contained in parallel edges.

Affine Symmetry: Like symplectic billiards, symplectic tiling billiards is affinely natural. If T is an affine transformation of the plane, then T maps the orbits relative to the pair (A, B) to the orbits relative to the pair (A', B') where $A' = T(A)$ and $B' = T(B)$. Also, if T is a dilation then the orbits relative to (A, B) are the same as the orbits relative to (A, B') .

It might be interesting to study symplectic tiling billiards on tilings which have affine symmetry. These are called affine crystallographic groups.

Half Translation Surfaces: Symplectic tiling billiards can also be played on a torus. This is equivalent to considering the game relative to a pair of doubly periodic tilings, and then considering the orbits on the quotient space.

More generally, one can play the game relative to a half-translation surface. Recall that a half-translation surface is a metric on a surface in which all but a discrete set of points are locally isometric to the Euclidean plane and the remaining points are cone points having cone angle πk for various integers k . One additional requirement for these surfaces is that there is a global parallel line field. (This is not quite implied by the other conditions.)

Let's say that a tiling of a translation surface is a decomposition of the surface into convex polygons such that every cone point appears as a vertex. Other points might be vertices as well. Choices of globally parallel line fields would give a way to line up the two tilings.

2.3 Rotated Square Grids

Here I will report on some preliminary experiments I did with I did not experiment too much with symplectic tiling billiards in the case when A and B are both square grids. In this case, the only parameter is the way A and B are rotated with respect to each other. The program I wrote uses exact rational arithmetic, so I will restrict my attention to that case.

Given $t \in \mathbf{Q}$ define

$$z_t = \left(\frac{1-t^2}{1+t^2}, \frac{2t}{1+t^2} \right). \quad (1)$$

This is the usual rational parametrization of the unit circle. We normalize so that A is the usual square grid. Let A_t denote the result of multiplying the usual square grid by z_t .

For the pair (A, A_0) , which is the same as just playing the game on A , the orbits are just rays. The choice $t = 1/3$ yields $z_t = (4/5) + (3/5)i$. This is the simplest non-trivial rational case. It is based on the $(3, 4, 5)$ right triangle. Let's take a look.

Figure 4 shows a picture of an orbit with respect to $(A, A_{1/3})$. The picture on the right shows a close-up of the most complicated part of the orbit on the right. (I have also continued the orbit a bit further on the right.) The slightly thicker segments on the right are actually unions of extremely close parallel segments.




Figure 4: The left half of an orbit on $(A, A_{1/3})$.

The orbit seems to be bounded and aperiodic. I did not formally prove the aperiodicity, but my rational calculations reveal that the numerators and denominators of the coordinates of the vertices are tending to ∞ . Geometrically, the orbit has an attracting limit cycle. All this would not be hard to prove; one would look at the first return map to a suitable interval, get an interval exchange transformation, and check that it had an attracting fixed point. After playing around for a while, I conjecture that all orbits on $(A, A_{1/3})$ are bounded and get attracted to a limit cycle. (The limit cycle itself is a periodic orbit.)

Figure 5 shows a picture of an unbounded orbit on $(A, A_{7/11})$. I picked this parameter somewhat randomly. After making two big spirals, the orbit starts heading southeast in a periodic pattern with a drift. The orbit is clearly unbounded, though I did not write down a formal proof.




Figure 5: A symplectic billiard orbit on $(A, A_{7/11})$.

Some parameters seem to support both unbounded orbits and bounded orbits. I intend to write a program which colors the points in the phase space of the system according to whether they have a bounded or unbounded orbit. I bet this will be a pretty picture.

For irrational parameters, one can sometimes get entirely periodic orbits. For instance, if we play on (A, B) , where B is obtained by rotating A $\pi/4$ radians, then all orbits are periodic, and they make squares in each factor.

3 A Local Result

3.1 Statement of Results

Say that an N -sunburst is a union of N -rays emanating from the origin such that the convex hull of the rays is the whole plane. In other words, the rays do not lie in any halfplane. An N -sunburst defines a tiling in the plane in which the tiles are unbounded acute sectors based at the origin. In this section we will consider symplectic billiards with respect to two N -sunbursts A and B . The number N is supposed to be the same in both cases.

Let ρ be some rotation about the origin. Let $A' = \rho(A)$. We call the system (A, B') a *phase modification* of (A, B) . This idea is similar to the idea we considered in the previous chapter, in connection with square grids.

Figure 6 illustrates an example of what we call a left-convex orbit. Here (A, B) is a pair of 5-sunbursts and the shown outer billiards orbit is periodic when projected onto A . When projected onto B the orbit is very likely not-periodic; it would make an infinite spiral. I have colored in the convex region bounded by the orbit. Comparing the colors of the sectors on the left with the colors of the rays on the right, you can see how the orbit works. The pink ray on the right is parallel to the boundary of the pink sector, etc.




Figure 6: A left-convex orbit with respect to a pair of 5-sunbursts.

The shown orbit is what I call *left-convex*: Its projection onto A is the boundary of a convex polygon. One could define *right-convex* orbits in the

same way, but this extra definition does not really add anything to the discussion. Notice that our concept is invariant with respect to dilation about the origin. Thus, if one orbit is left-convex, they all are.

We think of A as a collection $A_1, A_3, \dots, A_{2N-1}$ unit vectors, oriented counterclockwise, so that one rotates A_{i-1} counterclockwise less than π radians to arrive at A_{i+1} . In other words A_{i-1} and A_{i+1} make a positively oriented basis. Likewise we define B_2, B_4, \dots, B_{2N} .

We say that (A, B) is an *oriented weave* if, for all k , the vector B_{2k} lies in the interior of the sector defined by A_{2k+1} and $-A_{2k-1}$. There is no typo in the indices. What I mean is that B_2 lies in the sector defined by A_3 and $-A_1$, etc. From Figure 6 it is pretty clear that if (A, B) has a left-convex orbit then (A, B) is an oriented weave. Here is the main result.

Theorem 3.1 *Let (A, B) be an oriented weave. Then there is a unique phase modification (A, B') of (A, B) such that all the orbits with respect to (A, B') are left-convex.*

A slightly more general version of Theorem 3.1 would say that if (A, B) has a phase modification which is an oriented weave then (A, B) has a unique phase modification which has left-convex orbits.

To make this result useful we give a sufficient criterion for (A, B) to have a phase modification which is an oriented weave. We call A *regular* if A_1, \dots, A_{2N-1} form a regular N -gon. We call B *balanced* if the center of mass of the unit vectors B_2, \dots, B_{2N} is the origin.

Theorem 3.2 *If A is regular and B is balanced then (A, B) has a phase modification which is an oriented weave.*

Theorem 3.2 applies equally well to the case when A is balanced and B is regular. In case this is not clear already, the argument given in Lemma 3.4 below will make it clear. So, let us state the corollary we want to use.

Corollary 3.3 *If A is balanced and B is regular then (A, B) has a unique phase modification (A, B') with left-convex orbits.*

In the next chapter we will apply Corollary 3.3 to the study of convex equilateral polygons.

3.2 Proof of Theorem 3.1

As we have already mentioned, the symplectic tiling billiards orbit of a pair of particles might not be defined in case the tiling has some unbounded regions. The next lemma deals with this detail. To start our orbit, we pick $a_1 \in A_1$ and $b_2 \in B_2$. We make a_1 into a particle by choosing a vector which points into the sector bounded by A_1 and A_3 . Likewise for b_2 we choose a vector which points into the sector bounded by B_2 and B_4 .

Lemma 3.4 *The orbit of (a_1, b_2) is forever defined and consists of particles $a_1, b_3, a_3, b_4, \dots$ with $a_{2k-1} \in A_{2k-1}$ and $b_{2k} \in B_{2k}$ for all $k = 1, 2, 3, \dots$*

Proof: We take our indices mod N . It is easier to think about the argument when $N \geq 5$ but actually all we need is $N \geq 3$.

Let us call (A, B) a *weave* if, for all k , the line through B_{2k} does not lie in the closed acute sector bounded by A_{2k-1} and A_{2k+1} . If (A, B) is an oriented weave then (A, B) is a weave.

The condition that (A, B) is an oriented weave implies that the vectors $-A_3, -B_2, +A_1$ and $+A_5, +B_4, -A_3$ are in counterclockwise order (with acute gaps between them). But then notice that the vectors $+B_4, -A_3, -B_2$ occur in counterclockwise order. The same goes when we shift the indices by $2k$ for any k . This says that $(B, -A)$ is an oriented weave. Hence (B, A) is also a weave.

In short (A, B) and (B, A) are both weaves. But now we see inductively that the orbit of (a_1, b_2) is well defined and has the desired properties. The weave conditions, applied one after the other in alternating fashion, guarantee that the relevant line intersects both boundaries of the relevant sector.

♠

Now we discuss the phase modifications. For each index, there is an open interval of rotations which keep that part of the oriented weave condition. The intersection I of all these intervals is an open interval containing B .

We define a function $h : I \rightarrow \mathbf{R}_+$ which we call the *holonomy*. For simplicity we think of I as directly parametrizing the N -sunbursts B' such that (A, B') is an oriented weave. For each $B' \in I$ there is a dilation by $h(B')$ which maps a_k to a_{2k+N} . In other words, the orbit makes a spiral and $h(B')$ measures how much the spiral expands when it goes once around. We orient I so that as we move along I the sunburst B rotates clockwise.

The pair (A, B') has left-convex orbits if and only if $h(B') = 1$. The next two lemmas below show respectively that h is strictly monotone increasing and that h achieves all positive value on I' . These two results therefore combine to prove Theorem 3.1.

Lemma 3.5 *The function h is strictly monotone increasing on I .*

Proof: We consider the orbits $a_1, b_2, a_3, b_4, \dots$ and $a_1, b_2, a'_3, b'_4, \dots$ with respect to B and B' . Here B' is obtained by rotating B a tiny bit clockwise. Figure 5 shows the beginning of the orbit. We can see from the picture that a'_3 is further from the origin than a_3 .




Figure 7: Monotonicity of the holonomy

We can write $a'_3 = \lambda_3 a_3$, with $\lambda_3 > 1$. Shifting the indices and applying the same argument, we get constants $\lambda_5, \lambda_7, \dots$, all greater than 1, such that $h(B')/h(B) = \lambda_3, \dots, \lambda_{2k+1}$. ♠

Lemma 3.6 *The function h achieves every value in \mathbf{R}_+ .*

Proof: Consider what happens as we move to the left endpoint of I . One of the vectors B_{2k} , is becoming nearly equal to $-A_{2k-1}$. But this means that the point a_{2k+1} converges to the origin. Following this, the remaining points a_{2k+3}, \dots remain near the origin. In other words, $h \rightarrow 0$ in this case.

Consider what happens as we move to the right endpoint of I . One of the vectors B_{2k} is becoming nearly equal to A_{2k+1} . But then the point a_{2k+1} tends to ∞ . The remaining points a_{2k+3}, \dots also tend to ∞ . Hence $h \rightarrow \infty$ in this case.

Our result now follows from the continuity of h and the Intermediate Value Theorem. ♠

This completes the proof of Theorem 3.1.

3.3 A Calculus Interlude

Before proving Theorem 3.2 we prove two useful inequalities.

Lemma 3.7 *Let $f(t) = \sin(\pi t) - t/(1-t)$. Then f is positive on $(0, 1/2)$.*

Proof: Let $I = (0, 1/2)$. Note that we have $f(0) = f(1/2) = 0$. We write $f = f_1 + f_2$ where $f_1(t) = \sin(\pi t)$ and $f_2(t) = -t/(1-t)$. For $t \in I$ we have

$$f_1''(t) = -\pi^2 \sin(t) < 0, \quad f_2''(t) = -\frac{2}{(1-t)^2} - \frac{2t}{(1-t)^3} < 0.$$

Hence $f'' < 0$ on $(0, 1/2)$. But these conditions guarantee that $f > 0$ on I . ♠

Lemma 3.8 *Let $N \geq 3$ and let $j = 2, 3, \dots$ be such that $j-1 < N/2$. Then*

$$\sin(\pi(j-1)/N) > \frac{j-1}{N-j+1}.$$

Proof: Put $t = (j-1)/N$ and apply Lemma 3.7. ♠

3.4 Proof of Theorem 3.2

We first recall a special case of Helly's Theorem. If we have a finite collection of open intervals on the line which pairwise intersect, then they all intersect.

Now we launch into the proof of Theorem 3.2. Let (A, B) be a pair of N -sunbursts. Let S^1 be the set of all phase modifications (A, B') of (A, B) . As above, we think of S^1 as directly parametrizing B . For each index $k = 1, 2, 3, \dots$ there is an interval $I_k \subset S^1$ with the following property: All $B' \in I_k$ are such that B'_{2k} lies in the acute sector bounded by A_{2k+1} and $-A_{2k-1}$. The intersection $I = \bigcap I_k$ is the set of phase modifications which are oriented weaves. We just need to show that this intersection is nonempty when A is regular and B is balanced.

Let us first consider the case of a pair (A, B^*) where both are regular. It is convenient to introduce the constant

$$\mu_N = \pi - \frac{2\pi}{N}. \tag{2}$$

In this case, we have $I_1^* = I_2^* = I_3^* \dots$, by symmetry, and all these intervals have (angular) length μ_N . The point here is that the angle between A_{2k-1}^* and A_{2k+1}^* is $2\pi/N$ and hence the angle between A_{2k+1}^* and $-A_{2k-1}^*$ is μ_N .

Now we replace B^* by B , a balanced N -sunburst. The new intervals I_1, \dots, I_n have the same length as I_1^*, \dots, I_n^* but their positions may have moved. We will show for all indices i, j , that the relative notion of I_i and I_j is less than μ_N . In other words, if we pick a frame of reference in which one of the intervals does not move at all, the other one moves by less than μ_N . Once we have these pairwise intersections, the existence of a global intersection point follows from Helly's Theorem, applied to lifts of these intervals to the universal cover of S^1 . To be careful, we take a homotopy from B^* to B , so that the lifts are initially identical, and then vary the lifts continuously.

After cyclically relabelling, it suffices to consider the case when $i = 0$ and $j = 1, \dots, N - 1$. Let

$$\theta_j^* = 2j\pi/N$$

denote the angle between B_0^* and B_{2j}^* . Likewise let θ_j denote the angle between B_0 and B_{2j} . The relative notion we are talking about will be less than μ_N provided that

$$|\theta_j - \theta_j^*| < \mu_N.$$

Since $\theta_j^* \leq \mu_N$ and $\theta_j > 0$ we have

$$\theta_j^* - \theta_j < \mu_N.$$

The other direction is harder.

Lemma 3.9 $\theta_j - \theta_j^* < \mu_N$.

Proof: We will assume that $\theta_j - \theta_j^* \geq \mu_N$ and derive a contradiction.

Our last assumption implies that

$$\theta_j \geq \pi + 2(j-1)\pi/N.$$

We first treat two special cases.

- When $j = 1$ we have $\theta_1 \geq \pi$. This is an immediate contradiction because the angle θ_1 measures the acute angle between B_0 and B_2 .
- When $j - 1 \geq N/2$ we have $\theta > 2\pi$. This is a contradiction.

So, we only need consider the case $j = 2, 3, \dots$ and we can assume that $j - 1 < N/2$. These are the same conditions that arise in Lemma 3.8.

We rotate the picture so that B_0 and B_{2j} have the same Y -coordinate and B_0 lies to the right of B_{2j} . Given our angle bounds, the common Y -coordinate is at most

$$-\sin(\pi(j-1)/N).$$

Going counterclockwise from B_0 to B_{2j} we encounter $j - 1$ points of B strictly between them. The Y -coordinates of all these points sum to at most

$$u_{N,j} = j - 1.$$

The remaining $N - j + 1$ points have Y -coordinate at most that of B_0 and B_{2j} . The sum of the Y -coordinates of these remaining points is at most

$$v_{N,j} = -(N - j + 1) \sin(\pi(j-1)/N).$$

To finish our proof we just have to prove that $u_{N,j} + v_{N,j} < 0$ for all $N \geq 3$ and all $j = 2, 3, \dots$ with $j - 1 < N/2$. The inequality we seek is

$$\sin(\pi(j-1)/N) > \frac{j-1}{N-j+1}.$$

This is exactly Lemma 3.8. ♠

This completes the proof of Theorem 3.2.

4 Linkages and Hyperbolic Geometry

4.1 Equiangular and Equilateral Polygons

We consider convex planar polygonal linkages in which all the links have the same length. These are precisely the same thing as convex unit equilateral polygons. Let P be such an N -gon. We can associate to P a balanced N -sunburst. The rays are generated by the unit vectors $V_1, V_3, \dots, V_{2N-1}$ parallel to the sides of P when it is oriented counterclockwise. These vectors sum to 0 and hence the corresponding sunburst A_P is balanced.

We let B_P be the regular N -sunburst. By Corollary 3.3 we may rotate B_P in a unique way, arriving at B'_P , so that the pair (A, B'_P) has left-convex orbits. By construction, any left-convex orbit P^* is a convex N -gon whose interior angles are all the same. Thus, the assignment $P \rightarrow P^*$ is a canonical transformation from the set of convex equilateral N -gons mod isometry to the set of convex equiangular N -gons modulo similarity.

We can say more succinctly that P^* is an equiangular N -gon inscribed in the sunburst generated by the direction vectors of P , and that P^* is unique up to dilating about the origin.

Remark: As I mentioned in the introduction, J. Millson and M. Kapovich [KM] also give a correspondence between equilateral N -gons and equiangular N -gons. Their construction is more general than the one above, but also transcendental. They start with an equiangular N -gon, and then take the Riemann map to the unit disk. This gives them N unit complex numbers u_1, \dots, u_N . They then use the fact (Springborn's Theorem) that there is a unique point in the hyperbolic plane such that a Moebius transformation M mapping this point to the origin carries u_1, \dots, u_n to unit complex numbers whose sum is 0. The numbers u'_1, \dots, u'_n are then interpreted as the direction vectors of a unit equilateral N -gon. Here we have set $u'_k = M(u_k)$. This construction produces a unique equilateral N -gon up to scale.

The Millson-Kapovich correspondence is invertible, and the inverse would play the same role as our correspondence. Our correspondence is much more elementary, and does not rely on things like the Schwarz-Christoffel transformation.

4.2 Hyperbolic Structure

William Thurston's paper [T] constructs complex hyperbolic structures on spaces of flat cone spheres. A special case of a flat cone sphere is the double of a convex equiangular polygon. The corresponding subspace sits as a totally real slice of the convex hyperbolic moduli space. This imparts a real hyperbolic structure on the space of convex equiangular N -gons that is easy to compute. In the next section we will give an elementary (and well known) account of a direct construction which goes not through complex hyperbolic space.

Using our correspondence above we give a real hyperbolic structure to the space of equilateral N -gons modulo isometry. This method produces a more elementary hyperbolic structure than what one would get from the Kapovich-Millson correspondence. It also produces a more elementary structure, in the case of pentagons, that Richter-Gebert described in his talk [R-G]. The correspondence there also has a highly transcendental construction and is therefore difficult to compute.

All the constructions identify the space of unit equilateral convex pentagons with the interior of a regular right-angled hyperbolic pentagon. The whole space of unit equilateral convex pentagons is a genus 4 surface and we can use symmetry (which amounts to permuting the direction vectors in all ways) to tile this surface by 24 hyperbolic pentagons. The completion is a hyperbolic genus 4 surface. The degenerate pentagons are parametrized by points which lie in the boundaries of the hyperbolic pentagonal tiles.

The hexagonal case is also quite pretty. We start with a regular ideal octahedron and cone the 4 faces to the center of mass. This gives us 4 smaller tetrahedra, each with 3 ideal vertices and one interior vertex. We take 2 of these small tetrahedra and glue them together along their ideal faces. This gives us a hyperbolic triangular bi-pyramid with 2 interior vertices and 3 ideal vertices. Thurston's hyperbolic structure is the interior of this hyperbolic triangular bi-pyramid. I guess that 120 of these pieces can be glued together to give a hyperbolic structure to the full space of unit equilateral hexagons, though I did not think this through.

I think that the cases $N = 7, 8, 9, \dots$ are also pretty, but I have not thought much about them. We will see below that the corresponding hyperbolic polytopes have many right-angled sides. Given Vinberg's Theorem on reflection groups, however, they do not generate hyperbolic lattices in general.

4.3 A Concrete Description

The description given above of Thurston's hyperbolic structure is rather complicated. Various authors, including Ghys and Bavard, have given an elementary account which works directly with polygons. For the sake of exposition, I will sketch the elementary construction. I hope that this sketch makes clear the readily computible nature of the structure. Again, we are working with convex equiangular N -gons. The construction does not require convexity, though at the end we will specially discuss the convex case.

We start with N parallel families of lines, with each family being parallel to a different N th root of unity. These families are cyclically ordered, according to the roots of unity. We interpret an equiangular N -gon as a selection ℓ_1, \dots, ℓ_n of lines, one from each family. The vertices of the N -gon are given by $\ell_1 \cap \ell_2, \ell_2 \cap \ell_3$, etc. This interpretation gives a natural identification of the space of equiangular N -gons with \mathbf{R}^N . To get a concrete coordinatization we could pick some line L in the plane, not parallel to any of the families, and then use the intersection $\ell_1 \cap L, \dots, \ell_N \cap L$ give N linear coordinates. In other words, we are identifying \mathbf{R}^N with $L \times \dots \times L$. A different choice of L would give us a linear change of coordinates.

We now mod out by translations. This identifies the space of equiangular N -gons mod isometry with \mathbf{R}^{N-2} . Figure 8 shows, in the pentagon case, how we can introduce concrete coordinates on \mathbf{R}^{N-2} which are linear functions of the coordinates discussed above. Any other construction like this would result in a linear change of variables.




Figure 8: Coordinates on the space of equiangular pentagons

It is important to emphasize that these coordinates A, B, C are signed distances. They look nice in the convex case but they are defined even in the non-convex cases. For instance,

$$C = ((\ell_1 \cap \ell_3) - (\ell_1 \cap \ell_4)) \cdot (1, 0).$$

Now we consider the signed area in these coordinates. We have

$$\text{area} = -\alpha A^2 - \beta B^2 + \gamma C^2,$$

where α, β, γ are positive constants that do not depend on the choice of pentagon. In other words, the area of $V = (A, B, C)$ is $Q(V, V)$ where Q is a quadratic form of signature $(1, 2)$. When we restrict to the unit area pentagons, we get a hyperboloid in $\mathbf{R}^{2,1}$, and this is the usual Lorenz model of the hyperbolic plane.

Now let us explain how all this interacts with convexity. Figure 9 shows the *butterfly move* B_2 for pentagons.




Figure 9: The butterfly move B_2 .

In general, if $P = \ell_1, \dots, \ell_N$ then $B_k(P) = \ell'_1, \dots, \ell'_N$ where $\ell'_j = \ell_j$ for all $j \neq k$ and the two lines ℓ_k, ℓ'_k are equidistant from the intersection $\ell_{k-1}\ell_{k+1}$. The operation B_k is linear and preserves signed area. Hence B_k is a Lorenz transformation and induces a hyperbolic isometry on the hyperbolic structure we have explained. Moreover B_k is an involution. The fixed point set of B_k is the set of all degenerate polygons in which $\ell_{k-1}, \ell_k, \ell_{k+1}$ have a common point. This is a codimension one set. Hence B_k is a hyperbolic reflection. Note that B_a and B_b commute as long as a, b are not cyclically consecutive. The corresponding fixed point sets are perpendicular hyperplanes.

For pentagons, the 5 fixed lines in the hyperbolic plane f_1, f_3, f_5, f_2, f_4 are consecutively perpendicular in the cyclic sense. By symmetry these are the extensions of the edges of a regular right-angled pentagon. The interior of this pentagon is the space of convex unit equiangular pentagons modulo similarity.

5 References

[**ABSST**] P. Albers, G. Banhatti, P. Sadlo, R. Schwartz, S. Tabachnikov, *Polygonal Symplectic Billiards*, preprint 2019.

[**BDFI**] P Baird-Smith, D. Dais, E. Fromm, S Iyer, *Tiling Billiards on Triangle Tilings, and Interval Exchange Transformations*, Bulletin of the London Math Society, 2020

[**KM**] M. Kapovich and J. Millson, *On the Moduli Space of Polygons in the Euclidean Plane*,

[**R-G**] J. Richter-Gebert, *Mercator Workshop Lecture*, Heidelberg, 2023

[**T**] W. Thurston, *Shapes of Polyhedra*.