Technische Universität Berlin

Bachelorarbeit für die Prüfung zum Bachelor of Science im Studiengang Mathematik

The secondary fan of a punctured Riemann surface

Verfasser:
Joseph Helfer
Matrikelnummer 334487

Erstgutachter:
Prof. Boris Springborn

Zweitgutachter:
Prof. Yuri Suris

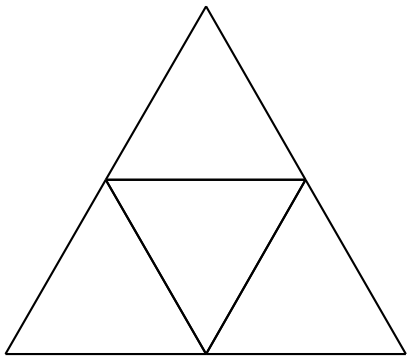

30. Oktober, 2013

1 Introduction

In 1994, Gelfand, Kapranov, and Zelevinsky [4] described a construction which associates to a finite set of points in the plane a certain polytope, called the "secondary polytope", such that the different faces of the polytope correspond to different ("nice") triangulations of the point set, and the combinatorial structure of the polytope represents the relations between these different triangulations.

Now, their purpose in defining this object is rooted in some very complicated algebraic geometry, which is far beyond the scope of this paper. Nevertheless, mathematicians love to classify things, and so a nice geometric classification of triangulations of a point set has an appeal in its own right.

Moreover, triangulations of punctured Riemann surfaces are of significant interest, for instance in their application by Penner [5] to the parametrization of the surface's Teichmüller space.

It turns out that the aforementioned concept of "nice" triangulations generalizes in a natural way to triangulations of punctured Riemann surfaces, and we similarly get a geometric object (the "secondary fan") which classifies all the "nice" triangulations of a punctured Riemann surface and their relationship to one another.

Here, we will present this construction, as well as a program which explicitly computes this secondary fan for a given Riemann surface, and take a look at a couple of examples which the program has produced.

2 "Nice" triangulations

To begin, we will describe the class of triangulations which Gelfand, Kapranov, and Zelevinsky managed to classify, albeit in a manner very different from how they described them.

We start by considering so-called "Delaunay triangulations".

By a triangulation of a polygon, we mean any simplicial complex whose underlying space is that polygon. A triangulation of a point set is a triangulation of its convex hull, whose vertex set is the given point set.

There are several equivalent ways of defining Delaunay triangulations. The simplest one is probably the following: a triangulation of a point set is *Delaunay* if the interior of each triangle's circumcircle contains no vertices of the triangulation. Roughly speaking, the main appeal of Delaunay triangulations is that they avoid "skinny" triangles, which turn out to be for various reasons undesirable. Known applications of Delaunay triangulations include the generation of meshes in computer graphics, the finite element method for numerically solving partial differential equations, and the visualization of landscapes [3].

In addition, aside from their applications, Delaunay triangulations are nice in the sense that they have a lot of useful characteristics which are also provable (as opposed to mesh algorithms which are only heuristically good).

One of the first things one might like to prove is that a Delaunay triangulation always exists and is effectively computable. To do this, we first introduce another

characterization of Delaunay triangulations which will be useful to us in the sequel. Here, we describe a condition not on the triangles of the triangulation (or their circumcircles), but rather on edges lying between two triangles. Any such edge defines a quadrilateral (see figure 1).

We call such an edge *Delaunay* if the sum of the two angles in the quadrilateral at the ends of the edge is greater than or equal to the sum of the two angles opposite the edge. (In figure 1 we must have $\angle kil + \angle kjl \geq \angle ikj + \angle ilj$). The edge is called *strongly Delaunay* if this inequality is strict. Otherwise, it is *weakly Delaunay*.

It can be shown (see, for example, [3]), that a triangulation is precisely Delaunay when all of its edges are Delaunay.

The nice thing about this characterization is its relation to "flipping".

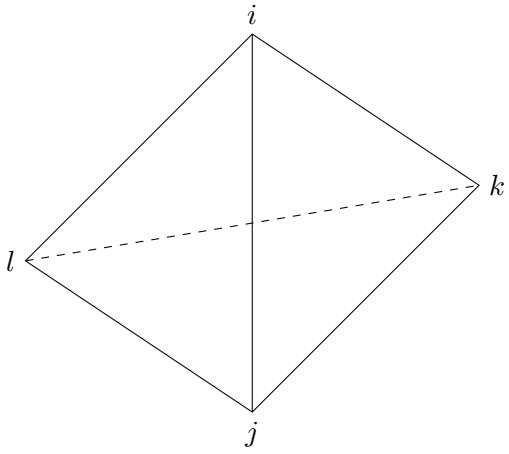

Figure 1: A pair of triangles before and after flipping

A *flip* is an operation with which triangulations can be modified to produce new triangulations. We again consider the quadrilateral formed by the two triangles adjoining an edge. If the quadrilateral is convex, we can remove the separating edge, and replace it with the quadrilateral's other diagonal (see figure 1).

It is then easy to see that an edge is strongly Delaunay precisely when its flipped counterpart is not Delaunay, and is weakly Delaunay precisely when its flipped counterpart is as well. Edges which cannot be flipped (i.e., are the diagonal of a non-convex quadrilateral) can easily be seen to be strongly Delaunay.

Given this, a very simple-minded algorithm for producing Delaunay triangulations – called the "flip algorithm" – presents itself: we start with an arbitrary triangulation of our point-set – it can easily be shown that one always exists – and then we repeatedly search for edges which are not Delaunay and flip them. Fortunately, this simple-minded algorithm *works*: we can prove that it terminates (in $O(\#vertices^2)$ time) and produces a Delaunay triangulation.

One of the ways to prove that this works is by using a powerful tool for characterizing Delaunay triangulations called the *parabolic lifting map*. Here, we place an elliptic paraboloid above the plane containing the triangulation, and lift each vertex onto it

(that is, the point (x, y) in the plane gets mapped onto the point $(x, y, x^2 + y^2)$ on the paraboloid). We then take the convex hull of the lifted point set, and project the 1-skeleton of its underside back down onto the plane. This produces a partition of our point set into polygonal regions. We then obtain a triangulation by arbitrarily refining the partition in any regions which aren't already triangles. It turns out, remarkably, that the triangulations obtained in this way are *precisely* the Delaunay triangulations (this, by the way, immediately gives us a proof for the existence of Delaunay triangulations, because we can always take a convex hull). See [3] for a proof.

One then shows that the Flip Algorithm terminates because each flip turns a concavity into a convexity, and thus lowers the lifted triangulation, so no edge which has been flipped once can reappear. That means, each edge can be flipped at most once, giving the $O(n(n-1)/2) = O(n^2)$ bound on the run-time ($n(n-1)/2$ is the maximal number of edges in a triangulation of n vertices). See [3] for more details.

As it happens, we are actually not particularly interested in Delaunay triangulations themselves, but rather a generalization thereof, known as "weighted" Delaunay triangulations. Here, we assign each vertex a real number (its "weight") and then, after lifting it to the paraboloid, pull the point back down by that amount. That is, if the vertex at (x, y) has the weight w, it gets lifted to $(x, y, x^2 + y^2 - w)$. We then call a triangulation "weighted Delaunay" (for given weights on each vertex) if it is the projection of the underside of the convex hull of this modified lifted set of points (again, possibly after a refinement of non-triangular regions). Note here that some vertices might lie *inside* the convex hull, in which case they will not appear in the resulting projection. We still allow this, so that weighted Delaunay triangulations are actually triangulations of *subsets* of the given point set.

We will return later on to this description of weighted Delaunay triangulations, but first we move onto the secondary polytope, which we are now ready to describe.

We can see that for the description of weighted Delaunay triangulations, the paraboloid is actually immaterial. Instead of assigning a weight to each point, by which we displace it from the paraboloid, we can just go right ahead and assign an arbitrary height to each point.

A weighted Delaunay triangulation is then any which is the projection of the underside of the convex hull of *some* set of points lying above the plane. This is, by the way, what Gelfand, Kapranov, and Zelevinsky call a *coherent* triangulation (although they do it the other way around, considering the upper side of the convex hull), and what we referred to earlier as a "nice" triangulation.

The construction of the secondary polytope (of a given point set in the plane) now proceeds as follows. First we enumerate the points. We then associate to each triangulation of (a subset of) the point set a point in $\mathbb{R}^n$ (where n is the number of points) by assigning to the i-th coordinate the sum of the areas of all triangles containing the i-th point as a vertex. We do this for all possible triangulations, and then take the convex hull of the resulting set of points in $\mathbb{R}^n$. The convex polytope that we obtain is the secondary polytope.

Gelfand, Kapranov and Zelvinsky showed that the triangulations corresponding to vertices of the resulting polytope are precisely the coherent triangulations. They ad-

ditionally described the relation between any two triangulations whose vertices in the secondary polytope are connected by an edge. This is however somewhat involved, and not of great importance to us; the corresponding relation between neighbouring triangulations of a punctured Riemann surface will be considerably simpler. We should, however, mention, that in the simplest case, their more complicated relation reduces to our simpler relation.

The object of which we will be producing an analogue for Riemann surfaces is not the secondary polytope, but its "normal fan", the *secondary fan*.

A *polyhedral cone* in $\mathbb{R}^n$ is any region which is the intersection of finitely many closed half-spaces containing the origin. A *polyhedral fan* is a collection of polyhedral cones which cover the space, and such that each face of each cone is in the collection, as well as the intersection of any two cones.

Any bounded polytope gives rise to a polyhedral fan, called its *normal fan*. To each facet of the polytope, we take the cone of all $v \in \mathbb{R}^n$ such that the function $\langle v, \cdot \rangle$ is maximized on that facet (where $\langle \cdot, \cdot \rangle$ is the standard scalar product on $\mathbb{R}^n$).

It's not hard to see that this construction associates to each k-dimensional face an $(n-k)$-dimensional cone.

Thus the secondary fan has a maximal cone (that is, of dimension n) for each coherent triangulation of the point set, and triangulations corresponding to neighbouring cones are related in some nice way which we haven't described.

3 Punctured Riemann surfaces

We now move on to the construction of a secondary fan for a punctured Riemann surface.

Here, we consider a closed orientable surface with a finite number of punctures – that is, removed points – with a prescribed hyperbolic structure – that is, a complete metric of constant negative curvature and finite area. Here, the set of punctures is what we will be triangulating, and it is analogous to the point set we started with in the plane.

It should be noted that the prescription of a hyperbolic structure is equivalent to the prescription of a conformal structure on the surface – that is, this makes the topological surface into a Riemann surface. This is why we speak of the "secondary fan of a Riemann surface".

By a *triangulation* of such a Riemann surface, we mean a choice of non-intersecting geodesics connecting the punctures, which partitions the surface into topological triangles with the corners removed.

When we refer to a triangulation on a punctured *topological* surface, we mean the same thing but without the restriction that the curves be geodesics. Where there is possible ambiguity, we will refer to the latter concept as a *topological triangulation*.

Note that we are not interested in punctured topological surfaces which can't be triangulated or which don't admit a complete hyperbolic structure, and whenever we refer to a "surface", we mean one for which both of these things are possible.

Here, we are again only interested in "nice" triangulations. To decide whether a triangulation is "nice", we must first give it some additional structure, analogous to the

weights in a weighted Delaunay triangulation. It is well known that each puncture on a hyperbolic surface forms a "cusp", a kind of infinitely long horn which gets thinner and thinner. By cutting off these cusps at a certain height, we truncate the sides of our triangles (which are currently infinitely long) and give them a finite length. A choice of "weights" for a triangulation is then essentially a choice of heights at which we cut off the cusps. A triangulation together with a choice of cusp truncations will then be "nice" if its truncated triangles satisfy a certain property (which we will again call the Delaunay condition, and with good reason).

3.1 Delaunay triangulations of punctured Riemann surfaces

To proceed, we'll need to pass to the universal cover of our surface. Now, since our surface is by assumption hyperbolic, the uniformization theorem tells us that its universal cover is the hyperbolic plane. That is, the surface is the quotient space of the hyperbolic plane by a discrete group of hyperbolic isometries. Thus, we can visualize the surface as a tiling of the hyperbolic plane (much as wee can see the torus as a tiling of the Euclidean plane by rectangles).

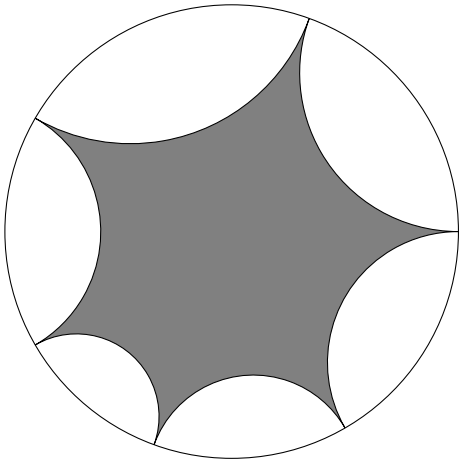

Figure 2: Surface displayed in the hyperbolic plane. We regain the original surface either by gluing the sides of this figure, or by tiling the plane with it and then taking the quotient of the plane by the corresponding discrete group of isometries.

One "tile" in the Poincaré disk model of the hyperbolic plane (in which geodesics are arcs of circles intersecting the disk's circumference orthogonally) would look something like figure 2. Note that the punctures must lie on the circle at infinity. We see now that in this picture, the triangles of our triangulation are actually *ideal hyperbolic triangles* – that is, a topological triangle (minus the vertices) in the hyperbolic plane bounded by three geodesics meeting pairwise in the circle at infinity (as in figure 3). Note that all the angles of such a triangle are zero, and each side is infinitely long.

Next, we recall that a *horocycle*, in the Poincaré disk model of the hyperbolic plane, is a (Euclidean) circle in the disk which is tangent to the circle at infinity (that is, the

circumference of the disk) at one point. We will now see that the "truncation of cusps" referred to in the previous section is actually a choice of a horocycle at each cusp.

Let us consider, for example, an ideal hyperbolic triangle, together with a horocycle at each vertex (see figure 3). If we consider the part of any side of the triangle lying between the two horocycles, we see that it has some finite length. For the side between vertices i and j, we denote this length by λ_{ij}, as in figure 3. We should note at this point that we also allow the two horocycles to overlap, in which case λ_{ij} is still defined as the length of the part of the side lying between the two horocycles, and is taken in this case to be negative. We will primarily make use of the quantities $\ell_{ij} := e^{\frac{1}{2}\lambda_{ij}} > 0$. Perhaps it would be most fitting to call this quantity the edge's *exponential length*, but as we will not have occasion to talk about any other sorts of edge-lengths, we will just refer to it as the edge's *length*.

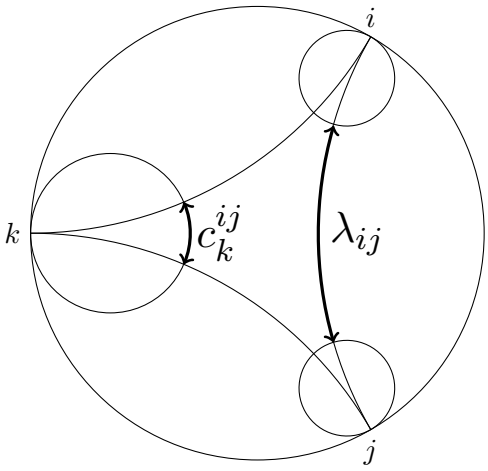

Figure 3: An ideal hyperbolic triangle with horocycles

We can also see that the choice of a horocycle gives our triangle something like angles (that is, some positive real number assigned to each corner), which we call *pseudo-angles*; we simply take the length of that part of the horocycle which lies within the triangle. In a triangle with vertices i, j, and k, we denote the pseudo-angle at k by c_k^{ij}, as in figure 3.

The quantities c and ℓ are related by the simple formula

$$c_i^{jk} = \frac{\ell_{jk}}{\ell_{ij}\ell_{ik}} \tag{1}$$

(a very simple proof of this is given in [2]).

Before we move on, let us gather one more important fact about these newly-defined quantities: the behavior of the lengths ℓ with respect to "flips". Here, we consider an ideal hyperbolic quadrilateral divided by an edge into two triangles (see figure 4). We want to derive a formula for ℓ_{jl} in terms of the lengths ℓ of the other sides.

Now, according to (1), we have

$$\ell_{jl} = c_k^{jl} \ell_{kl} \ell_{jk} \tag{2}$$

On the other hand, the pseudo-angle c_k^{jl} is clearly the sum of the two pseudo-angles which it is split into by the edge ik:

$$c_k^{jl} = c_k^{il} + c_k^{ij} \tag{3}$$

But, again by (1), these are given by

$$c_k^{ij} = \frac{\ell_{ij}}{\ell_{ik}\ell_{jk}} \quad \text{and} \quad c_k^{il} = \frac{\ell_{il}}{\ell_{ik}\ell_{kl}} \tag{4}$$

Putting (2), (3) and (4) together, we obtain

$$\ell_{ik}\ell_{jl} = \ell_{ij}\ell_{kl} + \ell_{il}\ell_{jk} \tag{5}$$

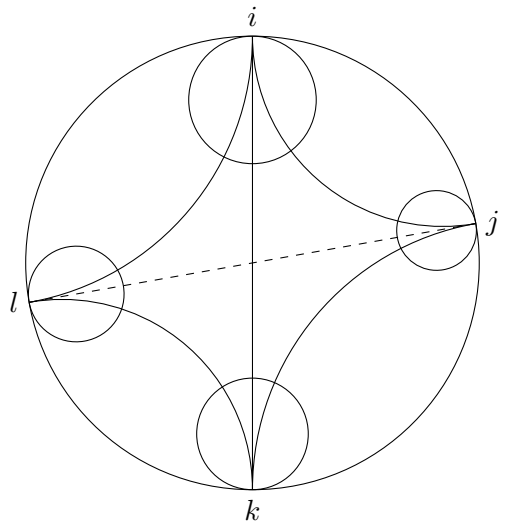

Figure 4: A hyperbolic flip

We now define a *weighted triangulation* of our surface as a triangulation, together with a choice of a particular horocycle at each cusp.

Unfortunately, this definition doesn't actually make any sense, because the choice of a horocycle at a cusp depends on the particular representation of the surface in the hyperbolic plane. But: given one such horocycle at a cusp in one particular "tile", we get one at every other tile by taking the Γ-orbit of that horocycle (where Γ is the discrete group of isometries defining the surface). That is, from the one horocycle $h \subset \mathbb{H}^2$ we obtain the family of horocycles $\Gamma h = \{\gamma(h) \subset \mathbb{H}^2 \mid \gamma \in \Gamma\}$[1].

[1]There is one technical difficulty with this construction: for it to be well-defined, Γh cannot have two horocycles centered at the same point at infinity. That is, any $\gamma \in \Gamma$ fixing a point at infinity corresponding to a cusp must also fix horocycles at that point at infinity (such an element is called "parabolic"). However, [5] states this without proof or further comment, so presumably it's obvious.

So we *actually* define a *weighted triangulation* as a triangulation, together with a choice of a Γ-orbit of horocycles for each cusp. We call these Γ-orbits of horocycles *weights*.

Note that, although the specification of a Γ-orbit of horocycles sounds like a pretty daunting task, we can actually parametrize these choices in quite a concrete way: we simply add up all the pseudo-angles of corners of triangles at a given vertex, which we call the *pseudo-angle-sum* at that vertex. (Intuitively, this is just the circumference of the circle bounding the truncated cusp). Clearly, there is exactly one horocycle for each such pseudo-angle-sum, and so a choice of Γ-orbits of horocycles at each cusp is really nothing more than a choice of positive real numbers.

The next task is to describe when such a weighted triangulation is *Delaunay*.

For this, we turn to the hyperboloid model of the hyperbolic plane. Recall that the hyperboloid model represents the hyperbolic plane as the hyperboloid $x_1^2 + x_2^2 - x_3^2 = 1, x_3 > 0$ in $\mathbb{R}^3$ with the metric induced by the Minkowski inner product in $\mathbb{R}^3$:

$$\left\langle \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}, \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} \right\rangle = x_1 y_1 + x_2 y_2 - x_3 y_3$$

The hyperboloid model is related to the Poincaré disk model by stereographic projection through the origin onto the disk $x_1^2 + x_2^2 \leq 1, x_3 = 1$.

The *light-cone* L^+ in the hyperboloid model is the set $\langle x, x \rangle = 0, x_3 > 0$. This is just the cone with slope 1 surrounding the hyperboloid.

Each vector in the light-cone lies above a point at infinity in the disk model and corresponds to a horocycle centered at that point (we say that a horocycle is *centered* at the corresponding point at infinity because each geodesic perpendicular to the horocycle passes through that point). Specifically, for a vector $x \in L^+$, the corresponding horocycle is the set of all points y in the hyperboloid with $\langle x, y \rangle = -1$.

Thus, our choice of "weighted vertices" – that is, of (Γ-orbits of) horocycles – corresponds to a choice of points in the light-cone lying under the hyperboloid. When we were considering plane weighted Delaunay triangulations, our choice of weighted vertices corresponded to a choice of points lying under the paraboloid. There, we obtained a weighted Delaunay triangulation by taking the convex hull of those points, and considering the induced polygonal complex (and, after an arbitrary refinement, triangulation) on its underside.

In the same vein, we can take the convex hull of the points representing our horocycles. The underside of this convex hull again induces a polygonal complex on the points in the light cone. By connecting the corresponding points at infinity with geodesics, we then obtain a partition of the hyperbolic plane into ideal polygons. Taking the quotient under Γ to obtain the surface, this descends to a polygonal complex on the surface (see [5] for a proof). We then, again, arbitrarily refine the partition in any non-triangular regions to obtain a triangulation.

We therefore call a weighted triangulation of the surface *Delaunay* if it arises as a triangulation obtained by the above construction for the choice of horocycles corresponding to that triangulation's weights.

Of course, with this definition, we have no practical way of checking whether a given weighted triangulation is Delaunay.

Fortunately, the above condition is equivalent to a much simpler one. We recall that one characterization of plane Delaunay triangulations was that at each edge adjoining two triangles, the sum of the two angles of the quadrilateral which touch the edge must be greater than or equal to the sum of the other two angles (see figure 1 and the text pertaining to it). The condition here is *exactly the same*, simply with the angles of plane triangles replaced with the pseudo-angles of our truncated hyperbolic triangles.

That is, for any two bordering triangles ijk and ikl (as in figure 4), we must have

$$c_k^{ij} + c_l^{ij} \leq c_i^{jk} + c_i^{jl} + c_j^{ik} + c_j^{il} \tag{6}$$

A proof of the equivalence of these conditions can again be found in [5].

We again call any edge ik for which the condition (6) holds *Delaunay*, and if the inequality is strict, then *strongly Delaunay* (and otherwise *weakly Delaunay*). We again have that an edge is strongly Delaunay precisely when its flipped counterpart is not Delaunay, and is weakly Delaunay precisely when its flipped counterpart is as well.

We note that a triangulation is obtained directly from the convex hull construction (that is, rather than being the refinement of a polygonal decomposition) precisely when all its edges are strongly Delaunay (which is, again, proved in [5]).

3.2 The secondary fan of a punctured Riemann surface

We are now ready to construct the secondary fan of a punctured Riemann surface.

We start with some surface with n punctures (labeled 1 to n) and some fixed triangulation on it, together with some arbitrary weights at each cusp. We now modify these weights, and would like to see for which modifications the given triangulation becomes (or remains) Delaunay.

Specifically, we modify the weight at i-th vertex by scaling all of the light-cone vectors in the corresponding orbit of horocycles by some positive factor x_i. Now, it can be shown that the edge-lengths ℓ_{ij} are proportional to the lengths of these vectors (see [5] for a proof). That is, after scaling, we have new side-lengths

$$\tilde{\ell}_{ij} = x_i x_j \ell_{ij}$$

Correspondingly, we have new pseudo-angles, given by (1):

$$\tilde{c}_i^{jk} = \frac{\tilde{\ell}_{jk}}{\tilde{\ell}_{ij}\tilde{\ell}_{ik}} = \frac{1}{x_i^2} c_i^{jk}$$

We recall that the Delaunay condition was given by the inequalities (6) for the pseudo-angles, one inequality for each edge. The condition on the new, modified pseudo-angles is then

$$\frac{1}{x_k^2} c_k^{ij} + \frac{1}{x_l^2} c_l^{ij} \leq \frac{1}{x_i^2}(c_i^{jk} + c_i^{jl}) + \frac{1}{x_j^2}(c_j^{ik} + c_j^{il})$$

This is a homogeneous linear inequality in the variables $u_i := \frac{1}{x_i^2}$. That is, it defines a closed half-space in the space $\mathbb{R}^n_{>0}$ of all values for the u_i's. All the inequalities together therefore define a *polyhedral cone* in this space. (Note that above, we only defined polyhedral cones in $\mathbb{R}^n$, but the definition is identical in $\mathbb{R}^n_{>0}$).

At any boundary of this cone, one or more inequalities become equalities; that is, the corresponding edges are weakly Delaunay. By passing to the other side of the boundary, the original triangulation is no longer Delaunay, but the one obtained by flipping all the edges which were weakly Delaunay at the border, is. This new triangulation then defines a new cone bordering the old one.

By considering all possible triangulations, we then get a partition of $\mathbb{R}^n_{>0}$ into cones (that is, a *fan*), such that each cone corresponds to a triangulation (or – as we'll see – several triangulations), and two neighbouring triangulations are related by one or more flips. This fan is, of course, what we call the *secondary fan* of the Riemann surface.

Note that we really do obtain a partition of the whole space, because for any choice of weights, there is a triangulation which is Delaunay with respect to those weights – namely, the one produced by the convex hull construction with the corresponding horocycles.

As a final remark before we move on, our definition of the weight modifications x_i depended on our starting with an initial arbitrary choice of weights, which is somewhat undesirable; for each such choice we will get a different fan (it will be scaled by some factor in each coordinate), although combinatorially they will all be the same. To "normalize" the fans, we therefore choose the convention that the initial weights are all chosen such that the pseudo-angle-sum (as described in the previous section) at each vertex is 1. Note that, in this case, the variables u_i are precisely the pseudo-angle-sums.

4 Computing the secondary fan

4.1 A short aside about the Penner coordinates for the decorated Teichmüller space of a punctured surface

We now move on to describe the program the present author has written to compute the secondary fan of a given punctured Riemann surface.

The first question is: what do we give this program as input? That is, how can we describe an arbitrary punctured Riemann surface in some finite way?

Now, describing the topology of the surface is simple enough; we can do this simply by giving a topological triangulation. The hard part is describing the geometric structure, that is, the metric.

The solution comes to us from R.C. Penner.

In general, the space of all hyperbolic structures on a topological surface is called its *Teichmüller Space*. This space is in general a smooth manifold, each point of which corresponds to a hyperbolic structure on the surface, and whose own smooth structure encodes the way these structures can vary in a smooth way.

Penner, in [5], introduces a Teichmüller space for punctured surfaces, where each point corresponds to a hyperbolic structure *plus* a choice of weights at each cusp. He

calls this the "decorated Teichmüller space" of the punctured surface. We observe that, given a fixed topological triangulation of the surface with m edges (one can easily convince oneself using the Euler characteristic that this number is the same for any triangulation), each point in the decorated Teichmüller space (that is, each hyperbolic structure on the surface) gives us a point in $\mathbb{R}^m_{>0}$, by taking the lengths of all the edges in the triangulation. Penner goes on to show that this correspondence is a bijection (in fact, it is a homeomorphism, but that needn't concern us).

What this means for us is that (i) any punctured Riemann surface can be completely specified by giving a topological triangulation of it and a choice of positive numbers for each edge, and furthermore (ii) each such specification corresponds to a punctured Riemann surface.

4.2 The program

We come to the description of the program.

As discussed, the program takes as input a triangulated surface with assigned edge lengths. The triangulation is specified by so-called *half-edges*. If we think of an unoriented edge as being two oriented edges glued together so that the orientations cancel out (something like $\rightleftharpoons$), then a half-edge is one of these two oriented halves. By specifying, for each half-edge in the triangulation, its "partner" half-edge, along with the "successor" half-edge in the triangle on its left-hand side (which makes sense because its oriented), we get a complete description of the triangulation.

It should also be mentioned at this point that the lengths are given as rational numbers, and all subsequent computations are performed with exact arithmetic.

We know that this input determines, along with the surface, a certain choice of weights. As mentioned in the previous section, we want these weights to be normalized, in order to get consistent secondary fans. The program performs this normalization immediately. That is, for each vertex, it calculates the sum of all the pseudo-angles $c_i^{jk} = \frac{\ell_{jk}}{\ell_{ij}\ell_{ik}}$ at that vertex and then scales the length ℓ_{ij} of each edge at that vertex by the same factor, so that this sum becomes 1.

The next task is to generate the inequalities (6), which is straightforward; we iterate through all the edges and compute the six quantities c_i^{jk} according to the formula (1). The "inequality" is then really nothing more than an array with the coefficients in (6).

The real heavy lifting is then provided by a very nice library called *cddlib* developed by Komei Fukuda (and the equally nice python wrapper *pycddlib*). This library converts between the so-called *H-representation* ("half-space") and *V-representation* ("vertex") of a convex polytope. The former describes a polytope in terms of its bounding planes – which is the description given by our generated inequalities – and the latter describes the same polytope by means of the vertices of which it is the convex hull – which is better suited for visualizing the polytope, or performing further analysis, like generating the face lattice.

Applying this algorithm to our set of inequalities, we obtain the coordinates of the first cone – that is the one corresponding to the original triangulation. It remains to find the rest.

We know we can obtain all the neighbouring triangulations by flipping the edges which are weakly Delaunay at each facet.

We obtain the facets of the cone by finding, for each inequality, all the vertices at which that inequality becomes an equality (that is, all vertices which lie on the plane defined by the inequality). To find the edges which are weakly Delaunay at a given facet, we evaluate each inequality at each vertex of the facet (remembering that each of our inequalities corresponds to an edge). We then take all the edges (that is, inequalities) which are weakly Delaunay (that is, are equalities) at every vertex of the facet.

We now flip all of these and obtain a new triangulation, and with it, a new cone. The flipping is simple to implement: we just change the "partner" and "successor" pointers in each half-edge in the corresponding quadrilateral to reflect the new configuration, and then set the length of the newly flipped edge according to (5).

The program then proceeds to do a depth-first search, applying the flips for each facet of each cone, to find the rest of the cones. There is also an interactive mode, wherein the user can manually specify the flips.

One might ask why this search must terminate – that is, why there might not be infinitely many cones. For a point set in the plane, there are only a finite number of triangulations (bounded, for example, by the number of graphs on the vertex set), but on a surface, two vertices can be connected in many, even infinitely many, ways, so this doesn't hold. Fortunately, it has been proven that there are only a finite number of triangulations which can be obtained by the convex-hull construction, see [1].

5 Some secondary fans

We will now look at a couple of secondary fans, as generated by the program, and see if we can observe anything interesting about them. Note that we will only be looking at fans of thrice-punctured surfaces, because these are the ones most easily visualized. These are namely the three-dimensional fans, which, when projected from $\mathbb{R}^3_{>0}$ onto the standard 2-simplex (which we can do, because our cones are all invariant under scaling) give us a partition of an equilateral triangle by polygons.

5.1 The thrice-punctured sphere

We begin with the thrice punctured sphere. The program, we recall, needs a specific triangulation as input, along with preassigned edge-lengths.

The most obvious triangulation of a sphere with three vertices consists of two triangles and looks like a sort of pin-cushion. That is, we place the three vertices evenly spaced around a great circle and connect them with geodesics.

Topologically, the triangulation looks like figure 5.

When choosing the edge-lengths, we have to ensure that the resulting triangulation can be made Delaunay for some weights. Otherwise, the initial cone will be empty, and the algorithm doesn't work. One easy way to make a given triangulation Delaunay is to set all the edge-lengths to 1. In this case, all the terms in (6) become 1, and the inequality becomes $1 \leq 2$, which is true. We call the secondary fan of a topological

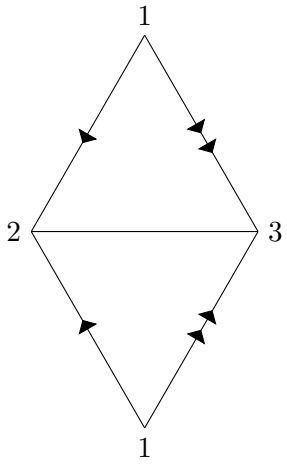

Figure 5: A topological picture of the pin-cushion triangulation of the sphere. The arrow-heads indicate the gluing and orientation of the edges.

surface which arises from this choice of edge-lengths the *standard fan* of that surface (where it is implied that we've chosen some canonical starting triangulation for the surface, which is the case for all surfaces we will investigate here).

The standard fan of the thrice punctured sphere is shown in figure 6. We see that there are three other triangulations. Each of these triangulations is obtained by flipping one of the three edges of the pin-cushion sphere.

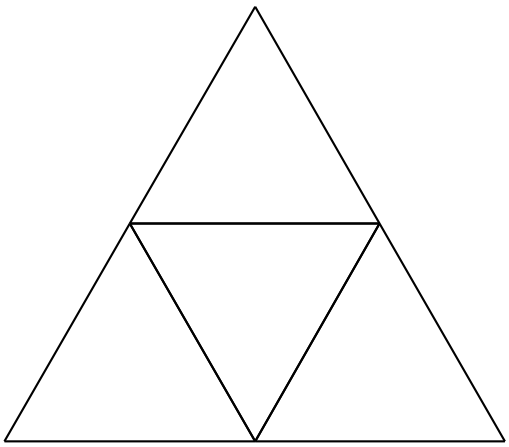

Figure 6: The standard fan of the thrice punctured sphere

The resulting triangulation of the sphere consists of two cones glued to each other along their bases. Topologically, the triangulation is as in figure 7. The upper triangle is shown on its own in 8 (this can be seen as a cone being viewed from the top). Note that the two triangles here are "degenerated", in the sense that two of their edges coincide and they only have two distinct vertices.

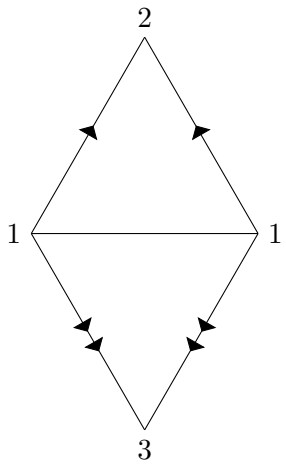

Figure 7: A topological picture of the two-cone triangulation of the sphere

Flipping the central edge in figure 7 (that is, the one between the two cones), we then return to the pin-cushion. One might ask why it isn't possible to flip the other two edges in figure 7, that is, the "degenerate" edge in figure 8 connecting the point of the cone to its base. Well, it isn't immediately obvious what this would mean. The two triangles lying on either side of one of these edges are actually *the same* triangle.

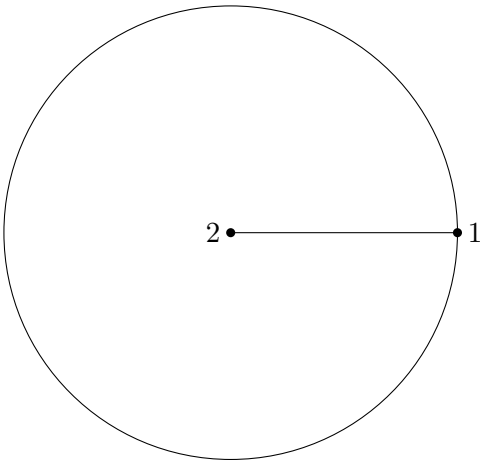

Figure 8: A triangle with one edge glued to itself

Fortunately, such an edge is *always* strictly Delaunay, so we have no need to flip it anyway, and we mustn't spend any more time trying to figure out whether this operation could be reasonably defined.

To see that an edge gluing a triangle to itself is always strictly Delaunay, let us simply write down the corresponding inequality. Consider, for example, the triangle $\triangle 1, 1, 2$ from figures 7 and 8. There are two pseudo-angles at the vertex 1, but they

both have to be equal, as they are both given by the formula

$$c_1^{12} = \frac{\ell_{12}}{\ell_{11}\ell_{12}}$$

The angle at 2 is similarly c_2^{11}.

The edge in question is the one between 1 and 2. The two pseudo-angles in the "quadrilateral" which are "opposite" to this edge are both c_1^{12}. Two of the four pseudo-angles adjacent to this edge are c_1^{12} and the other two are c_2^{12}. Altogether our inequality (6) is then

$$2c_1^{12} \leq 2c_1^{12} + 2c_2^{11}$$

which is obviously true (with strict inequality), because pseudo-angles are always positive.

We can actually see that the secondary fan of *any* thrice punctured sphere will look (combinatorially) like the one in figure 6.

To this end, let us consider the pin-cushion triangulation, but with arbitrary edge-lengths – recall, by considering all possible edge-lengths for this triangulation, we achieve all possible Riemann surfaces on the thrice-punctured sphere – and compute the corresponding cone "by hand".

Labeling the vertices as in figure 5, our three edge lengths are $\ell_{12}, \ell_{23}, \ell_{13} > 0$. Our three pseudo-angles are then

$$c_1^{23} = \frac{\ell_{23}}{\ell_{12}\ell_{13}}, \quad c_2^{13} = \frac{\ell_{13}}{\ell_{12}\ell_{23}}, \quad \text{and} \quad c_3^{12} = \frac{\ell_{12}}{\ell_{13}\ell_{23}}$$

(there are actually six pseudo-angles, three per triangle, but because the two triangles share the same sides, their corresponding pseudo-angles are equal).

To get our inequalities, we note that for any of the three edges, its surrounding quadrilateral consists of the two triangles of the triangulation, the two pseudo-angles "opposite" the edge are both the same and the four "adjacent" pseudo-angles consist of the other two pseudo-angles in the triangulation, each counted twice.

Our inequalities in the scaling variables u_1, u_2, u_3 (here we needn't worry about normalization, because we just want a combinatorial picture of the fan) are then

$$2\frac{\ell_{jk}}{\ell_{ij}\ell_{ik}}u_i \leq 2\frac{\ell_{ik}}{\ell_{ij}\ell_{jk}}u_j + 2\frac{\ell_{ij}}{\ell_{ik}\ell_{jk}}u_k$$

for $\{i,j,k\} = \{1,2,3\}$. Rearranging (and setting $l_i := \ell_{jk}^2, \{i,j,k\} = \{1,2,3\}$ for readability), we obtain

$$\begin{aligned} 0 &\leq l_3u_3 + l_1u_1 - l_2u_2 \\ 0 &\leq l_1u_1 + l_2u_2 - l_3u_3 \\ 0 &\leq l_2u_2 + l_3u_3 - l_1u_1 \end{aligned} \tag{7}$$

Setting $u_3 = 0$, we obtain

$$0 \leq l_1 u_1 - l_2 u_2$$
$$0 \leq l_1 u_1 + l_2 u_2$$
$$0 \leq l_2 u_2 - l_1 u_1$$

Setting u_1 and u_2 in turn to 0, we obtain similar inequalities. Each restriction $u_i = 0$ is one of the coordinate planes, and we see that in each case, one of the inequalities holds everywhere, and the other two become equalities simultaneously. That is, the planes corresponding to the inequalities (7) coincide pairwise on the coordinate planes, the third inequality holding everywhere on the coordinate plane where the other two coincide.

Projecting onto the 2-simplex $u_1 + u_2 + u_3 = 0$, we see that the region defined by (7) is then precisely a triangle like the central one in figure 6.

Again, the three facets of the central cone each correspond to a flip which produces a two-cone sphere, and since the resulting triangulation doesn't have any other flips, the corresponding cone must be the entire region lying opposite that facet. This ends the proof.

Another property of the thrice-punctured sphere which makes it particularly simple is that the pin-cushion and the two cones are its *only* triangulations. To see why, let's consider an arbitrary triangulation. First of all, it is easy to see that any triangulation has to have three edges and two faces. So to produce an arbitrary triangulation, we start with two triangles and glue them together along one edge. This gives us a square. We now have to pair off the other four edges and glue them together. It's now just a matter of checking all the possibilities, and doing so, we see that the pin-cushion and the two-cones are the only ones which produce a sphere.

So to recapitulate: the thrice punctured sphere has precisely 4 triangulations, and for every hyperbolic structure on it, each of these corresponds to a maximal cone, and they always have (combinatorially) the same fan.

We will see that with other surfaces, the situation is not so simple.

5.2 Hexagonal torus

It is also possible to triangulate a torus with three vertices. This can be seen most easily by representing the torus as a hexagon with its sides glued together. We then just triangulate the hexagon in the obvious way, giving us figure 9.

Again, we set the length of all the edges to 1. Generating the fan, we find, surprisingly, that it's the same fan that we got from the thrice punctured sphere (figure 6). This is rather disappointing, because the torus is considerably more complex; there are 9 edges to be flipped, instead of just three, so we would expect more cones in the fan.

Taking a closer look (with the program), we see that each of the three facets in this fan corresponds to *three* flips. Looking back at our torus, we see that this is due to the symmetries of the surface. There are, for example, three different edges connecting the vertices 1 and 2. As we've assigned the same weights to each edge, everything is

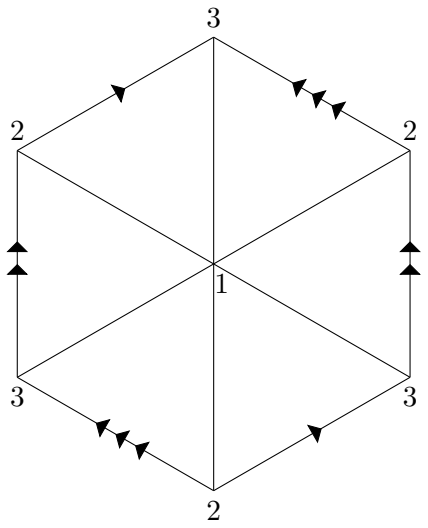

Figure 9: The hexagonal torus, triangulated. The vertices are numbered and identified appropriately.

perfectly symmetrical with respect to these three edges. When we reach a facet of the central cone where one of them is weakly Delaunay, for instance, it follows that they are all weakly Delaunay.

Suppose we flip only one or two of these edges. What is the cone of the resulting triangulation? Well, it has all the same inequalities as the central cone, except for one or two which define the half space on the other side of the facet. That is, its cone is precisely this facet.

We see, therefore, that triangulations can have cones which aren't maximal. Furthermore, each of the 6 triangulation obtained by flipping one or two of the edges corresponding to a facet all have that facet as a cone. So we also see that non-maximal cones can correspond to more than one triangulation.

By slightly perturbing the lengths we can remove some of the degeneracies in the fan. For example, by changing the length of a single edge (specifically, by changing the length of topmost vertical edge in figure 9 from 1 to 1.2), we get the fan in figure 10.

Notice that the basic structure of the first fan is preserved, but the bottom three cones have each been refined into multiple cones.

We discover another interesting phenomenon if we run the program on this new torus in interactive mode. We start off in cone A in figure 10. By performing the flip which results in configuration in figure 11-1, we get to the cone B. Another flip (producing the configuration in figure 11-2) brings us to the cone C.

We now want to continue to the cone D. The program tells us that we have to flip the two dashed edges in figure 11-2. That we have to flip more than one edge is clearly due to the symmetry of the configuration, as in the above case. What is different here, however, is that the flips don't *commute*. That is, when we had to simultaneously flip three edges above, we could flip each of the edges independently without affecting

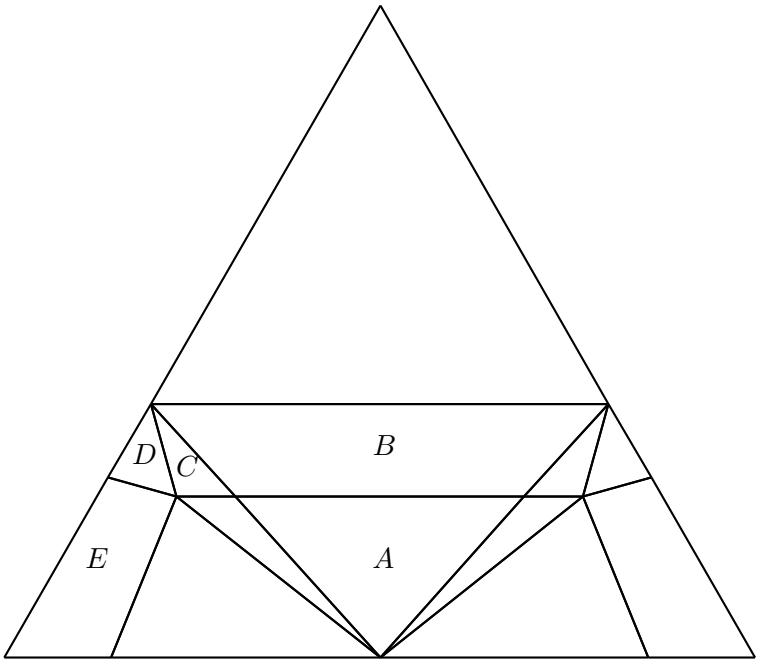

Figure 10: The standard fan of the thrice-punctured torus, with a perturbation on one edge length

the others. But here, we achieve two different triangulations depending on the order of flipping.

Performing these flips bring us to figure 11-3 and then to figure 11-4 and to the desired cone D in the fan. We see that if we had performed the flips in the opposite order, we would have arrived at the mirror image of the current configuration, which has the same cone due to symmetry. This shows us that also *maximal* cones can correspond to more than one triangulation!

The two triangulations in question are related by flipping the lower dotted edge in figure 11-4. This must mean that in this cone, that edge is *always* weakly Delaunay, because otherwise flipping it would produce a non-Delaunay triangulation.

But this seems impossible. Normally, an edge is weakly Delaunay on a plane, so it can't be weakly Delaunay on this entire maximal cone. The answer is, of course, that the inequality corresponding to this edge is trivial, so that it actually defines a whole-space (so to speak) instead of a half-space.

We can actually see that the inequality generated by this edge is trivial by looking at the diagram and remembering that all the edges have the same lengths except for the upper dotted edge. In the quadrilateral surrounding the edge in question, the two upper pseudo-angles are the same size $= c_2^{23}$, and the two lower ones (c_3^{23}) are the same as well. Also, both vertices on top are the same vertex 2, and both on bottom are the vertex 3. The generated inequality will then have the same terms on both sides:

$$c_2 \cdot u_2 + c_3 \cdot u_3 \leq c_2 \cdot u_2 + c_3 \cdot u_3$$

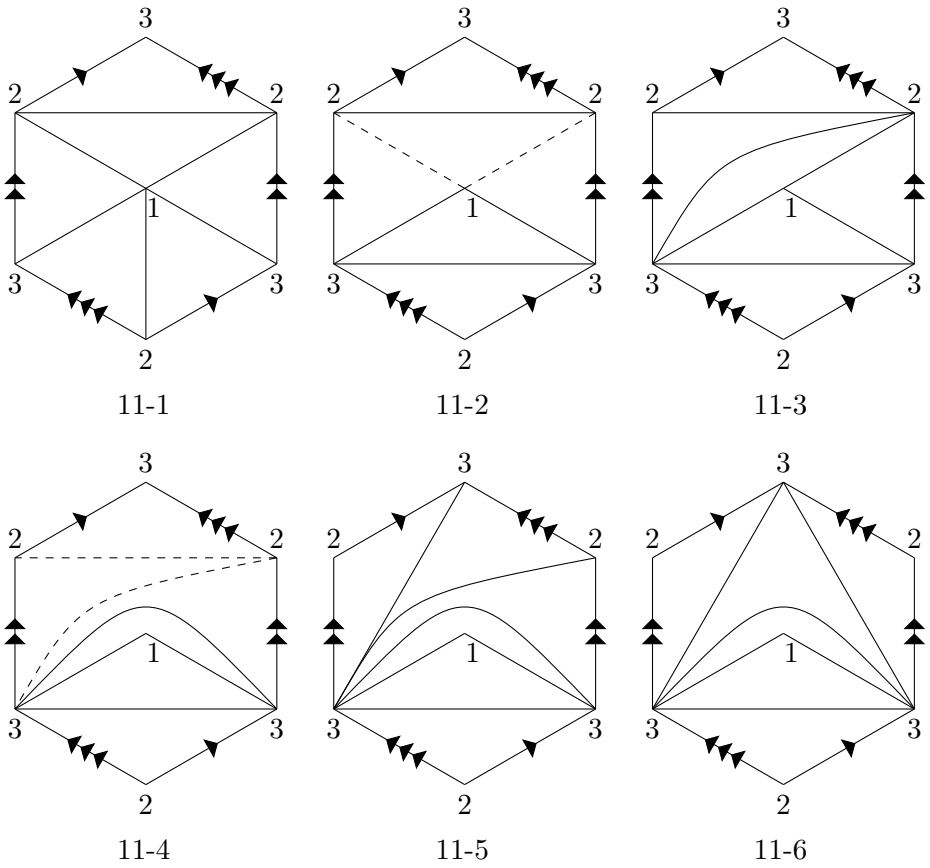

Figure 11: A sequence of flips on the torus

To move on to the next cone E in 10, we're instructed to flip the two dotted edges in figure 11-4. Now, we know that flipping the lower edge won't get us anywhere, so we flip the other one, giving us figure 11-5. Flipping the lower one then brings us to figure 11-6 and to the desired cone.

Here, again, changing the order of the flips led to different triangulations, but only one of them was the desired one. In fact, it wasn't a priori obvious that the second edge would remain weakly Delaunay in figure 11-5, so that we would still have to flip it to get to the next cone. Actually, we saw the same phenomenon above, when we made two non-commuting flips. In general, we can ask the question: given a facet of a cone where two or more non-commuting edges are all weakly Delaunay, will we always have to flip all of them to get to the next cone, or can the situation change after having flipped some of them? Unfortunately, we cannot, at present, answer this question.

It is clear that these unpleasant phenomena are all due to excess symmetry in the edge lengths. Presumably, by introducing enough asymmetry, we could avoid ever needing multiple flips at one facet, commuting or otherwise. A little experimentation confirms this. Figure 12 shows a fan with just this property (each facet corresponds to exactly one

flip). It is obtained by setting the edges connecting the vertices 1 and 2 in figure 10 to 1, 1.1 and 1.2, and setting the rest of the edges so that the configuration is symmetrical with respect to permutation of vertices.

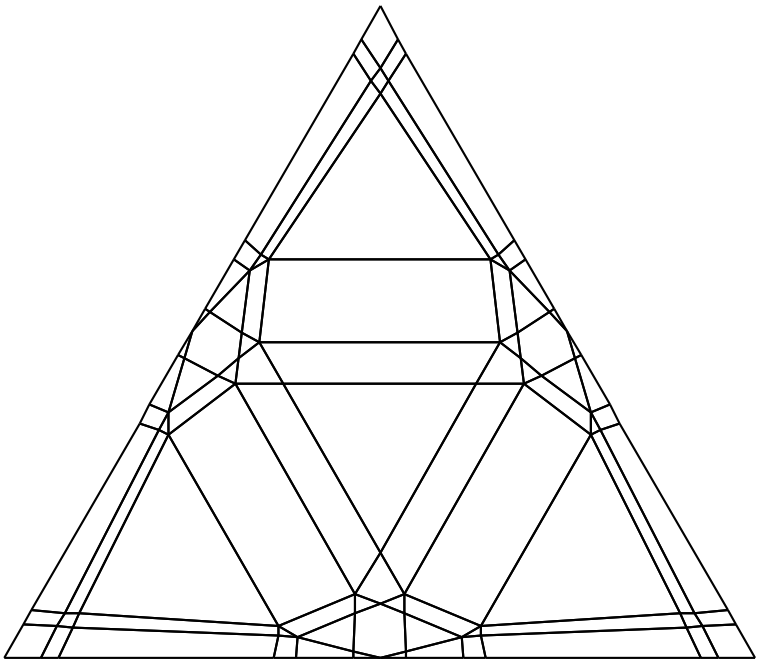

Figure 12: A secondary fan of the thrice-punctured torus with no facet-degeneracies

We also learn from the thrice-punctured torus, through experimentation, that there can be triangulations whose cones are empty, which, again, we didn't see with the thrice-punctured sphere. Indeed, if we set the lengths of two edges corresponding to the same facet in the standard fan to 1000, leaving all others 1, we get an empty cone.

In general, we have no clear picture of what the fan of the thrice-punctured torus looks like – in contrast to the thrice-punctured sphere. We don't know what all the triangulations are, and for a given hyperbolic structure, we don't know which of them will have maximal cones, which will have non-maximal cones, and which will have empty cones.

5.3 Higher-genus surfaces

In general, we can triangulate a surface of any genus n with three vertices by representing it as a $(4n+2)$-gon. For surfaces like the torus where n is odd, we simply glue together opposing edges with opposite orientations. If n is even, we glue together one pair of opposing edges (with opposite orientations), and glue the rest of the edges in groups of four according to the pattern $aba^{-1}b^{-1}$ (as in figure 13).

If we set all the edge lengths to 1, we will always get, for surfaces with an odd genus, the same kind of symmetry we had with the torus, and we obtain the same standard fan. In fact, with the surface of genus 3, if we perturb one or two edges, we get the

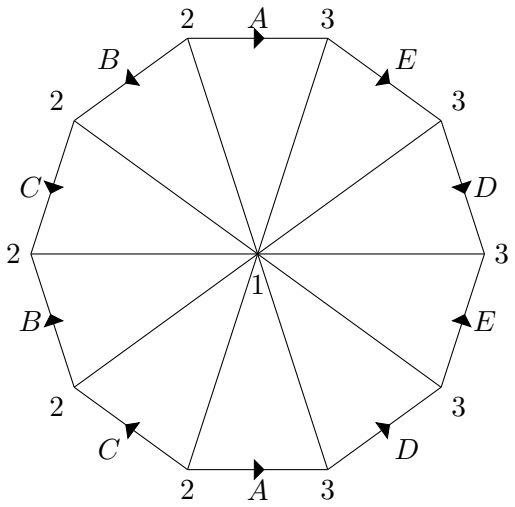

Figure 13: A topological picture of the surface of genus 2. The letters and arrows indicate which edges are glued together, and with which orientations.

exact same fans as with the torus. The difference is, of course, that many more flips are needed per facet.

Surfaces of even genus are a different story. They don't share the same symmetries. Figure 14 shows the standard fans of a genus 2 and 4 surface. They are quite similar but not identical.

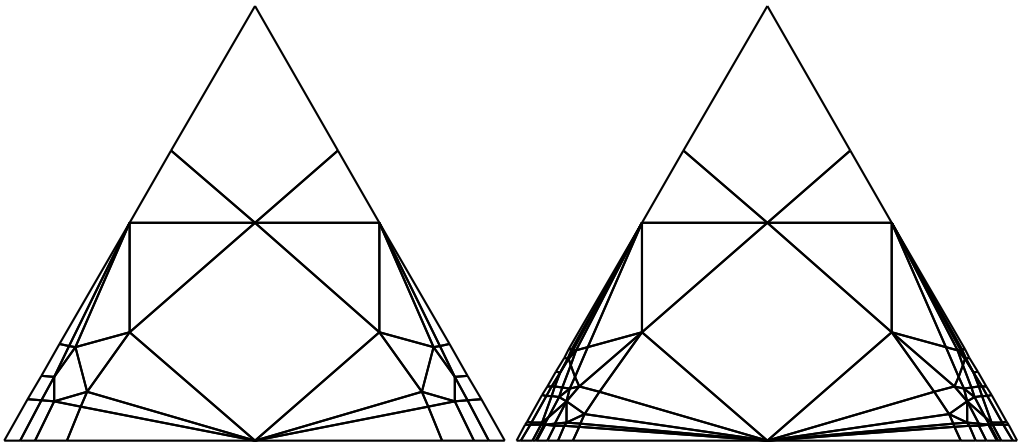

Figure 14: The standard fans of the thrice-punctured surfaces of genus two and four

We might wonder if we can get a fan free of degenerate facets (that is, facets corresponding to multiple simultaneous flips) like the one in figure 12 for higher genus surfaces. This is at least possible for surfaces of genus three, as shown in 15.

Here, the hyperbolic edge-lengths were set according to the same principle as in figure 12. That is, the edges connecting two of the vertices were set arbitrarily (to 1.0, 1.1, 1.2,

etc.), and then the rest of the edges were set so that the configuration is symmetrical with respect to permutation of the vertices. As we can see from the fan, however, the configuration is clearly *not* symmetric with respect to arbitrary permutation of the vertices, but rather only with respect to switching the two vertices represented by the lower two corners of the fan.

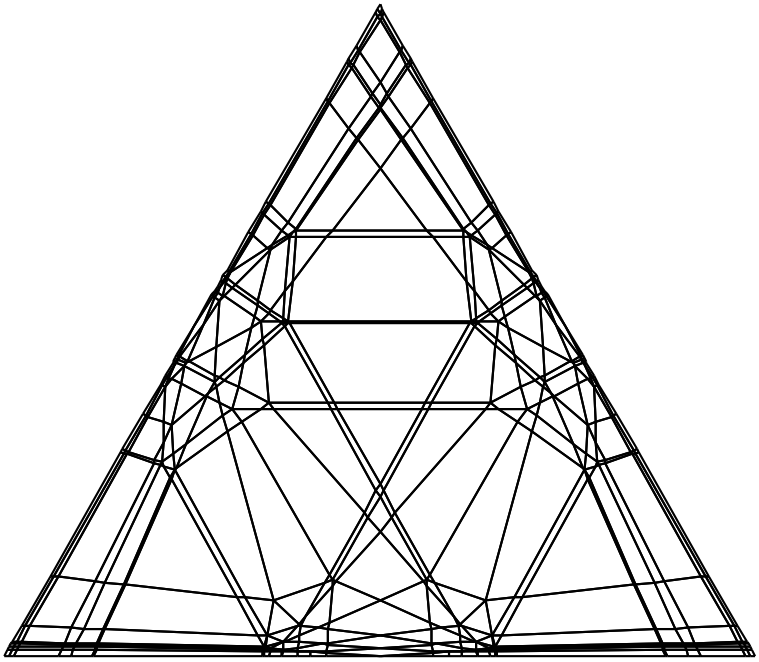

Figure 15: A secondary fan of the thrice-punctured surface of genus three with no facet-degeneracies

The problem is that the representation of the surface of genus three by a 14-gon with opposite sides glued together simply isn't as symmetrical as the representation of the torus by a hexagon with opposite sides glued together.

In the latter case, if we redraw the surface centered around one of the other vertices, we again obtain a hexagon with opposite sides glued together. However, if we redraw the surface of genus three centered around one of the other vertices, we get a different gluing.

We can similarly get a fan for the surface of genus two free of degeneracies by perturbing all edges lengths. Figure 16 shows the fan produced by setting the edges D and E in figure 13 to 1.1 and 1.2, respectively, setting the edges connecting vertices 1 and 3 to 1.1, 1.2, 1.3, etc., and setting the rest of the edges to maintain symmetry between the vertices 2 and 3.

The same pattern produces a fan with no degeneracies for the surface of genus four.

The reason that no fans for surfaces of higher genus are shown here is simply that it becomes increasingly time-consuming to write the input files for surfaces of higher genus. In principle, these could be generated automatically according to the recipe described

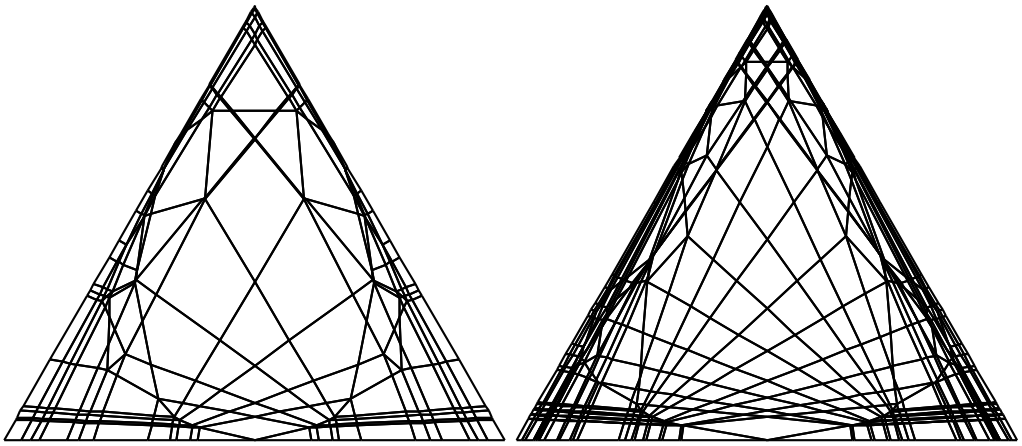

Figure 16: Secondary fans of the thrice-punctured surfaces of genus two and four with no facet-degeneracies

at the beginning of this section, but this has yet to be done.

It would, however, be interesting, for example, to see if the patterns used to produce figures 15 and 16 would also produce fans for the surfaces of genus five and higher which are free of degenerate facets (the present author suspects that it would).

5.4 More than three punctures

In principle, although they cannot be properly visualized, one could try to understand the higher-dimensional fans arising from surfaces with more than three punctures by looking at the coordinates of their cones. However, already in the simplest case – a sphere with four punctures (i.e. a tetrahedron) with all edge lengths set to one – the fan already has 47 cones, making this task quite intimidating.

We can, of course, obtain a crude visualization of such a four-dimensional fan, which is a partition of a 3-simplex after projection, by projecting it onto one of the faces of the 3-simplex. The result for the sphere with four punctures is shown in figure 17. Note that, due to symmetry, the same picture arises, regardless of which face we project onto.

It's hard to say much about such a picture, as pretty as it is. We can see a cube in the center, corresponding to the initial triangulation, each of whose six sides corresponds to one of the six edges of the triangulation.

Another interesting fact which can be mentioned for this case, is that none of the facets are degenerate (i.e. require multiple flips). We might be eager to conclude from this that degenerate facets are a special feature of three-dimensional fans.

Unfortunately, the standard fan (show in figure 18) of the torus with four punctures (a triangulation of which is obtained by representing it as a square with glued sides and drawing edges from the center to each corner and the center of each side) does have degenerate faces. Also, projection onto different faces produces different pictures, unlike the above case.

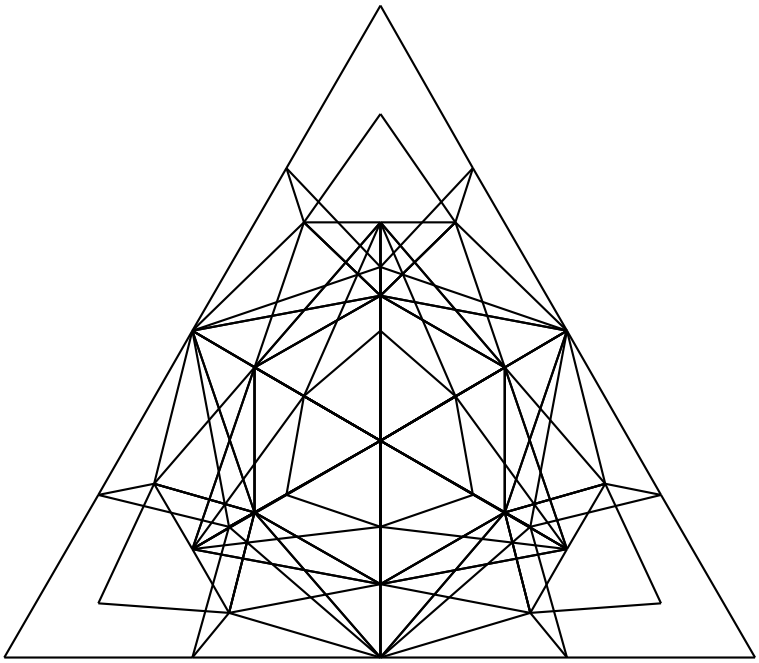

Figure 17: The standard fan of the sphere with four punctures

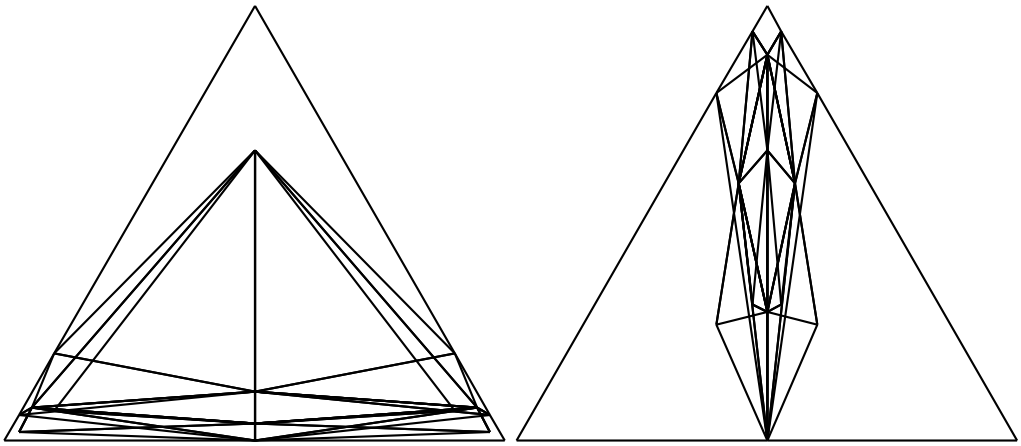

Figure 18: The standard fan of the torus with four punctures. Here, two projections are shown. The other two are the same as these.

6 The future

This investigation into the secondary fans of Riemann surfaces has generated more questions than it has answered. One question is whether any surfaces allow such a complete characterization of their secondary fans as the thrice-punctured sphere. Another is the question above regarding multiple non-commuting flips at a facet – whether they always all need to be performed, or whether performing certain flips can remove the necessity

of others (or necessitate new flips).

Another question is whether enough asymmetry can always be introduced to remove the need for multiple simultaneous flips, and whether the corresponding fan is then essentially unique to the surface, all other fans being obtained by "degenerating" certain of its cones.

There is also the question as to whether the secondary fan of a Riemann surfaces is always the normal fan to some polyhedron, as in the case of secondary fans of plane point sets.

Additionally, there may still be valuable knowledge to be gained from the software. The reader is heartily encouraged to try it out and experiment with and improve it. A three-dimensional graphical visualization for four- or higher-dimensional fans would be neat, for instance. To obtain the source code, clone the git repository at `http://www-pool.math.tu-berlin.de/~helfer/bachelor.git` or contact the author at `helfer@math.tu-berlin.de`.

7 Deutsche Zusammenfassung

Hier werden die wesentlichen Begriffe und Ergebnisse der Arbeit in deutscher Sprache zusammengefasst.

Ins Deutsche übersetzt ist der Titel "Der Sekundärfächer einer punktierten riemannschen Fläche". Dies ist eine Verallgemeinerung einer anderen Konstruktion von Gelfand, Kapranov und Zelevinsky, die auch "Sekundärfächer" heißt. Der Inhalt der Arbeit besteht darin, die ursprüngliche Konstruktion und diese Verallgemeinerung zu erläutern und ein Programm zu beschreiben, das diesen verallgemeinerten Sekundärfächer berechnet und anschließend einige Ergebnisse des Programms zu untersuchen.

Ein "Fächer" (oder genauer, "polygonaler Fächer") ist eine Menge von "Kegeln" (das heißt, skalierungsinvariante Gebiete) in $\mathbb{R}^n$, die den Raum überdecken und bestimmten Bedingungen erfüllen, analog zu denen eines Simplizialkomplexes.

Der Sekundärfächer von Gelfand, Kapranov und Zelevinsky assoziiert zu einer gegebenen endlichen Punktmenge in der Ebene einen bestimmten Fächer. Dies geschieht indirekt: zuerst wird der Punktmenge ein bestimmtes Polytop, das "Sekundärpolytop", zugeordnet, . Aus dem Polytop entsteht dann durch eine Standardkonstruktion (der "Normalfächer"zu einem Polytop) das Sekundärpolytop.

Der Sekundärfächer dient der Klassifizierung bestimmter Triangulieren der Punktmenge, der sogenannten "kohärenten Triangulierungen". Es wird nämlich gezeigt, dass jeder maximale Kegel im Fächer (d.h., ein Kegel der Dimension n, wobei der Fächer in $\mathbb{R}^n$ ist) entspricht einer kohärenten Triangulierung der Punktmenge, und zwar so, dass benachbarte maximale Kegel Triangulierungen entsprechen, die sich durch eine bestimmte einfache Transformation unterscheiden.

Zunächst betrachten wir punktierte riemannschen Flächen, d.h. solche, die (als topologisch Flächen) entstehen durch Entfernung endlich vieler Punkte von einer kompakten Fläche. Für unsere Zwecke betrachten wir nur solche Flächen, von denen genug Punkte entfernt sind, dass sie eine Metrik überall negativer Krümmung zulassen.

Hierbei wird die Punktmenge in der Ebene durch die Punktierungen der Fläche ersetzt und wir bekommen dann einen Fächer (hier, eigentlich in einem $\mathbb{R}^n_{>0}$), deren maximalen Kegel bestimmten Triangulierungen der Punktierungen der Fläche entsprechen. Hier sind die betroffenen Triangulierungen eine Verallgemeinerung von kohärente Triangulierungen, die "Delaunay Triangulierungen" heißen. In diesem Fall lässt sich die Relation zwischen Triangulierungen, deren Kegel benachbart sind, sehr einfach beschreiben: die eine entsteht aus der anderen, in dem man eine endliche Anzahl an "Flips" macht; d.h., man nimmt eine Kante aus der Triangulierung, und ersetzt sie durch die andere Diagonale des entstandenen Vierecks.

Das Programm zur Berechnung des Sekundärfächers nimmt als Eingabe eine punktierte riemannschen Fläche, zusammen mit vorgegebener Triangulierung. Dabei ist das theoretische Resultat von R. C. Penner ausschlaggebend, dass alle punktierten riemannschen Flächen vollständig durch Angabe einer Triangulierung zusammen mit bestimmten "Gewichten" (also, positiven reellen Zahlen) auf jeder Kante zu beschreiben sind.

Die Berechnung des Fächers erfolgt dann folgendermaßen: zuerst wird der Kegel

berechnet, der der Anfangstriangulierung entspricht. Die benachbarten Triangulierungen werden dann gefunden, indem die entsprechenden "Flips" gemacht werden. Eine rekursive Suche findet dann alle Kegel im Fächer.

Anschließend werden ein paar Beispielsfächer gezeigt und untersucht. Dabei werden nur Fächer für dreimal punktierte Flächen gezeigt, weil sie diejenigen sind, die man in zwei Dimensionen darstellen kann.

Für die dreimal punktierte Sphäre wird ein Fächer gezeigt, und es wird bewiesen, dass die Fächer für *alle* dreimal punktierte Sphären im wesentlich die gleiche sehr einfache und vollkommen zu verstehende Struktur haben.

Zunächst wird ein Fächer für den dreimal punktierten Torus gezeigt, der dem der Sphäre identisch ist. Es wird dann gezeigt, dass dreimal punktierte Tori auch deutlich kompliziertere Fächer haben können, mit vielen nicht trivialen Eigenschaften. Insgesamt kann man nicht so vollkommene Aussagen über den Torus machen, wie über die Sphäre.

Es werden einige Fächer von Flächen höheren Geschlechts gezeigt. Über diese ist aber noch weniger auszusagen, als über den Torus. Allerdings kann man sehen, dass alle Flächen ungeraden Geschlechts eine einfache und sehr symmetrische Metrik zulassen, so dass der Sekundärfächer derselbe einfache Fächer ist, der bei der Sphäre und dem Torus auftauchte.

Zum Schluß wird ein bisschen zu Fächern höherer Dimension gesagt. Die kann man näherungsweise in zwei Dimensionen Anschauen, aber über sie ist abermals wenig auszusagen.

References

[1] Hirotaka Akiyoshi. Finiteness of polyhedral decompositions of cusped hyperbolic manifolds obtained by the epstein-penner's method. *Proceedings of the American Mathematical Society*, 2000.

[2] Alexander Bobenko, Ulrich Pinkall, and Boris Springborn. Discrete conformal maps and ideal hyperbolic polyhedra, 2010.

[3] Siu-Wing Cheng, Tamal Krishna Day, and Jonathan Richard Shewchuk. *Delaunay Mesh Generation*. CRC Press, 2012.

[4] I. M. Gelfand, M. M. Kapranov, and A.V. Zelevinsky. *Discriminants, Resultants and Multidimensional Determinants*. Birkhäuser Boston, 1994.

[5] R. C. Penner. The decorated teichmüller space of punctured surfaces. *Communications in Mathematical Physics*, 1987.