

Discrete uniformization of finite branched covers over the Riemann sphere via hyper-ideal circle patterns

Alexander Bobenko Nikolay Dimitrov Stefan Sechelmann

Abstract

With the help of hyper-ideal circle pattern theory, we have developed a discrete version of the classical uniformization theorems for surfaces represented as finite branched covers over the Riemann sphere as well as compact polyhedral surfaces with non-positive curvature. We show that in the case of such surfaces discrete uniformization via hyper-ideal circle patterns always exists and is unique. We also propose a numerical algorithm, utilizing convex optimization, that constructs the desired discrete uniformization.

1 Introduction

In the current paper, we have constructed discrete conformal maps by applying the theory of *hyper-ideal circle patterns*. Informally speaking, a circle pattern is a discrete collection of overlapping circles together with intersection angles between adjacent circles, following the combinatorics of a given polygonal cell complex on a polyhedral surface [33, 21, 6]. In the case of *hyper-ideal circle patterns*, the collection is divided into two subsets – one assigned to the vertices of the complex and one assigned to the faces, so that the face circles are orthogonal to their vertex neighbors [26, 31, 11]. We have chosen this theory because of its existence and uniqueness theorems which hold true for fixed underlying combinatorial cell complexes, contrary to some other discrete conformal theories. The conceptual geometric rationale behind our choice is based on the combination of two equivalent characterizing properties of smooth conformal maps. First, at every point, a smooth conformal map stretches or shrinks the surface equally in all directions. Hence, at each point infinitesimal circles are mapped to infinitesimal circles. Second, smooth conformal maps preserve angles. Since circles and angles are the main ingredients of circle patterns, one can discretize the smooth theory by promoting a finite number of infinitesimal circles into circles on the surface, while keeping track of the intersection angles between pairs of neighboring ones.

The idea of using discrete collections of circles as a discrete model of conformal transformations comes from Thurston’s conjecture [35] that the Riemann mapping of a simply connected planar region can be approximated by a sequence of finer and finer regular, hexagonal circle packings [30] (a collection of touching circles with the combinatorics of the hexagonal grid on the plane). The conjecture was proved by Rodin and Sullivan [23]. After that, circle packings with more general combinatorics were used to

define various discrete analogues of holomorphic maps on the plane [30] as well as discrete analogue of the uniformization theorem [2]. Circle packings however are determined by combinatorics. In order to allow for more freedom, one can turn towards standard (Delaunay) circle patterns which provide the opportunity to incorporate the intrinsic geometry of the surfaces one works with, and more precisely, the intersection angles between adjacent circles. For example, some holomorphic transformations in the complex plane have been discretized using orthogonal circle patterns with the combinatorics of the square grid [28, 7], exploring links with integrable systems. Nevertheless, when handling polyhedral surfaces with cone singularities, standard circle patterns do not possess the necessary degrees of freedom that allow manipulation of the cone angles of the underlying geometry. This obstacle is evident in [16], where an attempt is made to construct discrete conformal maps via Delaunay circle patterns. The problem with such approach is that the definition of conformal equivalence does not actually form an equivalence relation but is more of an approximation. In contrast, *Hyper-ideal circle patterns* [26, 31] possess the necessary flexibility and, as we have shown in the current article, can be used to construct well structured discrete uniformizing conformal maps for surfaces with non-positive cone singularities as well as surfaces given as ramified covers over the sphere.

The classical uniformization theorem is a central result in the theory of complex analytic functions, conformal maps and Riemann surface theory. The uniformization theorem states that every simply connected Riemann surface is conformally isomorphic to either the Riemann sphere $\hat{\mathbb{C}}$, the plane \mathbb{C} or the unit disc \mathbb{D} , uniquely up to conformal automorphisms. In the case of a closed Riemann surface of genus one or greater, the theorem implies that, up to conformal automorphisms, its universal covering map can be seen as a unique conformal transformation from either the plane \mathbb{C} (in the case of genus one) or from the unit disc \mathbb{D} (in the case of genus two or higher) with a fundamental covering group represented by a subgroup of the conformal automorphism group of \mathbb{C} or \mathbb{D} respectively. Since \mathbb{C} has a natural Euclidean metric and \mathbb{D} is a conformal model of the hyperbolic plane \mathbb{H}^2 with a natural metric of constant Gaussian curvature -1 , the uniformization theorem can be reinterpreted as the fact that each closed Riemann surface has a natural homogeneous metric of constant Gaussian curvature.

Consequently, when discretizing the uniformization theorem, one could attempt to obtain the same properties for the discrete uniformizing map as the ones discussed above. Also, in general, the goal should be to start from a polyhedral metric (or some other general structure that determines a conformal structure) on the surface and obtain the discretely conformally equivalent smooth metric of constant curvature. For instance in [2], a circle packing approximation of the universal covering map of a Riemann surface is presented, but the question is more about inducing a metric with a circle packing on a topological triangulated surface with given combinatorics by laying it out on the universal cover. In [5, 14, 15] a flexible and constructive discrete uniformization is obtained based on the method of *discrete conformal factors* for triangulated surfaces [19, 29]. It is based on the discrete analog of conformally equivalent metrics. In the case of polyhedral surfaces with fixed combinatorics, there is a uniqueness theorem but there is no existence theorem. In other words, it is not always guaranteed that the method will produce a

discretely conformally equivalent image if the combinatorics is preserved. However, if one is allowed to perform edge-flips and thus alter the combinatorial nature of the triangular mesh, a solution is guaranteed [14, 15].

Another interesting approach to discrete conformal mappings is the method of *inverse distance circle packings*. It has been introduced in [9] and tested as a possible discretization of conformal mappings in [8]. Inverse distance packings are somewhat similar to hyper-ideal circle patterns, except that instead of intersection angles between adjacent face circles, inverse distances between vertex-circles are prescribed. In [18], Luo has proved a uniqueness theorem (i.e. a rigidity theorem) for inverse distance packings in the Euclidean and hyperbolic cases. There is no uniqueness in the spherical case since a counter-example has been constructed in [20]. No existence theorem for inversive distance packings is known.

Our discrete uniformization method describes a discrete conformal map in terms of hyper-ideal circle patterns (see Definition 5). It features a well-defined conformal equivalence relation and has the desired properties. In contrast to the methods of conformal factors and inverse distance packings, it possesses the existence and uniqueness theorems for fixed polygonal combinatorics. Furthermore, it is constructive and it has been algorithmically implemented. In the last section we have included some computer generated examples of discrete uniformization via hyper-ideal circle patterns. However, our approach is restricted to surfaces with polyhedral metrics of non-positive curvature, i.e. the cone singularities have angles greater than 2π , as well as to finite ramified covers over the sphere. All algebraic curves fall in the latter category.

2 Definitions and notations

We start with some terminology and notations. Assumed that S is an orientable compact topological surface with no boundary. Denote by d a metric of constant negative or zero Gaussian curvature on S with finitely many cone singularities $\text{sing}(d)$. Depending on whether the curvature away from singularities is negative or zero, the metric will be referred to as a *hyperbolic cone-metric* or a *Euclidean cone-metric* respectively. We will use \mathbb{F}^2 as a common notation for both the Euclidean plane \mathbb{E}^2 and the hyperbolic plane \mathbb{H}^2 . In this paper, we mostly use the Poincaré disc model and the upper half-plane model of \mathbb{H}^2 . Both of them are conformal and they have the advantage that circles in hyperbolic geometry appear in both models as circles in the underlying Euclidean geometry (for details, see [33, 34, 3]). The only particularity is that, in general, the hyperbolic centers of the circles do not necessarily coincide with their Euclidean ones. Throughout this article, we will also use the notation $\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ to denote the Riemann sphere, which could also be thought of as the projective complex line $\mathbb{C}\mathbb{P}^1$. The global conformal automorphisms of $\hat{\mathbb{C}} \cong \mathbb{C}\mathbb{P}^1$ are the Möbius transformations (i.e. the linear fractional transformations) which form the group $\mathbb{P}SL(2, \mathbb{C})$. From now on, V will be a finite set of points on S containing the cone singularities of d . Thus V splits into two disjoint subsets $V_1 = \text{sing}(d)$ and $V_0 = V \setminus V_1$. By $\mathcal{C} = (V, E, F)$ we will denote a topological cell complex on S , where V are its vertices, partitioned into two disjoint subsets V_0 and V_1 ,

E are the edges and F are the faces of \mathcal{C} (see for example Figure 2 below). All three sets are finite. Furthermore, without loss of generality, we will always assume that the cell complexes we work with, their dual complexes, and the various subdivisions of the former and the latter, are nicely embedded in the surface. Moreover, they will be assumed to be always *strongly regular* [37, 13, 32] which means that (i) closed cells (of any dimension) are attached without identifications on their boundaries and (ii) the intersection of any pair of closed cells is either a closed cell or empty.

Definition 1. A geodesic cell complex on (S, d) is a cell complex $\mathcal{C}_d = (V, E_d, F_d)$ whose edges, with endpoints removed, are open geodesic arcs embedded in $S \setminus V$.

In other words, we can think of a geodesic cell-complex \mathcal{C}_d on a geometric surface (S, d) as a two dimensional manifold, obtained by gluing together geodesic polygons along their edges. The edges that are being identified should have the same length and the identification should be an isometry. Observe the difference between a geodesic cell-complex \mathcal{C}_d and its topological (combinatorial) counterpart \mathcal{C} . While \mathcal{C}_d is made of geodesic polygons and thus provides the underlying surface S with a cone-metric d , the cell-complex \mathcal{C} is just a purely topological (and hence combinatorial) object. In many cases in this article \mathcal{C} will be obtained from \mathcal{C}_d by forgetting about the geometry of \mathcal{C}_d and focusing entirely on its combinatorics and topology.

Assume three disjoint circles c_i, c_j and c_k with centers i, j and k respectively, lie in the geometric plane \mathbb{F}^2 .

Then, there exists a unique fourth circle c_Δ orthogonal to c_i, c_j and c_k . Furthermore, let $\Delta = ijk$ be the geodesic triangle spanned by the centers i, j and k . Then Δ , together with the circles c_i, c_j, c_k and c_Δ , is called a *decorated triangle* (see Figure 1). The circles c_i, c_j and c_k are called the *vertex circles* of Δ , while c_Δ is called the *face circle* of Δ . We emphasize that it is allowed for one, two or all three vertex circles to have zero radii, i.e. to be points. Even in this more general set up, everything said above still applies. Whenever the vertex circles have zero radius, we will call them *degenerate vertex circles*.

Now, assume two non-overlapping decorated triangles, like $\Delta_1 = jis$ and $\Delta_2 = uis$ from Figure 1, share a common edge is . As usual, denote by c_i, c_j, c_s and c_u the vertex circles (some of which may be degenerate), and by c_{Δ_1} and c_{Δ_2} the corresponding face circles of the triangles. Although generically the two face circles c_{Δ_1} and c_{Δ_2} are different, sometimes it may happen that they coincide, i.e. $c_{\Delta_1} = c_{\Delta_2} = c_q$. In that case all four vertex circles c_i, c_j, c_s and c_u are orthogonal to c_q . Thus, we can erase the edge is and obtain a decorated geodesic quadrilateral $q = ijsu$ with vertex circles c_i, c_j, c_s and c_u , and a face circle c_q orthogonal to the vertex ones. Observe, that in this case the quadrilateral is convex. If we continue this way, we can obtain decorated polygons with arbitrary number of edges, like for instance the decorated pentagon $f' = ijsu$ from Figure 1.

Definition 2. A decorated polygon is a convex geodesic polygon p in \mathbb{F}^2 , with vertices i_1, i_2, \dots, i_n , together with:

- a set of disjoint circles $c_{i_1}, c_{i_2}, \dots, c_{i_n}$ such that each c_{i_s} is centered at the vertex i_s for $s = 1, \dots, n$. Some or all of the circles are allowed to be degenerate, i.e. circles of radius zero;

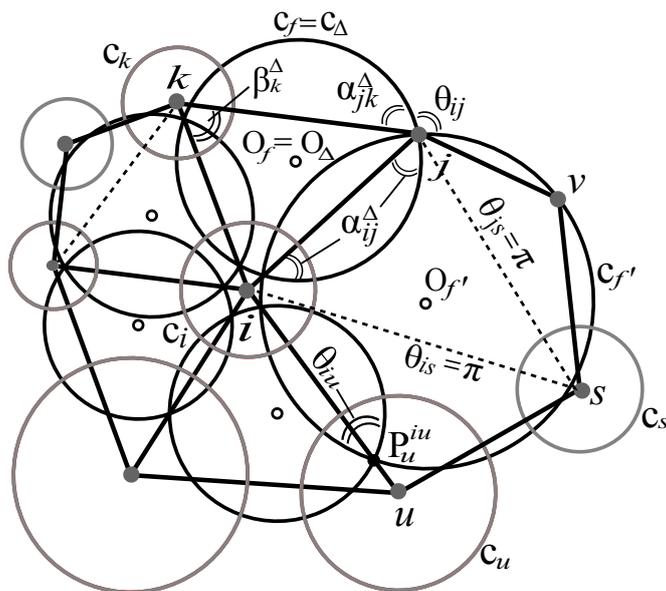


Figure 1: A hyper-ideal circle pattern of decorated polygons together with labels.

- another circle c_p orthogonal to c_{i_1}, \dots, c_{i_n} .

The circles c_{i_1}, \dots, c_{i_n} are called vertex circles and the additional orthogonal circle c_p is called the face circle of the decorated polygon p .

Two faces of a cell complex that share a common edge will be called *adjacent to each other*. Furthermore, assume two decorated polygons p_1 and p_2 share a common geodesic edge ij . Then, the decorated polygons p_1 and p_2 are called *compatibly adjacent* whenever the vertex circles c_i^1, c_j^1 of p_1 and c_i^2, c_j^2 of p_2 coincide respectively, that is $c_i^1 \equiv c_i^2$ and $c_j^1 \equiv c_j^2$. Furthermore, whenever two decorated polygons are compatibly adjacent, we will say that their face circles are *adjacent to each other*. A situation like that is depicted on Figure 1 for the edge ij and the two faces $f \equiv \Delta$ and f' with face circles c_f and $c_{f'}$.

Definition 3. Let p_1 and p_2 be two compatibly adjacent decorated polygons in \mathbb{F}^2 . Let ij be their common geodesic edge. Furthermore, let c_{p_1} and c_{p_2} be the face circles of p_1 and p_2 respectively.

- We say that the edge ij satisfies the local Delaunay property whenever each vertex circle of the decorated polygon p_2 is either (i) disjoint from the interior of the face circle c_{p_1} of p_1 , or (ii) if it is not, the intersection angle between the vertex circle in question and the face circle c_{p_1} , measured between the circular arcs that bound the region of common intersection of their discs, is less than $\pi/2$. See for instance edge ij on Figure 1.

- For the edge ij , which satisfies the local Delaunay property, $\theta_{ij} \in (0, \pi)$ denotes the intersection angle between the two adjacent face circles c_{p_1} and c_{p_2} , measured between

the circular arcs that bound the region of common intersection of their discs. (See for example angles θ_{ij} and θ_{iu} from Figure 1.)

It is not difficult to see that the definition of local Delaunay property is symmetric in the sense that if the condition of Definition 3 holds for the face circle c_{p_1} and the vertex circles of p_2 , then it also holds for the face circle c_{p_2} and the vertex circles of p_1 .

Definition 4. A hyper-ideal circle pattern on a given surface S (Figure 1) is a hyperbolic or Euclidean cone metric d on S together with a geodesic cell complex $\mathcal{C}_d = (V, E_d, F_d)$ whose faces are decorated geodesic polygons such that any two adjacent faces are compatibly adjacent and each geodesic edge of \mathcal{C}_d has the local Delaunay property. Whenever d is flat on $S \setminus V$, we call the circle pattern Euclidean, and whenever d is hyperbolic on $S \setminus V$, we call the pattern hyperbolic. Whenever all vertex radii are zero, the pattern will be called a Delaunay circle pattern (see left-hand side of Figure 2).

Intuitively speaking, a hyper-ideal circle pattern on a surface S is a surface homeomorphic to S , obtained by gluing together decorated geodesic polygons along pairs of corresponding edges. The edges that are being identified should have the same length, the identification should be an isometry and the vertices that get identified should have vertex-circles with same radii.

Observe that a hyper-ideal circle pattern on S consists of (i) a cone-metric d on S , (ii) a finite set of vertices $V \supseteq \text{sing}(d)$, (iii) an assignment of vertex radii r on V , and (iv) a geodesic cell complex \mathcal{C}_d together with (v) a collection of vertex circles and (vi) a collection of face circles. However, the geometric data (S, d, V, r) is enough to further identify uniquely the geodesic cell complex \mathcal{C}_d and the collections of vertex and face circles. This is done via the weighted Delaunay cell decomposition construction [12]. More precisely, given (i) a geometric surface (S, d) , (ii) a finite set of points $V \supset \text{sing}(d)$ on S and (iii) an assignment of disjoint vertex circle radii $r : V \rightarrow [0, \infty)$, one can uniquely generate (obtain) the corresponding r -weighted Delaunay cell complex \mathcal{C}_d , where each edge satisfies the local Delaunay property. In the process, the families of vertex and face circles naturally appear as part of the construction [12, 26, 31]. Thus, hyper-ideal circle patterns are the same as a certain type of weighted Delaunay cell decompositions. Generically, a weighted Delaunay cell complex is in fact a triangulation, also known as a weighted Delaunay triangulation [12].

Discussion on planar weighted Delaunay tessellations. As a brief illustration, we discuss the latter in the case when the geometric surface S, d is the plane \mathbb{F}^2 . So we are given a finite set of points V in \mathbb{F}^2 together with weights $r : V \rightarrow [0, \infty)$. This is equivalent to actually having a finite set of closed circular (vertex) disks in \mathbb{F}^2 with centers V and radii r . Assume they are disjoint. As already discussed in the paragraphs preceding Definition 2, for every triple of such disks there exists a circle that intersects the boundaries of the disks orthogonally. The weighted Delaunay triangulation, induced by the vertex disks, consists of those geodesic triangles whose vertices are the centers of triples of disks for which the orthogonal circle of this triple intersects no other disk more than orthogonally (see Figure 1 when $\mathbb{F}^2 = \mathbb{E}^2$ and the right-hand side of Figure

2 when $\mathbb{F}^2 = \mathbb{H}^2$). In some cases, an orthogonal circle can intersect orthogonally more than three circles, thus forming a polygon rather than a triangle, resulting in a more general cell complex than a triangulation, usually referred to as a *weighted Delaunay tessellation* and a *weighted Delaunay cell complex* (an example depicted on Figures 1 and 2). Consequently, the vertex circles and the orthogonal (face) circles form a hyper-ideal circle pattern on \mathbb{F}^2 with convex geodesic boundary. The boundary geodesic edges are technically also circles which in some cases, such as the case of $\mathbb{F}^2 = \mathbb{E}^2$ for instance, pass through the point of infinity and in the case of $\mathbb{F}^2 = \mathbb{H}^2$ are orthogonal to the ideal boundary of the hyperbolic plane, which can be thought of as a vertex circle centered at the point of infinity. Alternatively, one can obtain the r -weighted Delaunay cell complex as the geodesic dual to the r -weighted Voronoi diagram [12]. A Voronoi cell in the Euclidean case is defined as $W_r(i) = \{x \in \mathbb{E}^2 \mid d_{\mathbb{E}^2}(x, i)^2 - r_i^2 \leq d_{\mathbb{E}^2}(x, j)^2 - r_j^2 \text{ for all } j \in V\}$, and a Voronoi cell in the hyperbolic case is defined as $W_r(i) = \{x \in \mathbb{H}^2 \mid \cosh(r_j) \cosh d_{\mathbb{H}^2}(x, i) \leq \cosh(r_i) \cosh d_{\mathbb{H}^2}(x, j) \text{ for all } j \in V\}$. The weighted Voronoi complex and the weighted Delaunay complex are dual to each other and are geodesic and embedded. The vertices of the Voronoi complex are the centers of the Delaunay (face) circles of the pattern. For a rough illustration of a complex and its dual, see Figure 3. Both the weighted-Delaunay construction and the Voronoi construction can be performed on the surface S, d in an analogous manner.

Intuitively speaking, given a topological surfaces S with a finite number of points on it V , one can introduce a discrete conformal structure on (S, V) by assigning (i) either a hyperbolic or Euclidean cone metric d such that $\text{sing}(d) \subseteq V$ together with (ii) an appropriate vertex radii assignment $r : V \rightarrow [0, \infty)$. In short, (S, V, d, r) could be regarded as a surface with a discrete conformal structure, i.e. a discrete Riemann surface. Alternatively, instead of a cone-metric on S , one could also have a projective (i.e. conformal) $\mathbb{C}\mathbb{P}^1 \cong \hat{\mathbb{C}}$ structure with cone singularities, where the latter are assumed to be among the points from $V \subset S$. What we mean is that away from the points from $V_1 \subset V$, the surface has an atlas with transition functions given exclusively by Möbius transformations from the conformal group $\mathbb{P}SL(2, \mathbb{C})$. Since the elements of $\mathbb{P}SL(2, \mathbb{C})$ send circles to circles and preserve angles on $\hat{\mathbb{C}}$, both the notions of circles and angles (but not centers and radii of circles) are well defined on the surface itself. Consequently, hyper-ideal circle patterns on a surface with $\hat{\mathbb{C}}$ cone-structure make perfect sense and can be regarded as discrete conformal structures. Notice that the metric definitions above (Definition 4) are special cases of such cone $\hat{\mathbb{C}}$ structures because both the Euclidean and the hyperbolic isometry groups are subgroups of $\mathbb{P}SL(2, \mathbb{C})$.

Discussion on hyper-ideal circle patterns on $\hat{\mathbb{C}}$. In this article, somewhat implicitly, we will encounter Delaunay circle patterns (but not hyper-ideal ones) on surfaces with cone $\mathbb{C}\mathbb{P}^1$ structures (Section 4). The only explicit Delaunay and hyper-ideal circle patterns with $\hat{\mathbb{C}}$ structure we actually encounter in this paper, are the ones on $\hat{\mathbb{C}}$ itself. They can be acted upon by $\mathbb{P}SL(2, \mathbb{C})$, which preserves the combinatorics and the intersection angles, but does not preserve the notion of circle centers and radii. Hence, patterns on $\hat{\mathbb{C}}$ are defined up to $\mathbb{P}SL(2, \mathbb{C})$ Möbius transformations and for that reason the circles

from such patterns do not have naturally defined centers and the patterns themselves do not have canonically defined geometric cell complexes on $\hat{\mathbb{C}}$ associated to them, although they do have canonically defined underlying combinatorial complexes. However, Delaunay [6, 21, 33] and hyper-ideal circle patterns [1, 24, 25] naturally correspond to convex ideal and hyper-ideal hyperbolic polyhedra in \mathbb{H}^3 respectively (e.g. see as well as Section 10). These polyhedra have natural geodesic cell decompositions, determined by their vertices, faces and edges, and represent the combinatorial cell complexes arising from the circle patterns. Alternatively, one can directly use Definition 4 without resorting to hyperbolic polyhedra. For example, this can be done by performing a stereographic projection of the hyper-ideal circle pattern on $\hat{\mathbb{C}}$ to the plane and thus obtaining a hyper-ideal circle pattern with polygonal boundary either in the Euclidean or in the hyperbolic plane, just as described in the Discussion on planar weighted Delaunay tessellations two paragraphs above (see also Figures 1 and 2). Whenever one performs a stereographic projection from a vertex circle shrunk to a point, one obtains a Euclidean hyper-ideal circle pattern with convex polygonal boundary. Alternatively, one can also pick a point from the interior of a vertex circle (the interior disjoint from all other vertex circles) and then the stereographic projection with respect to that point leads to a hyper-ideal circle pattern with convex polygonal geodesic boundary in \mathbb{H}^2 , where $\partial\mathbb{H}^2$ is the image of the vertex circle whose interior contains the point. As a result, one obtains a geodesic representation of the cell complex defining the combinatorics of the pattern in the convex geodesic polygon, as depicted on Figures 1 and 2. To complete the cell decomposition to the sphere, one can add (i) in the Euclidean case straight rays from the vertices of the boundary going away to the point at infinity, for example following the angle bisectors at the vertices; (ii) in the hyperbolic case one can draw the straight (Euclidean) rays connecting the vertices of the hyperbolic boundary polygon to the point at infinity, each ray starting from a vertex and passing through its inverse image with respect to $\partial\mathbb{H}^2$ (when the polygon is in general position).

Remark: Just like in the Euclidean and the hyperbolic cases, one could define hyper-ideal circle patterns with underlying cone metrics of constant positive curvature (spherical cone metrics). For instance, one could fix such a metric on $\hat{\mathbb{C}}$ and use the metric Definition 4 to define hyper-ideal circle patterns with a spherical metric on $\hat{\mathbb{C}}$ and even on S (with cone singularities). Consequently, one can obtain a spherical geodesic Delaunay cell complex and its Voronoi dual representing the combinatorics of the pattern. However, it seems more natural to use the conformal structure on the Riemann sphere, rather than just one fixed metric. For that reason, we are going to skip the spherical case in this article.

Given a hyper-ideal circle pattern, one can extract from it the finite combinatorial data $(\mathcal{C}, \theta, \Theta)$, where (i) \mathcal{C} is the geodesic cell-complex \mathcal{C}_d on (S, d) , regarded as a purely combinatorial object (i.e. we forget all geometric information), (ii) $\theta : E \rightarrow (0, \pi)$ is the assignment of intersection angles of all pairs of adjacent face circles of the pattern and (iii) $\Theta : V \rightarrow (0, +\infty)$ are the cone angles around the points from V . In this case we will say that the hyper-ideal circle pattern *realizes the (combinatorial) angle data*

$(\mathcal{C}, \theta, \Theta)$. As pointed out in the definition of the set V , the cone angle $\Theta_k = 2\pi$ whenever $k \in V_0$ and $\Theta_k \neq 2\pi$ when $k \in V_1 = \text{sing}(d)$.

Definition 5. *Two hyper-ideal circle patterns on S are considered discretely conformally equivalent whenever their underlying (geodesic) cell-complexes are combinatorially isomorphic and the corresponding intersection angles between pairs of adjacent face circles are equal. Consequently, a discrete conformal map between two conformally equivalent hyper-ideal circle patterns is the pairing of face and vertex circles from the first pattern with the corresponding face and vertex circles from the second pattern.*

3 Discrete uniformization of polyhedral surfaces with non-positive curvature.

Consider the following input data on the surface S :

- (a) either Euclidean or hyperbolic cone-metric d on S ;
- (b) a finite set of point $V \subset S$ such that $V = V_0 \cup V_1$, where $V_1 = \text{sing}(d)$ and $V_0 \cap V_1 = \emptyset$;
- (c) the cone angle $\Theta_k > 2\pi$ when $k \in V_1$ and $\Theta_k = 2\pi$ when $k \in V_0$.

To put it shortly, the surface S is provided with a cone metric of non-positive Gaussian curvature. We will use the notation (S, d, V) to denote the aforementioned input data. One could interpret it as the geometric data $(S, d, V, 0)$ that gives rise to the classical Delaunay circle pattern, where all vertex radii are equal to zero.

Next, from the geometry of (S, d, V) , form the combinatorial angle data $(\mathcal{C}, \theta, 2\pi)$, where

1. $\mathcal{C} = (V, E, F)$ is the unique Delaunay cell complex of (S, d) with respect to V , regarded as a purely topological (combinatorial) complex.
2. $\theta : E \rightarrow (0, \pi)$ are the angles between the pairs of adjacent Delaunay circles of the Delaunay circle pattern.
3. $\Theta_k = 2\pi$ for all $k \in V$. We use the notation 2π to denote this constant angle assignment.

Theorem 1. *There exists a hyper-ideal circle pattern on S , with underlying smooth hyperbolic metric h , that realizes the combinatorial angle data $(\mathcal{C}, \theta, 2\pi)$ (see Figure 2). The pattern is unique up to isometry isotopic to identity (i.e. isometry preserving the combinatorics of the cell complex and its labelling). In other words, the Delaunay circle pattern corresponding to (S, d, V) is discretely conformally equivalent to a unique, up to label-preserving isometry, hyper-ideal circle pattern with a smooth hyperbolic metric h on S .*

Observe that Theorem 1 implies the existence of a unique discrete conformal map (see Definition 5 and Figure 2) between the classical Delaunay circle pattern on (S, d, V) and a uniquely defined (up to isometry) hyper-circle pattern with a smooth hyperbolic metric on S . Consequently, one can develop the geodesic cell-complex corresponding to

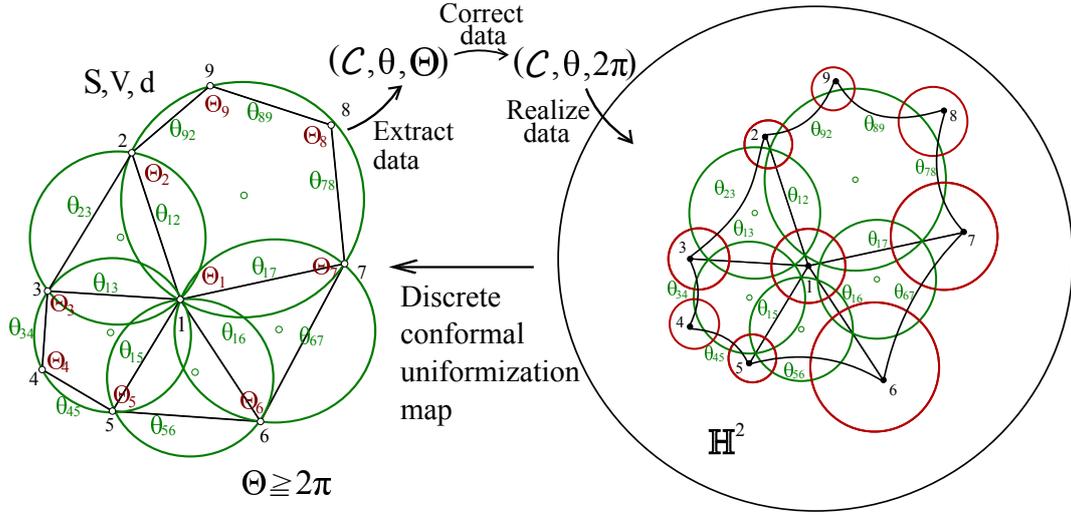


Figure 2: Discrete uniformization according to Theorem 1

this newly obtained hyper-ideal circle pattern in the hyperbolic plane and thus obtain a fundamental domain whose isometric gluing defines a Fuchsian group Γ such that \mathbb{H}^2/Γ is isometric to S, h . Moreover, one ends up with a discrete conformal universal covering map from the Γ -invariant hyper-ideal circle pattern on \mathbb{H}^2 to the Delaunay circle pattern on S, d (Figure 2). This is reminiscent of the classical uniformization theorem for Riemann surfaces and therefore we consider Theorem 1 as a discrete version of it.

4 Discrete uniformization of branch covers over the Riemann sphere

In relation to surfaces with discrete non-positive curvature, we propose a circle pattern uniformization of Riemann surfaces represented as finite branch covers over the Riemann sphere.

Let $p : S \rightarrow \hat{\mathbb{C}}$ be a finite topological branch cover over the Riemann sphere $\hat{\mathbb{C}}$. For example, one can think of S as a smooth compact algebraic curve in $\mathbb{C}\mathbb{P}^2$ whose affine part is given by a complex polynomial equation $P(x, y) = 0$. The branch covering map p can be defined as $p : (x, y) \in S \mapsto x$ for example, although any meromorphic function on S would do. Define by V_1 all ramification points of p on S and let $V_1(\hat{\mathbb{C}}) = p(V_1)$ be the branch points of p on $\hat{\mathbb{C}}$. Observe that $p^{-1}(V_1(\hat{\mathbb{C}}))$ contains V_1 but does not necessarily coincide with it. Denote by $N \in \mathbb{N}$ the number of sheets of p .

As input data we consider:

- a finite topological branch cover $p : S \rightarrow \hat{\mathbb{C}}$ with ramification points $V_1 \subset S$ and branch points $V_1(\hat{\mathbb{C}}) \subset \hat{\mathbb{C}}$;
- a finite set of points $V_{\hat{\mathbb{C}}} = V_0(\hat{\mathbb{C}}) \cup V_1(\hat{\mathbb{C}})$ on $\hat{\mathbb{C}}$ where $V_0(\hat{\mathbb{C}}) \cap V_1(\hat{\mathbb{C}}) = \emptyset$.
- a finite set $V = p^{-1}(V_{\hat{\mathbb{C}}})$ on S with $V_0 = V \setminus V_1$.

This input data can be denoted by $p : (S, V) \rightarrow (\hat{\mathbb{C}}, V_{\hat{\mathbb{C}}})$. Observe that we do not need to specify any metric on $\hat{\mathbb{C}}$ because of its natural conformal structure in which, as already mentioned in the Discussion on $\hat{\mathbb{C}}$ hyper-ideal patterns and the paragraph preceding it in Section 2, the notions of circles and angles are invariant under the action of the conformal automorphism group $\mathbb{PSL}(2, \mathbb{C})$, so Delaunay circle patterns naturally exist on $\hat{\mathbb{C}}$. Consequently, let us generate the unique Delaunay circle pattern on $\hat{\mathbb{C}}$ with respect to the points $V_{\hat{\mathbb{C}}}$. Recall, that in the Discussion of planar weighted Delaunay tessellations of Section 2 we have explained how, after stereographic projection, one can interpret hyper-ideal circle patterns (and Delaunay ones in particular) from either Euclidean or hyperbolic point of view, as specified in Definition 4, and also obtain the corresponding geodesic weighted Delaunay cell complexes and their dual geodesic Voronoi cell complexes. Consequently, we have an embedding of the topological complex $\mathcal{C}_{\hat{\mathbb{C}}} = (V_{\hat{\mathbb{C}}}, E_{\hat{\mathbb{C}}}, F_{\hat{\mathbb{C}}})$ and its dual $\mathcal{C}_{\hat{\mathbb{C}}}^* = (V_{\hat{\mathbb{C}}}^*, E_{\hat{\mathbb{C}}}^*, F_{\hat{\mathbb{C}}}^*)$ in $\hat{\mathbb{C}}$. By definition, a pair of Delaunay circles (a special case of face circles) are adjacent to each other exactly when there is a corresponding edge $ij \in E_{\hat{\mathbb{C}}}$ of $\mathcal{C}_{\hat{\mathbb{C}}}$, and equivalently a dual edge $ij^* \in E_{\hat{\mathbb{C}}}^*$ of $\mathcal{C}_{\hat{\mathbb{C}}}^*$. Thus, the intersection angle between them $\hat{\theta}_{ij} \in (0, \pi)$ gives rise to an angle assignment $\hat{\theta} : E_{\hat{\mathbb{C}}} \rightarrow (0, \pi)$ (and by duality $\hat{\theta} : E_{\hat{\mathbb{C}}}^* \rightarrow (0, \pi)$). Since the branch points of $p : S \rightarrow \hat{\mathbb{C}}$ are among the vertices $V_{\hat{\mathbb{C}}}$, we can lift $\mathcal{C}_{\hat{\mathbb{C}}}$ and its dual $\mathcal{C}_{\hat{\mathbb{C}}}^*$ to the embedded complexes $\mathcal{C} = (V, E, F)$ and $\mathcal{C}^* = (V^*, E^*, F^*)$, which are also dual to each other, i.e. $\mathcal{C} = p^{-1}(\mathcal{C}_{\hat{\mathbb{C}}})$ and $\mathcal{C}^* = p^{-1}(\mathcal{C}_{\hat{\mathbb{C}}}^*)$. The ramification points of the covering are by construction among the vertices of \mathcal{C} and thus lie in the interiors of their corresponding dual faces from \mathcal{C}^* . Furthermore, we can lift the angle assignment $\hat{\theta}$ to the covering angle assignment $\theta : E \rightarrow (0, \pi)$ by $\theta_{ij} = \hat{\theta}_{p(ij)}$ for all $ij \in E$. By duality, we also have $\theta : E^* \rightarrow (0, \pi)$.

With the preceding constructions in mind, one can extract from $(\hat{\mathbb{C}}, V_{\hat{\mathbb{C}}})$ the combinatorial angle data $(\mathcal{C}, \theta, \Theta)$, where

1. $\mathcal{C} = (V, E, F)$ is the topological cell-complex that branch covers the Delaunay cell complex $\mathcal{C}_{\hat{\mathbb{C}}}$ via the covering map p , as described above.
2. $\theta : E \rightarrow (0, \pi)$ are the lifts of the intersection angles between the pairs of adjacent Delaunay circles, also defined above.
3. $\Theta_k = 2\pi$ for all $k \in V_0$ and $\Theta_k = 2\pi N_k$ for all $k \in V_1$, where $N_k = 2, 3, \dots, N$ is the index of the ramification point k , i.e. the number of sheets meeting at the ramification point.

To discretely uniformize, correct the data (See Figure 2) to $(\mathcal{C}, \theta, 2\pi)$.

Theorem 2. *There exists a hyper-ideal circle pattern with an underlying either (i) hyperbolic cone metric h on S , when $\text{genus}(S) \geq 2$, or (ii) a Euclidean cone metric h on S , when $\text{genus}(S) = 1$, which realizes the combinatorial angle data $(\mathcal{C}, \theta, 2\pi)$. The pattern is unique up to isometry isotopic to identity (i.e. isometry preserving the combinatorics of the cell complex and its labelling) as well as scaling in the Euclidean case. In other words, the lift of the standard Delaunay circle pattern of $(\hat{\mathbb{C}}, V_{\hat{\mathbb{C}}})$ via p on S is discretely conformally equivalent to a unique (up to isometry isotopic to identity and scaling when Euclidean) hyper-ideal circle pattern with either a hyperbolic ($\text{genus} \geq 2$) or a Euclidean*

(genus = 1) metric h on S .

We state separately the case of S being homeomorphic to a sphere.

Theorem 3. *There exists a hyper-ideal circle pattern on $\hat{\mathbb{C}}$, unique up to label preserving $\mathbb{P}SL(2, \mathbb{C})$ transformation, which realizes the combinatorial angle data $(\mathcal{C}, \theta, 2\pi)$. In other words, the lift of the standard Delaunay circle pattern of $(\hat{\mathbb{C}}, V_{\hat{\mathbb{C}}})$ via p on the topological sphere S is discretely conformally equivalent to a unique, up to label respecting Möbius transformation, hyper-ideal circle pattern on the Riemann sphere $\hat{\mathbb{C}}$.*

Theorems 2 and 3 combined lead to a discrete analogue of the uniformization theorem for closed Riemann surfaces represented as branch covers over the Riemann sphere. In particular, these two theorems can provide a discrete uniformization method for smooth complex algebraic curves.

Just like in the case of the classical uniformization theorem, when the genus of S is two or greater, one can develop the final hyper-ideal circle patterns from Theorems 1 and 2 in the hyperbolic plane (e.g. the upper-half plane model or the Poincaré disk model) and obtain a Fuchsian group Γ , unique up to conjugation by a hyperbolic isometry, as well as a Γ -invariant hyper-ideal circle pattern on \mathbb{H}^2 whose factor \mathbb{H}^2/Γ is isometric to (S, h) together with the hyper-ideal circle pattern described in Theorems 1 or 2. When the genus of S is one, the same is true but for the Euclidean plane together with a lattice group of translations on it.

5 Existence and uniqueness of hyper-ideal patterns

In this section we introduce the main tool used in the proof of all theorems stated in the previous section.

As already discussed in the paragraph preceding Definition 5 from Section 2, given any hyper-ideal circle pattern one can always extract from it the combinatorial data $(\mathcal{C}, \theta, \Theta)$, where \mathcal{C} is a cell-complex representing the combinatorics of the pattern, $\Theta : V \rightarrow (0, \infty)$ is the assignment of cone-angles at the vertices of the complex and $\theta : E_d \rightarrow (0, \pi)$ is the assignment of intersection angles between adjacent face circles of the pattern. The proofs of the discrete uniformization theorems heavily rely on the solution to the following problem: *Find a hyperbolic or flat cone-metric, together with a hyper-ideal circle pattern on the given surface that realizes the data $(\mathcal{C}, \theta, \Theta)$.*

The solution to this problem was first presented by Schlenker in [26]. However, we will utilize the approach and the notations used in [11]. We start with the constructions and the definitions necessary for the formulation of the main result in this section.

Assume a cell complex $\mathcal{C} = (V, E, F)$ is fixed on the surface S (Figure 3a). Denote by $\mathcal{C}^* = (V^*, E^*, F^*)$ the cell complex dual to \mathcal{C} , where V^* are the dual vertices, E^* are the dual edges and F^* are the dual faces (see Figure 3b). On Figure 3b the elements of the original complex \mathcal{C} are drawn in grey, while the elements of the dual complex \mathcal{C}^* are in black.

Next, define the subdivision $\hat{\mathcal{T}} = (\hat{V}, \hat{E}, \hat{F})$ of \mathcal{C}^* , shown on Figure 4, where

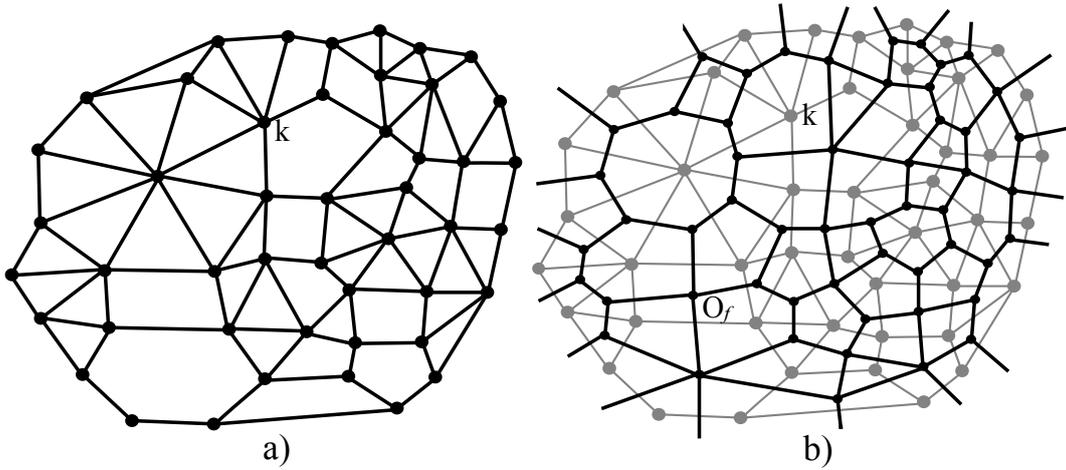


Figure 3: a) The cell complex $\mathcal{C} = (V, E, F)$ and b) its dual $\mathcal{C}^* = (V^*, E^*, F^*)$

- $\hat{V} = V \cup V^*$, i.e. the vertices of $\hat{\mathcal{T}}$ consist of all vertices of \mathcal{C} and all dual vertices. These are all black and grey vertices from Figure 4;
- $\hat{E} = E^* \cup \{iO_f \mid O_f \in V^* \text{ and } i \text{ is a vertex of } f\}$, i.e. the edges of $\hat{\mathcal{T}}$ consist of all dual edges and all edges, obtained by connecting a dual vertex $O_f \in f$ to all the vertices of the face $f \in F$ it belongs to. The latter type of edges will be called *corner edges*. The dual edges can be seen on both Figures 3b and 4 painted solid black, while the corner edges are the black dashed edges from Figure 4.
- $\hat{F} = \{iO_fO_{f'} \mid ij \in E \text{ common edge for } f \text{ and } f' \text{ from } F\}$, i.e. the faces of $\hat{\mathcal{T}}$ are the topological triangles obtained by looking at the connected components of the complement of the topological graph (\hat{V}, \hat{E}) on S . On Figure 4 these are the triangles with one solid black and two dashed black edges. They also have two black (dual) vertices and one grey vertex.

The next important notion to be defined is the *open star* of a vertex from $\hat{\mathcal{T}}$.

Definition 6. Let $\hat{v} \in \hat{V}$ be an arbitrary vertex of $\hat{\mathcal{T}}$. Then its open star $\text{OStar}(\hat{v})$ is defined as the open interior of the union of all closed triangles from $\hat{\mathcal{T}}$ which contain \hat{v} .

Whenever $\hat{v} = k \in V$ is a vertex of \mathcal{C} , then its open star is simply the open interior of the face from \mathcal{C}^* dual to k . An example denoted by $\text{OStar}(k)$ and colored in grey is shown on Figure 4. The boundary of $\text{OStar}(k)$ consists entirely of dual edges from E^* . If we denote by E_k the set of all edges of \mathcal{C} which have vertex k as an endpoint, then $\partial \text{OStar}(k) = \cup \{ik^* \in E^* \mid ik \in E_k\}$. If $\hat{v} = O_f \in V^*$ is a vertex from the dual complex \mathcal{C}^* , then the boundary of its open star consists entirely of corner edges from $\hat{\mathcal{T}}$ (e.g. the grey region $\text{OStar}(O_f)$ on Figure 4).

Recall that the vertex set V of the cell-complex \mathcal{C} is always partitioned into two subsets V_0 and V_1 .

Following the terminology of [26] (see also [11]), one can define what Schlenker calls an admissible domain.

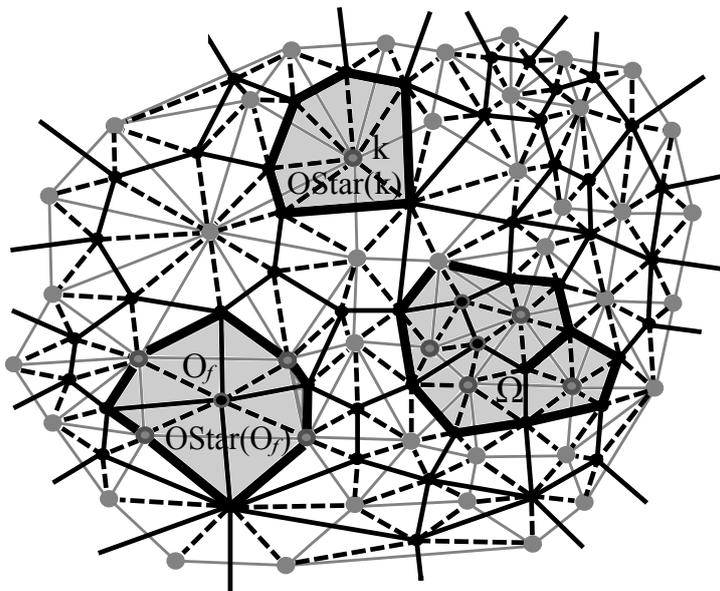


Figure 4: The triangulation $\hat{\mathcal{T}} = (\hat{V}, \hat{E}, \hat{F})$ together with two examples of open stars and one admissible domain Ω .

Definition 7. An open connected subdomain Ω of the surfaces S is called an admissible domain of (S, \mathcal{C}) whenever the following conditions hold:

1. There exists a subset $\hat{V}_0 \subseteq \hat{V}$, such that $\Omega = \cup\{\text{OStar}(\hat{v}) \mid \hat{v} \in \hat{V}_0\}$;
2. $\Omega \neq \emptyset$ and $\Omega \neq S$ and $\Omega \cap V \neq \emptyset$.
3. Ω is not punctured, i.e. if $\text{OStar}(\hat{v}) \setminus \{\hat{v}\} \subseteq \Omega$ then $\hat{v} \in \Omega$ for any $\hat{v} \in \hat{V}$. Consequently, the boundary of Ω is a nonempty set of edges of $\hat{\mathcal{T}}$, i.e. dual and/or corner edges.

A special example of an admissible domain is the open star of a vertex of \mathcal{C} . The open star of a dual vertex however is not an admissible domain because it is disjoint from V . An example of an admissible domain can be seen on Figure 4, denoted by the symbol Ω and shaded in grey. On this picture Ω is simply connected but in general it doesn't have to be.

The boundary of an admissible domain Ω is a disjoint union of topological curves on S , consisting entirely of edges together with their vertices, belonging to the triangulation $\hat{\mathcal{T}}$. In other words, the boundary of Ω consists of dual edges and/or corner edges from \hat{E} (and their vertices), but all of its connected components are interpreted as closed paths in the one-skeleton of $\hat{\mathcal{T}}$, so that some of the edges could be traced (counted) twice (see Figure 4) as well as some of the vertices could be counted twice or more times. Edges are traced twice exactly when an edge of $\hat{\mathcal{T}}$ is disjoint from Ω , but the interiors of the two topological triangles from $\hat{\mathcal{T}}$, lying on both sides of the edge, are contained in Ω . We denote this version of the boundary of Ω by $\partial\Omega$.

Theorem 4. (Schlenker [26, 11]) Let S be a closed surface with a topological cell complex $\mathcal{C} = (V, E, F)$ on it. Assume also that $V = V_0 \sqcup V_1$. Then the combinatorial angle data

$(\mathcal{C}, \theta, \Theta)$ is realized by a hyperbolic or Euclidean hyper-ideal circle pattern on S if and only if the assignment of angles $(\theta, \Theta) \in \mathbb{R}^E \times \mathbb{R}^V$ satisfies the following conditions:

- 1) $\theta_{ij} \in (0, \pi)$ for any $ij \in E$;
- 2) $\Theta_k > 0$ for all $k \in V_1$ and $\Theta_k = \sum_{ik \in E_k} (\pi - \theta_{ik})$ for all $k \in V_0$. The latter can be also written as $\Theta_k = \sum_{ik^* \subset \partial\Omega} (\pi - \theta_{ik})$ for $\Omega = \text{OStar}(k)$;
- 3) $\sum_{k \in V} (2\pi - \Theta_k) > 2\pi\chi(S)$ in the hyperbolic case and $\sum_{k \in V} (2\pi - \Theta_k) = 2\pi\chi(S)$ in the Euclidean case;
- 4) For any admissible domain Ω of (S, \mathcal{C}) , such that $\Omega \neq \text{OStar}(k)$ for some $k \in V_0$,

$$\sum_{ij^* \subset \partial\Omega} (\pi - \theta_{ij}) + \sum_{k \in \Omega \cap V} (2\pi - \Theta_k) + \pi|\partial\Omega \cap V| > 2\pi\chi(\Omega). \quad (1)$$

The notations $\chi(S)$ and $\chi(\Omega)$ are the Euler characteristics of S and Ω respectively.

Furthermore, whenever the hyper-ideal circle pattern in question exists, it is unique up to isometry isotopic to identity, as well as scaling in the Euclidean case. Finally, the hyper-ideal circle pattern can be reconstructed from the unique critical point of a strictly convex functional defined on an open convex set, bounded only by coordinate hyperplanes of \mathbb{R}^N for some suitable $N \in \mathbb{N}$.

Conditions 1 to 4 from Theorem 4 describe a convex polytope, which in the hyperbolic case we denote by $\mathcal{P}_{S, \mathcal{C}}^h$ and in the Euclidean case by $\mathcal{P}_{S, \mathcal{C}}^e$. We will also use the common notation $\mathcal{P}_{S, \mathcal{C}}$ whenever we do not need to specify the underlying geometry. We call these polytopes *angle data polytopes*.

To optimize the conditions of Theorem 4 a bit more, one can define the so called *strict admissible domain*.

Definition 8. An open connected subdomain Ω on the surfaces S is called a *strict admissible domain* of (S, \mathcal{C}) whenever Ω is admissible and $\partial\Omega \cap V_0 = \emptyset$.

As shown in [11], the angle data polytopes can be described via strict admissible domains instead of admissible domains. Simply, the admissible domains which are not strict do not add more restrictions to the angle data.

Corollary 1. The statements of Theorem 4 still hold even when the expression “admissible domain” in point 4) of Theorem 4 is replaced by the expression “strict admissible domain”.

6 Proof of Theorem 1

Let $(\mathcal{C}, \theta, \Theta)$ be the combinatorial angle data extracted from the Delaunay circle pattern on S, d with respect to the finite set of points V (refer to the explanations towards the end of Section 2. See also Figure 2). This circle pattern is a special case of a hyper-ideal circle pattern with all vertex circles of radius zero. Hence, Theorem 4 applies to it. Then the angles (θ, Θ) satisfy the conditions

- (i) $\theta_{ij} \in (0, \pi)$ for any $ij \in E$;
- (ii) $\Theta_k = \sum_{ik \in E_k} (\pi - \theta_{ik}) \geq 2\pi$ for any $k \in V$. Recall that V has been split into V_0 containing all k such that $\Theta_k = 2\pi$, and V_1 consisting of all k for which $\Theta_k > 2\pi$. By assumption, $V = V_0 \cup V_1$.
- (iii) $\sum_{k \in V} (2\pi - \Theta_k) \geq 2\pi\chi(S)$, where in the case of Euclidean cone-metric d we have an equality, and in the case of a hyperbolic cone-metric we have a strict inequality.
- (iv) For any admissible domain Ω of (S, \mathcal{C}) , such that $\Omega \neq \text{OStar}(k)$ for some $k \in V_0$,

$$\sum_{ij^* \subset \partial\Omega} (\pi - \theta_{ij}) + \sum_{k \in \Omega \cap V} (2\pi - \Theta_k) + \pi|\partial\Omega \cap V| > 2\pi\chi(\Omega). \quad (2)$$

To prove Theorem 1 all we have to do is check whether the angle data $(\theta, 2\pi)$ belongs to the polytope $\mathcal{P}_{S, \mathcal{C}}^h$, whose description is given by points 1, 2, 3 and 4 from Theorem 4. If we manage to confirm that, Theorem 4 guarantees the existence and the uniqueness of the sought hyper-ideal circle pattern with a smooth underlying hyperbolic metric h on S . So let us check conditions 1 to 4 of Theorem 4. We point out that if the metric d is already smooth, then we have nothing to prove, as d is actually the metric h and the hyper-ideal circle pattern in question is in fact the original Delaunay circle pattern. Therefore, for the rest of this proof we assume that d has at least one cone-singularity with cone angle greater than 2π , which means that $V_1 \neq \emptyset$.

Since the data $(\mathcal{C}, \theta, \Theta)$ comes from a Delaunay circle pattern, it is clear that condition 1 is immediately true. Condition 2 is also true because if $k \in V_0$, then $\Theta_k = 2\pi = \sum_{ik \in E_k} (\pi - \theta_{ik})$ follows from (ii) above. The case when $k \in V_1$ implies that $\Theta_k > 2\pi$ and therefore the already established in (ii) above strict inequality $\sum_{ik \in E_k} (\pi - \theta_{ik}) > 2\pi$ can be rewritten as $\sum_{ik^* \in \partial\Omega} (\pi - \theta_{ik}) + (2\pi - 2\pi) > 2\pi = 2\pi\chi(\Omega)$ which is a special case of condition 4 with $\Omega = \text{OStar}(k)$.

Since $V_1 \neq \emptyset$ and since we have assumed in Theorem 1 that $\Theta_k \geq 2\pi$ holds for each $k \in V$, then there is at least one $k \in V$ for which $\Theta_k > 2\pi$. Therefore $0 = \sum_{k \in V} (2\pi - 2\pi) > \sum_{k \in V} (2\pi - \Theta_k) \geq 2\pi\chi(S)$ by (iii) above. Thus, condition 3 is confirmed.

Finally, in order to verify point 4) from Theorem 4, let us fix an arbitrary admissible domain Ω (see Definition 7). Using again the assumption that $\Theta_k \geq 2\pi$, for all $k \in V$, and combining it with observation (iv) above, we can deduce that

$$\sum_{ij^* \subset \partial\Omega} (\pi - \theta_{ij}) + \sum_{k \in \Omega \cap V} (2\pi - 2\pi) \geq \sum_{ij^* \subset \partial\Omega} (\pi - \theta_{ij}) + \sum_{k \in \Omega \cap V} (2\pi - \Theta_k) > 2\pi\chi(\Omega) - \pi|\partial\Omega \cap V|.$$

The proof of Theorem 1 is complete.

7 Realization theorems for circle patterns on the sphere

One of the central tools in the proof of Theorems 2 and 3, alongside Theorem 4, are the Theorem of Rivin [22] and the Theorem of Bao and Bonahon [1]. Originally, both results

are stated as a geometric characterization of (i) convex ideal polyhedra (in the case of Rivin) and (ii) convex hyper-ideal polyhedra (in the case of Bao and Bonahon) in the hyperbolic three-space. However, due to the natural one-to-one correspondence between such polyhedra and (i) Delaunay circle patterns (Rivin [22]) and (ii) hyper-ideal circle patterns (Bao and Bonahon [1]) on the Riemann sphere (see also [33, 6, 25] as well as Section 10), we restate them as theorems about circle patterns.

We begin with some notations and definition. Let $\mathcal{C} = (V, E, F)$ be a topological cell complex and let $\mathcal{C}^* = (V^*, E^*, F^*)$ be its dual on a surface S (Figure 3). Denote by $\mathcal{C}^*(1)$ the one-skeleton of \mathcal{C}^* , which is the subcomplex of \mathcal{C}^* composed only of its vertices and edges (so faces excluded). In other words, $\mathcal{C}^*(1) = (V^*, E^*)$ and so it is a graph embedded in S . A *path in the one-skeleton* of \mathcal{C}^* is a sequence of dual edges from E^* whose union is a topological path on S . Analogously, a *loop in the one-skeleton* of \mathcal{C}^* is a sequence of dual edges from E^* whose union is a topological loop on S . A path or a loop in $\mathcal{C}^*(1)$ is called *simple* if it is respectively a simple topological path or a simple loop on S . This means that no two edges in the sequence repeat and each dual vertex that the simple path or loop goes through is adjacent to no more than two edges from the path (exactly two in the case of a simple loop). The same terminology applies to the triangulation $\hat{\mathcal{T}}$ (Figure 4).

In the case when S is a topological sphere we have the following two powerful theorems.

Theorem 5. (Rivin [22]) *Let S be a topological sphere with a strongly regular cell complex $\mathcal{C} = (V, E, F)$ on it. Let \mathcal{C}^* be the dual complex of \mathcal{C} . Then the combinatorial data $(\mathcal{C}, \theta, 2\pi)$ is realized by a Delaunay circle pattern on $\hat{\mathbb{C}}$ if and only if the following conditions are satisfied by the angle assignment $\theta : E \rightarrow \mathbb{R}$:*

- 1) $\theta_{ij} \in (0, \pi)$ for any $ij \in E$;
- 2) For each simple loop δ in the one-skeleton of \mathcal{C}^*

$$\sum_{ij^* \subset \delta} (\pi - \theta_{ij}) \geq 2\pi$$

where equality holds if and only if δ is the boundary of a face from \mathcal{C}^* .

Furthermore, whenever the hyper-ideal circle pattern in question exists, it is unique up to a $\mathbb{P}SL(2, \mathbb{C})$ transformation respecting the labelling of the complex.

Rivin's Theorem is a special case of the more general Bao and Bonahon's Theorem.

Theorem 6. (Bao and Bonahon [1]) *Let S be a topological sphere with a strongly regular cell complex $\mathcal{C} = (V, E, F)$ on it. Assume the vertices are partitioned into $V = V_0 \sqcup V_1$. Furthermore, let \mathcal{C}^* be the dual complex of \mathcal{C} . Then the combinatorial data $(\mathcal{C}, \theta, 2\pi)$ is realized by a hyper-ideal circle pattern on $\hat{\mathbb{C}}$ with true vertex circles corresponding to the vertices from V_1 and circles shrunk to points corresponding to the vertices from V_0 if and only if the following conditions are satisfied by the angle assignment $\theta : E \rightarrow \mathbb{R}$:*

- 1) $\theta_{ij} \in (0, \pi)$ for any $ij \in E$;

2) For each simple loop δ in the one-skeleton of \mathcal{C}^*

$$\sum_{ij^* \subset \delta} (\pi - \theta_{ij}) \geq 2\pi$$

where equality holds if and only if δ is the boundary of a face from \mathcal{C}^* dual to a vertex from V_0 .

3) For each simple path δ in the one-skeleton of \mathcal{C}^* joining two distinct dual vertices of the same dual face so that δ is not contained in the boundary of any dual face from \mathcal{C}^*

$$\sum_{ij^* \subset \delta} (\pi - \theta_{ij}) > \pi.$$

Furthermore, whenever the hyper-ideal circle pattern in question exists, it is unique up to a $\mathbb{P}SL(2, \mathbb{C})$ transformation respecting the labelling of the complex.

8 Proof of Theorem 2

Let $\hat{\mathcal{T}}_{\hat{\mathcal{C}}} = (\hat{V}_{\hat{\mathcal{C}}}, \hat{E}_{\hat{\mathcal{C}}}, \hat{F}_{\hat{\mathcal{C}}})$ be the triangular subdivision of the dual complex \mathcal{C}^* as defined in Section 5 and shown on Figure 4. Lift $\hat{\mathcal{T}}_{\hat{\mathcal{C}}}$ to the triangulation $\hat{\mathcal{T}}$ on the surface S via the branch covering map p . By construction, $\hat{\mathcal{T}}$ is the subtriangulation of the lifted dual complex \mathcal{C}^* on S described in Section 5 and depicted on Figure 4. Thus, $p(\hat{\mathcal{T}}) = \hat{\mathcal{T}}_{\hat{\mathcal{C}}}$ and $p(\mathcal{C}^*) = \mathcal{C}_{\hat{\mathcal{C}}}^*$. Just like before, in order to prove Theorem 2 we simply have to check whether the angle data $(\theta, 2\pi)$ belongs to the polytope $\mathcal{P}_{S, \mathcal{C}}$, whose description is given by points 1, 2, 3 and 4 of Theorem 4. However, in this case we are going to apply Theorem 4 in the setting of Corollary 1, which means we are going to use the description of the polytope $\mathcal{P}_{S, \mathcal{C}}$ in terms of strict admissible domains instead of admissible ones.

8.1 Verification of conditions 1, 2 and 3 of Theorem 4

Verification of conditions 1 and 3. Since the angles $\theta = p^* \hat{\theta} \in \mathbb{R}^E$ are lifts of the angles of a Delaunay circle pattern on $\hat{\mathcal{C}}$, condition 1 of Theorem 4 is automatically satisfied. Denoting by $g(S)$ the genus of the surface S , if we assume that $g(S) \geq 1$ then condition 3 is also satisfied because we have taken $\Theta_k = 2\pi$ for all $k \in V$ and thus

$$0 = \sum_{k \in V} (2\pi - \Theta_k) \geq 2\pi \chi(S) = 4\pi(1 - g(S)).$$

The inequality is strict when S has genus at least two and becomes an equality when S has genus 1, i.e. the case of the torus.

Verification of condition 2. From now on, whenever k is a vertex of \mathcal{C} or $\mathcal{C}_{\hat{\mathcal{C}}}$, we denote by $\text{Star}(k)$ the closure of its open star $\text{OStar}(k)$ (for the definition of $\text{OStar}(k)$ see Section 5 and Figure 4). This simply means that we add to the open star all dual

edges lying on its boundary. By definition, it is equivalent to saying that the open star of k is the interior of the face dual to k and so its closure is the closure of that dual face. Since we have assumed that the complexes are strongly regular, $\text{Star}(k)$ is an embedded closed disk.

Recall that since p is a branch covering map, the map $p : S \setminus p^{-1}(V_1(\hat{\mathbb{C}})) \rightarrow \hat{\mathbb{C}} \setminus V_1(\hat{\mathbb{C}})$ defines a regular cover. Assume that $k \in V_0$. First, if $k \in V_0 \setminus p^{-1}(V_1(\hat{\mathbb{C}}))$ then k lies on the regular cover and by construction the restricted map $p|_{\text{Star}(k)} : \text{Star}(k) \rightarrow \text{Star}(p(k))$ is a homeomorphism. Second, if $k \in V_0 \cap p^{-1}(V_1(\hat{\mathbb{C}}))$ then $p|_{\text{Star}(k)} : \text{Star}(k) \rightarrow \text{Star}(p(k))$ is a branch covering map between closed disks with only one ramification point k in the domain's interior with ramification index $N_k = 1, 2, 3, \dots, N$. However, since k is not an actual branch point, the index is $N_k = 1$ and so the restricted map $p|_{\text{Star}(k)}$ is again a homeomorphism. Consequently, since the boundaries $\partial \text{OStar}(k)$ and $\partial \text{OStar}(p(k))$ are homeomorphic with the same number of dual edges, and also $\theta_{ij} = \hat{\theta}_{p(ij)}$,

$$\sum_{ij^* \subset \partial \text{OStar}(k)} (\pi - \theta_{ij}) = \sum_{uv^* \subset \partial \text{OStar}(p(k))} (\pi - \hat{\theta}_{uv}) = 2\pi,$$

where the last equality follows from condition 2 of Rivin's Theorem 5 (as well as condition 2 of Bao and Bonahon's Theorem 6). Thus, condition 2 of Schlenker's Theorem 4 is verified.

8.2 Verification of condition 4 of Theorem 4

8.2.1 The case of admissible domains with general topology.

Next, we focus on the verification of condition 4 of Theorem 4. Assume Ω is a strict admissible domain of S, \mathcal{C} which is not the open star of a vertex from V_0 , but it can be the open star of a vertex from V_1 . By Definitions 7 and 8, Ω is homeomorphic to the interior of a compact surface with boundary. Therefore its Euler characteristic is $\chi(\Omega) = 2 - 2H - B$, where H is the number of handles and B is the number of boundary components of Ω . Observe that $B \geq 1$ because $\partial\Omega \neq \emptyset$. Consequently, $\chi(\Omega) = 2 - 2H - B$ yields the restriction $\chi(\Omega) \leq 1$ where $\chi(\Omega) = 1$ exactly when Ω has one boundary component and no handles, which means that it is an open topological disk. In all other cases $\chi(\Omega) < 0$. Since $\partial\Omega \neq \emptyset$, either $\partial\Omega \cap E^* \neq \emptyset$ or $\partial\Omega \cap V_1 \neq \emptyset$ or both. Therefore whenever $\chi(\Omega) < 1$,

$$\sum_{ij^* \subset \partial\Omega} (\pi - \theta_{ij}) + \pi |\partial\Omega \cap V_1| > 0 \geq 2\pi\chi(\Omega).$$

The latter inequality is condition 4 of Theorem 4 for strict admissible domains that are not open topological disks. Hence, we focus on the case when Ω is an open topological disk and so $\chi(\Omega) = 1$.

8.2.2 The case of admissible topological disks.

Denote by $\gamma = \partial\Omega$ the loop in the one-skeleton of $\hat{\mathcal{T}}$ that traverses the boundary of the topological disk Ω . Recall that we interpret $\gamma = (\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n, \hat{e}_1)$ as a cyclic sequence

of edges of $\hat{\mathcal{T}}$, so an edge $\hat{e}_s \in \hat{E}$ could appear twice or the loop may pass through the same vertex several times. Let us denote by $|\gamma|$ the one dimensional subcomplex of the one-skeleton of $\hat{\mathcal{T}}$ formed by the edges and vertices on γ , so repetitions of edges and vertices are ignored. Hence, while γ is a sequence of edges (i.e. ordered), $|\gamma|$ is a one-complex with vertices and edges (unordered). Absolutely the same notation we use in the case when γ is a path in $\hat{\mathcal{T}}(1)$.

The image of γ under the branch covering map p will be denoted by $\delta = p(\gamma)$ and it will be interpreted as $\delta = (p(\hat{e}_1), p(\hat{e}_2), \dots, p(\hat{e}_n), p(\hat{e}_1))$ which is a loop in the one-skeleton of $\hat{\mathcal{T}}_{\hat{\mathbb{C}}}$. Notice that the number of edges in γ and δ (counting repetitions) is the same.

Case 1. Let $\partial\Omega \cap V_1 = \emptyset$. Then $\gamma = \partial\Omega$ consists entirely of dual edges, i.e. it is a path in the one-skeleton of the dual complex \mathcal{C}^* . Consequently, γ lies on the regular cover $p : S \setminus p^{-1}(V_1(\hat{\mathbb{C}})) \rightarrow \hat{\mathbb{C}} \setminus V_1(\hat{\mathbb{C}})$. Furthermore, its image $\delta = p(\gamma)$ is a loop in the one-skeleton of the dual complex $\mathcal{C}_{\hat{\mathbb{C}}}^*$ on the Riemann sphere.

Subcase 1.1. Let $|\delta|$ be a simple closed curve. Recall the difference between δ and $|\delta|$. While the one-dimensional subcomplex $|\delta|$ of the one-skeleton of $\mathcal{C}_{\hat{\mathbb{C}}}^*$ is a simple closed curve on $\hat{\mathbb{C}}$, the sequence of edges δ is not necessarily simple and it may traverse $|\delta|$ several times. However, the assumption that $|\delta|$ is simple implies that there exists a simple closed loop δ' in the one skeleton of \mathcal{C}^* (i.e. a cyclic subsequence of δ without edge repetitions), such that $|\delta'| = |\delta|$. Then $\hat{\mathbb{C}} \setminus |\delta| = \Omega_1 \cup \Omega_2$ where Ω_1 and Ω_2 are disjoint open topological discs on the Riemann sphere such that $\partial\Omega_1 = \partial\Omega_2 = \delta'$.

Subcase 1.1.A. Assume that neither Ω_1 nor Ω_2 is an open star of a vertex from $V_{\hat{\mathbb{C}}}$. Then by condition 2 of Theorem 5, or alternatively Theorem 6,

$$\sum_{uv^* \in \delta'} (\pi - \hat{\theta}_{uv}) > 2\pi.$$

Since the map p is onto and $\delta = p(\gamma)$ as well as $\theta_{ij} = \hat{\theta}_{p(ij)}$ for all $ij \in E$, it is immediate to conclude that

$$\sum_{ij^* \in \gamma} (\pi - \theta_{ij}) = \sum_{uv^* \in \delta} (\pi - \hat{\theta}_{uv}) \geq \sum_{uv^* \in \delta'} (\pi - \hat{\theta}_{uv}) > 2\pi.$$

Subcase 1.1.B. We claim that under the assumptions of condition 4 of Theorem 4, neither Ω_1 nor Ω_2 can be the open star of a vertex from $V_0(\hat{\mathbb{C}})$. Indeed, assume that one of the two domains, say Ω_2 , is an open star $\text{OStar}(\tilde{k}_0)$, where $\tilde{k}_0 \in V_0(\hat{\mathbb{C}})$. Then its closure is $\bar{\Omega}_2 = \text{Star}(\tilde{k}_0)$ and it lies in the target space of a regular cover. Consequently, due to the contractibility of $\text{Star}(\tilde{k}_0)$ and the lifting property of covering spaces, the preimage $p^{-1}(\text{Star}(\tilde{k}_0))$ is a disjoint union of closed stars $\text{Star}(k_s)$ for $s = 1 \dots N$, where $k_s \in V_0$. The restriction of p on each $\text{Star}(k_s)$ is a homeomorphism. Consequently, the full preimage of δ is the disjoint union of the boundaries $\partial\text{Star}(k_s)$ for $s = 1 \dots N$. Since $\delta = p(\gamma)$, where γ is a connected closed loop in $\mathcal{C}^*(1)$ which bounds an admissible domain on S , γ must be among the loops $\partial\text{Star}(k_s)$, i.e. $\gamma = \partial\text{Star}(k_t)$ for some specific

$t \in \{1, \dots, N\}$ with $k_t \in V_0$. Consequently the loop γ is simple and it splits the surface S into two open subdomains $\text{OStar}(k_t)$ and $S \setminus \text{Star}(k_t)$. On the one hand, we have assumed that the admissible domain Ω is not the open star of a vertex from V_0 which means that Ω is not $\text{OStar}(k_t)$. On the other hand, we have assumed that Ω is a topological disc, while $S \setminus \text{Star}(k_t)$ is definitely not a disc but a surface of genus at least one with one closed disk removed (so it has at least one handle). Thus, we have arrived at a contradiction due to the assumption that Ω_2 is the open star of $\tilde{k}_0 \in V_0(\hat{\mathbb{C}})$. Hence, this situation cannot occur.

Subcase 1.1.C. One of the two open discs, say Ω_1 , is the open star $\text{OStar}(\tilde{k}_1)$ of a vertex $\tilde{k}_1 \in V_1(\hat{\mathbb{C}})$. Then, by construction of the complexes \mathcal{C}^* and $\hat{\mathcal{T}}$, the preimage of the closed star $p^{-1}(\text{Star}(\tilde{k}_1))$ is a disjoint union of closed stars $\text{Star}(k_s) : s = 1..M$ where $M < N$. In contrast with the case from the previous paragraph, this time the restriction of p on each closed disk $\text{Star}(k_s)$ is a branch covering map with exactly one ramification point k_s with index N_s and one branch point $p(k_s) = \tilde{k}_1$. Consequently, the full preimage of $|\delta|$ is given by the disjoint union of boundary loops $\partial\text{Star}(k_s)$ for $s = 1..M$. Just like in the previous paragraph, since $\delta = p(\gamma)$ where γ is a connected closed loop in $\mathcal{C}^*(1)$ which bounds an admissible domain on S , γ must be among the loops $\partial\text{Star}(k_s)$, i.e. $\gamma = \partial\text{Star}(k_t)$ for some specific $t \in \{1, \dots, M\}$. Consequently the loop γ is simple and it splits the surface S into two open subdomains $\text{OStar}(k_t)$ and $S \setminus \text{Star}(k_t)$ one of which should be the admissible domain Ω . Just like before, $S \setminus \text{Star}(k_t)$ is not a topological disc, but a surface with at least one handle, so the only option left is $\Omega = \text{OStar}(k_t)$. As we have assumed that k_t cannot be from V_0 it has to belong to the set of ramification points V_1 and its index of ramification should be $N_{k_t} > 1$. Therefore γ , and thus its projection δ , cover the simple loop δ' a number of N_{k_t} -times. Therefore, by condition 2 of Theorem 5

$$\sum_{ij^* \in \gamma} (\pi - \theta_{ij}) = \sum_{uv^* \in \delta} (\pi - \hat{\theta}_{uv}) = N_{k_t} \sum_{uv^* \in \delta'} (\pi - \hat{\theta}_{uv}) = 2\pi N_{k_t} > 2\pi.$$

So far we have concluded that whenever $|\delta|$ is simple closed curve, condition 4 of Schlenker's Theorem 4 holds.

Subcase 1.2. Assume $|\delta|$ is not a simple closed curve on the Riemann sphere, but as a one dimensional connected subcomplex of $\mathcal{C}_{\hat{\mathbb{C}}}^*(1)$ it is not simply-connected. Equivalently, $|\delta|$ has a non-trivial fundamental group. For that reason there exists a cyclic sequence δ' of edges of $|\delta|$ that defines a simple closed loop in the one-skeleton of $\mathcal{C}_{\hat{\mathbb{C}}}^*$. In particular, δ' is a cyclic subsequence of δ and because $|\delta|$ is not simple while $|\delta'|$ is, $\delta \setminus \delta' \neq \emptyset$. Consequently, by applying again condition 2 of Theorem 5

$$\begin{aligned} \sum_{ij^* \in \gamma} (\pi - \theta_{ij}) &= \sum_{uv^* \in \delta} (\pi - \hat{\theta}_{uv}) \\ &= \sum_{uv^* \in \delta'} (\pi - \hat{\theta}_{uv}) + \sum_{uv^* \in \delta \setminus \delta'} (\pi - \hat{\theta}_{uv}) > \sum_{uv^* \in \delta'} (\pi - \hat{\theta}_{uv}) \geq 2\pi. \end{aligned}$$

Subcase 1.3. $|\delta|$ is a simply-connected subcomplex of the one-skeleton of the dual complex $\mathcal{C}_{\hat{\mathbb{C}}}^*$. This is equivalent to saying that $|\delta|$ is a tree in $\mathcal{C}_{\hat{\mathbb{C}}}^*(1)$ and as such it

is contractible to a point. Therefore, by the lifting properties of covering maps, the full preimage of $|\delta|$ via p is a disjoint union of N homeomorphic copies of $|\delta|$ in the one-skeleton of the dual complex \mathcal{C}^* on S . Since γ is a lift of δ under p and is connected, $|\gamma|$ should be one of these copies. Hence, $|\gamma|$ is a tree and furthermore, $\Omega = S \setminus |\gamma|$. Since $|\gamma|$ is contractible, Ω is homeomorphic to S with a closed disc removed, which, as already pointed out, is not a topological disc. Hence, we conclude that this case cannot occur either.

To summarize, we have verified condition 4 of Theorem 4 whenever the boundary of Ω consists entirely of dual edges, i.e. $\partial\Omega \cap V_1 = \partial\Omega \cap V = \emptyset$.

Case 2. Let $|\gamma \cap V_1| \geq 2$, which means that γ passes through at least two different vertex points from V_1 or at least twice through the same point from V_1 . We claim that since γ is the boundary of the admissible domain Ω , the inequalities $|\gamma \cap V_1| \geq 3$ or $\gamma \cap E^* \neq \emptyset$ (or both) hold. Indeed, assume that this is not the case. Then $|\gamma \cap V_1| = 2$ and $\gamma \cap E^* = \emptyset$. There are only two ways this can happen. Either γ consists of exactly four corner edges or it consists of two corner edges repeated twice.

In the first case there are exactly two vertices i and $j \in \gamma \cap V_1$, two dual vertices O_f and $O_{f'} \in V^*$ which form the loop of four corner edges $\gamma = (O_f i, i O_{f'}, O_{f'} j, j O_f)$. Both triangles $\hat{\Delta} = i O_f O_{f'}$ and $\hat{\Delta}' = j O_f O_{f'}$ are two faces of $\hat{\mathcal{T}}$ (see Section 5) that have two points in common, namely O_f and $O_{f'}$. By strong regularity of $\hat{\mathcal{T}}$ the triangles $\hat{\Delta}$ and $\hat{\Delta}'$ share a common dual edge $O_f O_{f'}$ and thus γ is the boundary of the topological disc $\hat{\Delta} \cup \hat{\Delta}'$. Therefore γ separates S into two open subdomains, namely the open interior of $\hat{\Delta} \cup \hat{\Delta}'$ and $S \setminus (\hat{\Delta} \cup \hat{\Delta}')$. However, neither of them can be Ω because (i) the former is not an admissible domain as it does not contain any points from V and (ii) the latter is not a topological disc.

In the second case, $\gamma = (O_f i, i O_{f'}, O_{f'} i, i O_f)$. Then Ω can only be $S \setminus \gamma$, which is not possible since $S \setminus \gamma$ is not a topological disc.

Thus, we conclude that $|\gamma \cap V_1| \geq 3$ or $\gamma \cap E^* \neq \emptyset$, which yields the inequality

$$\sum_{ij^* \subset \gamma} (\pi - \theta_{ij}) + \pi |\gamma \cap V_1| > 2\pi.$$

Case 3. Let $|\gamma \cap V_1| = 1$, which means that γ passes through exactly one vertex point k_1 from V_1 exactly once. Then by projecting down to $\hat{\mathcal{C}}$ via p , we obtain $\delta = p(\gamma)$ which is a loop in $\hat{\mathcal{T}}_{\hat{\mathcal{C}}}(1)$ that passes only once through only one point from $V_1(\hat{\mathcal{C}})$ denoted by $\tilde{k}_1 = p(k_1)$. Remove from γ the point k_1 together with the two corner edges on γ attached to k_1 in order to obtain a path γ_1 in the one-skeleton of \mathcal{C}^* . Let $\delta_1 = p(\gamma_1)$ which is a path in the one-skeleton of $\mathcal{C}_{\hat{\mathcal{C}}}^*$ obtained by removing \tilde{k}_1 and its two adjacent corner edges (which may also be only one adjacent corner edge repeated twice) from δ .

As already discussed in Subcase 1.1.C, the restriction of p on the closed star $\text{Star}(k_1)$ is a branch covering map onto the closed star $\text{Star}(\tilde{k}_1)$ with one ramification point k_1 of ramification index N_{k_1} . Let us denote by O_{f_1} and O_{f_2} the two different dual vertices on the boundary of $\text{Star}(k_1)$ connected by the path γ_1 . Then their images $p(O_{f_1}) = O_{\tilde{f}_1}$

and $p(O_{f_2}) = O_{\tilde{f}_2}$ are the two dual vertices on the boundary of $\text{Star}(\tilde{k}_1)$ connected by δ_1 . Just like in the case of loops before, we are going to look at different cases for the topology of the one dimensional subcomplex $|\delta_1|$ of $\mathcal{C}_{\mathbb{C}}^*(1)$.

Subcase 3.1. Let $|\delta_1|$ be non-simply connected. This means that it has a non-trivial fundamental group and so there exists a cyclic sequence δ'_1 of edges of $|\delta_1|$ that defines a simple closed loop in $\mathcal{C}_{\mathbb{C}}^*(1)$. Therefore, by condition 2 of Theorem 5

$$\begin{aligned} \sum_{ij^* \subset \gamma_1} (\pi - \theta_{ij}) &= \sum_{uv^* \subset \delta_1} (\pi - \hat{\theta}_{ij}) = \sum_{uv^* \subset \delta'_1} (\pi - \hat{\theta}_{ij}) + \sum_{uv^* \subset \delta_1 \setminus \delta'_1} (\pi - \hat{\theta}_{ij}) \\ &\geq 2\pi + \sum_{uv^* \subset \delta_1 \setminus \delta'_1} (\pi - \hat{\theta}_{ij}) \geq 2\pi > \pi, \end{aligned}$$

and so $\sum_{ij^* \subset \gamma} (\pi - \theta_{ij}) + \pi > 2\pi$.

Subcase 3.2. Assume that $|\delta_1|$ is simply connected. This means that it is contractible and so it is a tree in the one-skeleton of $\mathcal{C}_{\mathbb{C}}^*$. Then, by the lifting properties of the regular covering map p on $S \setminus p^{-1}(V_1(\hat{\mathbb{C}}))$, the complex $|\gamma_1|$ is also a tree and is homeomorphic to $|\delta_1|$ via p . Consequently, since the path γ_1 traverses the tree $|\gamma_1|$, its image $\delta_1 = p(\gamma_1)$ traverses $|\delta_1|$ in the same way. Adding back the two corner edges $O_{f_1}k_1$ and $k_1O_{f_2}$ to γ_1 and $O_{\tilde{f}_1}\tilde{k}_1$ and $\tilde{k}_1O_{\tilde{f}_2}$ to δ_1 restores the loops γ and δ respectively, showing that the restricted map $p|_{|\gamma_1|} : |\gamma_1| \rightarrow |\delta_1|$ is a homeomorphism, due to the fact that it is a homeomorphism between $|\gamma_1|$ and $|\delta_1|$. Observe that the two corner edges are either different for each of the two loops γ and δ , or they coincide for each of these two loops.

Subcase 3.2.A. Assume $O_{\tilde{f}_1} \equiv O_{\tilde{f}_2}$. This is true exactly when $O_{f_1} \equiv O_{f_2}$. Then $|\gamma|$ is contractible, i.e. it is a tree, so $\Omega = S \setminus |\gamma|$ is not a topological disc when $g(S) \geq 1$. Hence, this scenario is impossible.

Subcase 3.2.B. Let $O_{\tilde{f}_1} \neq O_{\tilde{f}_2}$. This is true exactly when $O_{f_1} \neq O_{f_2}$. A *true leaf* of the tree $|\gamma_1|$ is a leaf which is neither O_{f_1} nor O_{f_2} . A *true leaf-edge* is the unique edge of the tree attached to a true leaf. The same terminology applies to $|\delta_1|$. Remove all true leaves and leaf-edges of $|\gamma_1|$. Perform the same operation on the homeomorphic tree $|\delta_1|$. Every time we remove a true leaf together with its corresponding true leaf-edge from $|\gamma_1|$, we actually add them to the admissible domain Ω that γ bounds, obtaining a new admissible domain. As each time we add one vertex with one edge attached to it, the Euler characteristic of the newly obtained domain is preserved, i.e. we obtain an open topological disc again. After removing the true leaves and leaf-edges from $|\gamma_1|$ and $|\delta_1|$, we end up with a pair of smaller trees, again homeomorphic via p . If these new trees have any true leaves, we repeat the procedure. We keep repeating until there are no true leaves left and the only leaves left are O_{f_1} and O_{f_2} from γ_1 and $O_{\tilde{f}_1}$ and $O_{\tilde{f}_2}$ from δ_1 . On the level of admissible domains on S , this procedure enlarges the initial admissible domain Ω to the admissible domain $\Omega' \supset \Omega$, where the latter is obtained by adding to the former all removed true leaves and leaf-edges. In the end, what is left from $|\gamma_1|$ and $|\delta_1|$ is a pair of homeomorphic simple paths in the one-skeletons of the dual complexes

\mathcal{C}^* and $\mathcal{C}_{\hat{\mathbb{C}}}^*$ respectively. Denote these two paths by γ'_1 and δ'_1 . Furthermore, γ'_1 lies on the boundary of Ω' so that if we add to γ'_1 the two corner edges $O_{f_1}k_1$ and $k_1O_{f_2}$ then we obtain the full boundary, call it γ' , of Ω' . Consequently, γ' is a simple loop in the one-skeleton of \mathcal{C}^* and $\gamma' = \partial\Omega'$. Recall that by construction $\delta'_1 = p(\gamma'_1)$ is a simple path and so $\delta' = p(\gamma')$ is a simple loop.

Assume that δ'_1 lies on the boundary of a closed star $\text{Star}(\tilde{k})$ for some $\tilde{k} \in V_{\hat{\mathbb{C}}}$. Since, by construction, the cell complexes on S are the lifts of the cell complexes on $\hat{\mathbb{C}}$ via p , the preimage of $\text{Star}(\tilde{k})$ is a disjoint union of stars and thus γ'_1 lies in the boundary of a star $\text{Star}(k)$ for some $k \in V$ such that $p(k) = \tilde{k}$.

First, if $k \neq k_1$, then the endpoints O_{f_1} and O_{f_2} of γ'_1 lie simultaneously in $\text{Star}(k)$ and $\text{Star}(k_1)$. As these two endpoints are by assumption different, the two stars share two different vertices and so, by strong regularity of $\hat{\mathcal{T}}$, they share exactly one common dual edge $O_{f_1}O_{f_2} \in E^*$. Then, if we denote by $\hat{\Delta}$ the (closed) triangular face $k_1O_{f_1}O_{f_2}$ of $\hat{\mathcal{T}}$, we conclude that the simple loop γ' splits the surface S into two open domains, one of which is the open interior of $\text{Star}(k) \cup \hat{\Delta}$ and the other is $S \setminus (\text{Star}(k) \cup \hat{\Delta})$. None of them can be Ω' because the latter is not an open topological disc, whenever $g(S) \geq 1$, while the former is a topological disc, but it is not an admissible domain, as it cannot be represented as the union of open stars of $\hat{\mathcal{T}}$ due to the presence of the additional triangle $\hat{\Delta}$. Therefore, this situation cannot occur.

Second, if $k = k_1$, then on one side γ' bounds on S a strict subdomain of $\text{Star}(k_1)$ which does not even contain an open star, so the domain cannot be admissible, while on the other side of γ' we have S with a closed topological disc removed, which, as before, is not an open topological disc. Therefore, this situation cannot occur either.

Consequently, δ'_1 cannot lie on the boundary of any star, which is equivalent to saying that the simple path of dual edges δ'_1 does not lie in the boundary of a dual face of $\mathcal{C}_{\hat{\mathbb{C}}}^*$. By Bao and Bonahon's Theorem 6 $\sum_{ij^* \subset \gamma'_1} (\pi - \theta_{ij}) = \sum_{uv^* \subset \delta'_1} (\pi - \hat{\theta}_{uv}) > \pi$. Returning to the original loop $\gamma = \partial\Omega$ which contains γ'_1

$$\begin{aligned} \sum_{ij^* \subset \gamma} (\pi - \theta_{ij}) + \pi &= \sum_{ij^* \subset \gamma'_1} (\pi - \theta_{ij}) + \sum_{ij^* \subset \gamma \setminus \gamma'_1} (\pi - \theta_{ij}) + \pi \\ &> \pi + \sum_{ij^* \subset \gamma \setminus \gamma'_1} (\pi - \theta_{ij}) + \pi \geq 2\pi. \end{aligned}$$

Finally, we have concluded that condition 4 of Theorem 4 holds for all possible cases. Thus, the proof of Theorem 2 is complete.

9 Proof of Theorem 3

The proof of Theorem 3 is analogous to the proof of Theorem 2 from the preceding section. The major difference is that instead of Schlenker's Theorem 4, our main tool is Bao and Bonahon's Theorem 6 because S is a topological sphere. Consequently, instead of working with strict admissible domains with different topologies, we check the necessary and sufficient conditions of Theorem 6 for simple loops γ in the one-skeleton of the dual

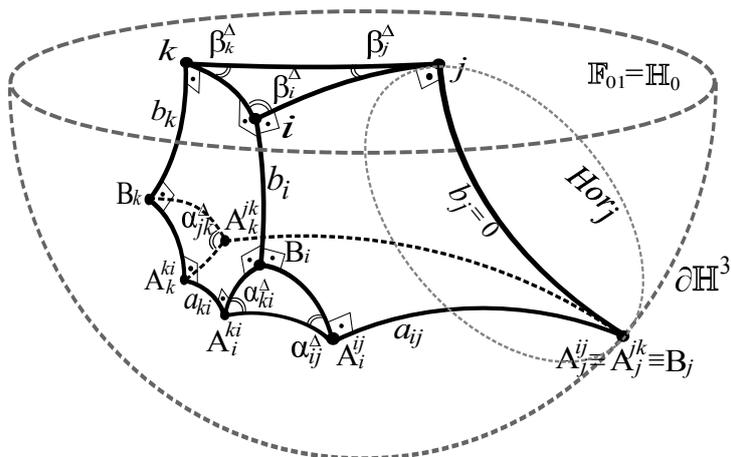


Figure 5: A hyper-ideal tetrahedron τ_Δ .

complex \mathcal{C}^* on S , as well as for simple paths γ_1 in $\mathcal{C}^*(1)$ joining two dual vertices from the same dual face without being contained in the boundary of any dual face. In the case of loops, we simply repeat the arguments from Subcases 1.1. and 1.2 in Section 8, observing that due to the simplicity of γ , Subcase 1.3 cannot occur. In the case of paths, we repeat the arguments from Subcases 3.1 and 3.2 with the simplification that neither $|\gamma_1|$ nor $p(|\gamma_1|)$ have true leaves, so the procedure of removing them is unnecessary in this case.

10 Spaces of decorated triangles and circle patterns

In order to outline the algorithmic recipe for discrete uniformization we need to provide some background on the spaces of decorated triangles and hyper-ideal circle patterns. In sections 3 and 4 we have explained how one can obtain in a natural and systematic way combinatorial angle data $(\mathcal{C}, \theta, \Theta)$ from either a negatively curved or a hyper-elliptic Riemann surface. Theorem 4 guarantees the unique geometric realizability of the data $(\mathcal{C}, \theta, \Theta)$. In the following two sections 10 and 11 we describe the variational principle which plays a central role in the verification of Theorem 4 (see [11]) and provides a method for the construction of hyper-ideal circle patterns. All constructions and notations follow closely the exposition of [11]. The interested reader can check the necessary details there. In the current article, we provide a sketch.

10.1 The space of decorated triangles

First, we describe the space of (labelled) decorated triangles in \mathbb{H}^2 considered up to hyperbolic isometries, with predetermined fixed splitting of the triangles' vertices $V_\Delta = \{i, j, k\}$ into V_Δ^1 and V_Δ^0 . Then, up to \mathbb{H}^2 -isometry, a decorated triangle $\Delta = ijk$ can be uniquely represented in three different ways [11]. The most natural way is to give the

triangles edge lengths and vertex radii $(l, r)_\Delta = (l_{ij}, l_{jk}, l_{ki}, r_k, r_i, r_j) \in \mathcal{ER}_\Delta$, where the first three are positive and satisfy all three strict triangle inequalities $l_{uv} < l_{vw} + l_{wu}$, as well as $l_{uv} > r_u + r_v$ for all $u \neq v \neq w \in \{i, j, k\}$, while the last three satisfy $r_u > 0$ for $u \in V_\Delta^1$ and $r_u = 0$ for $u \in V_\Delta^0$. The second way of uniquely representing a decorated triangle is by its six angles $(\alpha^\Delta, \beta^\Delta) = (\alpha_{ij}^\Delta, \alpha_{jk}^\Delta, \alpha_{ki}^\Delta, \beta_k^\Delta, \beta_i^\Delta, \beta_j^\Delta) \in \mathcal{A}_\Delta$ [26, 31, 11]. Here α_{uv}^Δ is the angle between the geodesic edge uv of $\Delta = ijk$ and its face circle c_Δ measured inside c_Δ and outside Δ , and β_v^Δ is the interior angle of the triangle at its vertex $v \in V_\Delta$ (refer to Figure 1). The third type of decorated triangle description is the one that is most crucial to our later constructions. It draws upon the link between decorated triangles in \mathbb{H}^2 and hyper-ideal tetrahedra in \mathbb{H}^3 .

10.2 Hyper-ideal tetrahedra

Definition 9. A hyper-ideal tetrahedron (see [26, 31] and Figure 5) is a geodesic polyhedron in \mathbb{H}^3 that has the combinatorics of a tetrahedron with some (possibly all) of its vertices truncated by triangular truncating faces. Each truncating face is orthogonal to the faces and the edges it truncates. Furthermore, a pair of truncating faces do not intersect. Finally, the non-truncated vertices are all ideal.

The polyhedron depicted on Figure 5 is an example of a hyper-ideal tetrahedron with three truncated *hyper-ideal vertices* and one ideal vertex. This terminology comes from the interpretation that in the Klein projective model or the Minkowski space-time model of \mathbb{H}^3 [34, 3, 31] the hyper-ideal tetrahedron can be represented by an actual tetrahedron with some (or all) vertices lying outside \mathbb{H}^3 (hence the term *hyper-ideal vertices*). The dual (projective polar) to each hyper-ideal vertex is the orthogonal truncating plane. However, in this article, we mostly use the two standard conformal models of \mathbb{H}^3 - the Poincaré ball model and the upper half-space model [33, 34, 3], which are the three dimensional analogs of the disk model and the upper half-plane model of \mathbb{H}^2 respectively. We call a *principal edge* a geodesic edge of a hyper-ideal tetrahedron that is not contained in a truncating plane of the hyper-ideal tetrahedron. Each hyper-ideal tetrahedron has exactly six of them. The rest of the edges lie on the truncating planes. We can call them *auxiliary edges*. A principal edge has one of the following three properties: (i) it goes from one truncating face to another, being by definition orthogonal to both of them (think of it as the two-sided truncation of the edge connecting two hyper-ideal vertices); (ii) it goes from one ideal vertex to a truncating face, being by definition orthogonal to the truncating face (think of it as the one-sided truncation of the edge connecting an ideal vertex to a hyper-ideal one); (iii) it connects two ideal vertices. See Figure 5.

10.3 Link between decorated triangles and hyper-ideal tetrahedra

To obtain a hyper-ideal tetrahedron τ_Δ from a decorated triangle $\Delta = ijk$, first think of \mathbb{H}^2 , together with the decorated triangle Δ drawn on it, as a hyperbolic plane $\mathbb{H}_0 \cong \mathbb{H}^2$ lying in \mathbb{H}^3 . This situation is depicted on Figure 6. The idea is that we can project \mathbb{H}_0 down on $\partial\mathbb{H}^3$ using a natural geometric map $F_{proj} : \mathbb{H}_0 \rightarrow \partial\mathbb{H}^3$ (see also Figure 6).

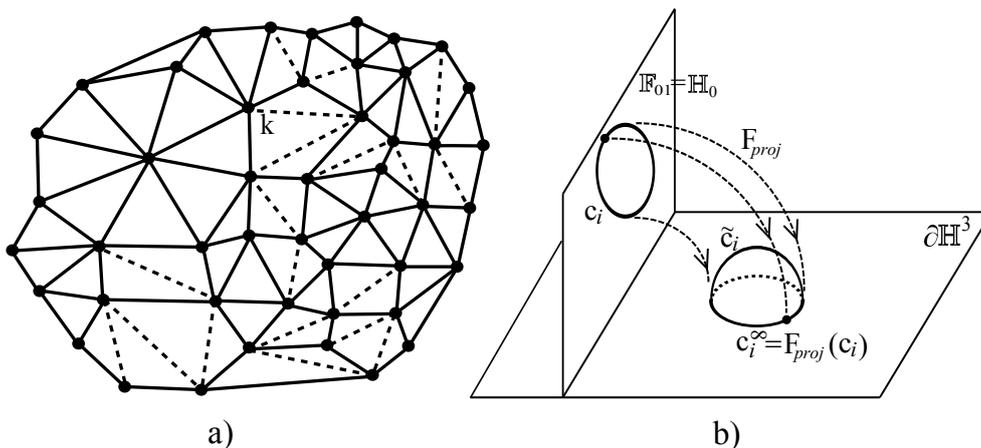


Figure 6: a) The subtriangulation $\mathcal{T} = (V, E_T, F_T)$ of \mathcal{C} whose dashed edges are the auxiliary edges from E_π ; b) The projection F_{proj} from the hyperbolic plane $\mathbb{H}_0 \subset \mathbb{H}^3$ to the ideal boundary $\partial\mathbb{H}^3$.

For any $x \in \mathbb{H}_0$ take the unique geodesic in \mathbb{H}^3 passing through x and perpendicular to \mathbb{H}_0 and follow it, only on one side of \mathbb{H}_0 , all the way down to $\partial\mathbb{H}^3$ reaching the ideal point $F_{proj}(x)$. Then the image $F_{proj}(\Delta)$ of the decorated triangle Δ determines a set of circles (and straight lines, depending on the model) in $\partial\mathbb{H}^3$ which when extended to hyperbolic planes cut out a hyper-ideal tetrahedron. Figure 6 shows how a vertex circle c_i is being mapped to a circle $F_{proj}(c_i) = c_i^\infty$ on $\partial\mathbb{H}^3$ and then extended to a hyperbolic plane \tilde{c}_i of \mathbb{H}^3 . The converse construction also holds in the sense that a (labelled) hyper-ideal tetrahedron τ gives rise to a unique decorated triangle Δ_τ . The face labelled ijk determines the hyperbolic plane \mathbb{H}_0 . Then, we can take all the ideal circles (and possibly straight lines, depending on the \mathbb{H}^3 model) on $\partial\mathbb{H}^3$ of all hyperbolic planes that the faces of τ determine, and map them back to \mathbb{H}_0 via F_{proj}^{-1} . Notice that since all constructions utilize only the geometry of \mathbb{H}^3 , they are invariant with respect to hyperbolic congruences (that respect the labelling of the triangles and the tetrahedra). Furthermore, because the models of both $\mathbb{H}^2 \cong \mathbb{H}_0$ and \mathbb{H}^3 are conformal, the six angles $(\alpha^\Delta, \beta^\Delta)$ of the decorated triangle become the corresponding dihedral angles at the six principal edges of the constructed hyper-ideal tetrahedron τ_Δ . This can be seen on Figure 5.

10.4 Tetrahedral edge-length variables

Now, after we have explained the one-to-one correspondence between (labelled) decorated triangles in \mathbb{H}^2 and (labelled) hyper-ideal tetrahedra, we can define the variables $(a, b)_\Delta = (a_{ij}, a_{jk}, a_{ki}, b_k, b_i, b_j) \in \mathcal{TE}_\Delta$ as the ‘‘hyperbolic lengths’’ of the principal geodesic edges of τ_Δ . For principal edges of type (i), as described above (i.e. edges perpendicular to two truncating faces), the geodesic length makes sense and is a positive real number. But edges of type (ii) and (iii) have actually infinite hyperbolic length. In order to fix this

issue, we *decorate* the hyper-ideal tetrahedron τ_Δ , which was obtained from the decorated triangle $\Delta = ijk$, with one horosphere per ideal vertex so that the horosphere touches $\partial\mathbb{H}^3$ at that ideal vertex and the plane \mathbb{H}_0 at the endpoint of the geodesic edge emanating from the ideal vertex in question. On Figure 5 B_j is an ideal vertex and the decorating horosphere Hor_j is tangent to $\partial\mathbb{H}^3$ at B_j and to \mathbb{H}_0 at j . Consequently, the length of a type (ii) geodesic edge of τ_Δ is the oriented hyperbolic distance between the truncating face on one side of the edge and the decorating horosphere on the other side, measured along the edge itself. The length is positive if the truncating face and the decorating horosphere are disjoint, zero if they are tangent and negative if they intersect. Similarly, the length of a type (iii) geodesic edge of τ_Δ is the oriented hyperbolic distance between the two decorating horospheres, one on each side of the edge, measured along the edge itself. As before, the length is positive if the two decorating horospheres are disjoint, zero if they are tangent and negative if they intersect. Edges of type (i) and (ii) can be seen on Figure 5. For instance, a_{ki}, b_k and b_i are the hyperbolic lengths of geodesic edges of type (i), while a_{ij} is the length of the edge $A_i^{ij}B_j$ defined as the (oriented) distance between the truncating face $A_i^{ki}A_i^{ij}B_i$ and the decorating horosphere Hor_j . Notice that $b_j = 0$ since Hor_j is by construction tangent to \mathbb{H}_0 at the point j .

10.5 Transition formulas between different sets of variables

It is very important to find how the angles $(\alpha^\Delta, \beta^\Delta) \in \mathcal{A}_\Delta$ of a decorated triangle Δ , which are also the six principal dihedral angles of the corresponding hyper-ideal tetrahedron τ_Δ , depend on the principal edge-lengths $(a, b)_\Delta \in \mathcal{TE}_\Delta$ of τ_Δ . It is also useful to know how the edge-lengths and vertex radii $(l, r)_\Delta \in \mathcal{ER}_\Delta$ of Δ depend on the parameters $(a, b)_\Delta$. By using various combinations of hyperbolic trigonometric formulas [10] applied to the faces of τ_Δ one can derive the necessary expressions. In what follows, we present in detail the formulas for hyperbolic decorated triangles, as this is the more general case of higher genus surfaces. The Euclidean case is analogous and can be worked out by applying the appropriate formulas given in [11] for instance. For the edge-lengths and vertex radii of a hyperbolic $\Delta = ijk$ the following formulas apply [11]:

$$r_v = \sinh^{-1} \left(\frac{1}{\sinh b_v} \right) \text{ if } v \in V_\Delta^1 \quad \text{and} \quad r_v = b_v = 0 \text{ if } v \in V_\Delta^0 \quad (3)$$

$$l_{uv} = \cosh^{-1} \left(\frac{\cosh b_u \cosh b_v + \cosh a_{uv}}{\sinh b_u \sinh b_v} \right) = f_3(b_u, b_v, a_{uv}) \text{ if } u, v \in V_\Delta^1 \quad (4)$$

$$l_{uv} = \cosh^{-1} \left(\frac{\cosh b_u + e^{a_{uv}}}{\sinh b_u} \right) = f_2(b_u, a_{uv}) \text{ if } u \in V_\Delta^1, \text{ and } v \in V_\Delta^0 \quad (5)$$

$$l_{uv} = 2 \sinh^{-1} (e^{a_{uv}/2}) = f_1(a_{uv}) \text{ if } u, v \in V_\Delta^0, \quad (6)$$

where $uv \in E_\Delta = \{ij, jk, ki\}$ is an edge of $\Delta = ijk$. Having computed the edge-lengths of Δ one can immediately find, using for example the hyperbolic law of cosines, the angles

$$\beta_v^\Delta = \arccos \left(\frac{\cosh l_{uv} \cosh l_{vw} - \cosh l_{wu}}{\sinh l_{uv} \sinh l_{vw}} \right) = g(l_{uv}, l_{vw}, l_{wu}) \text{ for all } v \in V_\Delta. \quad (7)$$

In order to find the angles α_{uv}^Δ we examine all four combinatorial types of decorated triangle. There are different ways and different hyperbolic trigonometric formulas one can use, but we show only some of them (all leading to the same result).

Case 1. $V_\Delta^1 = \{i, k\}$ and $V_\Delta^0 = \{j\}$. The hyper-ideal tetrahedron τ_Δ corresponding to this combinatorics is depicted on Figure 5. We work with the notations introduced there for the current case as well as for the rest of the cases. Compute the geodesic edge-lengths of the triangular truncating face $A_i^{ki} A_i^{ij} B_i$

$$\begin{aligned}\sigma_i^{ki} &= l_{\mathbb{H}^3}(A_i^{ki} B_i) = \cosh^{-1} \left(\frac{\cosh a_{ki} \cosh b_i + \cosh b_k}{\sinh a_{ki} \sinh b_i} \right) = f_3(a_{ki}, b_i, b_k) \\ \sigma_i^{ij} &= l_{\mathbb{H}^3}(A_i^{ij} B_i) = \cosh^{-1} \left(\frac{\cosh b_i + e^{-a_{ij}}}{\sinh b_i} \right) = f_2(b_i, -a_{ij}) \\ \sigma_i &= l_{\mathbb{H}^3}(A_i^{ki} A_i^{ij}) = \cosh^{-1} \left(\frac{\cosh a_{ki} + e^{a_{jk} - a_{ij}}}{\sinh a_{ki}} \right) = f_2(a_{ki}, a_{jk} - a_{ij}).\end{aligned}$$

Then $\alpha_{ij}^\Delta = g(\sigma_i, \sigma_i^{ij}, \sigma_i^{ki})$ and $\alpha_{ki}^\Delta = g(\sigma_i, \sigma_i^{ki}, \sigma_i^{ij})$ (see formula (7)). Alternatively, one can also use the hyperbolic law of sines

$$\alpha_{ij}^\Delta = \arcsin \left(\frac{\sinh \sigma_i^{ki}}{\sinh \sigma_i} \sin \beta_i^\Delta \right) \quad \text{and} \quad \alpha_{ki}^\Delta = \arcsin \left(\frac{\sinh \sigma_i^{ij}}{\sinh \sigma_i} \sin \beta_i^\Delta \right)$$

The easiest way to compute the last angle is $\alpha_{jk}^\Delta = \pi - \alpha_{ij}^\Delta - \beta_j^\Delta$.

Case 2. $V_\Delta^1 = V_\Delta = \{i, j, k\}$ while $V_\Delta^0 = \emptyset$. One way of computing the angles is to apply the hyperbolic law of cosines to two out of the three truncating faces different from the truncating face ijk . For instance, compute the geodesic edge-lengths of the triangular truncating faces $A_i^{ki} A_i^{ij} B_i$ and $A_j^{ij} A_j^{jk} B_j$. Obtain

$$\begin{aligned}\sigma_i^{ij} &= f_3(a_{ij}, b_i, b_j), & \sigma_i^{ki} &= f_3(a_{ki}, b_i, b_k), & \sigma_i &= f_3(a_{ki}, a_{ij}, a_{jk}) & \text{for } A_i^{ki} A_i^{ij} B_i \\ \sigma_j^{ij} &= f_3(a_{ij}, b_j, b_i), & \sigma_j^{jk} &= f_3(a_{jk}, b_j, b_k), & \sigma_j &= f_3(a_{ij}, a_{jk}, a_{ki}) & \text{for } A_j^{ij} A_j^{jk} B_j.\end{aligned}$$

Then $\alpha_{ij}^\Delta = g(\sigma_i, \sigma_i^{ij}, \sigma_i^{ki}) = g(\sigma_j, \sigma_j^{ij}, \sigma_j^{jk})$, as well as $\alpha_{jk}^\Delta = g(\sigma_j, \sigma_j^{jk}, \sigma_j^{ij})$ and $\alpha_{ki}^\Delta = g(\sigma_i, \sigma_i^{ki}, \sigma_i^{ij})$ (see formula (7)). Alternatively, just like in Case 1, one can also use the hyperbolic law of sines, applied to $A_i^{ki} A_i^{ij} B_i$ and $A_j^{ij} A_j^{jk} B_j$ whose angles β_i^Δ and β_j^Δ we already know.

Case 3. $V_\Delta^1 = \{i\}$ and $V_\Delta^0 = \{j, k\}$. In this case, for example, compute the edge-lengths of the truncating face $A_i^{ki} A_i^{ij} B_i$

$$\sigma_i^{ij} = f_2(b_i, -a_{ij}), \quad \sigma_i^{ki} = f_2(b_i, -a_{ki}), \quad \sigma_i = f_1(a_{ki} + a_{ij} + a_{jk})$$

and then calculate $\alpha_{ij}^\Delta = g(\sigma_i, \sigma_i^{ij}, \sigma_i^{ki})$. After that $\alpha_{jk}^\Delta = \pi - \beta_j - \alpha_{ij}^\Delta$ and $\alpha_{ki}^\Delta = \alpha_{ij}^\Delta + \beta_j - \beta_k = g(\sigma_i, \sigma_i^{ki}, \sigma_i^{ij})$.

Case 4. $V_\Delta^1 = \emptyset$ and $V_\Delta^0 = V_\Delta = \{i, j, k\}$. Then

$$\alpha_{ij}^\Delta = \frac{\pi + \beta_k^\Delta - \beta_i^\Delta - \beta_j^\Delta}{2}, \quad \alpha_{jk}^\Delta = \frac{\pi + \beta_i^\Delta - \beta_j^\Delta - \beta_k^\Delta}{2}, \quad \alpha_{ki}^\Delta = \frac{\pi + \beta_j^\Delta - \beta_k^\Delta - \beta_i^\Delta}{2}.$$

10.6 The space of generalized hyper-ideal circle patterns

As already mentioned, we construct the hyper-ideal circle pattern that realizes the data $(\mathcal{C}, \theta, \Theta)$ as the unique critical point of a convex functional defined on a suitable space of patterns. In what follows, we define this space.

Description in terms of edge-lengths and vertex radii of decorated triangles.

Let $\mathcal{T} = (V, E_T, F_T)$ be a subtriangulation of \mathcal{C} obtained by adding a maximal number of diagonals with non-intersecting interiors in each non-triangular face of \mathcal{C} (see Figure 6). Consequently, the elements of F_T are all combinatorial triangles and $E_T = E \cup E_\pi$, where E_π is the set of all diagonals we have added in the process of subtriangulation of \mathcal{C} . On Figure 6 these are all the dashed edges, while all solid edges are the elements of E . We refer to the edges from E_π as *redundant edges*. With this new combinatorics at hand, define the space of all *generalized hyper-ideal circle patterns* with combinatorics \mathcal{T} on the surface S , considered up to isometry isotopic to identity as follows: the vector $(l, r) \in \mathbb{R}^{E_T} \times \mathbb{R}^V$ belongs to \mathcal{ER} if and only if

- $l_{ij} > 0$ for $ij \in E_T$, as well as $r_k > 0$ for $k \in V_1$ and $r_k = 0$ for $k \in V_0$;
- $l_{ij} > r_i + r_j$ for $ij \in E_T$;
- $l_{ij} < l_{jk} + l_{ki}$, $l_{jk} < l_{ki} + l_{ij}$, $l_{ki} < l_{ij} + l_{jk}$ for $\Delta = ijk \in F_T$.

By definition, the space \mathcal{ER} is clearly a convex polytope in $\mathbb{R}^{E_T} \times \mathbb{R}^V$ of dimension $\dim \mathcal{ER} = |E_T| + |V_1|$. We can consider the vector space $\mathbb{R}^{E_T} \times \mathbb{R}^{V_1}$ as the vector subspace of $\mathbb{R}^{E_T} \times \mathbb{R}^V$ defined by setting $r_k = 0$ for all $k \in V_0$. Then it is clear that \mathcal{ER} is an open convex polytope of $\mathbb{R}^{E_T} \times \mathbb{R}^{V_1}$. The term *generalized hyper-ideal circle pattern* is used because the hyper-ideal circle patterns that satisfy the conditions of \mathcal{ER} do not necessarily satisfy the local Delaunay property from Definition 3.

Description in terms of edge-lengths of hyper-ideal tetrahedra. Furthermore, one can define the set $\mathcal{TE} \subset \mathbb{R}^{E_T} \times \mathbb{R}^V$, which is in fact an open subset of $\mathbb{R}^{E_T} \times \mathbb{R}^{V_1}$, as the domain of a real-analytic map $\Psi : \mathcal{TE} \rightarrow \mathcal{ER}$ defined by formulas (3) to (6) so that $\Psi(a, b) = (l, r)$ and $\Psi(\mathcal{TE}) = \mathcal{ER}$. It is straightforward to verify that formulas (3) to (6) can be inverted and an inverse map $\Psi^{-1} : \mathcal{ER} \rightarrow \mathcal{TE}$ can be obtained, which is also real-analytic. Thus, one sees that Ψ is a real-analytic diffeomorphism between the open subsets \mathcal{TE} and \mathcal{ER} of $\mathbb{R}^{E_T} \times \mathbb{R}^{V_1}$. Therefore, \mathcal{TE} also defines the space of all generalized hyper-ideal circle patterns with combinatorics \mathcal{T} on S , considered up to hyperbolic isometries isotopic to identity.

The hyper-ideal circle pattern that realizes the data $(\mathcal{C}, \theta, \Theta)$ can be subtriangulated by adding all geodesic redundant edges from E_π so that now it has combinatorics $\mathcal{T} = (V, E_T, F_T)$ instead of $\mathcal{C} = (V, E, F)$, where recall that $E_T = E \cup E_\pi$. Consequently,

$ij \in E_\pi$ if and only if the two compatibly adjacent decorated triangles that share ij as a common geodesic edge have coinciding face circles, i.e. they share the same face circle. This is equivalent to the fact that the intersection angle between the two adjacent face circles is $\theta_{ij} = \pi$. Consequently, one can extend $\theta : E \rightarrow (0, \pi)$ to $\tilde{\theta} : E_T \rightarrow (0, \pi]$ by setting $\tilde{\theta}_{ij} = \theta_{ij}$ whenever $ij \in E$ and $\tilde{\theta}_{ij} = \pi$ whenever $ij \in E_\pi$. Then the pattern which realizes $(\mathcal{T}, \tilde{\theta}, \Theta)$ is exactly the pattern which realizes the original data $(\mathcal{C}, \theta, \Theta)$ after erasing the redundant edges E_π . Thus the pattern we are looking for can be seen as a special pattern lying inside the space \mathcal{TE} .

11 Variational principle for construction of circle patterns

In this section we construct the functional whose only critical point is the unique pattern that realizes the combinatorial angle data $(\mathcal{T}, \tilde{\theta}, \Theta)$, which as discussed in the previous Section 10, is also the pattern that realizes $(\mathcal{C}, \theta, \Theta)$ after the removal of the redundant edges.

Take any $\Delta \in F_T$. For each $(a, b)_\Delta \in \mathcal{TE}_\Delta$ define the function

$$\begin{aligned} \mathbf{U}_\Delta(a, b) &= \sum_{ij \in E_\Delta} \alpha_{ij}^\Delta a_{ij} + \sum_{k \in V_\Delta} \beta_k^\Delta b_k + 2\mathbf{V}(\alpha^\Delta, \beta^\Delta) \\ &= \sum_{ij \in E_\Delta} \alpha_{ij}^\Delta a_{ij} + \sum_{k \in V_\Delta^1} \beta_k^\Delta b_k + 2\mathbf{V}(\alpha^\Delta, \beta^\Delta) \end{aligned} \quad (8)$$

where, as already discussed in Section 10 (cases 1 to 4), the angles $\alpha_{ij}^\Delta = \alpha_{ij}^\Delta(a, b)$ for $ij \in E_\Delta$ and $\beta_k^\Delta = \beta_k^\Delta(a, b)$ for $k \in V_\Delta$ are real-analytic functions depending on the tetrahedral edge-length variables $(a, b)_\Delta \in \mathcal{TE}_\Delta$. Recall that E_Δ is the set of edges of Δ and V_Δ is the set of its vertices with V_Δ^1 being the subset of those vertices of Δ that are supposed to have vertex circles of positive radius. The vertices from its complement V_Δ^0 satisfy the restriction $b_k = 0$, $k \in V_\Delta^0$, which leads us to the second sum in (8). The function \mathbf{V} , which depends analytically on the angles $(\alpha^\Delta, \beta^\Delta)$, is the hyperbolic volume of the hyper-ideal tetrahedron τ_Δ with principal edge-lengths $(a, b)_\Delta \in \mathcal{TE}_\Delta$ and corresponding dihedral angles $(\alpha^\Delta, \beta^\Delta)$.

As a function of the dihedral angles, \mathbf{V} is strictly concave [26, 27, 31] and because of that, as shown in [11], \mathbf{U}_Δ is a locally strictly convex function on \mathcal{TE}_Δ . It is straightforward to verify that each real-analytic angle function $\alpha_{ij}^\Delta = \alpha_{ij}^\Delta(a, b)$ and $\beta_k^\Delta = \beta_k^\Delta(a, b)$ can be continuously extended by $\alpha_{ij}^\Delta \equiv \pi \equiv \beta_k^\Delta$ whenever $l_{ij} \geq l_{jk} + l_{ki}$, and $\alpha_{ij}^\Delta \equiv 0 \equiv \beta_k^\Delta$ whenever either $l_{jk} \geq l_{ki} + l_{ij}$ or $l_{ki} \geq l_{ij} + l_{jk}$. Let us partition the set E_Δ of edges of Δ into

$$E_\Delta^1 = \{ij \in E_\Delta \mid i, j \in V_\Delta^1\} \text{ and } E_\Delta^0 = E_\Delta \setminus E_\Delta^1.$$

Then the angle functions α^Δ and β^Δ are continuous on $\mathbb{R}^{E_\Delta^0} \times \mathbb{R}_+^{E_\Delta^1 \cup V_\Delta^1}$. Furthermore, outside \mathcal{TE}_Δ , the volume \mathbf{V} is constantly zero. Consequently, as explained for example in [5], the function \mathbf{U}_Δ , which is real-analytic and locally strictly convex in \mathcal{TE}_Δ , can

be extended to a continuously differentiable convex function on the whole convex set $\mathbb{R}^{E_\Delta^0} \times \mathbb{R}_+^{E_\Delta^1 \cup V_\Delta^1} \supset \mathcal{TE}_\Delta$.

The reason for which \mathbf{U}_Δ is suitable for applications is that there exist formulas for the volume \mathbf{V} of a hyper-ideal tetrahedron with given principal dihedral angles $(\alpha^\Delta, \beta^\Delta)$. For instance, due to Springborn [31], in the case of a hyper-ideal tetrahedron with at least one ideal vertex there is a fairly nice explicit expression for its volume in terms of Lobachevsky's functions. This formula however does not work for a hyper-ideal tetrahedron with exactly four hyper-ideal vertices. Nevertheless, a formula in that case also exists due to Ushijima [36]. Moreover, it covers all possible combinatorial types of hyper-ideal tetrahedra. However, it is quite more complicated and in our computer realizations of the discrete uniformization procedure, whose results can be found in Section 14, we have mostly used Springborn's formula (and its simplifications), while Ushijima's version has been used only for the cases of tetrahedra with four hyper-ideal vertices. All of these formulas are real-analytic in nature.

With the data $(\mathcal{T}, \tilde{\theta}, \Theta)$ at hand, one can construct the functional

$$\mathbf{U}_{\theta, \Theta}(a, b) = \sum_{\Delta \in F_T} \mathbf{U}_\Delta(a, b) - \sum_{ij \in E} \theta_{ij} a_{ij} - \pi \sum_{ij \in E_\pi} a_{ij} - \sum_{k \in V^1} \Theta_k b_k, \quad (9)$$

which is real-analytic and strictly locally convex on the open domain \mathcal{TE} (see [11]), as well as convex and continuously differentiable on the open convex set $\mathbb{R}^{E_T^0} \times \mathbb{R}_+^{E_T^1 \cup V^1} \supset \mathcal{TE}$ since it is a sum of convex continuously differentiable functions minus a linear function. The set $E_T^1 = \{ij \in E_T \mid i, j \in V^1\}$ and $E_T^0 = E_T \setminus E_T^1$.

The constructibility of the pattern we are after comes from the following result, which is proved in [11] and which plays a central role in the proof of Theorem 4.

Theorem 7. *The unique up to isometry hyper-ideal circle pattern which realizes the data $(\mathcal{T}, \tilde{\theta}, \Theta)$, and hence the original data $(\mathcal{C}, \theta, \Theta)$, is represented by a unique minimum, located inside \mathcal{TE} , of the continuously differentiable convex functional*

$$\mathbf{U}_{\theta, \Theta} : \mathbb{R}^{E_T^0} \times \mathbb{R}_+^{E_T^1 \cup V^1} \rightarrow \mathbb{R}$$

defined by formula (9). The existence of the minimum is guaranteed if and only if the angle data (θ, Θ) satisfies conditions 1 to 4 of Theorem 4.

With all these tools at hand, we can move on the next section in which an algorithmic recipe for discrete uniformization is outlined.

12 Algorithm for discrete uniformization via hyper-ideal circle patterns

Start with a closed surface S of genus one or greater, together with one of the following two geometric data on it:

Type 1 data.

- Either a hyperbolic or Euclidean cone metric d on S with cone singularities $\text{sing}(d)$ whose cone-angles are greater than 2π .
- A finite set of points V on S such that $\text{sing}(d) = V_1 \subseteq V$ and $V_0 = V \setminus V_1$.

Type 2 data.

- A finite topological branch cover $p : S \rightarrow \hat{\mathbb{C}}$ with ramification points $V_1 \subset S$ and branch points $V_1(\hat{\mathbb{C}}) \subset \hat{\mathbb{C}}$.
- A finite set of points $V_{\hat{\mathbb{C}}} = V_0(\hat{\mathbb{C}}) \cup V_1(\hat{\mathbb{C}})$ on $\hat{\mathbb{C}}$ where $V_0(\hat{\mathbb{C}}) \cap V_1(\hat{\mathbb{C}}) = \emptyset$.
- A finite set $V = p^{-1}(V_{\hat{\mathbb{C}}})$ on S with $V_0 = V \setminus V_1$.

Step 1. Generate the Delaunay circle pattern either on S, d with respect to V if given Type 1 data, or on $\hat{\mathbb{C}}$ with respect to $V_{\hat{\mathbb{C}}}$ if given Type 2 data.

Step 2. In the case of Type 1 data, the Delaunay circle pattern from Step 1 gives rise to a combinatorial cell complex $\mathcal{C} = (V, E, F)$ and an angle assignment $\theta : E \rightarrow (0, \pi)$ of intersection angles between adjacent Delaunay circles.

In the case of Type 2 data, the Delaunay circle pattern from Step 1 gives rise to a combinatorial cell complex $\mathcal{C}_{\hat{\mathbb{C}}} = (V_{\hat{\mathbb{C}}}, E_{\hat{\mathbb{C}}}, F_{\hat{\mathbb{C}}})$ and an angle assignment $\hat{\theta} : E_{\hat{\mathbb{C}}} \rightarrow (0, \pi)$ of intersection angles between adjacent Delaunay circles.

Step 3. Only in the case of Type 2 data, lift the complex $\mathcal{C}_{\hat{\mathbb{C}}}$ to a cell complex $\mathcal{C} = (V, E, F)$ on S via the branch covering map p . Thus, $p^{-1}(V_{\hat{\mathbb{C}}}) = V$, $p^{-1}(E_{\hat{\mathbb{C}}}) = E$, $p^{-1}(F_{\hat{\mathbb{C}}}) = F$. Furthermore, define the lifted angle assignment $\theta : E \rightarrow (0, \pi)$ as $\theta_{ij} = \hat{\theta}_{p(ij)}$ for all $ij \in E$.

Step 4. Subtriangulate \mathcal{C} and obtain the combinatorial triangulation $\mathcal{T} = (V, E_T, F_T)$ by adding a maximal number of diagonals with non-intersecting interiors in each non-triangular face of \mathcal{C} (see Section 10). Define E_π as the set of all added diagonals, also called redundant edges in Section 10. Thus, $E_T = E \cup E_\pi$.

Step 5. Form the set E_T^1 of all edges from E_T both of whose endpoints are vertexes from V_1 . Let $E_T^0 = E_T \setminus E_T^1$.

Step 6. Form the functional

$$\mathbf{U}_\theta(a, b) = \sum_{\Delta \in F_T} \mathbf{U}_\Delta(a, b) - \sum_{ij \in E} \theta_{ij} a_{ij} - \pi \sum_{ij \in E_\pi} a_{ij} - 2\pi \sum_{k \in V} b_k, \quad (10)$$

for all (a, b) from the convex set $\mathbb{R}^{E_T^0} \times \mathbb{R}_+^{E_T^1 \cup V_1}$. The functions \mathbf{U}_Δ are defined by formula (8) relying on the real analytic expressions for the angles α_{ij}^Δ and β_k^Δ in terms of (a, b) , given in the second half of Section 10. Recall that the angle functions can be continuously extended as linear functions outside the domains of their initial (real analytic) definition (see Section 11). Furthermore, also in Section 11, it was commented that there exist analytic formulas for the volume function \mathbf{V} in terms of dihedral angles $(\alpha^\Delta, \beta^\Delta)$.

Step 7. As explained in Section 11, the functional \mathbf{U}_θ is convex and continuously differentiable on $\mathbb{R}^{E_T^0} \times \mathbb{R}_+^{E_T^1 \cup V_1}$, and locally strictly convex and real-analytic on its open subdomain \mathcal{TE} . Find the unique minimum $(a^*, b^*) \in \mathcal{TE}$ of \mathbf{U}_θ whose existence is guaranteed by Theorems 1 and 2.

Step 8. Compute the edge-lengths and vertex radii $(l^*, r^*) = \Psi(a^*, b^*)$ using formulas (3) to (6).

Step 9. Following the combinatorics of \mathcal{T} , lay out in \mathbb{H}^2 the hyperbolic triangles determined by the edge-lengths $l^* : E_T \rightarrow \mathbb{R}_+$. Thus, a fundamental domain of a Fuchsian group [10, 33, 34] is obtained and if one computes the generators of this Fuchsian group, a discrete equivalent of the classical uniformization theorem is obtained. If desired, erase the redundant edges E_π to represent accurately the geodesic realization of the complex \mathcal{C} in \mathbb{H}^2 .

Step 10 (Optional). One could also draw the resulting hyper-ideal circle pattern, by first drawing all vertex circles given by $r^* : V_1 \rightarrow \mathbb{R}_+$. Then, the presence of the vertex circles uniquely determines the orthogonal face-circles.

In Section 14 we present some results from the computer implementation of the algorithm.

13 Realization of hyper-ideal circle patterns on the Riemann sphere

13.1 Constructions

We start with a topological gluing construction. Let $\mathcal{C} = (V, E, F)$ be a strongly regular complex on the two-sphere S^2 . Fix $k_\infty \in V$ and define $\bar{D}(k_\infty) = \cup F_{k_\infty}$ to be the union of all closed faces of \mathcal{C} attached to k_∞ . By strong regularity, $\bar{D}(k_\infty)$ is a closed topological disk embedded in S^2 . Let $\partial\bar{D}(k_\infty) = \sigma$ be its boundary, which is composed of edges of \mathcal{C} , and let $D(k_\infty)$ be its open interior. Next, take two copies of (S^2, \mathcal{C}) and remove $D(k_\infty)$ from both of them. Then glue the two copies together along the two copies of boundary σ , identifying pairs of twin edges. We obtain the connected sum $S^2 \#_\sigma S^2 \cong S^2$ together with a strongly regular complex $\mathcal{C}_\sigma = (V_\sigma, E_\sigma, F_\sigma)$ on it. Notice that \mathcal{C}_σ has a topological symmetry, which is an involution fixing point-wise the simple closed loop σ .

Now, assume that our strongly regular complex \mathcal{C} on S^2 comes equipped with an angle assignment $\theta : E \rightarrow (0, \pi)$ which satisfies the conditions of Bao and Bonahon's Theorem 6. Then there exists a unique, up to $\mathbb{P}SL(2, \mathbb{C})$ automorphism, hyper-ideal circle pattern on $\hat{\mathbb{C}}$ which realizes the combinatorial angle data (\mathcal{C}, θ) . Let us assume that the pattern has at least one true vertex circle, i.e. $V = V_0 \sqcup V_1$ with $V_1 \neq \emptyset$. As already explained in the Discussion on hyper-ideal circle patterns on $\hat{\mathbb{C}}$ from Section 2, we can choose a point k_∞ from the interior of the vertex circle and stereographically project the pattern on \mathbb{C} so that the vertex circle becomes the ideal boundary of the hyperbolic plane and the rest of the pattern becomes a hyper-ideal circle pattern on a convex geodesic polygon P in \mathbb{H}^2 with combinatorics $\mathcal{C} \setminus D(k_\infty)$. The interior angle at a vertex i of ∂P

is equal to θ_{ik_∞} and the angle between a geodesic edge ij of ∂P and the face circle of the decorated polygon from the pattern attached to ij is θ_{ij} . Take two copies of P and glue them together along ∂P identifying isometrically the pairs of twin edges. The result is a sphere S^2_P together with a hyperbolic metric with cone singularities at the vertices that were once boundary vertices of P . Moreover, there is a (generalized) hyper-ideal circle pattern on S^2_P which is symmetric with respect to an isometric involution which fixes point-wise the former boundary ∂P . Inside each copy of P , the intersection angles between adjacent face circles are equal to the intersection angles from the original circle pattern on $\hat{\mathbb{C}}$. Hence, for a non-boundary edge ij of a decorated polygon in one of the two copies of P the intersection angle is $\tilde{\theta}_{ij} = \theta_{ij}$. The angle between adjacent face circles at a former boundary edge $ij \subset \partial P$ is $\tilde{\theta}_{ij} = 2\theta_{ij}$. For a vertex i inside a copy of P the cone angle is $\Theta_i = 2\pi$ while for a vertex at the former boundary ∂P the cone angle is $\Theta_i = 2\theta_{ik_\infty}$. Observe that by construction, the circle pattern on S^2_P has combinatorics \mathcal{C}_σ . Therefore, it realizes the combinatorial angle data $(\mathcal{C}_\sigma, \tilde{\theta}, \Theta)$. Furthermore, it is unique, up to isometry, due to its isometric involutive symmetry and the uniqueness of its two components P guaranteed by Bao and Bonahon's Theorem. Later we will see that uniqueness also follows from a variational principle. With these constructions and notations at hand, we are ready to proceed to the algorithm.

13.2 Algorithm for realization of hyper-ideal circle patterns on $\hat{\mathbb{C}}$

Start with a topological sphere S^2 together with the following data on it:

- A finite topological branch cover $p : S^2 \rightarrow \hat{\mathbb{C}}$ with ramification points $V_1 \subset S^2$ and branch points $V_1(\hat{\mathbb{C}}) \subset \hat{\mathbb{C}}$.
- A finite set of points $V_{\hat{\mathbb{C}}} = V_0(\hat{\mathbb{C}}) \cup V_1(\hat{\mathbb{C}})$ on $\hat{\mathbb{C}}$ where $V_0(\hat{\mathbb{C}}) \cap V_1(\hat{\mathbb{C}}) = \emptyset$.
- A finite set $V = p^{-1}(V_{\hat{\mathbb{C}}})$ on S^2 with $V_0 = V \setminus V_1$.

Step 1. Generate the Delaunay circle pattern on $\hat{\mathbb{C}}$ with respect to $V_{\hat{\mathbb{C}}}$.

Step 2. The Delaunay circle pattern from Step 1 gives rise to a combinatorial cell complex $\mathcal{C}_{\hat{\mathbb{C}}} = (V_{\hat{\mathbb{C}}}, E_{\hat{\mathbb{C}}}, F_{\hat{\mathbb{C}}})$ and an angle assignment $\hat{\theta} : E_{\hat{\mathbb{C}}} \rightarrow (0, \pi)$ of intersection angles between pairs of adjacent Delaunay circles.

Step 3. Lift the complex $\mathcal{C}_{\hat{\mathbb{C}}}$ to a cell complex $\mathcal{C} = (V, E, F)$ on S^2 via the branch covering map p . Thus, $p^{-1}(V_{\hat{\mathbb{C}}}) = V$, $p^{-1}(E_{\hat{\mathbb{C}}}) = E$, $p^{-1}(F_{\hat{\mathbb{C}}}) = F$. Furthermore, define the lifted angle assignment $\theta : E \rightarrow (0, \pi)$ as $\theta_{ij} = \hat{\theta}_{p(ij)}$ for all $ij \in E$. As a result of this, there is a strongly regular complex \mathcal{C} on S^2 together with angle assignment $\theta : E \rightarrow (0, \pi)$.

Remark: All steps from here on are independent of what the origin of the data (\mathcal{C}, θ) is, as long as it satisfies the conditions of Theorem 6.

Step 4. Take $k_\infty \in V_1$, remove the open disc $D(k_\infty)$ from \mathcal{C} , as described above, and form the symmetric (connected sum) cell complex $\mathcal{C}_\sigma = (V_\sigma, E_\sigma, F_\sigma)$ on the connected

sum $S^2 \#_{\sigma} S^2 \cong S^2$ over the boundary $\sigma = D(k_{\infty})$. The vertex set is naturally split into $V_{\sigma} = V_{\sigma,1} \sqcup V_{\sigma,0}$ inherited from the splitting of V .

Step 5. For an edge $ij \in E_{\sigma}$, if ij does not lie entirely on σ , then define $\tilde{\theta}_{ij} = \theta_{ij}$. If ij lies on σ , then define $\theta_{ij} = 2\theta_{ij}$. In both cases, ij is also interpreted as a former edge of $\mathcal{C} \setminus D(k_{\infty})$.

Step 6. For a vertex $i \in V_{\sigma}$, if i is not on σ , then define $\Theta_i = 2\pi$. If $i \in \sigma$, then define $\Theta_i = 2\theta_{ik_{\infty}}$. In both cases, i is also interpreted as a former vertex of $\mathcal{C} \setminus D(k_{\infty})$.

Step 7. Subtriangulate $\mathcal{C} \setminus D(k_{\infty})$ and obtain the combinatorial triangulation $\mathcal{T}_{\sigma} = (V_{\sigma}, E_T, F_T)$ by adding a maximal number of diagonals with non-intersecting interiors in each non-triangular face of \mathcal{C}_{σ} (see Section 10). The triangulation can be constructed so that it inherits the involutive symmetry of \mathcal{C}_{σ} . Indeed, one can first subtriangulate $\mathcal{C} \setminus D(k_{\infty})$ and then glue together two identical copies along σ . Define E_{π} as the set of all added diagonals, also called redundant edges in Section 10. Thus, $E_T = E_{\sigma} \sqcup E_{\pi}$.

Step 8. Form the set E_T^1 of all edges from E_T both of whose endpoints are vertexes from $V_{\sigma,1}$. Let $E_T^0 = E_T \setminus E_T^1$.

Step 9. Form the functional

$$\mathbf{U}_{\tilde{\theta}, \Theta}(a, b) = \sum_{\Delta \in F_T} \mathbf{U}_{\Delta}(a, b) - \sum_{ij \in E_{\sigma}} \tilde{\theta}_{ij} a_{ij} - \pi \sum_{ij \in E_{\pi}} a_{ij} - \sum_{k \in V} \Theta_k b_k, \quad (11)$$

for all (a, b) from the convex set $\mathbb{R}^{E_T^0} \times \mathbb{R}_+^{E_T^1 \cup V_{\sigma,1}}$. The functions \mathbf{U}_{Δ} are defined by formula (8) relying on the real analytic expressions for the angles α_{ij}^{Δ} and β_k^{Δ} in terms of (a, b) , given in the second half of Section 10. Recall that the angle functions can be continuously extended as linear functions outside the domains of their initial (real analytic) definition (see Section 11). Furthermore, also in Section 11, it was commented that there exist analytic formulas for the volume function \mathbf{V} in terms of dihedral angles $(\alpha^{\Delta}, \beta^{\Delta})$.

Step 10. As explained in Section 11, the functional $\mathbf{U}_{\tilde{\theta}, \Theta}$ is convex and continuously differentiable on $\mathbb{R}^{E_T^0} \times \mathbb{R}_+^{E_T^1 \cup V_{\sigma,1}}$, and locally strictly convex and real-analytic on its open subdomain \mathcal{TE} . Theorem 6 guarantees the existence of a critical point in \mathcal{TE} . By strict convexity, the critical point is unique (and it is a minimum). Hence, this is another justification for the uniqueness of the hyper-ideal circle pattern we are looking to construct. Find the unique minimum $(a^*, b^*) \in \mathcal{TE}$ of $\mathbf{U}_{\tilde{\theta}, \Theta}$ using convex optimization.

Step 11. Compute the edge-lengths and vertex radii $(l^*, r^*) = \Psi(a^*, b^*)$ using formulas (3) to (6).

Step 12. By uniqueness and symmetry of the data, the pattern represented by (l^*, r^*) has an isometric involution and thus it splits into two identical hyper-ideal circle patterns with convex geodesic boundary. Each of them has a smooth hyperbolic metric, so it is realizable in \mathbb{H}^2 . Following the combinatorics of \mathcal{T}_{σ} only on one side of the simple loop σ , lay out in \mathbb{H}^2 the hyperbolic triangles determined by the edge-lengths $l^* : E_T \rightarrow \mathbb{R}_+$.

Thus, a hyper-ideal circle patterns with convex geodesic boundary is obtained. If desired, erase all redundant edges from E_π to represent accurately the geodesic realization of the complex $\mathcal{C} \setminus D(k_\infty)$ in \mathbb{H}^2 .

Step 14. Using the assignment of vertex radii $r^* : V_\sigma \rightarrow \mathbb{R}_+$, first draw the vertex circles centered at the vertices of the realized geodesic complex in \mathbb{H}^2 . Then, the presence of the vertex circles uniquely determines the orthogonal face-circles. By adding the circle at infinity $\partial\mathbb{H}^2$ to the collection of vertex circles, and by adding the circles representing the hyperbolic geodesics of the boundary of the complex to the collection of face circles, we obtain the hyper-ideal circle pattern on $\hat{\mathbb{C}}$ which realizes the combinatorial angle data (\mathcal{C}, θ) .

Remark: The same procedure can be carried out by choosing a vertex $k_\infty \in V_0$ instead of V_1 . Since in this case the vertex circle at k_∞ is a point, the cone-metric on the doubled sphere will be Euclidean and uniquely defined up to scaling.

14 Numerical examples

We use the numerical optimization package TAO [4] to perform the minimization of the functional for certain examples. We use the BLMVM method which is a quasi-Newton method that approximates the Hessian matrix using a fixed number of explicit gradient evaluations. We configure the bounded domain—see Steps 5 and 6 of Section 12—using the API for bounded minimization in the TAO application.

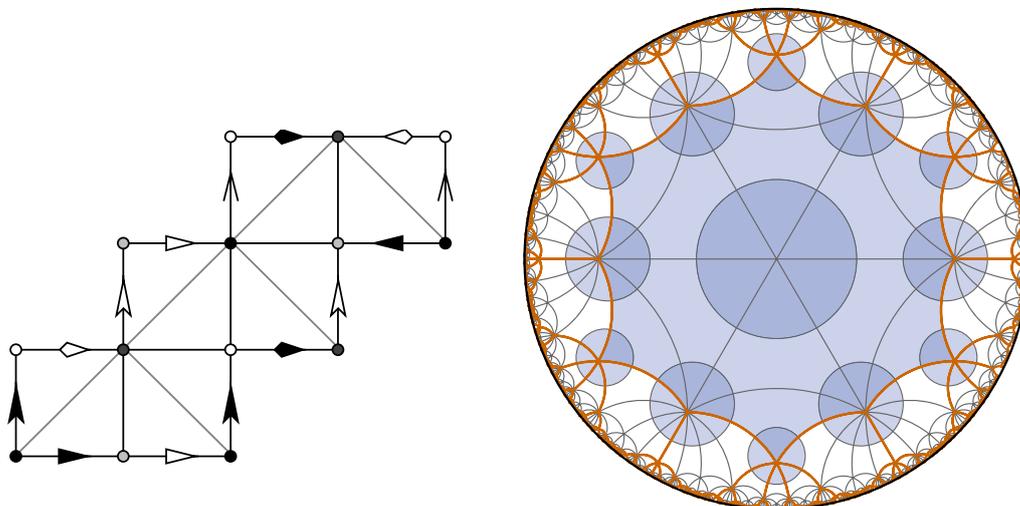


Figure 7: Square-tiled representation of the Riemann surface of Lawson's minimal surface (left). Universal cover, Fuchsian uniformization, fundamental domain and vertex circles (right), see Subsection 14.1.

14.1 Square-tiled Reimann surfaces

Example 1. As a first example we use a square-tiled surface which is conformally equivalent to the Riemann surfaces associated to the Lawson's genus 2 minimal surface in \mathbb{S}^3 [17], see Figure 7, left. It has four vertices (white, light-grey, dark-grey, black), 18 edges, and 6 faces. The 6 diagonal edges are redundant edges in the sense that the face circle intersection angle is π at these edges, i.e. the two circles on both sides of a redundant edge actually coincide. As a result of this, in the hyperbolic realization Figure 7, right, all pairs of decorated triangles sharing redundant edges are merged to form decorated quadrilaterals. Hence the conformal structure on the grey diagonals is given by $\theta_{\text{grey}} = \pi$ and on the black edges by $\theta_{\text{black}} = \frac{\pi}{2}$. Boundary edges are identified as indicated by the arrows. In this case, all vertices have cone angles greater than 2π so $V = V_1$ while $V_0 = \emptyset$. All of them become vertex circles after uniformization. On Figure 7, right, the vertex circles are depicted, while the face circles are not for the sake of better clarity of the picture.

In this example all variables a_i, b_i are strictly positive. We minimize the functional using the options described above. The solver converges after 17 iterations at a solution with a gradient norm less than 10^{-10} .

We choose a rotationally symmetric fundamental domain where each square is incident to a central vertex, see Figure 7, right.

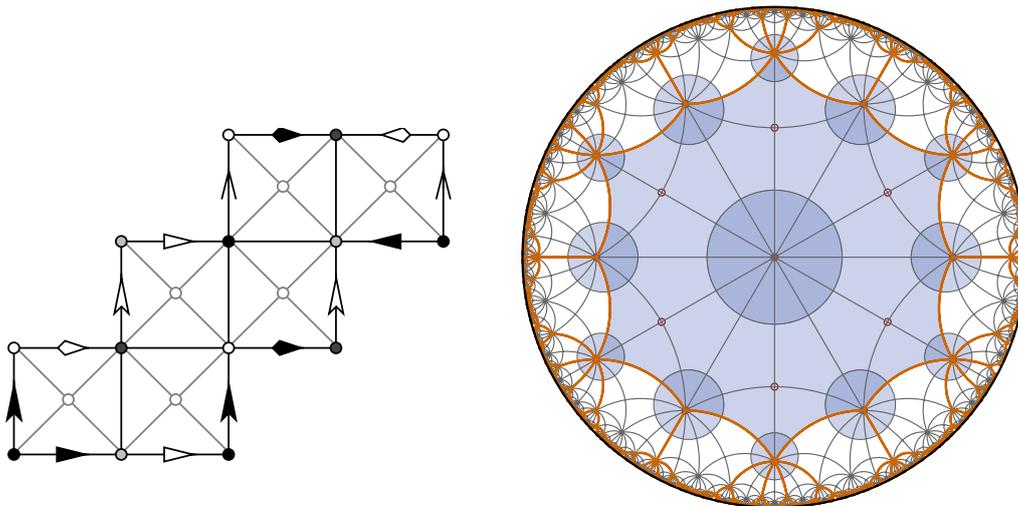


Figure 8: Square-tiled representation of the Riemann surface of Lawson's minimal surface including ideal vertices in the centers of the squares (left). Universal cover, Fuchsian uniformization, fundamental domain and vertex circles (right). See Subsection 14.1.

Example 2. For the second example, we add six new vertices to the square-tiled surface from Figure 7, left. Namely, in the center of each square we insert a vertex and triangulate

the quadrilaterals as shown on Figure 8, left. Note that these vertices correspond to the ramification points of the Riemann surface when represented as an algebraic curve in \mathbb{C}^2 , see Subsection 14.2.

The conformal structure is given by $\theta_{\text{black}} = \pi$ on black edges and $\theta_{\text{grey}} = \frac{\pi}{2}$ on grey edges. The black vertices are redundant, so we obtain decorated quadrilaterals after uniformization, Figure 8, right. As the new white vertices are flat to begin with we exclude the corresponding variables from the functional, i.e., $b_{\text{white}} \equiv 0$. In the corresponding hyper-ideal tetrahedron this corresponds to an ideal vertex of the tetrahedron. At the same time when introducing ideal vertices we change the domain of optimization for certain variables, i.e., for edges incident with at least one ideal vertex we have $a_{ij} \in \mathbb{R}$. In this case, the white vertices form the set vertex set V_0 while the rest of the vertices form the set V_1 . As one can see on Figure 7, right, all vertices from V_1 become vertex circles, while the ones from V_0 stay points (circles of radius zero).

Just like in the preceding example, the solver converges after 26 iterations to a solution with a gradient norm less than 10^{-7} .

We choose the same fundamental domain as in the previous example and calculate the generators of the Fuchsian uniformization group and the corresponding universal cover, see Figure 8, right.

14.2 Hyperelliptic Riemann surfaces

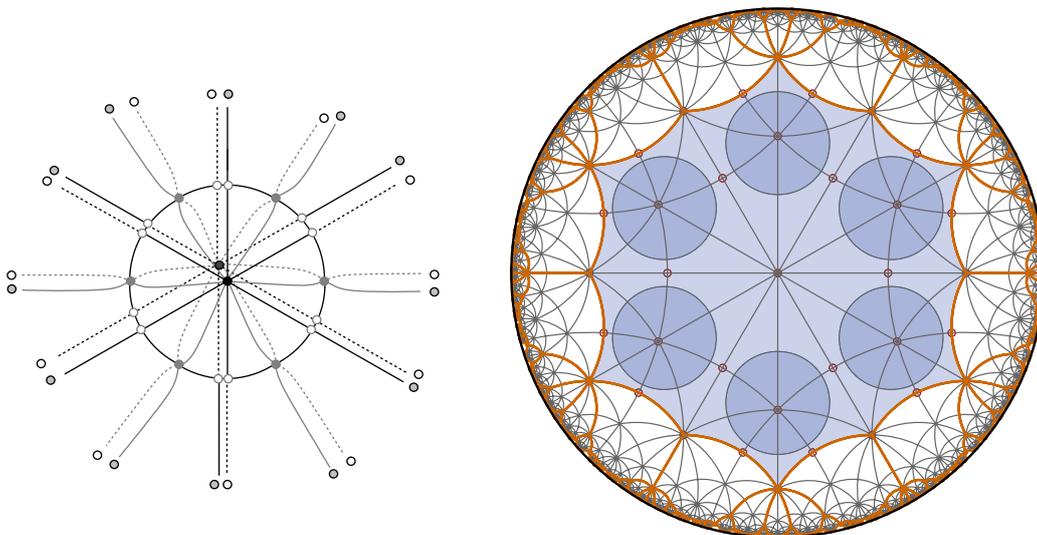


Figure 9: The Riemann surface of Lawson's genus 2 minimal surface in \mathbb{S}^3 represented as doubly covered polyhedral surface over $\hat{\mathbb{C}}$ (left). Universal cover, Fuchsian uniformization, fundamental domain and vertex circles (right). See Section 14.2.

Example 3. In this example we discretely uniformize the complex algebraic curve $\mu^2 = \lambda^6 - 1$. The latter is the Riemann surface associated to Lawson's minimal surface in the three sphere [17] and it is represented as a branched cover over $\hat{\mathbb{C}}$. We generate a Delaunay triangulation on $\hat{\mathbb{C}}$ that includes the ramification points of the algebraic curve. Hence the triangulation includes the six roots of unity as vertices. We add the mid-points between these vertices as well as the north and the south poles. Then we lift the corresponding Delaunay triangulation to the algebraic curve creating a two-sheeted cover of $\hat{\mathbb{C}}$ branched around the six vertices at the roots of unity, see Figure 9, left.

The conformal structure is calculated by measuring the intersection angles of the face circumcircles in $\hat{\mathbb{C}} \cong \mathbb{S}^2$. We first construct the Delaunay triangulation on \mathbb{S}^2 using a convex hull algorithm. Then we use a suitable stereographic projection to measure circle intersection angles in the plane.

Using this procedure we end up with six positive variable vertices, i.e., the vertices at branch points of the curve. The edges are all adjacent to at least one ideal vertex. Hence all edge variables are real, $a_{ij} \in \mathbb{R}$.

The solver converges after 16 iterations to an accuracy less than 10^{-8} gradient norm.

The four vertices on the north and south pole correspond to the four vertices of the quadrilaterals on the Riemann surfaces used in the examples from Subsection 14.1. We choose a fundamental domain with the same cuts on the surface to produce the universal cover presented in Figure 9, right.

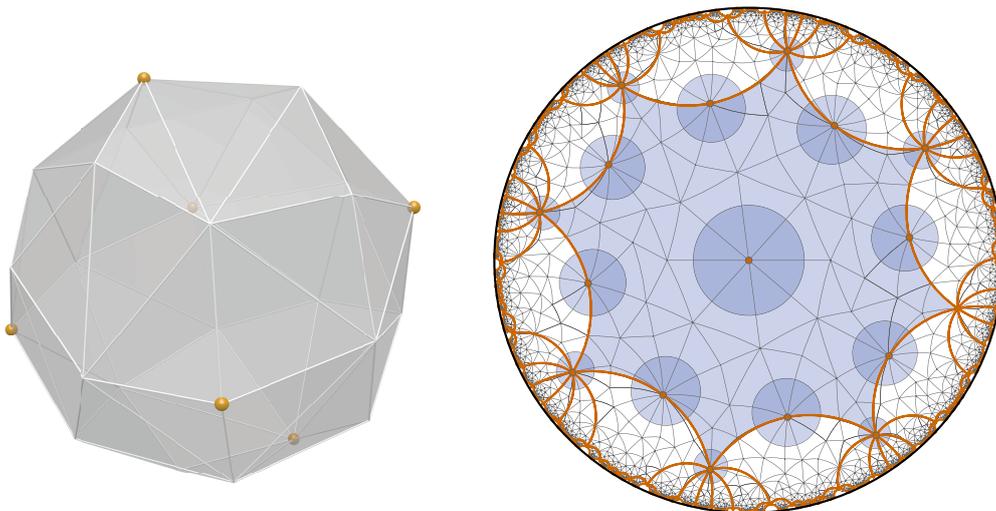


Figure 10: General hyperelliptic surface of genus 2. Two-sheeted cover of a polyhedron with vertices on the sphere (left). Large vertices are the branch points of the corresponding algebraic curve. Universal cover, Fuchsian uniformization, fundamental domain and vertex circles (right). See Section 14.2.

Example 4. In this example we calculate the discrete uniformization of a more general hyperelliptic curve of genus 2. To achieve a better visual representation, we choose the branch data so that the surface admits an approximately regular fundamental domain, i.e., the branch points form approximately a regular octahedron, see Figure 10.

The surface is constructed in a way similar to the first example of this Subsection. The triangulation includes the branch vertices $V_1(\hat{\mathbb{C}})$ and additional points $V_0(\hat{\mathbb{C}})$ chosen randomly on the sphere. Just like in the previous example, no edge connects two branch points, hence all edge variables are real.

The conformal structure is calculated by stereographic projection. The solver converges after 16 iteration with an accuracy of less than 10^{-8} .

We choose a fundamental domain that is almost a regular polygon where opposite sides are identified.

14.3 Discussion

We expect to increase the accuracy and speed of the solver if we implement the Hessian matrix explicitly and use Newton's method as implemented in TAO. Furthermore, by using a doubling (pillow) construction, one can apply the so far developed methods and algorithms to the case of a topological sphere S , in the spirit of Theorem 3. In other words, given the combinatorial angle data, we can construct hyper-ideal circle patterns corresponding to the hyper-ideal polyhedra form Bao and Bonahon's Theorem 6.

Acknowledgements

This research is supported by DFG (Deutsche Forschungsgemeinschaft) in the frame of Sonderforschungsbereich/Transregio 109 "Discretization in Geometry and Dynamics".

References

- [1] Bao, X., Bonahon, F.: Hyperideal polyhedra in hyperbolic 3-space, Bull. Soc. Math. Fr. 130(3), 457–491, (2002)
- [2] Beardon, A.F., Stephenson, K.: The uniformization theorem for circle packings, Indiana University Math. J., 39(4), 1383–1425,(1990)
- [3] Benedetti, R., Petronio, C.: Lectures on Hyperbolic Geometry. Universitext, Springer-Verlag, Berlin, (1992)
- [4] Benson, S., McInnes, L.C., Moré, J., Munson, T., Sarich, J.: TAO user manual (version 1.9), <http://www.mcs.anl.gov/tao>, (2007)
- [5] Bobenko, A.I., Pinkall, U., Springborn, B.: Discrete conformal maps and ideal hyperbolic polyhedra, Geom. Topol. 19(4), 2155–2215, (2015)

- [6] Bobenko, A.I., Springborn, B.A.: Variational principles for circle patterns and Koebe's theorem, *Trans. Amer. Math. Soc.* 356(2), 659-689, (2004)
- [7] Bobenko, A.I., Suris, Y.B.: *Discrete Differential Geometry: Integrable Structure*, Graduate Studies in Mathematics, Vol. 98, AMS, (2008)
- [8] Bowers, P.L.; Hurdal, M.K.: Planar conformal mappings of piecewise flat surfaces, *Visualization and Mathematics III*, 334, *Math. Vis.*, Springer-Verlag, Berlin, 3–34, (2003)
- [9] Bowers, P.L., Stephenson, K.: Uniformizing dessins and Belyi (maps via circle packing), *Mem. Amer. Math. Soc.* 170, no. 805, (2004)
- [10] Buser, P.: *Geometry and Spectra of Compact Riemann Surfaces*, *Prog. in Math.* 106, Birkhäuser, Boston, (1992)
- [11] Dimitrov, N.: Hyper-ideal circle patterns with cone singularities, *Results. Math.*, DOI: 10.1007/s00025-015-0453-3, (2015), extended version arXiv: math.MG/14066741
- [12] Edelsbrunner, H.: *Geometry and topology for mesh generation*, Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, (2001)
- [13] Grünbaum, B.: *Convex Polytopes*. Volume 221 of Graduate Texts in Mathematics, Springer-Verlag, Berlin, second edition, (2003)
- [14] Gu, X., Luo, F., Sun, J., Wu, T.: A discrete uniformization theorem for polyhedral surfaces, preprint arXiv:1309.4175
- [15] Gu, X., Guo, R., Luo, F., Sun, J., Wu, T.: A discrete uniformization theorem for polyhedral surfaces II, preprint arXiv:1401.4594
- [16] Kharevych, L., Springborn, B., Schröder, P.: Discrete conformal maps via circle patterns. *ACM Transactions on Graphics* 25(2), 412–438, (2006)
- [17] Lawson Jr., H. B.: Complete Minimal Surfaces in S^3 , *Annals of Math.* 92(3), 335–374, (1970)
- [18] Luo, F.: Rigidity of polyhedral surfaces, III, *Geom. Top.* 15(4), 2299-2319, (2011)
- [19] Luo, F.: Combinatorial Yamabe flow on surfaces, *Commun. Contemp. Math.* 6(5), 765–780, (2004)
- [20] Ma, J., Schlenker, J-M.: Non-rigidity of spherical inversive distance circle packings, *Discrete Comput. Geom.* 47(3), 610–617, (2012)
- [21] Rivin, I.: Euclidean structures of simplicial surfaces and hyperbolic volume, *Ann. of Math.* 139, 553-580, (1994)

- [22] Rivin, I.: A characterization of ideal polyhedra in hyperbolic 3-space, *Ann. of Math.* 2, 143(1), 51–70, (1996)
- [23] Rodin, B., Sullivan, D.: The convergence of circle packing to Riemann mapping, *J. Differ. Geom.* 26(2), 349–360, (1987)
- [24] Rousset, M.: Sur la rigidité de polyèdres hyperboliques en dimension 3: cas de volume fini, cas hyperidéale, cas fuchsien *Bull. Soc. Math. France* 132(2), 233-261, (2004)
- [25] Schlenker, J.-M.: Hyperideal circle patterns, *Math. Res. Lett.* 12(1), 85–112, (2005)
- [26] Schlenker, J.-M.: Circle patterns on singular surfaces, *Discrete Comput. Geom.* 40(1), 47-102, (2008)
- [27] Schlenker, J.-M.: Rigidity criterion for non-convex polyhedra, *Discrete Comput. Geom.* 33(2), 207–221, (2005)
- [28] Schramm, O.: Circle patterns with the combinatorics of the square grid, *Duke Math. J.* 86(2), 347–389, (1997).
- [29] Springborn, B., Schröder, P., Pinkall, U.: Conformal equivalence of triangular meshes, *ACM Trans. Graph.* 27(3), article 77, 11 pages, (2008)
- [30] Stephenson, K.: *Introduction to Circle Packing. The Theory of Discrete Analytic Functions.* Cambridge University Press, Cambridge, 2005
- [31] Springborn, B.A.: A variational principle for weighted Delaunay triangulations and hyperideal polyhedra, *J. Differ. Geom.* 78(2), 333–367, (2008)
- [32] Springborn, B.A.: A unique representation of polyhedral types. Centering via Möbius transformations, *Math. Z.* 249(3), 513–517, (2005)
- [33] Thurston, W.P.: *The Geometry and Topology of Three-Manifolds*, Electronic library of MSRI, available at <http://library.msri.org/books/gt3m/>, (2002)
- [34] Thurston, W.P.: *Three-dimensional geometry and topology*, Edited by Silvio Levy, Princeton Mathematical Series, 35 Princeton University Press, Princeton, NJ, (1997)
- [35] Thurston, W.P.: The finite Riemann mapping theorem, Invited address, International Symposium in Celebration of the Proof of the Bieberbach Conjecture, Purdue University, (1985)
- [36] Ushijima, A.: A volume formula for generalized hyperbolic tetrahedra, *Non-Euclidean Geom.* 581, 249–265, (2006)
- [37] Ziegler, G.M.: *Lectures on Polytopes.* Springer-Verlag, Berlin, (1995)