

Discrete conformal maps: Boundary value problems, circle domains, Fuchsian and Schottky uniformization

Alexander I. Bobenko, Stefan Sechelmann, Boris Springborn

1 Introduction

The notion of discrete conformal equivalence for polyhedral surfaces is based on a simple definition: Two polyhedral surfaces are discretely conformally equivalent if the edge lengths are related by scale factors assigned to the vertices. It leads to a surprisingly rich theory [17, 4, 9, 10]. In this article, we investigate different aspects of this theory.

We extend the notion of discrete conformal equivalence from triangulated surfaces to polyhedral surfaces with faces that are inscribed in circles. The basic definitions and their immediate consequences are discussed in Section 2.

In Section 3, we generalize a variational principle for discretely conformally equivalent triangulations [4] to the polyhedral setting. This variational principle is the main tool for all our numerical calculations. It is also the basis for our uniqueness proof for discrete conformal mapping problems (Theorem 3.9).

Section 4 is concerned with the special case of quadrilateral meshes. We discuss the emergence of orthogonal circle patterns, a peculiar necessary condition for the existence of solutions for boundary angle problems, and we extend the method of constructing discrete Riemann maps from triangulations to quadrangulations.

In Section 5, we briefly discuss discrete conformal maps from multiply connected domains to circle domains, and special cases in which we can map to slit domains.

Section 6 deals with conformal mappings onto the sphere. We generalize the method for triangulations to quadrangulations, and we explain how the spherical version of the variational principle can in some cases be used for numerical calculations although the corresponding functional is not convex.

Section 7 is concerned with the uniformization of tori, i.e., the representation of Riemann surfaces as a quotient space of the complex plane modulo a period lattice. We consider Riemann surfaces represented as immersed surfaces in \mathbb{R}^3 , and as elliptic curves. We conduct numerical experiments to test the conjectured convergence of discrete conformal maps. We consider the difference between the true modulus

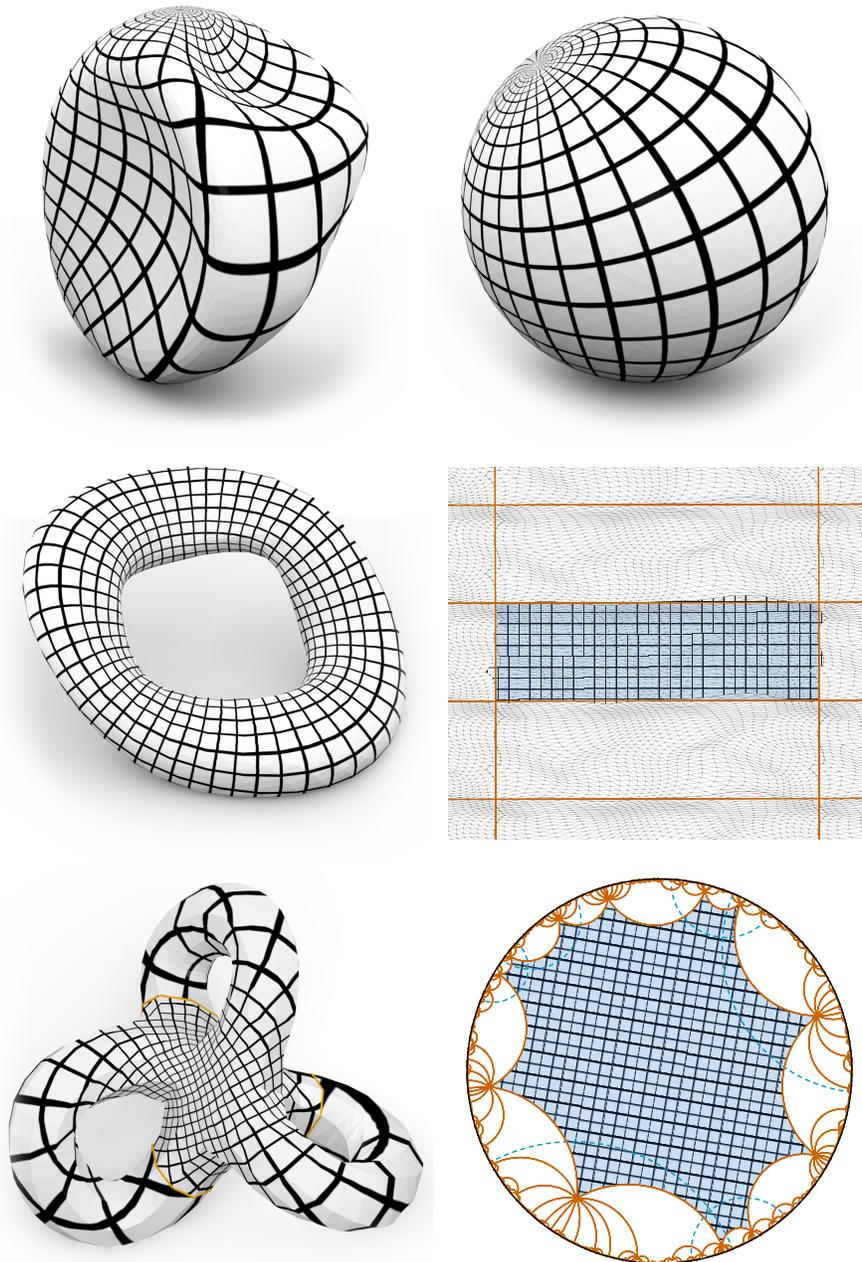


Fig. 1 Uniformization of compact Riemann surfaces. The uniformization of spheres is treated in Section 6. Tori are covered in Section 7, and Section 8 is concerned with surfaces of higher genus.

of an elliptic curve (which can be calculated using hypergeometric functions) and the modulus determined by discrete uniformization, and we estimate the asymptotic dependence of this error on the number of vertices.

In Section 8, we consider the Fuchsian uniformization of Riemann surfaces represented in different forms. We consider immersed surfaces in \mathbb{R}^3 (and S^3), hyperelliptic curves, and Riemann surfaces represented as a quotient of $\hat{\mathbb{C}}$ modulo a classical Schottky group. That is, we convert from Schottky uniformization to Fuchsian uniformization. The Section ends with two extended examples demonstrating, among other things, a remarkable geometric characterization of hyperelliptic surfaces due to Schmutz Schaller.

2 Discrete conformal equivalence of cyclic polyhedral surfaces

2.1 Cyclic polyhedral surfaces

A *euclidean polyhedral surface* is a surface obtained from gluing euclidean polygons along their edges. (A *surface* is a connected two-dimensional manifold, possibly with boundary.) In other words, a euclidean polyhedral surface is a surface equipped with, first, an intrinsic metric which is flat except at isolated points where it has cone-like singularities, and, second, the structure of a CW complex with geodesic edges. The set of vertices contains all cone-like singularities. If the surface has a boundary, the boundary is polygonal and the set of vertices contains all corners of the boundary.

Hyperbolic polyhedral surfaces and *spherical polyhedral surfaces* are defined analogously. They are glued from polygons in the hyperbolic and elliptic planes, respectively. Their metric is locally hyperbolic or spherical, except at cone-like singularities.

We will only be concerned with polyhedral surfaces whose faces are all cyclic, i.e., inscribed in circles. We call them *cyclic polyhedral surfaces*. More precisely, we require the polygons to be cyclic before they are glued together. It is not required that the circumcircles persist after gluing; they may be disturbed by cone-like singularities. A polygon in the hyperbolic plane is considered cyclic if it is inscribed in a curve of constant curvature. This may be a circle (the locus of points at constant distance from its center), a horocycle, or a curve at constant distance from a geodesic.

A *triangulated surface*, or *triangulation* for short, is a polyhedral surface all of whose faces are triangles. All triangulations are cyclic.

2.2 Notation

We will denote the sets of vertices, edges, and faces of a CW complex Σ by V_Σ , E_Σ , and F_Σ , and we will often omit the subscript when there is no danger of confusion. For notational convenience, we require all CW complexes to be *strongly regular*. This means that we require that faces are not glued to themselves along edges or at vertices, that two faces are not glued together along more than one edge or one vertex, and that edges have distinct end-points and two edges have at most one endpoint in common. This allows us to label edges and faces by their vertices. We will write $ij \in E$ for the edge with vertices $i, j \in V$ and $ijkl \in F$ for the face with vertices $i, j, k, l \in V$. We will always list the vertices of a face in the correct cyclic order, so that for example the face $ijkl$ has edges ij, jk, kl , and li . The only reason for restricting our discussion to strongly regular CW complexes is to be able to use this simple notation. Everything we discuss applies also to general CW complexes.

2.3 Discrete metrics

The *discrete metric* of a euclidean (or hyperbolic or spherical) cyclic polyhedral surface Σ is the function $\ell : E_\Sigma \rightarrow \mathbb{R}_{>0}$ that assigns to each edge $ij \in E_\Sigma$ its length ℓ_{ij} . It satisfies the polygon inequalities (one side is shorter than the sum of the others):

$$\left. \begin{array}{l} -\ell_{i_1 i_2} + \ell_{i_2 i_3} + \dots + \ell_{i_{n-1} i_n} > 0 \\ \ell_{i_1 i_2} - \ell_{i_2 i_3} + \dots + \ell_{i_{n-1} i_n} > 0 \\ \vdots \\ \ell_{i_1 i_2} + \ell_{i_2 i_3} + \dots - \ell_{i_{n-1} i_n} > 0 \end{array} \right\} \text{ for all } i_1 i_2 \dots i_n \in F_\Sigma \quad (1)$$

In the case of spherical polyhedral surfaces, we also require that

$$\ell_{i_1 i_2} + \ell_{i_2 i_3} + \dots + \ell_{i_{n-1} i_n} < 2\pi. \quad (2)$$

The polygon inequalities (1) are necessary and sufficient for the existence of a unique cyclic euclidean polygon and a unique cyclic hyperbolic polygon with the given edge lengths. Together with inequality (2) they are necessary and sufficient for the existence of a unique cyclic spherical polygon. For a new proof of these elementary geometric facts, see [14]. Thus, a discrete metric determines the geometry of a cyclic polyhedral surface:

Proposition and Definition 2.1. *If Σ is a surface with the structure of a CW complex and a function $\ell : E_\Sigma \rightarrow \mathbb{R}_{>0}$ satisfies the polygon inequalities (1), then there is a unique euclidean cyclic polyhedral surface and also a unique hyperbolic cyclic polyhedral surface with CW complex Σ and discrete metric ℓ . If ℓ also satisfies the inequalities (2), then there is a unique spherical cyclic polyhedral surface with CW complex Σ and discrete metric ℓ .*

We will denote the euclidean, hyperbolic, and spherical polyhedral surface with CW complex Σ and discrete metric ℓ by $(\Sigma, \ell)_{euc}$, $(\Sigma, \ell)_{hyp}$, and $(\Sigma, \ell)_{sph}$, respectively.

2.4 Discrete conformal equivalence

We extend the definition of discrete conformal equivalence from triangulations [17, 4] to cyclic polyhedral surfaces in a straightforward way (Definition 2.2). While some aspects of the theory carry over to the more general setting (e.g., Möbius invariance, Proposition 2.5), others do not, like the characterization of discretely conformally equivalent triangulations in terms of length cross-ratios (Section 2.5). We will discuss similar characterizations for polyhedral surfaces with 2-colorable vertices and the particular case of quadrilateral faces in Sections 2.7 and 2.8.

We define discrete conformal equivalence only for polyhedral surfaces that are combinatorially equivalent (see Remark 2.4). Thus, we may assume that the surfaces share the same CW complex Σ equipped with different metrics $\ell, \tilde{\ell}$.

Definition 2.2. *Discrete conformal equivalence* is an equivalence relation on the set of cyclic polyhedral surfaces defined as follows:

- Two euclidean cyclic polyhedral surfaces $(\Sigma, \ell)_{euc}$ and $(\Sigma, \tilde{\ell})_{euc}$ are *discretely conformally equivalent* if there exists a function $u : V_\Sigma \rightarrow \mathbb{R}$ such that

$$\tilde{\ell}_{ij} = e^{\frac{1}{2}(u_i + u_j)} \ell_{ij}. \quad (3)$$

- Two hyperbolic cyclic polyhedral surfaces $(\Sigma, \ell)_{hyp}$ and $(\Sigma, \tilde{\ell})_{hyp}$ are *discretely conformally equivalent* if there exists a function $u : V_\Sigma \rightarrow \mathbb{R}$ such that

$$\sinh\left(\frac{\tilde{\ell}_{ij}}{2}\right) = e^{\frac{1}{2}(u_i + u_j)} \sinh\left(\frac{\ell_{ij}}{2}\right). \quad (4)$$

- Two spherical cyclic polyhedral surfaces $(\Sigma, \ell)_{sph}$ and $(\Sigma, \tilde{\ell})_{sph}$ are *discretely conformally equivalent* if there exists a function $u : V_\Sigma \rightarrow \mathbb{R}$ such that

$$\sin\left(\frac{\tilde{\ell}_{ij}}{2}\right) = e^{\frac{1}{2}(u_i + u_j)} \sin\left(\frac{\ell_{ij}}{2}\right). \quad (5)$$

We will also consider mixed versions:

- A euclidean cyclic polyhedral surface $(\Sigma, \ell)_{euc}$ and a hyperbolic cyclic polyhedral surface $(\Sigma, \tilde{\ell})_{hyp}$ are discretely conformally equivalent if

$$\sinh\left(\frac{\tilde{\ell}_{ij}}{2}\right) = e^{\frac{1}{2}(u_i + u_j)} \ell_{ij}. \quad (6)$$

- A euclidean cyclic polyhedral surface $(\Sigma, \ell)_{euc}$ and a spherical cyclic polyhedral surface $(\Sigma, \tilde{\ell})_{sph}$ are discretely conformally equivalent if

$$\sin\left(\frac{\tilde{\ell}_{ij}}{2}\right) = e^{\frac{1}{2}(u_i+u_j)} \ell_{ij}. \quad (7)$$

- A hyperbolic cyclic polyhedral surface $(\Sigma, \ell)_{hyp}$ and a spherical cyclic polyhedral surface $(\Sigma, \tilde{\ell})_{sph}$ are discretely conformally equivalent if

$$\sin\left(\frac{\tilde{\ell}_{ij}}{2}\right) = e^{\frac{1}{2}(u_i+u_j)} \sinh\left(\frac{\ell_{ij}}{2}\right) \quad (8)$$

Remark 2.3. Note that relation (5) for spherical edge lengths is equivalent to relation (3) for the euclidean lengths of the chords in the ambient \mathbb{R}^3 of the sphere (see Figure 2, left). Likewise, relation (4) for hyperbolic edge lengths is equivalent to (3)

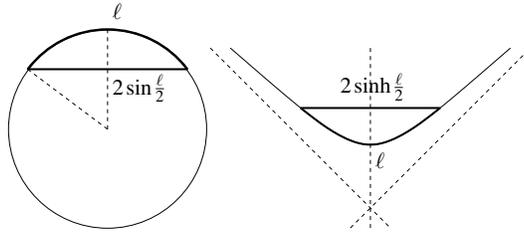


Fig. 2 Spherical and hyperbolic chords.

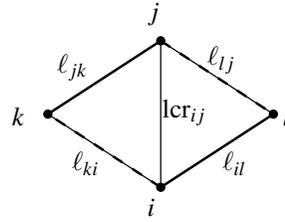


Fig. 3 Length cross-ratio

for the euclidean lengths of the chords in the ambient $\mathbb{R}^{2,1}$ of the hyperboloid model of the hyperbolic plane (see Figure 2, right).

Remark 2.4. For triangulations, the definition of discrete conformal equivalence has been extended to meshes that are not combinatorially equivalent [4, Definition 5.1.4] [9, 10]. It is not clear whether or how the following definitions for cyclic polyhedral surfaces can be extended to combinatorially inequivalent CW complexes.

The discrete conformal class of a cyclic polyhedral surface embedded in n -dimensional euclidean space is invariant under Möbius transformations of the ambient space:

Proposition 2.5 (Möbius invariance). *Suppose P and \tilde{P} are two combinatorially equivalent euclidean cyclic polyhedral surfaces embedded in \mathbb{R}^n (with straight edges and faces), and suppose there is a Möbius transformation of $\mathbb{R}^n \cup \{\infty\}$ that maps the vertices of P to the corresponding vertices of \tilde{P} . Then P and \tilde{P} are discretely conformally equivalent.*

Note that only vertices are related by the Möbius transformation, not edges and faces, which remain straight. The simple proof for the case of triangulations [4] carries over without change.

2.5 Triangulations: Characterization by length cross-ratios

For euclidean triangulations, there is an alternative characterization of conformal equivalence in terms of length cross-ratios [4]. We review the basic facts in this section.

For two adjacent triangles $ijk \in F$ and $jil \in F$ (see Figure 3), the *length cross-ratio* of the common interior edge $ij \in E$ is defined as

$$\text{lcr}_{ij} = \frac{\ell_{il}\ell_{jk}}{\ell_{ij}\ell_{ki}}. \quad (9)$$

(If the two triangles are embedded in the complex plane, this is just the modulus of the complex cross-ratio of the four vertices.) This definition of length cross-ratios implicitly assumes that an orientation has been chosen on the surface. For non-orientable surfaces, the length cross-ratio is well defined on the oriented double cover.

The product of length cross-ratios around an interior vertex $i \in V$ is 1, because all lengths cancel:

$$\prod_{ij \ni i} \text{lcr}_{ij} = 1. \quad (10)$$

Proposition 2.6. *Two euclidean triangulations $(\Sigma, \ell)_{\text{euc}}$ and $(\Sigma, \tilde{\ell})_{\text{euc}}$ are discretely conformally equivalent if and only if for each interior edge $ij \in E_{\Sigma}^{\text{int}}$, the induced length cross-ratios agree.*

Remark 2.7. Analogous statements hold for spherical and hyperbolic triangulations. Equation (9) has to be modified by replacing ℓ with $\sin \frac{\ell}{2}$ or $\sinh \frac{\ell}{2}$, respectively (compare Remark 2.3).

2.6 Triangulations: Reconstructing lengths from length cross-ratios

To deal with Riemann surfaces that are given in terms of Schottky data (Section 8.2) we will need to reconstruct a function $\ell : E_{\Sigma} \rightarrow \mathbb{R}_{>0}$ satisfying (9) from given length cross-ratios. (It is not required that the function ℓ satisfies the triangle inequalities.) To this end, we define auxiliary quantities c_{jk}^i attached to the angles of the triangulation. The value at vertex i of the triangle $ijk \in F$ is defined as

$$c_{jk}^i = \frac{\ell_{jk}}{\ell_{ij}\ell_{ki}}. \quad (11)$$

Then (9) is equivalent to

$$\text{lcr}_{ij} = \frac{c_{jk}^i}{c_{ij}^i}. \quad (12)$$

Now, given a function $\text{lcr} : E^{\text{int}} \rightarrow \mathbb{R}_{>0}$ defined on the set of interior edges E^{int} and satisfying the product condition (10) around interior vertices, one can find parameters c_{jk}^i satisfying (11) by choosing one value at each vertex and then successively multiplying length cross-ratios. The corresponding function ℓ is then determined by

$$\ell_{ij} = \frac{1}{\sqrt{c_{jk}^i c_{ki}^j}} = \frac{1}{\sqrt{c_{lj}^i c_{il}^j}}. \quad (13)$$

2.7 Bipartite graphs: Characterization by length multi-ratios

A different characterization of discrete conformal equivalence in terms of length multi-ratios holds if the 1-skeleton of the polyhedral surface is bipartite, i.e., if the vertices can be colored with two colors so that no two neighboring vertices share the same color.

Proposition 2.8. (i) *If two combinatorially equivalent euclidean cyclic polyhedral surfaces $(\Sigma, \ell)_{\text{euc}}$ and $(\Sigma, \tilde{\ell})_{\text{euc}}$ with discrete metrics ℓ and $\tilde{\ell}$ are discretely conformally equivalent, then the length multi-ratios for even cycles*

$$i_1 i_2, i_2 i_3, \dots, i_{2n} i_1$$

are equal:

$$\frac{\ell_{i_1 i_2} \ell_{i_3 i_4} \cdots \ell_{i_{2n-1} i_{2n}}}{\ell_{i_2 i_3} \ell_{i_4 i_5} \cdots \ell_{i_{2n} i_1}} = \frac{\tilde{\ell}_{i_1 i_2} \tilde{\ell}_{i_3 i_4} \cdots \tilde{\ell}_{i_{2n-1} i_{2n}}}{\tilde{\ell}_{i_2 i_3} \tilde{\ell}_{i_4 i_5} \cdots \tilde{\ell}_{i_{2n} i_1}}. \quad (14)$$

(ii) *If the 1-skeleton of Σ is bipartite, i.e., if all cycles are even, then this condition is also sufficient: If the length multi-ratios are equal for all cycles, then the polyhedral surfaces are discretely conformally equivalent.*

Proof. (i) This is obvious, because all scale factors e^u cancel. (ii) It is easy to see that equations (3) can be solved for the scale factors $e^{u/2}$ if the length multi-ratios are equal. Note that the scale factors are not uniquely determined: they can be multiplied by λ and $1/\lambda$ on the two vertex color classes, respectively. To find a particular solution, one can fix the value of $e^{u/2}$ at one vertex, and find the other values by alternately dividing and multiplying by $\tilde{\ell}/\ell$ along paths. The equality of length multi-ratios implies that the obtained values do not depend on the path. \square

Remark 2.9. If a polyhedral surface is simply connected, then its 1-skeleton is bipartite if and only if all faces are even polygons. If a polyhedral surface is not simply connected, then having even faces is only a necessary condition for being bipartite.

A polyhedral surface with bipartite 1-skeleton has even faces. If a polyhedral surface has even faces and is simply connected, then its 1-skeleton is bipartite, and the face boundaries generate all cycles. Thus, Proposition 2.8 implies the following corollary.

Corollary 2.10. *Two simply connected combinatorially equivalent euclidean cyclic polyhedral surfaces with even faces and with discrete metrics ℓ and $\tilde{\ell}$ are discretely conformally equivalent if and only if the multi-ratio condition (14) holds for every face boundary cycle.*

Remark 2.11. Analogous statements hold for spherical and hyperbolic cyclic polyhedral surfaces. In the multi-ratio condition, one has to replace non-euclidean lengths ℓ with $\sin \frac{\ell}{2}$ or $\sinh \frac{\ell}{2}$, respectively (compare Remark 2.3).

2.8 Quadrangulations: the cross-ratio system on quad-graphs

The case of cyclic quadrilateral faces is somewhat special (and we will return to it in Section 4), because equal length cross-ratio implies equal complex cross-ratio:

Proposition 2.12. *If two euclidean polyhedral surfaces with cyclic quadrilateral faces are discretely conformally equivalent, then corresponding faces $ijkl \in F$ have the same complex cross-ratio (when embedded in the complex plane):*

$$\frac{(z_i - z_j)(z_k - z_l)}{(z_j - z_k)(z_l - z_i)} = \frac{(\tilde{z}_i - \tilde{z}_j)(\tilde{z}_k - \tilde{z}_l)}{(\tilde{z}_j - \tilde{z}_k)(\tilde{z}_l - \tilde{z}_i)}$$

Proof. This follows immediately from Proposition 2.8: The length multi-ratio of a quadrilateral is the modulus of the complex cross-ratio. If the (embedded) quadrilaterals are cyclic, then their complex cross-ratios are real and negative, so their arguments are also equal. \square

For planar polyhedral surfaces, i.e., for quadrangulations in the complex plane, Proposition 2.12 connects discrete conformality with the cross-ratio system on quad-graphs. A *quad-graph* in the most general sense is simply an abstract CW cell decomposition of a surface with quadrilateral faces. Often, more conditions are added to the definition as needed. Here, we will require that the surface is oriented and that the vertices are bicolored black and white. For simplicity, we will also assume that the CW complex is strongly regular (see Section 2.2). The *cross-ratio system* on a quad-graph Σ imposes equations (15) on variables z_i that are attached to the vertices $i \in V_\Sigma$. There is one equation per face $ijkl \in F_\Sigma$:

$$\frac{(z_i - z_j)(z_k - z_l)}{(z_j - z_k)(z_l - z_i)} = Q_{ijkl}, \quad (15)$$

where we assume that i is a black vertex and the boundary vertices $ijkl$ are listed in the positive cyclic order. (Here we need the orientation). On the right hand side of the equation, $Q : F_\Sigma \rightarrow \mathbb{C} \setminus \{0, 1\}$ is a given function. In particular, it is required that the values z_i, z_j, z_k, z_l on a face are distinct.

By Proposition 2.12, two discretely conformally equivalent planar quadrangulations correspond to two solutions of the cross-ratio system on the same quad-graph

with the same cross-ratios Q . The following proposition says that in the simply connected case, one can find complex factors w on the vertices whose absolute values $|w| = e^{u/2}$ govern the length change of edges according to (3), and whose arguments govern the rotation of edges. Note that (3) is obtained from (16) by taking absolute values.

Proposition 2.13. *Let Σ be a simply connected quad-graph. Two functions $z, \tilde{z} : V_\Sigma \rightarrow \mathbb{C}$ are solutions of the cross-ratio system on Σ with the same cross-ratios Q if and only if there is a function $w : V_\Sigma \rightarrow \mathbb{C}$ such that for all edges $ij \in E_\Sigma$*

$$\tilde{z}_j - \tilde{z}_i = w_i w_j (z_j - z_i). \quad (16)$$

Proof. As in the proof of Proposition 2.8, it is easy to see that the system of equations (16) is solvable for w if and only if the complex multi-ratios for even cycles are equal. Because Σ is simply connected, this is the case if and only if the complex cross-ratios of corresponding faces are equal. \square

Remark 2.14. The cross-ratio system on quad-graphs (15) is an integrable system (in the sense of 3D consistency [5, 6]) if the cross-ratios Q “factor”, i.e., if there exists a function on the set of edges, $a : E_\Sigma \rightarrow \mathbb{C}$, that satisfies the following conditions for each quadrilateral $ijkl \in F$:

(i) It takes the same value on opposite edges,

$$a_{ij} = a_{kl}, \quad a_{jk} = a_{li}. \quad (17)$$

(ii)

$$Q_{ijkl} = \frac{a_{ij}}{a_{jk}}. \quad (18)$$

In Adler, Bobenko & Suris’ classification of integrable equations on quad-graphs [2], the integrable cross-ratio system is called $(Q1)_{\delta=0}$. It is also known as the discrete Schwarzian Korteweg–de Vries (dSKdV) equation, especially when it is considered on the regular square lattice [21] with constant cross-ratios.

If the cross-ratios Q have unit modulus, the cross-ratio system on quad-graphs is connected with circle patterns with prescribed intersection angles [5, 6].

Remark 2.15. The system of equations (16) is also connected with an integrable system on quad-graphs. Let $b_{ij} = z_j - z_i$, so b is a function on the oriented edges with $b_{ij} = -b_{ji}$. Let us also assume that the quad-graph Σ is simply connected. Then the system (16) defines a function $z : V \rightarrow \mathbb{C}$ (uniquely up to an additive constant) if and only if the complex scale factors $w : V_\Sigma \rightarrow \mathbb{C}$ satisfy, for each face $ijkl \in F$ the closure condition

$$b_{ij} w_i w_j + b_{jk} w_j w_k + b_{kl} w_k w_l + b_{li} w_l w_i = 0. \quad (19)$$

This system for w is integrable if, for each face $ijkl \in F$,

$$b_{ij} + b_{kl} = 0 \quad \text{and} \quad b_{jk} + b_{li} = 0.$$

In this case, (19) is known as discrete modified Korteweg–de Vries (dmKdV) equation [21], or as Hirota equation [5, 6].

3 Variational principles for discrete conformal maps

3.1 Discrete conformal mapping problems

We will consider the following discrete conformal mapping problems. (The notation $(\Sigma, \ell)_g$ was introduced in Definition 2.1.)

Problem 3.1 (prescribed angle sums). Given

- A euclidean, spherical, or hyperbolic cyclic polyhedral surface $(\Sigma, \ell)_g$, where $g \in \{euc, hyp, sph\}$,
- a desired total angle $\Theta_i > 0$ for each vertex $i \in V_\Sigma$,
- a choice of geometry $\tilde{g} \in \{euc, hyp, sph\}$,

find a discretely conformally equivalent cyclic polyhedral surface $(\Sigma, \tilde{\ell})_{\tilde{g}}$ of geometry \tilde{g} that has the desired total angles Θ around vertices.

For interior vertices, Θ prescribes a desired cone angle. For boundary vertices, Θ prescribes a desired interior angle of the polygonal boundary. If $\Theta_i = 2\pi$ for all interior vertices i , then Problem 3.1 asks for a flat metric in the discrete conformal class, with prescribed boundary angles if the surface has a boundary.

More generally, we will consider the following problem, where the logarithmic scale factors u (see Definition 2.2) are fixed at some vertices and desired angle sums Θ are prescribed at the other vertices. The problems to find discrete Riemann maps (Section 4.2) and maps onto the sphere (Section 6.1) can be reduced to this mapping problem with some fixed scale factors.

Problem 3.2 (prescribed scale factors and angle sums). Given

- A euclidean, spherical, or hyperbolic cyclic polyhedral surface $(\Sigma, \ell)_g$, where $g \in \{euc, hyp, sph\}$,
- a partition $V_\Sigma = V_0 \cup V_1$
- a prescribed angle $\Theta_i > 0$ for each vertex $i \in V_1$,
- a prescribed logarithmic scale factor $u_i \in \mathbb{R}$ for each vertex $i \in V_0$,
- a choice of geometry $\tilde{g} \in \{euc, hyp, sph\}$,

find a discretely conformally equivalent cyclic polyhedral surface $(\Sigma, \tilde{\ell})_{\tilde{g}}$ of geometry \tilde{g} that has the desired total angles Θ around vertices in V_1 and the fixed scale factors u at vertices in V_0 .

Note that for $V_0 = \emptyset$, $V_1 = V$, Problem 3.2 reduces to Problem 3.1.

3.2 Analytic formulation of the mapping problems

We rephrase the mapping Problem 3.2 analytically as Problem 3.4. The sides of a cyclic polygon determine its angles, but practical explicit equations for the angles as functions of the sides exist only for triangles, e.g., (21). For this reason it makes sense to triangulate the polyhedral surface. For the angles in a triangulation, we use the notation shown in Figure 4. In triangle ijk , we denote the angle at vertex i by α_{jk}^i .

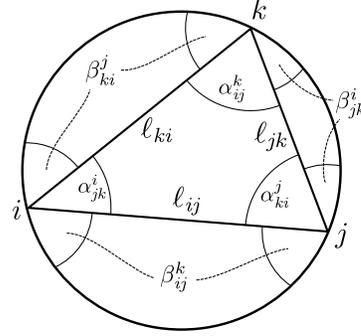


Fig. 4 Notation of lengths and angles in a triangle $ijk \in F$.

We denote by β_{ij}^i the angle between the circumcircle and the edge jk . The angles α and β are related by

$$\alpha_{jk}^i + \beta_{ki}^j + \beta_{ij}^k = \pi,$$

so betas determine alphas and vice versa:

$$2\beta_{jk}^i = \pi + \alpha_{jk}^i - \alpha_{ki}^j - \alpha_{ij}^k, \quad \dots \quad (20)$$

For euclidean triangles,

$$\alpha_{jk}^i + \alpha_{ki}^j + \alpha_{ij}^k = \pi, \quad \beta_{jk}^i = \alpha_{jk}^i.$$

The half-angle equation can be used to express the angles as functions of lengths:

$$\tan\left(\frac{\alpha_{jk}^i}{2}\right) = \begin{cases} \left(\frac{(-\ell_{ij} + \ell_{jk} + \ell_{ki})(\ell_{ij} + \ell_{jk} - \ell_{ki})}{(\ell_{ij} - \ell_{jk} + \ell_{ki})(\ell_{ij} + \ell_{jk} + \ell_{ki})}\right)^{\frac{1}{2}} & (euc) \\ \left(\frac{\sinh((\ell_{ij} - \ell_{jk} + \ell_{ki})/2) \sinh((\ell_{ij} + \ell_{jk} - \ell_{ki})/2)}{\sinh((- \ell_{ij} + \ell_{jk} + \ell_{ki})/2) \sinh((\ell_{ij} + \ell_{jk} + \ell_{ki})/2)}\right)^{\frac{1}{2}} & (hyp) \\ \left(\frac{\sin((\ell_{ij} - \ell_{jk} + \ell_{ki})/2) \sin((\ell_{ij} + \ell_{jk} - \ell_{ki})/2)}{\sin((- \ell_{ij} + \ell_{jk} + \ell_{ki})/2) \sin((\ell_{ij} + \ell_{jk} + \ell_{ki})/2)}\right)^{\frac{1}{2}} & (sph) \end{cases} \quad (21)$$

Lemma 3.3 (analytic formulation of Problem 3.2). *Let*

- the polyhedral surface $(\Sigma, \ell)_g$,
- the partition $V_0 \dot{\cup} V_1$,
- Θ_i for $i \in V_1$,
- u_i for $i \in V_0$,
- the geometry $\tilde{g} \in \{euc, hyp, sph\}$

be given as in Problem 3.2. Let Δ be an abstract triangulation obtained by adding non-crossing diagonals to non-triangular faces of Σ . (So $V_\Sigma = V_\Delta$, $E_\Sigma \subseteq E_\Delta$, and the set of added diagonals is $E_\Delta \setminus E_\Sigma$.) For $ij \in E_\Sigma$, define λ_{ij} by

$$\lambda_{ij} = \begin{cases} 2 \log \ell_{ij} & \text{if } g = euc \\ 2 \log \sinh \frac{\ell_{ij}}{2} & \text{if } g = hyp \\ 2 \log \sin \frac{\ell_{ij}}{2} & \text{if } g = sph \end{cases} \quad (22)$$

Then solving Problem 3.2 is equivalent to solving Problem 3.4 with $E_0 = E_\Sigma$ and $E_1 = E_\Delta \setminus E_\Sigma$.

Problem 3.4. Given

- an abstract triangulation Δ ,
- a partition $V_\Delta = V_0 \dot{\cup} V_1$,
- $u_i \in \mathbb{R}$ for $i \in V_0$
- $\Theta_i \in \mathbb{R}_{>0}$ for $i \in V_1$,
- a partition $E_\Delta = E_0 \dot{\cup} E_1$,
- λ_{ij} for $ij \in E_0$,
- $\tilde{g} \in \{euc, hyp, sph\}$,

find $u_i \in \mathbb{R}$ for $i \in V_1$ and λ_{ij} for $ij \in E_1$ such that

$$\tilde{\ell} : E_\Delta \rightarrow \mathbb{R}_{>0}$$

defined by

$$\tilde{\lambda}_{ij} = u_i + u_j + \lambda_{ij}, \quad (23)$$

and

$$\tilde{\ell}_{ij} = \begin{cases} e^{\frac{1}{2}\tilde{\lambda}_{ij}} & \text{if } \tilde{g} = euc \\ 2 \operatorname{arsinh} e^{\frac{1}{2}\tilde{\lambda}_{ij}} & \text{if } \tilde{g} = hyp \\ 2 \arcsin e^{\frac{1}{2}\tilde{\lambda}_{ij}} & \text{if } \tilde{g} = sph \end{cases} \quad (24)$$

satisfies for all $ijk \in F_\Delta$ the triangle inequalities

$$\tilde{\ell}_{ij} < \tilde{\ell}_{jk} + \tilde{\ell}_{ki}, \quad \tilde{\ell}_{jk} < \tilde{\ell}_{ki} + \tilde{\ell}_{ij}, \quad \tilde{\ell}_{ki} < \tilde{\ell}_{ij} + \tilde{\ell}_{jk}, \quad (25)$$

and for $\tilde{g} = sph$ also

$$\tilde{\ell}_{ij} + \tilde{\ell}_{jk} + \tilde{\ell}_{ki} < 2\pi, \quad (26)$$

and such that

$$\sum_{jk:ijk \in F_\Delta} \tilde{\alpha}_{jk}^i = \Theta_i \quad \text{for all } i \in V_1, \quad (27)$$

$$\tilde{\beta}_{ij}^k + \tilde{\beta}_{ji}^l = \pi \quad \text{for all } ij \in E_1, \quad (28)$$

where $\tilde{\alpha}$ and $\tilde{\beta}$ are defined by (21) and (20) (with α, β, ℓ replaced by $\tilde{\alpha}, \tilde{\beta}, \tilde{\ell}$). Note that for $\tilde{g} = sph$ it is also required that $\tilde{\lambda} < 0$ for $\tilde{\ell}$ to be well-defined.

Proof (of Lemma 3.3). Note that (27) says that the angle sums at vertices in V_1 have the prescribed values, and (28) says that neighboring triangles of $(\Delta, \tilde{\ell})_{\tilde{g}}$ belonging to the same face of Σ share the same circumcircle. So deleting the edges in $E_\Delta \setminus E_\Sigma$, one obtains a cyclic polyhedral surface $(\Sigma, \tilde{\ell}|_{E_\Sigma})_{\tilde{g}}$. \square

3.3 Variational principles

Definition 3.5. For an abstract triangulation Δ and a function $\Theta \in \mathbb{R}_{>0}^{V_\Delta}$, define the three functions

$$\begin{aligned} E_{\Delta, \Theta}^{euc}, E_{\Delta, \Theta}^{hyp}, E_{\Delta, \Theta}^{sph} : \mathbb{R}^{E_\Delta} \times \mathbb{R}^{V_\Delta} &\longrightarrow \mathbb{R}, \\ (\lambda, u) &\longmapsto E_{\Delta, \Theta}^{\tilde{g}}(\lambda, u) \end{aligned}$$

by

$$E_{\Delta, \Theta}^{\tilde{g}}(\lambda, u) = \sum_{ijk \in F_\Delta} (f^{\tilde{g}}(\tilde{\lambda}_{ij}, \tilde{\lambda}_{jk}, \tilde{\lambda}_{ki}) - \frac{\pi}{2}(\tilde{\lambda}_{jk} + \tilde{\lambda}_{ki} + \tilde{\lambda}_{ij})) + \sum_{i \in V_\Delta} \Theta_i u_i, \quad (29)$$

where $\tilde{g} \in \{euc, hyp, sph\}$, $\tilde{\lambda}$ is defined as function of λ and u by (23), and the functions $f^{euc}, f^{hyp}, f^{sph}$ are defined in Section 3.4.

We will often omit the subscripts and write simply $E^{euc}, E^{hyp}, E^{sph}$ when this is unlikely to cause confusion.

Definition 3.6. We define the *feasible regions* of the functions $E_{\Delta, \Theta}^{\tilde{g}}$ as the following open subsets of their domains:

- The feasible region of E^{euc} and E^{hyp} is the set of all $(\lambda, u) \in \mathbb{R}^{E_\Delta} \times \mathbb{R}^{V_\Delta}$ such that $\tilde{\ell} \in \mathbb{R}_{>0}^{E_\Delta}$ defined by (23) and (24) satisfies the triangle inequalities (25)
- The feasible region of E^{sph} is the set of all $(\lambda, u) \in \mathbb{R}^{E_\Delta} \times \mathbb{R}^{V_\Delta}$ such that $\tilde{\lambda}$ defined by (23) is negative, and $\tilde{\ell}$, which is then well-defined by (24), satisfies the triangle inequalities (25) and the inequalities (26).

Theorem 3.7 (Variational principles). *Every solution $(\Sigma, \tilde{\ell})_{\tilde{g}}$ of Problem 3.2 corresponds via (23) and (24) to a critical point $(\lambda, u) \in \mathbb{R}^{E_\Delta} \times \mathbb{R}^{V_\Delta}$ of the function $E_{\Delta, \Theta}^{\tilde{g}}$ under the constraints that λ_{ij} and u_i are fixed for $ij \in E_0$ and $i \in V_0$, respectively. (The triangulation Δ , and $E_0 = E_\Sigma$ and $E_1 = E_\Delta \setminus E_\Sigma$ are as in Lemma 3.3, and the given function Θ is extended from V_1 to V by arbitrary values on V_0 .)*

Conversely, if $(\lambda, u) \in \mathbb{R}^{E_\Delta} \times \mathbb{R}^{V_\Delta}$ is a critical point of the function $E_{\Delta, \Theta}^{\tilde{g}}$ under the same constraints, and if (λ, u) is contained in the feasible region of $E_{\Delta, \Theta}^{\tilde{g}}$, then $(\Sigma, \tilde{\ell})_{\tilde{g}}$ defined by (23) and (24) is a solution of Problem 3.2.

Proof. This follows from the analytic formulation of Problem 3.2 (see Section 3.2) and Proposition 3.8. \square

Proposition 3.8 (First derivative of $E^{\tilde{g}}$). *The partial derivatives of $E^{\tilde{g}}$ are*

$$\frac{\partial E^{\tilde{g}}}{\partial u_i}(\lambda, u) = \Theta_i - \sum_{ijk \ni i} \tilde{\alpha}_{jk}^i \quad (30)$$

$$\frac{\partial E^{\tilde{g}}}{\partial \lambda_{ij}}(\lambda, u) = \tilde{\beta}_{ij}^k + \tilde{\beta}_{ij}^l - \pi. \quad (31)$$

Here $\tilde{\alpha}, \tilde{\beta}$ are defined by (21) and (20) (with α, β, ℓ replaced by $\tilde{\alpha}, \tilde{\beta}, \tilde{\ell}$) if (λ, u) is contained in the feasible region of $E^{\tilde{g}}$. For (λ, u) not contained in the feasible region, the definition of $\tilde{\alpha}, \tilde{\beta}$ is extended like in Definition 3.12.

Proof. Equations (30) and (31) follow from the definition of $E^{\tilde{g}}$ and Proposition 3.14 on the partial derivatives of $f^{\tilde{g}}$. \square

Theorem 3.9 (Uniqueness for mapping problems). *If Problem 3.2 with target geometry $\tilde{g} \in \{euc, hyp\}$ has a solution, then the solution is unique — except if $\tilde{g} = euc$ and $V_0 = \emptyset$ (the case of Problem 3.1). In this case, the solution is unique up to scale.*

The critical point $(\lambda, u) \in \mathbb{R}^{E_\Delta} \times \mathbb{R}^{V_\Delta}$ that corresponds, via (23) and (24), to a solution $(\Sigma, \tilde{\ell})_{\tilde{g}}$ of Problem 3.2 with $\tilde{g} \in \{euc, hyp\}$ is a minimizer of $E_{\Delta, \Theta}^{\tilde{g}}$ under the constraints described in Theorem 3.7. The minimizer is unique except in the following cases. If $\tilde{g} = euc$ and $V_0 = \emptyset$, then $E_{\Delta, \Theta}^{\tilde{g}}$ is constant along all lines in the “scaling direction” $(0, 1_{V_\Delta}) \in \mathbb{R}^{E_\Delta} \times \mathbb{R}^{V_\Delta}$. If the 1-skeleton of Σ is bipartite and $V_0 = \emptyset$, then $E_{\Delta, \Theta}^{\tilde{g}}$ is constant in the direction that is ± 1 on the two color classes of V_Δ , respectively, and takes appropriate values on $E_\Delta \setminus E_\Sigma$ so that $\tilde{\lambda}_{ij}$ defined by (23) remains constant for all $ij \in E_\Delta$. (In both exceptional cases, one can make the minimizer unique by adding the constraint of fixing u_i for some $i \in V_\Delta$.)

Proof. The theorem follows from Theorem 3.7 and the following observations.

- (1) If the point $(\lambda, u) \in \mathbb{R}^{E_\Delta} \times \mathbb{R}^{V_\Delta}$ corresponds to a solution of Problem 3.2, it is contained in the feasible region of $E_{\Delta, \Theta}^{\tilde{g}}$.
- (2) By (29) and Proposition 3.16, the functions E^{euc} and E^{hyp} are convex.
- (3) For (λ, u) in the feasible region, the second derivative $D^2 E^{hyp}(\lambda, u)$ is a positive definite quadratic form of $d\tilde{\lambda}$, i.e., $D^2 E^{hyp}(\lambda, u)(\hat{\lambda}, \hat{u}) \geq 0$ for all $(\hat{\lambda}, \hat{u}) \in \mathbb{R}^{E_\Delta} \times \mathbb{R}^{V_\Delta}$ and $D^2 E^{hyp}(\lambda, u)(\hat{\lambda}, \hat{u}) = 0$ if and only if

$$\hat{\lambda}_{ij} + \hat{u}_i + \hat{u}_j = 0 \quad \text{for all } ij \in E_\Delta.$$

- (4) Similarly, for (λ, u) in the feasible region, the second derivative $D^2E^{euc}(\lambda, u)$ is a positive semidefinite quadratic form with $D^2E^{euc}(\lambda, u)(\dot{\lambda}, \dot{u}) = 0$ if and only if

$$\dot{\lambda}_{ij} + \dot{u}_i + \dot{u}_j = c \quad \text{for all } ij \in E_\Delta, \text{ for some } c \in \mathbb{R}.$$

□

In the following proposition, we collect explicit formulas for the second derivatives of the functions $E^{\tilde{g}}$. They are useful for the numerical minimization of E^{euc} and E^{hyp} , and even for finding critical points of E^{sph} , as explained in Section 6.2.

Proposition 3.10 (Second derivative of $E^{\tilde{g}}$). *The second derivatives of E^{euc} , E^{hyp} , and E^{sph} are the quadratic forms*

$$D^2E^{\tilde{g}}(\lambda, u) = \frac{1}{2} \sum_{ijk \in F_\Delta} (q_{ij}^k(\lambda, u) + q_{jk}^i(\lambda, u) + q_{ki}^j(\lambda, u)),$$

where $q_{ij}^k(\lambda, u) = 0$ if $\tilde{\ell}_{ij}, \tilde{\ell}_{jk}, \tilde{\ell}_{ki}$ defined by (23), (24) violate the triangle inequalities (25), or, in the case of $\tilde{g} = sph$, inequality (26). Otherwise, the quadratic forms $q_{ij}^k(\lambda, u)$ are defined by

$$q_{ij}^k = \begin{cases} \cot \tilde{\alpha}_{ij}^k (d\lambda_{ki} - d\lambda_{jk} + du_i - du_j)^2 & (euc) \\ \cot \tilde{\beta}_{ij}^k ((d\lambda_{ik} - d\lambda_{kj} + du_i - du_j)^2 + \tanh^2(\frac{\tilde{\ell}_{ij}}{2})(d\lambda_{ij} + du_i + du_j)^2) & (hyp) \\ \cot \tilde{\beta}_{ij}^k ((d\lambda_{ik} - d\lambda_{kj} + du_i - du_j)^2 - \tan^2(\frac{\tilde{\ell}_{ij}}{2})(d\lambda_{ij} + du_i + du_j)^2) & (sph) \end{cases}$$

where $\tilde{\alpha}, \tilde{\beta}$ are defined by (21) and (20) (with α, β, ℓ replaced by $\tilde{\alpha}, \tilde{\beta}, \tilde{\ell}$).

Proposition 3.10 follows from (29) and Proposition 3.15 about the second derivatives of f^g .

3.4 The peculiar triangle functions

This section is concerned with three real valued functions $f^{euc}, f^{hyp}, f^{sph}$ of three variables that are the main building blocks for the action functions $E^{euc}, E^{hyp}, E^{sph}$ of the variational principles. Since we consider single triangles in this section, not triangulations, we can use simpler notation. For $\{i, j, k\} = \{1, 2, 3\}$, let

$$\lambda_i = \lambda_{jk}, \quad \ell_i = \ell_{jk}, \quad \alpha_i = \alpha_{jk}^i, \quad \beta_i = \beta_{jk}^i.$$

The terminology introduced in the following definition makes Definition 3.12 easier to state.

Definition 3.11. Let the *feasible region* of f^{euc} and f^{hyp} be the open subset of all $\lambda \in \mathbb{R}^3$ such that $\ell \in \mathbb{R}_{>0}^3$ determined by (22) satisfies the triangle inequalities, i.e.,

$$\ell_k < \ell_i + \ell_j \tag{32}$$

for $\{i, j, k\} = \{1, 2, 3\}$.

Let the *feasible region* of f^{sph} be the open subset of all $\lambda \in \mathbb{R}^3$ such that $\lambda < 0$, and such that $\ell \in \mathbb{R}_{>0}^3$, which is then well-defined by (22), satisfies the triangle inequalities (32) and

$$\ell_1 + \ell_2 + \ell_3 < 2\pi. \tag{33}$$

Definition 3.12. We define the three functions

$$f^{euc}, f^{hyp}, f^{sph} : \mathbb{R}^3 \rightarrow \mathbb{R}$$

by

$$f^g(\lambda_1, \lambda_2, \lambda_3) = \beta_1 \lambda_1 + \beta_2 \lambda_2 + \beta_3 \lambda_3 + \mathbb{I}(\alpha_1) + \mathbb{I}(\alpha_2) + \mathbb{I}(\alpha_3) + \mathbb{I}(\beta_1) + \mathbb{I}(\beta_2) + \mathbb{I}(\beta_3) + \mathbb{I}\left(\frac{1}{2}(\pi - \alpha_1 - \alpha_2 - \alpha_3)\right), \tag{34}$$

where $g \in \{euc, hyp, sph\}$, $\mathbb{I}(x)$ denotes Milnor's Lobachevsky function [18]

$$\mathbb{I}(x) = - \int_0^x \log |2 \sin(t)| dt, \tag{35}$$

and,

- if λ is in the feasible region of f^g , then the angles α, β are defined as the angles (shown in Figure 4) in a euclidean, hyperbolic, or spherical triangle (depending on g) with sides ℓ_1, ℓ_2, ℓ_3 determined by (22). That is, α and β are defined by (21) and (20).

- Otherwise, if $g = sph$, and if either at least two λ s are non-negative or $\lambda < 0$ and inequality (33) is violated, let

$$\alpha_k = \alpha_i = \alpha_j = \pi, \quad \beta_k = \beta_i = \beta_j = 0.$$

- Otherwise, if the triangle inequality (32) is violated, or if $g = sph$ and $\lambda_k \geq 0$, let

$$\alpha_k = \beta_k = \pi, \quad \alpha_i = \alpha_j = \beta_i = \beta_j = 0.$$

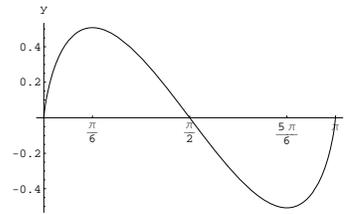


Fig. 5 Graph of Milnor's Lobachevsky function, $y = \mathbb{I}(x)$.

Figure 5 shows a graph of Milnor's Lobachevsky function. It is continuous, π -periodic, odd, has zeros at the integer multiples of $\pi/2$, and is real analytic except at integer multiples of π , where the derivative tends to $+\infty$.

Remark 3.13. In the euclidean case, (34) simplifies to

$$f^{\text{euc}}(\lambda) = \alpha_i \lambda_i + \alpha_j \lambda_j + \alpha_k \lambda_k + 2\text{JI}(\alpha_i) + 2\text{JI}(\alpha_j) + 2\text{JI}(\alpha_k). \quad (36)$$

This follows immediately from $\alpha_1 + \alpha_2 + \alpha_3 = \pi$, $\alpha = \beta$, and $\text{JI}(0) = 0$.

Proposition 3.14 (first derivative). *The functions f^g , $g \in \{\text{euc}, \text{hyp}, \text{sph}\}$, are continuously differentiable and*

$$\frac{\partial f^g}{\partial \lambda_i} = \beta_i. \quad (37)$$

Proof. Note that the angles α , β are continuous functions of λ on \mathbb{R}^3 . Hence f^g defined by (34) is also continuous. We will show that f^g is continuously differentiable with derivative (37) on an open dense subset of the domain, namely, the union of (a) the feasible region and (b) the interior of its complement. Since f^g is continuous and df^g extends continuously to \mathbb{R}^3 , the claim follows.

(a) First, suppose λ is contained in the feasible region of f^g . By symmetry, it suffices to consider the derivative with respect to λ_1 . From (34) and (35) one obtains

$$\begin{aligned} \frac{\partial f^g}{\partial \lambda_1} &= \beta_1 + \\ &\sum_{i=1}^3 \left((\lambda_i - \log(2 \sin \beta_i)) \frac{\partial \beta_i}{\partial \lambda_1} + (-\log(2 \sin \alpha_i) + \frac{1}{2} \log |2 \sin(\frac{\pi - \alpha_1 - \alpha_2 - \alpha_3}{2})|) \frac{\partial \alpha_i}{\partial \lambda_1} \right) \end{aligned} \quad (38)$$

For hyperbolic and spherical triangles, one derives from the respective cosine rules

$$\begin{aligned} \sinh^2 \frac{\ell_i}{2} &= \frac{\sin \beta_i \sin \frac{\pi - \alpha_1 - \alpha_2 - \alpha_3}{2}}{\sin \alpha_2 \sin \alpha_3} && \text{(hyperbolic),} \\ \sin^2 \frac{\ell_i}{2} &= \frac{\sin \beta_i \sin \frac{\alpha_1 + \alpha_2 + \alpha_3 - \pi}{2}}{\sin \alpha_2 \sin \alpha_3} && \text{(spherical).} \end{aligned}$$

In both cases, expand the fraction on the right hand side by four and take logarithms to find

$$\lambda_i = \log(2 \sin \beta_i) + \log \left| 2 \sin \frac{\pi - \alpha_1 - \alpha_2 - \alpha_3}{2} \right| - \log(2 \sin \alpha_j) - \log(2 \sin \alpha_k).$$

Substitute this expression for λ_i in (38) and use $d\beta_i = \frac{1}{2}(d\alpha_i - d\alpha_j - d\alpha_k)$ to see that all terms on the right hand side of (38) cancel, except β_1 .

For euclidean triangles, (38) simplifies to

$$\frac{\partial f^g}{\partial \lambda_1} = \beta_1 + \sum_{i=1}^3 (\lambda_i - 2 \log(2 \sin \alpha_i)) \frac{\partial \alpha_i}{\partial \lambda_1},$$

where

$$\lambda_i - 2 \log(2 \sin \alpha_i) = 2 \log \frac{\ell_i}{2 \sin \alpha_i} = 2 \log R$$

does not depend on i . (R denotes the circumradius.) Equation (37) follows because the angle sum is constant.

(b) Now suppose λ is contained the interior of the complement of the feasible region of f^g . Since $\beta_1, \beta_2, \beta_3$ are constant on each connected component of the complement of the feasible region, and since

$$f^g(\lambda_1, \lambda_2, \lambda_3) = \beta_1 \lambda_1 + \beta_2 \lambda_2 + \beta_3 \lambda_3,$$

outside the feasible region, equation 37 holds also in this case. This completes the proof. \square

Proposition 3.15 (second derivative). *For $g \in \{euc, hyp, sph\}$ the function f^g is twice continuously differentiable on its feasible set and the second derivative is*

$$D^2 f^{euc} = \frac{1}{2} \sum_{i=1}^3 \cot \alpha_i (d\lambda_j - d\lambda_k)^2, \quad (39)$$

$$D^2 f^{hyp} = \frac{1}{2} \sum_{i=1}^3 \cot \beta_i ((d\lambda_j - d\lambda_k)^2 + \tanh^2(\frac{\ell_i}{2}) d\lambda_i^2), \quad (40)$$

$$D^2 f^{sph} = \frac{1}{2} \sum_{i=1}^3 \cot \beta_i ((d\lambda_j - d\lambda_k)^2 - \tan^2(\frac{\ell_i}{2}) d\lambda_i^2). \quad (41)$$

On each component of the complement of its feasible set, the function f_g is linear so the second derivative vanishes.

A proof of (39) is contained in [4] (Proposition 4.2.3), see Remark 3.17 below. Equations (40) and (41) can be derived by lengthy calculations.

Proposition 3.16. (i) *The function f^{euc} is convex. On its feasible set, the second derivative $D^2 f^{euc}$ is positive semidefinite with one-dimensional kernel spanned by the “scaling direction” $(1, 1, 1)$.*

(ii) *The function f^{hyp} is convex. On its feasible set, the second derivative $D^2 f^{hyp}$ is positive definite, so the functions is locally strictly convex.*

Part (i) is proved in [4] (Propositions 4.2.4, 4.2.5, note the following remark) directly from (39). We do not know a similarly straightforward proof of part (ii). The proof in [4] (Section 6.2) is based on a connection with 3-dimensional hyperbolic geometry: f^{hyp} is the Legendre dual of the volume of an ideal hyperbolic prism considered as a function of the dihedral angles. This volume function is strictly concave, as shown by Leibon [16]. His argument uses the decomposition of an ideal prism into three ideal tetrahedra.

Remark 3.17. The functions f and \hat{V}_h defined in [4] (equations (4-3), (6-4)) are related to the functions f^{euc} and f^{hyp} by

$$f^{euc}(\lambda_1, \lambda_2, \lambda_3) = 2f\left(\frac{\lambda_1}{2}, \frac{\lambda_2}{2}, \frac{\lambda_3}{2}\right), \quad (42)$$

$$f^{hyp}(\lambda_1, \lambda_2, \lambda_3) = 2\hat{V}_h(\lambda_1, \lambda_2, \lambda_3, 0, 0, 0). \quad (43)$$

4 Conformal maps of cyclic quadrangulations

Having introduced the mapping problems and variational principles, we return to conformal maps of cyclic quadrangulations. Some basic facts were already discussed in Section 2.8. Here, in Section 4.1, we consider a simple experiment that demonstrates the somewhat unexpected appearance of orthogonal circle patterns, and also a necessary condition for the boundary angles. In Section 4.2, we discuss a discrete version of the Riemann mapping problem for quadrangulations.

4.1 Emerging circle patterns and a necessary condition

Consider the two discrete conformal maps shown in the two rows of Figure 6. The domains (shown left) are a square and a rectangle, subdivided into 6×6 and 6×5 squares, respectively. We solve the mapping Problem 3.1 by minimizing E^{euc} as explained in Section 3.3, prescribing boundary angles to obtain maps to parallelograms: $\Theta = 50^\circ$ and 130° for the corner vertices, $\Theta = 180^\circ$ for the other boundary vertices, and $\Theta = 360^\circ$ for interior vertices. The resulting quadrangulations are shown in the middle.

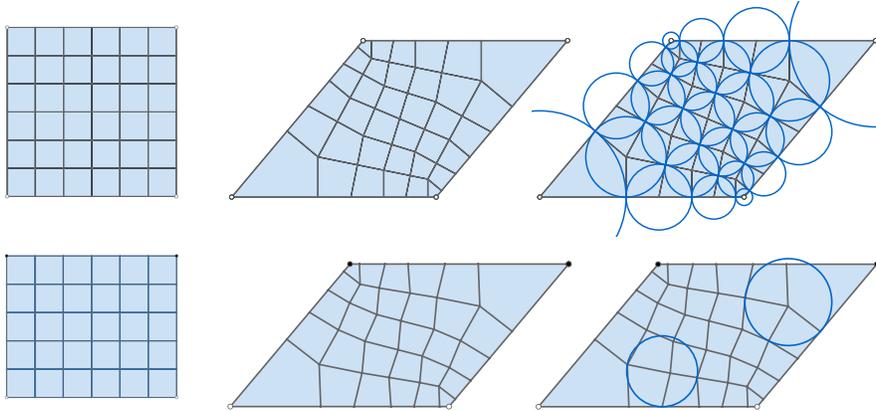


Fig. 6 Mapping a rectangle to a parallelogram. Note the orthogonal circle pattern in top row and the wiggly vertical lines in the bottom row.

On first sight, the 6×6 example shown in the top row behaves rather like one would expect from a conformal map. The horizontal and vertical “coordinate lines” of the domain are mapped to polygonal curves that look more or less like they could be discretizations of reasonable smooth curves. In the 6×5 example shown in the bottom row, the images of the vertical lines zigzag noticeably.

A closer look at the 6×6 example reveals a remarkable phenomenon. Let us bi-color the vertices black and white so that neighboring vertices have different colors, with the corners colored white. Then, in the image quadrangulation, the edges incident with a black vertex meet at right angles, and the edges incident with a white vertex have the same length. One can therefore draw a circle around each white vertex through the neighboring black vertices as shown in Figure 6 (top right). At the black vertices, these circles touch and intersect orthogonally. Such circle patterns were studied by Schramm [25] as discrete analogs of conformal maps.

Given such a circle pattern with orthogonally intersecting circles, the quadrangulation formed by drawing edges between circle centers and intersection points consists of quadrilaterals that are right-angled kites. Such kites have complex cross-ratio -1 . Hence, the quadrangulation coming from an orthogonal circle pattern is discretely conformally equivalent (in our sense) to a combinatorially equivalent quadrangulation consisting of squares.

The conformal map shown in the top row of Figure 6 “finds” the orthogonal circle pattern because that circle pattern exists and the conformal map is unique (by Theorem 3.9). For 6×5 example shown in the bottom row, a corresponding orthogonal circle pattern does not exist. No matter which coloring one chooses, there are two black vertices at which the total angle changes (from 90° to 50° and 130° , respectively). The neighbors of a vertex do not lie on a circle. Figure 6 (bottom right) shows two circles drawn through three out of four neighbors.

If we map an $m \times n$ square grid to a parallelogram like in Figure 6, an orthogonal circle pattern will appear if m and n are even. No such pattern will appear if one of the numbers is even and the other is odd. What happens if both m and n are odd? In this case, the conformal map does not exist. The corners with increasing angle and the corners with decreasing angle would have different colors. This violates the necessary condition expressed in the following theorem.

Theorem 4.1 (Necessary condition for the existence of conformal map). *Let Σ be an abstract quadrangulation of the closed disk, and let*

$$z, \tilde{z}: V_\Sigma \rightarrow \mathbb{C}$$

determine two discretely conformally equivalent immersions of Σ into the complex plane. Denote their angle sums at boundary vertices $v \in V_\Sigma$ by Θ_v and $\tilde{\Theta}_v$, respectively. Since the 1-skeleton of Σ is bipartite, we may assume the vertices are colored black and white. Let V_b^∂ and V_w^∂ denote the sets of black and white boundary vertices of Σ . Then

$$\sum_{v \in V_b^\partial} (\tilde{\Theta}_v - \Theta_v) \equiv 0 \pmod{2\pi}, \quad (44)$$

$$\sum_{v \in V_w^\partial} (\tilde{\Theta}_v - \Theta_v) \equiv 0 \pmod{2\pi}. \quad (45)$$

(Since $\sum_{V_b^\partial \cup V_w^\partial} (\tilde{\Theta}_v - \Theta_v) = 0$, equations (45) and (44) are equivalent.)

Proof. Since z and \tilde{z} are two solutions of the cross-ratio system on Σ with the same cross-ratios (see Section 2.8), there exists by Proposition 2.13 a function $w : V_\Sigma \rightarrow \mathbb{C}$ such that (16) holds for all edges $ij \in E_\Sigma$. Now suppose $v_0, \dots, v_{2n-1} \in V_\Sigma$ are the boundary vertices in cyclic order (with indices taken modulo $2n$). Then

$$e^{i(\tilde{\Theta}_{v_k} - \Theta_{v_k})} = \frac{(\tilde{z}_{v_{k+1}} - \tilde{z}_{v_k})(z_{v_{k-1}} - z_{v_k})}{(\tilde{z}_{v_{k-1}} - \tilde{z}_{v_k})(z_{v_{k+1}} - z_{v_k})} = \frac{w_{v_{k+1}}}{w_{v_{k-1}}},$$

so

$$\prod_{k=0}^{n-1} e^{i(\tilde{\Theta}_{v_{2k}} - \Theta_{v_{2k}})} = \prod_{k=0}^{n-1} e^{i(\tilde{\Theta}_{v_{2k+1}} - \Theta_{v_{2k+1}})} = 1.$$

Equations (44) and (45) follow. \square

4.2 Riemann maps with cyclic quadrilaterals

Consider the following discrete version of the Riemann mapping problem: Map a cyclic polyhedral surface that is topologically a closed disk discretely conformally to a planar polygonal region with boundary vertices on a circle. An example is shown in Figure 7, top row. This type of problem can often be reduced to Problem 3.2. Then, by the variational principle, if a solution exists, it can be found by minimizing a convex function. For triangulations, the reduction of the discrete Riemann mapping problem to Problem 3.2 is explained in [4] (Section 3.3). Here, we consider the case of quadrangulations. (The arguments can be extended to even polygons with more than four sides. We restrict our attention to quadrilaterals because the combinatorial restrictions discussed in the following paragraph become even more annoying for surfaces with hexagons, octagons, etc.)

The basic idea is the same as for triangulations: First, map the polyhedral surface to the half plane with one boundary vertex at infinity. Then apply a Möbius transformation. This leads to a combinatorial restriction: No face may have more than one edge on the boundary. (The face would degenerate when the boundary is mapped to a straight line.) For triangulations, this means that no triangle may be connected to the surface by only one edge. If this condition is violated, cutting off such “ears” often leads to an admissible triangulation. For quadrangulations, this fix does not work in typical situations. Instead, if a quadrilateral contains two consecutive edges on the boundary, cut off a triangle. The resulting polyhedral surface

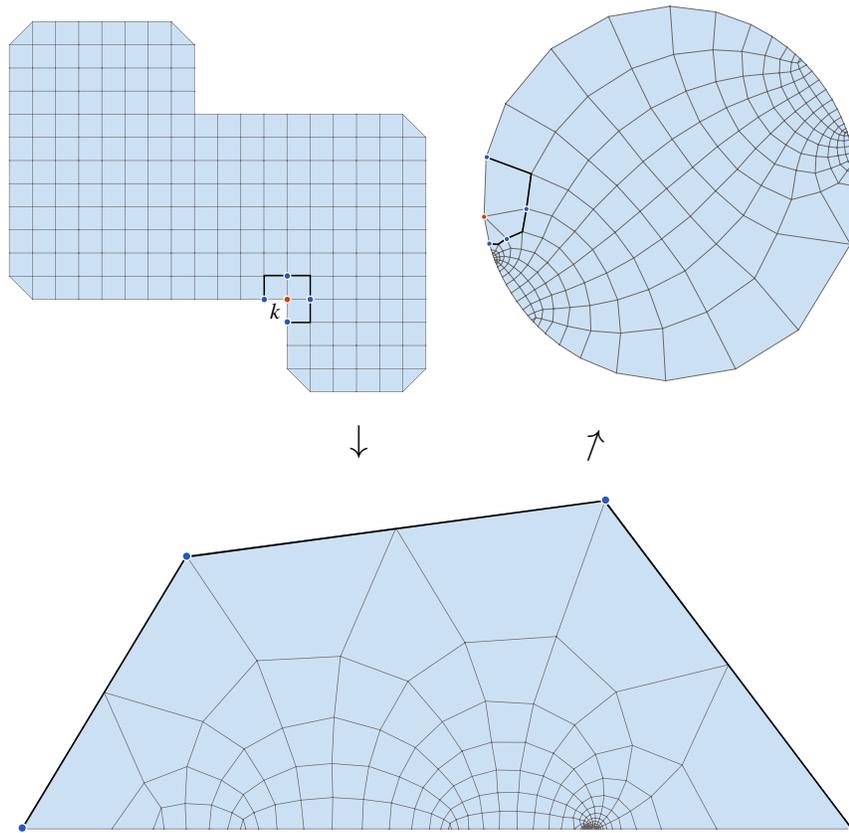


Fig. 7 Riemann mapping with cyclic quadrilaterals

will consist mostly of quadrilaterals with some triangles on the boundary, as in the example shown in Figure 7.

Suppose $(\Sigma, \ell)_{euc}$ is a euclidean cyclic polyhedral surface that is homeomorphic to the closed disk and consists mostly of quadrilaterals. (For the following construction we really only need a boundary vertex that is incident with quadrilateral faces.) To map it to a polygonal region inscribed in a circle, proceed as follows (see Figure 7):

- (1) Choose a vertex k on the boundary of Σ such that all incident faces are quadrilaterals.
- (2) Apply a discrete conformal change of metric (3) such that all edges incident with k have the same length. One may choose $u = 0$ for all vertices except the neighbors of k . It does not matter if polygon inequalities are violated after this step.

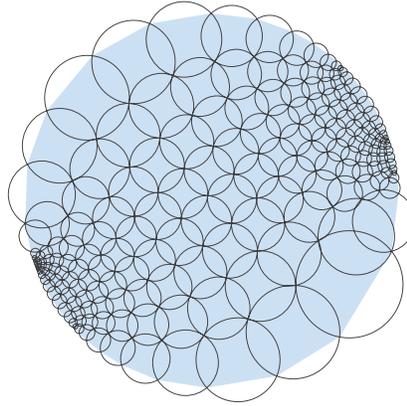


Fig. 8 Here we show the face circumcircles of the solution to the Riemann mapping problem of Figure 7. It looks conspicuously like an orthogonal circle pattern. But the face circumcircles intersect only approximately but not exactly at right angles.

- (3) Let $(\Sigma', \ell')_{euc}$ be the cyclic polyhedral complex obtained by removing vertex k and all incident quadrilaterals.
- (4) Solve Problem 3.2 for $(\Sigma', \ell')_{euc}$ with prescribed total angles $\Theta_i = 2\pi$ for interior vertices of Σ' , $\Theta_i = \pi$ for boundary vertices of Σ' that were not neighbors of k in Σ , and fixed logarithmic scale factors $u_i = 0$ for those that were neighbors of k . The result is a planar polyhedral surface as shown in Figure 7, bottom. The boundary consists of one straight line segment containing all boundary edges of Σ' that were also boundary edges of Σ , and two or more straight line segments, each consisting of two edges that were incident with a removed quadrilateral.
- (5) Apply a Möbius transformation (e.g., $z \mapsto 1/z$) to the vertices that maps the boundary vertices of Σ to a circle and the other vertices to the inside of this circle. Reinsert k at the image point of ∞ under this Möbius transformation. Each face $ijmk \in \Sigma$ incident with k is cyclic because the three vertices i, j , and m are contained in a line before transformation.
- (6) Optionally apply a 2-dimensional version of the Möbius normalization described in Section 6.3.

Proposition 4.2. *The result of this procedure is a planar cyclic polyhedral surface that is discretely conformally equivalent to $(\Sigma, \ell)_{euc}$ and has its boundary polygon inscribed in a circle.*

Proof. That the boundary polygon is inscribed in a circle is obvious from the construction. Using the Möbius invariance of discrete conformal equivalence (Proposition 2.5), it is not difficult to see that the surfaces without quadrilaterals incident with k are discretely conformally equivalent. To show that the whole surfaces are equivalent, it suffices to show that corresponding quadrilaterals incident with k have the same complex cross-ratio.

After step (2), the length cross-ratio of a quadrilateral incident with k is equal to the simple length ratio of the two edges that are not incident with k .

After step (4), the length cross-ratio of these edges is unchanged due to the fixed logarithmic scale factors $u = 0$ on the neighbors of k . Also, these edges are now collinear because of the prescribed angle $\Theta = \pi$ between them.

After applying the Möbius transformation in step (5), the image of the point at infinity and the other three vertices of our quadrilateral incident with k form again a cyclic quadrilateral with the same complex cross-ratio as in the beginning. \square

5 Multiply connected domains

5.1 Circle domains

Koebe's generalization of the Riemann mapping theorem says that multiply connected domains are conformally equivalent to domains bounded by circles, and the uniformizing map to such a circle domain is unique up to Möbius transformations. A method to construct discrete Riemann maps is described in [4] (Section 3.3) for triangulations and for mostly quadrilateral meshes in the previous Section 4.2. Having generalized the notion of discrete conformal equivalence from triangulations to cyclic polyhedral surfaces, it is straightforward to adapt this method to construct discrete maps to circle domains:

- (1) Fill holes by gluing faces to all but one boundary component, so that the resulting surface is homeomorphic to a disk.
- (2) Construct the discrete Riemann map.
- (3) Remove the faces that were added in step (1).

Figure 9 shows an example.

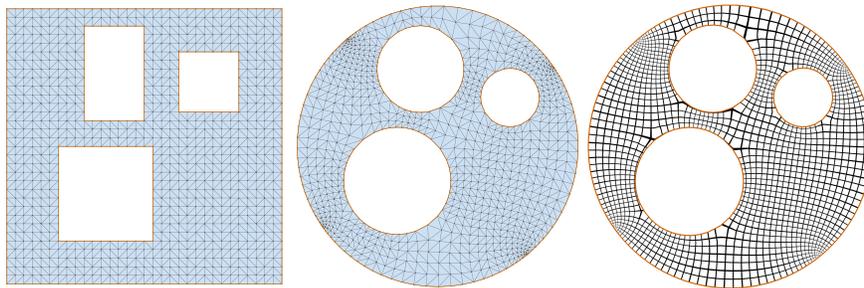


Fig. 9 Discrete conformal map of a multiply-connected domain (left) to a circle domain (middle). The images of vertical and horizontal “parameter lines” are shown on the right.

5.2 *Special slit domains*

Any multiply connected domain can be mapped to the complex plane with parallel slits [20]. In principle, it is possible to construct discrete conformal maps that map holes to slits by solving Problem 3.1. On each boundary component that should be mapped to a slit, set the desired total angle $\Theta = 2\pi$ for the two vertices that should be mapped to the endpoints of the slit, and set $\Theta = \pi$ for all other vertices on that boundary component. However, this will not work in general. While the resulting surface will be flat, the developing map to the plane will in general have translational monodromy for a cycle around the hole. The surface will only close up in the plane if the vertices that should be mapped to the endpoints of the slit are chosen exactly right. (This will in general require modifying the original mesh.)

Sometimes, the symmetry of the problem determines the right positions of the end-vertices, so that discrete conformal maps to slit surfaces can be computed. The first two rows of Figure 10 show examples. The bottom row visualizes a discrete conformal map where circular holes are mapped to slits. Here, we use the following trick: We start with the slit surface and map it to a surface with circular holes as described in Section 5.1.

6 Uniformization of spheres

This section is concerned with discrete conformal maps of polyhedral surfaces of genus 0 onto the round sphere. For triangulations, this is described in [4] (Section 3.2). In Section 6.1, we adapt this method to quadrangulations. This is similar to the discrete Riemann mapping with quadrilaterals described in Section 4.2. Effectively, this method reduces the problem to minimizing the convex euclidean functional E^{euc} . The spherical version of the variational principle (Theorem 3.7) involves the non-convex function E^{sph} . It is not as practical for calculations, because one has to find a saddle point instead of a minimum. Nevertheless, the spherical functional can often be used to calculate maps to the sphere. This is explained in Section 6.2.

6.1 *Uniformizing quadrangulations of the sphere*

Suppose $(\Sigma, \ell)_{euc}$ is a cyclic polyhedral surface with quadrilateral faces that is homeomorphic to the sphere.

- (1) Choose a vertex $k \in V_\Sigma$.
- (2) Apply a discrete conformal change of metric (3) such that all edges incident with k have the same length. One may choose $u = 0$ for all vertices except the neighbors of k . It does not matter if polygon inequalities are violated after this step.

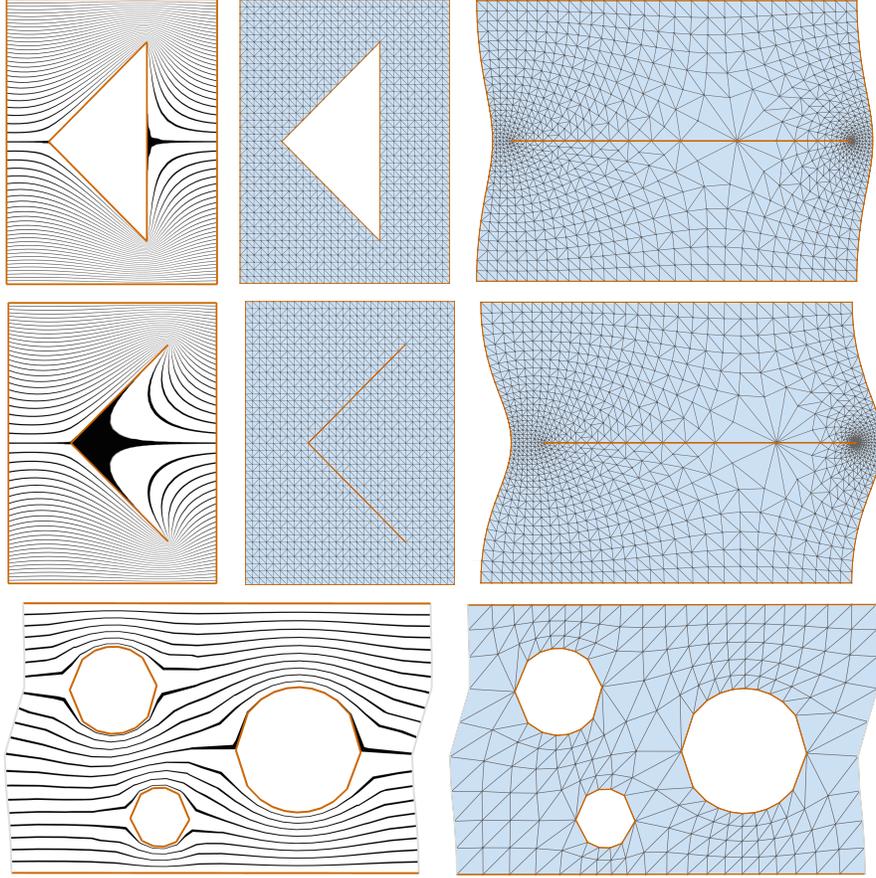


Fig. 10 Mapping surfaces with holes to slit surfaces. In all images, the left and right parts of the boundary are identified by a horizontal translation. Preimages of horizontal lines visualize the flow of an incompressible inviscid fluid around the hole in a channel with periodic boundary conditions. Top row: A cylinder with a triangular hole is mapped to a cylinder with a slit. One vertex of the triangle and the midpoint of the opposite side are mapped to the endpoints of the slit. Middle row: An arrow shaped slit is mapped to a straight slit. The two vertices at the arrows tip, on either side of the slit, are mapped to the endpoints of the straight slit. Bottom row: Three circular boundary components are mapped to horizontal slits. (The slit surface is not shown.)

- (3) Let (Σ', ℓ') _{euc} be the complex obtained by removing vertex k and all incident quadrilaterals.
- (4) Solve Problem 3.2 for (Σ', ℓ') _{euc} with prescribed total angles $\Theta_i = 2\pi$ for interior vertices of Σ' , $\Theta_i = \pi$ for boundary vertices of Σ' that were not neighbors of k in Σ , and fixed scale factors $u_i = 0$ for vertices that were neighbors of k in Σ . The result is a planar polyhedral surface with cyclic quadrilaterals. Consecutive boundary edges that belonged to a face incident with vertex k in Σ are contained in a straight lines.

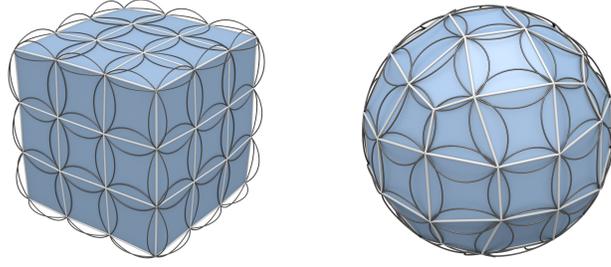


Fig. 11 Discrete conformal map from the cube to the sphere, calculated with the method described in Section 6.1. We apply Möbius normalization (Section 6.3) to the polyhedral surface with vertices on the sphere to achieve rotational symmetry.

- (5) Map the vertices to the unit sphere by stereographic projection and reinsert the vertex k at the image point of ∞ .
- (6) Optionally apply Möbius normalization, see Section 6.3.

Proposition 6.1. *The result is a cyclic polyhedral surface with vertices on the unit sphere that is discretely conformally equivalent to $(\Sigma, \ell)_{euc}$.*

This can be seen in the same way as the corresponding statement about discrete Riemann maps with quadrilaterals (Proposition 4.2). Figure 11 shows a discrete conformal map calculated by this method.

6.2 Using the spherical functional

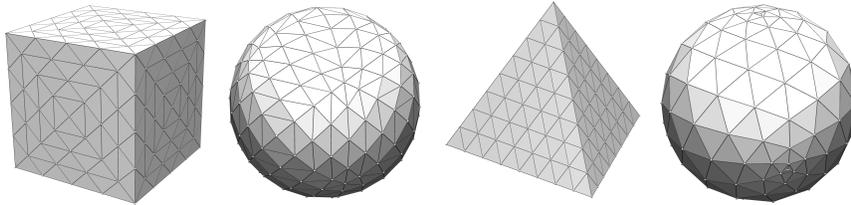


Fig. 12 Mapping conformally to the sphere using the spherical functional. The spherical surfaces are Möbius normalized to achieve rotational symmetry.

It is possible to use the spherical functional E^{sph} to calculate maps to the sphere even though it is not convex. For simplicity, we consider only triangulations, so all λ variables are fixed and we may consider E^{sph} as function of the logarithmic scale factors u only (see Section 3.3). A numerical method has to find a saddle point of $E^{sph}(u)$.

Note that the scaling direction $1_{V_\Delta} \in \mathbb{R}^{V_\Delta}$ is a negative direction of the Hessian at a critical point: Suppose $(\Delta, \ell)_{sph}$ is a spherical triangulation with the desired angle sum Θ_i at each vertex i . Then $0 \in \mathbb{R}^{V_\Delta}$ is a critical point of $E_{\Delta, \Theta}^{sph}(u)$. If we enlarge all edge lengths by a common factor $e^h > 1$, then all angles become larger, so every component (30) of the gradient of E^{sph} becomes negative. Following the negative gradient would make all lengths even larger.

The following minimax method works in many cases. Define the function \tilde{E} by *maximizing* the functional E^{sph} in the scaling direction,

$$\tilde{E}(u) = \max_{h \in \mathbb{R}} \left\{ E^{sph}(u + h1_{V_\Delta}) \right\} \quad (46)$$

Minimize functional \tilde{E} in a hyperplane of \mathbb{R}^{V_Δ} transverse to the direction 1_{V_Δ} .

Figure 1 (top) and Figure 12 show examples of discrete conformal maps to polyhedral surfaces inscribed in a sphere that were calculated using this method.

6.3 Möbius normalization

The notion of discrete conformal equivalence of euclidean polyhedral surfaces $(\Sigma, \ell)_{euc}$ in \mathbb{R}^3 is Möbius invariant (Proposition 2.5). If all vertices $v \in V_\Sigma$ are contained in the unit sphere $S^2 \subset \mathbb{R}^3$, then there is a Möbius transformation T of S^2 such that the center of mass of the transformed vertices is the origin [27],

$$\sum_{v \in V_\Sigma} T(v) = 0.$$

The Möbius transformation T is uniquely determined up to post-composition with a rotation around the origin.

The following method can be used to calculate such a Möbius transformation: Find the unique minimizer of the function δ defined below. Then choose for T a Möbius transformation that maps S^2 to itself and the minimizer to the origin. Here, we only provide explicit formulas for the function δ and its first two derivatives. For a more detailed account, we refer the reader to [27]. The function δ is defined on the open unit ball in \mathbb{R}^3 by

$$\delta(x) = \sum_{v \in V} \log \left(\frac{-\langle x, v \rangle}{\sqrt{-\langle x, x \rangle}} \right), \quad (47)$$

where

$$\langle x, y \rangle = x_1 y_1 + x_2 y_2 + x_3 y_3 - 1. \quad (48)$$

The gradient and Hessian matrix of δ are

$$\text{grad } \delta(x) = \sum_{v \in V} \left(\frac{v}{\langle x, v \rangle} - \frac{x}{\langle x, x \rangle} \right), \quad (49)$$

$$\text{Hess } \delta(x) = \sum_{v \in V} \left(2 \frac{x^T x}{\langle x, x \rangle^2} - \frac{v^T v}{\langle x, v \rangle^2} - \text{diag} \left(\frac{1}{\langle x, x \rangle} \right) \right). \quad (50)$$

7 Uniformization of tori

Every Riemann surface R of genus one is conformally equivalent to a flat torus, i.e., to a quotient space \mathbb{C}/Γ , where $\Gamma = \mathbb{Z}\omega_1 + \mathbb{Z}\omega_2$ is some two-dimensional lattice in \mathbb{C} . The biholomorphic map from R to \mathbb{C}/Γ , or from the universal cover of R to \mathbb{C} , is called a uniformizing map. For a polyhedral surface of genus one, constructing a discrete uniformizing map amounts to solving Problem 3.1 with prescribed total angle $\Theta = 2\pi$ at all vertices. This provides us with a method to calculate approximate uniformizing maps for Riemann surfaces of genus one given in various forms. We consider examples of tori immersed in \mathbb{R}^3 in Section 7.1 and elliptic curves in Section 7.2. (We will also consider tori in the form of Schottky uniformization in Section 8.2, as a toy example after treating the higher genus case.)

The belief that discrete conformal maps approximate conformal maps is not based on a proven theorem but on experimental evidence like the Wente torus example of Section 7.1 and the numerical experiments of Section 7.4.

7.1 Immersed tori

First we consider a simple example with quadrilateral faces. Figure 13 (left) shows a coarse discretization of a torus. The faces are isosceles trapezoids, so they are inscribed in circles. On the right, the figure shows the uniformization obtained by solving Problem 3.1 with prescribed total angle $\Theta = 2\pi$ at all vertices.

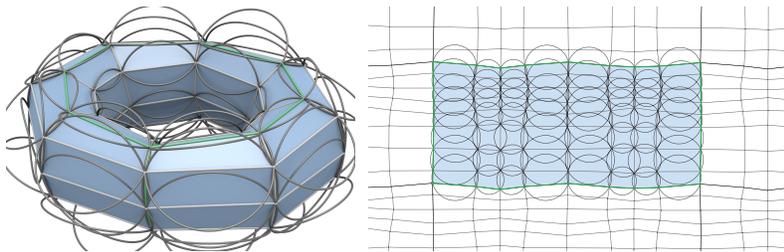


Fig. 13 Uniformization of an immersed torus with cyclic quadrilateral faces

To test the numerical accuracy of our discrete uniformizing maps, we consider the famous torus of constant mean curvature discovered by Wente [29]. Explicit doubly periodic conformal immersion formulas (i.e., formulas for the inverse of a uniformizing map) are known in terms of elliptic functions [1, 28, 3].

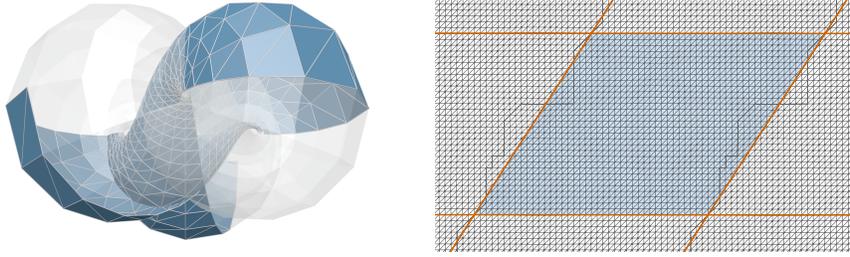


Fig. 14 Uniformization of the Wente torus

Figure 14 (left) shows a triangulated model of the Wente torus constructed by sampling an explicit immersion formula [3] on a nearly square lattice containing the period lattice Γ . On the right, the figure shows the discrete uniformization, which reproduces the regular lattice Γ to high accuracy. The modulus $\tau = \omega_2/\omega_1$ of the Wente torus has been determined numerically [11] as $\tau = 0.41300\dots + i0.91073\dots$. The modulus of the discrete uniformization of the discretized surface shown in the figure is $\tilde{\tau} = 0.41341\dots + i0.91061\dots$.

7.2 Elliptic curves

An algebraic curve of the form

$$\mu^2 = a \prod_{j=1}^k (\lambda - \lambda_j), \quad (51)$$

where the $\lambda_j \in \mathbb{C}$ are distinct and $k = 3$ (an elliptic curve) or $k = 4$ (with the singularity at infinity resolved), represents a Riemann surface of genus one as branched double cover of the λ -sphere \mathbb{CP}^1 , which we identify conformally with the unit sphere $S^2 \subset \mathbb{R}^3$. The branch points are $\lambda_1, \lambda_2, \lambda_3, \infty$ if $k = 3$ and $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ if $k = 4$. Every Riemann surface of genus one can be represented in this way.

We construct a discrete model for a double cover of S^2 branched at four points $\lambda_1, \dots, \lambda_4$ in the following way. Choose n other points $p_1, \dots, p_n \in S^2$ and let P be the boundary of the convex hull of the points $\{\lambda_1, \dots, \lambda_4, p_1, \dots, p_n\}$. Then P is a convex polyhedron with $n + 4$ vertices and with faces inscribed in circles. (Generically, the faces will be triangles. In Section 7.3 we explain the method we used to obtain “good” triangles.) Find two disjoint simple edge paths γ_1, γ_2 joining the

branch points λ_j in pairs. Take a second copy \hat{P} of the polyhedron P . Cut and glue P and \hat{P} along the paths γ_1, γ_2 to obtain a polyhedral surface of genus 1. Uniformize it by solving Problem 3.1. One obtains a discrete conformal map to a flat torus, whose inverse can be seen as a discrete elliptic function. Figure 15 shows examples. We will treat hyperelliptic curves in a similar fashion in Section 8.3.

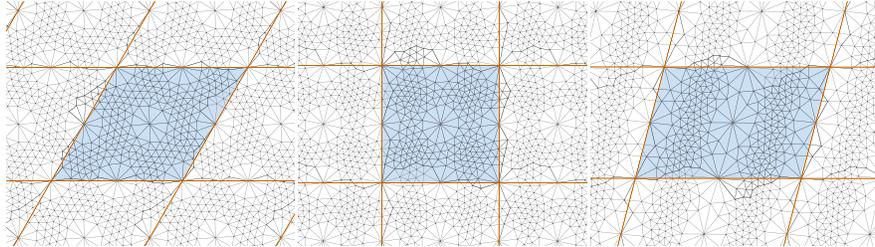


Fig. 15 Discrete uniformization of elliptic curves. Left: If the branch points in S^2 are the vertices of a regular tetrahedron, period lattice is very close to a hexagonal lattice. Middle: If the branch points form a square on the equator, the period lattice is very close to a square lattice. Right: an example with branch points in unsymmetric position.

Remark 7.1. Instead of constructing a doubly covered convex euclidean polyhedron with vertices on the unit sphere as described above, one could also construct a spherical triangulation of the doubly covered sphere that is invariant under the elliptic involution (exchanging sheets). These two approaches are in fact equivalent due to Remark 2.3.

Mapping a flat torus to an elliptic curve. We can also go the opposite way, mapping a flat torus to a double cover of S^2 . Start with a triangulated flat torus. The triangulation should be symmetric with respect to the elliptic involution, i.e., symmetric with respect to a half turn around one vertex (which is then also a half turn around three other vertices). The quotient space of the triangulated torus modulo the elliptic involution is then a triangulated sphere. Map it to the sphere by the procedure explained in [4] (Section 3.2), see also Section 6.1 of the present article.

Figures 16 and 17 show examples where we started with a hexagonal and a square torus respectively.

7.3 Choosing points on the sphere

The uniformization procedure for elliptic curves described in Section 7.2 requires choosing points on the sphere in addition to the four given branch points. For numerical reasons, these points should be chosen so that taking the convex hull leads to triangles that are close to equilateral. We obtained good triangulations by minimizing the following energy for n points in \mathbb{R}^3 while fixing the subset of branch

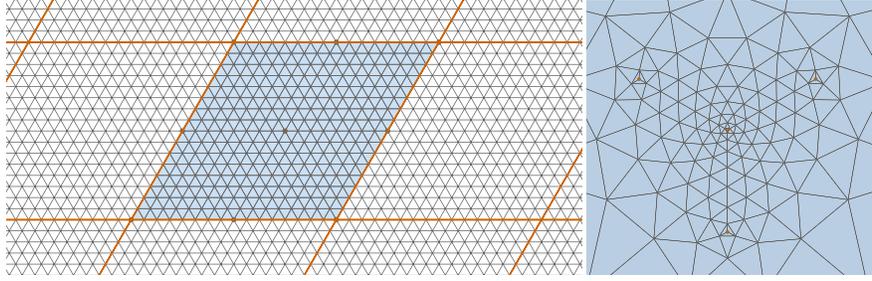


Fig. 16 Mapping the hexagonal torus $\mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$, $\tau = \frac{1}{2} + i\frac{\sqrt{3}}{2}$ (left) to a double cover of the sphere (right). Because the regular triangulation of the torus on the left is symmetric with respect to the elliptic involution, its image projects to a triangulation of the sphere seen on the right.

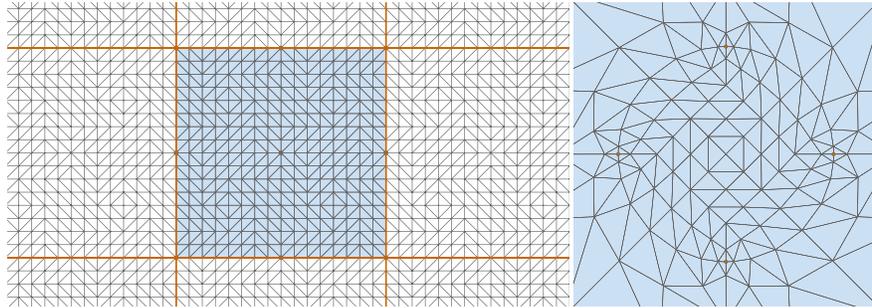


Fig. 17 Mapping the square torus $\mathbb{C}/(\mathbb{Z} + i\mathbb{Z})$ (left) to a double cover of the sphere (right). Again, the triangulation on the left is symmetric with respect to the elliptic involution, so the image on the right projects to a triangulation of the sphere.

points:

$$\mathcal{E} = n^2 \sum_{v \in V} (\langle v, v \rangle - 1)^2 + \sum_{\substack{v, w \in V \\ w \neq v}} \frac{1}{\langle w - v, w - v \rangle}, \quad (52)$$

where $\langle \cdot, \cdot \rangle$ denotes the standard euclidean scalar product of \mathbb{R}^3 . We do not enforce the constraint that the points should lie in the unit sphere S^2 . Instead, we simply project back to S^2 after the optimization.

As initial guess we choose points uniformly distributed in S^2 . To achieve this we choose points with normally distributed coordinates and project them to S^2 [19].

7.4 Numerical experiments

Given the branch points of an elliptic curve, the modulus τ can be calculated in terms of hypergeometric functions. In this section, we compare the theoretical value

of τ with the value $\hat{\tau}$ that we obtain by the discrete uniformization method explained in Section 7.2.

We consider elliptic curves in Weierstrass normal form

$$\begin{aligned}\mu^2 &= 4(z - \lambda_1)(z - \lambda_2)(z - \lambda_3) \\ &= 4z^3 - g_2z - g_3,\end{aligned}\tag{53}$$

so the branch points $\lambda_1, \lambda_2, \lambda_3, \infty$ satisfy $\lambda_1 + \lambda_2 + \lambda_3 = 0$, and

$$g_2 = -4(\lambda_1\lambda_2 + \lambda_2\lambda_3 + \lambda_3\lambda_1), \quad g_3 = 4\lambda_1\lambda_2\lambda_3.\tag{54}$$

We calculate the modulus τ with Mathematica using the built-in function `WeierstrassHalfPeriods[{g2, g3}]`. We normalize τ and the value $\hat{\tau}$ obtained by discrete uniformization so that they lie in the standard fundamental domain of the modular group, $|\tau| > 1$ and $|\operatorname{Re}(\tau)| < \frac{1}{2}$, and we consider the error $|\tau - \hat{\tau}|$. (We stay away from the boundary of the fundamental domain.)

Subdivided icosahedron. In this experiment we start with the twelve vertices of a regular icosahedron and choose the branch points $\lambda_1, \dots, \lambda_4$ among them. The remaining points act the role of p_1, \dots, p_n . To study the dependence of $|\tau - \hat{\tau}|$ on the number of points we repeatedly subdivide all triangles into four similar triangles and project the new vertices to S^2 . The number of vertices grows exponentially while the triangles remain close to equilateral. Figure 18 shows the result of this experiment. It suggests the error behaves like

$$|\tau - \hat{\tau}| = \mathcal{O}(n^\alpha), \quad \alpha \approx -0.88.\tag{55}$$

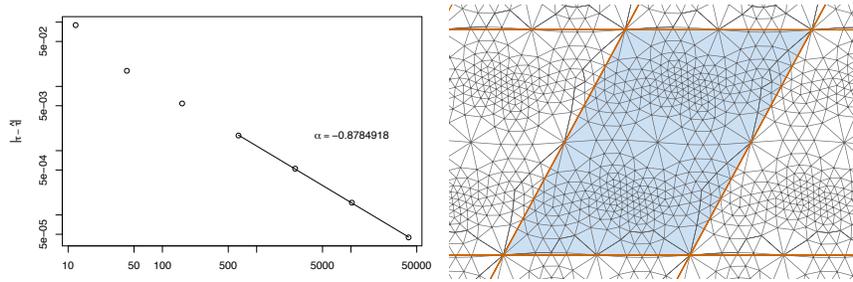


Fig. 18 Left: Error for zero to six subdivision steps. The log-log plot shows the error $|\tau - \hat{\tau}|$ against the number of vertices of the subdivided icosahedron (i.e., in one sheet of the doubly covered sphere). To estimate the asymptotic behavior of the error, we determine the slope $\alpha \approx -0.88$ of a line through the last four points by linear regression. Right: Result of the discrete uniformization after two subdivision steps.

Dependence on mesh quality. In the second experiment we choose the additional points randomly to analyze how the quality of the triangulation affects the approxi-

mation error. We use the following quantities to measure the quality of the triangulation based on length cross-ratios for edges:

$$Q_{\text{lcr}}(e) := \frac{1}{2} \left(\text{lcr}(e) + \frac{1}{\text{lcr}(e)} \right) - 1,$$

$$Q_{\text{lmr}}(f) := \frac{1}{2} \left(\text{lmr}(f) + \frac{1}{\text{lmr}(f)} \right) - 1,$$

where lcr denotes the length cross-ratio (9) of an edge, and lmr denotes the length multi-ratio defined for faces by $\text{lmr}(f) = \prod_{e \in f} \text{lcr}(e)$. If $Q_{\text{lcr}} = 0$ for all edges, then the mesh is discretely conformally equivalent to a mesh consisting of equilateral triangles. So less is better for these quality measures. To get enough “good” triangulations in our samples, we improve random meshes with the procedure described in Section 7.3.

Figure 19 (left) shows a plot of 2600 triangulations ranging from $n = 20$ to $n = 1500$ vertices. No clear convergence rate is discernible. The situation improves when only samples with a certain minimal mesh quality are considered. For the plot in Figure 19 (right) we selected only triangulations with $\max_e \{Q_{\text{lmr}}(e)\} < 0.3$. (The results are similar when using the quality measures $\max_e \{Q_{\text{lcr}}(e)\} < x$ or $\text{mean}_e \{Q_{\text{lcr}}(e)\} < x$.)

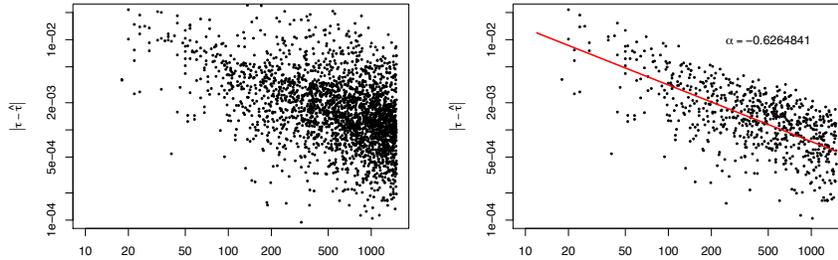


Fig. 19 Left: log-log plot of the error $|\tau - \hat{\tau}|$ against the number of vertices for a sample of optimized random triangulations with no quality constraint. Right: Only triangulations with $\max_e \{Q_{\text{lmr}}(e)\} < 0.3$ are considered. The regression line with slope $\alpha \approx -0.63$ is shown in red.

The results from these two experiments suggest that the error depends on the number of n of vertices asymptotically like n^α , where the exponent $\alpha < 0$ depends on the mesh.

7.5 Putting a square pattern on a spherical mesh

We can use a variant of the discrete uniformization of elliptic curves (Section 7.2) to put a square pattern on a surface that is homeomorphic to a sphere. Figure 20 shows an example.

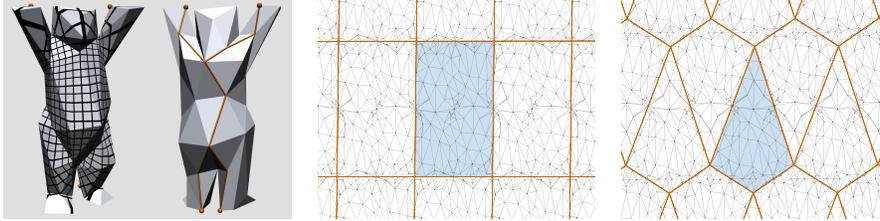


Fig. 20 The discrete “Berlin Buddy Bear”, a mascot of the SFB/Transregio 109 “Discretization in Geometry and Dynamics”. The square pattern is put on a bear model as described in Section 7.5. Four ramification vertices (marked in red) are chosen at the paws. The uniformization of the branched double cover is shown in the middle. Each fundamental domain covers the bear twice. Fundamental domains of the group generated by rotations around the branch points are shown on the right. Each covers the bear once.

Pick four vertices of the mesh as ramification points and create a two-sheeted branched cover of the mesh by gluing two copies along paths connecting the selected vertices. The resulting surface is a torus. It can be uniformized using the euclidean functional. The uniformizing group is generated by two translations. This group is a subgroup of the group generated by rotations around the branch vertices. Hence we can achieve the same result as follows. Instead of doubling the surface, prescribe total angles $\Theta = \pi$ at the ramification vertices and $\Theta = 2\pi$ at all other vertices. The result is a flat surface with four cone-like singularities of cone-angle π . The monodromy of the developing map is generated by half-turns. Avoiding the double cover is more efficient because one only has to minimize a function of (approximately) half the number of variables.

8 Uniformization of surfaces of higher genus

As in the case of tori (Section 7), we can find uniformizing maps for cyclic polyhedral surfaces of genus $g \geq 2$ by solving the hyperbolic version ($\tilde{g} = hyp$) of Problem 3.1 with prescribed total angle $\Theta = 2\pi$ at all vertices. (We will only consider triangulations in the following.) This allows us to calculate approximate uniformizations for Riemann surfaces of genus $g \geq 2$ given in various forms, by approximating them with polyhedral surfaces.

In Section 8.1, we briefly discuss how to construct fundamental polygons and group generators.

Not much needs to be said about the uniformization of immersed surfaces. Examples are shown in Figures 1 (bottom), 21, and 22. In Section 5.1 we discussed

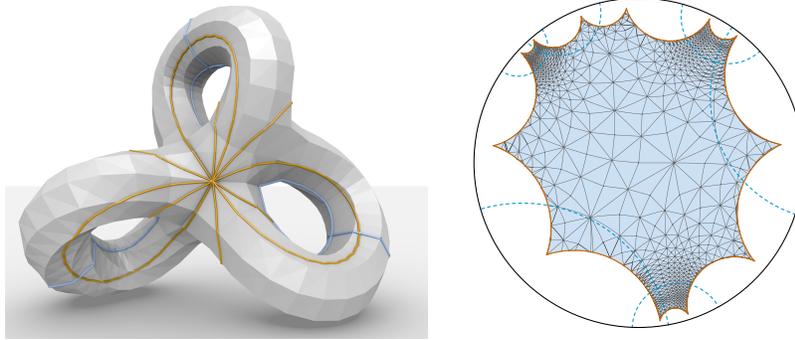


Fig. 21 Discrete uniformization of an embedded triangulated surface of genus 3. A fundamental polygon with “canonical” edge pairing is shown on the right together with the image mesh. The edges of the polygon (brown) and the axes of the edge-pairing translations (blue) are pulled back to the embedded surface shown on the left.

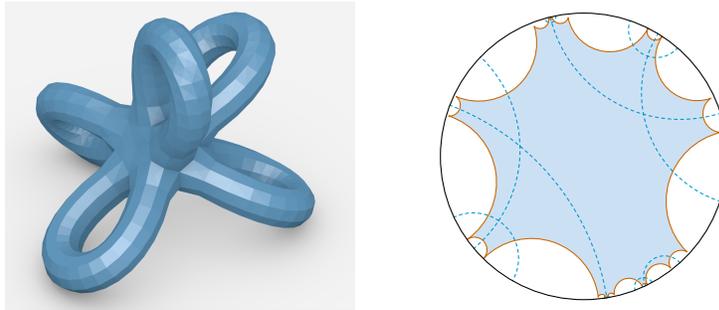


Fig. 22 Left: An embedded triangulated surface of genus 5. Right: Fundamental polygon with non-canonical edge-pairing. The axes of the edge pairing translations are shown in blue.

mappings from multiply connected domains to circle domains. Analogously, one can construct uniformizations of polyhedral surfaces of genus $g \geq 2$ with holes over the hyperbolic plane with circular holes. An example is shown in Figure 23. More precisely, the holes are bounded by hyperbolic polygons with vertices on a circle.

We explain how to calculate the Fuchsian uniformization for Riemann surfaces given in the form of a Schottky uniformization in Section 8.2. We discuss the uniformization of hyperelliptic curves in Section 8.3 and a geometric characterization of hyperelliptic Riemann surfaces in Section 8.4.

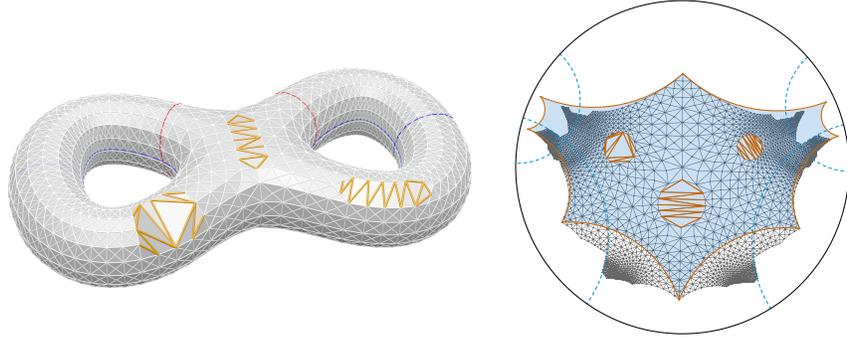


Fig. 23 Uniformization of a genus 2 surface with three boundary components over the hyperbolic plane with three circular holes. The three holes are filled with polygons, which are then triangulated during the calculation, see Sections 5.1 and 3.

8.1 Fundamental polygons and group generators

Basic facts and notation. Every compact Riemann R of genus $g \geq 2$ can be represented as the quotient of the hyperbolic plane H^2 modulo the action of a discrete group G of hyperbolic translations,

$$R = H^2/G. \quad (56)$$

Presentations of the group G play an important role. We will denote generators by capital letters and their inverses by primes,

$$A' := A^{-1} \in G. \quad (57)$$

The uniformization group G can be presented with a finite set of generators

$$A, B, C, D, \dots \in \text{Isom}(H^2)$$

subject to a single relation $r = 1$,

$$G = \langle A, B, C, D, \dots \mid r = 1 \rangle, \quad (58)$$

where r is a product in which all generators and their inverses appear exactly once. Such presentations are closely related with fundamental polygons: Every fundamental polygon in which all vertices are identified leads to such a presentation.

A fundamental domain of G is an open connected subset D of the hyperbolic plane such that the G -orbit of the closure \bar{D} covers H^2 , and $gD \cap D = \emptyset$ for all $g \in G \setminus \{1\}$. A fundamental polygon of G is a fundamental domain with polygonal boundary, i.e., the boundary consists of geodesic segments, the *edges* of the fundamental polygon, which are identified in pairs by the action of the group G . For each edge a , there is exactly one partner edge a' such that there exists a translation

$A \in G$ mapping a to a' . These edge gluing translations form a generating set of G . If all vertices of the fundamental polygon are identified (i.e., they belong to the same G -orbit), then the fundamental polygon has $4g$ edges. In this case there is only one relation for these generators. The relation can be determined from the edge labels, which we always list in counterclockwise order. For example, if the edges of an octagon are labeled “canonically”,

$$aba'b'cdc'd', \quad (59)$$

then the relation for the corresponding edge pairing translations is

$$DC'D'CB'A, \quad (60)$$

and if opposite edges are identified,

$$abcd'a'b'c'd', \quad (61)$$

then then relation is

$$DC'BA'D'CB'A = 1. \quad (62)$$

Computational aspects. Let (Σ, ℓ) be a closed (euclidean, spherical, or hyperbolic) triangulated surface of genus $g \geq 2$. We solve Problem 3.1 to obtain a combinatorially equivalent hyperbolic triangulated surface $(\Sigma, \tilde{\ell})_{hyp}$ with angle sum $\Theta = 2\pi$ at every vertex. We lay out the triangles in the hyperbolic plane one-by-one, following a breadth-first search of the the 1-skeleton of the dual cell complex of Σ . (Alternatively, one could use a shortest spanning tree of the 1-skeleton of the dual complex [7].) The result is a fundamental polygon with many vertices. An example is shown in Figure 24 (a).

We simplify this fundamental polygon by connecting vertices that are identified with more than one partner by geodesic arcs, as shown in Figure 24 (b). The resulting polygon has in general more than one vertex class.

Now we perform the standard algorithm involving cut-and-glue operations (see, e.g., [12]) to obtain a fundamental polygon with one vertex class and so-called canonical edge identification

$$aba'b'cdc'd' \dots . \quad (63)$$

During this process we maintain edge-identification transformations, which we represent as $SO^+(2, 1)$ matrices.

Hyperbolic translations tend to accumulate numerical errors quite fast when building products. The situation could be ameliorated somewhat by using the $PSL(2, \mathbb{R})$ -representation of hyperbolic isometries [8], but the fundamental problem remains. For this reason, it is desirable to perform the cut-and-glue algorithm in such a way that the number of matrix products required to maintain the gluing translations is small. We follow the following greedy approach. Repeatedly we have to choose labels x, y such that the order in the polygon is $x \cdots y \cdots x' \cdots y'$, and then perform a cut-and-glue sequence to bring the labels next to each other, $xyx'y'$. We

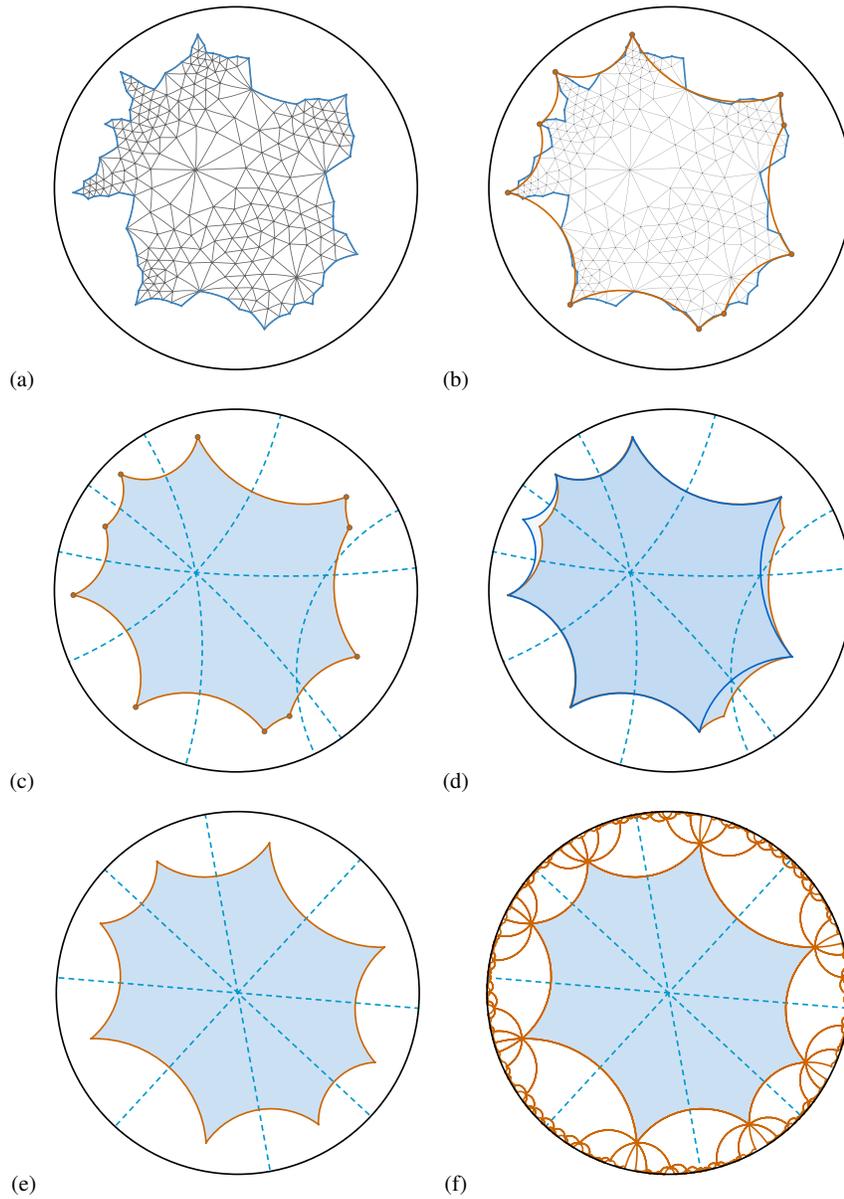


Fig. 24 Constructing a fundamental polygon with opposite edges identified. (a) Laying out hyperbolic triangles creates a fundamental polygon with many vertices. (b) Straighten the edges between vertices that are identified with more than one partner (shown in red). (c) Axes of the edge-pairing translations are shown in blue. (d,e) Two cut-and-paste operations lead to a fundamental polygon with one vertex class and opposite edges identified. The axes intersect in one point (see Section 8.4). We move this point to the origin. (f) Tiling the hyperbolic plane with fundamental polygons.

always choose a pair x, y , for which this requires the minimal number of matrix multiplications.

The polygon with canonical edge identifications may not be convex. Following Keen [13] we can transform this domain into a strictly convex fundamental polygon by choosing a different base vertex for the same group of transformations. Let $aba'b'cdc'd' \dots$ be a fundamental polygon and let

$$G = \langle A, B, C, D, \dots \in PSL(2, \mathbb{R}) \mid \dots DC'D' CBA'B'A = 1 \rangle \quad (64)$$

be the corresponding presentation of the uniformization group, see Figure 25 (left). Then the axes of the generators A and B intersect in a point p_0 . Choosing p_0 as the base point of a new fundamental polygon as shown in Figure 25 (right) renders it convex and uniquely determined for the given group and presentation.

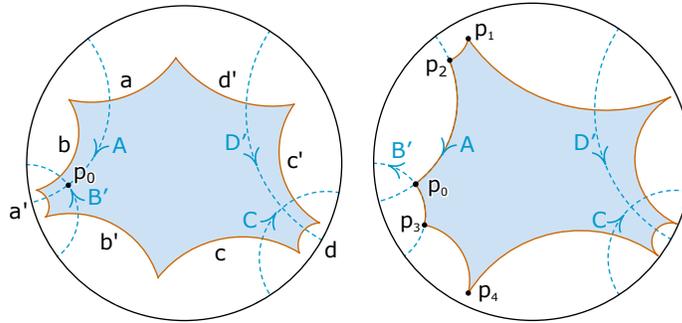


Fig. 25 The algorithm of Linda Keen to construct strictly convex fundamental polygons. Start with any canonical fundamental polygon $aba'b'cdc'd'$ with a corresponding relation $DC'D' CBA'B'A = 1$ (left). We choose the intersection p_0 of the axes of transformations A and B as base point for the new domain. The new vertices of the fundamental domain are calculated as $p_1 = A'Bp_0$, $p_2 = A'p_0$, $p_3 = Bp_0$, and $p_4 = BA'p_0$. The other vertices are obtained similarly from p_4 by applying C and D .

Fundamental polygons with opposite sides identified. When we consider the geometric characterization of hyperelliptic surfaces in Section 8.4, we want to transform fundamental polygons into fundamental polygons with opposite sides identified, i.e., polygons with edge labeling

$$abcd \dots a'b'c'd' \dots$$

Any fundamental polygon can be transformed into a fundamental polygon with opposite edges identified by cut-and-glue operations: First transform the polygon to canonical form $aba'b'cdc'd' \dots$ by the standard algorithm. Playing a sequence of steps that transforms a polygon with opposite edges identified to canonical form backwards, transforms the canonical polygon to a polygon with opposite edges identified.

This algorithm is not optimal with respect to the number of multiplications necessary to maintain the edge-gluing translations. Especially if the original polygon is already “close” to one with opposite sides identified, the detour via a canonical polygon is inefficient.

In all examples, we use a heuristic method based on the following idea: Find a longest sequence of different letters in the edge labeling (ignoring primes), and then try to move a different letter into this sequence by cutting and gluing.

8.2 From Schottky to Fuchsian uniformization

In this section, we consider Riemann surfaces presented as quotient spaces of classical Schottky groups.

Definition 8.1. Let $C_1, C'_1, \dots, C_g, C'_g$ be circles in $\hat{\mathbb{C}}$ that bound disjoint disks. A *classical Schottky group* is a Kleinian group generated by Möbius transformations $\sigma_1, \dots, \sigma_g$, where σ_j maps the outside of C_j onto the inside of C'_j .

Each generator σ_j has fixed points A_j, B_j inside C_j and C'_j , respectively. The limiting set A of G is the union of orbits of the fixed points A_j, B_j . G acts freely and properly discontinuously on the domain of discontinuity $\Omega = \hat{\mathbb{C}} \setminus A$. The quotient space $R = \Omega/G$ is a Riemann surface of genus g . The domain outside all of the circles is a fundamental domain of G . The identified pairs of circles form handles.

We discretize the Riemann surface $R = \Omega/G$ determined by a classical Schottky group G as follows. First, construct a triangulation of Ω whose vertex set and combinatorics are invariant under the action of G . (Ignore the fact that a Möbius transformation maps straight edges to circular arcs as in Proposition 2.5 on the Möbius invariance of conformal classes.) For example, the triangulation may be the Delaunay triangulation of a G -invariant point set. The following construction avoids Delaunay triangulations of infinite (but symmetric) point configurations:

If necessary, choose a Möbius normalization for which the fundamental domain is bounded in \mathbb{C} . For each pair of circles C_j, C'_j we construct polygons p_{1j}, \dots, p_{n_jj} inscribed in C_j and $p'_{1j}, \dots, p'_{n_jj}$ inscribed in C'_j such that $\sigma_j(p_{kj}) = p'_{kj}$. For example, we may choose a regular n -gon inscribed in C_j and map the vertices by σ_j to C'_j . Triangulate the compact region bounded by these polygons, adding vertices in the interior as wanted. (For example, use a constrained Delaunay triangulation.) The images of this triangulation under the action of G (again, considering only combinatorics and vertex positions) form a G -invariant triangulation $\hat{\Delta}$ of the universal cover of R , hence a triangulation Δ of R . More precisely, the triangulations $\hat{\Delta}$ and Δ are only defined up to isotopy fixing the vertices. The edge-lengths $\hat{\ell}$ (distances of vertices) do not project from $\hat{\Delta}$ to Δ , but the length cross-ratios \hat{lcr} calculated from these edge lengths do, because they are Möbius invariant. The projected length cross-ratios lcr determine a discrete conformal class for Δ (see Section 2.5).

To obtain a Fuchsian uniformization of R , construct edge lengths ℓ from the length cross-ratios lcr as described in Section 2.6. Then solve Problem 3.1 (or rather

the corresponding analytic version, Problem 3.4) for $(\Delta, \ell)_{euc}$ with $\tilde{g} = hyp$ and desired angle sums $\Theta = 2\pi$ at all vertices.

Note that the lengths ℓ calculated from length cross-ratios lcr may not satisfy all triangle inequalities. This does not matter for the corresponding analytic Problem 3.4 (with $V = V_1, E = E_1$). If Problem 3.4 has a solution, it is in the discrete conformal class determined by the length cross-ratios lcr . Also, whether or not Problem 3.4 has a solution does not depend on the choice of edge lengths ℓ provided they lead to the same length cross-ratios.

Figure 26 shows an example of the Fuchsian uniformization of a genus three surface presented by its Schottky uniformization.

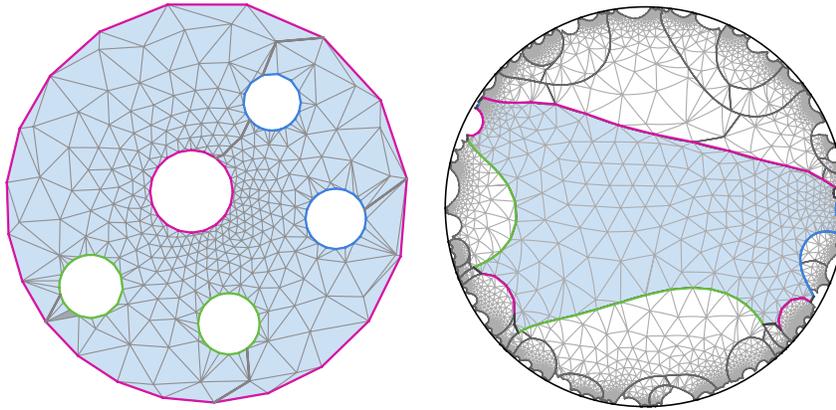


Fig. 26 Discrete Riemann surface of genus 3 given by Schottky data (left) and its Fuchsian uniformization (right). Circles with the same color are identified. The extra points of the triangulation are chosen so that the triangles are close to equilateral where possible. The shaded region in the right image corresponds to the fundamental domain of the Schottky group in the left image. Its boundary consist of curves corresponding to the circles and curves corresponding to lines connecting the circles (drawn in gray).

Tori given by Schottky data. For tori, the Schottky data consist of one generator

$$\frac{\sigma(z) - A}{\sigma(z) - B} = \mu \frac{z - A}{z - B} \tag{65}$$

and one pair of circles. To find a uniformization \mathbb{C}/Γ is elementary. It suffices to consider the case where $A = B = 0$ (and C, C' are concentric circles around 0 with radii i and μ). Figure 27 shows two examples where we apply the discrete method without adding extra points inside the fundamental domain of the Schottky group.

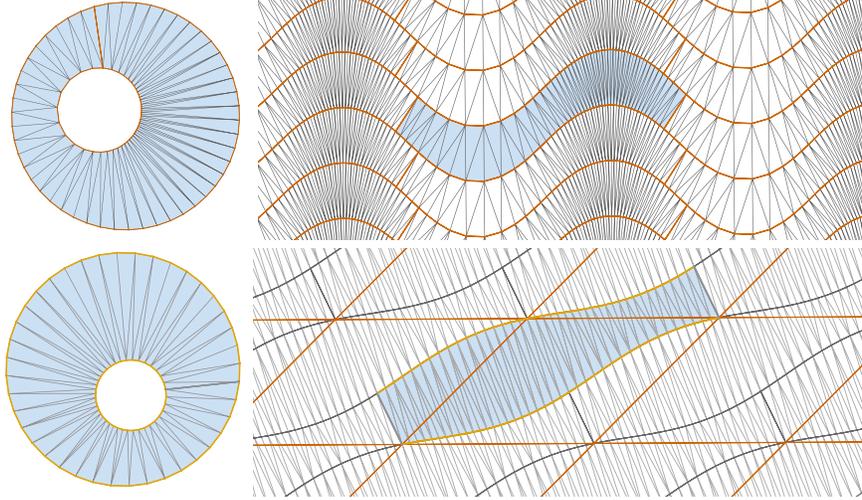


Fig. 27 Left: Fundamental domains of Riemann surfaces of genus 1 given by Schottky data C, C', A, B, μ (see (65)). The triangulations use only points on the circles C, C' . On purpose we have not chosen concentric circles with center $A = B$. Right: Representation of the same surfaces as \mathbb{C}/Γ for a lattice Γ . Top: For real $\mu = 0.3$ we get a rectangular lattice. Bottom: $\mu = 0.08 + 0.01i$ yields a parallelogram.

8.3 Hyperelliptic curves

A hyperelliptic curve is a complex algebraic curve of the form

$$\mu^2 = p(\lambda), \quad (66)$$

where p is a polynomial of degree $d \geq 5$ with d distinct roots. For $d = 2g + 2$ or $d = 2g + 1$, the hyperelliptic curve becomes a compact Riemann surface of genus g after singularities at infinity are resolved. For our purposes, a hyperelliptic curve is just a branched double cover of the λ -sphere with branch points $\lambda_1, \dots, \lambda_{2g+1}, \infty$ if $d = 2g + 1$ and branch points $\lambda_1, \dots, \lambda_{2g+2}$ if $d = 2g + 2$.

We construct a polyhedral approximation of a hyperelliptic curve in the same way as for elliptic curves (Section 7.2). We choose points p_1, \dots, p_n in addition to the $2g + 2$ branch points and take the convex hull. We cut the resulting polyhedron open along edge paths joining pairs of branch points and glue a second copy along the cuts.

Figure 28 shows uniformizations of the curves

$$\mu^2 = \lambda \prod_{k=1}^{2g} \left(\lambda - e^{\frac{ik\pi}{s}} \right). \quad (67)$$

for $g = 2, 3, 4$ that were obtained this way. The curves are branched at the $(2g)$ th roots of unity and at 0 and ∞ .

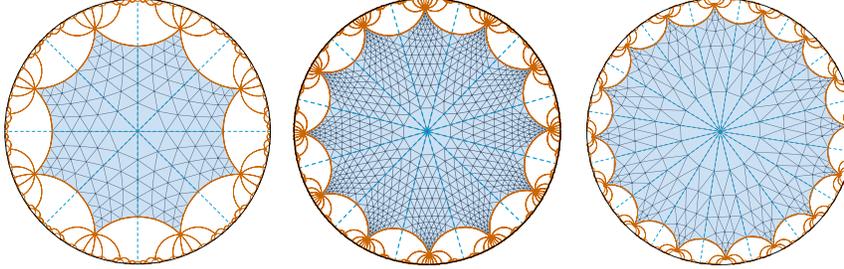


Fig. 28 Uniformizations of the hyperbolic curves (67) with genus 2, 3, and 4. The triangulation of the surfaces is a regular 1-to-4 subdivision of the convex hull of the branch points. Due to the symmetries of these curves, the fundamental domains are regular hyperbolic $4g$ -gons. Since the triangulation is as symmetric as the curves, and because the solution of the discrete uniformization problem is unique, the fundamental domains of the polyhedral surfaces are also exactly regular hyperbolic $4g$ -gons. I.e., any error in the domains is due to numerics, but not to the discretization.

Mapping a polyhedral surface to a hyperelliptic curve. We can also map a triangulated surface of genus g to a branched double cover of the sphere, provided it is symmetric with respect to a discrete conformal involution with $2g + 2$ fixed points, which are vertices. In the simplest case, the involution is an isometry. (Compare Section 7.2, where we map flat tori to elliptic curves.) Taking the quotient of the triangulation with respect to the involution, we get a triangulated sphere with a discrete conformal structure, which we map discretely conformally to the sphere. Figure 29 shows an example.

8.4 Geometric characterization of hyperelliptic surfaces

A Riemann surface R of genus $g \geq 2$ is called *hyperelliptic*, if one of the following equivalent conditions is true (and hence all are):

- (i) R is conformally equivalent to some hyperelliptic curve.
- (ii) R is conformally equivalent to a branched cover of the sphere with $2g + 2$ branch points.
- (iii) There is a conformal involution $\tau : R \rightarrow R$ with exactly $2g + 2$ fixed points.

The involution τ is called the *hyperelliptic involution* of R . By the Riemann-Hurwitz formula, the quotient surface R/τ is a sphere.

All Riemann surfaces of genus two are hyperelliptic, but for every genus greater than two, there are Riemann surfaces that are not hyperelliptic. The following geometric characterization of hyperelliptic Riemann surfaces is due to Schmutz Schaller [23, 24].

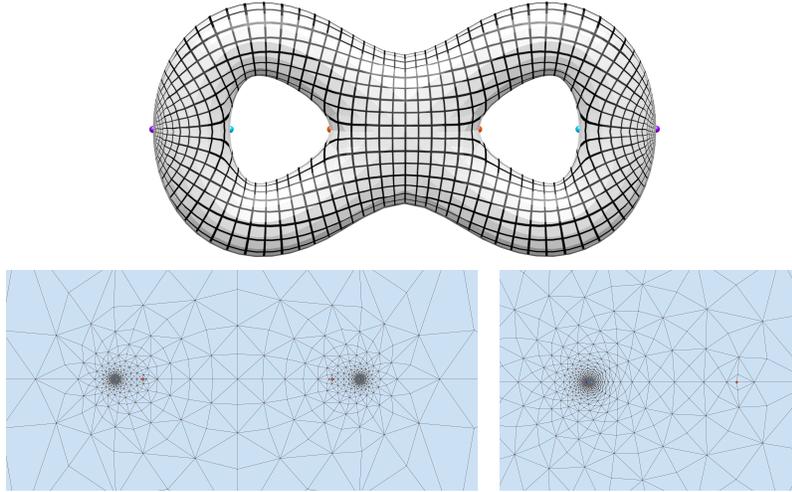


Fig. 29 A triangulated genus 2 surface is mapped to a branched cover of $\hat{\mathbb{C}}$. The 180° rotation about the horizontal symmetry axis is a discrete conformal involution with 6 fixed points, which are marked in red, blue, and purple. The texture is a square grid in the plane, pulled back to the doubly covered sphere by Mercator projection, then pulled back to the surface. Bottom: branched cover of $\hat{\mathbb{C}}$, and a closeup of three branch points.

Theorem 8.2. *Let R be a closed hyperbolic surface of genus g . Then the following statements are equivalent:*

- (i) *R is hyperelliptic.*
- (ii) *R has a set of $2g - 2$ simple closed geodesics which all intersect in one point and which intersect in no other point.*
- (iii) *R has a set of $2g$ simple closed geodesics which all intersect in one point and which intersect in no other point.*
- (iv) *R has a fundamental polygon that is a $4g$ -gon with opposite sides identified and equal opposite angles.*

The fundamental polygon of condition (iv) is symmetric with respect to a 180° rotation around its center, which corresponds to the hyperelliptic involution on R . The $2g + 2$ fixed points on R are the vertex of the polygon, its center, and the $2g$ edge midpoints. The axes of the $2g$ edge gluing translations all go through the center. They project to $2g$ simple closed geodesics on R which all intersect in one point and intersect in no other point.

8.5 Example: Deforming a hyperelliptic surface

We uniformize a hyperelliptic surface obtaining a centrally symmetric fundamental polygon with opposite edges identified as predicted by Theorem 8.2. The axes of

the generators meet in one point. Then we deform the surface slightly to a non-hyperelliptic surface to see how the fundamental polygon and the axes change. The result is shown in Figure 30.

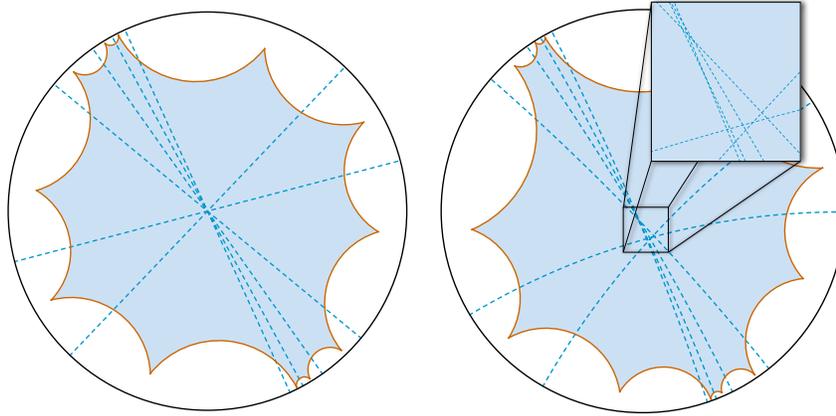


Fig. 30 Hyperelliptic vs. non-hyperelliptic. Left: Uniformization of a hyperelliptic surface with a centrally symmetric fundamental polygon. The axes of the generators meet in a common point. Right: Uniformization of the deformed surface, which is not hyperelliptic. The axes do not meet in one point.

For this example, we construct an elliptic-hyperelliptic triangulated surface with additional symmetry. A surface is called *elliptic-hyperelliptic* if it is conformally equivalent to a two-sheeted branched cover of the torus.

Take two regular tetrahedra (the faces of which are subdivided several times to obtain a finer mesh), cut them across pairs of opposite edges and glue them together to obtain a two-sheeted cover of a tetrahedron branched at the four vertices. Now choose two paths in one of the sheets that connect the centers of the tetrahedron's faces in pairs. Cut the surface along these paths, take another copy of this cut surface and glue corresponding cuts together to form an elliptic-hyperelliptic surface of genus three that is a four-fold cover of a regular tetrahedron. The surface possesses six anti-holomorphic involutions corresponding to the six reflectional symmetries of the tetrahedron, and three holomorphic involutions corresponding to the rotational symmetries of the tetrahedron of order two. Each of the holomorphic involutions has eight fixed points covering the midpoints of a pair of opposite edges. Thus, this elliptic-hyperelliptic surface is also hyperelliptic.

Figure 30 (left) shows a uniformization of the hyperelliptic elliptic-hyperelliptic surface. Destroying the symmetry by moving all points of the polyhedral surface in space by a small random offset destroys the hyperellipticity of the surface, see Figure 30 (right).

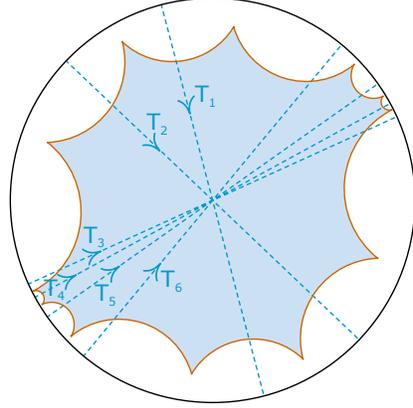


Fig. 31 Generator labels

Numerical Data. We list the numerical $SO^+(2, 1)$ matrices of the generators of the group

$$\langle T_1, T_2, T_3, T_4, T_5, T_6 \mid T_6 T_5^{-1} T_4 T_3^{-1} T_2 T_1^{-1} T_6^{-1} T_5 T_4^{-1} T_3 T_2^{-1} T_1 = 1 \rangle \quad (68)$$

representing the hyperelliptic elliptic-hyperelliptic surface constructed in this section (see Figure 31). The matrices satisfy the relation with error $\approx 10^{-7}$.

$$T_1 = \begin{bmatrix} 2.05443154523212 & -4.021591426903446 & -4.403849064057392 \\ -4.021591427085276 & 16.338309707059754 & 16.796236533536394 \\ -4.403849064222335 & 16.796236533484112 & 17.392741252301292 \end{bmatrix}$$

$$T_2 = \begin{bmatrix} 7.906334736200989 & -6.57792280760043 & -10.236171033333449 \\ -6.5779228079025245 & 7.265127613618063 & 9.749417813849163 \\ -10.236171033527825 & 9.749417813638956 & 14.171462349831586 \end{bmatrix}$$

$$T_3 = \begin{bmatrix} 933.210063638192 & 509.0929753776527 & 1063.0407708335915 \\ 509.09297492442374 & 279.0228056502974 & 580.5414569092936 \\ 1063.0407706165242 & 580.5414573067374 & 1211.2328692884857 \end{bmatrix}$$

$$T_4 = \begin{bmatrix} 47.8208492808903 & 21.282776040302117 & -52.33345184418173 \\ 21.28277609643665 & 10.67424906068982 & -23.788571865092973 \\ -52.333451867010325 & -23.788571814871467 & 57.49509834158029 \end{bmatrix}$$

$$T_5 = \begin{bmatrix} 933.2100574645401 & 509.09297238055706 & -1063.040763978619 \\ 509.092972765322 & 279.02280467565924 & -580.5414545474814 \\ -1063.0407641628826 & -580.5414542100707 & 1211.2328621402066 \end{bmatrix}$$

$$T_6 = \begin{bmatrix} 128.62265665383228 & 90.05086671584104 & -157.0093831621644 \\ 90.05086668827934 & 64.5401174556973 & -110.78621463208009 \\ -157.00938314635744 & -110.78621465448322 & 192.16277410952506 \end{bmatrix}$$

8.6 Example: Different forms of the same genus two surface

In this section we present Fuchsian uniformizations of the same Riemann surface represented in three different ways:

- as hyperelliptic curve $\mu^2 = \lambda^6 - 1$ (Figure 32),
- as Lawson’s genus 2 minimal surface in S^3 [15] (Figure 33),
- and as a surface glued from six squares (Figure 34).

For each representation we choose corresponding fundamental polygons that allow the comparison of the uniformization:

- an octagon with canonical edge pairing $abd'b'cd'c'd'$,
- an octagon with opposite sides identified, $abcd'a'b'c'd'$,
- a 12-gon that is adapted to the six-squares surface.

All data presented in this section is available on the DGD Gallery webpage [26].

Hyperelliptic curve. We uniformize the hyperelliptic curve $\mu^2 = \lambda^6 - 1$ as described in Section 7.2. The results are shown in Figure 32.

To understand the cuts on the hyperelliptic surface that lead to the 12-gon in the bottom row, imagine taking the canonical system of loops in the top row, meeting at the north pole, and deform them until they also meet at the south pole. This introduces a second vertex class in the fundamental polygon.

Lawson’s surface. Figure 33 shows Fuchsian uniformizations of Lawson’s minimal surface in S^3 . The triangulated surface model was kindly provided by Konrad Polthier [22].

This model of the Lawson surface realizes the hyperelliptic involution as a euclidean rotational symmetry. Its symmetry axis meets the surface in six points. These fixed points of the hyperelliptic involution correspond to the branch points of the hyperelliptic curve representation. This allows us to uniformize the model with corresponding fundamental domains.

Six-squares surface. Figure 34 (left) shows a surface glued from six squares, which is conformally equivalent to Lawson’s surface and the hyperelliptic curve. Edges with the same marking are glued together. We calculate a uniformization using the triangulation with vertices added in the centers of the squares as shown. An adapted fundamental domain for this square-tiled translational surface arranges all squares around a single vertex, see Figure 34 (right). By comparison with Figure 32 (bottom) we see that the vertices in the center of the squares correspond to the branch points of the hyperelliptic representation of the surface. The black, gray, and white vertices correspond to the north and south pole of the hyperelliptic representation.

Acknowledgement

This research was supported by DFG SFB/TR 109 “Discretization in Geometry and Dynamics”.

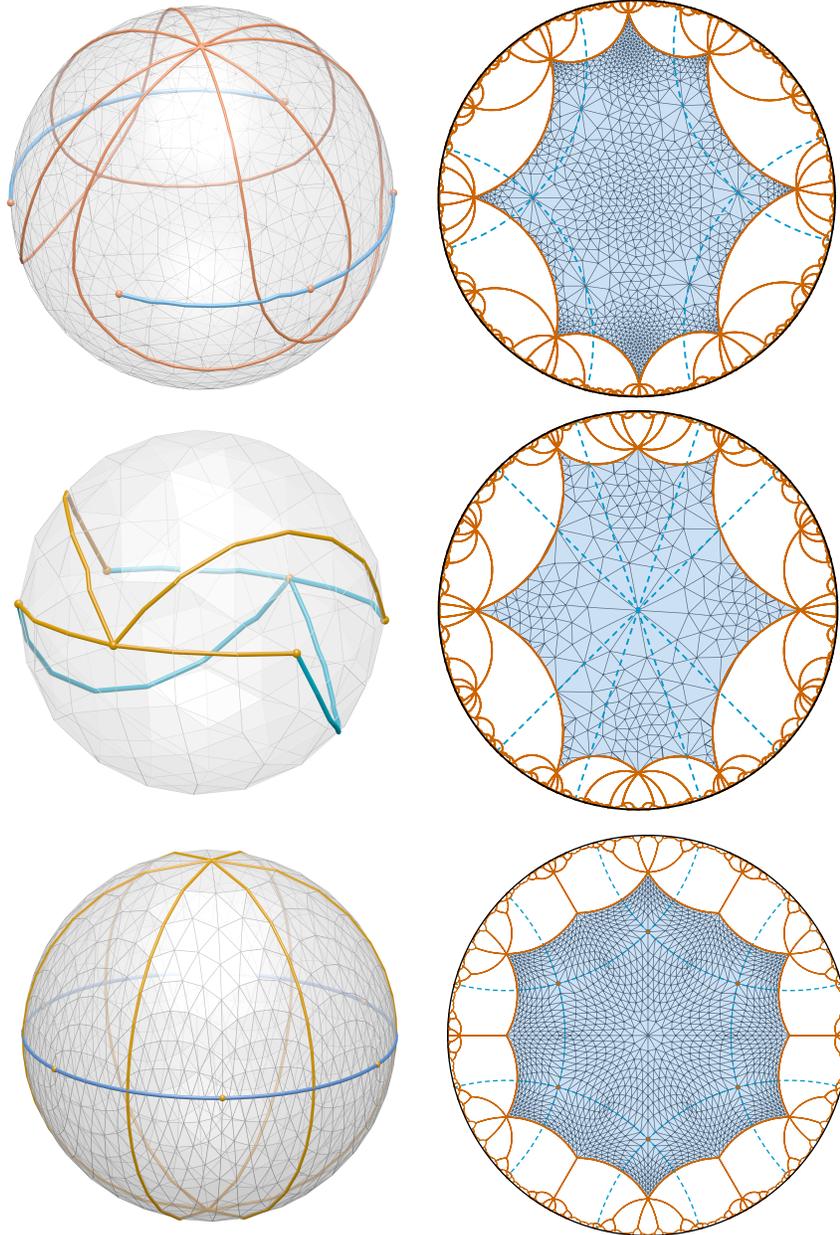


Fig. 32 Uniformization of the hyperelliptic curve $\mu^2 = \lambda^6 - 1$. Left: Triangulated double cover of the sphere branched at the 6th roots of unity, with the boundary of the fundamental domain shown in brown and the axes of generators shown in blue. Right: Fuchsian uniformization and fundamental polygons. Canonical polygon (top), polygon with opposite sides identified (middle), and 12-gon specially adapted to the six-squares surface (bottom).

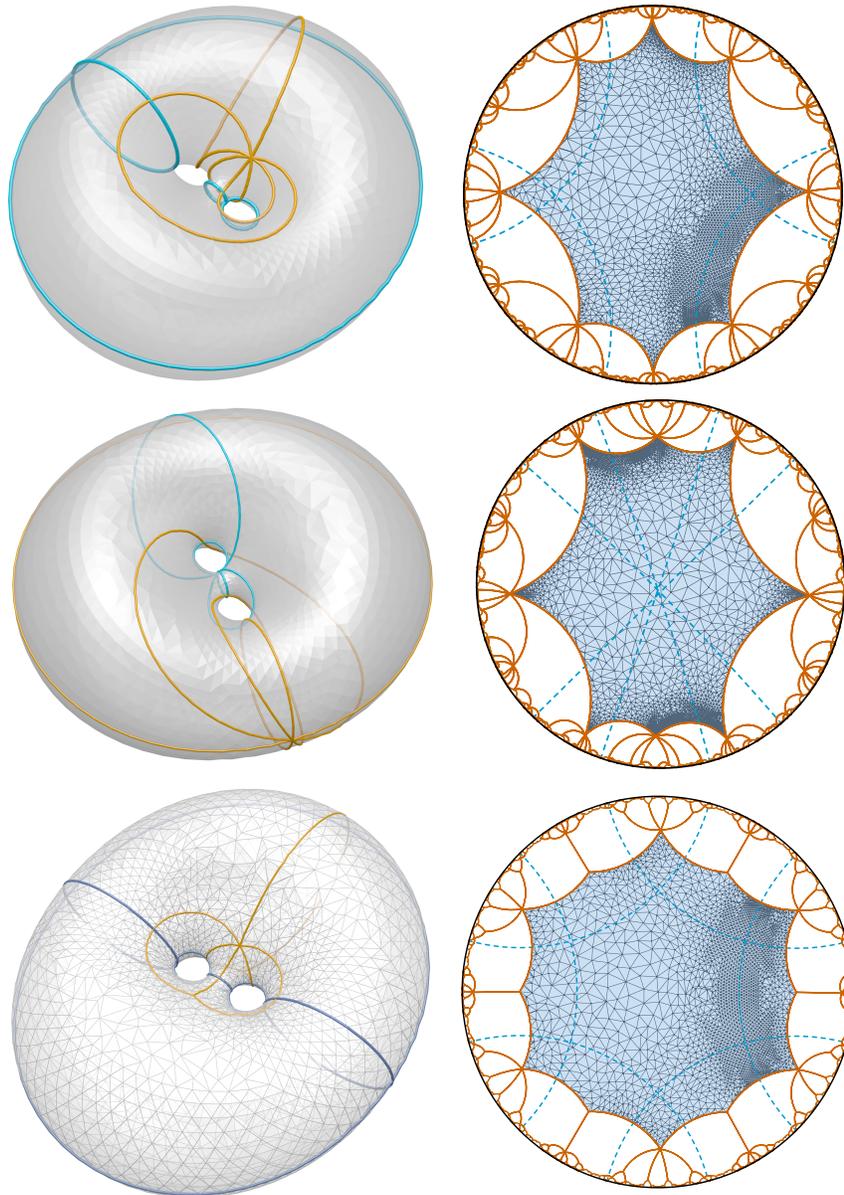


Fig. 33 Uniformization of Lawson's surface. Left: Triangulated model [22], with the boundary of the fundamental domain shown in brown and the axes of the generators shown in blue. Right: Fuchsian uniformizations and fundamental domains. Canonical domain (top), opposite sides domain (middle), and 12-gon (bottom).

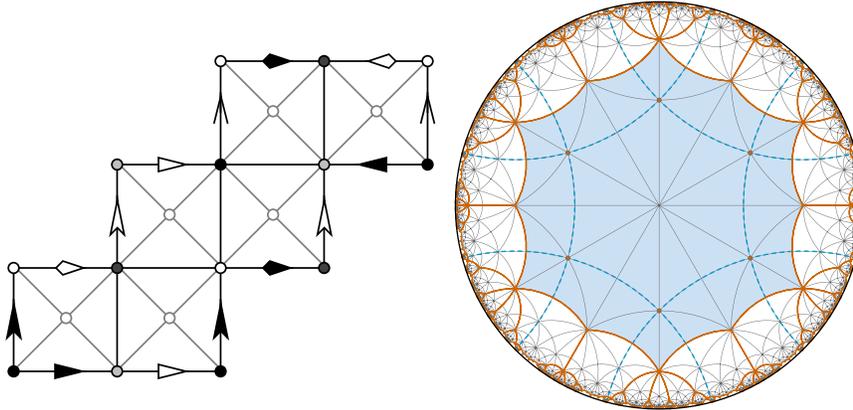


Fig. 34 Left: A surface glued from six squares. Right: Fuchsian uniformization and fundamental domain.

References

1. U. Abresch, *Constant mean curvature tori in terms of elliptic functions*, *J. Reine Angew. Math.* **374** (1987), 169–192.
2. V. E. Adler, A. I. Bobenko, and Yu. B. Suris, *Classification of integrable equations on quad-graphs. The consistency approach*, *Comm. Math. Phys.* **233** (2003), no. 3, 513–543.
3. A. I. Bobenko, *All constant mean curvature tori in \mathbf{R}^3 , S^3 , H^3 in terms of theta-functions*, *Math. Ann.* **290** (1991), no. 2, 209–245.
4. A. I. Bobenko, U. Pinkall, and B. Springborn, *Discrete conformal maps and ideal hyperbolic polyhedra*, *Geometry & Topology* **19** (2015), no. 4, 2155–2215.
5. A. I. Bobenko and Yuri B. Suris, *Integrable systems on quad-graphs*, *Int. Math. Res. Not.* (2002), no. 11, 573–611. MR 1890049 (2003d:37127)
6. ———, *Discrete differential geometry – Integrable structure*, *Graduate Studies in Mathematics*, vol. 98, American Mathematical Society, Providence, RI, 2008.
7. J. Erickson and S. Har-Peled, *Optimally cutting a surface into a disk.*, *Symposium on Computational Geometry*, 2002, pp. 244–253.
8. W. Floyd, B. Weber, and J. Weeks, *The Achilles’ heel of $O(3,1)$?*, *Experiment. Math.* **11** (2002), no. 1, 91–97.
9. X. Gu, F. Luo, J. Sun, and T. Wu, *A discrete uniformization theorem for polyhedral surfaces*, arXiv:1309.4175, 2013.
10. ———, *A discrete uniformization theorem for polyhedral surfaces II*, arXiv:1401.4594, 2014.
11. M. Heil, *Numerical tools for the study of finite gap solutions of integrable systems*, Ph.D. thesis, TU-Berlin, 1995.
12. J. Jost, *Compact Riemann surfaces*, third ed., Universitext, Springer-Verlag, Berlin, 2006.
13. L. Keen, *Canonical polygons for finitely generated Fuchsian groups*, *Acta Math.* **115** (1965), 1–16.
14. H. Kouřimská, L. Skuppin, and B. Springborn, *A variational principle for cyclic polygons with prescribed edge lengths*, *Advances in Discrete Differential Geometry* (A. I. Bobenko, ed.), Springer, 2016.
15. H. B. Lawson, Jr., *Complete minimal surfaces in S^3* , *Ann. of Math.* **92** (1970), no. 3, pp. 335–374.
16. G. Leibon, *Characterizing the Delaunay decompositions of compact hyperbolic surfaces*, *Geom. Topol.* **6** (2002), 361–391.

17. F. Luo, *Combinatorial Yamabe flow on surfaces*, Commun. Contemp. Math. **6** (2004), no. 5, 765–780.
18. J. Milnor, *Hyperbolic geometry: the first 150 years*, Bull. Amer. Math. Soc. (N.S.) **6** (1982), no. 1, 9–24.
19. M. E. Muller, *A note on a method for generating points uniformly on n -dimensional spheres*, Commun. ACM **2** (1959), no. 4, 19–20.
20. Z. Nehari, *Conformal mapping*, McGraw-Hill Book Co., Inc., New York, Toronto, London, 1952.
21. F. Nijhoff and H. Capel, *The discrete Korteweg-de Vries equation*, Acta Appl. Math. **39** (1995), no. 1-3, 133–158, KdV '95 (Amsterdam, 1995).
22. B. Oberknapp and K. Polthier, *An algorithm for discrete constant mean curvature surfaces*, Visualization and Mathematics (H.-Chr. Hege and K. Polthier, eds.), Springer-Verlag Berlin, 1997, pp. 141–161.
23. P. Schmutz Schaller, *Geometry of Riemann surfaces based on closed geodesics*, Bull. Amer. Math. Soc. (N.S.) **35** (1998), no. 3, 193–214.
24. ———, *Teichmüller space and fundamental domains of Fuchsian groups*, Enseign. Math. (2) **45** (1999), no. 1-2, 169–187.
25. O. Schramm, *Circle patterns with the combinatorics of the square grid*, Duke Math. J. **86** (1997), 347–389.
26. S. Sechelmann, A. I. Bobenko, and B. Springborn, *DGD Gallery, Lawson's surface uniformization*, https://gallery.discretization.de/models/lawsons_surface_uniformization, February 2015.
27. B. Springborn, *A unique representation of polyhedral types. Centering via Möbius transformations*, Math. Z. **249** (2005), no. 3, 513–517.
28. R. Walter, *Constant mean curvature tori with spherical curvature lines in non-Euclidean geometry*, Manuscripta Math. **63** (1989), no. 3, 343–363.
29. H. C. Wente, *Counterexample to a conjecture of H. Hopf.*, Pacific J. Math. **121** (1986), no. 1, 193–243.

Technische Universität Berlin, Institut für Mathematik, Straße des 17. Juni 136, 10623 Berlin, Germany

bobenko@math.tu-berlin.de

sechel@math.tu-berlin.de

boris.springborn@tu-berlin.de