



Ideal Hyperbolic Polyhedra and Discrete Uniformization

Boris Springborn¹

Received: 28 March 2018 / Revised: 7 August 2019 / Accepted: 22 August 2019 /
Published online: 5 September 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

We provide a constructive, variational proof of Rivin’s realization theorem for ideal hyperbolic polyhedra with prescribed intrinsic metric, which is equivalent to a discrete uniformization theorem for spheres. The same variational method is also used to prove a discrete uniformization theorem of Gu et al. and a corresponding polyhedral realization result of Fillastre. The variational principles involve twice continuously differentiable functions on the decorated Teichmüller spaces $\tilde{\mathcal{T}}_{g,n}$ of punctured surfaces, which are analytic in each Penner cell, convex on each fiber over $\mathcal{T}_{g,n}$, and invariant under the action of the mapping class group.

Keywords Decorated Teichmüller space · Penner coordinates · Horocycle · Discrete conformal equivalence · Triangulated surface

Mathematics Subject Classification 57M50 · 52B10 · 52C26

1 Introduction

This article is concerned with two types of problems that are in fact equivalent: realization problems for ideal hyperbolic polyhedra with prescribed intrinsic metric, and discrete uniformization problems. We develop a variational method to prove the respective existence and uniqueness theorems. Special attention is paid to the case of genus zero, because it turns out to be the most difficult one. In particular, we provide a constructive variational proof of Rivin’s realization theorem for convex ideal polyhedra with prescribed intrinsic metric:

Editor in Charge: Kenneth Clarkson

Boris Springborn
boris.springborn@tu-berlin.de

¹ Institut für Mathematik, Technische Universität Berlin, Strasse des 17. Juni 136, 10623 Berlin, Germany

Theorem 1.1 (Rivin [31]) *Every complete hyperbolic surface S of finite area that is homeomorphic to a punctured sphere can be realized as a convex ideal polyhedron in three-dimensional hyperbolic space H^3 . The realization is unique up to isometries of H^3 .*

The realizing polyhedron is allowed to degenerate to a *two-sided ideal polygon*. The uniqueness statement of Theorem 1.1 implies that this is the case if and only if S admits an orientation reversing isometry mapping each cusp to itself.

An analogous realization result for convex Euclidean polyhedra was proved by Alexandrov [2, pp. 99–100], and Rivin’s original proof of Theorem 1.1 follows the general approach introduced by Alexandrov: First, show that the realization is unique if it exists. Then use this rigidity result to show that the space of realizable metrics is open and closed in the connected space of all metrics. This topological argument does not provide a method of actually constructing a polyhedron with prescribed intrinsic metric, and to find such a method was posed as a problem for further research [31].

The proof of Theorem 1.1 presented here is variational in nature. It proceeds by transforming the realization problem into a finite dimensional nonlinear convex optimization problem with bounds constraints (see Theorem 7.18). This optimization problem is then shown to have an adequately unique solution (see Sect. 9). The number of variables is $n - 1$ for a sphere with n cusps. The target function $\tilde{\mathcal{E}}_{\Delta, \lambda}^{v, \infty}$ (see Definition 7.16) is twice continuously differentiable and piecewise analytic (see Proposition 7.17). The main work of proving the differentiability statement is done in Sect. 8.

Calculating a value of the target function involves Epstein and Penner’s convex hull construction [13, 26, 27] for surfaces. For the purposes of this article, it is necessary to translate this construction into the language of ideal Delaunay decompositions (see Sects. 4 and 5). An ideal Delaunay triangulation can be found using Weeks’s edge flip algorithm [38]. Once the Delaunay triangulation is known, the target function and its first and second derivatives are given by explicit equations (see Proposition 7.17).

A variational proof of Alexandrov’s realization theorem for Euclidean polyhedra was given by Bobenko and Izvestiev [5]. Their proof also provides a constructive method to produce polyhedral realizations, and there are some similarities between their approach and ours. The variational principles are analogous, and Delaunay triangulations play an important role, too. But there is one important difference: The variational principle of Bobenko and Izvestiev involves a non-convex target function, while the target function considered here is convex. This makes the case of ideal polyhedra actually simpler than the case of Euclidean polyhedra, for which no convex variational principle is known.

In Sect. 10 we turn to the other side of the theory, discrete conformal maps. The realization Theorem 1.1 is equivalent to a discrete uniformization theorem for spheres, Theorem 10.5. The equivalence of discrete conformal mapping problems and realization problems for ideal hyperbolic polyhedra was established in a previous article [7]. Previously, we treated conformal mapping problems and polyhedral realization problems with fixed triangulations. In this article, we require the variable triangulations to be Delaunay. The previously established variational principle extends to the setting of variable triangulations. For discrete conformal maps, this extension can be described

roughly as follows: Minimize the same function as in [7], but flip to a Delaunay triangulation before evaluating it. However, instead of using standard Euclidean edge flips that do not change the piecewise Euclidean metric of the triangulation, use *Ptolemy flips*: Update the length of the flipped edge using Ptolemy's relation (9). This does change the piecewise Euclidean metric, except if the adjacent triangles are inscribed in the same circle. Nevertheless, Ptolemy flips are the right thing to do because they do not change the induced hyperbolic metric (see Proposition 10.2 and Definition 10.3).

In Sect. 11 we use the variational approach to prove the uniformization theorem of Gu et al. [18] (see Theorem 11.1) and a polyhedral realization result for tori (see Theorem 11.2) that was proved by Fillastre [14, Thm. B] using Alexandrov's method. Note that Rivin's Theorems 1.1 and 10.5 on the uniformization spheres are not covered by the work of Gu et al. [17, 18].

A few other cases are known in which a variational principle reduces a polyhedral realization problem to convex optimization. In the Euclidean setting, Izmestiev [20] observed that the variational approach to Alexandrov's theorem [5] leads to a convex optimization problem if one considers the realization of convex caps instead of convex polyhedra.

In the hyperbolic setting, a very general realization result for polyhedral surfaces of arbitrary genus and with finite, ideal or hyperideal vertices is due to Fillastre [14], who applied Alexandrov's method, building on Schlenker's work [33] on an even more general result. In principle, it is possible to derive variational principles in this general setting from Schläfli's differential volume formula for hyperbolic tetrahedra (see, e.g., [7, Sect. 5.5]). But only in some special cases will this lead to *convex* optimization problems. This is due to the fact that the volume of a hyperbolic tetrahedron is a concave function of its dihedral angles only in some special cases.

A few of the cases allowing polyhedral realization by convex optimization have already been treated. Izmestiev and Fillastre [15] consider the case of hyperbolic polyhedral surfaces of genus one with finite vertices. (The volume of a hyperbolic tetrahedron with one ideal and three finite vertices is concave.) Most recently, Prosanov [29] has treated hyperbolic polyhedral surfaces of genus ≥ 2 with ideal vertices. (The building blocks are ideal tetrahedra with one hyperideal and three ideal vertices, truncated at the polar plane of the hyperideal vertex.) Prosanov's work also provides a variational proof of the uniformization theorem involving piecewise hyperbolic surfaces by Gu et al. [17]. While Prosanov does not provide explicit formulas for his Hilbert–Einstein functional in terms of triangulations and lengths, it is straightforward to adapt the known variational principle for a fixed triangulation [7, Sect. 6] to obtain such expressions (see also Remark 2.4). Such explicit formulas are useful if one wants to use the variational principle to solve the realization/uniformization problems numerically.

An interesting sub-case of Fillastre's realization theorem [14] that has not yet been treated by the variational method is the realization problem for hyperideal polyhedra with prescribed metric. Since the volume of a hyperideal tetrahedron is concave [34], this leads to a convex variational principle. Simultaneously, this would provide a constructive proof of an existence and uniqueness statement for hyperideal circle patterns [10], but with variable triangulation. Inversive distance circle patterns in the Euclidean plane correspond to polyhedra with hyperideal vertices except for one ideal vertex. Note that the related realization result for hyperideal polyhedra with given

combinatorial type and *prescribed dihedral angles* [4] has already been treated variationally [6,35].

Note that even if the variational method is applied to realize three-dimensional polyhedra, the problems are all essentially two-dimensional. Whether any truly three-dimensional problem can be treated in a similar fashion seems to be one of the most interesting questions raised by this approach. In particular, “deformations” of two-dimensional problems, like problems involving quasi-Fuchsian manifolds [25], might have a chance to be tractable.

2 Overview: Polyhedral Realizations from Realizable Coordinates

We want to show that the following problem has a unique solution up to isometries of H^3 :

Problem 2.1 Given a complete finite area hyperbolic surface S homeomorphic to a sphere with $n \geq 3$ punctures, find a realization as convex ideal polyhedron.

We assume the hyperbolic surface S is specified in Penner coordinates (Δ, λ) , consisting of a triangulation Δ of the sphere with n marked points and a function $\lambda: E_\Delta \rightarrow \mathbb{R}$ on the set E_Δ of edges (see Sect. 3). These coordinates determine the hyperbolic surface S together with an ideal triangulation Δ and a choice of horocycle at each cusp. For each edge $e \in E_\Delta$, the coordinate λ_e is the signed distance of the horocycles at the ends of e (see Fig. 2, Remark 3.1).

Remark 2.2 (Shear coordinates) Rivin [31] assumes that the surface S is specified by shear coordinates. This is not an issue because it is straightforward to convert between shear coordinates and Penner coordinates (see Proposition 3.5).

The basic idea of our approach is the following: For any polyhedral realization of S and any choice of a distinguished vertex v_∞ , there are special adapted coordinates $(\tilde{\Delta}, \tilde{\lambda})$ with certain characteristic properties. We call them *realizable coordinates with distinguished vertex* v_∞ (see Definition 6.1). The realizable coordinates $(\tilde{\Delta}, \tilde{\lambda})$ describe the same hyperbolic surface S as the given coordinates (Δ, λ) , but the ideal triangulation and the choice of horocycles are in general different. Conversely, if realizable coordinates for S are known, it is straightforward to reconstruct the polyhedral realization. Thus, solving Problem 2.1 turns out to be equivalent to the problem of finding realizable coordinates:

Problem 2.3 Given Penner coordinates (Δ, λ) of a complete finite area hyperbolic surface S homeomorphic to a sphere with $n \geq 3$ punctures and a chosen distinguished cusp v_∞ , find realizable coordinates $(\tilde{\Delta}, \tilde{\lambda})$ with distinguished vertex v_∞ for the same surface.

Definition 6.1 characterizes realizable coordinates in terms of the intrinsic geometry of the surface S . This characterization relies on Epstein and Penner’s convex hull construction [13,26,27] for cusped hyperbolic surfaces decorated with horocycles, and Akiyoshi’s [1] generalization, allowing partially decorated surfaces: Horocycles

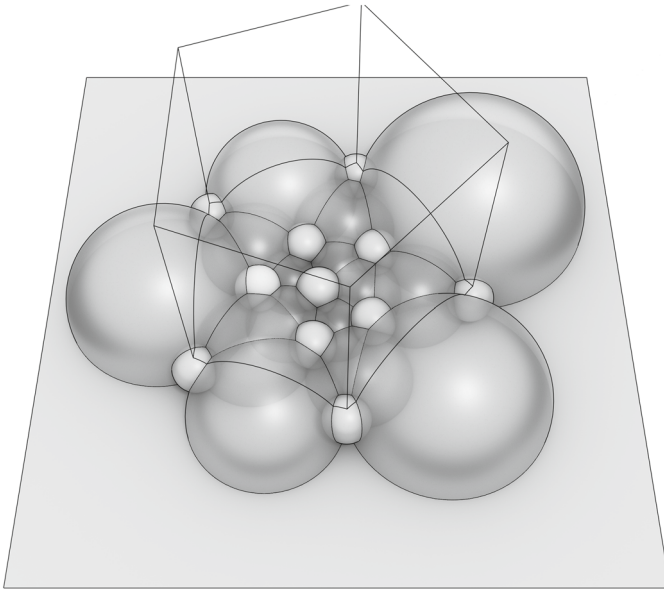


Fig. 1 Ideal polyhedron (bounded by the transparent hyperbolic planes) decorated with horospheres (white) at ideal vertices

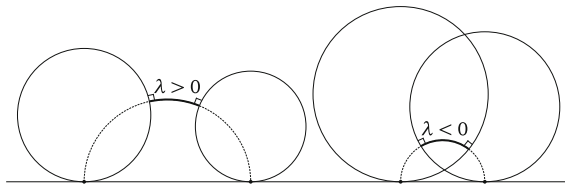


Fig. 2 Signed distance λ of disjoint and intersecting horocycles

may be missing at some (but not all) cusps. The necessary background will be reviewed in Sects. 4 and 5 in the language of ideal Delaunay decompositions.

While the intrinsic characterization of realizable coordinates requires some preparation, it is more straightforward to explain how a polyhedral realization gives rise to adapted Penner coordinates from which the realization can easily be reconstructed. Consider a convex ideal polyhedron P realizing the hyperbolic surface S in the half-space model of hyperbolic space (see Fig. 1). Assume that one ideal vertex, v_∞ , is the point at infinity in the half-space model. The faces of P are ideal polygons. Should any face have more than three sides, triangulate it by adding diagonal ideal arcs. The diagonals may be chosen arbitrarily, except in vertical faces incident with v_∞ , which should be triangulated by adding the vertical diagonals incident with v_∞ .

For every ideal vertex v of P , choose a horosphere s_v centered at v as follows: Choose an arbitrary horosphere s_{v_∞} at v_∞ (not shown in Fig. 1). For all other vertices $v \neq v_\infty$ let s_v be the horosphere centered at v that touches s_{v_∞} (white spheres in Fig. 1). For each edge e not incident with v_∞ let $\tilde{\lambda}_e$ be the signed distance (see Fig. 2) of the horospheres at the ends of e . For each edge e incident with v_∞ , let $\tilde{\lambda}_e = \infty$.

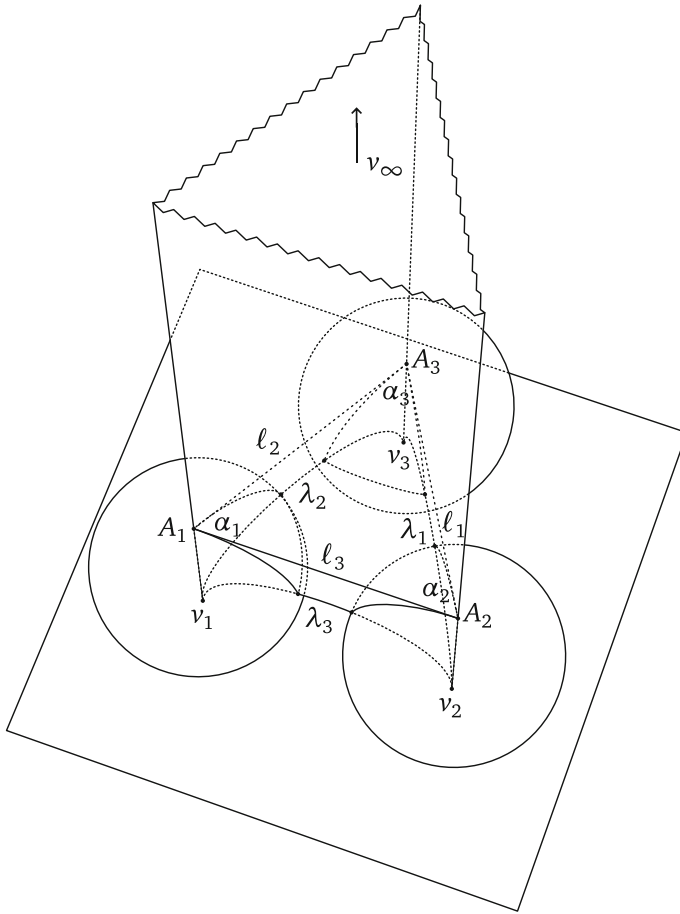


Fig. 3 Ideal tetrahedron decorated with horospheres at the vertices. The horosphere at v_∞ is the horizontal plane at height 1. The horospheres at v_1, v_2, v_3 touch the horosphere at v_∞ . The horosphere at v_∞ intersects the tetrahedron in a Euclidean triangle with sides $l_i = e^{\lambda_i/2}$, where λ_i are the signed distances between horospheres (see Fig. 4)

Intrinsically, the ideal polyhedron P is the hyperbolic surface S . The triangulated faces of P form an ideal triangulation $\tilde{\Delta}$ of S . The horospheres s_v at the vertices of P intersect the surface S in horocycles $h_v = s_v \cap P$. The signed distance $\tilde{\lambda}_e$ of horospheres in H^3 is also the intrinsic signed distance of the corresponding horocycles in S along the ideal arc e . By Proposition 6.2, $(\tilde{\Delta}, \tilde{\lambda})$ are realizable coordinates with distinguished vertex v_∞ as defined in Definition 6.1.

To reconstruct the realization P from $(\tilde{\Delta}, \tilde{\lambda})$, proceed as follows: For each triangle $t \in T_{\tilde{\Delta}}$ that is not incident with v_∞ , construct an ideal tetrahedron as shown in Fig. 3. These tetrahedra fit together to form the ideal polyhedron P . This construction of a realizing polyhedron works for any coordinates that are realizable in the sense of Definition 6.1 (see Proposition 6.3).

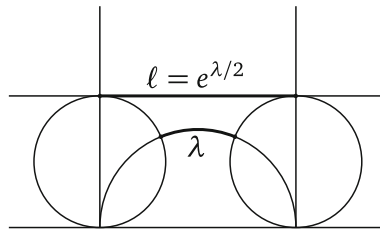


Fig. 4 Signed distance λ and horocyclic arc length ℓ in a configuration of two horocycles that are tangent to a third horocycle

The intrinsic Problem 2.3 of finding realizable coordinates turns out to be equivalent to a convex optimization problem (see Theorem 7.18). This leads to a proof of Theorem 1.1 (see Sect. 9).

Remark 2.4 (Hilbert–Einstein functional) Consider the tetrahedral building block shown in Fig. 3. In this paper, the truncated lengths of the base triangles are variable, while the truncated lengths of the vertical edges are fixed at zero. Alternatively, we could fix the truncated edge lengths of the base triangles and consider the truncated lengths of the vertical edges as the variables of the optimization problem. This approach, which was taken in [7, Sect. 5.4], may at first seem more intuitive. It also leads to an interpretation of the variational principle in terms of a discrete Hilbert–Einstein functional as in [5,15,20,29]. Nevertheless, we avoid this point of view in this article because it makes it harder to understand the role of horocyclic Delaunay triangulations (see, e.g., Theorem 4.14), which is difficult to explain and understand in any case. Moreover, the point of view taken here—fixing the vertical lengths at zero and varying the base lengths—reduces the three-dimensional realization problem to an intrinsic, two-dimensional problem.

3 Penner Coordinates

In Sects. 3–5, we review known results from Teichmüller theory. The aim is to fix notation, to collect required background material and equations for reference, and to translate the convex hull construction into the language of Delaunay decompositions. For the reader’s convenience, we indicate proofs whenever we see a way to do so by a short comment or a suggestive picture. For a more thorough treatment, we refer to the literature. In this section, we review Penner’s coordinates for the decorated Teichmüller spaces of punctured surfaces [26,27].

Let S_g be the oriented surface of genus g , let $V \subseteq S_g$ be a finite nonempty subset of $n = |V|$ points and let $S_{g,n} = S_g \setminus V$ be the oriented surface of genus g with n punctures. The pair (S_g, V) is the surface of genus g with n marked points. A triangulation of (S_g, V) is a triangulation with vertex set V . We will denote the set of edges by E_Δ and the set of triangles by T_Δ . We write $e_1(t), e_2(t), e_3(t)$ for the edges of a triangle t in cyclic order, $e_1(v), \dots, e_{\deg(v)}(v)$ for the edges emanating from a vertex v in cyclic order, and $v_1(e), v_2(e)$ for the vertices of an edge e (in arbitrary order).

The *Teichmüller space* $\mathcal{T}_{g,n}$ is the space of equivalence classes of complete finite area hyperbolic metrics on $S_{g,n}$, where two such metrics are considered equivalent if they are related by an automorphism of $S_{g,n}$ that is isotopic to the identity. Loosely speaking, $\mathcal{T}_{g,n}$ is the space of hyperbolic surfaces of genus g with n cusps. A *decorated hyperbolic surface with cusps* is a hyperbolic surface with cusps together with a horocycle at each cusp. The *decorated Teichmüller space* $\tilde{\mathcal{T}}_{g,n}$ is the space of decorated hyperbolic surfaces of genus g with n cusps. The decorated Teichmüller space $\tilde{\mathcal{T}}_{g,n}$ is a trivial \mathbb{R}^n -bundle over the ordinary Teichmüller space $\mathcal{T}_{g,n}$.

Penner coordinates (Δ, λ) on the decorated Teichmüller space $\tilde{\mathcal{T}}_{g,n}$ consist of a triangulation Δ of (S_g, V) and a function $\lambda \in \mathbb{R}^{E_\Delta}$. The Penner coordinates describe a complete hyperbolic surface with finite area of genus g with n cusps, marked by $S_{g,n}$, and decorated with a horocycle at each cusp, together with an ideal triangulation. The cusps correspond to the vertices $v \in V$ and the ideal triangulation corresponds to a triangulation in the isotopy class of Δ . For simplicity, we will identify cusps with vertices and the ideal triangulation with Δ . The value $\lambda_e = \lambda(e)$ for an edge $e \in E_\Delta$ is the signed distance (see Fig. 2) of the horocycles at its ideal vertices $v_1(e), v_2(e)$ as measured in a lift to the universal cover H^2 .

For two triangulations $\Delta, \tilde{\Delta}$, we denote the *chart transition function* by

$$\tau_{\Delta, \tilde{\Delta}} : \mathbb{R}^{E_\Delta} \rightarrow \mathbb{R}^{E_{\tilde{\Delta}}}. \tag{1}$$

That is, the Penner coordinates (Δ, λ) and $(\tilde{\Delta}, \tilde{\lambda})$ describe the same surface if and only if $\tilde{\lambda} = \tau_{\Delta, \tilde{\Delta}}(\lambda)$.

Remark 3.1 (Notation warning) Our λ s denote hyperbolic lengths, not Penner’s *lambda-lengths* [27]. The lambda-lengths are $e^{\lambda/2}$, and we denote them by ℓ in this paper. In earlier articles [26], the lambda-lengths were defined as $\sqrt{2} e^{\lambda/2}$.

The decorated Teichmüller space $\tilde{\mathcal{T}}_{g,n}$ is a fiber bundle over the ordinary Teichmüller space $\mathcal{T}_{g,n}$. The projection map $\pi : \tilde{\mathcal{T}}_{g,n} \rightarrow \mathcal{T}_{g,n}$ simply forgets about the decoration. The fibers, whose points correspond to choices of decorating horocycles, are naturally affine spaces. If a decorated surface with Penner coordinates (Δ, λ) is chosen as the origin in its fiber, then there is a natural parametrization of the fiber by \mathbb{R}^V :

Proposition 3.2 (Parametrizing a fiber of $\tilde{\mathcal{T}}_{g,n}$) *Let $\Lambda^{(\Delta, \lambda)}$ be the function*

$$\Lambda^{(\Delta, \lambda)} : \mathbb{R}^V \longrightarrow \mathbb{R}^{E_\Delta}, \quad u \longmapsto \Lambda^{(\Delta, \lambda)}(u),$$

where the value of $\Lambda^{(\Delta, \lambda)}(u)$ for $e \in E_\Delta$ is

$$\Lambda_e^{(\Delta, \lambda)}(u) = \lambda_e + u_{v_1(e)} + u_{v_2(e)}. \tag{2}$$

Then the decorated surfaces in the fiber of the surface with coordinates (Δ, λ) are precisely the surfaces with Penner coordinates $(\Delta, \Lambda^{(\Delta, \lambda)}(u))$ for some $u \in \mathbb{R}^V$, and u is uniquely determined by (Δ, λ) and the decorating horocycles. The horocycle of the

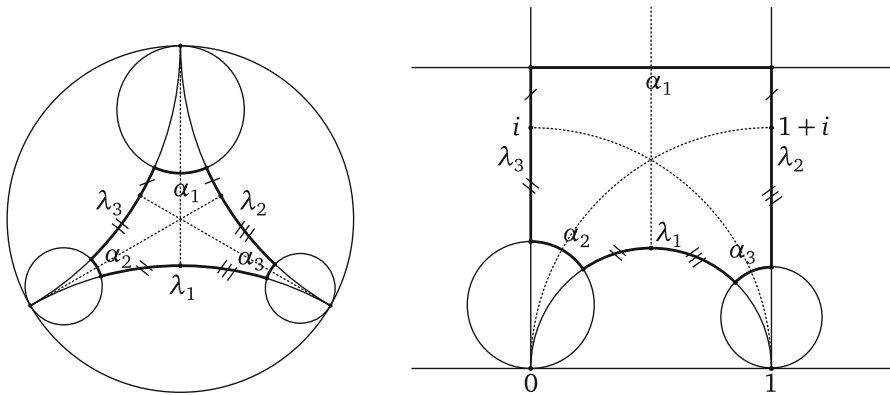


Fig. 5 Decorated ideal triangle in the Poincaré disk model (left) and in the half-plane model (right)

decorated surface $(\Delta, \Lambda^{\Delta, \lambda}(u))$ at a vertex v is a distance u_v away from the horocycle of (Δ, λ) at v , measured in the direction of the cusp.

The following two propositions provide relations between λ s and the lengths of horocyclic arcs.

Proposition 3.3 (Horocyclic arcs in a decorated triangle) *Consider a decorated ideal triangle with signed horocycle distances $\lambda_1, \lambda_2, \lambda_3$ as shown in Fig. 5. Each horocycle intersects the triangle in an arc of length α_k . The signed horocycle distances λ determine the horocyclic arc lengths α , and vice versa, via the relations*

$$\alpha_i = e^{(\lambda_i - \lambda_j - \lambda_k)/2}, \tag{3}$$

$$\lambda_i = -\log \alpha_j - \log \alpha_k, \tag{4}$$

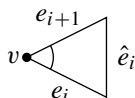
where (i, j, k) is a permutation of $(1, 2, 3)$.

Summing the horocyclic arc lengths around one vertex, one obtains the total horocycle length at a cusp:

Proposition 3.4 (Horocycle length at a cusp) *The total length of the horocycle at a cusp $v \in V$ of the decorated hyperbolic surface with Penner coordinates (Δ, λ) is*

$$c_v(\Delta, \lambda) = \sum_{i=0}^{d-1} e^{(\lambda_{\hat{e}_i(v)} - \lambda_{e_i(v)} - \lambda_{e_{i+1}(v)})/2}, \tag{5}$$

where $e_1(v), \dots, e_d(v) = e_0(v)$ are the edges emanating from v in cyclic order and $\hat{e}_1(v), \dots, \hat{e}_d(v) = \hat{e}_0(v)$ are the edges opposite v in cyclic order, so that the i -th triangle around v looks like this:



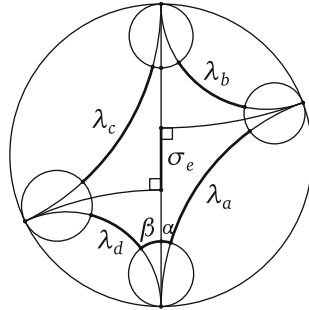


Fig. 6 Penner coordinates and shear coordinates

Shear coordinates (Δ, σ) on the Teichmüller space $\mathcal{T}_{g,n}$ also consist of a triangulation Δ of (S_g, V) and a function $\sigma \in \mathbb{R}^{E_\Delta}$. Shear coordinates describe a marked surface together with an ideal triangulation but without decorating horocycles. For each edge $e \in E_\Delta$, $\sigma_e = \sigma(e)$ is the shear with which the ideal triangles are glued together along e (see Fig. 6). The shear coordinates of a complete hyperbolic surface sum to zero around every cusp:

$$\sum_{i=1}^{\deg(v)} \sigma_{e_i(v)} = 0. \quad (6)$$

Proposition 3.5 (Penner coordinates and shear coordinates) (i) *The shear coordinates (Δ, σ) of a decorated surface with Penner coordinates (Δ, λ) are*

$$\sigma_e = \frac{1}{2} (\lambda_a - \lambda_b + \lambda_c - \lambda_d) = \log \beta - \log \alpha, \quad (7)$$

where a, b, c, d are the edges adjacent to edge e , and α, β are the horocycle arc lengths, as shown in Fig. 6.

(ii) *Conversely, if (Δ, σ) are the shear coordinates of a complete finite area hyperbolic surface S , then one obtains Penner coordinates (Δ, λ) for S , decorated with some choice of horocycles, as follows. First, note that the shear coordinates determine the length ratios of adjacent horocycle arcs α, β as shown in Fig. 6 by equation (7). Use the relations (7) to determine compatible arc lengths around each vertex. (Equations (7) are compatible by (6), but under-determined: At each vertex, exactly one arc length may be chosen arbitrarily, and this choice fixes the decorating horocycle.) Then use equations (4) to determine λ .*

The Penner coordinates of a decorated surface with respect to triangulations Δ and Δ' that differ by a single edge flip are related by Ptolemy's relation (9):

Proposition 3.6 (Ptolemy relation) *Consider a decorated ideal quadrilateral with edges a, b, c, d and diagonals e, f as shown in Fig. 7. Let $\lambda_a, \dots, \lambda_f$ be the respective signed horocycle distances. Then ℓ_a, \dots, ℓ_f defined by*

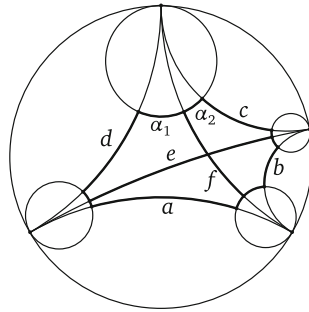


Fig. 7 Ptolemy relation

$$\ell = e^{\lambda/2} \tag{8}$$

satisfy Ptolemy’s relation

$$\ell_e \ell_f = \ell_a \ell_c + \ell_b \ell_d. \tag{9}$$

Proof This follows with (3) from

$$e^{(\lambda_a - \lambda_d - \lambda_f)/2} + e^{(\lambda_b - \lambda_c - \lambda_f)/2} = \alpha_1 + \alpha_2 = e^{(\lambda_e - \lambda_c - \lambda_d)/2}$$

by multiplying both sides with $e^{(\lambda_c + \lambda_d + \lambda_f)/2}$. □

4 Ideal Delaunay Triangulations

Epstein and Penner’s *convex hull construction* [13,26,27] is a fundamental tool for the polyhedral realization method presented here. The construction works for cusped hyperbolic manifolds of arbitrary dimension, but some aspects are simpler for surfaces (compare, e.g., the local condition (10) with the generalized tilt formula [32]), and other aspects (like the edge flip algorithm [38]) have no adequate counterpart in higher dimensions. In this section, we will review the relevant results focusing on the two-dimensional case.

Given a finite area hyperbolic surface with cusps, decorated with a horocycle at each cusp, the construction of Epstein and Penner returns an ideal cell decomposition of the given surface, called the *canonical cell decomposition* of the decorated surface. Roughly, the construction works as follows: Consider the universal cover of the given decorated surface in the hyperboloid model of the hyperbolic plane. Each decorated cusp is represented by an orbit of points in the positive light cone in a $(2 + 1)$ -dimensional Lorentz space. The convex hull of these points projects to an ideal cell decomposition of the surface. We refer to the original articles for a detailed exposition [13,26,27].

In order to apply the convex hull construction to the polyhedral realization problems at hand, we need to translate it into the language of Delaunay decompositions. Accordingly we will use the term “ideal Delaunay decomposition” (Definition 4.2) instead of “canonical cell decomposition”, but these terms are synonymous. The situation is

analogous to the classical case of power diagrams and weighted Delaunay decompositions in Euclidean n -space, which can be obtained by a convex hull construction in $(n + 1)$ -dimensional space [3].

For Epstein and Penner's convex hull construction, the translation into the language of Delaunay decompositions is a straightforward application of the pole-polar relationship of projective geometry and Proposition 4.1 below. It works as follows. In the hyperboloid model of the hyperbolic plane, circles are intersections of the hyperboloid with affine planes. Planes intersecting the hyperboloid correspond, via the pole-polar relationship, to points outside the hyperboloid. Two circles intersect orthogonally if the pole of one circle's plane lies in the other circle's plane. The vertices of the Epstein–Penner convex hull correspond to horocycles, while the faces correspond to orthogonally intersecting circles. (It may be necessary to scale the convex hull up to make all faces intersect the hyperboloid.) Convexity translates to the condition that a face-circle intersects all horocycles corresponding to non-incident vertices less than orthogonally. So far, everything is analogous to the standard theory of power diagrams and weighted Delaunay triangulations in Euclidean space. Proposition 4.1 allows us to express the Delaunay condition in terms of oriented contact instead of orthogonality. This only works for Delaunay triangulations in hyperbolic space and with horospheres as sites.

Proposition 4.1 (Orthogonality and contact) *Let h_1, \dots, h_n be horocycles with different centers in the hyperbolic plane. Of the following statements, any two labeled with the same letter are equivalent. If $n \geq 3$, statements labeled with different letters are mutually exclusive.*

- (a₁) *There is a circle that intersects every horocycle h_1, \dots, h_n orthogonally.*
- (a₂) *There is a circle that touches every horocycle h_1, \dots, h_n externally.*
- (b₁) *There is a horocycle that intersects every horocycle h_1, \dots, h_n orthogonally.*
- (b₂) *There is a horocycle that touches every horocycle h_1, \dots, h_n .*
- (c₁) *There is a hypercycle that intersects every horocycle h_1, \dots, h_n orthogonally.*
- (c₂) *There is a hypercycle or geodesic that touches every horocycle h_1, \dots, h_n on the same side.*

(A hypercycle is a complete curve at a constant nonzero distance from a geodesic.)

The following definitions and theorems summarize Epstein and Penner's results in the language of Delaunay decompositions.

Definition 4.2 (Ideal Delaunay decomposition) *Let S be a complete finite area hyperbolic surface with at least one cusp, decorated with a horocycle at each cusp. First, assume that the horocycles are small enough so that they bound disjoint cusp neighborhoods. An ideal cell decomposition D of S is an ideal Delaunay decomposition, if its lift \hat{D} to H^2 and the lifted horocycles satisfy the global Delaunay condition:*

- (gD) *For every face \hat{f} of \hat{D} there is a circle that touches all lifted horocycles centered at the vertices of \hat{f} externally and does not meet any other lifted horocycles.*

If the horocycles are not small enough to bound disjoint cusp neighborhoods, the cell decomposition D is an ideal Delaunay decomposition if condition (gD) holds

for smaller horocycles at equal distance to the original horocycles. The distance is arbitrary as long as it is large enough for the new horocycles to bound disjoint cusp neighborhoods.

Theorem 4.3 (Existence and uniqueness) *For each decorated complete finite area hyperbolic surface with at least one cusp there exists a unique ideal Delaunay decomposition.*

The faces of the ideal Delaunay decomposition are ideal polygons. An ideal Delaunay triangulation is obtained by adding ideal arcs to triangulate the non-triangular faces:

Definition 4.4 (Ideal Delaunay triangulation) *An ideal triangulation Δ of a decorated complete finite area hyperbolic surface is called an ideal Delaunay triangulation if it is a refinement of the ideal Delaunay decomposition D , that is, if $E_D \subseteq E_\Delta$. An edge $e \in E_\Delta$ is an essential edge of the Delaunay triangulation if $e \in E_D$. The edges in $E_\Delta \setminus E_D$ are nonessential.*

The decorated Teichmüller spaces $\tilde{\mathcal{T}}_{g,n}$ decompose into cells consisting of all decorated surfaces with the same ideal Delaunay decomposition:

Definition 4.5 (Penner cell) *For a triangulation Δ of (S_g, V) , $|V| = n$, let the Penner cell $\mathcal{C}(\Delta)$ be the set of decorated surfaces in $\tilde{\mathcal{T}}_{g,n}$ for which Δ is an ideal Delaunay triangulation.*

Theorem 4.6 (Canonical cell decomposition of $\tilde{\mathcal{T}}_{g,n}$) *The Penner cells $\mathcal{C}(\Delta) \subseteq \tilde{\mathcal{T}}_{g,n}$ are the top-dimensional closed cells of a cell decomposition of $\tilde{\mathcal{T}}_{g,n}$.*

Like in the standard Euclidean theory, the global Delaunay condition (gD) is equivalent to edge-local conditions:

Theorem 4.7 (Local Delaunay conditions) *Let S be a complete finite area hyperbolic surface with at least one cusp, decorated with small enough horocycles so they bound disjoint cusp neighborhoods. An ideal triangulation Δ is a Delaunay triangulation if and only if for every edge e , one and hence all of the equivalent conditions (ID1)–(ID3) are satisfied. Note that the directions left and right relative to e are only defined once an orientation is chosen for e , but the truth values of conditions (ID1) and (ID2) are independent of this choice.*

- (ID1) *The circle touching the three horocycles of the triangle to the left of e and the remaining horocycle of the triangle to the right of e are disjoint or externally tangent.*
- (ID2) *The center of the circle touching the three horocycles of the triangle to the left of e is to the left of, or coincides with, the center of the circle touching the horocycles of the triangle to the right of e .*
- (ID3) *The total length of the horocyclic arcs incident with e is not greater than the total length of the horocyclic arcs opposite to e , as shown in Fig. 8:*

$$\alpha + \alpha' \leq \beta + \beta' + \gamma + \gamma'. \tag{10}$$

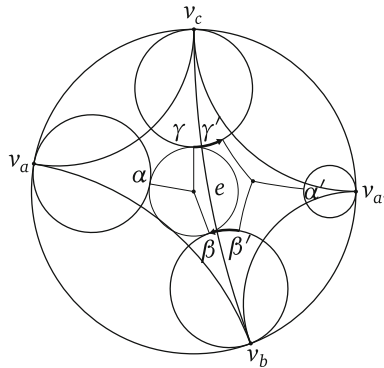


Fig. 8 Local Delaunay condition

Moreover, if Δ is a Delaunay triangulation, then an edge e is nonessential if and only if tangency holds in (ID1), or equivalently, if the circle centers coincide in condition (ID2), or equivalently, if equality holds in (10).

The global Delaunay condition (gD) obviously implies (ID1), and it is not difficult to see that (ID1) and (ID2) are equivalent. To see that (ID2) and (ID3) are equivalent, note that the oriented length of the thick horocyclic arcs in Fig. 8 is $-\alpha - \alpha' + \beta + \beta' + \gamma + \gamma'$. It remains to show that the local conditions imply the global condition (gD); see [26, Thm. 5.1, 27, p. 128].

Theorem 4.8 (Weeks' flip algorithm [38]) *An ideal Delaunay triangulation of the decorated surface (Δ, λ) can be found by the flip algorithm: Iteratively flip any edge violating the local condition (10) and update λ using Ptolemy's relation (9) until no such edge remains.*

- Remarks 4.9**
- (i) Issues of numerical instability that plague the Euclidean flip algorithm [12, 16] are also relevant in this setting.
 - (ii) No other algorithm for computing ideal Delaunay triangulations seems to be known.
 - (iii) Because the diameter of the flip-graph is infinite, the number of flips necessary to arrive at a Delaunay triangulation, depending on an arbitrary initial triangulation, is unbounded. (The only exception is the sphere with three punctures, which admits only three ideal triangulations.)
 - (iv) To analyze the complexity of a variational algorithm to solve the realization Problem 2.1, it would be important to bound the number of steps in the flip algorithm under the condition that the initial triangulation is a Delaunay triangulation for a different choice of horocycles.
 - (v) Little seems to be known about the following related question, except that the number is finite [1]: How many ideal Delaunay decompositions arise for a fixed hyperbolic surface as the decoration varies over all possible choices of horocycles? A recent result says that the lattice of Delaunay decompositions of a fixed punctured Riemann surface is the face lattice of an associated secondary polyhedron [21].

(vi) Recently, an analogous flip algorithm for surfaces with real projective structure was analyzed by Tillmann and Wong [37].

Definition 4.10 Let

$$\text{Del}: (\Delta, \lambda) \mapsto (\tilde{\Delta}, \tilde{\lambda})$$

be a function that maps the Penner coordinates (Δ, λ) of a decorated surface to the Penner coordinates $(\tilde{\Delta}, \tilde{\lambda})$ of the same decorated surface with respect to an ideal Delaunay triangulation $\tilde{\Delta}$.

Such a function Del can be computed using the flip algorithm (see Theorem 4.8). The function Del is not uniquely determined because an ideal Delaunay triangulation is in general not unique. We will use Del in situations in which it makes no difference which ideal Delaunay triangulation is chosen.

The last main point in this section is Theorem 4.14, which establishes a one-to-one correspondence between

- (a) ideal Delaunay triangulations of decorated complete hyperbolic surfaces, and
- (b) Delaunay triangulations of closed piecewise Euclidean surfaces with cone singularities.

Delaunay triangulations and Voronoi diagrams of piecewise Euclidean surfaces with cone singularities were invented and reinvented many times in the context of different applications, for example in [9,19,23,36]. It seems that Theorem 4.14 ought to be known, too, but we do not have a reference.

First we need the following observation, which is due to Penner [26, Lem. 5.2]:

Lemma 4.11 (Ideal Delaunay triangulations and triangle inequalities) *Let Δ be a triangulation of (S_g, V) , let $\lambda \in \mathbb{R}^{E_\Delta}$, and let $\ell \in \mathbb{R}_{>0}^{E_\Delta}$ be defined by (8). If the local Delaunay condition (10) is satisfied for every edge of the decorated hyperbolic surface with Penner coordinates (Δ, λ) , then ℓ satisfies the triangle inequalities for every triangle $t \in T_\Delta$:*

$$\begin{aligned} -\ell_{e_1(t)} + \ell_{e_2(t)} + \ell_{e_3(t)} &> 0, \\ \ell_{e_1(t)} - \ell_{e_2(t)} + \ell_{e_3(t)} &> 0, \\ \ell_{e_1(t)} + \ell_{e_2(t)} - \ell_{e_3(t)} &> 0. \end{aligned} \tag{11}$$

Remark 4.12 This is a global statement. The triangle inequalities (11) for any particular triangle t are satisfied if the local Delaunay conditions (10) are satisfied for *all* edges of the triangulation. It is not sufficient to assume the local Delaunay condition for, say, the edges of triangle t , or in some neighborhood.

If Δ is a triangulation of (S_g, V) and $\ell \in \mathbb{R}_{>0}^{E_\Delta}$ satisfies the triangle inequalities (11) for every triangle $t \in T_\Delta$, then there is a piecewise Euclidean metric on S_g that turns every triangle in T_Δ into a Euclidean triangle and every edge $e \in E_\Delta$ into a straight line segment of length ℓ_e .

Definition 4.13 We will refer to the surface S_g equipped with the piecewise Euclidean metric described in the previous paragraph and together with the triangulation Δ as the *triangulated piecewise Euclidean surface* (Δ, ℓ) .

Theorem 4.14 (Ideal and Euclidean Delaunay triangulations) *Let Δ be a triangulation of (S_g, V) , let $\lambda \in \mathbb{R}^{E_\Delta}$, and let $\ell \in \mathbb{R}_{>0}^{E_\Delta}$ be defined by (8).*

- (i) *If Δ is an ideal Delaunay triangulation of the decorated hyperbolic surface with Penner coordinates (Δ, λ) , then Δ is also a Euclidean Delaunay triangulation of the piecewise Euclidean surface (Δ, ℓ) , which exists by Lemma 4.11.*
- (ii) *Conversely, if ℓ satisfies the triangle inequalities for every triangle $t \in T_\Delta$, and if Δ is a Delaunay triangulation of the piecewise Euclidean surface (Δ, ℓ) , then Δ is also an ideal Delaunay triangulation of the decorated hyperbolic surface with Penner coordinates (Δ, λ) .*

Proof First note that a decorated ideal tetrahedron with horosphere distances λ_i as shown in Fig. 3 exists if and only if the ℓ_i defined by (8) satisfy the triangle inequalities. Now consider two decorated ideal tetrahedra as shown in Fig. 3 glued along two matching vertical faces. The local Delaunay condition (ID2) for the decorated ideal triangles in the base planes is equivalent to the local Delaunay condition for the Euclidean triangles in which the horocycle at v_∞ intersects the tetrahedra. To see this, note that the circumcenters of the Euclidean triangles project vertically down to the highest points of the hyperbolic base planes, and that these points are the centers of circles touching the incident horocycles. \square

5 Akiyoshi's Compactification

Every triangulation Δ of (S_g, V) with $n = |V| > 0$, $2 - 2g - n < 0$ is the Delaunay decomposition for some decorated complete finite area hyperbolic metric on $S_{g,n}$. For example, the decorated surface (Δ, λ) with Penner coordinates $\lambda = 0$ has Delaunay decomposition Δ . There are infinitely many triangulations of (S_g, V) , unless $g = 0$ and $n \leq 3$. However, a fixed hyperbolic surface supports only a finite number of ideal Delaunay triangulations:

Theorem 5.1 (Akiyoshi [1]) *Let S be a complete hyperbolic surface of finite area with at least one cusp. There are only finitely many ideal triangulations Δ of S such that there exists a decoration of S with horocycles such that Δ is a Delaunay triangulation of the decorated surface.*

In fact, Akiyoshi proved a more general result, the generalization of Theorem 5.1 to hyperbolic manifolds of arbitrary dimension. A simpler proof for the two-dimensional case can be found in [18].

Akiyoshi's proof is based on the observation that the convex hull construction of Epstein and Penner generalizes naturally to *partially decorated surfaces*, that is, complete hyperbolic surfaces of finite area, decorated with horocycles at some, but at least one, of the cusps. A missing horocycle should be interpreted as the limit of horocycles vanishing at infinity.

We need to consider partially decorated surfaces, their Delaunay decompositions, and Penner coordinates in some detail because the definition of *realizable coordinates* (Definition 6.1) is based on this: Roughly, realizable coordinates are Penner

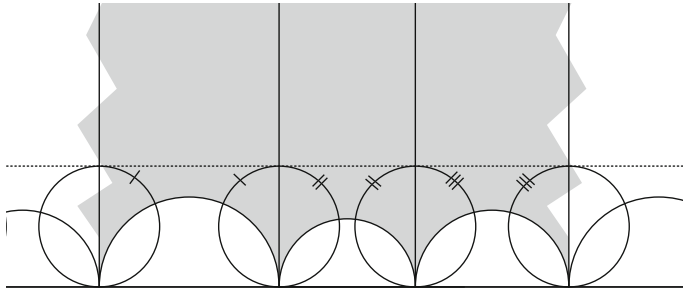


Fig. 9 An ideal polygon with a cusp (shaded region), which is a face of a Delaunay decomposition of a partially decorated surface. The left and right boundaries are identified by an isometry $z \mapsto z + c$. The vertical edges are nonessential edges of an adjusted Delaunay triangulation

coordinates with respect to a Delaunay triangulation of a partially decorated punctured sphere with exactly one horocycle missing (condition (r1) of Definition 6.1), that satisfy an additional condition (condition (r2)). In the remainder of this section, we collect the relevant definitions and facts. We state the results without proof because they are either straightforward translations of known facts into the language of ideal Delaunay decompositions, or they are relatively simple consequences.

Remark 5.2 The need to deal with partially decorated surfaces is one reason why the case of genus zero is more complicated. It is not necessary to consider partially decorated surfaces to treat the higher genus cases (Sect. 11). Readers interested only in the higher genus cases may safely skip the remainder of this section and everything in Sect. 7 that has to do with missing horocycles.

The existence and uniqueness Theorem 4.3 generalizes to partially decorated surfaces:

Theorem 5.3 (Existence and uniqueness) *Every partially decorated surface has a unique ideal Delaunay decomposition.*

However, Definition 4.2 of an ideal Delaunay decomposition has to be modified quite radically: an ideal Delaunay decomposition of a partially decorated surface with missing horocycles is not an ideal cell decomposition (see Definition 5.5). A cusp with missing horocycle is contained in a Delaunay face that is not an ideal polygon, but an ideal polygon with a cusp in the sense of the following definition:

Definition 5.4 *An ideal polygon with a cusp is a hyperbolic manifold-with-boundary f of finite area whose interior is homeomorphic to an open disk with one puncture such that a neighborhood of the puncture corresponds to a cusp neighborhood in f and the boundary is a union of complete geodesics (see Fig. 9).*

However, the Delaunay decomposition of a partially decorated surface is a cell decomposition when viewed as a decomposition of (S_g, V) , the closed surface of genus g with marked points. The faces are cells, but the marked points $v \in V$ that correspond to cusps without decorating horocycle are not vertices of the decomposition. A face

contains at most one marked point in its interior, which we call the *central vertex* of the face (although it is not a vertex of the decomposition). Let us call the regular vertices of an ideal polygon with a cusp its *peripheral vertices* to distinguish them from its central vertex.

Definition 5.5 (*Delaunay decomposition of a partially decorated surface*) Let D be a decomposition of a partially decorated surface into ideal polygons and ideal polygons with a cusp, called the *faces* of the decomposition. Suppose first that the decorating horocycles are disjoint. Then the decomposition D is called a *Delaunay decomposition*, if its lift \hat{D} to H^2 and the lifted horocycles satisfy the following conditions:

- (gD) If \hat{f} is the lift of a face of D that is an ideal polygon, then there is a circle that touches all lifted horocycles at the vertices of \hat{f} externally and does not meet any other lifted horocycles.
- (gD') If \hat{f} is the lift of a face f of D that is an ideal polygon with a cusp, then every peripheral vertex of f is decorated with a horocycle, and there exists a horocycle centered at the lifted central vertex of \hat{f} that touches the lifted horocycles at the lifted peripheral vertices of \hat{f} and does not meet any other lifted horocycles.

Extend this definition to the case of intersecting horocycles like in Definition 4.2.

Remark 5.6 Generically, the punctured Delaunay faces are punctured monogons. Such punctured monogons were already considered by Penner to construct coordinates on the Teichmüller space of partially decorated punctured surfaces [26, Addendum, 27, Sect. 5.1].

Even if a cusp is not decorated with a horocycle, differences of distances to horocycles at other cusps are well defined. This can be used to characterize the punctured Delaunay faces (see Fig. 9):

Proposition 5.7 (*Characterization of punctured Delaunay faces*) Let S be a partially decorated surface, let v be a cusp with missing horocycle, and let p be an ideal polygon with a cusp in S whose central vertex is v . Then p is a face of the Delaunay decomposition of S if and only if all peripheral vertices of p are decorated with horocycles, and these horocycles have all the same distance to v , and this distance is strictly smaller than the distance of any other horocycle to v .

Definition 4.4 of *ideal Delaunay triangulations*, and *essential* and *nonessential edges* remains valid without change. However, it will be useful to triangulate the punctured Delaunay faces in a canonical way:

Definition 5.8 (*Adjusted Delaunay triangulation*) An ideal triangulation Δ of a partially decorated surface is called an *adjusted Delaunay triangulation* if it refines the ideal Delaunay decomposition and every punctured face is triangulated by adding the ideal arcs connecting the central vertex with the peripheral vertices.

Theorem 4.7 on *local Delaunay conditions* remains valid after the following modifications: A triangle of a Delaunay triangulation has at most one vertex with missing horocycle. For an ideal triangle with missing horocycle at one vertex v , one has to consider the horocycle centered at v and touching the triangle's two horocycles, instead of

a circle touching three horocycles. If a horocycle is missing, the respective arc length in condition (10) is zero.

The flip algorithm (see Theorem 4.8) remains valid, but some details have to be modified (see Proposition 5.10) because Penner coordinates (Δ, λ) do not describe partially decorated surfaces. Instead, we may use the parametrization of a fiber of $\tilde{\mathcal{T}}_{g,n}$ described in Proposition 3.2, which extends nicely to partially decorated surfaces:

Definition 5.9 (*Parametrizing an extended fiber of $\tilde{\mathcal{T}}_{g,n}$*) Let (Δ, λ) be Penner coordinates for a decorated surface in $\tilde{\mathcal{T}}_{g,n}$, let

$$\bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}, \tag{12}$$

and

$$u \in \bar{\mathbb{R}}^V \setminus \{+\infty_V\},$$

where $+\infty_V$ denotes the constant function $+\infty$ on V . Let the *partially decorated surface* (Δ, λ, u) be the surface obtained from the decorated surface (Δ, λ) by moving, for each vertex v , the horocycle at v a distance $u(v)$ in the direction of the cusp v . If $u(v) = +\infty$, the horocycle at v is missing. In particular, for $u \in \mathbb{R}^V$, the surface (Δ, λ, u) is just the decorated surface with Penner coordinates $(\Delta, \Lambda^{\Delta, \lambda}(u))$.

If (10) is the local Delaunay condition at edge $e \in E_\Delta$ for the decorated surface (Δ, λ) , then the local Delaunay condition at e for the partially decorated surface (Δ, λ, u) is

$$\alpha e^{-u(v_a)} + \alpha' e^{-u(v_{a'})} \leq (\beta + \beta') e^{-u(v_b)} + (\gamma + \gamma') e^{-u(v_c)}, \tag{13}$$

where $e^{-\infty} = 0$.

Proposition 5.10 (Flip algorithm for partially decorated surfaces) *An ideal Delaunay triangulation of the partially decorated surface (Δ, λ, u) can be found by Weeks' flip algorithm (see Theorem 4.8) with the local condition (10) replaced by (13). An adjusted Delaunay triangulation can then be found by iteratively flipping all nonessential edges opposite an undecorated cusp until no such edge remains.*

With Proposition 5.11 below, the correctness of this algorithm for partially decorated surfaces can be deduced from the correctness of the original algorithm for decorated surfaces.

Proposition 5.11 *If $\tilde{\Delta}$ is an adjusted Delaunay triangulation for the partially decorated surface (Δ, λ, u) , then there is an $M \in \mathbb{R}$ such that $\tilde{\Delta}$ is also an ideal Delaunay triangulation for the decorated surfaces $(\Delta, \Lambda^{\Delta, \lambda}(\tilde{u}))$ with $\tilde{u} \in \mathbb{R}^V$ satisfying*

$$\begin{aligned} \tilde{u}(v) &= u(v) && \text{if } u(v) < +\infty, \\ \tilde{u}(v) &> M && \text{if } u(v) = +\infty. \end{aligned}$$

This is a corollary of the characterization of punctured Delaunay faces in terms of horocycle distances (see Proposition 5.7).

Definition 5.12 (*Generalized Penner coordinates*) If Δ is an arbitrary triangulation of a partially decorated surface, then the surface is in general not determined by Δ and the function $\lambda \in \overline{\mathbb{R}}^E$ of horocycle distances, which takes the value $+\infty$ if the horocycle at one or both ends is missing. However, if Δ is an adjusted Delaunay triangulation, then the pair (Δ, λ) determines the partially decorated surface uniquely (see Proposition 5.7), and we call such a pair *generalized Penner coordinates*.

Definition 5.13 Let Del also denote a function

$$\text{Del}: (\Delta, \lambda, u) \mapsto (\tilde{\Delta}, \tilde{\lambda})$$

that maps the parameters (Δ, λ, u) of a partially decorated surface to generalized Penner coordinates $(\tilde{\Delta}, \tilde{\lambda})$ of the same partially decorated surface, where $\tilde{\Delta}$ is an adjusted Delaunay triangulation. Hence

$$\tilde{\lambda} = \Lambda^{\tilde{\Delta}, \tau_{\Delta, \tilde{\Delta}}(\lambda)}(u) \in \overline{\mathbb{R}}^{E_{\tilde{\Delta}}},$$

where $\tau_{\Delta, \tilde{\Delta}}: \mathbb{R}^{E_{\Delta}} \rightarrow \mathbb{R}^{E_{\tilde{\Delta}}}$ is the chart transition function mapping Penner coordinates with respect to Δ to Penner coordinates with respect to $\tilde{\Delta}$.

Such a function Del can be computed by the modified flip algorithm of Proposition 5.10.

The bounds constraints in the variational principle of Theorem 7.18 involve signed distances of horocycles in a decorated surface, which are defined as follows:

Definition 5.14 (*Signed distance of horocycles*) Let $\delta_{\Delta, \lambda}(v_1, v_2)$ denote the signed distance of the horocycles at the vertices $v_1, v_2 \in V$ in the decorated surface with Penner coordinates (Δ, λ) . More precisely, let

$$\delta_{\Delta, \lambda}(v_1, v_2) = \min_{h_1, h_2} d(h_1, h_2), \quad (14)$$

where d denotes the signed distance of horocycles in H^2 and the minimum is taken over all pairs of horocycles $h_1 \neq h_2$ in H^2 that are lifts of the horocycles at v_1 and v_2 , respectively.

Remark 5.15 The distance $\delta_{\Delta, \lambda}(v_1, v_2)$ is well defined and non-trivial even for $v_1 = v_2$, but we will not need this.

Proposition 5.7 implies that the horocycle distances $\delta_{\Delta, \lambda}$ can be calculated using the flip algorithm:

Proposition 5.16 (Calculating $\delta_{\Delta, \lambda}$) Let (Δ, λ) be Penner coordinates for a decorated surface, let $v_1, v_2 \in V$, let

$$u \in \overline{\mathbb{R}}^V, \quad u(v) = \begin{cases} 0 & \text{if } v = v_2, \\ +\infty & \text{otherwise,} \end{cases}$$

and let $(\tilde{\Delta}, \tilde{\lambda}, u)$ be the output of the modified flip algorithm (see Proposition 5.10). Then v_1 is the central vertex of a punctured face of $\tilde{\Delta}$ and all peripheral vertices are v_2 . For any edge $e \in E_{\tilde{\Delta}}$ connecting v_1 and v_2 ,

$$\delta_{\Delta, \lambda}(v_1, v_2) = \tilde{\lambda}_e.$$

6 Realizable Coordinates

In this section, we characterize the Penner coordinates that may be obtained from an ideal polyhedron by the construction described in Sect. 2: They are realizable coordinates by Definition 6.1 (see Proposition 6.2). Conversely, for given realizable coordinates, a corresponding ideal polyhedron is uniquely determined by an explicit construction (see Proposition 6.3).

Definition 6.1 (*Realizable coordinates*) Let S be a complete finite area hyperbolic surface of genus 0 with $n \geq 3$ cusps. *Realizable coordinates for S with distinguished vertex $v_\infty \in V$* are a pair $(\tilde{\Delta}, \tilde{\lambda})$ consisting of a triangulation $\tilde{\Delta}$ of (S_0, V) and a function $\tilde{\lambda} = \overline{\mathbb{R}}^{E_{\tilde{\Delta}}}$, where $\overline{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$, satisfying the following conditions (r1) and (r2):

- (r1) $\tilde{\Delta}$ is an adjusted Delaunay triangulation for a decoration of S with exactly one missing horocycle at v_∞ , and $(\tilde{\Delta}, \tilde{\lambda})$ are the corresponding generalized Penner coordinates (see Definition 5.12).
- (r2) Let $\tilde{\Delta}^\circ$ be the subcomplex of $\tilde{\Delta}$ consisting of all closed cells not incident with v_∞ , i.e.,

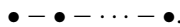
$$V_{\tilde{\Delta}^\circ} = V \setminus \{v_\infty\}, \tag{15}$$

$$E_{\tilde{\Delta}^\circ} = \{e \in E_{\tilde{\Delta}} \mid e \text{ not incident with } v_\infty\}, \tag{16}$$

$$T_{\tilde{\Delta}^\circ} = \{t \in T_{\tilde{\Delta}} \mid t \text{ not incident with } v_\infty\}. \tag{17}$$

Then either (r2a) or (r2b) are true:

- (r2a) $T_{\tilde{\Delta}^\circ} = \emptyset$, and $\tilde{\Delta}^\circ$ is a linear graph, that is, a graph of the form



- (r2b) $T_{\tilde{\Delta}^\circ} \neq \emptyset$ and $\tilde{\Delta}^\circ$ is a triangulation of a closed disk. Moreover,

$$\tilde{\Theta}_v = 2\pi \quad \text{if } v \text{ is an interior vertex of } \tilde{\Delta}^\circ, \tag{18}$$

$$\tilde{\Theta}_v \leq \pi \quad \text{if } v \text{ is a boundary vertex of } \tilde{\Delta}^\circ, \tag{19}$$

where

$$\tilde{\Theta}_v = \text{sum of angles around } v \text{ in } \tilde{\Delta}^\circ, \tag{20}$$

measured in the piecewise Euclidean metric that turns each triangle in $T_{\tilde{\Delta}^\circ}$ into a Euclidean triangle and each edge in $e \in E_{\tilde{\Delta}^\circ}$ into a Euclidean line segment of length $\tilde{\ell}_e$, where

$$\tilde{\ell} = e^{\tilde{\lambda}/2}. \quad (21)$$

Proposition 6.2 (Ideal polyhedron \rightarrow realizable coordinates) *Let P be a three-dimensional convex ideal polyhedron or a two-sided ideal polygon in H^3 realizing a surface $S \in \mathcal{T}_{0,n}$, which has Penner coordinates (Δ, λ) for some decoration. Let v_∞ be an ideal vertex of P . Let $\tilde{\Delta}$ be a triangulation obtained by adding diagonals to triangulate all non-triangular faces of P . The choice of diagonals is arbitrary except for non-triangular faces incident with v_∞ , in which the diagonals incident with v_∞ are be chosen. Let s_∞ be a horosphere centered at v_∞ . For any other vertex v , let s_v be the horosphere centered at v and touching s_∞ . For each edge $e \in E_\Delta$ not incident with v_∞ , let λ_e be the signed distance of the horospheres at the ends of e . If e is incident with v_∞ , let $\lambda_e = \infty$. Then (Δ, λ) are realizable coordinates for S .*

Proof In the case of a two-sided polygon, the statement follows easily from the characterization of punctured Delaunay faces, Proposition 5.7. It remains to consider the case of P being a three-dimensional polyhedron.

To show (r1), first note that for an edge e not incident with v_∞ , λ_e is also the intrinsic signed distance of the horocycles $h_v = s_v \cap P$ at the vertices of e . It remains to show that $\tilde{\Delta}$ is an adjusted Delaunay triangulation. First, the union of triangles incident with v_∞ is the Delaunay face around the vertex v_∞ with vanished horocycle. Indeed, the horocycle h_∞ touches the horocycles at adjacent vertices and does not meet the horocycles at all other vertices.

It remains to show that the local Delaunay conditions (see Theorem 4.7) are satisfied for all edges between two triangles that are not incident with v_∞ . We may assume without loss of generality that the horosphere s_∞ was chosen large enough so that the horospheres at the other vertices are pairwise disjoint. Consider a face f of P that is not incident with v_∞ . The point in the hyperbolic plane of f that is closest to s_∞ is the center of the circle in this plane that touches all horospheres at the vertices of f externally. (See Fig. 1: The points closest to s_∞ are the highest points in the hemispheres.) Using this fact, it is not difficult to see that the local convexity of the edge e in P is equivalent to the local Delaunay condition (ID2) of Theorem 4.7.

(r2) Since we assume that P is a three-dimensional polyhedron, $\tilde{\Delta}^\circ$ is a triangulation of a closed disk. Now (r2b) follows by decomposing P into ideal tetrahedra as shown in Fig. 3, one tetrahedron for each triangle in $\tilde{\Delta}^\circ$. The horosphere s_∞ intersects these tetrahedra in Euclidean triangles with side lengths $\tilde{\ell}$ determined by (21). This implies (18) and (19) because P intersects s_∞ in a convex Euclidean polygon. \square

Proposition 6.3 (Realizable coordinates \rightarrow ideal polyhedron) *Let $(\tilde{\Delta}, \tilde{\lambda})$ be realizable coordinates of a surface $S \in \mathcal{T}_{0,n}$. Let $\tilde{\Delta}^\circ$ be the subcomplex of $\tilde{\Delta}$ defined in Definition 6.1 (r2).*

- (i) *If $\tilde{\Delta}^\circ$ is a triangulation of a closed disk, one obtains a polyhedral realization of S as follows: Construct a decorated ideal tetrahedron as shown in Fig. 3 for each triangle of $\tilde{\Delta}^\circ$. These ideal tetrahedra fit together to form a polyhedron that realizes S .*
- (ii) *If $\tilde{\Delta}^\circ$ is a linear graph then $\tilde{\Delta}$ is a decomposition of S into partially decorated ideal triangles that fit together to form a realization of S as two-sided ideal polygon.*

We omit a detailed proof. The tetrahedra in (i) exist by Lemma 4.11. They fit together to form an ideal polyhedron realizing S by (r2b). That the polyhedron is convex follows from inequality (19) and from the fact that $\tilde{\Delta}$ is a Delaunay triangulation.

The realizable coordinates $(\tilde{\Delta}, \tilde{\lambda})$ with distinguished vertex v_∞ that are obtained from a polyhedron or two-sided polygon by Proposition 6.2 are not uniquely determined:

- Non-triangular faces that are not incident with v_∞ may be triangulated in different ways.
- A different choice of horosphere s_∞ leads to realizable coordinates $(\tilde{\Delta}, \tilde{\lambda} + h\mathbf{1}_{E_{\tilde{\Delta}}})$ for some $h \in \mathbb{R}$.

But these are the only sources of ambiguity: If $(\tilde{\Delta}, \tilde{\lambda})$ and $(\tilde{\Delta}', \tilde{\lambda}')$ are both realizable coordinates obtained from the same polyhedron or two-sided polygon with the same distinguished vertex v_∞ by Proposition 6.2, then

- (u1) $\tilde{\Delta}$ and $\tilde{\Delta}'$ are both adjusted Delaunay triangulations of the same ideal Delaunay decomposition,
 (u2) $\tilde{\lambda}' = \tau_{\tilde{\Delta}, \tilde{\Delta}'}(\tilde{\lambda} + h\mathbf{1}_{E_{\tilde{\Delta}}})$ for some $h \in \mathbb{R}$.

Conversely, the polyhedra obtained from different realizable coordinates $(\tilde{\Delta}, \tilde{\lambda})$ and $(\tilde{\Delta}', \tilde{\lambda}')$ with the same distinguished vertex v_∞ are congruent (as polyhedra marked by (S_0, V)) if and only if conditions (u1) and (u2) are satisfied.

Definition 6.4 (*Equivalent realizable coordinates*) Realizable coordinates $(\tilde{\Delta}, \tilde{\lambda})$ and $(\tilde{\Delta}', \tilde{\lambda}')$ with distinguished vertex v_∞ are *equivalent* if they satisfy conditions (u1) and (u2).

Realizable coordinates are equivalent if and only if they correspond to congruent realizations.

7 The Variational Principle

In this section we present a variational principle (see Theorem 7.18) for Problem 2.3 of finding realizable coordinates. The variational principle involves the function $\tilde{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}(u)$ (see Definition 7.16). The variables $u \in \mathbb{R}^{V \setminus \{v_\infty\}}$ parametrize a part of the extended fiber of $\tilde{\mathcal{T}}_{g, n}$ over the surface (Δ, λ) . More precisely, they parametrize the horocycle decorations of that surface with missing horocycle at v_∞ . The variables are subject to bounds constraints, which ensure that the horocycles do not intersect an arbitrary but fixed horocycle at v_∞ . The definition of $\tilde{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}(u)$ requires some preparation. After the following brief summary, a detailed account begins with Definition 7.1.

The function $f(x_1, x_2, x_3)$ (see Definition 7.1) provides a variational encoding of Euclidean trigonometry: If the variables x_i are the logarithmic side lengths of a Euclidean triangle, then the partial derivatives of f are its angles (see Proposition 7.2).

Using this building block, the function $\lambda \mapsto H_\Theta(\Delta, \lambda)$ (see Definition 7.3) is defined on an open subset A_Δ of the decorated Teichmüller space $\tilde{\mathcal{T}}_{g, n}$. The function $E_{\Theta, \Delta, \lambda}(u)$ (see Definition 7.5) is the restriction of $H_\Theta(\Delta, \cdot)$ to the intersection of A_Δ

with the fiber of $\tilde{\mathcal{T}}_{g,n}$ over the surface (Δ, λ) . This function was already used in a previous article (see Remark 7.6).

Next we extend the function $H_\Theta(\Delta, \lambda)$ to the whole decorated Teichmüller space and the function $E_{\Theta,\Delta,\lambda}(u)$ to a whole fiber by adapting the triangulation appropriately. To this end, consider the restriction of $H_\Theta(\Delta, \cdot)$ to the closed Penner cell of all $\lambda \in \mathbb{R}^{E_\Delta}$ for which Δ is a Delaunay triangulation of the decorated surface (Δ, λ) . For different triangulations Δ , these restrictions of $H_\Theta(\Delta, \cdot)$ fit together to define a C^2 function \mathcal{H}_Θ on the whole decorated Teichmüller space $\tilde{\mathcal{T}}_{g,n}$ (see Corollary 7.10). We denote by $\mathcal{H}_{\Theta,\Delta}(\lambda)$ the representation of \mathcal{H}_Θ in the global Penner coordinate system belonging to the ideal triangulation Δ (see Definition and Proposition 7.9). The convex function $\mathcal{E}_{\Theta,\Delta,\lambda}(u)$ (see Definition 7.11) is the restriction of $\mathcal{H}_{\Theta,\Delta}(\lambda)$ to the fiber of $\tilde{\mathcal{T}}_{g,n}$ over (Δ, λ) . Finally, $\bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}(u)$ is obtained by setting $\Theta_{v_\infty} = 0$ and $\Theta_v = 2\pi$ for $v \neq v_\infty$, and taking the limit $u_{v_\infty} \rightarrow +\infty$ (see Definition 7.16).

Definition 7.1 (*The triangle function f*) Let f be the function

$$f: \mathbb{R}^3 \supseteq \mathcal{A} \longrightarrow \mathbb{R},$$

$$f(x_1, x_2, x_3) = \alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 x_3 + \mathbb{J}(\alpha_1) + \mathbb{J}(\alpha_2) + \mathbb{J}(\alpha_3), \quad (22)$$

where

$$\mathcal{A} = \left\{ \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \in \mathbb{R}^3 \mid \begin{array}{l} e^{x_1} > e^{x_2} + e^{x_3} \\ e^{x_2} > e^{x_3} + e^{x_1} \\ e^{x_3} > e^{x_1} + e^{x_2} \end{array} \right\}$$

(see Fig. 10), \mathbb{J} is Milnor’s Lobachevsky function [24],

$$\mathbb{J}(\alpha) = - \int_0^\alpha \log |2 \sin(t)| dt \quad (23)$$

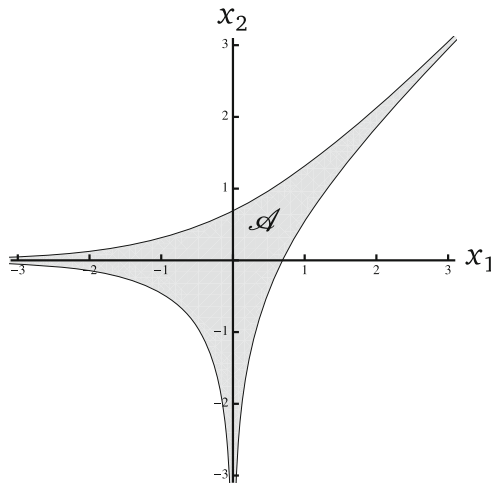


Fig. 10 Intersection of the domain \mathcal{A} with the plane $x_3 = 0$. The domain \mathcal{A} is invariant with respect to translations in the scaling direction $(1, 1, 1)$

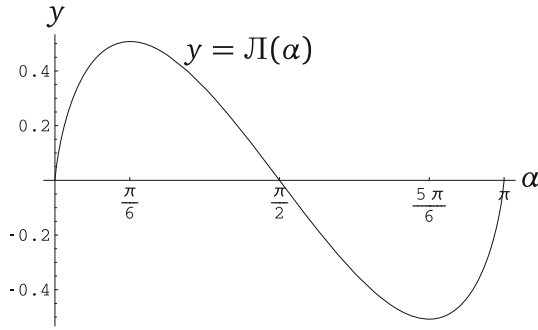


Fig. 11 Milnor’s Lobachevsky function Π is π -periodic, odd, and analytic except at $\alpha \in \pi \mathbb{Z}$, where the derivative tends to $+\infty$

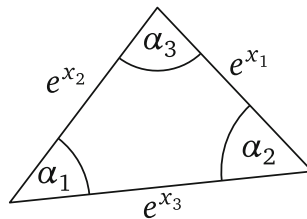


Fig. 12 Triangle with sides $e^{x_1}, e^{x_2}, e^{x_3}$

(see Fig. 11), and α_j are the angles in a triangle with sides e^{x_j} as shown in Fig. 12.

Proposition 7.2 (Properties of f)

(i) *The function f is analytic, and it satisfies the scaling relation*

$$f(x_1 + h, x_2 + h, x_3 + h) = f(x_1, x_2, x_3) + \pi h. \tag{24}$$

(ii) *The partial derivatives of f are*

$$\frac{\partial f}{\partial x_i} = \alpha_i. \tag{25}$$

(iii) *The second derivative of f is*

$$D^2 f|_x = \cot \alpha_1 (dx_2 - dx_3)^2 + \cot \alpha_2 (dx_3 - dx_1)^2 + \cot \alpha_3 (dx_1 - dx_2)^2. \tag{26}$$

(iv) *The second derivative $D^2 f|_x$ is positive semidefinite with one-dimensional kernel spanned by $(1, 1, 1)$. In particular, f is locally convex.*

See [7, Sect. 4.2] for proofs.

Definition 7.3 (The function H_Θ) For a triangulation Δ of (S_g, V) and $\Theta \in \mathbb{R}^V$, let

$$H_\Theta(\Delta, \cdot): A_\Delta \longrightarrow \mathbb{R},$$

$$\begin{aligned}
 H_\Theta(\Delta, \lambda) &= \sum_{t \in T_\Delta} 2f\left(\frac{\lambda_{e_1(t)}}{2}, \frac{\lambda_{e_2(t)}}{2}, \frac{\lambda_{e_3(t)}}{2}\right) - \pi \sum_{e \in E_\Delta} \lambda_e \\
 &\quad - \sum_{v \in V} \Theta_v \log c_v(\Delta, \lambda),
 \end{aligned}
 \tag{27}$$

where $A_\Delta \subseteq \mathbb{R}^{E_\Delta}$ is the subset

$$A_\Delta = \left\{ \lambda \in \mathbb{R}^{E_\Delta} \mid \frac{1}{2}(\lambda_{e_1(t)}, \lambda_{e_2(t)}, \lambda_{e_3(t)}) \in \mathcal{A} \text{ for all } t \in T_\Delta \right\},
 \tag{28}$$

and c_v is defined by (5).

Proposition 7.4 (Properties of H_Θ)

(i) *The function $H_\Theta(\Delta, \cdot)$ is analytic and satisfies the scaling relation*

$$H_\Theta(\Delta, \lambda + h \mathbf{1}_{E_\Delta}) = H_\Theta(\Delta, \lambda) + h \pi \left(|T_\Delta| - |E_\Delta| + \frac{1}{2\pi} \sum_{v \in V} \Theta_v \right).
 \tag{29}$$

(ii) *For $\Theta = 0$, the function $H_0(\Delta, \cdot)$ is convex, and the kernel of the positive semi-definite second derivative is one-dimensional and spanned by $\mathbf{1}_{E_\Delta}$.*

This follows immediately from the corresponding properties of f and the scaling behavior of c_v :

$$c_v(\Delta, \lambda + h \mathbf{1}_{E_\Delta}) = c_v(\Delta, \lambda) - \frac{1}{2} h.
 \tag{30}$$

Definition 7.5 (The function $E_{\Theta, \Delta, \lambda}$) Let $E_{\Theta, \Delta, \lambda}$ be the restriction of H_Θ to the fiber of $\tilde{\mathcal{T}}_{g,n}$ over (Δ, λ) parameterized by $\Lambda^{\Delta, \lambda}$ as defined by (2), that is,

$$\begin{aligned}
 E_{\Theta, \Delta, \lambda} &: \{u \in \mathbb{R}^V \mid \Lambda^{\Delta, \lambda}(u) \in A_\Delta\} \longrightarrow \mathbb{R}, \\
 E_{\Theta, \Delta, \lambda}(u) &= H_\Theta(\Delta, \Lambda^{\Delta, \lambda}(u)).
 \end{aligned}
 \tag{31}$$

Remark 7.6 The function $E_{\Theta, \Delta, \lambda}(u)$ is up to an additive constant equal to the function $E_{T, \Theta, \lambda}(u)$ (with $T = \Delta$) defined in the previous article [7, Eqs. (4–6)]. In that article, the domain of $E_{T, \Theta, \lambda}(u)$ is extended to the whole \mathbb{R}^V by exploiting the fact that the function f can be extended to a convex function on the whole \mathbb{R}^3 . Here, we do not need this extension. Instead, we will extend the functions $H_\Theta(\Delta, \cdot)$ and hence also $E_{\Theta, \Delta, \lambda}(u)$ by changing the triangulation (see Definition and Proposition 7.9 and Definition 7.11).

Proposition 7.7 (Properties of $E_{\Theta, \Delta, \lambda}$)

(i) *The function $E_{\Theta, \Delta, \lambda}$ is analytic and satisfies the scaling relation*

$$E_{\Theta, \Delta, \lambda}(u + h \mathbf{1}_V) = E_{\Theta, \Delta, \lambda}(u) + h 2\pi \left(|T_\Delta| - |E_\Delta| + \frac{1}{2\pi} \sum_{v \in V} \Theta_v \right).
 \tag{32}$$

(ii) For a vertex $v \in V$, the partial derivative of $E_{\Delta,\lambda}(u)$ with respect to u_v is

$$\frac{\partial}{\partial u_v} E_{\Theta,\Delta,\lambda}(u) = \Theta_v - \tilde{\Theta}_v, \tag{33}$$

where $\tilde{\Theta}_v$ is the angle sum around vertex v measured in the piecewise Euclidean metric that turns every triangle in T_Δ into a Euclidean triangle and every edge $e \in E_\Delta$ into a straight line segment of length $\tilde{\ell}_e$, where $\tilde{\ell}$ is defined by (21) with

$$\tilde{\lambda} = \Lambda^{\Delta,\lambda}(u).$$

(iii) The second derivative of $E_{\Theta,\Delta,\lambda}$ at u is

$$\sum_{v,v' \in V} \frac{\partial^2 E_{\Theta,\Delta,\lambda}}{\partial u_v \partial u_{v'}} du_v du_{v'} = \frac{1}{4} \sum_{e \in E_\Delta} (\cot \alpha_e(u) + \cot \alpha'_e(u))(du_{v_1(e)} - du_{v_2(e)})^2, \tag{34}$$

where $\alpha_e(u)$ and $\alpha'_e(u)$ are the angles opposite edge e in the piecewise Euclidean metric described in (ii).

(iv) The function $E_{\Theta,\Delta,\lambda}$ is locally convex. The second derivative is positive semidefinite with one-dimensional kernel spanned by $\mathbf{1}_V$.

Remark 7.8 The second derivative of $E_{\Theta,\Delta,\lambda}$ is the Dirichlet energy of the piecewise linear function taking the value u_v at vertex v [11, Eq. (8), 28]. A satisfactory explanation for this coincidence seems to be unknown.

Proof of Proposition 7.7 (i) and (iv) follow from the corresponding properties of H_Θ (see Proposition 7.4), the scaling behavior of $\Lambda^{\Delta,\lambda}$,

$$\Lambda^{\Delta,\lambda}(u + h \mathbf{1}_V) = \Lambda^{\Delta,\lambda}(u) + 2h \cdot \mathbf{1}_{E_\Delta}, \tag{35}$$

and the equation

$$\log c_v(\Delta, \Lambda^{\Delta,\lambda}(u)) = \log c_v(\Delta, \lambda) - u_v. \tag{36}$$

(ii) Follows from (25) and (36) by a direct calculation. Note that if (i, j, k) is a permutation of $(1, 2, 3)$ and

$$x_i = \frac{\lambda_i + u_j + u_k}{2}$$

then, with the notation of Fig. 5,

$$\frac{\partial}{\partial u_i} 2f(x_1, x_2, x_3) = \alpha_j + \alpha_k = \pi - \alpha_i. \tag{37}$$

(iii) Follows from (26), (27), and (36) by a direct calculation (see also [7, Prop. 4.1.6]). □

The functions $H_\Theta(\Delta, \cdot)$, restricted to the respective Penner cells of Δ , fit together to form a single C^2 function on the decorated Teichmüller space:

Definition and Proposition 7.9 (The function $\mathcal{H}_{\Theta, \Delta}$) *For a triangulation Δ of (S_g, V) , let $\mathcal{H}_{\Theta, \Delta}$ be the function*

$$\begin{aligned} \mathcal{H}_{\Theta, \Delta} &: \mathbb{R}^{E_\Delta} \longrightarrow \mathbb{R}, \\ \mathcal{H}_{\Theta, \Delta}(\lambda) &= H_\Theta(\text{Del}(\Delta, \lambda)). \end{aligned} \tag{38}$$

($\text{Del}(\Delta, \lambda)$ is defined in Definition 4.10.) *The function $\mathcal{H}_{\Theta, \Delta}$ is well defined, twice continuously differentiable, analytic in each open Penner cell of $\tilde{\mathcal{T}}_{g,n}$, and satisfies the scaling relations*

$$\mathcal{H}_{\Theta, \Delta}(\lambda + h \mathbf{1}_{E_\Delta}) = \mathcal{H}_{\Theta, \Delta}(\lambda) + h \pi \left(|T_\Delta| - |E_\Delta| + \frac{1}{2\pi} \sum_{v \in V} \Theta_v \right). \tag{39}$$

Corollary 7.10 *There is a C^2 function $\mathcal{H}_\Theta: \tilde{\mathcal{T}}_{g,n} \rightarrow \mathbb{R}$ on the decorated Teichmüller space, which is analytic on each open Penner cell, such that for each ideal triangulation Δ , the function $\mathcal{H}_{\Theta, \Delta}$ is the representation of \mathcal{H}_Θ in the global Penner coordinate chart belonging to Δ . The function \mathcal{H}_Θ is invariant under the action of the mapping class group.*

Proof of Proposition 7.9 The right hand side of (38) is well-defined for all $\lambda \in \mathbb{R}^{E_\Delta}$ because $(\tilde{\Delta}, \tilde{\lambda}) = \text{Del}(\Delta, \lambda)$ implies $\tilde{\lambda} \in A_{\tilde{\Delta}}$ by Lemma 4.11.

The function $\mathcal{H}_{\Theta, \Delta}$ is analytic on open Penner cells because the functions $H_\Theta(\tilde{\Delta}, \cdot)$ are analytic for all triangulations $\tilde{\Delta}$, and so are the chart transition functions $\tau_{\tilde{\Delta}, \tilde{\Delta}'}: \lambda \mapsto \tilde{\lambda}$ for Penner coordinates with respect to different triangulations Δ and Δ' .

If the Delaunay triangulation for (Δ, λ) is not unique, then the value of the right hand side of (38) and its first two derivatives are independent of the choice of Delaunay triangulation. This follows from Lemma 8.1, which we defer to Sect. 8. It also implies that $\mathcal{H}_{\Theta, \Delta}$ is twice continuously differentiable.

The scaling relation (39) follows from the scaling relation for H_Θ (see Proposition 7.4).

Definition 7.11 Let $\mathcal{E}_{\Theta, \Delta, \lambda}$ be the restriction of $\mathcal{H}_{\Theta, \Delta}$ to the fiber of $\tilde{\mathcal{T}}_{g,n}$ over (Δ, λ) parametrized by $\Lambda^{\Delta, \lambda}$, that is,

$$\begin{aligned} \mathcal{E}_{\Theta, \Delta, \lambda} &: \mathbb{R}^V \longrightarrow \mathbb{R}, \\ \mathcal{E}_{\Theta, \Delta, \lambda}(u) &= \mathcal{H}_{\Theta, \Delta}(\Lambda^{\Delta, \lambda}(u)). \end{aligned} \tag{40}$$

Proposition 7.12 (Properties of $\mathcal{E}_{\Theta, \Delta, \lambda}$)

- (i) *The function $\mathcal{E}_{\Theta, \Delta, \lambda}$ is twice continuously differentiable, analytic in the interior of each Penner cell, and satisfies the scaling relations*

$$\mathcal{E}_{\Theta, \Delta, \lambda}(u + h \mathbf{1}_V) = \mathcal{E}_{\Theta, \Delta, \lambda}(u) + h 2\pi \left(|T_\Delta| - |E_\Delta| + \frac{1}{2\pi} \sum_{v \in V} \Theta_v \right). \tag{41}$$

(ii) *The partial derivatives are*

$$\frac{\partial}{\partial u_v} \mathcal{E}_{\Theta, \Delta, \lambda}(u) = \Theta_v - \tilde{\Theta}_v, \tag{42}$$

where $\tilde{\Theta}_v$ is the total angle at v when S_g is equipped with the piecewise Euclidean metric that turns every triangle in $T_{\tilde{\Delta}}$ into a Euclidean triangle and every edge $e \in E_{\tilde{\Delta}}$ into a straight line segment of length $\tilde{\ell}_e$, where $\tilde{\ell}$ is defined by (21) and

$$(\tilde{\Delta}, \tilde{\lambda}) = \text{Del}(\Delta, \Lambda^{\Delta, \lambda}(u)). \tag{43}$$

(iii) *The second derivative of $\mathcal{E}_{\Theta, \Delta, \lambda}$ at u satisfies*

$$D^2 \mathcal{E}_{\Theta, \Delta, \lambda}|_u = D^2 E_{\Theta, \tilde{\Delta}, \tilde{\lambda}}|_u, \tag{44}$$

with $(\tilde{\Delta}, \tilde{\lambda})$ defined by (43).

(iv) *The function $\mathcal{E}_{\Theta, \Delta, \lambda}$ is convex. The second derivative is positive semidefinite with one-dimensional kernel spanned by $\mathbf{1}_V$.*

Remark 7.13 By (44), one can calculate the second derivative $D^2 \mathcal{E}_{\Theta, \Delta, \lambda}|_u$ by applying the flip algorithm (see Theorem 4.8) to determine $(\tilde{\Delta}, \tilde{\lambda})$ and then (34) for the second derivative of $E_{\Theta, \tilde{\Delta}, \tilde{\lambda}}$.

Proof The claims follow from the corresponding properties of $\mathcal{H}_{\Theta, \Delta}$ and $E_{\Theta, \Delta, \lambda}$. Note that the scaling action of $u \in \mathbb{R}^V$ on $\mathbb{R}^{E_{\Delta}}$ commutes with the chart transition functions $\tau_{\Delta, \tilde{\Delta}}$:

$$\tau_{\Delta, \tilde{\Delta}}(\Lambda^{\Delta, \lambda}(u)) = \Lambda^{\tilde{\Delta}, \tau_{\Delta, \tilde{\Delta}}(\lambda)}(u). \tag{45}$$

So with

$$(\tilde{\Delta}, \tilde{\lambda}) = \text{Del}(\Delta, \Lambda^{\Delta, \lambda}(u)) = (\tilde{\Delta}, \tau_{\Delta, \tilde{\Delta}} \Lambda^{\Delta, \lambda}(u)) = (\tilde{\Delta}, \Lambda^{\tilde{\Delta}, \tau_{\Delta, \tilde{\Delta}}(\lambda)}(u))$$

one obtains

$$\begin{aligned} \mathcal{E}_{\Theta, \Delta, \lambda}(u) &= \mathcal{H}(\Lambda^{\Delta, \lambda}(u)) = H_{\Theta}(\text{Del}(\Delta, \Lambda^{\Delta, \lambda}(u))) = H_{\Theta}(\tilde{\Delta}, \Lambda^{\tilde{\Delta}, \tau_{\Delta, \tilde{\Delta}}(\lambda)}(u)) \\ &= E_{\Theta, \tilde{\Delta}, \tau_{\Delta, \tilde{\Delta}}(\lambda)}(u). \end{aligned} \tag{46}$$

Statements (ii) and (iii) follow with (46) from the corresponding properties of $E_{\Theta, \Delta, \lambda}$ (see Proposition 7.7). \square

To formulate the variational principle of Theorem 7.18 and for the variational existence proofs (see Sects. 9 and 11) we need to consider limits of $\mathcal{E}_{\Theta, \Delta, \lambda}(u)$ as some variables u_v tend to infinity. It is enough to consider $\mathcal{E}_{0, \Delta, \lambda}$, that is, the case $\Theta = 0$, because by (36),

$$\mathcal{E}_{\Theta, \Delta, \lambda}(u) = \mathcal{E}_{0, \Delta, \lambda}(u) - \sum_{v \in V} \Theta_v (\log c_v(\Delta, \lambda) - u_v). \tag{47}$$

Lemma 7.14 (Limits of $\mathcal{E}_{0,\Delta,\lambda}$) *Let $\bar{u} \in \bar{\mathbb{R}}^V$, with $\bar{\mathbb{R}}$ defined by (12), and assume $\bar{u}_v < +\infty$ for at least one vertex $v \in V$. Then*

$$\lim_{u \rightarrow \bar{u}} \mathcal{E}_{0,\Delta,\lambda}(u) = \sum_{t \in T_{\tilde{\Delta}^\circ}} 2f\left(\frac{\tilde{\lambda}_{e_1(t)}}{2}, \frac{\tilde{\lambda}_{e_2(t)}}{2}, \frac{\tilde{\lambda}_{e_3(t)}}{2}\right) - \pi \sum_{e \in E_{\tilde{\Delta}^\circ}} \tilde{\lambda}_e, \tag{48}$$

where

$$(\tilde{\Delta}, \tilde{\lambda}) = \text{Del}(\Delta, \lambda, u),$$

and $\tilde{\Delta}^\circ$ is the subcomplex of the adjusted Delaunay triangulation $\tilde{\Delta}$ consisting of all closed cells that are not incident with an undecorated vertex, that is,

$$\begin{aligned} V_{\tilde{\Delta}^\circ} &= \{v \in V \mid u_v < +\infty\}, \\ E_{\tilde{\Delta}^\circ} &= \{e \in E_{\tilde{\Delta}} \mid \text{vertices of } e \text{ are contained in } V_{\tilde{\Delta}^\circ}\}, \\ T_{\tilde{\Delta}^\circ} &= \{t \in T_{\tilde{\Delta}} \mid \text{vertices of } t \text{ are contained in } V_{\tilde{\Delta}^\circ}\}. \end{aligned}$$

Corollary 7.15 *If $\Theta \geq 0$ and $\Theta_v > 0$ for at least one $v \in V$ with $\bar{u}_v = +\infty$, then*

$$\lim_{u \rightarrow \bar{u}} \mathcal{E}_{\Theta,\Delta,\lambda}(u) = +\infty.$$

Proof of Lemma 7.14 By Akiyoshi’s Theorem 5.1, only finitely many ideal Delaunay decompositions arise from different decorations of the surface (Δ, λ) . It is therefore enough to consider the limit of $\mathcal{E}_{0,\Delta,\lambda}(u)$ as u tends to \bar{u} in the subset

$$\{u \in \mathbb{R}^V \mid \bar{\Delta} \text{ is an ideal Delaunay triangulation for } (\Delta, \Lambda^{\Delta,\lambda}(u))\} \subseteq \mathbb{R}^V \tag{49}$$

for some fixed ideal triangulation $\bar{\Delta}$. Then $\bar{\Delta}$ is also a Delaunay triangulation of the partially decorated surface $(\Delta, \lambda, \bar{u})$, because the local Delaunay conditions (13) are non-strict inequalities, both sides of which extend continuously to $\bar{\mathbb{R}}^V$. In particular, $\bar{\Delta}$ and the adjusted Delaunay triangulation $\tilde{\Delta}$ are ideal Delaunay triangulations of the same ideal Delaunay decomposition. For u in the subset (49),

$$\begin{aligned} \mathcal{E}_{0,\Delta,\lambda}(u) &= H_0(\bar{\Delta}, \bar{\lambda}(u)) \\ &= \sum_{t \in T_{\bar{\Delta}}} 2f\left(\frac{\bar{\lambda}_{e_1(t)}(u)}{2}, \frac{\bar{\lambda}_{e_2(t)}(u)}{2}, \frac{\bar{\lambda}_{e_3(t)}(u)}{2}\right) - \pi \sum_{e \in E_{\bar{\Delta}}} \bar{\lambda}_e(u), \end{aligned} \tag{50}$$

where

$$\bar{\lambda}(u) = \tau_{\Delta,\bar{\Delta}} \circ \Lambda^{\Delta,\lambda}(u).$$

In particular, for each triangle $t \in T_{\bar{\Delta}}$ and all u in the subset (49),

$$\frac{1}{2} (\bar{\lambda}_{e_1(t)}(u), \bar{\lambda}_{e_2(t)}(u), \bar{\lambda}_{e_3(t)}(u)) \tag{51}$$

is contained in \mathcal{A} (see Definition 7.1). There are three possibilities:

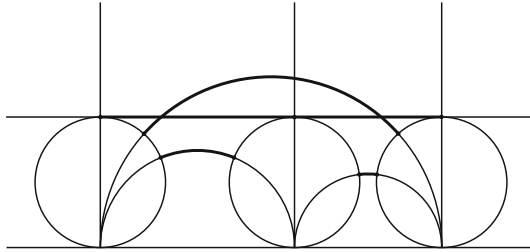


Fig. 13 Ideal triangle contained in punctured Delaunay face, but not incident with the central vertex

- (i) Triangle t is not contained in a punctured Delaunay cell of $(\Delta, \lambda, \bar{u})$. Then (51) converges to a point in \mathcal{A} .
- (ii) Triangle t is contained in a punctured Delaunay cell of $(\Delta, \lambda, \bar{u})$ and t is incident with the undecorated central vertex. Then (51) goes to infinity in \mathcal{A} . If $e_i(t)$ is the edge opposite the central vertex, and $e_j(t), e_k(t)$ are the edges incident with the central vertex, then $\bar{\lambda}_{e_i(t)}(u)$ has a finite limit while $\bar{\lambda}_{e_j(t)}(u)$ and $\bar{\lambda}_{e_k(t)}(u)$ tend to $+\infty$.
- (iii) Triangle t is contained in a punctured Delaunay cell of $(\Delta, \lambda, \bar{u})$ and t is not incident with the undecorated central vertex. Then (51) converges to a boundary point $(\bar{x}_1, \bar{x}_2, \bar{x}_3) \in \partial\mathcal{A}$, that is, for some permutation (i, j, k) of $(1, 2, 3)$,

$$e^{\bar{x}_i} = e^{\bar{x}_j} + e^{\bar{x}_k}$$

(see Figs. 4, 13).

Now (48) for the limit follows from (50) and the following limits of the function f .

- (a) As $(x_1, x_2, x_3) \rightarrow (\bar{x}_1, +\infty, +\infty)$ in \mathcal{A} ,

$$f(x_1, x_2, x_3) - \frac{\pi}{2}(x_2 + x_3) \rightarrow 0. \tag{52}$$

To see this, note that

$$(\alpha_1, \alpha_2, \alpha_3) \rightarrow \left(0, \frac{\pi}{2} + \delta, \frac{\pi}{2} - \delta\right)$$

for some $\delta \in [0, \frac{\pi}{2}]$. So

$$\alpha_1 x_1 \rightarrow 0 \quad \text{and} \quad \mathbb{J}(\alpha_1) + \mathbb{J}(\alpha_2) + \mathbb{J}(\alpha_3) \rightarrow 0$$

because $\mathbb{J}(0) = 0$ and $\mathbb{J}(\frac{\pi}{2} + \delta) + \mathbb{J}(\frac{\pi}{2} - \delta) = 0$. Now (52) follows from

$$\alpha_2 x_2 + \alpha_3 x_3 - \frac{\pi}{2}(x_2 + x_3) \rightarrow 0.$$

To see this, note that

$$\alpha_2 x_2 + \alpha_3 x_3 - \frac{\pi}{2} (x_2 + x_3) = -\frac{1}{2} \alpha_1 (x_2 + x_3) + \frac{1}{2} (\alpha_2 - \alpha_3) (x_2 - x_3).$$

The triangle inequalities imply $(x_2 - x_3) \rightarrow 0$, and using the sine rule one obtains $-\frac{1}{2} \alpha_1 (x_2 + x_3) \rightarrow 0$ from $\lim_{\alpha \rightarrow 0} \alpha \log \sin \alpha = 0$.

(b) As $(x_1, x_2, x_3) \rightarrow (\bar{x}_1, \bar{x}_2, \bar{x}_3) \in \partial \mathcal{A}$, where $e^{\bar{x}_1} = e^{\bar{x}_2} + e^{\bar{x}_3}$,

$$f(x_1, x_2, x_3) \rightarrow \pi \bar{x}_1.$$

This follows from $(\alpha_1, \alpha_2, \alpha_3) \rightarrow (\pi, 0, 0)$. □

The variational principle (see Theorem 7.18) involves the function $\mathcal{E}_{\Theta, \Delta, \lambda}$ with $\Theta_{v_\infty} = 0$, $\Theta_v = 2\pi$ for all other vertices, and $u_{v_\infty} \rightarrow +\infty$. We denote this function by $\bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}$ (see Definition 7.16) and collect its relevant properties (see Proposition 7.17).

Definition 7.16 ($\bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}$) For a triangulation Δ of (S_g, V) , a vertex $v_\infty \in V$, and $\lambda \in \mathbb{R}^{E_\Delta}$, let

$$V^\circ = V \setminus \{v_\infty\}, \tag{53}$$

let $\Theta \in \mathbb{R}^V$ be defined by

$$\Theta_v = \begin{cases} 0 & \text{if } v = v_\infty, \\ 2\pi & \text{if } v \in V^\circ, \end{cases} \tag{54}$$

and define

$$\begin{aligned} \bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty} : \mathbb{R}^{V^\circ} &\longrightarrow \mathbb{R}, \\ \bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}(u) &= \lim_{x \rightarrow +\infty} \mathcal{E}_{\Theta, \Delta, \lambda}(u|_{u_{v_\infty}=x}), \end{aligned} \tag{55}$$

where for $u \in \mathbb{R}^{V^\circ}$, we write $u|_{u_{v_\infty}=x}$ for the function in \mathbb{R}^V with value x at v_∞ and agreeing with u on V° .

Proposition 7.17 (Properties of $\bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}$)

(i) *The limit in (55) exists and is equal to*

$$\begin{aligned} \bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}(u) &= \sum_{t \in T_{\tilde{\Delta}^\circ}} 2f\left(\frac{\tilde{\lambda}_{e_1(t)}}{2}, \frac{\tilde{\lambda}_{e_2(t)}}{2}, \frac{\tilde{\lambda}_{e_3(t)}}{2}\right) - \pi \sum_{e \in E_{\tilde{\Delta}^\circ}} \tilde{\lambda}_e \\ &\quad - 2\pi \sum_{v \in V^\circ} (\log c_v(\Delta, \lambda) - u_v), \end{aligned} \tag{56}$$

where

$$(\tilde{\Delta}, \tilde{\lambda}) = \text{Del}(\Delta, \lambda, u|_{u_{v_\infty}=+\infty}) \tag{57}$$

(see Definition 5.13), and $E_{\tilde{\Delta}^\circ}$ and $T_{\tilde{\Delta}^\circ}$ are defined by (16) and (17).

- (ii) The function $\bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}$ is twice continuously differentiable and analytic in the interior of each Penner cell.
- (iii) The partial derivatives are

$$\frac{\partial}{\partial u_v} \bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}(u) = -\tilde{\Theta}_v + \pi(\text{deg}_2(v) - \text{deg}_1(v) + 2), \tag{58}$$

where $\tilde{\Theta}$ is defined as in Definition 6.1, (20), and

$$\text{deg}_1(v) = \text{edge-degree of } v,$$

that is, the number of edges emanating from v counted with multiplicity, and

$$\text{deg}_2(v) = \text{triangle-degree of } v,$$

that is, the number of triangles around v counted with multiplicity.

- (iv) The second derivative is

$$D^2 \bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty} \Big|_u = \frac{1}{4} \sum_{e \in E_{\tilde{\Delta}^\circ}} w_e(u) (du_{v_1(e)} - du_{v_2(e)})^2, \tag{59}$$

where, if edge e is not a boundary edge of $\tilde{\Delta}^\circ$,

$$w_e(u) = \cot \alpha_e(u) + \cot \alpha'_e(u)$$

and $\alpha_e(u), \alpha'_e(u)$ are the angles opposite e in the piecewise Euclidean metric defined in Definition 6.1 (r2b). If e is a boundary edge, then e has one or zero opposite angles and $w_e(u) = \cot \alpha_e(u)$ or $w_e(u) = 0$, respectively.

- (v) The function $\bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}$ is convex and satisfies the scaling relation

$$\bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}(u + h \mathbf{1}_{V^\circ}) = \bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}(u) + 2\pi h. \tag{60}$$

Proof Statement (i) follows from Lemma 7.14. By direct calculations, one obtains equations (58) and (59) in the interior of Penner cells. As for $\mathcal{H}_{\Theta,\Delta}$, one finds that the first and second derivatives are continuous at the boundaries of Penner cells. This implies the differentiability statement (ii).

Statement (iii) follows from the corresponding properties of $\mathcal{E}_{\Theta,\Delta,\lambda}$ (see Proposition 7.12), because convexity and the scaling relation survive taking the limit (55). Note that

$$\sum_{v \in V} \Theta_v = 2\pi(|V| - 1)$$

due to (54), and hence

$$|T_{\tilde{\Delta}}| - |E_{\tilde{\Delta}}| + \frac{1}{2\pi} \sum_{v \in V} \Theta_v = |T_{\tilde{\Delta}}| - |E_{\tilde{\Delta}}| + |V| - 1 = 1,$$

because $\tilde{\Delta}$ is a triangulation of a sphere. Alternatively, one can also deduce statement (v) directly from (56). □

Theorem 7.18 (Variational principle for Problem 2.3) *Let $S \in \mathcal{T}_{0,n}$ be a complete finite area hyperbolic surface of genus 0 with $n \geq 3$ cusps, and let (Δ, λ) be Penner coordinates for S , decorated with arbitrary horocycles. Let $V^\circ = V \setminus \{v_\infty\}$ for some distinguished vertex $v_\infty \in V$.*

(i) *If the function $\bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}$ attains its minimum under the constraints*

$$u_v \geq \delta_{\Delta,\lambda}(v, v_\infty) \tag{61}$$

at the point $u \in V^\circ$, then $(\tilde{\Delta}, \tilde{\lambda})$ defined by (57) are realizable coordinates with distinguished vertex v_∞ .

(ii) *Up to equivalence (see Definition 6.4), all realizable coordinates with distinguished vertex v_∞ correspond to constrained minima of $\bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}$ as in (i).*

Proof We will show (i) and omit the proof of the converse statement (ii) because it is easier and no new ideas are required. So assume $\bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}$ attains a minimum under the constraints (61) at $u \in \mathbb{R}^{V^\circ}$. We have to show conditions (r1) and (r2) of Definition 6.1. Since condition (r1) obviously holds by construction, it remains to show (r2).

First, note that the convex function $\bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}$ attains a minimum under the constraints (61) at $u \in \mathbb{R}^{V^\circ}$ if and only if for all $v \in V^\circ$:

$$\frac{\partial}{\partial u_v} \bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}(u) = 0 \quad \text{if } u_v > \delta_{\Delta,\lambda}(v, v_\infty), \tag{62}$$

$$\frac{\partial}{\partial u_v} \bar{\mathcal{E}}_{\Delta,\lambda}^{v_\infty}(u) \geq 0 \quad \text{if } u_v = \delta_{\Delta,\lambda}(v, v_\infty). \tag{63}$$

The scaling relation (60) implies that if a constrained minimum is attained at u then u satisfies at least one constraint (61) with equality. The vertices $v \in V^\circ$ for which the constraint (61) is satisfied with equality are precisely the vertices adjacent to v_∞ in $\tilde{\Delta}$ (see Proposition 5.7). Now let $v \in V$ be a vertex adjacent to v_∞ . By (58) and inequality (63), we have

$$-\tilde{\Theta}_v + \pi(\deg_2(v) - \deg_1(v) + 2) \geq 0. \tag{64}$$

Consider two cases separately:

(a) $\deg_2(v) = 0$: In this case $\tilde{\Theta}_v = 0$ because there are no triangles $t \in T_{\tilde{\Delta}^\circ}$ incident with v . Inequality (64) implies

$$\deg_1(v) \leq 2.$$

Using the following two observations, one deduces that the cell complex $\tilde{\Delta}^\circ = (V^\circ, E_{\tilde{\Delta}^\circ}, T_{\tilde{\Delta}^\circ})$ is a linear graph:

(1) Any vertex $v' \neq v_\infty$ adjacent to v also satisfies $\deg_2(v) = 0$.

(2) The cell complex $\tilde{\Delta}^\circ$ is connected because its complement in the sphere S_0 is an open disk.

This proves condition (r2a) of Definition 6.1.

(b) $\text{deg}_2(v) > 0$: In this case $\tilde{\Theta}_v > 0$, so inequality (64) implies

$$\text{deg}_1(v) < \text{deg}_2(v) + 2.$$

On the other hand, because $T_{\tilde{\Delta}^\circ}$ does not contain all triangles of $T_{\tilde{\Delta}}$ incident with v , we have

$$\text{deg}_1(v) \geq \text{deg}_2(v) + 1,$$

and therefore

$$\text{deg}_1(v) = \text{deg}_2(v) + 1. \tag{65}$$

Because the complement of $\tilde{\Delta}^\circ$ in the sphere S_0 is an open disk, this implies that the cell complex $\tilde{\Delta}^\circ$ is a triangulation of a closed disk. With (65), inequality (64) implies (19), and (62) implies (18). This proves condition (r2b).

This concludes the proof of (i). □

8 The Differentiability Lemma

In this section we treat Lemma 8.1, which proves the well-definedness and differentiability statement of Definition and Proposition 7.9.

Lemma 8.1 *Suppose Δ_1 and Δ_2 are both Delaunay triangulations for the decorated surface with Penner coordinates (Δ, λ^*) , and let $\tau_{12} = \tau_{\Delta_1, \Delta_2}$ be the chart transition function $\mathbb{R}^{E_{\Delta_1}} \rightarrow \mathbb{R}^{E_{\Delta_2}}$ mapping Penner coordinates with respect to Δ_1 to Penner coordinates with respect to Δ_2 . Then the function values and the first and second derivatives of $H_\Theta(\Delta_1, \cdot)$ and $H_\Theta(\Delta_2, \tau_{12}(\cdot))$ at λ^* are equal:*

$$H_\Theta(\Delta_1, \lambda^*) = H_\Theta(\Delta_2, \tau_{12}(\lambda^*)), \tag{66}$$

$$dH_\Theta(\Delta_1, \cdot)|_{\lambda^*} = d(H_\Theta(\Delta_2, \tau_{12}(\cdot)))|_{\lambda^*}, \tag{67}$$

$$D^2H_\Theta(\Delta_1, \cdot)|_{\lambda^*} = D^2(H_\Theta(\Delta_2, \tau_{12}(\cdot)))|_{\lambda^*}. \tag{68}$$

Remarks 8.2 (i) It is easy to check numerically that the analogous equations for higher derivatives are in general false. In particular,

$$D^3H_\Theta(\Delta_1, \cdot)|_{\lambda^*} \neq D^3(H_\Theta(\Delta_2, \tau_{12}(\cdot)))|_{\lambda^*},$$

so the third derivative of the function \mathcal{H}_Θ on $\tilde{\mathcal{T}}_{g,n}$ is generally discontinuous at the boundaries of Penner cells.

(ii) The proof of (68) given in this section consists of a straightforward but unilluminating calculation. A more conceptual argument would be desirable.

(iii) In the variational principles of Theorems 7.18 and 11.4, only the function $\mathcal{E}_{\Theta, \Delta, \lambda}(u)$ plays a role, which is the restriction of $\mathcal{H}_{\Theta, \Delta}(\lambda)$ to one fiber of the decorated Teichmüller space $\widetilde{\mathcal{T}}_{g,n}$. In the context of realization and discrete uniformization theorems, it would be enough to show that $\mathcal{E}_{\Theta, \Delta, \lambda}(u)$ is twice continuously differentiable. In view of Theorem 4.14 and Remark 7.8, this follows from Rippa’s Minimal Roughness Theorem [30] for the PL Dirichlet energy. (This is also proved by an unilluminating calculation.) Lemma 8.1 shows that the function $\mathcal{H}_{\Theta, \Delta}$ and hence the function \mathcal{H}_{Θ} of Corollary 7.10 is twice continuously differentiable on the whole decorated Teichmüller space. We believe this is of independent interest.

The rest of this section is devoted to the proof of Lemma 8.1.

1. Reduction to $\Theta = 0$ Since the total length of the decorating horocycle at a vertex v does not depend on the triangulation, we have

$$c_v(\Delta_1, \lambda) = c_v(\Delta_2, \tau_{12}(\lambda))$$

and hence trivially also

$$dc_v(\Delta_1, \cdot)|_{\lambda} = dc_v(\Delta_2, \tau_{12}(\cdot))|_{\lambda}, \tag{69}$$

$$D^2c_v(\Delta_1, \cdot)|_{\lambda} = D^2c_v(\Delta_2, \tau_{12}(\cdot))|_{\lambda} \tag{70}$$

for all ideal triangulations Δ_1, Δ_2 and all $\lambda \in \mathbb{R}^{E_{\Delta_1}}$. To prove Lemma 8.1, it is therefore enough to consider the function H_0 with $\Theta = 0$.

2. Reduction to a single edge flip Without loss of generality, we may assume that Δ_1 and Δ_2 differ by a single edge flip. Indeed, any two Delaunay triangulations for the same decorated surface are related by a finite sequence of flips of nonessential edges, and (66)–(68) have the necessary transitivity property. To be more specific, assume $\Delta_1, \Delta_2, \Delta_3$ are three ideal triangulations. To abbreviate, we write H_i for $H_{\Theta}(\Delta_i, \cdot)$ and τ_{ij} for $\tau_{\Delta_i, \Delta_j}$. Then, by a straightforward application of the chain rules for first and second derivatives, the equations

$$H_1(\lambda^*) = H_2 \circ \tau_{12}(\lambda^*), \quad H_2(\tau_{12}(\lambda^*)) = H_3 \circ \tau_{23}(\tau_{12}(\lambda^*)),$$

$$dH_1|_{\lambda^*} = d(H_2 \circ \tau_{12})|_{\lambda^*}, \quad \text{and} \quad dH_2|_{\tau_{12}(\lambda^*)} = d(H_3 \circ \tau_{23})|_{\tau_{12}(\lambda^*)},$$

$$D^2H_1|_{\lambda^*} = D^2(H_2 \circ \tau_{12})|_{\lambda^*}, \quad D^2H_2|_{\tau_{12}(\lambda^*)} = D^2(H_3 \circ \tau_{23})|_{\tau_{12}(\lambda^*)}$$

imply

$$H_1(\lambda^*) = H_3 \circ \tau_{13}(\lambda^*),$$

$$dH_1|_{\lambda^*} = d(H_3 \circ \tau_{13})|_{\lambda^*},$$

$$D^2H_1|_{\lambda^*} = D^2(H_3 \circ \tau_{13})|_{\lambda^*}.$$

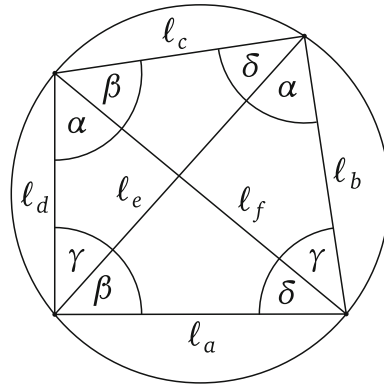


Fig. 14 Lengths and angles in a Euclidean cyclic quadrilateral

3. Notation In the following we assume that Δ_2 is the result of flipping edge e of Δ_1 , which is replaced by edge f in Δ_2 . Let $a, b, c, d \in E_{\Delta_1} \cap E_{\Delta_2}$ be the adjacent edges of e and f as in Fig. 7, and let ℓ be defined by (8). By Lemma 4.11 and Theorem 4.14, the Euclidean triangles with side lengths ℓ_a, ℓ_b, ℓ_e and ℓ_e, ℓ_c, ℓ_d form a cyclic quadrilateral as shown in Fig. 14. By Ptolemy’s theorem of Euclidean geometry, ℓ_f is the length of the other diagonal.

4. Equality of function values To show (66), we use the notation of Fig. 14 for the angles. Writing

$$V_{\text{tet}}(\alpha_1, \alpha_2, \alpha_3) = \mathbb{I}(\alpha_1) + \mathbb{I}(\alpha_2) + \mathbb{I}(\alpha_3), \tag{71}$$

we obtain from (27)

$$\begin{aligned} & H_0(\Delta_1, \lambda^*) - H_0(\Delta_2, \tau_{12}(\lambda^*)) \\ &= 2V_{\text{tet}}(\alpha, \beta, \gamma + \delta) + 2V_{\text{tet}}(\gamma, \delta, \alpha + \beta) + (\alpha + \beta + \gamma + \delta - \pi)\lambda_e^* \\ &\quad - 2V_{\text{tet}}(\beta, \gamma, \alpha + \delta) - 2V_{\text{tet}}(\delta, \alpha, \beta + \gamma) - (\alpha + \beta + \gamma + \delta - \pi)\lambda_f^*, \end{aligned}$$

which is zero because

$$\alpha + \beta + \gamma + \delta = \pi \tag{72}$$

and because (71) is Milnor’s formula [24] for the volume V_{tet} of an ideal tetrahedron with dihedral angles $\alpha_1, \alpha_2, \alpha_3$ as shown in Fig. 3. So

$$V_{\text{tet}}(\alpha, \beta, \gamma + \delta) + V_{\text{tet}}(\gamma, \delta, \alpha + \beta) \quad \text{and} \quad V_{\text{tet}}(\beta, \gamma, \alpha + \delta) + V_{\text{tet}}(\delta, \alpha, \beta + \gamma)$$

are two ways of writing the volume of an ideal quadrilateral pyramid as the sum of the volumes of two tetrahedra.

5. Equality of first derivatives To show (67), note that the chart transition function τ_{12} changes λ_e^* to λ_f^* as determined by Ptolemy’s relation (9) and leaves the values of λ^*

for all other edges unchanged. Now consider the partial derivatives of both sides of (66) at λ^* . From (25) and (27) one obtains by a straightforward calculation

$$\frac{\partial H_0(\Delta_1, \cdot)}{\partial \lambda_e} \Big|_{\lambda^*} = \alpha + \beta + \gamma + \delta - \pi = 0, \quad (73)$$

and similarly

$$\frac{\partial H_0(\Delta_2, \cdot)}{\partial \lambda_f} \Big|_{\tau_{12}(\lambda^*)} = \delta + \alpha + \beta + \gamma - \pi = 0, \quad (74)$$

which implies

$$\frac{\partial H_0(\Delta_2, \tau_{12}(\cdot))}{\partial \lambda_e} \Big|_{\lambda^*} = 0, \quad (75)$$

and hence equality of the partial derivatives with respect to λ_e . For the partial derivative with respect to λ_a one obtains

$$\frac{\partial H_0(\Delta_1, \cdot)}{\partial \lambda_a} \Big|_{\lambda^*} = \alpha = \frac{\partial H_0(\Delta_2, \tau_{12}(\cdot))}{\partial \lambda_a} \Big|_{\lambda^*} \quad (76)$$

and similarly for $\lambda_b, \lambda_c, \lambda_d$. For all other edges $\varepsilon \in E_{\Delta_1} \cap E_{\Delta_2}$, the difference of partial derivatives is zero because all terms depending on λ_ε in the difference $H_0(\Delta_1, \lambda) - H_0(\Delta_2, \tau_{12}(\lambda))$ cancel. This proves (67).

5. Some useful identities In the calculation proving (68), we will use the identities

$$\begin{aligned} -d\lambda_a - d\lambda_c + d\lambda_e + d\lambda_f &= \frac{\sin \beta \sin \delta}{\sin \alpha \sin \gamma + \sin \beta \sin \delta} (-d\lambda_a + d\lambda_b - d\lambda_c + d\lambda_d), \\ d\lambda_b + d\lambda_d - d\lambda_e - d\lambda_f &= \frac{\sin \alpha \sin \gamma}{\sin \alpha \sin \gamma + \sin \beta \sin \delta} (-d\lambda_a + d\lambda_b - d\lambda_c + d\lambda_d), \end{aligned} \quad (77)$$

which are valid at λ^* , as well as the simple trigonometric identities

$$\cot x \pm \cot y = \frac{\sin(\pm x + y)}{\sin x \sin y}. \quad (78)$$

To see (77), take derivatives on both sides of Ptolemy's relation

$$e^{(\lambda_e + \lambda_f)/2} = e^{(\lambda_a + \lambda_c)/2} + e^{(\lambda_b + \lambda_d)/2}$$

to obtain

$$\begin{aligned} d\lambda_e + d\lambda_f &= \frac{e^{(\lambda_a+\lambda_c)/2}}{e^{(\lambda_a+\lambda_c)/2} + e^{(\lambda_b+\lambda_d)/2}} (d\lambda_a + d\lambda_c) \\ &\quad + \frac{e^{(\lambda_b+\lambda_d)/2}}{e^{(\lambda_a+\lambda_c)/2} + e^{(\lambda_b+\lambda_d)/2}} (d\lambda_b + d\lambda_d) \\ &= \frac{1}{1 + e^{(-\lambda_a+\lambda_b-\lambda_c+\lambda_d)/2}} (d\lambda_a + d\lambda_c) \\ &\quad + \frac{1}{1 + e^{(\lambda_a-\lambda_b+\lambda_c-\lambda_d)/2}} (d\lambda_b + d\lambda_d), \end{aligned}$$

then by the law of sines

$$\begin{aligned} d\lambda_e + d\lambda_f &= \frac{\sin \alpha \sin \gamma}{\sin \alpha \sin \gamma + \sin \beta \sin \delta} (d\lambda_a + d\lambda_c) \\ &\quad + \frac{\sin \beta \sin \delta}{\sin \alpha \sin \gamma + \sin \beta \sin \delta} (d\lambda_b + d\lambda_d) \end{aligned} \tag{79}$$

and finally the identities (77).

6. Equality of second derivatives To show (68), first consider the right hand side. The chain rule for second derivatives says

$$\begin{aligned} D^2(H_0(\Delta_2, \tau_{12}(\cdot)))|_{\lambda^*}(v, w) &= D^2H_0(\Delta_2, \cdot)|_{\tau_{12}(\lambda^*)}(d\tau_{12}|_{\lambda^*}(v), d\tau_{12}|_{\lambda^*}(w)) \\ &\quad + \underbrace{dH_0(\Delta_2, \cdot)|_{\tau_{12}(\lambda^*)}(D^2\tau_{12}|_{\lambda^*}(v, w))}_{=0}, \end{aligned} \tag{80}$$

where the term involving $dH_0(\Delta_2, \cdot)$ vanishes due to (74), because $D^2\tau_{12}$ only has a component in the direction of $\frac{\partial}{\partial \lambda_f}$.

Now use (26), (27), and (72) to obtain

$$\begin{aligned} &2\left(D^2H_0(\Delta_1, \cdot) - D^2(H_0(\Delta_2, \tau_{12}(\cdot)))\right)|_{\lambda^*} \\ &= \cot \alpha \left((d\lambda_b - d\lambda_e)^2 - (d\lambda_f - d\lambda_d)^2 \right) \\ &\quad + \cot \beta \left((d\lambda_e - d\lambda_a)^2 - (d\lambda_c - d\lambda_f)^2 \right) \\ &\quad + \cot \gamma \left((d\lambda_d - d\lambda_e)^2 - (d\lambda_f - d\lambda_b)^2 \right) \\ &\quad + \cot \delta \left((d\lambda_e - d\lambda_c)^2 - (d\lambda_a - d\lambda_f)^2 \right) \\ &\quad + \cot(\alpha + \beta) \left((d\lambda_c - d\lambda_d)^2 - (d\lambda_a - d\lambda_b)^2 \right) \\ &\quad - \cot(\beta + \gamma) \left((d\lambda_d - d\lambda_a)^2 - (d\lambda_b - d\lambda_c)^2 \right) \\ &= \cot \alpha (d\lambda_b + d\lambda_d - d\lambda_e - d\lambda_f)(d\lambda_b - d\lambda_d - d\lambda_e + d\lambda_f) \\ &\quad + \cot \beta (-d\lambda_a - d\lambda_c + d\lambda_e + d\lambda_f)(-d\lambda_a + d\lambda_c + d\lambda_e - d\lambda_f) \\ &\quad + \cot \gamma (d\lambda_b + d\lambda_d - d\lambda_e - d\lambda_f)(-d\lambda_b + d\lambda_d - d\lambda_e + d\lambda_f) \end{aligned}$$

$$\begin{aligned}
& + \cot \delta (-d\lambda_a - d\lambda_c + d\lambda_e + d\lambda_f)(d\lambda_a - d\lambda_c + d\lambda_e - d\lambda_f) \\
& + \cot(\alpha + \beta)(d\lambda_a - d\lambda_b + d\lambda_c - d\lambda_d)(-d\lambda_a + d\lambda_b + d\lambda_c - d\lambda_d) \\
& - \cot(\beta + \gamma)(-d\lambda_a + d\lambda_b - d\lambda_c + d\lambda_d)(-d\lambda_a - d\lambda_b + d\lambda_c + d\lambda_d) \\
& = (d\lambda_b + d\lambda_d - d\lambda_e - d\lambda_f)((\cot \alpha - \cot \gamma)(d\lambda_b - d\lambda_d) \\
& \quad + (\cot \alpha + \cot \gamma)(-d\lambda_e + d\lambda_f)) + (-d\lambda_a - d\lambda_c + d\lambda_e + d\lambda_f) \\
& \quad ((\cot \beta - \cot \delta)(-d\lambda_a + d\lambda_c) + (\cot \beta + \cot \delta)(d\lambda_e - d\lambda_f)) \\
& \quad + (d\lambda_a - d\lambda_b + d\lambda_c - d\lambda_d)((\cot(\alpha + \beta) - \cot(\beta + \gamma))(d\lambda_b - d\lambda_d) \\
& \quad + (\cot(\alpha + \beta) + \cot(\beta + \gamma))(-d\lambda_a + d\lambda_c)) \\
& \stackrel{5.}{=} (-d\lambda_a + d\lambda_b - d\lambda_c + d\lambda_d) \left(\frac{1}{\sin \alpha \sin \gamma + \sin \beta \sin \delta} \right. \\
& \quad \times (\sin(-\alpha + \gamma)(d\lambda_b - d\lambda_d) + \sin(\alpha + \gamma)(-d\lambda_e + d\lambda_f) \\
& \quad \quad + \sin(-\beta + \delta)(-d\lambda_a + d\lambda_c) + \sin(\beta + \delta)(d\lambda_e - d\lambda_f)) \\
& \quad - \frac{1}{\sin(\alpha + \beta) \sin(\beta + \gamma)} (\sin(-\alpha + \gamma)(d\lambda_b - d\lambda_d) \\
& \quad \quad \left. + \sin(-\beta + \delta)(-d\lambda_a + d\lambda_c)) \right) \\
& = 0. \tag{81}
\end{aligned}$$

To see equality “ $\stackrel{5.}{=}$ ”, use the identities (77) and (78). To see the last equality, note that (72) implies

$$\sin(\alpha + \beta) \sin(\beta + \gamma) = \sin \alpha \sin \gamma + \sin \beta \sin \delta. \tag{82}$$

This proves (68) and completes the proof of Lemma 8.1.

Remark 8.3 In connection with Remark 8.2(iii), it is curious to note that for a vertical tangent vector in $T\tilde{\mathcal{T}}_{g,n}$, that is, a vector tangent to the fiber over a point in $\mathcal{T}_{g,n}$, the first factor in (81),

$$-d\lambda_a + d\lambda_b - d\lambda_c + d\lambda_d,$$

vanishes as well.

9 Proof of Theorem 1.1

In this section, we prove Theorem 1.1 using the variational principle of Theorem 7.18. By Propositions 6.2 and 6.3, Theorem 1.1 is equivalent to the following statement about the existence and uniqueness of realizable coordinate:

Proposition 9.1 *Problem 2.3 has a unique solution up to equivalence (see Definition 6.4).*

Proposition 9.1 is in turn equivalent to the following statement about the unique solvability of an optimization problem with bounds constraints (see Theorem 7.18):

Proposition 9.2 *Let Δ be a triangulation of (S_0, V) , the oriented surface of genus 0 with $n = |V| \geq 3$ marked points, let $\lambda \in \mathbb{R}^{E_\Delta}$, let $v_\infty \in V$, and let V° be defined by (53). Then there exists a unique solution to the minimization problem*

$$\begin{aligned} &\text{minimize } \bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}(u) \text{ for } u \in \mathbb{R}^{V^\circ}, \\ &\text{subject to the bounds constraints (61).} \end{aligned} \tag{83}$$

The rest of this section is concerned with proving Proposition 9.2, first the uniqueness statement, then the existence statement.

Uniqueness Assume the minimization problem (83) has a solution. To show the solution is unique, assume $u \in \mathbb{R}^{V^\circ}$ solves (83) and let $(\tilde{\Delta}, \tilde{\lambda})$ be defined by (57). By Theorem 7.18, $(\tilde{\Delta}, \tilde{\lambda})$ are realizable coordinates with distinguished vertex v_∞ .

Either the subcomplex $\tilde{\Delta}^\circ$ of cells not incident with v_∞ is a linear graph. In this case all constraints (61) are satisfied with equality, which determines u uniquely.

Or $\tilde{\Delta}^\circ$ is a triangulation of a closed disk. Then the second derivative of $\bar{\mathcal{E}}_{\tilde{\Delta}, \tilde{\lambda}}^{v_\infty}$ at u is positive semidefinite with one-dimensional kernel spanned by $\mathbf{1}_{V^\circ}$. This follows from (56) and the convexity of f (see Proposition 7.2 (iii)). Together with (62) and (63), this implies that

$$\bar{\mathcal{E}}_{\tilde{\Delta}, \tilde{\lambda}}^{v_\infty}(u') > \bar{\mathcal{E}}_{\tilde{\Delta}, \tilde{\lambda}}^{v_\infty}(u)$$

for any $u' \in \mathbb{R}^{V^\circ} \setminus \{u\}$ satisfying the constraints (61).

Existence To show that the continuous function $\bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}$ attains its minimum on the closed subset

$$D = \{u \in \mathbb{R}^{V^\circ} \mid u \text{ satisfies the bounds constraints (61)}\} \subseteq \mathbb{R}^{V^\circ},$$

it is enough to show that every unbounded sequence (u_n) in D has a subsequence (u_{n_k}) with $\bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}(u_{n_k}) \rightarrow +\infty$.

So let (u_n) be an unbounded sequence in D . Note that (u_n) is bounded from below by the constraints (61). Hence, after taking a subsequence if necessary, we may assume that for every $v \in V^\circ$ either $u_n(v)$ converges to a finite limit, or $u_n(v) \rightarrow +\infty$. Since $u_n(v) \rightarrow +\infty$ for at least one $v \in V^\circ$ and $\Theta_v = 2\pi$, Corollary 7.15 implies

$$\lim_n \bar{\mathcal{E}}_{\Delta, \lambda}^{v_\infty}(u_n) = +\infty.$$

This concludes the proof of Proposition 9.2, and hence of Theorem 1.1.

10 Discrete Conformal Equivalence and the Uniformization Spheres

In this section we recall the basic definitions of discrete conformal equivalence, and we discuss the equivalence of Rivin’s Theorem 1.1 and the discrete uniformization

theorem for spheres (see Theorem 10.5). The relation of discrete conformal equivalence was first defined for triangulated piecewise Euclidean surfaces. As explained in Definition 4.13, we use the notation (Δ, ℓ) for the triangulated piecewise Euclidean surface with triangulation Δ and edge lengths $\ell \in \mathbb{R}_{>0}^{E_\Delta}$.

Definition 10.1 (*Discrete conformal equivalence of triangulated surfaces*) The triangulated piecewise Euclidean surfaces (Δ, ℓ) and $(\Delta, \tilde{\ell})$ are *discretely conformally equivalent* if there is a function $u \in \mathbb{R}^{V_\Delta}$ such that for every edge $e \in E_\Delta$,

$$\tilde{\ell}(e) = e^{(u_{v_1(e)} + u_{v_2(e)})/2} \ell(e). \tag{84}$$

This definition is due to Luo [22]. It has the following interpretation in terms of hyperbolic geometry:

Proposition 10.2 (See [7, Thm. 5.1.2]) *On a triangulated piecewise Euclidean surface, one obtains a complete hyperbolic metric of finite area by equipping every triangle with the hyperbolic Klein metric induced by its circumcircle. Then the following statements are equivalent:*

- (i) *The triangulated piecewise Euclidean surfaces (Δ, ℓ) and $(\tilde{\Delta}, \tilde{\ell})$ are discretely conformally equivalent.*
- (ii) *The triangulated piecewise Euclidean surfaces (Δ, ℓ) and $(\tilde{\Delta}, \tilde{\ell})$ are isometric with respect to the induced hyperbolic metrics.*

The induced hyperbolic metric has Penner coordinates (Δ, λ) , where λ and ℓ are related by (8). Proposition 10.2 can also be seen by considering decorated ideal tetrahedra as shown in Fig. 3: The projection from the point v_∞ maps the Euclidean triangle $A_1 A_2 A_3$ in the horosphere centered at v_∞ to the ideal triangle $v_1 v_2 v_3$. The hyperbolic Klein metric induced on the Euclidean triangle $A_1 A_2 A_3$ by its circumcircle is the pullback of the hyperbolic metric of the ideal triangle $v_1 v_2 v_3$.

Note that Definition 10.1 requires the triangulations of both surfaces to be equal. Discrete conformal mapping problems based on this notion of discrete conformal equivalence can be solved using the variational principles introduced in [7]—provided a solution exists. The variational principle also implies strong uniqueness theorems for the solutions. But to prove any reasonable existence theorem for discrete conformal maps, it seems necessary to allow changing the triangulation. Proposition 10.2 motivates the following definition, which leads to strong uniformization theorems:

Definition 10.3 (*Discrete conformal equivalence of piecewise Euclidean surfaces*) Piecewise Euclidean metrics d and \tilde{d} on the oriented surface (S_g, V) of genus g with $n = |V|$ marked points are *discretely conformally equivalent* if the Delaunay triangulations of (S_g, V) with respect to the metrics d and \tilde{d} induce the same complete hyperbolic metric on the punctured surface $S_{g,n} = S_g \setminus V$.

Definition 10.3 is equivalent to the definition of Gu et al. [18]:

Proposition 10.4 *Two piecewise Euclidean metrics d and \tilde{d} on (S_g, V) are discretely conformally equivalent, if there is a sequence of triangulated piecewise Euclidean surfaces*

$$(\Delta_0, \ell_0), \dots, (\Delta_m, \ell_m)$$

such that:

- (i) The metric of (Δ_0, ℓ_0) is d and the metric of (Δ_m, ℓ_m) is \tilde{d} .
- (ii) Each Δ_i is a Delaunay triangulation of the piecewise Euclidean surface (Δ_i, ℓ_i) .
- (iii) If $\Delta_i = \Delta_{i+1}$, then (Δ_i, ℓ_i) and $(\Delta_{i+1}, \ell_{i+1})$ are discretely conformally equivalent in the sense of Definition 10.1.
- (iv) If $\Delta_i \neq \Delta_{i+1}$, then (Δ_i, ℓ_i) and $(\Delta_{i+1}, \ell_{i+1})$ are the same piecewise Euclidean surface with two different Delaunay triangulations Δ_i and Δ_{i+1} .

Proof This is a consequence of Proposition 10.2 and Theorems 4.6 and 4.14. □

The connection of realization problems for ideal polyhedra and discrete conformal equivalence in the sense of Definition 10.1 was observed in [7, Sect. 5.4]. With Definition 10.3, Rivin’s polyhedral realization Theorem 1.1 is equivalent to the following uniformization theorem for spheres:

Theorem 10.5 (Discrete uniformization of spheres) *For every piecewise Euclidean metric d on the 2-sphere (S_0, V) with $n = |V|$ marked points, there is a realization of (S_0, V) as a convex Euclidean polyhedron P with vertex set V , such that all vertices lie on the unit sphere and the induced piecewise Euclidean metric is discretely conformally equivalent to d . The polyhedron P is unique up to projective transformations of $\mathbb{RP}^3 \cong \mathbb{R}^3$ mapping the unit sphere to itself.*

The equivalence of both problems follows from Proposition 10.2, the Möbius invariance of discrete conformal equivalence [7, Sect. 2.5], the construction described in [7, Sect. 5.4], and Theorem 4.14.

The constructive variational proof of Theorem 1.1 (see Sect. 9) also shows that the uniformizing polyhedron of a piecewise Euclidean sphere with n vertices can be computed by solving a convex optimization problem with $n - 1$ variables.

11 Higher Genus and Prescribed Cone Angles

The variational method of proving Theorems 1.1 and 10.5 extends to other polyhedral realization and discrete uniformization problems. The following theorem was proved by Gu et al. [18]:

Theorem 11.1 *Let d be a piecewise Euclidean metric on (S_g, V) , the surface of genus g with $n = |V|$ marked points, and let $\Theta \in \mathbb{R}^V$ satisfy $\Theta > 0$ and the Gauss–Bonnet condition*

$$\frac{1}{2\pi} \sum_{v \in V} \Theta_v = 2g - 2 + n. \tag{85}$$

Then there exists a discretely conformally equivalent metric \tilde{d} on (S_g, V) such that the cone angle at each $v \in V$ is Θ_v . The metric \tilde{d} is uniquely determined up to scale.

The special case of $\Theta_v = 2\pi$ for all v is the uniformization theorem for tori:

Theorem 11.2 (Uniformization theorem for tori) *For every piecewise Euclidean metric d on (S_1, V) , the torus with $n = |V|$ marked points, there exists a flat metric \tilde{d} on (S_1, V) that is discretely conformally equivalent to d . The metric \tilde{d} is uniquely determined up to scale.*

Theorem 11.2 is equivalent to the following polyhedral realization theorem:

Theorem 11.3 *Every oriented complete hyperbolic surface of finite area that is homeomorphic to a punctured torus $S_{1,n}$ can be realized as a convex polyhedral surface in H^3 that is invariant under a faithful action of the fundamental group $\pi_1(S_1)$ on H^3 by parabolic isometries.*

Theorem 11.3 is a special case of a more general result of Fillastre [14, Thm. B], who used Alexandrov’s method to prove it. It seems the more general polyhedral realization theorem that is equivalent to Theorem 11.1 has not been treated. It would involve hyperbolic manifolds with one cusp, convex polyhedral boundary with ideal vertices, and with “particles”. Izestiev and Fillastre prove an analogous realization theorem for polyhedral surfaces with finite vertices instead of ideal ones [15]. They use a variational method that is analogous to the method presented in this article. Since the vertices are finite, they do not need the Epstein–Penner convex hull construction.

The following variational principle for Theorem 11.1 is simpler than the variational principle for the uniformization of spheres (see Theorem 7.18) because the minimization problem is unconstrained, no vertex is distinguished, and it involves the function $\mathcal{E}_{\Theta, \Delta, \lambda}$ instead of $\tilde{\mathcal{E}}_{\Delta, \lambda}^{v, \infty}$.

Theorem 11.4 (Variational principle for Theorem 11.1) *Let d be a piecewise Euclidean metric on (S_g, V) , the surface of genus g with $n = |V|$ marked points, and let $\Theta \in \mathbb{R}^V$. Let Δ be a straight triangulation of (S_g, V) and for each edge e let ℓ_e be length of edge e . Then the following statements are equivalent:*

- (i) *The metric of the piecewise flat surface $(\tilde{\Delta}, \tilde{\ell})$ is discretely conformally equivalent to d and has cone angle Θ_v at each vertex $v \in V$.*
- (ii) *The function $\mathcal{E}_{\Theta, \Delta, \lambda}$ attains its minimum at $u \in \mathbb{R}^V$, the realizable coordinates $(\tilde{\Delta}, \tilde{\lambda})$ are equivalent to $\text{Del}(\Delta, \Lambda^{\Delta, \lambda}(u))$ (see Definition 6.4), and $\tilde{\ell}$ satisfies (21).*

Proof This follows from Proposition 7.12 and Theorem 4.14. □

To prove Theorem 11.1 using the variational principle of Theorem 11.4, note the following:

- The Gauss–Bonnet condition (85) is equivalent to the scale invariance of $\mathcal{E}_{\Theta, \Delta, \lambda}$, that is,

$$\mathcal{E}_{\Theta, \Delta, \lambda}(u + h\mathbf{1}_V) = \mathcal{E}_{\Theta, \Delta, \lambda}(u). \tag{86}$$

This follows from (41).

- The uniqueness statement of Theorem 11.1 follows from the convexity of $\mathcal{E}_{\Theta, \Delta, \lambda}$, Proposition 7.12 (iii). If the Gauss–Bonnet condition is satisfied and $\mathcal{E}_{\Theta, \Delta, \lambda}$ attains its minimum at u and at u' , then $u - u' \in \mathbb{R}\mathbf{1}_V$.

- To prove the existence statement, proceed as in the proof of Theorem 1.1 (see Sect. 9). Note that due to the scale invariance (86) it is enough to consider unbounded sequences (u_n) in \mathbb{R}^V that are bounded from below.

A completely analogous theory of discrete conformal equivalence for triangulated piecewise hyperbolic surfaces, including a convex variational principle, was developed in [7, Sect. 6], see also [8]. A result analogous to Theorem 11.1 was proved by Gu et al. [17]. A corresponding realization result, analogous to Theorem 11.3 for higher genus, is also due to Fillastre [14, Thm. B']. To obtain a variational proof and a practical method for computation, one can translate the variational method developed here to the setting of piecewise hyperbolic surfaces. This is beyond the scope of this article.

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

References

1. Akiyoshi, H.: Finiteness of polyhedral decompositions of cusped hyperbolic manifolds obtained by the Epstein–Penner’s method. *Proc. Am. Math. Soc.* **129**(8), 2431–2439 (2001)
2. Alexandrov, A.D.: *Convex Polyhedra*. Springer Monographs in Mathematics. Springer, Berlin (2005)
3. Aurenhammer, F., Klein, R., Lee, D.-T.: *Voronoi Diagrams and Delaunay Triangulations*. World Scientific, Hackensack (2013)
4. Bao, X., Bonahon, F.: Hyperideal polyhedra in hyperbolic 3-space. *Bull. Soc. Math. France* **130**(3), 457–491 (2002)
5. Bobenko, A.I., Izmestiev, I.: Alexandrov’s theorem, weighted Delaunay triangulations, and mixed volumes. *Ann. Inst. Fourier (Grenoble)* **58**(2), 447–505 (2008)
6. Bobenko, A.I., Dimitrov, N., Sechelmann, S.: Discrete uniformization of polyhedral surfaces with non-positive curvature and branched covers over the sphere via hyper-ideal circle patterns. *Discrete Comput. Geom.* **57**(2), 431–469 (2017)
7. Bobenko, A.I., Pinkall, U., Springborn, B.A.: Discrete conformal maps and ideal hyperbolic polyhedra. *Geom. Topol.* **19**(4), 2155–2215 (2015)
8. Bobenko, A.I., Sechelmann, S., Springborn, B.: Discrete conformal maps: boundary value problems, circle domains, Fuchsian and Schottky uniformization. In: Bobenko, A.I. (ed.) *Advances in Discrete Differential Geometry*, pp. 1–56. Springer, Berlin (2016)
9. Bobenko, A.I., Springborn, B.A.: A discrete Laplace–Beltrami operator for simplicial surfaces. *Discrete Comput. Geom.* **38**(4), 740–756 (2007)
10. Bowers, J.C., Bowers, P.L., Pratt, K.: Rigidity of circle polyhedra in the 2-sphere and of hyperideal polyhedra in hyperbolic 3-space. *Trans. Am. Math. Soc.* **371**(6), 4215–4249 (2019)
11. Duffin, R.J.: Distributed and lumped networks. *J. Math. Mech.* **8**, 793–826 (1959)
12. Edelsbrunner, H.: *Geometry and Topology for Mesh Generation*. Cambridge Monographs on Applied and Computational Mathematics, vol. 7. Cambridge University Press, Cambridge (2001)
13. Epstein, D.B.A., Penner, R.C.: Euclidean decompositions of noncompact hyperbolic manifolds. *J. Differ. Geom.* **27**(1), 67–80 (1988)
14. Fillastre, F.: Polyhedral hyperbolic metrics on surfaces. *Geom. Dedicata* **134**, 177–196 (2008). Erratum **138**, 193–194 (2009)
15. Fillastre, F., Izmestiev, I.: Hyperbolic cusps with convex polyhedral boundary. *Geom. Topol.* **13**(1), 457–492 (2009)
16. Fortune, S.: Numerical stability of algorithms for 2D Delaunay triangulations. *Int. J. Comput. Geom. Appl.* **5**(1–2), 193–213 (1995)
17. Gu, X., Guo, R., Luo, F., Sun, J., Wu, T.: A discrete uniformization theorem for polyhedral surfaces II. *J. Differ. Geom.* **109**(3), 431–466 (2018)

18. Gu, X.D., Luo, F., Sun, J., Wu, T.: A discrete uniformization theorem for polyhedral surfaces. *J. Differ. Geom.* **109**(2), 223–256 (2018)
19. Indermitte, C., Liebling, T.M., Troyanov, M., Clémençon, H.: Voronoi diagrams on piecewise flat surfaces and an application to biological growth. *Theoret. Comput. Sci.* **263**(1–2), 263–274 (2001)
20. Izestiev, I.: A variational proof of Alexandrov’s convex cap theorem. *Discrete Comput. Geom.* **40**(4), 561–585 (2008)
21. Joswig, M., Löwe, R., Springborn, B.: Secondary fans and secondary polyhedra of punctured Riemann surfaces. *Exp. Math.* (2019). <https://doi.org/10.1080/10586458.2018.1477078>
22. Luo, F.: Combinatorial Yamabe flow on surfaces. *Commun. Contemp. Math.* **6**(5), 765–780 (2004)
23. Masur, H., Smillie, J.: Hausdorff dimension of sets of nonergodic measured foliations. *Ann. Math.* **134**(3), 455–543 (1991)
24. Milnor, J.: Hyperbolic geometry: the first 150 years. *Bull. Am. Math. Soc. (N.S.)* **6**(1), 9–24 (1982)
25. Moroianu, S., Schlenker, J.-M.: Quasi-Fuchsian manifolds with particles. *J. Differ. Geom.* **83**(1), 75–129 (2009)
26. Penner, R.C.: The decorated Teichmüller space of punctured surfaces. *Commun. Math. Phys.* **113**(2), 299–339 (1987)
27. Penner, R.C.: *Decorated Teichmüller Theory*. QGM Master Class Series. European Mathematical Society (EMS), Zürich (2012)
28. Pinkall, U., Polthier, K.: Computing discrete minimal surfaces and their conjugates. *Exp. Math.* **2**(1), 15–36 (1993)
29. Prosanov, R.: Ideal polyhedral surfaces in Fuchsian manifolds. [arXiv:1804.05893](https://arxiv.org/abs/1804.05893) (2018)
30. Rippa, S.: Minimal roughness property of the Delaunay triangulation. *Comput. Aided Geom. Des.* **7**(6), 489–497 (1990)
31. Rivin, I.: Intrinsic geometry of convex ideal polyhedra in hyperbolic 3-space. In: Gyllenberg, M., Persson, L.-E. (eds.) *Analysis, Algebra, and Computers in Mathematical Research*. Lecture Notes in Pure and Appl. Math., vol. 156, pp. 275–291. Dekker, New York (1994)
32. Sakuma, M., Weeks, J.R.: The generalized tilt formula. *Geom. Dedicata* **55**(2), 115–123 (1995)
33. Schlenker, J.-M.: Hyperbolic manifolds with polyhedral boundary. [arXiv:math/0111136](https://arxiv.org/abs/math/0111136) (2001)
34. Schlenker, J.-M.: A rigidity criterion for non-convex polyhedra. *Discrete Comput. Geom.* **33**(2), 207–221 (2005)
35. Springborn, B.A.: A variational principle for weighted Delaunay triangulations and hyperideal polyhedra. *J. Differ. Geom.* **78**(2), 333–367 (2008)
36. Thurston, W.P.: Shapes of polyhedra and triangulations of the sphere. In: Rivin, I., Rourke, C., Series, C. (eds.) *The Epstein Birthday Schrift*. Geometry & Topology Monographs, vol. 1, pp. 511–549. Geometry & Topology Publications, Coventry (1998)
37. Tillmann, S., Wong, S.: An algorithm for the Euclidean cell decomposition of a cusped strictly convex projective surface. *J. Comput. Geom.* **7**(1), 237–255 (2016)
38. Weeks, J.R.: Convex hulls and isometries of cusped hyperbolic 3-manifolds. *Topology Appl.* **52**(2), 127–149 (1993)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.