

RESEARCH ARTICLE

The moduli space of Harnack curves in toric surfaces

Jorge Alberto Olarte 

Institut für Mathematik, Technische Universität Berlin, Straße des 17. Juni 135, Berlin, Germany;
E-mail: olarte@math.tu-berlin.de.

Received: 17 October 2020; Revised: 1 April 2021; Accepted: 24 April 2021

2020 Mathematics Subject Classification: Primary – 14H50; Secondary – 14H10, 14M25, 52B20, 14T20

Abstract

In 2006, Kenyon and Okounkov [12] computed the moduli space of Harnack curves of degree d in $\mathbb{C}P^2$. We generalise their construction to any projective toric surface and show that the moduli space \mathcal{H}_Δ of Harnack curves with Newton polygon Δ is diffeomorphic to $\mathbb{R}^{m-3} \times \mathbb{R}_{\geq 0}^{n+g-m}$, where Δ has m edges, g interior lattice points and n boundary lattice points. This solves a conjecture of Crétois and Lang. The main result uses abstract tropical curves to construct a compactification of this moduli space where additional points correspond to collections of curves that can be patchworked together to produce a curve in \mathcal{H}_Δ . This compactification has a natural stratification with the same poset as the secondary polytope of Δ .

1. Introduction

Harnack curves are real algebraic plane curves inside a projective toric surface, introduced by Mikhalkin [16], with several remarkable properties. By definition, they have the maximum possible number of connected components for a given Newton polygon, and these components are arranged in a particular and unique way (see Theorem 2.1). Their amoebas are particularly special, since they are precisely the ones with maximal area [18]. Because of this, they have found applications in physics, where the dimer model is used to study crystal surfaces (see [13] for details). In this model, the limit of the shape of a crystal surface is given by the amoeba of a Harnack curve.

These curves are named after Axel Harnack, who constructed them in the projective plane to show that his upper bound on the number of connected components of a plane real algebraic curve is attained [9]. In the projective plane, the space of Harnack curves of degree d modulo the action of the torus $(\mathbb{C}^*)^2 \subseteq \mathbb{C}P^2$ was studied by Kenyon and Okounkov [12] to better understand the dimer model (equivalently, this is the space of amoebas of Harnack curves modulo translation). They show that this moduli space has global coordinates given by the areas of holes of the amoeba and the distances between consecutive tentacles. Therefore it is diffeomorphic to $\mathbb{R}_{\geq 0}^{(d+4)(d-1)/2}$. Crétois and Lang [3] generalised some of the techniques used in [12] to Harnack curves in any projective toric surface. They showed that given a lattice polygon Δ , the moduli space \mathcal{H}_Δ of Harnack curves with Newton polygon Δ is path connected, and they conjectured that it is also contractible.¹ We confirm this belief and further generalise the results of [12] to compute \mathcal{H}_Δ :

Theorem 1. *Let Δ be a lattice m -gon with g interior lattice points and n boundary lattice points. Then the moduli space \mathcal{H}_Δ of Harnack curves of Newton polytope Δ is diffeomorphic to $\mathbb{R}^{m-3} \times \mathbb{R}_{\geq 0}^{n+g-m}$.*

¹The conjecture appears as Remark 4.4 in the preprint version of [3]; the published version already cites our results.

The interior of \mathcal{H}_Δ corresponds to the smooth Harnack curves with transversal intersections with the axes of X_Δ . In the boundary, ovals may contract to double points, or the curves may not intersect with the axes transversally.

We further show that \mathcal{H}_Δ admits a compactification similar in spirit to the Deligne–Mumford compactification of $\mathcal{M}_{g,n}$. This compactification consists of what we call *Harnack meshes*. A Harnack mesh (see Theorem 5.4) consists of a regular subdivision of Δ and a Harnack curve with Newton polytope Δ_i for each facet Δ_i of the subdivision, with some gluing conditions that allow for the curves to be patchworked (using Viro’s method [22]) to produce a curve in \mathcal{H}_Δ . The space of Harnack meshes is naturally stratified into cells according to which regular subdivision is used in the patchworking recipe. The foregoing can be summed up in the following:

Theorem 2. *The space \mathcal{H}_Δ has a compactification $\overline{\mathcal{H}_\Delta}$ consisting of all Harnack meshes over Δ . Moreover, $\overline{\mathcal{H}_\Delta}$ has a cell complex structure whose poset is isomorphic to the face poset of the secondary polytope $\text{Sec}(\Delta \cap \mathbb{Z}^2)$.*

The structure of the paper is as follows. In section 2, we set notation and recall some background results on Harnack curves. Sections 3 and 4 are dedicated to proving Theorem 1 (Theorem 4.3), and section 5 is dedicated to proving Theorem 2 (Theorem 5.11).

Most of the proofs consist of showing that there are different parameters that can be taken as global coordinates for Harnack curves. In section 3 we consider the following diagram:

$$\left\{ \begin{array}{l} \text{Rational} \\ \text{Harnack} \\ \text{curves} \end{array} \right\} \hookrightarrow \left\{ \begin{array}{l} \text{Roots of} \\ \text{rational} \\ \text{parametrisation} \end{array} \right\} / \text{PSL}_2(\mathbb{R}) \xrightarrow{\tilde{\rho}} \left\{ \begin{array}{l} \text{Positions of} \\ \text{amoeba} \\ \text{tentacles} \end{array} \right\} / \mathbb{R}^2.$$

On the left we have the moduli space of rational Harnack curves, which we denote $\mathcal{H}_{0,\Delta}$; in the middle we have parametrisations $\phi : \mathbb{CP}^1 \rightarrow X_\Delta$ of Harnack curves modulo the action of $\text{PSL}_2(\mathbb{R})$ on \mathbb{CP}^1 ; and on the right we have the positions of the tentacles of the amoeba modulo translations of the amoeba. The main result of section 3 is that the map $\tilde{\rho}$ is a smooth embedding when restricted to the image of the first map.

In section 4 we show that the following are diffeomorphisms:

$$\left\{ \begin{array}{l} \text{Harnack curves} \\ \text{with fixed} \\ \text{tentacle positions} \end{array} \right\} \leftrightarrow \left\{ \begin{array}{l} \text{Bounded} \\ \text{Ronkin} \\ \text{intercepts} \end{array} \right\} \leftrightarrow \left\{ \begin{array}{l} \text{Areas of} \\ \text{holes of} \\ \text{the amoeba} \end{array} \right\}. \tag{1}$$

By putting together the two diagrams, we have

$$\mathcal{H}_\Delta \hookrightarrow \frac{\{\text{Tentacle positions}\} / \mathbb{R}^2}{\times \{\text{Bounded intercepts}\}} \rightarrow \frac{\{\text{All Ronkin}\}}{\{\text{intercepts}\}} / \mathbb{R}^3 \rightarrow \mathcal{M}_{g,n}^{\text{trop}}$$

where the \mathbb{R}^3 action in the the third space refers to translations of the graph of the Ronkin function. The first map is a smooth embedding by putting together the two previous diagrams.

In section 5 we look at the last two maps. The second map is a linear bijection between tentacle positions and unbounded Ronkin intercepts. The last map is given by what we call the *expanded spine*. Since $\mathcal{M}_{g,n}^{\text{trop}}$ is not a manifold (it is a tropical variety), this map can no longer be a diffeomorphism. However, we show that it is a piecewise linear embedding. We also show how Harnack meshes can be similarly embedded into the closure of the embedding $\mathcal{H}_\Delta \hookrightarrow \mathcal{M}_{g,n}^{\text{trop}}$, allowing us to construct the compactification $\overline{\mathcal{H}_\Delta}$.

We end the paper by suggesting some directions for future research in section 6. In particular, we conjecture that $\overline{\mathcal{H}_\Delta}$ is a regular CW-complex and we suggest a possible smooth structure on $\overline{\mathcal{H}_\Delta}$ as a manifold with generalised corners [11].

2. Preliminaries

2.1. Notation

We fix the following notation for the rest of the paper. As is usual in toric geometry, $M \cong N \cong \mathbb{Z}^2$ are the lattices of characters and 1-parametric subgroups of the algebraic 2-dimensional torus $(\mathbb{C}^*)^2$, respectively. (See [2] as a general reference for toric varieties.)

Let $\Delta \subset M \otimes \mathbb{R}$ be a convex lattice polygon. We write $\partial\Delta$ for the boundary of Δ and $\text{int}(\Delta)$ for the interior of Δ . We write Δ_M for the lattice points in Δ – that is, $\Delta_M = \Delta \cap M$. We use n and g to denote the number of lattice points in $\partial\Delta$ and $\text{int}(\Delta)$, respectively, and m for the number of edges of Δ . For any positive integer k , $[k]$ denotes the set $\{1, \dots, k\}$. We denote by $\Gamma_i, i = 1, \dots, m$, the edges of Δ in cyclic anticlockwise order. Let d_1, \dots, d_m be their respective integer lengths (i.e., $d_i = |\Gamma_i \cap M| - 1$). We say that a vector is primitive if it is the smallest integer vector in that direction. Similarly, we call a segment primitive if it does not contain integer points in its interior – that is, if orienting it produces a primitive vector. Let $u_i \in N$ be the primitive inner normal vector of Γ_i . We have the following equation:

$$\sum_{i=1}^m d_i u_i = 0. \tag{2}$$

To each $v = (v_1, v_2) \in M$ there is an associated Laurent monomial $x^v := x_1^{v_1} x_2^{v_2}$. The *Newton polygon* of a Laurent polynomial $f(x) = \sum_{v \in M} c_v x^v$ is the convex hull of $\{v \in M \mid c_v \neq 0\}$. For any subset $\Delta' \subseteq \Delta$, we write $f|_{\Delta'}(x) := \sum_{\substack{v \in M \\ v \in \Delta'}} c_v x^v$.

Given a lattice polygon Δ , there is an associated projective toric surface X_Δ whose geometry reflects the combinatorics of Δ . It contains a dense copy of the torus $(\mathbb{C}^*)^2$ where coordinate-wise multiplication extends to an action on all of X_Δ . For each edge Γ_i of Δ , there is a corresponding irreducible divisor L_i in X_Δ which is invariant under the action of the torus. We call these divisors the *axes* of X_Δ . Two axes intersect in a point if and only if they correspond to consecutive edges of Δ . We denote the real part of X_Δ as $\mathbb{R}X_\Delta$.

2.2. Harnack curves

Let f be a Laurent polynomial with real coefficients and Newton polygon Δ . The zeros of f define a curve $C^\circ \subset (\mathbb{C}^*)^2$. The closure of C° in X_Δ is a compact algebraic curve C . If C is smooth, its genus is equal to g [14]. The intersection of C with $\mathbb{R}X_\Delta$ is a real algebraic curve $\mathbb{R}C$. The intersection of C with an axis L_i is given by the restriction of f to Γ_i , which is, after a suitable change of variable, a polynomial of degree d_i . Therefore $L_i \cap C$ consists of exactly d_i points counted with multiplicities.

Definition 2.1 ([16, Definition 2]). Let Δ be a lattice polygon with g, m and the d_i s as already defined. A smooth real algebraic curve $\mathbb{R}C \subseteq \mathbb{R}X_\Delta$ is called a *smooth Harnack curve* if the following conditions hold:

- The number of connected components of $\mathbb{R}C$ is $g + 1$.
- Only one component of $\mathbb{R}C$ intersects $L_1 \cup \dots \cup L_m$. This component can be subdivided into m disjoint arcs, $\theta_1 \dots \theta_m$, in that order, such that $C \cap L_i = \theta_i \cap L_i$.

The components that are disjoint from $L_1 \cup \dots \cup L_m$ are called *ovals*.

Harnack curves were originally described as being in ‘cyclically maximal position’ in [16]. In the literature these curves are sometimes called ‘simple Harnack curves’. However, following [18, 13, 12] we omit the adjective ‘simple’ when referring to them (see [17, Remark 6.6]).

These curves are named after Axel Harnack because he showed in 1876 that smooth curves of genus g in the real projective plane have at most $g + 1$ connected components. To show that the bound was tight, he constructed the eponymous curves [9]. Curves which attain the maximum number of

components are called M -curves. These are the topic of the first part of Hilbert’s 16th problem, which seeks to classify all possible topological types of M -curves. When $\mathbb{R}X_\Delta = \mathbb{R}P^2$, Harnack curves are the M -curves such that only one component intersects the axes, and it does so in order. Mikhalkin proved that for any given Δ , if $\mathbb{R}C$ is a Harnack curve with Newton polygon Δ , then the topological type of $(\mathbb{R}X_\Delta, \mathbb{R}C, \mathbb{R}L_1 \cup, \dots, \cup \mathbb{R}L_n)$ is unique [16, Theorem 3].

Recall that a singular point in $\mathbb{R}C$ is an ordinary isolated double point if it is locally isomorphic to the singularity of $x_1^2 + x_2^2 = 0$.

Definition 2.2 ([18, Definition 3]). A (possibly singular) real algebraic curve $\mathbb{R}C \subseteq X_\Delta$ is a *Harnack curve* if

- the only singularities of $\mathbb{R}C$ are ordinary isolated double points away from the torus invariant divisors and
- replacing each singular point of $\mathbb{R}C$ by a small oval around it yields a curve $\mathbb{R}C'$ such that $(\mathbb{R}X_\Delta, \mathbb{R}C', \mathbb{R}L_1 \cup, \dots, \cup \mathbb{R}L_n)$ has the topological type of smooth Harnack curves.

Notice that any singular Harnack curve can be approximated by smooth Harnack curves. To see this, let f be a polynomial that vanishes on $\mathbb{R}C$ and define $g(x, y) := f(\lambda x, \lambda y)$ for a real number λ close to 1 but different from 1, so that the singular points of f and g are close but do not coincide. Then $f - \epsilon g$ vanishes on a smooth Harnack curve which approaches $\mathbb{R}C$ when ϵ tends to 0.

Let $\mathbb{R}[\Delta_M]$ be the vector space of real polynomials with Newton polygon contained in Δ . Since scaling all coefficients of f by the same constant does not change the curve $\mathbb{R}C$, we can identify the space of real curves with Newton polygon contained in Δ with $\mathbb{P}(\mathbb{R}[\Delta_M])$. The action of the torus $(\mathbb{R}^*)^2$ on $\mathbb{R}X_\Delta$ induces an action on $\mathbb{P}(\mathbb{R}[\Delta_M])$ given by $f(x_1, x_2) \mapsto f(r_1^{-1}x_1, r_2^{-1}x_2)$.

Definition 2.3. The *moduli space* \mathcal{H}_Δ of Harnack curves is the subspace of $\mathbb{P}(\mathbb{R}[\Delta_M]) / (\mathbb{R}^*)^2$ consisting of all (possibly singular) Harnack curves with Newton polygon Δ modulo the action of $(\mathbb{R}^*)^2$.

Given a curve $\mathbb{R}C \in \mathcal{H}_\Delta$, we say that a polynomial vanishes on $\mathbb{R}C$ if its zero locus is in the equivalence class given by $\mathbb{R}C$.

Remark. The notation \mathcal{H}_Δ was used in [3] to note the space of Harnack curves without taking them modulo the action of $(\mathbb{R}^*)^2$. It was defined with a more algebro-geometric language as follows: the space of curves with Newton polygon contained in Δ can be identified with the complete linear system $|D_\Delta|$ of the Cartier divisor D_Δ of X_Δ associated to Δ . Since X_Δ is a complete normal toric variety, $|D_\Delta|$ can be identified with the projectivisation of the space of global sections of the line bundle associated to Δ . Therefore \mathcal{H}_Δ can be defined as the subspace of $|D_\Delta|$ of Harnack curves, modulo the action of the torus $(\mathbb{R}^*)^2$ on $\mathbb{R}X_\Delta$. We emphasise that when we write \mathcal{H}_Δ , it is always modulo the torus action.

The case when Δ is the d th dilation of the unimodular triangle corresponds to degree d curves in $\mathbb{R}P^2$, and the moduli space is diffeomorphic to the closed orthant $\mathbb{R}_{\geq 0}^{(d+4)(d-1)/2}$ [12, Corollary 11].

2.3. Amoebas and the Ronkin function

Amoebas, which are essential to understanding Harnack curves, were defined in [7, Chapter 6], where details about them can be found.

Definition 2.4. Let $\text{Log} : (\mathbb{C}^*)^2 \rightarrow \mathbb{R}^2$ be the map

$$\text{Log}(z_1, z_2) := (\log |z_1|, \log |z_2|).$$

The *amoeba* of an algebraic curve C is $\mathcal{A}(C) := \text{Log}(C^\circ)$.

The amoebas of Harnack curves are specially well behaved:

Proposition 2.5 ([18]). Let $\mathbb{R}C$ be a real algebraic curve with Newton polygon Δ and $\mathcal{A} = \mathcal{A}(C)$ be its amoeba. The following are equivalent:

1. $\mathbb{R}C$ is Harnack curve.
2. The map $\text{Log}|_{C^\circ}$ is at most 2-to-1.
3. $\text{area}(\mathcal{A}) = \pi^2 \text{area}(\Delta)$.

For arbitrary curves, $\text{area}(\mathcal{A}) \leq \pi^2 \text{area}(\Delta)$ [20], so Harnack curves have amoebas with maximal area. Smooth Harnack curves are also characterised by having maximal curvature and by having totally real logarithmic Gauss maps [16, 19]. However, there are more general singular curves whose logarithmic Gauss maps are also totally real [15].

Each connected component of the complement of an amoeba is convex and has a point in Δ_M naturally associated to it, as we now show. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a Laurent polynomial. The Ronkin function $R_f : \mathbb{R}^2 \rightarrow \mathbb{R}$ of f , defined in [21], is

$$R_f(x) := \frac{1}{(2\pi\sqrt{-1})^2} \int_{\text{Log}^{-1}(x)} \frac{\log |f(z_1, z_2)|}{z_1 z_2} dz_1 dz_2.$$

The Ronkin function is convex [20]. Its gradient vector $\nabla R_f = (v_1, v_2)$ is given by

$$v_i(x) = \frac{1}{(2\pi\sqrt{-1})^2} \int_{\text{Log}^{-1}(x)} \frac{z_i \partial_{z_i} f(z_1, z_2)}{z_1 z_2 f(z_1, z_2)} dz_1 dz_2.$$

For any $x \in \mathbb{R}^2$ we have $\nabla R_f(x) \in \Delta$. If two points are in the same connected component of $\mathbb{R}^2 \setminus \mathcal{A}$, then their preimages under Log are homologous cycles in $(\mathbb{C}^*)^2 \setminus C^\circ$. This implies that ∇R_f is constant in each component, and it has integer coordinates by the residue theorem. Therefore, $\nabla R_f(x)$ induces an injection from the components of $\mathbb{R}^2 \setminus \mathcal{A}$ to Δ_M . The value that ∇R_f takes in a component of $\mathbb{R}^2 \setminus \mathcal{A}$ is called the *order* of that component, and we write E_ν for the component of order ν if it exists. For details of this construction, see [5].

To better understand amoebas, we review some facts about their behaviour [7, Section 6.1]. The component E_ν is bounded if and only if ν is in the interior of Δ . For each vertex ν of Δ , E_ν exists and contains a translation of $-\text{cone}(u_i, u_{i+1})$, where u_i and u_{i+1} are the inner normal vectors of the edges adjacent to ν . If ν is a lattice point in the relative interior of an edge Γ_i , E_ν is unbounded only in the direction $-u_i$. Parts of the amoeba extend to infinity in between the unbounded components of $\mathbb{R}^2 \setminus \mathcal{A}$, in direction u_i for some i . These are called the *tentacles* of the amoeba. Figure 1 serves as an illustration of what typical amoebas of Harnack curves look like.

For each $\nu \in \Delta_M$ such that E_ν exists, let $F_\nu : \mathbb{R}^2 \rightarrow \mathbb{R}$ be the affine linear function that coincides with R_f in E_ν . The *spine* of a curve C as defined in [20] is the corner locus of $\max F_\nu$, where \max is taken over all E_ν that exist. Notice that scaling f by a constant only changes R_f by an additive constant, so the spine of C is well defined.

The spine varies continuously for smooth curves. However, if E_ν vanishes for $\nu \in \text{int}(\Delta)$, then the spine changes abruptly. Fortunately, for Harnack curves there is an easy work-around. By the definition of singular Harnack curves, for each $\nu \in \text{int}(\Delta) \cap M$ such that $E_\nu = \emptyset$, there is an isolated double point p_ν in $\mathbb{R}C$ such that there is a smooth Harnack curve $\mathbb{R}C'$ arbitrarily close to $\mathbb{R}C$ with a component near p_ν with order ν . Therefore, $\nabla R_f(\text{Log}(p_\nu)) = \nu$. Let F_ν be the tangent plane of R_f at $\text{Log}(p_\nu)$.

Definition 2.6. Let $\mathbb{R}C$ be a Harnack curve. We call the corner locus of the piecewise affine linear convex function $\max_{\nu \in \Delta_M} F_\nu$ the *expanded spine* of C and denote it with $Y(C)$.

The expanded spine and the usual spine coincide if and only if $\mathbb{R}C$ is a smooth Harnack curve. The expanded spine varies continuously for Harnack curves, even singular ones. It has a cycle for each




Figure 1. The amoeba (blue), spine (black) and Newton polygon (green) of a Harnack curve.

$v \in \text{int}(\Delta) \cap M$. The bounded part of the expanded spine is a planar graph of genus g . This definition will be crucial in Section 5.2.

By definition, the expanded spine is a tropical plane curve. The intercepts c_v of the affine functions F_v are the coefficients of the tropical polynomial that vanishes on the expanded spine – that is,

$$\Upsilon(C) = \text{trop} \left(\bigoplus_v c_v \odot x^{\odot v} \right).$$

We call the numbers c_v the *Ronkin intercepts*. We call a Ronkin intercept c_v *bounded* if $v \in \text{int}(\Delta)$ and *unbounded* if $v \in \partial\Delta$. In other words, we say c_v is bounded if and only if whenever E_v exists, it is bounded.

2.4. Regular subdivisions

We review the basic definitions of regular subdivisions that are relevant to us; for more details, we refer the reader to [4].

A *subdivision* of Δ is a collection \mathcal{S} of subsets of Δ_M such that they satisfy the following:

- $\Delta = \bigcup_{B \in \mathcal{S}} \text{conv}(B)$.
- If $B \in \mathcal{S}$ and F is a face of $\text{conv}(B)$, then $F \cap B \in \mathcal{S}$.
- If $B, B' \in \mathcal{S}$, then $\text{relint}(\text{conv}(B)) \cap \text{relint}(\text{conv}(B')) = \emptyset$.

If the union $\bigcup \mathcal{S}$ of all elements of \mathcal{S} is equal to Δ_M , we call \mathcal{S} *full*.

We are interested only in a particular kind of subdivision, which can be obtained with the following construction. Consider a function $h : \Delta_M \rightarrow \mathbb{R}$, which we call a *height function*. The 3-dimensional polyhedron

$$P = \text{conv}(\{(v, t) \mid v \in \Delta_M \ t \geq h(v)\}) \subseteq M_{\mathbb{R}} \times \mathbb{R}$$

is unbounded only in the $(0, 0, 1)$ direction. For any face F of P , let $B_F := \{v \in \Delta_M \mid (v, h(v)) \in F\}$ be the collection of points in Δ_M which are lifted to lie in F . The *regular subdivision* of Δ_M induced by h is $\mathcal{S}(h) := \{B_F \mid F \text{ face of } P\}$.

Given a regular subdivision \mathcal{S} , consider the cone

$$\sigma(\mathcal{S}) := \{h \in \mathbb{R}^{\Delta_M} \mid \mathcal{S} = \mathcal{S}(h)\}.$$

If q is the restriction of any affine function $\mathbb{R}^2 \rightarrow \mathbb{R}$ to Δ_M , then $\mathcal{S}(h + q) = \mathcal{S}(h)$. This implies that this fan has a linearity space of dimension 3. The collection of all cones $\sigma(\mathcal{S})$ is a complete fan in \mathbb{R}^{Δ_M} called the *secondary fan* of Δ_M (see [7, Chapter 7]). The secondary fan happens to be the normal fan of a polytope $\text{Sec}(\Delta) \subseteq \mathbb{R}^{\Delta_M}$, called the *secondary polytope* of Δ_M .

Faces of the secondary polytope are naturally in bijection with cones in the secondary fan, which in turn are in bijection with regular subdivision. The dimension of the secondary polytope is $|\Delta_M| - 3$, since the linearity space of the secondary fan is 3-dimensional. More generally, given a face $F_{\mathcal{S}}$ of $\text{Sec}(\Delta_M)$ corresponding to a subdivision \mathcal{S} , we have $\dim(F_{\mathcal{S}}) = |\Delta_M| - \dim(\sigma(\mathcal{S}))$. The dimension of $\sigma(\mathcal{S})$ can be computed in the following way.

Given a height function $h : \Delta_M \rightarrow \mathbb{R}$, we can obtain another polyhedral complex $\mathcal{S}^*(h)$ which is a subdivision of $N_{\mathbb{R}}$ and is said to be a *dual subdivision* of $\mathcal{S}(h)$ of Δ . Two points $u_1, u_2 \in N_{\mathbb{R}}$ are in the same cell of $\mathcal{S}^*(h)$ if the linear functionals $(u_1, 1) : M_{\mathbb{R}} \times \mathbb{R} \rightarrow \mathbb{R}$ and $(u_2, 1) : M_{\mathbb{R}} \times \mathbb{R} \rightarrow \mathbb{R}$ are minimised in P by the same face of P . We should remark that we are considering $\mathcal{S}^*(h)$ as a set of open polyhedra inside $N_{\mathbb{R}}$, as opposed to $\mathcal{S}(h)$, which is a set of labels – that is, subsets of Δ_M – and that taking a different function $h' \in \sigma(\mathcal{S})$ may yield a different dual subdivision $\mathcal{S}^*(h')$.

The 1-skeleton of $\mathcal{S}^*(h)$ is a tropical curve, and its bounded part is a planar metric graph. The combinatorics of this graph depend only on \mathcal{S} ; disregarding its embedding into $N_{\mathbb{R}}$, we call it *the graph dual* to \mathcal{S} and denote it by $G_{\mathcal{S}}$. Vertices $V(G_{\mathcal{S}})$ correspond to facets of \mathcal{S} , the edges $E(G_{\mathcal{S}})$ correspond to interior edges of \mathcal{S} and the bounded facets $F(G_{\mathcal{S}})$ correspond to interior points in \mathcal{S} .

If \mathcal{S} is full, a height function in $\sigma(\mathcal{S})$ is determined by the length of the bounded edges of $\mathcal{S}^*(h)$ up to an affine transformation (adding an affine transformation to h affects $\mathcal{S}^*(h)$ only by translation). However, every bounded facet of $G_{\mathcal{S}}$ constitutes a cycle of edges and hence gives two relations (one for the x -coordinate and another for the y -coordinate) on the space of possible lengths of bounded edges of $\mathcal{S}^*(h)$. These relations are in fact independent and sufficient. So $\dim(\sigma(\mathcal{S})) = |E(G_{\mathcal{S}})| - 2|F(G_{\mathcal{S}})| + 3$. If \mathcal{S} is not full, then the heights of the missing points can be arbitrary, so long as they are sufficiently high. In general, we have

$$\dim(\sigma(\mathcal{S})) = |E(G_{\mathcal{S}})| - 2|F(G_{\mathcal{S}})| + \left| \Delta_M \setminus \bigcup \mathcal{S} \right| + 3. \tag{3}$$

2.5. Patchworking

We now give a basic overview of patchworking of real algebraic curves, a powerful tool developed by Viro [22] to construct curves with a prescribed topology.

The ingredients for (real) patchworking are a regular subdivision $\mathcal{S} = \mathcal{S}(h)$ of a polygon Δ and a real polynomial $f \in \mathbb{R}[\Delta_M]$. Let $\Delta_1, \dots, \Delta_s$ be the facets of \mathcal{S} . Then the polynomial $f|_{\Delta_i}$ defines a real curve $\mathbb{R}C_i \subseteq X_{\Delta_i}$. Suppose that every curve $\mathbb{R}C_i$ is smooth and intersects transversally the axes of X_{Δ_i} – that is, $\mathbb{R}C_i$ intersects each axis in d different points, where d is the integer length of the corresponding edge of Δ_i . Define

$$f_t(x) := \sum_{v \in \Delta_M} t^{h(v)} a_v x^v$$

and let $\mathbb{R}C_t \subseteq X_{\Delta}$ be the vanishing locus of f_t . The patchworking theorem by Viro [22] says that there exists $t_0 > 0$ small enough that for every $t \in (0, t_0]$, the topological type of $\mathbb{R}C_t$ can be computed from the topological type of each $\mathbb{R}C_i$ by gluing them in a certain way. We say $\mathbb{R}C_t$ is the result of patchworking the curves $\mathbb{R}C_i$. Given the curves $\mathbb{R}C_i$, there exist different f s such that $f|_{\Delta_i}$ vanishes on $\mathbb{R}C_i$. However, the topological type of the resulting curve $\mathbb{R}C_t$ depends only on the signs of each $f|_{\Delta_i}$ (real polynomials with the same zero locus differ only by scaling by a constant in \mathbb{R}^*).

We do not show in general how to do this computation; see [22] for that. We do, however, mention some important facts regarding Harnack curves. First, Mikhalkin showed that Harnack curves can be constructed using patchworking [16, Appendix]. There it is shown that Harnack curves are *T-curves* – that is, curves whose topological type can be obtained from patchworking using regular unimodular triangulations as regular subdivision. In that case, the signs of each coefficient of f contain all the relevant information, and this is known as combinatorial patchworking [10]. Consider the sign configuration $\Delta_M \rightarrow \{-1, 1\}$ given by $v \mapsto (-1)^{v_1 v_2}$. No matter the triangulation chosen, the result from patchworking with this sign configuration will always be a Harnack curve [16, Appendix]. Moreover, it is essentially (up to \mathbb{Z}_2^2) the only sign configuration whose patchwork is invariant under the chosen unimodular triangulation. These statements follow directly from the discussion in [7, Chapter 11 Section 5C].

Another important fact is that for any regular subdivision \mathcal{S} , if each curve $\mathbb{R}C_i$ is a Harnack curve, then there exists a choice of f_i such that the result from patchworking is a Harnack curve (see Theorem 5.6).

2.6. Cox coordinates

We now review Cox coordinates for toric surfaces, since it will help us understand the parametrisations of rational Harnack curves. For details, see [2, Chapter 5]. They are a generalisation of homogeneous coordinates in the projective space $\mathbb{C}P^d = (\mathbb{C}^{d+1} \setminus \{0\}) / \mathbb{C}^*$.

Let Δ be a Newton polygon and recall that u_1, \dots, u_m are the primitive inner normal vectors of Δ . Let $\alpha : (\mathbb{C}^*)^m \rightarrow (\mathbb{C}^*)^2$ be the group homomorphism

$$(z_1, \dots, z_m) \mapsto \left(\prod_{i=1}^m z_i^{u_{i1}}, \prod_{i=1}^m z_i^{u_{i2}} \right).$$

We have that $\ker(\alpha)$ is a subgroup of $(\mathbb{C}^*)^m$. Let Z be the subset of \mathbb{C}^m with at least three coordinates equal to 0 or at least two coordinates not cyclically consecutive that are equal to 0.

Proposition 2.7 ([2, Theorem 5.1.11]). *Let Δ be a lattice polygon. All $\ker(\alpha)$ -orbits of $\mathbb{C}^m \setminus Z$ are closed, and the quotient $(\mathbb{C}^m \setminus Z) / \ker(\alpha)$ is isomorphic to X_Δ as an algebraic variety.*

We write $[z_1 : \dots : z_m]_\Delta$ to denote the point in X_Δ corresponding to the orbit of $(z_1, \dots, z_m) \in \mathbb{C}^m \setminus Z$ under the action of $\ker(\alpha)$. We have

$$L_i = \{[z_1 : \dots : z_m]_\Delta \in X_\Delta \mid z_i = 0\}$$

and

$$\mathbb{R}X_\Delta = \{[z_1 : \dots : z_m]_\Delta \in X_\Delta \mid z_i \in \mathbb{R} \ \forall i\}.$$

Example 2.8. Let Δ be any rectangle with edges parallel to the \mathbb{R}^2 axes. The map $\alpha : M \rightarrow \mathbb{Z}^4$ is given by the matrix

$$\begin{pmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix}.$$

Then $\alpha : (\mathbb{C}^*)^4 \rightarrow (\mathbb{C}^*)^2$ is given by $(z_1, z_2, z_3, z_4) \mapsto (z_1 z_3^{-1}, z_2 z_4^{-1})$. The action of G consists of coordinate-wise multiplications by vectors of the form $(\lambda_1, \lambda_2, \lambda_1, \lambda_2)$, where $\lambda_1, \lambda_2 \in \mathbb{C}^*$. The set Z consists of the points where $z_1 = z_3 = 0$ or $z_2 = z_4 = 0$. So in this case X_Δ is isomorphic to $\mathbb{C}P^1 \times \mathbb{C}P^1$.

3. Rational Harnack curves

3.1. Parametrisations of rational Harnack curves

We start this section by describing a parametrisation of rational Harnack curves which was already used in [12] for $X_\Delta = \mathbb{CP}^2$ and more generally in [3]. We rewrite it using Cox homogeneous coordinates. Real rational curves in X_Δ with Newton polygon Δ can be parametrised by $\phi = [p_1 : \dots : p_m]_\Delta$, where each $p_i : \mathbb{CP}^1 \rightarrow \mathbb{C}$ is a homogeneous polynomial of degree d_i with real coefficients for $i \in [m]$ and no two polynomials have a common root. If the curve is Harnack, $\phi(\mathbb{RP}^1)$ is the 1-dimensional component of \mathbb{RC} . This implies that the roots of p_i are real and are ordered in the cyclic way according to Theorem 2.1 – that is, \mathbb{RP}^1 can be subdivided in m arcs $\theta_1, \dots, \theta_m$ ordered that way such that θ_i contains all the roots of p_i . In fact, this condition is sufficient for \mathbb{RC} to be a Harnack curve. This was shown in [12, Proposition 4] for $X_\Delta = \mathbb{CP}^2$, and it was noticed in [3, Equation (2)] that the same arguments work for any projective toric surface. So $\mathbb{RC}(\mathbb{CP}^1)$ is a Harnack curve if and only if, for some chart of \mathbb{CP}^1 , we have

$$\phi(t) = \left[b_1 \prod_{i=1}^{d_1} (t - a_{1,i}) : \dots : b_m \prod_{i=1}^{d_m} (t - a_{m,i}) \right]_\Delta, \tag{4}$$

where all $a_{i,j}$ are real, all b_i are real different from zero and

$$a_{1,1} \leq \dots \leq a_{1,d_1} < a_{2,1} \leq \dots \leq a_{2,d_2} < \dots < a_{m,1} \leq \dots \leq a_{m,d_m}. \tag{5}$$

We call $a_{1,1}, \dots, a_{m,d_m}$ the *roots* of ϕ . Composing ϕ with α yields a parametrisation for $C^\circ \subset (\mathbb{C}^*)^2$:

$$\alpha \circ \phi(t) = \left(\prod_{i=1}^m b_i^{u_{i1}} \prod_{j=1}^{d_i} (t - a_{i,j})^{u_{i1}}, \prod_{i=1}^m b_i^{u_{i2}} \prod_{j=1}^{d_i} (t - a_{i,j})^{u_{i2}} \right).$$

Let $\mathcal{H}_{0,\Delta}$ be the subspace of \mathcal{H}_Δ consisting of rational Harnack curves. The following generalises [12, Corollary 5]:

Proposition 3.1. *Let Δ be a lattice polygon with m sides and n lattice points in its boundary. Then $\mathcal{H}_{0,\Delta}$ is diffeomorphic to $\mathbb{R}^{m-3} \times \mathbb{R}_{\geq 0}^{n-m}$.*

Proof. The parametrisation given is unique up to the action of the projective special linear group $\text{PSL}_2(\mathbb{R})$ on the parameter t . This induces an action of $\text{PSL}_2(\mathbb{R})$ on the roots of ϕ . More explicitly, for $\psi \in \text{PSL}_2(\mathbb{R})$ the function $\phi \circ \psi^{-1}$ also parametrises C and has roots $\psi(a_{1,1}), \dots, \psi(a_{m,d_m})$. The roots $\psi(a_{1,1}), \dots, \psi(a_{m,d_m})$ are in the same cyclic order as in Eq. (5).

The action of $(\mathbb{R}^*)^2$ in $\mathbb{R}X_\Delta$ affects ϕ by changing the constants b_1, \dots, b_n but not the roots. The same is true for choosing different representatives of the Cox homogeneous coordinates of X_Δ in Eq. (4). So every rational Harnack curve is equivalent in \mathcal{H}_Δ to a curve with $b_1 = \dots = b_d = 1$. Therefore, the elements of \mathcal{H}_Δ corresponding to rational curves are uniquely determined by the roots $a_{1,1}, \dots, a_{m,d_m}$, up to the action of $\text{PSL}_2(\mathbb{R})$ on them.

The action of $\text{PSL}_2(\mathbb{R})$ can be fixed, for example, by considering the unique Möbius transformation $\psi \in \text{PSL}_2(\mathbb{R})$ such that $\psi(a_{1,1}) = 0, \psi(a_{2,1}) = 1$ and $\psi(a_{3,1}) = 2$. The map

$$\begin{aligned} \mathbb{RC} &\mapsto (\psi(a_{4,1}), \dots, \psi(a_{m,1})) \\ &\quad \times (\psi(a_{1,2}) - \psi(a_{1,1}), \dots, \psi(a_{m,d_m}) - \psi(a_{m,d_m-1})) \end{aligned}$$

is a diffeomorphism between $\mathcal{H}_{0,\Delta}$ and $\mathbb{R}^{m-3} \times \mathbb{R}_{\geq 0}^{n-m}$ (by identifying copies of \mathbb{R} with $(\phi(a_{i,d_i}), 0)$ for $i \geq 3$). □

The roots of ϕ are associated naturally to primitive segments of the edges of Δ . Theorem 3.1 says that we can take as global coordinates of $\mathcal{H}_{0,\Delta}$ the difference between two consecutive roots corresponding to the same edge of Δ , together with the first root of each edge except the first three edges.

3.2. The positions of the tentacles

Following [12], we now make a useful change of global coordinates in $\mathcal{H}_{0,\Delta}$. Instead of the roots of ϕ , we will use the positions of the tentacles of the amoeba which correspond to boundary points $C \setminus C^\circ$.

Let $J : N \rightarrow M$ be the 2×2 matrix that rotates vectors $\pi/2$ clockwise – that is, $J = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$. Observe that Ju_i is a character, which maps $[x_1, \dots, x_m]_\Delta \mapsto \prod_{k=1}^m x_k^{u_k \wedge u_i}$, where $u_j \wedge u_i \in \mathbb{R}$ is the determinant of the 2×2 matrix whose columns are u_j and u_i in that order. In other words, $u_i \wedge u_j = \langle Ju_j, u_i \rangle$. This is not well defined over all X_Δ . We take $0^0 = 1$ by convention, so Ju_i is well defined over L_i except the torus invariant points. In fact, the ring of functions of L_i (without the torus invariant points) consists of Laurent polynomials on Ju_i .

Definition 3.2. Let $\mathbb{R}C$ be a rational Harnack curve parametrised by ϕ as in Eq. (4). For $1 \leq i \leq m$ and $1 \leq j \leq d_i$, the position of the (i, j) tentacle of the amoeba is

$$\log \left| \phi(a_{i,j})^{Ju_i} \right|.$$

Explicitly,

$$\log \left| \phi(a_{i,j})^{Ju_i} \right| = \sum_{k \neq i}^m \sum_{l=1}^{d_k} u_k \wedge u_i (\log |b_i| + \log |a_{i,j} - a_{k,l}|).$$

Since the curve does not change with acting on the roots by $\text{PSL}_2(\mathbb{R})$, the positions of the tentacles of the amoeba do not change either. However, the tentacle positions are not invariant under the action of $(\mathbb{R}^*)^2$ on $\mathbb{R}X_\Delta$. Concretely, multiplying $\mathbb{R}C$ by $r \in (\mathbb{R}^*)^2$ translates the amoeba by $\text{Log}|r|$, thus changing the position of the (i, j) tentacle by $\langle Ju_i, \text{Log}|r| \rangle$. Thus, the $(\mathbb{R}^*)^2$ action on $\mathbb{R}X_\Delta$ induces an \mathbb{R}^2 action on the position of the tentacles by translations of the amoeba.

Consider the maps $\rho_{i,j} : \mathbb{R}^n \rightarrow \mathbb{R}$ given by

$$\rho_{i,j} := \sum_{k \neq i}^n \sum_{l=1}^{d_k} u_k \wedge u_i \log |a_{i,j} - a_{k,l}|.$$

This is almost the position of the (i, j) tentacle, except that we drop the $\log |b_i|$ terms. The maps $\rho_{i,j}$ are well defined on the space of roots satisfying Eq. (5). Together they form a map $\rho : \mathbb{R}^n \rightarrow \mathbb{R}^n$ invariant under the $\text{PSL}_2(\mathbb{R})$ action on the roots. Additionally, consider the action of \mathbb{R}^2 on the target space \mathbb{R}^n of ρ given by translations of the amoeba – that is,

$$r \cdot \rho_{i,j} := \rho_{i,j} + \langle Ju_i, r \rangle. \tag{6}$$

By construction, ρ descends to a map $\tilde{\rho}$ making the following diagram commutative:

$$\begin{array}{ccc} \mathbb{R}^n & \overset{\rho}{\dashrightarrow} & \mathbb{R}^n \\ \downarrow & & \downarrow \\ \mathbb{R}^n / \text{PSL}_2(\mathbb{R}) & \overset{\tilde{\rho}}{\dashrightarrow} & \mathbb{R}^{n-2}. \end{array}$$

The left downward arrow is the quotient of the space of roots by $\text{PSL}_2(\mathbb{R})$, and the right downward arrow is the quotient of \mathbb{R}^n by the action of \mathbb{R}^2 already described. By Theorem 3.1, we can identify $\mathcal{H}_{0,\Delta}$ with the space of roots that satisfy Eq. (5) modulo the $\text{PSL}_2(\mathbb{R})$ action, so $\tilde{\rho}$ is a well-defined map on $\mathcal{H}_{0,\Delta}$. The main objective of this section is to prove the following:

Theorem 3.3. *The restriction $\tilde{\rho}|_{\mathcal{H}_{0,\Delta}}$ is a smooth embedding.*

This is a generalisation of [12, Theorem 4], where the same statement is proven for the case where Δ is a dilated unit triangle. To prove that $\tilde{\rho} |_{\mathcal{H}_{0,\Delta}}$ is a diffeomorphism, we show that it is proper (Section 3.3) and that the differential is injective (Section 3.4).

Before going to the proof, let us show a concrete diffeomorphism between the positions of the tentacles and $\mathbb{R}^{m-3} \times \mathbb{R}_{\geq 0}^{n-m}$. For the semi-bounded components, consider the distance between two parallel tentacles

$$\rho_{i,j+1} - \rho_{i,j} \quad \text{for } 1 \leq i \leq m, 1 \leq j \leq d_i - 1.$$

Notice that $\rho_{i,j+1} - \rho_{i,j}$ is invariant under the \mathbb{R}^2 -action of translating the amoeba. For the unbounded components, take as coordinates

$$\tilde{\rho}_{i,1} \quad \text{for } 4 \leq i \leq m,$$

where $\tilde{\rho}_{i,1}$ is the position of the $(i, 1)$ -tentacle after translating the amoeba, so that the positions of the $(1, 1)$ - and $(2, 1)$ -tentacles are both 0. It is straightforward to see that the image satisfies

$$\sum_{i=1}^m \sum_{j=1}^{d_i} \rho_{i,j} = 0 \quad \text{and} \quad \rho_{i,j+1} - \rho_{i,j} \geq 0 \quad \text{for all } i, j. \tag{7}$$

In particular, $\tilde{\rho}_{3,1}$ is determined by the rest of the coordinates. The space described by the equation and inequalities is simply connected. This is important, since we will use the following global diffeomorphism theorem, which was known by Hadamard. A proof of it can be found in [8].

Proposition 3.4. *A local diffeomorphism between two manifolds which is proper and such that the image is simply connected is a diffeomorphism.*

3.3. Properness

To prove that $\tilde{\rho} |_{\mathcal{H}_{0,\Delta}}$ is proper, we make use of the following lemma:

Lemma 3.5. *Set $x_1, \dots, x_n, y_1, \dots, y_n \in \mathbb{R}$ such that $x_1 \leq \dots \leq x_n, \sum_{i=1}^n y_i = 0$ and there exists j such that $y_i < 0$ for $i < j$ and $y_i \geq 0$ for $i > j$. Then*

$$\sum_{i=1}^n x_i y_i \geq 0.$$

Proof. We do induction on n . For $n = 1$ we have $y_1 = 0$, so the sum is 0. Let $n > 1$ and suppose the inequality holds for fewer than n terms. Subtract from the left-hand-side $y_n(x_n - x_{n-1})$, which is a nonnegative number. The result is

$$\sum_{i=1}^{n-2} x_i y_i + x_{n-1}(y_{n-1} + y_n),$$

which is nonnegative by application of the induction hypothesis to x_1, \dots, x_{n-1} and $y_1, \dots, y_{n-2}, y_{n-1} + y_n$. □

Proposition 3.6. *The map $\tilde{\rho} |_{\mathcal{H}_{0,\Delta}}$ is proper when the codomain is restricted to the quotient of the space described by formula (7).*

Proof. We want to show that the preimage of a compact set is a bounded set in $\mathcal{H}_{0,\Delta} \cong \mathbb{R}^{m-3} \times \mathbb{R}_{\geq 0}^{n-m}$. First we bound the parameters corresponding to the $\mathbb{R}_{\geq 0}^{n-m}$ part of $\mathcal{H}_{0,\Delta}$. These correspond to the difference between roots along the same edge of Δ , which are trivially bounded from below by 0. Without loss of

generality, consider the roots $a_{1,j-1}$ and $a_{1,j}$ for $1 < j \leq d_1$. We fix the $\text{PSL}_2(\mathbb{R})$ action on the space of roots by setting $a_{1,j-1} = -1$, $a_{2,1} = 0$ and $a_{m,d_m} = 1$. To show that $a_{1,j} - a_{1,j-1}$ attains its upper bound, we show that $a_{1,j}$ cannot be arbitrarily close to $a_{2,1} = 0$. By assumption, the difference between the position of the tentacles

$$\rho_{1,j} - \rho_{1,j-1} = \sum_{i=2}^n \sum_{k=1}^{d_i} u_i \wedge u_1 (\log |a_{i,k} - a_{1,j}| - \log |a_{i,k} + 1|) \tag{8}$$

is bounded. We have $a_{1,j} \in [-1, 0)$ and $a_{i,k} \in [0, 1]$ if $i \geq 2$. The function $\log |x - a_{1,j}| - \log |x + 1|$ is increasing in $[0, 1]$. On the other hand,

$$\sum_{i=2}^n \sum_{k=1}^{d_i} u_i \wedge u_1 = \left(\sum_{i=1}^n d_i u_i \right) \wedge u_1 = 0$$

by Eq. (2). Notice that there exists $k \in \{2, \dots, n\}$ such that $u_i \wedge u_1 \geq 0$ for $i \leq k$ and $u_i \wedge u_1 \leq 0$ for $i > k$ ($u_i \wedge u_1 = 0$ if and only if Γ_i is parallel to Γ_1).

So applying Theorem 3.5 to the sequences $\log |a_{i,j} - a_{1,j}| - \log |a_{i,j} + 1|$ and $u_i \wedge u_1$ (repeated d_i times), we get that $\rho_{1,j} - \rho_{1,j-1}$ is nonnegative. More importantly, since $u_2 \wedge u_1 < 0 < u_n \wedge u_1$, we can subtract

$$\begin{aligned} & (\log |a_{m,d_m} - a_{1,j}| - \log |a_{m,d_m} + 1|) - (\log |a_{2,1} - a_{1,j}| - \log |a_{2,1} + 1|) \\ &= \log |1 - a_{1,j}| - \log |2| - \log |a_{1,j}| \\ &\geq -\log |2| - \log |a_{1,j}| \end{aligned}$$

from the right-hand side in Eq. (8), and again by Theorem 3.5 the result is still nonnegative. This implies $\rho_{1,j} - \rho_{1,j-1} \geq -\log |2| - \log |a_{1,j}|$, which is arbitrarily large if $a_{1,j}$ is arbitrarily close to 0. Since we assumed that $\rho_{1,j} - \rho_{1,j-1}$ is bounded, $a_{1,j}$ is not arbitrarily close to 0.

Now we turn our attention to the \mathbb{R}^{m-3} component, assuming $m > 3$. We choose a different representative of the $\text{PSL}_2(\mathbb{R})$ -orbit on the roots by setting $a_{1,1} = 0$, $a_{3,1} = 1$ and $a_{m,d_m} = 2$. We will show that $a_{2,1}$ cannot be arbitrarily close to $a_{1,1} = 0$. That implies that the parameter $a_{2,1}$ has a minimum in the preimage under $\tilde{\rho}$ of any compact set, and it attains it by continuity. Analogously, all other bounds regarding the \mathbb{R}^{m-3} component are achieved, and we conclude that the preimage is compact.

To fix the \mathbb{R}^2 action on the positions of the tentacles, we assume that the positions of the $(1, 1)$ -tentacle and the (m, d_m) -tentacle are both 0. This is translating the amoeba by the vector

$$w = \frac{-\rho_{m,d_m}}{u_1 \wedge u_m} u_1 + \frac{-\rho_{1,1}}{u_m \wedge u_1} u_m.$$

We show that if the position of the second tentacle after this translation,

$$\hat{\rho}_{2,1} = \rho_{2,1} - \rho_{m,d_m} \frac{u_1 \wedge u_2}{u_1 \wedge u_m} - \rho_{1,1} \frac{u_m \wedge u_2}{u_m \wedge u_1},$$

is bounded from below, then $a_{2,1}$ is not arbitrarily close to $a_{1,1} = 0$.

By Theorem 3.5, ρ_{m,d_m} is nonpositive, since $\log |x - 2|$ is a decreasing function in $[0, 2)$ and $u_m \wedge u_i$ is negative for the first values of i and positive for the latter. Since $u_1 \wedge u_2 < 0 < u_1 \wedge u_m$, we have that $-\rho_n \frac{u_1 \wedge u_2}{u_1 \wedge u_m}$ is nonpositive.

In both $\rho_{2,1}$ and $\rho_{1,1}$ there is a $\log |a_i|$ term which is arbitrarily large in absolute value if $a_{2,1}$ is close to 0. As $a_{i,j} > 1$ if $i \geq 3$, the only terms in $\rho_{1,1}$ and $\rho_{2,1}$ which can be arbitrarily large in absolute value are those corresponding to $a_{2,j}$ and $a_{1,j}$, respectively.

In other words, the part that could grow arbitrarily large in absolute value in $\rho_{1,1}$ is

$$\sum_{j=1}^{d_2} u_2 \wedge u_1 \log |a_{2,j}|,$$

and in $\rho_{2,1}$ it is

$$\sum_{j=1}^{d_1} u_1 \wedge u_2 \log |a_{2,1} - a_{1,j}|.$$

Notice that

$$|\log |a_{2,1} - a_{1,j}|| \geq |\log |a_{2,1}|| \geq |\log |a_{2,j}||. \tag{9}$$

Let c_1 be the real number such that

$$\sum_{j=1}^{d_1} u_1 \wedge u_2 \log |a_{2,1} - a_{1,j}| = c_1 u_1 \wedge u_2 \log |a_{2,1}|.$$

By Eq. (9) we have $c_1 \geq d_1$. Similarly, if c_2 is such that

$$\sum_{j=1}^{d_2} u_2 \wedge u_1 \log |a_{2,j}| = c_2 u_2 \wedge u_1 \log |a_{2,1}|,$$

then by Eq. (9), $c_2 \leq d_2$. So the part of $\hat{\rho}_{2,1}$ which grows in absolute value is

$$\begin{aligned} & c_1 u_1 \wedge u_2 \log |a_{2,1}| - c_2 u_2 \wedge u_1 \log |a_{2,1}| \frac{u_m \wedge u_2}{u_m \wedge u_1} \\ &= \log |a_{2,1}| \frac{u_1 \wedge u_2}{u_m \wedge u_1} (c_1 u_m \wedge u_1 + c_2 u_m \wedge u_2) \\ &= \log |a_{2,1}| \frac{u_1 \wedge u_2 \cdot u_m \wedge (c_1 u_1 + c_2 u_2)}{u_m \wedge u_1}. \end{aligned}$$

Notice that $-d_1 u_1 - d_2 u_2$ is the inner normal vector of the third side of the triangle formed by Γ_1 and Γ_2 , so $u_m \in \text{cone}(-d_1 u_1 - d_2 u_2, u_1)$, since $m > 3$. Thus, $u_m \wedge (d_1 u_1 + d_2 u_2) > 0$. Because $c_1 \geq d_1$ and $c_2 \leq d_2$, we have $c_1 u_1 + c_2 u_2 \in \text{cone}(u_1, d_1 u_1 + d_2 u_2)$. As $u_m \wedge u_1$ is also positive, we have $u_m \wedge (c_1 u_1 + c_2 u_2) > 0$. We conclude that

$$\frac{u_1 \wedge u_2 \cdot u_m \wedge (c_1 u_1 + c_2 u_2)}{u_m \wedge u_1} > 0,$$

which implies that $\hat{\rho}_{2,1}$ is negative and arbitrarily large in absolute value if $a_{2,1}$ is arbitrarily close to 0. □

3.4. The Jacobian of ρ

In this subsection we prove that $\tilde{\rho}|_{\mathcal{J}\mathcal{C}_{0,\Delta}}$ is a local diffeomorphism.

From now on, we slightly change the notation we have used so far to simplify the exposition and the computations. Instead of labelling the roots by pairs (i, j) with $i \in [m]$ and $j \in [d_i]$, we relabel them as a_1, \dots, a_n in the global cyclic order. Similarly, we relabel the u_i s to agree with the labelling of the roots; that is, we have vectors u_1, \dots, u_n , where u_i is the primitive inner normal vector of the edge of

Δ that corresponds to the axis in which $\phi(a_i)$ vanishes. With this notation, we have

$$\alpha(\phi)(t) = \prod_{i=1}^n (t - a_i)^{u_i},$$

and the a_i and u_i correspond to a parametrisation of a Harnack curve if

- $\sum_{i=1}^n u_i = 0$.
- $a_1 \leq \dots \leq a_n$.
- The u_1, \dots, u_n are ordered anticlockwise.

We write $a = (a_1, \dots, a_n)$ and $U = (u_1, \dots, u_n)$ for short, and we say that a and U are *cyclically ordered* if they satisfy these conditions.

Now we consider the Jacobian matrix D of ρ at a given point a . We have

$$D_{i,j} = \begin{cases} \frac{u_i \wedge u_j}{a_i - a_j} & \text{if } a_i \neq a_j, \\ 0 & \text{if } a_i = a_j \text{ but } i \neq j, \\ -\sum_{k \neq i} D_{i,k} & \text{if } i = j. \end{cases}$$

In general, D is a matrix that depends on a and U , so we denote it as $D(a, U) = D(a_1, \dots, a_n; u_1, \dots, u_n)$.

Proposition 3.7. *The map $\tilde{\rho}|_{\mathcal{H}_{c_0,\Delta}}$ is a local diffeomorphism.*

Proof. Let a and U be cyclically ordered. Let $T_a \text{PSL}_2(\mathbb{R})a$ be the tangent space of the orbit of a under the $\text{PSL}_2(\mathbb{R})$ action at a , and similarly let K be the kernel of the quotient $\mathbb{R}^n \rightarrow \mathbb{R}^n/\mathbb{R}^2$ by the \mathbb{R}^2 action defined in Eq. (6). In other words,

$$K = \{(r \wedge u_1, \dots, r \wedge u_n) \mid r \in \mathbb{R}^2\},$$

which is a linear space. Let us look at the relation of these spaces with D .

To compute the tangent space $T_a \text{PSL}_2(\mathbb{R})a$, recall that Möbius transformations are of the form

$$t \mapsto \frac{at + b}{ct + d}.$$

We see that

$$\left. \frac{\partial}{\partial \epsilon} t + \epsilon \right|_{\epsilon=0} = 1, \quad \left. \frac{\partial}{\partial \epsilon} (1 + \epsilon)t \right|_{\epsilon=0} = t, \quad \left. \frac{\partial}{\partial \epsilon} \frac{t}{\epsilon t + 1} \right|_{\epsilon=0} = -t^2,$$

so $T_a \text{PSL}_2(\mathbb{R})a$ is spanned by the vectors $(1, \dots, 1)$, (a_1, \dots, a_n) and (a_1^2, \dots, a_n^2) . Since

$$\sum_{j=1}^n D_{i,j} = D_{i,i} - D_{i,i} = 0$$

and

$$\sum_{j=1}^n a_j D_{i,j} = \sum_{j=1}^n \frac{a_j \cdot u_i \wedge u_j}{a_i - a_j} - \sum_{j=1}^n \frac{a_i \cdot u_i \wedge u_j}{a_i - a_j} = -\sum_{j=1}^n u_i \wedge u_j = 0,$$

both $(1, \dots, 1)$ and (a_1, \dots, a_n) are in the kernel of D . In Theorem 3.8 we will prove that the kernel of D is 2-dimensional, so the vector (a_1^2, \dots, a_n^2) is not in the kernel. However, since ρ descends to the map $\tilde{\rho}$, we have $D \cdot (a_1^2, \dots, a_n^2)^\top \in K$.

Since D is symmetric, its image is the orthogonal complement of its kernel. So the image of D is orthogonal to $(1, \dots, 1)$ and (a_1, \dots, a_n) . On the other hand, K is always orthogonal to $(1, \dots, 1)$. By a similar argument as in the proof of Theorem 3.6, we have $\sum_{i=1}^n a_i u_i \neq 0$, so (a_1, \dots, a_n) is never orthogonal to K . Thus, the intersection of the image of D with K is exactly 1-dimensional, so it is spanned by (a_1^2, \dots, a_n^2) . This implies that the Jacobian of $\tilde{\rho}$ is injective, since no vector outside $T_a \text{PSL}_2(\mathbb{R})a$ vanishes under the composition of D and the quotient $\mathbb{R}^n \rightarrow \mathbb{R}^n/\mathbb{R}^2$. \square

Now the only thing left to prove is the following:

Proposition 3.8. *If a and U are cyclically ordered, then the rank of $D(a, U)$ is $n - 2$.*

Proof. In [12, Theorem 4] the authors prove this for the case where Δ is the dilation of a the unit triangle. They do so by showing that D is a sum of 3×3 -block semidefinite positive matrices of rank 1, each corresponding to the Jacobian of the unimodular triangle case. We here generalise this for any polygon Δ . Let e_1, e_2, e_3 be the primitive normal vector of the standard unimodular triangle in clockwise order and let $T(a_i, a_j, a_k) = D(a_i, a_j, a_k, e_1, e_2, e_3)$ (see [12, Equation (4.8)]). We have that $T(a_i, a_j, a_k)$ is a rank 1 matrix with kernel generated by $(1, 1, 1)$ and (a_i, a_j, a_k) . We obtain the nonzero eigenvalue by computing the image under D of the cross-product of the two vectors in the kernel, $(1, 1, 1) \times (a_i, a_j, a_k)$. The eigenvalue is

$$\frac{(a_i - a_j)^2 + (a_j - a_k)^2 + (a_k - a_i)^2}{(a_i - a_j)(a_j - a_k)(a_k - a_i)},$$

which is always positive when $a_i < a_j < a_k$.

Let $T_{i,j,k}(a_i, a_j, a_k)$ be the $n \times n$ matrix that restricts to $T(a_i, a_j, a_k)$ in the 3×3 submatrix with indices $\{i, j, k\}$ and is zero elsewhere. We will show that if a and U are cyclically ordered, then $D(a, U)$ is a positive sum of matrices of the form $T_{i,j,k}(a_i, a_j, a_k)$. In other words, we want to show that

$$D(a, U) \in \text{cone}(\{T_{i,j,k}(a_i, a_j, a_k) \mid 1 \leq i < j < k \leq n\}).$$

To do so, we write each u_i in the unique form $u_i = x_i e_1 + y_i e_2 + z_i e_3$, where $x_i, y_i, z_i \geq 0$ and at most two are positive. With this notation we have

$$u_i \wedge u_j = x_i y_j + y_i z_j + z_i x_j - x_i z_j - y_i x_j - z_i y_j$$

and

$$\sum_{i=1}^n x_i = \sum_{i=1}^n y_i = \sum_{i=1}^n z_i =: c.$$

For $i < j < k$, let

$$q_{i,j,k} = \det \begin{pmatrix} x_i & y_i & z_i \\ x_j & y_j & z_j \\ x_k & y_k & z_k \end{pmatrix}.$$

For all $i < j < k$, we have $q_{i,j,k} \geq 0$. To see that, note that the vectors (x_i, y_i, z_i) are ordered cyclically along

$$\mathbb{R}^2 \times \{0\} \cup \mathbb{R} \times \{0\} \times \mathbb{R} \cup \{0\} \times \mathbb{R}^2,$$

since they project to U under $\begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \end{pmatrix}$. Therefore, by the right-hand rule, the determinant of that matrix is nonnegative. Now we claim that

$$\sum_{i < j < k} q_{i,j,k} T_{i,j,k}(a_i, a_j, a_k) = cD(a, U).$$

To verify this claim, look at the coefficient of $\frac{1}{a_i - a_j}$ in the (i, j) -entry of the left-hand side. It is equal to

$$\begin{aligned}
 & x_i y_j \left(\sum_{k=1}^n z_k \right) + y_i z_j \left(\sum_{k=1}^n x_k \right) + z_i x_j \left(\sum_{k=1}^n y_k \right) \\
 & - x_i z_j \left(\sum_{k=1}^n y_k \right) - y_i x_j \left(\sum_{k=1}^n z_k \right) - z_i y_j \left(\sum_{k=1}^n x_k \right) = c \cdot u_i \wedge u_j.
 \end{aligned}$$

Therefore, $D(a, U)$ is the sum of positive semidefinite matrices, so its kernel is the intersection of the kernels of all of the summands. This implies that the kernel is the span of $(1, \dots, 1)$ and (a_1, \dots, a_n) . \square

Proof of Theorem 3.3. By Theorems 3.6 and 3.7, $\tilde{\rho}$ is a proper local diffeomorphism whenever a and U are cyclically ordered. By Theorem 3.4, $\tilde{\rho}|_{\mathcal{H}_{0,\Delta}}$ is a diffeomorphism onto the space defined by Eq. (7) modulo \mathbb{R}^2 . \square

4. From $\mathcal{H}_{0,\Delta}$ to \mathcal{H}_Δ

The reason for the change of coordinates by $\tilde{\rho}$ is that we are fixing $C \setminus C^\circ$ by fixing the position of the tentacle. In other words, we are fixing $f|_{\partial\Delta}$ up to scaling by a constant. Polynomials using the remaining monomials $\text{int}(\Delta) \cap M$ were shown to be in correspondence with holomorphic differentials in [12, Section 2.2.4] for $\mathbb{C}P^2$ and in [3, Lemma 4.3] for smooth toric projective surfaces. The following lemma generalises this to X_Δ for any Δ :

Lemma 4.1. *Let Δ be any lattice polygon and let $C \subseteq X_\Delta$ be the vanishing set of a polynomial f with Newton polygon Δ such that $\mathbb{R}C$ is a Harnack curve. Then the space of holomorphic differentials on C is isomorphic to the space of polynomials with coefficients in $\text{int}(\Delta) \cap M$, via the map*

$$h(z_1, z_2) \mapsto \frac{h(z_1, z_2) dz_2}{\partial_{z_1} f(z_1, z_2) z_1 z_2}. \tag{10}$$

Proof. The map is injective, and both spaces have dimension g , so it remains to prove that the differentials from Eq. (10) are holomorphic. If Δ has a vertex at the origin with incident edges given by the coordinate axes, then the differentials from Eq. (10) are holomorphic over $\mathbb{C}^2 \cap C$ [3, Lemma 4.3].

Given any lattice-preserving affine map $A : M \otimes \mathbb{R} \rightarrow M \otimes \mathbb{R}$ that sends a polygon Δ' to Δ , there is a dual map $A^\vee : X_\Delta \rightarrow X_{\Delta'}$. In [3] it is shown that the pullback of A^\vee sends differentials of the form of Eq. (10) for Δ' to differentials of that form for Δ . For each vertex v of Δ , consider the lattice-preserving affine map that sends the positive orthant to the cone spanned by Δ from v . Then the differentials from Eq. (10) are holomorphic in the intersection of C with the affine chart corresponding to v . Doing that for every vertex, we get that they are holomorphic in all of C .

If Δ is not a smooth polygon, then such a lattice-preserving map does not exist. However, given any vertex v of Δ , there is a lattice-preserving map that sends the cone spanned by $(1, 0)$ and (p, q) to the cone spanned by Δ from v , for some suitable $p, q \in \mathbb{N}$. Let Δ' be the preimage of Δ under such a map. By the same arguments as in [3, Lemma 4.3], the differentials of the form of Eq. (10) for Δ' are holomorphic in $(\mathbb{C} \times \mathbb{C}^*) \cap C$. This implies that the pullback is holomorphic in $((\mathbb{C}^*)^2 \cup L) \cap C$, where L is the axis that corresponds to the edge contained in the image under A of the coordinate axis $\{x_1 = 0\}$. This can be done for every edge of Δ . Since C does not contain the intersection of any two axes, it follows that the differentials in Eq. (10) are holomorphic over all C . \square

Proposition 4.2. *The areas of the holes of the amoeba are global coordinates for the moduli space of Harnack curves with fixed Newton polygon Δ and fixed boundary points. Moreover, the moduli space of Harnack curves with fixed boundary is diffeomorphic to $\mathbb{R}_{\geq 0}^g$.*

Proof. Recall the diagram from formula (1). The first map sends a Harnack curve with fixed boundary (that is, we fix $f|_{\partial\Delta}$) to the bounded Ronkin intercepts. By Theorem 4.1, the differential of that map is the period matrix of C (see [12, Proposition 6] and [3, Theorem 3]). The second map, from the bounded intercepts to the areas of the holes in the amoeba, is also a local diffeomorphism because its differential is diagonally dominant [12, Proposition 10]. The areas of the holes of the amoeba are nonnegative, and the composition of the two maps is proper over $\mathbb{R}_{\geq 0}^g$ [12, Theorem 6]. All of these facts are proven in [12], and none of the arguments used there require that $X_\Delta = \mathbb{CP}^2$. Again, by Theorem 3.4 the composition of the maps is a diffeomorphism onto $\mathbb{R}_{\geq 0}^g$. \square

Notice that the positions of the tentacles are also well-defined numbers for nonrational Harnack curves: they are simply the evaluation of Ju_i on the points $C \cap L_i$. So by Theorem 3.3 and Theorems 3.1 and 4.2, we have that the positions of the tentacles of the amoeba modulo translation together with the areas of the holes of the amoeba are global coordinates for \mathcal{H}_Δ . Hence, we have proved the following:

Theorem 4.3. *Let Δ be a lattice polygon with m sides, g interior lattice points and n boundary lattice points. Then \mathcal{H}_Δ is diffeomorphic to $\mathbb{R}^{m-3} \times \mathbb{R}_{\geq 0}^{n+g-m}$.*

5. Compactification of \mathcal{H}_Δ

The goal of this section is to construct a natural compactification $\overline{\mathcal{H}_\Delta}$ of \mathcal{H}_Δ by collections of ‘patch-workable’ Harnack curves.

5.1. Abstract tropical curves

We begin with a review of abstract tropical curves and of $\mathcal{M}_{g,n}^{\text{trop}}$, the moduli space of tropical curves with n legs and genus g . For details of this construction, see [1].

A *weighted graph with n legs* G is a triple (V, E, L, w) , where

- (V, E) is a perhaps nonsimple connected graph – that is, we allow multiple edges and loops;
- $L : [n] \rightarrow V$ is a function which we think of as attaching n labelled *legs* at vertices of the graph;
- w is a function $V \rightarrow \mathbb{N}$ which we call the *weights* of the vertices.

The *genus* of G is the usual genus of (V, E) plus the sum of the weights on all vertices; that is,

$$\text{genus}(G) = \sum_{v \in V} w(v) - |V| + |E| + 1.$$

An *isomorphism* between two graphs $G_1 = (V_1, E_1, L_1, w_1)$ and $G_2 = (V_2, E_2, L_2, w_2)$ is a pair of bijections $\phi_V : V_1 \rightarrow V_2$ and $\phi_E : E_1 \rightarrow E_2$ such that

- for any edge $e \in E_1$ and any vertex $v \in V_1$, $\phi_E(e)$ is incident to $\phi_V(v)$ if and only if e is incident to v ;
- $L_2 = \phi_V(L_1)$;
- $w_1(v) = w_2(\phi_V(v))$.

Let G/e denote the usual contraction of G over an edge e with the following change of weights: if we contract a nonloop ab , then the contracted vertex gets weight $w(a) + w(b)$. If the contracted edge is a loop on a , then the weight of a is increased by 1. Observe that the genus is invariant under contraction.

We say that a weighted graph G is *stable* if every vertex with weight 0 has degree at least 3 and every vertex with weight 1 has positive degree. An *(abstract) tropical curve* is a pair (G, l) where G is a stable weighted graph and l is a function that assigns *lengths* to the edges of G – in other words, l is a function $l : E(G) \rightarrow \mathbb{R}_{\geq 0}^{|E(G)|}$. The *genus* of the tropical curve is the genus of G . An isomorphism between two abstract tropical curves (G_1, l_1) and (G_2, l_2) is an isomorphism ϕ of the weighted graphs G_1 and G_2 such that $l_1 = l_2 \circ \phi_E$, or such that one is the result of contracting an edge of length 0 from the other, or the transitive closure of these two relations.

Given a weighted stable graph G , one can identify the space of all tropical curves over G with $\mathbb{R}_{\geq 0}^{|E(G)|}$. We define $\mathcal{M}_{g,n}^{\text{trop}}(G) := \mathbb{R}_{\geq 0}^{|E(G)|}$ and

$$\mathcal{M}_{g,n}^{\text{trop}} := \left(\bigsqcup_{G \text{ stable}} \mathcal{M}_{g,n}^{\text{trop}}(G) \right) / \sim,$$

where \sim denotes isomorphism. This is a connected Hausdorff topological space which bijectively parametrises isomorphism classes of tropical curves. It is covered by the set of $\mathcal{M}_{g,n}^{\text{trop}}(G)$, where G runs over all 3-valent graphs with all vertices of weight 0. For such graphs we have that $\mathcal{M}_{g,n}^{\text{trop}}(G)$ is just \mathbb{R}^{3g+n-3} . However, $\mathcal{M}_{g,n}^{\text{trop}}$ is not a manifold, as there are triples of graphs of this form glued along codimension 1.

To compactify this space we allow lengths to be infinite. Let $\mathbb{R}_{\infty} = \mathbb{R}_{\geq 0} \sqcup \{\infty\}$ be the 1-point compactification of $\mathbb{R}_{\geq 0}$. An *extended* tropical curve (G, l) consists of a stable weighted graph G and a length function $l : E(G) \rightarrow \mathbb{R}_{\infty}^{|E(G)|}$. We define isomorphism classes of extended tropical curves in the same way as for tropical curves. This way we define $\overline{\mathcal{M}}_{g,n}^{\text{trop}} := \mathbb{R}_{\infty}^{|E(G)|}$ and

$$\overline{\mathcal{M}}_{g,n}^{\text{trop}} := \left(\bigsqcup_{G \text{ stable}} \overline{\mathcal{M}}_{g,n}^{\text{trop}}(G) \right) / \sim$$

This is a compact Hausdorff space with $\mathcal{M}_{g,n}^{\text{trop}}$ as an open dense subspace.

5.2. Ronkin intercepts

From now on, lightly abusing notation, we call elements of \mathcal{H}_{Δ} curves and denote one of them C , even though by definition they are equivalence classes of Harnack curves. We say that a polynomial vanishes on C if its zero locus is in the equivalence class C .

The expanded spines of different Harnack curves in the same equivalence class in \mathcal{H}_{Δ} differ only by translations. In particular, the combinatorial type and the lengths of the bounded edges remain the same. So given a curve $C \in \mathcal{H}_{\Delta}$, we have a well-defined abstract tropical curve structure for its expanded spine $\Upsilon(C) \in \overline{\mathcal{M}}_{g,n}^{\text{trop}}$: fix a labelling of the boundary segments of Δ by $[n]$ in a cyclical way and let $L(\Upsilon(C))(k)$ be the vertex incident to the ray corresponding to the segment labelled k . This induces a map

$$\Upsilon : \mathcal{H}_{\Delta} \rightarrow \overline{\mathcal{M}}_{g,n}^{\text{trop}}.$$

Recall that Ronkin intercepts are the coefficients of the tropical polynomial defining the expanded spine.

Proposition 5.1. *The Ronkin intercepts modulo translations of the graph of the Ronkin function can be taken as global coordinates for \mathcal{H}_{Δ} .*

Notice that translations of the Ronkin function are the same as translations of the amoeba.

Proof. Since we proved in Theorem 4.2 that composition of the maps in Eq. (1) is a diffeomorphism, each of the maps themselves is a diffeomorphism. This implies that the bounded Ronkin intercepts can be taken as global coordinates for Harnack curves with fixed boundary.

Now, Theorem 3.3 says that the positions of the amoeba tentacles can be taken as global coordinates for rational Harnack curves, and it is easy to recover the unbounded Ronkin intercepts from the positions of the tentacles as follows. Let ρ_i be the position of a tentacle. It corresponds to a segment in $\partial\Delta$ lying between two lattice points. Let c_i and c_{i+1} be the intercepts corresponding to those points. It is

straightforward that $\rho_i = c_i - c_{i+1}$. This implies that the map

$$\left\{ \begin{array}{l} \text{Positions of} \\ \text{amoeba} \\ \text{tentacles} \end{array} \right\} / \mathbb{R}^2 \rightarrow \left\{ \begin{array}{l} \text{Unbounded} \\ \text{Ronkin} \\ \text{intercepts} / \mathbb{R}^3 \end{array} \right\}$$

is a linear bijection when restricted to $\left\{ \sum_{i=1}^n \rho_i = 0 \right\}$. □

Proposition 5.2. *The map $\mathcal{H}_\Delta \rightarrow \mathcal{M}_{g,n}^{\text{trop}}$ is a piecewise linear topological embedding.*

Proof. By Theorem 5.1, we can take the Ronkin intercepts as global coordinates of \mathcal{H}_Δ . The Ronkin intercepts can be recovered from $\Upsilon(C)$ and Δ up to translations of the amoeba. Computing the lengths of the bounded edges of a tropical curve from the tropical polynomial is again solving a system of linear equations. So over every component $\mathcal{M}_{g,n}^{\text{trop}}((G))$ of the codomain, the map Υ is linear. □

5.3. Harnack meshes

Definition 5.3. Let $B \subseteq M$ be a finite set of affine dimension 2 (i.e., its convex hull is a 2-dimensional polygon; not all of them are in a line). We define \mathcal{H}_B to be the subset of $\mathcal{H}_{\text{conv}(B)}$ consisting of curves C such that for every $v \in \text{conv}(B)_M$, the corresponding component E_v in $\mathbb{R}^2 \setminus \mathcal{A}(C)$ exists if and only if $v \in B$.

This is well defined, since the existence of E_v depends only on the equivalence class of a Harnack curve. By Theorem 4.3, \mathcal{H}_B is diffeomorphic to $\mathbb{R}^{|B|-3}$. To see this, notice that all the $\mathbb{R}_{\geq 0}$ factors of $\mathcal{H}_{\text{conv}(B)}$ are replaced by either 0 for $v \notin B$ (as the corresponding component must vanish) or $\mathbb{R}_{>0} \cong \mathbb{R}$ for $v \in B$ (as the corresponding component cannot vanish).

Definition 5.4. Consider a regular subdivision \mathcal{S} of Δ with facets $\{B_1, \dots, B_s\}$ and let $\Delta_i = \text{conv}(B_i)$. A *Harnack mesh* over \mathcal{S} is a collection of curves (C_1, \dots, C_s) with $C_i \in \mathcal{H}_{B_i}$ such that there exists a polynomial f with $f|_{\Delta_i}$ vanishing on C_i . We write

$$\mathcal{H}_\Delta(\mathcal{S}) \subseteq \prod_{i=1}^s \mathcal{H}_{B_i}$$

for the space of all Harnack meshes over \mathcal{S} .

Notice that \mathcal{H}_Δ is equal to the disjoint union of all $\mathcal{H}_\Delta(\mathcal{S})$ where \mathcal{S} has exactly one face – that is, all \mathcal{S} of the form $\{B\}$ with $\Delta = \text{conv}(B)$.

The existence of such f is equivalent to the C_i s agreeing on their common boundary. That is, given Δ_i and Δ_j such that $\Gamma = \Delta_i \cap \Delta_j$, the distances between the tentacles of C_i corresponding to Γ are the same as the distances between the respective tentacles in C_j .

Given any Harnack mesh $(C_1, \dots, C_s) \in \mathcal{H}_\Delta(\mathcal{S})$, we can define its expanded spine as an extended tropical curve. Let Y_i be the expanded spines of C_i . For each edge $\Delta_i \cap \Delta_j$, glue the expanded spines Y_i and Y_j by removing the legs corresponding to that edge and placing instead an edge of infinite length for each primitive segment in $\Delta_i \cap \Delta_j$ between the two vertices that were incident to the corresponding leg. The remaining legs are labelled by the boundary segments of Δ . This way we have a map

$$\Upsilon_{\mathcal{S}} : \mathcal{H}_\Delta(\mathcal{S}) \rightarrow \overline{\mathcal{M}_{g,n}^{\text{trop}}}$$

which is an embedding, by Theorem 5.1.

Definition 5.5. Let Δ be a lattice polygon. The compactified moduli space of Harnack curves is

$$\overline{\mathcal{H}}_{\Delta} := \bigsqcup_{\mathcal{S}} \mathcal{H}_{\Delta}(\mathcal{S}),$$

where the union runs over all regular subdivisions \mathcal{S} of Δ_M . We give it the coarsest topology that makes continuous the map

$$\Upsilon : \overline{\mathcal{H}}_{\Delta} \rightarrow \overline{\mathcal{M}}_{g,n}^{\text{trop}}$$

defined by $\Upsilon|_{\mathcal{H}_{\Delta}(\mathcal{S})} := \Upsilon_{\mathcal{S}}$.

We will prove that this in fact is a compactification of \mathcal{H}_{Δ} – that is, a compact space where \mathcal{H}_{Δ} is dense.

Harnack meshes are essentially collections of Harnack curves that can be patchworked into another Harnack curve, except that we allow singularities and nontransversal intersection with the axes. Recall that we call a subdivision \mathcal{S} of Δ *full* if it uses all the points. We can only do patchworking whenever \mathcal{S} is full.

Proposition 5.6. *Let \mathcal{S} be a full regular subdivision of Δ with lifting function h and let $(C_1, \dots, C_s) \in \mathcal{H}_{\Delta}(\mathcal{S})$ be a Harnack mesh. Then there exist polynomials f_1, \dots, f_s vanishing on C_1, \dots, C_s such that the polynomial f_i obtained by patchworking them vanishes on a Harnack curve.*

Proof. \mathcal{S} being full implies that each C_i is smooth and nontransversal in the boundary, so that they can be patchworked together. The topological type of C_i can be obtained by patchworking using a unimodular triangulation of Δ_i and the Harnack sign pattern $v \mapsto (-1)^{v_1 v_2}$. Through taking a unimodular triangulation of Δ that refines \mathcal{S} and the Harnack sign pattern, the result of patchworking is a Harnack curve. It must have the same topological type as the patchwork using \mathcal{S} and the curves C_i when the polynomials f_i are chosen with the same sign as the polynomial that results from patchworking Δ_i with the Harnack sign configuration. \square

Notice that $\mathcal{H}_{\Delta}(\mathcal{S})$ is a linear subspace of $\prod_{B \text{ facet of } \mathcal{S}} \mathcal{H}_B \cong \mathbb{R}^c$, cut out by the linear relations $x_i = x_j$, where x_i and x_j correspond to the distance between opposite pairs of parallel tentacles and where c is a suitable integer $\left(c = \sum_{B \text{ facet of } \mathcal{S}} (|B| - 3) \right)$. So $\mathcal{H}_{\Delta}(\mathcal{S})$ is homeomorphic to an open ball. We will now show that this ball can be identified with $\text{relint}(F_{\mathcal{S}})$, where $F_{\mathcal{S}}$ is the face of $\text{Sec}(\Delta)$ corresponding to \mathcal{S} .

Proposition 5.7. *Let \mathcal{S} be a regular subdivision of Δ . Then*

$$\dim(\mathcal{H}_{\Delta}(\mathcal{S})) = \dim(F_{\mathcal{S}}).$$

Proof. Recall that in Section 2.4 we used the graph $G_{\mathcal{S}}$ dual to \mathcal{S} to compute the dimension of $\sigma(\mathcal{S})$ (see Eq. (3)). Now we will use it to compute the dimension of $\mathcal{H}_{\Delta}(\mathcal{S})$. For every face $B \in \mathcal{S}$, we have that $\mathcal{H}_B \cong \mathbb{R}^{|B|-3}$. However, for every interior edge $e \in \mathcal{S}$, there are $|e| - 2$ parameters corresponding to distances between tentacles orthogonal to e , which must agree for the two Harnack curves corresponding to the two faces containing e . Then we have

$$\dim(\mathcal{H}_{\Delta}(\mathcal{S})) = \sum_{B \in \mathcal{S} \text{ facet}} |B| - 3 - \sum_{e \in \mathcal{S} \text{ interior edge}} |e| - 2.$$

Now we consider how many times a point in Δ_M is taken into account in this equation. If a point is contained in a single facet of \mathcal{S} , then it cannot be contained in any interior edge. If a point is in the boundary of Δ , then it is contained in one more facet of \mathcal{S} than interior edges. The same is true if the point is in the interior of an edge in the interior of Δ . However, if the point is a vertex of \mathcal{S} in the interior

of Δ , it is contained in the same number of facets as interior edges, and the same is true for points not contained in any facet of \mathcal{S} . So we obtain

$$\begin{aligned} \dim(\mathcal{H}_\Delta(\mathcal{S})) &= |\Delta_M| - |\{\text{interior points of } \mathcal{S}\}| - \left| \Delta_M \setminus \bigcup \mathcal{S} \right| \\ &\quad - 3|\{\text{facets of } \mathcal{S}\}| + 2|\{\text{interior edges of } \mathcal{S}\}| \\ &= \left| \bigcup \mathcal{S} \right| - |F(G_{\mathcal{S}})| + 2|E(G_{\mathcal{S}})| - 3|F(G_{\mathcal{S}})|. \end{aligned} \tag{11}$$

If we add Eqs. (3) and (11) and use Euler’s formula for $G_{\mathcal{S}}$, we get

$$\begin{aligned} \dim(\mathcal{H}_\Delta(\mathcal{S})) + \dim(\sigma(\mathcal{S})) &= |\Delta_M| + 3(1 - |V(G_{\mathcal{S}})| + |E(G_{\mathcal{S}})| - |F(G_{\mathcal{S}})|) \\ &= |\Delta_M|. \end{aligned}$$

Hence $\dim(\mathcal{H}_\Delta(\mathcal{S})) = \dim(F_{\mathcal{S}}) = |\Delta_M| - \dim(\sigma(\mathcal{S}))$. □

Proposition 5.8. *Let Δ be a lattice polygon and \mathcal{S} be a regular subdivision of Δ_M . Then $Y_{\mathcal{S}}(\mathcal{H}_\Delta(\mathcal{S})) \subseteq \overline{Y(\mathcal{H}_\Delta)}$, where $\overline{Y(\mathcal{H}_\Delta)}$ is the closure of $Y(\mathcal{H}_\Delta)$ in $\mathcal{M}_{g,n}^{\text{trop}}$.*

Proof. Suppose \mathcal{S} is a full subdivision of Δ . Choose a height function $h \in \sigma(\mathcal{S})$. By Theorem 5.6, for any Harnack mesh $\mathcal{C} = (C_1, \dots, C_s)$ in $\mathcal{H}_\Delta(\mathcal{S})$ there exist polynomials f_1, \dots, f_s and $t_0 > 0$ such that for any $0 < t < t_0$, the curve C_t obtained by patchworking is in \mathcal{H}_Δ . So we have a path $(0, t_0) \rightarrow \mathcal{H}_\Delta$. For each facet B_i of \mathcal{S} there is a family of polynomials $\{f_t^i \mid t \in (0, t_0)\}$ with real coefficients such that f_t^i vanishes on C_t and every coefficient of f_t^i outside $\text{conv}(B_i)$ goes to 0 as t goes to 0. This follows from picking the height function h^i affinely equivalent to h such that points in B_i have height 0 and doing patchworking with h^i . The limit $\lim_{t \rightarrow \infty} f_t^i = f_i$ vanishes on C_i .

As t goes to 0, the lengths of the edges of $Y(C_t)$ that corresponds to the interior of B_i tend to the lengths of the edges of $Y(C_i)$. Doing this for every i , we have that all the finite lengths of $Y_{\mathcal{S}}(\mathcal{C})$ agree with the lengths of $\lim_{t \rightarrow 0} Y(C_t)$. The edges going to infinity are precisely those dual to primitive segments of \mathcal{S} . Then $Y(C_t)$ forms a path $(0, t_0) \rightarrow Y(\mathcal{H}_\Delta)$ such that the limit of this path when t goes to 0 is $Y(\mathcal{C})$. So $\mathcal{C} \in \overline{Y(\mathcal{H}_\Delta)}$.

If \mathcal{S} is not full, let \mathcal{S}' be the subdivision whose facets are $B' = \text{conv}(B)_M$ for each facet $B \in \mathcal{S}$. That is, \mathcal{S}' is the finest full subdivision that coarsens \mathcal{S} . It is regular, as we can take a height function in $\sigma(\mathcal{S})$ and linearly extrapolate in each Δ_i to make it full. As the expanded spine is continuous, even when ovals contract, we have

$$Y_{\mathcal{S}}(\mathcal{H}_\Delta(\mathcal{S})) \subseteq \overline{Y_{\mathcal{S}'}(\mathcal{H}_\Delta(\mathcal{S}'))} \subseteq \overline{Y(\mathcal{H}_\Delta)}.$$

□

Lemma 5.9. *Let Δ be a lattice polygon. Then $\overline{Y(\mathcal{H}_\Delta)} = Y(\overline{\mathcal{H}_\Delta})$.*

Proof. Proposition 5.8 implies that $Y(\overline{\mathcal{H}_\Delta}) \subset \overline{Y(\mathcal{H}_\Delta)}$.

For the other containment, let C_1, C_2, \dots be a sequence of curves in \mathcal{H}_Δ such that their expanded spines converge to a point in $G \in \mathcal{M}_{g,n}^{\text{trop}}$. The graph G is connected but with some edges of infinite length, so we call *connected components* of G the connected components after deleting all infinite edges. Notice that vertices of an expanded spine correspond to polygons inside Δ given by the regular subdivision dual to the expanded spine. This association is carried on to the limit, so the connected components G_1, \dots, G_s induce a regular subdivision $\mathcal{S} = \{B_1, \dots, B_s\}$ of Δ . For a connected component G_i of G , we can choose polynomials f_1^i, f_2^i, \dots vanishing on C_1, C_2, \dots such that they converge to a polynomial f^i which vanishes on a curve whose expanded spine is G_i . This can be done, for example, by picking a vertex of G_i and fixing it to be in the origin – that is, translating the amoebas of C_1, C_2, \dots so that the




Figure 2. $\overline{\mathcal{H}_\Delta}$ for $\Delta = \text{conv}((1, 0), (0, 1), (-1, 0), (0, -1))$.

corresponding vertex in the expanded spine is always at the origin. Since the limit of Harnack curves is a Harnack curve [18, Remark 2], f^i vanishes on a Harnack curve $C^i \in \mathcal{H}_{B_i}$. By construction, the tuple $\mathcal{C} = (C^1, \dots, C^s)$ is a Harnack mesh over \mathcal{S} , and we have $Y(\mathcal{C}) = G$. \square

Corollary 5.10. *Let Δ be a lattice polygon. Then*

$$\overline{\mathcal{H}_\Delta(\mathcal{S})} = \bigsqcup_{\mathcal{T} \leq \mathcal{S}} \mathcal{H}_\Delta(\mathcal{T}),$$

where the union runs over all subdivisions \mathcal{T} of Δ that refine \mathcal{S} .

Theorem 5.11. *Let Δ be a lattice polygon. The stratification of $\overline{\mathcal{H}_\Delta}$ by $\mathcal{H}_\Delta(\mathcal{S})$ is a cell complex with a poset isomorphic to the face poset of the secondary polytope $\text{Sec}(\Delta_M)$ given by its faces.*

Proof. The faces of $\text{Sec}(\Delta_M)$ are in correspondance with regular subdivisions. By Theorem 5.7, $\mathcal{H}_\Delta(\mathcal{S})$ has the same dimension as the face of $\text{Sec}(\Delta_M)$ corresponding to \mathcal{S} . By Theorem 5.10, the boundary of $\overline{\mathcal{H}_\Delta(\mathcal{S})}$ consists of $\mathcal{H}_\Delta(\mathcal{T})$ for every subdivision \mathcal{T} that refines \mathcal{S} . Similarly, the faces contained in the face of $\text{Sec}(\Delta_M)$ corresponding to \mathcal{S} are those corresponding to refinements of \mathcal{S} . \square

Example 5.12. Set $\Delta := \text{conv}((1, 0), (0, 1), (-1, 0), (0, -1))$. We have that $\text{Sec}(\Delta)$ is a triangle. Figure 2 shows the space $\overline{\mathcal{H}_\Delta}$ together with the subdivisions of the corresponding face in $\text{Sec}(\Delta)$ and the amoebas of the corresponding Harnack meshes. The horizontal coordinate represents the relative position of the tentacles. This is parametrised, for example, by $\rho_1 + \rho_3$. Going to the left stretches the amoeba vertically, whereas going to the right stretches it horizontally. The vertical coordinate corresponds to the area of the oval, where going downward decreases the area and going upward increases it. The bottom open segment corresponds to $\mathcal{H}_{0,\Delta}$, and that segment together with the interior face corresponds to $\mathcal{H}_\Delta \cong \mathbb{R} \times \mathbb{R}_{\geq 0}$.

Remark. We warn the reader that the regular subdivision corresponding to the open cell where Harnack mesh \mathcal{C} lies is not the same as the subdivision dual to $\Upsilon(\mathcal{C})$. In fact, the two subdivisions are not always comparable with the refinement order. For example, the subdivision corresponding to the bottom edge of Figure 2 is the square without the middle point, but the subdivision dual to the graph of a spine in that cell would be the subdivision with four triangles.

6. Questions and future directions

6.1. $\overline{\mathcal{H}_\Delta}$ as a CW-complex

We begin by suggesting the following strengthening of Theorem 5.11:

Conjecture 6.1. *The cell decomposition of the compactified moduli space $\overline{\mathcal{H}_\Delta} = \bigsqcup_{\mathcal{S}} \mathcal{H}_\Delta(\mathcal{S})$ is a regular CW-complex.*

Two regular CW-complexes with isomorphic cell posets are isomorphic – that is, there is a homeomorphism that maps each cell to the corresponding cell given by the poset isomorphism. By Theorem 5.11, if Theorem 6.1 holds, then $\overline{\mathcal{H}_\Delta}$ would be isomorphic to $\text{Sec}(\Delta_M)$. To show that $\overline{\mathcal{H}_\Delta}$ is a regular CW-complex, it is enough to show that for any regular subdivision \mathcal{S} of Δ , $\Upsilon_{\mathcal{S}}(\mathcal{H}_\Delta(\mathcal{S}))$ is a closed ball. Moreover, it is enough to prove the following:

Conjecture 6.2. *Let $\mathcal{C} \in \mathcal{H}_\Delta(\mathcal{S})$ be a Harnack mesh and \mathcal{S}' be a coarsening of \mathcal{S} . Then there is a neighbourhood of \mathcal{C} in $\overline{\mathcal{H}_\Delta(\mathcal{S}'})$ homeomorphic to a half space of dimension $\dim(\mathcal{H}_\Delta(\mathcal{S}'))$.*

Since the poset of $\overline{\mathcal{H}_\Delta}$ is Eulerian by Theorem 5.11, Theorem 6.2 implies that the closure of the cells of $\overline{\mathcal{H}_\Delta}$ are closed balls by (a reformulation of) Poincaré’s conjecture. This argument was recently used to prove that the positroid stratification of the totally nonnegative Grassmannian is a CW-complex [6]. It is worth remarking that Harnack curves enjoy several similarities with the total positivity phenomenon (see, e.g., [13, Section 5.2] or our proof of Theorem 3.8). In the next subsection we will see that Theorem 6.2 holds when \mathcal{S} is full.

6.2. $\overline{\mathcal{H}_\Delta}$ as a manifold with generalised corners

The foregoing discussion suggests a study of topological charts in \mathcal{H}_Δ . We can be more ambitious and try to endow \mathcal{H}_Δ with a smooth structure. Theorem 4.3 is already a description of \mathcal{H}_Δ as a smooth manifold with corners. A natural question is whether we can extend this smooth structure to $\overline{\mathcal{H}_\Delta}$. A desirable trait of such a smooth structure (besides being compatible with the chart given by Theorem 4.3) is that the cell complex structure from Theorem 5.11 be recoverable from it. However, secondary polytopes are not always simple polytopes, and manifolds with corners lack the capacity to describe nonsimple vertices. To mend this, we turn our attention to a wider category, namely that of manifolds with *generalised corners*, or *gc-manifolds*, as defined in [11].

Definition 6.3 ([11]). A *g-chart* of a topological space X is a triple (ϕ, \mathcal{L}, U) such that

- \mathcal{L} is a weakly toric monoid – that is, a semilattice of the form $\mathcal{L} = \mathbb{Z}^s \cap \sigma$, where s is a positive integer and $\sigma \subseteq \mathbb{R}^s$ is a rational polyhedral cone;
- U is an open subset of $\text{Hom}(\mathcal{L}, \mathbb{R}_{\geq 0})$ – that is, the space of monoid morphisms from \mathcal{L} to the monoid $(\mathbb{R}_{\geq 0}, \cdot)$ with the weakest topology that makes evaluation on a point $q \in \mathcal{L}$ continuous;
- $\phi : U \rightarrow X$ is a topological embedding to an open subset $\phi(U) \subseteq X$.

We call X a *gc-manifold* if it has a *g-atlas* – that is, a collection of *g-charts* covering X and satisfying certain compatibility conditions on the transition functions. These conditions depend on the monoids, but we refrain from explaining them in detail in this paper. Easy examples of *gc-manifolds* are $\text{Hom}(\mathbb{N}, \mathbb{R}_{\geq 0}) \cong \mathbb{R}_{\geq 0}$ and $\text{Hom}(\mathbb{Z}, \mathbb{R}_{\geq 0}) \cong \mathbb{R}$. So as a *gc-manifold*, $\mathcal{H}_\Delta \cong \text{Hom}(\mathbb{N}^{m-3} \times \mathbb{Z}^{n+g-m}, \mathbb{R}_{\geq 0})$ by Theorem 4.3.

Proposition 6.4. *Let $\mathcal{C} \in \mathcal{H}_\Delta(\mathcal{S})$ be a Harnack mesh where \mathcal{S} is a full subdivision. Then there exists a g -chart around \mathcal{C} . Moreover, all charts provided this way are compatible with each other.*

Proof. Since $(\mathbb{R} \cup \{\infty\}, +)$ is isomorphic as a monoid to $(\mathbb{R}_{\geq 0}, \cdot)$ with $x \mapsto e^{-x}$ as isomorphism, g -charts can be equivalently defined to be homeomorphisms from open subsets of ‘affine tropical toric varieties’ – that is, from open subsets $U \subseteq \text{Hom}(\mathcal{L}, \mathbb{R} \cup \{\infty\})$. Consider a graph G embedded in \mathbb{R}^2 that is dual to the subdivision \mathcal{S} . In particular, the edges of G have a prescribed slope. The lengths of these edges satisfy linear equations with integer coefficients given by the circuits of G (two for each circuit). These equations are binomial relations under the isomorphism $(\mathbb{R} \cup \{\infty\}, +) \cong (\mathbb{R}_{\geq 0}, \cdot)$. Thus, the edges of G (which correspond to edges of \mathcal{S} in the interior of Δ) generate a toric monoid $\mathcal{L}_\mathcal{S}$ under these relations.

Let \mathcal{C}' be a Harnack mesh close enough to \mathcal{C} . The spine $Y(\mathcal{C}')$ has a subgraph G_i which is very close to $Y(C_i)$ for each curve $C_i \in \mathcal{C}$. These subgraphs are glued together with edges of very large (possibly infinite) length. Contracting these subgraphs results in the graph dual to \mathcal{S} , so the distances between these graphs induce a homomorphism $\phi_{\mathcal{C}'} : \mathcal{L}_\mathcal{S} \rightarrow \mathbb{R} \cup \{\infty\}$.

The coordinates of a Harnack mesh in $\mathcal{H}_\Delta(\mathcal{S})$ encode the same information as the spines of each curve in the Harnack mesh. Since \mathcal{C}' is close enough to \mathcal{C} , there exists a mesh $\mathcal{H}_\Delta(\mathcal{S})$ that has a curve whose spine is isomorphic as metric graphs to G_i for each i (here we use the fact that \mathcal{S} is full). The coordinates of this mesh in $\mathcal{H}_\Delta(\mathcal{S})$ are a vector in \mathbb{R}^d , where $d = n + g - \dim(\sigma(\mathcal{S}))$, by Theorem 5.7. This vector induces a homomorphism $\psi_{\mathcal{C}'} : \mathbb{Z}^d \rightarrow \mathbb{R}_{\geq 0}$.

The Harnack mesh \mathcal{C}' is completely determined by $\phi_{\mathcal{C}'}$ and $\psi_{\mathcal{C}'}$, so we obtain an embedding from a neighbourhood of \mathcal{C} in $\overline{\mathcal{H}_\Delta}$ to $\text{Hom}(\mathcal{L}_\mathcal{S} \times \mathbb{Z}^d, \mathbb{R}_{\geq 0})$ given by

$$\mathcal{C}' \mapsto \left((x, y) \mapsto e^{-\phi_{\mathcal{C}'}(x)} \psi_{\mathcal{C}'}(y) \right),$$

where $x \in \mathcal{L}_\mathcal{S}$ and $y \in \mathbb{Z}^d$. Since $\text{Hom}(\mathcal{L}_\mathcal{S} \times \mathbb{Z}^d, \mathbb{R}_{\geq 0})$ is of the same dimension as $\overline{\mathcal{H}_\Delta}$, this mapping forms a g -chart.

That this g -chart is compatible with the g -chart of \mathcal{H}_Δ given by Theorem 4.3 is a consequence of Theorem 5.1. Similarly, charts constructed this way are compatible with each other. □

Corollary 6.5. *Theorem 6.2 holds when \mathcal{S} is full.*

Proof. For any \mathcal{L} , the space $\text{Hom}(\mathcal{L}, \mathbb{R}_{\geq 0})$ is stratified by its support. All of the strata are again of the form $\text{Hom}(\mathcal{L}', \mathbb{R}_{\geq 0})$ for some submonoid $\mathcal{L}' \leq \mathcal{L}$ and are topological manifolds with boundary. The g -charts constructed before respect the cell strata of $\text{Hom}(\mathcal{L}_\mathcal{S} \times \mathbb{Z}^d, \mathbb{R}_{\geq 0})$ and $\overline{\mathcal{H}_\Delta}$, so the result follows. □

Unfortunately, we do not know of a good way of constructing g -charts for points in cells corresponding to nonfull subdivisions. Length of edges is not a good parameter for the chart, since the edges of cycles in Γ that correspond to ovals contracting to a point are finite (that is, positive after applying $x \mapsto e^{-x}$), so the preimage in such a chart would not be open. One could expect a g -chart covering $\mathcal{H}_\Delta(\mathcal{S})$ for a nonfull subdivision \mathcal{S} to be defined over an open subset of $\text{Hom}(\mathcal{L}_\mathcal{S} \times \mathbb{Z}^d \times \mathbb{N}^k, \mathbb{R}_{\geq 0})$, where k is the number of missing points of \mathcal{S} . However, it is not clear what the coordinates corresponding to the copies of \mathbb{N} should be.

For example, using the area of ovals as coordinates, as we did in Theorem 4.3, does not work either. Consider the bottom right corner of Figure 2. We can take a continuous path along the interior of the triangle by stretching the amoeba horizontally but maintaining the area of the bounded component constant. Since the square bounded by the expanded spine is contained in the union of the amoeba with the bounded component of the complement, its area is bounded. Stretching the amoeba horizontally causes the length of the vertical edges of the square to tend to 0, which means that the path ends in the bottom right corner. This coordinate should be 0 at this point, so the continuity is broken.

Question 6.6. *Is there a natural way of completing a g -atlas on $\overline{\mathcal{H}_\Delta}$ with g -charts respecting the cell strata?*

A positive answer to this question implies a positive answer to the conjectures in Section 6.1.

6.3. A cell complex for T -curves

Harnack meshes can also be patchworked into non-Harnack curves by choosing polynomials with different sign patterns. The resulting curves are called T -curves. They can be thought of as the ‘neighbourhood’ of \mathcal{H}_Δ , which suggests the following question:

Question 6.7. *Given a lattice polygon Δ , are there other topological types of curves in X_Δ such that their moduli space can be given a cell complex structure similar to \mathcal{H}_Δ ? Can such moduli spaces be glued together to form a larger cell complex, or even a polytopal complex, where cells correspond to different topological types?*

Example 6.8. When Δ is the unit square, $\overline{\mathcal{H}_\Delta}$ is a segment. When the Harnack meshes of the extremes are patchworked in a non-Harnack way, we get a curve whose amoeba has a pinching [16, Example 1]. From one of the extremes, the resulting expanded spine has a bounded edge parallel to $\{x_1 = x_2\}$, and from the other extreme the edge is parallel to $\{x_1 = -x_2\}$. When the length of the bounded edge goes to 0, both cases degenerate to a reducible curve (the union of two axis-parallel lines). In this case the complex of Question 6.7 exists and is isomorphic to the boundary of a triangle.

Acknowledgements. The author was supported by the Einstein Foundation Berlin through Francisco Santos’ Visiting Professor Fellowship at Freie Universität Berlin and by the German Research Foundation and the Open Access Publication Fund of Technische Universität Berlin.

The author would like to thank Mauricio Velasco for proposing this problem and advising the master’s thesis that preceded this paper, as well as the rest of the committee members: Tristram Bogart, Felipe Rincón and Florent Schaffhauser; the organisers of the Master Class in Tropical Geometry in Stockholm 2017, where much of the inspiration for section 5 was obtained; and Grigory Mikhalkin, Timo de Wolff, Lionel Lang and the anonymous referees for helpful discussions and suggestions. Special thanks to Francisco Santos for all of his help and support in the development of this paper.

Conflict of Interest: None.

References

- [1] L. Caporaso, ‘Algebraic and tropical curves: Comparing their moduli spaces’, in *Handbook of Moduli. Vol. I*, Advanced Lectures in Mathematics (ALM) vol. 24 (International Press, Somerville, MA, 2013), 119–160.
- [2] D. A. Cox, J. B. Little and H. K. Schenck, *Toric Varieties*, Graduate Studies in Mathematics vol. 124 (American Mathematical Society, Providence, RI, 2011).
- [3] R. Crétois and L. Lang, ‘The vanishing cycles of curves in toric surfaces I’, *Compos. Math.* **154**(8) (2018), 1659–1697.
- [4] J. A. De Loera, J. Rambau and F. Santos, *Triangulations*, Algorithms and Computation in Mathematics vol. 25 (Springer-Verlag, Berlin, 2010).
- [5] M. Forsberg, M. Passare and A. Tsikh, ‘Laurent determinants and arrangements of hyperplane amoebas’, *Adv. Math.* **151**(1) (2000), 45–70.
- [6] P. Galashin, S. N. Karp and T. Lam, ‘Regularity theorem for totally nonnegative flag varieties’, Preprint, 2019, [arXiv:1904.00527](https://arxiv.org/abs/1904.00527).
- [7] I. M. Gelfand, M. M. Kapranov and A. V. Zelevinsky, *Discriminants, Resultants, and Multidimensional Determinants* Mathematics: Theory & Applications (Birkhäuser Boston, Inc., Boston, MA, 1994).
- [8] W. Gordon, ‘On the diffeomorphisms of euclidean space’, *Amer. Math. Monthly* **79**(7) (1972), 755–759.
- [9] A. Harnack, ‘Ueber die Vieltheiligkeit der ebenen algebraischen Curven’, *Math. Ann.* **10**(2) (1876), 189–198.
- [10] I. Itenberg and O. Viro, ‘Patchworking algebraic curves disproves the Ragsdale conjecture’, *Math. Intelligencer* **18**(4) (1996), 19–28.
- [11] D. Joyce, ‘A generalization of manifolds with corners’, *Adv. Math.* **299** (2016), 760–862.
- [12] R. Kenyon and A. Okounkov, ‘Planar dimers and Harnack curves’, *Duke Math. J.* **131**(3) (2006), 499–524.
- [13] R. Kenyon, A. Okounkov and S. Sheffield, ‘Dimers and amoebae’, *Ann. of Math. (2)* **163**(3) (2006), 1019–1056.
- [14] A. G. Khovanskii, ‘Newton polyhedra and the genus of complete intersections’, *Funct. Anal. Appl.* **12**(1) (1978), 38–46.
- [15] L. Lang, ‘A generalization of simple Harnack curves’, Preprint, 2015, [arXiv:1504.07256](https://arxiv.org/abs/1504.07256).

- [16] G. Mikhalkin, 'Real algebraic curves, the moment map and amoebas', *Ann. of Math. (2)* **151**(1) (2000), 309–326.
- [17] G. Mikhalkin and A. Okounkov, 'Geometry of planar log-fronts', *Mosc. Math. J.* **7**(3) (2007), 507–531, 575.
- [18] G. Mikhalkin and H. Rullgård, 'Amoebas of maximal area', *Int. Math. Res. Not. IMRN* **2001**(9) (2001), 441–451.
- [19] M. Passare and J.-J. Risler, 'On the curvature of the real amoeba', in *Proceedings of the Gökova Geometry-Topology Conference 2010* (International Press, Somerville, MA, 2011), 129–134.
- [20] M. Passare and H. Rullgård, 'Amoebas, Monge-Ampère measures, and triangulations of the Newton polytope', *Duke Math. J.* **121**(3) (2004), 481–507.
- [21] L. I. Ronkin, *Introduction to the Theory of Entire Functions of Several Variables*, Translations of Mathematical Monographs vol. **44** (American Mathematical Society, Providence, RI, 1974). Translated from the Russian by Israel Program for Scientific Translations.
- [22] O. Viro, 'Patchworking real algebraic varieties', Preprint, 2006, [arxiv:0611.382](https://arxiv.org/abs/0611.382).