UNIVERSITY OF CALIFORNIA, SAN DIEGO

Triangulating Teichmüller space using the Ricci flow

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Mathematics

by

Graham P. Hazel

Committee in charge:

Professor Justin Roberts, Chair
Professor Bennett Chow
Professor Benjamin Grinstein
Professor Kenneth Intriligator
Professor Hans Wenzl

2004

The dissertation of Graham P. Hazel is approved, and it is acceptable in quality and form for publication on microfilm:

Chair

University of California, San Diego

2004

To my mother Judith, who never had a chance to finish her academic career, and my father Philip, who was the first to tell me to RTFM.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

First of all, this thesis would not exist without the unqualified love and support of my family: Mum, Dad, Andrew and Colin. Nor would it exist without the mathematical guidance and patience of my supervisor Justin Roberts.

Secondly, in no particular order: Rich, Rich, Richie, Helen, Kate, Vince, Dan, Nick, Jess, Liz, Tom, Des, The Glick, Phonebooth, Tim, Rino, Tai, Ben, Larissa, Ryan, D.C., Brian, Matt, Brett, Chopper, Suzanne, Eugene, Sriram, Cameron, the SND crowd, the semi-regular frisbee game, and numerous other students, card-players and fellow-travellers I've forgotten... thank you for your various contributions to my well-being, whether through mathematical solidarity, good company, fine dining, sharing a game, a song or a pint or two, making sure each new year is stimulating, or just generally keeping me sane.

Finally, for providing a pleasant numbing sensation between bouts of mathematics (and sometimes during), I owe a debt of gratitude to, among others, The Castle, the Afton Hotel bar (RIP), the Caledonian Brewery, London's West End, and a number of distilleries from Islay to Jalisco.

VITA

1997 B.A., University of Cambridge

1998 Certificate of Advanced Study in Mathematics, University of Cambridge

1998–2000 Ph.D. student, University of Edinburgh

2000–2004 Teaching Assistant, Department of Mathematics, University of California, San Diego

2004 Ph.D., University of California, San Diego

ABSTRACT OF THE DISSERTATION

Triangulating Teichmüller space using the Ricci flow

by

Graham P. Hazel
Doctor of Philosophy in Mathematics

University of California San Diego, 2004

Professor Justin Roberts, Chair

The behaviour of the Ricci flow on hyperbolic surfaces is investigated via a combinatorial analogue first studied by Chow and Luo. The natural cell decompositions of Teichmüller and moduli spaces of hyperbolic metrics on decorated surfaces give parameters for and are shown to be compatible with the combinatorial flow. Hence elegant new proofs of the cell decompositions are obtained, as well as a practical algorithm for reconstructing a metric on a hyperbolic surface from a point in the corresponding cell complex.

Chapter 1

Introduction

The Ricci flow, first introduced by Hamilton [4], is a natural evolution equation on the space of metrics on a manifold, under which a metric g evolves by

$$\frac{\mathrm{d}}{\mathrm{d}t}g(t) = -2\mathrm{Ric}_{g(t)}$$

where $\mathrm{Ric}_{g(t)}$ is the Ricci curvature. Roughly speaking, the action of this flow is to "even the distribution of curvature" in a manifold—the fixed points of the (volume-normalised) flow are the geometries of constant Ricci curvature.

Remarkable new results of Perelman [8] [9] [10] show that, for a closed, oriented 3-manifold M equipped with an initial smooth metric g, the singularities (that is, degenerations of the metric) which appear as the flow evolves are directly related to the topological sphere and torus decompositions of M. The singularities can be resolved by performing surgery on the manifold according to these decompositions, with the flow then continuing on the newly separated pieces. It turns out that any piece which has no further singularities for all time must have a geometric structure—a complete, locally homogeneous Riemannian metric—and so must be a quotient of one of the eight standard 3-dimensional geometries: $\mathbb{H}^3$, $\mathbb{R}^3$, $\mathbb{S}^3$, $\mathbb{H}^2 \times \mathbb{R}$,

$\widetilde{SL(2,\mathbb{R})}$, *Nil*, *Sol* and $\mathbb{S}^2 \times \mathbb{R}$. Thus—assuming Perelman's results are verifed—the flow not only proves Thurston's celebrated geometrisation conjecture [11], that each component of the sphere and torus decomposition of M admits a geometric structure (which, as an aside, implies the Poincaré conjecture), but it identifies the decomposition of M and its unique geometric structure, the "best possible" metric, for you!

In two dimensions, of course, the story is much different. The classical uniformisation theorem holds that every closed surface admits a metric of constant Gauss curvature +1, 0 or -1, and thus is a quotient of either $\mathbb{R}^2$, $\mathbb{S}^2$ or $\mathbb{H}^2$ by a discrete group of isometries. This fact is reflected by the behaviour of the two-dimensional Ricci flow introduced by Hamilton in [5] and defined by

$$\frac{\mathrm{d}}{\mathrm{d}t} g(t) = -2Kg(t)$$

where K is the Gaussian curvature of the surface (equal to the Ricci curvature in two dimensions). Hamilton [5] and Chow [2] show that for any closed surface and any initial Riemannian metric, the Ricci flow exists for all time (that is, never degenerates), and—after normalising for constant area—the flow converges to a metric with constant curvature as time goes to infinity. Thus the uniformisation theorem says that in two dimensions, the topology is never interesting enough to require a decomposition of the surface into pieces of different geometric types; the Ricci flow, which never degenerates, agrees.

Sadly, though, this isn't the whole story, for it turns out that the topology of a surface isn't even interesting enough to nail down the "best possible" constant curvature metric uniquely. For example, a closed surface F of genus $g \geq 2$ has a $(6g-6)$-dimensional space of such metrics (in a suitable natural sense). So although the Ricci flow is guaranteed to converge to a metric of constant curvature, which one it heads towards depends on the choice of initial metric.

One way to try to understand this is to concentrate not on the Ricci flow itself, but on a combinatorial analogue. In [3], Chow and Luo study such an analogous flow on triangulations of closed surfaces, showing that, with suitable initial conditions, such a flow will converge to a circle packing metric.

However, when our surfaces admit hyperbolic structures and have some kind of "decoration" (for instance, geodesic boundary or marked points), help is at hand. There is a natural cell decomposition of the space of hyperbolic structures on such a surface, with coordinates coming from the intrinsic geometry of the metrics. This was first shown for the case of surfaces with at least one cusp, by Bowditch and Epstein [1]. (The cell complex itself was first introduced by Harer [6], but his coordinates come from quadratic differentials, not from geometry.)

In the thesis we show that this description of the space of hyperbolic structures is compatible with the Ricci flow, in the following sense. Each cell in the decomposition corresponds to a different combinatorial blueprint (analogous to a triangulation), while the simplicial coordinates within the cell give us parameters for the metric. We show that—when subjected to these constraints—the combinatorial Ricci flow of Chow and Luo converges to a unique hyperbolic metric, for any initial choice of the unconstrained variables.

Hence we obtain elegant new proofs of the cell decompositions of the spaces of hyperbolic metrics on decorated surfaces. Furthermore, our work suggests a practical algorithm (given a computer and enough time) for finding the "best possible" constant curvature metric for a surface, given its coordinates in the cell decomposition.

In chapter 2 we make definitions, fix notation and introduce concepts relevant to the whole account. Chapters 3–5 concentrate individually on the three separate cases (types of surface) we tackle with the combinatorial Ricci flow. Chapter 6 then deals with a further case where we identify a cell decomposition by building on our previous work. Finally, chapter 7 contains some comments about our failure to address the case of closed surfaces.

A good reference for the basic hyperbolic plane geometry used in the calculational lemmas (containing, for example, derivations of the hyperbolic sine and cosine rules for triangles) is Thurston's book [12].

Chapter 2

Some definitions and notation

2.1 Teichmüller space

Let F be a fixed topological surface which admits a hypberbolic structure under which the boundary of F (if non-empty) is geodesic. Let $\mathcal{S}F$ be the space of such structures.

Definition 2.1. The *Teichmüller space* $\mathcal{T}F$ of F is the quotient of $\mathcal{S}F$ by $\mathrm{Diff}_0 F$, the group of diffeomorphisms of F homotopic to the identity via a homotopy taking the boundary to itself at all times.

Definition 2.2. The *moduli space* $\mathcal{M}F$ of F is the quotient of $\mathcal{S}F$ by the whole diffeomorphism group $\mathrm{Diff}\, F$.

Definition 2.3. The discrete group $\mathcal{MCG}(F) = \frac{\mathrm{Diff}\, F}{\mathrm{Diff}_0 F}$ is called the *mapping class group* of F and is (in fact, finitely) generated by Dehn twists. A *Dehn twist* about an oriented simple closed curve $c \subset F$ is a diffeomorphism formed by cutting F along c, performing a full twist in an annular neighbourhood of one copy of c, then regluing.

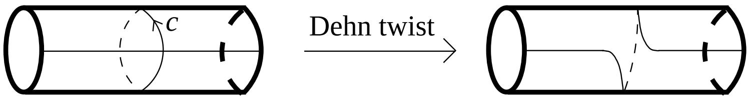

Figure 2.1: Dehn twist about c, illustrated by its action on a perpendicular curve

If F is a closed, orientable surface of genus $g \geq 2$, the standard method of finding coordinates for $\mathcal{T}F$ is as follows. (For a more complete discussion, see for instance Thurston's book [12].)

Pick a set of $n = 3g - 3$ disjoint, pairwise non-homotopic, essential simple closed curves $\Gamma = \{\gamma_1, ..., \gamma_n\}$ which cut F into $2g - 2$ pairs of pants. (A "pair of pants" is a sphere with three open discs removed.)

The Teichmüller space of hyperbolic structures on a pair of pants P is easy to compute. Given any hyperbolic metric on P, there are unique geodesic "seams" which cut the pants into two right-angled hexagons which are congruent under reflection. Since a hyperbolic right-angled hexagon is determined by the lengths of three alternating sides, the point in Teichmüller space corresponding to a hyperbolic metric on P is determined by the ordered triple of boundary lengths under the metric.

Note that twisting one of the legs of the pants does not alter the point in $\mathcal{T}P$, since it is a diffeomorphism homotopic to the identity, with boundary mapped to itself.

However, when we glue two hyperbolic pairs of pants P_1 and P_2 together along a boundary curve c, the twisting does matter—the only way to undo a twist in $P_1 \cup_c P_2$ is to cut along c, untwist, and then reglue (that is, perform a Dehn twist about c).

Consequently, when we glue the $2g - 2$ hyperbolic pairs of pants to recover a point in the Teichmüller space of F, we begin with $6g - 6$ parameters (the lengths of the boundary curves for each pair of pants), lose $3g-3$ of them (since the pairs of curves we identify must have equal lengths), but gain $3g-3$ new parameters to record the twisting information. These boundary lengths and twists are the *Fenchel-Nielsen coordinates* for Teichmüller space.

Although this approach makes it easy to see that $\mathcal{T}F$ is homeomorphic to the open ball $\mathbb{R}^{6g-6}$, it has some undesirable features.

First, the coordinates do not arise from intrinsic features of the metric—we have to choose the set of curves Γ (and in fact some extra data in order to accurately record the twisting), and the coordinates which result depend on these choices.

Secondly, it is not equivariant with respect to the mapping class group, so we do not obtain a description of $\mathcal{M}F$. The only Dehn twists whose action can easily be written down in Fenchel-Nielsen coordinates are those about one of the γ_i; an arbitrary Dehn twist, transverse to some of the curves in Γ, will have an action which is very difficult to describe.

2.2 What we do instead

In each of the cases we study, our goal is to describe the Teichmüller space of a certain type of hyperbolic surface via a cell decomposition. We start with an intrinsic method for constructing a distinguished set (a "spine") on the surface. Each different possible spine corresponds to a cell in our decomposition; to complete the description, we define coordinates measured relative to the spine within each cell.

2.2.1 Constructing the spine

Let F be a surface equipped with a hyperbolic metric. Assume that F has some kind of "decoration"—for instance, one or more marked points, or a geodesic boundary. Although in each case the nature of the set $B \subset F$ so distinguished is different, the procedure for obtaining the spine is the same: identify the set of points $\Sigma \subset F$ which have two or more distinct shortest geodesics to B. (In the case of hyperbolic cusps we do a little more preliminary work to make sense of this.)

The set Σ forms a graph embedded in F. The interiors of the edges are geodesic line segments, each of whose points has precisely two distinct shortest geodesics to B; they connect the vertices, which are the points with three or more such shortest paths.

Let S be the topological fat-graph—that is, a graph with a cyclic ordering of the edges at each vertex—corresponding to Σ, which we recover via the embedding $\epsilon : S \to F$. The *geometric spine* of F is the pair $(S, \tilde{\epsilon})$, where $\tilde{\epsilon}$ is the homotopy class of ϵ. Since the point in Teichmüller space is unchanged if the metric on F is altered by a diffeomorphism homotopic to the identity, the geometric spine is well-defined on $\mathcal{T}F$.

Note that we made no choices when building the spine—it really is an intrinsic feature of the metric on the surface. Furthermore, the construction is equivariant with respect to the mapping class group in the sense that the *combinatorial spine* S (that is, the purely combinatorial data from the geometric spine, without the embedding) is invariant under any diffeomorphism. So the cell decomposition of Teichmüller space which results will descend to moduli space in a completely natural way.

2.2.2 Further construction: the spinal triangulation

Taking the shortest geodesics from each vertex to B, along with Σ, we obtain a "spinal triangulation" of F. It is a genuine hyperbolic triangulation in the case where B consists of marked points; in other cases, F may be dissected into hyperbolic quadrilaterals, or ideal triangles. From each piece in this geometric decomposition we read off a coordinate α_e associated to the corresponding edge of the spine. In the case of an actual triangulation, we measure the angle subtended opposite the spine edge. It is always the case that the pieces on either side of an edge are reflections of each other (since the two shortest geodesics from a point in the interior of an edge to B must be interchanged by reflection in the edge), so given a metric on F these coordinates are well-defined; since they are unaffected by diffeomorphisms homotopic to the identity, they are well-defined on $\mathcal{T}F$ too.

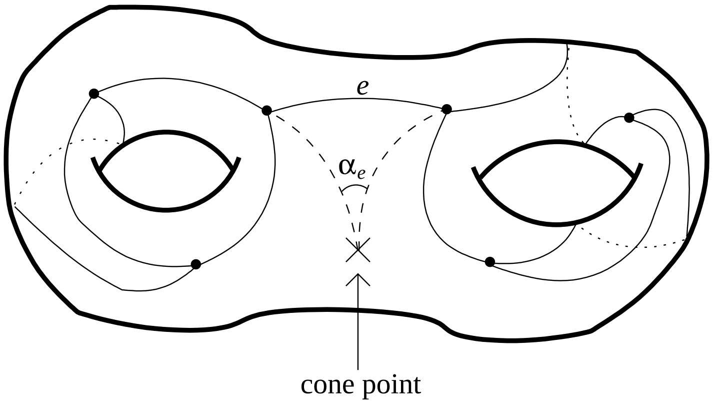

Figure 2.2: A spine on a surface with one marked point

A schematic example of a typical spine is shown in figure 2.2. One of the triangles in the spinal triangulation is illustrated, along with its angular coordinate α_e. Since each vertex in this example is trivalent (that is, has three incident edges), each has three distinct shortest geodesics to the marked point. If we were to add the rest of these to the picture, we would see the complete spinal triangulation, consisting of 18 triangles: 9 pairs of mutual reflections, one pair for each edge.

The α_e are coordinates on the cell in our decomposition corresponding to the geometric spine of F. Natural linear constraints—requiring, for example, that the coordinates around a marked point sum to a specified cone angle—ensure that the cells are all simplices and of the correct dimension. Note that the top-dimensional simplices correspond to graphs S with only trivalent vertices. Lower dimensional simplices have one or more vertices of valency at least four; the faces of a simplex with combinatorial spine S have combinatorial spines obtained from S by "elision" of an edge. (That is, the face $\alpha_e \to 0$ in a simplex with combinatorial spine S will have a combinatorial spine formed from S by removing the edge e and amalgamating its endpoints into a single vertex, as in figure 2.3.)

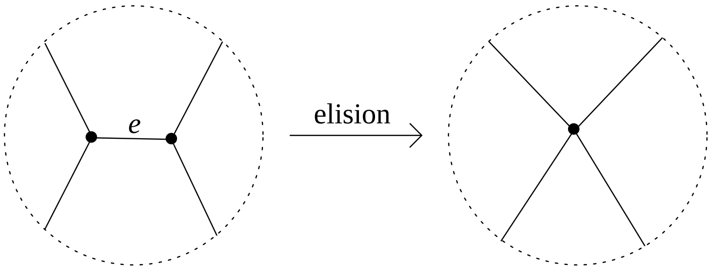

Figure 2.3: Elision of an edge e

2.2.3 Reconstruction and combinatorial curvature

The bulk of our work, to demonstrate that the decomposition we obtain really is the Teichmüller space, consists of showing that knowledge of a geometric spine $(S, \tilde{\epsilon})$ with coordinates α_e is enough to reconstruct a representative metric on F corresponding to a unique point in $\mathcal{T}F$.

Happily, the spine gives us all the combinatorial information of the spinal triangulation: how to glue up the hyperbolic pieces to obtain a metric on F. Unfortunately, the α_e do not, *a priori*, tell us the shapes of the pieces.

Let V be the set of vertices and E the set of edges of S. For each vertex $v_i \in V$, where $1 \leq i \leq n = |V|$, define r_i to be the length of the shortest geodesics from v_i to B. For each edge $e \in E$, we can build a model of the associated hyperbolic piece of the spinal triangulation if we know α_e, r_i and r_j, where v_i and v_j are the vertices joined by e.

Consequently, given S and the α_e, we can make an arbitrary choice for each r_i, construct model hyperbolic pieces and glue them up in the pattern specified by S. This will produce a "metric" on F with the correct combinatorial spine—except that the sum σ_i of angles around each vertex v_i may not equal 2π.

Define the *combinatorial curvature* K_i at the vertex v_i to be the defect in the angle sum: $2\pi - \sigma_i$. Our strategy is to show there is a unique choice of the r_i such that the combinatorial curvature at each vertex is zero (that is, each vertex is "flat"). This ensures that up to diffeomorphism, there is a unique way to construct a model metric on F with combinatorial spine S and associated coordinates α_e. To reconstruct a representative metric for a point in $\mathcal{T}F$, we simply extend the embedding $\epsilon : S \to F$ to a diffeomorphism $\eta : F \to F$ and apply it to our model metric.

Chapter 3

A new proof of an old theorem

In this chapter we present a new proof of a theorem of Bowditch and Epstein [1]. Their proof is arguably more direct—they simply show, without extra machinery, that the map taking an assignment of the variables r_i to the cone angles σ_i is a diffeomorphism onto a certain convex space containing the point $(2\pi, ..., 2\pi)$. However, it's also more mysterious; it doesn't tell you how to find the correct choice for the r_i. We take a more concrete approach, adapting the combinatorial Ricci flow argument of Chow and Luo in [3] to show the existence of a gradient flow on the r_i which converges to the unique solution with zero combinatorial curvature at each vertex in the spine.

3.1 Cusped hyberbolic surfaces

A *cusp* in a hyperbolic surface may be defined as a subset of the surface isometric to the region

$$\frac{\{z : \mathrm{Im} z \geq y\}}{[z \mapsto z + c]}$$

in the upper half-plane model of the hyperbolic plane, for some fixed $y \in \mathbb{R}^+$. Thus it can be thought of as "a point where the metric becomes infinite."

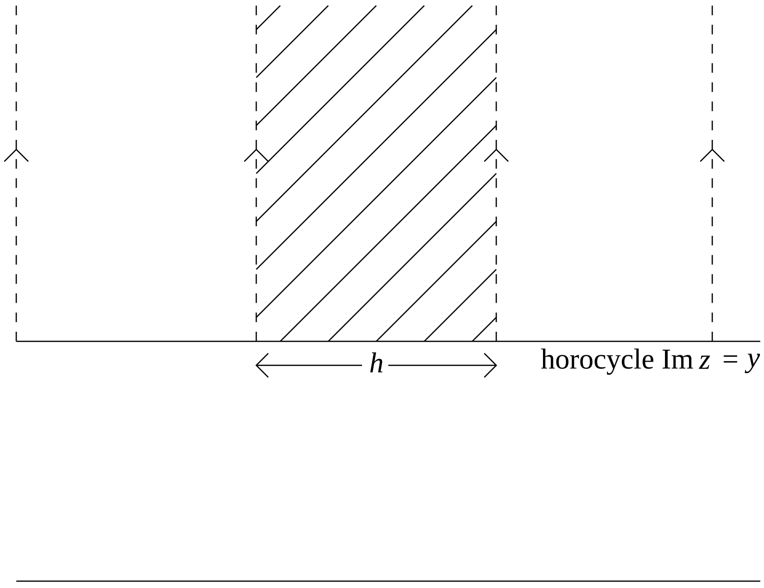

Figure 3.1: A cusp region in the upper half-plane model

Bowditch and Epstein [1] describe the moduli space of hyperbolic surfaces with at least one cusp. For simplicity, we shall assume here that F has exactly one cusp. In order to make sense of the spine, we need to work around the infinite length of geodesics heading into the cusp. We do this by picking a parameter $h > 0$ and considering $F \setminus C$, where the cusp C has horocyclic boundary length h. We then construct the spine Σ in the usual way; it is the set of points with at least two distinct shortest geodesics to the horocyclic boundary B. The coordinates α_e are the lengths subtended along the horocycle, and the constraint becomes $2\sum \alpha_e = h$. As long as we choose h sufficiently small, our choice will not affect the combinatorial type of the spine.

Lemma 3.1. *The spine Σ of a hyperbolic surface with cusps is a graph embedded* *in the surface, each of whose vertices is at least trivalent and each of whose edges is a segment of a geodesic.*

Proof. This is (contained in) lemma 2.2.1 in [1]. □

The following lemma is one of the key ingredients in our proof. The first three statements are used in [1] (see section 4.2)—they tell us in which direction the r_i should flow to head towards the critical point. The new part of the lemma is the fourth statement (compare [3], lemma 2.3)—this suggests to us how to normalise in order to obtain a gradient flow.

Lemma 3.2. *For a hyperbolic triangle with one ideal vertex, horocyclic coordinate α, edge lengths r_i and angles θ_i ($i = 1, 2$),*

(a) $\partial\theta_i/\partial r_i < 0$

(b) $\partial\theta_i/\partial r_j > 0$ *for* $i \neq j$

(c) $\partial(\theta_i + \theta_j)/\partial r_i < 0$

(d) $(\partial\theta_1/\partial r_2)\mathrm{e}^{r_2} = (\partial\theta_2/\partial r_1)\mathrm{e}^{r_1}$

Proof. We work in the upper half-plane model and put the ideal vertex at the point at infinity ($y = \infty$), so the sides of the triangle are parallel vertical lines.

For ease of calculation and without loss of generality, fix the horocycle to lie along the Euclidean line $y = 1$, with the sides of the triangle along $x = \pm d$, as in figure 3.2 (where the parenthetical values are Euclidean coordinates in the upper half-plane).

The hyperbolic distance between points with Euclidean coordinates (x, y) and (u, v) in this model is given by

$$\cosh^{-1}\left(1 + \frac{(x-u)^2 + (y-v)^2}{2yv}\right)$$

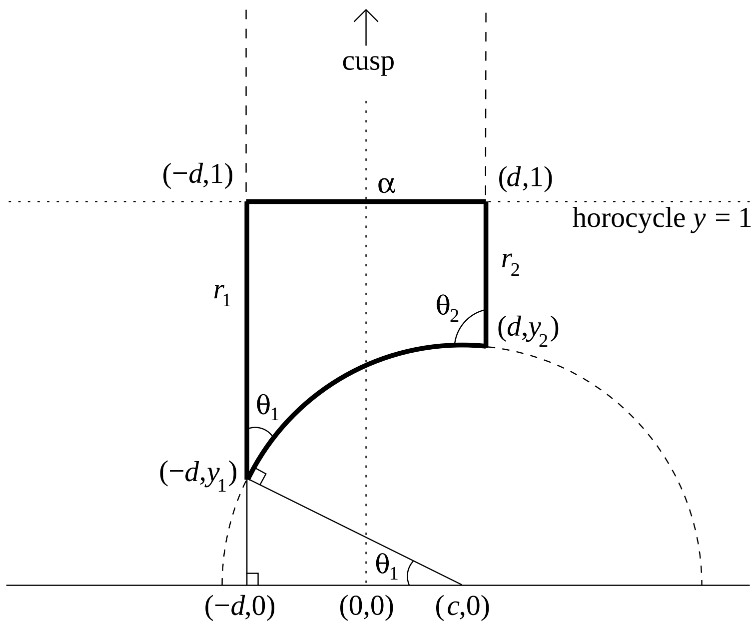

Figure 3.2: Hyperbolic triangle with one ideal vertex

Using this, we compute the Euclidean coordinates of the other two vertices of the triangle. The x-coordinates $\pm d$ depend only on α:

$$\alpha = \cosh^{-1}\left(1 + \frac{4d^2}{2}\right)$$

or

$$d = \sqrt{\frac{\cosh\alpha - 1}{2}}$$

while y_1 and y_2 depend only on r_1 and r_2, respectively:

$$r_i = \cosh^{-1}\left(1 + \frac{(y_i - 1)^2}{2y_i}\right)$$

or

$$0 = {y_i}^2 - 2y_i \cosh r_i + 1$$

which yields

$$y_i = \cosh r_i - \sinh r_i = \mathrm{e}^{-r_i}$$

(Note that this solution is valid for all values of r_i, positive and negative.)

Now, in the upper half-plane model, the third side of the triangle—the geodesic between these two vertices—is a circle with centre $(c, 0)$ on the x-axis. Since the vertices are (Euclideanly) equidistant from this centre, we have

$$(c+d)^2 + \mathrm{e}^{-2r_1} = (c-d)^2 + \mathrm{e}^{-2r_2}$$

which gives

$$c = \frac{\mathrm{e}^{-2r_2} - \mathrm{e}^{-2r_1}}{2\sqrt{2\cosh\alpha - 2}}$$

By Euclidean trigonometry in the triangle formed by $(c, 0)$, $(-d, y_1)$ and $(-d, 0)$, we obtain

$$\cot\theta_1 = \frac{c+d}{\mathrm{e}^{-r_1}} = \frac{\mathrm{e}^{-2r_2+r_1} - \mathrm{e}^{-r_1} + \mathrm{e}^{r_1}(2\cosh\alpha - 2)}{2\sqrt{2\cosh\alpha - 2}} \tag{3.1}$$

so

$$\frac{\partial \cot\theta_1}{\partial r_1} = \frac{\mathrm{e}^{-2r_2+r_1} + \mathrm{e}^{-r_1} + \mathrm{e}^{r_1}(2\cosh\alpha - 2)}{2\sqrt{2\cosh\alpha - 2}} > 0$$

and

$$\frac{\partial \cot\theta_1}{\partial r_2} = -\frac{\mathrm{e}^{-2r_2+r_1}}{\sqrt{2\cosh\alpha - 2}} < 0$$

which give us the first two statements in the lemma.

The third statement is also clear since lengthening a side increases the area of the triangle, which in hyperbolic geometry is equal to π minus the sum of its angles.

Finally, we observe that

$$\frac{\partial \cot\theta_1}{\partial \theta_1} = -\frac{1}{\sin^2\theta_1}$$

Via more trigonometry, we see

$$-\sin^2\theta_1 = -\frac{\mathrm{e}^{-2r_1}}{(c+d)^2+\mathrm{e}^{-2r_1}}$$

By the equation we used to find the centre of the circle, this can be rewritten

$$-\frac{\mathrm{e}^{-2r_1}}{\frac{1}{2}\left[(c+d)^2+(c-d)^2+\mathrm{e}^{-2r_1}+\mathrm{e}^{-2r_2}\right]}$$

and hence

$$\frac{\partial\theta_1}{\partial r_2} = \frac{\frac{\partial\cot\theta_1}{\partial r_2}}{\frac{\partial\cot\theta_1}{\partial\theta_1}} = \frac{2\mathrm{e}^{-2r_2-r_1}}{\sqrt{2\cosh\alpha-2}(c^2+d^2+\mathrm{e}^{-2r_1}+\mathrm{e}^{-2r_2})}$$

Since c^2 is symmetrical in r_1 and r_2, so is the whole denominator of this expression; if we multiply by a factor of e^{r_2} the numerator is too. The final statement in the lemma follows, and this completes the proof. □

Note that we have also shown that the partial derivatives are *elementary functions* in r_i (that is, obtained from the r_i via composition of polynomials, exponentials, logarithms, trigonometric functions and basic arithmetic operations).

Furthermore, the equation (3.1) used in the proof above is enough to establish the following result, which we shall need later.

Scholium 3.3. *In a triangle as in lemma 3.1, for any $\epsilon > 0$ there exists a number L so that when $r_1 > L$, $\theta_1 < \epsilon$.*

3.2 Constructing the gradient flow

We now construct the gradient flow as per [3] section 3.

Definition 3.4. The *cusped combinatorial Ricci flow* is given by

$$\frac{\mathrm{d}r_i(t)}{\mathrm{d}t} = -K_i\mathrm{e}^{r_i(t)}$$

Proposition 3.5. *The Ricci flow (3.4) is the negative gradient flow of a strictly convex function.*

Proof. Let $u_i = -\mathrm{e}^{-r_i}$, so that $\mathrm{d}u_i/\mathrm{d}r_i = \mathrm{e}^{-r_i}$, and let $U = (-\infty, 0)^n$. Under this change of variable, the Ricci flow becomes

$$\frac{\mathrm{d}u_i}{\mathrm{d}t} = -K_i$$

which is a gradient flow in u.

To see this, consider the final statement in lemma 3.2. The change of variable converts

$$\frac{\partial K_i}{\partial r_j}\mathrm{e}^{r_j} = \frac{\partial K_j}{\partial r_i}\mathrm{e}^{r_i}$$

to

$$\frac{\partial K_i}{\partial u_j} = \frac{\partial K_j}{\partial u_i}$$

Thus the 1-form $\sum_{i=1}^n K_i \mathrm{d}u_i$ is closed in the (simply connected) space U, allowing us to define a function $f(u) = \int_a^u \sum_{i=1}^n K_i \mathrm{d}u_i$, where a is any point in U. Clearly $\partial f(u)/\partial u_i = K_i$, so the Ricci flow is the negative gradient flow of the function f.

It remains to show that $f(u) : U \to \mathbb{R}$ is strictly convex. To this end, consider the Hessian of f. Let

$$a_{ij} = \frac{\partial^2 f}{\partial u_i \partial u_j} = \frac{\partial K_i}{\partial u_j}$$

Note that

$$a_{ii} = -\sum_{j\sim i} \frac{\partial \theta_i^j}{\partial r_i}\mathrm{e}^{r_i}$$

where the sum is over indices j such that v_j is adjacent to v_i, and θ_i^j is the sum of the angles at v_i in the triangles whose non-ideal vertices are v_i and v_j. Also

$$a_{ij} = -\sum_{e>i,j} \frac{\partial \theta_i^e}{\partial r_j}\mathrm{e}^{r_j} \quad (i \neq j)$$

where the sum is over all edges joining v_i and v_j, and θ_i^e is the sum of the angles at v_i in the triangles on the edge e.

By lemma 3.2, $a_{ii} > 0$ and $a_{ij} \leq 0$ ($a_{ij} < 0$ if v_i is adjacent to v_j). Also, $(\partial\theta_i^e/\partial r_i)\mathrm{e}^{r_i} + (\partial\theta_i^e/\partial r_j)\mathrm{e}^{r_j} < 0$ if $i, j < e$. Summing over all such inequalities at the vertex v_i, we conclude that $\sum_{j=1}^n a_{ij} > 0$. Since $a_{ii} > 0$ and $a_{ij} \leq 0$, it follows that $a_{ii} > \sum_{j\neq i} |a_{ij}|$.

Finally, to establish the proposition, we appeal to the following result from linear algebra.

Lemma 3.6. *Suppose $A = [a_{ij}]_{n\times n}$ is a symmetric matrix. If $a_{ii} > \sum_{j\neq i} |a_{ij}|$ for all indices i, then A is positive definite.*

A proof of this statement is given in [3] (the first part of lemma 3.11); we don't repeat it here. This completes the proof of the proposition. □

Corollary 3.7. *The map sending the side lengths $(r_1, ..., r_n)$ to the corresponding curvatures $(K_1, ..., K_n)$ is injective. Thus if the zero curvature solution exists, it is unique.*

The corollary follows from the strict convexity of f, since for a smooth strictly convex function defined on an open convex set, the map sending a point to its gradient is always injective.

3.3 Existence and convergence of the flow

We now investigate the evolution of the flow more thoroughly, again following [3] very closely.

Proposition 3.8. *Under the Ricci flow (3.4), the curvature $K_i(t)$ evolves according* *to*

$$\frac{\mathrm{d}K_i}{\mathrm{d}t} = \sum_{j\sim i} C_{ij}(K_j - K_i) - B_i K_i$$

where C_{ij} and B_i are positive elementary functions in $r_1, ..., r_n$. Furthermore, $C_{ij} = C_{ji}$.

Proof. We proceed by focusing on the angles θ_1 and θ_2 in one triangle. Consider a *single triangle flow* under which all other r_i are fixed, so K_1 and K_2 are functions of only r_1 and r_2, which still evolve according to the Ricci flow. We show that under the single triangle flow, the evolution of θ_1 satisfies

$$\frac{\mathrm{d}\theta_1}{\mathrm{d}t} = -A_{12}(K_2 - K_1) + A_1 K_1$$

(and similarly for θ_2), where A_1, A_2 and $A_{12} = A_{21}$ are positive-valued elementary functions in r_1 and r_2.

To this end, by the chain rule we have

$$\frac{\mathrm{d}\theta_1}{\mathrm{d}t} = \frac{\partial\theta_1}{\partial r_1}\frac{\mathrm{d}r_1}{\mathrm{d}t} + \frac{\partial\theta_1}{\partial r_2}\frac{\mathrm{d}r_2}{\mathrm{d}t}$$

$$= -\frac{\partial\theta_1}{\partial r_1}\mathrm{e}^{r_1}K_1 - \frac{\partial\theta_1}{\partial r_2}\mathrm{e}^{r_2}K_2$$

$$= -\frac{\partial\theta_1}{\partial r_2}\mathrm{e}^{r_2}(K_2 - K_1) - \left(\frac{\partial\theta_1}{\partial r_1}\mathrm{e}^{r_1} + \frac{\partial\theta_1}{\partial r_2}\mathrm{e}^{r_2}\right)K_1$$

By lemma 3.2 we can rewrite

$$\frac{\partial\theta_1}{\partial r_1}\mathrm{e}^{r_1} + \frac{\partial\theta_1}{\partial r_2}\mathrm{e}^{r_2} = \frac{\partial\theta_1}{\partial r_1}\mathrm{e}^{r_1} + \frac{\partial\theta_2}{\partial r_1}\mathrm{e}^{r_1} = \mathrm{e}^{r_1}\frac{\partial(\theta_1+\theta_2)}{\partial r_1}$$

and the last partial derivative is negative.

Define

$$A_i = -\mathrm{e}^{r_i}\frac{\partial(\theta_1+\theta_2)}{\partial r_i} \quad , \quad A_{ij} = \frac{\partial\theta_i}{\partial r_j}\mathrm{e}^{r_j} \quad (i \neq j)$$

We have $A_i > 0$. By lemma 3.2, A_{ij} is a positive elementary function and $A_{ij} = A_{ji}$. Hence we have our claim.

The proposition now follows, since the total flow is simply a sum of single triangle flows flows—that is

$$\frac{\mathrm{d}K_i}{\mathrm{d}t} = -\sum_{j\sim i} \frac{\mathrm{d}\theta_i^j}{\mathrm{d}t}$$

□

This proposition has an immediate consequence.

Corollary 3.9 (The maximum principle). *Let $r(t) = (r_1(t), ..., r_n(t))$ be a solution of the Ricci flow (3.4) in an interval. Define $M(t) = \max(K_1(t), ..., K_n(t), 0)$ and $m(t) = \min(K_1(t), ..., K_n(t), 0)$.*

Then $M(t)$ is non-increasing in t and $m(t)$ is non-decreasing in t.

First of all, we use the maximum principle to obtain the following.

Proposition 3.10. *For any initial assignment of edge lengths $r(0) \in \mathbb{R}^n$, the solution to the Ricci flow (3.4) exists for all time $t \geq 0$.*

Proof. We need to show that $r(t)$ remains bounded away from $\pm\infty$ for finite t. By the change of variable we used to construct the gradient flow, $u_i(t) = -\mathrm{e}^{-r_i(t)}$, we know that

$$\frac{\mathrm{d}(\mathrm{e}^{-r_i(t)})}{\mathrm{d}t} = K_i(t)$$

The angle sum at a vertex must be positive, so $K_i < 2\pi$ and we see immediately that

$$\mathrm{e}^{-r_i(t)} < c + 2\pi t$$

where c is a constant. Therefore $r_i(t)$ is bounded away from $-\infty$ for finite t.

On the other hand, suppose (seeking a contradiction) that $r_i(t)$ is not bounded above. By scholium 3.3, as $r_i \to \infty$, the angles in the triangles at v_i will become arbitrarily small, so $\sigma_i \to 0$ and $K_i \to 2\pi$. But this contradicts the maximum principle: the fact that $M(t) < 2\pi$ is non-increasing.

Hence for finite t, $r_i(t)$ is bounded away from $\pm\infty$ and the solution to the Ricci flow exists for all time. $\square$

The maximum principle describes the flow qualitatively—the flow "remains bounded." The following two propositions in combination describe the quantitative behaviour of the flow: that it converges exponentially quickly. First, we establish that the map taking edge lengths to the corresponding curvatures is proper (compare [1] proposition 4.5); then we use this boundedness in the space of edge lengths to prove the convergence.

Proposition 3.11. *If $r(t) \in \mathbb{R}^n$ for each $t \in [0, \infty)$ is an assignment of edge lengths r_i so that $\liminf_{t\to\infty} K_i(r(t)) \geq 0$ and $\limsup_{t\to\infty} K_i(r(t)) < 2\pi$ for all indices i, then the set $\{r(t) | t \in [0, \infty)\}$ lies in a compact region in $\mathbb{R}^n$.*

Proof. By scholium 3.3 and $\limsup_{n\to\infty} K_i(r(t)) < 2\pi$, $r_i(t)$ is bounded from above.

To see that $r_i(t)$ is bounded from below, suppose otherwise that it contains a subsequence $\{r_i(t_j) | j \in \mathbb{N}\}$ converging to $-\infty$. Let I be the non-empty subset of indices so that $\lim_{j\to\infty} r_i(t_j) = -\infty$ for $i \in I$ and $\lim_{j\to\infty} r_i(t_j)$ is finite for $i \notin I$. Since $\liminf_{j\to\infty} K_i(r(t_j)) \geq 0$ for all indices i, it follows that $\lim_{j\to\infty} \sum_{i\in I} K_i(r(t_j)) \geq 0$.

On the other hand, in triangles where $r_1 \to -\infty$ and r_2 remains finite, $\theta_1 \to \pi$; in triangles where r_1 and r_2 both $\to -\infty$, $\theta_1 + \theta_2 \to \pi$. So $\lim_{j\to\infty} \sum_{i\in I} \sigma_i(r(t_j)) = 2\pi |E_I|$, where E_I is the set of edges incident on v_i for some $i \in I$. Since $|E_I| \geq \frac{3}{2}|I|$ (every vertex in the spine is at least trivalent), we have $\lim_{j\to\infty} \sum_{i\in I} \sigma_i(r(t_j)) \geq$

$3\pi|I|$. Rewriting, using $K_i = 2\pi - \sigma_i$, we obtain $\lim_{j\to\infty} \sum_{i\in I} K_i(r(t_j)) \le -\pi|I|$, a contradiction. This completes the proof. $\square$

Proposition 3.12. *Suppose $r(t)$ for $t \in [0,\infty)$ is a solution to the Ricci flow so that the set $\{r(t)|t \in [0,\infty)\}$ lies in a compact region in $\mathbb{R}^n$. Then $r(t)$ converges exponentially fast to a limit R whose curvature at each vertex is zero.*

Proof. Since the coefficient B_i in proposition 3.8 is an elementary function in $r_1, ..., r_n$ and is always positive, it follows from the compactness assumption that there exist two positive constants c_1 and c_2 so that

$$c_1 \le B_i(r_1(t), ..., r_n(t)) \le c_2$$

for all time $t \ge 0$. By proposition 3.8 we obtain

$$\frac{\mathrm{d}M(t)}{\mathrm{d}t} \le -c_1 M(t)$$

and

$$\frac{\mathrm{d}m(t)}{\mathrm{d}t} \ge -c_2 m(t)$$

Thus there are two constants $c_3 > 0$ and $c_4 > 0$ so that

$$c_5 \mathrm{e}^{-c_3 t} \le m(t) \le M(t) \le c_6 \mathrm{e}^{-c_4 t}$$

and hence the curvature $K_i(t)$ converges exponentially fast to zero.

This in turn implies that $\int_0^t K_i(s)\mathrm{d}s$ converges exponentially fast to some constant. Since the Ricci flow can be integrated to

$$\mathrm{e}^{-r_i(t)} = \int_0^t K_i(s)\mathrm{d}s + \mathrm{e}^{-r_i(0)}$$

and r_i remains bounded in a compact region, it follows that the right-hand side remains positive, so $\lim_{t\to\infty} r_i(t) = R_i$ exists and $r_i(t)$ converges to R_i exponentially fast. $\square$

We now have everything we need to prove the following.

Theorem 3.13. *For any initial assignment $r(0)$ of the r_i, the combinatorial Ricci flow (3.4) converges exponentially quickly to the unique solution with zero combinatorial curvature.*

Corollary 3.14. *For a given combinatorial spine S and any choice of coordinates α_e satisfying $2\sum\alpha_e = h$ there is a unique way (up to diffeomorphism) to reconstruct a cusped hyperbolic metric on F with combinatorial spine S and associated coordinates α_e.*

Proof. First, we demonstrate the existence of the zero curvature solution.

Pick an initial assignment of the r_i, $r(0)$, with all $r_i(0)$ sufficiently large so that each K_i is close to 2π (possible by scholium 3.3). This ensures that $K_i(0) \geq 0$. Let $r(t)$ be a solution for $t \geq 0$ to the Ricci flow with initial value $r(0)$. By the maximum principle we have $K_i(t) \geq 0$ for all $t \geq 0$, and also $\limsup_{t\to\infty} K_i(t) < 2\pi$. Thus by proposition 3.11 the set $\{r(t)|t \in [0,\infty)\}$ lies in a compact region in $\mathbb{R}^n$. It then follows by proposition 3.12 that $r(t)$ converges exponentially fast to a solution with zero curvature. In particular, a zero curvature solution exists. This, combined with corollary 3.7—giving us the uniqueness of the zero curvature solution—is enough to establish corollary 3.14, the original theorem of Bowditch and Epstein.

Secondly, we show the convergence of the flow from any initial assignment of the r_i. Here we apply proposition 3.5—that the Ricci flow is the negative gradient flow of the strictly convex function f in the transformed variable u.

Recall that a continuously differentiable function $h : X \to \mathbb{R}$ satisfies the Palais–Smale condition (originally defined in [7]) if for each sequence x_n such that $h(x_n)$ is bounded and $|\nabla h(x_n)| \to 0$, there exists a convergent subsequence x_{n_k}. Furthermore, if h satisfies the Palais–Smale condition, is bounded from below and has non-degenerate critical points, all its negative gradient lines must converge to a critical point.

We already know, by the above existence argument, that f has a unique critical point, which must be a minimum by convexity. So f is bounded from below and the critical point is non-degenerate.

To verify the Palais–Smale condition, take a sequence of points $\{u(t_j)|j \in \mathbb{N}\}$ in $U = (-\infty, 0)^n$ so that the gradient of f at $u(t_j)$ converges to zero. Let the untransformed points corresponding to $u(t_j)$ be $r(t_j)$. Then we have that the curvature $K_i(r(t_j))$ tends to zero for all indices i as $j \to \infty$. By an argument analogous to proposition 3.11, this implies that the set $\{r(t_j)|j \in \mathbb{N}\}$ lies in a compact region in $\mathbb{R}^n$. This in turn implies that $\{u(t_j)\}$ contains a convergent subsequence in U. Hence the Palais–Smale condition holds, and all the gradient lines converge to the critical point.

Finally, since every gradient line converges to the unique minimum of f, it follows that for any initial $r(0) \in \mathbb{R}^n$, the set $\{r(t)|t \in (0, \infty)\}$ lies in a compact region in $\mathbb{R}^n$. Hence by proposition 3.12 again we see that the convergence from any initial choice of edge lengths is exponentially fast. □

Theorem 3.15. *For a given geometric spine $(S, \tilde{\epsilon})$ and any choice of coordinates α_e satisfying $2\sum \alpha_e = h$ there is a unique way (up to diffeomorphism homotopic to the identity) to reconstruct a cusped hyperbolic metric on F with geometric spine $(S, \tilde{\epsilon})$ and associated coordinates α_e.*

Proof. Construct the model metric with combinatorial spine S as in corollary 3.14, and pick an embedding $\epsilon : S \to F$ in the homotopy class $\tilde{\epsilon}$. Since we are only concerned about the resulting metric up to diffeomorphism homotopic to the identity, which embedding we choose is immaterial. The complement $F \setminus \epsilon(S)$ is homeomorphic to an open disc with a point (the cusp) removed, so up to diffeomorphism homotopic to the identity there is a unique way to extend ϵ to a diffeomorphism $\eta : F \to F$. Clearly, the metric we obtain by applying η to our model metric will have geometric spine $(S, \tilde{\epsilon})$ and coordinates α_e, and up to diffeomorphism homotopic to the identity it is unique. $\square$

3.4 Conclusion: the cell complexes

We have proved everything we need to completely describe the Teichmüller and moduli spaces $\mathcal{T}_{g,1}^{\text{cusp}}$ and $\mathcal{M}_{g,1}^{\text{cusp}}$ of hyperbolic surfaces of genus g with one cusp.

Definition 3.16. Let $C_{g,1}$ be a cell complex constructed as follows. Take one simplex $C_{(S,\tilde{\epsilon})}$ for each geometric spine $(S, \tilde{\epsilon})$ which can occur in a hyperbolic surface of genus g with one cusp. For a spine of combinatorial type S, a graph with vertices V_S and edges E_S, $C_{(S,\tilde{\epsilon})}$ is $(|E_S| - 1)$-dimensional, with simplicial coordinates equal to the horocyclic lengths α_e. The conditions $\alpha_e > 0$ and $2\sum_e \alpha_e = h$ ensure that $C_{(S,\tilde{\epsilon})}$ is an open simplex.

The Euler characteristic tells us $2 - 2g - 1 = |V_S| - |E_S|$, and $C_{(S,\tilde{\epsilon})}$ will have biggest dimension when all the vertices are trivalent. In this case $1 - 2g = -|E|/3$, so $|E| = 6g - 3$ and $C_{(S,\tilde{\epsilon})}$ is $(6g - 4)$-dimensional. (There will be one $(6g - 4)$-dimensional cell for each connected trivalent graph with $4g - 2$ vertices and $6g - 3$ edges which can be embedded in the surface so that no edge or pair of edges bounds a disc.) The smallest dimension will occur when $|V_S| = 1$, $|E| = 2g$ and $C_{(S,\tilde{\epsilon})}$ is $(2g - 1)$-dimensional.

Finally, the attaching maps for $C_{g,1}$ are defined as follows. A $(k-1)$-dimensional cell $C_{(S_1,\tilde{\epsilon}_1)}$ in $C_{g,1}$ is the face of a k-dimensional cell $C_{(S_2,\tilde{\epsilon}_2)}$ if S_1 can be obtained from S_2 by removing an edge "homotopically"—that is, removing an edge e_0 with two distinct endpoints, which we amalgamate into a single vertex, in a way which is compatible with the homotopy classes of embeddings $\tilde{\epsilon}_1$ and $\tilde{\epsilon}_2$ in some suitable natural sense. Then the cell $C_{(S_1,\tilde{\epsilon}_1)}$ is the face $\alpha_{e_0} = 0$ in the closure of $C_{(S_2,\tilde{\epsilon}_2)}$, and we glue it in place by identifying the remaining coordinates $\{\alpha_e | e \neq e_0\}$ with the simplicial coordinates on $C_{(S_1,\tilde{\epsilon}_1)}$.

Theorem 3.17. *The Teichmüller space $\mathcal{T}_{g,1}^{\text{cusp}}$ is homeomorphic to the cell complex $C_{g,1}$.*

Proof. Let the map $\Phi_{g,1} : \mathcal{T}_{g,1}^{\text{cusp}} \to C_{g,1}$ be defined by mapping a metric on F with geometric spine $(S, \tilde{\epsilon})$ and associated coordinates α_e to the point with simplicial coordinates α_e in the simplex $C_{(S,\tilde{\epsilon})} \subset C_{g,1}$. Transforming the metric on F by a diffeomorphism homotopic to the identity changes neither the combinatorial type of its spine S, nor the homotopy class of embedding $\epsilon : S \to F$, nor the coordinates α_e, so $\Phi_{g,1}$ is well-defined on Teichmüller space.

The existence part of theorem 3.15 shows that $\Phi_{g,1}$ is surjective; the uniqueness part shows that it is injective and $\Phi_{g,1}^{-1}$ is well-defined. This completes the proof. □

Theorem 3.18. *The moduli space $\mathcal{M}_{g,1}^{\text{cusp}}$ is homeomorphic to the quotient of $C_{g,1}$ by the relation ρ which identifies two points in $C_{g,1}$ if and only if the combinatorial spine types of their respective cells are the same, and their simplical coordinates are equal.*

Proof. The action of the mapping class group on Teichmüller space only affects the embedding of the combinatorial spine S of a metric into F, not S itself or the simplicial coordinates. Hence it is exactly the action of the relation ρ ("forgetting the embedding") on $C_{g,1}$, and dividing by this relation will produce the moduli space. □

Note that in the one-cusped case the parameter h is an overall scale factor, so it does not affect the homeomorphism type of the Teichmüller or moduli spaces.

When there are $c > 1$ cusps, the simplicial complex we get by imposing the overall scaling condition $2\sum_e \alpha_e = h$ is not homeomorphic to $\mathcal{T}_{g,c}^{\text{cusp}}$, but rather to $\mathcal{T}_{g,c}^{\text{cusp}} \times I^{c-1}$. To see the Teichmüller space we need to make further (arbitrary) choices and impose a length condition for each individual horocycle. Each such extra condition will be linear in a subset of the α_e, so we end up with a codimension $c - 1$ slice through the simplicial complex.

Chapter 4

Closed surfaces

We consider the case of a closed surface with marked points—these allow us to define the spine in the usual way, by considering the number of distinct shortest geodesics to a marked point. The Ricci flow works similarly to the previous chapter; the differences are that our variables r_i are now bounded by 0 and $+\infty$, and that there are non-zero angles at the marked points when we glue up triangles (as opposed to the zero angles of 1-ideal triangles at cusps)—consequently, we need to be more careful about whether or not the flow can ever reach a solution where the angle sum at each vertex of the spine is equal to 2π.

4.1 Cone manifolds

For our purposes, a *cone manifold* is a topological surface equipped with a hyperbolic *(cone)* metric so that all but a finite number of *cone points* $P = \{p_1, ..., p_m\}$ have a neighbourhood locally isometric to a piece of the hyperbolic plane $\mathbb{H}^2$. The cone points themselves have respective *cone angles* $\{\gamma_1, ..., \gamma_m\}$ so that a neighbourhood of p_i is isometric to a hyperbolic cone with angle sum γ_i

around the cone point.

Note that a cone manifold with all cone angles equal to 2π is a closed surface with marked points. We are simply considering a slightly more general situation.

We construct the spine Σ for a cone manifold in the usual way—it is precisely the set of points having at least two distinct shortest geodesics to a cone point.

Lemma 4.1. *The spine Σ of a cone manifold is a graph embedded in the surface, each of whose vertices is at least trivalent and each of whose edges is a segment of a geodesic.*

Proof. We prove the following:

(a) any point $x \notin \Sigma$—that is, having a unique shortest geodesic to a cone point—lies in an open neighbourhood $N_x \subset \Sigma$, all of whose points have unique shortest geodesics to a cone point (the same cone point which is closest to x).

(b) if $x \in \Sigma$ has precisely two shortest geodesics to (possibly the same) cone point(s), then it lies in an open neighbourhood N_x such that $N_x \cap \Sigma$ is an open geodesic arc, all of whose points have precisely two shortest geodesics to cone points.

(c) if $x \in \Sigma$ has more than two shortest geodesics to cone points, then it lies in an open neighbourhood N_x, every other point of which has at most two shortest geodesics to cone points.

Proof of (a). Suppose the unique shortest geodesic from x to a cone point c has length L, and the next shortest has length $L' > L$. Take an open horocyclic neighbourhood N_x of radius $\delta = \frac{1}{2}(L' - L) > 0$ about x. For any point $y \in N_x$, by the triangle inequality there is a geodesic path from y to c with length less than $L + \delta$. Again by the triangle inequality, any other geodesic path from y to a cone point has length greater than $L' - \delta$. But we chose δ such that $L + \delta = L' - \delta$, so this establishes statement (a).

Proof of (b). By the same argument as part (a), there is an open horocyclic neighbourhood N_x about x for which the only possible shortest geodesics to cone points are perturbations of the shortest geodesics from x. Now, at least locally, where there are no other cone points near enough to interfere, the locus of points equidistant from the two cone points closest to x (or possibly the same cone point in two different directions) is a geodesic—the fixed set of the reflection which interchanges the cone points. This proves part (b).

Proof of (c). Again, there is an open horocyclic neighbourhood N_x about x for which the only possible shortest geodesics to cone points are perturbations of those from x. Working locally, as in part (b), the set points equidistant to three or more cone points (counting the same point reached in different directions a suitable number of times) is a single point, since it has to be fixed by two different reflections. Hence we have part (c).

Finally, we observe that topologically, $F \setminus \Sigma$ is a disjoint union of m open discs, since a simple closed curve all of whose points have a unique shortest geodesic to the same cone point is necessarily null homotopic, via the straight line homotopy which retracts the curve to the cone point. This shows, in particular, that Σ is non-empty; parts (a), (b) and (c) show that it has the composition claimed in the statement of the lemma. □

This graph, along with the shortest geodesics from its vertices to the cone points, forms a triangulation of the surface; the coordinates we measure are the angles α_e subtended at the cone points. Therefore, the basic lemma we need (which functions exactly as lemma 3.1 did in the previous chapter) addresses triangles with a fixed angle α.

Lemma 4.2. *For a hyperbolic triangle with one fixed angle $\alpha > 0$ formed between two sides of lengths r_1 and r_2, with adjacent angles θ_1 and θ_2,* *(a)* $\partial\theta_i/\partial r_i < 0$

(b) $\partial\theta_i/\partial r_j > 0$ *for* $i \neq j$

(c) $\partial(\theta_i + \theta_j)/\partial r_i < 0$

(d) $(\partial\theta_1/\partial r_2)\sinh r_2 = (\partial\theta_2/\partial r_1)\sinh r_1$

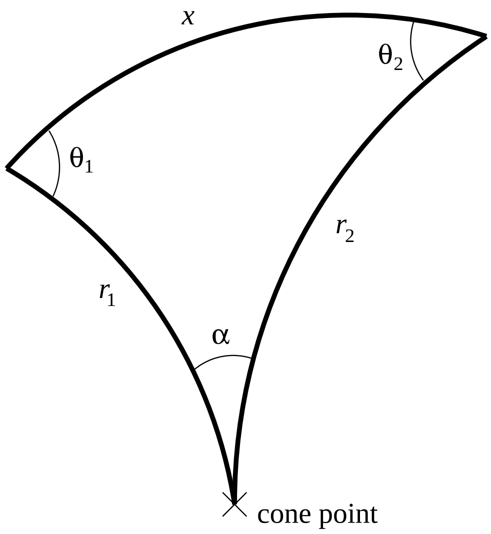

Figure 4.1: Hyperbolic triangle with one fixed angle

Proof. As in figure 4.1, let x be the length of the third side of the triangle. By the hyperbolic sine rule we have

$$\frac{\sin\alpha}{\sinh x} = \frac{\sin\theta_1}{\sinh r_2} \tag{4.1}$$

By the hyperbolic cosine rule,

$$\cosh x = \cosh r_1 \cosh r_2 - \sinh r_1 \sinh r_2 \cos\alpha \tag{4.2}$$

and

$$\cos\theta_1 = \frac{\cosh x \cosh r_1 - \cosh r_2}{\sinh x \sinh r_1} \tag{4.3}$$

Via (4.1), (4.2) and the identity $\cosh^2 - \sinh^2 = 1$, we see

$$\sin\theta_1 = \frac{\sin\alpha\sinh r_2}{\sinh x} = \frac{\sin\alpha\sinh r_2}{\sqrt{(\cosh r_1\cosh r_2 - \sinh r_1\sinh r_2\cos\alpha)^2 - 1}} \tag{4.4}$$

From (4.2), (4.3) and $\cosh^2 - \sinh^2 = 1$,

$$\begin{aligned}\cos\theta_1 &= \frac{\cosh^2 r_1\cosh r_2 - \cosh r_1\sinh r_1\sinh r_2\cos\alpha - \cosh r_2}{\sinh r_1\sqrt{(\cosh r_1\cosh r_2 - \sinh r_1\sinh r_2\cos\alpha)^2 - 1}}\\ &= \frac{\sinh r_1\cosh r_2 - \cosh r_1\sinh r_2\cos\alpha}{\sqrt{(\cosh r_1\cosh r_2 - \sinh r_1\sinh r_2\cos\alpha)^2 - 1}}\end{aligned}$$

Combining these, we obtain

$$\cot\theta_1 = \frac{\cos\theta_1}{\sin\theta_1} = \frac{\sinh r_1\cosh r_2 - \cosh r_1\sinh r_2\cos\alpha}{\sin\alpha\sinh r_2}$$

so

$$\frac{\partial(\cot\theta_1)}{\partial r_1} = \frac{\cosh r_1\cosh r_2 - \sinh r_1\sinh r_2\cos\alpha}{\sin\alpha\sinh r_2}$$

Since $\cos\alpha < 1$, the numerator of this expression is strictly greater than

$$\cosh r_1\cosh r_2 - \sinh r_1\sinh r_2 = \cosh(r_1 - r_2) \geq 1$$

The denominator is also greater than zero, so the quotient is likewise and part (a) follows.

We also have, after cancellation and $\cosh^2 - \sinh^2 = 1$,

$$\frac{\partial(\cot\theta_1)}{\partial r_2} = -\frac{\sinh r_1}{\sin\alpha\sinh^2 r_2} < 0$$

which establishes part (b).

Again, part (c) is clear since lengthening one of the r_i increases the area of the triangle, and thus decreases the sum of its angles.

Finally, as in lemma 3.1,

$$\frac{\partial(\cot\theta_1)}{\partial\theta_1} = -\frac{1}{\sin^2\theta_1}$$

so substituting equation (4.4) above and using the chain rule, we find that

$$\frac{\partial \theta_1}{\partial r_2} = \frac{\sin\alpha \sinh r_1}{(\cosh r_1 \cosh r_2 - \sinh r_1 \sinh r_2 \cos\alpha)^2 - 1}$$

If we multiply by $\sinh r_2$, this expression is symmetrical in r_1 and r_2, so part (d) follows. $\square$

Lemma 4.3. *In a triangle as in lemma 4.2, for any $\epsilon > 0$ there exists a number L so that when $r_1 > L$, $\theta_1 < \epsilon$.*

Proof. We have, as above,

$$\cot\theta_1 = \frac{\cos\theta_1}{\sin\theta_1} = \frac{\sinh r_1 \cosh r_2 - \cosh r_1 \sinh r_2 \cos\alpha}{\sin\alpha \sinh r_2}$$

$$= \frac{(\sinh r_1 \cosh r_2 - \cosh r_1 \sinh r_2) + (1 - \cos\alpha)\cosh r_1 \sinh r_2}{\sin\alpha \sinh r_2}$$

$$= \frac{\sinh(r_1 - r_2)}{\sin\alpha \sinh r_2} + \frac{(1-\cos\alpha)\cosh r_1}{\sin\alpha}$$

Noting that for $r_1, r_2 > 0$, $\sinh(r_1 - r_2)/\sinh r_2 > -1$, we complete the proof by taking L sufficiently large so that

$$\cosh L > \frac{\cot\epsilon \sin\alpha + 1}{1 - \cos\alpha}$$

$\square$

4.2 Gradient flow revisited

Definition 4.4. The *marked point combinatorial Ricci flow* is given by

$$\frac{\mathrm{d}r_i(t)}{\mathrm{d}t} = -K_i \sinh r_i(t)$$

Proposition 4.5. *The Ricci flow (4.4) is the negative gradient flow of a strictly convex function.*

Proof. Let $u_i = \ln\tanh(r_i/2)$, so that $\mathrm{d}u_i/\mathrm{d}r_i = \frac{1}{\sinh r_i}$. This change of variable maps $(0,\infty)^n$ to the same space as before, $U = (-\infty, 0)^n$. The proposition now follows by the same proof given in the last chapter—the change of variable converts the Ricci flow to a gradient flow in u:

$$\frac{\mathrm{d}u_i}{\mathrm{d}t} = -K_i$$

$\square$

Corollary 4.6. *The map sending the side lengths $(r_1, ..., r_n)$ to the corresponding curvatures $(K_1, ..., K_n)$ is injective.*

Proposition 4.7. *Under the Ricci flow (4.4), the curvature $K_i(t)$ evolves according to*

$$\frac{\mathrm{d}K_i}{\mathrm{d}t} = \sum_{j\sim i} C_{ij}(K_j - K_i) - B_i K_i$$

where C_{ij} and B_i are positive elementary functions in $r_1, ..., r_n$. Furthermore, $C_{ij} = C_{ji}$.

Proof. Once again, the proof given in the last chapter obtains here, with the appropriate normalising factor $\sinh r_i$ (rather than e^{r_i}). $\square$

Corollary 4.8 (The maximum principle). *Let $r(t) = (r_1(t), ..., r_n(t))$ be a solution of the Ricci flow (4.4) in an interval. Define $M(t) = \max(K_1(t), ..., K_n(t), 0)$ and $m(t) = \min(K_1(t), ..., K_n(t), 0)$.*

Then $M(t)$ is non-increasing in t and $m(t)$ is non-decreasing in t.

Proposition 4.9. *For any initial assignment of edge lengths $r(0) \in (0,\infty)^n$, the solution to the Ricci flow (4.4) exists for all time $t \geq 0$.*

Proof. We need to show that $r(t)$ is bounded away from 0 and ∞ for finite t. Here the change of variable $u_i = \ln\tanh(r_i/2)$ gives us

$$\frac{\mathrm{d}(\ln\coth(r_i(t)/2))}{\mathrm{d}t} = K_i(t)$$

so since $K_i < 2\pi$ we have

$$\coth r_i(t)/2 < c\mathrm{e}^{2\pi t}$$

for some constant c, and r_i is bounded away from 0 for finite t.

Now by lemma 4.3, if $r_i \to \infty$ then $\sigma_i \to 0$ and $K_i \to 2\pi$, contradicting the non-increasing nature of $M(t) < 2\pi$. Hence r_i is bounded above and $r(t)$ exists for all t. □

We now encounter the biggest difference between the cusped and marked point cases. The first of the two propositions we use to ascertain the convergence of the flow is more complicated—the problem being that if the α_e are collectively "too big" there will be too much angle concentrated at the marked points and not enough left over for the angle sums σ_i at the vertices to all be flat. Consequently we need an extra condition on the sums of the α_e—if this condition is violated the flow will not converge in $(0,\infty)^n$; instead, one or more of the r_i will tend to zero.

Proposition 4.10. *Let S be a combinatorial spine for a cone manifold, with associated coordinates α_e. Suppose that for any non-empty subset of indices $I \subset \{1,...,n\}$, the sum of the coordinates over edges incident at some v_i with $i \in I$ satisfies the condition $\sum_{e<I} \pi - \alpha_e > \pi|I|$.*

Then if $r(t) \in (0,\infty)^n$ for each $t \in [0,\infty)$ is an assignment of edge lengths r_i so that $\liminf_{t\to\infty} K_i(r(t)) \geq 0$ and $\limsup_{t\to\infty} K_i(r(t)) < 2\pi$ for all indices i, then the set $\{r(t)|t \in [0,\infty)\}$ lies in a compact region in $(0,\infty)^n$.

Proof. The proof works exactly as proposition 3.10; the extra condition in the statement ensures that we obtain the contradiction we need.

By lemma 4.3 and $\limsup_{n\to\infty} K_i(r(t)) < 2\pi$, $r_i(t)$ is bounded from above.

To see that $r_i(t)$ is bounded from below, suppose otherwise that it contains a subsequence $\{r_i(t_j)|j \in \mathbb{N}\}$ converging to 0. Let I be the non-empty subset of indices so that $\lim_{j\to\infty} r_i(t_j) = 0$ for $i \in I$ and $\lim_{j\to\infty} r_i(t_j) > 0$ for $i \notin I$. Since $\liminf_{j\to\infty} K_i(r(t_j)) \geq 0$ for all indices i, it follows that $\lim_{j\to\infty} \sum_{i\in I} K_i(r(t_j)) \geq 0$.

On the other hand, if the triangle on the edge e has $r_1 \to 0$ and $r_2 \nrightarrow 0$, $\theta_1 \to \pi - \alpha_e$; if instead r_1 and r_2 both $\to 0$, $\theta_1 + \theta_2 \to \pi - \alpha_e$. So $\lim_{j\to\infty} \sum_{i\in I} \sigma_i(r(t_j)) = \sum_{e<I} 2(\pi - \alpha_e)$ (since there are two copies of the triangle on an edge). Rewriting, using $K_i = 2\pi - \sigma_i$, we obtain $\lim_{j\to\infty} \sum_{i\in I} K_i(r(t_j)) = 2\pi|I| - 2\sum_{e<I} \pi - \alpha_e$. By the condition on the α_e, this gives us $\lim_{j\to\infty} \sum_{i\in I} K_i(r(t_j)) < 0$, a contradiction. This completes the proof. □

Proposition 4.11. *Suppose $r(t)$ for $t \in [0,\infty)$ is a solution to the Ricci flow so that the set $\{r(t)|t \in [0,\infty)\}$ lies in a compact region in $(0,\infty)^n$. Then $r(t)$ converges exponentially fast to a limit R whose curvature at each vertex is zero.*

Proof. The argument is identical to that in the previous chapter, with the marked point Ricci flow integrating to

$$\coth(r_i(t)/2) = c\mathrm{e}^{\int_0^t K_i(s)\mathrm{d}s}$$

□

Theorem 4.12. *If $\sum_{e<I} \pi - \alpha_e > \pi|I|$ for all sets of indices $I \subset \{1, ..., n\}$ then for any initial assignment $r(0)$ of the r_i, the combinatorial Ricci flow (4.4) converges exponentially quickly to the unique solution with zero combinatorial curvature.*

Corollary 4.13. *For a given combinatorial spine S and any choice of coordinates α_e, there is a way to reconstruct a hyperbolic cone metric on F with spine S and coordinates α_e if and only if $\sum_{e<I} \pi - \alpha_e > \pi|I|$ for all non-empty sets of indices*

$I \subset \{1, ..., n\}$. Furthermore, when this reconstruction exists it is unique (up to diffeomorphism).

Proof. The proof of the theorem and the "if" part of the corollary are exactly as in chapter 3. To see the "only if" part, suppose otherwise that there is a set of indices $I \subset \{1, ..., n\}$ with $\sum_{e<I} \pi - \alpha_e \leq \pi|I|$. Note that since $r_i \in (0, \infty)$, all the hyperbolic triangles have positive area and thus angle sum $\Delta_e < \pi$. It follows that $\sum_{i \in I} \sigma_i \leq 2 \sum_{e<I} \Delta_e - \alpha_e < 2 \sum_{e<I} \pi - \alpha_e \leq 2\pi|I|$, so the curvatures K_i can never all be zero. Thus there exists no cone manifold with the given spine type and coordinates. $\square$

Theorem 4.14. *For a given geometric spine $(S, \tilde{\epsilon})$ and any choice of coordinates α_e there is a way to reconstruct a hyperbolic cone metric on F with geometric spine $(S, \tilde{\epsilon})$ and associated coordinates α_e if and only if $\sum_{e<I} \pi - \alpha_e > \pi|I|$ for all non-empty sets of indices $I \subset \{1, ..., n\}$. Furthermore, when this reconstruction exists it is unique up to diffeomorphism homotopic to the identity.*

Proof. The theorem follows exactly as in the previous chapter: the complement $F \setminus \epsilon(S)$ is homeomorphic to a disjoint union of open discs, so again there is a unique extension to a diffeomorphism $\eta : F \to F$. Apply η to the model metric from corollary 4.13 to obtain a representative metric which is unique up to diffeomorphism homotopic to the identity. $\square$

4.3 The cell complex

We can now describe the Teichmüller and moduli spaces $\mathcal{T}^{\text{m.p.}}_{g,m,\Gamma}$ and $\mathcal{M}^{\text{m.p.}}_{g,m,\Gamma}$ of cone manifolds of genus g with m cone points and total cone angle Γ.

Definition 4.15. Let $C_{g,m,\Gamma}$ be a "truncated" cell complex constructed as follows. Take one simplex $C_{(S,\tilde{\epsilon})}$ for each $(S, \tilde{\epsilon})$ which can appear as the geometric spine of a cone manifold of genus g with m marked points. For a spine of combinatorial type S with vertices V_S and edges E_S, C_S is $(|E_S|-1)$-dimensional, with simplicial coordinates equal to the angles subtended at the marked points α_e. We impose an overall scaling condition $2\sum_e \alpha_e = \Gamma$, which, along with $\alpha_e > 0$, ensures that $C_{(S,\tilde{\epsilon})}$ is an open simplex.

We now "truncate" each simplex by imposing the condition $\sum_{e<I} \pi - \alpha_e > \pi|I|$ for each non-empty set of indices $I \subset \{1, ..., |V_S|\}$.

The attaching maps for $C_{g,m,\Gamma}$ are defined in the same way as before: a $(k-1)$-dimensional cell $C_{(S_1,\tilde{\epsilon})}$ in $C_{g,m,\Gamma}$ is the face of a k-dimensional cell $C_{(S_2,\tilde{\epsilon})}$ if S_1 can be obtained from S_2 by removing an edge e_0 "homotopically." Since the face corresponds to $\alpha_{e_0} = 0$, the truncating planes are compatible with the gluing: in the limit $\alpha_{e_0} \to 0$, if $e_0 < I$, $\sum_{e<I} \pi - \alpha_e > \pi|I|$ if and only if $\sum_{e<I, e\neq e_0} \pi - \alpha_e > \pi(|I|-1)$.

Theorem 4.16. *The Teichmüller space $\mathcal{T}^{\text{m.p.}}_{g,m,\Gamma}$ is homeomorphic to the truncated cell complex $C_{g,m,\Gamma}$.*

Proof. Let the map $\Phi_{g,m,\Gamma} : \mathcal{T}^{\text{m.p.}}_{g,m,\Gamma} \to C_{g,m,\Gamma}$ be defined by mapping a cone manifold F with geometric spine $(S, \tilde{\epsilon})$ and associated coordinates α_e to the point with simplicial coordinates α_e in the simplex $C_{(S,\tilde{\epsilon})} \subset C_{g,m,\Gamma}$. Again, transforming the metric on F by a diffeomorphism homotopic to the identity changes neither the combinatorial type of its spine S, nor the homotopy class of embedding $\epsilon : S \to F$, nor the coordinates α_e, so $\Phi_{g,m,\Gamma}$ is well-defined on Teichmüller space.

The existence part of theorem 4.14 shows that $\Phi_{g,m,\Gamma}$ is surjective; the uniqueness part shows that it is injective and $\Phi^{-1}_{g,m,\Gamma}$ is well-defined. This completes the proof. $\square$

Theorem 4.17. *The moduli space $\mathcal{M}^{\text{m.p.}}_{g,m,\Gamma}$ is homeomorphic to the quotient of $C_{g,m,\Gamma}$ by the relation ρ which identifies two points in $C_{g,m,\Gamma}$ if and only if the* *combinatorial spine types of their respective cells are the same, and their simplicial coordinates are equal.*

Finally, to obtain the Teichmüller and moduli spaces of cone manifolds with $m > 1$ marked points and individually prescribed cone angles, we apply the additional linear constraints on the spine coordinates to take codimension $m-1$ slices through the simplicial complexes $\mathcal{T}^{\text{m.p.}}_{g,m,\Gamma}$ and $\mathcal{M}^{\text{m.p.}}_{g,m,\Gamma}$.

Chapter 5

Surfaces with geodesic boundary

The third and final case we tackle with the combinatorial Ricci flow is that of hyperbolic surfaces with geodesic boundary. We construct the spine in the usual way: it is the set of points having at least two distinct shortest (necessarily geodesic) paths to the boundary. Since the locus of points equidistant to two geodesics is itself a geodesic (because—at least locally—there is a reflection which interchanges them), the geometric spine is once again an embedded graph whose edges are geodesic line segments.

Adding in the shortest paths from the vertices of the spine to the boundary we obtain not a triangulation, but rather a "quadrilateralation"—the pieces are hyperbolic quadrilaterals with two right-angles at the boundary. The coordinates α_e we measure are the side lengths along the boundary.

The only significant difference in the result from the cusped case of chapter 3 is that the overall scale factor is no longer an arbitrary choice (recall we had to choose the total horocyclic length of the truncated cusps), but the actual total length of the geodesic boundary.

5.1 How this case differs from the others

Rather than repeat a series of near-identical results, we instead note the differences between this case and those already proven. The biggest such is that we need a new lemma about the quadrilateral pieces, and another to serve the function of scholium 3.2.

Lemma 5.1. *For a hyperbolic quadrilateral with one fixed side of length $\alpha > 0$ meeting sides of lengths r_1 and r_2 at right angles, and with adjacent angles θ_1 and θ_2,*

(a) $\partial\theta_i/\partial r_i < 0$
(b) $\partial\theta_i/\partial r_j > 0$ *for* $i \neq j$
(c) $\partial(\theta_i + \theta_j)/\partial r_i < 0$
(d) $(\partial\theta_1/\partial r_2)\cosh r_2 = (\partial\theta_2/\partial r_1)\cosh r_1$

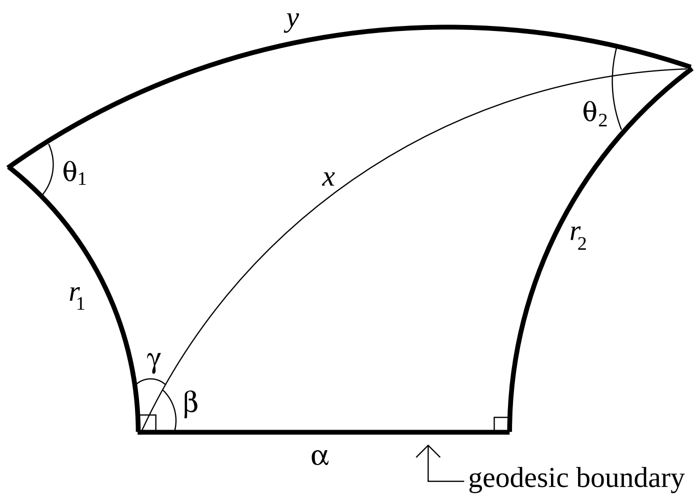

Figure 5.1: Hyperbolic quadrilateral with one fixed side

Proof. As in figure 5.1, let y be the length of the fourth side, and let x be the length of the diagonal forming a triangle with the fixed side and the side of length r_2. Further, let β be the angle in this triangle opposite r_2, and let $\gamma = \pi/2 - \beta$, the angle opposite the side of length y in the triangle of side lengths y, r_1, x.

We proceed by obtaining a "sine rule" and a "cosine rule" for this particular type of right-angled hyperbolic quadrilateral.

First we obtain the "sine rule." By the hyperbolic cosine rule for triangles,

$$\sin\gamma = \cos\beta = \frac{\cosh x \cosh\alpha - \cosh r_2}{\sinh x \sinh\alpha}$$

and thus by the hyperbolic sine rule for triangles,

$$\frac{\sin\theta_1}{\sinh x} = \frac{\sin\gamma}{\sinh y} = \frac{\cosh x \cosh\alpha - \cosh r_2}{\sinh x \sinh\alpha \sinh y}$$

From the hyperbolic cosine rule for a right-angled triangle, $\cosh x = \cosh\alpha \cosh r_2$. Substituting in for $\cosh x$ we obtain the "sine rule" for hyperbolic quadrilaterals (the second equality by symmetry):

$$\frac{\sinh y}{\sinh\alpha} = \frac{\cosh r_2}{\sin\theta_1} = \frac{\cosh r_1}{\sin\theta_2}$$

Now we prove the "cosine rule." By the sine rule for hyperbolic triangles,

$$\frac{\sin\beta}{\sinh r_2} = \frac{1}{\sinh x}$$

so

$$\cos\gamma = \frac{\sinh r_2}{\sinh x}$$

By the cosine rule for hyperbolic triangles, we have

$$\cosh y = \cosh r_1 \cosh x - \sinh r_1 \sinh x \cos\gamma$$

which together with the above expression for $\cos\gamma$ gives us the "cosine rule" for hyperbolic triangles:

$$\cosh y = \cosh r_1 \cosh r_2 \cosh\alpha - \sinh r_1 \sinh r_2$$

We now complete the proof by using these rules to obtain an expression for θ_1. By the hyperbolic cosine rule for triangles, we have

$$\cosh x = \cosh r_1 \cosh y - \sinh r_1 \sinh y \cos\theta_1$$

Using $\cosh x = \cosh\alpha \cosh r_2$ and the above formulae for $\cosh y$ and $\sinh y$, we end up with

$$\cot\theta_1 = \coth\alpha \sinh r_1 - \frac{\cosh r_1 \sinh r_2}{\sinh\alpha \cosh r_2}$$

Differentiating with respect to r_1, we get

$$\frac{\partial(\cot\theta_1)}{\partial r_1} = \coth\alpha \cosh r_1 - \frac{\sinh r_1 \sinh r_2}{\sinh\alpha \cosh r_2} > 0$$

which establishes part (a).

With respect to r_2, the partial derivative is

$$\frac{\partial(\cot\theta_1)}{\partial r_2} = -\frac{\cosh r_1}{\sinh\alpha \cosh^2 r_2} < 0$$

and part (b) follows.

Again, part (c) is clear since lengthening a side increases the area of the quadrilateral, so decreases the sum of its angles.

For part (d), we have by the "sine rule" that

$$\frac{1}{\frac{\partial(\cot\theta_1)}{\partial\theta_1}} = -\sin^2\theta_1 = -\frac{\cosh^2 r_2 \sinh^2\alpha}{\sinh^2 y}$$

By the "cosine rule" and $\cosh^2 y - \sinh^2 y = 1$,

$$\sinh^2 y = (\cosh r_1 \cosh r_2 \cosh\alpha - \sinh r_1 \sinh r_2)^2 - 1$$

so the final result, by the chain rule, is

$$\frac{\partial\theta_1}{\partial r_2} = \frac{\sinh\alpha\cosh r_1}{(\cosh r_1\cosh r_2\cosh\alpha - \sinh r_1\sinh r_2)^2 - 1}$$

This expression is symmetrical in r_1 and r_2, except for the factor of $\cosh r_1$ in the numerator, so the result follows. □

Lemma 5.2. *For a hyperbolic quadrilateral as in lemma 5.1, for any $\epsilon > 0$ there exists a number L so that when $r_1 > L$, $\theta_1 < \epsilon$.*

Proof. We have

$$\cot\theta_1 = \coth\alpha\sinh r_1 - \frac{\cosh r_1\sinh r_2}{\sinh\alpha\cosh r_2}$$

so

$$\cot\theta_1 > \frac{1}{\sinh\alpha}(\cosh\alpha\sinh r_1 - \cosh r_1) = \frac{\mathrm{e}^{r_1}(\cosh\alpha - 1) - 2\mathrm{e}^{-r_1}}{2\sinh\alpha}$$

Since $\alpha > 0$, $\cosh\alpha > 1$, so we can choose L sufficiently large that

$$\mathrm{e}^L > \frac{2\sinh\alpha\cot\epsilon + 2}{\cosh\alpha - 1}$$

and we are done. □

Now we define the new Ricci flow, using the normalising factor indicated by lemma 5.1.

Definition 5.3. The *geodesic boundary combinatorial Ricci flow* is given by

$$\frac{\mathrm{d}r_i(t)}{\mathrm{d}t} = -K_i\cosh r_i(t)$$

To prove that (5.3) is the negative gradient flow of a strictly convex function, we use the change of variable $u_i = \sin^{-1}(\tanh r_i)$. (Our space $U = (0, \pi/2)^n$ is different from before, but this is irrelevant.)

The rest of the argument of chapter 3 now obtains, allowing only that the normalising factor should be $\cosh r_i$ (not e^{r_i}), and that the r_i live in $(0,\infty)^n$, as in chapter 4 (not $\mathbb{R}^n$).

Skipping straight to the end, we obtain characterisations of $\mathcal{T}^{\text{g.b.}}_{g,b,\lambda}$ and $\mathcal{M}^{\text{g.b.}}_{g,b,\lambda}$, the Teichmüller and moduli spaces of hyperbolic surfaces of genus g with b closed geodesic boundary components and total boundary length λ.

Definition 5.4. Let $C_{g,b,\lambda}$ be a cell complex constructed as follows. Take one simplex $C_{(S,\tilde{\epsilon})}$ for each $(S,\tilde{\epsilon})$ which can appear as the geometric spine of a hyperbolic surface of genus g with boundary consisting of b closed geodesics. The conditions $\alpha_e > 0$ and $2\sum_e \alpha_e = \lambda$ ensure that $C_{(S,\tilde{\epsilon})}$ is an open simplex. Define the attaching maps in the usual way, by identifying a $(k-1)$-simplex corresponding to $\alpha_{e_0} = 0$ in the closure of a k-simplex as the appropriate face of that k-simplex.

Theorem 5.5. *The Teichmüller space $\mathcal{T}^{\text{g.b.}}_{g,b,\lambda}$ is homeomorphic to the cell complex $C_{g,b,\lambda}$. The moduli space $\mathcal{M}^{\text{g.b.}}_{g,b,\lambda}$ is homeomorphic to the quotient of $C_{g,b,\lambda}$ by the relation which identifies two points if and only if the geometric spines corresponding to their respective simplices have the same combinatorial type, and their coordinates within those simplices are equal.*

Chapter 6

Surfaces with polygonal geodesic boundary

In this chapter we do something a little different—we use the theorem about cone manifolds we proved in chapter 4 to investigate the Teichmüller and moduli spaces of hyperbolic surface metrics which have polygonal geodesic boundary.

Definition 6.1. A *hyperbolic metric with polygonal geodesic boundary* on a topological surface F (with non-empty boundary ∂F) is a hyperbolic metric under which each component of ∂F is composed of a finite number of geodesic arcs. Points where two ends of geodesic boundary arc meet need not have a neighbourhood locally isometric to a hyperbolic half-plane, only to a "hypberbolic half-cone"—that is, we allow the angle sum at those points to differ from π.

6.1 Polygonal boundaries and cone points

6.1.1 Reflection...

Suppose G is a closed, orientable topological surface. Choose a set of simple closed curves $C = \{c_1, ..., c_n\}$ embedded in G which disconnects G into two pieces, each of genus g. Also pick a topological reflection $\sigma : G \to G$ which acts by swapping the two pieces, keeping C fixed. Any hyperbolic metric or hyperbolic cone metric h on G has a well-defined induced reflection $\sigma^*(h)$, which is also a (cone) metric on G. Furthermore, if h_1 and h_2 are diffeomorphic via a diffeomorphism ϕ isotopic to the identity, then so are $\sigma^*(h_1)$ and $\sigma^*(h_2)$, via the conjugated map $\sigma^* \circ \phi \circ (\sigma^*)^{-1}$. Hence σ^* is a well-defined involution acting on the Teichmüller space of (cone) metrics on G.

6.1.2 ...and doubling

Fix a topological surface F of genus g with a non-empty boundary ∂F of b components. The *topological double* DF of F is formed by gluing two copies F_+ and F_- of F together, identifying their boundaries via the identity map. DF is a closed topological surface with genus $g' = 2g + b - 1$.

Now let $C = \{c_1, ..., c_b\}$ be the images of the boundary components of F_+ (and, necessarily, F_-, since the boundaries are identified) under the inclusions $i_\pm : F_\pm \to DF$. C is a set of simple closed curves embedded in DF. Let $\sigma : DF \to DF$ be the topological reflection which interchanges F_+ and F_- via the identity, keeping C fixed.

Now suppose F is equipped with h, a hyperbolic metric with polygonal geodesic boundary. The *geometric double* of the pair (F, h) is the pair (DF, Dh), where DF is the topological double of F and Dh is a hyperbolic cone metric on DF, formed from h as follows. On the interior of one copy of F, $F_+ \subset DF$ say, let Dh be the metric induced by the inclusion, $i_+^*(h)|_{\mathrm{Int}F_+}$. At each point in the boundary ∂F_+ (identified with ∂F_- in DF), we take the half-plane (or half-cone) neighbourhood given by h and glue it to its (canonical) reflection to obtain the metric Dh at that point. We now extend this reflection process, using the reflection σ to induce the metric on $\mathrm{Int}F_-$ from the metric we already have on $\mathrm{Int}F_+$.

We see that under Dh all but a finite number of points will have neighbourhoods locally isometric to the hyperbolic plane; the exceptions (the points which were formerly "half-cone points" in the boundary of F) will have neighbourhoods locally isometric to hyperbolic cones, and we see that Dh is indeed a cone metric.

6.2 A theorem

Let F be as above, which fixes also the double DF and the reflection $\sigma : DF \to DF$. Our strategy is to identify the Teichmüller space of hyperbolic metrics with polygonal geodesic boundary on F by describing how the Teichmüller space of their doubles sits inside the Teichmüller space of hyperbolic cone metrics on DF. To close, we shall then note that everything we do at the level of Teichmüller space is equivariant under the action of the mapping class group, so the same construction descends to moduli space.

Theorem 6.2. *The Teichmüller space of hyperbolic metrics with polygonal geodesic boundary and total half-cone angle* Γ *on* F*,* $\mathcal{T}^{\text{p.g.b.}}_{g,b,\Gamma}$*, is homeomorphic to the subset* T_σ *of* $\mathcal{T}^{\text{m.p.}}_{g',b,2\Gamma}(DF)$ *which is fixed by the action of* σ^**.*

Proof. Let $\phi : \mathcal{T}^{\text{p.g.b.}}_{g,b,\Gamma} \to \mathcal{T}^{\text{m.p.}}_{g',b,2\Gamma}(DF)$ be the map which sends a metric h with polygonal boundary on F to its double Dh on DF . If two such metrics h and h' are in the same Teichmüller class, they are related by a diffeomorphism of F

homotopic to the identity via a homotopy which fixes the boundary—so Dh and Dh' are certainly related by a diffeomorphism homotopic to the identity, and ϕ is well-defined.

Futhermore, by construction Dh is fixed by σ, so the Teichmüller class of Dh is fixed by σ^* and $\phi : \mathcal{T}_{g,b,\Gamma}^{\text{p.g.b.}} \to T_\sigma$.

We now want to define an inverse for ϕ—we want to know how to construct an (essentially unique) representative hyperbolic metric with polygonal boundary on F from a Teichmüller class of metric on DF which is fixed by σ^*.

Observe first that σ acts on the set of geometric spines for DF, and that $\sigma(S, \tilde{\epsilon}) = (S, \tilde{\epsilon})$ if and only if S comes equipped with an involution ι of its vertices (fixing none of them) such that there is an isomorphism of graphs $\iota S \cong S$, and ι is compatible with σ and the homotopy class of embedding, that is $\sigma(\epsilon S)$ is homotopic to $\epsilon(\iota S)$.

In other words, σ^* is a simplicial map on the cell decomposition of $\mathcal{T}_{g',b,2\Gamma}^{\text{m.p.}}(DF)$, and we need only concern ourselves with simplices corresponding to geometric spines with the above symmetry property—the set fixed by σ^* does not intersect any other simplices.

Within such a fixed simplex, the set T_σ is precisely those points with coordinates α_e fixed by ι, that is, satisfying $\alpha_{\iota(e)} = \alpha_e$ for all edges $e \in E$. This is a trivial condition on the set of edges $E_{\text{fix}} \subset E$ whose vertices are sent to each other under ι, but a codimension 1 condition for each pair of edges $e, e' \in E_{\text{swap}} \subset E$ which are interchanged by ι.

Now suppose that $(S, \tilde{\epsilon})$ and α_e are the geometric spine and associated coordinates for a point in T_σ. We know by the work done in chapter 4 that for any initial assignment of lengths $r_i(0)$ between the vertices and marked points, the combinatorial Ricci flow will converge to the unique solution with flat curvature at the vertices, and we obtain a metric which is unique up to diffeomorphism of

DF.

However, we also know that any point in T_σ has a spine and coordinates with the symmetry discussed above. So let us choose initial lengths $r_i(0)$ such that $r_i(0) = r_j(0)$ if $\iota v_i = v_j$. As a consequence of this choice, by the symmetry of the spine—and thus the gluing of the hyperbolic pieces to form a metric—we see that if $\iota v_i = v_j$, $K_i(0) = K_j(0)$. But by the definition of the Ricci flow, the evolution of r_i at time t only depends on $K_i(t)$ and $r_i(t)$, so this in turn implies that if $\iota v_i = v_j$, $r_i(t) = r_j(t)$ for all t. Finally, we see that the limiting solution—the assignment R_i of lengths which furnishes us with the required metric—satisfies $R_i = R_j$ for i, j with $\iota v_i = v_j$.

Once we have this symmetric model metric, act on it in the same way as before with a diffeomorphic extension $\eta : DF \to DF$ of ϵ, to obtain a representative of the required Teichmüller class.

Let us examine the metric we have constructed more closely. For each pair of edges $e, e' \in E_{\text{swap}}$ which are interchanged by ι, a hyperbolic triangle on e with opposite vertex p will have a reflection on e' incident at the same marked point. On the other hand, for each edge $e \in E_{\text{fix}}$ fixed by ι, a hyperbolic triangle on e with opposite vertex p will be equilateral, since the edge lengths framing the angle α_e must be equal. Hence the geodesic arc from p perpendicular to e splits the triangle into two pieces (mutual reflections), and hits p' (possibly equal to p), the vertex opposite e in the other direction.

It is clear that along with the cone points, these geodesic arcs for $e \in E_{\text{fix}}$ form the fixed set of the reflection σ on DF. (The compatibility of ϵ, ι and σ ensures that η doesn't affect this property.) Hence, by the definition of σ, they are (topologically) isotopic to the set of closed curves C, and thus disconnect DF into two pieces, each diffeomorphic to F. Furthermore, the boundary of each piece is polygonally geodesic, and by symmetry, the angle sum at the half-cone points is Γ

for each side. So the restriction of our metric on DF to either side is a hyperbolic metric with polygonal geodesic boundary and total half-cone angle Γ on F. In fact, since the metric is fixed by σ, which restricts to a diffeomorphism from F_+ to F_-, the metric is the same on both sides. Hence we have a well-defined point in $\mathcal{T}^{\text{p.g.b.}}_{g,b,\Gamma}$—this is $\phi^{-1}((S,\tilde{\epsilon}),\alpha_e)$.

Clearly $\phi \circ \phi^{-1}$ is the identity: ϕ is well-defined, so for a suitable metric in T_σ, $\phi(\phi^{-1}((S,\tilde{\epsilon}),\alpha_e))$ is represented by the double of the metric we obtain by the process above. But by the construction, this is precisely the metric characterised by $(S,\tilde{\epsilon})$ and α_e.

On the other hand, since ϕ is well-defined and the reconstruction process is unique, $\phi^{-1} \circ \phi$ is also the identity.

Furthermore, it is easy to see that both ϕ and ϕ^{-1} are continuous. Hence ϕ is a homeomorphism whose existence proves the theorem. $\square$

Chapter 7

Closed surfaces: an open question

Ideally, we would like to be able to use methods similar to those in previous cases to obtain a simplicial decomposition of the Teichmüller and moduli spaces of a closed surface—not least because the existence of such a decomposition is an open question. Unfortunately, for a closed surface there is no immediately apparent distinguished set B (no marked points or boundary), and consequently no obviously analogous way to construct a spine as an intrinsic feature of the metric.

Even more worryingly, a heuristic dimension count using the Euler characteristic suggests that the thing we are looking for should not exist. The Euler characteristic of a surface of genus g is $2-2g$, while the Euler characteristic of a graph is $v-e+f$, where v, e, and f are the numbers of vertices, edges and faces. In the case of a closed surface with m marked points, $f=m$, and, as usual, a spine in a top-dimensional simplex will have only trivalent vertices, so $v=\frac{2}{3}e$. Hence such a spine satisfies $2-2g=-\frac{1}{3}e+m$, and has $6g-6+3m$ edges. This is what we expect, since there is one linear constraint for each marked point (the sum of the incident angles), and the dimension of the Teichmüller space is $6g-6+2m$.

All of which suggests that for a closed surface, we need to construct a graph with no faces! Clearly that goal is hopeless.

However, we can make a few comments concerning the first step towards a possible solution.

7.1 A fibration

There is a natural projection from the space of hyperbolic metrics with one marked point on a surface to the space of hyperbolic metrics with no marking, whose action is simply to "forget the marked point." This gives us a projection between the corresponding Teichmüller spaces, so one way to try to investigate the case of no marking would be to study how the fibre of this projection over a metric on a closed surface intersects the cells of the Teichmüller space of metrics with one marked point.

At least locally, this process gives us sets of coordinates for the Teichmüller space of the closed surface, since it turns out that the fibre above a generic metric h must make a transverse intersection with a cell of codimension two upstairs. There will be a neighbourhood about h within which perturbed metrics h' intersect the same cell, so the simplicial coordinates on that cell are local coordinates on the Teichmüller space of the surface about the metric h.

Unfortunately, the good news ends there. The fibre above a metric will generally intersect many cells of codimension two (or higher), and we have a choice of local sets of coordinates—which transform to each other in a complicated way. Even worse, these cellular intersections glue together in a singular fashion—essentially because the cells of codimension at least two sit in the whole complex in a messy, twisted way—so there is no way to produce a simplicial complex from this process.

In short, there seems to be no way to parlay this approach into a satisfactory decompositon.

7.2 An alternative set of coordinates

Let us reconsider the case of a surface with marked points. Recall that we constructed the spinal triangulation on the surface and took as coordinates the angles incident at the marked points. We did so because they gave us nice simplicial coordinates on the cells of our decomposition, but they are not the unique set of coordinates we could have chosen.

Definition 7.1. Let the *perpendicular coordinate* a_e corresponding to an edge $e \in E$ of a spine be the length of the geodesic arc which is perpendicular to e and joins the (not necessarily unique) pair of marked points from which any point in e is equidistant.

Given the perpendicular coordinates corresponding to a spine, we can certainly reconstruct the metric uniquely. In fact the proof of the reconstuction is more-or-less trivial—but the coordinates we obtain are not simplicial, and the natural constraints we want to impose become ungainly.

7.3 An intrinsic spine?

We close with a speculative conjecture about how we might proceed.

Definition 7.2. Suppose F is a closed surface equipped with a hyperbolic metric. Define a function $f : F \to \mathbb{N}$ assigning to a point x the number of distinct shortest geodesic arcs (of positive length) from x to itself. Let B be the set of points b for which $f(b) \geq 3$.

This definition of B is intrinsic to the metric in the sense we require, and since the condition on f is a codimension two condition, B will be a set of discrete points.

Hence, exactly as in the case of a surface with marked points, we can construct an intrinsic spine, using the set B as our "decoration." The problem with this approach is that we end up with too many coordinates (edges) in the spine, and it is not clear how to express the extra constraints we have (that the points in B have three or more shortest geodesic arcs back to themselves) in a way which is compatible with the simplicial coordinates.

However, the constraints do seem to be at least partially compatible with perpendicular coordinates. In the generic case, a point $b \in B$ will have three shortest geodesic arcs to itself. Suppose that the edges perpendicular to these arcs all appear in the spine. Then the constraint on b is that the perpendicular coordinates corresponding to these three edges are equal—a very simple condition.

Furthermore, as the metric on F varies, we can imagine the spine type changing. One way in which it could change is that two points in B, each with three shortest geodesic arcs to itself, become closer together, eventually merging (instantaneously) in a "codimension one" metric having one point with four unique geodesic arcs to itself. This corresponds to the perpendicular coordinate of the edge equidistant from these two points tending to zero—again, a simple condition.

The challenge, then, would be to identify a natural set of coordinates. In a perfect world, such a description would combine the simplicial nature of the angular coordinates with the ease of expressing the constraints afforded by the perpendicular coordinates. If such a thing could be unearthed it would reveal a beautiful, hitherto unseen structure woven into the Teichmüller and moduli spaces of a closed surface.

Bibliography

[1] B. H. Bowditch and D. B. A. Epstein. Natural triangulations associated to a surface. *Topology*, 27(1):91–117, 1988.

[2] B. Chow. The Ricci flow on the 2-sphere. *Journal of Differential Geometry*, 33(2):325–334, 1991.

[3] B. Chow and F. Luo. Combinatorial Ricci flows on surfaces. *Journal of Differential Geometry*, 63:97–129, 2003. Preprint: math.DG/0211256.

[4] R. S. Hamilton. Three manifolds of positive Ricci curvature. *Journal of Differential Geometry*, 17:255–306, 1982.

[5] R. S. Hamilton. The Ricci flow on surfaces. In *Mathematics and general relativity (Santa Cruz, CA, 1986)*, volume 71 of *Contemporary Mathematics*, pages 237–262. American Mathematical Society, 1988.

[6] J. L. Harer. The virtual cohomological dimension of the mapping class groups of an orientable surface. *Inventiones Mathematicae*, 84(1):157–176, 1986.

[7] R. S. Palais and S. Smale. A generalized Morse theory. *Bulletin of the American Mathematical Society*, 70:165–172, 1965.

[8] G. Perelman. The entropy formula for the Ricci flow and its geometric applications. Feb 2002. Preprint: math.DG/0211159.

[9] G. Perelman. Finite extinction time for the solutions to the Ricci flow on certain three-manifolds. Jul 2003. Preprint: math.DG/0307245.

[10] G. Perelman. Ricci flow with surgery on three-manifolds. Mar 2003. Preprint: math.DG/0303109.

[11] W. P. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bulletin of the American Mathematical Society*, 6:357–381, 1982.

[12] W. P. Thurston. *Three-dimensional geometry and topology (volume 1)*, volume 35 of *Princeton Mathematical Series*. Princeton University Press, 1997.