



2460-1

Advanced School and Workshop in Real and Complex Dynamics

20 - 31 May 2013

Conformal Geometry and Dynamics of Quadratic Polynomials

LYUBICH Mikhail

State University of New York At Stony Brook NY 11794-3660 Stony Brook USA

Conformal Geometry and Dynamics of Quadratic Polynomials

Mikhail Lyubich

Abstract.

Contents

Chapter 0. Introduction 1. Preface 2. Background	5 5 8
Part 1. Conformal and quasiconformal geometry	17
 Chapter 1. Conformal geometry Riemann surfaces Holomorphic proper maps and branched coverings Extremal length and width Principles of the hyperbolic metric Uniformization Theorem Carathéodory boundary Appendix: Potential theory 	19 19 40 44 50 59 62 67 67
 Chapter 2. Quasiconformal geometry 11. Analytic definition and regularity properties 12. Geometric definitions 13. Further important properties of qc maps 14. Measurable Riemann Mapping Theorem 15. One-dimensional qs maps, quasicircles and qc welding 16. Removability 	79 79 86 90 94 101 104
 Chapter 3. Elements of Teichmüller theory 17. Holomorphic motions 18. Moduli and Teichmüller spaces of punctured spheres 19. Bers Embedding 	107 107 109 113
Part 2. Complex quadratic family	117
 Chapter 4. Dynamical plane 20. Glossary of Dynamics 21. Holomorphic dynamics: basic objects 22. Periodic motions 23. Hyperbolic maps 24. Parabolic maps 25. Misiurewicz maps 26. Quasiconformal deformations 27. Quadratic-like maps: first glance 28. Appendix: Expanding circle maps 	119 119 123 133 141 146 146 146 150 153

CONTENTS

Chapter 5. Remarkable functional equations	159
29. Linearizing coordinate in the attracting case	159
30. Existence of Siegel disks	161
31. Global leaf of a repelling point	161
32. Superattracting points and Böttcher coordinates	163
33. Parabolic points and Écale-Voronin cylinders	167
Chapter 6. Parameter plane (the Mandelbrot set)	169
28. Definition and first properties	169
29. Connectivity of M	172
30. The Multiplier Theorem	176
31. Structural stability	179
32. Notes	187
Chapter 7. Combinatorics of external rays	189
39. Dynamical ray portraits	189
40. Limbs and wakes of the Mandelbrot set	197
41. Geodesic laminations	197
42. Limbs and wakes of the Mandelbrot set	197
43. Misiurewisz wakes and decorations	202
44. Topological model	202
45. Renormalization	202
Chapter 8. Thurston theory	205
45. Rigidity Theorem	205
46. Rigidity of superattracting polynomials	206
47. Hubbard tree determines f	206
48. Realization of critically periodic maps	207
Part 3. Little Mandelbrot copies	211
Chapter 9. Quadratic-like maps	213
49. Straightening	213
50. External map	220
51. Uniqueness of the straightening	221
52. Weak q-l maps and Epstein class	224
Chapter 10. Quadratic-like families	225
42. Fully equipped families	225
43. QL families over the complex renormalization windows	234
44. Notes	234
Chapter 11. Yoccoz Puzzle	235
46. Combinatorics of the puzzle	235
47. Local connectivity of non-renormalizable Julia sets	240
48. Local connectivity of of M at non-renormalizable points	240
Part 4. Hints and comments to the exersices	241
49. Index	247

4

CONTENTS

Bibliography

5

251

CHAPTER 0

Introduction

1. Preface

In the last quarter of 20th century the complex and real quadratic family

 $f_c: z \mapsto z^2 + c$

was recongnized as a very rich and representative model of chaotic dynamics. In the complex plane it exhibits fractal sets of amazing beauty. On the real line, it contains regular and stochastic maps intertwined in an intricate fashion. It also has remarkable universality properties: its small pieces (if to look at the right place) look exactly the same as the whole family. Interplay between real and complex dynamics provide us with deep insights into both. This interplay eventually led to a complete picture of dynamics in the real quadratic family and a nearly compete picture in the complex family. In this book we attempt to present this picture beginning from scratch and supplying all needed background (beyond the basic graduate education).

Part 1 of the book contains a necessary background in conformal and quasiconformal geometry with elements of the Teichmüller theory. The main analytical



FIGURE 1. Mandelbrot set. It encodes in one picture all beauty and subtlety of the complex quadratic family.



FIGURE 2. Baby M-set.

tools of holomorphic dynamics are collected here in the form suitable for dynamical applications: principles of hyperbolic metric and extremal length, the classical Uniformization Theorem and Measurable Riemann Mapping, and various versions of the λ -lemma.

Part 2 begins with the classical Fatou-Julia theory (adatpted to the quadratic family): basic properties of Julia sets, classification of periodic motions, important special classes of maps. Then Sullivan's No Wandering Domains Theorem is proved, which completes description of the dynamics on the Fatou set. We proceed with a discussion of remarkable functional equations associated with the local dynamics (which were one of the original motivations for the classical theory).

In Chapter 2 we pass to the parameter plane, introducing the Mandelbrot set and proving two fundamental theorem about it: Connectivity and the Multiplier Theorem (by Douady and Hubbard). We proceed with the Structural Stability theory (by Mané-Sad-Sullivan and the author). We conclude this chapter with a proof of the Milnor-Thurston Entropy Monotonicity Conjecture that gives the first illustration of the power of complex methods in real dynamics.

The next chapter (3) is dedicated to the combinatorial theory of the quadratic family developed by Douady and Hubbard. It provides us with explicit combinatorial models for Julia sets and the Mandelbrot set. The problem of local connectivity of Julia sets and the Mandelbrot set (MLC) arises naturally in this context.

In the final chapter of this part we introduce a powerful tool of contemporary holomorphic dynamics: Yoccoz puzzle, – and prove local connectivity of nonrenormalizable quadratic polynomials.



FIGURE 3. Real quadratic family as a model of chaos. This picture presents how the limit set of the orbit $\{f_c^n(0)\}_{n=0}^{\infty}$ bifurcates as the parameter c changes from 1/4 on the right to -2 on the left. Two types of regimes are intertwined in an intricate way. The gaps correspond to the regular regimes. The black regions correspond to the stochastic regimes (though of course there are many narrow invisible gaps therein). In the beginning (on the right) you can see the cascade of doubling bifurcations. This picture became symbolic for one-dimensional dynamics.

One of the most fascinatining features of the Mandelbrot set, clearly observed on computer pictures, is the presence of the little copies of itself ("baby M-sets"), which look almost identically with the original set (except for possible absence of the main cusp). The complex renormalization theory is designed to explain this phenomenon. In part 3 we develop the the Douady-Hubbard theory of quadraticlike maps and complex renormalization that justifies presence of the baby M-sets, and classify them. (The geometric theory that explains why these babies have a universal shape will be developed later.)

This will roughly constitute the 1st volume of the book.

In the 2nd volume we plan to prove the Feigenbaum-Coullet-Tresser Renormalization Conjecture (by Sullivan, McMullen and the author), density of hyperbolic maps in the real quadratic family, and the Regular and Stochastic Theorem asserting that almost any real quadratic map is either regular (i.e., has an attracting cycle

0. INTRODUCTION

that attracts almost all orbits) or stochastic (i.e., it has an absolutely continuous invariant measure that governs behavior of almost all orbits)– by the author. These results were obtained in 1990's but recently new insights, particularly by Avila and Kahn, led to much better understanding of the phenomena.

We plan to dedicate the 3d volume to recent advances in the MLC Conjecture (based on the work of Kahn and the author).

The last volume (if ever written) will be devoted to the measure-theoretic theory of Julia sets and the Mandelbrot set. We will discuss the measure of maximal entropy and conformal measures, Hausdorff dimension and Lebesgue measure of Julia sets and the Mandelbrot set. It would culminate with a construction of examples of Julia sets of positive area, by Buff and Cheritat, and more recently, by Avila and the author (not announced yet).

This book can be used in many ways:

• For a graduate class in conformal and quasiconformal geometry illustrated with dynamical examples. This would cover Part I with selected pieces from Part 2.

• As the first introduction to the one-dimensional dynamics, complex and real. Then the reader should begin with Part 2 consulting the background material from Part 1 as needed.

• As an introduction to advanced themes of one-dimensional dynamics for the reader who knows basics and intends to do research in this field. Such a reader can go through selected chapters of Part 2 proceeding fairly fast to Part 3.

• Of course, the book can also be used as a monograph, for reference.

2. Background

In this section we collect some standing (usually, standard) notations, definitions, and properties. It can be consulted as long as the corresponding objects and properties appear in the text.

2.1. Complex plane and its affiliates. As usually, $\mathbb{N} = \{0, 1, 2, ...\}$ stands for the additive semigroup of natural numbers (with the French convention that zero is natural);

 \mathbb{Z} is the group of integers,

 \mathbb{Z}_+ and \mathbb{Z}_- are the sets of positive and negative integers respectively;

 \mathbb{R} stands for the real line;

 $\mathbb C$ stands for the complex plane,

and $\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ stands for the Riemann sphere;

 S^2 is a topological sphere, i.e., a topological manifold homeomorphic $\hat{\mathbb{C}}$;

We let $\mathbb{C}_{\mathbb{R}} \approx \mathbb{R}^2$ be the *decomplexified* \mathbb{C} (i.e., \mathbb{C} viewed as 2D real vector space).

For $a \in \mathbb{C}$, r > 0, let

$$\mathbb{D}(a, r) = \{ z \in \mathbb{C} : |z - a| < r \}; \quad \overline{\mathbb{D}}(a, r) = \{ z \in \mathbb{C} : |z - a| \le r \}.$$

$$\begin{split} \mathbb{D}_r &\equiv \mathbb{D}(0,r), \text{ and let } \mathbb{D} \equiv \mathbb{D}_1 \text{ denote the unit disk.} \\ \text{Let } \mathbb{T}(a,r) &= \partial \mathbb{D}(a,r), \ \mathbb{T}_r \equiv \mathbb{T}(0,r), \text{ and let } \mathbb{T} \equiv \mathbb{T}_1 \text{ denote the unit circle;} \\ S^1 \text{ is a topological circle;} \\ \mathbb{C}^* &= \mathbb{C} \smallsetminus \{0\}, \ \mathbb{D}^* &= \mathbb{D} \smallsetminus \{0\}. \\ \mathbb{A}(r,R) &= \{z : r < |z| < R\} \text{ is an open round annulus; The notations } \mathbb{A}[r,R] \text{ or} \end{split}$$

2. BACKGROUND

 $\mathbb{A}(r, R]$ for the closed or semi-closed annuli are self-explanatory. The *equator* of $\mathbb{A}(r, R)$ is the curve $|z| = \sqrt{Rr}$.

$$\begin{split} \mathbb{H} &\equiv \mathbb{H}_{+} = \{z : \ \mathrm{Im} \, z > 0\} \text{ is the upper half plane,} \\ \mathbb{H}_{h} &= \{z : \ \mathrm{Im} \, z > h\}, \\ \mathbb{H}_{-} &= \{z : \ \mathrm{Im} \, z < 0\} \text{ is the lower half plane;} \\ \mathbb{P} &= \{z : 0 < \mathrm{Im} \, z < \pi\}; \end{split}$$

A plane domain is a domain in \mathbb{C} .

2.2. Point set topology.

2.2.1. Spaces and maps. In what follows, all topological spaces (except L^{∞} Banach spaces) are assumed to satisfy the Second Countability Axiom, i.e., they have a countable base of neighborhoods. We also assume that all topological spaces in question are metrizable, unless otherwise is explicitly said. (An important exception will be the space of quadratic-like germs.) Recall hat a compact space is metrizable iff it satisfies tha Second Countability Axiom (and iff it is separable), so in the compact case our two conventions exactly match.

• \overline{X} denotes the closure of a set X; int X denotes its interior.

A neighborhood of a point x will mean an open neighborhood, unless otherwise is explicitly said. For instance, a closed neighborhood $P \ni x$ means a closed set such that int $P \ni x$.

• • $U \Subset V$ means that U is compactly contained in V, i.e., \overline{U} is a compact set contained in V.

We say that a sequence {z_n} in a locally compact space X escapes to infinity, z_n → ∞, if for any compact subset K ⊂ X, only finitely many point z_n belong to K. In other words, z_n → ∞ in the one-point compactification X̂ = X ∪ {∞} of X.
Similarly, a sequence of subsets E_n ⊂ X escapes to infinity, if for any compact subset K ⊂ X, only finitely many sets E_n intersect K. In other words, E_n → ∞ uniformly in X̂.

• An embedding $i : X \hookrightarrow Y$ is a homeomorphism onto the image. An immersion $i : X \to Y$ is a continuous (not necessarily injective) map which is locally an embedding.

• A function $f : X \to \mathbb{R}$ is called *upper semicontinuous* at $z \in X$ if $f(z) \ge \limsup_{\zeta \to f(\zeta)} f(\zeta)$. It is called *lower semicontinuous* if $f(z) \le \liminf_{\zeta \to f(\zeta)} f(\zeta)$.

• A continuous map $f: X \to Y$ between two locally compact spaces is called *proper* if for any compact set $K \subset Y$, its full preimage $f^{-1}(K)$ is compact. Equivalently, $f(z) \to \infty$ in Y as $z \to \infty$ in X, or in other words, f extends continuously to a map $\hat{f}: \hat{X} \to \hat{Y}$ between the one-point compactifications of X and Y.

• Full preimages of points under a proper map will also be called its *fibers*. Note that discrete fibers are finite.

EXERCISE 0.1. An injective proper map $i: X \hookrightarrow Y$ is an embedding.

In this case we say that X is properly embedded into Y.

• A connected compact space is called *continuum*. (We hope it will not be confused with the set-theoretical notion of continuum.) A singleton is a "trivial continuum"

0. INTRODUCTION

• A compact space is called *perfect* if it does not have isolated points. Perfect sets are always uncountable. (In particular, non-trivial continua are uncountable.)

• A *Cantor set* is a totally disconnected perfect set. All Cantor sets are homeomorphic.

2.2.2. Paths and arcs. A path and a curve in X mean the same: a continuous map γ of an interval (of any type) or a circle to X. In the latter case we also refer to it as a *closed* curve or a *loop*. Abusing terminology, we often refer to the image of γ as a path/curve as well.

An arc is an *injective* path parametrized by a *closed* interval. A *simple closed* curve is the embedding of the circle into X, i.e., $\gamma : S^1 \hookrightarrow X$.

LEMMA 0.2. Any path parametrizd by a closed interval contains an arc with the same endpoints.

PROOF. It can be done by the *loop erasing* procedure. Let $\gamma : [0,1] \to X$. A subloop of γ is a restriction of γ to an interval [a, b] such that $\gamma(a) = \gamma(b)$. Any loop can be erased by "restricting" γ to the connected sum $[0, a] \sqcup_{a=b} [b, 1]$, and rescaling the latter to the unit size.

More generally, if $\gamma : \sqcup I_k \to X$ is a disjoint union of loops, we can simultaneously erase all of them (the devil staircase). If $\cup I_k$ is a maximal set such that all $\gamma | I_k$ are loops, then the loop erasing leads to an arc.

Note finally that such a maximal set exists by Zorn's Lemma.

Thus, path connectivity of a space X is equivalent to its arc connectivity.

2.2.3. Metrics. We use notation B(x, r) for a ball in a metric space of radius r centered at x (recall that in \mathbb{C} we also use notation $\mathbb{D}(x, r)$).

For two sets X and Y in a metric space with metric d, let

$$\operatorname{dist}(X,Y) = \inf_{x \in X, \, y \in Y} d(x,y).$$

If one of these sets is a singleton, say $X = \{x\}$, then we use notation dist(x, Y) for the distance from X to Y.

$$\operatorname{diam} X = \sup_{x,y \in X} d(x,y).$$

The Hausdorff distance between two subsets Y and Y is defined as follows:

(2.1)
$$\operatorname{dist}_{\mathrm{H}}(Y, Z) = \max(\sup_{y \in Y} d(y, Z), \, \sup_{z \in Z} d(Y, z))$$

Note that $\operatorname{dist}_{\mathrm{H}}(Y, Z) < \epsilon$ means that Z is contained in an ϵ -neighborhood of Y and the other way around.

EXERCISE 0.3. Let \mathcal{Z} be the space of closed subsets in a metric space Z.

- (i) Show that that dist_H defines a metric on \mathcal{Z} ;
- (ii) If Z is complete then \mathcal{Z} is complete as well;

(iii) If Z is compact then \mathcal{Z} is compact as well.

• Notation (X, Y) stands for the pair of spaces such that $X \supset Y$. A pair (X, a) of a space X and a "preferred point" $a \in X$ is called a *pointed* space.





FIGURE 4. A *comb* is a typical cause for non-local-connectivity. To establish non-local-connectivity of a subset set in \mathbb{R}^2 , *it chercher le peigne*. Notice that this comb is path connected.

• Notation $f: (X, Y) \to (X', Y')$ means a map $f: X \to X'$ such that $f(Y) \subset Y'$. In the particular case of pointed spaces $f: (X, a) \to (X', a')$ we thus have: f(a) = a'. Similar notations apply to triples, (X, Y, Z), where $X \supset Y \supset Z$, etc.

• For a manifold M, $T_x M$ stand for its tangent space at x, and TM stands for its tangent bundle.

• If M is Riemannian and X is a path connected subset of M, then one can induce the metric from M to X in two ways:

• The *chordal metric* is obtained by restricting to X the global metric on M.

The path metric is obtained by restricting to X the Riemannian metric on M and then defining the path distance between x and y as the infimum of the length of paths $\gamma \subset X$ connecting x to y (which could be infinite). For instance, one can induce the Euclidean metric on \mathbb{R}^2 to the circle \mathbb{T} in these two ways leading to the chordal and length metrics on \mathbb{T} . A more interesting example is obtained by inducing the Euclidean metric to a Jordan domain $D \subset \mathbb{R}^2$ with fractal boundary.

2.3. Local connectivity. This notion is crucial in Holomorphic Dynamics.

A topological space X is called *locally connected* ("lc") at a point $x \in X$ if x has a local base of connected neighborhoods. A space X is called locally connected if it is locally connected at every point.

EXERCISE 0.4. A space X is locally connected iff connected components of any open subset $U \subset X$ are open.

There is a convenient weaker notion of local connectivity: A space X is called weakly locally connected at a point $x \in X$ if any neighborhood $U \ni x$ contains a connected set P such that $x \in \text{int } P$.

EXERCISE 0.5. If a space is weakly locally connected (at every point) then it is locally connected. However, a space can be weakly locally connected at some point x without being locally connected at this point.

For a metric space X, a *lc modulus* at $x \in X$ is a function $\omega : \mathbb{R}_+ \to \mathbb{R}_+$, $\omega(\epsilon) \to 0$ as $\epsilon \to 0$, such that if $d(x, y) < \epsilon$ then there exist a connected set Y containing both x and y such that diam $Y < \omega(\epsilon)$. If h works for all points $x \in X$ then it is called lc modulus for X. EXERCISE 0.6. Show that a space X is weakly lc at some point x iff it has an lc modulus at this point. Conclude that a compact space X is locally connected if and only if it has a lc modulus.

EXERCISE 0.7. a) Show that curves are locally connected.

b) More generally, the image of a lc continuum is an lc continuum.

EXERCISE 0.8. Lc continuum $K \subset \mathbb{R}^n$ is arc connected.

A space X is called a *path/arc locally connected at a point* $x \in X$ if there exists a modulus of continuity $\omega(\epsilon)$ such that any point $y \in X$ which is ϵ -close to x can be connected to x with a path/arc of diameter less than $\omega(\epsilon)$. As usual, path/arc lc of the whole space means path/arc lc at every point.

EXERCISE 0.9. Properties of being path lc, arc lc, and locally connected are all equivalent.

EXERCISE 0.10. Let K be a compact subset of \mathbb{R}^n , and let $J = \partial K$. If J is locally connected then so is K.

Quite remarkably, local connectivity gives a characterization of curves:

THEOREM 0.11 (Hahn-Mazurkevich). Let X be a compact space. Then X is a lc continuum if and only if there is a space-filling curve $\gamma : [0,1] \to X$ ("Peano curve").

EXERCISE 0.12. Prove this theorem.

2.4. Plane topology. A Jordan curve γ is a simple closed curve in the 2-sphere S^2 .

JORDAN THEOREM. The complement of a Jordan curve γ consists of two components D_1 and D_2 with the common boundary γ .

These components are called *(open)* Jordan disks. Their closures $\overline{D}_i = D_i \cup \gamma$ are called *closed Jordan disks.*

EXERCISE 0.13. Show that any open Jordan disk is simply connected.

When S^2 is realized as one-point compactification of \mathbb{R}^2 , $S^2 = \mathbb{R}^2 \cup \{\infty\}$ and a Jordan curve γ lies in \mathbb{R}^2 , then one of the corresponding Jordan disks is bounded in \mathbb{R}^2 , while another contains ∞ . They are called the *inner* and *outer* Jordan disks respectively. If a point z belongs to the inner Jordan disk, we say that " γ goes around z" or " γ surrounds z".

A compact subset K in \mathbb{R}^2 is called *full* if $\mathbb{R}^2 \smallsetminus K$ is connected. (Intuitively, K "does not have holes"). A full non-trivial continuum is called a *hull*.

EXERCISE 0.14. a) If K is a hull, then any component of int K is simply connected. b) Let J be a compact subset of \mathbb{R}^2 , and let U_i be the bounded components of $\mathbb{R}^2 \setminus J$. Then $K := J \cup \bigcup U_i$ is a hull.

This procedure is called *filling in the holes of* J.

LEMMA 0.15. Let $K \subset \mathbb{R}^2$ be a lc hull, and let U be a component of int K. Take a $z \in K \setminus \overline{U}$ and connect it with an arc $\alpha \subset K$ to some point in \overline{U} . Let $\pi_{\alpha}(z) \equiv \pi_{U,\alpha}(z)$ be the first point of intersection of α with \overline{U} . Then $\pi_{\alpha}(z)$ is independent of α .

2. BACKGROUND

PROOF. Assume we have two arcs α_1 and α_2 in K connecting z to U such that $\zeta_1 := \pi_{\alpha_1}(z) \neq \pi_{\alpha_2}(z) =: \zeta_2$. Without loss of generality we can assume that the α_i end at ζ_i . Let $(u, \zeta_1]$ be the maximal subarc of α_1 that contains ζ_1 and does not cross α_2 , and let $\alpha'_i = [u, \zeta_i]$ be the closed subarcs of the α_i bounded by u and ζ_i (i = 1, 2). Then u is the only common point of these arcs.

Let us take some points $w_1, w_2 \in U$ that are ϵ -close to ζ_1, ζ_2 respectively. By Lemma 0.2, w_i can be connected to the respective point ζ_i by an arc $\gamma_i \subset K$ with diam $\gamma_i < \omega(\epsilon)$. So, for ϵ small enough, γ_1 is disjoint from $\delta'_2 := \alpha'_2 \cup \gamma_2$ and γ_2 is disjoint from $\delta'_1 := \alpha'_1 \cup \gamma_1$.

Applying Lemma 0.2 again, we can straighten the curves δ'_i to arcs $\delta_i \subset \delta'_i$ connecting u to w_i . Then u is the only one common point of these arcs as well.

Let us now connect w_1 to w_2 with an arc $\sigma \subset U$ disjoint from $\delta_1 \cup \delta_2$ (except for the endpoints).

The union of three arcs, δ_1 , δ_2 and σ , form a Jordan curve in K. Let D be the open Jordan disk bounded by this curve. Since K is full, $D \subset K$. Moreover, D intersects U, and hence $U \cup D$ is contained in a component of int K, so $D \subset U$. On the other hand D also intersects $\mathbb{R}^2 \setminus \overline{U}$ since $u \notin \overline{U}$ – contradiction.

So, under the above circumstances we have a well defined projection:

(2.2)
$$\pi_U: K \to U.$$

EXERCISE 0.16. The projection π_U is continuous and locally constant on $K \smallsetminus \overline{U}$.

Together with Exercise 0.7 b), this implies:

COROLLARY 0.17. If $K \subset \mathbb{R}^2$ is a lc hull and U is a component of int K then \overline{U} is a hull as well.

EXERCISE 0.18. Let K be a lc hull in \mathbb{R}^2 whose complement has infinitely many components U_i . Then diam $U_i \to 0$.

Further study of plane hulls will require analytic methods (see $\S6$).

An external neighborhood of a hull $K \subset \mathbb{C}$ is a set $U \smallsetminus K$ where U is a neighborhood of K.

2.5. Group actions. SL(2, R) is the group of 2×2 matrices over a ring R with determinant 1 (we will deal with $R = \mathbb{C}$, \mathbb{R} , or \mathbb{Z});

 $PSL(2, R) = SL(2, R) / \{\pm I\}$, where I is the unit matrix;

 $SO(2) \approx \mathbb{T}$ is the group of plane rotations;

 $PSO(2) = SO(2)/\{\pm I\}$ (this group is actually isomorphic to SO(2), but it is naturally embedded into $PSL(2,\mathbb{R})$ rather than $SL(2,\mathbb{R})$). Sim(2) is the group of *similarites* of \mathbb{R}^2 , i.e., compositions of rotations and scalar operators.

 $\mathbb{C}_{\mathbb{R}}$ is naturally embedded into \mathbb{C}^2 by $z \mapsto (z, \overline{z})$ (as the reflector for the antiholomorphic involution $(z, \zeta) \mapsto (\overline{\zeta}, \overline{z})$). Linear operators of \mathbb{C}^2 preserving $\mathbb{C}_{\mathbb{R}}$ and the area therein have the form

$$\left(\begin{array}{cc} \alpha & \beta \\ \bar{\beta} & \bar{\alpha} \end{array}\right), \quad \alpha, \beta \in \mathbb{C}, \ |\alpha|^2 - |\beta|^2 = 1$$

We let $\mathrm{SL}^{\#}(2,\mathbb{R})$ be the group of these operators (it is another representation of $\mathrm{SL}(2,\mathbb{R})$ in $\mathrm{SL}(2,\mathbb{C})$). Note that it acts on $\mathbb{C}_{\mathbb{R}} \subset \mathbb{C}^2$ by transformations $z \mapsto \alpha z + \beta \overline{z}$. An action of a discrete group Γ on a locally compact space X is said to be *properly*

An action of a discrete group 1 on a locarly compact space X is said to be property discontinuous if any two points $x, y \in X$ have neighborhoods $U \ni x, V \ni y$ such

0. INTRODUCTION

that $\gamma(U) \cap V = \emptyset$ for all but finitely many $\gamma \in \Gamma$. The quotient of X by a properly discontinuous group action is a *Hausdorff* locally compact space.

The stabilizer (or, the isotropy group) $\operatorname{Stab}(X)$ of a subset $Y \subset X$ is the subgroup $\{\gamma \in \Gamma : \gamma(Y) = Y\}$. A set Y called *completely invariant* under some subgroup $G \subset \Gamma$ if $G = \operatorname{Stab}(Y)$ and $\gamma(Y) \cap Y = \emptyset$ for any $\gamma \in \Gamma \setminus G$.

A group element γ is called *primitive* if it generates a maximal cyclic group.

Isometries of a metric space are also called *motions* (e.g., Euclidean motions, hyperbolic motions etc.).

2.6. Coverings. In this section we summarize for reader's convenience necessary background in the theory of covering spaces.

Let E and B be topological manifolds (maybe with boundary), where B is connected. A continuous map $p: E \to B$ is called a *covering of degree* $d \in \mathbb{N} \cup \{\infty\}$ (with *base* B and covering space E) if any point $b \in B$ has a neighborhood V such that

$$p^{-1}(V) = \bigsqcup_{i=1}^{d} U_i,$$

where each U_i is mapped homeomorphically onto V. The preimages $p^{-1}(b)$ are called *fibers* of the covering. The inverse maps $p_i^{-1}: V \to U_i$ are called the *local branches* of p^{-1} . Let us make a coupe of simple observations:

• A covering of degree one is a homeomorphism;

• Restriction of a covering $p: E \to B$ to any connected component of E is also a covering.

• If V is a domain in B, $U = p^{-1}(V)$ then the restriction $p : U \to V$ is also a covering.

Coverings $p : E \to B$ and $p' : E' \to B'$ are called *equivalent* if there exist homeomorphisms $\phi : E \to E'$ and $\psi : B \to B'$ such that $\psi \circ p = p' \circ \phi$. Similarly, one defines a lift of a homotopy.

Given a continuous map $f: X \to B$, a continuous map $f: X \to E$ is called a lift of f if $p \circ \tilde{f} = f$. Theory of covering spaces is based on the following fundamental property:

Path Lifting Property. Let γ be a path in B that begins at $b \in B$, and let $e \in p^{-1}(b)$. Then there is a unique lift $\tilde{\gamma}$ of γ (i.e., $p \circ \tilde{\gamma} = \gamma$) that begins at e. If γ is homotopic to γ' (rel the endpoints) then the corresponding lifts $\tilde{\gamma}$ and $\tilde{\gamma}'$ are also homotopic rel the endpoints.

It implies, in particular, that the induced homomorphism

$$p_*: \pi_1(E, e) \to \pi_1(B, b)$$

is injective; let $G = G_p \subset \pi_1(B, b)$ be its image If E is connected then replacing e with another point in the fiber replaces G by a conjugate subgroup. In this way, to any covering p (with connected E) we associate a subgroup of the fundamental group, up to conjugacy.

The Path Lifting Property implies a general

Lifting Criterion. A continuous map $f : (X, x) \to (B, b)$ admits a lift $\tilde{f} : (X, x) \to (E, e)$, where $e \in \mathbf{p}^{-1}(b)$, if and only if $f_*(\pi_1(X, x) \subset p_*(\pi_1(E, e)))$.

In particular, if E is simply connected, then all maps $f: X \to B$ are liftable. In what follows we assume that E is also connected.

2. BACKGROUND

A covering is called *Galois* if there is a group G acting freely and properly discontinuously on E whose orbits are fibers of the covering. In this case $B \approx E/G$. The group G is called the *group of deck transformations*, or the *covering group* for p.

Vice versa, if a group G acts freely and properly discontinuous on a manifold E then the quotient B := E/G is a manifold, and the natural projection $f : E \to B$ is a covering.

A covering $u : \mathcal{U} \to B$ is called Universal if the space E is connected and simply connected. This covering is Galois, with the fundamental group $\pi_1(T)$ acting by deck transformations. Any manifold has a unique Universal Covering up to equivalence.

REMARK 0.1. We supress the base point in the notation for the fundamental group, unless it can lead to confusion. (On most occasions, our statements are invariant under conjugacies (inner automorphisms) in the fundamental group, and hence are base point independent.)

The Universal Covering allows us to recover any covering p from the a subgroup $\Gamma \subset \pi_1(B)$ as $p: \mathcal{U}/\Gamma \to B$. Moreover, the universal covering factors through p since $q \circ p = u$, where $q: \mathcal{U} \to \mathcal{U}/\Gamma$. This provides us with a natural one-toone correspondence between classes of conjugate subgroups of $\pi_1(T)$ and classes of equivalent coverings $p: E \to B$. Moreover, the covering p is Galois if and only if the corresponding subgroup Γ is normal. In this case, the group of deck transformations of p is $\pi_1(T)/\Gamma$.

In particular, a simply connected manifold T does not admit any non-trivial coverings: any covering $p: S \to T$ with connected S is a homeomorphism. Putting this together with the above observations, we obtain an important statement which is often referred in anlysis as the Monodromy Theorem:

If $p: E \to B$ is a covering and $V \subset B$ is a simply connected domain, then $p^{-1}(V)$ is a disjoint union of domains U_i , $i = 1, \ldots, d$, such that each restriction $p: U_i \to V$ is a homeomorphim. Thus, on any simply connected domain there exist d well defined inverse branches $p_i^{-1}: V \to U_i$.

Given a base point $b \in B$, there exists a natural monodromy action of the fundamental group $G = \pi_1^{-1}(B, b)$ on the fiber $F := \pi^{-1}(b)$. Namely, let an element $A \in G$ is represented by a loop α in T based at b. Lift α to a path $\tilde{\alpha}$ in E based at some $e \in F$. Then A(e) is defined as the endpoint of $\tilde{\alpha}$. The stabilizer G of this action is the subgoup Γ corresponding to p (well defined up to conjugacy), which gives yet another viepoint on the relation between coverings over B and subgoups of G.

EXERCISE 0.19. Let $p: E \to B$ be a covering of degree d. Then there exists a Galois covering $q: L \to B$ of degree at most d! that factors through p, i.e., $q = p \circ r$ (where $r: L \to E$ is also (automatically) a Galois covering).

One says that a connected submanifold $V \subset B$ (maybe with boundary) is essential in T if the induced (by the emebdding) homomorphism $\pi_1(V) \to \pi_1(B)$ is injective. In other words, any non-trivial loop in V remains non-trivial in B.

0. INTRODUCTION

PROPOSITION 0.20. Let $u : \mathcal{U} \to B$ be the universal covering. If $V \subset B$ is an essential submanifold then any component \hat{V} of $u^{-1}(V)$ is simply connected. Moreover, the stabilizer Γ of \hat{V} in the group G of deck transformations is the covering group for the restriction $p : \hat{V} \to V$ (and thus, Γ is isomorphic to $\pi_1(V)$).

PROOF. By the above observations, restriction $u : \hat{V} \to V$ is a covering, so we only need to show (for the first assertion) that \hat{V} is simply connected. But otherwise, there would be a non-trivial loop $\tilde{\alpha}$ in \hat{V} . Then the loop $\alpha = p_*(\tilde{\alpha})$ would be non-trivial in V (since p_* is injective) but trivial in B (since $\tilde{\alpha}$ is trivial in \mathcal{U}).

Since $p^{-1}(V)$ is invariant under G, each deck transformation $\gamma : \mathcal{U} \to \mathcal{U}$ permutes the components of $p^{-1}(V)$. Hence for any $\gamma \in G$, \hat{V} is either invariant under γ or else $\gamma(\hat{V}) \cap \hat{V} = \emptyset$. It follows that the stabilizer Γ of \hat{V} acts transitively on the fibers of $p \mid \hat{V}$, and the conclusion follows.

COROLLARY 0.21. Let γ be an essential simple closed curve in B. Then each lift $\tilde{\gamma}$ to the universal covering \mathcal{U} is a topological line whose stabilizer is an infinite cyclic group. Different lifts have conjugate stabilizers.

Thus, to each (oriented) simple closed curve in B we can associate a conjugacy class in the fundamental group (the generators of the above stabilizers).

EXERCISE 0.22. There is a natural one-to-one correspondence between classes of freely homotopic (oriented) closed curves (not necessarily simple) and conjugacy classes in $G = \pi_1(B)$.

LEMMA 0.23. Let V be an essential submanifold in B. Then there is a covering $q: E \to B$ with $\pi_1(E) = \pi_1(V)$ and such that one of the components U of $q^{-1}(V)$ projects homeomorphically onto V.

PROOF. In the notation of Lemma 0.20, let $E = \mathcal{U}/\Gamma$, $U = \hat{V}/\Gamma$.

Informally speaking, we unwide all the loops in B except those that are essentailly confined to V.

COROLLARY 0.24. Let $\gamma \subset B$ be an essential simple closed curve. Then there is a covering space E with $\pi_1(E) \approx \mathbb{Z}$ containing a simple closed curve $\hat{\gamma}$ that projects homeomorphically onto γ .

Part 1

Conformal and quasiconformal geometry

CHAPTER 1

Conformal geometry

1. Riemann surfaces

1.1. Topological surfaces.

1.1.1. Definitions and examples.

DEFINITION 1.1. A (topological) surface S (without boundary) is a two-dimensional topological manifold with countable base. It means that S is a topological space with a countable base and any $z \in S$ has a neighborhood $U \ni z$ homeomorphic to an open subset V of \mathbb{R}^2 . The corresponding homeomorphism $\phi : U \to V$ is called a (topological) local chart on S. Such a local chart assigns to any point $z \in U$ its local coordinates $(x, y) = \phi(z) \in \mathbb{R}^2$.

A family of local charts whose domains cover S is called a *topological atlas* on S.

Given two local charts $\phi: U \to V$ and $\tilde{\phi}: \tilde{U} \to \tilde{V}$, the composition

$$\tilde{\phi} \circ \phi^{-1} : \phi(U \cap \tilde{U}) \to \tilde{\phi}(U \cap \tilde{U})$$

is called the *transition map* from one chart to the other.

A surface is called *orientable* if it admits an atlas with orientation preserving transition maps. Such a surface can be oriented in exactly two ways. In what follows we will only deal with orientable (and naturally oriented) surfaces.

Unless otherwise is explicitly said, we will assume that the surfaces under consideration are *connected*. The simplest (and most important for us) surfaces are:

• The whole plane \mathbb{R}^2 (homeomorphic to the open unit disk $\mathbb{D} \subset \mathbb{R}^2$).

• The unit sphere S^2 in \mathbb{R}^3 (homeomorphic via the stereographic projection to the one-point compactification of the plane); it is also called a "closed surface of genus 0" (in this context "closed" means "compact without boundary").

• A cylinder or topological annulus $C(a, b) = \mathbb{T} \times (a, b)$, where $-\infty \leq a < b \leq +\infty$. It can also be represented as the quotient of the strip $P(a, b) = \mathbb{R} \times (a, b)$ modulo the cyclic group of translations $z \mapsto z + 2\pi n$, $n \in \mathbb{Z}$. All the cylinders C(a, b) are homeomorphic to any annulus $\mathbb{A}(r, R)$, to the punctured disk \mathbb{D}^* and to the punctured plane \mathbb{C}^*).

• The torus $\mathbb{T}^2 = \mathbb{T} \times \mathbb{T}$, also called a "closed surface of genus 1". It can also be represented as the quotient of \mathbb{R}^2 modulo the action of a rank 2 abelian group $z \mapsto z + \alpha m + \beta n$, $(m, n) \in \mathbb{Z}^2$, where α and β is an arbitrary basis in \mathbb{R}^2 .

It is intuitively obvious that (up to a homeomorphism) there are only two simply connected surfaces: the plane and the sphere.

If we have a certain standard surface S (say, the unit disk or the unit sphere), a "topological S" (say, a "topological disk" or a "topological sphere") refers to a surface homeomorphic to the standard one.

1. CONFORMAL GEOMETRY

One can also consider surfaces with boundary. The local model of a surface near a boundary point is given by a relative neighborhood of a point (x, 0) in the closed upper half-plane $\overline{\mathbb{H}}$. The orientation of a surface naturally induces an orientation of its boundary (locally corresponding to the positively oriented real line).

For instance, we can consider cylinders with boundary: $C[a, b] = \mathbb{T} \times [a, b]$ or $C[a, b) = \mathbb{T} \times [a, b)$. They will be still called "cylinders" or "topological annuli". Cylinders C(a, b) without boundary will be also called "open", while cylinders C[a, b] will be called "closed" (according to the type of the interval involved).

Cylinders (with or without boundary) are the only topological surfaces whose fundamental group is \mathbb{Z} .

1.1.2. Ends and ideal circles. A non-compact domain $E \subset S$ bounded by a simple closed curve γ in S is called *end-region* (and the same term is applied to the closure of E). An *end-nest* $\{E_n\}$ of S is a nest $E_0 \supset E_1 \supset \ldots$ of end-regions escaping to infinity in S. Two end-nests, $\{E_n\}$ and $\{E'_n\}$, are called equivalent if any E_n is contained in some E'_m , and the other way around. An *end e* of S is a class of equivalent end-nests. We let $\mathcal{E}(S)$ be the set of ends of S.

An end is called *tame* if eventually all the end-regions (in some and hence in any end-nest) are cylinders. Any of these cylinders uniquely determines the corresponding end.

We can compactify S by attaching one point ∞_e to each end $e \in \mathcal{E}(S)$ and declaring the representing nest-regions E_n the base of neighborhoods of ∞_e . We call it *one-point-per-end compactification of* S.

EXERCISE 1.2. Let D be a domain on the sphere S^2 such that each component of $S^2 \setminus D$ is a full. Describe the ends of D.

A tame end can be also compactified in a different way by attaching a topological circle at infinity called the *ideal circle* at infinity.

REMARK 1.1. On the topological (or smooth) level, attaching one point or the ideal circle to a tame end have equal footing: the end does not have a preference. However, a *conformal end* knows exactly what should be attached to it: see Proposition 1.19.

We say that a map $f: S \to S'$ properly maps an end e of S to an end e' of S' if $f(z) \to \infty_{e'}$ as $z \to \infty_e$. Such a map extends continuously to the one-point compactification of the ends by letting $\infty_e \mapsto \infty_{e'}$.

1.1.3. New surfaces from old ones. There are two basic ways of building new surfaces out of old ones: making holes and gluing their boundaries. Of course, any open subset of a surface is also a surface. In particular, one can make a (closed) hole in a surface, that is, remove a closed Jordan disk. A topologically equivalent operation is to make a *puncture* in a surface. By removing an open Jordan disk (open hole) we obtain a surface with boundary.

If we have two open holes (on a single surface or two different surfaces S_i) bounded by Jordan curves γ_i , we can glue these boundaries together by means of an orientation reversing homeomorphism $h : \gamma_1 \to \gamma_2$. (It can be also thought as attaching a cylinder to these curves.) We denote this operation by $S_1 \sqcup_h S_2$. For instance, by gluing together two closed disks we obtain a topological sphere: $\mathbb{D} \sqcup_h \mathbb{D} \approx S^2$.

Combining the above operations, we obtain operations of taking *connected sums* and attaching a *handle*. To take a connected sum of two surfaces S_1 and S_2 , make

an open hole in each of them and glue together the boundaris of these holes. To attach a handle to a surface S, make two open holes in it and glue together their boundaries.

If we attach a handle to a sphere, we obtain a topological torus. If we attach g handles to a sphere, we obtain a *closed surface of genus* g'. A Fundamental Theorem of Topology asserts that any closed orientable surface is homeomorphic to one of those. Thus closed orientable surfaces are topologically classified by a single number $g \in \mathbb{N}$, its genus.

One says that a surface S (with or without boundary) has a *finite topological* type if its fundamental group $\pi(S)$ is finitely generated (e.g., any compact surface is of finite type). It turns out that it is equivalent to saying that S is homeomorpic to a closed surface with finitely many open or closed holes. Clearly such a surface admits a decomposition

$$S = K \sqcup_{h_i} C_i \,,$$

where K is a compact surface and the C_i are cylinders. The set $K = K_S$ is called the *compact core* of S. Note that it is obviously a deformation retract for S. The cylinders C_i represent the ends of S: all of them are tame in this case.

1.1.4. Euler characteristic. Let S be a compact surface (with or without boundary) Its Euler characteristic is defined as

$$\chi(S) = f - e + v,$$

where f, e and v are respectively the numbers of faces, edges and vertices in any triangulation of S.

The Euler characteristic is obviously additive:

$$\chi(S_1 \sqcup_h S_2) = \chi(S_1) + \chi(S_2).$$

Since the cylinder $\mathbb{T} \times [0,1]$ has zero Euler characteristic, $\chi(\hat{S}) = \chi(K_S)$ for a surface S of finite type. We can use this as a definition of $\chi(S)$ in this case.

Making a hole in a surface drops its Euler characteristic by one; attaching a handle does not change it. Hence $\chi(S) = 2 - 2g - n$ for a surface of genus g with n holes.

Note that the above list of simplest surfaces is the full list of sufaces of finite type without boundary with *non-negative* Euler characteristic:

$$\chi(\mathbb{R}^2) = 1, \quad \chi(S^2) = 2, \quad \chi(\mathbb{T} \times (0, 1)) = \chi(\mathbb{T}^2) = 0.$$

1.1.5. Marking. A surface S can be marked with an extra topological data. It can be either several marked points $a_i \in S$, or several closed curves $\gamma_i \subset S$ up to homotopy (usually but not always they form a basis of $\pi_1(S)$), or a parametrization of several boundary components $\Gamma_i \subset \partial S$, $\phi_i : \mathbb{T} \to \Gamma_i$.

The marked objects may or may not be distinguished (for instance, two marked points or the generators of π_1 may be differently colored). Accordingly, the marking is called *colored* or *uncolored*.

A homeomorphism $h: S \to \tilde{S}$ between marked surfaces should respect the marked data: marked points should go to the corresponding points $(h(a_i) = \tilde{a}_i)$, marked curves γ_i should go to the corresponding curves $\tilde{\gamma}_i$ up to homotopy $(h(\gamma_i) \simeq \tilde{\gamma}_i)$, and the boundary parametrizations should be naturally related $(h \circ \phi_i = \tilde{\phi}_i)$.

1. CONFORMAL GEOMETRY

1.2. Analytic and geometric structures on surfaces. Rough topological structure can be refined by requiring that the transition maps belong to a certain "structural pseudo-group", which often means: "have certain regularity". For example, a *smooth* structure on S is given by a family of local charts $\phi_i : U_i \to V_i$ such that all the transition maps are smooth (with a prescribed order of smoothness). A surface endowed with a smooth structure is naturally called a *smooth surface*. A local chart $\phi : U \to V$ smoothly related to the charts ϕ_i (i.e., with smooth transition maps) is referred to as a "smooth local chart". A family of smooth local charts covering S is called a "smooth atlas" on S. A smooth structure comes together with affiliated notions of smooth functions, maps and diffeomorphisms.

There is a smooth version of the connected sum operation in which the boundary curves are assumed to be smooth and the boundary gluing map h is assumed to be an orientation reversing diffeomorphism. To get a feel for it, we suggest the reader to do the following exercise:

EXERCISE 1.3. Consider two copies D_1 and D_2 of the closed unit disk $\mathbb{D} \subset \mathbb{R}^2$. Glue them together by means of a diffeomorphism $h: \partial D_1 \to \partial D_2$ of the boundary circles. You obtain a topological sphere S^2 . Show that it can be endowed with a unique smooth structure compatible with the smooth structures on D_1 and D_2 (that is, such that the tautological embeddings $D_i \to S^2$ are smooth). The boundary circles ∂D_i become smooth Jordan curves on this smooth sphere. Show that this sphere is diffeomorphic to the standard "round sphere" in \mathbb{R}^3 .

 ${\it Real\ analytic\ structures\ would\ be\ the\ next\ natural\ refinement\ of\ smooth\ structures.}$

If \mathbb{R}^2 is considered as the complex plane \mathbb{C} with z = x + iy, then we can talk about *complex analytic* \equiv *holomorphic* transition maps and corresponding complex analytic structures and surfaces. Such surfaces are known under a special name of *Riemann surfaces*. A holomorphic diffeomorphism between two Riemann surfaces is often called an *isomorphism*. Accordingly a holomorphic diffeomorphism of a Riemann surface onto itself is called its *automorphism*.

Connected sum operation still works in the category of Riemann surfaces. In its simplest version the boundary curves and the gluing diffeomorphism should be taken real analytic. Here is a representative statement:

EXERCISE 1.4. Assume that in Exercise $1.3 \mathbb{R}^2 \equiv \mathbb{C}$ and the gluing diffeomorphism h is real analytic. Then S^2 can be supplied with a unique complex analytic structure compatible with the complex analytic structure on the disks $D_i \subset \mathbb{C}$. The boundary circles ∂D_i become real analytic Jordan curves on this "Riemann sphere".

More generally, we can attach handles to the sphere by means of real analytic boundary map, and obtain an example of a Riemann surface of genus g. It is remarkable that, in fact, it can be done with only smooth gluing map, or even with a singular map of a certain class. This operation (with a singular gluing map) has very important applications in Teichmiller theory, theory of Kleinian groups and dynamics (see ??).

If \mathbb{R}^2 is supplied with the standard Euclidean metric, then we can consider *conformal* transition maps, i.e., diffeomorphisms preserving angles between curves. The first thing students usually learn in complex analysis is that the class of orientation preserving conformal maps coincides (in dimension 2!) with the class of invertible complex analytic maps. Thus the notion of a *conformal structure* on an

oriented surface is equivalent to the notion of a complex analytic structure (though it is worthwhile to keep in mind their conceptual difference: one comes from geometry, the other comes from analysis).

One can go further to *projective*, *affine*, *Euclidean/flat* or *hyperbolic* structures. We will specify this discussion in a due course.

One can also go in the opposite direction and consider *rough structures* on a topological surface whose structural pseudo-group is bigger than the pseudo-group of diffeomorphisms, e.g., "bi-Lipschitz structurs". Even rougher, *quasi-conformal*, structures will play an important role in our discussion.

To comfort a rigorously-minded reader, let us finish this brief excursion with a definition of a *pseudo-group* on \mathbb{R}^2 (in the generality adequate to the above discussion). It is a family of local homeomorphisms $f: U \to V$ between open subsets of \mathbb{R}^2 (where the subsets depend on f) which is closed under taking inverse maps and taking compositions (on the appropriately restricted domains). The above structures are related to the pseudo-groups of all local (orientation preserving) homeomorphisms, local diffeomorphisms, locally biholomorphic maps, local isometries (Euclidean or hyperbolic) etc.

1.3. Flat (affine) geometry. Consider the complex plane \mathbb{C} . Holomorphic automorphisms of \mathbb{C} are complex affine maps $A : z \mapsto az + b$, $a \in \mathbb{C}^*$, $b \in \mathbb{C}$. They form a group Aff(\mathbb{C}) acting freely bi-transitively on the plane: any pair of points can be moved in a unique way to any other pair of points. Moreover, it acts freely transitively on the tangent bundle of \mathbb{C} .

Thus the complex plane \mathbb{C} is endowed with the *affine structure* canonically affiliated with its complex analytic structure. Of course, the plane can be also endowed with a Euclidean metric $|z|^2$. However, this metric can be multiplied by any scalar t > 0, and there is no way to normalize it in terms of the complex structure only. All these Euclidean structures have the same group $\text{Euc}(\mathbb{C})$ of *Euclidean motions* $A : z \mapsto az + b$ with |a| = 1. This group acts transitively on the plane with the group of rotations $z \mapsto e^{2\pi i \theta} z$, $0 \le \theta < 1$, stabilizing the origin. Moreover, it acts freely transitively on the unit tangent bundle of \mathbb{C} (corresponding to any Euclidean structure).

The group Aff has very few discrete subgoups acting freely on \mathbb{C} : rank 1 cyclic group actions $z \mapsto z + an$, $n \in \mathbb{Z}$, and rank 2 cyclic group actions $z \mapsto an + bm$, $(m,n) \in \mathbb{Z}^2$, where (a,b) is an arbitrary basis in \mathbb{C} over \mathbb{R} . All rank 1 actions are conjugate by an affine transformation, so that the quotients modulo these actions are all isomorphic. Taking a = 1 we realize these quotients as the bi-infinite cylinder \mathbb{C}/\mathbb{Z} . It is isomorphic to the puncured plane \mathbb{C}^* by means of the exponential map $\mathbb{C}/\mathbb{Z} \to \mathbb{C}^*$, $z \mapsto e^{2\pi z}$. The quotients of rank 2 are all homeomorphic to the torus. However, they may represent different Riemann surfaces (see below 1.6.2).

Note that the above discrete groups preserve the Euclidean structures on \mathbb{C} . Hence these structures can be pushed down to the quotient Riemann surface. Moreover, now they can be canonically normalized: in the case of the cylinder we can normalize the lengths of the closed geodesics to be 1. In the case of the torus we can normalize its total area. Thus, complex tori and the bi-infinite cylinder are endowed with a canonical Euclidean structure. For this reason, they are called *flat*.

1. CONFORMAL GEOMETRY

1.4. Spherical (projective) geometry. Consider now the Riemann sphere $\overline{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$. Its bi-holomorphic automorphisms are *Möbius* transformations

$$\phi: z \mapsto \frac{az+b}{cz+d}; \quad \det \begin{pmatrix} a & b \\ c & d \end{pmatrix} \neq 0.$$

We will denote this Möbius group by $M\"ob(\hat{\mathbb{C}})$. It acts freely triply transitive on the sphere: any (ordered) triple of points (a, b, c) on the sphere can be moved by a unique Möbius transformation to any other triple (a', b', c').

EXERCISE 1.5. Show that topology of $PSL(2, \mathbb{R})$ and topology of uniform convergence on the sphere coincide. Given an $\epsilon > 0$, let us consider the set of Möbius transformations ϕ such that the triple $(\phi^{-1}(0, 1, \infty)$ is ϵ -separated in the spherical metric (i.e., the three points stay at least distance ϵ apart one from another). Show that this set is compact in Möb $(\hat{\mathbb{C}})$.

Note that the Riemann sphere is isomorphic to the complex projective line \mathbb{CP}^1 . For this reason Möbius transformations are also called *projective*. Algebraicly the Möbius group is isomorphic to the linear projective group $\mathrm{PSL}(2,\mathbb{C}) = \mathrm{SL}(2,\mathbb{C})/\{\pm I\}$ of 2×2 matrices M with det M = 1 modulo reflection $M \mapsto -M$.

Any Möbius transformation M has a fixed point $\alpha \in \hat{\mathbb{C}}$, i.e. $M(\alpha) = \alpha$. Hence there are no Riemann surfaces whose universal covering is $\bar{\mathbb{C}}$. In fact, any nonidentical Möbius transformations has either one or two fixed points, and can be classified depending on their nature.

We would like to bring a Möbius transformation to a simplest normal form by means of a conjugacy $\phi^{-1} \circ f \circ \phi$ by some $\phi \in \text{Möb}(\hat{\mathbb{C}})$. Since $\text{Möb}(\hat{\mathbb{C}})$ acts double transitively, we can find some ϕ which sends one fixed point of f to ∞ and the other (if exists) to 0. This leads to the following classification:

(i) A hyperbolic Möbius transformation has an attracting and repelling fixed points with multipliers¹ λ and λ^{-1} , where $0 < |\lambda| < 1$. Its normal form is a global linear contraction $z \to \lambda z$ (with possible spiralling if λ is unreal².)

(ii) An *elliptic* Möbius transformation has two fixed points with multipliers λ and λ^{-1} where $\lambda = e^{2\pi i \theta}$, $\theta \in [0, 1)$. Its normal form is the rotation $z \to e^{2\pi i \theta} z$.

(iii) (ii) A *parabolic* Möbius transformation has a single fixed point with multiplier 1. Its normal form is a translation $z \mapsto z + 1$.

EXERCISE 1.6. Verify those of the above statements which look new to you.

1.5. Hyperbolic geometry.

1.5.1. Hyperbolic plane and its motions. Let us now consider a Riemann surface S conformally equivalent to the unit disk \mathbb{D} , or equivalently, to the upper half plane \mathbb{H} , or equivalently, to the strip \mathbb{P} (we refer to such a Riemann surface as a "conformal disk"). Using the isomorphism $S \approx \mathbb{D}$, S can be naturally compactified by adding to it the ideal boundary $\partial S \approx \mathbb{T}$ (see §1.1.3), also called the *circle at infinity* or the absolute.

26

¹The multiplier of a fixed point α is the derivative $f'(\alpha)$ calculated in *any* local chart around α , compare §22.

²Hyperbolic Möbius transformations with unreal λ are also called *loxodromic*

The group $\operatorname{Aut}(S)$ of conformal automorphisms of S in the upper half-plane model consist of Möbius transformations with real coefficients:

$$M: z \mapsto \frac{az+b}{cz+d}; \quad \left(\begin{array}{cc} a & b\\ c & d \end{array}\right) \in \mathrm{SL}(2,\mathbb{R}).$$

Hence $\operatorname{Aut}(S) \approx \operatorname{SL}(2, \mathbb{R})/\{\pm I\} = \operatorname{PSL}(2, \mathbb{R})$. In the unit disk model, it is realised as the group $\operatorname{PSL}^{\#}(2, \mathbb{R})$:

$$M: z \mapsto \frac{\alpha z + \beta}{\bar{\beta} z + \bar{\alpha}} = \lambda \frac{z - a}{1 - \bar{a} z}; \quad \left(\begin{array}{cc} \alpha & \beta \\ \bar{\beta} & \bar{\alpha} \end{array}\right) \in \mathrm{PSL}(2, \mathbb{R}),$$

where $\lambda = \alpha / \bar{\alpha} \in \mathbb{T}, \ z = -\beta / \alpha \in \mathbb{D}.$

The above classification of Möbius transformations $M \in PSL(2, \mathbb{R})$ has a clear meaning in terms of their action on S:

(i) A hyperbolic transformation $M \in PSL(2, \mathbb{R})$ has two fixed points on the ideal boundary ∂S (and does not have fixed points in S). Its normal form in the \mathbb{H} -model is a dialtion $z \mapsto \lambda z$ ($0 < \lambda < 1$), and is a translation $z \mapsto z + a$ in the \mathbb{P} -model, where $a = \log \lambda$.

(ii) A parabolic transformation has a single fixed point on ∂S (and does not have fixed points in S). Its normal form in the \mathbb{H} -model is a translation $z \mapsto z + 1$.

(iii) An elliptic transformation $M \neq id$ has a single fixed point $a \in S$ (and does not have fixed points on ∂S). Its normal form in the D-model is a rotation $z \mapsto e^{2\pi i \theta} z$.

A remarkable discovery by Poincaré is that a conformal disk S is endowed with the intrinsic hyperbolic structure, that is, there exists a Riemannian metric ρ_S on S of constant curvature -1 invariant with respect $PSL(2, \mathbb{R})$ -action. In the \mathbb{H} -, \mathbb{D} - and \mathbb{P} -models, the length element of ρ_S is given respectively by the following expressions:

$$d\rho_{\mathbb{D}} = \frac{2|dz|}{1-|z|^2}, \quad d\rho_{\mathbb{H}} = \frac{|dz|}{y}, \quad d\rho_{\mathbb{P}} = \frac{|dz|}{\sin y},$$

where z = x + iy. This metric is called *hyperbolic*.

EXERCISE 1.7. Verify that the above three expressions correspond to the same metric on S, which has curvature -1 and is invariant under $PSL(2, \mathbb{R})$. Show that the group of orientation preserving hyperbolic motions of S is equal to $Aut(S) \approx PSL(2, \mathbb{R})$.

A conformal disk S endowed with the hyperbolic metric is called the *hyperbolic* plane. In this way, $PSL(2, \mathbb{R})$ assumes the meaning of the group of (orientation preserving) hyperbolic motions of the hyperbolic plane. It acts freely transitively on the unit tangent bundle of \mathbb{H} , so the latter can be identified with $PSL(2, \mathbb{R})$. The isotropy group of $i \in \mathbb{H}$ coincides with the group PSO(2) of hyperbolic rotations

$$z \mapsto \frac{z\cos\theta - \sin\theta}{z\sin\theta + \cos\theta}, \quad \theta \in \mathbb{R}/\pi\mathbb{Z}.$$

Thus, the hyperbolic plane gets identified with the symmetric space

(1.1)
$$\operatorname{PSL}(2,\mathbb{R})/\operatorname{PSO}(2) \approx \mathbb{H}.$$

From this point of view, the hyperbolic metric on \mathbb{H} can be interpreted as follows. Let us consider the Lie algebra $sl(2,\mathbb{R})$ of trace free 2×2 real matrices.

It is endowed with the inner product $\langle a, b \rangle = \text{tr } ab$ (the *Killing form*) which is invariant under the adjoint action

$$a \mapsto gag^{-1}, \quad a \in \mathrm{sl}(2,\mathbb{R}), \ g \in \mathrm{SL}(2,\mathbb{R}),$$

of $SL(2,\mathbb{R})$ on $sl(2,\mathbb{R})$.

Viewed as the linear space, $sl(2, \mathbb{R})$ is just the tangent space to $SL(2, \mathbb{R})$ at the identity. By the left action of $SL(2, \mathbb{R})$ on itself, the Killing form can be promoted to a left-invariant Riemannian metric on $SL(2, \mathbb{R})$. In fact, since the Killing form is invariant under the adjoint action, this metric will also be right invariant. Hence it descends to a metric on the symmetric space $SL(2, \mathbb{R})/SO(2)$ invariant under $SL(2, \mathbb{R})$ -action.

EXERCISE 1.8. Verify that this metric coincides (via the identification (1.1)) with the hyperbolic metric on \mathbb{H} .

A Fuchsian group Γ is a discrete subgroup of $PSL(2, \mathbb{R})$ acting on S.

EXERCISE 1.9. Show that any Fuchsian group acts properly discontinuously on S.

Hence the quotient $X = S/\Gamma$ is a Hausdorff space. Moreover, if Γ acts freely on S, then the complex structure and the hyperbolic metric naturally descend from S to X, and we obtain a hyperbolic Riemann surface.

1.5.2. Hyperbolic geodesics and horocycles. Hyperbolic geodesics in the \mathbb{D} -model of the hyperbolic plane are Euclidean semicircles orthogonal to the absolute \mathbb{T} . For any hyperbolic unit tangent vector $v \in T\mathbb{D}$, there exists a unique oriented hyperbolic geodesic tangent to v. For any two points x and y on the absolute, there exists a unique hyperbolic geodesic with endpoints x and y. The group $PSL(2, \mathbb{R})$ acts freely and transitively on the space of oriented hyperbolic geodesics.

EXERCISE 1.10. Verify the above assertions if they are not familiar to you.

A horocycle in \mathbb{D} centered at $x \in \mathbb{T}$ is a Euclidean circle $\gamma \subset \mathbb{D}$ tangent to \mathbb{T} at x. A horodisk $D \subset \mathbb{D}$ is the disk bounded by the horocycle. In purely hyperbolic terms, horocycles centered at x form a foliations orthogonal to the foliation of geodesics ending at x. The stabilizer of any horocycle (and the corresponding horoball) is the parabolic group fixing its center.

In fact, the \mathbb{H} -model fits better for describing horocycles: in this model the horocycles centered at $x = \infty$ are horizontal lines $L_h = \{ \text{Im } z = h \}$, the corresponding horoballs are the upper half-planes \mathbb{H}_h , and their stabilizer is the group of translations $z \mapsto z + t, t \in \mathbb{R}$.

The quotient of a horoball \mathbb{H}_h by a discrete cyclic group of parabolic transformations $\mathbb{Z} = \langle z \mapsto z + n \rangle$ is called a *cusp*. Conformally it is the punctured disk \mathbb{D}^* , hyperbolically it is the pseudosphere (see Figure ??). Simple closed curves $L_t/\mathbb{Z} \subset \mathbb{H}_h/\mathbb{Z}, t > h$, are also called horocycles (in the cusp).

EXERCISE 1.11. Any cusp \mathbb{H}_h/\mathbb{Z} has infinite hyperbolic diameter but a finite hyperbolic area. The hyperbolic length of the horocylcle L_t/\mathbb{Z} goes to zero as $t \to \infty$.

Let us now consider a Fuchsian group Γ and the corresponding hyperbolic Riemann surface $S = \mathbb{D}/\Gamma$. Hyperbolic geodesics on S are (obviously) projections

of the hyperbolic geodesics on \mathbb{D} ; horocycles on S are (by definition) projections of the horocycles on \mathbb{D} .

Let γ be a non-trivial simple closed curve on S, and let $[\gamma]$ be the class of simple closed curves freely homotopic to γ . To this class corresponds a conjugacy class $A(\gamma)$ of deck transformations (see §2.6). Since deck transformations cannot be elliptic, the elements of $A(\gamma)$ are either all hyperbolic or all parabolic. Accordingly, we say that the class $[\gamma]$ itself is either hyperbolic or parabolic.

PROPOSITION 1.12. a) If the class $[\gamma]$ is hyperbolic then it is represented by a unique closed hyperbolic geodesic $\delta \in [\gamma]$. This geodesic minimizes the hyperbolic length of the closed curves in $[\gamma]$.

b) If the class $[\gamma]$ is parabolic then S contains a neighborhood U isometric to a cusp, and $[\gamma]$ is represented by any horocycle in it. In this case, the class contains arbitrary short curves.

PROOF. Let us consider a lift $\tilde{\gamma}$ of γ , and let $G = \langle \phi^n \rangle_{n \in \mathbb{Z}}$ be its stabilizer.

a) If ϕ is hyperbolic then it has two fixed points, x_{-} and x_{+} , on the absolute, and then the closure of $\tilde{\gamma}$ in \mathbb{D} is a topological interval with endpoints x_{1} and x_{+} . Let us consider the hyperbolic geodesic $\tilde{\delta}$ in \mathbb{D} with endpoints x_{\pm} . It is invariant under the action of the cyclic group G. In fact, it is completely invariant. Indeed, if $\psi(\tilde{\delta}) \cap \tilde{\delta} \neq \emptyset$ for some $\psi \in \Gamma \smallsetminus G$, then $\psi(\tilde{\gamma}) \cap \tilde{\gamma} \neq \emptyset$ as well, which is impossible since γ does not have self-intersections. Hence the projection of $\tilde{\delta}$ to S is equal to $\tilde{\delta}/G$, which is the desired simple closed geodesic representing $[\gamma]$.

b) If ϕ is parabolic then it has a single fixed point x on the absolute, and the closure of $\tilde{\gamma}$ in \mathbb{D} is a topological circle touching \mathbb{T} at x (a "topological horocycle centered at x").

Let U be the corresponding topological horoball bounded by $\tilde{\gamma}$. Let us show that it is completely invariant under G. Indeed, for $\psi \in \Gamma \smallsetminus G$, $\psi(\tilde{U})$ is a topological horoball centered at $\beta(x) \neq x$. But since γ is a simple curve, $\psi(\tilde{\gamma}) \cap \tilde{\gamma} = \emptyset$ for any $\beta \in \Gamma \smallsetminus G$. Since two topological horoballs with disjoint boundaries are disjoint, $\psi(\tilde{U}) \cap \tilde{U} = \emptyset$.

It follows that \tilde{U}/G is is isometrically embedded into $\mathbb{D}/\Gamma = S$. But \tilde{U}/G is a conformal punctured disk containing some standard cusp \mathbb{H}_h/\mathbb{Z} . Thus, this cusp isometrically embeds into S as well, and its horocycles give us desired representatives of $[\gamma]$.

We express part b) of the above statement by saying that the class $[\gamma]$ (or, the curve γ itself) is *represented by a horocycle*, or by a *puncture*, or by a *cusp*.

A simple closed curve on S is called *peripheral* if it is either trivial or is represented by a cusp. For instance, if $S = \hat{\mathbb{C}} \setminus \{x_i\}$ is a sphere with finitely many punctures then γ is non-peripheral iff each component of $\hat{\mathbb{C}} \setminus \gamma$ contains at least two punctures.

EXERCISE 1.13. Show that there is one-to-one correspondence between conjugacy classes of primitive parabolic transformations in a Fuchsian group Γ and cusps of the Riemann surface $S = \mathbb{D}/\Gamma$.

1. CONFORMAL GEOMETRY

1.5.3. *Limit set and ideal boundary*. Let us formulate the dynamical structural theorem for Fuchsian group actions:

THEOREM 1.14. Let Γ be a Fuchsian group acting on $\overline{\mathbb{D}}$. There is a non-empty closed Γ -invariant set $\Lambda = \Lambda(\Gamma) \subset \mathbb{T}$ equal to the limit set of any $\operatorname{orb}(z), z \in \overline{\mathbb{D}}$. Moreover, Λ is either a Cantor set or else, it consists of at most two points. The action of Γ on the complementary set, $\Omega(\Gamma) = \overline{\mathbb{D}} \smallsetminus \Lambda(\Gamma)$, is properly discontinuous.

The set $\Lambda(\Gamma)$ is naturally called the *limit set* of Γ . The complementary set $\Omega(\Gamma)$ is called the *set of discontinuity*. A Fuchsian group is called *elementary* if $|\Lambda(\Gamma)| \leq 2$.

EXERCISE 1.15. Classify elementary Fuchsian groups.

Since Γ acts properly discontinuous on Ω , the quotient space $\hat{S}^I := \Omega/\Gamma$ is Hausdorff. Its interior, $S := \mathbb{D}/\Gamma$, is an orbifold Riemann surface. Thus, \hat{S}^I gives a partial compactification of S called the *ideal compactification*. Accordingly,

 $\partial^I S := (\mathbb{T} \cap \Omega) / \Gamma$

is called the *ideal boundary* of S. (We will see momentarily that this notion is consistent with the topological one introduced in §1.1.3.)

PROPOSITION 1.16. Assume $|\Lambda(\Gamma)| > 1$. Let an interval $\tilde{C} \subset \mathbb{T}$ be a component of $\mathbb{T} \cap \Omega(\Gamma)$. Then the stabilizer of \tilde{C} in Γ is a cyclic group generated by a hyperbolic transformation M.

COROLLARY 1.17. Under the above circumstances, the quotient $C := \tilde{C}/\Gamma$ is a topological circle.

The hyperbolic transformation M from Proposition 1.16 fixes the endpoints of \tilde{C} , and hence its axis $\tilde{\gamma}$ shares the endpoint with \tilde{C} . Let $\tilde{E} \subset \mathbb{D}$ be the topological bigon bounded by I and J.

LEMMA 1.18. The quotient $E := \tilde{E}/\Gamma$ is a cylinder representing a tame end of S with the ideal circle at infinity $C = \tilde{C}/\Gamma$.

Let us summarize our discussion:

PROPOSITION 1.19. Let S be a hyperbolic Riemann surface \mathbb{H}/Γ . Any tame end of S is represented by either a cusp or by a hyperbolic geodesic. In the former case, it admits a one-point compactification with complex structure extended through ∞ . In the latter case, it admits an ideal boundary compactification with an ideal circle at infinity attached to the end. The type of compactification is totally determined by the conformal type of the end.

1.5.4. Convex core. A subset $X \subset \overline{\mathbb{D}}$ is called *(hyperbolically) convex* if for any two points $x, y \in X$, the hyperbolic geodesic arc connecting x and y is also contained in X. The hyperbolic convex hull \hat{X} of a subset $X \subset \overline{\mathbb{D}}$ is the smallest convex set containing X.

For instance, let X be a closed subset of \mathbb{T} , and let $\subset \mathbb{T}$ be the complementary intervals ("gaps") of X. Let us consider open (in $\overline{\mathbb{D}}$) hyperbolic half-planes $H_j \supset I_j$ based on the I_j (they are bounded the hyperbolic geodesics Γ_j that share the endpoints with I_j). Then

(1.2)
$$\hat{X} = \bar{\mathbb{D}} \setminus \left[H_j \right]$$

30

Note that \hat{X} is closed in $\overline{\mathbb{D}}$ and $\hat{X} \cap \mathbb{T} = X$.

In particular, let $\Lambda = \Lambda(\Gamma)$ be the limit set of a Fuchsian group Γ of second kind, and let $\pi : \mathbb{D} \to S$ be the projection onto the quotient Riemann surface. Since Λ is invariant under Γ , the convex hull $\hat{\Lambda}$ is Γ -invariant as well. Hence it covers a Riemann surface $C = C_S$ with boundary called the *convex core* of S.

PROPOSITION 1.20. The natural embedding $C \to S$ is a homotry equivalence.

PROPOSITION 1.21. The group Γ is convex co-compact if and only if the convex core C is compact.

1.5.5. Linking. Let X and Y be two disjoin closed non-sigleton subsets of the unit circle \mathbb{T} . We say that X and Y are unlinked if the geodesics connecting any two poins of X are disjoint from the analogous geodesics for Y.

PROPOSITION 1.22. For sets X and Y as above the following properties are equivalent:

(i) X and Y are unlinked;

(ii The convex hulls \hat{X} and \hat{Y} are disjoint;

(iii) X is contained in a single gap of Y (and the other way around);

(iv) There exist disjoint continua $X' \supset X$ and $Y' \supset Y$ in $\overline{\mathbb{D}}$.

1.5.6. Geodesic laminations. A geodesic lamination in \mathbb{D} is a closed set $Z \subset \mathbb{D}$ that is the disjoint union of complete hyperbolic geodesics. In other words, there is a unique complete geodesic³ $\gamma_z \subset Z$ passing through any point $z \in Z$, and these geodesics are either equal or disjoint. The set Z itself is called the *support* of the lamination.

LEMMA 1.23. For a geodesic lamination in \mathbb{D} , the geodesic γ_z depends continuously on z: If $z_n \to z$ then $\gamma_{z_n}(t) \to \gamma_z(t)$, $t \in \mathbb{R}$, uniformly on \mathbb{R} in the Euclidean metric of \mathbb{D} . In particular, the endpoints at infinity of the γ_{z_n} converge to the endpoint of γ_z . (Here $\gamma_z(t)$ is the natural parametrization of a geodesic γ_z with the origin at z.)

1.6. Annulus and torus.

1.6.1. Modulus of an annulus. Consider an open topological annulus A. Let us endow it with a complex structure. Then A can be represented as the quotient of either \mathbb{C} or \mathbb{H} modulo an action of a cyclic group $\langle \gamma \rangle$. As we have seen above, in the former case A is isomorphic to the flat cylider $\mathbb{C}/\mathbb{Z} \approx \mathbb{C}^*$. In the latter case, we obtain either the punctured disk \mathbb{D}^* (if γ is parabolic) or an annulus $\mathbb{A}(r, R)$ (if γ is hyperbolic). In the hyperbolic case we call A a conformal annulus.

EXERCISE 1.24. Write down explicitly the covering maps $\mathbb{H} \to \mathbb{D}^*$ and $\mathbb{H} \to A(r, R)$.

EXERCISE 1.25. Prove that two round annuli $\mathbb{A}(r, R)$ and $\mathbb{A}(r', R')$ are conformally equivalent if and only if R/r = R'/r'.

³i.e., γ_z fully extended to the ideal boundary in both directions

Let

$$\mod A = \frac{1}{2\pi}\log \frac{R}{r}$$

for a round annulus $A = \mathbb{A}(r, R)$. For an arbitrary conformal annulus A, define its *modulus*, mod(A), as the modulus of a round annulus $\mathbb{A}(r, R)$ isomorphic to A. According to the above exercise, this definition is correct and, moreover, mod A is the only conformal invariant of a conformal annulus.

If A is isomorphic to \mathbb{C}^* or \mathbb{D}^* then we let $\operatorname{mod} A = \infty$.

If A is a topological annulus with boundary whose interior is endowed with a complex structure, then mod(A) is defined as the modulus of the int(A).

The equator of a conformal annulus A is the image of the equator of the round annulus (see §2) under the uniformization $A(r, R) \rightarrow A$.

EXERCISE 1.26. (i) Write down the hyperbolic metric on a conformal annulus represented as the quotient of the strip $S_h = \{0 < \text{Im } z < h\}$ modulo the action of the cyclic group generated by $z \mapsto z + 2\pi$. (What is the relation between h and mod A?)

(ii) Prove that the equator is the unique closed hypebolic geodesic of a conformal annulus A in the homotopy class of the generator of $\pi_1(A)$.

(iii) Show that the hyperbolic length of the equator is equal to $1/ \mod(A)$. Relate it to the multiplier of the deck transformation of \mathbb{H} covering A.

Even if A is a hyperbolic annulus, it is possible to endow it with a flat, rather than hyperbolic, metric. To this end realize A as the quotient of a strip S_h modulo the cyclic group of translations (see the above exercise). Since the flat metric on S_h is translation invariant, it descends to A. In this case we call A a *flat cylinder*.

1.6.2. Modulus of the torus. Let us take a closer look at the actions of the group $\Gamma \approx \mathbb{Z}^2$ on the (oriented) affine plane $P \approx \mathbb{C}$ by translations (see §1.3). We would like to classify these actions up to affine conjugacy, i.e., two actions T and S are considered to be equivalent if there is an (orientation preserving) affine automorphism $A: P \to P$ and an algebraic automorphism $i: \Gamma \to \Gamma$ such that for any $\gamma \in \Gamma$ the following diagram is commutative:

(1.3)
$$\begin{array}{cccc} P & \xrightarrow{} & P \\ A \downarrow & & \downarrow A \\ P & \xrightarrow{} & F \\ \xrightarrow{S^i(\gamma)} & P \end{array}$$

This is equivalent to classifying the quotient tori P/T^{Γ} up to conformal equivalence (since a conformal isomorphism between the quotient tori lifts to an affine isomorphism between the universal covering spaces conjugating the actions of the covering groups).

The conjugacy A in the above definition will also be called *equivariant* with respect to the corresponding group actions.

The problem becomes easier if to require first that i = id in (11.2). Fix an *uncolored pair* of generators α and β of Γ . Since T acts by translations and since P is affine, the ratio

$$\tau = \tau(T) = \frac{T^{\beta}(z) - z}{T^{\alpha}(z) - z}$$

makes sense and is independent of $z \in P$. Moreover, by switching the generators α and β we replace τ with $1/\tau$. Thus, we can color the generators in such a way that $\operatorname{Im} \tau > 0$. (With this choice, the basis of P corresponding to the generators $\{\alpha, \beta\}$ is positively oriented.)

EXERCISE 1.27. Show that two actions T and S of $\Gamma = \langle \alpha, \beta \rangle$ are affinely equivalent with i = id if and only if $\tau(T) = \tau(\tilde{T})$.

According to the discussion in §1.1.5, the choice of generators of Γ means *(un-colored) marking* of the corresponding torus. Thus, the marked tori are classified by a single complex modulus $\tau \in \mathbb{H}$.

Forgetting the marking amounts to replacement one basis $\{\alpha, \beta\}$ in Γ by another, $\{\tilde{\alpha}, \tilde{\beta}\}$. If both bases are positively oriented then there exists a matrix

$$\left(\begin{array}{cc}a&b\\c&d\end{array}\right)\in\mathrm{SL}(2,\mathbb{Z})$$

such that $\tilde{\alpha} = a \alpha + b \beta$, $\tilde{\beta} = c \alpha + d \beta$. Hence

$$\tilde{\tau} = \frac{a\tau + b}{c\tau + d}.$$

Thus, the unmarked tori are parametrized by a point $\tau \in \mathbb{H}$ modulo the action of $\mathrm{SL}(2,\mathbb{Z})$ on \mathbb{H} by Möbius transformations. The kernel of this action consists of two matrices, $\pm I$, so that the quotient group of Möbius transformations is isomorphic to $\mathrm{PSL}(2,\mathbb{Z}) = \mathrm{SL}(2,\mathbb{Z})/ \mathrm{mod}\{\pm I\}$. This group is called *modular*. (In what follows, the modular group is identified with $\mathrm{PSL}(2,\mathbb{Z})$.)

Remark. Passing from $SL(2, \mathbb{Z})$ to $PSL(2, \mathbb{Z})$ has an underlying geometric reason. All tori \mathbb{C}/Γ have a conformal symmetry $z \mapsto -z$. It change marking $\{\alpha, \beta\}$ by $-I\{\alpha, \beta\}$. Thus, remarking by -I acts trivially on the space of marked tori.

The modular group has two generators, the translation $z \mapsto z+1$ and the second order rotation $z \mapsto -1/z$. The intersection of their fundamental domains gives the standard fundamental domain Δ for this action.

EXERCISE 1.28. a) Verify the last statement.

b) Find all points in Δ that are fixed under some transformation of $PSL(2,\mathbb{Z})$. What are the orders of their stabilizers?

c) What is the special property of the tori corresponding to the fixed points?

d) Show that by identifying the sides of Δ according to the action of the generators we obtain a topological plane $Q \approx \mathbb{R}^2$.

e) Endow the above plane with the complex structure so that the natural projection $\mathbb{H} \to Q$ is holomorphic. Show that $Q \approx \mathbb{C}$. (The corresponding holomorphic function $\mathbb{H} \to \mathbb{C}$ is called modular).

Thus, the unmarked tori are parametrized by a single modulus $\mu \in \mathbb{H}/\mathrm{PSL}(2,\mathbb{Z}) \approx \mathbb{C}$.

In the dynamical context we will be dealing with the intermadiate case of *partially marked* tori, i.e., tori with one marked generator α of the fundamental group. This space can be viewed as the quotient of the space of fully marked tori by means of forgetting the second generator, β . If we have two bases $\{\alpha, \beta\}$ and $\{\alpha, \tilde{\beta}\}$ in Γ with the same α , then $\tilde{\beta} = \beta + n\alpha$ for some $n \in \mathbb{Z}$. Hence $\tilde{\tau} = \tau + n$.

Thus, the space of partially marked tori is parametrized by \mathbb{H} modulo action of the cyclic group by translations $\tau \mapsto \tau + n$. The quotient space is identified with the punctured disk \mathbb{D}^* by means of the exponential map $\mathbb{H} \to \mathbb{D}^*$, $\tau \mapsto \lambda = e^{2\pi i \tau}$. So, the partially marked tori are parametrized by a single modulus $\lambda \in \mathbb{D}^*$. We will denote such a torus by \mathbb{T}^2_{λ} .

This modulus λ makes a good dynamical sense. Consider the covering $p: S \to \mathbb{T}_{\lambda}^2$ of the partially marked torus corresponding to the marked cyclic group. Its covering space S is obtained by taking the quotient of \mathbb{C} by the action of the marked cyclic group $z \mapsto z + n$, $n \in \mathbb{Z}$. By means of the exponential map $z \mapsto e^{2\pi i z}$, this quotient is identified with \mathbb{C}^* . Moreover, the action of the complementary cyclic group $z \mapsto z + n\tau$, $n \in \mathbb{Z}$, descends to the action $\zeta \mapsto \lambda^n \zeta$ on \mathbb{C}^* , where the multiplier $\lambda = e^{2\pi i \tau}$ is exactly the modulus of the torus!

Thus, the partially marked torus \mathbb{T}^2_{λ} with modulus $\lambda \in \mathbb{D}^*$ can be realized as the quotient of \mathbb{C}^* modulo the cyclic action generated by the hyperbolic Möbius transformation $\zeta \mapsto \lambda \zeta$ with multiplier λ .

1.7. Geometry of quadratic differentials.

1.7.1. Flat structures with cone singularities and boundary corners. Recall that a Euclidean, or flat, structure on a surface S is an atlas of local charts related by Euclidean motions. However, for topological reasons, many surfaces do not admit any flat structure: the Gauss-Bonnet Theorem bans such a structure on any compact surface except the torus (see below). On the other hand, if we allow some simple singularities, then these obstruction disappears.

Everybody is familiar with a Euclidean cone of angle $\alpha \in (0, 2\pi)$. To give a formal definition, just take a standard Euclidean wedge of angle α and glue its sides by the isometry. It is harder to define (and even harder to visualize) a cone of angle $\alpha > 2\pi$. One possible way is to partition α into several angles $\alpha_i \in (0, 2\pi)$, $i = 0, 1, \ldots n - 1$, to take wedges W_i of angles α_i , and paste W_i to W_{i+1} by gluing the sides isometrically (where *i* is taken mod *n*) (and then to check, by taking a "common subdivision", that the result is independent of the particular choice of the angles α_i).

But there is a more natural way. Consider a smooth universal covering exp: $\mathbb{H} \to \mathbb{D}^*, z \mapsto e^{iz}$, over the punctured disk, and endow \mathbb{H} with the pullback of the Euclidean metric, $e^{-y}|dz|$. Let us define the wedge $W = W(\alpha)$ of angle α as the strip $\{z : 0 \leq \operatorname{Re} z \leq \alpha\}$ completed with one point at $\operatorname{Im} z = +\infty$. If we isometrically glue the sides of this wedge, we obtain the cone $C = C(\alpha)$ of angle α . (We can also define $C(\alpha)$ as the one-point completion at $+\infty$ of the quotient $\mathbb{H}/\alpha\mathbb{Z}$.)

EXERCISE 1.29. Let γ be a little circle around a cone singularity of angle α . Check that the tangent vector γ' rotates by angle α as we go once around γ .

According to the discussion in Appendix 1.9, a cone singularity x with angle $\alpha = \alpha(x)$ carries curvature $2\pi - \alpha$.

Let us now consider a compact flat surface S with boundary. Assume that the boundary is piecewise linear with corners. It means that near any boundary point, S is isometric to a wedge $W(\alpha)$ with some $\alpha > 0$. Points where $\alpha \neq \pi$ are called *corners of angle* α (as the corners are isolated, there are only finitely many of them). The *rotation* $\rho(x)$ at a corner $x \in \partial S$ of angle $\alpha = \alpha(x)$ is defined as $\pi - \alpha$ (see Appendix 1.9).

1.7.2. Gauss-Bonnet Formula.

THEOREM 1.30. If S is a compact flat surface with cone singularities and piecewise linear boundy with corners then

$$\sum K(x) + \sum \rho(y) = 2\pi\chi(S),$$

where the first sum is taken over the cone singularities while the second sum is taken over the boundary corners.

This is certainly a particular case of the general Gauss-Bonnet formula (1.13) from Appendix 1.9, but in our special case we will give a direct combinatorial proof of it.

PROOF. Let us triangulate S by Euclidean triangles in such a way that all cone singularities and all boundary corners are contained in the set of vertices. Let α_i be the list of the angles of all triangles. Summing these angles over the triangles, we obtain:

$$\sum \alpha_i = \pi(\# \text{ triangles}).$$

On the other hand, summation over the vertices gives:

$$\sum \alpha_i = 2\pi (\# \text{ regular vertices}) + \sum_{cones} \alpha(x) + \sum_{corners} \alpha(y)$$
$$= 2\pi (\# \text{ vertices}) - \sum_{cones} K(x) - \sum_{corners} \rho(y) + \pi(\# \text{ corners}).$$

Hence

 $\sum K(x) + \sum \rho(y) = \pi \left(2(\# \text{ vertices}) + (\# \text{ corners}) - (\# \text{ triangles})\right) = 2\pi\chi(S),$

where the last equality follows from

$$3(\# \text{ triangles}) = 2(\# \text{ edges}) + (\# \text{ corners}).$$

1.7.3. Geodesics. Let S be a flat surface with cone singularities. A piecewise smooth curve $\gamma(t)$ in S is called a *geodesic* if it is locally shortest, i.e., for any $x = \gamma(t)$ there exists an $\epsilon > 0$ such that for any $t_1, t_2 \in [t - \epsilon, t + \epsilon], \gamma : [t_1, t_2] \to S$ is the shortest path connecting $\gamma(t_1)$ to $\gamma(t_2)$.

Obviously, any geodesic is piecewise linear: a concatenation of straight Euclidean intervals meeting at cone points. Moreover, both angles between two consecutive intervals in a geodesic must be at least π (in particular, the intervals cannot meet at a cone point with angle $< 2\pi$).

EXERCISE 1.31. Verify these assertions by exploring geodesics on a cone $C(\alpha)$.

THEOREM 1.32. Let S be a closed flat surface with only negatively curved cone singularities. Then for any path $\gamma : [0,1] \to S$, there is a unique geodesic $\delta : [0,1] \to S$ homotopic to γ rel the endpoints.

PROOF. Existence. Let L be the infimum of the lengths of smooth paths homotopic to γ rel the endpoints. We can select a minimizing sequence of picewise linear paths with the intervals of definite length. Such paths form a precompact sequece in S, so we can select a subsequence converging to a path δ in S of length L. Obviously, δ is a local minimizer, and hence is a geodesic.
Uniqueness. Let γ and δ be two geodesics on S homotopic rel the endpoints. They can be lifted to the universal covering \hat{S} to geodesics $\hat{\gamma}$ and $\hat{\delta}$ with common endpoints. We can assume without loss of generality that the endpoints a and b are the only intersection points of these geodesics (replacing them if needed by the arcs $\hat{\gamma}'$ and $\hat{\delta}'$ bounded by two consecutive intersection points). Then $\hat{\gamma}$ and $\hat{\delta}$ bound a polygon Π with vertices at a and b and some corner points x_i . Let y_j be the cone points in int Π . By the Gauss-Bonnet formula,

$$(\pi - \rho(a)) + (\pi - \rho(b)) + \sum (\pi - \rho(x_i)) + \sum K(y_j) = 2\pi.$$

But the first two terms in the left-hand side are less than π while the others are negative – contradiction.

1.7.4. Quadratic differentials and Euc(2)-structures. Let S^* stand for a flat surface S with its cone singularities punctured out.

A parallel line field on S is a family of tangent lines $l(z) \in T_z S$, $z \in S^*$, that are parallel in any local chart of S.

Let $j : \operatorname{Euc}(\mathbb{C}) \to U(2)$ be the natural projection that associates to a Euclidean motion its rotational part. Let $\operatorname{Euc}(n)$ stand for the *j*-preimage of the cyclic group of order *n* in U(2). In other words, motions $A \in \operatorname{Euc}(n)$ are compositions of rotations by $2\pi k/n$ and translations. (So, the complex coordinate, they assume the form $A: z \mapsto e^{2\pi k/n} z + c$.)

LEMMA 1.33. A flat surface S admits a parallel line field if and only if its Euclidean structure can be refined to a Euc(2)-structure.

PROOF. Let S be Euc(2)-surface and let $\theta \in \mathbb{R}/ \mod \pi\mathbb{Z}$. Given a local chart, we can consider the parallel line field in the θ -direction. Since the θ -direction is preserved (mod π) by the group Euc(2), we obtain a well defined parallel line field on S^* .

Vice versa, assume we have a parallel line field on S^* . Then we can rotate the local charts so that this line field becomes horizontal. The transit maps for this atlas are Euclidean motions preserving the horizontal direction, i.e., elements of Euc(2).

LEMMA 1.34. S admits a parallel line field if and only if all cone angles are multiples of π .

PROOF. Any tangent line can be parallely trnsported along any path in S^* . Since S is flat, the result is independent of the choice of a path within a certain homotopy class. S admits a parallel line field if and only if the holonomy of this parallel transport around any cone singularity is trivial, i.e., it rotates the line by a multiple of π . But the holonomy around a cone singularity of angle α rotates the line by angle α .

Next, we will relate flat geometry to complex geometry. Namely, any flat surface S is naturally a Riemann surface. Indeed, since Euclidean motions are conformal, the flat structure induces complex structure on S^* . To extend it through cone singularities, consider a conformal isomorphism $\phi : \mathbb{H}/\alpha\mathbb{Z} \to \mathbb{D}^*$, $z \mapsto e^{2\pi i z/\alpha}$. It extend to a homeomorphism $C(\alpha) \to \mathbb{D}$ that serves as a local chart on the cone $C(\alpha)$.

EXERCISE 1.35. Show that the extension of the conformal structures from S^* to S is unique.

1.7.5. Abelian differentials and translation surfaces.

1.8. Appendix 1 : Tensor calculus in complex dimesnion one.

1.8.1. General notion. For $(n,m) \in \mathbb{Z}^2$, an (n,m)-tensor on a Riemann surface S is an object τ that can be locally written as a differential form

(1.4)
$$\tau(z) \, dz^n d\bar{z}^m$$

Formally speaking, to any local chart $z = \gamma(x)$ on S corresponds a function $\tau_{\gamma}(z)$, and this family of functions satisfy the transformation rule: if $\zeta = \delta(x)$ is another local chart and $z = \phi(\zeta)$ is the transit map, then

(1.5)
$$\tau_{\delta}(\zeta) = \tau_{\gamma}(\phi(\zeta)) \, \phi'(\zeta)^n \overline{\phi'(\zeta)}^m.$$

The regularity of the tensor (e.g., τ can be measurable, smooth or holomorphic) is determined by the regularity of all its local representative τ_{γ} .

Even when dealing with globally defined tensor, we will often use local notaion (1.4), and we will usually use the same notation for a tensor and the representing local function.

Disregarding the regularity issue, tensors form a bigraded commutative semigroup: if τ and τ' are respectively (m, n)- and (m', n')-tensors, then $\tau \tau'$ is an (m + m', n + n')-tensor.

A holomorphic (1, 0)-tensor $\omega(z)dz$ is called an *Abelian differential*; a holomorphic (2, 0)-tensor $q(z)dz^2$ is called a *quadratic differential*. More generally, we can consider *meromorphic* (n, 0)-tensors, e.g., meromorphic quadratic differentials.

A (-1, 1)-tensor $\mu(z)d\bar{z}/dz$ is called a *Beltrami differential*. Notice that the absolute value of a Beltrami differential, $|\mu|$, is a global function on S. (In this book all Beltrami differentials under consideration are assumed measurable and bounded.)

A (1,1)-tensor $\rho = \rho(z) dz d\bar{z}$ with $\rho \ge 0$ is a conformal Riemannian metric $\rho(z)|dz|^2$ on S. Its area form

$$\frac{i}{2}\,\rho(z)\,dz\wedge d\bar{z}=\rho(z)\,dx\wedge dy$$

is a tensor of the same type (both are transformed by the factor $|\phi'(\zeta)|^2$). This allows us to integrate (1, 1)-tensors:

$$\int \rho = \frac{i}{2} \int \rho(z) \, dz \wedge d\bar{z}.$$

For instance, if q is a quadratic differential then |q| is a (1,1)-form, so that we can evaluate $\int |q|$ (at least locally). If q is a quadratic differential and μ is a Beltrami differential, then $q\mu$ is again a (1,1)-form, so the local integral $\int q\mu$ makes sense.

A (-1, 0)-tensor $\frac{v(z)}{dz}$ has the same type as a vector field. Indeed, in this case the tensor rule (1.5) assumes the form $v_{\gamma}(\phi(\zeta)) = \phi'(\zeta) v_{\delta}(\zeta)$ that coinides with the transformation rule for vector fields. EXERCISE 1.36. (i) Let v = v(z)/dz be a C^1 -smooth vector field near ∞ on $\hat{\mathbb{C}}$. Show that $v(z) = az^2 + bz + O(1)$. Moreover, $v(\infty) = 0$ iff a = 0.

(ii) A vector field v(z)/dz is holomorphic on the whole sphere $\hat{\mathbb{C}}$ iff

$$v(z) = az^2 + bz + c$$

EXERCISE 1.37. (i) Let $q = q(z) dz^2$ be a meromorphic quadratic differential near ∞ on $\hat{\mathbb{C}}$. If $q(z) \approx z^{-n}$ then q has a pole of order n + 4 at ∞ . In particular, q has at most a simple pole at ∞ iff q(z) vanishes to the second order at ∞ , i.e., $q(z) = O(|z|^{-3})$.

(ii) $q \in Q^1(\hat{\mathbb{C}})$ iff q(z) is a rational function with simple poles in \mathbb{C} and vanishing to the second order at ∞ .

1.8.2. ∂ and $\bar{\partial}$. The differential of a function $\tau(z)$ can be expressed in (z, \bar{z}) ccordinates as follows:

$$d\tau = \partial_x \tau \, dx + \partial_y \tau \, dy = \partial_z \tau \, dz + \partial_{\bar{z}} \tau \, d\bar{z},$$

where

(1.6)
$$\partial_z = \frac{1}{2}(\partial_x - i\partial_y), \quad \partial_{\bar{z}} = \frac{1}{2}(\partial_x + i\partial_y).$$

This suggests to introduce differential operators ∂ and $\overline{\partial}$ (acting from functions to (1,0)- and (0,1)-forms respectively):

$$\partial \tau = \partial_z \tau \, dz, \quad \bar{\partial} \tau = \partial_{\bar{z}} \tau \, d\bar{z}, \quad \text{so } d = \partial + \bar{\partial}.$$

EXERCISE 1.38. Check that $\partial \tau$ and $\bar{\partial} \tau$ are correctly defined (1,0)- and (0,1)-forms respectively.

We will sometimes use notation ∂ and $\overline{\partial}$ for the partial derivatives (1.6) as well, unless it can lead to a confusion.

Using the semigroup structure, we can extend these differential operators to arbitrry tensors:

$$\partial(\tau \, dz^n d\bar{z}^m) = \partial\tau \, dz^n d\bar{z}^m = \partial_z \tau \, dz^{n+1} d\bar{z}^m,$$
$$\bar{\partial}(\tau \, dz^n d\bar{z}^m) = \bar{\partial}\tau \, dz^n d\bar{z}^m = \partial_{\bar{z}}\tau \, dz^n d\bar{z}^{m+1}.$$

These operators increase the grading by (1,0) and (0,1) respectively.

For instance, if v is a vector field viewed as a (-1,0) tensor, then $\bar{\partial}v$ is a Beltrami differential.

REMARK 1.2. The above *commutative tensor* operators should not be confused with their *anti-commitative exterior* counterparts acting on differential forms. For instance, if $\omega = \omega(z) dz$ is a holomorphic (1,0)-form then in the tensor sense $\partial \omega = \omega'(z) dz^2$, while $\partial \omega = 0$ in the exterior sense. 1.8.3. Pullback and push-forward. Let $f: S \to T$ be a holomorphic map between two Riemann surfaces. Then any (n, m)-form τ on T can be pulled back to an (n, m)-form $f^*\tau$ on S, which in is given in local coordinates by the expression

$$f^*(\tau(w) \, dw^m d\bar{w}^m) = \tau(f(z)) f'(z)^n \bar{f}'(z)^m \, dz^n d\bar{z}^m.$$

Moreover, if τ is a holomorphic/meromorphic (n, 0)-form then so is $f^*(\tau)$.

If f is invertible then of course forms can be also pushed forward. For $\tau = \tau(z) dz^n d\bar{z}^m$, it looks as follows:

$$f_* \tau \equiv (f^{-1})^*(\tau) = \frac{\tau(z)}{f'(z)^n \bar{f}'(z)^m} dw^n d\bar{w}^m$$
 substituting $z = f^{-1}(w)$.

It is less standard that tensors can be also pushed forward by non-invertible maps (at least, by branched coverings of finite degree) by summing up the local pushforwards over the preimages:

$$f_*\tau = \sum (f_i)_*(\tau) = \sum_{z_i \in f^{-1}(w)} \frac{\tau(f(z_i))}{f'(z_i)^n \bar{f}'(z_i)^m} \, dw^n d\bar{w}^m \text{ substituting } z_i = f_i^{-1}(w).$$

where f_i is the local branch of f near $z_i \in f^{-1}(w)$. This expression is well defined outside the set V of critical values of f.

Moreover, if τ is a meromorphic (n, 0)-form with the polar set P then $f_*\tau$ is also meromorphic, with the polar set contained in $f(P) \cup V$. Indeed, outside $f(P) \cup V$, the push-forfard $f_*(\tau)$ is a holomorphic (n, 0)-form with at most power growth near $f(P) \cup V$.

This discussion applies directly to the case of meromorphic quadratic differentials $q = q(z) dz^2$, which will be the main case of our interest:

$$f_*q = \sum (f_i)_*q = \sum_{z_i \in f^{-1}(w)} \frac{q(z_i)}{f'(z_i)^2}.$$

In the case of an area form $\rho dz \wedge d\overline{z}$, the push-forward operation is actually standard as it corresponds to the push-forward of the measure with density ρ :

$$f_*(\rho \, dz \wedge d\bar{z}) = \sum_{z \in f^{-1}(w)} \frac{\rho(z)}{|f'(z)|^2}.$$

Since the area is conserved under invertible changes of variable, we have:

(1.7)
$$\int f^* \rho = d \int \rho, \quad \int f_* \rho = \int \rho$$

(assuming ρ has a finite total mass).

1.8.4. Push-forward is a contraction in \mathcal{Q}^1 . Intergability of a meromorphic quadratic differential q on a Riemann surface S means integrability of the corresponding area form |q|. Let $\mathcal{Q}^1(S)$ stand for the space of integrable meromorphic quadratic differential on S, and $\mathcal{Q}^1_{\text{loc}}(S)$ stand for the space of *locally integrable* ones. Note that $q \in \mathcal{Q}^1_{\text{loc}}$ if and only if it has only simple poles.

For $q \in Q^1(S)$, transformation rules (1.7) (together with the triangle inequality) imply:

(1.8)
$$\int |f_*q| \le \int f_*|q| = \int |q|$$

Thus, the push-forward operator is contracting in the space of integrable holomorphic quadratic differentials. This property plays a key role in the Thurston theory, see §48.

EXERCISE 1.39. Consider a holomorphic quadratic differntial $q = q(z)dz^2$ on the whole Riemann sphere $\hat{\mathbb{C}}$, so q(z) is a rational function.

(i) What is the condition that q has zero/pole at ∞ . If so, what is its order?

(ii) q is integrable if and only if all its poles (including at ∞) are simple;

(iii) For $f: z \mapsto z^d$ and $q = z^n dz^2$, calculate f^*q and f_*q .

LEMMA 1.40. Let $f: S \to T$ be a holomorphic covering of degree d, and let q be an integrable quadratic differential on S. Then

(1.9)
$$\int |f_*q| = \int |q|$$

if and only if $f^*(f_*q) = dq$.

PROOF. Equality (1.9) is equivalent to attaining equality in (1.8). Since both q and f_*q are continuous outside a finite set and $|f_*q| \leq f_*|q|$ everywhere, integral equality in (1.8) is equivalent to pointwise equality $|f_*q| = f_*|q|$. But equality in the triangle inequality is attained if and only if all the terms have the same phase, so

$$f_*q = c_i (f_i)_*q, \quad c_i > 0.$$

Being positive and holomorphic in z, the factors c_i must be constants. Applying the pullback f_i^* to the last equation, we obtain:

$$f^*(f_*q) = c_i q \quad \text{near } z_i \in f^{-1}z.$$

But the ratio $f^*(f_*q)/q$ is a global meromorphic function: if it is locally constant, it must be globally constant, so $f^*(f_*q) = cq$. Finally, by (1.7)

$$\int |f^*(f_*q)| = d \int |f_*q| = d \int |q|,$$

so c = d.

1.8.5. Duality.

LEMMA 1.41. Let $f: S \to T$ be a holomorphic covering of degree d. Consider a meromorphic quadratic differential $q \in Q^1(S)$ and a measurable essentially bounded Belrtami differential μ on T. Then

$$\int_{S} q \cdot f^* \mu = \int_{T} f_* q \cdot \mu.$$

PROOF. It is sufficient to check that

$$\int_U q \cdot f^* \mu = \int_V f_* q \cdot \mu$$

for a base of neighborhoods V on T and $U = f^{-1}(V)$. Since f is covering, we can choose the V so that

$$U = \bigsqcup_{i=1}^{a} U_i,$$

where the restrictions $f_i = (f : U_i \to V)$ are biholomorphic. Then

$$\int_{U} q \cdot f^* \mu = \sum \int_{U_i} q \cdot f^* \mu = \sum \int_{U_i} f^*((f_i)_* q \cdot \mu)$$
$$= \sum \int_{U_i} (f_i)^* q \cdot \mu = \int_{U} (f_i)^* q \cdot \mu.$$

REMARK 1.3. All the above statements concerning covering maps extend immediately to maps $f: S \to T$ that are coverings over $T \smallsetminus A$ where A is a discrete subset. This includes branched coverings (see §2).

1.9. Appendix 2: Gauss-Bonnet formula for variable metrics. Formally speaking, we can skip a discussion of this general version of the Gauss-Bonnet formula as we have verified it directly in all special cases that we need. However, it does give a deeper insight into the matter. The reader can consult, e.g., [] for a proof.

Let S be a compact smooth Riemannian surface, maybe with boundary. Let K(x) be the Gaussian curvature at $x \in S$, and let $\kappa(x)$ be the geodesic curvature at $x \in \partial S$. The Gauss-Bonnet formula related these gemeotric quantities to topology of S:

(1.10)
$$\int_{S} K d\sigma + \int_{\partial S} \kappa ds = 2\pi \chi(S),$$

where $d\sigma$ and ds are the area and length elements respectively.

In particular, if S is closed then

(1.11)
$$\int_{S} K d\sigma = 2\pi \chi(S)$$

which, in particular, implies that there are no flat structures on a closed surface of genus $g \neq 0$.

The boundary term in (1.10) admits a nice interpretation. Let us parametrize a closed boundary curve γ with the length parameter, so that $\gamma'(t)$ is the unit tangent vector to γ . Then for nearby points $\gamma(t)$ and $\gamma(\tau)$, where $\tau = t + \Delta t > t$, let $v(t, \tau)$ be the tangent vector $\gamma'(\tau)$ parallelly transported from $\gamma(\tau)$ back to $\gamma(t)$. Then let $\theta(t, \tau)$ be the angle between $\gamma'(t)$ and $v(t, \tau)$ (taking with positive sign if v points "into S". Summing these angles up over a partition of γ into small intervals, we

obtain the *rotation number* of the tangent vector. It coincides with $\int \kappa ds$.

Note that if ∂S consists of geodesics, the boundary term in (1.10) disappears, and it assumes the same form (1.11) as in the closed case.

If we allow the Riemannian metric to have an isolated singularity at some point $x \in S$ then using the Gauss-Bonnet formula for a small disk arround x, we can assign the Gaussian curvature to x:

(1.12)
$$K(x) = 2\pi - \lim_{\gamma \to x} \int_{\gamma} \kappa ds,$$

provided the limit exists. (Here γ is a small circle around x, and K(x) is assumed to be integrable.)

If we allow a corner of angle $\alpha \in (0, \infty)$ at a boundary point $y \in \partial S$ (see §1.7.1), we can assign the rotation number $\rho(y) = \pi - \alpha \in (\pi, -\infty)$ to it as the angle between the incoming and outgoing tangent vectors.

Then the Gauss-Bonnet formula is still valid for surfaces with singularities and boundary corners, assuming the following form:

(1.13)
$$\int_{S} K d\sigma + \sum_{sing} K(x) + \int_{\partial S} \kappa ds + \sum_{corners} \rho(y) = 2\pi \chi(S).$$

2. Holomorphic proper maps and branched coverings

2.1. First properties. Topological proper maps were defined in §2.

EXERCISE 1.42. Assume that S and T are precompact domains in some ambient surfaces and that $f: S \to T$ admits a continuous extension to the closure \overline{S} . Then f is proper if and only if $f(\partial S) \subset \partial T$.

EXERCISE 1.43. Let $V \subset T$ be a domain and $U \subset S$ be a component of $f^{-1}V$. If $f: S \to T$ is proper, then the restriction $f: U \to V$ is proper as well.

Let now S and T be topological surfaces, and f be a topologically holomorphic map. The latter means that for any point $a \in S$, there exist local charts ϕ : $(U, a) \to (\mathbb{C}, 0)$ and $\psi : (V, fa) \to (\mathbb{C}, 0)$ such that $\psi \circ f \circ \phi^{-1}(z) = z^d$, where $d \in \mathbb{N}$. This number $d \equiv \deg_a f$ is called the *(local)* degree of f at a. If $\deg_a f > 1$, then a is called a branched or critical point of f, and f(a) is called a branched or critical value of f. We also say that d is the multiplicity of a as a preimage of b = f(a).

EXERCISE 1.44. Show than any non-constant holomorphic map between two Riemann surfaces is topologically holomorphic.

EXERCISE 1.45. Show that a continuous map $f : (\mathbb{D}, 0) \to (\mathbb{D}, 0)$ that restricts to a covering $\mathbb{D}^* \to \mathbb{D}^*$ is topologically holomorphic.

A basic property of topologically holomorphic proper maps is that they have a global degree:

PROPOSITION 1.46. Let $f : S \to T$ be a topologically holomorphic proper map between two surfaces. Assume that T is connected. Then all points $b \in T$ have the same (finite) number of preimages counted with multiplicities. This number is called the degree of f, deg f.

PROOF. Since the fibers of a topologically holomorphic map are discrete, they are finite. Take some point $b \in T$, and consider the fiber over it, $f^{-1}b = \{a_i\}_{i=1}^l$. Let $d_i = \deg_{a_i} f$. Then there exists a neighborhood V of b and neighborhood U_i of a_i such that any point $z \in V$, $z \neq b$, has exactly d_i preimages in U_i , and all of them are regular.

Let us show that if V is sufficiently small then all preimages of $z \in V$ belong to $\cup U_i$. Otherwise there would exist sequences $z_n \to b$ and $\zeta_n \in S \setminus \cup U_i$ such that $f(\zeta_n) = z_n$. Since f is proper, the sequence $\{\zeta_n\}$ would have a limit point $\zeta \in S \setminus \cup U_i$. Then $f(\zeta) = b$ while ζ would be different from the a_i - contradiction.

Thus all points close to b have the same number of preimages counted with multiplicities as b, so that this number is locally constant. Since T is connected, this number is globally constant.

COROLLARY 1.47. Topologically holomorphic proper maps are surjective.

The above picture for proper maps suggests the following generalization. A topologically holomorphic map $f: S \to T$ between two surfaces is called a *branched* covering of degree $d \in \mathbb{N} \cup \{\infty\}$ if any point $b \in T$ has a neighborhood V with the following property. Let $f^{-1}b = \{a_i\}$ and let U_i be the components of $f^{-1}V$ containing a_i . Then these components are pairwise disjoint, and there exist maps $\phi_i: (U_i, a_i) \to (\mathbb{C}, 0)$ and $\psi: (V, b) \to (\mathbb{C}, 0)$ such that $\psi \circ f \circ \phi_i^{-1}(z) = z_i^d$. Moreover, $\sum d_i = d$. (A branched covering of degree 2 will be also called a *double branched* covering.)

We see that a topologically holomorphic map is proper if and only if it is a branched covering of finite degree. All terminilogy developed above for proper maps immediately extends to arbitrary branched coverings.

Note that if $V \subset T$ is a domain which does not contain any critical values, then the "map f is unbranched over V", i.e., its restriction $f^{-1}V \to V$ is a covering map. In particular, if V is simply connected, then $f^{-1}V$ is the union of d disjoint domains U_i each of which homeomorphically projects onto V. In this case we have d well-defined branches $f_i^{-1}: V \to U_i$ of the inverse map. (When it does not lead to confusion, we will often use notation f^{-1} for the inverse branches.)

2.2. Riemann-Hurwitz formula. This formula gives us a beautiful relation between topology of the surfaces S and T, and branching properties of f.

RIEMANN - HURWITZ FORMULA. Let $f : S \to T$ be a branched covering of degree d between two topological surfaces of finite type. Let C be the set of branched points of f. Then

$$\chi(S) = d \cdot \chi(T) - \sum_{a \in C} (\deg_a f - 1).$$

Let us define the *multiplicity* of $a \in C$ as a critical point to be equal to deg_a f-1 (in the holomorphic case it is the multiplicity of a as the root of the equation f'(a) = 0). Then the sum in the right-hand side of the Riemann-Hurwitz formula is equal to the *number of critical points of* f *counted with multiplicities*.

PROOF. Let us first assume that S and T are closed Riemann surfaces.

Let us consider a triangulation \mathcal{T} of T such that all critical values of f are vertices of \mathcal{T} . By the Euler formula,

$$\chi(T) = v(\mathcal{T}) - e(\mathcal{T}) + t(\mathcal{T}),$$

where v, e and t stand for the number of vertices, edges and faces (triangles) of \mathcal{T} . Let \mathcal{S} be the lift of this triangulation to S. Then

$$t(\mathcal{S}) = d \cdot t(\mathcal{T}), \quad e(\mathcal{S}) = d \cdot e(\mathcal{T}), \quad v(\mathcal{S}) = d \cdot v(\mathcal{T}) - \sum_{a \in C} (\deg_a f - 1),$$

and the conclusion follows.

To deal with non-closed case, consider the one-point-per-end compactifications \hat{S} and \hat{T} of our surfaces. If S and T are of finitre type then these surfaces are closed. Since f is proper, it continuously extends to a map $\hat{f}: \hat{S} \to \hat{T}$. This map is certainly proper. By Exersice 1.45), it is topologically holomorphic. Thus, it is a branched covering (of the same degree d). As in the above calulation, we have

$$|\mathcal{E}(S)| = d \cdot |\mathcal{E}(T)| - \sum_{e \in \mathcal{E}(S)} (\deg_e \hat{f} - 1).$$

Putting this together with the Riemann-Hurwitz formula for \hat{f} implies the desired.

REMARK 1.4. The formula also applies to surfaces with boundary, with the same proof (or by removing boundary, which does not change the Euler characteristic).

COROLLARY 1.48. Under the above circumstances, assume that T is a topological disk. Then S is a topological disc as well if and only if there are d-1 critical points in S (counted with multiplicities).

PROOF. A surface S is a topological disk if and only if $\chi(S) = 1$.

2.3. Topological Argument Principle. Consider the punctured plane $\mathbb{R}^2 \setminus \{b\}$. If $\gamma : S^1 \to \mathbb{R}^2 \setminus \{b\}$ is a smooth oriented Jordan curve then one can define the winding number of γ around b as

$$w_b(\gamma) = \int_{\gamma} d(\arg(x-b)).$$

Since the 1-form $d(\arg(x-b))$ is closed, the winding number is the same for homotopic curves. Hence we can define the winding number $w_b(\gamma)$ for any continuous Jordan curve $\gamma: S^1 \to \mathbb{R}^2 \setminus \{b\}$ by approximating it with a smooth Jordan curves.

Furthermore, the winding number can be linearly extended to any simplicial 1-cycle in $\mathbb{R}^2 \setminus \{b\}$ with integer coefficients (i.e., a formal combination of oriented Jordan curves in $\mathbb{R}^2 \setminus \{b\}$) and then factored to the first homology group. It gives an isomorphism

(2.1)
$$w: H_1(\mathbb{R}^2 \setminus \{b\}) \to \mathbb{Z}, \quad [\gamma] \mapsto w_b(\gamma).$$

EXERCISE 1.49. Prove the last statement.

PROPOSITION 1.50. Let D be a Jordan disc and let $f : \overline{D} \to \mathbb{R}^2$ be a continuous map that does not assume some value $b \in \mathbb{R}^2$ on ∂D . If $w_b(f|\partial \mathbb{D}) \neq 0$ then fassumes the value b in D.

PROOF. Obviously, the curve $\gamma = (f : \partial D \to \mathbb{R}^2)$ is contractible in $f(\overline{D})$. If $b \notin f(D)$ then γ would be contractible in $\mathbb{R}^2 \setminus \{b\}$, so it would have zero winding number around b.

Let $x \in D$ be an isolated preimage of b = fx. Then one can define the $\operatorname{ind}_x(f)$ as follows. Take a disk $V \subset D$ around x that does not contain other preimages of b = fx. Take a positively oriented Jordan loop $\gamma \subset V \setminus \{x\}$ around x whose image does not pass through b, and calculate the winding number of the curve $f : \gamma \to \mathbb{R}^2 \setminus \{b\}$ around b:

$$\operatorname{ind}_x(f) = w_{fx}(f \circ \gamma).$$

Clearly it does not depend on the loop γ , since the curves corresponding to different loops are homotopic without crossing b.

PROPOSITION 1.51. Let $D \subset \mathbb{R}^2$ be a domain bounded by a Jordan curve Γ , and let $f : \overline{D} \to \mathbb{R}^2$ be a continuous map such that the curve $f \circ \Gamma$ does not pass through some point $b \in \mathbb{R}^2$. Assume that the preimage of this point $f^{-1}b$ is discrete in D. Then

$$\sum_{x \in f^{-1}b} \operatorname{ind}_x(f) = w_b(f \circ \Gamma),$$

provided Γ is positively oriented.

PROOF. Note first that since $f^{-1}b$ is a discrete subset of a compact set \overline{D} , $f^{-1}x$ is actually finite, so that the above sum makes sense.

Select now small Jordan loops γ_i around points $x_i \in f^{-1}b$, and orient them anticlockwise. Since Γ and these loops bound a 2-cell, $[\Gamma] = \sum [\gamma_i]$ in $H_1(\bar{D} \smallsetminus f^{-1}b)$. Hence $f_*[\Gamma] = \sum f_*[\gamma_i]$ in $H_1(\mathbb{R}^2 \smallsetminus \{b\})$. Applying the isomorphism (2.1), we obtain the desired formula.

EXERCISE 1.52. Let $f: D \to \mathbb{R}^2$ be a continuous map, and let $a \in D$ be an isolated point in the fiber $f^{-1}b$, where b = f(a). Assume that $\operatorname{ind}_a(f) \neq 0$. Then f is locally surjective near a, i.e., for any $\epsilon > 0$ there exists $a \delta > 0$ such that $f(\mathbb{D}_{\epsilon}(a)) \supset \mathbb{D}_{\delta}(b)$.

Hint: For a small ϵ -circle γ around a, the curve $f \circ \gamma$ stays some positive distance δ from b. Then for any $b' \in \mathbb{D}_{\delta}(b)$ we have: $\operatorname{ind}_{b}(f \circ \gamma) = \operatorname{ind}_{b}(f \circ \gamma) \neq 0$. But if $b' \notin f(\mathbb{D}_{\epsilon}(a))$ then the curve $f \circ \gamma$ could be shrunk to b without crossing b'.

2.3.1. Degree of proper maps.

2.4. Lifts.

LEMMA 1.53. Let $f : (S, a) \to (T, b)$ and $\tilde{f} : (\tilde{S}, \tilde{a}) \to \tilde{T}, \tilde{b})$ be two double branched between topological disks (with or without boundary) coverings branched at points a and \tilde{a} respectively. Then any homeomorphism $h : (T, b) \to (\tilde{T}, \tilde{b})$ lifts to a homeomorphism $H : (S, a) \to (\tilde{S}, \tilde{a})$ which makes the diagram

$$\begin{array}{ccc} (S,a) & \longrightarrow & (S,\tilde{a}) \\ f \downarrow & & \downarrow \tilde{f} \\ (T,b) & \longrightarrow & (\tilde{T},\tilde{b}) \end{array}$$

commutative. Moreover, the lift H is uniquely determined by its value at any unbranched point $z \neq a$. Hence there exists exactly two lifts.

If the above surfaces are Riemann and the map h is holomorphic then then the lifts H are holomorphic as well.

PROOF. Puncturing all the surfaces at their preferred points, we obtain four topological annuli. The maps f and \tilde{f} restrict to the unbranched double coverings between respective annuli, while h restricts to a homeomorphism. The image of the fundamental group $\pi_1(S \setminus \{a\})$ under f consist of homotopy classes of curves with winding number 2 around b, and similar statement holds for \tilde{f} . Since the winding number is preserved under homeomorphisms,

(2.2)
$$h_*(f_*(\pi_1(S \setminus \{a\})) = f_*(\pi_1(S \setminus \{\tilde{a}\})).$$

By the general theory of covering maps, h admits a lift

$$H: S \smallsetminus \{a\} \to S \smallsetminus \{\tilde{a}\}$$

which makes the "punctured" diagram (2.2) commutative. Moreover, this lift is uniquely determined by the value of H at any point $z \in S \setminus \{a\}$.

Extend now H at the branched point by letting $H(a) = \tilde{a}$. It is clear from the local structure of branched coverings that this extension is continuous (as well as the inverse one), so that it provides us with the desired lift.

If all the given maps are holomorphic then the lift H is also holomorphic on the punctured disk $S \setminus \{a\}$. Since isolated singularities are removable for bounded holomorphic maps, the extension of H to the whole disk is also holomorphic. \Box

EXERCISE 1.54. Similar statement holds for branched coverings f and \tilde{f} with a single branched point (of any degree). Analyse the situation with two branched points.

EXERCISE 1.55. Assume that all the topological disks in the above lemma are \mathbb{R} -symmetric and that all the maps commute with the reflection σ with respect to \mathbb{R} . Assume also that $h(f(T \cap \mathbb{R})) = \tilde{f}(\tilde{T} \cap \mathbb{R})$. Then both lifts H also commute with σ (in particular, they preserve the real line).

3. Extremal length and width

3.1. Definitions. Let us now introduce one of the most powerful tools of conformal geometry. Given a family Γ of curves in a Riemann surface U, we will define a conformal invariant $\mathcal{L}(\Gamma)$ called the extremal length of Γ . Consider a measurable conformal metric $\rho|dz|$ on \mathbb{C} with finite total mass

$$m_{\rho}(U) = \int \int \rho^2 dx \wedge dy$$

(such metrics will be called *admissible*). Let

$$\rho(\gamma) = \int_{\gamma} \rho |dz|,$$

stand for the length of $\gamma \in \Gamma$ in this metric (with the convention $\rho(\gamma) = \infty$ if γ is non-rectifiable, or $\rho | \gamma$ is not measurable, or else it is not integrable⁴. Define the ρ -length of Γ as

$$\rho(\Gamma) = \inf_{\gamma \in \Gamma} \rho(\gamma).$$

Normalize it in the scaling invariant way:

$$\mathcal{L}_{\rho}(\Gamma) = \frac{\rho(\Gamma)^2}{m_{\rho}(U)},$$

and define the *extremal length* of Γ as follows:

$$\mathcal{L}(\Gamma) = \sup_{\rho} \mathcal{L}_{\rho}(\Gamma),$$

where the supremum is taken over all admissible metrics.

A metric ρ on which this supremum is attained (if exists) is called *extremal*.

EXERCISE 1.56. Show that the value of $\mathcal{L}(\Gamma)$ does not change if one uses only continuous admissible metrics ρ .

Let us summarize immediate consequences of the definition:

⁴For this to make sense, we should think of ρ as an actual *function* rather than a class of functions up to modification on null-sets. It is also convenient to assume that ρ is defined *everywhere*.

EXERCISE 1.57. • Extension of the family: If a family of curves Γ' contains a family Γ , then $\mathcal{L}(\Gamma') \leq \mathcal{L}(\Gamma)$.

• Overflowing: If Γ overflows Γ' (i.e., each curve of Γ contains some curve of Γ'), then $\mathcal{L}(\Gamma) \geq \mathcal{L}(\Gamma')$.

• Independence of the embient surface: If $U \subset U'$ and Γ is a family of curves in U then $\mathcal{L}(\Gamma) = \mathcal{L}(\Gamma')$. (This justifies skipping of "U" in the notation.)

The extremal width of the family Γ is defined as the inverse to its length: $\mathcal{W}(\Gamma) = \mathcal{L}(\Gamma)^{-1}$. One can also conveniently define it as follows:

EXERCISE 1.58.

$$\mathcal{W}(\Gamma) = \inf m_{\rho}(U),$$

where the infimum is taken over all admissible metrics with $\rho(\gamma) \ge 1$ for all curves $\gamma \in \Gamma$.

REMARK 1.5. One should think that a family is "big" if it has big extremal width. So, big families are short.

The extremal lenght and width are *conformal invariants*: If $\phi: U \to U'$ is a conformal isomorphism between two Riemann surfaces such that $\phi(\Gamma) = \Gamma'$, then $\mathcal{L}(\Gamma) = \mathcal{L}(\Gamma')$. This immediately follows from the observation that ϕ transfers the family of admissible metrics on U onto the family of admissible metrics on U'.

3.2. Electric circuits laws. We will now formulate two crucial properties of the extremal length and width which show that the former behaves like the resistance in electric circuits, while the latter behaves like conductance.

Let Γ_1 , Γ_2 and Γ be three families of curves on U. We say that Γ disjointly overflows Γ_1 and Γ_2 if any curve $\gamma \in \Gamma$ contains a pair of disjoint curves $\gamma_1 \in \Gamma_1$ and $\gamma_2 \in \Gamma_2$.

SERIES LAW. Assume that a family Γ disjointly overflows families Γ_1 and Γ_2 . Then

$$\mathcal{L}(\Gamma) \geq \mathcal{L}(\Gamma_1) + \mathcal{L}(\Gamma_2),$$

or equivalently,

$$\mathcal{W}(\Gamma) \leq \mathcal{W}(\Gamma_1) \oplus \mathcal{W}(\Gamma_2).$$

PROOF. Let ρ_1 and ρ_2 be arbitrary admissible metrics. By appropriate scalings, we can normalize them so that

$$\rho_i(\Gamma_i) = m_{\rho_i}(U) = \mathcal{L}_{\rho_i}(\Gamma_i), \quad i = 1, 2.$$

Let $\rho = \max(\rho_1, \rho_2)$. Since any $\gamma \in \Gamma$ contains two disjoint curves $\gamma_i \in \Gamma_i$, we have:

$$\rho(\gamma) \ge \rho_1(\gamma_1) + \rho_2(\gamma_2) \ge \rho_1(\Gamma_1) + \rho_2(\Gamma_2) = \mathcal{L}_{\rho_1}(\Gamma_1) + \mathcal{L}_{\rho_2}(\Gamma_2).$$

Taking the infimum over all $\gamma \in \Gamma$, we obtain:

$$\rho(\Gamma) \ge \mathcal{L}_{\rho_1}(\Gamma_1) + \mathcal{L}_{\rho_2}(\Gamma_2).$$

On the other hand,

$$m_{\rho}(U) \leq m_{\rho_1}(U) + m_{\rho_2}(U) = \mathcal{L}_{\rho_1}(\Gamma_1) + \mathcal{L}_{\rho_2}(\Gamma_2).$$

Hence

$$\mathcal{L}_{\rho}(\Gamma) \geq \mathcal{L}_{\rho_1}(\Gamma_1) + \mathcal{L}_{\rho_2}(\Gamma_2).$$

Taking the supremum over all normalized metrics ρ_1 and ρ_2 , we obtain the desired inequality.

We say that two families of curves, Γ_1 and Γ_2 , are *disjoint* if they are contained in disjoint measurable sets.

PARALLEL LAW. Let $\Gamma = \Gamma_1 \cup \Gamma_2$. Then

$$\mathcal{W}(\Gamma) \leq \mathcal{W}(\Gamma_1) + \mathcal{W}(\Gamma_2).$$

Moreover, if Γ_1 and Γ_2 are disjoint then

$$\mathcal{W}(\Gamma) = \mathcal{W}(\Gamma_1) + \mathcal{W}(\Gamma_2).$$

PROOF. This time, let us normalize admissible metrics ρ_1 and ρ_2 so that $\rho_i(\Gamma_i) \geq 1$, and let again $\rho = \max(\rho_1, \rho_2)$. Then $\rho(\Gamma) \geq 1$ as well, and hence

$$\mathcal{W}(\Gamma) \le m_{\rho}(U) \le m_{\rho_1}(U) + m_{\rho_2}(U).$$

Taking the infimum over the metrics ρ_i , we obtain the desired inequality.

Assume now that Γ_1 and Γ_2 are disjoint. Let X_1 and X_2 be two disjoint measurable sets supporting the respective families. Take any admissible metric ρ with $\rho(\Gamma) \geq 1$, and let $\rho_i = \rho | X_i$. Then $\rho_i(\Gamma_i) \geq 1$ as well, and hence

$$m_{\rho}(U) = m_{\rho_1}(U) + m_{\rho_2}(U) \ge \mathcal{W}(\Gamma_1) + \mathcal{W}(\Gamma_2).$$

Taking the infimum over ρ , we obtain the opposite inequality.

REMARK 1.6. Both laws obviously extend to the case of n families $\Gamma_1, \ldots, \Gamma_n$.

3.3. Modulus of an annulus revisited.

3.3.1. Modulus as the extremal length. We will now calculate the modulus of an annulus (see §1.6.1) in terms of the extremal length. Consider a flat cylinder $C = C[l, h] = (\mathbb{R}/l\mathbb{Z}) \times [0, h]$ with circumferance l and height h. Curves joining the top to the bottom of C will be called *vertical*. Among these curves there are genuinly vertical, that is, straight intervals perpendicular to the top and the bottom. Horizontal curves in C are closed curves homotopic to the top and the bottom of C. Among them there are genuinly horizontal, that is, the circles parallel to the top and the bottom. Genuinly verical and horizontal curves form the vertical and horizontal foliations respectively.

If A is an open conformal annulus, then it is isomorphic to a flat cylinder, $A \approx C(0, h)$, and we will freely identify them. In particular, curves in A corresponding to vertical/horizontal curves in the cylinder will be also referred to as vertical/horizontal.⁵

PROPOSITION 1.59. Let Γ be a family of vertical curves in the annulus A containing almost all genuinly vertical ones. Then $\mathcal{L}(\Gamma) = \operatorname{mod}(A)$.

PROOF. We will identify A with the cylinder C(l, h). Take first the flat metric e on the cylinder.⁶ Then $e(\gamma) \ge h$ for any $\gamma \in \Gamma$, so that, $e(\Gamma) = h$. On the other hand, $m_e(\Gamma) = lh$. Hence

$$\mathcal{L}(\Gamma) \ge \mathcal{L}_e(\Gamma) = h^2/lh = \mathrm{mod}(A).$$

⁵Notice that if $A \subset \mathbb{C}$ but ∂A is not locally connected, then vertical curves do not have to land at some points of ∂A .

 $^{^{6}}$ As we will see, e will happen to be the extremal metric.

Take now any admissible metric ρ on A. Let γ_{θ} be the genuinly vertical curve through $\theta \in \mathbb{R}/l\mathbb{Z}$. Then $\rho(\Gamma) \leq \rho(\gamma_{\theta})$ for any $\theta \in \mathbb{R}/l\mathbb{Z}$. Integrating this over $\mathbb{R}/l\mathbb{Z}$ (using that $\gamma_{\theta} \in \Gamma$ for a.e. $\theta \in \mathbb{R}/l\mathbb{Z}$) and applying the Cauchy-Schwarz inequality, we obtain:

(3.1)
$$(l \cdot \rho(\Gamma))^2 \le \left(\int_{\mathbb{R}/lZ} \rho(\gamma_\theta) \, d\theta \right)^2 = \left(\int_A \rho \, dm_e \right)^2 \le lh \, m_\rho(A).$$

Hence $\mathcal{L}_{\rho}(A) \leq \operatorname{mod}(A)$, and the conclusion follows.

There is also the "dual" way to evaluate the same modulus:

EXERCISE 1.60. Let Γ be a family of horizontal curves in A containing almost all genuinly horizontal curves. Then

$$\operatorname{mod}(A) = \mathcal{W}(\Gamma).$$

3.3.2. *Gröztsch Inequality*. The following inequality plays an outstanding role in holomorphic dynamics (the name we use for it is widely adopted in the dynamical literature):

PROPOSITION 1.61. Consider a conformal annulus A containing n disjoint conformal annuli $A_1, \ldots A_n$ homotopically equivalent to A. Then

$$\operatorname{mod}(A) \ge \sum \operatorname{mod} A_k.$$

PROOF. Let Γ_k be the horizontal family of A_k and Γ be the horizontal family in A. By the Parallel Law, $\mathcal{W}(\Gamma) \geq \sum \mathcal{W}(\Gamma_k)$, and the concusion follows from Exercise 1.60. (Dually, one can apply the Series Law to the extremal length of the vertical families.)

3.3.3. Euclidean geometry of an annulus. The length-area method allows one to relate mod(A) to the Euclidean geometry of A. As a simple illustration, let us show that mod(A) is bounded by the "distance between the inner and the outer complements of A rel the size of the inner complement":

LEMMA 1.62. Consider a topological annulus $A \subset \mathbb{C}$. Let K and Q stand for its inner and outer complements respectively. Then

define

$$\operatorname{mod}(A) \leq C \operatorname{dist}(K, Q) / \operatorname{diam} K.$$

PROOF. Let Γ be the family of horizontal curves in A. According to the last Exercise, we need to bound $\lambda(\Gamma)$ from below.

Take points $a \in K$ and $c \in Q$ on minimal distance $\operatorname{dist}(K, Q)$, and then select a point $b \in K$ such that $\operatorname{dist}(a, b) > \operatorname{diam} K/2$. Consider a family Δ of closed Jordan curves $\gamma \subset \mathbb{C} \setminus \{a, b, c\}$ with winding number 1 around a and b and winding number 0 around c. Since $\Gamma \subset \Delta$, $\lambda(\Gamma) \geq \lambda(\Delta)$.

Let us estimate $\lambda(\Delta)$ from below. Rescale the configuration $\{a, b, c\}$ (without changing notations) so that |a - b| = 1 and |a - c| = d, where

$$\frac{1}{2}\operatorname{dist}(K,Q)/\operatorname{diam} K \le d \le \operatorname{dist}(K,Q)/\operatorname{diam} K.$$

Consider a unit neighborhood B of the union $[a, b] \cup [a, c]$ of two intervals, and endow it with the Euclidean metric E (extended by 0 outside B). Then $l_E(\Delta) \ge 1$ while $m_E(B) \le Ad$. Hence $\lambda_E(\Delta) \ge 1/Ad$, and we are done.

EXERCISE 1.63. For an annulus A as above, prove a lower bound:

 $\operatorname{mod}(A) \ge \mu(\operatorname{dist}(K, Q) / \operatorname{diam}(K)) > 0.$

3.3.4. Shrinking nests of annuli. Let $X \subset \mathbb{C}$ be a compact connected set. Let us say that a sequence of disjoint annuli $A_n \subset \mathbb{C}$ is nested around X if for any any n, A_n separates both A_{n+1} and X from ∞ . (We will also call it a "nest of annuli around X".)

COROLLARY 1.64. Consider a nest of annuli A_n around X. If $\sum \text{mod } A_n = \infty$ then X is a single point.

PROOF. Only the first annulus, A_1 , can be unbounded in \mathbb{C} . Take some disk $D = \mathbb{D}_R$ containing A_2 , and consider the annulus $D \smallsetminus X$. By the Gröztsch Inequality,

$$\operatorname{mod}(D \smallsetminus X) \ge \sum_{n \ge 2} \operatorname{mod} A_n = \infty.$$

Hence X is a single point.

3.3.5. Quadrilaterals. Given a standard flat recatangle $\Pi[l, h] = [0, l] \times [0, h]$, we can naturally define (genuinly) vertical/horizontal curves in it. We let mod $\Pi = h/l$. Two rectangels Π and Π' are called *conformally equivalent* if there is a conformal isomorphism $\Pi \to \Pi'$ that maps the horizontal sides of Π to the horizontal sides of Π' .

EXERCISE 1.65. Two rectangles Π and Π' are conformally equivalent if and only if mod $\Pi = \mod \Pi'$.

EXERCISE 1.66. Let Γ be a family of vertical curves in $\Pi[l, h]$ that contains almost all genuinly vertical curves. Then $\mathcal{L}(\Gamma) = \text{mod}(\Pi)$.

A quadrilateral or a conformal rectangle Q(a, b, c, d) is a conformal disk Q with four marked points a, b, c, d on its ideal boundary. We will often let Q = Q(a, b, c, d)so that there is no notational difference between the quadrilateral and the underlying disk. (If the underlying disk is called, say, D then the corresponding quadrilateral is denoted accordingly, D = D(a, b, c, d).

A quadrilateral has four ideal boundary sides. *Marking* of a quadrilateral is a choice of pair of opposite sides called "horizontal" (and then the other pair is naturally called "vertical"). Any marked quadrilateral can be conformally mapped onto a rectangle $\Pi(l, h)$ so that the horizontal sides of Q go to the horizontal sides of $\Pi(l, h)$. At this point, we can naturally define (genuinly) vertical/horizontal curves in Q, and also let mod $Q = \mod \Pi(l, h)$. With this at hands, Exercises 1.65 and 1.66 immediately extend to general marked quadrilaterals.

As an important example, let us consider the quadrilateral $\Pi_R = \mathbb{H}(0, 1, R, \infty)$, R > 1, based on the upper half-plane. Let $\theta(R) := \operatorname{mod}(\Pi_R) = \mathcal{L}(\Gamma)$ where Γ is the path family in \mathbb{H} connecting [0, 1] to $[R, \infty]$.

EXERCISE 1.67. Show that $\frac{1}{4\pi} \log R \le \theta(R) \le -\frac{4\pi}{\log(1-1/R)}.$

(Here the left-hand estimate is good for big R, while the right-hand one is good for $R \approx 1$.)

3.3.6. Tori. Let us now consider a flat torus \mathbb{T}^2 . Given a non-zero homology class $\alpha \in H_1(\mathbb{T}^2)$, we let Γ_{α} be the family of closed curves on \mathbb{T}^2 representing α (we call them α -curves). Among these curves, there are closed geodesics, α -geodesics (they lift to straight lines in the universal covering \mathbb{R}^2). They form a foliation. All these geodesics have the same length, l_{α} .

EXERCISE 1.68. Let Γ be a family of α -curves containing all α -geodesics. Then

$$\mathcal{W}(\Gamma) = \frac{\operatorname{area} \mathbb{T}^2}{l_\alpha^2}$$

An annulus A embedded into \mathbb{T}^2 is called an α -annulus if its horizontal curves represent the class α . The following observation finds interesting applications in dynamics and geometry:

PROPOSITION 1.69. Let
$$A_1, \ldots, A_n$$
 be a family of disjoint α -annuli. Then

$$\sum \operatorname{mod} A_i \leq \frac{\operatorname{area} \mathbb{T}^2}{l_\alpha^2}$$

PROOF. Let Γ_i be the family of horizontal curves of the annulus A_i . Then by the Parallel Law, $\sum \mathcal{W}(\Gamma_i) \leq \mathcal{W}(\Gamma_{\alpha})$, and the result follows from Exercises 1.60 and 1.68.

3.4. Dirichlet integral.

3.4.1. Definition. Consider a Riemann surface S endowed with a smooth conformal metric ρ . The Dirichlet integral (D.I.) of a function $\chi : S \to \mathbb{C}$ is defined as

$$D(\chi) = \int \|\nabla \chi\|_{\rho} \, dm_{\rho},$$

where the norm of the gradient and the area form are evaluated with respect to ρ . However:

EXERCISE 1.70. The Dirichlet integral is independent of the choice of the conformal metric ρ . In particular, it is invariant under conformal changes of variable.

In the local coordinates, the Dirichlet integral is expressed as follows:

$$D(h) = \int (|h_x|^2 + |h_y|^2) dm = \int (|\partial h|^2 + |\bar{\partial} h|^2) dm.$$

In particular, for a conformal map $h: U \hookrightarrow \mathbb{C}$ we have the area formula:

$$D(h) = \int |h'(z)|^2 dm = \operatorname{area} h(U).$$

3.4.2. D.I. of a harmonic function.

EXERCISE 1.71. Consider a flat cylinder $A = S^1 \times (0, h)$ with the unit circumference. Let $\chi : A \to (0, 1)$ be the projection to the second coordinate (the "height" function) divided by h. Then $D(\chi) = 1/h$.

Note that the function χ in the exercise is a harmonic function with boundary values 0 and 1 on the boundary components of the cylinder (i.e., the solution of the Dirichlet problem with such boundary values).

EXERCISE 1.72. Such a harmonic function is unique up to switching the boundary components of A, which leads to replacement of χ by $1 - \chi$. Due to the conformal invariance of the Dirichlet integral (as well as the modulus of an annulus and harmonicity of a function), these trivial remarks immediately yield a non-trivial formula:

PROPOSITION 1.73. Let us consider a conformal annulus A. Then there exist exactly two proper harmonic function $\chi_i : A \to (0,1)$ (such that $\chi_1 + \chi_2 = 1$) and $D(\chi_i) = 1/ \mod(A)$.

3.4.3. Multi-connected case. Let S be a compact Riemann surface with boundary. Let $\partial S = (\partial S)_0 \sqcup (\partial S)_1$, where each $(\partial S)_i \neq \emptyset$ is the union of several boundary components of ∂S . Let us consider two families of curves: the "vertical family" Γ^v consisting of arcs joining $(\partial S)_0$ to $(\partial S)_1$, and the "horizontal family" Γ^h consisting of Jordan multi-curves separating $(\partial S)_0$ from $(\partial S)_1$. (A multicurve is a finite union of Jordan curves.)

Let $\chi: S \to [0,1]$ be the solution of the Dirichlet problem equal to 0 on $(\partial S)_0$ and equal to 1 on $(\partial S)_1$.

Theorem 1.74.

$$\mathcal{L}(\Gamma^v) = \mathcal{W}(\Gamma^h) = \frac{1}{D(h)}.$$

The modulus of S rel the boundaries $(\partial S)_0$ and $(\partial S)_1$ is defined as the above extremal length:

$$\operatorname{mod}((\partial S)_0, (\partial S)_1) = \mathcal{L}(\Gamma^v).$$

Remark. Physically, we can think of the pair $(\partial S)_0$ and $(\partial S)_1$ in S as an electric condensator. The harmonic function χ represents the potential of the electric field created by the uniformly distributed charge on $(\partial S)_1$. The Dirichlet integral $D(\chi)$ is the energy of this field. Thus, $mod((\partial S)_0, (\partial S)_1) = 1/D(\chi)$ is equal to the ratio of the charge to the energy, that is, to the capacity of the condensator.

4. Principles of the hyperbolic metric

4.1. Schwarz Lemma. In terms of the hyperbolic metric, the elementary Schwarz Lemma can be brought to a conformally invariant form that plays an outstanding role in holomorphic dynamics:

SCHWARZ LEMMA. Let $\phi:S\to S'$ be a holomorphic map between two hyperbolic Riemann surfaces. Then

either φ is a strict contraction, i.e., ||Dφ(z)|| < 1 for any z ∈ S, where the norm of the differential is evaluated with respect to the hyperbolic metrics of S and S';
or else, φ is a covering map, and then it is a local isometry: ||Dφ(z)|| = 1 for any z ∈ S.

PROOF. Given a point $z \in S$, let $\pi : (\mathbb{D}, 0) \to (S, z)$ and $\pi' : (\mathbb{D}, 0) \to (S', \phi(z))$ be the universal coverings of the Riemann surfaces S and S' respectively. Then ϕ can be lifted to a holomorphic map $\tilde{\phi} : (\mathbb{D}, 0) \to (\mathbb{D}, 0)$. By the elementary Schwarz Lemma, $|\tilde{\phi}'(0)| < 1$ or else $\tilde{\phi}$ is a conformal automorphism of \mathbb{D} (in fact, rotation). This yields the desired dichotomy for ϕ .

In particular, if $S \subset S'$ then $\rho_S \geq \rho_{S'}$ ("a smaller Riemann surface is more hyperbolic"). Moreover, if $S \neq S'$ then $d\rho_S(z) > d\rho_{S'}(z)$ for any $z \in S$.

COROLLARY 1.75. Let $S \subset S'$ be a nest of two hyperbolic Riemann surfaces, $S \neq S'$, and let $f : S \to S'$ be an (unramified) covering map. Then for $z \in S$ we have

$$||Df(z)||_{S'} > 1,$$

where the norm of Df(z) is evaluated in the hyperbolic metric of S' (in both the domain and the target).

PROOF. Since f is a local isometry from the hyperbolic metric of S to that of S', we have

(4.1)
$$\|Df(z)\|_{S'} = \frac{d\rho}{d\rho'}(z),$$

and the desired estimate follows from the remark preceding this Corollary. \Box

4.2. Hyperbolic metric blows up near the boundary. For a domain $U \subset \overline{\mathbb{C}}$, let d(z) stand for the *spherical distance* from $z \in U$ to ∂U .

EXERCISE 1.76. Show that $d\rho_{\mathbb{D}^*}(z) = -\frac{|dz|}{|z|\log|z|}$.

LEMMA 1.77. Let **S** be a Riemann surface, $x \in \mathbf{S}$, and assume that the punctured surface $S = \mathbf{S} \setminus \{x\}$ is hyperbolic with the hyperbolic metric ρ . Then

$$d\rho(z) \asymp -\frac{|dz|}{|z|\log|z|},$$

where z is a local coordinate on **S** with z(x) = 0.

PROOF. By Proposition 1.12, a standard cusp \mathbb{H}_h/\mathbb{Z} is isometrically embedded into S so that its puncture corresponds to x. On the other hand, by means of the exponential maps $\mathbb{H} \to \mathbb{D}^*$, $z \mapsto e^{2\pi i z}$, the cusp \mathbb{H}_h/\mathbb{Z} is isometric to the punctured disk \mathbb{D}_r^* , $r = e^{-2\pi h}$, in the hyperbolic metric of \mathbb{D}^* . By the previous Excercise, the latter has the desired form in the plane coordinate of \mathbb{D}_r^* (which extends to a local coordinate on **S** near x). Hence it has the desired form in any other local coordinate on **S** near x.

PROPOSITION 1.78. For any hyperbolic plane domain $U \subset \hat{\mathbb{C}}$, there exists $\kappa = \kappa(U) > 0$ such that:

$$\frac{d\rho_U}{d\sigma}(z) \ge -\frac{\kappa}{d(z)\log d(z)}, \quad z \in U,$$

where σ is the spherical metric.

PROOF. Take some point $z \in U$, and find the closest to it point $a \in \partial U$. Since ∂U consists of at least three points, we can find two more points, $b, c \in \partial U$, such that the points a, b, c are ϵ -spearated on $\overline{\mathbb{C}}$, where $\epsilon > 0$ depends only on U. Let us consider the Möbius transformation ϕ that moves (a, b, c) to $(0, 1, \infty)$. By Exercise 1.5, these transformations are uniformly bi-Lipschitz in the spherical metric, which reduces the problem to the case when $(a, b, c) = (0, 1, \infty)$. But in this case, $\rho_U(z)$ dominates the hyperbolic metric on $\mathbb{U} = \mathbb{C} \setminus \{0, 1\}$, and the desired estimate follows from Lemma 1.77. EXERCISE 1.79. More generally, let **S** be a Reimann surface endowed with a conformal Riemannian metric σ , and let K be a compact subset of **S** such that $\mathbf{S} \setminus K$ is a hyperbolic Riemann surface with hyperbolic metric ρ . Then there exists a $\kappa = \kappa(\mathbf{S}, K) > 0$ such that

$$\frac{d\rho}{d\sigma}(z) \ge -\frac{\kappa}{d(z)\log d(z)}, \quad z \in \mathbf{S} \smallsetminus K,$$

where $d(z) = \operatorname{dist}(z, K)$.

4.3. Normal families and Montel's Theorem. Let U be a Riemann surface, and let $\mathcal{M}(U)$ be the space of meromorphic functions $\phi: U \to \overline{\mathbb{C}}$. Supply the *target* Riemann sphere $\overline{\mathbb{C}}$ with the spherical metric d_s and the space $\mathcal{M}(U)$ with the topology of uniform convergence on compact subsets of U. Thus $\phi_n \to \phi$ if for any compact subset $K \subset U$, $d_s(\phi_n(z), \phi(z)) \to 0$ uniformly on K. Since locally uniform limits of holomorphic functions are holomorphic, $\mathcal{M}(U)$ is *closed* in the space C(U) of continuous functions $\phi: U \to \overline{\mathbb{C}}$ (endowed with the topology of uniform convergence on compact subsets of U).

EXERCISE 1.80. Endow $\mathcal{M}(U)$ with a metric compatible with the above convergence that makes $\mathcal{M}(U)$ a complete metric space.

It is important to remember that the target should be supplied with the *spherical* rather than *Euclidean* metric even if the original family consists of *holomorphic* functions. In the limit we can still obtain a meromorphic function, though of a very special kind:

EXERCISE 1.81. Let $\phi_n : U \to \mathbb{C}$ be a sequence of holomorphic functions converging to a meromorphic function $\phi : U \to \overline{\mathbb{C}}$ such that $\phi(z) = \infty$ for some $z \in U$. Then $\phi(z) \equiv \infty$, and thus $\phi_n(z) \to \infty$ uniformly on compact subsets of U.

A family of meromorphic functions on U is called *normal* if it is precompact in $\mathcal{M}(U)$.

EXERCISE 1.82. Show that normality is the local property: If a family is normal near each point $z \in U$, then it is normal on U.

THEOREM 1.83 (Little Montel). Any bounded family of holomorphic functions is normal.

PROOF. It is because the derivative of a holomorphic function can be estimated via the function itself. Indeed by the Cauchy formula

$$|\phi'(z)| \le \frac{\max_{\zeta \in U} |\phi(\zeta)|}{\operatorname{dist}(z, \partial U)^2}.$$

Thus, if a family of holomorphic functions ϕ_n is uniformly bounded, their derivatives are uniformly bounded on compact subsets of U. By the Arzela-Ascoli Criterion, this family is precompact in the space C(U) of continuous functions. Since $\mathcal{M}(U)$ is closed in C(U), we see that the original family is precompact in the space $\mathcal{M}(U)$. \Box

EXERCISE 1.84. If a domain $U \subset \mathbb{C}$ is supplied with the Euclidean metric |dz|while the target $\overline{\mathbb{C}}$ is supplied with the spherical metric $|dz|/(1 + |z|^2)$, then the corresponding "ES norm" of the differential $D\phi(z)$ is equal to $|\phi'(z)|/(1 + |\phi(z)|^2)$, $z \in U$. Show that a family of meromorphic functions $\phi_n : U \to \overline{\mathbb{C}}$ is normal if and only if the ES norms $||D\phi_n(z)||$ are uniformly bounded on compact subsets of U. EXERCISE 1.85. A sequence of holomorphic functions is normal if and only if one can extract from any subsequence a further subsequence which is either locally bounded or divergent (locally uniformly) to ∞ .

THEOREM 1.86 (Montel). If a family of meromorphic functions $\phi_n : U \to \mathbb{C}$ does not assume three values then it is normal.

PROOF. Since normality is a local property, we can assume that U is a disk. Let us endow it with the hyperbolic metric ρ . Let a, b, c be omitted values on $\overline{\mathbb{C}}$, and let ρ' be the hyperbolic metric on the thrice punctured sphere $\overline{\mathbb{C}} \setminus \{a, b, c\}$.

By the Schwarz Lemma, all the functions ϕ_n are contractions with respect to these hyperbolic metrics. By Proposition 1.78 (iii), the spherical metric is dominated by ρ' , so the ϕ_n are uniformly Lipschitz from metric ρ to the spherical metric. Normality follows.

THEOREM 1.87 (Refined Montel). Let $\{\phi_n : U \to \overline{\mathbb{C}}\}\$ be a family of meromorphic functions. Assume that there exist three meromorphic functions $\psi_i : U \to \overline{\mathbb{C}}\$ such that for any $z \in U$ and $i \neq j$ we have: $\psi_i(z) \neq \psi_j(z)$ and $\phi_n(z) \neq \psi_i(z)$. Then the family $\{\phi_n\}$ is normal.

PROOF. Let us consider the holomorphic family of Möbius transformations $h_z: \overline{\mathbb{C}} \to \overline{\mathbb{C}}$ depending on $z \in U$ as a parameter such that

$$h_z: (\psi_1(z), \psi_2(z), \psi_3(z)) \mapsto (0, 1, \infty).$$

Then the family of functions $\Phi_n(z) = h_z(\phi_n(z))$ omits value $0, 1, \infty$, and hence is normal by Theorem 1.86. It follows that the original family is normal as well. \Box

EXERCISE 1.88. Show that the theorem is still valid if the functions ψ_j are different but $\psi_i(z) = \psi_j(z)$ is allowed for some $z \in U$.

Given a family $\{\phi_n\}$ of meromorphic functions on U, we can define its *set of* normality as the maximal open set $F \subset U$ on which this family is normal.

4.4. Koebe Distortion Theorem. We are now going to discuss one of the most beautiful and important theorems of the classical geometric functions theory.

The inner radius $r_D(a) \equiv \operatorname{dist}(a, \partial D)$ of a pointed disk (D, a) is as the biggest round disk $\mathbb{D}(a, \rho)$ contained in D. The outer radius $R_D(a) \equiv \operatorname{dist}_H(a, \partial D)$ is the radius of the smallest disk $\overline{\mathbb{D}}(a, \rho)$ containing D. (If a = 0, we will simply write r_D and R_D .) The shape of a disk D around a is the ratio $R_D(a)/r_D(a)$.

THEOREM 1.89 (Koebe Distortion). Let $\phi : (\mathbb{D}, 0) \to (D, a)$ be a conformal isomorphism, and let $k \in (0, 1)$, $D_k = \phi(\mathbb{D}_k)$. Then there exist constants C = C(k) and L = L(k) (independent of a particular ϕ !) such that

(4.2)
$$\frac{|\phi'(z)|}{|\phi'(\zeta)|} \le C(k) \text{ for all } z, \zeta \in \mathbb{D}_k$$

and

(4.3)
$$L(k)^{-1}|\phi'(0)| \le r_{D_k,a} \le R_{D_k}(a) \le L(k) |\phi'(0)|$$

In particular, the inner radius of the image $\phi(\mathbb{D})$ around a is bounded from below by an absolute constant times the derivative at the origin:

(4.4)
$$r_{\phi(D)}(a) \ge \rho |\phi'(0)| > 0.$$

The expression in the left-hand side of (4.2) is called the *distortion* of ϕ . Thus, estimate (4.2) tells us that the function ϕ restricted to \mathbb{D}_k has a *uniformly bounded distortion* (depending on κ only). Estimate (4.3) tells that the shape of the domain D_k around *a* is uniformly bounded. This shape is also called the *dilatation* of *h* on \mathbb{D}_{κ} . So, univalent functions have *uniformly bounded dilatation* on any disk \mathbb{D}_{κ} . Note that since *any* proper topological disk in \mathbb{C} can be uniformized by \mathbb{D} , there could be no possible bounds on the distortion and dilatation of ϕ in the whole unit disk \mathbb{D} . However, once the disk is slightly shrunk, the bounds appear!

The Koebe Distortion Theorem is equivalent to the normality of the space of normalized univalent functions:

THEOREM 1.90. The space \mathcal{U} of univalent functions $\phi : (\mathbb{D}, 0) \to (\mathbb{C}, 0)$ with $|\phi'(0)| = 1$ is compact (in the topology of uniform convergence on compact subsets of \mathbb{D}).

Let us make a simple but important observation:

LEMMA 1.91. Let $\phi : (\mathbb{D}, 0) \to (\mathbb{C}, 0)$ be a univalent function normalized so that $|\phi'(0)| = 1$. Then the image $\phi(\mathbb{D})$ cannot contain the whole unit circle \mathbb{T} .

PROOF. Otherwise the inverse map ϕ^{-1} would be well defined on some disk \mathbb{D}_r with r > 1, and the Schwarcz Lemma would imply $|D\phi^{-1}(0)| \leq 1/r < 1$, contrary to the normalization assumption.

Proof of Theorem 1.90. By Lemma 1.91, for any $\phi \in \mathcal{U}$ there is a $\theta \in \mathbb{R}$ such that the rotated function $e^{i\theta}\phi$ does not assume value 1. Since the group of rotation is compact, it is enough to prove that the space $\mathcal{U}_0 \subset \mathcal{U}$ of univalent functions $\phi \in \mathcal{U}$ which do not assume value 1 is compact.

Let us puncture \mathbb{D} at the origin, and restrict all the functions $\phi \in \mathcal{U}_0$ to the punctured disk \mathbb{D}^* . Since all the ϕ are univalent, they do not assume value 0 in \mathbb{D}^* . By the Montel Theorem, the family \mathcal{U}_0 is normal on \mathbb{D}^* .

Let us show that it is normal at the origin as well. Take a Jordan curve $\gamma \subset \mathbb{D}^*$ around 0, and let Δ be the disk bounded by γ . Restrict all the functions $\phi \in \mathcal{U}_0$ to γ . By normality in \mathbb{D}^* , the family \mathcal{U}_0 is either uniformly bounded on γ , or admits a sequence which is uniformly going to ∞ . But the latter is impossible since all the curves $\phi_n(\gamma)$ intersect the interval [0, 1] (as they go once around 0 and do not go around 1). Thus, the family \mathcal{U}_0 is uniformly bounded on γ . By the Maximum Principle, it is is uniformly bounded, and hence normal, on Δ as well.

Thus, the family \mathcal{U}_0 is precompact. What is left, is to check that it contains all limiting functions. By the Argument Principle, limits of univalent functions can be either univalent or constant. But the latter is not possible in our situation because of normalization $|\phi'(0)| = 1$. \Box

EXERCISE 1.92. (a) Show that a family \mathcal{F} of univalent functions $\phi : \mathbb{D} \to \mathbb{C}$ is precompact in the space of all univalent functions if and only if there exists a constant C > 0 such that

$$|\phi(0)| \leq C$$
 and $C^{-1} \leq |\phi'(0)| \leq C$ for all $\phi \in \mathcal{F}$.

b) Let (Ω, a) be a pointed domain in \mathbb{C} and let C > 0. Consider a family \mathcal{F} of univalent functions $\phi : \Omega \to \mathbb{C}$ such that $|\phi(a)| \leq C$. Show that this family is normal if and only if there exists $\rho > 0$ such that each function $\phi \in \mathcal{F}$ omits some value ζ with $|\zeta| < \rho$.

Proof of the Koebe Distortion Theorem. Compactness of the family \mathcal{U} immediately yields that functions $\phi \in \mathcal{U}$ and their derivatives are uniformly bounded on any smaller disk \mathbb{D}_k , $k \in (0, 1)$. Combined with the fact that all functions of \mathcal{U} are univalent, compactness also implies a lower bound on the inner radius $r_{\phi(D_k)}$ and on the derivative $\phi'(z)$ in \mathbb{D}_k . These imply estimates (4.2) and (4.3) on the distortion and shape by normalizing a univalent function $\phi : \mathbb{D} \to \mathbb{C}$, i.e., considering

$$\tilde{\phi}(z) = \frac{\phi(z) - a}{\phi'(0)} \in \mathcal{U}.$$

(Note that this normalization does not change either distortion of the function or its dilatation.)

Estimate (4.4) is an obvious consequence of the left-hand side of (4.3). \Box

We have given a qualitative version of the Koebe Distortion Theorem, which will be sufficient for all our purposes. The quantitative version provides sharp constants C(k), L(k), and ρ , all attained for a remarkable extremal *Koebe function* $f(z) = z/(1-z)^2 \in \mathcal{U}$. The sharp value of the constant ρ is particularly remarkable:

KOEBE 1/4-THEOREM. Let $\phi : (\mathbb{D}, 0) \to (\mathbb{C}, 0)$ be a univalent function with $\phi'(0) = 1$. Then $\phi(\mathbb{D}) \supset \mathbb{D}_{1/4}$, and this estimate is attained for the Koebe function.

We will sometimes refer to the Koebe 1/4-Theorem rather than its qualitative version (4.4), though as we have mentioned, the sharp constants never matter for us.

EXERCISE 1.93. Find the image of the unit disk under the Koebe function.

Let us finish with an invariant form of the Koebe Distortion Theorem:

THEOREM 1.94. Consider a pair of conformal disks $\Delta \subseteq D$. Let $\operatorname{mod}(D \setminus \Delta) \ge \mu > 0$. Then any univalent function $\phi : D \to \mathbb{C}$ has a bounded (in terms of μ) distortion on Δ :

$$\frac{|\phi'(z)|}{|\phi'(\zeta)|} \le C(\mu) \text{ for all } z, \zeta \in \Delta.$$

The proof will make use of one important property of the modulus of an annulus: if an annulus is getting pinched, then its modulus is vanishing:

LEMMA 1.95. Let $0 \in K \subset \mathbb{D}$, where K is compact. If

$$mod(D \smallsetminus K) \ge \mu > 0$$

then $K \subset \mathbb{D}_k$ where the radius $k = k(\mu) < 1$ depends only on μ .

PROOF. Assume there exists a sequence of compact sets K_i satisfying the assumptions but such that $R_i \to 1$, where R_i is the outer radius of K_i around 0. Let us uniformize $D \setminus K_i$ by a round annulus, $h_i : \mathbb{A}(\rho_i, 1) \to \mathbb{D} \setminus K_i$. Then $\rho_i \leq \rho \equiv e^{-\mu} < 1$. Thus, the maps h_i are well-defined on a common annulus $A = \mathbb{A}(\rho, 1)$. By the Little Montel Theorem, they form a normal family on A, so that we can select a converging subsequence $h_{i_n} \to h$.

Let $\gamma \subset A$ be the equator of A. Then $h(\gamma)$ is a Jordan curve in \mathbb{D} which separates the sets K_{i_n} (with sufficiently big n) from the unit circle - contradiction.

Remark. The extremal compact sets in the above lemma (minimizing k for a given μ) are the straight intervals $[0, ke^{i\theta}]$.

Proof of Theorem 1.94. Let us uniformize D by the unit disk, $h : \mathbb{D} \to D$, in such a way that $h(0) \in \Delta$. Let $\tilde{\Delta} = h^{-1}\Delta$ and $\tilde{\phi} = \phi \circ h$. By Lemma ??, $\tilde{\Delta} \subset \mathbb{D}_k$, where $k = k(\mu) < 1$. By the Koebe Theorem, the distortion of the functions hand $\tilde{\phi}$ on $\tilde{\Delta}$ is bounded by some constant C = C(k). Hence the distortion of ϕ is bounded by C^2 . \Box

We will often use the following informal formulation of Theorem 1.94: "If $\phi : D \to \mathbb{C}$ is a univalent function and $\Delta \subset D$ is well inside D, then ϕ has a bounded distortion on Δ ". Or else: "If a univalent function $\phi : \Delta \to \mathbb{C}$ has a definite space around Δ , then it has a bounded distortion on Δ ".

Let us summarize some of the above results in a very useful comparison of the derivative of a univalent function with the inner radius of its image:

COROLLARY 1.96. For any univalent function $\phi : (\mathbb{D}, 0) \to (D, a)$, we have:

$$r_D(A) \le |\phi'(0)| \le 4 r_D(a).$$

PROOF. The left-hand side estimate follows from Lemma 1.91 by normalizing ϕ . The Koebe 1/4-Theorem implies the right-hand side one: $r_D(a) \geq \frac{1}{4} |\phi'(0)|$. \Box

4.5. Hyperbolic metric on simply connected domains. For simply connected plane domains, the hyperbolic metric can be very well controlled:

LEMMA 1.97. Let $D \subset \mathbb{C}$ be a conformal disk endowed with the hyperbolic metric ρ_D . Then

$$\frac{1}{4} \frac{|dz|}{\operatorname{dist}(z, \partial D)} \le d\rho_D(z) \le \frac{|dz|}{\operatorname{dist}(z, \partial D)}$$

Remark. Of course, particular constants in the above estimates will not matter for us.

PROOF. Let $r = \text{dist}(z, \partial D)$; then $\mathbb{D}(z, r) \subset D$. Consider a linear map $h : \mathbb{D} \to \mathbb{D}(z, r)$ as a map from \mathbb{D} into D. By the Schwarz Lemma, it contracts the hyperbolic metric. Hence

$$d\rho_D(z) \le h_*(d\rho_{\mathbb{D}}(0)) = h_*(|d\zeta|) = |dz|/r.$$

To obtain the opposite inequality, consider the Riemann mapping $\psi : (\mathbb{D}, 0) \to (D, z)$. By definition of the hyperbolic metric,

$$d\rho_D(z) = \psi_*(d\rho_{\mathbb{D}}(0)) = \psi_*(|d\zeta|) = \frac{|dz|}{|\psi'(0)|}.$$

But by the Koebe 1/4-Theorem, $r \leq |\psi'(0)|/4$, so that $d\rho_D(z) \geq |dz|/4r$.

The 1/d-metric on a plane domain U is a continuous Riemannian metric with the length element |dz|/d(z). The previous lemma tells us that the hyperbolic metric on a simply connected domain is equivalent to the 1/d-metric.

4.6. Definitive Schwarz Lemma.

4.6.1. *Definitive contraction*. Montel's compactness allows one to turn the Schwarz Lemma into a definitive form. Let us begin with an elementary version:

LEMMA 1.98. Let $\phi : (\mathbb{D}, 0) \to (\mathbb{D}, 0)$ be a holomorphic map that omits a point z with $|z| \leq \rho < 1$. Then $|\phi'(0)| \leq \sigma(\rho) < 1$.

PROOF. By the Little Montel Theorem and the Hurwitz Theorem, the space of maps in question is compact (for a given $\rho < 1$). Hence the Schwarz Lemma becomes definitive on this space.

Now the Uniformization Theorem immediately turns this elementary fact into an invariant geometric property:

LEMMA 1.99. Let $\phi : (S, a) \to (S', a')$ be a holomorphic map between hyperbolic Riemann surfaces. If dist $(a', \partial(\phi S)) \leq \rho$ then $\|\mathbb{D}\phi(a)\| \leq \sigma(\rho) < 1$, where the norm is evaluated with resepct to the hyperbolic metrics.

PROOF. Following the proof of the Schwarz Lemma given in §4.1, lift ϕ to a holomorphic map $\tilde{\phi} : (\mathbb{D}, 0) \to (\mathbb{D}, 0)$. By assumption, there is a point $z \in \partial(\phi S)$ such that $\operatorname{dist}_{S'}(a', z) \leq \rho$. Then $\tilde{\phi}$ omits a point \tilde{z} such that

$$\operatorname{dist}_{\mathbb{D}}(\tilde{z}, 0) = \operatorname{dist}_{S'}(z, a') \le \rho.$$

By Lemma 1.98,

$$\|D\phi(a)\| = |\tilde{\phi}'(0)| \le \sigma(\rho) < 1.$$

COROLLARY 1.100. For a nest of two hyperbolic Riemann surfaces $S \subset S'$ and any $z \in S$ such that $\operatorname{dist}_{S'}(z, \partial S) \leq \rho$ we have:

$$\frac{d\rho'}{d\rho}(z) \le \sigma(\rho) < 1.$$

COROLLARY 1.101. Let $S \subset S'$ be a nest of two hyperbolic Riemann surfaces, and let $f: S \to S'$ be an (unramified) covering map. Then for $z \in S$ we have

$$||Df(z)||_{S'} \ge \lambda(\rho) > 1$$
, provided dist_{S'} $(z, \partial S) \le \rho$

PROOF. It follows from (4.1) and Corollary 1.100.

EXERCISE 1.102. Let $A' \supset A \supset \mathbb{T}$ be a nest of two annuli symmetric with respect to the unit circle \mathbb{T} such that

$$0 < \mu' \le \operatorname{mod} A' \le 1/\mu', \quad 0 < \mu \le \operatorname{mod} A \le 1/\nu.$$

Then for any $z \in \mathbb{T}$ we have:

$$\frac{d\rho'}{d\rho}(z) \le \sigma(\mu, \mu') < 1.$$

Moreover, if $g: A \to A'$ is a holomorphic double covering then

$$\|Dg(z)\|_{A'} \ge \lambda(\mu') > 1.$$

4.7. Hyperbolic metric on the thick part.

 \square

4.8. Carathéodory convergence. Let us consider the space \mathcal{D} of all pointed conformal disks (D, a) in the complex plane. This space can be endowed with a natural topology called *Carathéodory*. We will describe it it terms of convergence:

DEFINITION 1.103. A sequence of pointed disks $(D_n, a_n) \in \mathcal{D}$ converges to a disk $(D, a) \in \mathcal{D}$ if:

(i) $a_n \to a$;

(ii) Any compact subset $K \subset D$ is eventually contained in all disks D_n :

$$\exists N: K \subset D_n \ \forall n \ge N;$$

(iii) If U is a topological disk contained in infinitely many domains D_n then U is contained in D.

Note that this definition allows one to pinch out big bubbles from the domains D_n (see Figure ...).

EXERCISE 1.104. a) Define a topology on \mathcal{D} that generates the Carathéodory convergence.

b) Show that if ∂D_n converges to ∂D in the Hausdorff metric then the disks D_n converge to D in the Carathéodory sense.

The above purely geometric definition can be reformulated in terms of the uniformizations of the disks under consideration. Let us uniformize any pointed disk $(D, a) \in \mathcal{D}$ by a conformal map $\phi : \mathbb{D} \to D$ positively normalized so that $\phi(0) = a$ and $\phi'(0) > 0$.

PROPOSITION 1.105. A sequence of pointed disks $(D_n, a) \in \mathcal{D}$ converges to a pointed disk $(D, a) \in \mathcal{D}$ if the corresponding sequence of normalized uniformizations $\phi_n : \mathbb{D} \to D_n$ converges to the positively normalized uniformization $\phi : \mathbb{D} \to D$ uniformly on compact subsets of \mathbb{D} .

PROOF. Assuming $\phi_n \to \phi$, let us check properties (i)-(iii) of Definition 1.103. The first one is obvious. To verify (ii), take a compact subset K of D. Then $\phi(\mathbb{D}_r) \supset K$ for some r < 1. Hence $\operatorname{dist}(\phi(\mathbb{T}_r), K) > 0$ and the curve $\phi : \mathbb{T}_r \to \mathbb{C}$ has winding number 1 around any point of K. Since $\phi_n \to \phi$ uniformly on \mathbb{T}_r , eventually all the curves $\phi_n : \mathbb{T}_r \to \mathbb{C}$ have winding number 1 around all points of K. Then $\phi_n(\mathbb{D}_r) \supset K$.

Let us now veryfy (iii). It is enough to check that any disk $V \Subset U$ is contained in D. For such a disk, we have: $\operatorname{mod}(D_n, V) \ge \mu > 0$ for all n. Let $W_n = \phi_n^{-1}(V_n)$. By the conformal ivariance, $\operatorname{mod}(\mathbb{D}, W_n) \ge \mu$ as well. Hence $W_n \subset \mathbb{D}_{1-2\epsilon}$ for some $\epsilon > 0$ (by Lemma 1.95 or 1.62). Using conformal invariance of moduli and Lemma 1.62 once again, we conclude that $\operatorname{dist}(\phi_n(\mathbb{T}_{1-\epsilon}), V) \ge \rho > 0$. Since eventually $|\phi(z) - \phi_n(z)| < \rho/2$ on $\mathbb{T}_{1-\epsilon}$, the curve $\phi : \mathbb{T}_{1-\epsilon} \to \mathbb{C}$ has the same winding number around any point of V as $\phi_n : \mathbb{T}_{1-\epsilon} \to \mathbb{C}$, and the latter is equal to 1 (for n sufficiently big). Hence $\phi(\mathbb{D}_{1-\epsilon}) \supset V$, as required.

Vice versa, assume $(D_n, a_n) \to (D, a)$ in the Carathéodory topology. By Property (ii) of Definition 1.103, the domains D_n eventually contain the disc $\mathbb{D}(a, r_D(a)/2)$ (where $r_D(a)$ stands for the inner radius of the domain D with respect to $a \in D$, see §4.4). By Corollary 1.96, $|\phi_n(0)| \ge r_D(a)/2$.

On the other hand, by Property (iii), the domains D_n do not eventually contain the disc $\mathbb{D}(a, 2r_D(a))$. By Corollary 1.96, $|\phi'_n(0)| \leq 8r_D(a)$.

Thus, $|\phi_n(0)| \approx 1$. By the Koebe Distortin Theorem (see Exercise 1.92), the family $\{\phi_n\}$ is precompact in the space of univalent functions. But by the first part of this lemma, any limit function $\phi = \lim \phi_{n(k)}$ is the positively normalized uniformization of (D, a) by $(\mathbb{D}, 0)$. It follows that the ϕ_n converge to this uniformization.

For $r \in (0, 1)$, let \mathcal{D}_r stand for the family of pointed disks $(D, a) \in \mathcal{D}$ with

 $r \le r_D(a) \le 1/r.$

COROLLARY 1.106. The space \mathcal{D}_r is compact.

PROOF. Let $\phi_D : (\mathbb{D}, 0) \to (D, a)$ be the positively normalized uniformization of D. By Corollary 1.96, $r \leq \phi'_D(0) \leq 4/r$ By the Koebe Distortion Theorem (see Exercise 1.92), the family of univalent functions ϕ_D , $D \in \mathcal{D}_r$, is compact. By Proposition 1.105, the space \mathcal{D}_r is compact as well.

With thes notions in hands, we can define convergence of a sequence of functions $\psi_n : (D_n, a_n) \to (\mathbb{C}, b_n)$ on varying domains. Namely, the functions ψ_n converge to a function $\psi : (D, a) \to (\mathbb{C}, b)$ if the pointed domains (D_n, a_n) converge to (D, a), and $\psi_n \to \psi$ uniformly on compact subsets of D. (This makes sense since for any $K \subseteq D$, all but funitely many functions ψ_n are well defined on K.)

REMARK 1.7. We will often supress mentioning of the base points a_n , as long as it would not lead to a confusion.

We can now naturally define *normality* of a family of functions $\psi_n : D_n \to \mathbb{C}$ with varying domains of definition. In cases when the D_n converge to some domain D, we also say that "the family $\{\psi_n\}$ is normal on D".

The statement of the Montel Theorem admits an obvious adjustment in this setting: If the family of domains D_n is Carath'eodory precompact and the functions $\psi_n : D_n \to \hat{\mathbb{C}}$ omit three values on the Riemann sphere, then the family $\{\phi\}_n$ is normal.

5. Uniformization Theorem

5.1. Statement. The following theorem of Riemann and Koebe is the most fundamental result of complex analysis:

THEOREM 1.107. Any simply connected Riemann surface is conformally equivalent to either the Riemann sphere $\hat{\mathbb{C}}$, or to the complex plane \mathbb{C} , or the unit disk \mathbb{D} .

5.2. Classification of Riemann surfaces. Consider now any Riemann surface S. Let $\pi : \hat{S} \to S$ be its universal covering. Then the complex structure on S naturally lifts to \hat{S} turning S into a simply connected Riemann surface which holomorphically covers S. Thus, we come up with the following classification of Riemann surfaces:

THEOREM 1.108. Any Riemann surface S is conformally equivalent to one of the following surfaces:

• The Riemann sphere $\hat{\mathbb{C}}$ (spherical case);

- The complex plane \mathbb{C} , or the punctured plane \mathbb{C}^* , or a torus \mathbb{T}^2_{τ} , $\tau \in \mathbb{H}$ (parabolic case);
- The quotient of the hyperbolic plane \mathbb{H}^2 modulo a discrete group of isometries (hyperbolic case).

Thus, *any* Riemann surface comes endowed with one of the three geometries described in §??: projective, affine, or hyperbolic.

5.3. Uniformization of simply connected plane domains. For dynamical applications, we will not need the full strength of the Uniformization Theorem: only uniformization of plane domains will be relevant. Let us start with the most classical case:

RIEMANN MAPPING THEOREM. Any simply connected domain $D \subset \mathbb{C}$ whose complement contains more than one point is conformally equivalent to the unit disk \mathbb{D} . The conformal isomorphism $\phi : \mathbb{D} \to D$ is unique up to pre-composition with a Möbius transformation $M \in \operatorname{Aut}(\mathbb{D})$.⁷

PROOF. The uniqueness part is obvious, so let us focus on the existence.

First, notice that D can be conformally mapped onto a bounded domain in \mathbb{C} . Indeed, since $\hat{\mathbb{C}} \setminus D$ contains more than one point and D is simply connected, $\hat{\mathbb{C}} \setminus D$ is in fact a continuum. Let us take two points $a_1, a_2 \in \hat{\mathbb{C}} \setminus D$, and move them to $0, \infty$ by a Möbius transformation. This turns D into a domain in \mathbb{C}^* .⁸ Since Dis simply connected, the square root map $Q : z \mapsto \sqrt{z}$ has a single-valued branch on D. Applying it, we obtain a domain whose complement has non-empty interior (the image of the other branch of Q). Moving ∞ to this complement by a Möbius transformation, we make D a bounded domain in \mathbb{C} .

Let us now take a point $a \in D$, and consider the space \mathcal{C} of conformal embeddings $\psi : D \to \mathbb{D}$ normalized so that $\psi(a) = 0$. Note that $\mathcal{C} \neq \emptyset$ since D can be embedded into \mathbb{D} by an affine map. By the Little Montel Theorem, \mathcal{C} is normal. Hence we can find a conformal map $\psi_0 \in \mathcal{C}$ that maximizes the derivative $|\psi'(a)|$ over the class \mathcal{C} .

We claim that ψ_0 conformally maps D onto \mathbb{D} . The only issue is surjectivity. Assume there is a point $a \in \mathbb{D} \setminus \psi_0(D)$. Let $B : (\mathbb{D}, 0) \to (\mathbb{D}, 0)$ be a double branched covering with critical point at a.

EXERCISE 1.109. Write down B explicitly.

Since $\psi_0(D)$ is simply connected, there is a single-valued branch $B^{-1}: \psi_0(D) \to \mathbb{D}$. By the Schwarz Lemma, |B'(0)| < 1, and hence the embedding $B^{-1} \circ \psi_0 : (D, a) \to (\mathbb{D}, 0)$ has a bigger derivative at a than ψ – contradiction. \Box

5.4. Thrice punctured sphere and modular function λ . Let us now consider the case of the biggest hyperbolic plane domain, the thrice punctured sphere $\mathbb{U} = \mathbb{C} \setminus \{0, 1\}^9$. In this case, there is a simple explicit construction of the universal covering. Namely, let us consider an ideal triangle Δ in the hyperbolic

⁷For instance, it is uniquely determined by its value at 0 and the image of the tangent vector $1 \in T_0 \mathbb{D}$ under $D\phi(0)$.

⁸We will keep notation D for various domains conformally equivalent to D.

 $^{^{9}\}text{Note}$ that all thrice punctured spheres are equivalent under the action of the Möbius group Möb($\hat{\mathbb{C}}).$

plane, that is, the geodesic triangle with vertices on the absolute¹⁰ (see Figure ??). By the Riemann Mapping Theorem, it can be conformally mapped onto the upper half plane \mathbb{H} so that its vertices go to the points 0, 1 and ∞ . By the Schwarz Reflection Principle, this conformal map can be extended to the three symmetric ideal triangles obtained by reflection of Δ in its edges. Each of these symmetric rectangles will be mapped onto the lower half-plane \mathbb{H}^- . Then we can extend this map further to the six symmetic rectangles each of which will be mapped onto \mathbb{H} again, etc. Proceeding in this way, we obtain the desired universal covering $\lambda : \mathbb{D} \to \mathbb{U}$ called a *modular function*.

EXERCISE 1.110. Verify the following properties:

a) The union of these triangles tile the whole disk \mathbb{D} ;

b) The modular function λ is the desired universal covering;

c) Its group of deck transformations is the congruent group Γ_2 , that is, the subgroup of $PSL(2,\mathbb{Z})$ consisting of matrices congruent to I mod 2.

5.5. Annuli. We will now pass to non-plane domains beginning with annuli:

PROPOSITION 1.111. Let $A \in S$ be a topological annulus on a Riemann surface S with piecewise smooth boundary. Then A is conformally equivalent to a standard annulus $\mathbb{A}(r, R)$.

PROOF. Let us call one of the boundary components of A "inner", $\partial^i A$, and the other one "outer", $\partial^o A$. Let us consider the "harmonic measure" of the outer component, i.e. a harmonic function u(z) on A vanishing on $\partial^i A$ and $\equiv 1$ on $\partial^0 A$ (see §7.8). Let u^* be its harmonic conjugate, This function is not single valued, but rather gets changed by the period

$$p = \int_{\gamma} *du$$

under the monodromy along a non-trivial cycle γ in A (see §7.1). Hence the holomorphic function

$$f = \exp\frac{2\pi}{p}(u + iu*)$$

is single valued. Moreover, it properly maps A onto the round annulus $\mathbb{A}(1, e^{1/p})$ and has degree one (since f homeomorphically maps the equipotentials of A onto the round circles. The conclusion follows.

5.6. Simply connected domains.

PROPOSITION 1.112. Let $D \subseteq S$ be a simply connected domain on a Riemann surface S with piecewise smooth boundary. Then D is conformally equivalent to the unit disk \mathbb{D} .

PROOF. Take a base point $z_0 \in D$, and let D_{ϵ} be a coordinate disk of radius $\epsilon > 0$ centered at z_0 . Then $U \smallsetminus D_{\epsilon}$ is a topological annulus with piecewise smooth bundary, so by Proposition 1.111 there is a conformal map $\phi_{\epsilon} : D \smallsetminus D_{\epsilon} \to \mathbb{A}(r(\epsilon), 1)$ onto a round annulus. By the Little Montel Theorem, the family of maps ϕ_{ϵ} is normal on $D \smallsetminus \{z_0\}$.¹¹ Let us select a converging subsequence $\phi_{\epsilon_k} \to \phi$ as $\epsilon_k \to 0$,

¹⁰Note that all these triangles are equivalent under the action of $PSL(2, \mathbb{R})$.

 $^{^{11}\}mathrm{On}$ normality with varying domains of definition, see §4.8.

where $\phi : D \setminus \{z_0\} \to \mathbb{D}$ is a holomorphic map. By Removability of isolated singularities, ϕ holomorphically extends through z_0 .

Let us show that $\phi : D \to \mathbb{D}$ is proper. It is sufficient to check that for any $r \in (0,1)$, the preimage $\phi^{-1}(\mathbb{D}_r)$ is compactly contained in D. Indeed, take any $R \in (r,1)$. By invariance of the modulus,

$$\operatorname{mod}(\phi_{\epsilon}^{-1}(\mathbb{A}(R,1)) = \frac{1}{2\pi}\log\frac{1}{R} > 0$$
, for any $\epsilon > 0$ sufficiently small.

By Lemma 1.95, $\operatorname{dist}(\phi_{\epsilon}^{-1}(\mathbb{T}_R), \partial D) \ge \rho > 0$ for some $\rho = \rho(R) > 0$. Letting $\epsilon \to 0$, we conclude that $\operatorname{dist}(\phi^{-1}(\mathbb{T}_r), \partial D) \ge \rho > 0$, and properness of ϕ follows.

So, ϕ has a well defined degree. Since degree is stable under perturbations, deg $\phi = \deg \phi_{\epsilon_k}$ for all k sufficiently large. Thus deg $\phi = 1$, and hence ϕ is a conformal isomorphism.

6. Carathéodory boundary

6.1. Prime ends. Let us consider a conformal disk $D \in \hat{\mathbb{C}}$. Its *cross-cut* is a path $\gamma : [0,1] \to \overline{D}$ such that $\operatorname{int} \gamma := \gamma(0,1) \subset D$ while $\partial \gamma := \gamma\{0,1\} \subset \partial D$. It divides D into two domains, $D^+(\gamma)$ and $D^-(\gamma)$.

A sequence $\bar{\gamma}$ of disjoint crass-cuts γ_n , $n \in \mathbb{N}$, form a *nest* if for all $n \in \mathbb{Z}_+$, int γ_n separates int γ_{n-1} from int γ_{n+1} in D. In this case we let $D_n^+(\bar{\gamma})$ be the component of $D \smallsetminus \gamma_n$ containing γ_{n+1} (and hence all further γ_i , i > n).

A nest of cross-cuts is *shrinking* if length(γ_n) $\rightarrow 0$.

Let us say that a nest $\bar{\gamma}'$ is *n*-subordinated to a nest $\bar{\gamma}$ if there exists m_0 such that int $\gamma'_m \subset D_n^+(\bar{\gamma})$ for all $m \geq m_0$. If this happens for every $n \in \mathbb{N}$, we say that $\bar{\gamma}'$ is subordinated to $\bar{\gamma}$. Two nests $\bar{\gamma}$ and $\bar{\gamma}'$ are *equivalent* if each of them is subordinated to the other one.

Now we are ready to give the main definition: a *prime end* P of D is the equivalence class of shrinking nests of cross-cuts.

Let $\partial^C D$ denote the space of all prime ends (the Carathéodory boundary of D), and $\mathrm{cl}^C D = D \cup \partial^C D$ (the Carathéodory compactification of D). Endow $\mathrm{cl}^C D$ with the following topology. Let us consider a prime end $x \in \partial^C D$ represented by a shrinking nest $\bar{\gamma}$. Given an $n \in \mathbb{N}$, let $\mathcal{U}_n(x)$ be the union of $D_n^+(\bar{\gamma})$ and all the prime-ends that are *n*-subordinated to $\bar{\gamma}$. This is the base of neighborhoods of x.

EXERCISE 1.113. Show that $\operatorname{cl}^C \mathbb{D} \approx \overline{\mathbb{D}}$.

EXERCISE 1.114. Let us consider two conformal disks $D, D' \subset \hat{\mathbb{C}}$ and a homeomorphism $h : \overline{D} \to \overline{D}'$. Then $h : D \to D'$ continuously extends to a homeomorphism $\hat{h} : cl^C D \to cl^C D'$.

The *impression* of the prime end P is defined as

$$I(P) = \bigcap_{n} \operatorname{cl} D_{n}^{+}(\bar{\gamma})$$

(which is easily seen to be independent of the choice of the nest $\bar{\gamma}$ representing P). When the impression is a singleton, $I(P) = \{x\}$, the prime end can be identified with the corresponding point $x \in \partial D$ (like in Exersice 1.113). **6.2. Extension of the Riemann mapping.** We are ready to formulate a fundamental result of the classical boundary values theory:

FIRST CARATHÉODORY THEOREM. The Riemann mapping $\phi : \mathbb{D} \to D$ extends to a homeomorphism $\hat{\phi} : \overline{\mathbb{D}} \to \text{cl}^C D$.

PROOF. Let $a \in \mathbb{T}$ and $P_{\mathbb{D}}(a)$ be the corresponding prime end of \mathbb{D} . It is represented by shrinking nests $\bar{\delta}$ of circular cross-cuts $\delta_r := \mathbb{T}(a, r) \cap \mathbb{D}$ around a (see Exercise 1.113).

By the Fatou Theorem, almost all images $\phi(\delta_r)$ are cross-cuts of D – let us call such δ_r (and also r) "good". They represent a class of equivalent nests, not necessarily shrinking. However, as we will see in a moment, some of them shrink, and these will represent the prime end of D corresponding to the image of a under $\hat{\phi}$.

It will be slightly more convenient to replace \mathbb{D} with the upper-half plane \mathbb{H} and to put a at the origin. Let us consider half-circles $S(r) = \mathbb{T}_r \cap \mathbb{H}$ around 0. We will show that there is a sequence of good radii $r_i \to 0$ such that the cross-cuts $\phi(S(r_i))$ of D shrink. To this end, let us consider half-annuli $A_r = \mathbb{A}(r/2, r) \cap \mathbb{H}$ viewed as rectangles whose horizontal sides are the semi-circles. Let \mathcal{F}_r be the horizontal foliation of A_r by the good half-circles S_ρ , $r/2 < \rho < r$. The extremal length of this foliation is equal to $1/ \mod A_r = \pi/\log 2$. By the conformal invariance, foliation $\phi(\mathcal{F}_r)$ has the same extremal length.

Let l_r be the minimal spherical length of the curves of $\phi(S_{\rho})$, $r/2 < \rho < r$. By definition of the extremal length,

$$\frac{l_r^2}{\operatorname{area}(\phi(A_r))} \le \mathcal{L}(\phi(\mathcal{F}_r)) = \frac{\pi}{\log 2},$$

where the "area" stands for the spherical area. Since $\operatorname{area}(\phi(A_r)) \to 0$ as $r \to 0$, we conclude that $l_r \to 0$ as well, which gives us the desired nest of shrinking cross-cuts.

So, we have constructed an extension $\hat{\phi} : \overline{\mathbb{D}} \to \operatorname{cl}^C D$. It easily follows from the definitions that $\hat{\phi}$ is injective and continuous. To check surjectivity, notice that any $x \in \partial^C D$ is the limit of some sequence $\zeta_n = \phi(z_n) \in D$. Selecting a limit point $a = \lim z_{n_k} \in \partial \mathbb{D}$, we see that $x = \hat{\phi}(a)$.

Hence ϕ is a homeomorphism.

Thus, the Carathédory boundary gets canonically identified with the ideal boundary of D, and we will freely use this identification in what follows.

6.3. Local connectivity and the Schönflis Theorem.

EXERCISE 1.115. Show that the Riemann map $\phi : \mathbb{D} \to D$ extends continuously to a point $a \in \partial \mathbb{D}$ if and only if the corresponding impression $I(\hat{\phi}(a))$ is a singleton.

The next classical theorem will motivate some central problems of holomorphic dynamics:

SECOND CARATHÉODORY THEOREM. The following properties are equivalent:

(i) The Riemann mapping $\phi : \mathbb{D} \to D$ extends to a continuous map $\overline{\mathbb{D}} \to \overline{D}$;

(ii) ∂D is locally connected;

(iii) $\hat{\mathbb{C}} \smallsetminus D$ is locally connected.

PROOF. (i) \implies (ii) by Exercise 0.7.

(ii) \implies (iii) by Exercise 0.10.

(iii) \implies (i). Assume ϕ does not admit a continuous extension to \mathbb{D} . Then there is a point $a \in \partial \mathbb{D}$ such that the corresponding prime end $\hat{\phi}(a)$ has a non-singleton impression $I = I(\hat{\phi}(a))$. Let us consider a nest of semi-circles δ_n shrinking to a whose images $\gamma_n := \phi(\delta_n)$ form a nest $\bar{\gamma}$ of cross-cuts representing the prime end $\hat{\phi}(a)$ (see the proof of the First Carathéodory Theorem). By selecting a subsequence, we can assume that the cross-cuts γ_n shrink to some point $y \in \partial D$.

Since I is not a singleton, diam $D_n^+(\bar{\gamma}) \not\rightarrow 0$. Hence there exist $\epsilon > 0$ and a sequence of points $\zeta_n = \phi(z_n) \in D_n^+(\bar{\gamma})$ such that $\operatorname{dist}(\zeta_n, \gamma_n) > \epsilon$. Let us connect z_n to 0 by the straight interval $[0, z_n]$; it crosses δ_n at some point b_n . As the distance $d(\phi(0), \phi(b_n))$ stays away from 0, we can assume it is bigger than ϵ as well.

Thus, both arcs, $\phi[0, b_n]$ and $\phi[b_n, z_n]$ must intersect the circle of radius $\epsilon/2$ around y (for n sufficiently big). Then there is a subarc

$$\omega_n \subset \mathbb{D}(y, \epsilon/2) \cap \phi[0, z_n] \subset D$$

with endpoint on this circle that crosses γ_n at a single point $\phi(b_n)$. This arc separates the endpoints of γ_n in $\mathbb{D}(y, \epsilon/2) \smallsetminus D$, contradicting local connectivity of $\hat{\mathbb{C}} \searrow D$ at y.

As a consequence, we obtain:

CONFORMAL SCHÖNFLIS THEOREM. Let $\gamma \subset \hat{\mathbb{C}}$ be a Jordan curve and D be a component of $\hat{\mathbb{C}} \smallsetminus \gamma$. Then the Riemann mapping $\phi : \mathbb{D} \to D$ extends to a homeomorphism $\overline{\mathbb{D}} \to \overline{D}$.

6.4. Proper Ends.

PROPOSITION 1.116. Let $S \subset S'$ be a nest of two Riemann surfaces. Let e and e' be tame ends of S and S' respectively such that the embedding $i : S \hookrightarrow S'$ properly maps e to e'. Then i continuously extends to a homeomorphism $\partial^{I} e \to \partial^{I} e'$ between the ideal circles at infinity of the ends.

Under these circumstances, we identify e and e'' by means of the homeomorphism i.

More generally, we have:

EXERCISE 1.117. Let $f: S \to S'$ be a holomorphic map between two Riemann surfaces. Let e and e' be tame ends of S and S' respectively such that f induces a proper map $e \to e'$. Then i continuously extends to a covering $\partial^{I} e \to \partial^{I} e'$ between the ideal circles at infinity of the ends.

There is a useful local version of the above result:

EXERCISE 1.118. Let us consider two domains $U, U' \subset \mathbb{C}$ whose boundaries contain open Jordan arcs $\gamma \subset \partial U$ and $\gamma' \subset U'$. Let $\phi : U \to U'$ be a holomorphic map which is proper near γ in the sense that $\phi(z) \to \gamma'$ as $z \to \gamma$. Then ϕ extends continuously to a map $\gamma \to \gamma'$.

6.5. Landing rays and cut-points. Let $K \subset \mathbb{C}$ be a hull and $J = \partial K$. Given a point $a \in J$, the connected components of $K \setminus \{a\}$ are called *unrooted limbs* of K (with the *root* at a). The *limbs* at a are obtained by adding a to the urooted limbs.

If \mathcal{R}_1 and \mathcal{R}_2 are two rays landing at some point $a \in K$, then $\mathcal{R}_1 \cup \mathcal{R}_2 \cup \{a\}$ is a simple curve whose both ends go to ∞ . By the Jordan Theorem, it divides \mathbb{C} into two domains called the *sectors bounded by* \mathcal{R}_1 and \mathcal{R}_2 .

LEMMA 1.119. If two rays \mathcal{R}_1 and \mathcal{R}_2 land at the same point $a \in J$ then each sector bounded by these rays contains an unrooted limb of K at a.

PROOF. Assume one of the sectors, S, does not contain any points of K. This sector is the impage of a circular sector

$$\Delta := \{ r e^{2\pi i \theta} : r > 1, \ \theta_2 < \theta < \theta_2 \} \subset \mathbb{D}$$

under the Riemann map $\Phi : \mathbb{C} \setminus \overline{D} \to \mathbb{C} \setminus K$. Since *a* is the only point of *K* in \overline{S} , we have $\Phi(z) \to a$ as $z \to \mathbb{T}$ within Δ . But this is impossible for a non-constant holomorphic function.

EXERCISE 1.120. Can you justify this assertion?

For this reason, landing points of at least two rays are called *cut-points* of K.

LEMMA 1.121. If a hull K is locally connected then any two limbs at $a \in J$ can be separated by a pair of rays landing a,

Let us consider all the rays \mathcal{R}_{θ} landing at a cut-point $a \in J$. Assume that $\cup \mathcal{R}_{\theta} \cup \{a\}$ is closed subset of \mathbb{C} . (For instance, this is the case when there are only finitely many rays, or when K is locally connected.) Then the components of $\mathbb{C} \setminus (\cup \mathcal{R}_{\theta} \cup \{a\})$ are called the *wakes* of K at a.

6.6. Components of locally connected hulls. We will now use the Carathé-

odory Theorem for further study of the topology of lc hulls.

Let K be a hull, and let (D, b) be a pointed component of int K. Since it is simply connected, it can be uniformized by the unit disk, $\Phi_D : (\mathbb{D}, 0) \to (D, b)$. Internal rays \mathcal{R}_{θ} of (D, b) are defined to be the images of the straight rays $\{re^{2\pi i\theta}: 0 \leq r < 1\}$ under Φ_D .

PROPOSITION 1.122. Let $K \subset \mathbb{C}$ be a lc hull. Then any component D of int K is a Jordan disk.

PROOF. Let us consider the projection $\pi_D : K \to \overline{D}$ (2.2). Since it is continuous and K is lc, \overline{D} is lc as well (Exercise 0.7,b)). By the Second Carathéodory Theorem, the boundary ∂D is lc as well and the uniformization $\Phi_D : \mathbb{D} \to D$ extends continuously to the boundary.

This shows that ∂D is a curve. We just need to show that it is simple. If not, then there are two internal rays \mathcal{R}_1 and \mathcal{R}_2 in D that land at the same point $a \in \partial D$. Then by Lemma 1.119 (applied to the hull $\hat{\mathbb{C}} \setminus D$), the Jordan curve $\gamma := \mathcal{R}_1 \cup \mathcal{R}_2 \cup \{a\}$ surrounds a point $b \in \partial D \subset \partial K$. On the other hand, since K is full, the open Jordan disk bounded by γ is contained in int K; in particular, $b \in \text{int } K$ – contradiction.

EXERCISE 1.123. For any two components D_1 and D_2 of a lc hull K, the closures \overline{D}_1 and \overline{D}_2 are either disjoint or touch at a single point.

6.7. Pinched toplogical model of a lc hull. The Carathéodory Theorem allows one to represent any hull $K \subset \mathbb{C}$ as a quotient of the unit disc \mathbb{D} by a special equivalence relation $\underset{K}{\sim}$. Namely, this theorem provides us with the continuous extension $\phi : \mathbb{C} \setminus \mathbb{D} \to (\mathbb{C} \setminus K) \cup \partial K$. Now, the equivalence classes of $\underset{K}{\sim}$ on the unit circle \mathbb{T} are defined as the fibers $\phi^{-1}(\cdot)$ of $\phi | \mathbb{T}$. Obviously, ∂K is homeomorphic to the quotient $\mathbb{T}/\underset{K}{\sim}$.

We will now extend it to \mathbb{D} . Given a non-singleton class X of $\underset{K}{\sim}$, let \hat{X} stand for the hyperbolic convex hull of X, see (1.2). (For any singleton class $X = \{x\}$, we let $\hat{X} = X$.

LEMMA 1.124. Given a lc hull K, the convex hulls \hat{X} are pairwise disjoint.

PROOF. Let us compactify the complex plane \mathbb{C} with the circle \mathbb{T}_{∞} at infinity. Convergence of points $z_n \in \mathbb{C}$ to $\theta \in \mathbb{T}_{\infty}$ means that $z_n \to \infty$ and $\arg z_n \to \theta$. It is easy to check that this compactification, $\overline{\mathbb{C}}$, is homeomorphic to $\overline{\mathbb{D}}$.

The Riemann mapping $\phi : \mathbb{C} \setminus \overline{\mathbb{D}} \to \mathbb{C} \setminus K$ extends to a homeomorphism $\overline{\mathbb{C}} \setminus \overline{\mathbb{D}} \to \overline{\mathbb{C}} \setminus K$ in an obvious way. Since K is locally connected, it further extends to a continuous map $\overline{\mathbb{C}} \setminus \mathbb{D} \to \overline{\mathbb{C}} \setminus \text{int } K$ by the Second Carathéodory Theorem. (We will keep notation ϕ for all these extensions.)

Given an $\underset{K}{\sim}$ equivalence class $X = \phi^{-1}(x) \subset \mathbb{T}_{\infty}, x \in \partial K$, let

$$\tilde{X} = \{ re^{2\pi i\theta} : r \in [0,\infty], \ \theta \in X \} \subset \overline{\mathbb{C}} \smallsetminus \mathbb{D} :,$$

and let

$$X' = \phi(\tilde{X}) = X \cup \bigcup_{\theta \in X} \mathcal{R}_{\theta} \cup \{x\} \subset \mathbb{C} \setminus \operatorname{int} K.$$

This is a compact set intersecting \mathbb{T}_{∞} by X and intersecting K by $\{x\}$.

Consider now another equivalence class, $Y = \{\phi^{-1}(y)\}, y \in \partial K, y \neq x$. Then $X \cap Y = \emptyset$, and hence $\tilde{X} \cap \tilde{Y} = \emptyset$. Since $\phi : \bar{\mathbb{C}} \setminus \bar{\mathbb{D}} \to \bar{\mathbb{C}} \setminus K$ is a homeomorphism, the sets $X' \setminus K$ and $Y' \setminus K$ are disjoint. But the intersections $X' \cap K = \{x\}$ and $Y' \cap K = \{y\}$ are also disjoint. Thus, $X' \cap Y' = \emptyset$.

By Proposition 1.22, the sets X and Y are unlinked on $\mathbb{T}_{\infty} \approx \mathbb{T}$, so their convex hulls \hat{X} and \hat{Y} are disjoint in $\overline{\mathbb{D}}$.

The sets \hat{X} are declared to be non-singleton equivalence classes of \sim_{K} . All other equivalence classes are singletons. (This equivalence relation can be considered not only on $\overline{\mathbb{D}}$ but on the whole plane \mathbb{C} .)

THEOREM 1.125. A locally connected hull $K \subset \mathbb{C}$ is homeomorphic to the quotient $\mathbb{C}/\underset{\nu}{\sim}$.

PROOF. Let $\hat{\mathbb{T}} := \bigsqcup \hat{X}$, where the union is taken over all equivalence classes on $X \subset \mathbb{T}$.

Step 1: The set $\hat{\mathbb{T}}$ is closed. Let $z_n \to z \in \mathbb{D}$ and $z_n \in \hat{X}_n = \phi^{-1}(\zeta_n)$ with $\zeta_n \in \partial K$. Passing to a subsequence, we can assume that $\zeta_n \to \zeta \in \partial K$. By continuity of $\phi | \mathbb{T}$, we have

$$\limsup_{n \to \infty} X_n \subset X := \phi^{-1}(\zeta).$$

It easily implies that

$$\limsup_{n \to \infty} \hat{X}_n \subset \hat{X}$$

so $z \in \hat{X}$.

Step 2: The map $\phi : \mathbb{T} \to \partial K$ extends to a continuous map $\hat{\phi} : \hat{\mathbb{T}} \to \partial K$ by declaring $\hat{\phi}(\hat{X}) = \phi(X)$.

Let $z_n \to z \in \mathbb{D}$, $z_n \in \hat{X}_n$. Without loss of generality, we can assume that the X_n are pairwise disjoint. Then there exist points $z'_n \in \partial X_n$ converging to z as well, so we can assume that $z_n \in \partial X_n$ in the first place. But $\bigcup \partial X_n$ is a geodesic lamination, so

Step 3: Any gap Q in $\hat{\mathbb{T}}$ (i.e, a connected componet of $\overline{\mathbb{D}} \setminus \hat{T}$) is a convex set, and a map $\hat{\phi}$ continuously extends to a homeomorphism $\overline{Q} \to \overline{D}$, where D is a component of int K.

The ideal boundary of the gap Q, $\partial^I Q := Q \cap \mathbb{T}$, is a closed subset of \mathbb{T} . Let I_j be its componentary intervals, and let Γ_j be the hyperbolic geodesics sharing the endpoints with the I_j . Then the interior of the gap, int $Q := Q \cap \mathbb{D}$, is a convex subset of \mathbb{D} bounded (in \mathbb{D}) by the geodesics Γ_j .

It is clear from this picture that Q is a closed Jordan disc: its boundary (in $\overline{\mathbb{D}}$) can be homeomorpically mapped onto \mathbb{T} by projecting the geodesics Γ_j onto the ideal intervals I_j . The quotient $Q/\underset{K}{\sim}$ obtained by collapsing the Γ_j to singletons is also a closed Jordan disc.

EXERCISE 1.126. Check it.

Any homeomorphism between the boundaries of two Jordan discs extends continuously to the whole discs (e.g., radially). In particular, the embedding $\hat{\phi} : (\partial Q / \underset{K}{\sim}) \to \partial K$ extends to a homeomorphism $(Q / \underset{K}{\sim}) \to \overline{D}$, where D is the (open) Jordan disc bounded by $\hat{\phi}(\partial Q)$. This Jordan disc is contained in int Dsince K is full. Since $\partial D \subset \partial K$, D is a component of int K.

Step 4: The map $\hat{\phi} : \overline{\mathbb{D}} \to K$ is continuous.

Given $z_n \to z \in \overline{\mathbb{D}}$, we want to show that $\hat{\phi}(z_n) \to \hat{\phi}(z)$. By the above discussion (Steps 2-3), we only need this check it in case $z_n \in Q_n$ where the Q_n are distinct gaps. Since area $Q_n \to 0$, there exist points $z'_n \in \partial Q_n \subset \widehat{\mathbb{T}}$ such that $\operatorname{dist}(z_n, z'_n) \to 0$, so $z'_n \to z$ as well. By Step 2, $\hat{\phi}(z'_n) \to \hat{\phi}(z)$. But $\operatorname{dist}(\hat{\phi}(z_n), \hat{\phi}(z'_n) \leq \operatorname{diam} \hat{\phi}(Q_n) \to 0$ by Proposiation 1.124. The conclusion follows.

Step 5: The map $\hat{\phi} : \overline{\mathbb{D}} \to K$ is onto. Here we will make use of the exterior of $\overline{\mathbb{D}}$. Let us consider some circle \mathbb{T}_R with R > 1 and the corresponding equipotential $\mathcal{E}_R = \phi(\mathbb{T}_R)$. It goes once around K, so by the Topological Argument Principle (Proposition 1.50) all values in K must be assumed by $\hat{\phi}$

7. Appendix: Potential theory

Harmonic and subharmonic functions is a very important subject on its own right that penetrates deeply into analysis, geometry, and probability theory. From our perspective, their outstanding role comes from the fact that they lay down a foundation for a proof of the Uniformization Theorem. For readers' convenience, here we will briefly review needed basics of the theory.

7.1. Harmonic functions and differentials. Recall that a function $u: U \to \mathbb{R}$ on a domain $U \subset \mathbb{C}$ is called *harmonic* if $u \in C^2(U)$ and $\Delta u = 0$ where $\Delta = \partial_x^2 + \partial_y^2$ is the usual Euclidean Laplacian. The real and imaginary parts of any holomorphic function f = u + iv on U are harmonic, which is readily seen from the *Cauchy-Riemann equations*

$$\partial_x u = \partial_y v, \quad \partial_y u = -\partial_x v.$$

They are called *conjugate harmonic functions*.

Vice versa, any harmonic function u can locally be represented as the real part of a holomorphic function. Indeed, $\Delta u = 0$ gives the integrability condition for the Cauchy–Rieamann equations that allow one to recover locally the conjugate function v.

This can be nicely expressed in terms of the Hodge * operator. Let $V \approx \mathbb{R}^2$ be the oriented 2D Euclidean space. By self-duality, we identify vector fields $\tau = a\partial_x + b\partial_y$ with 1-forms $\omega = adx + bdy$. The Hodge *-operator is defined as $\pi/2$ -rotation of ω or τ , i.e. $*\omega = -bdx + ady$.

Then the Cauchy–Riemann equations can be written as

(7.1) $dv = d^c u$, where $d_c := *d$, while $dd^c u = \Delta u \, dz \wedge dy$.

So, u is harmonic if and only if the form $d^{c}u$ is closed, and then (7.1) can be locally integrated:

(7.2)
$$v(z) = \int_{z_0}^{z} d^c u = \int_{\gamma} \frac{\partial u}{\partial n} ds,$$

where γ is a smooth (oriented) path connecting z_0 to z (within a small disk), ds is the length element on γ and n is the unit normal vector to γ rotated clockwise from the corresponding tangent vector to γ .

Globally, the integral (7.2) depends on the homotopy class of the path γ (rel the endpoints), so it defines a multi-valued harmonic function v and the corresponding multivalued holomorphic function f = u + iv. The monodromy for this function along a cycle γ depends only on the homology class of γ and is given by the *periods* of $d^c u$:

$$f_{\gamma}(z) - f(z) = i \int_{\gamma} d^{c}u = i \int_{\gamma} \frac{\partial u}{\partial n} ds,$$

where f_{γ} is the result of analytic continuation of f along along γ . In particular, if a is an isolated singularity for u, then the monodromy if f as we go around a little circle $\gamma = S_r := \{|z - a| = r\}$ is equal to

$$f_{\gamma}(z) - f(z) = i \int_{S_r} \frac{\partial u}{\partial r}(\zeta) \, d\theta.$$

Relation between harmonic and holomorphic functions makes the notion of harmonicity manifestly invariant under holomorphic changes of variable: if u is harmonic then so is $u \circ \phi$ for any holomorphic map ϕ . Thus, harmonicity is well-defined on an arbitrary Riemann surface S. This can also be seen from the original definition by expressing the Laplacian in terms of the differential operators ∂ and $\overline{\partial}$ (see §1.8). Indeed, we have:

(7.3)
$$\partial = \frac{1}{2}(d+id^c), \quad \bar{\partial} = \frac{1}{2}(d-id^c).$$
so,

$$\Delta u \, dx \wedge dy = dd^c \, u = 2i \, \partial \bar{\partial} \, u.$$

REMARK 1.8. Expressions (7.3) show that d and d^c are (twice) the real and imaginary parts of the operators ∂ and $\overline{\partial}$.

A C^1 differential 1-form $\omega = adx + bdy$ is called *harmonic* if it is locally the differential of a harmonic function. It is called *co-closed* if $d(*\omega) = 0$. It is straightforward to check that a form ω is harmonic if and only if it is closed and co-closed.

Another characterization is that harmonic 1-forms are real part of Abelian differentials. Namely, the differential $\alpha = \omega + i\eta$ is holomorphic if and only if ω is is harmonic and $\eta = *\omega$. (Note the unlike the case of functions, this relation is global.)

7.2. Basic properties. Given a domain U on a Riemann surface S, let $\mathcal{H}(U)$ stand for the space of harmonic functions in U, and let $\mathcal{H}(\bar{U})$ stand for the subspace of $\mathcal{H}(U)$ consisting of functions that admit continuous extension to \bar{U} .

MEAN VALUE PROPERTY. A C^2 function u on a domain $U \subset \mathbb{C}$ is harmonic is and only if for any disk $\mathbb{D}(a, r) \subset U$, we have

$$u(a) = M_u(a, r) := \frac{1}{2\pi} \int_0^{2\pi} h(a + re^{i\theta}) \, d\theta;$$

PROOF. The mean value property for harmonic functions immediatedly follows from the corresponding property for holomorphic ones. The inverse follows from the second order Taylor expansion at z averaged over a little circle:

(7.4)
$$M_u(z,r) - u(z) = \frac{1}{4}\Delta u(z)r^2 + o(r^2).$$

The Mean Value Property implies in a standard way (as for holomorphic functions):

MAXIMUM/MINIMUM PRINCIPLE. If a harmonic function u on a Riemann surface U has a local maximum or minimum in U then it is constant.

COROLLARY 1.127. Let $U \in S$ be a compactly embedded domain in a Riemann surface S, and let $u \in \mathcal{H}(\overline{U})$. Then u attains its maximum and minimum on ∂U .

COROLLARY 1.128. Under the above circumstances, u is uniquely determined by its boundary values, $u | \partial U$.

7.3. Poisson Formula. The Poisson Formula allows us to recover a harmonic function $h \in H(\overline{\mathbb{D}})$ from its boundary values:

PROPOSITION 1.129. For any harmonic function $h \in H(\overline{\mathbb{D}})$ in the unit disk, we have: formula: the following Poisson representation:

$$h(z) = \frac{1}{2\pi} \int_0^{2\pi} h(\zeta) P(z,\zeta) \, d\theta, \quad z \in \mathbb{D}, \ \zeta = e^{i\theta} \in \mathbb{T},$$

with the the Poisson kernel

$$P(z,\zeta) = \frac{1-|z|^2}{|z-\zeta|^2}.$$
PROOF. For z = 0, this formula amounts to the Mean Value Property:

$$h(0) = \frac{1}{2\pi} \int_0^{2\pi} h(e^{i\theta}) d\theta.$$

It implies the formula at any point $z \in \mathbb{D}$ by making a Möbius change of variable

$$\phi_z: \overline{\mathbb{D}} \to \overline{\mathbb{D}}, \quad \zeta \mapsto \frac{\zeta - z}{1 - \overline{z}\zeta}$$

that moves z to 0. Since $h \circ \phi_z^{-1} \in H(\overline{\mathbb{D}})$, we obtain:

$$h(z) = (h \circ \phi_z^{-1})(0) = \frac{1}{2\pi} \int_0^{2\pi} h \circ \phi_z^{-1} d\theta = \frac{1}{2\pi} \int_0^{2\pi} h \, d\theta_z,$$

where

$$d\theta_z = (\phi_z)^* (d\theta) = |(\phi_z)'(\theta)| \, d\theta,$$

and the latter derivative is equal to the Poisson kernel $P(z,\zeta)$ (check it!).

Uniqueness of the extension follows from the Maximum Principle.

The Dirichlet problem (in some domain $D \subset \hat{\mathbb{C}}$) is the problem of recovery of a harmonic function $h \in \mathcal{H}(\overline{D})$ from its boundary values on ∂D . The Poisson formula provides us with an explicit solution of this problem in the unit disk:

PROPOSITION 1.130. Any continuous function $g \in C(\mathbb{T})$ on the unit circle admits a unique harmonic extension $h \in \mathcal{H}(\overline{\mathbb{D}})$ to the unit disk (so that $g = h | \mathbb{T})$. This extension is given by the Poisson formula:

$$h(z) = \frac{1}{2\pi} \int_0^{2\pi} g(\zeta) P(z,\zeta) \, d\theta, \quad z \in \mathbb{D}, \ \zeta = e^{i\theta} \in \mathbb{T}.$$

PROOF. The Poisson kernel $P(z,\zeta)$ as a function of $\zeta \in \mathbb{T}$ and $z \in \mathbb{D}$ has the following properties:

(i) $P(z,\zeta) > 0$, and for any $z \in \mathbb{D}$, we have $\frac{1}{2\pi} \int_{\mathbb{T}} P(z,\zeta) \, d\theta = 1$;

(ii) For any $\zeta \in \mathbb{T}$, the kernel $P(z, \zeta)$ is harmonic in $z \in \mathbb{D}$;

(iii) For any $\zeta_0 \in \mathbb{T}$ and any $\epsilon > 0$, we have:

$$\mathbb{P}_{z}(\zeta) \to 0$$
 as $z \to \zeta_{0}$ uniformly in $\zeta \in \mathbb{T} \setminus \mathbb{D}(\zeta_{0}, \epsilon)$.

Property (i) follows from the Poisson representation of the function $h(z) \equiv 1$ in $\overline{\mathbb{D}}$.

To check (ii), notice that $P(\cdot, \zeta)$ is the pullback of the function $\operatorname{Im} u$ on the upper half-plane to the unit disk under the Móbius transformation

$$\phi_{\zeta} : \mathbb{D} \to \mathbb{H}, \quad \phi_{\zeta} : z \mapsto i \frac{\zeta + z}{\zeta - z}$$

EXERCISE 1.131. Check this using that ϕ_z is a hyperbolic isometry.

The last property is obvious (it corresponds to the fact the function $\operatorname{Im} u$ vanishes on \mathbb{R}).

Properties (i) and (iii) imply that $P(z, e^{i\theta}) d\theta$, viewed as measures on \mathbb{T} weakly converge to δ_{ζ_0} . This implies that g gives the boundary values of h. Property (ii) implies harmonicity of h in \mathbb{D} .

7.4. Harnak Inequality and normality. This inequality allows one to control a positive harmonic function by its value at one point. Let us begin with the case of disk:

LEMMA 1.132. For any $r \in (0,1)$, there exists a constant C(r) > 1 such that for any positive harmonic function $u \in \mathcal{H}(\overline{\mathbb{D}})$, we have:

$$C(r)^{-1} u(0) \le h(z) \le C(r) u(0), \quad |z| \le r$$

PROOF. It immediately follows from the Poisson representation since

$$C(r)^{-1} \le P(z,\zeta) \le C(r)$$
 $(|\zeta| = 1, |z| \le r)$ with $C(r) = \frac{1+r}{1-r}$

and the Mean Value Property.

Let us now consider the general case. By a *coordinate disk* $D(a, \epsilon)$ we mean a domain lying within some local chart and equal to the disk $\mathbb{D}(z(a), \epsilon)$ in this cordinate.

THEOREM 1.133. Let S be a (connected) Riemann surface, and let $z_0 \in U$, $K \Subset U$. Then there exists a constant $C_K > 1$ such that for any positive harmonic function $u \in \mathcal{H}(U)$, we have:

$$C_K^{-1} u(z_0) \le u(z) \le C_K u(z_0), \text{ for any } z \in K.$$

PROOF. We can find finitely many coordinate disks $D(z_i, \epsilon_i)$ whose union $\cup D(z_i, \epsilon_i/2)$ is connected and covers $K \cup \{z_0\}$. Applying the Lemma 1.132 consecutively to these disks, we obtain the desired inequlities.

Similarly to holomorphic functions, bounded families of harmonic functions are *normal* (i.e., precompact in the topology of uniform convergence on compact subsets):

PROPOSITION 1.134. A bounded family of harmonic functions on U is normal.

PROOF. The Poisson formula gives a bound on the partial derivatives of $u \in \mathcal{H}(u)$ on a compect subset $K \Subset U$ in terms of the bound on u (and the set K). By the Ascoli-Arcela, our family is precompact in the space of consinuous functions on U (in topology of uniform convergence on compact subsets). But the Mean Value Property survives under taking locally uniform limits. Hence harmonicity survives as well.

COROLLARY 1.135. Let $u_n \in \mathcal{H}(U)$ be an increasing sequence of harmonic functions, and let $u_n(z_0) \leq C$ at some point $z_0 \in U$. Then the u_n converge, uniformly on compact subsets of U, to a harmonic function $u \in \mathcal{H}(\mathcal{U})$.

PROOF. Subtracting u_0 from the u_n , we see that our functions can be assumed positive. By the Harnak Inequality, the u_n are uniformly bounded on compact subsets. So, their pointwise limit u(z) is finite. Moreover, by Proposition 1.134, they form a normal sequence, and hence u is harmonic.

1. CONFORMAL GEOMETRY

7.5. Subharmonic functions. Harmonic functions are analytic and hence rigid: they cannot be locally modified. Subharmonic functions are much more flexible, but at the same time, they still possess good compactness properties (an *a priori* upper bound is sufficient). This combination makes them very useful.

The basic example of a subharmonic function is $u = \log |f(z)|$ where f is a holomorphic function. In fact, this function is harmonic everwhere except for zeros of f where it assumes value $-\infty$ ("poles" of u). This suggests that in general subharmonic functions should also be allowed to have poles. Of course, $[-\infty, \infty)$ is naturally endwed with topology of a half-open interval.

DEFINITION 1.136. A function $u: D \to [-\infty, \infty)$ on a domain $D \subset \mathbb{C}$ is called *subharmonic* if it is not identically equal to $-\infty^{12}$ and satisfies the following two conditions:

• Mean Value Property (subharmonic): For any disk $\mathbb{D}(z,r) \subseteq D$,

(7.5)
$$u(z) \le M_u(z,r)$$

• *u* is upper-semicontinuous.

REMARK 1.9. Notice that the two conditions in the above definition make the *value* of a subharmonic function well determined at a point by its values nearby. In fact, below we will be dealing only with *continuous* subharmonic functions, and mostly, assuming only *finite* values. However, the following basic subharmonic function does have a pole:

EXAMPLE 1.1. Let $u(z) = \log |z|$. This function is harmonic in \mathbb{C}^* , so the MVP is satisfied on an any disk $\mathbb{D}(a,r) \in \mathbb{C}^*$. It is also obviously satisfied on \mathbb{D}_r as $-\infty < M_u(0,r)$.

Let us check it for the disk $\mathbb{D}(a, r) \ni 0$. Making an affine change of variable, we can consider instead the Mean Value Property on \mathbb{D} for a function $v(z) = \log |z - c|$, $c \in \mathbb{D}^*$. Then we have:

$$\frac{1}{2\pi} \int_{\mathbb{T}} u(z) \, d\theta = \frac{1}{2\pi} \int_{\mathbb{T}} \left(u(z) + \log \left| \frac{1 - \bar{c}z}{z - c} \right| \right) \, d\theta$$
$$= \frac{1}{2\pi} \int_{\mathbb{T}} \log |1 - \bar{z}c| \, d\theta = 0 > \log |c| = v(0).$$

For the disk $\mathbb{D}(a, |a|)$ whose boundary passes through 0, MVP follows by continuity.

REMARK 1.10. The above estimate is a particular case of the Jenssen formula:

$$\frac{1}{2\pi} \int_{\mathbb{T}} \log |f(\zeta)| d\,\theta = \log |f(0)| + \sum \log \frac{1}{|a_i|},$$

where f is a holomorphic function in \mathbb{D} , continuous up to the boundary, that does not vanish on \mathbb{T} and at 0.

We let $\mathcal{SH}(U)$ stand for the space of continuous subharmonic functions in U.

Obviously, the set of subharmonic functions is invariant under addition and multiplication by positive numbers, so it is a *cone*. Also, *Maximum of finitely* many subharmonic functions is subharmonic. For instance, the function $\log^+ |z| = \max\{\log |z|, 0\}$ is subharmonic.

As for harmonic functions, the Subharmonic Mean Value Property implies:

¹²This convention is not competely standartized.

MAXIMUM PRINCIPLE. If a subharmonic function u on a Riemann surface U has a local maximum in U then it is constant.

However, the *Minimum Principle* is not any more valid for subharmonic functions.

More generally, we can majorant a subharmonic function by a harmonic one:

LEMMA 1.137. Let D be a bounded domain in \mathbb{C} , and let u and h be respectively harmonic and a continuous subharmonic functions on D, both admitting a continuous extensions to \overline{D} . If $u \leq h$ on ∂D then $u \leq h$ in D.

Vice versa, if a function u is continuous in a domain $U \subset \mathbb{C}$ and the above property is satisfied for any domain $D \subseteq U$ and any harmonic $h \in \mathcal{H}(U)$, then u is subharmonic.

PROOF. To check the former assertion, apply the Maximum Principle to u - h. To check the latter, let us consider a coordinate disk D and let h solves the Dirichlet Problem in D with the boundary values $h | \partial D = u | \partial D$. Then $u | D \le h | D$. Evaluating it at the center of D, we obtain the Mean Value Property for subharmonic functions.

This lemma shows that the notion of subharmonicity is bi-holomorphically invariant (at least for continuous functions¹³, and hence is well defined on an arbitrary Riemann surface.

Also, let us consider a function

$$\tilde{u}_D(z) = u(z)$$
 for $z \in U \setminus D$, and $u(z) = h(z)$ for $z \in D$,

where h is a harmonic function in D defined in the second part of Lemma 1.137. We call \tilde{u}_d the harmonic majorant of u rel ∂D . The first part of Lemma 1.137 implies that the harmonic majorant of u is subharmonic.

A function u is called *superhamonic* if -u is subharmonic. Properties of such functions follow immediately form the corresponding properties of subharmonic ones.

7.6. Perron method. A (non-empty) family \mathcal{P} of continuous subharmonic functions on a Riemann surface U is called *Perron* if it satisfies the following properties:

- (i) If $u, v \in \mathcal{P}$ then $\max(u, v) \in \mathcal{P}$;
- (ii) For any $u \in \mathcal{P}$ and any coordinate disk $D \Subset U$, the harmonic majorant \tilde{u}_D also belongs to \mathcal{P} .

PROPOSITION 1.138. If \mathcal{P} is a Perron family on U then the function

$$h(z) := \sup_{\mathcal{D}} u(z)$$

is either harmonic or identically equal to ∞ .

PROOF. Since harmonicity is a local property, it is enough to check it within coordinate disks $D \in U$. Fix such a disk D. Since $u \leq \tilde{u} \in \mathcal{P}$, we have $h(z) := \sup_{\mathcal{P}} \tilde{u}(z)$. So, without loss of generality we can assume that all the functions $u \in \mathcal{P}$ are harmonic in D.

 $^{^{13}}$ It is still true in general, but we will not need it

1. CONFORMAL GEOMETRY

Take a countable dense subset $X \subset D$. By means of the diagonal procedure, we can select a sequence of functions $u_n \in \mathcal{P}$ such that $h(z) = \sup u_n(z)$ for any $z \in X$. Let v_n be the harmonic majorant (rel ∂D) of the function $\max(u_1, \ldots, u_n)$, $n \in \mathbb{Z}_+$. This is a monotonically inreasing sequence of functions of the family \mathcal{P} , harmonic on D, and such that $v_n(z) \to h(z)$ on X. By Corollary 1.135, $v_n \to \phi$ locally uniformly on D, where ϕ is either harmonic, or else $\phi \equiv \infty$. In either case, we have:

$$\phi(z) = h(z) \ge u(z)$$
 for any $z \in X$, $u \in \mathcal{P}$.

Since both ϕ and u are continuous, we conclude that $\phi \geq u$ everywhere on D; hence $\phi \geq h$ everywhere on D. On the other hand, since $\phi = h$ on the dense set X and h is upper semicontinuous (as sup of a family of continuous functions), we conclude that $\phi \leq h$ everywhere on D. Thus $\phi \equiv h$ on D.

7.7. Dirichlet barriers. We will now apply the Perron method to solve the Dirichlet problem in an arbitrary domain (for which it is solvable at all).

Let $U \subseteq S$ be a domain in a Riemann surface S, and let g be a continuous function on ∂U . Let us consider the following Perron family of subharmonic functions:

$$\mathcal{P} \equiv \mathcal{P}_U(g) = \{ u \in \mathcal{SH}(U) : \limsup_{\zeta \to z} u(\zeta) \le g(z) \quad \forall z \in \partial U \}.$$

By Proposition 1.138, the function $h_g := \sup_{\mathcal{P}} u$ is harmonic in U. To study its boundary values, we will introduce the following notions:

A barrier b_a at a boundary point $a \in \partial U$ is a subharmonic function $b_a(z)$ defined on a relative neighborhood D of a in U, continuous up to ∂D ,¹⁴ and such that $b_a(a) = 0$ while $b_a(z) < 0$ for any $z \neq a$. A point $a \in \partial U$ is called *Dirichlet regular* if it has a barrier.

EXAMPLE 1.2. If ∂U near a is an arc of a smooth curve then a is regular. Indeed, then there is a wedge

$$W = \{ |\arg(z - a) - \alpha| < \epsilon, \ 0 < |z| < 2\pi \epsilon \}$$

which is disjoint from \overline{U} . The complementary wedge can be mapped conformally onto the lower half-plane (by a branch of the power function $\phi(z) = e^{i\theta}(z-a)^{\gamma}$ with appropriate $\gamma \in (0,1)$ and θ . The function $b = \text{Im } \phi(z)$ restricts to a barrier at a on U.

EXERCISE 1.139. Show that the same is true is ∂U near a is a Jordan arc.

THEOREM 1.140. Let $U \in S$ be a domain in a Riemann surface S, and let g be a continuous function on ∂U . Let us consider the harmonic function $h = h_g$ constructed above by means of the Perron method. Then for any Dirichlet regular point $a \in \partial U$, we have: $h(z) \to g(a)$ as $z \to a$.

PROOF. Without loss of generality, we can assume that g(a) = 0. Let us first show that

(7.6)
$$\liminf_{z \to z} h(z) \ge 0$$

Take a small r > 0 such that the barrier $b(z) = b_a(z)$ is well defined in $D_r := \mathbb{D}(a, 2r) \cap U$. Let ξ be the supremum of b on $S_r := \{|z - a| = r\} \cap U$. By definition of the barrier, $\xi < 0$.

 $^{^{14}{\}rm This}$ condition can be relaxed, but it is sufficient for our purposes. In fact, harmonic barriers would also be good enough for us.

The function $\hat{b}(z) := \max(b(z), \xi)$ is a continuous subharmonic function in $\mathbb{D}(a, r) \cap U$ equal to ξ on S_r . Hence it extends to a continuous subharmonic function in in U by letting $\hat{b} \equiv \xi$ in $U \smallsetminus D_r$.

Let now

$$\eta = \inf\{g(z): z \in \partial U \smallsetminus D_r\}, \quad -\epsilon = \inf\{g(z): z \in \partial U \cap D_r\} < 0,$$

and consider

$$\beta(z) = \frac{\eta}{\xi} \, \hat{b}(z) - \epsilon.$$

This is a subharmonic function in U with

$$\lim_{z\to p}\beta(z)=\eta\quad\text{for }p\in\partial U\smallsetminus D_r;\quad \liminf_{z\to p}\beta(z)\leq -\epsilon\quad\text{for }p\in\partial U\cap D_r,$$

so β belongs to the Perron family \mathcal{P} .

It follows that $h \ge \beta$ and hence

$$\liminf_{z \to a} h(z) \ge -\epsilon.$$

Since $\epsilon \to 0$ as $r \to 0$, we obtain (7.6).

To obtain the opposite estimate, let us consider the negative barrier -b(z). It allows us to construct, for any $\epsilon > 0$, a superharmonic function α in U such that

$$\liminf_{z \to p} \alpha(z) \ge g(p) \ \forall p \in \bigcup \quad \text{and} \quad \limsup_{z \to p} \alpha(z) \le \epsilon.$$

By the Maximum Principle, $u \leq \alpha$ for any $u \in \mathcal{P}$, and hence $h \leq \alpha$ as well. It flows that

$$\limsup_{z \to a} h(z) \le \epsilon$$

and we are done.

We say that a domain $U \subseteq S$ has a *Dirichlet regular boundary* if ∂U is nonempty and all points of ∂U are regular.

COROLLARY 1.141. Let $U \subseteq S$ be a domain with Dirichlet regular boundary. Then the Dirichlet problem is solvable in U for any continuous boundary values.

7.8. Harmonic measure. Let $U \in S$ be a domain with Dirichlet regular boundary. Then any continuous function $g \in C(\partial U)$ admits a harmonic extension $\hat{g} \in \mathcal{H}(\bar{U})$ to U. Endow $\mathcal{H}(\bar{U})$ with uniform topology on the whole \bar{U} . It is a Banach space ismorphic to $C(\partial U)$ by means of the natural restriction and the above extension operators.

For a given $z \in U$, evaluation $\hat{g}(z)$ is a bounded linear functional on $C(\partial U)$ and hence it is represented by a Borel measure μ_z on ∂U :

$$\hat{g}(z) = \int_{\partial U} g \, d\mu_z.$$

This measure is called the *harmonic measure* for U at z. For instance, in the unit disk, we have $d\mu_z = \mathcal{P}(z,\zeta) d\theta$ where \mathcal{P} is the Poisson kernel.

If ∂U is disconnected and $K \subset \partial U$ is a clopen subset then $\mu_z(K)$ is a harmonic function on U with boundary values 1 on K and 0 on $\partial U \smallsetminus K$. This function itself is sometimes referred to as the "harmonic measure of K" (which may sound confusing).

7.9. Green function. We will restrict our discussion to domains $U \Subset S$ with Dirichlet regular boundary. The *Green function* $G = G_p$ on U with pole at $p \in U$ is a harmonic function such that

(Gr1) $G(z) \to 0$ as $z \to \partial U$;

(Gr2) In a local coordinate z near p such that z(p) = 0, we have:

$$G(z) = \log \frac{1}{|z|} + O(1) \quad \text{near } p.$$

For instance, the Green function in \mathbb{D} with pole at 0 is $-\log |z|$.

REMARK 1.11. Obviously, existence of such a function G implies the Dirichlet regularity of U as -G provides a barrier at any boundary point. In the non-regular case, condition (Gr1) can be relaxed so that the Green function still exists as long as ∂U has positive capacity.

REMARK 1.12. The Green function has a clear electrostatical meaning as the potential of the unit charge placed at p in a domain bounded by a conducting material with the ground potential 0.

The level sets of the Green function G_p are called *equipotentials*, its gradient lines are called *rays* (eminated from p). They form two orthogonal foliations on $U \setminus \{p\}$ with singularities at the critical points of G_p

THEOREM 1.142. Let $U \subseteq S$ be a domain in a Riemann surface S with Dirichlet regular boundary. Then for any $p \in U$, there exists a unique Green function G_p with pole at p.

PROOF. Let us consider the following family $\mathcal{P} = \mathcal{P}_U[p]$ of functions on $U \setminus \{p\}$: (i) $\limsup u(z) \leq 0$;

(ii) In a local coordinate z near p such that z(p) = 0, we have:

1

$$\frac{1}{1}$$

$$u(z) = \log \frac{1}{|z|} + O(1).$$

Obviously, it is a Perron family, so the function $G = \sup_{\mathcal{P}} u$ is harmonic in $U \setminus \{p\}$ unless it is identically equal to ∞ . We will show that this function is actually finite, and it is the desired Green function.

First, \mathcal{P} is non-empty. Indeed, for a small r > 0, the function $u_0 := \log^+(r/|z|)$ (equal to $\log(|z|/r)$ on the coordinate disk D(p,r) and extended by 0 the whole U) is in \mathcal{P} . Thus,

(7.7)
$$G(z) \ge \log^+ \frac{|z|}{r} \ge 0.$$

Let us show that G is finite. Let S_r be the coordiate circle centered at p of radius r, and let $||u||_r$ be the sup-norm of a function u on S_r . Let us fix two small radii 0 < r < R and compare $||u||_r$ and $||u||_R$ for $u \in \mathcal{P}$.

First, let us look at u from "inside". Take a small $\epsilon>0$ and let

$$u_{\epsilon}(z) = u(z) + (1+\epsilon) \log |z|.$$

This function is subharmonic in $D(p, R) \setminus \{p\}$ and equal to $-\infty$ at p (by property (ii) of the family \mathcal{P}). Hence it is subharmonic on the whole disk D(p, R). By the Maximum Principle, $||u_{\epsilon}||_{r} \leq ||u_{\epsilon}||_{R}$, so

$$||u||_r \le ||u||_R + (1+\epsilon)\log\frac{R}{r}$$

Letting $\epsilon \to 0$, we obtain

(7.8)
$$||u||_r \le ||u||_R + \log \frac{R}{r}$$

On the other hand, we can look at u from "outside". The Maximum Principle in $S\smallsetminus D(z,r)$ implies that for any $u\in\mathcal{P}$

(7.9)
$$||u||_R < ||u||_r,$$

but we want to have a definite drop:

$$(7.10) ||u||_R \le \lambda ||u||_r$$

with some $\lambda < 1$ independent of u. Together with (7.8), this would imply

$$\|u\|_r \leq \frac{1}{1-\lambda}\log\frac{R}{r}$$

that would prove finitness of G on S_r and hence everywhere on U.

To prove (7.10), let us consider the solution v of the Dirichlet problem in $U \\ (z, r)$ with boundary values 1 on S_r and 0 on ∂U (the "harmonic measure" of S_r). Since the boundary of ∂U is regular by assumption and S_r is regular as a smooth curve, such a v exists (Corollary 1.141). By the Maximum Principe, $\lambda := \|v\|_R < 1$.

Furthermore, the function u(z) is asymptotically majorated by $||u||_r v(z)$ near the boundary of $S \setminus D(p, r)$. By the Maximum Principle,

(7.11)
$$u(z) \le \|u\|_r v(z), \quad z \in \mathbb{V}D(p,r).$$

Taking its sup on S_R , we obtain (7.10).

The required properties of the Green function also follow from the above estimates. Indeed, (7.7) and (7.11) imply (Gr1), while (7.7) and (7.8) imply (Gr2).

Notice in conclusion that the Green function extends subharmonically to the whole Riemann surface S by letting $G \equiv 0$ on $S \smallsetminus U$.

EXERCISE 1.143. The Green function has a critical point in U if and only if U is not simply connected.

CHAPTER 2

Quasiconformal geometry

11. Analytic definition and regularity properties

11.1. Linear discussion.

11.1.1. Teichmüller metric on the space of conformal structures. Let $V \approx \mathbb{R}^2$ be a real two-dimensional vector space. A conformal structure μ on V is a Euclidean structure $(v, w)_{\mu}$ up to scaling. Equivalently, it is an ellipse $E_{\mu} = \{ \|w\|_{\mu} = 1 \}$ centered at the origin, up to scaling (here $\|w\|_{\mu}$ is the associted Euclidean norm). Let Conf(V) stand for the space of conformal structures on V.

Let us consider two Euclidean structures, $(v, w)_{\mu}$ and $(v, w)_{\nu}$ representing conformal structures μ and ν . We define the *Teichmüller distance* between μ and ν as the distortion of one Euclidean norm with respect to the other:

$$\operatorname{dist}_{\mathrm{T}}(\mu,\nu) = \log\left(\max_{w \in V^*} \frac{\|w\|_{\mu}}{\|w\|_{\nu}} : \min_{w \in V^*} \frac{\|w\|_{\mu}}{\|w\|_{\nu}}\right) \text{ where } V^* = V \setminus \{0\}.$$

Note that it is independent of the the choice of Euclidean structures representing μ and ν .

EXERCISE 2.1. Check that dist_{T} is a metric on $\operatorname{Conf}(V)$.

If we simultaneously diagonalize the Euclidean structures so that

$$\|w\|_{\nu}^{2} = x^{2} + y^{2}, \ \|w\|_{\mu}^{2} = x^{2}/a^{2} + y^{2}/b^{2}, \ \text{where} \ w = (x,y), \ a \geq b > 0,$$

then

$$\operatorname{dist}_{\mathrm{T}}(\mu, \nu) = \log(a/b) \equiv \log K.$$

The ratio K = a/b of the axes of the ellipse E_{μ} is called the *dilatation* of μ relative ν . Informally we can say that the Teichmiller distance measures the *relative shape* of the ellipses representing our conformal structures.

An invertible linear operator $A: V' \to V$ induces a natural *pullback* operator $A^*: \operatorname{Conf}(V) \to \operatorname{Conf}(V')$: If $(v, w)_{\mu}$ is the Euclidean structure representing $\mu \in \operatorname{Conf}(V)$ then the pullback $A^*\mu$ is represented by $(Av, Aw)_{\mu}$. It follows immediately from the definitions, that the *Teichmüller metric is preserved by the pullback transformations*.

In particular, the group GL(V) of invertible linear automorphisms of V isometrically acts on Conf(V) on the right: $\mu A = A^*\mu$. Let us restrict this action to the group $GL_+(V)$ of orientation preserving automorphisms. Since this action is transitive, it turns Conf(V) into a $GL_+(V)$ -homogeneous space.

To understand this space, let us fix some reference conformal structure σ and select coordinates (x, y) on V that bring it to the standard form $x^2 + y^2$. Then $GL_+(V)$ gets identified with $GL_+(2, \mathbb{R})$, and the isotropy group of σ gets identified

with the group Sim(2) of similarities. Hence

(11.1) $\operatorname{Conf}(V) \approx \operatorname{Sim}(2) \backslash \operatorname{GL}_+(2, \mathbb{R}) = \operatorname{SO}(2) \backslash \operatorname{SL}(2, \mathbb{R}).$

Recall that in §1.5 we endowed the symmetric space $SO(2) \setminus SL(2, \mathbb{R})$ with an invariant metric.

EXERCISE 2.2. This invariant metric coincides with the Teichmüller metric on $\operatorname{Conf}(V)$.

But according to Exercise 1.8, the hyperbolic plane $\mathbb H$ is naturally isometric to the symmetric space

$$PSL(2,\mathbb{R})/PSO(2) \approx SL(2,\mathbb{R})/SO(2).$$

Since the left and right symmetric spaces are equivariantly isometric by the inversion $A \mapsto A^{-1}$, we conclude:

PROPOSITION 2.3. The space $\operatorname{Conf}(V)$ endowed with the Teichmüller metric is equivariantly isometric to the hyperbolic plane \mathbb{H} .

In concusion, let us give one more interpretation of the isomorphism (11.1). It is obtained by associating to an operator $A \in GL_+(2, \mathbb{R})$ the conformal structure μ represented by the Euclidean structure $(v, w)_{\mu} = (Av, Aw)$ (where (v, w) is the standard Euclidean structure on \mathbb{R}^2). The corresponding ellipse E_{μ} is the pullback of the standard round circle: $E_{\mu} = A^{-1}(\mathbb{T})$.

Making use of the polar decompositions of linear operators, we can uniquely represent A as a product of a positive self-adjoint operator P and a rotation O, $A = O \cdot P$. Let $\sigma_{\max} \ge \sigma_{\min} > 0$ stands for the eigenvalues of P. The operator A is a similarity if and only if P is scalar, i.e., $\sigma_{\max} = \sigma_{\min}$. Otherwise we have two orthogonal (uniquely defined) eigenlines l_{\max} and l_{\min} corresponding to σ_{\max} and σ_{\min} respectively. These lines give the directions of maximal and minimal expansion for the operator A. Moreover, the ellipse $E_{\mu} = A^{-1}(\mathbb{T}) = P^{-1}(\mathbb{T})$ has the big axis of length $1/\sigma_{\min}$ on l_{\min} and the small axis of length $1/\sigma_{\max}$ on l_{\max} . The dilatation of this ellipse (equal to $\sigma_{\max}/\sigma_{\min}$) will be also called the *dilatation* of A, Dil A.

EXERCISE 2.4. Show that $\text{Dil}(AB) \leq \text{Dil} A \text{ Dil} B$ with equality attained iff the eigenlines of A and B^{-1} coincide.

11.1.2. Beltrami coefficients. Let now $V = \mathbb{C}_{\mathbb{R}}$ be the decomplixified \mathbb{C} . It is endowed with the standard conformal structure σ (represented by the Euclidean metric $|z|^2$) and with the standard orientation (such that $\{1, i\}$ is positively oriented). Let $A : \mathbb{C}_{\mathbb{R}} \to \mathbb{C}_{\mathbb{R}}$ be an invertible \mathbb{R} -linear operator (which can be also viewed as a \mathbb{C} -valued \mathbb{R} -linear form on V).

Let us describe the conformal structure $A^*\sigma$ in coordinates z, \overline{z} of $\mathbb{C}_{\mathbb{R}}$. The operator A can be represented as

(11.2)
$$z \mapsto az + b\bar{z} = az(1 + \mu \frac{z}{z}),$$

where $\mu = b/a$ is called the *Beltrami coefficient* of A. Let $\mu = |\mu|e^{2i\theta}$, where $\theta \in \mathbb{R}/\pi\mathbb{Z}$.

EXERCISE 2.5. A is conformal iff $\mu = 0$. A is invertible iff $|\mu| \neq 1$. A is orientartion preserving iff $|\mu| < 1$.

picture

In what follows we assume that A is an invertible orientation preserving operator, i.e., $|\mu| < 1$. If we have another form $A' = a'z + b'\bar{z}$ on V then $A/A \equiv \text{const}$ iff $\mu = \mu'$. Thus, the *conformal structures* $A^*\sigma$ are in one-to-one correspondence with the Beltrami coefficient $\mu \in \mathbb{D}$, so $\text{Conf}(V) \approx \mathbb{D}$.

Let us now describe the shape of the ellipse $A^{-1}(\mathbb{T})$ in terms of μ . The maximum of |Az| on the unit circle $\mathbb{T} = \{z = e^{i\phi}\}$ is attained at the direction $\phi = \theta \mod \pi \mathbb{Z}$, while the minimum is attained at the orthogonal direction $\theta + \pi/2 \mod \pi \mathbb{Z}$. These are the eigenlines l_{\max} and l_{\min} of the positive part P of A. The corresponding eigenvalues are equal to

$$\sigma_{\max} = |a|(1+|\mu|) = |a|+|b|, \quad \sigma_{\min} = |a|(1-|\mu|) = |a|-|b|.$$

Thus

(11.3)
$$\operatorname{Dil} A = \frac{1+|\mu|}{1-|\mu|}, \quad \det A = |a|^2 - |b|^2 = \sigma_{\min} \operatorname{Dil} A.$$

This gives is a description of the dilatation and orientation of the ellipse $E = A^{-1}(\mathbb{T})$ are described in terms $|\mu|$ and $\arg \mu$ respectively.

EXERCISE 2.6. Show that the correspondence $\operatorname{Conf}(V) \approx \mathbb{D}$ is a hyperbolic isometry equivariant with respect to the standard actions of $\operatorname{SL}^{\#}(2,\mathbb{R})$ on $\mathbb{C}_{\mathbb{R}}$ and \mathbb{D} .

Under conformal changes of variable, $z = A\zeta = \alpha\zeta$ ($\alpha \in \mathbb{C}^*$) the Beltrami coefficients is rotated: $\nu := A^*\mu = (\bar{\alpha}/\alpha)\mu$, while the (-1,1)-form

$$\mu \frac{\bar{z}}{z} = \nu \frac{\zeta}{\zeta}$$

does not change. It shows that the Betrami coefficients in various conform coordinates represent a single (-1,1)-"Beltrami form".

In what follows we will feel free to identify confomal structures with the corresponding Beltrami forms (and in a particual coordinate, with the corresponding Beltrami coefficients). We will often use the same notation for these objects.

11.1.3. Infinitesimal notation. Let us now interpret the above discussion in infinitesimal terms. Consider a map $h: U \to \mathbb{C}$ on a domain $U \subset \mathbb{C}$ differentiable at a point $z \in U$, and apply the above considerations to its differential Dh(z): $T_z U \to T_{hz} \mathbb{C}$. In the (dz, dz)-coordinates of the tangent spaces, it assumes the form

$$\partial h + \bar{\partial} h = \partial_z h \, dz + \partial_{\bar{z}} h \, \bar{dz},$$

where the partial derivatives ∂_z and $\overline{\partial}_z$ and the operators ∂ and $\overline{\partial}$ are defined in §1.8. Moreover,

$$Dh(z) = \partial_z h(z) dz \left(1 + \mu_h(z) \frac{dz}{d\bar{z}} \right),$$

where $\mu_h = \partial_{\bar{z}} h / \partial_z h$ is the *Beltrami coefficient* of h at z. In fact, as was explained above, these coefficients represent a (-1, 1)-form

$$\bar{\partial}h/\partial h=\mu_h\frac{d\bar{z}}{dz}$$

called the *Beltrami differential* of h at z. However, in what follows we will not make a notational difference between the Beltrami differential and the coefficient (and will usually use notation ∂ , $\bar{\partial}$ for the partial derivatives ∂_z , $\partial_{\bar{z}}$).

Assume that Dh(z) is non-singular and orientation preserving, i.e., $|\mu_h| < 1$. The map h is conformal at z if and only if $\mu_h(z) = 0$, which is equivalent to the Cauchy-Riemann equation $\bar{\partial}h(z) = 0$.

Let us consider an infinitesimal ellipse

(11.4)
$$E_h(z) \equiv Dh(z)^{-1}(\mathbb{T}_{hz}) \subset \mathcal{T}_z U,$$

where \mathbb{T}_{hz} is a round circle in the tangent space $\mathbb{T}_{hz}U$. If h is not conformal at z, then $E_h(z)$ is a genuine (not round) ellipse with the small axis in the direction $\arg(\mu_h(z))/2 \mod \pi$ and the shape

(11.5)
$$\operatorname{Dil}(h, z) = \frac{1 + |\mu_h(z)|}{1 - |\mu_h(z)|}.$$

Moreover, by the second formula of (11.3), we have:

(11.6)
$$\operatorname{Jac}(h, z) = |\partial h(z)|^2 - |\bar{\partial} h(z)|^2 = \sigma_{\min}(z) \operatorname{Dil}(h, z)$$

where $\operatorname{Jac}(h, z) \equiv \det Dh(z)$ and $\sigma_{\min}(z) = \inf_{|v|=1} Dh(z) v$.

11.2. Measurable conformal structures. A (measurable) conformal structure on a domain $U \subset \mathbb{C}$ is a measurable family of conformal structures in the tangent planes T_zU , $z \in U$. In other words, it is a measurable family \mathcal{E} of infinitesimal ellipses $E(z) \subset T_zU$ defined up to scaling by a measurable function $\rho(z) > 0$, $z \in U$. (As always in the measurable category, all the above objects are defined almost everywhere.) According to the linear discussion, any conformal structure is determined by its *Beltrami coefficient* $\mu(z)$, $z \in U$, a measurable function in z assuming its values in \mathbb{D} , and vice versa. Thus, conformal structures on U are described analytically as elements μ from the unit ball of $L^{\infty}(U)$. We say that a conformal structure has a *bounded dilatation* if the dilatations of the ellipses E(z)are bounded almost everywhere. In terms of Beltrami coefficients, it means that $\|\mu\|_{\infty} < 1$ since

$$\text{Dil}\,\mu = \|\,\text{Dil}\,E(z)\|_{\infty} = \frac{1+\|\mu\|_{\infty}}{1-\|\mu\|_{\infty}}$$

The standard conformal structure σ is given by the family of infinitesimal circles. The corresponding Beltrami coefficient vanishes almost everywhere: $\mu = 0$ in $L^{\infty}(U)$.

The space of conformal structures on S with bounded dilatation is endowed with the Teichmüller metric:

$$\operatorname{dist}_{\mathrm{T}}(\mu,\nu) = \|\operatorname{dist}_{\mathrm{T}}(\mu(z),\nu(z))\|_{\infty}.$$

Denote by $\mathcal{D}^+(U, V)$ (standing for "differentiable homeomorphisms") the space of orientation preserving homeomorphisms $f: U \to V$, which are differentiable almost everywhere with a non-singular differential Df(z) measurably depending on z.¹ Consider some homeomorphism $h \in \mathcal{D}^+(U, V)$ between two domains in \mathbb{C} . Then by the above linear discussion we obtain a measurable family \mathcal{E} of infinitesimal ellipses $E_h(z) = Dh(z)^{-1}(\mathbb{T}_{hz}) \subset \mathbb{T}_z U$ that determines a (measurable) conformal structure $\mu_h = h^* \sigma$ on U. Analytically this structure can be described as the

picture

¹If we do not need to specify the domain and the range of h we write simply $h \in D^+$; if we do not assume that f is orientation preserving, we skip "+".

Beltrami coefficient $\mu_h(z) = \partial h(z)/\partial h(z)$ of h. We say that h has a bounded dilatation if the corresponding conformal structure $h^*\sigma$ does. In this case we let

Dil
$$h := \| \text{Dil}(h, z) \|_{\infty} = \frac{1 + \|\mu_h\|_{\infty}}{1 - \|\mu_h\|_{\infty}} = \text{dist}_{\mathrm{T}}(h^* \sigma, \sigma).$$

Obviously, the pullback structure $h^*\sigma$ does not change if we *postcompose* h with a conformal map ϕ . If we *precompose* h with a conformal map ϕ then the Beltrami coefficient will be transformed as follows:

$$\mu_{h\circ\phi} = \frac{\overline{\phi'}}{\phi'}\,\mu_h,$$

so that the Beltrami coefficients in various local charts represent a single (-1,1)-form $\mu d\bar{z}/dz$ called the Beltrami differential of h (compare §11.1.3).

This allows us to generalize the above discussion to arbitrary Riemann surfaces. A (measurable) conformal structure on a Riemann surface S is a measurable family \mathcal{E} of infinitesimal ellipses E(z) defined up to scaling. Analytically it is described as a measurable *Beltrami differential* μ with $|\mu(z)| < 1$ a.e. To any homeomorphism $h \in D^+(S, S')$ between two Riemann surfaces corresponds the pullback structure $h^*\sigma$ represented by the field of ellipses $E_h(z) = Dh(z)^{-1}(\mathbb{T}_r)$.² The corresponding Beltrami differential $\mu_h = \bar{\partial}h/\partial h$ (where $\bar{\partial}h$ and ∂h are now viewed as 1-forms).

Remark. A key problem is whether any conformal structure \mathcal{E} is associated to a certain map h. This problem has a remarkable positive solution in the category of quasiconformal maps (see §2.34 below).

Let us consider a smaller class $AC^+(S, S') \subset D^+(S, S')$ of absolutely continuous orientation preserving homeomorphisms between Riemann surfaces S and S'.³ Then we can naturally pull back any measurable conformal structure μ' on S' to obtain a conformal structure $\mu = h^*(\mu')$ on S. If h^{-1} is also absolutely continuous then we can push forward the structures: $\mu' = h_*(\mu)$.

More generally, let us consider a (non-invertible) map $f: U \to V$ which locally belongs to class AC⁺ outside a finite set of "critical points". For such maps the push-forward operation is not well-defined, but the pullback $\nu = f^*\mu$ is still welldefined. The fact that f has critical points does not cause any troubles since we need to know μ only almost everywhere. The property that $\text{Dil}(f^*\mu) \leq \text{Dil}(f) \cdot \text{Dil}(\mu)$ is obviously valid in this generality.

11.3. Distributional derivatives and absolute continuity on lines. Let U be a domain in $\mathbb{C} \equiv \mathbb{C}_{\mathbb{R}}$. All functions below are assumed to be complex valued. A *test function* ϕ on U is an infinitely differentiable function with compact support. One says that a locally integrable function $f : U \to \mathbb{C}$ has *distributional partial derivatives of class* L^1_{loc} if there exist functions h and g of class L^1_{loc} on U such that for any test function ϕ ,

$$\int_{U} f \cdot \partial \phi dm = -\int_{U} h \phi dm; \quad \int_{U} f \cdot \bar{\partial} \phi dm = -\int_{U} g \phi dm,$$

where m is the Lebesgue measure. In this case h and g are called ∂ and $\overline{\partial}$ derivatives of f in the sense of distributions. Clearly this notion is invariant under smooth

²Note that the ellipses $E_h(z)$ are defined only up to scaling since the round circles \mathbb{T}_r on S' are (as there is no preferred metric on S').

 $^{{}^3\!\}mathrm{Reminder}:\,h$ is absolutely continuous if for any set X of zero Lebesgue measure, the preimage $h^{-1}X$ has also zero measure.

changes of variable, so that it makes sense on any smooth manifold (and for all dimensions).

EXERCISE 2.7. Prove that a function f on the interval (0,1) has a destributional derivative of class L^1_{loc} if and only if it is absolutely continuous. Moreover, its classical derivative f'(x) coincides with the distributional derivative.

There is a similar criterion in the two-dimensional setting. A continuous function $f: U \to \mathbb{C}$ is called *absolutely continuous on lines* if for any family of parallel lines in any disk $D \in U$, f is absolutely continuous on almost all of them. Thus, taking a typical line l of the above family, the curve $f: l \to \mathbb{C}$ is rectifiable. Clearly such functions have classical partial derivatives almost everywhere.

PROPOSITION 2.8. Consider a homeomorphism $f: U \to V$ between two domains in the complex plane. It has distributional partial derivatives of class L^1_{loc} if and only if it is absolutely continuous on lines.

In fact, in the proof of existence of distributional partial derivatives (the easy direction of the above Proposition), just two transversal families of parallel lines are used. Thus one can relax the definition of absolutele continuity on lines by taking any two directions ("horizontal" and "vertical").

PROPOSITION 2.9. Consider a homeomorphism $f: U \to V$ which is absolutely continuous on lines. Then for almost any $z \in U$, f is differentiable at z in the classical sense, i.e., $f \in D$.

This result can be viewed as a measurable generalization of the elementary fact that existence of continuous partial derivatives implies differentiability.

11.4. Definition. We are now ready to give a definition of quasiconformality. An orientation preserving homeomorphism $f: S \to S'$ between two Riemann surfaces is called quasi-conformal if

- It has locally integrable distributional partial derivatives;
- It has bounded dilatation.

Note that the second property makes sense because the first property implies that f is differentiable a.e. in the classical sense (by the results of §11.3).

We will often abbreviate "quasiconformal" as "qc". A qc map f is called K-qc if $\text{Dil}(f) \leq K$.

A map $f: S \to S'$ is called K- quasiregular if for any $z \in S$ there exist K-qc local charts $\phi: (U, z) \to (\mathbb{C}, 0)$ and $\psi: (V, f(z)) \to (\mathbb{C}, 0)$ such that $\psi \circ f \circ \phi^{-1}: z \mapsto z^d$. Sometimes we will abbreviate K-quasiregular maps as "K-qr". A map is called quasiregular if it is K-qr for some K.

EXERCISE 2.10. Show that any quasiregular map $f: S \to S'$ can be decomposed as $g \circ h$, where $h: S \to T$ is a qc map to some Riemann surface T and $g: T \to S'$ is holomorphic. In particular, if $S = S' = \hat{\mathbb{C}}$ then also $T = \mathbb{C}$ and $g: \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ is a rational map.

11.5. Absolute continuity and Sobolev class H. We will now prove several important regularity properties of quasi-conformal maps. Let us define a Sobolev class H = H(U) as the space of uniformly continuous functions $f: U \to \mathbb{C}$ whose distributional partial derivatives on U belong to $L^2(U)$. The norm on H is the maximum of the uniform norm of f and L^2 -norm of its partial derivatives.

Infinitely smooth functions are dense in H. This can be shown by the standard regularization procedure: convolute f with a sequence of functions $\phi_n(x) = n^2 \phi(n^{-1}x)$, where ϕ is a non-negative test function on U with $\int \phi \, dm = 1$ (see [?, Ch V, §2.1] or [**LV**, Ch. III, Lemma 6.2]).

PROPOSITION 2.11. Quasiconformal maps are absolutely continuous with respect to the Lebesgue measure, and thus for any Borel set $X \subset U$,

$$m(fX) = \int_X \operatorname{Jac}(f, z) \, dm.$$

The partial derivatives ∂f and $\bar{\partial} f$ belong to L^2_{loc} .

PROOF. Since both statements are local, we can restrict ourselves to homeomorphisms $f: U \to U'$ between domains in the complex plane. Consider the pull-back of the Lebesgue measure on U', $\mu = f^*m$. It is a Borel measure defined as follows: $\mu(X) = m(fX)$ for any Borel set $X \subset U$. Let us decompose it into absolutely continuous and singular parts: $\mu = h \cdot m + \nu$. By the Lebesgue Density Points Theorem, for almost all $z \in U$, we have:

$$\frac{1}{\pi\epsilon^2}\int_{\mathbb{D}(z,\epsilon)}h\;dm\to h(z);\quad \frac{1}{\pi\epsilon^2}\nu(\mathbb{D}(z,\epsilon))\to 0\quad \text{as}\quad \epsilon\to 0.$$

Summing up we obtain:

$$\frac{m(f(\mathbb{D}(z,\epsilon))}{m(\mathbb{D}(z,\epsilon))} = \frac{\mu(\mathbb{D}(z,\epsilon)}{m(\mathbb{D}(z,\epsilon))} \to h(z) \quad \text{as} \quad \epsilon \to 0.$$

But if f is differentiable at z then the left hand-side of the last equation goes to $\operatorname{Jac}(f, z)$. Hence $\operatorname{Jac}(f, z) = h(z)$ a.e. It follows that for any Borel set X,

(11.7)
$$\int_X \operatorname{Jac}(f, z) \, dm = \int_X h \, dm \le \mu(X) = m(fX).$$

But $Jac(f, z) = |\bar{\partial}f(z)|^2 - |\partial f(z)|^2 \ge (1 - k^2) |\partial f(z)|^2$, where $k = \|\mu_f\|_{\infty}$. Thus

(11.8)
$$\int_X |\partial f|^2 \, dm \le \frac{1}{1-k^2} \, m(fX); \quad \int_X |\bar{\partial} f|^2 \, dm \le \frac{k^2}{1-k^2} \, m(fX),$$

and we see that the partial derivatives of f are locally square integrable.

What is left is to prove the opposite to (11.7). As we have just shown, f locally belongs to the Sobolev class H. Without loss of generality we can assume that it is so on the whole domain U, i.e., $f \in H(U)$. Let us approximate f in H(U)by a sequence of C^{∞} functions f_n . Take a domain $D \subseteq U$ with piecewise smooth boundary (e.g., a rectangle).

Let $V_n \subset f_n D$ be the set of regular values of f_n . By Sard's Theorem, it has full measure in $f_n D$. Let $R = f_n^{-1} V_n \cap D$. Note that the $\int_{R_n} \operatorname{Jac} f_n dm$ is equal to the area of the image of $f_n | R_n$ counted with multiplicities:

$$\int_{R_n} \operatorname{Jac}(f_n, z) \, dm = \int_{V_n} \operatorname{card}(f_n^{-1}\zeta) \, dm \ge m(V_n) = m(f_n D)$$

Since $f_n \to f$ uniformly on D, $\liminf m(f_n D) \ge m(fD)$. Since $\operatorname{Jac}(f_n) \to \operatorname{Jac}(f)$ in $L^1(U)$,

$$\int_{R} \operatorname{Jac}(f_{n}, z) \, dm \to \int_{R} \operatorname{Jac}(f, z) \, dm \le \int_{D} \operatorname{Jac}(f, z) \, dm.$$

Putting the last estimates together, we obtain the desired estimate for D.

For an arbitrary Borel set $X \subset U$, the result follows by a simple approximation argument using a covering of X by a union of rectangles D_i with disjoint interiors such that $m(\bigcup D_i \setminus X) < \epsilon$.

12. Geometric definitions

Besides the analytic definition given above, we will give two geometric definitions of quasi-conformality, in terms of quasi-invariance of moduli, and in terms of bounded circular dilatation (or, "quasi-symmetricity").

12.1. Quasi-invariance of moduli. In this section we will show, by the length-area method, that the moduli of annuli are quasi-invariant under qc maps. This will follow from a more general result on quasi-invariance of extremal length:

LEMMA 2.12. Let $h: U \to \tilde{U}$ be a K-qc map. Let Γ and $\tilde{\Gamma} = f(\Gamma)$ be two families of rectifiable curves in the respective domains such that h is absolutely continuous on all curves of Γ . Then $\mathcal{L}(\Gamma) \leq K\mathcal{L}(\tilde{\Gamma})$.

PROOF. To any measurable metric ρ on U, we are going to associate a metric $\tilde{\rho}$ on \tilde{U} such that $h^*(\tilde{\rho}) \geq \rho$ while $h^*(m_{\tilde{\rho}}) \leq Km_{\rho}$ (so, the map h is expanding with respect to these metrics, with area expansion bounded by K). Then $\rho(\tilde{\gamma}) \geq \rho(\gamma)$ for any $\gamma \in \Gamma$ and $\tilde{\gamma} = f(\gamma) \in \tilde{\Gamma}$, while $m_{\tilde{\rho}}(\tilde{U}) \leq Km_{\rho}(U)$. Hence $\mathcal{L}_{\tilde{\rho}}(\tilde{\Gamma}) \geq K^{-1}\mathcal{L}_{\rho}(\Gamma)$. Taking the supremum over all metrics ρ , we obtain the desired estimate.

To define correspondence $\rho \mapsto \tilde{\rho}$, recall formula (11.6) relating the Jacobian and the minimal expansion. Letting $\tilde{\rho}(hz) = \rho(z)/\sigma_{\min}(z)$, we obtain for a.e. $z \in U$ and any unit tangent $v \in T_z U$:

$$|h^*(d\tilde{\rho})v| = \tilde{\rho}(hz)|Dh(z)v| \ge d\rho(v)$$

and

$$h^*(dm_{\tilde{\rho}}) = \tilde{\rho}(hz)^2 \operatorname{Jac} h(z) \, dx dy = K(z)\rho(z)^2 dx dy \le K dm_{\rho},$$

which are the required properties of the metrics.

PROPOSITION 2.13. Consider a K-qc map $h: A \to \tilde{A}$ between two topological annuli. Then

$$K^{-1} \operatorname{mod}(\tilde{A}) \le \operatorname{mod}(A) \le K \operatorname{mod}(\tilde{A}).$$

PROOF. Let $\tilde{\Gamma}$ be the family of genuinely vertical paths on \tilde{A} on which h^{-1} is absolutely continuous, and let $\Gamma = h^{-1}(\tilde{\Gamma})$. By Proposition 1.59, mod $\tilde{A} = \mathcal{L}(\tilde{\Gamma})$, while mod $A \leq \mathcal{L}(\Gamma)$. By Lemma 2.12, $\mathcal{L}(\Gamma) \leq K\mathcal{L}(\tilde{\Gamma})$, which yields the desired right hand-side estimate. The left-hand side estimate is obtained by replacing hwith h^{-1} .

EXERCISE 2.14. Show that the moduli of rectangles are quasi-invariant in the same sense as for the annuli.

EXERCISE 2.15. Prove that \mathbb{C} and \mathbb{D} are not qc equivalent.

12.2. Macroscopic and upper dilatation. According to the original analytic definition of qc maps, they have bounded infinitesimal dilatation a.e. It turns out that this property can be substantially strengthened: in fact, qc maps have bounded macroscopic dilatation in sufficiently small scales *everywhere*.

Let $h: U \to V$ be a homeomorphism between two domains, and let $D := \mathbb{D}(z,\rho) \subset U$. Then we can define the *macroscopic dilatation* $\mathrm{Dil}(h,z,\rho)$ as the shape of h(D) around h(z) (as for conformal maps in §4.4). Recall also from §4.4 the definitions of the inner and outer radii of a pointed domain.

LEMMA 2.16. Let $h: U \to V$ be a K-qc homeomorphism. Let $D = \mathbb{D}(z, \rho) \subset U$ and $\mathbb{D}(h(z), R) \subset V$, where R is the outer radius of h(D). Then

$$\operatorname{Dil}(h, z, \rho) \le \exp CK,$$

where C an absolute constant.

PROOF. For notational convenience, let us normalize h so that z = h(z) = 0, and let r be the inner radius of h(D). Let a and b be two points on the circle \mathbb{T}_{ρ} for which |h(a)| = r and |h(b)| = R. Let us consider the annulus $A' = \mathbb{A}(r, R) \subset V$ and let $A = h^{-1}(A')$. The inner component of $\mathbb{C} \smallsetminus A$ contains points 0 and $a \in \mathbb{T}_{\rho}$, while its outer component of $\mathbb{C} \backsim A$ contains $b \in \mathbb{T}_{\rho}$. By Lemma 1.62, mod A is bounded by an absolute constant C. By Lemma 2.13,

$$\frac{1}{2\pi}\log\frac{R}{r} = \mod A' \le K \mod A \le KC,$$

and we are done.

The upper dilatation of h at z is defined as

$$\overline{\mathrm{Dil}}(h,z) = \limsup_{\rho \to 0} \mathrm{Dil}(h,z,\rho).$$

(Of course, if h is differentiable at z then $\overline{\text{Dil}}(h, z) = \text{Dil}(h, z)$.) We define the upper dilatation of h as

$$\overline{\mathrm{Dil}}(h) = \sup_{z \in U} \overline{\mathrm{Dil}}(h, z).$$

Lemma 2.16 immediately implies:

PROPOSITION 2.17. Any K-qc map $U \rightarrow V$ has a bounded upper dilatation:

$$\overline{\mathrm{Dil}}(h) \le \exp CK,$$

where C is an absolute constant.

12.3. Quasisymmetry.

12.3.1. Generalities. We will now give a characterization of qc maps that can be applied in a very general setting. For a triple of points (x, y, z) in a metric space X, let the brackets

$$[y, z]_x := \frac{\operatorname{dist}(z, x)}{\operatorname{dist}(y, x)}$$

denote the distance ratio centered at x.

Let $\eta : \mathbb{R}_+ \to \mathbb{R}_+$ be a function such that $\eta(t) \to as \ t \to 0$. An embedding $h : X \to Y$ between two metric spaces is called η -quasisymmetric (" η -qs") if for any triple of points (x, y, z) in X we have:

(12.1)
$$[y, z]_x \le t \implies [h(y), h(z)]_{h(x)} \le \eta(t).$$

A map h is called *quasisymmetric* if it is η -qs for some η . Such an h distorts the ratios in a controlled way.

The function $\eta(t)$ is called the *qs dilatation* of *h*.

EXERCISE 2.18. Show that the dilatation function η can be selected as a homeomorphism $\mathbb{R}_+ \to \mathbb{R}_+$.

For instance, bi-Lipschitz homeomorphisms with constant L are η -qs with linear dilatation $\eta(t) = L^2 t$. However, the class of qs maps is much bigger:

EXERCISE 2.19. The power homeomorphisms of \mathbb{R} , $x \mapsto \operatorname{sign}(x)|x|^{\delta}$, are quasisymmetric. What are their qs dilatations?

Qs maps can surve as *morphisms* of the category of metric spaces:

EXERCISE 2.20. The compositions and the inverse of qs maps are qs. Calculate their qs diltations.

In particular, qs maps $X \to X$ form a group.

12.3.2. Qc vs qs. The most important value of the dilatation function $\eta(t)$ is $\eta(1)$ that controls macroscopic dilatation of h on the balls and (as we will see momentarily) often controls the full $\eta(t)$.

LEMMA 2.21. An embedding $h : \mathbb{R}^m \to \mathbb{R}^n$ is η -qs if and only if (quantitatively) it has L-bounded macroscopic dilatation: $\operatorname{Dil}(h, z, \rho) \leq L$ for all discs $\mathbb{D}(z, \rho)$.

PROOF. Obviously, quasisymmetry implies that macroscopic dilatation is bounded by $L = \eta(1)$. Vice versa, bounded macroscopic dilatation implies (12.1) with a function $\eta(t) = O(t^{\alpha})$ as $t \to \infty$, where the exponent $\alpha \ge 1$ depends only on the dilatation.

EXERCISE 2.22. Prove this assertion and calculate $\eta(t)$ in terms of $L = \eta(1)$.

What is more subtle is to show that $\eta(t) \to 0$ as $t \to 0$.

Let us take a triple of points x, y, z, and let x', y', z' stand for their images under h (in what follows, the images of points under h will be marked with the "prime" as well). Property (12.1) implies:

(12.2)
$$[z,y]_x \ge 1 \implies [z',y']_{x'} \ge \epsilon = 1/L > 0$$

By making affine changes of variable in the domain and the target, we can normalize the situation so that x = x' = 0, |y| = |y'| = 1, $z = R \in \mathbb{R}$, $z' = r \in \mathbb{R}$. Of course, we can assume that R > 1. We want to show that $r \to \infty$ as $R \to \infty$. Let us partition the interval [0, z] by points $z_n = z/2^n$, $n = 0, 1, \ldots, N$, where N is selected so that $z_N \in [1, 2)$. So, $N \ge \log_2 R - 1 \to \infty$ as $R \to \infty$.

Applying (12.2) to the triple of points $(0, 1, z_N)$, we obtain: $|z'_N| \ge \epsilon$. Then applying it inductively (backwards) to the triples $(z_n, 0, z_{n-1})$ (centered at z_{n-1}), we conclude that

$$|z'_n - z'_{n-1}| \ge \epsilon |z'_{n-1}| \ge \epsilon^2$$
,

so the net of points z'_n is ϵ^2 -separated. On the other hand, applying (12.2) to the triple $(0, z_n, z)$, we conclude that $|z'| \ge \epsilon |z'_n|$, so that all the points z'_n belong to the disc $\mathbb{D}_{r/e}$. Hence the discs of radius $\epsilon^2/2$ centered at the z_n are pairwise dosjoint and are contained in the disc $\mathbb{D}_{2r/e}$. It follows that

$$N \leq \frac{\operatorname{area} \mathbb{D}_{2r/\epsilon}}{\operatorname{area} \mathbb{D}_{\epsilon^2/2}} = \frac{16}{\epsilon^6} r^2,$$

and hence $r \ge c\sqrt{\log R}$ with c > 0 depending only on L.

In the light of the above result, embeddings $h : \mathbb{R}^w \to \mathbb{R}^n$ with *L*-bounded macroscopic dilatation will also be referred to as "*L*-qs". (We hope that this slight terminological inconsistency will not cause confusion).

Putting together Propositions 2.16 and 2.21, we obtain:

PROPOSITION 2.23. There is an L depending only on K such that:

- (i) Any K-qc homeomorphism $h : \mathbb{C} \to \mathbb{C}$ is L-qs (in the Euclidean metric);
- (ii) Any K-qc homeomorphism $h : \overline{\mathbb{C}} \to \overline{\mathbb{C}}$ fixing 0,1 and ∞ is L-qs (in the spherical metric).

EXERCISE 2.24. Without normalization, the last assertion would fail.

12.4. Back to the analytic definition.

PROPOSITION 2.25. If a homeomorphism $h: U \to V$ between domains U and V has an L-bounded circular dilatation then it is L-qc.

PROOF. Since the L-bounded circular dilatation implies the L-bounded infinitesimal dilatation at any point of differentiability, all we need to show is that h has the required regularity, i.e., it is absolutely continuous on almost all parallel lines. Since this is a local property, we can assume that U us the unit square, and that the parallel lines in question are horizontal.

Let $U_b = \{z \in U : \text{ Im } z \leq b\}$. Since the area function

 $\mu: b \mapsto \operatorname{area}(h(U_b))$

is monotonic, it is differentiable for a.e. b. Let us take such a point b where μ is differentiable, and prove absolute continuity of h on the corresponding line $\gamma_b = \{z : \text{Im } z = b\}.$

For $K \in \mathbb{N}$, let $X_K = \{z \in \gamma_b : L_h(x, \epsilon) \leq K/2 \text{ for } \epsilon \leq 1/K\}$. Since the dilatation of h is bounded⁴, we have: $\bigcup X_K = \gamma_b$. Hence it is enough to prove that $h \mid X_K$ is absolutely continuous.

Let $Q \subset X_K$ be a set of zero length. We want to show that h(Q) has zero length as well. By approximation, it is sufficient to show for closed sets. Then Qcan be covered with finitely many disks $D_i = \mathbb{D}(z_i, \epsilon)$ $(z_i \in \gamma_b, i = 1, ..., n)$ with intersection multiplicity at most 2 and an arbitrary small total length. Hence for any $\delta > 0$, we have $n\epsilon \leq \delta$ once ϵ is sufficiently small.

Let $M_i = M_h(z_i, \epsilon)$ and $m_i = m_h(z_i, \epsilon)$. Then $M_i \leq km_i$, $l(hX) \leq \sum M_i$, and by the Cauchy-Bunyakovsky inequality,

$$l(hX)^2 \le n \sum M_i^2 \le nK^2 \sum m_i^2 \le \frac{K^2\delta}{\pi} \cdot \frac{\operatorname{area}(h \cup D_i))}{\epsilon}.$$

But the last ratio is bounded by

$$\frac{\mu(b+\epsilon)-\mu(b-\epsilon)}{\epsilon} \to \frac{1}{2}\,\mu'(b) \quad \text{as $\epsilon \to 0$},$$

and the desired conclusion follows.

 $^{{}^{4}}$ For the regularity purposes, it is sufficient to assume that the circular dilatation is finite everywhere.

2. QUASICONFORMAL GEOMETRY

13. Further important properties of qc maps

13.1. Weyl's Lemma. This lemma asserts that a 1-qc map is conformal. In other words, if a qc map is infiniesimally conformal on the set of full measure (i.e., $\bar{\partial}h(z) = 0$ a.e.), then it is conformal in the classical set. Since $\bar{\partial}h(z) = 0$ is just the Cauchy-Riemann equation, this statement is classical for smooth maps.

Let us formulate a more general version of Weyl's Lemma:

WEYL'S LEMMA (Weyl's Lemma). Assume that a continuous function $h: U \to \mathbb{C}$ has distributional derivatives of class L^1_{loc} . If $\bar{\partial}h(z) = 0$ a.e., then h is holomorphic.

PROOF. By approximation, Weyl's Lemma can be reduced to the classical statement. Since the statement is local, we can assume without loss of generality that the partial derivatives of h belong to $L^1(U)$. Convoluting h with smooth bump-functions we obtain a sequence of smooth functions $h_n = h * \theta_n$ converging to h uniformly on U with derivatives converging in $L^1(U)$. Let us show that $\bar{\partial}h_n = 0$. For a test function η on U, we have:

$$\int \bar{\partial}h_n(z) \eta(z) dm(z) = -\int h_n(z) \bar{\partial}\eta(z) dm(z)$$
$$= -\int h(\zeta) dm(\zeta) \int \theta_n(z-\zeta) \bar{\partial}\eta(z) dm(z)$$
$$= \int h(\zeta) dm(\zeta) \int \bar{\partial}\theta_n(z-\zeta)\eta(z) dm(z)$$
$$= \int \eta(z) dm(z) \int h(\zeta) \bar{\partial}\theta_n(z-\zeta) dm(\zeta)$$
$$= \int \eta(z) dm(z) \int \bar{\partial}h(\zeta) \theta_n(z-\zeta) dm(\zeta) = 0$$

Here the first and the third equalities are the classical integration by parts, the next to the last one comes from the definition of the distributional derivative, and the intermediate ones come from the Fubini Theorem.

It follows that the smooth functions h_n satisfy the Cauchy-Riemann equations and hence holomorphic. Since uniform limits of holomorphic functions are holomorphic, h is holomorphic as well.

13.2. Devil Staircase. The following example shows that Weyl's Lemma is not valid for homeomorphisms of class D (i.e., differentiable a.e.). The technical assumption that the classical derivative can be understood in the sense of distributions (which allows us to integrate by parts) is thus crucial for the statement.

Take the standard Cantor set $K \subset [0,1]$ and construct a devil staircase $h : [0,1] \rightarrow [0,1]$, i.e., a continuous monotone function which is constant on the complementary gaps to K.

EXERCISE 2.26. Do the construction. (Topologically it amounts to showing that by collapsing the gaps to points we obtain a space homeomorphic to the interval.)

Consider a strip $S = [0,1] \times \mathbb{R}$ and let $f : (x,y) \mapsto (x,y+h(x))$. This is a homeomorphism on S which is a rigid translation on every strip $G \times R$ over a gap $G \subset [0,1] \setminus K$. Since $m(K \times \mathbb{R}) = 0$, this map is conformal a.e. However it is obviously not conformal on the whole strip P.

Clearly f in not absolutely continuous on the horizontal lines: it translates them to devil staircases.

13.3. Quasiconformal Removability and Gluing. A closed set $K \subset \mathbb{C}$ is called *qc removable* if any homeomorphism $h: U \to \mathbb{C}$ defined on an neighborhood U of K, which is quasiconformal on $U \smallsetminus K$, is quasiconformal on U.

Remark. We will see later on (§??) that qc removable sets have zero measure and hence $\text{Dil}(f|U) = \text{Dil}(f|U \smallsetminus K)$.

EXERCISE 2.27. Show that isolated points are removable.

PROPOSITION 2.28. Smooth Jordan arcs are removable.

PROOF. Let us consider a smooth Jordan arc $\Gamma \subset U$ and a homeomorphism $f: U \to \mathbb{C}$ which is quasi-conformal on $U \smallsetminus \Gamma$. We should check that f is absolutely continuous on lines near any point $z \in \Gamma$. Take a small box B centered at z whose sides are not parallel to $T_z\Gamma$. Then any interval l in B parallel to one of its sides intersects Γ at a sinle point ζ . Since for a typical l, f is absolutely continuous on the both sides of $l \smallsetminus \{\zeta\}$, it is absolutely continuous on the whole interval l as well.

Moreover, Dil(f) is obviously bounded since it is so on $U \smallsetminus \Gamma$ and Γ has zero measure.

The above statement is simple but important for holomorphic dynamics. It will allow us to construct global qc homeomorphisms by gluing together different pieces without spoiling dilatation.

Let us now state a more delicate gluing property:

BERS' LEMMA. Consider a closed set $K \subset \overline{\mathbb{C}}$ and two its neighborhoods U and V. Assume that we have two quasi-conformal maps $f: U \smallsetminus K \to \widehat{\mathbb{C}}$ and $g: V \to \widehat{\mathbb{C}}$ that match on ∂K , i.e., the map

$$h(z) = \begin{cases} f(z), & z \in U \smallsetminus K \\ g(z), & z \in K \end{cases}$$

is continuous. Then h is quasi-conformal and $\mu_h(z) = \mu_q(z)$ for a.e. $z \in K$.

PROOF. Consider a map $\phi = f^{-1} \circ h$. It is well-defined in a neighborhood Ω of K, is identity on K and is quasi-conformal on $\Omega \setminus K$. Let us show that it is quasi-conformal on Ω . Again, the main difficulty is to show that h is abosultely continuous on lines near any point $z \in K$.

Take a little box near some point $z \in K$ with sides parallel to the coordinate axes. Without loss of generality we can assume that $z \neq \infty$ and ϕB is a bounded subset of \mathbb{C} . Let ψ denote the extension of $\partial \phi / \partial x$ from $B \setminus K$ onto the whole box B by 0. By (11.8), ψ is square integrable on B and hence it is square integrable on almost all horizontal sections of B. All the more, it is integrable on almost all horizontal sections. Take such a section I, and let us show that ϕ is absolutely continuous on it.

Let $I_j \subset I$ be a finite set of disjoint intervals; $\Delta \phi_j$ denote the increment of ϕ on I_j . We should show that

(13.1)
$$\sum |\Delta \phi_j| \to 0 \text{ as } \sum |I|_j \to 0.$$

Take one interval I_j and decompose it as $L \cup J \cup R$ where $\partial J \subset K$ and int L and int R belong to $B \setminus K$. Then

$$|\Delta \phi_j| \le |J| + \int_{L \cup R} g \, dx \le |I_j| + \int_{I_j} g \, dx.$$

Summing up the last estimates over j and using integrability of g on I_j , we obtain (13.1).

Absolute continuity on the vertial lines is treated in exactly the same way. \Box

13.4. Weak topology in L^2 . Before going further, let us briefly recall some background in functional analysis. Consider the space $L^2 = L^2(X)$ on some measure space (X, m). A sequence of functions $h_n \in L^2$ weakly converges to some function $h \in L^2$, $h_n \xrightarrow[]{w} h$, if for any $\phi \in L^2$, $\int h_n \phi \, dm \to \int h \phi \, dm$. The main advantage of this topology is the property that the balls of L^2 are weakly compact (see e.g., [?,]). Note also that vice versa, any weakly convergent sequence belongs to some ball in L^2 (Banach-Schteinhaus [?,]).

However, one should handle the weak topology with caution: for instance, product is not a weakly continuous operation:

EXERCISE 2.29. Show that $\sin nx \xrightarrow{w} 0$ in $L^2[0, 2\pi]$, while $\sin^2 nx \xrightarrow{w} 1/2$.

At least, the weak topology respects the order:

EXERCISE 2.30. Let $h_n \xrightarrow{w} h$.

- If $h_n \ge 0$ then $h \ge 0$;
- If $h_n = 0$ a.e. on some subset $Y \subset X$, then h = 0 a.e. on Y;
- After selecting a further subsequence,

$$(h_n)_+ \xrightarrow{w} h_+$$
 and $(h_n)_- \xrightarrow{w} h_-$, so that $|h_n| \xrightarrow{w} |h|$.

Here $h_+(z) = \max(h(z), 0), h_(z) = \min(h(z), 0).$

13.5. Compactness. We will proceed with the following fundamental property of qc maps:

THEOREM 2.31. The space of K-qc maps $h : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ fixing 0,1 and ∞ is compact in the topology of uniform convergence on $\hat{\mathbb{C}}$

PROOF. It will be more convenient to consider the space \mathcal{X} of K-qc maps h such that $h\{0, 1, \infty\} = \{0, 1, \infty\}$. First, we will show that the family of maps $h \in \mathcal{X}$ is equicontinuous. Otherwise we would have an $\epsilon > 0$, a sequence of maps $h_n \in \mathcal{X}$, and a sequence of points $z_n, \zeta_n \in \hat{\mathbb{C}}$ such that such that $d(z_n, \zeta_n) \to 0$ while $d(h_n z_n, h_n \zeta_n) \ge \epsilon$, where d stands for the special metric. By compactness of $\hat{\mathbb{C}}$, we can assume that the $z_n, \zeta_n \in \hat{\mathbb{C}}$ converge to some point a and the $h_n(a)$ converge to some b. Postcomposing or/and precomposing if necessary the maps h_n 's with $z \mapsto 1/z$, we can make $|a| \le 1$, $|b| \le 1$.

Consider a sequence of annuli $A_n = \{z : r_n < |z - a| < 1/2\}$ where $r_n = \max(|z_n - a|, |\zeta_n - a|) \to 0$. Since the disk $\mathbb{D}(a, 1/2)$ does not contain one of the points 0 or 1, its images $h_n(\mathbb{D}(a, 1/2))$ have the same property. Hence the Euclidean distance from the point $h_n(a)$ (belonging to the inner complement of $h_n(A_n)$) to the outer complement of that annulus is eventually bounded by 3. On the other hand, the diameter of the inner complement of $h_n(A_n)$ is bounded from below by $\epsilon > 0$.

By Lemma 1.62, $\operatorname{mod}(h_n(A_n))$ is bounded from above. But $\operatorname{mod}(A_n) = 1/r_n \to 0$ contradicting quasi-invariance of the modulus (Proposition 2.13).

Hence \mathcal{X} is precompact in the space of continuous maps $\hat{\mathbb{C}} \to \hat{\mathbb{C}}$. Since \mathcal{X} is invariant under taking the inverse $h \mapsto h^{-1}$, and the composition is a continuous operation in the uniform topology, \mathcal{X} is precompact in Homeo($\hat{\mathbb{C}}$). Since Homeo⁺($\hat{\mathbb{C}}$) is closed in Homeo($\hat{\mathbb{C}}$), \mathcal{X} is precompact in the former space as well.

To complete the proof, we should show that the limit functions are also K-qc homeomorphisms. Let a sequence $h_n \in \mathcal{X}$ uniformly converges to some h. Given a point $a \in \hat{\mathbb{C}}$, we will show that in some neighborhood of a, f has distributional derivatives of class L^2 . Without loss of generality we can assume that $a \in \mathbb{C}$. Take a neighborhood $B \ni a$ such that h(B) is a bounded subset of \mathbb{C} . Then the neighborhoods $h_n(B)$ are eventually uniformly bounded. By (11.8), the partial derivatives ∂h_n and $\bar{\partial} h_n$ eventually belong to a fixed ball of $L^2(D)$. Hence they form weakly precompact sequences, and we can select limits along subsequences (without changing notations):

$$\partial h_n \xrightarrow{} \phi \in L^2(D); \qquad \bar{\partial} h_n \xrightarrow{} \psi \in L^2(D).$$

It is straightforward to show that ϕ and ψ are the distributional partial derivatives of h. Indeed, for any test functions η we have:

(13.2)
$$\int h \,\partial\eta \,dm = \lim \int h_n \,\partial\eta \,dm = -\lim \int \partial h_n \,\eta \,dm = -\int \phi \,\eta \,dm,$$

and the similarly for the ∂ -derivative.

What is left is to show that $|\phi(z)| \leq k|\psi(z)|$ for a.e. z, where k = (K - 1)/(K + 1). To see this, select a further subsequence in such a way that $|\partial h_n| \xrightarrow{w} |\phi|$, $|\bar{\partial}h_n| \xrightarrow{w} |\psi|$ and use the fact that the weak topology respects the order (see Exercise 2.30).

EXERCISE 2.32. Fix any three points a_1, a_2, a_3 on the sphere \mathbb{C} . A family \mathcal{X} of K-qc maps $h: \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ is precompact in the space of all K-qc homeomorphisms of the sphere (in the uniform topology) if and only if the reference points are not moved close to each other (or, in formal words: there exists a $\delta > 0$ such that $d(ha_i, ha_j) \geq \delta$ for any $h \in \mathcal{X}$ and $i \neq j$, where d is the spherical metric). Consider first the case K = 0.

We will also need a disk version of the above Compactness Theorem:

COROLLARY 2.33. The space of K-qc homeomorphisms $f : \mathbb{D} \to \mathbb{D}$ fixing 0 is compact in the topology of uniform convergence on \mathbb{D} .

PROOF. Let \mathcal{Y} be the space of K-qc homeomorphisms $h : \mathbb{D} \to \mathbb{D}$ fixing 0, and \mathcal{X} be the space of \mathbb{T} -symmetric K-qc homeomorphisms $H : \mathbb{C} \to \mathbb{C}$ fixing 0 and ∞ . (To be \mathbb{T} -symmetric means to commute with the involution $\tau : \mathbb{C} \to \mathbb{C}$ with respect to the circle.) Clearly maps $H \in \mathcal{X}$ preserve the unit circle (the set of fixed points of τ); in particular, they do not move 1 close to 0 and ∞ . By Theorem 2.31 (and the Exercise following it), \mathcal{X} is compact.

Let us show that \mathcal{X} and \mathcal{Y} are homeomorphic. The restriction of a map $H \in \mathcal{X}$ to the unit disk gives a continuous map $i : \mathcal{X} \to \mathcal{Y}$. The inverse map $i^{-1} : \mathcal{Y} \to \mathcal{X}$ is given by the following extension procedure. First, extend $h \in \mathcal{Y}$ continuously to the closed disk \mathbb{D} (Theorem ??), and then reflect it symmetrically to the exterior of the

disk, i.e., let $H(z) = \tau \circ h \circ \tau(z)$ for $z \in \mathbb{C} \setminus \overline{\mathbb{D}}$. Since τ is an (orientation reversing) conformal map, H is K-qc on $\overline{\mathbb{C}} \setminus \mathbb{T}$. By Lemma 2.28, it is K-qc everywhere, and hence belongs to \mathcal{X} .

Hence \mathcal{Y} is compact as well.

14. Measurable Riemann Mapping Theorem

We are now ready to prove one of the most remarkable facts of analysis: any measurable conformal structure with bounded dilatation is generated by a quasiconformal map:

THEOREM 2.34 (Measurable Riemann Mapping Theorem). Let μ be a measurable Beltrami differential on $\overline{\mathbb{C}}$ with $\|\mu\|_{\infty} < 1$. Then there is a quasi-conformal map $h: \overline{\mathbb{C}} \to \overline{\mathbb{C}}$ which solves the Beltrami equation: $\overline{\partial}h/\partial h = \mu$. This solution is unique up to post-composition with a Möbius automorphism of $\widehat{\mathbb{C}}$. In particular, there is a unique solution fixing three points on $\widehat{\mathbb{C}}$ (say, 0, 1 and ∞).

The local version of this result sounds as follows:

THEOREM 2.35 (Local integrability). Let μ be a measurable Beltrami differential on a domain $U \subset \mathbb{C}$ with $\|\mu\|_{\infty} < 1$. Then there is a quasi-conformal map $h: U \to \mathbb{C}$ which solves the Beltrami equation: $\bar{\partial}h/\partial h = \mu$. This solution is unique up to post-composition with a conformal map.

The rest of this section will be occupied with a proof of these two theorems.

14.1. Uniqueness. Uniqueness part in the above theorems is a consequence of Weyl's Lemma. Indeed, if we have two solutions h and g, then the composition $\psi = g \circ h^{-1}$ is a qc map with $\bar{\partial}\psi = 0$ a.e. on its domain. Hence it is conformal.

14.2. Local vs global. Of course, the global Riemann Measurable Riemann Theorem immediately yields the local integrability (e.g., by zero extantion of μ from U to the whole sphere). Vice versa, the global result follows from the local one and the classical Uniformization Theorem for the sphere . Indeed, by local integrability, there is a finite covering of the sphere $S^2 \equiv \hat{\mathbb{C}}$ by domains U_i and a family of qc maps $\phi_i : U_i \to \mathbb{C}$ solving the Beltrami equation on U_i . By Weyl's Lemma, the gluing maps $\phi_i \circ \phi_j^{-1}$ are conformal. Thus the family of maps $\{\phi_i\}$ can be interpreted as a complex analytic atlas on S^2 , which endows it with a new complex analytic structure m (compatible with the original qc structure). But by the Uniformization Theorem, all complex analytic structures on S^2 are equivalent, so that there exists a biholomorphic isomorphism $h : (S^2, m) \to \hat{\mathbb{C}}$. It means that the maps $h \circ \phi_i^{-1}$ are conformal on $\phi_i U_i$. Hence h is quasi-conformal on each U_i and $h_*(\mu) = (h \circ \phi_i^{-1})_* \sigma$ over there. Since the atlas is finite, h is a global quasi-conformal solution of the Beltrami equation.

14.3. Strategy. The further strategy of the proof will be the following. First, we will solve the Beltrami equation locally assuming that the coefficient μ is real analytic. It is a classical (and elementary) piece of the PDE theory. By the Uniformization Theorem, it yields a global solution in the real analytic case. Approximating a measurable Beltrami coefficient by real analytic ones and using compactness of the space of normalized K-qc maps, we will complete the proof.

14.4. Real analytic case. Assume that μ is a real analytic Beltrami coefficient in a neighborhood of 0 in $\mathbb{R}^2 \equiv \mathbb{C}_{\mathbb{R}}$ with $|\mu(0)| < 1$. Then it admits a complex analytic extension to a neighborhood of 0 in the complexification \mathbb{C}^2 . Let (x, y) be the standard coordinates in \mathbb{C}^2 , and let u = x + iy, v = x - iy. In these coordinates the complexified Beltrami equation assumes the form:

(14.1)
$$\frac{\partial h}{\partial v} - \mu(u, v)\frac{\partial h}{\partial u} = 0$$

This is a linear equation with variable coefficients, which can be solved by the standard method of characteristics. Namely, let us consider a vector field $W(u, v) = (1, -\mu(u, v))$ near 0 in \mathbb{C}^2 . Since the left-hand side of (14.1) is the derivative of h along X, we come to the equation Wh = 0. Solutions of this equation are the first integrals of the ODE $\dot{w} = W$. But since W is non-singular at 0, this ODE has a non-singular local first integral h(u, v). Restricting h to \mathbb{R}^2 , we obtain a local solution $h: (\mathbb{R}^2, 0) \to \mathbb{C}$ of the original Beltrami equation. Since h is non-singular at 0, it is a local (real analytic) diffeomorphism.

By means of the Uniformization Theorem, we can now pass from local to global solutions of the Beltrami equation with a real analytic Beltrami differential $\mu(z)d\bar{z}/dz$ on the sphere (see §14.2). Note that the global solution is real analytic as well since the complex structure generated by the local solutions is compatible with the original real analytic structure of the sphere (as local solutions are real analytic).

EXERCISE 2.36. For a real analytic Beltrami coefficient

$$\mu(z) = \sum a_{n,m} z^n \bar{z}^m$$

on \mathbb{C} , find the condition of its real analyticity at ∞ .

There is also a "semi-local" version of this result:

If μ is a real analytic Beltrami differential on the disk \mathbb{D} with $\|\mu\|_{\infty} < 1$, then there is a quasi-conformal (real analytic) diffeomorphism $h : \mathbb{D} \to \mathbb{D}$ solving the Beltrami equation $\overline{\partial}h/\partial h = \mu$.

To see it, consider the complex structure m on the disk generated by the local solutions of the Beltrami equation. We obtain a simply connected Riemann surface $S = (\mathbb{D}, m)$. By the Uniformization Theorem, it is conformally equivalent to either the standard disk (\mathbb{D}, σ) or to the complex place \mathbb{C} . But S is quasi-conformally equivalent to the standard disk via the identical map id : $(\mathbb{D}, m) \to (\mathbb{D}, \sigma)$. By Exercise 2.15, it is then conformally equivalent to the standard disk, and this equivalence $h : (\mathbb{D}, m) \to (\mathbb{D}, \sigma)$ provides a desired solution of the Beltrami equation.

By 14.1 Such a solution is unique up to a postcomposition with a Möbius automorphism of the disk.

14.5. Approximation. Let us consider an arbitrary measurable Beltrami coefficient μ on a disk \mathbb{D} with $\|\mu\| < \infty$. Select a sequence of real analytic Beltrami coefficients μ_n on \mathbb{D} with $\|\mu_n\|_{\infty} \leq k < 1$, converging to μ a.e.

EXERCISE 2.37. Construct such a sequence (first approximate μ with continuous Beltrami coefficients).

Applying the results of the previous section, we find a sequence of quasiconformal maps $h_n : (\mathbb{D}, 0) \to (\mathbb{D}, 0)$ solving the Beltrami equations $\bar{\partial}h_n/\partial h_n = \mu_n$. The dilatation of these maps is bounded by K = (1 + k)/(1 - k). By Corollary 2.33, they form a precompact sequence in the topology of uniform convergence on the disk. Any limit map $h : \mathbb{D} \to \mathbb{D}$ of this sequence is a quasi-conformal homeomorphism of \mathbb{D} . Let us show that its Beltrami coefficient is equal to μ .

By (11.8), the partial derivatives of the h_n belong to some ball of the Hilbert space $L^2(\mathbb{D})$. Hence we can select weakly convergent subsequences $\partial h_n \to \phi$, $\bar{\partial} h_n \to \psi$. We have checked in (13.2) that $\phi = \partial h$ and $\psi = \bar{\partial} h$. What is left is to check that $\psi = \mu \phi$. To this end, it is enough to show that $\mu_n \partial h_n \to \mu \phi$ weakly (to appreciate it, recall that the product is not weakly continuous, see Exercise 2.29). For any test function $\eta \in L^2(\mathbb{D})$, we have:

$$\left| \int (\eta \mu \phi - e t a \mu_n \, \partial h_n) \, dm \right| \leq \\ \leq \left| \int \eta \mu (\phi - \partial h_n) \, dm \right| + \int |\eta (\mu - \mu_n) \, \partial h_n| \, dm$$

The first term in the last line goes to 0 since the ∂h_n weakly converge to ϕ . The second term is estimated by the Cauchy-Schwarz inequality by $\|\eta(\mu-\mu_n)\|_2 \|\partial h_n\|_2$, which goes to 0 since $\mu_n \to \mu$ a.e. and the ∂h_n belong to some Hilbert ball. This yields the desired.

It proves the Measurable Riemann Mapping Theorem on the disk \mathbb{D} , which certainly implies the local integrability. Now the global theorem on the sphere follows from the local integrability by §14.2. This completes the proof.

14.6. Conformal and complex structures. Let us discuss the general relation between the notions of complex and conformal structures. Consider an oriented surface S endowed with a *qs structure*, i.e., supplied with an atlas of local charts $\psi_i : V_i \to \mathbb{C}$ with uniformly qc transit maps $\psi_i \circ \psi_j^{-1}$ ("uniformly qc" means "with uniformly bounded dilatation"). Note that a notion of a measurable conformal structure with bounded dilatation makes perfect sense on such a surface (in what follows we call it just a "conformal structure").

Endow S with a complex structure compatible with its qs structure. By definition, it is determined by an atlas $\phi_i : U_i \to \mathbb{C}$ on S of uniformly qc maps such that the transit maps are complex analytic. Then the conformal structures $\mu_i = \phi_i^*(\sigma)$ on U_i coincide on the intersections of the local charts and have uniformly bounded dilatations. Hence they glue into a global conformal structure on S.

Vice versa, any conformal structure μ determines by the Local Integrability Theorem a new complex structure on the surface S compatible with its qc structure (see §14.2).

Thus the notions of conformal and complex structures on a qc surface are equivalent. In what follows we will not distinguish them either conceptually or notationally.

Fixing a reference complex structure on S (so that S becomes a Riemann surface), complex/conformal structures on S get parametrized by measurable Beltrami differentials μ on S with $\|\mu\|_{\infty} < 1$.

14.7. Moduli spaces. Consider some qc surface S (with or without boundary, possibly marked or partially marked).

The moduli space $\mathcal{M}(S)$, or the deformation space of S is the space of all conformal structures on S compatible with the underlying qc structure, up to the action of qc homeomorphisms perserving the marked data. In other words, $\mathcal{M}(S)$

is the space of all Riemann surfaces qc equivalent to S, up to conformal equivalence relation (respecting the marked data).

If we fix a reference Riemann surface S_0 , then its deformations are represented by qc homeomorphisms $h: S_0 \to S$ to various Riemann surfaces S. Two such homeomorphisms h and \tilde{h} represent the same point of the moduli space if there exists a conformal isomorphism $A: S \to \tilde{S}$ such that the composition $H = \tilde{h}^{-1} \circ A \circ h: S_0 \to S_0$ respects all the marked data. In particular, H = id on the marked boundary. In the case when the whole fundamental group is marked, H must be homotopic to the id relative to the marked boundary.

For instance, if S has a finite conformal type, i.e., S is a Riemann surface of genus g with n punctures (without marking), then $\mathcal{M}(S)$ is the classical moduli space $M^{g,n}$. If S is fully marked then $\mathcal{M}(S)$ is the classical Teichmüller space $T^{g,n}$. This space has a natural complex structure of complex dimension 3g - 3 + n for g > 1. For g = 1 (the torus case), dim $T^{1,0} = 1$ (see §1.6.2) and dim $T^{1,n} = n - 1$ for $n \ge 1$. For g = 0 (the sphere case), dim $T^{0,n} = 0$ for $n \le 3$ (by the Riemann-Koebe Uniformization Theorem and 3-transitivity of the Möbius group action) and dim $T^{0,n} = n - 3$ for n > 3.

EXERCISE 2.38. What is the complex modulus of the four punctured sphere?

There is a natural projection (fogetting the marking) from $T^{g,n}$ onto $M^{g,n}$. The fibers of this projection are the orbits of the so called "Teichmüller modular group" acting on $T^{g,n}$ (it generalizes the classical modular group $PSL(2,\mathbb{Z})$, see §1.6.2).

By the Riemann Mapping Theorem, the disk \mathbb{D} does not have moduli. However, if we mark its boundary \mathbb{T} , then the space of moduli, $\mathcal{M}(\mathbb{D}, \mathbb{T})$, becomes infinitely dimensional! By definition, $\mathcal{M}(\mathbb{D}, \mathbb{T})$ is the space of all Beltrami differentials μ on \mathbb{D} up to the action of the group of qc homeomorphisms $h: \mathbb{D} \to \mathbb{D}$ whose boundary restrictions are Möbius: $h|\mathbb{T} \in \mathrm{PSL}(2,\mathbb{R})$. It is called the *universal Teichmüller space*, since it contains all other deformation spaces. This space has several nice descriptions, which will be discussed later on. It plays an important role in holomorphic dynamics.

14.8. Dependence on parameters. It is important to know how the solution of the Beltrami equation depends on the Beltrami differential. It turns out that this dependence is very nice. Below we will formulate three statements of this kind (on continuous, smooth and holomorphic dependence).

PROPOSITION 2.39. Let μ_n be a sequence of Beltrami differentials on \mathbb{C} with uniformly bounded dilatation, converging a.e. to a differential μ . Consider qc solutions $h_n : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ and $h : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ of the corresponding Beltrami equations fixing $0, 1 \text{ and } \infty$. Then the h_n converge to h uniformly on \mathbb{C} .

PROOF. By Theorem 13.5, the sequence h_n is precompact. Take any limit map g of this sequence. By the argument of §14.5, its Beltrami differential is equal to μ . By uniqueness of the normalized solution of the Beltrami equation, g = h. The conclusion follows.

Consider a family of Beltrami differentials μ_t depending on parameters $t = (t_1, \ldots, t_n)$ ranging over a domain $U \subset \mathbb{R}^n$. This family is said to be *differentiable* at some $t \in U$ if there exist Beltrami differentials α_t^i of class $L^{\infty}(\mathbb{C})$ (but not necessarily

in the unit ball of this space) such that for all sufficiently small $\epsilon = (\epsilon_1, \ldots, \epsilon_n) \in \mathbb{R}^n$, we have:

$$\mu_{t+\epsilon} - \mu_t = \sum_{i=1}^n \alpha_t^i \epsilon_i + |\epsilon| \,\beta(t,\epsilon),$$

where the norm $\|\beta_{t,\epsilon}\|_{\infty}$ stays bounded and $\beta_{t,\epsilon}(z) \to 0$ a.e. on \mathbb{C} as $\epsilon \to 0$.

Assume additionally that the family μ_t is differentiable at all points $t \in U$, that the norms $\|\alpha_t^i\|$ are locally bounded, and that the $\alpha_t^i(z)$ continuously depend on t in the sense of the convergence a.e. Then the family μ_t is said to be *smooth*.

Let us now consider a family of qc maps $h_t : \mathbb{C} \to \mathbb{C}$ depending on parameters $t \in U$. Considering these maps as elements of the Sobolev space H, we can define differentiability and smoothness in the usual way. This family is differentiable at some point $t \in U$ if there exist vector fields v_t^i on \mathbb{C} of Sobolev class H such that

$$h_{t+\epsilon} - h_t = \sum_{i=1}^n \epsilon_i v_t^i + |\epsilon| g_{t,\epsilon},$$

where $g_{t,\epsilon} \to 0$ in the Sobolev norm as $\epsilon \to 0$ (in particular $g_{t,\epsilon} \to 0$ uniformly on the sphere). If additionally the v_t^i depend continuously on t (as elements of H), then one says that h_t smoothly depends on t. Of course, in this case, any point $z \in \mathbb{C}$ smoothly moves as parameter t changes, i.e., $h_t(z)$ smoothly depends on t.

THEOREM 2.40. If μ_t , $t \in U \subset \mathbb{R}^n$, is a smooth family of Beltrami differentials, then the normalized solutions $h_t : \mathbb{C} \to \mathbb{C}$ of the corresponding Beltrami equations smoothly depend on t.

Let us finally discuss the holomorphic dependence on parameters. Let U be a domain in \mathbb{C}^n and let \mathcal{B} be a complex Banach space. A function $f: U \to \mathcal{B}$ is called *holomorphic* if for any linear functional $\phi \in \mathcal{B}^*$, the function $\phi \circ f: U \to \mathbb{C}$ is holomorphic. Beltrami differentials are elements of the complex Banach space L^{∞} , while qc maps $h: \mathbb{C} \to \mathbb{C}$ are elements of the complex Sobolev space H. So, it makes sense to talk about holomorphic dependence of these objects on complex parameters $t = (t_1, \ldots, t_n) \in U$. Note that if h_t depends holomorphically on t, then any point $z \in \mathbb{C}$ moves holomorphically as t changes (in fact, holomorphic dependence on parameters is often understood in this weak sense).

THEOREM 2.41. If the Beltrami differential μ_t holomorphically depends on parameters $t \in U$, then so do the normalized solutions $h_t : \mathbb{C} \to \mathbb{C}$ of the corresponding Beltrami equations.

The proofs of the last two theorems can be found in [**AB**]. 14.8.1. *Simple conditions.*

LEMMA 2.42. Let \mathcal{B} be a Banach space, and let $\{f_{\lambda}\}, \lambda \in \mathbb{D}_{\rho}$, be a uniformly bounded family of linear functionals on \mathcal{B} such that for a dense linear subspace Lof points $x \in \mathcal{B}$, the function $\lambda \mapsto f_{\lambda}(x)$ is holomorphic in λ . Then $\{f_{\lambda}\}$ as an element of the dual space \mathcal{B}^* depends holomorphically on λ .

PROOF. For $x \in L$, we have a power series expansion

$$f_{\lambda}(x) = \sum a_n(x)\lambda^n$$

convergent in \mathbb{D}_{ρ} . By the Cauchy estimate,

$$|a_n(x)| \le \frac{C||x||}{\rho^n}, \quad x \in L,$$

100

be careful here!

where C is an upper bound for the norms $||f_{\lambda}||, \lambda \in \mathbb{D}_{\rho}$. Clearly, the $a_n(x)$ linearly depend on $x \in L$. Hence, a_n are bounded linear functionals on L; hence they admit an extension to bounded linear functionals on \mathcal{B} . Moreover, $||a_n|| \leq C\rho^{-n}$. It follows that the power series $\sum a_n \lambda^n$ converges in the dual space \mathcal{B}^* uniformly in λ over any disk \mathbb{D}_r , $r < \rho$. Hence it represents a holomorphic function $D_{\rho} \mapsto \mathcal{B}^*$, which, of course, coincides with $\lambda \mapsto f_{\lambda}$.

For further applications, let us formulate one simple condition of holomorphic dependence:

LEMMA 2.43. Let $\rho > 0$ and let $U \subset \mathbb{C}$ be an open subset in \mathbb{C} of full measure. Let $\mu_{\lambda} \in L^{\infty}(\mathbb{C}), \ \lambda \in \mathbb{D}_{\rho}$, be a family of Beltrami differentials with $\|\mu_{\lambda}\|_{\infty} \leq 1$ whose restriction to U is smooth in both variables (λ, z) and is holomorphic in λ . Then $\{\mu_{\lambda}\}$ is a holomorphic family of Beltrami differentials.

PROOF. Let us first assume that $U = \hat{\mathbb{C}}$. Then

$$\mu_{\lambda}(z) = \sum a_n(z)\lambda^n, \quad \lambda \in \mathbb{D}_{\rho}$$

where the a_n are smooth functions on $\hat{\mathbb{C}}$, and the series converges uniformly over $\hat{\mathbb{C}} \times \mathbb{D}_r$ for any $r < \rho$. It follows that the series $\sum a_n \lambda^n$ in L^{∞} converges uniformly over \mathbb{D}_r and hence represents a holomorphic function $\mathbb{D}_r \to L^{\infty}$.

Let us now consider the general case; put $K = \hat{\mathbb{C}} \setminus U$. Consider a sequence of smooth functions $\chi_l : \hat{\mathbb{C}} \to [0,1]$ such that $\chi_l = 0$ on K and for any $z \in U$, $\chi_l(z) \to 1$ as $l \to \infty$.

Consider smooth Beltrami differentials $\mu_{\lambda}^{l} = \chi_{l}\mu_{\lambda}$. By the above consideration, they depend holomorphically on λ . Moreover, since K has zero area, $\chi_{l}\mu_{\lambda} \to \mu_{\lambda}$ a.e. as $l \to \infty$. Note also that $\|\mu_{\lambda}^{l}\|_{\infty} \leq 1$.

Take any smooth test function ϕ on $\hat{\mathbb{C}}$ and let

$$g_l(\lambda) = \int \mu_\lambda^l \phi dA; \quad g(\lambda) = \int \mu_\lambda \phi dA,$$

where dA is the (normalized) area element on $\hat{\mathbb{C}}$. The family $\{g_l\}$ is uniformly bounded: $|g_l(\lambda)| \leq ||\phi||_{\infty}$ By the Lebesgue Bounded Convergence Theorem, $g_l(\lambda) \rightarrow g(\lambda)$ as $l \rightarrow \infty$

By the previous discussion, functions g_l are holomorphic functions on \mathbb{D}_{ρ} . By the Little Montel Theorem, this family is normal. Hence we can select a subsequence conveging to g uniformly on compact subsets of $\hat{\mathbb{C}}$. It follows that g is holomorphic on \mathbb{D}_{ρ} .

Since smooth functions are dense in L^1 , Lemma 2.42 can be applied. It implies the assertion.

By Lemma ?? from Appendix 14.9, we have:

COROLLARY 2.44. A family of Beltrami differentials μ_{λ} with $\|\mu\|_{\infty} < 1$ is holomorphic if and only if for a.e. z, the function $\lambda \mapsto mu_{\lambda}(z)$ is holomorphic.

COROLLARY 2.45. If $h_{\lambda} : U \to U_{\lambda}$ is a holomorphic motion of a domain XU then the Beltrami differential $\mu_{\lambda} = \bar{\partial}h_{\lambda}/dih_{\lambda}$ depends holomorphically on λ .

EXERCISE 2.46. Let $f: S \to T$ be a holomorphic map between two Riemann surfaces, and let $\{\mu_{\lambda}\}$ be a holomorphic family of Beltrami differentials on T. Then $f^*(\mu_{\lambda})$ is a holomorphic family of Beltrami differentials on S.

2. QUASICONFORMAL GEOMETRY

14.9. Appendix: Holomorphic maps between Banach spaces.

14.9.1. Definition. In this section, all Banach spaces are assumed to be complex. Given a Banach space \mathcal{B} , let $\mathcal{B}(x, r)$ stand for the ball in \mathcal{B} of radius r centered at x, and let $\mathcal{B}_r \equiv \mathcal{B}(0, r)$. The dual space to \mathcal{B} is denoted by \mathcal{B}^* .

Let \mathcal{U} be a domain in \mathcal{B} , and let \mathcal{B}' be another Banach space. A continuous map $f : \mathcal{B} \to \mathcal{B}'$ is called *holomorphic* if for any complex line $\mathcal{L} = \{x + \lambda v\}_{\lambda \in \mathbb{C}}$ (where $x, v \in \mathcal{B}$) and any linear functional $\phi \in (\mathcal{B}')^*$, the restrictin $\phi \circ f | \mathcal{L} \cap \mathcal{U}$ is holomorphic in λ . (As we see, this is essentially one-dimensional notion.)

14.9.2. Smoothness.

LEMMA 2.47. Any holomorphic map is differentiable:

$$f(z + v) = f(z) + Df_z(v) + o(||v||).$$

COROLLARY 2.48. A map $f : (\mathbb{D}_r, 0) \to (\mathcal{B}, 0)$ is holomorphic if and only it admits a power series representation

$$f(\lambda) = \sum_{n=0}^{\infty} a_n \lambda^n,$$

where $a_n \in \mathcal{B}$ and $||a_n|| \leq C\rho^n$ for any $\rho > 1/r$ (with C depending on ρ).

14.9.3. Space L^{∞} .

LEMMA 2.49. Let μ_{λ} be a family in $L^{\infty}(S)$ over a domain $D \subset \mathbb{C}$. It is holomorphic in λ if and only if it is locally bounded and the functions $\lambda \mapsto \mu_{\lambda}(z)$ are holomorphic in λ for a.e. z.

PROOF. Without loss of generality we can assume that $D = \mathbb{D}$ is the unit disk. Assume $\lambda \mapsto \mu_{\lambda}$ is holomorphic over \mathbb{D} . Then by Corollary 2.48, it admits a power series representation

(14.2)
$$\mu_{\lambda}(z) = \sum_{n=0}^{\infty} \mu_n(z) \lambda^n,$$

where $\mu_n \in L^{\infty}$ and $\|\mu_n\|_{\infty} \leq C\rho^n$ for any $\rho > 1$. Hence there exists a subset $X \subset S$ of full measure such that for any $\rho > 1$ we have:

$$\mu_n(z) \le C\rho^n$$

It follows that for any $z \in X$, the function $\lambda \mapsto \mu_{\lambda}(z)$ is holomorphic over \mathbb{D} .

Vice versa, assume that for a.e. $z \in S$, the function $\lambda \mapsto \mu_{\lambda}(z)$ is holomorphic over \mathbb{D} . Then (14.2) holds for a.e. $z \in S$, with

$$\mu_n(z) = \frac{1}{2\pi i} \int_{|\lambda|=r} \frac{\mu_\lambda(z) d\lambda}{(-\lambda)^{n+1}}$$

(for any $r \in (0,1)$). But since $\lambda \mapsto \mu_{\lambda}$ is continuous over \mathbb{D} , it is bounded on the circle $\{|\lambda| = r\}$, implying that

$$\mu_n(z) \le \frac{C}{r^n},$$

with C = C(r) independent of z. Hence $\|\mu_n\|_{\infty} = O(r^{-n})$, and the map $\lambda \mapsto \mu_{\lambda}$ is holomorphic by Corollary 2.48.

LEMMA 2.50. If h_{λ} is a holomorphic motion then the corresponding Beltrami differential $\mu_{\lambda} = \bar{\partial}h_{\lambda}/\partial h_{\lambda}$ depends holomorphically in λ .

15. One-dimensional qs maps, quasicircles and qc welding

15.1. Quasisymmetric 1D maps.

15.1.1. Qs maps of the line. Let us first consider the rel line \mathbb{R} in the Euclidean metric. According to Lemma 2.21, L-qs maps $h : \mathbb{R} \to \mathbb{R}$ can be defined as in terms of bounded macroscopic dilatation. Namely, for any two adjacent intervals $I, J \subset \mathbb{R}$ of equal length, we require:

(15.1)
$$\frac{|f(I)|}{|f(J)|} \le L.$$

It looks at first glance that the class of 1D qs maps is a good analogue of the class of 2D qc maps. However, this impression is superficial: two-dimensional qc maps are fundamentally better than one-dimensional qs maps. For instance, qc maps can be glued together without any loss of dilatation (Lemma 2.28), while qs maps cannot:

EXERCISE 2.51. Consider a map $h : \mathbb{R} \to \mathbb{R}$ equal to id on the negative axis, and equal to $x \mapsto x^2$ on the positive one. This map is not quasi-symmetric, though its restrictions to the both positive and negative axes are.

Another big deficiency of one-dimensional qs maps is that they can well be singular (and typically are in the dynamical setting - see ??), while 2D qc maps are always absolutely continuous (Proposition 13.1).

These advantages of qc maps makes them much more efficient tool for dynamics than one-dimensional qs maps. This is a reason why complexification of one-dimensional dynamical systems is so powerful.

15.1.2. Qs circle maps. Of course, an L-qs circle homeomorphism $h : \mathbb{T} \to \mathbb{T}$ can be defined is the same way as in the case of \mathbb{R} , with understanding of (??) in terms of the circle metric. However, there is a subtle difference between these two cases. Namely, in the line case, the group of 1-qs maps coincides with the group of affine maps $x \mapsto ax + b$, which equal to the group of Möbius automorphisms of \mathbb{R} . On the other hand, in the circle case, only rotations are 1-qs, and in fact,

EXERCISE 2.52. The group of Möbius automorphisms of the circle \mathbb{T} is not uniformly qs.

EXERCISE 2.53. A metric space is called geodesic if any two points in it can be joined with an isometric image of a real interval [x, y]. Assume that X is geodesic and $h: X \to Y$ is κ -qs. Then for any L > 0 there exists an $M = M(\kappa, L) > 0$ such that

 $\operatorname{dist}(a,c) \leq L \operatorname{dist}(a,b) \Rightarrow \operatorname{dist}(h(a),h(c)) \leq M \operatorname{dist}(h(a),h(b)).$

15.2. Ahlfors-Beurling Extension.

15.2.1. Extension from \mathbb{R} . As we know, the class of orientation preserving qs maps on the plane coincides with the class of qc maps (Propositions 2.23 and 2.25). In particular, if we consider a quasi-conformal map $h : \mathbb{C} \to \mathbb{C}$ preserving the real line \mathbb{R} , it restricts to a quasi-symmetric map on the latter. Remarkably, the inverse is also true:

THEOREM 2.54. Any L-qs orientation preserving map $h : \mathbb{R} \to \mathbb{R}$ extends to a K(L)-qc map $H : \mathbb{C} \to \mathbb{C}$. Moreover, this extension can be selected to be affinely

equivariant (i.e, so that it commutes with the action of the affine group $z \mapsto az + b$, $a \in \mathbb{R}_+, b \in \mathbb{R}$).

Proof.

15.2.2. Extension from \mathbb{T} .

15.2.3. Interpolation in \mathbb{A} . Let us now state an Interpolation Lemma in an annulus:

LEMMA 2.55. Let us consider two round annuli $A = \mathbb{A}[1,r]$ and $\tilde{A} = \mathbb{A}[1,\tilde{r}]$, with $0 < \epsilon \leq \mod A \leq \epsilon^{-1}$ and $\epsilon \leq \mod \tilde{A} \leq \epsilon^{-1}$. Then any κ -qs map $h : (\mathbb{T}, \mathbb{T}_r) \to (\tilde{\mathbb{T}}, \tilde{\mathbb{T}}_{\tilde{r}})$ admits a $K(\kappa, \epsilon)$ -qc extension to a map $H : A \to \tilde{A}$.

PROOF. Since A and \tilde{A} are ϵ^2 -qc equivalent, we can assume without loss of generality that $A = \tilde{A}$. Let us cover A by the upper half-plane, $\theta : \mathbb{H} \to A$, $\theta(z) = z^{\frac{-\log ri}{\pi}}$, where the covering group generated by the dilation $T : z \mapsto \lambda z$, with $\lambda = e^{\frac{2\pi^2}{\log r}}$. Let $\bar{h} : (\mathbb{R}, 0) \to (\mathbb{R}, 0)$ be the lift of h to \mathbb{R} such that $\bar{h}(1) \in [1, \lambda) \equiv I_{\lambda}$ and $\bar{h}(1) \in (-\lambda, -1]$ (note that \mathbb{R}_+ covers \mathbb{T}_r , while \mathbb{R}_- covers \mathbb{T}). Moreover, since deg h = 1, it commutes with the deck transformation T.

A direct calculation shows that the dilatation of the covering map θ on the fundamental intervals I_{λ} and $-I_{\lambda}$ is comparable with $(\log r)^{-1}$. Hence \bar{h} is $C(\kappa, r)$ -qs on this interval. By equivariance it is $C(\kappa, r)$ -qc on the rays \mathbb{R}_+ and \mathbb{R}_- .

It is also quasi-symmetric near the origin. Indeed, by the equivariance and normalization,

$$(1+\lambda)^{-1}|J| \le |h(J)| \le (1+\lambda)|J|$$

for any interval J containing 0, which easily implies quasi-symmetry.

Since the Ahlfors-Börling extension is affinely equivariant, the map h extends to a $K(\kappa, r)$ -qc map $\overline{H} : \mathbb{H} \to \mathbb{H}$ commuting with T. Hence \overline{H} descends to a $K(\kappa, r)$ -qc map $H : A \to A$.

15.3. Quasicircles.

15.3.1. *Geometric definition*. Let us start with an intrinsic geometric definition of quasicircles:

DEFINITION 2.56. A Jordan curve $\gamma \subset \mathbb{C}$ is called a κ -quasicircle if for any two points $x, y \in \gamma$ there is an arc $\delta \subset \gamma$ bounded by these points such that

(15.2)
$$\operatorname{diam} \delta \le \kappa |x - y|.$$

A curve is called a quasicircle if it is a κ -quasicircle for some κ . The best possible κ in the above definition is called the *geometric dilatation* of the quasicircle. Let us emphasice that this notion is *global* in the sense that (15.2) should be satisfied in all scales. However, it can be localized as follows:

EXERCISE 2.57. If (15.2) is satisfied for all pairs of points with $|x-y| \leq \epsilon$, then γ is a κ' -quasicircle with κ' depending only on κ and N, where N is the number of arcs of diam ϵ needed to cover γ .

A Jordan disk (either open or closed) is called $(\kappa$ -)quasidisk if it is bounded by a $(\kappa$ -)quasicircle.

EXERCISE 2.58. A Jordan disk D is a κ -quasidisk if and only if the Eulidean path metric on \overline{D} is κ -Lipcshitz equivalent to the Euclidean chordal metric.

105

A *C*-quasicenter of a Jordan curve γ (or, of the corresponding Jordan disk *D*) is a point $a \in D$ such that *D* has a *C*-bounded shape around *a*:

$$\frac{R_D(a)}{r_D(a)} \le C.$$

(Here $R_D(a)$ and $r_D(a)$ are outer and inner radii of D around a, see §4.4.)

EXERCISE 2.59. Any κ -quasidisk has a $C(\kappa)$ -quasicenter.

The shape bound $C(\kappa)$ will often be implicit in our discussion, and sometimes we will even say that D is "centered at a".

15.3.2. Quasirectangles and the cross-ratios. Given four points a, b, c, d on a Jorda curve γ , let $\Pi_{\gamma}(a, b, c, d)$ stand for the corresponding quadrilateral. In case when γ is a quasicircle, this quadrilateral will be called a *quasirectangle*.

LEMMA 2.60. The modulus of a quasirectangle, $mod(\Pi_{\gamma}(a, b, c, d))$, is controlled by the cross-ratio R := [a, b, c, d]. More precisely,

$$0 < \theta_1(R) \le \mod \prod_{\gamma}(a, b, c, d) \le \theta_2(R),$$

where the functions θ_i depend only on the geometric dilatation of γ , and $\theta_1(R) \to \infty$ as $R \to \infty$.

15.3.3. Quasitriangles and ratios. A Jordan domain D with four marked points a, b, c, d such that $a, b, c \in \gamma = \partial D$ while $d \in \operatorname{int} D$ is called a *pointed topological triangle* $\Delta_{\gamma}(a, b, c; d)$. Let as define $\operatorname{mod} \Delta_{\gamma}(a, b, c; d)$ as the extremal length of the family of proper paths $\gamma \subset D$ connecting [a, b] to [c, a] and separating d from [b, c]. In case when γ is a quasicircle centered at d, $\Delta_{\gamma}(a, b, c; d)$ will be called *pointed quasitriangle*.

LEMMA 2.61. The modulus of a quasitriangle, $\operatorname{mod} \Delta_{\gamma}(a, b, c; d)$, is controlled by the ratio R := |b - c|/|b - a|, in the same sense as above.

15.3.4. The Riemann mapping.

PROPOSITION 2.62. A pointed domain (D, a) is a κ -quasicicle centered at a if and only if the Riemann mapping $\phi : (\mathbb{D}, 0) \to (D, a)$ is L-qs.

15.3.5. The main criterion. What makes quasicircles so important is their characterization as qc images of the circle:

THEOREM 2.63. Let a be a quasicenter of a κ -quasidisk D, and let $\phi : (\mathbb{D}, 0) \rightarrow (D, a)$ be the normalized Riemann mapping. Then ϕ admits a $K(\kappa)$ -qc extension to the whole complex plane.

Vice versa, let (D, a) be a pointed Jordan disk such that there exists a K-qc map $h : (\mathbb{C}, \mathbb{D}, 0) \to (\mathbb{C}, D, a)$. Then D is a κ -quasidisk with a quasicenter a.

PROOF. The last assertion follows immediated from the fact that h has L(K)-bounded macroscopic dilatation (by Proposition 2.16).

15.3.6. Compactness in the space of quasicircles. Let $\mathcal{QD}_{\kappa,r}$, r > 0, denote the space of pointed κ -quasidisks (D,0) with $r \leq r_{D,0} \leq R_{D,0} \leq 1/r$, endowed with the Carathéodory topology.

PROPOSITION 2.64. The space $QD_{\kappa,r}$ is compact.

PROOF. Consider a quasidisk $(D,0) \in \mathcal{QD}_{\kappa,r}$. By Theorem 2.63, the normalized Riemann mapping $h : (\mathbb{D}, 0) \to (D, 0)$ admits a K-qc extension to the whole complex plane \mathbb{C} , where K depends only on κ and r. Moreover, $r \leq |h(1)| \leq 1/r$. By the Compactness Theorem (see Exercise 2.32), this family of qc maps is compact in the uniform topology on \mathbb{C} . Since uniform limits of κ -quasidisks are obviously κ -quasidisks, the conclusion follows.

A set is called "0-symmetric" if it is invariant under the reflection with respect to the origin.

EXERCISE 2.65. Let γ be a 0-symmetric κ -quasicircle. Then the eccentricity of γ around 0 is bounded by $2\kappa + 1$.

15.4. Douady-Earle Extension.

16. Removability

16.1. Conformal vs quasiconformal. Similarly to the notion of qc removability introduced in §13.3 we can define conformal removability:

DEFINITION 2.66. A compact subset $X \subset \mathbb{C}$ is called *conformally removable* if for any open sets $U \supset X$ in \mathbb{C} , any homeomorphic embedding $h : U \hookrightarrow \mathbb{C}$ which is conformal on $U \smallsetminus X$ is conformal/qc on U.

In fact, these two properties are equivalent:

PROPOSITION 2.67. Conformal removability is equivalent to qc removability.

Thus, we can unambiguously call a set "removable".

It is classical that isolated points and smooth Jordan curves are conformally removable. Proposition 2.67 implies that they are qc removable as well (which was also shown directly in §13.3 of Ch. 2). Since qc removability is invariant under qc changes of variable, we obtain:

COROLLARY 2.68. Quasicircles are removable.

16.2. Removability and area. The Measurable Riemann Mapping Theorem yields:

PROPOSITION 2.69. Removable sets have zero area.

PROOF. Assume that m(X) > 0. Then there exists a non-trivial Beltrami differential μ supported on X. Let $h : \mathbb{C} \to \mathbb{C}$ be a solution of the corresponding Beltrami equation. Then h is conformal outside X but is not conformal on X. \Box

The reverse is false:

Example 2.1.

16.3. Divergence property.

DEFINITION 2.70. Let us say that a compact set $X \subset \mathbb{C}$ satisfies the *divergence* property if for any point $z \in X$ there exists a nest of annuli $A^n(z)$ around z such that

$$\sum A^n(z) = \infty.$$

Without loss of generality we can assume (and we will always do so) that each annulus in this definition is bounded by two Jordan curves.

LEMMA 2.71. Compact sets satisfying the divergence property are Cantor.

PROOF. Consider any connected component X_0 of X, and let $z \in X_0$. Then the annuli $A^n(z)$ are nested around X_0 . By Corollary 1.64 of the Grötzsch Inequality, X_0 is a single point.

LEMMA 2.72. Let $X \subset \mathbb{C}$ be a compact set satisfying the divergence property. Then for any neighborhood $U \supset X$, any qc embedding $h : U \smallsetminus X \hookrightarrow \mathbb{C}$ admits a homeomorphic extension through X.

PROOF. Let $h: U \smallsetminus X \hookrightarrow \mathbb{C}$ be a K-qc embedding. If $X \subset U' \Subset U$ then h(U') is bounded in \mathbb{C} . So, without loss of generality we can assume that h(U) is bounded in \mathbb{C} .

For $z \in X$, let us consider the nest of annuli $h(A^n(z))$. Since h is quasiconformal,

$$\sum \mod h(A^n(z)) \ge K^{-1} \sum \mod A^n(z) = \infty.$$

Let $\Delta^n(z)$ be the bounded component of $\mathbb{C} \smallsetminus h(A^n(z))$, and let

$$\Delta^{\infty}(z) = \bigcap_{n} D^{n}(z).$$

By Corollary 1.64 of the divergence property, $\Delta^{\infty}(z)$ is a single point $\zeta = \zeta(z)$. Let us extend h through X by letting $h(z) = \zeta$.

This extension is continuous. Indeed, let $D^n(z)$ be the bounded component of $\mathbb{C} \smallsetminus A^n(z)$. Then by Corollary 1.64, diam $D^n(z) \to 0$, so that $D^n(z)$ is a base of (closed) neighborhoods of z. But

diam
$$h(D^n(z)) = \operatorname{diam} \Delta^n(z) \to 0$$
,

which yields continuity of h at z.

Switching the roles of (U, X) and (h(U), h(X)), we conclude that h^{-1} admits a continuous extension through h(X). Hence the extension of h is homeomorphic. \Box

It is worthwhile to note that, in fact, general homeomorphisms extend through Cantor sets:

EXERCISE 2.73. (i) Let us consider two Cantor sets X and \tilde{X} in \mathbb{C} and their respective neighborhoods U and \tilde{U} . Then any homeomorphism $h: U \smallsetminus X \to \tilde{U} \smallsetminus \tilde{X}$ admits a homeomorphic extension through X.

(ii) It was essential to assume that both sets X and \tilde{X} are Cantor! For any compact set $X \subset \mathbb{C}$, give an example of an embedding $h : C \setminus X \hookrightarrow \mathbb{C}$ which does not admit a continuous extension through X.

LEMMA 2.74. Compact sets satisfying the divergence property have zero area.
We will show now that sets satisfying the divergence property are removable, and even in the following stronger sense:

THEOREM 2.75. Let $X \subset \mathbb{C}$ be a compact set satisfying the divergence property. Then for any neighborhood $U \supset X$, any conformal/qc embedding $h: U \smallsetminus X \hookrightarrow \mathbb{C}$ admits a conformal/qc extension through X.

PROOF. Let $h: U \smallsetminus X \hookrightarrow \mathbb{C}$ be a K-qc embedding. By Lemma 2.72, h extends to an embedding $U \hookrightarrow \mathbb{C}$, which will be still denoted by h. Let us show that h belongs to the Sobolev class H(U).

Since X is a Cantor set, it admits a nested base of neighborhoods U^n such that each U^n is the union of finitely many disjoint Jordan diks. Take any $\mu > 0$. By the Gritzsch Inequality, for any $n \in \mathbb{N}$ there is $k = k(\mu, l) > 0$ such that $\operatorname{mod}(\partial U^{n+k}, \partial U^n) \geq \mu > 0$. Let χ_n be the solution of the Dirichlet problem in $U^n \smallsetminus U^{n+k}$ vanishing on ∂U^{n+k} and equal to 1 on ∂U^n . By Theorem 1.74, $D(\chi_n) \leq 1/\mu$.

Let us continuously extend χ to the whole plane in such a way that it vanishes on U^{n+k} and identically equal to 1 on $\mathbb{C} \setminus U^n$. We obtain a piecewice smooth function $\chi : \mathbb{C} \to [0, 1]$, with the jump of the derivative on the boundary of the domains U^n and U^{n+k} .

Let $h_n = \chi_n h$. These are piecewise smooth functions with bounded Dirichlet integral. Indeed,

$$D(h_n) = \int (|\nabla \chi_n|^2 |h|^2 + |\chi_n|^2 |\nabla h|^2) dm \le \operatorname{diam}(h(U))/\mu + C(K)m(h(U)),$$

where $C(K) = (1+k^2)/(1-k^2)$ comes from the area estimate (area estimate). By weak compactness of the unit ball in $L^2(U)$, we can select a converging subsequence $\partial h_n \to \phi$, $\bar{\partial} h_n \to \psi$. But $h_n \to h$ pointwise on $U \smallsetminus X$, so that by Lemma 2.74, $h_n \to h$ a.e. It follows that ϕ and ψ are distributional partial derivatives of h (see (13.2)).

Finally, if h is conformal on $U \smallsetminus X$ then by Weyl's Lemma it is conformal on U.

CHAPTER 3

Elements of Teichmüller theory

17. Holomorphic motions

17.1. Definition. Let $(\Lambda, *)$ be a pointed complex Banach manifold¹ and let $X \subset \overline{\mathbb{C}}$ be an arbitrary subset of the Riemann sphere (can be non-measurable). A holomorphic motion **h** over $(\Lambda, *)^2$ is a family of injections $h_{\lambda} : \overline{\mathbb{C}} \to \overline{\mathbb{C}}, \lambda \in \Lambda$, depending holomorphically on λ (in a weak sense that the functions $z \mapsto h_{\lambda}(z)$ are holomorphic in λ for all $z \in X$) and such that $h_* = \text{id}$. In this situation, we let $X_{\lambda} = h_{\lambda}(X_*)^3$.

Holomorphic functions $\phi_z : \Lambda \to \hat{\mathbb{C}}, \lambda \mapsto h_\lambda(z)$, are called *orbits* of the holomorphic motion. Since the functions h_λ are injective, the orbits do not collide, or equivalently, their graphs $\Gamma_z \subset \Lambda \times \hat{\mathbb{C}}$ are disjoint. Thus, a holomorphic motion provides us with a family of disjoint holomorphic graphs over Λ . We refer to such a family as a *(trivial) holomorphic lamination*. Of course, the above reasoning can be reversed, so that, trivial holomorphic laminations give us an equivalent (dual) way of describing holomorphic motions.

A regularity of a holomorphic motion is the regularity of the maps h_{λ} on X. For instance, a holomorphic motion is called *continuous*, qc, smooth or *bi*-holomorphic if all the maps h_{λ} , $\lambda \in \Lambda$, have the corresponding regularity on X (to make sense of it in some cases we need extra assumptions on X, e.g., opennes).

Notice that a priori we do not impose any regularity on the maps h_{λ} (not even measurability!). A remarkable property of holomorphic motions is that they automatically have nice regularity properties and that they automatically extend to motions of the whole Riemann sphere. This set of properties are usually referred to as the λ -lemma. It will be the subject of the rest of this section.

While dealing with a holomorphic motion of a set X, Y, etc., we let $X_{\lambda} = h_{\lambda}(X)$, $Y_{\lambda} = h_{\lambda}(Y)$, etc. We will refer to the z-variable of a holomorphic motion as the *dynamical variable* (though in general, there is no dynamics in the z-plane). The λ -variable is naturally referred to as the *parameter*.

17.2. Extension to the closure and continuity.

LEMMA 3.1. A holomorphic motion \mathbf{h} of any set $X \subset \hat{\mathbb{C}}$ extends to a continuous holomorphic motion of its closure \overline{X} .

¹We will eventually deal with infinite dimensional parameter spaces, so we need to prepare the background in this generality. However, in the first reading the reader can safely assume that the space Λ is a one-dimensional disk (which is the main case to consider anyway).

 $^{^{2}}$ We will often make a point * implicit in the notation and terminology.

³we will sometimes say briefly that "the sets X_{λ} move holomorphically" or "the set X_* moves holomorphically" without mentioning explicitly the maps h_{λ}

PROOF. If X is finite, there is nothing to prove, so assume it is infinite (or, at least, contains more than two points).

Let us show that the family of orbits ϕ_z , $z \in X$, of our holomorphic motion is normal. To this end, let us remove from X three points $z_i \in X$; let $X' = X \setminus \{z_i\}$ and let ψ_i be the orbits of the points z_i . Since the orbits of a holomorphic motion do not collide, the family of orbits of points $z \in X'$ satisfies the condition of the Refined Montel Theorem (1.87) with exceptional functions ψ_i , and the normality follows.

Let Φ be the closure of the family of orbits in the space $\mathcal{M}(\Lambda)$ of meromorphic functions on Λ . By the Hurwitz Theorem (see §??) the graphs of these functions are disjoint, so they form a holomorphic lamination representing a holomorphic motion of \overline{X} .

Let us keep notation h_{λ} for the extended holomorphic motion, and notation $\phi_z, z \in \overline{X}$, for its orbits.

Let us show that this motion is continuous. Let $\lambda \in \Lambda$, let $z_n \to z$ be a converging sequence of points in \overline{X} , and let $\phi_n \in \Phi$ and $\phi \in \Phi$ be their respective orbits. We want to show that $h_{\lambda}(z_n) \to h_{\lambda}(z)$, or equivalently $\phi_n(\lambda) \to \phi(\lambda)$. But otherwise, the sequence ϕ_n would have a limit point $\psi \in \mathcal{M}(\Lambda)$ such that $\psi(*) = \phi(*)$ while $\psi(\lambda) \neq \phi(\lambda)$, which would contradict to the laminar property of the family Φ .

17.3. Extension of smooth holomorphic motions. In thi short section we will prove a simple extension lemma for smooth holomorphic motions.

LEMMA 3.2 (Local extension). Let us consider a compact set $Q \subset \mathbb{C}$ and a smooth holomorphic motion h_{λ} of a neighborhood U of Q over a Banach domain $(\Lambda, 0)$. Then there is a smooth holomorphic motion H_{λ} of the whole complex plane \mathbb{C} over some neighborhood $\Lambda' \subset \Lambda$ of 0 whose restriction to Q coincides with h_{λ} .

PROOF. We can certainly assume that \overline{U} is compact. Take a smooth function $\phi : \mathbb{C} \to \mathbb{R}$ supported in U such that $\phi | Q \equiv 1$, and let

$$H_{\lambda} = \phi h_{\lambda} + (1 - \phi) \operatorname{id}$$

Clearly H is smooth in both variables, holomorphic in λ , and identical outside U. As $H_0 = \text{id}, H_{\lambda} : \mathbb{C} \to \mathbb{C}$ is a diffeomorphism for λ sufficiently close to 0, and we are done.

17.4. Transverse quasiconformality.

SECOND λ -LEMMA. Let $h_{\lambda} : X \to \mathbb{X}_{\lambda}$ be a holomorphic motion of a set $X \subset \mathbb{C}$ over the disk \mathbb{D} . Then for $|\lambda| \leq r < 1$, the maps h_{λ} are η -quasisymmetric with dilatation η depending only on r.

Given two complex one-dimensional transversals S and T to the lamination \mathcal{F} in $\mathcal{B}_1 \times \mathbb{C}$, we have a holonomy $S \to T$. We say that this map is locally quasiconformal if it admits local quasi-conformal extensions near any point.

Given two points $\lambda, \mu \in \mathcal{B}_1$, let us define the hyperbolic distance $\rho(\lambda, \mu)$ in \mathcal{B}_1 as the hyperbolic distance between λ and μ in the one-dimensional complex slice $\lambda + t(\mu - \lambda)$ passing through these points in \mathcal{B}_1 .

LEMMA 3.3. Holomorphic motion h_{λ} of a set X over a Banach ball \mathcal{B}_1 is transversally quasi-conformal. The local dilatation K of the holonomy from p =

 $(\lambda, u) \in S$ to $q = (\mu, v) \in T$ depends only on the hyperbolic distance ρ between the points λ and μ in \mathcal{B}_1 . Moreover, $K = 1 + O(\rho)$ as $\rho \to 0$.

PROOF. If the transversals are vertical lines $\lambda \times \mathbb{C}$ and $\mu \times \mathbb{C}$ then the result follows from 17.4 by restricting the motion to the complex line joining λ and μ .

Furthermore, the holonomy from the vertical line $\lambda \times \mathbb{C}$ to the transversal S is locally conformal at point p. To see this, let us select a holomorphic coordinates (θ, z) near p in such a way that p = 0 and the leaf via p becomes the parameter axis. Let $z = \psi(\theta) = \epsilon + \ldots$ parametrizes a nearby leaf of the foliation, while $\theta = g(z) = bz + \ldots$ parametrizes the transversal S.

Let us do the rescaling $z = \epsilon \zeta$, $\theta = \epsilon \nu$. In these new coordinates, the above leaf is parametrized by the function $\Psi(\nu) = \epsilon^{-1} \psi(\epsilon \nu)$, $|\nu| < R$, where R is a fixed parameter. Then $\Psi'(\nu) = \psi'(\epsilon \nu)$ and $\Psi''(\nu) = \epsilon \psi''(\epsilon \nu)$. By the Cauchy Inequality, $\Psi''(\nu) = O(\epsilon)$. Moreover, ψ uniformly goes to 0 as $\psi(0) \to 0$. Hence $|\Psi'(0)| =$ $|\psi'(0)| \le \delta_0(\epsilon)$, where $\delta_0(\epsilon) \to 0$ as $\epsilon \to 0$. Thus $\Psi'(\nu) = \delta_0(\epsilon) + O(\epsilon) \le \delta(\epsilon) \to 0$ as $\epsilon \to 0$ uniformly for all $|\nu| < R$. It follows that $\Psi(\nu) = 1 + O(\delta(\epsilon)) = 1 + o(1)$ as $\epsilon \to 0$.

On the other hand, the manifold S is parametrized in the rescaled coordinates by a function $\nu = b\zeta(1+o(1))$. Since the transverse intersection persists, S intersects the leaf at the point $(\nu_0, \zeta_0) = (1, b)(1 + o(1))$ (so that R should be selected bigger than ||b||). In the old coordinates the intersection point is $(\theta_0, z_0) = (\epsilon, b\epsilon)(1+o(1))$.

Thus the holonomy from $\lambda \times \mathbb{C}$ to S transforms the disc of radius $|\epsilon|$ to an ellipse with small eccentricity, which means that this holonomy is asymptotically conformal. As the holonomy from $\mu \times \mathbb{C}$ to \mathcal{T} is also asymptotically conformal, the conclusion follows.

Quasiconformality is apparently the best regularity of holomorphic motions which is satisfied automatically.

18. Moduli and Teichmüller spaces of punctured spheres

18.1. Definitions. Let us consider the Riemann sphere with a tuple of n marked points $\mathcal{P} = (z_1, \ldots, z_n)$ (or, equivalently, n punctures). The punctures are considered to be "coloured", or, in other words, the set \mathcal{P} is ordered. Two such spheres $(\mathbb{C}, \mathcal{P})$ and $(\mathcal{C}, \mathcal{P}')$ are considered to be equivalent if there is a Möbius transformation $\phi : (\mathbb{C}, \mathcal{P}) \to (\mathbb{C}, \mathcal{P}')$ (preserving the colors of the punctures, i.e., $\phi(z_i) = z'_i$). The space of equivalence classes is called the *moduli space* \mathcal{M}_n .

If $n \leq 3$ then the moduli space \mathcal{M}_n is a single point. If $n \geq 4$, we can place the last three points to $(0, 1, \infty)$ by means of a Möbius transformation. With this normalization $(\mathbb{C}, \mathcal{P}) \sim (\mathbb{C}, \mathcal{P}')$ if and only if $\mathcal{P} = \mathcal{P}'$, and we see that

$$\mathcal{M}_n = \{ (z_1, \dots, z_{n-3}) : z_i \neq 0, 1; z_i \neq z_j \}.$$

This shows that \mathcal{M}_n an (n-3)-dimensional complex manifold.

Let us fix some reference normalized tuple $\mathcal{P}_{o} = (a_{1}, \ldots a_{n-3}, 0, 1, \infty)$. Then we can also define \mathcal{M}_{n} as the space of homeomorphisms $h : (\mathbb{C}, \mathcal{P}_{o}) \to \mathbb{C}$ normalized by h(0) = 0, h(1) = 1, up to equivalence: $h \sim h'$ if $h(\mathcal{P}_{o}) = h'(\mathcal{P}_{o})$.

Let us now refine this equivalence relation by declaring that $h \simeq h'$ if h is homotopic (or, equivalently, isotopic) to h' rel \mathcal{P}_{o} , and let [h] stand for the corresponding equivalence classes. It inherits the quotient topology from the space of homeomorphisms (endowed with the uniform topology). This quotient space is called the *Teichmüller space* \mathcal{T}_n . Since the equivalence relation \simeq is obviously stronger than \sim we have a natural projection $\pi : \mathcal{T}_n \to \mathcal{M}_n$.

18.2. Spiders. The homotopy class [h] can be visualised as the punctured sphere marked with a "spider". A spider S on the punctured sphere $(\mathbb{C}, \mathcal{P})$ is a family of disjoint paths σ_i connecting z_i to ∞ , $i = 1, \ldots n - 1$. We let [S] be the class of isotopic spiders (rel \mathcal{P}).

LEMMA 3.4. There is a natural one-to-one correspondence between points of \mathcal{T}_n and classes of isotopic spiders, $(\mathbb{C}, \mathcal{P}, [S])$.

PROOF. Let us fix a reference spider $(\mathbb{C}, \mathcal{P}_{o}, \mathcal{S}_{o})$. Then to each homeomorphism $h \in \mathcal{T}_{n}$ we can assossiate a spider $\mathcal{S} = h(\mathcal{S}_{o})$. Isotopy h_{t} rel \mathcal{P}_{o} induces isotopy of the corresponding spiders rel \mathcal{P} . Hence we obtain a map $[h] \mapsto [\mathcal{S}]$.

Vice versa, let us have a spider $(\mathbb{C}, \mathcal{P}, \mathcal{S})$. Then there exists a homeomorphism $h : (\mathbb{C}, \mathcal{P}_{o}, \mathcal{S}_{o}) \to (\mathbb{C}, \mathcal{P}, \mathcal{S})$. If $(\mathbb{C}, \mathcal{P}, \mathcal{S}')$ is an isotopic spider then the isotopy \mathcal{S}_{t} rel $\mathcal{P}, 0 \leq t \leq 1$, lifts to an isotopy h_{t} rel \mathcal{P}_{0} . Given any parametrizing homeomorphism $h' : \mathcal{S}_{o} \to \mathcal{S}'$, we can isotopy h_{1} so that it will coincide with h' on \mathcal{S}_{o} . Since two homeomorphisms of a topological disk coinciding on the boundary are isotopic rel the boundary, we are done.

18.3. Universal covering. The spiders can be labeled by tuples of n-1 elments of the fundamental group $\pi_1(\mathbb{C} \setminus \mathcal{P}) \approx \mathbb{F}_{n-1}$ (where the latter stands for the free group in n-1 generators). Indeed, let us consider a bouquet of circles $\wedge_{i=1}^{n-1}C_i$ in $\mathbb{C}_i \setminus \mathcal{P}$ based at some point $a \in \mathbb{C} \setminus \mathcal{P}$ and such that the circle C_i surrounds z_i but not the other points of \mathcal{P} . These circles oriented anti-clockwise represent generators of the fundamental group $\pi_1(\mathbb{C} \setminus \mathcal{P}, a)$. Accordingly, any loop in $\wedge C_i$ is homotopic to a concatenation of the loops C_i and their inverse. Let us select an arc γ_{∞} connecting a to ∞ in the complement of $\wedge C_i$, and n-1 arcs γ_i in the disks bounded by the C_i . Since $\wedge C_i$ is a homotopy retract for $\mathbb{C} \setminus \mathcal{P}$, any arc connecting z_i to ∞ is homotopic to the concatenation of the γ_i , a loop in $\wedge C_i$, and γ_{∞} . Thus, any spider leg is labeled by an element of $\pi_1(\mathbb{C} \setminus \mathcal{P}, a)$.

PROPOSITION 3.5. The natural projection $\pi : \mathcal{T}_n \to \mathcal{M}_n$ is the universal covering over \mathcal{M}_n .

PROOF. Let us first show that π is a covering. Take some base tuple $\mathcal{P}_{\circ} = (z_1^{\circ}, \ldots z_{n-1}^{\circ}) \in \mathcal{M}_n$, and consider a bouquet of circles C_i and the paths $\gamma_i^{\circ}, \gamma_{\infty}^{\circ}$ in $\mathbb{C} \smallsetminus \mathcal{P}$ as above. Consider a neighborhood $U_1 \times \cdots \times U_{n-3}$ of \mathcal{P}_{\circ} in \mathcal{M}_n , where the U_i are little round disks around z_i° fully surrounded by the circle C_i . Let us connect any point $z_i \in U_i$ to z_i° with a straight interval. Concatentating them with γ_i° , we obtain a path γ_i connecting z_i to a and continuously depending on $z_i \in U_i$. Select now any element $\tau \in \pi(\mathbb{C} \smallsetminus \mathcal{P}, a)$.

18.4. Infinitesimal theory. A tangent vector to the moduli space \mathcal{M}_n at point $z = (z_1 \dots, z_{n-3}, 0, 1, \infty)$ can be represented as a tuple

$$v = (v(z_1), \dots, v(z_{n-3}))$$

of tangent vectors to \mathbb{C} at points z_i . Since the natural projection $\mathcal{T}_n \to \mathcal{M}_n$ is a covering, tangent vectors to \mathcal{T}_n can be represented in the same way.

Any such tuple of vectors admits an extension to a smooth vector field v vanishing at points $(0, 1, \infty)$ (such vector field will be called "normalized"). So, we can

view the tangent space to \mathcal{M}_n (and \mathcal{T}_n) as the space Vect of smooth normalized vector fields modulo equivalence relation: $v \sim w$ if $v(z_i) = w(z_i)$, $i = i, \ldots, n-3$.

With this in mind, we can give a nice description of the cotangent space to \mathcal{M}_n and \mathcal{T}_n . Let us consider the space $\mathcal{Q} = \mathcal{Q}(\hat{\mathbb{C}} \smallsetminus \mathcal{P})$ of integrable quadratic differentials $\phi = \phi(z)dz^2$ on $\hat{\mathbb{C}} \smallsetminus \mathcal{P}$. Such differentials must have at most simple poles at the punctures (at ∞ it amounts to $\phi(z) = O(1/|z^3|)$).

EXERCISE 3.6. Show that this space \mathcal{Q} of quadratic differentials has complex dimension n-3. Moreover, the map $\phi \mapsto (\lambda_1, \ldots, \lambda_{n-3})$, where $\lambda_i = \operatorname{Res}_{z_i} \phi$, is an isomorphism between \mathcal{Q} and \mathbb{C}^{n-3} .

It turns out that it is not an accident that $\dim \mathcal{Q} = \dim \mathcal{M}_n$.

PROPOSITION 3.7. The space $\mathcal{Q}(\hat{\mathbb{C}} \setminus \mathcal{P})$ of quadratic differentials is naturally identified with the cotangent space to \mathcal{M}_n (and \mathcal{T}_n). The pairing between a cotangent vector $\phi \in \mathcal{Q}$ and a tangent vector $v \in \text{Vect}$ is given by the formula:

(18.1)
$$\langle \phi, v \rangle = \frac{1}{2\pi i} \int \int \phi \,\overline{\partial} v.$$

PROOF. Let us first note that this pairing is well defined. Indeed, as we saw in §1.8, $\bar{\partial}v$ can be interpreted as a Beltrami differential, and the product $\phi \bar{\partial}v$ as an area form. Moreover, this area form is integrable since ϕ is integrable and ∂v is bounded.

Let us now calculate this integral. Since ϕ is holomorphic, we have:

$$\phi \,\partial_{\bar{z}} v \,dz \wedge d\bar{z} = \partial_{\bar{z}}(\phi \,v) \,dz \wedge d\bar{z} = -\partial(\phi \,v \,dz) = -d(\phi \,v \,dz)$$

Let $\gamma_{\epsilon}(z_i)$ be the ϵ -circles centered at finite points of \mathcal{P} , $i = 1, \ldots, n-1$, and let Γ_{ϵ} be the ϵ^{-1} -circle centered at 0 (where all the circles are anti-clockwise oriented), and let D_{ϵ} be the domain of \mathbb{C} bounded by these circles. Then by the Stokes formula

$$-\frac{1}{2\pi i} \int \int_{D_{\epsilon}} d(\phi \, v dz) = \frac{1}{2\pi i} \sum \int_{\gamma_{\epsilon}(z_i)} \phi \, v dz - \frac{1}{2\pi i} \int_{\Gamma_{\epsilon}} \phi \, v dz$$

But near any $z_i \in \mathbb{C}$ we have:

$$\phi v = \frac{\lambda_i v(z_i)}{z - z_i} + O(1),$$

where $\lambda_i = \operatorname{Res}_{z_i} \phi$. Hence

$$\frac{1}{2\pi i}\int_{\gamma_\epsilon(z_i)}\phi\,v\,dz\to\lambda_iv(z_i)\,\,\mathrm{as}\,\,\epsilon\to 0.$$

Note that these integrals asymptotically vanish at $z_{n-2} = 0$ and $z_{n-1} = 1$ since v vanishes at these points. The integral over Γ_{ϵ} asymptotically vanishes as well since $\phi(z) = O(|z|^{-3})$ while $v(z) = o(|z|^2)$ near ∞ (as the vector field v/dz vanishes at ∞).

Finally, we obtain:

$$\frac{1}{2\pi i} \int \int \phi \,\partial_{\bar{z}} v \,dz \wedge d\bar{z} = \sum_{i=1}^{n-3} \lambda_i v(z_i)$$

So, the pairing (18.1) depends only on the values of v at the points z_1, \ldots, z_{n-3} , and hence defines a functional on tangent space $T\mathcal{M}_n$. This gives an isomorphism between \mathcal{Q} and the cotangent space $T^*\mathcal{M}_n$ since $(\lambda_1, \ldots, \lambda_{n-3})$ are global coordinates on the both spaces (see Exercise 3.6).

3. ELEMENTS OF TEICHMÜLLER THEORY

18.5. General Teichmüller spaces.

18.5.1. Marked Riemann surfaces. The previous discussion admits an extension to an arbitrary qc class \mathcal{QC} of Riemann surfaces that we will outline in this section. Take some base Riemann surface $S_0 \in \mathcal{QC}$ (without boundary), and let \bar{S}_0 be the ideal boundary compactification of S_0 . Given another Riemann surface $S \in \mathcal{QC}$ (with compactification \bar{S}), a marking of S is a choice of a qc homeomorphism $\phi: \bar{S}_0 \to \bar{S}$ (parametrization by S_0) up to the following equivalence relation. Two parametrized surfaces (S, ϕ) and (S', ϕ') are equivalent if there is a conformal isomorphism $h: S \to S'$ that makes the following diagram homotopically commutative rel the ideal boundary (i.e., there is a qc homeomorphism $\tilde{\phi}: S_0 \to S$ homotopic to ϕ rel $\partial \bar{S}_0$ such that $h \circ \tilde{\phi} = \phi'$). A marked Riemann surfaces is an equivalence class $\tau = [S, \phi]$ of this relation. The space of all marked Riemann surfaces is called the Teichmüller space $\mathcal{T}(S_0)$.

REMARK 3.1. Fixing a set Δ_0 of generators of $\pi_1(S_0)$ and parametrizations of the boundary components of $\partial \bar{S}_0$ by the standard circle \mathbb{T} , we naturally endow any marked Riemann surface $[S, \phi]$ with a set of generators of $\pi_1(S)$ (up to an inner automolphism of $\pi_1(S)$) and with a parametrization of the components ∂S by \mathbb{T} . Thus, we obtain a marked surface in the sense of §1.1.5.

18.5.2. Representation variety. Let us now uniformize the base Riemann surface S_0 by a Fuchsian group Γ_0 . The (Fuchsian) representation variety $\operatorname{Rep}(\Gamma_0)$ is the space of faithful⁴ Fuchsian representations $i : \Gamma_0 \to \operatorname{PSL}(2, \mathbb{R})$ up to conjugacy in $\operatorname{PSL}(2, \mathbb{R})$ endowed with the algebraic topology. In this topology $i_n \to i$ if after a possible replacement of the i_n with conjugate representations, we have: $i_n(\gamma) \to i(\gamma)$ for any $\gamma \in \Gamma_0$.

LEMMA 3.8. There is a natural embedding $e: \mathcal{T}(S_0) \to \operatorname{Rep}(S_0)$.

PROOF. Let $\phi: S_0 \to S$ be a qc parametrization of some Riemann surface $S \in QC$, and let Γ be a Fuchsian group uniformizing S. Then ϕ lifts to an equivariant qc homeomorphism $\Phi: (\mathbb{H}, \Gamma_0) \to (\mathbb{H}, \Gamma)$, so there is an isomorphism $i: \Gamma_0 \to \Gamma$ such that $\Phi \circ \gamma_0 = \gamma \circ \Phi$ for any $\gamma_0 \in \Gamma_0$ and $\gamma = i(\gamma_0)$.

If we replace Φ with another lift $T \circ \Phi$, where $T \in \Gamma$, then *i* will be replaced with a conjugate representation $\gamma_0 \mapsto T^{-1} \circ i(\gamma_0) \circ T$.

If we replace ϕ with a homotopic parametrization $\tilde{\phi} : S_0 \to S$ then the induced representation $\Gamma_0 \to \Gamma$ will not change. Indeed, a homotopy ϕ_t connecting ϕ to $\tilde{\phi}$ lifts to an equivariant homotopy $\Phi_t : (\mathbb{H}, \Gamma_0) \to (\mathbb{H}, \Gamma)$ inducing a path of representations $i_t : \Gamma_0 \to \Gamma$. Then for any $\gamma_0 \in \Gamma_0$, the image $i_t(\gamma_0) \in \Gamma$ moves continuously with t. Since Γ is discrete, $i_t(\gamma)$ cannot move at all.

If we further replace ϕ with $h \circ \phi$, where $h : S \to S'$ is a conformal isomorphism then the representation $i : \Gamma_0 \to \Gamma$ will be replaced with a conjugate by $T : \mathbb{H} \to \mathbb{H}$ where $T \in \text{PSL}(2, \mathbb{R})$ is a lift of h.

Thus, we obtain a well defined map $e : \mathcal{T}(S_0) \to \operatorname{Rep}(S_0)$ that associates to a marked surface $[S, \phi]$ the induced representation $i : \Gamma_0 \to \Gamma$ up to conjugacy in PSL(2, \mathbb{R}).

Let us now show that e is injective. Let $\phi : S_0 \to S$ and $\phi' : S_0 \to S'$ be two parametrizations whose lifts Φ and Φ' to \mathbb{H} induce two representations i and

⁴i.e., injective

i' of Γ_0 that are conjugate by $T \in PSL(2, \mathbb{R})$. Then Φ and $\Psi = T^{-1} \circ \Phi$ are two equivariant homeomorphisms $(\mathbb{H}, \Gamma_0) \to (\mathbb{H}, \Gamma)$ that induce the same representation $i: \Gamma_0 \to \Gamma$. We need to show that they are equivariantly homotopic.

To this end let us consider the following diagram encoding equivariance of Φ and $\Psi:$

Let $\delta(x)$ be the hyperbolic geodesic connecting $\Phi(x)$ to $\Psi(x)$. Since γ is a hyperbolic isometry, it isometrically maps $\delta(x)$ to $\delta(\gamma_0 x)$. Let $t \mapsto \Phi_t(x)$ be a uniform motion along $\delta(x)$ from $\Phi(x)$ to $\Psi(x)$ with such a speed that at time t = 1we reach the destination (in other words, $\Phi_t(x)$ is the point on $\delta(x)$ on hyperbolic distance $t \operatorname{dist}_{hyp}(\Phi(x), \Psi(x))$ from $\Phi(x)$). Then $\gamma(\Phi_t x) = \Phi_t(\gamma_0 x)$, and we obtain a desired equivariant homotopy.

18.5.3. Teichmüller metric. Let us endow the space $\mathcal{T}(S_0)$ with the following Teichmüller metric. Given two marked surfaces $\tau = [S, \phi]$ and $\tau' = [S', \phi']$, we let $\operatorname{dist}_{\mathrm{T}}(\tau, \tau')$ be the infimum of dilatations of qc maps $h: S \to S'$ that make diagram (??) homotopically commutative.

LEMMA 3.9. dist_T is a metric.

PROOF. Triangle ineaquality for dist_T follows from submultiplicativity of the dilatation under composition. So, dist_T is a pseudo-metric. Let us show that it is a metric, Indeed, if dist_T(τ, τ') = 0 then there exists a sequence $h_n : S \to S'$ of qc maps in the right homotopy class with $\text{Dil}(h_n) \to 0$. Let $H_n : \mathbb{H} \to \mathbb{H}$ be the lifts of the h_n that induce the same isomorphism between Γ and Γ' . Then the H_n is a sequence of qc maps with uniformly bounded dilatation whose extensions to $\mathbb{R} = \partial \mathbb{H}$ all coincide. Now Compactness Theorem 2.31 implies that the H_n uniformly converge to an equivariant conformal isomorphism $T : (\mathbb{H}, \Gamma_0) \to (\mathbb{H}, \Gamma)$. It descends to a conformal isomorphism $h: S \to S'$ in the samehomotopy class as the h_n .

EXERCISE 3.10. Show that the embedding $e : \mathcal{T}(S_0) \to \operatorname{Rep}(\Gamma_0)$ is continuous. (from the Teichmúller metric to the algebraic topology).

19. Bers Embedding

19.1. Schwarzian derivative and projective structures.

19.1.1. Definition. The fastest way to define the Schwarzian derivative Sf is by means of a mysterious formula:

(19.1)
$$Sf = \frac{f'''}{f'} - \frac{3}{2} \left(\frac{f''}{f'}\right)^2.$$

However, there is a bit longer but much better motivated approach.

Let us try to measure how a function f at a non-critical point z deviates from a Möbius transformation. Möbius transformations depend on three complex parameters. So, one expects to find a unique Möbius transformation A_z that coincides with f to the second order. Then

$$f(\zeta) - A_z(\zeta) \sim \frac{b}{6} (\zeta - z)^3$$

near z, and we let Sf(z) = b/f'(z).

REMARK 3.2. Division by f'(z) ensures scaling invariance of the Schwarzian derivative: $S(\lambda f) = Sf$. Coefficient 1/6 provides a convenient normalization suggested by the Taylor formula: it makes Sf = f''' for a normalized map $f(\zeta) = \zeta + O(|\zeta - z|^3)$.

The best Möbius approximation to f is easy to write down explicitly. Let $f(\zeta) = a_0 + a_1(\zeta - z) + a_2(\zeta - z)^2 + \dots$ near z with $a_1 = f'(z) \neq 0$. Then

$$A_z(\zeta) = a_0 + \frac{a_1(\zeta - z)}{1 - \beta(\zeta - z)} \quad \text{with} \quad \beta = \frac{a_2}{a_1},$$

the 3d Taylor coefficient of $f - A_z$ is $(a_3 - a_2^2/a_1)$, and (19.1) follows.

19.1.2. Chain rule.

LEMMA 3.11. Let f be a holomorphic function on a domain U. Then $Sf \equiv 0$ on U if and only if f is a restrictin of a Möbius map to U.

PROOF. Sufficiency is obvious from the definition: If f is a Möbius map then $A_z = f$ at any point z, and Sf(z) = 0.

Vice versa, assume $Sf \equiv 0$ on U. Then f is a solution of a 3d order analytic ODE

$$f''' = \frac{3}{2} \frac{(f'')^2}{f'}$$

on $U \smallsetminus C_f$, where C_f is the critical set of f. Such a solution is uniquely determined by its 2-jet⁵ at any point $z \in U \smallsetminus C_f$. Hence $f = A_z$.

Similarly, one can prove:

EXERCISE 3.12. Let f and g be two holomorphic functions on a domain U. Then $Sf \equiv Sg$ on U if and only if $f = A \circ g$ for some Möbius map A.

LEMMA 3.13 (Chain Rule).

(19.2)
$$S(f \circ g) = (Sf \circ g) \cdot (g')^2 + Sg.$$

PROOF. Since the Schwarzian derivative is translationally invariant on both sides (i.e., $S(T_1 \circ f \circ T_2) = Sf$ for any translations T_1 and T_2), it is sufficient to check (19.2) at the origin and to assume that g(0) = f(0) = 0. Furthermore, by Exercise 3.12, postcomposition of f with a Möbius transformation would not change either side of (19.2). In this way, we can bring f to a normalized form:

(19.3)
$$f(\zeta) = \zeta + \frac{Sf(0)}{6}\zeta^3 + \dots$$

and then painlessly check (19.2) by composing (19.3) with the 3-get of g.

In particular, for a Möbius transformation A, we have:

(19.4)
$$S(f \circ A) = (Sf \circ A) \cdot (A')^2,$$

which coincides with the transformation rule for quadratic differentials. It suggests that the Schwarzian should be viewed not as a function but rather as a quadratic differential $Sf(z)dz^2$. This point of view is not quite right on Riemann surfaces, but it becomes exactly correct on projective surfaces.

⁵Recall that a *n*-jet of a function f at z is its Taylor approximant of order n at z.

19. BERS EMBEDDING

19.1.3. Projective surfaces. A projective structure on a Riemann surface S is an atlas of holomorphic local charts with Möbius transit maps. A surface endowed with a projective structure is called a *projective surface*. Projective morphisms are defined naturally, so that we can refer to isomorphic projective surfaces.

Of course, the Riemann sphere $\overline{\mathbb{C}}$ has a natural projective structure, and any domain $U \subset \overline{\mathbb{C}}$ inherits it. If we have a group Γ of Möbius transformations acting properly discontinuously and freely on U then the quotient Riemann surface $V = U/\Gamma$ inherits a unique projective structure that makes the quotient map $\pi : U \to V$ projective. In particular, any Riemann surface S is endowed with the *Fuchsian* projective structure coming from the uniformization $\pi : \mathbb{H} \to V$.

Given a meromorphic function f on a projective surface V, the Chain Rule (19.4) tells us that the local expressions $Sf(z)dz^2$ determine a global quadratic differential on V.

EXERCISE 3.14. Check carefully this assertion.

More generally, let us consider two projective structures f and g on a Riemann surface V given by atlases $\{f_{\alpha}\}$ and $\{g_{\beta}\}$ respectively. Then the Chain Rule (more specifically, Exercises 3.12 and 3.14) tell us that the local expressions $S(f_{\alpha} \circ g_{\beta}^{-1})(z) dz^2$ determine a global quadratic differential on V endowed with the g-structure. This differential is denoted $S\{f,g\}$. It measures the distance between f and g.

In particular, given a holomorphic map $f: V \to W$ between two projective surfaces, we obtain a quadratic differential $S\{f^*(W), V\}$ on⁶ V. Writing f in projective local coordinates $(\zeta = f(z))$, we obtain the familiar expression, $Sf(z) dz^2$, for this differential.

 $^{^{6}}$ Here we notationally identify surfaces with their projective structures

Part 2

Complex quadratic family

CHAPTER 4

Dynamical plane

20. Glossary of Dynamics

This glossary collects some basic notions of dynamics. Its purpose is to fix terminology and notation and to comfort a reader who has a little experience in dynamics.

20.1. Orbits and invariant sets. Consider a continuous endomorphism $f: X \to X$ of a topological space X. The *n*-fold iterate of f is denoted by $f^n, n \in \mathbb{N}$. A topological dynamical system (with discrete positive time) is the \mathbb{N} -action generated by $f, n \mapsto f^n$. The orbit or trajectory of a point $x \in X$ is $\operatorname{orb}(x) = \{f^n x\}_{n \in \mathbb{N}}$. The subject of topological dynamics is to study qualitative behavior of orbits of a topological dynamical system.

Here is the simplest possible behavior: a point α is called *fixed* if $f\alpha = \alpha$. More generally, a point α is called *periodic* if it has a finite orbit, i.e., there exists a $p \in \mathbb{Z}_+$ such that $f^p \alpha = \alpha$. The smallest p with this property is called the *period* of α . The orbit of α (consisting of p permuted points) is naturally called a *periodic orbit* or a *cycle* (of period p). We will write periodic orbits in bold: $\alpha = \operatorname{orb} \alpha$, $\beta = \operatorname{orb} \beta$, etc. The sets of fixed and periodic points are denoted $\operatorname{Fix}(f)$ and $\operatorname{Per}(f)$ respectively.

A subset $Z \subset X$ is called *(forward) invariant* under f if $f(Z) \subset Z$ (or equivalently, $f^{-1}(Z) \supset Z$). It is called *backward invariant* if $f^{-1}(Z) \subset Z$. If Z is simultaneously forward and backward invariant (so that $f^{-1}(Z) = Z$), it is called *completely invariant*.

A set Z is called wandering if $f^n Z \cap f^m Z = \emptyset$ for any $n > m \ge 0$. It is called weakly wandering¹ if $f^{-n}(Z) \cap Z = \emptyset$ for any n > 0.

EXERCISE 4.1. Show that wandering sets are weakly wandering but not the other way around (in general). Show that Z is weakly wandering if and only if either of the following properties is satisfies:

- $f^{-n}(Z) \cap f^{-m}(Z) = \emptyset$ for any $n > m \ge 0$.
- No point $z \in Z$ returns back to Z under iterates of f.

The asymptotical behavior of an orbit can be studied in terms of its limit set. The ω -limit set $\omega(x)$ of a point x is the set of all accumulation points of $\operatorname{orb}(x)$. It is a closed forward invariant subset of X, so its complement is an open backward invariant subset. If X is compact then $\omega(x)$ is a non-empty compact subset of X. We say that the orbit of x converges to a cycle (of a periodic point α) if $\omega(x) = \operatorname{orb}(\alpha)$.

A point x is called *recurrent* if $\omega(x) \ni x$. Existence of non-periodic recurrent points is a feature of non-trivial dynamics.

¹Actually, in literature these sets are usually referred to as just "wandering".

The set

$$\operatorname{Orb}(z) = \bigcup_{n \ge 0} f^{-n}(\operatorname{orb}(z))$$

is called the *grand orbit* of a point z. It is an equivalence class of the following equivalence relation:

$$z \sim \zeta$$
 if $f^n z = f^m \zeta$ for some $m, n \in \mathbb{N}$.

Note that the usual forward orbits orb z are not classes of any equivalence relation. In fact, the grand orbits relation is the minimal one generated by the forward orbits.

There is a smaller equivalence relation

$$z \sim \zeta$$
 if $f^n z = f^n \zeta$ for some $n \in \mathbb{N}$.

These equivalence classes are called *small orbits* of f.

Given a connected set U and a point z such that $f^n z \in \text{int } U$ for some $n \ge 0$, let V be the connected component of $f^{-n}U$ containing z. It is called i the *pullback* of U (along the n-rbit of z).

20.2. Equivariant maps. Two dynamical systems $f: X \to X$ and $g: Y \to Y$ are called *topologically conjugate* (or topologically equivalent) if there exists a homeomorphism $h: X \to Y$ such that $h \circ f = g \circ h$, i.e., the following commutative diagram is valid:

$$\begin{array}{cccc} X & \xrightarrow{f} & X \\ h \downarrow & & \downarrow h \\ Y & \xrightarrow{g} & Y \end{array}$$

Classes of topologically equivalent dynamical systems (within an *a priori* specified family) are called *topological classes*. If X and Y are endowed with an extra structure (e.g., smooth, conformal, quasi-conformal etc.) respected by h, then f and g are called *smoothly/conformally/quasi-conformally conjugate* (or equivalent). The corresponding equivalence classes are called smooth/conformal/quasi-conformal classes.

Topological conjugacies respect all properties which can be formulated in terms of topological dynamics: orbits go to orbits, cycles go to cycles of the same period, ω -limit sets go to ω -limit sets, converging orbits go to converging orbits etc.

A homeomorphism $h: X \to X$ commuting with a dynamical system $f: X \to X$ (i.e., conjugating f to itself) is called an *automorphism* of f.

A continuous map which makes the above diagram commutative is called *equivariant* (with respect to the actions of f and g). A surgective equivariant map is called a *semi-conjugacy* between f and g. In this case g is also called a *quotient* of f.

It will be convenient to extend the above terminology to partially defined maps. Let f and g be partially defined maps on the spaces X and Y respectively (i.e., f maps its domain $\text{Dom}(f) \subset X$ to X, and similarly does g). Let $A \subset X$. A map $h: A \to Y$ is called equivariant (with respect to the actions of f and g) if for any $x \in A \cap \text{Dom}(f)$ such that $fx \in A$ we have: $hx \in \text{Dom}(g)$ and h(fx) = g(hx). (Briefly speaking, the equivariance equation is satisfied whenever it makes sense.)

20.3. Elements of ergodic theory.

20.4. Invertible one-dimensional maps.

20.4.1. Invertibel interval maps. are the simplest dynamical examples:

EXERCISE 4.2. Let $f : I \to I$ be a continuous monotone map of an interval. Its set of fixed points, Fix(f), is a non-empty closed set.

(i) If f is increasing than any orbit converges to a fixed point.

(ii) If f is decreasing than Fix(f) is a sigleton, $Fix(f) = \{\alpha\}$, and any orbit either converges to α or it converges to a cycle of period 2.

Here is *zigzag* pictures illustrating the above types of behavior:

20.4.2. Circle homeomorphisms.

20.5. Bernoulli map.

20.6. Doubling map. The doubling map is just the squaring map $D: z \mapsto z^2$ on the unit circle \mathbb{T} . Passing to the annular coordinate $\theta \in \mathbb{R}/\mathbb{Z}$, where $z = e^{2\pi i \theta}$, we obtain the map $\theta \mapsto 2\theta \mod \mathbb{Z}$, which justifies the term "doubling". We can also view it as a map $D: \theta \mapsto 2\theta \mod 1$ on the unit interval [0, 1], i.e.,

$$D(\theta) = 2\theta$$
, for $\theta \in [0, 1/2]$ and $D(\theta) = 2\theta - 1$, for $\theta \in [1/2, 1]$,

with understuding that the endpoints must be undentified. We will use the same notation D for all these models.

The doubling map has a unique fixed point z = 1, i.e., $\theta = 0$. The preimages of this point under D^n are diadic rationals $\theta = p/2^n$, $p = 0, 1, \ldots, 2^n - 1$. They divide the circle into (closed) *diadic intervals*

$$J_{i_0\dots i_{n-1}} = \left\{ \theta = \frac{i_0}{2} + \frac{i_1}{4} + \dots + \frac{i_{n-1}}{2^n} + \frac{\theta_n}{2^{n+1}}, \quad \text{where } \theta_n \in [0,1] \right\},\$$

consisting of angles whose diadic expansion begins with $[i_0, \ldots, i_{n-1}]$ (and may end with the infinite number of "1"'s, to make the interval closed). Note that

 $J_{i_0\dots i_{n-1}} = J_{i_0\dots i_{n-1},0} \cup J_{i_0\dots i_{n-1},1}, \text{ and } D(J_{i_0\dots i_{n-1}}) = J_{i_1\dots i_{n-1}}.$

It follows that $D^n(J_{\overline{i}}^n) = \mathbb{T}$ for any diadic interval, and this map is one-to-one, except that it glues the endpoints of $J_{\overline{i}}^n$ to z = 1.

EXERCISE 4.3. The map $\phi: \Sigma_2^+ \to [0,1]$ that associates to a diadic sequence $\overline{i} = [i_0, i_1, \ldots,]$ the angle θ with this diadic expansion, is a semiconjugacy between the Bernoulli shift $\sigma: \Sigma_2^+ \to \Sigma_2^+$ and the doubling map.

EXERCISE 4.4. Show that $J_{i_0...i_{n-1}}^n = \{\theta : D^k \theta \in J_{i_k}^1, k = 0, ..., n-1\}.$

EXERCISE 4.5. Periodic points of the doubling map are rationals $\theta = p/q$ with odd denominator. Pre-periodic points are rationals $\theta = p/q$ with even denominator (in the irreducible representation).

PROPOSITION 4.6. There are no non-trivial orientation preserving hoemomorphisms $h: \mathbb{T} \to \mathbb{T}$ commuting with the doubling map.

PROOF. Since $\theta = 0$ is the unique fixed point of D, it must be also fixed by h. Since $\theta = 1/2$ is the only D-preimage of 0 different from 0, it must be fixed by h as well. Hence the diadic intervals $J_0^1 = [0, 1/2]$ and $J_1^1 = [1/2, 1]$ are either h-invariant or are permuted by h (fixing the endpoints). But in the latter case, h would be orientation reversing, so both intervals are invariant.

Each of them contains one D-preimage of 1/2, respectively $\theta = 1/4$ and $\theta = 3/4$, so these points must also be fixed by h. Hence all the diadic intervals $J_{i_0i_1}^2$ of level 2 are h-invariant (with the endpoints fixed).

Assume inductively that all the diadic intevals $J_{i_0...i_{n-1}}^n$ of level n are h-invariant (with the endpoints fixed). Since each of them contains one diadic point of next level, $p/2^{n+1} \in D^{-n+1}(0)$ with odd p, all these points must be fixed, and hence all the diadic intervals of level n + 1 are h-invariant.

We concude by induction that all the diadic points $p/2^n$, $n \in \mathbb{N}$, are fixed by h. By continuity, h = id.

The doubling map, and its quotients, will serve as the main dynamical model for quadratic polynomials on their Julia sets.

20.7. Markov chains.

20.8. Holomorphic equivalence relations. Holomorphic equivalence relations provide an adequate general set-up for various situations we will face. However, we will not make a serious use of it exploiting it only as a convenient language. And even in this capacity, it will not be used until §26.2.

Let S be a Riemann surface, and let \mathcal{R} be an equivalence relation on S with countable classes. We say that \mathcal{R} is *holomorphic* if there exists a countable family Φ of holomorphic functions $\phi_n(z,\zeta)$ in two variables such that two points $z, \zeta \in S$ are \mathcal{R} -equivalent if and only if $\phi_n(z,\zeta) = 0$ for some $\phi_n \in \Phi$. The functions ϕ_n are called the *local sections* of \mathcal{R} . We will assume that local sections ϕ_n are *primitive* in the sense that they are not powers of other holomorphic functions, $\phi_n \neq \psi^k$ for $k \geq 2$.

For instance, orbits of a discrete subgroup $\Gamma \subset \text{Aut } S$ (e.g., consider a Fuchsian group acting on \mathbb{D}) form a holomorphic equivalence relation. For the context of this book, the most important type of a holomorphic equivalence relation is the grand orbit relation generated by a holomorphic map f, e.g., by a rational endomorphism of the Riemann sphere $\hat{\mathbb{C}}$ (but we will also deal with partially defined maps).

REMARK 4.1. More generally, one can consider relations generated by holomorphic pseudo-groups or pseudo-semigroups. One can also consider algebraic equivalence relations.

A critical point of \mathcal{R} is a point z_0 such that $\partial_z \phi_n(z_0) = 0$ for some local section $\phi_n \in \Phi$. A critical equivalence class is a class containing a critical point. Since the local sections are primitive, the critical points of any section are isolated, and hence altogether there are at most countably many critical points. Non-critical points are called *regular*.

Any equivalence relation on S can be restricted to a subset $D \subset S$. An open subset D is called the *fundamental domain* for \mathcal{R} if the restriction \mathcal{R} to D is trivial (in other words, D contains at most one point of any equivalence class) while its restriction to the closure \overline{D} is complete (i.e., any equivalence class crosses \overline{D}). Under these circumstances, the closure \overline{D} will also be referred as a "(closed) fundamental domain" for \mathcal{R} as long as meas $\partial D = 0$.

The \mathcal{R} -saturation D of a set $D \subset S$ is the union of all equivalence classes that cross D.

Terms "fundamental domain for f", "f-saturation of D" etc. mean the corresponding objects for the grand orbit equivalence relation generated by f. For instance, f-saturation of a set X is its grand orbit $\bigcup_{n \in \mathbb{N}} f^{-n}(\operatorname{orb} X)$.

EXERCISE 4.7. A domain $D \subset S$ is a fundamental domain for a map f restricted to the f-saturation of D if and only if D is wandering and the iterates $f^n | D$ are injective, $n \in \mathbb{N}$.

EXERCISE 4.8. Show that the saturation of an open subset by a holomorphic equivalence relation is open.

21. Holomorphic dynamics: basic objects

Below

$$f \equiv f_c : z \mapsto z^2 + c$$

unless otherwise is stated. Dynamical objects will be labeled by either f or c whatever is more convenient in a particular situation (for instance, $D_f(\infty) \equiv D_c(\infty)$), $J(f_c) \equiv J_c$ by default). Moreover, the label can be skipped altogether if f is not varied.

21.1. Critical points and values. First note that f^n is a branched covering of \mathbb{C} over itself of degree 2^n . Its critical points and values have a good dynamical meaning:

EXERCISE 4.9. The set of finite critical points of f^n is $\bigcup_{k=0}^{n-1} f^{-k}(0)$. The set of critical values of f^n is $\{f^k 0\}_{k=1}^n$.

Note that there are much fewer critical values than critical points! We let

(21.1)
$$\operatorname{Crit}(f) = \bigcup_{k=0}^{n-1} f^{-k}(0)$$

be the set of all critical points of iterated f.

Thus, f^n is an unbranced covering over the complement of $\{f^k 0\}_{k=1}^n$.

COROLLARY 4.10. Let V be a topological disk which does not contain points $f^k_{0,k}$ $k = 1, 2, \ldots, n$. Then the inverse function f^{-n} has 2^n single-values branches f_i^{-n} that univalently map V onto pairwise disjoint topological disks U_i , $i = 1, 2, ..., 2^n$.

These simple remarks explain why the forward orbit of 0 plays a very special role. We will have many occasions to see that this single orbit is responsible for the complexity and variety of the global dynamics of f.

However, f has one more critical point overlooked so far:

21.2. Looking from infinity. Extend *f* to an endomorphism of the Riemann sphere $\hat{\mathbb{C}}$. This extension has a critical point at ∞ fixed under f. We will start exploring the dynamics of f from there. The first observation is that $\hat{\mathbb{C}} \smallsetminus \overline{\mathbb{D}}_R$ is finvariant for a sufficiently big R, and moreover $f^n z \to \infty$ as $n \to \infty$ for $z \in \hat{\mathbb{C}} \setminus \overline{\mathbb{D}}_R$.

This can be expressed by saying that $\mathbb{C} \setminus \mathbb{D}_R$ belongs to the *basin of infinity* defined as the set of all *escaping* points:

$$D_f(\infty) = \{ z \in \hat{\mathbb{C}} : f^n z \to \infty \text{ as } n \to \infty \} = \bigcup_{n=0}^{\infty} f^{-n} (\hat{\mathbb{C}} \setminus \bar{\mathbb{D}}_R).$$

PROPOSITION 4.11. The basin of infinity $D_f(\infty)$ is a completely invariant domain containing ∞ .

PROOF. The only non-obvious statement to check is connectivity of $D_f(\infty)$. To this end let us show inductively that the sets $U_n = f^{-n}(\hat{\mathbb{C}} \setminus \bar{\mathbb{D}}_R)$ are connected. Indeed, assume that U_n is connected while U_{n+1} is not. Consider a bounded component V of U_{n+1} . Then the restriction $f: V \to U_n$ is proper and hence surjective (Corollary 1.47). In particular f would have a pole in V – contradiction.

21.3. Basic Dichotomy for Julia sets. We can now introduce the fundamental dynamical object, the *filled Julia set*

$$K(f) = \hat{\mathbb{C}} \smallsetminus D_f(\infty).$$

Proposition 4.11 implies that K(f) is a completely invariant compact subset of \mathbb{C} . Moreover, it is *full*, i.e., it does not separate the plane (since $D_f(\infty)$ is connected).

EXERCISE 4.12. (i) The filled Julia set consists of more than one point. (ii) Each component of int K(f) is simply connected.

The filled Julia set and the basin of infinity have a common boundary, which is called the Julia set, $J(f) = \partial K(f) = \partial D_f(\infty)$. Figures in this section show several representative pictures of the Julia sets $J(f_c)$ for different parameter values c.

*** Figures inserted***

Generally, topology and geometry of the Julia set is very complicated, and it is hard to put a hold on it. However, the following rough classification will give us some guiding principle:

THEOREM 4.13 (Basic Dichotomy). The Julia set (and the filled Julia set) is either connected or Cantor. The latter happens if and only if the critical point escapes to infinity: $f^n(0) \to \infty$ as $n \to \infty$.

PROOF. As in the proof of Proposition 4.11, let us consider the increasing sequence of domains $U_n = f^{-n}(\hat{\mathbb{C}} \setminus \bar{\mathbb{D}}_R)$ exhausting the basin of infinity. Assume first that the critical point does not escape to ∞ . Then $f: U_{n+1} \to U_n$ is a branched double covering with the only branched point at ∞ . By the Riemann-Hurwitz formula, if U_n is simply connected then U_{n+1} is simply connected as well. We conclude inductively that all the domains U_n are simply connected. Hence their union, $D_f(\infty)$, is simply connected as well, and its complement, K(f), is connected. But the boundary of a full connected compact set is connected. Hence J(f) is connected, too.

Assume now that the critical point escapes to infinity. Then 0 belongs to some domain U_n . Take the smallest n with this property. Adjust the radius R in such a way that the orbit of 0 does not pass through $\mathbb{T}_R = \partial U_0$. Then $0 \notin \partial U_{n-1}$, and hence ∂U_{n-1} is a Jordan curve. Let us consider the complementary Jordan disk $D \equiv D^0 = \mathbb{C} \setminus \overline{U}_{n-1}$. Since $f(0) \in U_{n-1}$, f is unbranched over D. Hence $f^{-1}D = D_0^1 \cup D_1^1$, where the $D_i^1 \Subset D$ are disjoint topological disks conformally mapped onto D.



FIGURE 1. The filled Julia set of $z^2 + \epsilon$ for small ϵ is a quasidisk. It has an attracting fixed point.

Take now the *f*-preimages of $D_0^1 \cup D_1^1$ in D_0^1 . We obtain two Jordan disks D_{00}^2 and D_{01}^2 with disjoint closures conformally mapped by *f* onto D_0^1 and D_1^1 repsectively. Similar disks, D_{10}^2 and D_{11}^2 , we find in D_1^1 (see Figure).

picture

Iterating this procedure, we will find that $f^{-n}D$ is the union of 2^n Jordan disks $D^n_{i_0i_1...i_{n-1}}$ such that $D^n_{i_0...i_{n-1}}$ is compactly contained in $D^{n-1}_{i_0...i_{n-2}}$ and is conformally mapped by f onto $D^{n-1}_{i_1...i_{n-1}}$.

Since $D_0^1 \cup D_1^1$ is compactly contained in D, the branches of the inverse map, $f^{-1}: D_j^1 \to D_{ij}^2$, are uniformly contracting in the hyperbolic metric of D (by the Schwarz Lemma). Since the domains $D_{i_0i_1...i_{n-1}}^n$ are obtained by iterating these branches, they uniformly exponentially shrink as $n \to \infty$. Hence the filled Julia set $K(f) = \bigcap f^{-n}D$ is a Cantor set. Of course, the Julia set J(f) coincides with K(f)in this case.

REMARK 4.2. The above proof also shows that in the case when J(f) is a Cantor set, the map f is expanding on it: $||Df(z)|| \ge \lambda > 1$ for any $z \in J(f)$, with respect to the hyperbolic metric on the disk $D \subset J(f)$.

The Basic Dichotomy gives us the first illustration of how the behavior of the critical point influences the global dynamics.

Since neither Cantor nor connected sets (containing more than one point) can have isolated points, we conclude:

COROLLARY 4.14. Both K(f) and J(f) are perfect sets.



FIGURE 2. Cauliflower: the filled Julia set of $z^2 + 1/4$. It has a parabolic fixed point.



FIGURE 3. Basilica: the filled Julia set of $z^2 - 1$. It has a superattracting cycle of period 2.



FIGURE 4. Douady rabbit: his head and two ears form the immediate basin for a superattracting cycle of period 3.

21.4. Bernoulli shift. When the Julia set is Cantor, there is an explicit symbolic model for the dynamics of f on it. Consider the space $\Sigma \equiv \Sigma_2^+$ of one-sided sequences $(i_0i_1...)$ of zeros and ones. Supply it with the weak topology (convergence in this topology means that all coordinates eventually stabilize). We obtain a Cantor set. Define the shift β on this space as the map of forgetting the first coordinate,

$$\beta:(i_0i_1\ldots)\mapsto(i_1i_2\ldots).$$

It is called the *(one-sided)* Bernoulli shift with two states.

EXERCISE 4.15. Show that:

- For any open set $U \subset \Sigma$, there exists an $n \in \mathbb{N}$ such that $\beta^n(U) = \Sigma$;
- β is topologically transitive;
- Periodic points of β are dense in Σ .

EXERCISE 4.16. Show that the only non-trivial automorphism of the one-sided Bernoulli shift with two states is induced by the relabeling $0 \leftrightarrow 1$.

If some endomorphism $f: X \to X$ of a compact space is topologically conjugate to a one-sided Bernoulli shift with two states, then X can be partitioned into two pieces X_0 and X_1 corresponding to sequences which begin with 0 and 1 respectively. This partition is called a *Bernoulli generator* for f. The statement of Exercise 4.16 is equivalent to saying that a Bernoulli generator is unique. For a Cantor Julia



FIGURE 5. Dendrite: the Julia set of $z^2 + i$. Here the critical point is preperiodic with period 2.

set $J(f_c)$, the Bernoulli generator was constructed in the course of the proof of Theorem 4.13:

EXERCISE 4.17. If J(f) is a Cantor set, then the restriction of f onto J(f) is topologically conjugate to the one-sided Bernoulli shift with two states.

21.5. Real quadratic family.

21.5.1. Two shift loci. In the case of real parameter values c, the Bernoulli coding of $J(f_c)$ becomes particularly nice:

EXERCISE 4.18. Consider a quadratic polynomial $f_c : z \mapsto z^2 + c$ with a real c. Let $J \equiv J(f_c)$.

(i) If c < -2 then J is a Cantor set on the real line. In this case the Bernoulli generator for f_c consists of

 $J^0 = J \cap \{z : \operatorname{Re} z < 0\} \text{ and } J^1 = J \cap \{z : \operatorname{Re} z > 0\}.$

picture

(ii) If c > 1/4 then J is a Cantor set disjoint from the real line. In this case the Bernoulli generator for f_c consists of

 $J^0 = J \cap \{z : \operatorname{Im} z > 0\} \text{ and } J^1 = J \cap \{z : \operatorname{Im} z < 0\}.$

We see, in particular, that for $c \in \mathbb{R} \setminus [-2, 1/4]$, the map f_c has no invariant intervals on the real line.

21.5.2. Invariant interval and fixed points. We say that a fixed point β of a real map f is orientation preserving if it has a positive multiplier: $f'(\beta) > 0$. Similarly, it is orientation reversing if the multiplier is negative.

EXERCISE 4.19. For $c \in [-2, 1/4]$, the map f_c has an invariant interval. The maximal invariant interval has a form $I_c = [-\beta_c, \beta_c]$, where β_c is an orientation preserving fixed of f_c . Moreover, this point is either repelling or neutral, i.e. $f'_c(\beta_c) \geq 1$ (it is neutral only for c = 1/4). The orbits of $x \in \mathbb{R} \setminus I_c$ escape to ∞ . Finally, I_c is the real slice of the Julia set: $I_c = J_c \cap \mathbb{R}$.

The quadratic maps $f_c: I_c \to I_c, c \in [-2, 1/4]$, give us examples of unimodal interval maps. A unimodal interval map $f: I \to I$ is a continuous map that has exactly two intervals of monotonicity (and hence, it has exactly one extremum $o \in \text{int } I$). In this book we will assume, unless otherwise is explicitly stated that a unimodal map f under consideration is smooth, o is its only critical point, and a is non-degenerate: $f''(o) \neq 0$.

We say that a unimodal map f is *proper* if $f(\partial I) \subset \partial I$, i.e., one of the boundary points of I is fixed, while the other is its preimage. (More pedantically, this means that the map $f : \operatorname{int} I \to \operatorname{int} I$ is proper in the usual sense). Note that the fixed boundary point must have a positive multiplier.

A proper unimodal map has a dynamical symmetry: $\sigma : x \mapsto x'$, where f(x) = f(x'). We will be mostly dealing with even unimodal maps, so o = 0 and σ is the usual central symmetry $x \mapsto -x$. In fact, in what follows we will assume that f is even unless otherwise is explicitly stated.

Let us now return to the quadratic family. The boundary parameter values c = 1/4 and c = -2 play a special role in one-dimensional dynamics (both real and complex).

EXERCISE 4.20. The map $f \equiv f_{1/4} : z \mapsto z^2 + 1/4$ is singled out among the quadratic maps $f_c, c \in \mathbb{C}$, by the property that it has a multiple fixed point $\alpha = \beta = 1/2$, i.e., $f(\beta) = \beta$, $f'(\beta) = 1$. In this case, $f^n x \to \beta$ for all $x \in I$ (where $I = I_c$ is the in

The Julia set of $f_{1/4}$ is a Jordan curve depicted on Figure ... (see §?? for an explanation of some features of this picture). It is called the *cauliflower*, and the map $f_c: z \mapsto z^2 + 1/4$ itself is sometimes called the *cauliflower map*.

Let take a look at what happens as c crosses 1/4:

EXERCISE 4.21. For any c < 1/4, the map f_c has two fixed points $\alpha_c < \beta_c$. The point α_c is attracting for $c \in (-3/4, 1/4)$, and repelling for c < -3/4. It is orientation preserving for $c \in (0, 1/4)$, and orientation reversing for c < 0.

Moreover, the multiplier $\sigma : c \mapsto f'_c(\alpha)$ is an orientation preserving diffeomorphism $[-3/4, 1/4] \to [-1, 1]$.

One says that the saddle-node bifurcation occured at c = 1/4 and the superattracting bifurcation occured at c = 0.

21.5.3. Chebyshev map and the saw-like map. The latter map (c = -2) is specified by the property that the second iterate of the critical point is fixed under $f_c: 0 \mapsto -2 \mapsto 2 \mapsto 2$ (see Figure ...). This map is called *Chebyshev* or *Ulam-Neumann*. The Julia set of this map is unusually simple:

EXERCISE 4.22 (Chebyshev map). Let $f \equiv f_{-2} : z \mapsto z^2 - 2$.

- The interval I = [-2, 2] is completely invariant under f, i.e., $f^{-1}I = I$.
- J(f) = I. (To show that all points in $\mathbb{C} \setminus I$ escape to ∞ , use Montel's Theorem.)
- Consider the the saw-like map

 $g: [-1,1] \to [-1,1], \quad g: x \mapsto 2|x| - 1.$

Show that the map $h: x \mapsto 2\sin \frac{\pi}{2}x$ conjugates g to f|I.

• The map f|I is nicely semi-conjugate to the one-sided Bernoulli shift σ : $\Sigma \to \Sigma$. Namely, there exists a natural semi-conjugacy $h: \Sigma \to I$ such that card $h^{-1}x = 1$ for all $x \in I$ except countable many points (preimages of the fixed point $\beta = 2$ under iterates of f). For these special points, card $h^{-1}(x) = 2$.

21.5.4. *Real Basic Dichotomy*. Let us finish with a statement which will complete our discussion of the Basic Dichotomy for real parameter values:

EXERCISE 4.23. For $c \in (-\infty, 1/4)$, the map f_c has two real fixed points $\alpha_c < \beta_c$. (We have already observed that these two points collide at 1/2 when c = 1/4.) Point β_c is always repelling. Moreover, for $c \in [-2, 1/4]$ we have:

- (i) The interval $I_c = [-\beta_c, \beta_c]$ is invariant under f_c , and it is the maximal f_c -invariant interval on the real line.
- (ii) The critical point is non-escaping and hence the Julia set J_c is connected. Moreover, the interval I_c is the real slice of the Julia set: I_c = J_c ∩ ℝ.
- (iii) For c < 0, the interval $T_c = [c, f(c)]$ is the minimal f-invariant interval containing the critical point.

The above fixed points, α_c and β_c , will be called α - and β -fixed points respectively. As one can see from the second item of the above Exercise, they play quite a different dynamical role. In §?? a similar classification of the fixed points will be given for any quadratic polynomial with connected Julia set.

Let us summarize Exercises 4.18 and 4.23:

PROPOSITION 4.24. For real c, the Julia set J_c is connected if and only if $c \in [-2, 1/4]$.

21.5.5. Period doubling bifurcation.

EXERCISE 4.25. For parameters c < -3/4 near -3/4, the map f_c has an attracting cycle γ_c of period 2. This attracting cycle persits on the parameter interval (-5/4, -3/4) and its multiplier is an orientation preserving diffeomorphism from this interval onto (-1, 1).

EXERCISE 4.26. For $c \in [-5/4, 0)$, the interval $T_c = [-\alpha_c, \alpha_c]$ is a periodic interval of period 2, i.e., $f_c^2(T_c) \subset T_c$, and moreover, $f_c(T_c) \cap T_c = \{\alpha_c\}$. The restriction $f_c^2 | T_c$ is a unimodal map. Any orbit in I_c eventually lands in T_c , i.e., for any $x \neq \pm \beta_c$, we have $f^n x \in T_c$ for some $n \in \mathbb{N}$. [In fact, the periodic interval T_c persists until much smaller parameter $c_* < -5/4$ – which one?]

This gives us the first glimples of a fundamental phenomenon that would play a central role throughout this book.

Let us conclude with a useful remark:

EXERCISE 4.27. Assume a proper unimodal map f has neither attracting fixed points in int I, nor attracting cycles of period two with non-negative multiplier. Then the critical point 0 lies in between f(0) and $f^2(0)$, and the interval $[f(0), f^2(0)]$ is invariant. Thus, this is the smallest invariant interval containing 0.

21.5.6. Formation of monotnicity intervals. Let us consider a real quadratic polynomial $f \equiv f_c$ with $c \in [-2, 1/4]$, and let $I \equiv I_c$ be it invariant interval from Exercise 4.23. For $n \in \mathbb{Z}_+$ and $x \in I \setminus \operatorname{Crit}(f^n)$, let $L_n(x) \subset \mathbb{R}$ be the maximal interval containing x on which f^n is monotone. The boundary points of $L_n(x)$ belong to $\operatorname{Crit}(f^n) \cup \partial I$. By (21.1), for each endpoint $a \in \partial L_n(x)$, there exists an integer $k \in [0, n-1]$ such that $f^k a = 0$, so the interval $f^k L$ "grabs" the critical point 0 and "carries it forward" to the image $M_n(x) := f^n L_n(x)$.

21.5.7. Inverse branches. For an interval $M \subset \mathbb{R}$ we let I° be its interior rel the real line.

$$\mathbb{C}M_n(x) := \mathbb{C} \smallsetminus (\mathbb{R} \smallsetminus M_n^{\circ}(x))$$

be the complex plane slit by two rays that complement $M_n(x)$. By Corollary 4.10, there is a well defined inverse branch $f^{-n} : \mathbb{C}M_n(x) \to \mathbb{C}$ that map $M_n(x)$ to $L_n(x)$.

LEMMA 4.28. The image of the half-plane \mathbb{H}_+ under the above branch of f^{-n} is contained in one of the half planes \mathbb{H}_+ or \mathbb{H}_- (depending on whether $f^n : L_n(x) \to M_n(x)$ is orientation preserving or reversing). Similarly for the half-plane \mathbb{H}_- .

PROOF. Since $f^n(\mathbb{R}) \subset \mathbb{R}$, we have $f^{-n}(\mathbb{C} \setminus \mathbb{R}) \subset \mathbb{C} \setminus \mathbb{R}$. All the more, any inverse branch of f^{-n} maps the half-plane \mathbb{H}_+ inside $\mathbb{C} \setminus \mathbb{R}$. The orinetation rule comes from the fact that f^n preserves orientation of \mathbb{C} .

There is a nice way to visualize these branches. Let us color the half-plane \mathbb{H}_+ in black while keeping \mathbb{H}_- white. Then the complement $\mathbb{C} \smallsetminus f^{-n}(\mathbb{C} \smallsetminus \mathbb{R})$ assumes the chess-board cloloring as on Figure ??.

21.6. Fatou set. The Fatou set is defined as the complement of the Julia set:

$$F(f) = \widehat{\mathbb{C}} \smallsetminus J(f) = D_f(\infty) \cup \operatorname{int} K(f).$$

Since K(f) is full, all components of int K(f) are simply connected. Only one of them can contain the critical point. Such a component (if exists) is called *critical*.

Let U be one of the components of int K. Since int K is invariant, it is mapped by f to some other component V. Moreover, $f(\partial U) \subset \partial V$ since the Julia set is also invariant. Hence $f: U \to V$ is proper, and thus surjective. Moreover, since V is simply connected, $f: U \to V$ is either a conformal isomorphism (if U is not critical), or is a double branched covering (if U is critical).

The Fatou set can be also characterized as *the set of normality* (and was actually classically defined in this way):

PROPOSITION 4.29. The Fatou set F(f) is the maximal set on which the family of iterates f^n is normal.

PROOF. On $D_f(\infty)$, the iterates of f locally uniformly converge to ∞ , while on int K(f) they are uniformly bounded. Hence they form a normal family on F(f). On the other hand, if $z \in J(f)$, then the orbit of z is bounded while there are nearby points escaping to ∞ . Hence the family of iterates is not normal near z. \Box

21.7. Postcritical set. Let $O_f = cl\{f^n(0)\}_{n=1}^{\infty}$ stand for the *postcritical set* of f. It is forward invariant and contains the critical value c of f. The map f is tremendously contracting near the critical point 0, and under iteration this contraction propagates through the postcritical set. The following lemma is the first indication that otherwise, the map f tends to be expanding:

LEMMA 4.30. Let $c \neq 0$. Then the complement $\mathbb{C} \setminus O_f$ is hyperbolic². Let Ω be a component of $\mathbb{C} \setminus O_f$ that intersects $f^{-1}(O_f) \setminus O_f$. Then f on Ω is strictly expanding with respect to this hyperbolic metric, i.e, for any $z \in \Omega \setminus f^{-1}(O_f)$, $\|Df(z)\|_{\text{hyp}} > 1$.

PROOF. If $\mathbb{C} \setminus O_f$ is not hyperbolic, then O_f consists of a single point, c. But then f(c) = c and hence c = 0.

Let ρ and ρ' be the hyperbolic metrics on $\mathbb{C} \smallsetminus O_f$ and $\mathbb{C} \smallsetminus f^{-1}(O_f)$ respectively. Since the map $f : \mathbb{C} \smallsetminus f^{-1}O_f \to \mathbb{C} \smallsetminus O_f$ is a covering, it is a local isometry from ρ' to ρ .

Let Ω' be the component of $\mathbb{C} \smallsetminus f^{-1}(O_f)$ containing z. Since O_f is forward invariant, $\Omega' \subset \Omega$, and by the assumption of the lemma, Ω' is strictly smaller than Ω . By the Schwarz Lemma, the natural emebdding $i : \Omega' \to \Omega$ is strictly contracting from ρ' to ρ . Thus, the inverse map $i^{-1} | \Omega'$ is strictly expanding from ρ to ρ' , and we conclude that the composition $f \circ i^{-1} : \Omega' \to \mathbb{C} \smallsetminus O_f$ is strictly expanding with respect to ρ .

21.8. Preimages of points.

PROPOSITION 4.31. Let $f: z \mapsto z^2 + c$. If $c \neq 0$, then for any neighborhood U intersecting J(f) we have:

$$\operatorname{orb} U := \bigcup_{n=0}^{\infty} f^n U = \mathbb{C}.$$

If c = 0 and $U \not\supseteq 0$ then orb $U = \mathbb{C}^*$.

PROOF. By the Montel Theorem, $\mathbb{C} \setminus \operatorname{orb} U$ contains at most one point. If there is one, a, then $f^{-1}a = \{a\}$. Hence a is the critical point of f, i.e., a = 0. Moreover, f(a) = a, so c = 0.

This result immediately yields:

COROLLARY 4.32. For any point $z \in \mathbb{C}$, except z = 0 in case $f : z \mapsto z^2$, we have:

$$\operatorname{cl}\bigcup_{n=0}^{\infty}f^{-n}z\supset J(f).$$

COROLLARY 4.33. If $J' \subset J(f)$ is a non-empty backward invariant closed subset of J(f) then J' = J(f). If $K' \subset K(f)$ is a non-empty full backward invariant closed subset of K(f) then K' = K(f).

²Meaning that each component of $\mathbb{C} \smallsetminus O_f$ is hyperbolic.

21.9. Higher degree polynomials. The above basic definitions and results admit a straightforward extension to higher degree polynomials

$$f: z \mapsto a_0 z^d + a_1 z^{d-1} + \dots + a_d, \quad d \ge 2, \ a_0 \ne 1.$$

The following point should be kept in mind though: the Basic Dichotomy is not valid any more in the higher degree case. Instead, there is the following partial description of the topology of the Julia set:

• The Julia set J(f) (and the filled Julia set K(f)) is connected if and only all the critical points c_i are non-escaping to ∞ , i.e., $c_i \in K(f)$.

• If all the critical points escape to ∞ , then J(f) is a Cantor set.

However, the Basic Dichotomy is still valid in the case of *unicritical* polynomials, that is, the ones that have a single critical point. (Note that any such polynomial is affinely conjugate to $z \mapsto z^d + c$.)

Note also that the exceptional cases in Proposition 4.31 are polynomials affinely conjugate to $z \mapsto z^d$.

EXERCISE 4.34. Work out the basic dynamical definitions and results in the case of higher degree polynomials.

In the theory of quadratic maps f_c , higher degree polynomials still appear as the iterates of f_c . It is useful to know that they have the same Julia set:

EXERCISE 4.35. Show that $K(f^n) = K(f)$ for any polynomial f.

In the higher degree, we will keep notation

(21.2)
$$\operatorname{Crit}_{f} = \bigcup_{k=0}^{i} nfty \bigcup_{i} f^{-k}(c_{i})$$

for the set of critical points of all iterates of f.

22. Periodic motions

"Peirodic solutions is the only openning through which we can try to penetrate to the domain that was viewed unaccessible" (Poincaré [**Poi**, §36].

22.1. Rough classification of periodic points by the multiplier. Consider a periodic point α of period p. The local dynamics near its cycle $\alpha = \{f^n \alpha\}_{n=0}^{p-1}$ depends first of all on its *multiplier*

$$\sigma = (f^p)'(z) = \prod_{n=0}^{p-1} f'(f^n \alpha).$$

The point (and its cycle)³ is called *attracting* if $|\sigma| < 1$, A particular case of an attracting point is a *superattracting* one when $\sigma = 0$. In this case, the critical point 0 belongs to the cycle α . (When we want to emphasize that an attracting periodic point is not superattracting, we call it *simply* attracting.)

A periodic point is called *repelling* if $|\sigma| > 1$, and *neutral* if $\sigma = e^{2\pi i\theta}$, $\theta \in \mathbb{R}/\mathbb{Z}$. In latter case, θ is called the *rotation number* of α . Local dynamics near a neutral cycle depends delicately on the arithmetic of the rotation number. A neutral point is called *parabolic* if the rotation number is rational, $\theta = r/q$, and is called *irrational*

 $^{^{3}\}mathrm{All}$ the terminology introduced for periodic points applies to their cycles, and vice versa.

othewise. An irrational periodic point can be of *Siegel* and *Cremer* type, to be defined below.

Let us consider these cases one by one.

22.2. Attracting cycles. Let α be an attracting cycle. The orbits of all nearby points uniformly converge to α and, in particular, are bounded. It follows that attracting cycles belong to F(f). The rate of convergence is exponential in the simply attracting case and superexponential in the superattracting case.

For a simply attracting periodic point α , we say that a piecewise smooth (open) disk $P \ni \alpha$ is a *petal* of α if f|P is univalent and $f(P) \Subset P$. (For instance, one can take a small round disk $\mathbb{D}(\alpha, \epsilon)$ as a petal.) Then the annulus $A = \overline{P} \smallsetminus f^p(P)$ is called a *fundamental annulus* of α .

In the superattracting case, a *petal* is a smooth disk $P \ni \alpha$ such that $f^p : P \to f^p(P)$ is a branched covering of degree d (with a single critical point at α), and $f(P) \subseteq P$. (For instance, one can let P be the component of $f^{-p}(\mathbb{D}(\alpha, \epsilon))$ containing α .) The corresponding fundamental annulus is $\bar{P} \smallsetminus f^p(P)$.

The basin of attraction of an attracting cycle α is the set of all points whose orbits converge to α :

$$D(\boldsymbol{\alpha}) = D_f(\boldsymbol{\alpha}) = \{z: f^n z \to \boldsymbol{\alpha} \text{ as } n \to \infty.\}$$

EXERCISE 4.36. Show that the basin $D(\alpha)$ is a completely invariant union of components of int K(f).

The union of components of $D(\alpha)$ containing the points of α is called the *immediate basin* of attraction of the cycle α . We will denote it by $D^0 = D_f^0(\alpha)$. The component of $D^0(\alpha)$ containing α will be denoted $D^0(\alpha) = D_f^0(\alpha)$.

EXERCISE 4.37. (i) The immediate basin of an attracting cycle consists of exactly p components, where p is the period of α .

(ii) Show that it can be constructed as follows. Let P_0 be a petal of α and let P_n be defined inductively as the component of $f^{-p}(P_{n-1})$ containing α . Then $P_0 \subset P_1 \subset P_2 \subset \ldots$, and

$$D^0(\alpha) = \bigcup_{n=0}^{\infty} P_n.$$

We will now state one of the most important facts of the classical holomorphic dynamics:

THEOREM 4.38. The immediate basin of attraction $D_f^0(\alpha)$ of an attracting cycle α contains the critical point 0. Moreover, if α is simply attracting then the critical orbit orb(0) crosses any fundamental annulus A.

REMARK 4.3. Of cource, the assertion is trivial when α is superattracting as $0 \in \alpha$ in this case.

PROOF. Otherwise f^p would conformally map each component D of the immediate basin onto itsef. Hence it would be a hyperbolic isometry of D, despite the fact that $|f'(\alpha)| < 1$.

To prove the second assertion (which would also give another proof of the first one), let us consider a petal P_0 containing some point α of α , and let us define P_n inductively as the component of $f^{-p}(P_{n-1})$ containing α (compare with Exercise 4.37 above). Then $P_0 \subset P_1 \subset P_2 \subset \ldots$ If non of these domains contains

a critical point of f^p , then the all the maps $f^p: P_n \to P_{n-1}$ are isomorphisms and all the P_n are topological disks. Hence their union, P_{∞} , is a topological disk as well, and $f^p: P_{\infty} \to P_{\infty}$ is an automorphism. Hence it is a hyperbolic isometry contradicting the fact that α is attracting.

Hence some P_n contains a critical point of f^p . Take the first such n (obviously, $n \geq 1$). Then $P_n \setminus P_{n-1}$ contains a critical value of f^p , which is contained in orb(0). Applying further iterates of f^p , we will bring it to the fundamental annulus. \Box

COROLLARY 4.39. A quadratic polynomial can have at most one attracting cycle. If it has one, all other cycles are repelling.

PROOF. The first assertion is immediate. For the second one, notice that under the circumstances, the postcritical set O_f is a discrete set accumulating on the attracting cycle α . Hence it does not divide the complex plane, and $0 \in \mathbb{C} \setminus O_f$. Applying Lemma 4.30, we conclude that $|\sigma(\beta)| = ||Df^q(\beta)||_{\text{hyp}} > 1$ for any other periodic point β of period q.

Of course, the period of the attracting cycle can be arbitrary big. A quadratic polynomial is called *hyperbolic* if it either has an attracting cycle, or if its Julia set is Cantor. (The unifying property is that for hyperbolic maps, orb(0) coverges to an attracting cycle in the Riemann sphere.) For instance, polynomials $z \mapsto z^2$, $z \mapsto z^2 - 1$,... (see Figure ...) are hyperbolic. Though dynamically non-trivial, it is a well understood class of quadratic polynomials (see §23).

22.3. Parabolic cycles.

22.3.1. Leau-Fatou Flower. Let us consider a parabolic germ

$$f: z \mapsto e^{2\pi r/q} z + a z^2 + \dots$$

with rotation number r/q near the origin.

EXERCISE 4.40. The first non-vanishing term of the expansion $f^q(z) = z + b_k z^k + \ldots$ has order k = ql + 1 for some $l \in \mathbb{N}$.

We call *l* the order of degeneracy of f at 0. In the case l = 1, the parabolic germ f is called *non-degenerate*.

An open Jordan disk P is called an *attracting petal* for f if:

- $0 \in \partial P;$
- $f^q(P) \subset P$ and $f^q|P$ is univalent;
- 0 is the only point where ∂P and $\partial (f^q P)$ touch;
- If $z \in P$ then $f^{qn}z \to 0$ as $n \to \infty^4$.

Given such a petal, the set $\overline{P} \setminus \overline{f(P)}$ is called an attracting *fundamental crescent*. We say that a petal P has wedge γ at 0 if both local branches of the boundary

 $\partial P \smallsetminus \{\alpha\}$ have tangent lines at 0 that meet at angle γ .

Two attracting petals are called *equivalent* if they overlap⁵.

 $^{{}^{4}}$ It is convenient to impose this condition, though in fact it can be derived from the the other properties.

 $^{{}^{5}}$ At the moment, it is not evident that this is an equivalence relation, but the following theorem shows that it is.

THEOREM 4.41. There is a choice of disjoint lq petals P_i (one in each class) with wedge $2\pi/ql$ at 0 such that the flower $\Phi = \bigcup P_i$ is invariant under rotation by $2\pi/ql$ and under f. The orbits is $z \in \Phi$ converge to 0 locally uniformly. Vice versa, if some orb(z) converges to 0 without direct landing at 0 then eventually it lands in the flower Φ .

PROOF. The proof will be split in several cases. The main analysis happens in the following one:

The germ f is non-degenerate with zero rotation number. Thus $f: z \mapsto z + az^2 + \ldots, a \neq 0$. Conjugating f by complex scaling $\zeta = az$ we make a = 1.

Let us move the fixed point to ∞ by inversion $Z = -\frac{1}{z}$. It brings f to the form

(22.1)
$$F: Z \mapsto Z + 1 + O(\frac{1}{Z})$$

near ∞ . It is obvious from this asymptotical expression that any right half-plane

with t > 0 sufficiently big is invariant under F, and in fact

$$(22.3) F(Q_t) \subset Q_{t+1-\epsilon},$$

where $\epsilon = \epsilon(t) \to 0$ as $t \to \infty$. So, such a half-plane provides us with a petal with wedge π at ∞ . Moreover,

(22.4)
$$\operatorname{Re}(F^n Z) \ge \operatorname{Re} Z + (1 - \epsilon)n,$$

so the orbits in Q_t converge to ∞ locally uniformly.

Vice versa, if $F^n Z \to \infty$ without direct landing at it, then due to asymptotical expression (22.1) we eventually have $\operatorname{Re}(F^{n+1}Z) \ge \operatorname{Re}(F^n Z) + 1 - \epsilon$. Hence $\operatorname{Re}(F^n Z) \to +\infty$ and orb z eventually lands in the halh-plane Q_t .

Now we would like to enlarge Q_t to a petal P with wedge 2π at ∞ . To this end let us consider two logarithmic curves

$$\Gamma_{\pm} = \{Y = \pm C \log(t - X + 1) + R)\}, X \le t, \text{ where } Z = X + iY.$$

If R is big enough then Γ_{\pm} lie in the domain where the asymptotics (22.1) applies. If C is big enough then the half-slope of these curves is bigger (in absolute value) than the slope of the displaycement vector F(Z) - Z. It follows that F moves the curves Γ_{\pm} to the right, and the region P bounded by these curves and the segment of the vertical line $\operatorname{Re} Z = t$ in between is mapped univalently into itself. This is the desired petal.

Let f be a general parabolic germ with zero rotation number, $f: z \mapsto z+bz^{k+1}+$... with $k \ge 1$, $b \ne 0$. Again, conjugating f by a complex scaling $\zeta = \lambda z$, where $\lambda^k = b$, we make b = 1.

Let us now use a non-invertible change of variable $\zeta = z^k$. A formal calculation shows that it conjugates f to a multi-valued germ $g: \zeta \mapsto \zeta + \zeta^2 + O(|\zeta|^{2+1/k})$, where the residual term is given by a power series in $\zeta^{1/k}$. Making now a change of variable Z = -1/z we come up with a multi-valued germ $G: Z \mapsto Z + 1 + O(1/Z^{1/k})$ near ∞ . Let us take any single-valued branch of this germ in the slit plane $\mathbb{C} \setminus \mathbb{R}_0$. Then the same considerations as in the non-degenerate case show that G has a petal P with wedge 2π at ∞ . Lifting this petal to the z-plane, gives us k petals of f with wedge $2\pi/k$.

Let us now consider a parabolic periodic point α with period p and rotation number r/q. As the following exercise implies, $\alpha \in J(f)$:

EXERCISE 4.42. Show that $(f^{pqn})'(\alpha) \to \infty$.

The basin of attraction of a parabolic cycle α is defined as follows:

 $D_f(\alpha) = \{z : f^n z \to \alpha \text{ as } n \to \infty \text{ but } f^n z \notin \alpha \text{ for any } n \in \mathbb{N}^*\}.$

EXERCISE 4.43. The basin of attraction $D_f(\alpha)$ is a completely invariant union of components of int K. Moreover, among these components there are pql components permuted by f, while all others are preimages of these. (Here l comes from Theorem 4.41 applied to f^{pq} ; compare also Exercise 4.40.)

The union of these pql components is called the *immediate basin* of attraction of α . It will also be denoted as $D_f^0(\alpha)$. Each of these components is periodic with period pq. So, the immediate basin comprises l cycles of periodic components.

Let us take a component D of $D_f^0(\alpha)$ and a petal $P \subset D$. The map $g = f^{pq}$: $D \to D$ maps one boundary component of the fundamental crescent $\overline{P} \smallsetminus f(P)$ to the other, so the quotient $C_f := D/g = P/g$ is a topological cylinder called *Ecalle-Voronin cylinder*. A priori, there are several options for the conformal type of C_f : it can be isomorphic to an annulus $\mathbb{A}(r, R)$, or to the punctured disc $\mathbb{H}/\langle z+1 \rangle \approx \mathbb{D}^*$, or to the bi-infinite cylinder $\mathbb{C}/\langle z+1 \rangle \approx \mathbb{C}^*$. In fact, the latter happens:

LEMMA 4.44. The Ecale-Voronin cylinder C_f is isomorphic to the bi-infinite cylinder.

PROOF. Notice first that the cylinder $C = C_f := D/\langle g \rangle$ is independent of the petal P, so we can make any convenient choice. Let us use for this purpose the half-plane Q_t (22.2) near ∞ . Then the fundamental crescent $\bar{Q} \setminus F(Q)$ becomes a vertical strip S whose boundary curves stay ditance ~ 1 apart, by (22.1). Moreover, the right-hand boundary curve $Y \mapsto F(it + Y)$ is almost vertical, so each straight interval $I_Y := [t + iY, F(t + iY)]$ cuts S into two half-strips. Projecting these intervals to the cylinder C = S/F, we obtain a horizontal foliation Γ on S by circles that we also denote I_Y .

Let C^{\pm} be the half-cylinders obtained by cutting C by the circle I_0 . It is enough to show that

(22.5)
$$\operatorname{mod} C^{\pm} \equiv \mathcal{W}(\Gamma | C^{\pm}) = \infty.$$

Let us deal with C^+ for definiteness. Let us further cut the cylinder on some big height H > 0 by the circle I_H), and call the corresponding cylinder C_H^+ . Put any conformal metric $\rho = \rho(z)|dz|$ on C_H^+ with

(22.6)
$$\int_{I_Y} \rho \, dl_Y = l_\rho(I_Y) \ge \mathcal{L}_\rho(\Gamma | C_H^+) \ge 1, \quad Y \in [0, H],$$

where dl_Y is the Euclidean length element along I_Y . Since the circles I_Y are almost horizontal, we have for the Euclidean area form $d\sigma = dl_Y \wedge dY \geq (1/2) dl_Y dY$. Hence, integrating (22.6) over dY gives us:

$$\int_{C_H^+} \rho \, d\sigma \ge cH,$$

with c > 0 independent of ρ . By the Cauchy-Schwarz Inequality (compare (3.1)), we obtain:

$$H m_{\rho}(C_H^+) \asymp \operatorname{area}(C_H^+) \int \rho^2 \, d\sigma \ge \left(\int \rho \, d\sigma\right)^2 \ge c^2 H^2.$$

So, $\operatorname{mod}(C_H^+) \ge c^2 H \to \infty$ as $H \to \infty$, and we are done.

As in the attracting case, we have:

THEOREM 4.45. The immediate basin $D_f^0(\boldsymbol{\alpha})$ of a parabolic cycle contains the critical point. In fact, each cycle of components of $D_f^0(\boldsymbol{\alpha})$ contains the critical point.

PROOF. Let D be a component of $D_f^0(\alpha)$. If it does not contain critical points of $g := f^{pq}$, then $g : D \to D$ is an (unbranched) covering, and hence an automorphism of D (since D is simply connected). Since the orbits of g in D escape to infinity (of D) and $D \approx \mathbb{D}$ is hyperbolic, the quotient $D/\langle g \rangle$ is isomrphic to either an annulus $\mathbb{A}(r, R)$ (if g is hyperbolic) or to the punctured disc \mathbb{D}^* (if g is parabolic), contradicting Lemma 4.44.

As in the hyperbolic case, we now conclude:

COROLLARY 4.46. A quadratic polynomial f can have at most one parabolic cycle. Moreover, this cycle is non-degenerate: it comprises a unique cycle of periodic components (i.e. l = 1). If f has a parabolic cycle, then all other cycles are repelling.

Such a quadratic polynomial is naturally called *parabolic*.

22.4. Repelling cycles. Let us now consider a repelling cycle $\boldsymbol{\alpha} = \{f^k \alpha\}_{k=0}^{p-1}$. Nearby points escape (exponentially fast) from a small neighborhood of $\boldsymbol{\alpha}$, which implies that the family of iterates f^n is not normal near $\boldsymbol{\alpha}$. Hence repelling periodic points belong to the Julia set. In fact, as we are about to demonstrate, they are dense in the Julia set, so that the Julia can be alternatively defined as the closure of repelling cycles. It gives us a view of the Julia set "from inside".

But first, let us now show that almost all cycles are repelling:

LEMMA 4.47. A quadratic polynomial may have at most two non-repelling cycles.

PROOF. Let α_{o} be a neutral periodic point of period p with multiplier σ_{o} of a quadratic polynomial $f_{o}: z \mapsto z^{2} + c_{o}$. Due to Lemma 4.46, we can assume that $\sigma_{o} \neq 1$. Then by the Implicit Function Theorem, the equation $f^{p}(z) = z$ has a local holomorphic solution $z = \alpha_{c}$ assuming value α_{o} at c_{o} . The multiplier of this periodic point, $\sigma_{c} = (f_{c}^{p})'(\alpha_{c})$ is also a local holomorphic function of c. In fact, it is a global algebraic function. So, if it was locally constant then it would be globally constant, and the map $f_{0}: z \mapsto z^{2}$ would have a neutral cycle. Since this is not the case, the multiplier is not constant, and hence near c_{o} it assumes all values in some neighborhood of σ_{o} . In particular, it assumes values with $|\sigma| < 1$. Moreover, if near c_{o}

$$\sigma(c) = \sigma_{o} + a(c - c_{o})^{k} + \dots, \quad a \neq 0,$$

then the set $\{c : |\sigma(c)| < 1\}$ is the union of k sectors that asymptotically occupy 1/2 of the area of a small disk $\mathbb{D}(c_{o}, \epsilon)$. It follows that if we take three of

140

such multiplier functions, then two of them must have overlapping sectors, so that the corresponding two cycles can be made simultaneously attracting, contradicting Corollary 4.39. $\hfill\square$

THEOREM 4.48. The Julia set is the closure of repelling cycles.

PROOF. Let us first show that any point of the Julia set can be approximated by a periodic point. Let $z \in J(f)$ be a point we want to approximate. Since the Julia set does not have isolated points (see Corollary 4.14), we can assume that z is not the critical value. Then in a small neighborhood $U \ni z$, there exist two branches of the inverse function, $\phi_1 = f_1^{-1}$ and $\phi_2 = f_2^{-1}$. Since the family of iterates is not normal in U, one of the equations, $f^n z = z$, $f^n z = \phi_1(z)$, or $f^n z = \phi_2(z)$, has a solution in U for some $n \ge 1$ (by the Refined Montel Theorem). If it is an equation of the first series, we find in U a periodic point of period n (maybe, not the least one). Otherwise, we find a periodic point of period n + 1.

Since by Lemma 4.47, almost all periodic points are repelling, we come to the desired conclusion. $\hfill \Box$

22.5. Siegel and Cremer cycles. Irrational periodic points may or may not belong to the Julia set (depending primarily on the Diophantine properties of its rotaion number). Irrational periodic points lying in the Fatou set are called *Siegel*, and those lying in the Julia set are called *Cremer*. The component of F(f) containing a Siegel point is called a *Siegel disk*. Local dynamics on a Siegel disk is quite simple:

PROPOSITION 4.49. Let U be a Siegel disk of period p containing a periodic point α with rotation number θ . Then $f^p|U$ is conformally conjugate to the rotation of \mathbb{D} by θ .

PROOF. Consider the Riemann map $\phi : (U, \alpha) \to (\mathbb{D}, 0)$. Then $g = \phi \circ f^p \circ \phi^{-1}$ is a holomorphic endomorphism of the unit disk fixing 0, with $|g'(0)| = |\lambda| = 1$. By the Schwarz Lemma, $g(z) = \lambda z$.

We will see later on that a quadratic polynomial can have at most one nonrepelling cycle (see Theorem 4.80). If it has one, it can be non-contradictory classified as either hyperbolic, or parabolic, or Siegel, or Cremer.

Let us show that Cremer cycles indeed exist:

PROPOSITION 4.50. In the family $f_{\theta} : z \mapsto e^{2\pi i \theta} z + z^2$, $\theta \in \mathbb{Z}/2\pi Z$, the origin 0 is the Cremer fixed point for a generic rotation number θ .

PROOF. Let us consider the set $\Lambda \subset \mathbb{R}/2\pi\mathbb{Z}$ of rotation numbers θ for which $0 \in J(f_{\theta})$. We have $\Lambda = \Lambda_p \sqcup \Lambda_C$, where Λ_p is the set of parabolic (i.e., is rational), rotation numbers, while Λ_C is the set of Cremer (i.e., irrational) numbers.

We will show that Λ is of type G_{δ} , i.e., it is the countable intersection of open sets. Since Λ_p is dense, Λ is a dense G_{δ} , so rotation numbers $\theta \in \Lambda$ are generic by definition. Of course, removing a countable subset preserves genericity, so the conclusion would follow.

To prove that Λ is G_{δ} , let us consider a function $\rho : \mathbb{R}/2\pi\mathbb{Z} \to \mathbb{R}_{\geq 0}$, $\rho(\theta) = \text{dist}(0, J(f_{\theta}))$. Then Λ is the set of zeros of ρ . We will show that ρ is uppersemicontinuous. Since the set of zeros of a non-negative upper semicontinuous function is of type G_{δ} , it will complete the proof.

22.6. Periodic components. The notions of a periodic component of F(f) and its cycle are self-explanatory. It is classically known that such a component is always associated with a non-repelling periodic point:

THEOREM 4.51. Let $\mathbf{U} = \{U_i\}_{i=1}^p$ be a cycle of periodic components of int K(f). Then one of the following three possibilities can happen:

- U is the immediate basin of an attracting cycle;
- U is the immediate basin of a parabolic cycle $\alpha \subset \partial U$ of some period q|p;
- U is the cycle of Siegel disks.

Proof Take a component U of the cycle \mathbf{U} , and let $g = f^p$. By the Schwarz-Pick Lemma, g|U is either a conformal automorphism of U, or it strictly contracts the hyperbolic metric dist_h on U. In the former case, it is either elliptic, or otherwise. If g is elliptic then U is a Siegel disk. Otherwise the orbits of g converge to the boundary of U.

Let us show that if an orbit $\{z_n = g^n z\}, z \in U$, converges to ∂U , then it converges to a g-fixed point $\beta \in \partial U$. Join z and g(z) with a smooth arc γ , and let $\gamma_n = f^n \gamma$. By the Schwarz-Pick Lemma, the hyperbolic length of the arcs γ_n stays bounded. Hence they uniformly escape to the boundary of U. Moreover, by the relation between the hyperbolic and Euclidean metrics (Lemma 1.97), the Euclidean length of the γ_n shrinks to 0. In particular,

(22.7)
$$|g(z_n) - z_n| = |z_{n+1} - z_n| \to 0$$

as $n \to \infty$. By continuity, all limit points of the orbit $\{z_n\}$ are fixed under g. But g being a polynomial has only finitely many fixed points. On the other hand, (22.7) implies the ω -limit set of the orbit $\{z_n\}$ is connected. Hence it consists of a single fixed point β .

Moreover, the orbit $\{\zeta_n\}$ of any other point $\zeta \in U$ must converge to the same fixed point β . Indeed, the hyperbolic distance between z_n and ζ_n stays bounded and hence the Euclidean distance between these points shrink to 0.

Thus either U is a Siegel disk, or the g-orbits in U converge to a g-fixed point β , or the map $g: U \to U$ strictly contracts the hyperbolic metric and its orbits do not escape to the boundary ∂U . Let us show that in the latter case, g has an attracting fixed point α in U.

Take a g-orbit $\{z_n\}$, and let $d_n = \text{dist}_h(z_0, z_n)$. Since g is strictly contracting,

$$\operatorname{dist}_{h}(z_{n+1}, z_n) \le \rho(d_n) \operatorname{dist}_{h}(z_n, z_{n-1}),$$

where the contraction factor $\rho(d_n) < 1$ depends only on $\operatorname{dist}_h(z_n, z_0)$. Since the orbit $\{z_n\}$ does not escape to ∂U , this contraction factor is bounded away from 1 for infinitely many moments n, and hence $\operatorname{dist}_h(z_{n+1}, z_n) \to 0$. It follows that any ω -limit point of this orbit in U is fixed under g.

By strict contraction, g can have only one fixed point in U, and hence any orbit must converge to this point. Strict contraction also implies that this point is attracting.

We still need to prove the most delicate property: in the case when the orbits escape to the boundary point $\beta \in \partial U$, this point is parabolic. In fact, we will show that $g'(\beta) = 1$. Of course, this point cannot be either repelling (since it attracts some orbits) or attracting (since it lies on the Julia set). So it is a neutral point with some rotation number $\theta \in [0, 1)$. The following lemma will complete the proof.

LEMMA 4.52 (Necklace Lemma). Let $f: z \mapsto \lambda z + a_2 z^2 + \ldots$ be a holomorphic map near the origin, and let $|\lambda| = 1$. Assume that there exists a domain $\Omega \subset \mathbb{C}^*$ such that all iterates f^n are well-defined on Ω , $f(\Omega) \cap \Omega \neq \emptyset$, and $f^n(\Omega) \to 0$ as $n \to \infty$. Then $\lambda = 1$.

PROOF. Consider a chain of domains $\Omega_n = f^n \Omega$ convergin to 0. Without loss of generality we can assume that all the domains lie in a small neighborhood of 0 and hence the iterates $f^n | \Omega$ are univalent. Fix a base point $a \in \Omega$ such that $f(a) \in \Omega$, and let

$$\phi_n(z) = \frac{f^n(z)}{f^n(a)} \,.$$

These functions are univalent, normalized by $\phi_n(a) = 1$, and do not have zeros. By the Koebe Distortion Theorem (the version given in Exercise 1.92,b), they form a normal family. Moreover, any limit function ϕ of this family is non-constant since $\phi(fa) = \lambda \neq 1 = \phi(a)$. Hence the derivatives $\phi'_n | \Omega$ are bounded away from 0 and dist $(1, \partial \Omega_n) \geq \epsilon > 0$ for all $n \in \mathbb{N}$. It follows that

$$\operatorname{dist}(f^n a, \partial \Omega_n) \ge \epsilon \ r_n, \quad n \in \mathbb{N},$$

where $r_n = |f^n a|$. On the other hand, f acts almost as the rotation by θ near 0, where $\theta = \arg \lambda \in (0, 1)$. Since this rotation is recurrent and $\theta \neq 0$, there exists an l > 0 such that

$$\operatorname{dist}(f^{n+l}a, f^n a) = o(r_n) \quad \text{as } n \to \infty$$

The last two estimates imply that $\Omega_{n+l} \cap \Omega_n \neq \emptyset$ for all sufficiently big n.

Hence the chain of domains $\Omega_n, \ldots, \Omega_{n+l}$ closes up, and their union form a "necklace" around 0. Take a Jordan curve γ in this necklace, and let D be the disk bounded by γ . Then $f^n(\gamma) \to 0$ as $n \to \infty$. By the Maximum Principle, $f^N(D) \subseteq D$ for some N. By the Schwarz Lemma, $|\lambda| < 1$ – contradiction.

23. Hyperbolic maps

23.1. Definition revisited. A compact f-invariant set $Z \subset \mathbb{C}$ is called *expanding* or *hyperbolic* (and also, f is called expanding/hyperbolic on Z) if there exist constants C > 0 and $\lambda > 1$ such that

$$|Df^n(z)| \ge C\lambda^n$$
, for any $z \in Z$, $n \in \mathbb{N}$.

Of course, we can define the expanding property with respet to another Riemannian metric $\|\cdot\|$ on Z. The property is independent of a particular choice of the metric as long as the two metrics are equivalent (which is the case if both metrics are defined in a neighborhood of X).

For instance, a *Cantor* Julia set J(f) of a quadratic polynomial f is always expanding, see remark 4.2.

THEOREM 4.53. Let f be a quadratic polynomial with connected Julia set. Then the Julia set J(f) is expanding on its Julia set if and only if f has an attracting cycle α . Moreover, in this case all points $z \in \text{int } K(f)$ are attracted to the cycle α .

PROOF. Assume f has an attracting cycle $\boldsymbol{\alpha} = \{\alpha_k\}_{k=0}^{p-1}$. Take a small invariant neigborhood $U = \bigcup U_k \Subset F(f)$ of $\boldsymbol{\alpha}$. Let n be the first moment when $f^n(0)$ lands in U, and moreover, let $f^n(0) \in U_k$. Let V_i be the pullback (see §20) of U_k containing $f^i(0), k = 0, 1, \ldots, n-1, V = \bigcup V_k$, and let

$$\Omega' = \mathbb{C} \smallsetminus (\bar{U} \cup \bar{V}), \quad \Omega = f^{-1}(\Omega').$$
Then $\Omega \subset \Omega'$, $\Omega \neq \Omega'$, and $f : \Omega \to \Omega'$ is a covering map. By Corollary 1.75, $\|Df(z)\| > 1$ for any $z \in \Omega$, in the hyperbolic metric of Ω' . Since J(f) is compactly contained in Ω' , there exists $\lambda > 1$ such that $\|Df(z)\| \geq \lambda$, $z \in \Omega$, so f is expanding on J(f) with respect to the hyperbolic metric. Since the hyperbolic metric and the Euclidean metrics over the Julia set $J(f) \Subset \Omega$ are equivalent, f is expanding with respect to the latter as well.

We see that a quadratic polynomial f is hyperbolic in the sense of §22.2 if and only if its Julia set J(f) is hyperbolic – so, the terminology is consistent.

23.2. Local connectivity of the Julia set.

PROPOSITION 4.54. Let f be a polynomial. If the Julia set J(f) is connected then it is locally connected. Moreover, the uniformization $B^{-1}: \mathbb{C} \setminus \overline{\mathbb{D}} \to \mathbb{C} \setminus K(f)$ admits a Hölder continuous extension to the boundary.

PROOF. Of course, the second assertion implies the first. And for the second, it is enough to show that the uniformization itself is Hölder continuous.

Recall that B conjugates $f \equiv f_c$ on $\mathbb{C} \smallsetminus K(f)$ to $f_0 : z \mapsto z^2$ on $\mathbb{C} \smallsetminus \overline{\mathbb{D}}$. Let us use the diadic grids $\Delta_{\overline{i}}^n$ and $D_{\overline{i}}^n$ for f_0 and f respectively, where $B(D_{\overline{i}}^n) = \Delta_{\overline{i}}^n$ (see §32.4). Then

diam
$$\Delta_{\overline{i}}^n \simeq 2^{-n}$$
, diam $D_{\overline{i}}^n = O(\lambda^{-n})$,

where λ is the expanding factor for f on the Julia set, so

diam
$$D_{\overline{i}}^n = O((\operatorname{diam} \Delta_{\overline{i}}^n)^{\alpha})$$
 with $\alpha = \frac{\log \lambda}{\log 2}$.

Take now any two points $z, z' \in \mathbb{C} \setminus \overline{\mathbb{D}}$ on distance $\epsilon > 0$ apart. Assume first they are $\leq 2\epsilon$ -close to \mathbb{T} . Then they fit into the union of two adjacent diadic boxes $\Delta_{\overline{i}}^{n}$ with $2^{-(n+2)} \leq \epsilon < 2^{-(n+1)}$. Ther images, $\zeta = B^{-1}z$ and $\zeta' = B^{-1}z'$, fit into the union of two corresponding boxes $D_{\overline{i}}^{n}$ of size $O(\lambda^{-n})$, so $\operatorname{dist}(\zeta, \zeta') = O(\epsilon^{\alpha})$.

Assume now that one of the points E_i of the 2ϵ -away from \mathbb{T} , so both of them stay distance $\delta \geq \epsilon$ -away from \mathbb{T} .⁶ Let $\delta \asymp 2^{-k}$, $k \leq n$. Then $f_0^k(z)$ and $f_0^k(z')$ are two points on a distance $\asymp 2^{-(n-k)}$ apart and on a distance of order 1 from \mathbb{T} (which of course, follows from the Koebe Distorion Theorem, but is also eaily seen by noting that f_0 is just the doubling map in the logarithmic coordinate). Since B^{-1} is bi-Lipschitz on any compact subset of $\mathbb{C} \smallsetminus \overline{\mathbb{D}}$, the same is true for the points $B^{-1}(f_0^n z) = f^n \zeta$ and $B^{-1}(f_0^n z') = f^n \zeta'$. But then

dist
$$(\zeta, \zeta') = O(2^{-(n-k)}\lambda^{-k}) = O(2^{-\alpha n})$$
 with the same $\alpha = \frac{\log \lambda}{\log 2}$.

23.3. Blaschke model for the immediate basin. We can now refer to general properties of lc hulls (Proposition 1.122) to conclude:

COROLLARY 4.55. Let f be a hyperbolic quadratic polynomial, and let D_i be the components on int K (arbitrary ordered). Then any D_i is a Jordan disk, and diam $D_i \rightarrow 0$.

 $^{^{6}}$ To conclude that B^{-1} is just continuous up to the boundary, we do not actually need this consideration.

In particular, the immediate basin $D_0 \ni 0$ of α_0 is a Jordan disk. Let us uniformaze it by the unit disk, $\phi : (D_0, \alpha_0) \to (\mathbb{D}, 0)$. By the Conformal Schönflis Theorem, ϕ extends to a homeomorphism $\phi : \overline{D}_0 \to \overline{\mathbb{D}}$. Let

$$g = \phi \circ f^p \circ \phi^{-1} : \overline{\mathbb{D}} \to \overline{\mathbb{D}}$$

PROPOSITION 4.56. If the uniformization ϕ is appropriately normalized, then

$$g(z) = z \frac{z + \sigma}{1 + \bar{\sigma}z},$$

where $\sigma \in \mathbb{D}$ is the multiplier of the attracting cycle α .

PROOF. Consider first an arbitrary uniformization $\phi : (D_0, \alpha_0) \to (\mathbb{D}, 0)$. The map $g : \overline{\mathbb{D}} \to \overline{\mathbb{D}}$ is a double branched covering of the disk fixing 0 and preserving the unit circle \mathbb{T} . A general form of such a map is

$$g(z) = \lambda z \frac{z-a}{1-\bar{a}z}$$
, with $|\lambda| = 1$.

Replacing ϕ with $\lambda \phi$ results in replacing g(z) with $\lambda g(z/\lambda)$ killing the coefficient λ in front of the Blaschke product.

Since the multiplier is invariant under conformal conjugacies, we have:

$$\sigma = g'(0) = -a.$$

REMARK 4.4. In fact, we did not need to know that D_0 is a Jordan disk to concude that g is the above Blaschke product. It would follow from the property that $g: \mathbb{D} \to \mathbb{D}$ is proper.

23.4. Hubbard trees. We will now attach to any superattracting quadratic polynomial a combinatorial object called a Hubbard tree. It will be eventually shown that such polynomials (and there are countable many of them) are fully classified by their Hubbard trees.

23.4.1. Definition and first properties. Let $f = f_c$ be a superattracting quadratic polynomial, so its critical point 0 is periodic with some period $p \in \mathbb{Z}_+$. Let $c_k = f^k(0)$, $\mathbf{c} = \{c_k\}$, and let D_k be the component of int K(f) containing c_k . By the last part of Theorem 4.53, int K(f) coincides with the basin of $\boldsymbol{\alpha}$, so for any component D of int K(f) there exists a unique $n = n(D) \in \mathbb{N}$ such that f^n univalently maps D onto D_0 . Let us mark in D the preimage of 0 under this map. This allows us to consider *legal paths* in K(f), see §??.

Since K(f) is locally connected, any two points z, ζ in it can be connected by a unique legal path $[z, \zeta]$. Let us consider the legal hull $H = H_f$ of the points c_k . When 0 is a fixed point (i.e., p = 1), then $H = \{0\}$ is trivial, so in what follows we will assume that p > 1. Then H is a topological tree called the Hubbard tree. Let us mark on H the points c_k and all the branch points b_j , and let $\mathbf{b} = \{b_j\}$.

PROPOSITION 4.57. We have:

(i) The marked Hubbard tree $(H, \alpha \cup \beta)$ is unvariant under f; hence all branch points of H are (pre)periodic;

(ii) The critical value c is a vertex of H; the critical point 0 is not a branch point of H;

PROOF. (i) First note that the image $f(\gamma)$ of any legal path $\gamma \subset K(f)$ is a legal curve since internal rays go to internal rays under the dynamics. However, $f(\gamma)$ is not necessarily a path since it can "backtrack" if int γ passes through the critical point. But otherwise, $f(\gamma)$ is a legal path.

Let us take the legal path $\gamma_k = [0, c_k] \subset H$ connecting 0 to any other c_k , $k = 1, \ldots, p-1$. Since int $\gamma_k \not\supseteq 0$, $f(\gamma_k)$ is a legal path connecting c_1 to c_{k+1} (where the index is taken mod p). Since H is "legally convex", $f(\gamma_k) \subset H$. Since $H = \cup \gamma_k$, the first assertion follows.

(ii) Since $H = \bigcup [c_k, c_j]$, all the vertices of H are containes in the cycle **c**. So, one of the points c_k must be a vertex. But if c_k with $k \neq 0 \mod p$ is not a vertex then c_{k+1} is not either, since the map f | H near any non-critical point is a local diffemorphism onto the image. Thus, if $c \equiv c_1$ is not a vertex then non of the c_k are – contradiction.

REMARK 4.5. This argument shows that there is an $l \in [1, p]$ such that the vertices of H are exactly the points $c_k, k = 1, \ldots, l$.

PROPOSITION 4.58. Dividing H by the marked points into intervals J_k we obtain a recurrent Markov chain.

23.5. Dynamical quasi-self-similarity. We will now show that hyperbolic Julia sets in small scales look roughly the same as they do in moderate scales:

LEMMA 4.59. Let f be a hyperbolic quadratic map. Then there is an $\epsilon_0 > 0$ such that for any $\epsilon \in (0, \epsilon_0)$, $z \in J(f)$ and $\rho \in (0, \epsilon)$ there is n such that f^n univalently and with bounded distortion maps the disk $\mathbb{D}(z, \rho)$ onto an oval of radius of order ϵ and of bounded shape around $f^n z$ (where the distortion and shape bounds are absolute, while the other constants depend on f).

PROOF. Since f is hyperbolic, the postcritical set stays away from the Julia set. Let $\epsilon_0 > 0$ be the distance between these two sets.

Fix some $z \in J(f)$, and let $z_n = f^n(z)$. Then there is a univalent branch f^{-n} on the disk $\mathbb{D}(z_n, \epsilon_0)$ such that $f^{-n}(z_n) = z$. By the Koebe Distortion Theorem, restriction of this branch to the half-disk $D(z_n, \epsilon_0/2)$ is a univalent map with an absolutely bounded distortion. Hence the pullback $U_n := f^{-n}(D(z_n, \epsilon_0/2))$ is an oval centered at z of bounded shape. The size of this oval is comparable with $\rho_n = |Df^n(z)|^{-1}$. Since $\rho_n \to 0$ and $\rho_{n+1} \simeq \rho_n$, the conclusion easily follows. \Box

23.6. Porousity and area of the Julia set. Let us say that a compact set $J \subset \mathbb{C}$ is *porous* (in all scales) if there is a $\kappa \in (0, 1)$ such that any disk $\mathbb{D}(z, \rho)$ with $z \in J$ contains a disk $\mathbb{D}(\zeta, \kappa \rho) \subset \mathbb{D}(z, \rho) \smallsetminus J$. Informally speaking, J has definite gaps in all scales.

By the Lebesgue Density Points Theorem, porous sets have zero area.

Note that all nowhere dense compact sets are porous in moderate scales:

EXERCISE 4.60. Let J be a nowhere dense compact subset of \mathbb{C} . Then for any $\epsilon > 0$ there is a $\kappa = \kappa(\epsilon) \in (0,1)$ such that any disk $\mathbb{D}(z,r)$ with $z \in J$ and $r \geq \epsilon$ contains a gap $\mathbb{D}(\zeta, \kappa r) \subset \mathbb{D}(z, r) \setminus J$.

PROPOSITION 4.61. Any hyperbolic Julia set J(f) is porous and hence it has zero area.

PROOF. By Lemma 4.59, any small scale disk $\mathbb{D}(z, \rho)$ with $z \in J$ can be mapped with bounded distortion onto a moderate scale oval of bounded shape. Since the latter contains a definite gap, the former contains one as well.

23.7. Dynamical qc removability of J.

LEMMA 4.62. Let f and \tilde{f} be two quadratic polynomials, let U and \tilde{U} be neighborhoods of their Julia sets, and let $h : (U, J) \to (\tilde{U}, \tilde{J})$ be a homeomorphism conjugating f to \tilde{f} near the Julia sets such that h is qc on $U \smallsetminus J$. Then h is qc.

PROOF. We will use definition of quasiconformality in terms of the circular dilatation, see Proposition 2.25. It is enough to check that the image h(D) of a sufficiently small disk $D := \mathbb{D}(z, \rho), z \in J$, has a bounded shape around h(z). To this end, we will make use of the quasi-self-similarity of J and \tilde{J} (Lemma 4.59). According to that lemma, for all sufficiently small $\epsilon > 0$ (how small is independent of z and ρ) there exists an n (depending on z and ρ) such that f^n maps D univalently onto an oval V of size of order ϵ and bounded shape around $z_n = f^n z$. Since h is a homeomorphism, h(V) is an oval whose inner and outer radii (around $h(z_n)$) are squeezed in between $r(\epsilon) > 0$ and $R(\epsilon) \to 0$ as $\epsilon \to 0$. If $R(\epsilon)$ is sufficiently small then there exists a inverse branch \tilde{f}^{-n} on h(V) with bounded Koebe distortion such that $\tilde{f}^{-n}(h(z_n)) = h(z)$. Hence $\tilde{f}^{-n}(V) = h(D)$ has a bounded shape around h(z).

23.8. Rigidity. In this section we will prove that the superattracting parameter value c is uniquely determined by its Hubbard map. It is the first occasion of the Rigidity Phenomenon which is a central theme of this book.

A reference to a "Hubbard map $F: T \to T$ " will mean an abstract Hubbard tree with the marked points and piecewise affine dynamics that models $f_c: T_c \to T_c$.

RIGIDITY THEOREM FOR SUPERATTRACTING MAPS. Two superattracting parameters c and \tilde{c} with the same Hubbard map $F: T \to T$ must coincide: $c = \tilde{c}$.

We let $f = f_c$, $\tilde{f} = f_{\tilde{c}}$; $K = K_c$, $\tilde{K} = K_{\tilde{c}}$, etc.

As we know from ??, two superattracting maps f and \tilde{f} with the same Hubbard map are topologically conjugate by a homeomorphism $h : (\mathbb{C} K) \to (\mathbb{C}, \tilde{K})$ which is conformal on the basin of infinity, $\mathbb{C} \setminus K$. If we showed that h is actually conformal on the whole plane (and hence is affine), we would be done, since different quadratic maps f_c are not affinely equivalent. We will do it in two steps:

Step 1. The map h is conformal on int K_c , and hence it is conformal outside the Julia set J.

Step 2. The map $h : \mathbb{C} \to \mathbb{C}$ is quasiconformal.

Since the Julia set J has zero area (Theorem ??), Weyl's Lemma assures that the map h is, indeed, conformal on the whole plane.

So, let us go through the above two steps.

Proof of Step 1. Let D_0 be the immediate basin of 0. We know that it is a Jordan disk (Corollary 4.55), so the Riemann mapping $\phi : (D_0, 0) \to (\mathbb{D}, 0)$ extends to a homeomorphism $\operatorname{cl} D_0 \to \overline{\mathbb{D}}$ (denoted by ϕ as well). Moreover, ϕ conjugates the return map $f^p : \operatorname{cl} D_0 \to \operatorname{cl} D_0$ to $z \mapsto e^{i\theta} z^2$ on $\overline{\mathbb{D}}$. Such a map has a unique fixed point on the boundary, so f^p has a unique fixed point γ on ∂D_0 . Let us normalize ϕ so that $\phi(\gamma) = 1$. Then $\theta = 0$, so ϕ conjugates $f^p | \operatorname{cl} D_0$ to $f_0 : z \mapsto z^2$ on $\overline{\mathbb{D}}$.

Similarly, the normalizaed Riemann mapping $\tilde{\phi} : \operatorname{cl} D_0 \to \overline{\mathbb{D}}$ conjugates $\tilde{f}^p | \operatorname{cl} \tilde{D}_0$ to $z \mapsto z^2$ on $\overline{\mathbb{D}}$. Hence the composition $h_0 = \tilde{\phi}^{-1} \circ \phi : \operatorname{cl} D_0 \to \operatorname{cl} \tilde{D}_0$ conjugates $f^p | \operatorname{cl} D_0$ to $\tilde{f}^p | \operatorname{cl} \tilde{D}_0$.

We claim that this map h_0 continuously matches on ∂D_0 with the conjugacy $h: \mathbb{C} \setminus \operatorname{int} K \to \mathbb{C} \setminus \operatorname{int} \tilde{K}$. Indeed, both of them conjugate $f^p | \partial D_0$ to $\tilde{f}^p | \partial \tilde{D}_0$. Hence the composition $h^{-1} \circ h_0 : \partial D_0 \to \partial D_0$ commutes with with $f^p | \partial D_0$. But as we know, the latter map is topologically equivalent to $z \mapsto z^2$ on \mathbb{T} , which has the trivial commutator (Proposition 4.6). Hence $h^{-1} \circ h_0 | \partial D_0 = \operatorname{id}$, and the claim follows.

Let us now consider another component D of int K. Since int K is equal to the basin of **c** (Theorem 4.53), there is $n = n_D \in \mathbb{Z}_+$ such that f^n homeomorphically maps $\operatorname{cl} D$ onto $\operatorname{cl} D_0$. Let $f^{-n} : \operatorname{cl} D \to \operatorname{cl} D_0$ stand for the inverse map. Then we let

(23.1)
$$h_D = f^n \circ h_0 \circ f^{-n} : \operatorname{cl} D \to \operatorname{cl} D.$$

Obviously, this map conjugates $f^p | \operatorname{cl} D$ to $\tilde{f}^p | \operatorname{cl} \tilde{D}$.

Moreover, h_D matches continuously on ∂D with h. Indeed, since h is a conjugacy on the whole Julia set, we have

$$h|\partial D = \tilde{f}^n \circ (h|\partial D_0) \circ f^{-n} : \partial D \to \partial \tilde{D}.$$

Comparing this with (30.2), taking into account that $h | \partial D_0 = h_0$, yields $h | D = h_D$.

Thus, we have externded h conformally and equivariantly to all the components D_i of int K. Since diam $D_i \to 0$, this extension is a global homeomorphism (??), and Step 1 is accomplished.

Proof of Step 2.

24. Parabolic maps

25. Misiurewicz maps

26. Quasiconformal deformations

26.1. Idea of the method.

26.1.1. Pullbacks. Consider a K-quasi-regular branched covering $f: S \to S'$ between Riemann surfaces (see §11.4). Then any conformal structure μ on S' can be pulled back to a structure $\nu = f^*(\mu)$ on S. Indeed, quasi-regular maps are differentiable a.e. on S with non-degenerate derivative so that we can let $\nu(z) = (Df(z)^{-1})_*(\mu)$ for a.e. $z \in S$. This structure has a bounded dilatation:

$$\frac{\|\nu\|_{\infty} + 1}{\|\nu\|_{\infty} - 1} \le K \frac{\|\mu\|_{\infty} + 1}{\|\mu\|_{\infty} - 1}.$$

If f is holomorphic then in any conformal local charts near z and f(z) we have:

$$f^*\mu(z) = \frac{f'(z)}{f'(z)}\mu(fz)$$

(since the critical points of f are isolated, this expression makes sence a.e.). An obvious (either from this formula or geometrically) but crucial remark is that holomorphic pull-backs preserve dilatation of conformal structures.

26.1.2. Qc surgeries and deformations. Consider now a qr map $f : \mathbb{C} \to \mathbb{C}$ preserving some conformal structure μ on $\hat{\mathbb{C}}$. By the Measurable Riemann Mapping Theorem, there is a qc homeomorphism $h_{\mu} : \mathbb{C} \to \mathbb{C}$ such that $(h_{\mu})_*(\mu) = \sigma$. Then $f_{\mu} = h_{\mu} \circ f \circ h_{\mu}^{-1}$ is a quasi-regular map preserving the standard structure σ on $\hat{\mathbb{C}}$. By Weil's Lemma, f_{μ} is holomorphic outside its critical points. Since the isolated singularities are removable, f_{μ} is holomorphic everywhere, so that it is a rational endormorphism of the Riemann sphere. Of course, $\deg(f_{\mu}) = \deg(f)$. Since h_{μ} is unique up to post-composition with a Möbius map, $f = f_{\mu}$ is uniquely determined by μ up to conjugacy by a Möbius map.

Thus, a qc invariant view of a rational map of the Riemann sphere is a quasiregular endomorphism $f : (S^2, \mu) \to (S^2, \mu)$ of a qc sphere S^2 which preserves some conformal structure μ . This provides us with a powerful tool of holomorphic dynamics: the method of qc surgery. The recepie is to cook by hands a quasi-regular endomorphism of a qc sphere with desired dynamical properties. If it admits an invariant conformal structure, then it can be realized as a rational endomorphism of the Riemann sphere.

It may happen that f itself is a rational map preserving a non-trivial conformal structure μ . Then f_{μ} is called a qc deformation of f. If f is polynomial, then let us normalize h_{μ} so that it fixes ∞ . Then $f_{\mu}^{-1}(\infty) = \infty$ and hence the deformation f_{μ} is polynomial as well. If $f: z \mapsto z^2 + c$ is quadratic then let us additionally make h_{μ} fix 0. Then 0 is a critical point of f_{μ} , so that

(26.1)
$$f_{\mu}(z) = t(\mu)z^2 + b(\mu), \quad t \in \mathbb{C}^*.$$

Composing h_{μ} with complex scaling $z \mapsto t(\mu)z$, we turn this quadratic polynomial to the normal form $z \mapsto z^2 + c(\mu)$.

26.1.3. Holomorphic dependence. Assume now that $\mu = \mu_{\lambda}$ depends holomorphically on parameter λ . By Theorem 2.41, the map $h_{\lambda} \equiv h_{\mu(\lambda)}$ is also holomorphic in λ . However, the inverse map h_{λ}^{-1} is not necessarily holomorphic in λ .

EXERCISE 4.63. Give an example.

It is a miracle that despite it, the deformation $f_{\lambda} \equiv f_{\mu(\lambda)}$ is still holomorphic in λ !

LEMMA 4.64. Let $f_{\lambda} = h_{\lambda} \circ f \circ h_{\lambda}^{-1}$, where f and f_{λ} are holomorphic functions and h_{λ} is a holomorphic motion (of an appropriate domain). Then f_{λ} holomorphically depends on λ .

PROOF. Taking $\partial_{\bar{\lambda}}$ -derivative of the expression $f_{\lambda} \circ h_{\lambda} = h_{\lambda} \circ f_0$, we obtain:

$$0 = \partial_{\bar{\lambda}} h_{\lambda} \circ f_0 = f_{\lambda}' \circ \partial_{\bar{\lambda}} h_{\lambda} + \partial_{\bar{\lambda}} f_{\lambda} \circ h_{\lambda} = \partial_{\bar{\lambda}} f_{\lambda} \circ h_{\lambda}.$$

COROLLARY 4.65. Consider a quadratic map $f : z \mapsto z^2 + c_0$. Let μ_{λ} be a holomorphic family of f-invariant Beltrami differentials on \mathbb{C} . Normalize the solution $h_{\lambda} : \mathbb{C} \to \mathbb{C}$ of the corresponding Beltrami equiation so that the qc deformation $f_{\lambda} = h_{\lambda} \circ f \circ h_{\lambda}^{-1}$ has a normal form $f_{\lambda} : z \mapsto z^2 + c(\lambda)$. Then the parameter $c(\lambda)$ depends holomorphically on λ .

PROOF. Consider first the solution $H_{\lambda} : \mathbb{C} \to \mathbb{C}$ of the Beltrami equation which fixes 0 and 1. It conjugates f to a quadratic polynomial of form (26.1).

By Lemma 4.64, its coefficients $t(\lambda)$ and $b(\lambda)$ depend holomorphically on λ . The complex rescaling $T_{\lambda} : z \mapsto t(\lambda)z$ reduces this polynomial to the normal form with $c(\lambda) = t(\lambda)b(\lambda)$, and we see that $c(\lambda)$ depends holomorphically on λ as well. \Box

26.1.4. Invariant extensions of conformal structures. In applications, we usually start with an invariant conformal structure on a smaller Riemann surface and extend it to an invariant conformal structure on an embient one. It can be done under very general circumstances.

Let S be a Riemann surfaces endowed with a holomorphic equivalence relation \mathcal{R} , and let U be an open subset of S. Let \tilde{U} stand for the \mathcal{R} -saturation of U (see §20.8). For all practical purposes, the reader can think of the grand orbit equivalence relation for a holomorphic map f, so \tilde{U} is just the grand orbit of U.

LEMMA 4.66. Any \mathcal{R} -invariant conformal structures μ on $U \cup (S \setminus \tilde{U})$ admits a unique \mathcal{R} -invariant extension $\tilde{\mu}$ to S. Moreover $\text{Dil } \mu = \text{Dil } \nu$.

In particular, if U is a fundamental domain for U, then any conformal structures μ on U admits a unique \mathcal{R} -invariant extension $\tilde{\mu}$ to S, and $\text{Dil}\,\mu = \text{Dil}\,\nu$.

PROOF. Since the set of critical points of \mathcal{R} is at most countable, while the desired conformal structure has to be only measurable, we do not need to define it at the critical equivalence classes.

Let $\zeta_{o} \in \tilde{U}$ be a point in a regular equivalence class. By definition of the saturation \tilde{U} , it has an \mathcal{R} -equivalent point $z_{o} \in U$, hence there exists a local section ϕ of \mathcal{R} such that $\phi(z_{o}, \zeta_{o}) = 0$. Since z_{o} is regular, we can locally express z as $\psi(\zeta)$ with a holomorphic ψ , and let $\tilde{\mu} = \psi^{*}(\mu)$ near ζ . This definition is independent of the choice of the local section ϕ since $\mu | U$ is \mathcal{R} -invariant.

COROLLARY 4.67. Any \mathcal{R} -invariant conformal structure μ on U admits a unique \mathcal{R} -invariant extension $\tilde{\mu}$ to S such that μ coincides with the standard structure σ on $S \setminus \tilde{U}$. Moreover, $\operatorname{Dil} \tilde{\mu} = \operatorname{Dil} \mu$.

In particular, if U is a fundamental domain for \tilde{U} , then any conformal structure μ on U admits a unique \mathcal{R} -invariant extension $\tilde{\mu}$ to S such that $\mu = \sigma$ on $S \setminus \tilde{U}$, and Dil $\tilde{\mu} = \text{Dil } \mu$.

We will refer to the extension given in this Corollary as *canonical*.

COROLLARY 4.68. Let $X \subset \hat{\mathbb{C}}$ be a wandering measurable set for a rational map $f : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ such that all the iterates $f^n | X, n \in \mathbb{N}$, are injective. Then any conformal structure μ on X admits the canonical f-invariant extension $\tilde{\mu}$ to the whole sphere, and $\operatorname{Dil} \tilde{\mu} = \operatorname{Dil} \mu$.

26.2. Sullivan's No Wandering Domains Theorem. Consistently with the general terminology of §20.1, a component D of the Fatou set F(f) is called wandering if $f^n D \cap f^m D = \emptyset$ for any natural n < m. Such components will also be referred to as "wandering domains".⁷

THEOREM 4.69. A quadratic polynomial f has no wandering domains.

⁷One can consider more general "wandering domains", not necessarily full components of F(f). Such domains can certainly exist (in the basins of attracting and parabolic points). We hope this slight terminological inconsistency will not cause a problem.

The rest of the section will be devoted to the proof of this theorem. The idea is to endow a wadering domain D with a 3-parameter family of conformal structures $\mu_{\lambda}, \lambda \in \mathbb{R}^3$, then to promote it to a family of f-invariant conformal structures on the whole Riemann sphere $\hat{\mathbb{C}}$, and to consider the corresponding qc deformation f_{λ} of f. With some care this deformation can be made *efficient*, i.e., the map $\lambda \mapsto f_{\lambda}$ can be made injective. But this is certainly impossible since a 3D parameter domain cannot be embedded into \mathbb{C} .

Let us now supply the details. Since D is wandering, only one domain $D_n = f^n D$, $n \in \mathbb{N}$, can contain the critical point 0. By replacing D with $f^{n+1}D$, we can eliminate this possibility.

So, assume orb D does not contain 0. Then all the maps $f: D_n \to D_{n+1}$ are conformal isomorphisms (being unbranched coverings over simply connected domains, see Exercise 4.12). Hence D is a fundamental domain for its saturation Orb D by the grand orbit equivalence relation.

Let us now consider an arbitrary conformal structure μ_0 on D (as always, μ_0 is assumed to be measurable with bounded dilatation). By Corollary 4.68, μ_0 canonically extends to an invariant conformal structure μ on the whole sphere $\hat{\mathbb{C}}$, and moreover $\text{Dil } \mu = \text{Dil } \mu_0$.

EXERCISE 4.70. Work out details of this canonical extension.

By the Measurable Riemann Mapping Theorem, there exists a qc map h_{μ} : $\hat{\mathbb{C}} \to \hat{\mathbb{C}}$ such that $\mu = h_{\mu}^* \sigma$. Let

$$f_{\mu} = h_{\mu} \circ f \circ h_{\mu}^{-1} : \hat{\mathbb{C}} \to \hat{\mathbb{C}}.$$

By Corollary 4.65, h_{μ} can be normalized so that $f_{\mu} : z \mapsto z^2 + c_{\mu}$ is a quadratic polynomial holomorphically depending on μ .

We will now make a special choice of a 3-parameter family $\mu = \mu_{\lambda}$ of the initial conformal structures on D to ensure that the qc deformation f_{λ} is efficient. Namely, we let $\mu_{\lambda} = (\psi_{\lambda})_* \sigma$, where $\psi_{\lambda} : D \to D$ is a smooth 3-parameter family of diffeomorphisms that extend to the ideal boundary $\partial^i D$, and the family $\lambda \mapsto \psi_{\lambda}$ is efficient in $\operatorname{Aut}(D) \setminus \operatorname{Diff}_+(\partial^i D)$.

EXERCISE 4.71. Construct such a family of diffeomorphisms.

Since dimension of the parameter space is bigger than 2, by the Implicit Function Theorem, there exists a one-parameter family of conformal structures μ_t (within our 3-parameter family) such that $c_t \equiv \text{const.}$ Let us take a base point τ in this family. Then $f_t = f_{\tau}$ for all t, and hence the homeomrphisms $H_t = h_t \circ h_{\tau}^{-1}$ commute with f_{τ} .

EXERCISE 4.72. Let $H_t : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ be a one-parameter family of homeomorphisms commuting with a quadratic polynomial f such that $f_{\tau} = \text{id}$ for some parameter τ . Then $H_t | J(f) = \text{id}$ for all t.

Since $\partial D \subset J(f_{\tau})$ (here ∂D is the ordinary boundary, not the ideal one), we conclude that $H_t | \partial D = \text{id}$ for all t. Hence $h_t | \partial D = h_{\tau} | \partial D$ and, in partial, $h_t(D) = h_{\tau}(D) := \Delta$.

⁸In other words, $\tilde{\psi}_{\lambda} = h^{-1} \circ \psi_{\lambda} \circ h$ where $\phi : \mathbb{D} \to D$ is the Riemann mapping and $\tilde{\psi}_{\lambda}$ is a family diffeomnorphisms $\mathbb{D} \to \mathbb{D}$ such that $\psi_t | \mathbb{T} \neq M \circ \psi_{\lambda} | \mathbb{T}$ for any $t \neq \lambda$, $M \in \text{PSL}^{\#}(2, \mathbb{R})$.

Note now that since $(\psi_t)_* \sigma = \mu_t | D = (h_t)^* \sigma$, the map $h_t \circ \psi_t : D \to \Delta$ is conformal. Hence the map

$$\psi_t^{-1} \circ h_t^{-1} \circ h_\tau \circ \psi_\tau : D \to D$$

is a conformal automorphism of D. But on ∂D it coincides with $\psi_t^{-1} \circ \psi_\tau$. By Exercise 1.114, these two maps have the same extension to the ideal boundary $\partial^i D$ contradicting the efficiency of ψ_t .

The theorem is proved.

26.3. Complete picture of the dynamics on the Fatou set. Putting together the No Wandering Domains Theorem and Theorem 4.51, we obtain:

THEOREM 4.73. For any point $z \in F(f)$, the orb z either converges to an attracting or parabolic cycle, or else lands in a Siegel disk.

27. Quadratic-like maps: first glance

27.1. The concept.

27.1.1. *Definition and first properties.* The notion of a quadratic-like map is a fruitful generalization of the notion of a quadratic polynomial.

DEFINITION 4.74. A quadratic-like map $f: U \to U'$ (abbreviated as "q-l map") is a holomorphic double branched covering between two conformal disks U and U' in \mathbb{C} such that $U \in U'$.

By the Riemann-Hurwitz Theorem, any quadratic-like map has a single critical point, which is of course non-degenerate. We normalize f so that the critical point sits at 0 (unless otherwise is explicitly stated). Note that any quadratic polynomial $f = f_c$ restricts to a quadratic-like map $f : f^{-1}(\mathbb{D}_R) \to \mathbb{D}_R$ whose range is a round disk of radius R > |f(0)|. More canonically, for any $r > |B_f(0)|$ (recall that B_f is the Böttcher function for f), the restriction of f to the subpotential domain $\Omega_f(r)$ (see §32.2) provides us with a quadratic-like map $f : \Omega_f(r) \to \Omega_f(r^2)$.

From now on (unless otherwise is explicitly stated) we will make the following *Technical Conventions*: For any quadratic-like map $f: U \to U'$, we assume that the domains U and U' are 0-symmetric and that f is even, i.e., f(z) = f(-z) for all $z \in U$. Moreover, we assume that both domains are quasidisks.

Note that the last assumption can be secured by the following *adjustment* of f:

EXERCISE 4.75. Take any 0-symmetric topological disk $V' \ni f(0)$ such that $U \subset V' \subset U'$, and let $V = f^{-1}(V')$. Then the map $f : V \to V'$ is quadratic-like. (Of course, V' can be chosen so that its boundary is real analytic.)

Sometime we will refer to a q-l map satisfying the above Technical Conventions as *conventional*. Such a map f extends continuously to \overline{U} , so we can assume this without loss of generality.

The annulus $A = \overline{U}' \setminus U$ is called the *fundamental annulus* of f. (We will refer in the same way to the corresponding open and semi-open annuli.)

The notion of quadratic-like map does not fit to the canonical dynamical framework, where the phase space is assumed to be invariant under the dynamics. In the quadratic-like case, some orbits *escape* through the fundamental annulus (i.e., $f^n z \in A$ for some $n \in \mathbb{N}$), and we cannot iterate them any further. However, there are still a plenty of non-escaping points, which form a dynamically significant

object. The set of all non-escaping points is called the *filled Julia set* of f and is denoted in the same way as for polynomials:

$$K(f) = \{ z : f^n z \in U, \ n = 0, 1, \dots \}$$

By definition, the Julia set of f is the boundary of the filled Julia set: $J(f) = \partial K(f)$. Dynamical features of quadratic-like maps are very similar to those of quadratic maps (in §49.2 we will see a good reason for it):

EXERCISE 4.76. Check that all dynamical properties of quadratic polynomials established in in §§21 - 22 are still valid for quadratic-like maps. In particular,

(i) The filled Julia set K(f) is a completely invariant full compact subset of U.

(ii) Basic dichotomy: J(f) and K(f) are either connected or Cantor; the former holds if and only if the critical point is non-escaping: $0 \in K(f)$.

(iii) Any periodic component of int K(f) is either in the immediate basin of an attracting/parabolic cycle, or is a Siegel disk.

(iv) f can have at most one attracting cycle.

EXERCISE 4.77. Show that adjustments from Exercise 4.75 do not change the filled Julia set.

Let us consider a quadratic-like map $f: U \to U'$ with real symmetric domains U and U'. Since these domains are simply connected, their real slices

$$I := U \cap \mathbb{R} \quad and \quad I' := U' \cap \mathbb{R}$$

are open intervals; moreover, $I \in I'$. If additionally, the map f is real, i.e. $f(I) \subset I'$, then it is naturally referred to as a *real-symmetric* (or just *real*) quadratic-like map. Note also that according to our Conventions, f extends continuously to ∂U , in particular to ∂I , and we have $f(\partial I) \subset \partial I'$.

For a real-symmetric quadratic-like map f, we let $K^{\mathbb{R}} \equiv K^{\mathbb{R}}(f) := K(f) \cap \mathbb{R}$ be the real slice of its filled Julia set.

EXERCISE 4.78. For a real-symmetric quadratic-like map $f : U \to U'$ with connected Julia set K(f), the real slice $K^{\mathbb{R}}(f)$ is a closed interval compactly contained in I. Moreover, the restriction $f : K^{\mathbb{R}} \to K^{\mathbb{R}}$ is a proper unimodal map (see §21.5.2), and its boundary fixed point $\beta \in \partial K^{\mathbb{R}}$ is either repelling or parabolic with positive multiplier: $f'(\beta) \geq 1$.

27.2. Uniqueness of a non-repelling cycle. We will now give the first illustration of how useful the notion of a quadratic-like map is. It exploits the flexibility of this class of maps: small perturbations of a quadratic-like map are still quadratic-like (on a slightly adjusted domain):

EXERCISE 4.79 (compare Exercise 4.75). Let $f: U \to U'$ be a quadratic-like map with the fundamental annulus A. Take a 0-symmetric smooth Jordan curve $\gamma' \subset A$ generating $H_1(A)$, and let V' be the domain bounded by γ' . Let ϕ be a bounded holomorphic function on U with $\|\phi\|_{\infty} < \operatorname{dist}(\gamma, \partial U')$. Let $g = f + \phi$ and $V = g^{-1}V'$. Then $g: V \to V'$ is a quadratic-like map.

THEOREM 4.80. Any quadratic-like map (in particular, any quadratic polynomial) can have at most one non-repelling cycle.

PROOF. Assume that a quadratic-like map $f: U \to U'$ has two non-repelling cycles $\boldsymbol{\alpha} = \{\alpha_k\}_{k=0}^{p-1}$ and $\boldsymbol{\beta} = \{\beta_k\}_{k=0}^{q-1}$. Let μ and ν be their multipliers. Take two numbers a and b to be specified below.

Using the Interpolation formulas, find a polynomial ϕ (of degree 2p + 2q - 1) vanishing at points α_k and β_k , such that $\phi'(\alpha_0) = a$, $\phi'(\beta_0) = b$, while the derivatives of ϕ at all other points α_k and β_k (k > 0) vanish.

Let $f_{\epsilon} = f + \epsilon \phi$, where $\epsilon > 0$. Then α and β are periodic cycles for f_{ϵ} with multipliers

$$\lambda_{\epsilon} = \lambda + a\epsilon \prod_{k>0} f'(\alpha_k) \text{ and } \mu_{\epsilon} = \mu + b\epsilon \prod_{k>0} f'(\beta_k)$$

respectively. Since $|\lambda| \leq 1$ and $|\mu| \leq 1$, parameters a and b can be obviously selected in such a way that $|\lambda_{\epsilon}| < 1$ and $|\mu_{\epsilon}| < 1$ for all sufficiently small $\epsilon > 0$. Thus, the cycles α and β become attracting for f_{ϵ} . But for a sufficiently small ϵ , the map f_{ϵ} is quadratic-like on a slightly adjusted domain containing both cycles (see Exercise 4.79). As such, it is allowed to have at most one attracting cycle (Exercise 4.76) – contradiction.

This result together with Exercise 4.76 (iii) immediately yields:

COROLLARY 4.81. Any quadratic-like map (in particular, any quadratic polynomial) can have at most one cycle of components of int K(f).

27.3. Concept of renormalization.

27.3.1. *Complex renormalization*. The primarily motivation for introducing quadratic-like maps comes from the idea of renormalization, which is a central idea in contemporary theory of dynamical systems.

A quadratic-like map $f: U \to U'$ is called *renormalizable with period* p if there is a topological disk $V \ni 0$ such that all the domains f^iV , $i = 0, 1, \ldots, p-1$, are contained in U, the map $g := (f^p: V \to f^p(V))$ is quadratic-like with connected Julia set K(g) (see Figure ??), and the following technical "almost disjointness" property is satisfied: The images $K_i(g) := f^i(K(g)), i = 1, \ldots, p-1$, can touch $K(g) \equiv K_0(g)$ at most at one point, and this point is not a cut-point. (The meaning of these technical assumptions will become clear later, see §??.)

The sets $K_i(g)$, i = 0, 1, ..., p-1, are called the *little (filled) Julia sets*. If they are actually disjoint, then the renormalization is called *primitive*. Otherwise it is called *satellite*.

The quadratic-like map $g: V \to V'$ is called the *pre-renormalization* of f. (The *renormalization* will be defined later by allowing to adjust and to rescale the domains of g, see §49.5.)

A quadratic polynomial f_c is called renormalizable if it restricts to a renormalizable quadratic-like map.

27.3.2. Real renormalization. Let $f: I \to I$ be a proper unimodal map with non-attracting boundary fixed point $\beta \in \partial I$. Such a map is called *renormalizable* if there is a periodic interval $T \ni 0$ of period p such that the $int(f^i(T))$, $i = 0, 1, \ldots, p - 1$, are disjoint, and the return map $g = f^p : T \to T$ is a proper unimodal map with non-attracting boundary fixed point.

The following simple statement shows that for real-symmetric q-l maps, real and complex renormalizations match.

EXERCISE 4.82. Let $f: U \to U'$ be a real-symmetric quadratic-like map which is complex renormalizable so that its pre-renormalization $g: V \to V'$ is also realsymmetric. Then its restriction to the real line, $f: I \to I'$, is real renormalizable, with pre-renormalization $g: T \to T'$ where $T := V \cap \mathbb{R}$.

The inverse statement is also true except when g has a parabolic fixed point of period p, see §??.

REMARK 4.6. Consider quadratic maps f_c with -3/4 < c < 0 (see Exercise 4.25). These maps have a periodic interval T of period 2 but they are not renormalizable as the boundary fixed point of $g = f^2 : T \to T$ is attracting.

27.4. Global measure-theoretic attractor. The following result is a manifistation of the leading role of the critical point in global holomorphic dynamics:

THEOREM 4.83. For almost all $z \in K(f)$ either $f^n z \to \omega(0)$ or else orb z lands in the Siegel disk.

28. Appendix: Expanding circle maps

28.1. Definition. Recall that $\mathbb{T} \subset \mathbb{C}$ stands for the unit circle (endowed with the induced real analytic structure and Riemannian metric). Symmetry with respect to \mathbb{T} is understood in the sense of the anti-holomorphic reflection $\tau : z \mapsto 1/\overline{z}$.

Let us say that $g: \mathbb{T} \to \mathbb{T}$ is a (degree two) expanding circle map of class \mathcal{E} if it satisfies the following properties:

(i) g is an orientation preserving double covering of the circle over itself;

(ii) g is real analytic;

(iii) g is expanding, i.e, there exist constants C > 0 and $\lambda > 1$ such that for any $z \in \mathbb{T}$,

(28.1)
$$||Dg^n(z)|| \ge C\lambda^n, \quad n = 0, 1, \dots$$

The simplest example is provided by the quadratic circle map $f_0 : z \mapsto z^2$. Slightly more generally, we have the Blyaschke circle maps:

EXERCISE 4.84. Let $g : \mathbb{D} \to \mathbb{D}$ be a holomorphic double covering of the unit disk over itself which has a fixed point $\alpha \in \mathbb{D}$. Then

(i) α is attracting;

- (ii) g extends analytically to a degree two rational function.
- (iii) This function is a degree two Blyaschke product which is Möbius conjugate to the following normal form:

(28.2)
$$B_a(z) = z \frac{z-a}{1-\bar{a}z}, \quad |a| < 1;$$

(iv) $B_a | \mathbb{T}$ is an expanding circle map of class \mathcal{E} .

Calculate the multiplier of B_a at the origin.

This Blyaschke map is an example of a hyperbolic rational function.

To state some results in adequately general form, we will also consider a bigger class \mathcal{E}^1 of C^1 -smooth expanding circle maps and a class $\mathcal{E}^{1+\delta}$ of C^1 -smooth maps whose derivative satisfies the Hölder condition with exponent $\delta \in (0, 1)$. (However, for applications to holomorphic dynamics we will only need real analytic maps, so the reader can always assume it.)

EXERCISE 4.85. For any $g \in \mathcal{E}^1$, there exists a smooth Riemannian metric ρ on \mathbb{T} such that

 $||Dg(z)||_{\rho} \ge \lambda > 1 \text{ for all } z \in \mathbb{T}.$

This metric is called Lyapunov.

EXERCISE 4.86. Show that any expanding circle map $g \in \mathcal{E}^1$ has a unique fixed point $\beta \equiv \beta_q \in \mathbb{T}$.

Conjugating g be a rotation, we can always normalize it so that g(1) = 1.

Any expanding circle map $g: \mathbb{T} \to \mathbb{T}$ lifts to a homeomorphism $\tilde{g}: \mathbb{R} \to \mathbb{R}$ of the universal covering of \mathbb{T} satisfying $\tilde{g}(x+1) = \tilde{g}(x) + 2$ (as it is assumed to have degree two). Moreover, if g is normalized then \tilde{g} can be chosen so that $\tilde{g}(0) = 0$. In what follows, we will often identify the circle \mathbb{T} with the quotient \mathbb{R}/\mathbb{Z} , and view gas $\tilde{g} \mod \mathbb{Z}$ without making notational difference between these maps.

28.2. Symbolic model. Let us consider a symbolic sequence $\bar{k} = (k_0, k_1, ...) \in \Sigma$ of zeros and ones. Each such a sequence represents some number

$$\theta(\bar{k}) = \sum_{n=0}^{\infty} \frac{k_n}{2^{n+1}} \in [0,1]$$

in its diadic expansion. As everybody learns in the school (in the context of decimal expansions), all numbers except those of the form $m/2^n$ admit a unique diadic expansion. The numbers of the form $m/2^n$ with odd m admit exactly two diadic expansions:

$$\frac{k_0}{2} + \dots + \frac{k_{n-2}}{2^{n-1}} + \frac{1}{2^n} = \frac{k_0}{2} + \dots + \frac{k_{n-2}}{2^{n-1}} + \sum_{m=n+1}^{\infty} \frac{1}{2^m}.$$

Thus the corresponding symbolic sequences viewed as representations of numbers should be identified. If we consider the numbers mod 1, then we should also identify the sequence **0** of all zeros to the sequence **1** of all ones. Let us call these identifications on Σ "arithmetic" and the space Σ modulo these identifications arithmetic quotient of Σ . Of course, this quotient is in a natural one-to-one correspondence with the unit interval with identified endpoints, i.e., with the circle.

EXERCISE 4.87. Show that the projection

$$\pi_0: \Sigma \to \mathbb{T}, \quad \bar{k} \mapsto \exp(2\pi i \,\theta(\bar{k}))$$

(continuously) semi-conjugates the Bernoulli shift $\sigma : \Sigma \to \Sigma$ (see §21.4) to the circle endomorphism $f_0 : z \mapsto z^2$. Thus $f_0 : \mathbb{T} \to \mathbb{T}$ is topologically conjugate to the arithmetic quotient of the Bernoulli shift.

It turns out that the same is true for all expanding circle maps $g \in \mathcal{E}^1$:

LEMMA 4.88. Any circle expanding map $f \in \mathcal{E}^1$ is topologically conjugate to the arithmetic quotient of the Bernoulli shift.

PROOF. Let $g \in \mathcal{E}^1$. Consider its fixed point β . It has a single perimage β^1 different from $\beta \equiv \beta^0$. These two points, β and β^0 , divide the circle into two (open) intervals intervals, I_0^1 and I_1^1 (counting anti-clockwise starting from β). Moreover, g homeomorphically maps each I_k^1 onto $\mathbb{T} \smallsetminus \beta$. Hence each I_k^1 contains a preimage β_k^2 of β^1 . This point divides I_k^1 into two open intervals, I_{k0}^2 and I_{k1}^2 (counting anti-clockwise). We obtain four intervals, I_{kj}^2 , $k, j \in \{0, 1\}$ such that g homeomorphically maps each I_{kj}^2 onto I_k^1 .

Continuing inductively, we see that

$$\mathbb{T} \smallsetminus g^{-n}\beta = \bigcup_{k_s \in \{0,1\}} I^n_{k_0 k_1 \dots k_{n-1}},$$

where:

(i) the anti-clockwise order of the intervals $I_{\bar{k}}^n$ (starting from β) corresponds to

the lexicographic order on the symbolic strings $\tilde{\vec{k}} = (k_0 k_1 \dots k_{n-1});$ (ii) the map g homeomorphically maps $I^n_{\vec{k}}$ onto $I^{n-1}_{\sigma(\vec{k})}$, where the string $\sigma(\vec{k}) =$ $(k_1 \ldots k_{n-1})$ is obtained from \bar{k} by erasing the first symbol.

(iii) any interval $I_{\bar{k}}^n$ contains a point $\beta_{\bar{k}}^{n+1} \in g^{-(n+1)}\beta$ which divides it into two intervals $I_{\bar{k}0}^{n+1}$ and $I_{\bar{k}1}^{n+1}$ of the next level. Thus g^n homeomorphically maps each interval $I_{\bar{k}}^n$ onto the punctured circle

 $\mathbb{T} \setminus \{\beta\}$. Since g is expanding, the lengths of these intervals shrink exponentially fast:

$$|I_{\bar{k}}^n| \le \frac{2\pi}{C} \lambda^{-n},$$

where C > 0 and $\lambda > 1$ are constants from (28.1). It follows that for any infinite sequence $\bar{k} = (k_0 k_1 \dots) \in \Sigma$ of zeros and ones, the closed intervals $\bar{I}_{k_0 \dots k_{n-1}}^n$ form a nest shrinking to a single point $z = \pi(\bar{k})$. Thus we obtain a map $\pi: \Sigma \to \mathbb{T}$.

Under this map, the cylinders of rank n are mapped to the intervals of rank n. Since the latter shrink, π is continuous.

The above property (ii) implies that π is equivariant. Thus q is a quotient of the Bernoulli shift.

We only need to describe the fibers of π . If z is not an iterated preimage of β , then it belongs to a *single* interval of any rank. Hence $\operatorname{card}(\pi^{-1}(z)) = 1$. Obviously the fiber $\pi^{-1}(\beta)$ consists of two extremal sequences, (0) and **1**. Otherwise $z = \beta_{k_0...k_{n-1}}^{n+1} \in g^{-(n+1)}\beta$ for some $n \ge 0$ (except that for n = 0, the point β^1 does not have subsripts). Then it is a boundary point for exactly two intervals of each order $m \ge n+1$. For m = n+1, the corresponding symbolic sequences differ by the last symbol only: $(k_0 \dots k_{n-1} 0)$ and $(k_0 \dots k_{n-1} 1)$. For all further levels, we should add symbol 1 to the first sequence and symbol 0 to the second one. Thus:

$$\pi(k_0 \dots k_{n-1} \, 0 \, 1 \, 1 \, 1 \dots) = z = \pi(k_0 \dots k_{n-1} \, 1 \, 0 \, 0 \, \dots),$$

which are exactly the arithmetic identifications on Σ .

Thus all expanding circle maps of class \mathcal{E}^1 are topologically the same:

PROPOSITION 4.89. Any two expanding circle maps of class \mathcal{E}^1 are topologically conjugate by a unique orientation preserving circle homeomorphism. In particular, expanding circle maps do not admit non-trivial orientation preserving automorphisms.

PROOF. Lemma 4.88 gives the same standard model for any expanding circle map of class \mathcal{E}^1 . In this model, the anti-clockwise order on $\mathbb{T} \setminus \{\beta\}$ corresponds to the lexicographic order on Σ . Hence the corresponding conjugacy h between two circle maps, g and \tilde{g} , is orientation preserving.

Such a conjugacy is unique. Indeed, it must carry the points of $g^{-n}(\beta)$ to $\tilde{g}^{-1}(\tilde{\beta})$ preserving their anti-clockwise order starting from the corresponding fixed points, β and $\tilde{\beta}$. Hence h is uniquely determined on the iterated preimages of β . Since these preimages are dense in \mathbb{T} (by the previous lemma), h is uniquely determined on the whole circle. \square

Remarks. 1. Expanding circle maps have one orientation reversing automorphism. In the case of $z \mapsto z^2$ it is just $z \mapsto \overline{z}$ (compare with Exercise 4.16).

2. The above discussion can be generalized in a straightforward way to expanding circle maps of degree d > 2. There is one difference though: if d > 2 then the group of orientation preserving automorphisms of g is not trivial any more but rather the cyclic group of order d - 1 (consider $z \mapsto z^d$).

28.3. Complex extensions of circle maps. In this section we will take a closer look at the holomorphic extensions of expanding cicle maps of class \mathcal{E} .

EXERCISE 4.90. (i) For any $g \in \mathcal{E}$, there exist two \mathbb{T} -symmetric topological annuli $V \subseteq V'$ (bounded by smooth Jordan curves) such that g admits a holomorphic extension to V and maps it onto V' as a double covering.

Hint: Extend the Lyapunov metric from Exercise 4.85 to a neighborhood of \mathbb{T} . (ii) Show that vice versa, property (i) imlies that $g \in \mathcal{E}$. Hint: Use the hyperbolic metric in V'.

(iii) Show that all points $z \in V \setminus \mathbb{T}$ escape, i.e., $g^n z \in V' \setminus V$ for some $n \in \mathbb{N}$.

Thus property (i) can be used as a definition of an expanding circle map of class \mathcal{E} . In fact, only exterior part of the above extension is needed to reconstruct the circle map (it will be useful in what follows):

LEMMA 4.91. Let $\Omega \subset \Omega' \subset \mathbb{C}$ be two open conformal annuli whose inner boundaries coincide with the unit circle \mathbb{T} . Let $g : \Omega \to \Omega'$ be a holomorphic double covering. Then g admits an extension to a holomorphic double covering $G : V \to V'$, where $V \subseteq V'$ are \mathbb{T} -symmetric annuli such that $\Omega = V \setminus \overline{D}$ and $\Omega' = V' \setminus \overline{D}$. If the outer boundary of Ω is contained in Ω' , then $V \subseteq V'$ and the restriction $G|\mathbb{T}$ is an expanding cicle map of class \mathcal{E} .

PROOF. First show that g continuously extends to \mathbb{T} (apply boundary properties of confomal maps to inverse branches of g??). Then use the Schwarz Reflection Principle.

Consider a holomorphic extension $g: V \to V'$ of a map $g \in \mathcal{E}$ given by Exercise ??. Thus $V \Subset V'$ are two T-symmetric annuli neighborhoods of the circle. Let $A = (\bar{V}' \setminus V) \setminus \mathbb{D}$ be the "outer" fundamental annulus for g.

Given another map $\tilde{g}: \tilde{V} \to \tilde{V}'$ as above, we will mark the corresponding objects with "tilde".

PROPOSITION 4.92. Any two expanding circle maps $g: V \to V'$ and $\tilde{g}: V \to \tilde{V}'$ are conjugate by a qc map $h: (V', V, \mathbb{T}) \to (\tilde{V}', \tilde{V}, \mathbb{T})$ commuting with the reflection τ about the circle. In fact, any equivariant qc map $H: A \to \tilde{A}$ between the fundamental annuli admits a unique extension to a qc conjugacy h as above. Moreover Dil(h) = Dil(H).

PROOF. Consider an equivariant qc map H as above with dilatation K. By Lemma ?? it can be uniquely lifted to an equivariant K-qc homeomorphism h: $V' \setminus \overline{\mathbb{D}} \to \widetilde{V}' \setminus \overline{D}$. By ??, h admits a continuous extension to the unit circle. Reflecting it to the interior of the circle (and then exploiting Proposition 2.28) we obtain a desired K-qc conjugacy $h: V' \to \widetilde{V}'$.

Let us endow the exterior $\mathbb{C} \setminus \overline{\mathbb{D}}$ of the unit disk, with the hyperbolic metric $\rho \equiv \rho_{\mathbb{C} \setminus \overline{D}}$. The hyperbolic length of a curve γ will be denoted by $l_{\rho}(\gamma)$, while it Euclidean length will be denoted by $|\gamma|$.

LEMMA 4.93. Let $g: V \to V'$ be an expanding circle map of class \mathcal{E} . Let Ω and Ω' be two (open) annuli whose inner boundary is the circle \mathbb{T} . Let $h: \Omega \to \Omega'$ be a homeomorphism commuting with g. Then h admits a continuous extension to a map $\Omega \cup \mathbb{T} \to \tilde{\Omega} \cup \mathbb{T}$ identical on the circle.

Lifting $g: V \to V'$ to the universal covering, we obtain a conformal map $\tilde{g}: \tilde{V} \to \tilde{V}'$, where \tilde{V} and \tilde{V}' are \mathbb{Z} -invariant \mathbb{R} -symmetric neighborhoods of \mathbb{R} , and f(z+1) = f(z) + 2.

f(x) = f(x) + 2.

PROOF. Given a set $X \subset A$, let \tilde{X} denote its image by ω . Let us take a configuration consisting of a round annulus $L^0 = \mathbb{A}[r, r^2]$ contained in A, and an interval $I_0 = [r, r^2]$. Let $L^n = P_0^{-n} L^0$, and I_k^n denote the components of $P_0^{-n} I^0$, $k = 0, 1, \ldots, 2^n - 1$. The intervals I_k^n subdivide the annulus L^n into 2^n "Carleson boxes" Q_k^n .

Since the (multi-valued) square root map P_0^{-1} is infinitesimally contracting in the hyperbolic metric, the hyperbolic diameters of the boxes \tilde{Q}_k^n are uniformly bounded by a constant C.

Let us now show that ω is a hyperbolic quasi-isometry near the circle, that is, there exist $\epsilon > 0$ and A, B > 0 such that

(28.3)
$$A^{-1}\rho(z,\zeta) - B \le \rho(\tilde{z},\tilde{\zeta}) \le A\rho(z,\zeta) + B,$$

provided $z, \zeta \in \mathbb{A}(1, 1+\epsilon), |z-\zeta| < \epsilon.$

Let γ be the arc of the hyperbolic geodesic joining z and ζ . Clearly it is contained in the annulus $\mathbb{A}(1,r)$, provided ϵ is sufficiently small. Let t > 1 be the radius of the circle \mathbb{T}_t centered at 0 and tangent to γ . Let us replace γ with a *combinatorial geodesic* Γ going radially up from z to the intersection with \mathbb{T}_t , then going along this circle, and then radially down to ζ . Let N be the number of the Carleson boxes intersected by Γ . Then one can easily see that

$$p(z,\zeta) = l_{\rho}(\gamma) \asymp l_{\rho}(\Gamma) \asymp N$$

provided $\rho(z,\zeta) \ge 10 \log(1/r)$ (here $\log(1/r)$ is the hyperbolic size of the boxes Q_k^n). On the other hand

$$\rho(\tilde{z},\zeta) \le l_{\rho}(\Gamma) \le CN,$$

so that $\rho(\tilde{z}, \tilde{\zeta}) \leq C_1 \rho(z, \zeta)$, and (28.3) follows.

But quasi-isometries of the hyperbolic plane admit continuous extensions to \mathbb{T} (see, e.g., [**Th**]). Finally, it is an easy exercise to show that the only homeomorphism of the circle commuting with P_0 is identical.

We will show next that "outer automorphisms" of circle maps move points bounded hyperbolic distance:

LEMMA 4.94. Let $g: V \to V'$ be a map of class \mathcal{E} . Let Ω and Ω' be two open annuli in $V \setminus \overline{\mathbb{D}}$ with inner boundary \mathbb{T} , and let $h: \Omega \to \Omega'$ be an automorphism of g. Then for any $\delta > 0$ there exists an $R = R(\delta) > 0$ such that $\rho(z, hz) \leq R$ for all points $z \in \Omega$ whose distance from the outer boundary of Ω is at least δ .

PROOF. By Proposition 4.92, g is qc conjugate to the quadratic circle map $f_0: z \mapsto z^2$. Of course, this conjugacy can be extended to a global qc homeomorphism of \overline{C} (e.g., by ??). Since qc homeomorphisms of $\mathbb{C} \setminus \overline{\mathbb{D}}$ are hyperbolic quasi-isometries

(??), it is enough to prove the assertion for f_0 . So, let us assume from now on that $g = f_0$.

Of course, the assertion is true for any compact subset of Ω . Hence we need to check it only near to the unit circle.

By 4.93, h admits a continuous extension to the unit circle. Of course, it still commutes with g on the circle. By Proposition 4.89, $h|\mathbb{T} = \text{id}$. Hence for any $\epsilon > 0$ there exists an r > 1 such that $\mathbb{A}(1, r] \subseteq \Omega$ and

$$|z - hz| < \epsilon$$
 for $z \in \mathbb{A}(1, r]$.

Consider a fundamental annulus A of g compactly contained in $\mathbb{A}(1, r]$. By compactness, there exists an R > 0 such that

$$\rho(z, hz) \le R \quad \text{for} \quad z \in A.$$

Let $A^n = g^{-n}A$. Take some $z \in A^1$. Since $|z - hz| < \epsilon$, these points are obtained by applying the same local branch of the square root map g^{-1} to the points gz and g(hz) = h(gz). Since the local branches of g^{-1} preserve the hyperbolic distance on $\mathbb{C} \setminus \overline{\mathbb{D}}$, we have: $\rho(z, hz) = \rho(gz, h(gz)) \leq R$.

Replacing A by A^1 , we obtain the same bound for any $z \in A^2$, etc. The conclusion follows.

28.4. Notes. Classical Theorem 4.38 (due to Fatou and Julia) plays a fundamental role in the field. It is valid for a general rational function f of degree d and implies that f may have at most 2d - 2 (the number of critical points) attracting cycles. In particulr, a polynomials of degree d may have at most d - 1 finite (in \mathbb{C}) attracting cycles.

Lemma 4.47 is due to Fatou: it gives twice bigger bound on the number of non-repelling cycles than was anticipated.

The notion of a quadratic-like map was introduced by Douady and Hubbard in their fundamental paper [**DH3**]. The application to the sharp bound on the number of finite non-repelling cycles for polynomials (by d - 1, see Theorem 4.80) was given in [**D1**]. An analogous bound (by 2d - 2) for rational maps is much harder to prove; it was later established by Shishikura [**Sh1**].

A simple proof for existence of Siegel disks (Proposition 5.5) is due to Yoccoz.

The No Wandering Domains Theorem for rational functions appeared in [S1, S2, S3]. Since then, it appeared in every basic text book on the subject. The above exposition extends without changes to the case of higher degree polynomials. For rational functions, the proof is exactly the same for simply connected components of F(f) but some extra analysis is needed to rule out the multiply-connected domains: (this can be actually done by a direct geometric argument that avoids qc deformations, see [Ba]). Let us also mention that there is a class of transcendental functions "of finite type" (including λe^z and $\lambda \sin z$) that enjoy similar description of the dynamics on F(f) as their rational counterparts, see [BR, EL, GK].

The global measure-theoretic attractor appeared in [?]

CHAPTER 5

Remarkable functional equations

29. Linearizing coordinate in the attracting case

Study of certain functional equations was one of the main motivations for the classical work in holomorphic dynamics. By means of these equations the local dynamics near periodic points of different types can be reduced to the simplest normal form. But it turns out that the role of the equations goes far beyond local issues: global solutions of the equations play a crucial role in understanding the dynamics.

We will start with the local analysis and then globalize it (though sometimes one can go the other way around). For the local analysis we put the fixed point at the origin and consider a holomorphic map

$$(29.1) f: z \mapsto \sigma z + a_2 z^2 + \dots$$

near the origin.

29.1. Attracting points and linearizing coordinates.

29.1.1. Local linearization. Let us start with the simplest case of an attracting fixed point. In turns out that such a map can always be linearized near the origin:

THEOREM 5.1. Consider a holomorphic map (29.1) near the origin. Assume $0 < |\sigma| < 1$. Then there exists an f-invariant Jordan disk $V \ni 0$, an r > 0, and a conformal map $\phi : (V, 0) \to \mathbb{D}_r$ with $\phi'(0) = 1$ satisfying the equation:

(29.2)
$$\phi(fz) = \sigma\phi(z)$$

The above properties determine uniquely the germ of ϕ at the origin.

The above function ϕ is called the *linearizing coordinate* for f near 0 or the Königs function. The linearizing equation (29.2) is also called the Schröder equation. It locally conjugates f to its linear part $z \mapsto \sigma z$.

PROOF. The linearizer ϕ can be given by the following explicit formula:

(29.3)
$$\phi(z) = \lim_{n \to \infty} \sigma^{-n} f^n z.$$

To see that the limit exists (uniformly near the origin), let $z_n = f^n z$, $z_0 \equiv z$, notice that $z_n = O(|z\sigma|^n)$ uniformly near the origin, and take the ratio of the two consecutive terms in (29.3):

$$\frac{\sigma^{-n-1}z_{n+1}}{\sigma^{-n}z_n} = \sigma^{-1}\frac{\sigma z_n(1+O(|z_n|))}{z_n} = 1 + O(|z\sigma^n|).$$

Hence

$$\phi(z) = z \prod_{n=0}^{\infty} \frac{\sigma^{-n-1} z_{n+1}}{\sigma^{-n} z_n} = z(1 + O(|z|))$$

uniformly near the origin, and the conclusion follows.

Obviously, ϕ is a linearizer. Its uniqueness follows from the exercise below. \Box

EXERCISE 5.2. Show that if a holomorphic germ f near the origin commutes with the linear germ $z \mapsto \sigma z$, $0 < |\sigma| < 1$, then f is itself linear.

REMARK 5.1. We see that the conjugacy ϕ is constructed by going forward by the iterates of f and then returning back by the iterates of the corresponding linear map. This method of constructing a conjugacy between two maps will be used on several other occassions, see (32.2) and (??).

Let us note in conclusion that the Königs function $\phi = \phi_f$ depends holomorphically on f:

LEMMA 5.3. Let

$$f_{\lambda}(z): z \mapsto \sigma(\lambda)z + a_2(\lambda)z^2 + \dots$$

be a holomorphic family of local maps with attracting fixed point 0. Then the Königs function $\phi_{\lambda}(z)$ depends holomorphically on λ .

PROOF. The above proof shows that convergence in Königs formula (29.3) is locally uniform over λ . Hence the limit is holomorphic in (λ, z) .

29.1.2. *Extension to the immediate basin.* Next, we will extend the Königs function to the immediate basin of attraction:

PROPOSITION 5.4. Let f be a polynomial with anttracting periodic point α . Then the Königs function ϕ analytically extends to the immediate basin $D = D^0(\alpha)$, and it satisfies there Schröder functional equation (29.2). Moreover, the map ϕ : $D \to \mathbb{C}$ is a branched covering of infinite degree branched on $D \cap \operatorname{Crit}_f$. The fibers of ϕ are small orbits of $f \mid D$.

PROOF. We can assume without loss of generality that α is fixed, $f(\alpha) = \alpha$. The immediate basin D is exahausted by an increasing nest of domains

$$P_0 \subset P_1 \subset \cdots \subset P_n,$$

where P_0 is a domain for the local solution of (29.2) and P_{n+1} is the component of $f^{-1}(P_n)$ containing α (compare proof of Theorem 4.38). Then we can consecutively extend ϕ from P_n to P_{n+1} by means of the Schröder equation:

$$\phi(z) = \sigma^{-1}\phi(fz), \quad z \in P_{n+1}.$$

Since the maps $f : P_{n+1} \to P_n$ are branched coverings, all the extensions $\phi_n = \phi : P_n \to \sigma^n P_0$ are branched coverings, and hence the limiting map $\phi : D \to \mathbb{C}$ is a branched covering as well. As $\deg(f|P_n) > 1$ eventually for all n (once the P_n contain a critical point of f), we have $\deg \phi_n \to \infty$.

Moreover, any critical point of ϕ_{n+1} is either a critical point of f or else an f-preimage of a critical point ϕ_n .

The last assertion is also easily supplied consecutively for the maps ϕ_n .

So, in case of quadratic $f : z \mapsto z^2 + c$, the map ϕ branches on $(f|D)^{-pm}(0)$, where p is the period of α . Moreover, the critical points of ϕ are simple in this case, and its critical values are $\sigma^{-n} \phi(0)$, $n = 0, 1, \ldots$

30. Existence of Siegel disks

We will now give a simple proof of existence of Siegel disks in the quadratic family. Here it will be convenient to put a fixed point at the origin and to normalize the quadratic term so that $f_{\lambda}(z) = \lambda z + z^2$.

PROPOSITION 5.5. In the quadratic family $f_{\lambda}(z) = \lambda z + z^2$, $\lambda = e^{2\pi i\theta}$ with $\theta \in \mathbb{R}/\mathbb{Z}$, the map f_{λ} is linearizable for Lebesgue almost all rotation umbers θ .

PROOF. The idea is to construct Siegel disks as limits of attracting petals. To this end we need to control the size of the latter. By Proposition 5.4, the Königs map ϕ_{λ} is unbranched over the disk \mathbb{D}_r , where $r = r_{\lambda} = |\phi_{\lambda}(-\lambda/2)|$. Hence there exists a petal $D_{\lambda} \ni 0$ containing the critical point $-\lambda/2$ on its boundary which is univalently mapped by ϕ_{λ} onto \mathbb{D}_r .

By Lemma 5.3, the function $\lambda \mapsto \phi_{\lambda}(-\lambda/2)$ is holomorphic on the unit disc \mathbb{D} . Let us show that it is also bounded, and in fact $r_{\lambda} < 2$. Indeed, it is trivial to check that the filled Julia set $K(f_{\lambda})$ is contained in the disc $\overline{\mathbb{D}}_2$. Hence

$$D_{\lambda} \subset \operatorname{int} K(f_{\lambda}) \subset \mathbb{D}_2$$

But then $r_{\lambda} < 2$ by the Schwarz Lemma applied to the inverse function

(30.1)
$$\psi_{\lambda} = \phi_{\lambda}^{-1} : (\mathbb{D}_{r_{\lambda}}, 0) \to (D_{\lambda}, 0), \quad \psi'(0) = 1.$$

By classical results of Complex Analysis (Fatou and Privalov), the function $g(\lambda) := \phi_{\lambda}(-\lambda/2)$ has non-vanishing radial limits

$$\bar{g}(\theta) = \lim_{\rho \to 1} g(\rho e^{2\pi i \theta})$$
 for almost all $\theta \in \mathbb{R}/\mathbb{Z}$.

Let us finally show that for such a θ , the map f_{λ} with $\lambda = e^{2\pi i\theta}$ is linearizable on the disk of radius $\bar{r} := |\bar{g}(\theta)|/2 > 0$. Indeed, the family of functions ψ_{λ} (30.1) with $\lambda = \rho e^{2\pi i\theta}$ is well defined and normal (by the Little Montel) on the disk of radius \bar{r} (as long as ρ is sufficiently close to 1). Then any limit function ψ linearizes f_{λ} on \bar{r} .

31. Global leaf of a repelling point

Taking the local inverse of f, we conclude that repelling maps are also locally linearizable:

COROLLARY 5.6. Consider a holomorphic map (29.1) near the origin. Assume $|\sigma| > 1$. Then there exist Jordan disks $V \supseteq V' \supseteq 0$ such that f(V') = V, an r > 0, and a conformal map $\phi : (V, 0) \to \mathbb{D}_r$ with $\phi'(0) = 1$ satisfying the equation:

(31.1)
$$\phi(fz) = \sigma\phi(z), \quad z \in V'$$

The above properties determine uniquely the germ of ϕ at the origin.

Assume now that $f : \overline{\mathbb{C}} \to \overline{\mathbb{C}}$ is a polynomial with a repelling fixed point a. Let us consider the inverse linearizing function $\psi : (\mathbb{D}_r, 0) \to (V, a), \ \psi = \phi^{-1}$. It satisfies the functional equation

(31.2)
$$\psi(\sigma z) = f(\psi(z)), \quad z \in V'.$$

It allows us to extend ψ holomorphically to the disk $\mathbb{D}_{|\sigma|r}$ by letting $\psi(\zeta) = f(\psi(\zeta/\lambda))$ for $\zeta \in \mathbb{D}_{|\sigma|r}$. Repeating this procedure, we can consecutively extend f to the disks $\mathbb{D}_{|\sigma|^n r}$, $n = 1, 2, \ldots$, so that in the end f we obtain an entire function $\psi : \mathbb{C} \to \mathbb{C}$ satisfying (31.2).

We will now construct in a dynamical way the Riemann surface of the inverse (multivalued) function $\phi = \psi^{-1}$. The construction below is a special case of a general *natural extension* or *inverse limit* construction. Let us consider the space of inverse orbits of f converging to the fixed point 0:

$$\mathcal{L} = \{ \hat{z} = (z_{-n})_{n=0}^{\infty} : f(z_{-n-1}) = z_{-n}, \ z_{-n} \to 0 \}.$$

Define $\pi_{-n} : \mathcal{L} \to \mathbb{C}$ as the natural projections $\hat{z} \mapsto z_{-n}$. Let $z \equiv z_0$ and $\pi \equiv \pi_0 : \hat{z} \mapsto z$. The map f lifts to an invertible map $\hat{f} : \mathcal{L} \to \mathcal{L}$, $\hat{f}(\hat{z}) = (fz_{-n})_{n=0}^{\infty}$ such that $\hat{f}^{-1}(\hat{z}) = (z_{-n})_{n=1}^{\infty}$. Moreover, the projection π is equivariant: $\pi \circ \hat{f} = f \circ \pi$.

For a neighborhood U of z let $\hat{U} = \hat{U}(\hat{z}) = (U_{-n})_{n=0}^{\infty}$, where U_{-n-1} is defined inductively as the component of $f^{-1}(U_{-n})$ containing z_{-n-1} . We call \hat{U} the *pullback* of U along \hat{z} . Let us call a pullback \hat{U} regular if the maps $f : U_{-n} \to U_{-n-1}$ are eventually univalent. Since $z_{-n} \to 0$, $z_{-n} \in V$ for all $n \ge N$. Selecting U so small that $U_{-N} \subset V$, we see that $U_{-n} \subset V$ for all $n \ge N$, and hence all the maps $f : U_{-n-1} \to U_{-n}$ are univalent for $n \ge N$. Thus, \hat{U} is regular for a sufficiently small U.

We define topology on \mathcal{L} by letting all the regular pullbacks $\hat{U}(\hat{z})$ be the basis of neighborhoods of $\hat{z} \in \mathcal{L}$. Moreover, if $f: U_{-n-1} \to U_{-n}$ are univalent for $n \geq N$, the projection $\pi_{-N}: \hat{U} \to U_{-N}$ is homeomorphic, and we take it as a local chart on $\hat{\mathcal{L}}$. Transition maps between such local charts are given by iterates of f, so that, they turn \mathcal{L} into a Riemann surface.

EXERCISE 5.7. Show that the projections $\pi_{-n} : \mathcal{L} \to \mathbb{C}$ are holomorphic. Show that the critical points of π are the orbits $\hat{z} = (z_{-n})_{n=0}^{\infty}$ passing through a critical point of f (such orbits are called critical). Find the degree of branching of π at \hat{z} .

Let $\hat{a} = (a a \dots) \in \mathcal{L}$ be the fixed point lift of a. The following statement shows that $\hat{\mathcal{L}}$ is the indeed the Riemann surface for ϕ :

PROPOSITION 5.8. The maps ψ and ϕ lift to mutually inverse conformal isomorphisms $\hat{\psi} : (\mathbb{C}, 0) \to (\mathcal{L}, \hat{a})$ and $\hat{\phi} : \mathcal{L} \to \mathbb{C}$ conjugating $z \mapsto \sigma z$ to \hat{f} and such that $\pi \circ \hat{\psi} = \psi$.

PROOF. For $u \in \mathbb{C}$, we let $\hat{\psi}(u) = (\psi(u/\sigma^n)_{n=0}^{\infty})$.

Vice versa, if $\hat{z} = (z_{-n})_{n=0}^{\infty}$ then eventually $z_{-n} \in V$, so that the local linearizer ϕ is well defined on all z_{-n} , $n \geq N$. Let now $\hat{\psi}(\hat{z}) = \sigma^n \psi(z_{-n})$ for any $n \geq N$. It does not depend on the choice of n since $\phi | V$ conjugates f to $z \to \sigma z$.

We leave to the reader to check all the properties of these maps.

LEMMA 5.9. Let $C_f = \pi^{-1}(\bar{C}_f)$. Then the map $\mathcal{L} \smallsetminus C_f \to \mathcal{L} \smallsetminus \bar{C}_f$ is a covering. PROOF. Let $z \in \mathbb{C} \smallsetminus \bar{C}_f$ and let $U \subset \mathbb{C} \smallsetminus \bar{C}_f$ be a little disk around z. Then

$$\pi^{-1}(U) = \bigcup_{\hat{z} \in \pi^{-1}z} \hat{U}(\hat{z}),$$

and each \hat{U} projects univalently onto U.

Let $\hat{K}(f) = \pi^{-1}(K(f)).$

COROLLARY 5.10. Assume K(f) is connected. Let U be a component of $\mathcal{L} \setminus \hat{K}(f)$. Then U is simply connected, so that, the projection $\pi : U \to D_f(\infty)$ is a universal covering.

PROOF. Since K(f) is connected, $\overline{C}_f \subset K(f)$. By Lemma 5.9, $U \to D_f(\infty)$ is a covering map. Since $D_f(\infty)$ is conformally equivalent to \mathbb{D}^* , U is either conformally equivalent to \mathbb{D}^* or is simply connected. But in the former case U would be a neighborhood of ∞ in $\mathcal{L} \approx \mathbb{C}$, so that, $\hat{K}(f)$ would be bounded. It is impossible since $\hat{K}(f)$ is \hat{f} -invariant, where by Proposition 5.8 \hat{f} is conjugate to $z \mapsto \sigma z$ with $|\sigma| > 1$.

32. Superattractng points and Böttcher coordinates

THEOREM 5.11. Let $f: z \mapsto z^d + a_{d+1}z^{d+1} + \ldots$ be a holomorphic map near the origin, $d \geq 2$. Then there exists an f-invariant Jordan disk $V \ni 0$, $r \in (0,1)$, and a conformal map $B: (V,0) \to (\mathbb{D}_r, 0)$ satisfying the equation:

$$(32.1) B(fz) = B(z)^d.$$

The above properties determine uniquely the germ of B at the origin, up to postcomposition with rotation $z \mapsto e^{2\pi i/(d-1)}z$ (so, it is unique in the quadratic case d = 2). Moreover, it can be normalized so that B'(0) = 1.

The map B is called the *Böttcher function*, or the *Böttcher coordinate* near 0. Equation (32.1) is called the *Böttcher equation*. In the Böttcher coordinate the map f assumes the normal form $z \mapsto z^d$.

PROOF. The Böttcher function can be given by the following explicit formula: (32.2) $B(z) = \lim_{n \to \infty} \sqrt[d^n]{f^n z},$

where the value of the d^n th root is selected so that it is tangent to the id at ∞ . Obviously, this function, if exists, satisfied the Böttcher equation. So, we only need to check that the limit exists.

Let
$$z_n = f^n z$$
, where $z_0 \equiv z$. Then

$$\frac{\frac{d^{n+1}\sqrt{z_{n+1}}}{2^n \sqrt{z_n}}}{\frac{d^n \sqrt{z_n}}{d^n \sqrt{z_n}}} = \frac{\frac{d^{n+1}\sqrt{z_n^d(1+O(z_n))}}{d^n \sqrt{z_n}}}{d^n \sqrt{z_n}} = 1 + O\left(\frac{z_n}{d^{n+1}}\right)$$

Hence

$$B(z) = \lim_{n \to \infty} \sqrt[d^n]{z_n} = z \prod_{n=0}^{\infty} \frac{\sqrt[d^{n+1}]{z_{n+1}}}{\sqrt[d^n]{z_n}} = z \prod_{n=0}^{\infty} \left(1 + O\left(\frac{z_n}{d^{n+1}}\right) \right) = z(1 + O(z)),$$

where the last product is convergant uniformly at a superexponential rate.

Finally, uniqueness of the Böttcher function follows from the exercise below. $\hfill\square$

EXERCISE 5.12. Let $d \ge 2$. Show that there are no holomorphic germs commuting with $g: z \mapsto z^d$ near the origin, except rotations $z \mapsto e^{2\pi i/(d-1)}z$.

32.1. Böttcher vs Riemann. Let us now consider a quadratic polynomial f_c near ∞ . Since ∞ is a superattracting fixed point of f of degree 2, the map f_c near ∞ can be reduced in the Böttcher coordinate to the map $z \mapsto z^2$ (Theorem 5.11). Thus, there is a Jordan disk $V = V_c \subset \mathbb{C}$ whose complement $\mathbb{C} \setminus V$ is f_c -invariant, some R > 1, and a conformal map $B_c : \mathbb{C} \setminus V \to \mathbb{C} \setminus \mathbb{D}_R$ satisfying the Böttcher equation:

(32.3)
$$B_c(f_c z) = B_c(z)^2$$

Moreover, $B_c(z) \sim z$ as $z \to \infty$.

We will now globalize the Böttcher function.

).

32.1.1. Connected case.

THEOREM 5.13. Let $f_c : z \mapsto z^2 + c$ be a quadratic polynomial with connected Julia set. Then the Böttcher function admits an analytic extension to the whole basin of ∞ . Moreover, it conformally maps $D_c(\infty)$ onto the complement of the unit disk and globally satisfies (32.3).

PROOF. We will skip label c from the notations. Let, as usual, $f_0(z) = z^2$.

Let $U^n = \hat{\mathbb{C}} \smallsetminus f^{-n}\bar{V}$. Then $U^0 \subset U^1 \subset U^2 \subset \ldots$ and $\bigcup U^n = D_f(\infty)$. Since the filled Julia set K(f) is connected, the domains U^n are topological disks and the maps $f : U^{n+1} \to U^n$ are double coverings branched point at ∞ (recall the proof of Theorem 4.13).

Let $\Delta^n = \overline{\mathbb{C}} \smallsetminus \overline{\mathbb{D}}_{R^{1/2^n}}$. By Lemma 1.53, the Böttcher function $B : U^0 \to \Delta^0$ admits a lift $\tilde{B} : U^1 \to \Delta^1$ such that $f_0 \circ \tilde{B} = B \circ f$. But the Böttcher equation tells us that $B : U^0 \to \Delta^0$ is a lift of its restriction $B : f(U^0) \to f_0(\Delta^0)$. If we select \tilde{B} so that $\tilde{B}(z) = B(z)$ at some finite point $z \in U^0$, then these two lifts must coincide on U^0 : $\tilde{B}|U^0 = B$. Thus, \tilde{B} is the analytic extension of ϕ to U^1 . Obviously, it satisfies the Böttcher equation as well.

In the same way, the Böttcher function can be consecutively extended to all the domains U^n and hence to their union, $D_f(\infty)$.

Thus, the Böttcher function gives the uniformization of $\mathbb{C} \setminus K(f)$ by the unit disk. Given the intricate fractal structure of the Julia set, this is quite remarkable that its complement can be uniformized in this explicit way!

One can also go the other way around and costruct the Böttcher function by means of uniformization:

EXERCISE 5.14. Let $f = f_c$ be a quadratic polynomial with connected Julia set. Then the basin of infinity $\overline{D}_f(\infty)$ is a conformal disk. Uniformize it by the complement of the unit disk; $\psi : (\mathbb{D}, \infty) \to (D_f(\infty), \infty)$, normalized at ∞ so that $\psi(z) \sim \lambda z$ with $\lambda > 0$. Prove (without using the Böttcher theorem) that ψ conjugates $f_0: z \mapsto z^2$ on $\mathbb{C} \setminus \mathbb{D}$ to f on the basin of ∞ (and that $\lambda = 1$).

EXERCISE 5.15. Prove that $D_c(\infty)$ is the maximal domain of analyticity of the Böttcher function.

Let us finish with a curious consequence of Theorem 5.13. The *capacity* of a connected compact set $K \subset \mathbb{C}$ rel ∞ is defined as 1/R, where R is the radius of the disk \mathbb{D}_R such that the domain $\mathbb{C} \setminus K(f_c)$ can be conformally mapped onto $\mathbb{C} \setminus \overline{\mathbb{D}}_R$ by a map tangent to the id at ∞ .

COROLLARY 5.16. Let $f_c : z \mapsto z^2 + c$. Then the capacity of the filled Julia set $K(f_c)$ is equal to 1.

32.1.2. Cantor case. In the disconnected case the Böttcher function B_c cannot be any more extended to the whole basin of ∞ , as it branches at the critical point 0. However, B_c can still be extended to a big invariant region Ω_c containing 0 on its boundary.

THEOREM 5.17. Let $f_c: z \mapsto z^2 + c$ be a quadratic polynomial with disconnected Julia set. Then the Böttcher function B_c admits the analytic extension to a domain Ω_c bounded by a "figure eight" curve branched at the critical point 0. Moreover, B_c maps Ω_c conformally onto the complement of some disk $\overline{\mathbb{D}}_R$ with R > 1. The inverse map extends continuously to a map $\mathbb{C} \setminus \mathbb{D}_R \to \overline{\Omega}_c$ which is -one-to-one except that it maps two antipodal points $\pm Re^{2\pi i\theta} \in \mathbb{T}_R$ to 0.

PROOF. Again, we skip the label c.

Since $0 \in D_f(\infty)$, the orb(0) lands at the domain V of the Böttcher function near ∞ . By shrinking V, we can make $f^n 0 \in \partial V$ for some n > 0. Then there are no obstructions for consecutive extensions of B to the domains $U^k = \overline{\mathbb{C}} \setminus f^{-k}\overline{V}$, $k = 0, 1, \ldots, n$ (in the same way as in the connectef case). All these domains are bounded by real analytic curves except the last one, U^n , which is bounded by a figure eight curve branched at 0. This is the desired domain Ω .

For $c \in M$, we let Ω_c be the whole basin of infinity, $\Omega_c(\infty)$.

For a point $z \in \mathbf{\Omega}_c$, the polar coordinates (r, theta) of $B_c(z)$ are called the *external coordinates* of z.

32.1.3. Böttcher position of the critical value. Since the critical value $c \in \partial U^{n-1}$ belongs to the domain of B_c , the expression $B_c(c)$ is well-defined (provided the Julia set $J(f_c)$ is disconnected). It gives the Böttcher position of the critical value as a function of the parameter c. This function will play a crucial role in what follows.

32.2. External rays and equipotentials. The map $f_0: z \mapsto z^2$ on $\mathbb{C} \setminus \overline{\mathbb{D}}$ has two invariant foliations, by the straight rays going to ∞ and by round circles centered at the origin. (Note that the first foliation is dynamically defined: see the hint to Exercise 5.12.) We will label the rays by their angles $\theta \in \mathbb{R}/2\pi\mathbb{Z}$ and the circles by their radia r > 1 or by their "heights" $t = \log r \in \mathbb{R}_+$. So,

 $\mathcal{R}_0^{\theta} = \{ re^{i\theta} : r \in \mathbb{R}_+ \}, \quad \mathcal{E}_0^r \equiv \mathcal{E}_0^t = \{ re^{i\theta} : \theta \in \mathbb{R}/2\pi\mathbb{Z} \}, \ t = \log r,$

where the subscript 0 suggests affiliation to the map f_0 . Note that

$$f_0(\mathcal{R}_0^\theta) = \mathcal{R}_0^{2\theta}$$
 and $f_0(\mathcal{E}_0^t) = \mathcal{E}_0^{2t}$.

If we now take an arbitrary quadratic polynomial f_c , then by means of the Böttcher function B_c , the above two foliations can be transferred to the domain $\Omega_c \subset D_c(\infty)$, supplying us with the foliation by *external rays* and *equipotentials*. The rays naturally labeled by the corresponding *external angles* θ , while the equipotentials are labeled by the *equipotential radii* r or heights t. Let $\mathcal{R}^{\theta} \equiv \mathcal{R}^{\theta}_c$ stand for the external ray of angle θ and let $\mathcal{E}^r \equiv \mathcal{E}^r_c$ or $\mathcal{E}^t \equiv \mathcal{E}^t_c$ stand for the equipotential of height $t = \log r$. (We will also use notation $\mathcal{R}^{\theta}(t) \equiv \mathcal{R}^{\theta}(r)$ for the point on the ray \mathcal{R}^{θ} whose equipotential level is equal to $t = \log r$.)

If $K(f_c)$ is connected then $\Omega_c = D_c(\infty)$, so that the whole basin of infinity is foliated by the external rays and equipotentials.

In the disconnected case, we can pull the two foliations in Ω_c back by the iterates of f to obtain to singular foliations on the whole basin of ∞ . They have singularities at the critical points of iterated f, i.e., at 0 and all its preimages under the iterates of f.

In this context external rays will be understood as the non-singular leaves of these foliaitons that go to ∞ (i.e., the maximal non-singular extensions of the rays in Ω_c). Countably many rays land at the preimages of 0. All other rays are properly embedded into the basin; they will be called *proper rays*. Two (improper) rays landing at the critical point 0 will be called the *critical rays*. The particularly important ray going through the critical value will be called the *charactersistic ray*

(its external angle will be also called *characteristic*). Of cource, it contains the (coinciding) images of the critical rays.

The figure-eight that bounds Ω_c will be called the *critical equipotential*.

For $r > B_c(0)$, we let $\Omega_c(r) \equiv \Omega_c(t)$ be the Jordan disk bounded by the equipotential of height $t = \log r$. We will refer to it as a *subpotential* disk of height t (or, of radius r).

32.3. Dynamical Green function. The *Green function* of a quadratic polynomial $f = f_c$ is defined as follows:

(32.4)
$$G_c(z) = \log |B_c(z)|,$$

where B_c is the Böttcher function of f_c . The Green function is harmonic wherever the Böttcher function is defined (since the Böttcher function never vanishes) and has a logarithmic singularity at ∞ :

$$G(z) = \log|z| + o(1).$$

In the connected case, (32.4) defines the Green function in the whole basin $D(\infty)$. In the disconnected case definition (32.4) can be used only in the domain Ω . However, in either case the Green function satisfies the equation:

$$(32.5) G(fz) = 2G(z).$$

This equation can be obviously used in order to extend the Green function harmonically to the *whole basin of* ∞ . Let us summarize simple properties of this extension:

EXERCISE 5.18. a) In the connected case the Green function does not have critical points. In the disconnected case, its critical points coincide with the critical points of iterated f.

b) Equipotentials are the level sets of the Green function, while external rays (and their preimages) are its gradient curves.

c) The Brolin formula holds:

(32.6)
$$G(z) = \lim_{n \to \infty} \frac{1}{2^n} \log |f^n z|, \quad z \in D(\infty).$$

d) Extention of the Green function by 0 through the filled Julia set K(f), gives a continuous subharmonic function on the whole complex plane.

e) The Julia set is Dirichlet regular.

These properties show that the dynamical Green function G is indeed the Green function of $D(\infty)$ with the pole at ∞ as was defined in the general context in §7.9. Moreover, the dynamical notion of external rays and equipotentials matches with the general one.

EXERCISE 5.19. Assume that the Julia set J(f) is connected. Endow its basin $D(\infty) \setminus \{\infty\}$ with the hyprolci metric ρ . Then for any external ray \mathcal{R}^{θ} we have:

$$\rho(z,\zeta) = \left| \log \frac{G(z)}{G(\zeta)} \right|, \quad z,\zeta \in \mathcal{R}^{\theta}.$$

32.4. Diadic grid. Let us fix some $r = e^t > 1$, and consider the annulus $\mathbb{A}(1,r]$ cut along the real line, $\Delta^0 = \mathbb{A}(1,r] \smallsetminus (0,r]$. Let us pull it back by the dynmics of $f_0: z \mapsto z^2$; let $\Delta^n_{\overline{i}} \equiv \Delta^n_{i_0...i_{n-1}}$ be the pullback under the branch of $f_0^{-n}: \mathbb{C} \smallsetminus \mathbb{R}_+ \to \mathbb{C}$ that maps $\mathbb{T} \smallsetminus \{1\}$ to the diadic interval $J_{\overline{i}}$. It provides us with the tiling of each annulus $\mathbb{A}(1, r^{1/2^n})$ by 2^n rectangles $\Delta^n_{\overline{i}}$

such that

(32.7)
$$\Delta_{i_0\dots i_n}^n \subset \Delta_{i_0\dots i_{n-1}}^n \quad \text{and} \quad f_0(\Delta_{i_0\dots i_n}^n) = \Delta_{i_1\dots i_n}^n.$$

Let now f be a quadratic polynomial with connected Julia set. Taking the pullback of the above grid under the Böttcher map, $D_{\overline{i}}^n = B^{-1}(\Delta_{\overline{i}}^n)$, we obtain the corresponding tilings of the external annuli neighborhood of the Julia set. Since Bis equivariant, the behavior of this grid under the the dynamics (and the inclusion) is the same as in (32.7).

This grid give a useful dynamical picture for f in the external neighborhood of the Julia set.

33. Parabolic points and Écale-Voronin cylinders

CHAPTER 6

Parameter plane (the Mandelbrot set)

28. Definition and first properties

28.1. Notational convention. We will label the objects corresponding to a map f_c by c, e.g., $J_c = J(f_c)$, $Per(f_c) = Per_c$. We often use notation $c_o \equiv o$ for a base parameter, so that $f_o = f_{c_o}$, $J_o = J_{c_o}$, etc.

28.2. Connectedness locus and polynomials $c \mapsto f_c^n(0)$. The Mandelbrot set presents at one glance the whole dynamical diversity of the complex quadratic family $f_c: z \mapsto z^2 + c$. Figure ... shows this set and its blow-ups in several places. It is remarkable that all this intricate structure is hidden behind the following one-line definition.

Recall the Basic Dichotomy for the quadratic maps: the Julia set $J(f_c)$ is either connected or Cantor (Theorem 4.13). By definition, the *Mandelbrot set* M consists of those parameter values $c \in \mathbb{C}$ for which the Julia set J_c is connected. It is equivalent to saying that the orbit of the critical point

(28.1)
$$0 \mapsto c \mapsto c^2 + c \mapsto (c^2 + c)^2 + c \mapsto \dots$$

is not escaping to ∞ . Let us denote the *n* th polynomial in (28.1) by $v_n(c)$, so that $v_0(c) \equiv 0, v_1(c) \equiv c$, and recursively

(28.2)
$$v_{n+1}(c) = v_n(c)^2 + c.$$

Note that deg $v_n = 2^{n-1}$.

Though the polynomials v_n are not iterates of a single polynomial, they behave in many respects similarly to the iterated polynomials:

EXERCISE 6.1 (Simplest properties of M). Prove the following properties:

- (i) If $|v_n(c)| > 2$ for some $n \in \mathbb{N}$ then $v_n(c) \to \infty$ as $n \to \infty$. In particular, $M \subset \overline{\mathbb{D}}_2$.
- (ii) $v_n(c) \to \infty$ locally uniformly on $\mathbb{C} \setminus M$. Hence M is compact.
- (iii) C \ M is connected. Hence M is full and all components of int M are simply connected.
- (iv) The set of normality of the sequence $\{v_n\}$ coincides with $\mathbb{C} \smallsetminus \partial M$.

One can see a similarity between the Mandelbrot set (representing the whole quadratic family) and a fillied Julia set of a particular quadratic map. It is just the first indication of a deep relation between dynamical and parameter objects.

Note that Proposition 4.24 describes the real slice of the Mandelbrot set:

$$M \cap \mathbb{R} = [-2, 1/4].$$

172 6. PARAMETER PLANE (THE MANDELBROT SET)

28.3. Dependence of periodic points on c. What immediately catches the eye in the Mandelbrot set is the main cardioid C with a cusp at c = 1/4. The cardioid bounds a domain of parameter values c such that f_c has an attracting fixed point.

EXERCISE 6.2. Show that the main cardioid is given by the equation

$$c = \frac{1}{2}e^{2\pi i\theta} - \frac{1}{4}e^{4\pi i\theta}, \quad 0 \le \theta < 2\pi,$$

where $\sigma = e^{2\pi i\theta}$ is the multiplier of the neutral fixed point of f_c .

Let us now take a look at how periodic points move with parameter:

LEMMA 6.3. Let f_{\circ} has a cycle $\{\alpha_k\}_{k=0}^{p-1}$ of period p with multiplier $\sigma_{\circ} \neq 1$. Then for nearby c, the maps f_c have a cycle $\{\alpha_k(c)\}_{k=0}^{p-1}$ holomorphically depending on c. Its multiplier $\sigma(c)$ holomorphically depends on c as well.

PROOF. Consider an algebraic equation $f_c^p(z) = z$. For $c = c_0$ it has roots $z = \alpha_k, k = 0, \ldots, p-1$ (and maybe others). Since

$$\frac{d(f_c^p(z)-z)}{dz}\bigg|_{c=c_0,\,z=\alpha_k} = \sigma_0 - 1 \neq 0,$$

the Implicit Function Theorem yields the first assertion. The second assertion follows from the formula for the multiplier:

$$\sigma(c) = 2^p \prod_{k=0}^{p-1} \alpha_k(c).$$

Thus periodic points of f_c as functions of the parameter are algebraic functions branched at parabolic points only.

28.4. Hyperbolic components. A parameter value $c \in \mathbb{C}$ is called *hyperbolic/parabolic/Siegel* etc. if the corresponding quadratic polynomial f_c is such.

PROPOSITION 6.4 (Hyperbolic components). The set \mathcal{H} of hyperbolic parameter values is the union of $\mathbb{C} \smallsetminus M$ and some set of components of int M.

PROOF. By definition, $\mathbb{C} \setminus M \subset \mathcal{H}$. Aslo, the property to have an attracting cycle is stable (see Lemma 6.3), hence $\mathcal{H} \cap M \subset \operatorname{int} M$.

Take now some hyperbolic parameter $c_o \in M$ and let H_o be the component of int M containing c_o . Let us show that $H_o \subset \mathcal{H}$. The map f_o has an attracting cycle of some period p. By Theorem 4.38, this cycle contains a point α_0 such that

$$v_{pn}(c_{o}) \equiv f_{o}^{pn}(0) \to \alpha_{0} \text{ as } n \to \infty.$$

It is easy to see (Exercise!) that for nearby $c \in H$ we have:

$$v_{pn}(c) \equiv f_{o}^{pn}(0) \to \alpha_0(c) \text{ as } n \to \infty,$$

where $\alpha_0(c)$ is the holomorphically moving attracting periodic point of f_c (Lemma 6.3). But the sequence of polynomials $v_{pn}(c)$, $n = 0, 1, \ldots$, is normal in H (Exercise 6.1, (iv)). Hence it must converge in the *whole domain* H to some holomorphic function $\tilde{\alpha}(c)$ coinciding with $\alpha_0(c)$ near c_0 . By analytic continuation, $\tilde{\alpha}(c)$ is a periodic point of f_c with period dividing p. Moreover, the cycle of this point attracts the critical orbit persistently in H. It is impossible if this cycle is repelling somewhere. Indeed, a repelling cycles can only attract an orbit which eventually lands at it. This property is not locally persistent since otherwise it would hold for all $c \in \mathbb{C}$ (while it is violated, say, for c = 1).

If $\tilde{\alpha}(c)$ were parabolic for some $c \in H$, then it could be made repelling for a nearby parameter value. Thus $\tilde{\alpha}(c)$ is attracting for all $c \in H$, so that $H \subset \mathcal{H}$. \Box

COROLLARY 6.5. Neutral parameters lie on the boundary of M.

PROOF. Let c_{\circ} be a neutral parameter, i.e., the map f_{\circ} has a neutral cycle. This parameter can be perturbed to make the cycle attracting. If c_{\circ} belonged to int M then by Proposition 6.4 it would be hyperbolic itself – contradiction.

EXERCISE 6.6. (i) Any parameter $c \in \partial M$ can be approximated by superattracting parameters; (ii) Misiurewicz parameters form a countable dense subset of ∂M .

A component Λ of int M is called *hyperbolic* if it consists of hyperbolic parameter values. Otherwise Λ is called *queer*. The reason for the last term is that it is generally believed that *there are no queer components*. In fact, it is a central conjecture in contemporary holomorphic dynamics:

CONJECTURE 6.7 (Density of hyperbolicity). There are no queer components. Hyperbolic parameters are dense in \mathbb{C} .

Because of Exersice 6.6 (i), the second part of the conjecture would follow from the first one. It is sometimes referred to as *Fatou's Conjecture*.

28.5. Primitive and satellite hyperbolic components.

PROPOSITION 6.8. Let H be a hyperbolic component of period n of M, let $p/q \neq 0 \mod 1$, and let $r_{p/q} \in \partial H$ be a parabolic parameter with rotation number p/q. Then there is a hyperbolic component H' of period nq attached to H at $r_{p/q}$.

PROOF. We let $c_{\circ} \equiv r_{p/q}$, $f_{\circ} \equiv f_{c_{\circ}}$, and $g_c = f_c^n$. Let α_{\circ} be a parabolic fixed point for g_{\circ} Since $g'_{\circ}(\alpha_{\circ}) \neq 1$, nearby maps g_c have a fixed point α_c depending holomorphically on c. Making a change of variable $z \mapsto z - \alpha_c$, we obtain a holomorphic family of quadratic polynomials that fix 0; we keep the same notation f_c for this family and its *n*-fold iterate g_c .

By Corollary 4.46, g_{\circ} has q parabolic petals attached to 0 that are cyclically permutted by g_{\circ} . Hence near the origin we have:

 $g_{\mathsf{o}}^{q}(z) - z = b_{q+1}z^{q+1} + \dots, \quad b_{q+1} \neq 0.$

So 0 is a fixed point of multiplicity q + 1 for g_o , and hence nearby maps g_c have q + 1 simple fixed points. One of them is 0 which is also fixed by f_c . Others are permuted by f_c . In fact, they form a single cycle of order q since $f'_c(0) \approx e^{2\pi i p/q}$ and hence f_c cannot have small cycles of order less than q.

The multiplier σ_c of this cycle is a non-constant algebraic function of c equal to 1 at c_o . Hence there is a parameter domain attached to c_o in which our cycle is attracting. It is contianed in the desired hyperbolic component H'.

A hyperbolic component H' that was born from another hyperbolic component by the period increasing bifurcation described in Proposition is called *satellite*. All other hyperbolic components of M are called *primitive*. They appear as a result of a *saddle-node* bifurcation.

Parabolic points on ∂H with multiplier 1 are called the *roots* of H. (In fact, we will see below (Theorem 6.12) that any hyperbolic component has a single root.) In particular, the bifucation point $r_{p/q}$ is the root of the satellite component H'.

The type of the component can be easily recognized geometrically:

PROPOSITION 6.9. Satellite components are bounded by smooth curves, while primitive components have cusps at their roots.

29. Connectivity of M

29.1. Uniformization of $\mathbb{C} \setminus M$. In this section we will prove the first nontrivial result about the Mandelbrot set. The strategy of the proof is quite remarkable: it is based on the explicit uniformization of the complement $\mathbb{C} \smallsetminus M$ by $\mathbb{C} \smallsetminus \mathbb{D}$. Recall from Theorem 5.17 that for $c \in \mathbb{C} \setminus M$, we have a well-defined function

(29.1)
$$\Phi_M(c) := B_c(c),$$

where B_c is the Böttcher function for f_c extended to the domain Ω_c bounded by the critical figure-eight equipotential.

THEOREM 6.10. The Mandelbrot set M is connected. The function Φ_M conformally maps $\mathbb{C} \setminus M$ onto $\mathbb{C} \setminus \mathbb{D}$. Moreover, it is tangent to the identity at ∞ :

$$\Phi_M(c) \sim c \ as \ c \to \infty.$$

We immediately obtain the parameter analogue of Corollary 5.16:

COROLLARY 6.11. The Mandelbrot set has capacity 1.

29.2. Phase-parameter relation. Formula (29.1) reveals a remarkable relation between the dynamical and parameter planes of the quadratic family: The Riemann position $\Phi_M(c)$ of a parameter $c \in \mathbb{C} \setminus M$ coincides with the Böttcher position $B_c(c)$ of the corresponding critical value $c \in \mathbb{C} \setminus J(f_c)$.

Reacall from §32.1.2 that the polar coordinates of $B_c(z)$ are called the (dynamical) external coordinates of a point $z \in \Omega_c$. Similarly, the *(parameter) external coordiantes* of a point $c \in \mathbb{C} \setminus M$ are defined as the polar coordinates of $\Phi_M(c)$.

We see that the parameter external coordinates of a point $c \in \mathbb{C} \setminus M$ coincide with its dynamical external coordinates (in the f_c -dynamical plane).

Similarly to the dynamical situation (see $\S32.2$), we can now introduce parameter equipotentials $\mathcal{E}_{par}^r \equiv \mathcal{E}_{par}^t$ (where $t = \log r$) and parameter external rays $\mathcal{R}_{par}^{\theta}$ by pulling back round circles (of radius r) and radial rays (of angle θ) by means of Φ_M . We obtain two (non-singular) foliations in $\mathbb{C} \smallsetminus M$. We concude that

- For c ∈ R^θ_{par} we have c ∈ R^θ_c;
 For c ∈ E^r_{par} we have c ∈ E^r_c.

29.3. An elementary proof of Theorem 6.10. We will give two proofs of this theorem. The first proof is short and elementary. The second proof, though longer and more demanding, illuminates the deeper meaning of formula (29.1) and the idea of qc deformations.

It is based upon the explicit formula (32.2) for the Böttcher coordinate near ∞ ,

(29.2)
$$B_c(z) = \lim_{n \to \infty} (f_c^n(z))^{1/2^n},$$

where the root in the right-hand side is selected in such a way that it is tangent to the identity at ∞ . The statedy is to show that Φ_{MM} is a holomorphic branched covering of degree 1.

Step 1: analyticity. Let us consider the set $\Omega = \{(c, z) \in \mathbb{C}^2 : z \in \mathbb{C} \setminus K_c\}$. It is easy to see that this set is open. Indeed, for any c_o , there exist an R > 0 and $\epsilon > 0$ such that $|f_c(z)| > 2|z|$ for all $c \in \mathbb{D}(c_o, \epsilon)$ and |z| > R. Hence $\mathbb{C} \setminus \overline{\mathbb{D}}_R \subset \mathbb{C} \setminus K_c$ for all (c, z) as above.

Now, if $\zeta_{o} \in \mathbb{C} \setminus K_{o}$ then $f_{o}^{n}(\zeta_{o}) \in \mathbb{C} \setminus \overline{\mathbb{D}}_{R}$ for some *n*. By continuity, $f_{c}^{n}(\zeta) \in \mathbb{C} \setminus \overline{\mathbb{D}}_{R}$ for all (c, ζ) sufficiently close to (c_{o}, ζ_{o}) , and the openness follows.

We also see that the orbits of $\{f_c^n z\}_{n \in \mathbb{N}}$, $(c, z) \in \Omega$, escape to ∞ at a locally uniform rate, which implies that convergence in the Brolin formula (32.6),

$$G_c(z) = \lim_{n \to \infty} \frac{1}{2^n} \log |f_c^n z|,$$

is locally uniform on Ω . Hence $(c, z) \mapsto G_c(z)$ is a continuos function on Ω ,¹ so that, the set

$$\mathbf{\Omega}' = \{ (c, z) \in \mathbb{C}^2 : z \in \Omega_c \} = \{ (c, z) \in \mathbf{\Omega} : G_c(z) > G_c(0) \}$$

is also open. (Recall that $\Omega_c \subset D_c(\infty)$ is the maximal domain of analyticity of the Böttcher function B_c foliated by (non-singular) equipotentials, see §32.1).

But for the same reason, convergence in the Böttcher formula (29.2) is locally uniformly on Ω' . Hence the Böttcher function $(c, z) \mapsto B_c(z)$ is holomorphic on Ω' . We conclude that the function $\Phi_M : \mathbb{C} \setminus M \to \mathbb{C} \setminus \overline{\mathbb{D}}, \ \Phi_M(c) = B_c(c)$, is holomorphic on $\mathbb{C} \setminus M$.

Step 2: behavior at ∞ . Let $v_n = f_c^n(c)$. Then $v_{n+1} = v_n^2(1+O(1/v_n))$, so there is a function $\delta(v) = O(1/v)$ such that

$$((1 - \delta(v_n))v_n)^2 \le (1 - \delta(v_n))v_{n+1} < (1 + \delta(v_n))v_{n+1} \le ((1 + \delta(v_n))v_n)^2,$$

Iterating these estimates backwards, we see that

$$\sqrt[2^n]{v_n} = c(1 + O(1/c)) \quad \text{as } c \to \infty,$$

It follows that $\Phi_M(c) = c(1 + O(1/c)) \sim c$ as $c \to \infty$, so Φ_M exclusion holomorphically to ∞ , and is tangent to id at ∞ .

Step 3: properness. Let us show that the map $\Phi_M : \mathbb{C} \setminus M \to \mathbb{C} \setminus \overline{\mathbb{D}}$ is proper:

$$|\Phi_M(c)| \to 1 \text{ as } c \to \partial M.$$

Let us define $n(c) \in \mathbb{N} \cup \{\infty\}$ as the last moment n such that $v_n(c) \in \overline{\mathbb{D}}_3$. By Exercise 6.1(i), $n(c) = \infty$ iff $c \in M$. Moreover, $n(c) \to \infty$ as $c \to M$. Otherwise there would exist $N \in \mathbb{N}$ and a sequence $c_k \to c \in M$ such that $v_N(c_k) \in \mathbb{C} \setminus \overline{\mathbb{D}}_3$, implying that $v_N(c) \in \mathbb{C} \setminus \overline{\mathbb{D}}_3$ – contradiction.

Let us take a small neighborhood U of M such that $K_c \subset \overline{\mathbb{D}}_3$ for $c \in \overline{U}$ (equivalently, n(c) > 0 for $c \in \overline{U}$). Since the Green function is continuous on Ω ,

$$L := \sup\{G_c(z) : (c, z) \in \overline{U} \times \mathbb{T}_3\} < \infty.$$

¹It also follows that this function is *pluriharmonic* on Ω , i.e., its restrictions to onedimensional holomorphic curves in Ω are harmonic.

Since $z \mapsto G_c(z)$ is subharmonic on the whole plane \mathbb{C} for any c, by the Maximal Principle we have $G_c(z) \leq L$ for $(c, z) \in \overline{U} \times \overline{\mathbb{D}}_3$. Hence

$$G_c(c) = \frac{1}{n(c)} G_c(v_{n(c)}(c)) \le \frac{1}{n(c)} L \to 0 \text{ as } c \to M.$$

It follows that $|B_c(c)| = e^{G_c(c)} \to 1$ as $c \to M$ ($c \in \mathbb{C} \setminus M$) as was asserted.

Conclusion. Thus, the map $\Phi_M : \mathbb{C} \smallsetminus M \to \mathbb{C} \smallsetminus \overline{\mathbb{D}}$ is a branched covering, so that, it has a well-defined degree. But $\Phi_M^{-1}(\infty) = \{\infty\}$, and by Step 2, Φ_M has local degree 1 at ∞ . Hence deg $\Phi_M = 1$, and we are done.

29.4. Second proof.

29.4.1. Step 1: Qc deformation. The idea is to deform the map by moving around the Böttcher position of its critical value. To this end let us consider a two parameter family of diffeomorphisms $\psi_{\omega,q} : \mathbb{C} \setminus \mathbb{D} \to \mathbb{C} \setminus \mathbb{D}$ written in the polar coordinates as follows:

$$\psi = \psi_{\omega,q}(r,\theta) = (r^{\omega}, \theta + q\log r), \quad \omega > 0, \ q \in \mathbb{R}.$$

In terms of complex variable $a = re^{i\theta} \in \mathbb{C} \setminus \mathbb{D}$ and complex parameter $\lambda = \omega + iq$, Re $\lambda > 0$, this family can be expressed in the following concise form:

(29.3)
$$\psi_{\lambda}(a) = |a|^{\lambda - 1} a.$$

This family commutes with $f_0 : a \mapsto a^2$: $\psi(a^2) = \psi(a)^2$, and acts transitively on $\mathbb{C} \setminus \mathbb{D}$, i.e., for any a_* and a in $\mathbb{C} \setminus \mathbb{D}$, there exists a λ , such that $\psi_{\lambda}(a_*) = a$. (Note also that ψ_{λ} are automorphisms of $\mathbb{C} \setminus \mathbb{D}$ viewed as a multiplicative semigroup.)

Take now a quadratic polynomial $f_{\star} \equiv f_{c_{\star}}$ with $c_{\star} \in \mathbb{C} \smallsetminus M$. Let us consider its Böttcher function $\phi_{\star} : \Omega_{\star} \to \mathbb{C} \smallsetminus \mathbb{D}_{\star}$, where $\Omega_{\star} \equiv \Omega_{c_{\star}}$ is the complement of the figure eight equipotenial (see §??) and $\mathbb{D}_{\star} \equiv \mathbb{D}_{R_{\star}}$ is the corresponding round disk, $R_{\star} > 1$. Take the standard conformal structure σ on $\mathbb{C} \smallsetminus \mathbb{D}$ and pull it back by the composition $\psi_{\lambda} \circ \phi_{\star}$. We obtain a conformal structure $\mu = \mu_{\lambda}$ in Ω_{\star} . Since ψ_{λ} commute with f_0 while the Böttcher function conjugates f_{\star} to f_0 , the structure μ is invariant under f_{\star} .

Let us pull this structure back to the preimages of Ω_{\star} :

$$\mu^n \,|\, \Omega^n = (f^n_\star)^*(\mu),$$

where $\Omega_{\star}^{n} = f_{\star}^{-n} \Omega_{\star}$. Since μ is invariant on Ω_{\star} , the structures μ^{n+1} and μ^{n} coincide on Ω_{\star}^{n} , so that they are organized in a single conformal structure on $\cup \Omega_{\star}^{n} = \mathbb{C} \setminus J(f_{\star})$. Extend it to the Julia set $J(f_{\star})$ as the standard conformal structure.

We will keep notation $\mu \equiv \mu_{\lambda}$ for the conformal structure on \mathbb{C} we have just constructed. By construction, it is invariant under f_{\star} . Moreover, it has a bounded dilatation since holomorphic pullbacks preserve dilatation: $\|\mu_{\lambda}\|_{\infty} = \|(\psi_{\lambda})^*(\sigma)\|_{\infty} < 1$.

By the Measurable Riemann Mapping Theorem, there is a qc map $h_{\lambda} : (\mathbb{C}, 0) \to (\mathbb{C}, 0)$ such that $(h_{\lambda})_{\star}(\mu_{\lambda}) = \sigma$. By Corollary ??, $h_l a$ can be normalized so that it conjugates f_{λ} to a quadratic map $f_c \equiv f_{c(\lambda)} : z \mapsto z^2 + c(\lambda)$. Of course, the Julia set f_c is also Cantor, so that $c \in \mathbb{C} \setminus M$.

This family of quadratic polynomials is the desired qc deformation of f_{\star} .

29.4.2. Step 2: Analyticity. We have to check three propertices of the map $\Phi_M : \mathbb{C} \smallsetminus M \to \mathbb{C} \smallsetminus \mathbb{D}$: analyticity, surjectivity, and injectivity. Let us take them one by one.

It is obvious from formula (29.3) that the Beltrami differential

$$\nu_{\lambda} = (\psi_{\lambda})^*(\sigma) = \partial \psi_{\lambda} / \partial \psi_{\lambda}$$

depends holomorphically on λ . Hence the Beltrami differential $(f_{\star})^{*}(\nu_{\lambda})$ on Ω_{\star} also depends holomorphically on λ (see Exercise 2.46). Pulling it back by the iterates of f_{\star} and extending it in the standard way to J(f), we obtain by Lemma 2.43 a holomorphic family of Beltrami differentials μ_{λ} on \mathbb{C} . By Corollary 4.64, $c(\lambda)$ is holomorphic on λ as well.

29.4.3. Step 3: Surjectivity. Note that the map $\psi_{\lambda} \circ \phi_{\star} \circ h_{\lambda}^{-1}$ conformally conjugates the polynomial $f_c \equiv f_{c(\lambda)}$ near ∞ to $f_0 : z \mapsto z^2$. By Theorem 5.11, these properties determine uniquely the Böttcher map ϕ_c of f_c , so that $\phi_c = \psi_{\lambda} \circ \phi_{\star} \circ h_{\lambda}^{-1}$ with $c = c(\lambda)$. Since h_{λ} conjugates f_{\star} to f_c , we have: $h_{\lambda}(c_{\star}) = c$ and hence

$$\Phi_M(c) = \phi_c(c) = \psi_\lambda \circ \phi_\star(c_\star) = \psi_\lambda(a_\star),$$

where a_{\star} is the Böttcher position of the critical value of f_{\star} . Since the family $\{\psi_{\lambda}\}$ acts transitively on $\mathbb{C} \setminus \mathbb{D}$, any point $a \in \mathbb{C} \setminus \mathbb{D}$ can be relasized as $\Phi_M(c)$ for some $c = c(\lambda)$.

29.4.4. Step 4: Injectivity. We have to check that if

(29.4)
$$\phi_c(c) = a = \phi_{\tilde{c}}(\tilde{c})$$

for two parameter values c and \tilde{c} in $\mathbb{C} \setminus M$, then $c = \tilde{c}$. We let $f \equiv f_c$, $\phi \equiv \phi_c$, $f \equiv f_{\tilde{c}}$, $\tilde{\phi} \equiv \phi_{\tilde{c}}$. Similarly, we will mark with "tilde" the dynamical objects associated with \tilde{f} that naturally correspond to dynamical objects associated with f.

Let $R = \sqrt{|a|}$. Then the maps ϕ^{-1} and $\tilde{\phi}^{-1}$ map $\mathbb{C} \setminus \overline{\mathbb{D}}_R$ onto the domains $\Omega \equiv \Omega_c$ and $\tilde{\Omega} \equiv \Omega_{\tilde{c}}$ respectively. Moreover, they extend continuously to the boundary circle mapping it onto the boundary figures eight $\Gamma = \partial \Omega$ and $\tilde{\Gamma} = \partial \tilde{\Omega}$, and this extension if one-to-one except that

$$\phi^{-1}(\pm\sqrt{a}) = 0 = \tilde{\phi}^{-1}(\pm\sqrt{a}).$$

Hence the conformal map $h = \tilde{\phi}^{-1} \circ \phi : \Omega \to \tilde{\Omega}$ admits a homeomorphic extension to the closure of its domain:

$$h: (\operatorname{cl}(\Omega), 0) \to (\operatorname{cl}(\tilde{\Omega}), 0).$$

Consider a domain $\Omega^0 = f(\Omega)$ (exterior of the equipotential passing through c) and the complementary Jordan disk $\Delta^0 = \mathbb{C} \smallsetminus \Omega^0$. We will describe a hierarchical decomposition of Δ^0 into topological annuli A_i^n , $n = 1, \ldots, i = 1, 2, \ldots, 2^n$. Let $\Omega^n = f^{-n}\Omega^0$ (so that $\Omega \equiv \Omega^1$). The boundary $\partial\Omega^n$ consists of 2^{n-1} disjoint figures eight. The loops of these figures eight bound 2^n (closed) Jordan disks Δ_i^n . The map f conformally maps Δ_i^n onto some Δ_j^{n-1} , $n \ge 1$. Let $A_i^n = \Delta_i^n \cap cl(\Omega^{n+1})$. These are closed topological annuli each of which is bounded by a Jordan curve and a figure eight. They tile $\Delta^0 \smallsetminus J(f)$. The map f conformally maps A_i^n onto some A_i^{n-1} , $n \ge 1$.

Let us lift $h \equiv h_1$ to conformal maps $H_i: A_i^1 \to \tilde{A}_i^1$:

(29.5)
$$H_i | A_i^1 = (\tilde{f} | \tilde{A}_i^1)^{-1} \circ h \circ (f | A_i^1)$$

Since h is equivariant on the boundary of $\Omega^1 \setminus \Omega^0$, it matches with the H_i on $\partial \Delta_i^1$. Putting these maps together, we obtain an equivariant homeomorphism h_2 : $\operatorname{cl}(\Omega^2) \to \operatorname{cl}(\tilde{\Omega}^2)$ conformal in the complement of the figure eight Γ :

$$h_2(z) = \begin{cases} h(z), & z \in \Omega^1, \\ H_i(z), & z \in A_i^1. \end{cases}$$

Since smooth curves are removable (recall §16), h_2 is conformal in $\Omega^2 \setminus \{0\}$. Since isolated points are removable, h_2 is conformal in Ω^2 . Thus h admits an equivariant conformal extension to Ω^2 .

In the same way, h_2 can be lifted to four annuli A_i^2 . This gives an equivariant conformal extension of h to Ω^3 . Proceeding in this way, we will consecutively obtain an equivariant conformal extension of h to all the domains Ω^n and hence to their union $\cup \Omega^n = \mathbb{C} \setminus J(f)$.

Since the Julia set J(f) is removable (Theorem 2.75), this map admits a conformal extension through J(f). Thus, f and \tilde{f} are conformally equivalent, and hence $c = \tilde{c}$.

This completes the second proof of Theorem 6.10.

30. The Multiplier Theorem

30.1. Statement. Let us pick a favorite hyperbolic component H of the Mandelbrot set M. For $c \in H$, the polynomial f_c has a unique attracting cycle $\alpha_c = \{\alpha_k(c)\}_{k=0}^{p-1}$ of period p. By Lemma 6.3, the multiplier $\lambda(c)$ of this cycle holomorphically depends on c, so that we obtain a holomorphic map $\lambda : H \to \mathbb{D}$. It is remarkable that this map gives an explicit uniformization of H by the unit disk:

THEOREM 6.12. The multiplier map $\lambda : H \to \mathbb{D}$ is a conformal isomorphism.

This theorem is in many respects analogous to Theorem 6.10 on connectivity of the Mandelbrot set. The latter gives an explicit dynamical uniformization of $\mathbb{C} \setminus M$; the former gives the one for the hyperbolic component. The ideas of the proofs are also similar.

We already know that λ is holomorphic, so we need to verify that it is surjective and injective. The first statement is easy:

EXERCISE 6.13. The multiplier map $\lambda : H \to \mathbb{D}$ is proper and hence surjective. In particular, H contains a superattracting parameter value.

30.2. Qc deformation. Let $Z \subset H$ be the set of superattracting parameter values in H. Take some point $c_0 \in H \setminus Z$, and let $\lambda_0 \in \mathbb{D}^*$ be the multiplier of the corresponding attracting cycle. We will produce a qc deformation of $f_* \equiv f_{c_0}$ by deforming the associated fundamental torus.

30.2.1. Fundamental torus. Take a little topological disk $D = \mathbb{D}(a_0, \epsilon)$ around the attracting periodic point a_0 of f_* . It is invariant under $g_0 \equiv f_*^p$ and the quotient of D under the action of f_* is a conformal torus T_0 . Its fundamental group has one marked generater corresponding to a little Jordan curve around α_0 .

By the Linearization Theorem (??), the action of g_0 on D is conformally equivalent to the linear action of $\zeta \mapsto \lambda_0 \zeta$ on \mathbb{D}^* . Hence the partially marked torus T_0 is conformally equivalent to $\mathbb{T}^2_{\lambda_0}$, so that λ_0 is the modulus of T_0 (see §1.6.2).

Let us select a family of deformations $\psi_{\lambda} : \mathbb{T}^2_{\lambda_0} \to \mathbb{T}^2_{\lambda}$ of T_{λ_0} to nearby tori. For instance, ψ_{λ} can be chosen to be linear in the logarithmic coordinates $(x, y) = \log \zeta$,

$$\tau = \log \lambda:$$

$$x + y\tau_0 \mapsto x + y\tau; \quad x \in \mathbb{R}, \ y \ge 0$$

This gives us a complex one-parameter family of Beltrami differentials $\nu_{\lambda} = \psi_{\lambda}^*(\sigma)$ on $T_0 \approx \mathbb{T}^2_{\lambda_0}$ (in what follows we identify T_0 with $\mathbb{T}^2_{\lambda_0}$).

EXERCISE 6.14. Calculate ν_{λ} explicitly (for the linear deformation).

30.2.2. Qc deformation of f_* . We can lift ν_{λ} to the disk D and then pull it back by iterates of f_* . This gives us a family of f_* -invariant Beltrami differentials μ_{λ} on the attracting basin of α . These Beltrami differentials have a bounded dilatation since the pull-backs under holomorphic maps preserve dilatation. Extend the μ_{λ} by 0 outside the attracting basin (keeping the notation). We obtain a family of measurable f_* -invariant conformal structures μ_{λ} on the Riemann sphere. Solving the Beltrami equation $(h_{\lambda})_*(\mu_{\lambda}) = \sigma$ (with an appropriately normalization) we obtain a qc deformation of f_* (see Corollary 4.65):

(30.1)
$$f_{c(\lambda)} = h_{\lambda} \circ f_* \circ h_{\lambda}^{-1} : z \mapsto z^2 + c(\lambda).$$

Moreover, note that this deformation is conformal on the basin of ∞ .

Let us show that the multiplier of the attracting fixed point of $f_{c(\lambda)}$ is equal to λ . Consider the torus T_{λ} associated with the attracting cycle of $f_{c(\lambda)}$. Then h_{λ} descends to a homeomorphism $H_{\lambda}: T_0 \to T_{\lambda}$ such that $(H_{\lambda})_*(\nu_{\lambda}) = \sigma$. Since $(\psi_{\lambda})_*(\nu) = \sigma$ as well, the map

$$\psi_{\lambda} \circ H_{\lambda}^{-1} : T_{\lambda} \to \mathbb{T}_{\lambda}^{2}$$

is conformal. Hence the partially marked torus T_{λ} has the same modulus as \mathbb{T}^2_{λ} , which is λ . But as we know, this modulus is equal to the multiplier of the corresponding attracting cycle.

This deformation immediately leads to the following important conclusion:

LEMMA 6.15. All maps f_c , $c \in H \setminus Z$, are qc equivalent (and the conjugacy is conformal on the basin of ∞). Moreover, card Z = 1.

PROOF. Take some $c_0 \in H \setminus Z$. By Proposition 2.39, the deformation parameter $c(\lambda)$ in (30.1) depends continuously on λ . Hence $c : \lambda \mapsto c(\lambda)$ is the local right inverse to the multiplier function. But holomorphic functions do not have continuous right inverses near their critical points. Consequently, c_0 is not a critical point of the multiplier function λ and, moreover, c is the local inverse to λ . It follows that any c near c_0 can be represented as $c(\lambda)$, and hence f_c is qc equivalent to f_{c_0} .

Let us decompose the domain $H \smallsetminus Z$ into the union of disjoint qc classes (with conformal conjugacy on the basin of ∞). We have just shown that each qc class in this decomposition is open. Since $H \smallsetminus Z$ is connected, it consists of a single qc class.

Furthermore, we have shown that λ does not have critical points in $H \smallsetminus Z$. Hence $\lambda : H \smallsetminus Z \to \mathbb{D}^*$ is an unbranched covering. By the Riemann-Hurwitz formula (for the trivial case of unbranched coverings), the Euler characteristic of $H \searrow Z$ is equal to 0, i.e., $1 - \operatorname{card} Z = 0$.

Thus, every hyperbolic component H contains a unique superattracting parameter value c_H . It is called the *center* of H. We let $H^* = H \setminus \{c_H\}$.
30.3. Injectivity. The following lemma will complete the proof of the Multiplier Theorem:

LEMMA 6.16. Consider two parameter values c and \tilde{c} in $H \setminus Z$. If $\lambda(c) = \lambda(\tilde{c})$ then the quadratic maps f_c and $f_{\tilde{c}}$ are conformally equivalent on \mathbb{C} .

The idea is to turn the qc conjugacy from Lemma 6.15 into a conformal conjugacy. To this end we need to modify the conjugacy on the basin of the attracting cycle. Let us start with the component D_0 of the basin containing 0.

30.4. Second Proof of the Multiplier Theorem. Let us give an alternative proof of the following key lemma that immediately implies the Multiplier Theorem. It is one of the first manifistations of the Rigidity Phenomenon and one more application of the Surgery techniques (gluing a map from model pieces).

LEMMA 6.17. Let f_c and $f_{\tilde{c}}$ be two hyperbolic quadratic maps. Assume that the Böttcher conjugacy

$$h: D_c(\infty) \to D_{\tilde{c}}(\infty), \quad h = B_{\tilde{c}}^{-1} \circ B_c,$$

extends to a homeomorphism $D_c(\infty) \cup J_c \to D_{\tilde{c}}(\infty) \cup J_{\tilde{c}}$. If the attracting cycles of these maps have the same multiplier then $c = \tilde{c}$.

PROOF. We let $f = f_c$, $\tilde{f} = f_{\tilde{c}}$; $K = K_c$, $\tilde{K} = K_{\tilde{c}}$, etc.

Let D_0 be the immediate basin of the attracting point α_0 (for f) containing 0. We know that it is a Jordan disk (Corollary 4.55), so the Riemann mapping $\phi : (D_0, \alpha_0) \to (\mathbb{D}, 0)$ extends to a homeomorphism $\operatorname{cl} D_0 \to \overline{\mathbb{D}}$ (denoted by ϕ as well). Moreover, ϕ (appropriately normalized) conjugates the return map f^p : $\operatorname{cl} D_0 \to \operatorname{cl} D_0$ to the Blaschke map g from Proposition 4.56.

Since α_0 and $\tilde{\alpha}_0$ have the same multipliers, the normalizaed Riemann mapping $\tilde{\phi} : \operatorname{cl} D_0 \to \overline{\mathbb{D}}$ conjugates $\tilde{f}^p | \operatorname{cl} \tilde{D}_0$ to the same Blaschke map g. Hence the composition $h_0 := \tilde{\phi}^{-1} \circ \phi : \operatorname{cl} D_0 \to \operatorname{cl} \tilde{D}_0$ conjugates $f^p | \operatorname{cl} D_0$ to $\tilde{f}^p | \operatorname{cl} \tilde{D}_0$.

This map h_0 continuously matches on ∂D_0 with the Böttcher conjugacy

$$h: D(\infty) \cup J \to D(\infty) \cup J$$

Indeed, the composition $h^{-1} \circ h_0 : \partial D_0 \to \partial D_0$ commutes with with $f^p | \partial D_0$, and hence $h^{-1} \circ h_0 | \partial D_0 = \text{id by Proposition 4.89}$.

We can now easily lift h_0 to all other components of int K. Consider a component D. Since int K is equal to the basin of α (Theorem 4.53), there is $n = n_D \in \mathbb{Z}_+$ such that f^n homeomorphically maps cl D onto cl D_0 . Let f^{-n} : cl $D \to$ cl D_0 stand for the inverse map. Then we let

(30.2)
$$h_D = \tilde{f}^n \circ h_0 \circ f^{-n} : \operatorname{cl} D \to \operatorname{cl} \tilde{D}.$$

Obviously, this map conjugates $f^p | \operatorname{cl} D$ to $\tilde{f}^p | \operatorname{cl} \tilde{D}$.

Moreover, h_D matches continuously on ∂D with h. Indeed, since h is a conjugacy on the whole Julia set, we have

$$h|\,\partial D = \tilde{f}^n \circ (h|\,\partial D_0) \circ f^{-n} \,:\, \partial D \to \partial \tilde{D}.$$

Comparing this with (30.2), taking into account that $h | \partial D_0 = h_0$, yields $h | D = h_D$.

Thus, we have extended h conformally and equivariantly to all the components D_i of int K. Since diam $D_i \to 0$ and diam $\tilde{D}_i \to 0$, this extension is a global homeomorphism (Exersice ??).

By the dynamical qc removability of J (Lemma 4.62), this homeomorphism is quasiconformal. Moreover it is conformal outside the Julia set, while area J = 0. By Weyl's Lemma, it is conformal.

COROLLARY 6.18. Let C be a hyperbolic component of the Mandelbrot set M, and let $c, \tilde{c} \in C$. If $\sigma(c) = \sigma(\tilde{c})$ then $c = \tilde{c}$.

30.5. Internal angles. For $c \in \overline{H}$, $\arg \lambda(c)$ is called the *internal angle* of c.

31. Structural stability

31.1. Statement of the result. A map $f_{\circ} : z \mapsto z^2 + c_{\circ}$ (and the corresponding parameter $c_{\circ} \in \mathbb{C}$) is called *structurally stable* if for any $c \in \mathbb{C}$ sufficiently close to c_{\circ} , the map f_c is topologically conjugate to f_{\circ} , and moreover, the conjugacy $h_c : \mathbb{C} \to \mathbb{C}$ can be selected continuously in c (in the uniform topology). By definition, the set of structurally stable parameters is open. In this section we will prove that it is dense:

THEOREM 6.19. The set of structurally unstable parameters is equal to the boundary of the Mandelbrot set together with the centers of hyperbolic components. Hence the set of structurally stable parameters is dense in \mathbb{C} . Moreover, any structurally stable map f_{\circ} is quasi-conformally conjugate to all nearby maps f_c .

Notice that parameters $c_b ase \in \partial M$ are obviously unstable since the Julia set J_o is connected, while the Julia sets J_c for nearby $c \in \mathbb{C} \setminus M$ are disconnected. The centers of hyperbolic components are also unstable since the topological dynamics near a superattracting cycle is different from the topological dynamics near an attracting cycle (the grand orbits on the basin of attraction are discrete in the latter case and are not in the former).

The proof of stability of other parameters will occupy $\S{31.2} - \S{31.5}$. The desired conjugacies will be constructed as equivariant holomorphic motions.

A holomorphic motion $h_c: X_{o} \to X_c$ of a set $X \subset \mathbb{C}$ is called *equivariant* if

(31.1)
$$h_c(f_o(z)) = f_c(h_c(z))$$

whenever both points z and $f_{o}(z)$ belong to X_{o} . If the X_{c} are f_{c} -invariant, this just means that the maps h_{c} conjugate $f_{o}|X_{o}$ to $f_{c}|X_{c}$. (Of course, we can apply this terminology not only to the quadratic family).

Notice that the equivariance property (31.1) means that the associated lamination (see §17.1) is invariant under the map

(31.2)
$$\mathbf{f}: (\lambda, z) \mapsto (\lambda, f_{\lambda}(z)).$$

Since by the Second λ -lemma, holomorphic motions are automatically quasiconformal in the dynamical variable, the last assertion of Theorem 6.19 will follow automatically.

31.2. *J*-stability. Let us first show that the Julia set J_c moves holomorpically outside the boundary of M. (Strictly speaking, this step is not needed for the proof of Theorem 6.19 given below, but it gives a good illustration of the method.)

A map $f_{\circ} : z \mapsto z^2 + c_{\circ}$ (and the corresponding parameter $c_{\circ} \in \mathbb{C}$) is called *J-stable* if for any $c \in \mathbb{C}$ sufficiently close to c_{\circ} , the map $f_c | J_c$ is topologically conjugate to $f_{\circ} | J_{\circ}$, and moreover the conjugacy $h_c : J_{\circ} \to J_c$ depends continuously on c. THEOREM 6.20. The set of J-stable parameters is equal to $\mathbb{C} \setminus \partial M$ and hence is dense in \mathbb{C} . Moreover, the corresponding conjugacies $h_c : J_o \to J_c$ form a holomorphic motion of the Julia set over the component of $\mathbb{C} \setminus \partial M$ containing \circ .

PROOF. Let C be the component of int M containing c_{\circ} . By Corollary 6.5, C does not contain neutral parameters, and hence all periodic points are persistently hyperbolic over C, either repelling or attracting. Hence they depend holomorphically on $c \in C$. Since C is simply connected (Exercise 6.1 (iii)), these holomorphic functions $c \mapsto \alpha(c)$ are single valued. Moreover, they cannot collide since collisions could occur only at parabolic parameters. Thus, they provide us with a holomorphic motion h_c : Per_o \rightarrow Per_c of the set of periodic points.

This holomorphic motion is equivariant. Indeed, if

$$c \mapsto \alpha(c) = h_c(\alpha)$$

is a holomorphically moving periodic point then $c \mapsto f_c(\alpha(c))$ is also a holomorphically moving periodic point. Hence $f_c(\alpha(c)) = h_c(f_o(\alpha))$ and we obtain:

$$f_c(h_c(\alpha)) = f_c(\alpha(c)) = h_c(f_o(\alpha)).$$

By the First λ -lemma (3.1), this holomorphic motion extends to a continuous equivariant holomorphic motion of the closure of periodic points, which contains the Julia set. Moreover, this motion is automatically continuous in both variables (λ, z) , and hence provides us with a family of topological conjugacies between J_{o} and J_{c} continuously depending on c.

EXERCISE 6.21. An equivariant holomorphic motion of the Julia set is unique.

31.3. Böttcher motion: connected case. In this section, we will show that the basin of infinity, $D_c(\infty)$, moves bi-holomorphically over any component of int M.

PROPOSITION 6.22. Let C be a component of int M with a base point \circ . Then there exists an equivariant bi-holomorphic motion $h_c: D_{\circ}(\infty) \to D_c(\infty)$ of the basin of infinity over C.

PROOF. Let $\phi_c : D_c(\infty) \to \mathbb{C} \setminus \overline{\mathbb{D}}$ be the Böttcher-Riemann uniformization of the basin of infinity (see Theorem 5.13). It is a holomorphic function in two variables on the domain $\{(c, z) : c \in C, z \in D_c(\infty)\}$ (see Step 1 of §29.3). It follows that $h_c = \phi_c^{-1} \circ \phi_o$ is a bi-holomorphic motion of D_c over C. Since the ϕ_c conjugate f_c to $z \mapsto z^2$, this motion is equivariant. \Box

EXERCISE 6.23. Show that an equivariant bi-holomorphic motion of the basin of ∞ over C is unique.

Now the first λ -lemma implies:

COROLLARY 6.24. For any component C of int M, there is a unique equivariant holomorphic motion $h_c : \overline{D}_c(\infty) \to \overline{D}_c(\infty)$ over C which is bi-holomorphic on $D_c(\infty)$.

If Q is a queer component then $\mathbb{C} = \overline{D}_c(\infty)$ for any $c \in Q$, and so, we obtain the Structural Stability Theorem in this case:

COROLLARY 6.25. For a queer component Q of int M, there is a unique equivariant holomorphic motion $h_c : \mathbb{C} \to \mathbb{C}$ over Q which is bi-holomorphic on $D_c(\infty)$. Hence all parameters $c \in H$ are structurally stable. **31.4.** Motion of an attracting basin. For a hyperbolic parameter c, let α_c stand for the corresponding attracting cycle, and let $D(\alpha_c)$ be its basin.

PROPOSITION 6.26. Let H be a hyperbolic component of int M, and let $c_o \in H^*$. Then there is an equivariant smooth holomorphic motion of the attracting basin $D(\boldsymbol{\alpha}_c)$ over some neighborhood of c_o .

To prove this assertion, we need three simple lemmas. The first one is concerned with local extension of smooth motions.

LEMMA 6.27. Let us consider a compact set $Q \subset \mathbb{C}$ and a smooth holomorphic motion h_{λ} of a neighborhood U of Q over a parameter domain (Λ, \circ) . Then there is a smooth holomorphic motion H_{λ} of the whole complex plane \mathbb{C} over some neighborhood Λ' of \circ whose restriction to Q coincides with h_{λ} .

PROOF. We can certainly assume that U is compact. Take a smooth cut-off function $\eta : \mathbb{C} \to \mathbb{R}$ supported in U such that $\eta | Q \equiv 1$, and let

$$H_{\lambda} = \eta h_{\lambda} + (1 - \eta) \operatorname{id}$$

Clearly H is smooth in both variables, holomorphic in λ , coinsides with h on Q and with the identity outside U. As $H_0 = \mathrm{id}, H_\lambda : \mathbb{C} \to \mathbb{C}$ is a diffeomorphism for λ sufficiently close to \circ , and we are done.

The second lemma is concerned with lifts of holomorphic motions.

LEMMA 6.28. Let $h_{\lambda}: V_{o} \to V_{\lambda}$ be a holomorphic motion of a domain $V_{o} \subset \mathbb{C}$ over a simply connected parameter domain Λ . Let $f_{\lambda}: U_{\lambda} \to V_{\lambda}$ be a holomorphic family of proper maps with critical points c_{λ}^{k} such that the critical values $v_{\lambda}^{k} = f_{\lambda}(c_{\lambda}^{k})$ form orbits of h_{λ} .² Then h_{λ} uniquely lifts to a holomorphic motion $H_{\lambda}: U_{o} \to U_{\lambda}$ such that

(31.3)
$$f_{\lambda} \circ H_{\lambda} = h_{\lambda} \circ f_{\circ}.$$

PROOF. Notice that (31.3) means that the lamination associated with the motion **H** is the pullback of the lamination associated with the motion **h** under the map **f** (31.2). Clearly, such a pullback unique if exists.

Let us take any regular value $\zeta_{o} = f_{o}(z_{o}) \in V_{o}$, and let $\phi(\lambda) = h_{\lambda}(\zeta_{o})$ be its orbit. We would like to lift this orbit to a desired orbit of z_{o} , so we are looking for a holomorphic solution $z = \psi(\lambda)$ of an equation

(31.4)
$$f_{\lambda}(z) = \phi(\lambda)$$

with $\psi(z_0) = \zeta_0$. Since $\phi(\lambda)$ is a regular point of f_{λ} for any $\lambda \in \Lambda$, the Implicit Function Theorem implies that near any point (λ', z') satisfying (31.4), it admits a unique local analytic solution $z = \psi(\lambda)$. Since the maps f_{λ} are proper, this continuation along any path compactly contained in Λ cannot escape the domain U_{λ} . Since Λ is simply connected, $\psi(\lambda)$ extends to the whole domain Λ as a single valued holomorphic function.

Two different orbits $\lambda \mapsto \psi(\lambda)$ obtained in this way do not collide, for (31.4) would have two different solutions near the collision point. Hence they form a holomorphic motion of $V_{o} \smallsetminus \{v_{o}^{k}\}$ over Λ . By the First λ -lemma, this motion extends to the whole domain V_{o} .

²In particular, any holomorphic family of univalent maps $f_{\lambda} : U_{\lambda} \to V_{\lambda}$ is allowed.

Finally,

$$f_{\lambda}(H_{\lambda}(z_{o})) = f_{\lambda}(\psi(\lambda)) = \phi(\lambda) = h_{\lambda}(\zeta_{o}) = h_{\lambda}(f_{o}(z_{o}))$$

holds for any point $z_o \in U_o$ except perhaps finitely many exceptions (preimages of the critical values of f_o). By continuity, it holds for all $z_o \in U_o$.

The last lemma is concerned with dependence of the linearizing coordinate (the Königs function) on parameters

LEMMA 6.29. Let f_{λ} be a holomorphic family of germs near the origin over a parameter domain Λ such that 0 is a simply attracting point. Then the normalized linearizing coordinate ϕ_{λ} depends holomorphically on λ .

PROOF. The linearizing coordinate ϕ_{λ} is given by an explicit Königs formula (29.3):

(31.5)
$$\phi_{\lambda}(z) = \lim_{n \to \infty} \sigma_{\lambda}^{-n} f_{\lambda}^{n}, \quad \text{where } \sigma_{\lambda} = f_{\lambda}'(0).$$

Since analyticity is a local property, we need to verify the assertion near an arbitrary parameter $\lambda_{o} \in \Lambda$. There exist $\epsilon > 0$ and $\rho < 1$ such that $f_{o}(\mathbb{D}_{\epsilon}) \Subset \mathbb{D}_{\rho\epsilon}$. Then the same is true for λ in some neighborhood Λ' of λ_{o} . By the Schwarz Lemma, the orbits $\{f_{\lambda}^{n}(z)\}_{n=0}^{\infty}$ of points $z \in \mathbb{D}_{\epsilon}$ converge to 0 at a uniformly exponential rate: $|f_{\lambda}^{n}(z)| \leq \rho^{n}$ for $\lambda \in \Lambda'$. This implies (by examining the proof of (29.3)) that convergence in (31.5) is uniform on $\Lambda' \times \mathbb{D}_{\epsilon}$. Hence $\phi_{\lambda}(z)$ is holomorphic on $\Lambda' \times \mathbb{D}_{\epsilon}$.

Proof of Proposition 6.26. Let $\alpha_c = \{f_c^k(\alpha)\}_{k=0}^{p-1}$ be the attracting cycle of f_c , and let us consider the maps f_c^p near their fixed points α_c . Lemma 6.29 implies that there is a neighborhood $H' \subset H^*$ of c_0 and an $\epsilon > 0$ such that the inverse linearizing coordinate $\phi_c^{-1}(z)$ for f_c^p is holomorphic on $\Lambda' \times \mathbb{D}_{\epsilon}$. Let $V_c = \phi_c^{-1}(\mathbb{D}_{\epsilon}) \ni \alpha_c$, and let us consider a fundamental annulus $A_c = \operatorname{cl}(V_c \setminus f_c(V_c))$.

By Theorem 4.38, the critical orbit $\operatorname{orb}_c(0)$ must $\operatorname{cross} A_c$. By adjusting ϵ and shrinking H' if needed, we can ensure that it does not $\operatorname{cross} \partial A_c$. Then it crosses A_c at a single point $v_n(c) = f_c^n(0) \in \operatorname{int} A_c$, where $n \in \mathbb{N}$ is independent of c. Its position in the linearizing coordinate, $a_c = \phi_c(v_n(c)) \in \mathbb{A}(\epsilon, \sigma_c \epsilon) \equiv \mathbb{A}_c$, depends holomorphically on c (here σ_c is the multiplier of $\boldsymbol{\alpha}_c$).

Let $Q_c = \partial A_c \cup \{a_c\}$. Let us define a smooth equivariant holomorphic motion **h** of a small neighborhood of Q_c over H' as follows: $h_c = \text{id}$ near the outer boundary of A_c , $h_c : z \mapsto \sigma_c z / \sigma_0$ near the inner boundary of A_c , and $h_c : z \mapsto a_c z / a_0$ near a_c . By Lemma 6.27, this motion extends to a smooth motion of the whole plane over some neighborhood of c_0 (we will keep the same notation H' for this neighborhood). Let us restrict the motion to the fundamental annulus A_c (keeping the same notation h_c for it). By Lemma 6.28 (in the simple case when there are no critical points), this motion can be first extended to the forward orbit of A_c , (providing us with an equivariant holomorphic motion of \mathbb{D}_{ϵ}). Then we can transfer it using the linearizing coordinates to a holomorphic motion of V_c , then extend it to an invariant neighborhood $\mathbf{V}_c = \bigcup_{k=0}^{p-1} f_c^k(V_c)$ of $\boldsymbol{\alpha}$, and finally we can use Lemma 6.28 to pull this motion back to all preimages of \mathbf{V}_c (the assumption of Lemma 6.28 on the critical values is secured by the property that a_c is an orbit of the motion **h**). It provides us with the desired equivariant holomorphic motion of the basin $D(\boldsymbol{\alpha}_c)$. \Box

COROLLARY 6.30. Let H be a hyperbolic component of int M, and let $c_o \in H^*$. Then there is an equivariant holomorphic motion of the whole plane \mathbb{C} over some neighborhood of c_o . Hence all parameters $c \in H^*$ are structurally stable.

PROOF. Since for $\mathbb{C} = \operatorname{cl}(D_c(\infty) \cup D(\boldsymbol{\alpha}_c))$ for $c \in H$, Propositions 6.22 and 6.26, together with the First λ -lemma yield the desired.

31.5. Böttcher motion: Cantor case. Let us finally deal with the complement of M.

PROPOSITION 6.31. Let $c_{\circ} \in \mathbb{C} \setminus M$. Then there is an equivariant smooth holomorphic motion of the basin of infinity, $D_{c}(\infty)$, over some neighborhood of c_{\circ} .

The proof is similar to the one given in the attracting case, using the Böttcher coordinate in place of the linearizing coordinate. To implement it, we need a rotationally equivariant Extension Lemma:

LEMMA 6.32. Let R > r > 1 and let $z \in \mathbb{A}(r, R)$. Let ϕ be a holomorphic function on a domain (Λ, \circ) with $\phi(\circ) = z$. Then there is a smooth holomorphic motion H_{λ} of the whole complex plane \mathbb{C} over some neighborhood Λ' of \circ such that

(i) $H_{\lambda}(z) = \phi(\lambda);$

(ii) $H_{\lambda} = \text{id } on \mathbb{C} \smallsetminus \mathbb{A}(r, R) = \text{id};$

(iii) The H_{λ} commute with the rotation group $\zeta \mapsto e^{i\theta}\zeta$.

PROOF. Let $\tau(\lambda) = \phi(\lambda)/z$, and let $h_{\lambda}(\zeta) = \tau(\lambda)\zeta$. This motion satisfies requirements (i) and (iii). To make it satisfy (ii) as well, we will use a smooth cut-off function $\phi : \mathbb{R} \to \mathbb{R}$ supported on a small neighborhood of |z|. Then the motion

$$H_{\lambda}(\zeta) = \phi(|\zeta|) h_{\lambda}(\zeta) + (1 - \phi(|\zeta|))\zeta$$

satisfies all the requirements.

Proof of Proposition 6.31. Let us consider the Böttcher coordinate B_c of f_c near ∞ . Since it depends holomorphically on c, there is a neighborhood $U \subset \mathbb{C} \setminus M$ of c_o and an R > 1 such that the function $(c, z) \mapsto B_c^{-1}(z)$ is holomorphic on $U \times (\mathbb{C} \setminus \overline{D}_R)$.

Let $V_c = B_c^{-1}(\mathbb{C} \setminus \overline{\mathbb{D}}_R)$. By adjusting R and U if necessary, we can ensure that the $\operatorname{orb}_c(0)$ does not cross the boundary of the fundamental annulus $A_c = V_c \setminus f_c(V_c)$. Then there is a unique n such that $v_n(c) = f_c^n(0) \in \operatorname{int} A_c$. Let us mark the corresponding point $a_c = B_c(v_n(c))$ in the annulus $\mathbb{A} = \mathbb{A}(R, R^2)$.

Applying lemma 6.32, we find a rotationally equivariant holomorphic motion $H_c : \mathbb{A} \to \mathbb{A}$ such that $H_c(a_0) = a_c$ and $H_c = \text{id on } \partial \mathbb{A}$. Let us show that

EXERCISE 6.33. Show that this holomorphic motion extends to a holomorphic motion $H_c: \mathbb{C} \setminus \mathbb{D}_R \to \mathbb{C} \setminus \mathbb{D}_R$ commuting with $z \mapsto z^2$.

Let us now transfer H_c by means of the Böttcher coordinate to a holomorphic motion $h_c: V_c \to V_c$, $h_c = B_c^{-1} \circ H_c \circ B_o$. This motion is equivarinat, $B_c \circ f_o = f_c \circ B_c$, and has $v_n(c)$ as one of its orbits. By Lemma 6.28, it can be lifted to a holomorphic motion of $f_c^{-1}(V_c)$ that has $v_{n-1}(c)$ as its orbit. Moreover, by the uniqueness of the lift, it coincides on V_o with the original motion h_c , which implies that it is equivariant. Then we can lift it further to $f^{-2}(V_c)$, and so on: in this way we will exhaust the whole basin of ∞ . \Box Since $\mathbb{C} = \overline{D}_c(\infty)$ for $c \in \mathbb{C} \setminus M$, Proposition 6.31 (together with the First λ -lemma) yields:

COROLLARY 6.34. Let $c_o \in \mathbb{C} \setminus M$. Then there is an equivariant holomorphic motion of the whole plane \mathbb{C} over some neighborhood of c_o . Hence all parameters $c \in \mathbb{C} \setminus M$ are structurally stable.

Corollaries 6.25, 6.30 and 6.34 cover all types of components of $\mathbb{C} \setminus \partial M$, and together prove the Structural Stability Theorem (6.19).

31.6. Invariant line fields and queer components.

31.6.1. Definition. Informally speaking, a line field on \mathbb{C} is a family of tangent lines $l(z) \in T_z \mathbb{C}$ depending measurably on $z \in \mathbb{C}$.

Here is a precise definition. Any line $l \in \mathbb{C}$ passing through the origin is uniquely represented by a pair of centrally symmetric points $e^{\pm 2\pi i\theta} \in \mathbb{T}$ in the unit circle, or by a single number

(31.6)
$$\nu = e^{4\pi i\theta} \in \mathbb{T}, \quad \theta \in \mathbb{R}/(\mathbb{Z}/2).$$

The space of these lines form, by definition, the one-dimensional projective line \mathbb{PR}^1 , and (31.6) provides us with its parametrization by the angular coordinate (and shows that $\mathbb{PR}^1 \approx \mathbb{T}$).

Let us now consider the projective tangent bundle over \mathbb{C} ,

$$\operatorname{PT}(\mathbb{C}) = \mathbb{C} \times \mathbb{PR}^1$$

parametrized by $\mathbb{C} \times (\mathbb{R}/(\mathbb{Z}/2))$. A line field on \mathbb{C} is a measurable section of $\operatorname{PT}(\mathbb{C})$ defined on some set $X \subset \mathbb{C}$ of positive area called its *(measurable) support*. In terms of the angular coordinate, we obtain a measurable function $X \to \mathbb{R}/(\mathbb{Z}/2)$, $z \mapsto \theta(z)$.³ In the circular coordinate ν , we obtain a measurable function $X \to \mathbb{R}/(\mathbb{Z}/2)$. In what follows, we will always extend ν by 0 to the whole plane.

EXERCISE 6.35. Show that a line field on a Riemann surface S is given by a Beltrami differential $\nu(z)\frac{d\bar{z}}{dz}$ with $|\nu(z)| \in \{0,1\}$.

A line field on a set $J \subset \mathbb{C}$ is a line field on \mathbb{C} whose support is contained in J. If such a line field exists (with a non-empty support) then area J > 0.

A line field is called *invariant* (under a holomorphic map f) if it is invariant under the natural action of f on the projective line bundle: l(fz) = Df(z)l(z), or in the angular coordinate, $\theta(fz) = \theta(z) + \arg f'(z)$, or in the Beltrami coordinate, $f^*\nu = \nu$ (where the pullback is understood in the sense of Beltrami differentials).

If an invariant line field l is supported on a set X then we can pull it back by the dynamics to obtain an invariant line field supported on the set $\tilde{X} = \bigcup_{n=0}^{\infty} f^{-n}(X)$.⁴

Hence we can assume in the first place that l is supported on a completely invariant set: this will be our standing assumption.

186

³As always, a measurable function is considered up to an arbitrary change on null-sets.

⁴The pullback would fail at the critical point but we can always remove its grand orbit (as any other completely invariant null-set) from \tilde{X} .

31.6.2. Existence criterion.

PROPOSITION 6.36. Let Q be a queer component of int M. Then any map f_c , $c \in Q$, has an invariant line field on its Julia set. In particular, area $J(f_c) > 0$.

Vice versa, if f_c has an invariant line field on its Julia set then c belongs to a queer component of int M.

PROOF. Take some $c_0 \in Q$. By Corollary 6.25, there is an equivariant holomorphic motion h_c over (Q, c_0) which is bi-holomorphic on $D_c(\infty)$. Let us consider the corresponding Belrtami differentials $\mu_c = \bar{\partial}h_c/\partial h_c$, $c \in Q$. Each μ_c vanishes on $D_0(\infty)$, however $\mu_c \neq 0$ for $c \neq c_0$ (for otherwise, by Weyl's Lemma the map h_c would be affine, contrary to the fact the quadratic maps f_c and f_o are not affinely conjugate). Hence area($\supp \mu_c$) > 0 for any $c \neq c_0$, and all the more, area $J_0 > 0$. Moreover, since μ_c is f_0 -invariant, the normalized Beltrami differential $\nu_c = \mu_c/|\mu_c|$ (where we let $\nu_c = 0$ outside $\supp \mu_c$) is also f_0 -invariant, and hence determines an invariant line field on the Julia set J_0 .

Vice versa, assume f_{\circ} has an invariant line field on J_{\circ} given by an invariant Beltrami differential ν_{\circ} . For any $\lambda \in \mathbb{D}$, the Beltrami differential $\lambda\nu_{\circ}$ is also finvariant. Let $h_{\lambda} : (\mathbb{C}, 0) \to (\mathbb{C}, 0)$ be the solution of the corresponding Beltrami equation tangent to the identity at infinity. Then the map $h_{\lambda} \circ f_{\circ} \circ h_{\lambda}^{-1}$ is a quadratic polynomial $f_{\sigma(\lambda)} : z \mapsto z^2 + \sigma(\lambda)$ (see §26.1.2). By Corollary 4.65, the map $\sigma : \mathbb{D} \to \mathbb{C}$ is holomorphic. Since the line field is non-trivial, it is not identically constant. Hence its image covers a neighborhood of c_{\circ} contained in int M. So, it is contained in some component of int M. By Theorem ??, this component cannot be hyperbolic, so it must be queer.

Thus, Fatou's Conjecture (32) is equivalent to the following one:

CONJECTURE 6.37 (No Invariant Line Fields). No quadratic polynomial has an invariant line field on its Julia set.

31.6.3. Uniqueness and ergodicity. As a line field l(z) is rotated by angle $2\pi\alpha$ with $\alpha \in \mathbb{R}/(\mathbb{Z}/2)$, the corresponding Beltrami differential is multiplied by $\lambda = e^{4\pi i\alpha} \in \mathbb{T}$. Of course, if the original line field was *f*-invariant then so is the rotated one.

LEMMA 6.38. A quadratic polynomial can have at most one, up to rotation, invariant line field on its Julia set.

This will follow from the ergodicity of the action of f on the support of any invariant line field. Recall that a map $f: X \to X$ of a measure space is called *ergodic* if X cannot be decomposed into a disjoint unnion of two invariant (and hence completely invariant) subsets of positive measure. Equivalently, there are no non-constant measurable functions $\phi: X \to \mathbb{R}$ invariant under f, i.e., such that $\phi \circ f = \phi$.

LEMMA 6.39. Let f be a quadratic polynomial, and let l(z) be an invariant line field on J(f). Then the action of f on supp l is ergodic.

PROOF. Assume that supp l admits a disjoint decomposition $X_1 \sqcup X_2$ into two measurable invariant subsets of positive measure. Then the restriction of lto these sets gives us two invariant line fields l_i with disjoint supports. Let ν_i be the corresponding Beltrami differentials. Then we can consider a complex twoparameter family of Beltrami differentials $\nu_{\lambda} = \lambda_1 \nu_1 + \lambda_2 \nu_2$, where $\lambda = (\lambda_1, \lambda_2) \in \mathbb{D}^2$. Since $\|\nu_{\lambda}\|_{\infty} < 1$ for each λ , we can solve the corresponding Beltrami equations and obtain a two parameter family of qc maps $h_{\lambda} : (\mathbb{C}, 0) \to (\mathbb{C}, 0)$ tangent to the identity at infinity. Then the maps $h_{\lambda} \circ f_{\circ} \circ h_{\lambda}^{-1}$ form a family of quadratic polynomials $f_{\sigma(\lambda)} : z \mapsto z^2 + \sigma(\lambda)$ (see §26.1.2).

By Proposition 2.39, the map $\sigma : \mathbb{D}^2 \to \mathbb{C}$ we have obtained this way is continuous (in fact, by Corollary 4.65, it is holomorphic). Hence it cannot be injective: there exist $\lambda \neq \kappa$ in \mathbb{D}^2 such that $\sigma(\lambda) = \sigma(\kappa)$. Then the map $\phi = h_{\kappa}^{-1} \circ h_{\lambda}$ commutes with f_{o} . But the only conformal automorphism of $D_{o}(\infty)$ commuting with f_{o} is the identity (see Exercise 5.12). Hence $h_{\lambda} = h_{\kappa}$ implying that $\lambda = \kappa$ – contradiction.

Proof of Lemma 6.38. Assume we have two invariant line fields given by Beltrami differentials ν_i . Let $X_i = \operatorname{supp} \nu_i$. Notice that due to our convention, both differences, $X_1 \smallsetminus X_2$ and $X_2 \smallsetminus X_1$, are completely invariant sets. If $\operatorname{area}(X_2 \smallsetminus X_1) > 0$ then an invariant Beltrami differential ν which is equal to ν_1 on X_1 and is equal to ν_2 on $X_2 \smallsetminus X_1$ has a non-ergodic support, contradicting Lemma 6.39. Hence $\operatorname{area}(X_2 \smallsetminus X_1) = 0$, and for the same reason $\operatorname{area}(X_1 \smallsetminus X_2) = 0$, so that the set $Y = X_1 \cap X_2$ can be taken as a measurable support of both differentials.

By Lemma 6.39, f acts ergodically on Y. But the ratio ν_2/ν_1 is an invariant function on Y. By ergodicity, it is equal to const a.e. on Y, and we are done. \Box

31.6.4. Dynamical uniformization of queer components. We can now construct a dynamical uniformization of any queer component Q by a Beltrami disk. (Compare with the uniformizations of hyperbolic components of $\mathbb{C} \setminus \partial M$ given by Theorems 6.10 and 6.12.)

For a base map f_{\circ} , let us select an invariant line field on J_{\circ} given by an f-invariant Beltrami differential ν_{\circ} . Then the Beltrami disk $\{\lambda\nu_{\circ}\}_{\lambda\in\mathbb{D}}$ generates a holomorphic family of quadratic polynomials $f_{\sigma(\lambda)}: z \mapsto z^2 + \sigma(\lambda)$ (see the proof of the second part of Lemma 6.36). This is the desired uniformization:

PROPOSITION 6.40. The map $\sigma : (\mathbb{D}, 0) \to (Q, \circ)$ is the Riemann mapping.

PROOF. The map σ is a holomorphic embedding for the same reason as in the proof of Lemma 6.39. Let us show that it is surjective. Let $c \in Q$. By Corollary 6.25, the map f_c is conjugate to f_o by a qc homeomorphism h_c which is conformal outside J_o . Let $\mu_c = \bar{\partial}h_c/\partial h_c$ be the Beltrami differential of h_c , and let $\nu_c = \mu_c/|\mu_c|$. Since the latter differential determines an invariant line field on J_o , Lemma 6.38 yields:

$\operatorname{supp} \mu_c = \operatorname{supp} \nu_c = \operatorname{supp} \nu_o.$

Since the differential μ_c is f_{\circ} -invariant, the ratio μ_c/ν_{\circ} is an f-invariant function. By ergodicity, it is const a.e., so that $\mu_c = \lambda \nu_c$ for some $\lambda \in \mathbb{D}$. It follows that $c = \sigma(\lambda)$, and we are done.

31.7. Quasiconformal classification of the quadratic maps. We can now give a complete classification of the quadratic maps up to qc conjugacy:

PROPOSITION 6.41. Any qc class in the parameter plane \mathbb{C} of the quadratic family is one on the following list:

• the complement of the Mandelbrot set;

32. NOTES

- a hyperbolic component of int M punctured at the center;
- a queer component of int M;
- the center of a hyperbolic component;
- a single point of the boundary of M.

The first three types of maps are *deformable*, the last two are *qc rigid*.

PROOF. By the Structural Stability Theorem (6.19), each of the above listed sets is contained in some qc class. What we need to show that they belong to different qc classes.

Assume it is not the case: let c_o and c be two parameters in different sets but in the same qc class. Then the quadratic polynomials f_o and f_c are conjugate by a qc map h. Let $\mu = \bar{\partial}h/\partial h$ be the Beltrami differential of h, and let $r = 1/||\mu||_{\infty}$. Let us consider the Beltrami disk $\{\lambda \mu : |\lambda| < r\}$ and the corresponding qc deformation

$$f_{\sigma(\lambda)}: z \mapsto z^2 + \sigma(\lambda)$$

of f_{\circ} (see Corollary 4.65). Then $\sigma : \mathbb{D}_r \to \mathbb{C}$ is a hololmorphic map such that $\sigma(0) = c_{\circ}$ and $\sigma(1) = c$. In particular, it is not identically constant and hence its image U is a domain in \mathbb{C} . But U is not contained in a single component of int M, so it must intersect ∂M , and hence it must intersect $\mathbb{C} \setminus M$. Thus, U contains quadratic maps of both dichotomy types: with connected as well as Cantor Julia sets, which is impossible as all the maps in U are topologically conjugate. \Box

32. Notes

Fatou's Conjecture on Density of Hyperbolicity () is an interpretation of several remarks that Fatou made on page of $[\mathbf{F}]$. First, Fatou observes that hyperbolicity is preserved under perturbations. Then he conjectures that any rational map can be approximated by a stable one (compare with Theorem 6.19. He also suggests that unstable maps form some kind of algebraic set: apparently, he did not give a real thought to this issue. The general theory of hyperbolicity and structural stability developed in 1960's by Smale, Anosov, and many other people, greatly clarified how Fatou's remarks should be interpreted.

Theorem 6.10 on the connectivity of M was proved by Douady and Hubbard ??, and independently by Sibony (unpublished: see a remark in ??), in early 1980's. The elementary proof given in §29.3 reproduces the original argument from ??.

CHAPTER 7

Combinatorics of external rays

39. Dynamical ray portraits

39.1. Motivaing problems. Consider a quadratic polynomial $f = f_c$ with connected Julia set. As we know (§??), its basin of infinity is uniformized by the Böttcher map $\phi : D_f(\infty) \to \mathbb{C} \setminus \mathbb{D}$, which conjugates f to $z \mapsto z^2$. If the Julia set was locally connected then by the Carathéodory theorem the inverse map would ϕ^{-1} extend continuously to the unit circle \mathbb{T} . This would give a representation of $f \mid J(f)$ as a quotient of the the doubling map $\theta \mapsto 2\theta \mod 1$ of the circle $\mathbb{R}/Z \approx \mathbb{T}$. This observation immediately leads to the following problems:

1) Describe explicitly equivalence realtions on the circle corresponding to all possible Julia sets;

2) Study the problem of local conectivity of the Julia sets.

It turns out that the first problem can be addressed in a comprehensive way. The second problem is very delicate. However, even non-locally connected examples can be partially treated due to the fact that many external rays always land at some points of the Julia set. This is the main theme of the following discussion.

39.2. Landing of rational rays. We say that an external ray \mathcal{R}^{θ} lands at some point z of the Julia set if $\mathcal{R}^{\theta}(t) \to z$ as $t \to 0$. Two rays $\mathcal{R}^{\theta/2}$ and $\mathcal{R}^{\theta/2+1/2}$ will be called "preimages" of the ray \mathcal{R}^{θ} . Obviously, if some ray lands, then its image and both its preimages land as well.

An external ray \mathcal{R}^{θ} is called *rational* if $\theta \in \mathbb{Q}$, and *irrational* otherwise. Dynamically the rational rays are characterized by the property of being either periodic or preperiodic:

EXERCISE 7.1. Let $\mathcal{R} = \mathcal{R}^{\theta}$.

points z_n and z_{n+1} . Then $g(\gamma_n) = \gamma_{n-1}$.

a) If θ is irraional then the rays $f^n(\mathcal{R})$, $n = 0, 1, \ldots$, are all distinct.

Assume θ is rational: $\theta = q/p$, where q and p are mutually prime. Then

(i) If p is odd then \mathcal{R} is periodic: there exists an l such that $f^{l}(\mathcal{R}) = \mathcal{R}$.

(ii) If p is even then \mathcal{R} is preperiodic: there are l and r > 0 such that $f^r(\mathcal{R})$ is a periodic ray of period l, while the rays $f^k(\mathcal{R})$, k = 0, 1, ..., r - 1, are not periodic. How to calculate l and r?

THEOREM 7.2. Let f be a polynomial with connected Julia set. Then any

periodic ray $\mathcal{R} = \mathcal{R}_f^{p/q}$ lands at some repelling or parabolic point of f. PROOF. Without loss of generality we can assume that the ray \mathcal{R} is periodic and hence invariant under some iterate $g = f^l$. Let $d = 2^l$. Consider a sequence of points $z_n = \mathcal{R}(1/d^n)$, and let γ_n be the sequence of arcs on \mathcal{R} bounded by the Endow the basin $D = D_f(\infty)$ with the hyperbolic metric ρ . Since $g: D \to D$ is a covering map, it locally preserves ρ . Hence the hyperbolic length of the arcs γ_n are all equal to some L.

But all the rays accumulate on the Julia set as $t \to 0$. By the relation between the hyperbolic and Euclidean metrics (Lemma 1.97), the Euclidean length of these arcs goes to 0 as $n \to \infty$. Hence the limit set of the sequence $\{z_n\}$ is a connected set consisting of the fixed points of g. Since g has only finitely many fixed points, this limit set consists of a single fixed point β . It follows that the ray \mathcal{R} lands at $\beta \in J(f)$ (compare with the proof of Theorem 4.51).

Since $\beta \in J(f)$, it can be either repelling, or parabolic, or Cremer. But the latter case is excluded by the Necklace Lemma 4.52.

39.3. Inverse Theorem: periodic points are landing points. It is much harder to show that, vice versa, any repelling or parabolic point is a landing point of at least one ray:

39.3.1. Repelling case.

THEOREM 7.3. Let f be a polynomial with connected Julia set. Then any repelling point a is the landing point of at least one periodic ray.

PROOF. Replacing f with its iterate, we can assume without loss of generality that a is a fixed point. We will consider the linearizing coordinates ϕ and ψ near a based on the discussion and notation of §31. Let \tilde{U}_i be the components of $\psi^{-1}(D(\infty))$. These components are permuted by the map $g: z \mapsto \lambda z$. The main step of the proof is to show that each component \tilde{U}_i is periodic under this action. It will be done by studying the rate of escape of hyperbolic geodesics in \tilde{U} to infinity.

Let us consider the Green function $G : \mathbb{C} \to \mathbb{R}_{\geq 0}$ of f (see §32.3). Recall that it is a continuous subharmonic function satisfying the functional equation G(fz) = dG(z). Let us lift it to the dynamical plane of g. We obtain a continuous subharmonic function $\tilde{G} = G \circ \psi$ on \mathbb{C} satisfying the functional equation $\tilde{G}(\lambda z) =$ $d\tilde{G}(z)$. Letting $M_n = \max_{|z|=|\lambda|^n} \tilde{G}(z), M \equiv M_1$, we see that $M_n \leq d^n M$.

By Lemma 5.10, the domains \tilde{U}_i are simply connected and the restrictins ψ : $\tilde{U}_i \to D(\infty)$ are the universal coverings. Let us fix one of these domains, $\tilde{U} = \tilde{U}_i$ and endow it with the hyperbolic metric ρ . Let γ be the hyperbolic geodesic in \tilde{U} that begins at a point $u_0 \in \mathbb{T}_1$ and goes to ∞ (i.e., γ is the pullback of a straight ray in $\mathbb{C} \setminus \overline{\mathbb{D}}$ by the $B \circ \psi : U \to \mathbb{C} \setminus \overline{\mathbb{D}}$). This geodesic must cross all the circles $\mathbb{T}_{|\lambda|^n}$; let u_n stand for the first crossing point, and let $\rho_n = \rho(u_0, u_n)$.

By Excersice 5.19 and the above estimate on M_n , we have:

(39.1)
$$\rho_n = \log \frac{G(u_n)}{G(u_0)} \le \log \frac{M_n}{G(u_0)} = n \log d + O(1).$$

Thus, the points u_n escape to infinity no faster than at linear rate.

The above discussion is generally applied, no matter whether the domain U is periodic under g or not. Assuming now that it is *aperiodic*, we will argue that the points u_n must escape to infinity at a superlinear rate.

If \tilde{U} is aperiodic then the action of the cyclic group $\langle g \rangle$ on the orbit of \tilde{U} is faithful, so that, \tilde{U} is embedded into the quotient torus $\mathbb{T}^2 = \mathbb{C}^* / \langle g \rangle$ under the natural projection $\mathbb{C}^* \to \mathbb{T}^2$. Let $W \subset \mathbb{T}^2$ be the image under this embedding.

It is now convenient to make the logarithmic change of variable on \mathbb{C}^* that turns it to the cylinder \mathbb{C}/\mathbb{Z} . Then the complex scaling g becomes the translation $z \mapsto z + \tau$, where $\tau = \log i\lambda/2\pi \mod \mathbb{Z}$, the circles $\mathbb{T}_{|\lambda|^n}$ become the circles $\mathbf{T}_n =$ $\{v : \operatorname{Im} v = n \operatorname{Im} \tau\}$, U becomes a domain $\mathbf{U} \subset \mathbb{C}/\mathbb{Z}$, the geodesic γ in U becomes a geodesic γ in \mathbf{U} , and points u_n turn into points $\mathbf{u}_n \in \gamma \cap \mathbf{T}_n$. Let us parametrize γ by the length parameter $t \in \mathbb{R}_{>0}$ so that $\gamma(0) = \mathbf{u}_0$.

Let us endow the cylinder \mathbb{C}/\mathbb{Z} and the torus \mathbb{T}^2 with the flat Euclidean metric so that the natural projection $\pi : \mathbb{C}/\mathbb{Z} \to \mathbb{T}^2$ is locally isometric. Then

(39.2)
$$\operatorname{dist}(\boldsymbol{\gamma}(t), \partial \mathbf{U}) \to \mathbf{0} \quad \text{as} \quad \mathbf{t} \to \infty.$$

Otherwise there would exist $\epsilon > 0$ and a sequence of points $\mathbf{x}_n \in \mathbb{C}/\mathbb{Z}$ such that Im $\mathbf{x}_{n+1} > \text{Im } \mathbf{x}_n + 2\epsilon$ and $D_n \equiv D(\mathbf{x}_n, \epsilon) \subset \mathbf{U}$. Since $\pi : \mathbf{U} \to \mathbb{T}^2$ is a locally isometric embedding, the images $\pi(D_n)$ would be disjoint ϵ -disks in \mathbb{T}^2 , which is impossible by compactness of \mathbb{T}^2 .

Let
$$d\rho(\boldsymbol{\gamma}(t)) = \sigma(t)|dz|$$
. By Lemma 1.97 and (39.2),

(39.3)
$$\sigma(t) \approx \frac{1}{\operatorname{dist}(\boldsymbol{\gamma}(t), \partial \mathbf{U})} \to \infty \quad \text{as} \quad t \to \infty$$

Let l_n stands for Euclidean length of the arc of γ bounded by \mathbf{u}_{n+1} and \mathbf{u}_n , and let $\sigma_n = \inf \sigma(t)$ on that arc. Then

$$\rho(\mathbf{u}_{n+1}, \mathbf{u}_n) \ge \sigma_n l_n \ge \sigma_n \operatorname{Im} \tau.$$

Hence

$$\rho(\mathbf{u}_{n+1},\mathbf{u}_0) \ge \operatorname{Im} \tau \sum_{k=0}^{n-1} \sigma_k,$$

and by (39.3)

$$\frac{1}{n}\rho(\mathbf{u}_{n+1},\mathbf{u}_0)\to\infty$$
 as $n\to\infty$,

contradicting (39.1).

Thus, the domain \tilde{U} is periodic under the action of g with some period q. Hence the image W of \tilde{U} in \mathbb{T}^2 is the quotient of \tilde{U} by the cyclic $\langle g^q \rangle$. It follows that it is conformally equivalent to either an annulus $\mathbb{A}(1,r)$ or to the punctured disk \mathbb{D}^* depending on whether g^q is hyperbolic or parabolic. In fact, the first option is realized. Indeed,

$$\Psi \equiv B \circ \psi : \cup \tilde{U}_i \to \mathbb{C} \smallsetminus \bar{\mathbb{D}}$$

is a covering map conjugating g to $z \mapsto z^d$. Hence it semi-conjugates $g^q : \tilde{U} \to \tilde{U}$ to $z \mapsto z^{d^q}$. Since \tilde{U} is simply-connected, Ψ lifts to a conformal isomorphism $\hat{\Psi} : \tilde{U} \to \mathbb{H}$ conjugating g^q to $\tau : \zeta \to d^q \zeta$. But the latter is a hyperbolic map of \mathbb{H} , so that, $W \approx \mathbb{H}/\langle \tau \rangle$ is an annulus (with modulus $\pi/(q \log d)$).

To complete the proof, let us consider the simple closed hyperbolic geodesic Γ in W. It lifts to a hyperbolic geodesic $\tilde{\Gamma}$ in \tilde{U} invariant under the action of $\langle g^q \rangle$. Let δ be a fundamental arc on $\tilde{\Gamma}$ bounded by some point u and $g^{-q}u$. Then $g^{-qn}(\delta) \to 0$ as $n \to \infty$, so that, $\tilde{\Gamma}$ lands at 0 (in "negative" time).

Since $\psi : \tilde{U} \to D(\infty)$ is a covering map semi-conjugating g^q to f^q , $\psi(\tilde{\Gamma})$ is a hyperbolic geodesic in $D(\infty)$ invariant under f^q and hence escaping to infinity in positive time. But hyperbolic geodesics in $D(\infty)$ escaping to infinity are exactly the external rays of f.

Finally, $\psi(\Gamma)$ lands at a in negative time since ψ is continuous at 0.

PROBLEM 7.4. a) Modify the above proof to show that there are only finitely many domains \tilde{U}_i .

b) Conclude that there are only finitely many rays landing at any repelling point, and all of these rays are periodic.

39.3.2. Parabolic case. A point $a \in K(f)$ is called *dividing* if $K(f) \setminus \{a\}$ is disconnected.

EXERCISE 7.5. Assume K(f) is connected. Show that a repelling or parabolic periodic point $a \in K(f)$ is dividing if and only if there are more than one ray landing at a.

39.3.3. Cantor case.

PROPOSITION 7.6. Let $f_c : z \mapsto z^2 + c$ be a quadratic polynomial with Cantor Julia set, i.e., $c \in \mathbb{C} \setminus M$. Then any external ray \mathcal{R}^{θ} that does not crash at a precritical point lands at some point of $J(f_c)$.

39.4. Rotation sets on the circle. We will now briefly deviate from the complex dynamics to study "rotation cycles" on the circle.

The oriented cicle $\mathbb{T} \approx \mathbb{R}/\mathbb{Z}$ is certainly not ordered, but it rather *cyclically* ordered. Namely, any finite subset $\Theta \subset \mathbb{T}$ is ordered, $\Theta = (\theta_1 \dots, \theta_n)$, up to cyclic permutation of its points and this order is compatible with the inclusions of the sets. We say that a tuple of points $(\theta_1, \dots, \theta_n)$ of \mathbb{T} is correctly ordered if their order is compatible with the cyclic order of \mathbb{T} .

Given two points $\theta_1, \theta_2 \in \mathbb{T}$, we let (θ_1, θ_2) be the (open) arc of \mathbb{T} that begins at θ_1 and ends at θ_2 (which makes sense since \mathbb{T} is oriented).

A tuple of two points (θ_1, θ_2) of Θ is called neighbors in Θ if the corresponding arc $\overline{(\theta_1, \theta_2)}$ does not contain other points of Θ . (Note that this relation is asymmetric.)

Given a subset Θ and an injection $g : \Theta \to \mathbb{T}$, we say that g is monotone if it preserves the cyclic order of finite subsets of Θ (i.e., if $(\theta_1, \ldots, \theta_n)$ is a correctly ordered tuple of points of Θ , then the tuple of points $(g(\theta_1), \ldots, g(\theta_n))$ is also correctly ordered).

EXERCISE 7.7. Show that g is monotone on a finite set $\Theta \subset \mathbb{T}$ iff it maps any tuple of neighbors in Θ to a tuple of neighbors in $g(\Theta)$.

Monotone bijections $g: \Theta \to \Theta$ are called *rotations of* Θ , and Θ is correspondingly called a *rotation set* for g.

Any finite rotation set $\Theta \subset \mathbb{T}$ has a well defined rational rotation number $p/q \in \mathbb{Q}/\mathbb{Z}$. Namely, take a point $\theta \in \Theta$ and let q be the period of θ , while p be the number of points in the orb(θ) contained in the semi-open arc $\overline{[\theta, g(\theta)]}$.

EXERCISE 7.8. Check that q and p are independent of the choice of θ .

We will now analyze rotation cycles for the doubling map $q: \theta \mapsto 2\theta \mod 1$.

LEMMA 7.9. Let Θ be a rotation cycle for g with rotation number p/q. Then complementary arcs to Θ (counted according to the action of g starting with the shortest one) have lengths $2^{k-1}/(2^q-1)$, $k = 1, \ldots q$.

PROOF. Let $\omega_i = \overline{(\theta_i, \kappa_i)}$ be the complementary arcs to Θ , where $g(\theta_i) = \theta_{i+1}$, $g(\kappa_i) = \kappa_{i+1}$ $(i \in \mathbb{Z}/q\mathbb{Z})$. If some ω_i is shorter that half-circle then g maps it

homeomorphically onto the arc ω_{i+1} of length $|\omega_{i+1}| = 2|\omega_i|$. So, if all the arcs ω_i were shorter than half-circle then we would arrive at the basic logical contradiction:

$$1 = \sum_{i \in \mathbb{Z}/q\mathbb{Z}} |\omega_{i+1}| = 2 \sum_{i \in \mathbb{Z}/q\mathbb{Z}} |\omega_i| = 2.$$

Thus, one of the arcs ω_i must be longer than half-circle. Let us call it ω_0 , and let $|\omega_0| = (1 + \epsilon)/2$. This arc is the union of the half-circle $\xi = \overline{[\kappa'_0, \theta_0]}$ and the arc $\eta = \overline{(\theta_0, \kappa'_0)}$ of length $\epsilon/2$, where $\kappa'_0 = \kappa_0 - 1/2$ is the point symmetric with κ_0 . Moreover, under g, the arc ξ is bijectively mapped onto the whole circle \mathbb{T} , while η is homeomorphically mapped onto $\overline{(\theta_1, \kappa_1)} = \omega_1$. We see that $|\omega_1| = \epsilon$.

Since each arc ω_i , $i = 1, \ldots, q - 1$, is shorter than half-circle, it is mapped homeomorphically onto the arc ω_{i+1} , and $|\omega_{i+1}| = 2|\omega_i|$. Hence $|\omega_i| = 2^{i-1}\epsilon$, $i = 1, \ldots, q$. Since q = 0 in $\mathbb{Z}/q\mathbb{Z}$, we obtain the equation:

$$\frac{1+\epsilon}{2} = |\omega_0| = |\omega_q| = 2^{q-1}\epsilon$$

which gives us the desired value of ϵ .

The arc ω_0 of $\mathbb{T} \setminus \Theta$ which is longer than half-circle is called *critical*. The shortest arc ω_1 is called *characteristic*.

PROPOSITION 7.10. For the doubling map $g: \theta \mapsto 2\theta$ on \mathbb{T} and any rational $p/q \in \mathbb{Q}/\mathbb{Z}$, there exists a unique rotation cycle $\Theta_{p/q} \subset \mathbb{T}$ with rotation number p/q.

PROOF. Let Θ be a rotation cycle with rotation number p/q. Let us consider its characteristic arc $\xi_1 = \overline{(\theta, \kappa)}$. Since κ is the neighbor of θ , we have: $\kappa = 2^l \theta \mod 1$, where l = 1/p in $\mathbb{Z}/q\mathbb{Z}$. On the other hand, $\kappa = \theta + 1/(2^q - 1)$ by Lemma 7.9. Hence

(39.4)
$$(2^l - 1) \theta \equiv 1/(2^q - 1) \mod 2$$

Since θ is g-periodic with period q, $2^{q}\theta = \theta \mod 1$, so that, $\theta = t/(2^{q}-1)$. Plugging it into (39.4), we come up with the equation

(39.5)
$$(2^l - 1) t \equiv 1 \mod 2^q - 1.$$

Since l and q are mutually prime, so are $2^{l} - 1$ and $2^{q} - 1$, and hence (39.5) has a unique solution mod $2^{q} - 1$. This prove uniqueness of the rotation cycle.

Going backwards, we take the solution of (39.4), let $\kappa = g^l \theta$ and $\xi_1 = \overline{(\theta, \kappa)}$. Then $\xi_2 = g^l(\xi)$ is the arc of length $2/(2^q - 1)$ adjacent to ξ_1 ; $\xi_3 = g^{2l}(\xi)$ is the arc of length $2/(2^q - 1)$ adjacent to ξ_2 , etc., up to the arc $\xi_q = g^{lq}(\xi)$ of length $2^{q-1}/(2^q - 1) > 1/2$. Since the total length of these arcs is equal to 1, their closures tile the whole circle \mathbb{T} , so that, $\Theta = \operatorname{orb}(\theta)$ is a rotation cycle of g^l with rotation number 1/q. Since $pl = 1 \mod q$, we have: $g|\Theta = (g^l)^p|\Theta$, and hence $g|\Theta$ has rotation number p/q.

EXERCISE 7.11. Derive the uniqueness part of the last proposition directly from Lemma 7.9, without finding the rotation cycle explicitly.

EXERCISE 7.12. Analyse the structure of rotation sets on \mathbb{T} with irrational rotation number. Prove that for any $\eta \in \mathbb{R}/\mathbb{Z}$, there exists a unique closed rotation set Θ_{η} on \mathbb{T} with rotation number η .

7. COMBINATORICS OF EXTERNAL RAYS

EXERCISE 7.13. Analyze the structure of rotation cycles for the map $g_d: \theta \mapsto d\theta$. Prove that there are at most d-1 cycles with a given rotation number.

39.5. Fixed points and their combinatorial rotaion number.

39.5.1. Combinatorial rotation number. Let us now consider a polynomial f of degree d with connected Julia set. Let a be its repelling or parabolic fixed point, and let $\mathcal{R}_i \equiv \mathcal{R}^{\theta_i}$ be the rays landing at a. The set of angles $\Theta(a) = \{\theta_i\} \subset \mathbb{T}$ is called the ray portrait of a.

LEMMA 7.14. The ray portrait $\Theta(a)$ is a rotation set for the map $g_d: \theta \mapsto d\theta$.

PROOF. Let S_i be the complementary sectors to the rays, i.e., the connected components of $\mathbb{C} \setminus \bigcup \mathcal{R}_i$. Each sector S is bounded by a pair of rays $(\mathcal{R}, \mathcal{R}')$, which can be ordered so that \mathcal{R} is positively oriented rel S. Thus, we can order the rays, $(\mathcal{R}_0, \mathcal{R}_1, \ldots, \mathcal{R}_{n-1})$, so that, \mathcal{R}_i and \mathcal{R}_{i+1} are neighbors, and this ordering is well defined up to cyclic permutation of the rays. So, the rays are cyclically ordered.

The map f preserves this cyclic order. Indeed, it is a local homeomorphism near a, and hence it permutes the local sectors. It follows that neighboring rays are mapped by f to neighboring ones, which is equivalent to preserving the cyclic order (compare Exercise 7.7).

But the cylic order of the rays \mathcal{R}^{θ_i} coincides with the cyclic order of their slopes θ_i at ∞ . Since these slopes are permuted by the map g_d , the conclusion follows. \Box

So, we have a well defined *combinatorial rotation number* of f at a. In fact, it can be defined in terms of any periodic curve landing at a:

EXERCISE 7.15. Let γ be a periodic path in $\mathbb{C} \setminus K(f)$ landing at a. Then it has period \mathbf{q} , and f cylically permutes the curves $f^k(\gamma)$, $k = 0, 1, \ldots, \mathbf{q} - 1$, with combinatorial rotation number \mathbf{p}/\mathbf{q} .

For a periodic point a of period p, the combinatorial rotation number is defined by considering it as a fixed point of f^p .

39.5.2. The α - and β - fixed points of a quadratic polynomial. Let us now assume that $f = f_c$ is a quadratic polynomial $z \mapsto z^d + c$ with connected Julia set. It turns out that the two fixed points of f (which use statically undistinguishable) play very different dynamical role.

The polynomail f has only one invariant ray, \mathcal{R}^0 . By Theorem 7.2, this ray lands at some fixed point called β ; moreover, this point is either repelling or parabolic with multiplier 1. (In the latter case, c = 1/4 is the cusp of the Mandelbrot set, and the two fixed points coincide.) The ray \mathcal{R}^0 is the only ray landing at β (for any other ray would be also invariant by Lemma 7.14). Thus β is the *non-dividing* fixed point (see Exercise 7.5).

Outside the cusp c = 1/4, f_c has the second fixed point called α . It is either attracting (for c in the main hyperbolic component $H_0 \subset M$ bounded by the main cardioid C – see §28) or neutral (for c on the main cardioid C), or repelling. If α is repelling or parabolic then by Theorem 7.3 it is a landing point of some periodic ray $\mathcal{R} = \mathcal{R}^{\theta}$. Since $\theta \neq 0 \mod 1$, the period q of this ray is greater than 1. Of course, all the rays $\mathcal{R}_i = f^i(\mathcal{R}), n = 0, 1, \ldots, q-1$, also land at α , so that, α is the *dividing* fixed point.

By Lemma 7.14 the ray portrait $\Theta(\alpha) \subset \mathbb{T}$ is a rotation set for the doubling map $\theta \mapsto 2\theta$. By Proposition 7.10, it is in fact, a single rotation cycle. Hence the rays \mathcal{R}_i are cyclically permuted by f with a combinatorial rotation number p/q.

196

This rotation number, $\rho(f_c) \equiv \rho(c)$, is also called the *combinatorial rotation number* of f (or of the corresponding parameter c).

The rays \mathcal{R}_i divide the plane into q sectors S_i , $i = 0, \ldots, q-1$, which cut off arcs ω_i at the circle at infinity. We studied these arcs in Lemma 7.9. Recall that the longest of these arcs, labeled $\omega_0 \equiv \omega_q$, is called critical, while the shortest, ω_1 is called characteristic. The corresponding sectors, $S_0 \equiv S_q$ and S_1 , will be called in the same way.

LEMMA 7.16. For i = 1, ..., q - 1, the map f univalently maps the sectors S_i onto S_{i+1} . The critical sector S_0 contains the critical point 0, while the characteristic sector S_1 contains the critical value c.

PROOF. Let \bar{S}_i be the compactification of the sector S_i at infinity obtained by adding the arc ω_1 to S_i . This is a topological triangle. For $i = 1, \ldots, q-1$, the boundary of S_i is homeomorphically mapped onto the boundary of S_{i+1} . By the Argument Principle, the whole triangle \bar{S}_i is homeomorphically mapped onto \bar{S}_{i+1} . Hence there are no critical points in these S_i , so that, $0 \in S_0$.

Let $\alpha' = -\alpha$; this is the second preimage of the fixed point α . By symmetry, there are q rays \mathcal{R}'_i landing at α' symmetric to the rays \mathcal{R}_i , so that, $f(\mathcal{R}'_i) = \mathcal{R}_{i+1}$, $i \in \mathbb{Z}/q\mathbb{Z}$. Altogether, the rays \mathcal{R}_i and \mathcal{R}'_i partition the plane into q-1 pairs of symmetric sectors $S_i, S'_i, i = 1, \ldots, q-1$ (bounded by two rays each) and a central domain $\Omega_0 \ni 0$ bounded by two pairs of symmetric rays.

LEMMA 7.17. The central domain Ω_0 is mapped onto the characteristic sector S_1 as a double branched covering.

PROOF. Each pair of symmetric rays that bound Ω_0 is mapped homeomorphically onto a characteristic ray that bound S_1 , so we have a 2-to-1 map $\partial \Omega_0 \rightarrow \partial S_1$.

Let Ω_0 be the compactification of Ω_0 by two symmetric arcs η and η' at infinity (where the arc η appeared in the proof of Lemma 7.9). Each of these arcs is mapped homeomorphically onto the characteristic arc ω_1 .

We see that the boundary of $\overline{\Omega}_0$ is mapped to the boundary of \overline{S}_1 as a double covering, and the conclusion follows.

Below we will describe the set of parameters with a given combinatorial rotation number.

39.6. Hubbard trees revisited. 39.6.1. The α -fixed point.

LEMMA 7.18. The legal path γ connecting the critical point $c_0 = 0$ to the critical value $c_1 = c$ contains a fixed point α of f.

PROOF. The image $\delta := f(\gamma)$ is the legal path connecting c to c_2 . Let us orient γ from 0 to c, and respectively, orient δ from c to c_2 . Since c is a vertex of T, these orientations are opposite near c.

Note that the inclusion $\delta \subset \gamma$ is impossible, since in this case a topological interval, γ , is mapped by f to itself, so it would contain a fixed point which is attracting at least on one side (as long as we know that the fixed points are isolated, which is certainly the case for polynomials).

If $\delta \supset \gamma$ then the inverse branch $f^{-1} : \delta \to \gamma$ maps a topological interval, δ , to a smaller interval γ , so it has a fixed point $\alpha \in \gamma$.

Otherwise, $\delta \cap \gamma = [c, \alpha]$, where α is a branch point of T. Let us show that this point is fixed under f. If this is not the case, then γ contains a point $\alpha_{-1} \neq \alpha$ such that $f(\alpha_{-1}) = \alpha$. Then the image $\alpha_1 := f(\alpha) \neq \alpha$ is a point of the path $[c, c_2]$. Let us consider the topological interval $I_0 = [\alpha, \alpha_1] \subset [c, c_2]$. It is oriented away from the critical point 0 (in the sense that the path $[0, \alpha]$ is disjoint from I_0).

Then the interval $I_1 \equiv [\alpha_1, \alpha_2] := f(I_0) \subset T$ extends I_0 beyond α_1 to an interval $[\alpha, \alpha_2]$ oriented away from 0 as well. Attachning to it the next iterate $I_3 \equiv [\alpha_2, \alpha_3] := f^2[\alpha, \alpha_1]$, we obtain a bigger interval $[\alpha, \alpha_3]$ with the same property, etc. Note that we will never hit the critical point (since the intervals $[\alpha, \alpha_n]$ grow away from it in the Hubbard tree), so the intervals in question will never be folded under f. In this way, we obtain infinitely many branch points $\alpha_n = f^n(\alpha)$ of T, which is impossible.

39.6.2. Spine and the β -fixed point. Since the zero-ray \mathcal{R}_0 is f-invariant, it lands at a fixed point. This point is usually called β . Let $\beta' \equiv -\beta$ be the corresponding *co-fixed* point. The legal arc $\sigma \equiv \sigma_f := [\beta, \beta']$ is called the *spine* of K(f). It turns out that the spine captures all cut-points in K.

LEMMA 7.19. A point $a \in J$ belongs to the interior of the spine if and only of there are two rays, \mathcal{R}_{θ_0} and θ_1 landing at a such that the diadic expansions for θ_0 and θ_1 begin with different digits (0 and 1 respectively).

PROOF. Of course, interior point of any path in K that belongs to ∂K is a cut-point.

Let us cosider the arc Γ composed of the spine and two rays, \mathcal{R}_0 and \mathcal{R}_{π} . It divides \mathbb{C} into two symmetric topological half-planes, Π_0 and P_1 , such that rays \mathcal{R}_{θ} with $\theta = [0\epsilon_2 \dots] \in (0, 1/2)$ lie in the upper half-plane P_0 , while rays \mathcal{R}_{θ} with $\theta = [1\epsilon_1 \dots] \in (1/2, 1)$ lie in the lower half-plane. It follows that \mathcal{R}_{θ_0} cannot land in Π_1 , while \mathcal{R}_{theta_1} cannot land in Π_0 . Hence they both land on the spine. \Box

PROPOSITION 7.20. A point $a \in J$ is a cut-point if it only if it belongs to some preimage of the interior of the spine, i.e., $f^n a \in \text{int } \sigma = (\beta, \beta')$ for some $n \in \mathbb{N}$.

PROOF. The "only if" part follows directly from Lemma 7.19.

Vice versa, if $a \in J$ is a cut-point then by definiton there are at least two rays, $\mathcal{R}_{\theta^{\pm}}$, landing at a. The diadic expansions of the corresponding angles, θ^{\pm} , differ at some place n. Then the diadic expansions for $2^{n}\theta^{+}$ and $2^{n}\theta^{-} \mod 1$ differ at the first place. By Lemma 7.19, $f^{n}a$ is a cut point, and hence a is, too.

The extended Hubbard tree T_f^e is the legal hull of the cycle **c** and the points β , β' .

EXERCISE 7.21. Show that $T_f^e = T_f \cup \sigma_f$ and that T_f^e is f-invariant.

One can describe exactly how the spine σ is located with respect to the Hubbard tree T.

39.7. Characteristic pair of rays and the Combinatorial Model.

198

40. Limbs and wakes of the Mandelbrot set

41. Geodesic laminations

42. Limbs and wakes of the Mandelbrot set

42.1. Stability of landing.

42.1.1. Repelling case. If a is a repelling periodic point of period p for a polynomial f then by the Implicit Function Theorem, a nearby polynomial \tilde{f} has a unique repelling periodic point \tilde{a} near a. We will refer to this point as the perturbed a.

LEMMA 7.22 (Stability Lemma). Assume that a periodic ray $\mathcal{R} = \mathcal{R}^{\theta}(f)$ lands at a repelling periodic point a for a polynomial f. Then for \tilde{f} sufficiently close to f, the corresponding ray $\tilde{\mathcal{R}}^{\theta}$ lands at the perturbed repelling periodic point \tilde{a} .

REMARK 7.1. Let us emphasize that this lemma applies to both connected and disconnected cases.

PROOF. Without loss of generality we can assume that the point a is fixed. and the ray \mathcal{R} is invariant.

Let us take a small disk $D = D(a, 2\epsilon)$ such that the local inverse branch g of f^{-1} is well defined in D and $g(D) \in D$. Then the same is true for \tilde{f} sufficiently close to f.

Let us fix some equipotential level t > 0 such that $\mathcal{R}(\tau) \subset D(a, \epsilon)$ for $\tau \leq dt$. Let γ be the closed arc of \mathcal{R} in between the potential levels t and dt.

Let us consider the inverse Böttcher function

$$\tilde{B}^{-1} \equiv B_{\tilde{f}}^{-1} : \mathbb{C} \smallsetminus \bar{\mathbb{D}}_{\rho(f)} \to \Omega_f$$

on its maximal domain fo definition. Let

$$\mathcal{R}_{>t} = B^{-1} \{ e^{\tau} e^{2\pi i \theta} : t \le \tau < \infty \}.$$

The notations $\mathcal{R}_{>t}$ and similarly notations for $\tilde{\mathcal{R}}$ (whenever they are well defined) are self-explanatory. The Böttcher formula (32.2) implies that \tilde{B}^{-1} depends continuously on \tilde{f} in the closed-open topology. Hence if \tilde{f} is sufficiently close to f, then the ray $\tilde{\mathcal{R}}_{\geq t}$ (parametrized by the potential level) is well defined and ϵ -close to the ray $\mathcal{R}_{\geq t}$. Let $\tilde{\gamma} = [\tilde{a}, \tilde{b}]$ be the arc of $\tilde{\mathcal{R}}$ between the potential levels t and dt. It follows that $\tilde{\gamma} \subset D$, so that, the inverse branch \tilde{g} is well defined on $\tilde{\gamma}$.

But $b = f(\tilde{a})$, so that, $\tilde{a} = g(b)$. Thus the arc $\tilde{g}(\tilde{\gamma}) \subset D$ gives an extension of the ray $\tilde{\mathcal{R}}_{\geq t}$ to the ray $\tilde{\mathcal{R}}_{\geq t/d}$. Repeating this argument, we conclude that the arcs $\tilde{g}^{-n}(\tilde{\gamma})$ give an extension of $\tilde{\mathcal{R}}_{>t}$ to the full ray $\tilde{\mathcal{R}}_{t>0}$.

42.1.2. Parabolic case.

LEMMA 7.23. Let f be a polynomial with connected Julia set. For a parabolic periodic point a with multimplier $\lambda = e^{2\pi i p/q}$, the combinatorial rotation number coincides with its rotation number p/q.

PROOF. Without loss of generality we can assume that a is a fixed point.

By Lemma ??, the rays landing at a are tangent to the bisectors L_i of the repelling petals, which are permuted by the differential Df(a) with rotation number p/q. Hence these rays are organized in q groups $\mathcal{G}_i = (\mathcal{R}_{ij})_j$, $i = 1, \ldots q$, so that the rays in \mathcal{G}_i are tangent to L_i . The rays within one group are naturally ordered:

 $\mathcal{R}_k \succ \mathcal{R}_j$ if \mathcal{R}_j is positively oriented relatively to the local sector S of zero angle bounded by these rays (in other words, \mathcal{R}_k is obtained from \mathcal{R}_j by the anti-clockwise "rotation" in S). Since f is a local orientation preserving diffeomorphism near a, it permutes these groups preserving the order of the rays. It follows that under f^q each ray is mapped back onto itself, and hence it is permuted by f with rotation number p/q.

Putting this together with Proposition 7.10, we obtain:

LEMMA 7.24. Assume that the α -fixed point of the quadratic polynomial $f_c : z \mapsto z^2 + c$ is parabolic with rotation number p/q. Then it is a landing point of q rays that are permuted with the same rotation number.

Recall that $r_{p/q} \in \mathcal{C}$ is the parabolic parameter with rotation number p/q. Let $H_{p/q}$ stand for the satellite hyperbolic component attached to the main cardioid \mathcal{C} at $r_{p/q}$ (see Proposition 28.5).

LEMMA 7.25. For any rotation number $p/q \neq 0$, there exists a curve c(t), $t \in [0, \epsilon)$ such that $c(0) = r_{p/q}$, $c(t) \in H_{p/q}$ for t > 0, and $\rho(c(t)) = p/q$.

42.2. Limbs and wakes.

42.2.1. Limbs. Let $\mathcal{L}_{p/q}^*$ be the connected component of $M \smallsetminus \{r_{p/q}\}$ containing $H_{p/q}$, and let $\mathcal{L}_{p/q} = \mathcal{L}_{p/q}^* \cup \{r_{p/q}\}$. This set is called the p/q-limb of the Mandelbrot set, while $\mathcal{L}_{p/q}^*$ is called the "unrooted p/q-limb".

PROPOSITION 7.26. For any $c \in \mathcal{L}_{p/q}$, the combinatorial rotation number $\rho(c)$ is equal to p/q.

PROOF. By Lemma 7.24, it is true at the root $r_{p/q}$. By Lemma 7.25, it is also true on some curve $\gamma \in H_{p/q}$ landing at $r_{p/q}$. By stability of ray portraits at repelling points (Lemma 7.22), the combinatorial rotation number $c \mapsto \rho(c)$ is a continuous function of $c \in \mathcal{L}_{p/q}^*$. Since the unrooted limb $\mathcal{L}_{p/q}$ is connected, while ρ can assume only rational values, it is constant on the whole limb $\mathcal{L}_{p/q}$. \Box

42.2.2. Characteristic parameter rays. Since the Stability Lemma 7.22 applies to the disconnected case as well, the p/q-ray portrait at the α -fixed point persists at some open set containing the unrooted limb $\mathcal{L}_{p/q}^*$. Below we will give the precese description of this open set.

For a parameter c with a well defined finite ray portrait, let $\mathcal{R}^-_{dyn}(c)$ and $\mathcal{R}^+_{dyn}(c)$ be the characteristic rays landing at the α -fixed point α_c of f_c , and let $S_{char}(c)$ be the characteristic sector bounded by these rays. For $c \in \mathcal{L}_{p/q}$, let $\theta^-_{p/q} < \theta^+_{p/q}$ be the angles of the characteristic rays (which are independent of c by by Proposition 7.26).

The corresponding objects in the parameter plane are the rays $\mathcal{R}_{par}^-(p/q)$ and $\mathcal{R}_{par}^+(p/q)$ with angles $\theta_{p/q}^-$ and $\theta_{p/q}^+$, and the p/q-wake $\mathcal{W}_{p/q}$, the component of $\mathbb{C} \setminus \operatorname{cl}(\mathcal{C} \cup \mathcal{R}_{par}^- \cup \mathcal{R}_{par}^+)$ containing the satellite hyperbolic component $H_{p/q}$.¹

In what follows we sometimes suppress the label p/q and c, as long as this cannot lead to a confusion.

 $^{^{1}}$ This definition is convenient to start with, but eventually it will be simplified (see Theorem 7.29).

LEMMA 7.27 (Key Observation). For $c \in \mathcal{R}_{par}^{\pm}(p/q)$, the dynamical characteristic rays $\mathcal{R}_{dyn}^{\pm}(c)$ do not land on $J(f_c)$.

PROOF. Assume for definiteness that $c \in \mathcal{R}_{par}^{-}(p/q)$. Then by the Basic Phase-Parameter relation, $c \in \mathcal{R}_{dyn}^{-}$.

Let $\Theta = \{\theta_i\}_{i=0}^{q-1} \subset \mathbb{T}$ be the cycle of θ^- under the doubling map, where $\theta_1 = \theta^-$. Then $0 \in \mathcal{R}_{dyn}^{\theta_0}$, so, the ray $\mathcal{R}_{dyn}^{\theta_0}$ does not land on J(f) but rather crashes at the critical point 0.

Going backwards along the cycle of rays $\mathcal{R}_{dyn}^{\theta_i}$, we see that all the rays of this cycle crash at some precritical point. In particular, the characteristic rays do. \Box

LEMMA 7.28. The wake $\mathcal{W}_{p/q}$ contains the unrooted limb $\mathcal{L}_{p/q}^*$ and some component Ω of $(\mathbb{C} \setminus M) \setminus (\mathcal{R}_{par}^-(p/q) \cup \mathcal{R}_{par}^+(p/q))$. All the points in the wake have combinatorial rotation number p/q.

PROOF. By the Stablility Lemma 7.22 and the Key Observation, the parameter rays $\mathcal{R}_{par}^{\pm}(p/q)$ cannot accumulate on a point $c \notin \mathcal{C}$ with rotation number p/q. In particular, they do not accumulate on the unrooted limb $\mathcal{L}_{p/q}^{*}$, which implies the first assertion.

It follows that the wake $\mathcal{W}_{p/q}$ intersects $\mathbb{C} \smallsetminus M$, and hence it contains the component Ω of $\mathbb{C} \smallsetminus M \smallsetminus (\mathcal{R}^-_{par}(p/q) \cup \mathcal{R}^+_{par}(p/q))$ such that $\mathcal{L}^*_{p/q} \subset \partial \Omega$. (Notice that $(\mathbb{C} \smallsetminus M) \smallsetminus (\mathcal{R}^-_{par}(p/q) \cup \mathcal{R}^+_{par}(p/q))$ consists of two components.)

Let us prove the last assertion. Assume there is a parameter $c_1 \in \mathcal{W}_{p/q}$ with $\rho(c_1) \neq p/q$. Let us fix a reference point $c_0 \in H_{p/q}$ and connect it to c_1 with a curve $c_t \subset \mathcal{W}_{p/q}$, $0 \leq t \leq 1$.

By the Stability Lemma 7.22, there is a maximal interval $[0, \tau)$ such that $\rho(c_t) = p/q$ for $t \in [0, \tau)$. By Proposition 7.26, $c(\tau) \notin \mathcal{L}_{p/q}^*$, so $c(\tau) \in \Omega$. Then by Proposition 7.6 only two events can happen:

(i) The characateristic ray $\mathcal{R}^+_{dyn}(c_{\tau})$ lands at some periodic point $a \neq \alpha$ of $J(f_{c_{\tau}})$. But then by the Stability Lemma, this would also be the case for $c_{\tau-\epsilon}$ for $\epsilon > 0$ sufficiently small, contradicting definition of τ .

(ii) The characteristic ray $\mathcal{R}_{dyn}^-(c_{\tau})$ crashes at some precritical point. But then the critical value c_{τ} would belong to one of the ray $\mathcal{R}_{dyn}^{\theta_i}(c_{\tau})$ of the cycle of $\mathcal{R}_{dyn}^-(c_{\tau})$. Since for $c_{\tau-\epsilon}$, the critical value c belongs to the characteristic sector $S_{char}(c)$, this can only be one of the characteristic rays $\mathcal{R}_{dyn}^{\pm}(c_{\tau})$. But then by the Basic Phase-Parameter relation, $c_{\tau} \in \mathcal{R}_{par}^{\pm}$ contradicting the definition of Ω .

THEOREM 7.29. Both parameter rays $\mathcal{R}_{par}^{\pm}(p/q)$ land at the root $r_{p/q}$. The wake $\mathcal{W}_{p/q}$ coincides with the domain bounded by the curve $\mathcal{R}_{par}^{-}(p/q) \cup \mathcal{R}_{par}^{+}(p/q) \cup r_{p/q}$ and containing $H_{p/q}$. The combinatorial rotation number is equal to p/q throughout the wake.

PROOF. We know from the proof of Lemma 7.28 that the rays $\mathcal{R}_{par}^{\pm}(p/q)$ cannot accumulate on a point $c \in M \setminus \mathcal{C}$ with rotation number p/q. Let us show that they can neither accumulate on other points $c \in M \setminus \mathcal{C}$.

Let $\rho(c) = r/s \neq p/q$. By the Stability Lemma, $\rho(\tilde{c}) = r/s$ for all $\tilde{c} \in D(c, \epsilon)$, provided $\epsilon > 0$ is sufficiently small. But if \mathcal{R}_{par}^- accumulates on c then all nearby parameter rays $\mathcal{R}_{par}^{\theta}$ enter the disk $D(c, \epsilon)$. Take such a parameter ray in the domain Ω , and let $\tilde{c} \in D(c, \epsilon) \cap \mathcal{R}^{\theta}_{par}$. Since $\tilde{c} \in W$, $\rho(\tilde{c}) = p/q$ by Lemma 7.28, and we have arrived at a contradiction.

Hence the rays \mathcal{R}_{par}^{\pm} can accumulate only the points of main cardioid \mathcal{C} . Let $\omega^{\pm} \subset \mathcal{C}$ be the limit sets of the rays \mathcal{R}^{\pm} . If one of then, say, ω^{-} , was not a single point, then we could find a rational point $p/q \in \operatorname{int} \omega^{-}$, and the ray \mathcal{R}_{par}^{-} would have to cross the satellite component $H_{p/q}$. Since it is certainly impossible, the limit sets ω^{\pm} are, in fact, single points, so that both rays land at some points c_{\pm} of the main cardioid.

If $c_+ \neq c_-$ then the wake $\mathcal{W}_{p/q}$ would contain, besides $H_{p/q}$, some other satellite hyperbolic domain $H_{r/s}$. But the combinatorial rotation number in $H_{r/s}$ is equal to $r/s \neq p/q$ contradicting Lemma 7.28. This shows that the rays \mathcal{R}_{par}^{\pm} land at the root $r_{p/q}$, and the rest of the lemma easily follows.

The angles $\theta_{p/q}^{\pm}$ of the rays $\mathcal{R}_{par}^{\pm}(p/q)$ landing at $r_{p/q}$ are also called the *external* angels of $r_{p/q}$.

42.3. Limbs and wakes attached to other hyperbolic components. One can generalize the above discussion to limbs attached to any hyperbolic component H in place of the main one, H_0 . Let $a_c = a_c^H$, $c \in H$, be the attracting periodic point of f_c , and let $\lambda_c = \lambda_c^H$ be its multiplier. On the boundary of H the point a_c becomes neutral. By the Multiplier Theorem 6.12, for $c \in \partial H$, the rotation number $\rho(a_c)$ assumes once every value $\theta \in \mathbb{R}/\mathbb{Z}$. Let $r_{p/q}(H) \in \partial H$ stand for the parabolic parameter with rotatin number p/q, i.e., $\rho(a_c) = p/q$.

THEOREM 7.30. Let $p/q \neq 0 \mod 1$. Then there are two parameter rays $\mathcal{R}_{par}^{\pm}(p/q, H)$ landing at $r_{p/q}(H)$ such $\rho(a_c^H) = p/q$ in the wake $\mathcal{W}_{p/q}(H)$ bounded these rays, and moreover, this wake is a maximal region with this property.

REMARK 7.2. Note that at this moment we do not claim that there are no other rays landing at $r_{p/q}(H)$ since we will use Theorem 7.30 to show this.

The limb $\mathcal{L}_{p/q}(H)$ of M attached to the parabolic point $r_{p/q}(H)$ is defined as in the case of the main component H_0 .

In the case when H is itself a satellite hyperboplic component attached to H_0 , we call $\mathcal{L}_{p/q}(H)$ and $\mathcal{W}_{p/q}(H)$ secondary limbs and wakes respectively.

In what follows, we will also need to know that the external angles $\theta_{p/q}^{\pm}(H)$ of a root point depend continuously on the internal angle p/q.

LEMMA 7.31. • Let $p/q \neq 0$. Then $\theta^{\pm}(r/s)(H) \rightarrow \theta^{-}_{p/q}(H)$ as $r/s \nearrow p/q$, and $\theta^{\pm}(r/s)(H) \rightarrow \theta^{+}_{p/q}(H)$ as $r/s \searrow p/q$.

- Let $H = H_0$. Then $\theta_{p/q}^{\pm}(H) \to 0$ as $p/q \to 0 \mod 1$.
- Let H be a satellite hyperbolic component and $\theta_0^{\pm}(H)$ be the characteristic rays landing at the root of H. Then $\theta_{p/q}^{\pm}(H) \to \theta_0^-(H)$ a $p/q \searrow 0$ and $\theta_{p/a}^{\pm}(H) \to \theta_0^+(H)$ as $p/q \nearrow 1$.

REMARK 7.3. The only reason why in the last statement we assume that H is satellite is that we do not know yet that there are rays landing at the root of a primitive hyperbolic component $H \neq H_0$.

202

42.4. No fake limbs. A *fake limb* of M is a component of $M \setminus H_0$ different from any limb $\mathcal{L}^*_{p/q}$.

LEMMA 7.32. There are no fake limbs.

PROOF. Let X be such a limb. Notice first that $\overline{X} \cap \mathcal{C} \neq \emptyset$, for otherwise M would be disconnected. Also, since X is connected, the combinatorial rotation number $\rho(c)$ is independent of $c \in X$, so we have a well defined number $\rho(X) = p/q$.

Obviously $X \cap \partial \mathcal{W}_{p/q} = \emptyset$, so that, X is either contained in $\mathcal{W}_{p/q}$ or lies outside its closures. Let us first assume the latter. Then for $r/s \notin \{p/q, 0\}$, any parameter $c_0 \in X$ can be connected to any parameter $c_1 \in H_{r/s}$ by a path $c_t \in \mathbb{C} \setminus \overline{H}_0$, $t \in [0, 1]$, that does not cross $\partial \mathcal{W}_{p/q}$. But then by the Stability Lemma, $\rho(c_t) = p/q$ for all $t \in [0, 1]$, contradicting to $\rho(c_1) = r/s$.

Assume now that $X \subset \mathcal{W}_{p/q}$. Let us then consider the periodic cycle $\{f^n(a_c)\}_{n=0}^{q-1}$ of period q obtained by analytic continuation of the attracting cycle bifurcated from the α -fixed point at $r_{p/q}$. At this moment we do not know yet that the multiplier λ_c of this cycle is different from 1 throughout the wake $\mathcal{W}_{p/q}$, so that, the function $c \mapsto a_c$ can be multi-valued. Let $Z = \{c \in \mathcal{W}_{p/q} \cup \{r_{p/q}\} : |\lambda_c| \leq 1\}$. Since this is a finite union of disjoint Jordan disks, $\mathcal{W}_{p/q} \smallsetminus Z$ is connected. Notice also that Xis not contained in Z since there are always satellite components attached to each compenent of Z. Let $c_0 \in X \smallsetminus Z$, and let k/l be the combinatorial rotation number of the periodic point a_{s_0} .

Let us consider the secondary wakes $\mathcal{W}_{r/s}^2$ attached to the satellite component $H_{p/q}$. Again, we have the anlternative: either $X \subset \mathcal{W}_{k/l}^2$ or $X \cap \overline{\mathcal{W}}_{k/l}^2 = \emptyset$. But the former option is actually impossible since $\overline{\mathcal{W}}_{k/l}^2$ does not touch \mathcal{C} . The latter option is ruled out in the same way as above by taking a different rotation number $r/s \neq k/l$ and connecting c_0 to a seconday satellite component $H_{r/s}^2$ attached to $H_{p/q}$ with a path $c_t \in \mathcal{W}_{p/q} \smallsetminus Z$.

COROLLARY 7.33. The Mandelbrot set admits the following partition:

$$M = H_0 \cup \mathcal{C} \cup \bigcup_{p/q \in \mathbb{Q}/\mathbb{Z} \setminus 0} \mathcal{L}_{p/q}^*.$$

COROLLARY 7.34. The rays $\mathcal{R}_{par}^{\pm}(p/q)$ are the only parameter rays landing at $r_{p/q}$.

PROOF. If there was an extra ray \mathcal{R} landing at $r_{p/q}$ then by Lemma 1.119 there would be an extra component of $M \setminus \overline{H}_0$ attached to \mathcal{C} at $r_{p/q}$.

42.5. The α -rays and their holomorphic motion. Let us fix some combinatorial rotation number p/q. For $c \in W_{p/q}$, let $\mathcal{R}_c^{\theta_i}$ be the dynamical rays landing at the α -fixed point α_c , and let

$$\mathcal{I}_{c}^{(0)} = \bigcup_{i} \mathcal{R}_{\mathrm{dyn}}^{\theta_{i}}(c) \cup \{\alpha_{c}\}.$$

This configuration of rays partition the plane into q sectors S_i described in Lemma 7.16.

Let $h_c : X_* \to X_c$ be a holomorphic motion of some dynamical set over a pointed parameter domain $(\Lambda, *)$ of the quadratic family $z \mapsto z^2 + c$. We say that it respects the Böttcher marking if for any point $z \in X_* \setminus J(f_*)$ we have:

$$B_c(h_c(z)) = B_*(z), \quad c \in \Lambda$$

FIGURE 1. Almost renormalization for the elephant eye.

(so that, the Böttcher coordinate B_c is the "first integral" of the motion).

PROPOSITION 7.35. There is a holomorphic motion of the configuration $\mathcal{I}_c^{(0)}$ over the parabolic wake $\mathcal{W}_{p/q}$ that respects the Böttcher marking.

PROOF. Let us select an arbitrary base point $* \in \mathcal{W}_{p/q}$.

By definition, $B_c(\mathcal{R}^{\theta}_c(t)) = e^{t+i\theta}$, where $t \in \mathbb{R}_+$ and $\theta \in \mathbb{R}/\mathbb{Z}$ are the Böttcher coordinates of the point $\mathcal{R}^{\theta}_c(t)$. Hence for $B_*(\mathcal{R}^{\theta_i}_*(t)) = B_c(\mathcal{R}^{\theta_i}_c(t))$, so that, $h_c(z) = B_c^{-1} \circ B_*(z)$ determines a motion of the external rays $\mathcal{R}^{\theta_i}_c$ over $\mathcal{W}_{p/q}$ respecting the Böttcher marking. This motion is holomorphic since the Böttcher function B_c depends holomorphically on c.

On the other hand, the point α_c obviously moves holomorphically over $\mathcal{W}_{p/q}$ as well, and we obtain the desired motion of the whole configuration $\mathcal{I}_c^{(0)}$.

Let

(42.1)
$$\mathcal{I}_{c}^{(n)} = f^{-n}(\mathcal{I}_{c}^{(0)}).$$

42.6. MLC on the main cardioid.

THEOREM 7.36. The Mandelbrot set is locally connected at any point of the main cardioid C.

PROPOSITION 7.37. For any irrational $\theta \in \mathbb{R}/\mathbb{Z}$, there is a single parameter ray \mathcal{R}^{η} landing at the point $c(\theta) \in \mathcal{C}$ with internal angle θ .

43. Misiurewisz wakes and decorations

44. Topological model

45. Renormalization

45.1. Hyperbolic maps.

45.1.1. Canonical almost renormlization. Let us consider a hyperbolic quadratic polynomial $f = f_c$ with an attracting cycle $\boldsymbol{\alpha} = \{\alpha_k\}_{k=0}^{p-1}$ of period p > 1. Here we will show that it is renormalizable with period p.

Let us consider the immediate basin $\sqcup D_i$ of $\boldsymbol{\alpha}$, where $D_i \ni \alpha_i$ and $D_0 \ni 0$. As we know, the return maps $f^p | \partial D_i$ are topologically conjugate to $z \mapsto z^2$ on the unit circle \mathbb{T} ; hence each boundary ∂D_i contains a unique fixed point γ_i of f^p and a unique "co-fixed" point $\gamma'_i \neq \gamma_i$ (such that $f^p(\gamma'_i) = \gamma_i$). Note that the points γ_i form a cycle of period l which is a diviser of p, so $\gamma_i = \gamma_j$ if $i \equiv j \mod l$. Note also that $\gamma'_0 = -\gamma_0$.

Let \mathcal{R}_1^{\pm} be the pair of characteristic rays (see §??), i.e., the pair of rays landing at γ_1 that bound the sector containing the critical value c. Let θ_{\pm} be their angles. Let us consider an arc

$$\Gamma_1 = \mathcal{R}_1^+ \cup \mathcal{R}_1^- \cup \{\gamma_1\}.$$

Let Γ_i and Γ'_i be its lifts by f through the points γ_0 and γ'_0 respectively. Its further pullbacks by the iterates f^{p-i} (i = 1, ..., p - 1) passing through the point γ_i and γ'_i will be called Γ_i and Γ'_i respectively. (Note that there are p/l arcs Γ_j passing through each γ_i .) See Figure ??.

FIGURE 2. Thickening.

Let $\pm \gamma_0'' \in \partial D_0$ be two preimages of γ_0' under f^p , and let $\pm \Gamma_0''$ be the corresponding pullbacks of the arc Γ_1' .

Select your favorite t > 0, and truncate this configuration of rays by equipotential \mathcal{E}_{2t} of level 2t > 0. Then the subpotential domains $\Omega(2t)$ get tiled by several topological disks D'_j . Let $V'_1 \ni c$ be the critical value tile in $\Omega(2t)$, let $V'_0 = f^{-1}(V'_1)$, and let $V_0 \ni 0$ be the pullback of V'_0 under f^p .

EXERCISE 7.38. (i) $V_0 \subset V'_0$;

(ii) The disk V_0 is bounded by arcs of Γ_0 , Γ'_0 , $\pm \Gamma''_0$, and four equipotential arcs of \mathcal{E}_t .

(iii) $f: V_0 \to V_1$ is a double branched covering.

Let $V'_0 := f^p(V_0)$. It follows that $V_0 \subset V'_0$ and $f^p : V_0 \to V'_0$ is a double branched covering. It is called the *canonical almost renormalization of* f (of period p).

EXERCISE 7.39. Draw the canonical almost renormalization picture for the rabbit, airplane, and other favourite hyperbolic maps of yours.

45.1.2. Thickening. The double covering $f^p: V_0 \to V'_1$ described above is not a quadratic-like map since the domains V_0 and V'_1 have a common boundary (arcs of Γ_0 and Γ'_0). To fix this problem, let us slight; *thicken* these domains, see Figure ??. Namely, one can replace Γ_0 with a nearby arc $\tilde{\Gamma}_0$ comprising pieces of two nearby rays and a little circle around γ_0 . Since the periodic angles θ_{\pm} are repelling under the *p*-fold iterate of the doubling map $\mathbb{T} \to \mathbb{T}$, and γ_0 is a repelling fixed point for f^p , the arc $\tilde{\Gamma}_0$ will be "pushed farther away" from Γ_0 under f^p . Similarly, replace Γ'_0 with the arc $\tilde{\Gamma}'_0$ which is symmetric to $\tilde{\Gamma}_0$. Truncating this pair of arcs by the equipotential $\mathcal{E}(2^pt)$ we obtain a domain $\tilde{V}'_0 \supseteq V'_0$. Pulling it back under f^p , we obtain a domain $\tilde{V}_0 \supseteq V_0$.

LEMMA 7.40. The map $f^p: \tilde{V}_0 \to \tilde{V}'_0$ is a quadratic-like renormalization of f.

45.1.3. Real renormalization. Let us now consider a real hyperbolic quadratic map $f = f_c$, $c \in [-2, 1/4]$, with an attracting cycle $\boldsymbol{\alpha} = \{\alpha_i\}_{i=0}^{p-1}$ of period p > 1. This cycle is real, and it has the real immediate attracting basin consisting of the intervals $I_j = (\gamma_j, \gamma'_j) = D_j \cap \mathbb{R}, j = 0, \ldots, p-1$. Moreover, $I_0 \ni 0$. The return map $f^p: \overline{I_0} \to \overline{I_0}$ is the real renormalization of f. Since on the real line we do not need to create a quadratic-like picture, no thickening is needed in this construction.

45.2. Renormalization windows.

45.2.1. Complex windows. Let us go back to a complex hyperbolic map $f_{\circ} = f_{c_{\circ}}$ of period p > 1. The ray portrait used to construct its canonical almost renormalization survives on a parameter region much bigger that the hyperbolic component h_{\circ} of f_{\circ} . Namely, let us consider the open parameter domain $W_{\circ} \equiv W_{c_{\circ}}$ corresponding to the dynamical domain D_1 around the critical value, see Figure ??. This domain is bounded by two parameter rays of angle θ_{\pm} and two parameter rays of angle ψ_{\pm} truncated be the equipotential lof level 2t. The former rays land at the root r_{\circ} of H_{\circ} , while the latter land at its tip t_{\circ} .

PROPOSITION 7.41. The configuration of arcs Γ_i , Γ'_i , $\pm \Gamma''_i$, and the equipotentials that creates the canonical almost renormalization moves holomorphically over the parameter window W_{\circ} .

PROPOSITION 7.42. All the maps f_c , $c \in W_o$, are renormalizable with period p. Furthermore, the map f_{tip} is also renormalizable, while the map f_{root} is renormalizable if and only if the hyperbolic component H_o is primitive.

45.2.2. Real windows. Let us take the real slice $W_{\circ}^{\mathbb{R}} := W_{\circ} \cap \mathbb{R}$ of the complex renormalization window. It is an open interval bounded be the root and tip of W_{\circ} . Proposition 7.41 implies:

COROLLARY 7.43. Let f_{\circ} be a real hyperbolic map of period p > 1. Then the configuration of arcs Γ_i , Γ'_i , $\pm \Gamma''_i$, and the equipotentials that creates the canonical almost renormalization moves continuously over the real renormalization window $W^{\mathbb{R}}_{\circ}$.

Similarly, we have the real counterpart of Proposition 7.44:

COROLLARY 7.44. All the maps f_c , $c \in \overline{W}^{\mathbb{R}}_{\circ}$, are really :) renormalizable with period p. (Note that both the tip and the root are included to $\overline{W}^{\mathbb{R}}_{\circ}$.)

CHAPTER 8

Thurston theory

45. Rigidity Theorem

45.1. Thurston equivalence. Let us consider two quadratic polynomials, $f = f_c$ and $\tilde{f} = f_{\tilde{c}}$, with postcitical sets \mathcal{O} and $\tilde{\mathcal{O}}$. They are called *Thurston equivalent* if there exists a homeomorphism $h : (\mathbb{C}, \mathcal{O}) \to (\mathbb{C}, \tilde{\mathcal{O}})$ which is a conjugacy on the postrcitical sets and that can be lifted (via f and \tilde{f}) to a homeomorphism $h_1 : (\mathbb{C}, \mathcal{O}) \to (\mathbb{C}, \tilde{\mathcal{O}})$ homotopic to h rel \mathcal{O} .

A couple of remarks are due. First, recall that h_1 is a lift of h via f and \tilde{f} if $\tilde{f} \circ h_1 = h \circ f$. A homeomorphism $h : \mathbb{C} \to \mathbb{C}$ is liftable if and only if $h(c) = \tilde{c}$. If so, there are two lifts determined by whether $h_1(c) = \tilde{c}$ or $h_1(c) = -\tilde{c}$.

Next, since $h: \mathcal{O} \to \tilde{\mathcal{O}}$ is a conjugacy, we have $h(c) = \tilde{c}$, hence h is liftable in two ways. The above definition requires that if the lift is selected so that $h_1(c) = \tilde{c}$, then $h_1 | \mathcal{O} = h | \mathcal{O}$, and h_1 is homotopic to h rel \mathcal{O} .

Finally, we can specify an extra regularity (qc, smooth, etc.) of a Thurston equivalence h.

45.2. Pullback Argument.

LEMMA 8.1. If two quadratic polynomials $f = f_c$ and \tilde{f} are K-qc Thurston equivalent, then the Böttcher conjugacy $D_c(\infty) \to D_{\tilde{c}}(\infty)$ admits a K-qc extension to the whole complex plane which is a conjugacy on the postcritical sets (and is automatically a conjugacy on the Julia sets).

PROOF. Let $D(t) = \mathbb{C} \setminus \Omega(t)$, where $\Omega(t)$ is the subpotential domain for f of level t > 0 (see §32.2). As usual, the corresponding objects for \tilde{f} are marked with tilde. Let $\phi : D(\infty) \to \tilde{D}(\infty)$ be the Böttcher conjugacy between f and \tilde{f} on the basins of ∞ .

Take a quasidisk Δ containing \mathcal{O} , let $\tilde{\Delta} = h(\Delta)$, and select an equipotential level t > 0 so that $\Delta \in \Omega(t)$ and $\tilde{\Delta} \in \tilde{\Omega}(t)$. We can modify h and h_1 on $\mathbb{C} \setminus \Delta$ so that on D(t) they are both equal to the Böttcher conjugacy $\phi : D(t) \to \tilde{D}(t)$ (by means the quasiconformal interpolation into the annulus $\Omega(t) \setminus \Delta$, see ??). This modification does not change the homotopy type of h rel \mathcal{O} since it does not change h on Δ (see the Alexander trick ??). Moreover, the modification can be done so that h_1 is homotopic to h rel D(t). Indeed, since $h|D(t) = \phi$, the lift $h_1|D(t/2)$ is equal to either ϕ or $-\phi$. In the latter case we replace the homotopy type of h on the annulus $\Omega(t) \setminus \Delta$ rel its boundary to get the correct lift (see Exersice ??).

By the Lift Homotopy Theorem, h_1 admits a lift h_2 homotopic to h_1 (which is a lift of h) rel $f^{-1}(\mathcal{O}) \supset \mathcal{O}$. In particular, $h_2 | \mathcal{O} = h_1 | \mathcal{O}$.

On the domain D(t/2), the lift h_2 is either equal to the is a Böttcher conjugacy . Moreover:

• h_2 is holomorphic on D(t/2), and hence it is the Böttcher conjugacy on this larger domain.

• h_2 is K-qc (since f and \tilde{f} are holomorphic).

Repeating this lifting procedure we obtain a sequence of K-qc homeomorphisms $h_n : (\mathbb{C}, \mathcal{O}) \to (\mathbb{C}, \tilde{\mathcal{O}})$ in the same homotopy class rel \mathcal{O} and such that h_n is the Böttcher conjugacy on $D(t/2^n)$.

By compactness of the space of normalized K-qc maps, there exists a subsequence $h_{n(k)}$ uniformly converging to a K-qc map H. Moreover, H coincides with the Böttcher conjugacy on the whole basin ∞ , and the conclusion follows.

COROLLARY 8.2. Let f and \tilde{f} be two quadratic polynomials with nowhere dense Julia sets. If they are K-qc Thurston equivalent then they are K-qc conjugate.

46. Rigidity of superattracting polynomials

THEOREM 8.3. Let $f = f_c$ and $\tilde{f} = f_{\tilde{c}}$ be two superattracting quadratic polynomials. If they are Thurston equivalent then $c = \tilde{c}$.

PROOF. Since the postcritical set is finite, the Thurston equivalence can be assumed smooth, and hence qc, on the whole Riemann sphere. By the Pullback Argument, the Böttcher conjugacy between f and \tilde{f} extends continuously to the Julia sets. Since the attracting cycles of our maps have the same multipliers (equal to 0), the conclusion follows from Lemma 6.17.

REMARK 8.1. Instead of using Lemma 6.17, one could adjust the Pullback Argument so that it would directly imply the statement. Namely, one can modify the Thurston equivalence so that is becomes a conformal conjugacy near the superattracting cycles (similarly to the adjustment near ∞ carried in the proof of Lemma 8.1). Then the Pullback Argument will turn it into a qc conjugacy which is conformal outside the Julia set. Since the latter has zero area, it is conformal on the whole plane.

47. Hubbard tree determines f

We say that two maps f and \tilde{f} "have the same Hubbard trees", \mathcal{T} and $\tilde{\mathcal{T}}$ if there is a conjugacy $h_{\mathcal{T}} : (\mathcal{T}, 0) \to (\tilde{\mathcal{T}}, 0)$ between the restriction of the maps to their trees. Such a conjugacy should necessarily respect the marking of the trees.

REMARK 8.2. Note that the condition $h_{\mathcal{T}}(0) = 0$ is satisfied automatically if 0 is not an extremity.

LEMMA 8.4. If two superattracting quadratic polynomials have the same Hubbard trees, then they are Thurston equivalent.

PROOF. Let us partition the plane by the rays \mathcal{R}_i landing at the marked points of the Hubbard trees, and let $\mathcal{R} = \bigcup \mathcal{R}_i$. Then the conjugacy $h_{\mathcal{T}} : \mathcal{T} \to \tilde{\mathcal{T}}$ can be extended to the whole plance so that it is the Böttcher conjugacy \mathcal{R} . It further lifts to a homeomorphism

$$h_1: (\mathbb{C}, f^{-1}(\mathcal{T} \cup \mathcal{R}) \to (\mathbb{C}, \tilde{f}^{-1}(\tilde{\mathcal{T}} \cup \tilde{\mathcal{R}}))$$

which is Böttcher on $f^{-1}(\mathcal{R})$. As $h_1 = h$ on \mathcal{R} , these two maps are homotopic rel \mathcal{R} , all the more rel \mathcal{O} . It gives us a desired Thurston equivalence.

208

Putting this together with Theorem 8.3, we obtain:

COROLLARY 8.5. If two superattracting quadratic polynomials, f_c and $f_{\tilde{c}}$, have the same Hubbard tree then $c = \tilde{c}$.

47.1. Growth of entropy. Let us now consider a real superattracting quadratic polynomial $f = f_c$, $c \in [-2, 1/4]$ of period p > 0. Its Hubbard tree \mathcal{T}_f is the interval [c, f(c)] with the marked postrcritical set \mathcal{O} . As this information is equivalent to presribing the (finite) kneading sequence of f, we obtain:

COROLLARY 8.6. An finite kneading sequence uniquely determines a real superattracting parameter.

With the Intermediate Value Theorem in hands, we can promote this resul to infinite sequences as well.

THEOREM 8.7. An infinite kneading sequence uniquely determines a real parameter.

For real maps, *topological entropy* can be defined as the growth rate for the number of periodic points:

$$h(f) = \lim \frac{1}{n} \log |\operatorname{Per}_n|,$$

where $\operatorname{Per}_n = \{x : f^n x = x\}.$

COROLLARY 8.8. As c moves from 1/4 to -2, the topological entropy $h(f_c)$ monotonically changes from 0 to $\log 2$.

REMARK 8.3. The entropy function $c \mapsto h(f_c)$ is an example of the Devil Staircase: it is constant on the hyperbolic windows, and grows on the compementary Cantor set (of positive length).

48. Realization of critically periodic maps

48.1. Statement. Let $g : \mathbb{C} \to \mathbb{C}$ be a topological double branched covering of the plane with the critical point at 0. We say that it is *postcritically periodic* if $f^p(0) = 0$ for some $p \ge 1$. (As we know, quadratic polynomials with this property are called "supertracting", but this term could be misleading in the topological setting since a periodic critical point can be even repelling.) Let $\mathcal{O} = \{f^n(0)\}_{n=0}^{p-1}$.

Similarly to the actual quadratic maps, we can define *Thurston equivalence* between such maps, see §??. We say that a map g in question is *realizable* if there exists a quadratic polynomial f_c in the Thurston class of g. By the Uniqueness Theorem, this realization is unique.

THURSTON REALIZATION THEOREM. Any postcritically periodic topological double branched covering $g : \mathbb{C} \to \mathbb{C}$ is realizable.

48.2. Proof.

48.2.1. The Teichmüller and Moduli spaces. Without loss of generality, we can assume that g is quasiregular. All the conformal structures on \mathbb{C} (or, on the punctured \mathbb{C}) will be assumed to have a bounded dilatation.

Let us consider the Teichmüler space $\mathcal{T} = \mathcal{T}_g \approx \mathcal{T}_{p+1}$ of the punctured plane $\mathbb{C} \setminus \mathcal{O} = \hat{\mathbb{C}} \setminus \mathcal{Z}$, where $\mathcal{Z} = \mathcal{O} \cup \{\infty\}$. By definition, it is the space of conformal structures μ on $\mathbb{C} \setminus \mathcal{O}$ up to homotopy. More precisely, let $h_{\mu} : (\mathbb{C}, \mathcal{O}) \to (\mathbb{C}, \mathcal{O}_{\mu})$ be

the the solution of the Beltrami equation for the structure μ , where $\mathcal{O}_{\mu} = h_{\mu}(\mathcal{O})$. Two structures μ and μ' are equivalent if there is a complex affine transformation $\phi : (\mathbb{C}, \mathcal{O}_{\mu}) \to (\mathbb{C}, \mathcal{O}_{\mu'})$ such that $\phi \circ h_{\mu}$ is homotopic $h_{\mu'}$ rel \mathcal{O} . A class $\tau = [h_{\mu}]$ of equivalent maps represents a point of \mathcal{T} .

The moduli space $\mathcal{M} = \mathcal{M}_g$ is the space of embeddings $\mathcal{O} \to \mathbb{C}$ up to affine transformation. The natural projection $\pi : \mathcal{T} \to \mathcal{M}$ associates to a class $[h_\mu] \in \mathcal{T}_g$ the class of embeddings $[h_\mu | \mathcal{O}] \in \mathcal{M}$.

48.2.2. Pullback operator and its fixed points. By an affine conjugacy and homotopy rel \mathcal{O} , we can normalize g so that $c_0 = 0$ is its critical point and $g(z) = z^2$ near ∞ . Let $c_k = f^k(0), k = 0, \ldots, p-1$.

Let us now define the pullback operator $g^*: \mathcal{T} \to \mathcal{T}$ induced by the pullback $\mu \mapsto \mu' = g^*(\mu)$ of the compex structures on \mathbb{C} . More precisely, let a point $\tau \in \mathcal{T}$ be represented by a homeomorphism $h: (\mathbb{C}, \mathcal{O}, 0) \to (\mathbb{C}, \mathcal{O}_{\tau}, 0)$. Then $\mu = h^*(\sigma)$ and $\mu' = (h \circ g)^*(\sigma)$. Let $h': (\mathbb{C}, \mathcal{O}, 0) \to (\mathbb{C}, h'(\mathcal{O}, 0))$ be the solution of the Beltrami equation with the conformal structure μ' . Then h' represents the point $\tau' = g^*(\tau)$. If \tilde{h} is homotopic to h rel \mathcal{O} then by the Lift Homotopy Theorem ensures that \tilde{h}' is homotopic to \tilde{h} rel \mathcal{O} , so the operator is well defined. In particular, $h'(\mathcal{O})$ (up to rescaling) depends only on τ' , and it can be called $\mathcal{O}_{\tau'}$.

Let $h(\mathcal{O}) = (z_0 = 0, z_1, \dots, z_{p-1}), h'(\mathcal{O}) = (z'_0 = 0, z'_1, \dots, z'_{p-1}),$ where $z_i = h(c_i), z'_i = h'(c_i).$

Composition $h' \circ g \circ h^{-1}$ is a holomorphic double branched covering

$$f: (\mathbb{C}, \mathcal{O}_{\tau'}) \to (\mathbb{C}, \mathcal{O}_{\tau}),$$

so it is a quadratic polynomial. Notice, however, that f does not have a dynamical meaning as $h'(\mathcal{O}) \neq h(\mathcal{O})$. Moreover, we have two independent scaling factors to normalize the maps h and h', and there are several useful ways to do so. For instance,

N1: Let $z_1 = 1$ while $z'_{p-1} = i$; then $f(z) = z^2 + 1$. N2: Let $z_1 = z'_1 = 1$; then $f(z) = (z_2 - 1) z^2 + 1$.

(48.1)
$$\begin{array}{ccc} (\mathbb{C}, \mathcal{O}, \tau') & \longrightarrow & (\mathbb{C}, \mathcal{O}_{\tau'}, \sigma) \\ g \downarrow & & \downarrow f \\ (\mathbb{C}, \mathcal{O}, \tau) & \longrightarrow & (\mathbb{C}, \mathcal{O}_{\tau}, \sigma) \end{array}$$

48.2.3. Ambiguity in the Moduli space. The operator g^* does not descend to the moduli space \mathcal{M}_{p+1} : the Riemann surface $(\mathbb{C}, \mathcal{O}_{\tau'})$ is diagram (48.1) is not uniquely determined by $(\mathbb{C}, \mathcal{O}_{\tau})$. However, the ambiguity is bounded:

LEMMA 8.9. For a given Riemann surface $(\mathbb{C}, \mathcal{O}_{\tau'})$, ther exists a bounded (in terms of $p = |\mathcal{O}|$) number of Riemann surfaces $(\hat{\mathbb{C}}, \mathcal{O}_{\tau'})$. Moreover, if $(\mathbb{C}, \mathcal{O}_{\tau'})$ belongs to a compact subset \mathcal{K} of \mathcal{M}_p then $(\mathbb{C}, \mathcal{O}_{\tau'})$ belongs to a compact subset $\mathcal{K}' \Subset \mathcal{M}_p$ as well.

PROOF. Let us use normalization N1 for h and h', so $f(z) = z^2 + 1$. Then the set $f^{-1}(\mathcal{O}_{\tau})$ is uniquely defined by \mathcal{O}_{τ} :

(48.2)
$$z'_i = \pm \sqrt{z_{i+1} - 1}.$$

Since $\mathcal{O}_{\tau'} \subset f^{-1}(\mathcal{O}_{\tau})$, we have only finitely many (at most 2^{p-2}) options for $\mathcal{O}_{\tau'}$. The choice of \pm signs in (48.2) is determined by the marking $z'_i = h'(c_i)$, and formulas (48.2) express the pulback operator g^* in the local coordinates:

$$g^*: (z_2, \dots, z_{p-1}) \mapsto (z'_1, \dots, z'_{p-2}).$$

Let $\mathcal{Z}_{\tau} = \mathcal{O}_{\tau} \cup \{\infty\}$. If $(\mathbb{C}, \mathcal{O}_{\tau}) \in \mathcal{K}$ then the points of \mathcal{Z}_{τ} are ϵ -separated in the spherical metric for some $\epsilon = \epsilon(\mathcal{K}) > 0$ (see Lemma ??). But then the points of $\mathcal{Z}_{\tau'}$ (i.e., $z'_i = \pm \sqrt{z_i - 1}$ and ∞) are ϵ' -separated for some $\epsilon' > 0$ depending only on ϵ , and the conclusion follows.

48.2.4. Fixed points of g^* .

PROPOSITION 8.10. A branched covering $g : (\mathbb{C}, \mathcal{O}) \to (\mathbb{C}, \mathcal{O})$ is realizable if and only if the pullback operator $g^* : \mathcal{T} \to \mathcal{T}$ has a fixed point.

PROOF. If g is realizable then by definition there is a superattracting quadratic polynomial f_c with the postcritical set \mathcal{O}_c and homeomorphisms h and h' homotopic rel \mathcal{O} such that the diagram is valid:

$$\begin{array}{cccc} (\mathbb{C}, \mathcal{O}) & \longrightarrow & (\mathbb{C}, \mathcal{O}_c) \\ g \downarrow & & \downarrow f_c \\ (\mathbb{C}, \mathcal{O}) & \longrightarrow & (\mathbb{C}, \mathcal{O}_c) \end{array}$$

Comparing it with diagram (48.1) we see that $[h'] = g^*[h]$. But [h] = [h'] by definition of a point in \mathcal{T} . So, [h] is a fixed point of g^* .

Vice versa, assume that a homeomorphism h in diagram (48.1) represents a fixed point of f^* . Then [h'] = [h], which means by definition that after postcomposing h' with a scaling $z \mapsto \lambda z$ we have $h' \simeq h$ rel \mathcal{O} . But then the quadratic polynomial f is Thurston equivalent to g.

48.2.5. Infinitesimal Contraction. Recall from §?? that the cotangent space $T^*_{\tau}\mathcal{T}$ to the Teichmüller space \mathcal{T} is isometric to the space $\mathcal{Q}^1(\mathbb{C} \setminus \mathcal{O})$ of integrable meromorphic quadratic differentials. So, the codifferential

$$Dg(\tau)^* : \mathrm{T}^*_{\tau'}\mathcal{T} \to \mathrm{T}^*_{\tau}\mathcal{T}$$

can be viewed as an operator

$$Dg(\tau)^*: \mathcal{Q}^1(\mathbb{C} \smallsetminus \mathcal{O}_{\tau'}) \to \mathcal{Q}^1(\mathbb{C} \smallsetminus \mathcal{O}_{\tau}).$$

On the other hand, the quadratic map f from diagram (48.1) induces the pushforward operator between the same spaces:

$$f_*: \mathcal{Q}^1(\mathbb{C} \smallsetminus \mathcal{O}_{\tau'}) \to \mathcal{Q}^1(\mathbb{C} \smallsetminus \mathcal{O}_{\tau})$$

(see $\S1.8.3$). It turns out that these two operators are the same:

LEMMA 8.11. Up to the above isometries, the codifferential of $Dg(\tau)^*$ is equal to the push-forward operator f_* .

8. THURSTON THEORY

PROOF. Let us take a smooth vector field v on $\hat{\mathbb{C}}$ representing a tangent vector to \mathcal{T} at τ . The differential $Dg^*(\tau)$ acts on v as the pullback f^*v . Since the $\bar{\partial}$ operator behaves naturally under holomorphic pullbacks, we have $\bar{\partial}(f^*v) = f^*(\bar{\partial}v)$. Using now duality between the pullback and push-forward (Lemma 1.41), we obtain for any quadratic differential $q \in \mathcal{Q}^1(\hat{\mathbb{C}} \setminus \mathcal{O}_{\tau'})$:

$$\langle q, f^*v \rangle = \int q \cdot \bar{\partial}(f^*v) = \int q \cdot f^*(\bar{\partial}v) = \int f_*q \cdot \bar{\partial}v = \langle f_*q, v \rangle.$$

48.2.6. Non-escaping creates the fixed point. The previous discussion implies that the pullback operator g^* is uniformly contracting depending only on the location of the Riemann surface $(\mathbb{C}, \mathcal{O}_{\tau})$ in the moduli space:

LEMMA 8.12. For any compact subset \mathcal{K} in \mathcal{M}_p there exists $\rho = \rho(\mathcal{K}) < 1$ such that $\|Dg^*(\tau)\| \leq \rho$ for any $\tau \in \mathcal{T}_p$ such that $(\mathbb{C}, \mathcal{O}_\tau) \equiv \pi(\tau) \in \mathcal{K}$.

48.2.7. Escaping creates an invariant multicurve.

48.2.8. Structure of an invariant multicurve. Let us consider an invariant multicurve Γ . Let γ be a component of Γ that surrounds the fewest number, say $l \geq 2$, of postcritical points. Then $g^*(\gamma)$ must be a single curve, γ_{-1} , surrounding the same number, l, postcritical points. For the same reason, the pullback $g^*(\gamma_{-1})$, is a single curve, $\gamma_{-2} \in \Gamma$, also surrounding l postcritical points. Continuing this way, we will construct a cycle of curves $\gamma_i \in \Gamma$, $i \in \mathbb{Z}/p\mathbb{Z}$, such that $\gamma_{i-1} = f^*(\gamma_i)$. Of course, one of these curves surrounds 0.

If there is a curve $\gamma' \in \Gamma$ that does not belong to the above cycle, we can construct a new cycle of curves γ'_{-i} , $i \in \mathbb{A}/p\mathbb{Z}$, etc., until we will exhaust all the curves.

We see that Γ is decomposed into a disjoint union of cycles of curves, each of which surrounds the critical point.

48.2.9. Improvement. Under $(f^*)^p$, all the annuli in the cycles

48.3. Realization of abstract Hubbard trees.

212

Part 3

Little Mandelbrot copies

CHAPTER 9

Quadratic-like maps

49. Straightening

49.1. Adjustments. The notion of a quadratic-like map with the fixed domain is too rigid, so we allow *adjustments* of the domains which do not effect the essential dynamics of the map (see Exercises 4.75, 4.77). An appropriate adjustment allows one to improve the geometry of the map:

LEMMA 9.1. Consider a quadratic-like map $f: U \to U'$ with

$$(49.1) \qquad \qquad \text{mod} \ A \ge \mu > 0$$

and $f(0) \in U$. Then there is an adjustment $g: V \to V'$ such that:

(i) The new domains V and V' are bounded by real analytic κ -quasicircles γ and γ' with κ depending only on μ . Moreover, these curves have a bounded (in terms of μ) eccentricity around the origin.

(ii) $\operatorname{mod}(V' \smallsetminus \overline{V}) \ge \mu/2 > 0.$

(iii) g admits a decomposition

(49.2)
$$g = h \circ f_0,$$

where $f_0(z) = z^2$ and h is a univalent function on $W = f_0(V)$ with distortion bounded by some constant $C(\mu)$.

PROOF. Let us uniformize the fundamental annulus A of f by a round annulus, $\phi : \mathbb{A}(1/r, r) \to A$, where $r \ge e^{\mu/2} \equiv r_0$. Then $\gamma' = \phi(\mathbb{T})$ is the equator of A. Consider the disk V' bounded by γ' , and let $V = f^{-1}V'$. Since $f(0) \in V'$, V is a conformal disk and the restriction $f : V \to V'$ is a quadratic-like adjustment of f (see Exercesie 4.75).

Restrict ϕ to the annulus $\mathbb{A}(1/r_0, r_0)$. Take an arc $\alpha = [a, b]$ on \mathbb{T} of length at most $(1-1/r_0)/2$. By the Koebe Distortion and 1/4 Theorems in the disk $\mathbb{D}_{2\delta}(a)$,

$$|\phi(b) - \phi(a)| \ge \frac{|b-a|}{4} |\phi'(a)|; \quad l(\phi(\alpha)) \le K |\phi'(a)| \, l(\alpha),$$

where l stands for the arc length, and K is an absolute constant. Hence $\gamma' = \phi(\mathbb{T})$ is a quasi-circle with the dilatation depending only on $r_0 = r_0(\mu)$ (see Excercise ??).

Applying the same argument to the uniformization of $f^{-1}A$, we conclude that its equator $\gamma = \partial V$ is a quasicircle with bounded dilatation as well.

Since γ and γ' are 0-symmetric κ -quasicircle, the eccentricity of these curves around 0 is bounded by some constant $C(\kappa)$ (see Exercise 2.65). This proves (i).

Property (ii) is obvious since $\operatorname{mod}(V' \smallsetminus \overline{V}) \ge \operatorname{mod} \mathbb{A}(1, r_0) \ge \mu/2$.
Since g is even, it admits decomposition (49.2). Moreover, h admits a univalent extension to the disk $\tilde{W} = f_0(U)$, and

$$\operatorname{mod}(\tilde{W} \smallsetminus W) = 2 \operatorname{mod}(U \smallsetminus V) \ge \mu/2.$$

The Koebe Distortion Theorem (in the invariant form 1.94) completes the proof. \Box

If a q-l map g admits decomposition (49.2), we call it "quadratic up to a bounded distortion".

49.2. Straightening Theorem. If the reader attempted to extend the basic dynamical theory from quadratic polynomials to quadratic-like maps, quite likely he/she had a problem with the No Wandering Domains Theorem. The only known proof of this theorem crucially uses the fact that a polynomial of a given degree depends on finitely many parameters. The flexibility offered by the infinite dimensional space of quadratic-like maps looks at this moment like a big disadvantage. It turns out, however, that the theorem is still valid for quadratic-like maps, and actually there is no need to prove it independently (as well as to repeat any other pieces of the topological theory). In fact, quadratic-like maps do not exibit any new features of topological dynamics, since all of them are topologically equivalent to polynomials (restricted to appropriate domains)!

The proof of this theorem was historically the first application of the so called *quasiconformal surgery* technique. The idea of this technique is to cook by hands a quasiregular map with desired dynamical properties which topologically looks like a polynomial. If you then manage to find an invariant conformal structure for this map, then by the Measurable Riemann Mapping Theorem it can be realised as a true polynomial.

To state the result precisely, we need a few definitions. Two quadratic-like maps f and g are called *topologically conjugate* if they become such after some adjustments of their domains. Thus there exist adjustments $f : U \to U'$ and $g : V \to V'$ and a homeomorphism $h : (U', U) \to (V', V)$ such that the following diagram is commutative:

$$\begin{array}{cccc} U & \longrightarrow & U' \\ h \downarrow & & \downarrow h \\ V & \longrightarrow & V' \end{array}$$

In case when one of the maps is a global polynomial, we allow to take any quadraticlike restriction of it.

If the homeomorphism h in the above definition can be selected quasiconformal (respectively: conformal or affine) then the maps f and g are called quasiconformally (respectively: conformally or affinely) conjugate. Two quadratic-like maps are called *hybrid* equivalent if they are qc conjugate by a map h with $\bar{\partial}h = 0$ a.e. on the filled Julia set K(f).

Remark. The last condition implies that h is conformal on the int K(f). On the Julia set J(f) it gives an extra restriction only if J(f) has positive area.

The equivalence classes of hybrid (respectively: qc, topological etc.) conjugate quadratic-like maps are called *hybrid* (respectively: qc, topological etc.) *classes*.

THEOREM 9.2. Any (conventional) quadratic-like map g is hybrid conjugate to a quadratic polynomial f_c . If J(f) is connected then the corresponding polynomial f_c is unique.

This polynomial f_c is called the *straightening* of g.

COROLLARY 9.3. If g is a quadratic-like map, then:

- (i) There are no wandering components of int K(q);
- (ii) Repelling periodic points are dense in J(g);
- (iii) If all periodic points of g are repelling then K(g) is nowhere dense.

Remark. If J(g) is a Cantor set, then the straightening is not unique. Indeed, by Corollary 6.34 and the Second λ -Lemma (§17.4), all quadratic polynomials f_c , $c \in \mathbb{C} \setminus M$, are qc equivalent. Since their filled Julia sets have zero area, they are actually hybrid equivalent. Hence all of them are "straightenings" of g. We will see however that sometimes there is a preferred choice (see §??).

Existence of the straightening will be proven in the next section, while uniqueness will be postponed until the end of §28.

49.3. Construction of the straightening. The idea is to "mate" g near K(g) with $f_0: z \mapsto z^2$ near ∞ .

Take some r > 1. Consider two closed disks: the disk \overline{U}' endowed with the map $g: \overline{U} \to \overline{U}'$ and the disk $\mathbb{C} \setminus \mathbb{D}_r$ endowed with the map $f_0: \mathbb{C} \setminus \mathbb{D}_r \to \mathbb{C} \setminus \mathbb{D}_{r^2}$. Let us view them as two hemi-spheres, $S_0^2 \equiv \overline{U}'$ and $S_\infty^2 \equiv \mathbb{C} \setminus \mathbb{D}_r$ (see Figure ??). Glue them together by an orientation preserving *equivariant* qc homeomorphism $B: A \to \mathbb{A}[r, r^2]$ between the closed fundamental annuli. Here "equivariance" means that h respects the boundary dynamical relation:

(49.3)
$$B(gz) = f_0(Bz) \text{ for } z \in \partial U.$$

Such a map $B = B_g$ is called a *tubing* of g.

EXERCISE 9.4. Construct a tubing B (using that f is conventional). Do it so that Dil B is bounded in terms of mod A and qc dilatation of the quasicircles U, U'.

In this way we obtain an oriented qc sphere

$$S^2 = S_0^2 \sqcup_B S_\infty^2 \equiv \overline{U}' \sqcup_B (\widetilde{\mathbb{C}} \smallsetminus \mathbb{D}_r)$$

with the atlas of two local charts given by the identical maps $\phi_0 : S_0^2 \to \overline{U}'$ and $\phi_\infty : S_\infty^2 \to \widehat{\mathbb{C}} \setminus \mathbb{D}_r$. Moreover, these hemispheres are quasidiscs in S^2 . For instance, in the local chart ϕ_0 the curve $\gamma := \partial S_\infty^2$ becomes $\phi_0(\gamma) = \partial U$ which is a quasicirlce since f is conventional.

Define now a map $F: S^2 \to S^2$ by letting

$$F(z) = \begin{cases} \phi_0^{-1} \circ g \circ \phi_0(z) & \text{for } z \in \phi_0^{-1} \overline{U} \\ \phi_\infty^{-1} \circ f_0 \circ \phi_\infty(z) & \text{for } z \in \overline{S}_\infty^2 \end{cases}$$

(It is certainly quite a puritan way of writing since the maps ϕ_{-} and ϕ_{+} are in fact identical.) Since *B* is equivariant (49.3), these two formulas match on γ . Hence *F* is a continuous endomorphism of S^2 . Moreover, it is a double branched covering of the sphere onto itself (with two simple branched points at "0" $\equiv \phi_0^{-1}(0)$ and " ∞ " $\equiv \phi_{\infty}^{-1}(\infty)$).

Since $F: S^2 \to S^2$ is holomorphic in the local charts ϕ_0 and ϕ_{∞} , it is quasiregular on $S^2 \smallsetminus \gamma$. Since γ is a quasicircle, it is removable (Corollary 2.68). Hence F is quasiregular on the whole sphere.

EXERCISE 9.5. Let us adjust f so that ∂U is smooth. Then the gluing map B can be chosen so that S^2 is a smooth sphere and the map F is smooth.

We will now construct an *F*-invariant conformal structure μ on S^2 (with a bounded dilatation with respect to the qc structure of the sphere S^2). Start in a neighborhood of ∞ : $\mu | S_{\infty}^2 = (\phi_{\infty})^* \sigma$. Since σ is f_0 -invariant, $\mu | S_{\infty}^2$ is *F*-invariant. Since ϕ_{∞} is qc, $\mu | S_{\infty}^2$ has a bounded dilatation.

Next, look at this structure in the local chart $\phi_0 : S_0^2 \to U'$, and by means of Corollary 4.67 extend it canonically to an invariant structure on the whole sphere S^2 with the same dilatation. We will keep the extension the same notation μ for the extension.

EXERCISE 9.6. Work out details of this canonical extension.

We obtain an *F*-invariant measurable conformal structure μ with bounded dilatation on the whole sphere S^2 . By the Measurable Riemann Mapping Theorem, there exists a qc map $H: (S^2, \mu) \to \hat{\mathbb{C}}$ normalized so that $H(0) = 0, H(\infty) = \infty$ and $H \circ \phi_{\infty}^{-1}(z) \sim z$ as $z \to \infty$. Then the map $f = H \circ F \circ H^{-1}$ is a quadratic polynomial (see §26.1.2) with the critical point at the origin and asymptotic to z^2 at ∞ . Hence $f = f_c: z \mapsto z^2 + c$ for some *c*.

EXERCISE 9.7. Show that $K(f) = H(\phi_0^{-1}K(g))$.

The qc map $H \circ \phi_0^{-1}$ conjugates $g: U \to U'$ to a quadratic-like restriction of f. Moreover, restricting it to K(g), we see that

$$(H \circ \phi_0^{-1})_* \sigma = H_* \mu = \sigma,$$

so that H is a hybrid conjugacy between g and the restriction of f. Thus, f is a straightening of g.

REMARK 9.1. The straightening construction of f_c was uniquely determined by the choice of the tubing $B: A \to \mathbb{A}[r, r^2]$. In fact, one can do better: in the case of connected Julia set, the straightening is independent of the choice of tubing; in the disconnected case, it depends only on the tubing position of the critical value (see §51).

49.4. Addendum to the straightening construction. Here we will refine the straightening construction in several ways. In particular, we will extend the tubing to a bigger annulus, through a series of liftings (similarly to the extension of the Böttcher function carried in §32.1).

49.4.1. Tubing equipotentials and rays. The tubung $B : A \to \mathbb{A}[r, r^2]$ plays a role of the Böttcher coordinate for the quadratic-like map g. In particular, we can use it to define equipotentials and rays for f as pullbacks by B of the round circles and radial intervals in $\mathbb{A}[r, r^2]$. In this way we obtain two foliations in the fundamental annulus A. There are natural radia/levels assigned to the equipotentials and external angles assigned to the rays. (For instance, the the boundary equipotential ∂U has radius r and level $t = \log r$.)

By means of the dynamics, we can now extend these foliations to invariant (singular) foliations in $U' \smallsetminus K(g)$. If K(g) is connected then these foliations are

in fact non-singular. In the disconnected case, they have simple cross-singularities at the critical point 0 and its iterated preimages. In this case, the figure-eight equipotential passing through 0 is called *critical*. We let $\Omega_g \subset U'$ be the topological annulus bounded by this equipotential and the external boundary $\partial U'$. In the connected case, we let $\Omega_g = U' \smallsetminus K(g)$ (everything is similar to the polynomial case).

49.4.2. Equivariant extension of the tubing. Similarly to the Böttcher coordinate, the tubing can be equivariantly extended to the domain Ω_g (compare §32.1). It is based on a simple lifting step:

LEMMA 9.8. Let us consider a nest of two Riemann surfaces $\Omega \subset \Omega'$ with boundary. We assume that $\Gamma' := \partial \Omega'$ and $\Gamma := \partial \Omega$ are quasicircles, and that $A := \Omega' \setminus \operatorname{int} \Omega$ is a closed annulus bounded by Γ and Γ' (its "inner" and "outer" boundary components respectively). Let $g : \Omega \to \Omega'$ be a holomorphic double covering map such that $g(\Gamma) = \Gamma'$.

Consider also another map $\tilde{g} : \tilde{\Omega} \to \tilde{\Omega}'$ with the same properties (all corresponding objects for \tilde{g} are marked with "tilde"). Let $h : A \to \tilde{A}$ be an equivariant K-qc homeomorphism, i.e., $h(gz) = \tilde{g}(hz)$ for $z \in \Gamma$.

Assume A and \tilde{A} do not contain the critical values of g and \tilde{g} (respectively). The $A^1 := g^{-1}(A)$ and $\tilde{A}^1 := g^{-1}(\tilde{A})$ are annuli attached to A and \tilde{A} , and h extends to an equivariant K-qc homeomorphism $H : A \cup A^1 \to \tilde{A} \cup \tilde{A}^1$.

PROOF. Since A does not contain the critical values of g, A^1 is an annulus. Since $g(\Gamma) = \Gamma'$, A^1 is attached to A along Γ , so together they form an annulus $A \cup A^1$.

By the general lifting theory, h lifts to a homeomorphism $h^1: A^1 \to \tilde{A}^1$ in two ways determined by the choice of value of h_1 at one point. But since $h: A \to \tilde{A}$ is equivariant, $h|\Gamma$ is a lift of $h|\Gamma'$. Hence the lift h_1 can be selected so that it coincides with h on Γ , and we obtain a single equivariant homeomorphism H: $A \cap A^1 \to \tilde{A} \cup \tilde{A}^1$.

Since g is holomorphic, h_1 is K-qc. Since Γ is a quasicircle, H is K-qc as well (by Lemma 2.68).

By iterating the Lifting Construction, we obtain:

COROLLARY 9.9. Let $g: U \to U'$ and $\tilde{g}: \tilde{U} \to \tilde{U}'$ br two quadratic-like maps with fundamental annuli A and \tilde{A} . Let $h: A \to \tilde{A}$ be a homeomorphism between fundamental annuli equivariant on the boundary. If K(g) and $K(\tilde{g})$ are connected then h extends uniquely to an external conjugacy $U' \smallsetminus K(f) \to \tilde{U}' \searrow K(\tilde{g})$. Moreover, if h is K-qc then so is the extension.

COROLLARY 9.10. Let $B : A \to \mathbb{A}[r, r^2]$ be a K-qc tubing for g. Then it extends to an equivariant K-qc map $\Omega_g \to \mathbb{D}_{r^2} \setminus \mathbb{D}_R$, where R = 1 in the connected case and R > 1 in the Cantor case.

This exension will be denoted B and called a "tubing" as well. Its equvariance means that $B(gz) = B(z)^2$ for $z \in \Omega_g \cap U$.

Note that in the Cantor case, we have $g(0) \in \Omega_g$, so the point B(g(0)) is well defined. We call it the *tubing position of the critical value*.

9. QUADRATIC-LIKE MAPS

49.4.3. Böttcher coordinate for the straightening. The map $B \equiv \phi_{\infty} \circ H^{-1}$ in the above construction is the Böttcher coordinate for f on $\Omega := H(S_{\infty}^2)$. Indeed, $B|\Omega$ is conformal (since both ϕ_{∞} and H transfer the conformal structure $\mu|S_{\infty}^2$ to σ) and B conjugates f to $f_0: z \mapsto z^2$.

Since $B(\partial\Omega) = \mathbb{T}_r$, $\partial\Omega = E_r$ is the equipotential of radius r for f. Thus, we have conjugated $g: U \to U'$ to $f: \Omega_c(r) \to \Omega_c(r^2)$ where $\Omega_c(r)$ is the subpotential disk of radius r for f_c (see §32.2).

In the Cantor case, Figure 49.3 (with the extended tubing) shows that the tubing position of the critical value for a polynomial-like map g coincides with Böttcher position of the critical value for its straightening f_c :

(49.4)
$$B_g(g(0)) = B_c(c).$$

49.4.4. Dilatation. Finally, let us dwell on an important issue of a bound on the dilatation of the qc homeomorphism that straightens g.

LEMMA 9.11. Let $g: U \to U'$ be a quadratic-like map with $\text{mod } A \ge \delta > 0$. Then g is hybrid conjugate to a straightening f_c by a K-qc map whose dilatation K depends only on δ .

PROOF. Let us first adjust g according to Lemma 9.1 (keeping the same notations for the domains U and U').

Let us now follow the proof of the Straightening Theorem. Look at the conformal structure $\mu | S_0^2$ in the local chart ϕ_0 , i.e., consider the conformal structure $\nu = (\phi_0)_*(\mu | S_0^2)$ on U'. On $U' \smallsetminus K(g)$, it is obtained by pulling back (by the conformal g-dynamics) the structure $B^*(\sigma)$ from the fundamental annuus A. On K(g) it is equal to the standard structure σ . Hence the dilatation of ν is equal to the dilatation of the tubing B.

The qc map $H \circ \phi_0^{-1}$ conjugating $g: U \to U'$ to $f: D_r \to D_{r^2}$ transfers ν to σ . Hence its dilatation is also equal to Dil(B). But by the latter is bounded in terms of δ (see Exercise 9.4).

49.4.5. Standard equipment of q-l maps. Due to the Straightening Theorem, we can equip q-l maps with the strandard amunition of quadratic polynomials. Notice first that the α - and β - fixed points are well defined as long as the Julia set J(g) is connected (the α -fixed point is either non-repelling or the dividing repelling one). Moreover, as pointed out in §49.4.1, once we select a tubing, we obtain the external foliations of rays and equipotentials in $U' \smallsetminus K(g)$. Under the straightening conjugacy they are mapped to the corresponding foliations for the polynomial f_c . Since all the conjugacies agree on the Julia set (see Corollary 9.16 below), the landing properties of the rays are independent of the particular choice of the tubing. In particular, the β -fixed point is always the landing point for the 0-ray.

49.5. Concept of renormalization. The primarily motivation for introducing quadratic-like maps comes from the idea of renormalization, which is a central idea in contemporary theory of dynamical systems.

A quadratic-like map $f: U \to U'$ is called *renormalizable with period* p if there is a topological disk $V \ni 0$ such that all the domains f^iV , $i = 0, 1, \ldots, p-1$, are contained in U, the map $g := (f^p: V \to f^p(V))$ is quadratic-like with connected Julia set (see Figure ??), and a technical "almost disjointness" property explained below is satisfied.

Let $V'_i = f^i V$, i = 1, ..., p; then f^{p-i} maps V'_i univalently onto $V'_p = f^p V$. Let $\phi_i : V_i \to V'_p$ stands for this univalent map. Let $V_i \Subset V'_i$ be the pullback of Vunder ϕ_i . Then the map $g_i := (f^p : V_i \to V'_i)$ is a double covering, and thus, it is a quadratic-like map (note that $g_p = g$). Moreover, ϕ_i conjugates g_i to $g : V \to V'$. Hence the Julia sets $J(f_i)$ and $\mathcal{K}_i = K(g_i)$ are connected and their α and β -fixed points are well defined.¹

These Julia sets are called *little Julia sets* (associated with a particular renormalization scheme under consideration). Almost disjointness property requires that the little Julia sets can touch only at their β -fixed points (see Figure ??). This completes the definition of a renormalizable map.

The quadratic-like map $g: V \to V'$ is called the *pre-renormalization* of f. In fact, all pre-renormalizations define the same quadratic-like germ:

LEMMA 9.12. Let $f: U \to U'$ and $\tilde{f}: \tilde{U} \to \tilde{U}'$ be two quadratic-like maps representing the same ql germ. If f is renormalizable with period p, then so is \tilde{f} , and the corresponding pre-renormalizations represent the same ql germ. The little Julia sets \mathcal{K}_i are canonically defined.

Thus, we can promote the above pre-renormalization to the *renormalization* $R = R_p$ acting on ql germs (considered up to rescaling).

The above discussion applies to quadratic polynomials by considering the corresponding quadratic-like germs.

The renormalization is called *primitive* if the little Julia sets \mathcal{K}_i are disjoint and is called *satellite* otherwise. In the latter case, a union of little Julia sets that share a common point *a* is called a *bouquet* of little Julia sets ("centered" at *a*).

Each renormalization comes together with certain combinatorial data. It accounts for the renormalization period p and the "positions" of the little Julia sets \mathcal{K}_i on the big one, K(f). More precisely, let us consider a graph whose vertices are the little Julia sets \mathcal{K}_i and whose edges are defined as follows: \mathcal{K}_i is connected to \mathcal{K}_j if these little Julia sets belong to the same component of $K(f) \setminus \bigcup_{m \neq i,j} \mathcal{K}_m$. It turns out that this graph is a tree. It will be called the *Hubbard tree* of the renormalization in question.

We can now define the *combinatorics* of the renormalization as its Hubbard tree. In §?? we will show that the renormalizable quadratic maps f_c with a given combinatorics assemble a little Mandelbrot copy M' (see Figure ??).

Let $p_0 < p_1 < \ldots$ be the sequence of all renormalization periods of a map f. If this sequence has length at least n then f is n times renormalizable. In particular, if it has infinite length then f is *infinitely renormalizable*. If it has zero length (no periods) then f is *non-renormalizable*.

LEMMA 9.13. If (p_n) is the sequence of all renormalization periods of a map f, then p_{n+1} is a multiple of p_n , the map $f_n := R_{p_n} f$ is renormalizable with relative period $q_n = p_{n+1}/p_n$, and $R_{q_n} f_n = f_{n+1}$.

We will refer to $R_{p_0}f$ as the first renormalization of f, and we will usually reserve notation Rf for this one. Then $R_{p_{n-1}}f = R^n f$ is the *n*-fold renormalization of f.

¹As long as it does not cause a confusion, we will skip "filled" when referring to the \mathcal{K}_i .

9. QUADRATIC-LIKE MAPS

49.6. Higher degree case. A polynomial-like map of degree d is a branched covering $f: U \to U'$ of degree d between two nested topological discs $U \subseteq U'$. The basic theory of quadratic-like maps developed above extends to the higer degree case in the straigtforward way.

50. External map

Before passing to the uniqueness part of the Straightening Theorem, let us dwell on an important relation between quadratic-like and circle maps.

50.0.1. Connected case. To any quadratic-like map $f: U \to U'$ one can naturally associate an expanding circle map g of class \mathcal{E} which captures dynamics outside the Julia set. For this reason g is called the *external map* of f.

The construction is very simple if the Julia set J(f) is connected. In this case the basin of infinity $D_f(\infty) = \mathbb{C} \setminus K(f)$ is simply connected and can be conformally mapped onto the complement of the unit disk:

$$R: \mathbb{C} \smallsetminus K(f) \to \mathbb{C} \smallsetminus \overline{\mathbb{D}}.$$

Let $\Omega = R(V \setminus K(f))$, $\Omega' = R(V' \setminus K(f))$. These are two conformal annuli with smooth boundary. Moreover, the have a common inner boundary, the unit circle \mathbb{T} , while the outer boundary of Ω is contained in Ω' . Conjugating f by R we obtain a holomorphic double covering

$$g: \Omega \to \Omega', \quad g(z) = R \circ f \circ R^{-1}(z) \quad \text{for} \quad z \in \Omega.$$

By Lemma 4.91, g can be extended to an expanding circle map of class \mathcal{E} .

In fact, this map is not uniquely defined since the Riemann map R is defined up to post-composition with rotation $z \mapsto e^{2\pi i \theta} z$, $0 \leq \theta < 2\pi$. A natural way to *normalize* g is to put its fixed point β to $1 \in \mathbb{T}$.

Note also that if f is replaced by an affinely conjugate map $A^{-1} \circ f \circ A$, where $A: z \mapsto \lambda z, \lambda \in \mathbb{C}^*$, then the Riemann map R is replaced by $R \circ A$, and the external map g remains the same. Thus, to any quadratic-like map f (with connected Julia set) prescribed up to an affine conjugacy corresponds an expanding circle map g well-defined up to rotation conjugacy.

 $50.0.2.\ General \ case.$ In the case of disconnected Julia set the construction is more subtle.

Take a fundamental annulus $A = U' \setminus U$ with real analytic boundary curves $E = \partial U'$ and $I = \partial U$. Then $f : I \to E$ is a real analytic double covering.

Let $\mu = \mod A$. Let us consider an abstract double covering $\xi_1 : A_1 \to A$ of an annulus A_1 of modulus $\mu/2$ over A. Let I_1 and E_1 be the "inner" and "outer" boundary of A_1 , i.e., ξ_1 maps I_1 onto I and E_1 onto E. Then there is a real analytic diffeomorphism $\theta_1 : E_1 \to I$ such that $\xi_1 = f \circ \theta_1$. This allows us to stick the annulus A_1 to the domain $\mathbb{C} \smallsetminus U$ bounded by I. We obtain a Riemann surface $T_1 = (\mathbb{C} \smallsetminus U) \cup_{\theta_1} A_1$. Moreover, the maps f on A and ξ_1 on A_1 match to form an analytic double covering $f_1 : A_1 \to A$.

This map f_1 restricts to a real analytic double covering of the inner boundary of A_1 onto its outer boundary. This allows us to repeat this procedure: we can attach to the inner boundary of T_1 an annulus A_2 of modulus $\frac{1}{4}\mu$, and extend f_1 to the new annulus T_2 . Proceeding in this way, we will construct a Riemann surface

(50.1)
$$T \equiv T^{A}(f) = \lim T_{n} = (\mathbb{C} \setminus U) \cup_{\theta_{1}} A_{1} \cup_{\theta_{2}} A_{2} \dots$$

and an analytic double covering $F: \bigcup_{n\geq 1}A_n \to \bigcup_{n\geq 0}A_n$ extending f.

The inner end of T can be represented by a puncture or by an ideal circle. But in the former case, after filling that pucture we would obtain a superattracting fixed point α (since the map F is a double covering near α). This would contradict to the property that the trajectories of F are repelled from the inner end of T.

Thus, the inner end of T is not a puncture but a circle. Hence T can be uniformized by $\mathbb{C} \setminus \mathbb{D}$ (with the inner ideal boundary uniformized by the unit circle \mathbb{T}). Now by the reflection principle, this conformal representation of F can be extended to an analytic expanding endomorphism $g \equiv g_A : \mathbb{T} \to \mathbb{T}$.

For a given choice of the fundamental annulus A, the map $g_A : V \to V'$ (which comes together with the domains (V, V')) is well-defined up to rotation. Indeed, for two such maps g_A and \tilde{g}_A , by construction there is a conformal isomorphism $h : \mathbb{C} \setminus \mathbb{D} \to \mathbb{C} \setminus \mathbb{D}$ conjugating them on an outer neighborhood of the circle. Reflecting h to the unit disk, we conclude that h is a rotation conjugating g_A and g'_A near the circle.

EXERCISE 9.14. Show that in the connected case this construction leads to the same result as the construction of $\S 50$.

51. Uniqueness of the straightening

51.0.3. Connected case. Let us first show that an "external automorphisms" of a quadratic-like map admits a continuous extension to the Julia set by the identity (compare with Lemma 4.93).

LEMMA 9.15. Let $f: U \to U'$ be a quadratic-like map with connected Julia set. Let $W \subset U$ and $W' \subset U$ be two (open) annuli whose inner boundary is J(f). Let $h: W \to W'$ be an automorphism of f. Then h admits a continuous extension to a map $W \cup J(f) \to W' \cup J(f)$ identical on the Julia set.

PROOF. Consider the Riemann mapping $\phi : \mathbb{C} \setminus K(f) \to \mathbb{C} \setminus \overline{\mathbb{D}}$ and the external circle map $g : V \to V', g | V \setminus \overline{\mathbb{D}} = \phi \circ f \circ \phi^{-1}$. Transfer the annuli W and W' to the g-plane. We obtain two annuli $\Omega = \phi(W)$ and $\Omega' = \phi(W')$ in $V \setminus \overline{\mathbb{D}}$ attached to the unit circle \mathbb{T} . Of course, the homeomorphism $k : \Omega \to \Omega', k = \phi \circ h \circ \phi^{-1}$, commutes with g.

By Lemma 4.94, k moves points near \mathbb{T} bounded hyperbolic distance:

$$\rho_{\mathbb{C} \smallsetminus \bar{\mathbb{D}}}(k(z), z) \le R.$$

Since the Riemann mapping $\phi : \mathbb{C} \setminus \overline{\mathbb{D}} \to \mathbb{C} \setminus K(f)$ is a hyperbolic isometry, the same is true for h:

$$\rho_{\mathbb{C}\smallsetminus K(f)}(h(z), z) \le R \quad \text{for } z \in W \text{ near } J(f).$$

By Proposition 1.78, the Euclidean distance |z - h(z)| goes to 0 as $z \to J(f)$. It follows that the extension of h by the identity to the Julia set is continuous.

COROLLARY 9.16. Let f and \tilde{f} be two quadratic-like maps, and let h and h' be two homeomorphisms conjugating f to \tilde{f} in some neighborhoods of the filled Julia sets. Then h = h' on J(f).

PROBLEM 9.17. Assume that quadratic polynomials f and \tilde{f} are conjugate on the Julia sets only. Is the conjugacy unique?

9. QUADRATIC-LIKE MAPS

Let us now put together the above results:

THEOREM 9.18. Let us consider two quadratic-like maps $f : U \to U'$ and $\tilde{f}: \tilde{U} \to \tilde{U}'$ with connected Julia sets. Assume that they are topologically conjugate near their Julia sets by a homeomorphism $\psi: V \to \tilde{V}$. Assume also that we are given an equivariant homeomorphism $H: A \to \tilde{A}$ between the (closed) fundamental annuli of f and \tilde{f} .

Then there exists a unique homeomorphism $h: U' \to \tilde{U}'$ conjugating f to \tilde{f} , coinciding with ψ on the Julia set J(f), and coinciding with H on A.

If H is qc, then $h|U \setminus K(f)$ is also qc with the same dilatation. If both H and ψ are qc, then h is qc, and

$$\operatorname{Dil}(h) = \max(\operatorname{Dil} H, \operatorname{Dil}(\psi | K(f))).$$

In particular, if ψ is a hybrid equivalence, then Dil(h) = Dil(H).

PROOF. By the Lifting Construction of Corollary 9.9, H admits a unique equivariant extension to a homeomorphism $h: U \smallsetminus K(f) \to \tilde{U} \backsim K(\tilde{f})$. This extension continuously matches with ψ on the filled Julia set. Indeed, $\psi^{-1} \circ h$ commutes with f on some external neighborhood of K(f). By Corollay 9.15, this map continuously extends to the filled Julia set as the identity. Hence h continuously extends to the filled Julia set as ψ .

If H is qc then $h|U \smallsetminus K(f)$ is qc with the same dilatation (Corollary 9.9). All the rest follows from Bers' Lemma.

Of course, we can always construct an equivariant qc map H between the fundamental annuli. Hence if two quadratic-like maps are topologically equivalent, then the conjugacy can be selected quasi-conformal outside the filled Julia set. If they are hybrid equivalent, then the dilatation of the conjugacy is completely controlled by the dilatation of H, which is in turn controlled by the geometry of the fundamental annuli (see Lemma 9.11). In the polynomial case we can do even better:

COROLLARY 9.19. Consider two quadratic polynomials $f : z \mapsto z^2 + c$ and $\tilde{f} : z \mapsto z^2 + \tilde{c}$ with connected Julia sets. If they are topologically conjugate near their filled Julia sets by a map h_0 , then there exists a unique global conjugacy $h : \mathbb{C} \to \mathbb{C}$ that coincides with h_0 on K(f) and is conformal on the basin of ∞ . If h_0 is qc then so is h, and Dil $h = \text{Dil}(h_0 | K(f))$. If h_0 is hybrid then h = id and $f = \tilde{f}$.

PROOF. By Theorem 5.13, the Riemann-Böttcher map $B_f: D_f(\infty) \to \mathbb{C} \setminus \overline{\mathbb{D}}$ conjugates f to $z \mapsto z^2$, and similarly for \tilde{f} . Hence the composition

(51.1)
$$H = B_{\tilde{f}}^{-1} \circ B_f : D_f(\infty) \to D_{\tilde{f}}(\infty)$$

conformally conjugates f to \tilde{f} on their basins of ∞ . By the previous theorem, this conjugacy matches with the topological conjugacy on the filled Julia set giving us a desired global conjugacy h.

Moreover, If f and \overline{f} are hybrid equivalent, then Dil(h) = 0 a.e. on \mathbb{C} . By Weyl's Lemma, h is conformal and hence affine. As h(0) = 0 and $h(z) \sim z$ near ∞ , we conclude that h = id.

The uniqueness of h follows from the fact that id is the only conformal automorphism $\mathbb{C} \smallsetminus \overline{\mathbb{D}} \to \mathbb{C} \smallsetminus \overline{\mathbb{D}}$ commuting with $z \mapsto z^2$ (and hence (51.1) is the only conformal isomorphism $D_f(\infty) \to D_{\tilde{f}}(\infty)$ conjugating f to \tilde{f}). \Box

The last statement of the above Corollary gives the uniqueness part of the Straightening Theorem in the connected case.

51.0.4. Disconnected case.

PROPOSITION 9.20. For a q-l map g with disconnected Julia set, the tubing position of the critical value, $B_g(g(0))$, determines the straightening $f_c: z \mapsto z^2 + c$.

PROOF. By (49.4), the tubing position of the critical value is equal to the Böttcher position $B_c(c)$ of the critical value for f_c . But by Theorem 6.10, the latter is equal to the Riemann position $\Phi_M(c)$ of the parameter c. (Recall that $\Phi_M : \mathbb{C} \setminus M \to \mathbb{C} \setminus \overline{\mathbb{D}}$ is the Riemann mapping on the complement of the Mandelbrot set.) As $\Phi_M(c)$ determines c, the conclusion follows.

51.1. Mating of f_c with $g \in \mathcal{E}$. We have associated to any quadratic-like map f with connected Julia set its straightening $f_c : z \mapsto z^2 + c$ and its external map $g : \mathbb{T} \to \mathbb{T}$. Recall that quadratic-like maps are considered up to affine conjugacy, while expanding circle maps are considered up to rotation. These maps can be normalized so that

$$f(z) = c + z^2 + \mathrm{h.o.t.}$$

near the origin, while g has his fixed point at 1. Now we will reverse the above construction constructing a *mating* between f_c and g:

PROPOSITION 9.21. Given a parameter $c \in M$ and an expanding circle map $g: V \to V'$, there exists a unique quadratic-like map f (up to affine conjugacy) such that f_c and g are the straightening and the external map of f respectively.

PROOF. The proof is similar to the proof of the Straigtening Theorem, so we will just sketch it.

Existence. Let us consider a quadratic-like restriction $f_c: U \to U'$ of our quadratic polynomial (e.g., we can select U as a disk bounded by by some equipotential of f_c). Take some equivariant diffeomorphism $h_0: U' \smallsetminus U \to V' \smallsetminus V$ and extend it by Lemma ?? to an equivariant diffeomorphism $h: U' \backsim K(f) \to V' \backsim \overline{\mathbb{D}}$. Now glue two hemi-spheres $S^2_+ := U'$ and $S^2_- := \hat{\mathbb{C}} \smallsetminus \overline{\mathbb{D}}$ by means of h to obtain a qc sphere S^2 . Define a map

$$F: U \sqcup_h (V \smallsetminus \overline{\mathbb{D}}) \to U' \sqcup_h (V' \smallsetminus \overline{\mathbb{D}})$$

as f_c on $U \subset S^2_+$ and as g on $V \setminus \overline{\mathbb{D}} \subset S^2_-$. It is a well defined quasiregular double branched covering. Moreover, it preserves the conformal structure μ which is standard on $K(f) \subset S^2_+$ and on S^2_- . My means of the Measurable Riemann Mapping Thorem, F can be turned into the desired quadratic-like map f.

Uniqueness. Assume that two quadratic-like maps f and \tilde{f} have the same straightenings and the same normalized external maps. Then they are hybrid conjugate by a qc map $h: U \to U'$ near their filled Julia sets, and are conformally conjugate by a map $\phi: \mathbb{C} \setminus K(f) \to \mathbb{C} \setminus K(\tilde{f})$. By Theorem 9.18, these two conjugacies match on the Julia set and glue together into a global conformal (and hence affine) map $\mathbb{C} \to \mathbb{C}$.

EXERCISE 9.22. Supply the missing details in the above proof.

9. QUADRATIC-LIKE MAPS

EXERCISE 9.23. A quadratic-like map is a quadratic polynomial if and only if its external map is $z \mapsto z^2$.

52. Weak q-l maps and Epstein class

52.1. Weak quadratic-like maps. Let us now introduce a slightly generalized notion of quadratic-like map.

DEFINITION 9.24. A holomorphic double branched covering $f: U \to U'$ between two nested conformal disks $U \subset U'$ is called a *weak quadratic-like map*.

The only difference compared with standard q-l maps is that U may not be compactly contained in U' (so there may be no space between the domain U and the range U'). Of course, such a map also has a unique critical point which will be placed at the origin, unless otherwise is explicitly assumed. The filled Julia set $K(f) \subset U$ is defined in the same way as in the standard situation, as the set of points that never escape from U. It is obviously completely invariant, so the map $f: K(f) \to K(f)$ is two-to-one: every point $z \in K(f)$ except the critical value c = f(0) has two preimages in K(f). However, K(f) may be non-compact (and even non-closed). In fact, this is essentially the only difference between the two settings:

PROPOSITION 9.25. Let $f: U \to U'$ be a weak q-l map whose Julia set K(f) is a compact continuum. Then f admits a q-l restriction $V \to V'$ with the same Julia set. Moreover, if $mod(U' \setminus K(f)) \ge \nu > 0$ then $mod(V' \setminus V) \ge \mu(\nu) > 0$.

PROOF. Since K(f) is a compact continuum, we can uniformize its compenent by the Riemann map $\phi : \mathbb{C} \setminus \overline{\mathbb{D}} \to \mathbb{C} \setminus K(f)$, and then construct the external circle map $g : (\Omega, \mathbb{T}) \to (\Omega', \mathbb{T})$ exactly as in the standard situation (see §50). Here Ω' is the \mathbb{T} -symmetric annulus obtained by symmetrization of $\phi^{-1}(U' \setminus K(f))$.

By the Definitive Schwarz Lemma (see Corollary 1.101 and Exercise 1.102), g is strictly expanding in the hyperbolic metric of Ω' . Let $W' \Subset \Omega'$ be the hyperbolic 1-neighborhood of \mathbb{T} . Since $W' \Subset \Omega'$, g | W' is strongly expanding. Since the circle \mathbb{T} is g-invariant, the full preimage $W := g^{-1}(W')$ is contained in a ρ -neighborhood of \mathbb{T} with $\rho < 1$. Hence $W \Subset W'$ and $g : W \to W'$ is a double covering.

Let now $V = \phi(W) \cup K(f)$ and $V' = \phi(W') \cup K(f)$. Then $f : V \to V'$ is the desired q-l restriction of f.

Let us now quantify this construction. Without loss of generality, we can assume that $\mu \leq 1$ and $\mu \leq \operatorname{mod}(U' \smallsetminus K(f) \leq 2\mu$ (by restricting f to an appropriate preimage $f^{-n}(U)$). Then mod $\Omega' = 2 \operatorname{mod}(U' \smallsetminus K(f)) \in [2\mu, 4\mu]$ and

$$\operatorname{mod} \Omega = \frac{1}{2} \operatorname{mod} \Omega' \in [\mu, 2\mu].$$

These conditions determine a compat family of nested pairs of annuli $\Omega' \supset \Omega \supset \mathbb{T}$ (in the Carathéodory topology, see §??). Since the hyperbolic metrics depend continuously on the annuli, the embeddings $\Omega...\Omega'$ are uniformly contracting. This implies that the map g|W' is uniformly expanding, and all the conclusions follow.

CHAPTER 10

Quadratic-like families

42. Fully equipped families

42.1. Definitions. Let $\Lambda \subset \mathbb{C}$ be a domain in the complex plane. A quadraticlike family **g** over Λ is a family of quadratic-like maps $g_{\lambda} : U_{\lambda} \to U'_{\lambda}$ depending on $\lambda \in \Lambda$ such that:

• The tube $\mathbb{U} = \{(\lambda, z) : \lambda \in \Lambda, z \in U_{\lambda}\}$ is a domain in \mathbb{C}^2 ;

• $g_{\lambda}(z)$ is holomorphic in two variables on \mathbb{U} .

As usual, we assume that the critical point of each f_{λ} is located at the origin 0, and that U_{λ} and U'_{λ} are 0-symmetric quasidisks.

We will now make several additional assumptions. The first of them is minor. We say that **g** extends beyond \mathbb{U} if there exists a domain $\Lambda' \supseteq \Lambda$ and a quadratic-like family $G_{\lambda} : V_{\lambda} \to V'_{\lambda}$ over Λ' such that for $\lambda \in \Lambda$, g_{λ} is an adjustment (see §49.1) of G_{λ} .

We call a quadratic-like family $\mathbf{g}: U_{\lambda} \to U'_{\lambda}$ over Λ proper if

- \mathbf{g} admits an extension beyond \mathbb{U} ;
- For $\lambda \in \partial \Lambda$, $g_{\lambda}(0) \in \partial U'_{\lambda}$.

Obviously $g_{\lambda}(0) \neq 0$ for $\lambda \in \partial \Lambda$, so that we have a well defined winding number of the curve $\lambda \mapsto g_{\lambda}(0), \ \lambda \in \partial \Lambda$, around 0. We call it the winding number of **g** and denote $w(\mathbf{g})$. A proper family **g** is called *unfolded* if $w(\mathbf{g}) = 1$. By the Argument Principle, any proper unfolded quadratic-like family has a unique parameter value \circ such that f_{\circ} has a superattracting fixed point, i.e., $f_{\circ}(0) = 0$. We will select \circ as the *base point* in Λ .

Finally, we want the fundamental annulus $A_{\lambda} = \overline{U}'_{\lambda} \smallsetminus U_{\lambda}$ of g_{λ} to move holomorphically with λ . So, assume that there is an equivariant holomorphic motion $h_{\lambda} : A_{\circ} \to A_{\lambda}$, i.e., such that

$$h_{\lambda}(g_{\circ}z) = g_{\lambda}(h_{\lambda}(z)) \quad \text{for} \quad z \in \partial U_{\circ}.$$

Moreover, we will make a technical

Assumption H: The motion of any compact subset $Q \subset \overline{U}'_{o} \setminus \overline{U}_{o}$ extends to a slightly bigger disk $\Lambda_Q \supseteq \Lambda$.

REMARK 10.1. Note that the motion of ∂U_{\circ} cannot be extended beyond Λ since for $\lambda \in \partial \Lambda$ the boundary curve ∂U_{λ} pinches at the critical point 0 (becoming a figure-eight curve).

Denote this holomorphic motion by \mathbf{h} . We say that the quadratic-like family \mathbf{g} is *equipped* with the holomorphic motion \mathbf{h} . Sometimes we will use notation (\mathbf{g}, \mathbf{h}) for an equipped quadratic-like family.

For equipped families, there is a natural choice of tubing (see §49.3) holomorphically depending on λ . Namely, select any tubing $B_{\circ}: A_{\circ} \to \mathbb{A}[r, r^2]$ for the base

point, and then let

$$B_{\lambda} = B_{\circ} \circ h_{\lambda}^{-1}$$

These are tubings since the holomorphic motion h_λ is equivariant.

The Mandelbrot set of the quadratic-like family is defined as

$$M(\mathbf{g}) = \{\lambda \in \Lambda : J(g_{\lambda}) \text{ is connected}\}.$$

If **g** is proper, then $M(\mathbf{g})$ is compactly contained in Λ .

Let us finish with a few terminological and notational remarks. Let $\pi : \mathbb{C}^2 \to \mathbb{C}$ stand for the projection onto the first coordinate. We call a set $\mathbb{U} \subset \mathbb{C}^2$ a *tube over* $\Lambda = \pi(\mathbb{U}) \subset \mathbb{C}$ if it is a fiber bundle over Λ whose *fibers* $U_{\lambda} = \mathbb{U} \cap \pi^{-1}\lambda$ are Jordan disks (either open or closed). For $X \subset \Lambda$, we let $\mathbb{U}|X = \mathbb{U} \cap \pi^{-1}X$.

42.2. Restricted quadratic family. In this section we will show that the quadratic family $\{f_c\}_{c \in \mathbb{C}}$ can be naturally restricted to a proper unfolded equipped quadratic-like family.

Fix some r > 1. Restrict the parameter plane \mathbb{C} to the subpotential disk $D \equiv D_{r^2}$ bounded by the parameter equipotential of radius r^2 (see §42.4). According to formula (29.1), this parameter domain is specified by the property that

$$f_c(0) \in \Omega_c(r^2) \equiv \Omega_c'$$

(Recall that $\Omega_c(\rho)$ stands for the dynamical subpotential disk of radius ρ , see §32.2). Hence for $c \in D$, f_c restricts to a quadratic-like map $f_c : \Omega_c \to \Omega'_c$, where $\Omega_c \equiv \Omega_c(r)$. These quadratic-like maps obviously form a quadratic-like family over D, which we will call a *restricted quadratic family*.

The restricted quadratic family is proper. The first property of the definition is obvious. The main property, $f_c(0) \in \partial \Omega'_c$ for $c \in \partial D$, follows from formula (29.1). The winding number of this family is equal to 1. Indeed, when the *parameter* c runs once along the boundary ∂D , the *critical value* $c = f_c(0)$ runs once around $0 \in D$.

The restricted quadratic family is equipped with the holomorphic motion of the fundamental annulus given by the Böttcher maps. Select 0 as the base point in D and let

(42.2)
$$B_c^{-1} : \mathbb{A}[r, r^2] \to \overline{\Omega}'_c \smallsetminus \Omega_c$$

(note that $\mathbb{A}[r, r^2] = \overline{\Omega}'_0 \smallsetminus \Omega_0$). Since the Böttcher function $B_c^{-1}(z)$ is holomorphic it two variables (Step 1 in §29.3), $\{B_c^{-1}\}_{c \in D}$ is a biholomorphic motion (not only in c but also in z). We call it the Böttcher motion.

Finally, note that for any slightly smaller annulus $\mathbb{A}[\rho, r^2]$, $\rho > r$, the Böttcher motion (42.2) extends to a slightly bigger subpotentia domain, $D_{\rho^2} \subseteq D$.

Thus, the restricted quadratic family satisfies all the properties required for an equipped proper unfolded quadratic-like family.

42.3. Straightening of quadratic-like families. The Mandelbrot set $M(\mathbf{g})$ of any quadratic-like family \mathbf{g} can be canonically mapped to the genuine Mandelbrot set M. Namely, by the Straightening Theorem, for any $\lambda \in M(\mathbf{g})$ there is a unique quadratic polynomial $f_{c(\lambda)} : z \mapsto z^2 + c(\lambda), c(\lambda) \in M$, which is hybrid equivalent to g_{λ} . The map $\chi : \lambda \mapsto c(\lambda)$ is called the *straightening* of $M(\mathbf{g})$.

We know that the straightening is not canonically defined outsed the Mandelbrot set but rather depends on the choice of the tubing. But for equipped families

228

(42.1)

there is a natural choice given by (42.1). With this choice, the straightening χ admits an extension to the whole parameter domain Λ , which well be still denoted by χ .

Recall that D_r stands for the parameter subpotential disk of radius r (in the quadratic family). We can now formulate a fundamental result of the theory of quadratic-like families:

THEOREM 10.1. Let \mathbf{g} be a proper unfolded equipped quadratic-like family over Λ . Endow it with the tubing given by (42.1). Then the corresponding straightening χ is a homeomorphism from Λ onto D_{r^2} mapping $M(\mathbf{g})$ onto M.

The proof of this theorem will be split into several pieces each of which is important on its own right.

42.4. The critical value moves transversally to h. We say that a holomorphic curve $\Gamma \subset \mathbb{C}^2$ is a *global transversal* to a holomorphic motion h if it transversally intersects each leaf of h at a single point.

LEMMA 10.2. Under the assumptions of Theorem 10.1, the graph of the function $\lambda \mapsto g_{\lambda}(0), \lambda \in \Lambda$, is a global transversal to the holomorphic motion \mathbf{h} on $\mathbb{U}' \setminus \mathbb{U}$.

We will also express it by saying that the critical value moves transversally to **h**. The moral of this lemma is that in the complex setting the transversality can be achieved for purely topological reasons.

PROOF. Take a point $z \in A_{o} = \overline{U}'_{o} \smallsetminus U_{o}$ and consider its orbit

$$\psi_z: \lambda \mapsto h_\lambda(z)$$

under the motion **h**. By Assumption H of §42.1, for $z \in \overline{U}'_{o} \setminus \overline{U}_{o}$ the function ψ_{z} admits a holomorphic extension to a slightly bigger parameter domain $\Lambda_{z} \supseteq \Lambda$. For $z \in \partial U_{o}$, equivariance equation

$$f_{\lambda}(\psi_z(\lambda)) = \psi_{f_0 z}(\lambda)$$

implies that ψ_z admits an extension to the domain $\Lambda_{f_{\circ}z} \supseteq \Lambda$ (note that $f_{\circ}z \in \partial U'_{\circ}$) as a multiply valued holomorphic function. Such a function is continuous up to the boundary of Λ .

Thus, for any $z \in A_0$, the function ψ_z admits a continuous extension to $\overline{\Lambda}$. Moreover,

(42.3)
$$\psi_z(\lambda) \in U'_{\lambda}$$
 for any $z \in U'_{0} \setminus U_{0}$ and $\lambda \in \overline{\Lambda}$

For $z \in U_{o} \setminus \overline{U}_{o}$, it follows immediately from Assumption H. To see it for $z \in \partial U_{o}$, let us take any intermedite Jordan disk, $U_{o} \Subset W_{o} \Subset U'_{o}$, and let W_{λ} be the Jordan disk bounded by $h_{\lambda}(\partial W_{o}), \lambda \in \overline{\Lambda}$. Then we have:

• $W_{\lambda} \Subset U_{\lambda}$ for any $\lambda \in \overline{\Lambda}$ (by Assumption H);

• $\psi_z(\lambda) \in V_\lambda$ for any $\lambda \in \Lambda$, and by continuity, $\psi_z(\lambda) \in \overline{V}_\lambda$ for $\lambda \in \partial \Lambda$,

and (42.3) follows.

In what follows we fix $z \in U'_{o} \setminus U_{o}$ and let $\psi \equiv \psi_{z}$. Since the tube $\mathbb{V} \equiv \mathbb{U}' | \partial \Lambda$ is homeomorphic to the solid torus $\partial \Lambda \times \mathbb{D}$ over $\partial \Lambda$, the curve $\lambda \mapsto \psi(\lambda), \lambda \in \partial \Lambda$, is homotopic to the zero curve $\lambda \mapsto 0$ in \mathbb{V} , i.e., these two curves can be joined by a continuous family of curves $\psi^{t} : \partial \Lambda \to \mathbb{V}, 0 \leq t \leq 1$.

Consider now the curve $\phi : \lambda \mapsto g_{\lambda}(0), \lambda \in \partial \Lambda$. Since **g** is proper, $\phi(\lambda) \in \partial V_{\lambda}$. Hence $\phi(\lambda) - \psi^{t}(\lambda) \neq 0$ for any $t \in [0, 1], \lambda \in \partial \Lambda$. It follows that the curves $\lambda \mapsto \phi(\lambda) - \psi(\lambda)$ and $\lambda \mapsto \phi(\lambda)$, $\lambda \in \partial \Lambda$, have the same winding number around 0. But the latter number is equal to 1, since **g** is unfolded. Hence the former number is equal to 1 as well. By the classical Argument Principle, the graphs of the functions ϕ and ψ have a single transverse intersection, as asserted.

42.5. External uniformization. In this section we will construct a dynamical (locally quasiconformal) uniformization of $\Lambda \setminus M(\mathbf{g})$ which generalizes the uniformization of $\mathbb{C} \setminus M$ constructed in §29.1. This construction will illustrate how to relate the parameter and dynamical planes by means of holomorphic moions ("phase-parameter relation").

Let us consider a set $P^0 = \{\lambda \in \Lambda : g_{\lambda}(0) \in U'_{\lambda} \setminus U_{\lambda}\}$ (i.e., the set of parameters for which the critical point escapes under the first iterate through the "half-closed" fundamental annulus $A^0_{\lambda} := U'_{\lambda} \setminus U_{\lambda}$). Since **g** is proper, all points in Λ sufficiently close to $\partial \Lambda$ belong to P^0 . We will show that P is an annulus naturally homeomorphic to the base fundamental annulus A^0_0 .

To this end consider the graph of the function $\phi : \lambda \mapsto g_{\lambda}(0)$,

$$\Gamma = \{ (\lambda, z) \in \mathbb{C}^2 : \lambda \in \Lambda, z = g_{\lambda}(0) \}.$$

By Lemma 10.2, this graph is a global transversal to the holomorphic motion \mathbf{h} of A_o^0 . Hence there is a well defined holonomy $\gamma^0 : A_o^0 \to \Gamma$ along the leaves of \mathbf{g} , and it maps A_o^0 homeomorphically onto a topological annulus $B^0 \subset \Gamma$. Obviously, $\pi(B^0) = P^0$. Altogether, we have a homeomorphism $\pi \circ \gamma^0$ from A_o^0 onto P^0 . It follows, in particular that P^0 is a topological annulus, whose inner boundary is a Jordan curve $\pi \circ \gamma^0(\partial U_o)$ in Λ and the outer boundary is $\partial \Lambda$.

Let us consider the domain $\Lambda^1 = \Lambda \smallsetminus P^0$. The restriction of our quadratic-like family to this parameter domain is not proper any more. To restore this property, we have to restrict the dynamical domains as well. Let $U_{\lambda}^1 = g_{\lambda}^{-1}U_{\lambda}$. For any $\lambda \in \Lambda^1$, $g_{\lambda}(0) \in U_{\lambda}$; hence U_{λ}^1 is a topological disk and $g_{\lambda} : U_{\lambda}^1 \to U_{\lambda}$ is a quadratic-like map. This gives us a quadratic-like family over Λ^1 .

It is proper since by construction $g_{\lambda}(0) \in U_{\lambda}$ for $\lambda \in \partial \Lambda^{1}$. It has winding number one since the function $\phi : \lambda \mapsto g_{\lambda}(0)$ does not have zeros in the annulus \bar{P}^{0} . It follows that the boundary curves $\phi : \partial \Lambda \to \mathbb{C}^{*}$ and $\phi : \partial \Lambda' \to \mathbb{C}^{*}$ are homotopic (after parametrizing $\partial \Lambda$ and $\partial \Lambda^{1}$ by the standard circle) and hence they have the same winding number around 0.

Let us now equip this family with a holomorphic motion $h_{\lambda}^{1}: A_{o}^{1} \to A_{\lambda}^{1}$ of the fundamental annulus $A_{\lambda}^{1}:=U_{\lambda} \setminus U_{\lambda}^{1}$. This motion is obtained by lifting the motion h_{λ} by means of the double coverings $g_{\lambda}: A_{\lambda}^{1} \to A_{\lambda}^{0}$ (see Lemma 6.28):

$$\begin{array}{cccc} A^1_{\mathbf{o}} & \xrightarrow{} & A^1_{\lambda} \\ g_{\mathbf{o}} \downarrow & & \downarrow g_{\lambda} \\ A_{\mathbf{o}} & \xrightarrow{} & h_{\lambda} \end{array}$$

By the First λ -lemma, the original holomorphic motion **h** matches with **h'** on the common boundary $\partial^i A^0_{\lambda} = \partial^o A^1_{\lambda}$, so that together they provide a single holomorphic motion of the union $A^0_{\lambda} \cup A^1_{\lambda}$ over Λ^1 . Let $P^1 = \{\lambda \in \Lambda^1 : g_{\lambda}(0) \in A^1_{\lambda}\}$. Applying the above result to the re-

Let $P^1 = \{\lambda \in \Lambda^1 : g_{\lambda}(0) \in A^1_{\lambda}\}$. Applying the above result to the restricted quadratic-like family, we obtain a homeomorphism $\pi \circ \gamma^1 : A^1_{o} \to P^1$, where $\gamma^1 : A^1_{o} \to \Gamma$ is the holonomy along \mathbf{h}^1 . Since γ^1 matches with γ^0 on the common boundary of the annuli, they give us a homeomorphism of the union of the dynamical annuli onto the union of parameter annuli, $A^0 \cup A^1 \to P^0 \cup P^1$.

Proceeding in the same way, we construct: • A nest of parameter annuli P^n attached one to the next and the corresponding parameter domains $\Lambda^n = \Lambda^{n-1} \smallsetminus \cup P^{n-1}$ (where $\Lambda^0 \equiv \Lambda$). Moreover, $\cup P^n = \Lambda \smallsetminus M(\mathbf{g})$.

• A sequence of proper unfolded quadratic-like families

$$g_{n,\lambda} \equiv g_{\lambda} : U_{\lambda}^n \to U_{\lambda}^{n-1} \text{ over } \Lambda^n,$$

where $U_{\lambda}^{n} = g_{\lambda}^{-n} U_{\lambda}$ (thus $U_{\lambda}^{0} \equiv U_{\lambda}, \ U_{\lambda}^{-1} \equiv U_{\lambda}'$).

• A sequence of holomorphic motions $h_{n,\lambda}$ of the fundamental annulus $A_{\lambda}^{n} := U_{\lambda}^{n-1} \smallsetminus U_{\lambda}^{n}$ over Λ^{n} that equip $g_{n,\lambda}$; moreover $h_{n+1,\lambda}$ is obtained by lifting $h_{n,\lambda}$ by means of the coverings $g_{\lambda} : A_{\lambda}^{n} \to A_{\lambda}^{n-1}$. These holomorphic motions match on the common boundaries of the fundamental annuli.

Let $\gamma_n : A^n_{o} \to \Gamma$ be the holonomy along \mathbf{h}_n . Since the holomorphic motions match on the common boundaries, these holonomies also match, and determine a continuous injection

$$\gamma: U'_{o} \smallsetminus K(f_{o}) \to \Gamma$$

Composing it with the projection π , we obtain a homeomorphism

(42.4)
$$\pi \circ \gamma : U'_{o} \smallsetminus K(f_{o}) \to \Lambda \smallsetminus M(\mathbf{g})$$

between the dynamical and parameter annuli. Note that the inverse map is equal to $\gamma^{-1} \circ \phi$. This is the *phase-parameter relation* we alluded earlier.

Composing the above homeomorphism with the tubing (42.1), we obtain a "uniformization" $\Phi_{M(\mathbf{g})} \equiv \Phi_{\mathbf{g}}$ of $\Lambda \smallsetminus M(\mathbf{g})$ by a round annulus:

(42.5)
$$\Phi_{\mathbf{g}} = B_{\circ} \circ (\pi \circ \gamma)^{-1} = B_{\lambda} \circ \phi : \Lambda \smallsetminus M(\mathbf{g}) \to \mathbb{A}(1, r^2), \quad \Phi_{\mathbf{g}}(\lambda) = B_{\lambda}(g_{\lambda}(0)).$$

We see that this uniformization is given by the tubing position of the critical value of g_{λ} (see §49.2).

COROLLARY 10.3. The Mandelbrot set $M(\mathbf{g})$ is connected and full.

The above uniformization of $\Lambda \setminus M(\mathbf{g})$ is generally not conformal. However, in the case of a restricted quadratic family (see §42.2), it is a restriction of the Riemann map $\Phi_M : \mathbb{C} \setminus M \to \mathbb{C} \setminus \overline{\mathbb{D}}$. Indeed, in this case, the tubing B_{λ} turns into the Böttcher maps B_c (see (42.2)), the critical value $g_{\lambda}(0)$ turns into c, and formula (42.5) turns into formula (29.1) for the Riemann map Φ_M .

42.6. External straightening. We are now ready to prove that the straightening is a homeomorphism outside the Mandelbrot sets.

LEMMA 10.4. Under the assumptions of Theorem 10.1, the straightening

$$\chi:\Lambda\smallsetminus M(\mathbf{g})\to D_{r^2}\smallsetminus M$$

is a homeomorphism.

PROOF. Let us consider the uniformizations $\Phi_{\mathbf{g}} : \Lambda \smallsetminus M(\mathbf{g}) \to \mathbb{A}(1, r^2)$ and $\Phi_M : D_{r^2} \smallsetminus M \to \mathbb{A}(1, r^2)$ constructed above. Then

(42.6)
$$\chi = \Phi_M^{-1} \circ \Phi_{\mathbf{g}}$$

Indeed, let $\lambda \in \Lambda \setminus M(\mathbf{g})$ and $c = \chi(\lambda) \in D_{r^2} \setminus M$. Putting together (29.1), (42.5) and (49.4), we obtain:

$$\Phi_{\mathbf{g}}(\lambda) = B_{\lambda}(g_{\lambda}(0)) = B_{c}(c) = \Phi_{M}(c),$$

which is exactly (42.6). Since $\Phi_{\mathbf{g}}$ and Φ_M are both homeomorphisms, χ is a homeomorphism as well.

42.7. Quasiconformality. Given a holomorphic motion \mathbf{h} over Λ , let

$$\operatorname{Dil}(\mathbf{h}) = \sup_{\lambda \in \Lambda} \operatorname{Dil}(h_{\lambda})$$

(which can be infinite). We say that the holomorphic motion \mathbf{h} is K-qc if

 $\operatorname{Dil}(\mathbf{h}) \leq K.$

LEMMA 10.5. Under the assumptions of Theorem 10.1, assume that the tubing $B_{\circ}: A_{\circ} \to \mathbb{A}[r, r^2]$ and the holomorphic motion **h** are K-qc. Then the uniformization $\Phi_{\mathbf{g}}: \Lambda \setminus M(\mathbf{g}) \to \mathbb{A}(1, r^2)$ (42.5) is K-qc as well.

In fact, we can make the dilatation depend only on mod A_{\circ} and mod $(\Lambda \setminus \Lambda')$, after an appropriate adjustment of the family **g**:

LEMMA 10.6. Let us consider a quadratic-like family \mathbf{g} over Λ satisfying the assumptions of Theorem 10.1. This family can be adjusted to a family $\tilde{\mathbf{g}}$ over $\tilde{\Lambda}$ in such a way that the dilatation of the straightening $\tilde{\chi} : \tilde{\Lambda} \setminus M(\tilde{\mathbf{g}}) \to D \setminus M$ will depend only on mod A_{\circ} and mod $(\Lambda \setminus \Lambda')$.

42.8. Miracle of continuity. We will now show that the straightening is continuous on the boundary of $M(\mathbf{g})$:

LEMMA 10.7. Under the assumptions of Theorem 10.1, the straightening χ is continuous at any point $\lambda \in \partial M(\mathbf{g})$ and moreover $\chi(\lambda) \in \partial M$.

PROOF. First we will show that $\chi | \partial M(\mathbf{g})$ is a continuous extension of $\chi | \Lambda \setminus M(\mathbf{g})$. Let $\lambda_n \in \Lambda \setminus M(\mathbf{g})$ be a sequence of parameter values converging to some $\lambda \in \partial M$. Let $c_n = \chi(\lambda_n)$ and $c = \chi(\lambda) \in M$. We should show that $c_n \to c$. Let $g_{\lambda} : U \to U', f_c : \Omega \to \Omega'$.

By Lemma 10.4, the map $\chi : \Lambda \smallsetminus \operatorname{int} M(\mathbf{g}) \to D \smallsetminus \operatorname{int} M$ is proper, and hence any limit point d of $\{c_n\} \subset D \smallsetminus M$ belongs to ∂M . We assert that $g_{\lambda} : U \to U'$ is qc conjugate to $f_d : V \to V'$. Indeed, the $g_{\lambda_n} : U_n \to U'_n$ are hybrid equivalent to the $f_{c_n} : \Omega_n \to \Omega'_n$ by means of some qc maps $\psi_n : U'_n \to \Omega'_n$. By the straightening construction (see the proof of Lemma 9.11), the dilatation of ψ_n is equal to the dilatation of the tubing $B_{\lambda_n} = B_0 \circ h_{\lambda}^{-1}$, which is locally bounded by the λ -lemma. By Theorem 2.31, the sequence ψ_n is precompact in the topology of uniform convergence on compact subsets of U'. Take a limit map $\psi : U' \to \Omega'$. Since $g_{\lambda_n} \to g_{\lambda}$ uniformly on compact subsets of U and $f_{c_n} \to f_d$ (along a subsequence) uniformly on compact subsets of Ω , the map ψ conjugates g_{λ} to f_d , as was asserted.

But g_{λ} is also hybrid equivalent to f_c . Thus f_c and f_d are qc conjugate in some neighborhoods of their filled Julia sets. By Corollary 9.19, they are qc conjugate on the whole complex plane. Since $d \in \partial M$, Proposition 6.41 implies the desired: c = d (and, in particular, $c \in \partial M$). The above argument implies that χ continuously maps $\Lambda \setminus \operatorname{int} M(\mathbf{g})$ into $D \setminus \operatorname{int} M$. We still need to show that χ is continuous at any point $\lambda \in \partial M(\mathbf{g})$ even if it is approached from the interior of $M(\mathbf{g})$. The argument is similar to the above except one detail. So, let now $\{\lambda_n\}$ be any sequence in Λ converging to λ . Let c_n , c and d be as above. Then the above argument shows that f_c is qc equivalent to f_d . But now we already know that $c \in \partial M$ (though this time we do not know this for d). Applying Proposition 6.41 once again, we conclude that c = d.

"Only by miracle can one ensure the continuity of straightening in degree 2", said Adrien Douady As we have seen, a reason behind this miracle is quasiconformal rigidity of the quadratic maps f_c with $c \in \partial M$ (Proposition 6.41). Another reason is the λ -lemma. All these reasons are valid only for one-parameter families. There are no miracles in the polynomial families with more parameters.

42.9. Hyperbolic components. As in the case of the genuine Mandelbrote set, a component H of int $M(\mathbf{g})$ is called hyperbolic if it contains a hyperbolic parameter value.

- EXERCISE 10.8. Show that:
- (i) All parameter values in a hyperbolic component of $M(\mathbf{g})$ are hyperbolic;
- (ii) Neutral parameter values belong to $\partial M(\mathbf{g})$.

LEMMA 10.9. If P is a hyperbolic component of int $M(\mathbf{g})$ then there exists a hyperbolic component Q of int M such that $\chi : P \to Q$ is a proper holomorphic map.

PROOF. Obviously the straightening of a hyperbolic map is hyperbolic. Hence $\chi(P)$ belongs to some hyperbolic component Q of int M. Moreover, since the hybrid conjugacy is conformal on the interior of the filled Julia set, it preserves the multiplies of attracting cycles. Hence

$$\mu_P(\lambda) = \mu_Q(c)$$
 for $c = \chi(\lambda)$,

where μ_P and μ_Q are the multiplier functions on the domains P and Q respectively. By the Implicit Function Theorem, both these functions are holomorphic. Moreover, by the Multiplier Theorem, μ_Q is a conformal isomorphism onto \mathbb{D} . Hence $\chi = \mu_Q^{-1} \circ \mu_P$ is holomorphic as well.

By Lemma 10.7, the map $\chi: P \to Q$ is continuous up to the boundary and $\chi(\partial P) \subset \partial Q$. Hence it is proper.

42.10. Queer components. As in the quadratic case, a non-hyperbolic component Q of int $M(\mathbf{g})$ is called queer. In this section we will prove, using the dynamical uniformization of queer components (§31.6.4), that the straightening χ is holomorphic on Q. Let us begin with an extention of Corollary 6.25 to quadratic-like families:

LEMMA 10.10. Let Q be a queer component of $M(\mathbf{g})$. Take a base point $o \in Q$. Then there is a holomorphic motion $H_{\lambda}: U'_{o} \to U'_{\lambda}$ conjugating g_{o} to g_{λ} .

PROOF. Since $M(\mathbf{g})$ is equipped, there is an equivariant holomorphic motion $h_{\lambda} : A_{\circ} \to A_{\lambda}$. Let $A_{\lambda}^{n} = g_{\lambda}^{-n}A_{\lambda}$. Since the critical point is non-escaping under the iterates of g_{λ} , the A_{λ}^{n} are annuli and the maps $g_{\lambda}^{n} : A_{\lambda}^{n} \to A_{\lambda}$ are double coverings. By Lemma 6.28, **h** can be consequtively lifted to holomorphic motions $h_{n,\lambda} : A_{*}^{n} \to A_{\lambda}^{n}$. By the First λ -lemma, they automatically match on the common

boundaries of the annuli, so that we obtain an equivariant holomorphic motion $H_{\lambda}: U'_{o} \smallsetminus K(g_{o}) \to U'_{\lambda} \smallsetminus K(g_{\lambda})$. Since the sets $K(g_{\lambda})$ are nowhere dense (see Corollary 9.3), the First λ -lemma implies that the H_{λ} extends to an equivariant holomorphic motion $U'_{o} \to U'_{\lambda}$.

EXERCISE 10.11. Let H_{λ} be the holomorphic motion constructed in the previous lemma. Then the Beltrami differential

(42.7)
$$\mu_{\lambda}(z) = \begin{cases} \frac{\partial H_{\lambda}(z)}{\partial H_{\lambda}(z)}, & z \in K(g_{o}), \\ 0, & z \in \mathbb{C} \smallsetminus K(g_{o}) \end{cases}$$

holomorphically depends on $\lambda \in Q$.

We can now prove an analogue of Lemma 10.9 for queer components:

LEMMA 10.12. The straightening χ is holomorphic on any queer component Q of int $M(\mathbf{g})$.

PROOF. Select a base point $o \in Q$, and let $\phi : U' \to \Omega'$ denote the hybrid conjugacy between $g_o : U \to U'$ and its straightening $f_o \equiv f_{c_o} : \Omega \to \Omega'$. Let μ_{λ} be the holomorphic family of conformal structures on $K(g_o)$ considered in the previous Excercise. Push it forward to the f_o -plane: let ν_{λ} be the f_o -invariant Beltrami differential equal to $\phi_*(\mu_{\lambda})$ on $K(f_o)$ and vanishing on $\mathbb{C} \setminus K(f_o)$. Since ϕ is confomal a.e. on the Julia set,

$$\nu_{\lambda} = \left(\mu_{\lambda} \, \phi' / \, \overline{\phi'}\right) \circ \phi^{-1},$$

which is obviously holomorphic in $\lambda \in Q$. Let $h_{\lambda} : (\mathbb{C}, 0) \to (\mathbb{C}, 0)$ be the solution of the Beltrami equation for μ_{λ} tangent to the identity at ∞ . Then the map $f_{\lambda} := h_{\lambda} \circ f_{\circ} \circ h_{\lambda}^{-1}$ is a quadratic polynomial $z \mapsto z^2 + c(\lambda)$, and by Corollary 4.65, $c(\lambda)$ is holomorphic in λ . Finally, f_{λ} is the straightening of g_{λ} by means of the hybrid conjugacy $h_{\lambda} \circ \phi \circ H_{\lambda}^{-1}$.

42.11. Discreteness of the fibers.

LEMMA 10.13. For any $c \in M$, the fiber $\chi^{-1}(c)$ is finite.

PROOF. Since $M(\mathbf{g})$ is compact, it is enough to show that the fibers are discrete. Assume that there exists some $c \in M$ with an infinite fiber $\chi^{-1}(c)$. Then this fiber contains a sequence of distinct parameter values $\lambda_n \in \chi^{-1}(c)$ converging to some point $\lambda_{\infty} \in \chi^{-1}(c)$. Let $g \equiv g_{\infty} : U \to U'$.

Without loss of generality, we can assume that $\lambda_{\infty} \in \partial M$. [Otherwise, consider the component U of int M containing λ_{∞} . Since χ is holomorphic on U and continuous on \overline{U} , we conclude that $\chi | U \equiv \text{const.}$ But then we can replace λ_{∞} by any boundary point of U.]

Let us select λ_{∞} as the base point in Q. Since the quadratic-like family $g_{\lambda} : U_{\lambda} \to U'_{\lambda}$ is equipped, there exists an equivariant holomorphic motion $h_{\lambda} : A \to A_{\lambda}$ of the closed fundamental annulus $A_{\lambda} = \overline{U}'_{\lambda} \setminus U_{\lambda}$ over Λ (where $A \equiv \overline{U}' \setminus U$). Extend it by the Elementary λ -lemma to a holomorphic motion $h_{\lambda} : \mathbb{C} \setminus U \to \mathbb{C} \setminus U_{\lambda}$ over a neighborhood $Q' \subset Q$ of λ_{∞} (keeping the same notation for the extension). We will now construct a holomorphic family of hybrid deformations G_{λ} of g over Q' naturally generated by this holomorphic motion.

To this end let us first pull back the standard conformal structure to $\mathbb{C} \smallsetminus U$, $\mu_{\lambda} = h_{\lambda}^*(\sigma)$. Then extend μ_{λ} to a *g*-invariant conformal structure on $\mathbb{C} \smallsetminus K(g)$ by pulling it back by iterates of g. Finally, extend it to K(g) as the strandard structure. This gives us a holomorphic family of g-invariant conformal structures on \mathbb{C} . We will keep the same notation μ_{λ} for these structures. Solving the Beltrami equations, we obtain a holomorphic family of qc maps $H_{\lambda} : \mathbb{C} \to \mathbb{C}$ such that $\mu_{\lambda} = (H_{\lambda})^*(\sigma)$, and in particular, $\bar{\partial}H(z) = 0$ a.e. on K(g). Conjugating g by these maps, we obtain a desired hybrid deformation $G_{\lambda} = H_{\lambda} \circ g \circ H_{\lambda}^{-1}$, $\lambda \in Q'$.

On the other hand, for maps $g_n := g_{\lambda_n}$, we can construct the Beltrami differentials $\mu_n \equiv \mu_{\lambda_n}$ in a different way. Namely, since the map g_n is hybrid equivalent to g, the equivariant map $h_n := h_{\lambda_n}$ uniquely extends to a hybrid conjugacy (Theorem 9.18). Let us keep the same notation h_n for this conjugacy.

The above two constructions naturally agree: $(h_n)^*\sigma = \mu_n$. Indeed, it is true on $\mathbb{C} \setminus U$ by definition. It is then true on $U \setminus K(f)$, since both Beltrami differentials are *g*-invariant. Finally, it is true on the filled Julia set K(g) since h_n is conformal a.e. on it.

Thus the qc maps $H_n : \mathbb{C} \to \mathbb{C}$ and $h_n : \mathbb{C} \to \mathbb{C}$ satisfy the same Beltrami equation. They also coincide at two points, e.g., at the critical point and at the β -fixed point of g (in fact, by Corollary 9.16 they coincide on the whole Julia set of g). By uniqueness of the solution of the Beltrami equation, $H_n = h_n$. Hence $G_n = g_n$. Returning to the original notations, we have

(42.8)
$$G_{\lambda_n}(z) = g_{\lambda_n}(z)$$

Take an $\epsilon > 0$ such that both functions $G_{\lambda}(z)$ and $g_{\lambda}(z)$ are well-defined in the bidisk $\{(\lambda, z) \in \mathbb{C}^2 : |\lambda - \lambda_{\star}| < \epsilon, z \in V \equiv g^{-1}U\}$. For any $z \in V$, consider two holomorphic functions of λ :

$$\Phi_z(\lambda) = G_\lambda(z) \quad \text{and} \quad \phi_z(\lambda) = g_\lambda(z), \quad |\lambda - \lambda_\infty| < \epsilon.$$

By (42.8), they are equal at points λ_n converging to λ_∞ . Hence they are identically equal.

Thus for $|\lambda| < \epsilon$, two quadratic-like maps, G_{λ} and g_{λ} , coincide on V. But it is impossible since the Julia set of G_{λ} is always connected, while the Julia set of g_{λ} is disconnected for some λ arbitrary close to λ_{∞} (recall that we assume that $\lambda_{\infty} \in \partial M(\mathbf{g})$).

COROLLARY 10.14. $\chi(\operatorname{int} M(\mathbf{g})) \subset \operatorname{int} M$.

Remark. Of course, it is not obvious only for queer components.

PROOF. Take a component P of int M. We have proven that $\chi|P$ is a nonconstant holomorphic function. Hence the image $\chi(P)$ is open. Since it is obviously contained in M, it must be contained in int M.

42.12. Bijectivity. What is left to show is that the map $\chi : M(\mathbf{g}) \to M$ is bijective. By §42.5, the winding number of the curve $\chi : \partial \Lambda \to \mathbb{C}$ around any point $c \in D_{r^2}$ is equial to 1. By the Topological Argument Principle (§2.3),

(42.9)
$$\sum_{a \in \chi^{-1}c} \operatorname{ind}_a(\chi) = w_c(\chi, \partial \Lambda) = 1, \quad c \in D_{r^2}.$$

It immediately follows that the map $\chi : \Lambda \to \mathbb{D}_{r^2}$ is *surjective* (for otherwise the sum in the left-hand side would vanish fo some $c \in D_{r^2}$).

Let us show that χ is injective on the interior of $M(\mathbf{g})$. Indeed, if $a_0 \in \operatorname{int} M$, then by Corollary 10.14 $c = \chi(a_0) \in \operatorname{int} M$, and by Lemma 10.7 $\chi^{-1}(c) \subset \operatorname{int} M$.

But since $\chi | \text{int } M$ is holomorphic (see §42.9 and 42.10), we have $\text{ind}_a(\chi) > 0$ for any $a \in \text{int } M$. It follows that the sum in the left-hand side of (42.9) actually contains only one term, so that c has only one preimage, a_0 .

Finally, assume that there is a point $c \in \partial M$ with more than one preimage. By the Topological Argument Principle, χ has a non-zero index at one of those preimages, say, a_1 . Take another preimage a_2 . Both a_1 and a_2 belong to ∂M .

Take a point $a'_2 \notin \partial M(\mathbf{g})$ near a_2 , and let $c' = \chi(a'_2)$. By Exercise 1.52, χ is locally surjective near a_1 , so that c' has a preimage a'_1 over there. This contradicts injectivity of χ on $\Lambda \setminus \partial M(\mathbf{g})$.

This completes the proof of Theorem 10.1.

43. QL families over the complex renormalization windows

Let us go back to the quadratic family $f_c: z \mapsto z^2 + c$. Take some superattracting parameter c_o of period p > 1. It is the center of the renormalization window $\Lambda = \Lambda_{c_0}$ described in §45.2. Recall that $\partial \Lambda$ intersects the Mandelbrot set M in two points called the *root* and the *tip*.

For any polynomial $f_c, c \in \Lambda$, we have constructed a quadratic-like return map $g_c = f_c^p : V_c \to V'_c$ around the critical point. If the Julia set $J(g_c)$ of this map is connected then f_c is renormalizable with combinatorics c_o . Let

 $M_{\circ} = \{c \in \Lambda : f_c \text{ is renormalizable with combinatorics } c_{\circ}\} \cup \{root, tip\}.$

THEOREM 10.15. The set M_{\circ} is canonically homeomorphic to the Mandelbrot set M.

Theorem 10.1 is designed to imply this result. However, it does not do it (at, least, not immediately) since the quadratic-like family g_c over Λ is not full: it misses the root and the tip of M_{\circ} . This problem can be fixed for the tip. In case of primitive renormalization, it can also be fixed for the root. However, in the satellite case, it is not fixable: in fact, in this case the root of M_{\circ} is not renormalizable with period p.

In this section we will give a proof of Theorem 10.15, which will produce for us all little M-copies. For primitive copies, an alternative proof will be given in §?? where we will construct the corresponding full quadratic-like families to which Theorem 10.1 can be applied.

44. Notes

The notion of a polynomial-like map was introduced by Douady and Hubbard in their fundamental work [**DH3**]. Basic theory of these maps, including the Straightening Theorem, was developed in the same paper. This theorem was the first application of the method of *quasiconformal surgery*: you cook by hands some topological object that looks like a polynomia, and then you realize it as an actual polynomial by means of the MRMT.

Miracle of continuity: **[D2**]

No miracle in higher degrees: ??, §...

CHAPTER 11

Yoccoz Puzzle

46. Combinatorics of the puzzle

Kids know well the "puzzle game" of cutting a picture into small pieces and then trying to put them back together. Such a game can be played with dynamical pictures like Julia sets and the Mandelbrot set as well. It turned out to be a very efficient way to describe the combinatorics of the corresponding dynamical systems and to control their geometry.

Our standing assumption will be that both fixed points of f are repelling.

46.1. Description of the puzzle. Let us fix some parameter wake $\mathcal{W}_{\mathbf{p}/\mathbf{q}}$ (see Theorem 7.29), and let $c \in \mathcal{W}_{\mathbf{p}/\mathbf{q}}$, $f = f_c$. The puzzle game starts by cutting the complex plane with the α -rays \mathcal{R}_i , $i = 1, \ldots, \mathbf{q}$, landing at the α -fixed point (see §39.5.2). They are cyclically permuted by the dynamics, and divide the plane into \mathbf{q} sectors S_i as described in Lemma 7.16. (Recall that $S_0 \ni 0$ is the critical sector, and $S^1 \ni f(0)$ is the characteristic one.)

Let us select some equipotential $\mathcal{E} = \mathcal{E}^t$ of height t > 0 surrounding the critical value f(0). Let $U^0 \ni f(0)$ be the (open) Jordan disk bounded by \mathcal{E} . Its closure \overline{U} is tiled by \mathbf{q} (closed) "triangles"

$$Y_i^{(0)} = U^0 \cap S_i$$

called *puzzle pieces of depth* 0 (see Figure ??). The puzzle piece $Y_0^{(0)} \equiv Y_0^{(0)} \ni 0$ is naturally called critical, while $Y_1^{(0)} \ni c$ is called characteristic. We denote this initial puzzle $\mathcal{Y}^{(0)}$.

Consider now the preimage $\mathcal{Y}^{(2)}$ of $\mathcal{Y}^{(0)}$ under f. Let $U^1 = f^{-1}(U^0)$ be the disk bounded by the equipotential $\mathcal{E}^{(1)} = \mathcal{E}^{t/2}$. Cut it by $2\mathbf{p}$ external rays landing at the points α and $\alpha' = -\alpha$. We obtain a tiling of U^1 by $2\mathbf{p} - 1$ closed topological disks $Y_i^{(1)}$ called *puzzle pieces of depth 1* ($2\mathbf{p} - 2$ lateral triangles and one central 6-gone). We label them in such a way that $Y_i^{(1)} \subset Y_i^{(0)}$, $i = 0, 1, \ldots, \mathbf{q}$, and we let

$$Z_i := Y_{-i}^{(0)} = -Y_i^{(1)}, \quad i = 1, \dots, \mathbf{q}.$$

Again, $Y^{(1)} \equiv Y_0^{(1)} \ni 0$ is called critical, while $Y_1^{(0)} \ni f(0)$ is called characteristic. Lemmas 7.16 and 7.17 imply that

(46.1)
$$f(Y^{(1)}) = Y_1^{(0)}; \quad f(Y_i^{(1)}) = f(Z_i) = Y_{j+1}^{(0)} \quad (i = 1, \dots, \mathbf{q} - 1)$$

(where $Y_q^{(0)}$ is understood as $Y^{(0)}$). Moreover, the map $f: Y^{(1)} \to Y_1^{(0)}$ is a double branched covering, while all other maps, $f: Y_i^{(1)} \to Y_{i+1}^{(0)}$ and $f: Z_i \to Y_{i+1}^{(0)}$ $(i = 1, ..., \mathbf{q} - 1)$ are univalent.

If $f(0) \in U^1$, we can take the preimage of $\mathcal{Y}^{(1)}$ to obtain puzzle $\mathcal{Y}^{(2)}$ of depth 2, etc. In general, if $f^n(0) \in U^0$ then we define puzzle $\mathcal{Y}^{(n)}$ as the *n*-fold preimage

11. YOCCOZ PUZZLE

of $\mathcal{Y}^{(0)}$. It is a tiling of the disk $U^n = f^{-n}(U^0)$ bounded by the equipotential $E^{(n)} = E^{1/2^n}$ obtained by cutting U^1 by the external rays comprising $f^{-n}(\cup \mathcal{R}_i)$ (i.e., the external rays landing at the points of $f^{-n}\alpha$). The tiles $Y_i^{(n)}$ of $\mathcal{Y}^{(n)}$ are called *puzzle pieces of depth n*. If $f^n(0) \neq \alpha$ then among these puzzle pieces there is one, $Y^{(n)} \equiv Y_0^{(n)}$, containing the critical point 0. It is called *critical*, while the puzzle piece $Y_1^{(n)}$ containing the critical value f(0) is called characteristic.

The following lemma summarizes obvious but crucial properties of the puzzle pieces (that can be viewed as axioms of the puzzle):

LEMMA 11.1. (i) Puzzle pieces are closed Jordan disks with piecewise analytic boundary ("polygons") that meet the Julia set at points of $f^{-n}\alpha$.

(ii) Under f, every puzzle piece $Y_i^{(n)}$ of depth n > 0 is mapped onto some puzzle piece $Y_j^{(n-1)}$ of depth n-1. This map is univalent if $Y_i^{(n)}$ is off-critical, and is a double covering if $Y_i^{(n)}$ is critical (i.e., if i = 0).

(iii) Any two puzzle pieces are either nested or have disjoint interiors.

(iv) Markov Property: If $f(Y_i^{(n)})$ intersects the interior of $Y_j^{(n)}$ then $f(Y_i^{(n)}) \supset Y_i^{(n)}$.

PROOF. (i) By definition, any $\operatorname{int} Y_i^{(n)}$ is a component of some $f^{-n}(\operatorname{int} Y_j^{(0)})$. But for a polynomial map, the full preimage of an open Jordan disk is a disjoint union of Jordan disks. Since each $Y_j^{(0)}$ is a piecewise analytic triangle, $Y_i^{(n)}$ is a piecewise analytic polygon.

(ii) Since $\operatorname{int} Y_i^{(n)}$ is a component of some $f^{-1}(\operatorname{int} Y_j^{(n-1)})$, the map $f: Y_i^{(n)} \to Y_j^{(n-1)}$ is a branched covering. Since both pieces are simply connected, the conclusion follows from the Riemann-Hurwitz formula.

(iii) Since $f^{-n}(\cup \mathcal{R}_i) \supset f^{-(n-1)}(\cup \mathcal{R}_i)$, the tiling $\mathcal{Y}^{(n)}$ is a refinement of $\mathcal{Y}^{(n-1)}|U^n$.

(iv) It is obvious for n = 0, so let n > 0. Then by property (ii), $f(Y_i^{(n)}) = Y_k^{(n-1)}$ for some k. By property (iii), $Y_k^{(n-1)}$ contains $Y_j^{(n)}$.

If the Julia set is connected, then all puzzles $\mathcal{Y}^{(n)}$ are well defined, forming finer and finer tilings of nested neighborhoods U^n of the filled Julia set K(f) that nicely behave under the dynamics. In the rest of the section, we will describe how these puzzles capture the recurrence of the critical orbit.

If we consider below $f^n(0)$ or puzzle $\mathcal{Y}^{(n)}$, we assume without mentioning that $f^n(0) \in U^0$, so that $\mathcal{Y}^{(n)}$ is well defined. (Not to be destracted by these details, we suggest the reader to assume in the first reading that the Julia set is connected, so the above assumptions hold automatically.)

46.2. Immediately renormalizable maps. By (46.1), $f^{\mathbf{q}}(0) \in Y^{(0)}$. So, if $f^{\mathbf{q}}(0) \in U^1$, it has two options: either $f^{\mathbf{q}}(0) \in Y^{(1)}$ (central return) or $f^{\mathbf{q}}(0) \in Z_{\boldsymbol{\kappa}}$ for some $\boldsymbol{\kappa} \in \{1, \ldots, \mathbf{q}\}$ (non-central return). In the former case, if $f^{2\mathbf{q}}(0) \in U^1$, it has the same options: either $f^{2\mathbf{q}}(0) \in Y^{(1)}$ or $f^{2\mathbf{q}}(0) \in Z_{\boldsymbol{\kappa}}$ for some $\boldsymbol{\kappa}$, etc. So, either the critical point always returns to $Y^{(1)}$,

(46.2)
$$f^{n\mathbf{q}}(0) \in Y^{(1)}, \quad n = 0, 1, \dots,$$

or else there exists the escaping moment $\mathbf{n} \in \mathbb{Z}_+$ such that

(46.3) $f^{n\mathbf{q}}(0) \in Z_{\kappa}$ for some $\kappa \in \{1, \dots, \mathbf{q}\}$

(provided $f^{\mathbf{nq}}(0) \in U^1$).

The map f is called *immediately renormalizable* if option (46.2) takes place.

PROPOSITION 11.2. If the map f is immediately renormalizable then it is satellite renormalizable with period \mathbf{q} in the sense of §49.5, and

(46.4)
$$\mathcal{K} = \{ z : f^{n\mathbf{q}} z \in Y^1, \ n = 0, 1, 2, \dots \},\$$

is the corresponding little Julia set. The point α is the non-dividing fixed point of $R_{\mathbf{q}}f$, and the little Julia sets $\mathcal{K}_i = f^i\mathcal{K}$, $i = 0, 1..., \mathbf{q} - 1$, form a bouquet centered at α . Moreover, \mathbf{q} is the smallest renormalization period of f.

PROOF. Let V be the puzzle piece Y^1 truncated by the equipotential $\mathcal{E}^{t/2^{\mathbf{q}}}$, and let $V' = Y^0$. Then $V \subset V'$ and the map $g = (f^{\mathbf{q}} : V \to V')$ is a double covering. However, the map g is not quadratic-like since ∂V touches $\partial V'$ along arcs of two external rays, \mathcal{R}^{θ} and \mathcal{R}^{γ} , landing at α (where $\theta \in (\frac{1}{4}, \frac{1}{2}), \gamma \in (\frac{1}{2}, \frac{3}{4})$). It is not a big problem, though, as a little *thickening* of V and V' turn g into

It is not a big problem, though, as a little *thickening* of V and V' turn g into a quadratic-like map (see Figure ??). Namely, let S be a little circle centered at α which is mapped by f onto a bigger circle. Let $\mathcal{R}^{\theta'}$ and $\mathcal{R}^{\gamma'}$ be two external rays close to \mathcal{R}^{θ} and \mathcal{R}^{γ} respectively that do not intersect U (thus $\theta' > \theta$, $\gamma' < \gamma$). Since the doubling map $\omega \mapsto 2\omega \mod 1$ is expanding on the circle \mathbb{R}/\mathbb{Z} , the external rays $\mathcal{R}^{\theta'}$ and $\mathcal{R}^{\gamma'}$ are "pushed away" from \mathcal{R}^{θ} and \mathcal{R}^{γ} (respectively) under the map $f^{\mathbf{q}}$.

Let $\tilde{\mathcal{R}}^{\theta'}$ and $\tilde{\mathcal{R}}^{\gamma'}$ stand for the shortest arcs of these rays connecting the equipotential $\mathcal{E}^{t/2^{\mathbf{q}}}$ to the circle S, and let Γ be the path composed of these two external arcs and the arc of S connecting them that does not intersect U. Then $\Gamma' = f^{\mathbf{q}}(\Gamma)$ is a path with endpoints on \mathcal{E}^t that does not cross Γ and lies "farther" from α than Γ .

Let $-\Gamma$ be the 0-symmetric path, and let $\tilde{U} \supset U$ be the Jordan disk bounded by Γ , $-\Gamma$ and two arcs of the equipotential $\mathcal{E}^{t/2^{\mathbf{q}}}$ connecting them. It is mapped under $f^{\mathbf{q}}$ onto the Jordan disk $\tilde{U}' \supset U'$ bounded by Γ' and the appropriate arc of the equipotential \mathcal{E}^t . Moreover, $\tilde{U} \in \tilde{U}'$ and the map $f^{\mathbf{q}} : \tilde{U} \to \tilde{U}'$ is a double branched covering. This is the desired pre-renormalization of f (we will keep notation g for it).

Obviously, $\mathcal{K} \subset K(g)$. To see the inverse inclusion, notice that $g : \mathcal{K} \to \mathcal{K}$ is two-to-one map. Indeed, $f^{\mathbf{q}} : \overline{U} \to \overline{U}'$ is a double branched covering, and obviously \mathcal{K} is completely invariant under this map.

Hence, \mathcal{K} is completely invariant under the pre-renormalization g. It is also closed and full. By Corollary 4.33 (and the Straightening Theorem), \mathcal{K} must coincide with the whole filled Julia set K(Rf).

Since \mathcal{R}^{γ} is an *g*-invariant curve in the complement of K(g) landing at α , the combinatorial rotation number of *g* at α is 0 (see Exercise 7.15). Hence α is the non-dividing fixed point of *g*.

Of course, α is a common point of all little Julia sets \mathcal{K}_i . Since α is the only point where the limbs $K(f) \cap Y_i^0$ touch one another, it is the only point where the little Julia sets $\mathcal{K}_i \subset K(f) \cap Y_i^0$ do.

Let us now prove the last assertion. Assume $p \in (2, \mathbf{q})$ is a smaller renormalization period of f; let h be the corresponding pre-renormalization. Then $f^p(0) \in Y^0_p \cap K(h)$. It follows that $K(h) \ni \alpha$, for otherwise K(h) would not intersect the curve $\mathcal{R}^{\theta} \cup \mathcal{R}^{\gamma} \cup \{\alpha\}$ that separates 0 from $f^p(0)$ (contradincting connectivity of K(h)).

Thus, α is a separating fixed point of K(h), and $f(K(h)) \cap K(h) \neq \emptyset$ – contradicting the almost disjointness property of §49.5.

Let us consider two symmetric triangles

(46.5)
$$L = \bigcup_{i=1}^{\mathbf{q}-1} Y_i^{(1)}, \quad R = \bigcup_{i=1}^{\mathbf{q}-1} Z_i.$$

Notice that the above little Julia set \mathcal{K} of the satellite renormalization is obtained from K(f) by chopping off infinitely many triangles: the preimages of R under all iterates of the double covering $f^{\mathbf{q}}: V \to V'$ (where V and V' are defined in the beginning of the above proof).

We will show in Chapter ?? that the set of parameters $c \in \mathcal{W}_{\mathbf{p}/\mathbf{q}}$ for which f_c is immediately renormalizable assemble a little copy of the Mandelbrot set attached to the main cardioid (see Figure ??).

46.3. Principal nest. Consider a puzzle piece P of depth n and a point z such that $f^m z \in \text{int } P$ for some $n \ge 0$. The puzzle piece Q of depth n+m containing z is called the *pullback* of P along the orbit $\{f^k z\}_{k=0}^m$. Clearly, the map $f^m : Q \to P$ is a branched covering of degree 2^l , where l is the number of critical puzzle pieces among $f^k Q$, $k = 0, 1, \ldots, m-1$. In particular, if there are no critical puzzle pieces among them, then $f^m : Q \to P$ is univalent. This yields:

LEMMA 11.3. Let P be a critical puzzle piece and let Q be the pull-back of P along $\{f^k z\}_{k=0}^m$.

If $f^m z$ is the first landing of the orb z at int $P, m \ge 0$, then $f^m : Q \to P$ is univalent.

If $z \in \text{int } P$ and $f^m z$ is the first return of the orb z to int P, m > 0, then $f^m : Q \to P$ is univalent or a double covering depending on whether Q is offcritical or otherwise.

We are now ready to introduce the *principal nest* of critical puzzle pieces,

(46.6)
$$V^0 \supset V^1 \supset V^2 \supset \cdots \ni 0,$$

and associated double coverings $g_n: V^n \to V^{n-1}$.

Assume f is not immediately renormalizable, and let **n** be the first escaping moment (46.3). We let $V_0 \ni 0$ be the pullback of Z_{κ} along the orbit $\{f^{n\mathbf{q}}\}_{n=0}^{\mathbf{n}}$.

Let $V^0 = P_0^{(0)}$. Assume inductively that we have defined the nest up to V^{n-1} . If the orb(0) never returns to int V^{n-1} then the construction stops here. Otherwise consider the first return $f^{l_n}0$ of the critical point back to V^{n-1} . Let V^n be the pullback of V^{n-1} along this orbit and let $g_n = f^{l_n} : V^n \to V^{n-1}$. By Lemma 11.3, this map is a double covering. This completes the construction.

We call V^n the principal puzzle piece of *level* n (pay attention to the difference between the "level" and the "depth").

A map f is called *combinatorially recurrent* if the critical orbit visits all critical puzzle pieces. In this (and only this) case, the principal nest is infinite.

46.4. Central returns and primitive renormalization. There are two different combinatorial possibilities on every level which are important to distinguish. The return of the critical point to level n-1 (and the level itself) is called *central* if $g_n 0 \in V^n$ (see Figure ??). In this case, the critical orbit returns to level n-1 at the same time as to level n, so that $l_n = l_{n+1}$ and $g_{n+1} : V^{n+1} \to V^n$ is just the restriction of g_n to V^{n+1} . Central returns indicate the *fast* recurrence of the critical orbit.

If N consecutive levels, $m-1, m, \ldots, m+N-2$, are central then the nest

$$(46.7) V^{m-1} \supset V^m \supset \cdots \supset V^{m+N-1}$$

is called a *central cascade* of length N + 1. In this case, $g^{l_m} 0 \in V^{m+N-1}$ and the maps

$$g_{m+k}: V^{m+k} \to V^{m+k-1}, \ k = 1, \dots, N,$$

are just the restrictions of g_m to the corresponding puzzle pieces.

If this cascade is maximal then the levels m-2 and m+N-1 are non-central. In this case, the length N+1 is equal to the *escaping time* it takes for the critical orbit to escape V^m under the iterates of g_m .

If the return to level m-1 is non-central, we will formally consider $\{V^{m-1}\}$ to be a "central cascade" of length 1. With this convention, the whole principal nest is decomposed into consecutive maximal central cascades. In fact, one of these cascades, the last one, can have infinite length:

PROPOSITION 11.4. A map f is renormalizable if and only if its principal nest ends up with an infinite central cascade $V^{m-1} \supset V^m \supset \ldots$. Moreover, in this case the map $g_m : V^m \to V^{m-1}$ is the renormalization of f.

PROOF. We will explain the "if" direction of this assertion.

Assume that we immediately observe an infinite central cascade $V^0 \supset V^1 \supset \ldots$. In this case we say that f is *immediately renormalizable*. One can show that this corresponds to parameters in the satellite M-copies attached to the main cardioid (compare §II.??).

In the immediately renormalizable case, the critical orbit never escapes V^1 under the iterates of $g_1 = f^p : V^1 \to V^0$ (where p is the number of α -rays). The map g_1 is a double covering of a smaller domain onto a bigger one but it is not a quadratic-like map, since the domains V^1 and V^0 have a common boundary (consisting of four external arcs). To turn this map into a quadratic-like, one should "thicken" the domains V^0 and V^1 a little bit (see Figure ...).

Assume that f is not immediately renormalizable. One can show that in this case, $V^m \Subset V^{m-1}$, so that $g_m : V^m \to V^{m-1}$ is a quadratic-like map with non-escaping critical point, which can be identified with the first renormalization of f.

Let us define the *height* of f as the number of the maximal central cascades in the principal nest. We see that f is renormalizable if and only if it has finite height.

Thus, the principal nest provides an algorithm to decide whether the map in question is renormalizable, whether this renormalization is of satellite type or otherwise, and to capture this renormalization.

On the negative side, the puzzle provides us with dynamical information only up to the first renormalization level. If we wish to penetrate deeper, we need to cut the Julia set of the renormalization into pieces and to go through its principal nest. Since the renormalization is a quadratic-like map rather than a quadratic polynomial, this motivates the need of the puzzle for quadratic-like maps. It will be discussed in §??.

46.5. Puzzle associated with periodic orbits.

47. Local connectivity of non-renormalizable Julia sets

THEOREM 11.5. Assume that all periodic points of f are hyperbolic and f is not infinitely renormalizable. Then the Julia set of f is locally connected.

Here is the main particular case of this result:

THEOREM 11.6. Assume both fixed point of f are repelling, and f is non-renormalizable. Then the Julia set of f is locally connected.

48. Local connectivity of of M at non-renormalizable points

Part 4

Hints and comments to the exersices

Chapter ??

0.5. Use Exercise 0.4. For a counterexample, see [?], Fig. 25.1.

0.7. a) Fix some $t \in [0, 1]$, and let $\gamma(t) = x$. Consider the decomposition of $\gamma^{-1}(B(x, \epsilon))$ into connected components I_n and J_k , where $I_n \ni t$ while $J_k \not\supseteq t$. Show that the paths $\gamma(J_n)$ do not accumulate on x and conclude that γ is weakly lc at x.

0.8. Construct a sequence of polygonal curves γ_n in \mathbb{R}^n connecting x to y such that:

• The vertices of the γ_n belong to K;

• γ_{n+1} is a *refinement* of γ_n , i.e., the vertices of γ_n are also vertices of γ_{n+1} ;

• $\|\gamma_n - \gamma_{n+1}\| \leq 1/2^n$, where $\|\cdot\|$ stands for the uniform norm.

REMARK 11.1. In fact, this is true without assuming that K is emebdded into \mathbb{R}^n . Indeed, any compact metric space X embeds into a Banach space (for instance, by associating to $x \in X$ the distance function $y \mapsto d(x, y)$), where one can repeat the above argument.

0.9. One direction: arc lc is stronger than weak lc. The other (non-trivial) direction: use the argument for Exercise 0.8.

??. Use that J is path lc and show that K is such.

??. Have fun!

0.18. Otherwise there is a sequence of arcs $\gamma_n \subset U_{i_n}$ whose diameter is bounded away from 0. Take an accumulation point *a* for the "mid-points" a_n of the γ_n . Then *K* is not lc at *a*.

??.

0.19. Consider the covering corresponding to the Ker of the monodromy action.

Chapter 1

1.5. The space of ϵ -separated triples of points is compact. The Möbius transformation ϕ depends continuously on the triple $(\alpha, \beta, \gamma) = \phi^{-1}(0, 1, \infty)$ as obvious from the explicit formula

$$\phi(z) = \frac{z - \alpha}{z - \gamma} \cdot \frac{\beta - \gamma}{\beta - \alpha}.$$

(This can also be used to verify equivalence of the two topologies.)

ŀ

1.7. The curvature of a metric $\rho(z)|dz|$ can be calculated by the formula:

$$\kappa(z) = -\frac{\Delta \log \rho(z)}{\rho(z)^2}$$

 $PSL(2, \mathbb{R})$ -invariance of the hyperbolic metric in the \mathbb{H} -model amounts to the identity:

$$\operatorname{Im} \phi(z) = \frac{\operatorname{Im} z}{|cz+d|^2}, \quad \phi(z) = \frac{az+b}{cz+d}.$$

Smooth isometries preserve angles between tangent vectors, and so conformal. In fact, one does not need to impose smoothness *a priori*. Any isometry is quasi-conformal (e.g., by the Pesin criterion, Theorem ??), and hence conformal by Weyl's Lemma (13.1).

1.9. (It is a generality about discrete groups of isometries of locally compact spaces.) If proper discontinuity (see the definition in §2) was violated, then there would exist as sequence of distinct motions $\gamma_n : \mathbb{D} \to \mathbb{D}$, and sequence of points $x_n \to x \in \mathbb{D}$ such that $\gamma(x_n) \to y \in \mathbb{D}$. Then, since the γ_n are isometries, for any neighborhood $U \Subset \mathbb{D}$, the family of maps $\gamma_n : U \to \mathbb{D}$ would be uniformly bounded and equicontinuous. Hence it would be pre-compact, contradicting discreteness.

1.38. Formal rules of differentiation with respect to (z, \bar{z}) look as if they are independent variables (these rules are particularly clear on the level of formal power series). For instance:

$$\partial_{\bar{z}}(\tau \circ \phi) = ((\partial_z \tau) \circ \phi) \cdot \partial_{\bar{z}} \phi + ((\partial_{\bar{z}} \tau) \circ \phi)) \cdot \partial_{\bar{z}} \bar{\phi}.$$

(Here one should think of $\tau \circ \phi$ as $\tau(\phi, \overline{\phi})$.) In case of holomorphic ϕ , we have

$$\partial_{\bar{z}}\bar{\phi} = \overline{\partial_z\phi} = \overline{\phi'}, \quad \partial_{\bar{z}}\phi = 0.$$

1.80. Let $U_n \in U$ be an increasing sequence of domains exhausting U, and let

$$\operatorname{dist}(\phi, \psi) = \sum \frac{1}{2^n} \sup_{z \in U_n} d_s(\phi(z), \psi(z)).$$

1.81. Consider a sequence of holomorphic functions $1/\phi_n(z)$ (which are the original functions written in terms of the local chart 1/z near ∞ in the target Riemann sphere). Apply the Hurwitz Theorem on the stability of roots of holomorphic functions.

1.76. Push the hyperbolic metric on \mathbb{H} forward to \mathbb{D}^* by the universal covering map $\mathbb{H} \to \mathbb{D}^*$, $z \mapsto e^{iz}$.

1.110. (iii) An ideal quadrilateral consisting of two adjacent triangles of the tiling gives us a fundamental domain of λ . In the \mathbb{H} -model, we can normalize it so that it is bounded by two vertical lines $x = \pm 1$ and two half-circles $|z \pm 1/2| = 1/2$. Then the boundary identifications are given by two parabolic deck transformations $z \mapsto z + 2$ and $z \mapsto z/(2z + 1)$. They generate the group of deck transformations, on the one hand, and the group Γ_2 , on the other.

1.88. Without loss of generality, we can assume that $U = \mathbb{D}$, the functions ψ do not collide in \mathbb{D}^* , $\psi_1 \equiv \infty$ and $\psi \equiv \psi_2$ has a pole at 0. Then the functions ϕ_n are holomorphic on \mathbb{D} and form a normal family on \mathbb{D}^* . By Exercise 1.85, we can assume that the ϕ_n are either uniformly bounded on each \mathbb{T}_r , $r \in (0, 1)$, or

(48.1)
$$\phi_n \to \infty$$
 uniformly on \mathbb{T}_r .

In the first case, the Maximal Principle completes the proof, so assume (48.1) occurs. If $\phi_{n(k)}(0) \neq 0$ for a subsequence n(k), then by the Minimum Principle $\phi_{n(k)} \to \infty$ uniformly on \mathbb{D}_r , and we are done. So, we can assume that $\phi_n(0) = 0$ for all n. Then the winding number of the curve $\phi_n : \mathbb{T}_r \to \mathbb{C}^*$ around 0 is positive. But by (48.1), the curve $\phi_n - \psi : \mathbb{T}_r \to \mathbb{C}^*$ eventually has the same winding number around 0 (r should be selected so that ψ does have poles on \mathbb{T}_r) and hence the equation $\phi_n(z) = \psi(z)$ has a solution in \mathbb{D}_r .

1.67. The path family Γ overflows the half-annulus $\mathbb{A}(1, R) \cap \mathbb{H}$, which implies the lower estimate for $\theta(R)$. Similarly, one can obtain the lower estimate for the dual path family Γ' (connecting $(-\infty, 0]$ to [1, R] in \mathbb{H} . This yields the upper estimate for $\theta(R) = 1/\mathcal{L}(\Gamma')$.

1.113 Associate to a point $a \in \mathbb{T}$ the prime end of \mathbb{D} represented by nests of cross-cuts shrinking to a.

1.120. Use the Schwarz Reflection Principle.

1.143 Apply the Index Formula to the gradient vector field ∇G in a region $\{z: 0 < \epsilon < G(z) < R\}$. (Or apply the Morse theory.)

Chapter 2

2.5. If $|\mu| < 1$ then A can be deformed to $z \mapsto az$ through invertible operators.

2.6. Start with equivariance. Let $T \in SL^{\#}(2, \mathbb{R})$. It acts on $\mathbb{C}_{\mathbb{R}}$ as $z \mapsto \alpha z + \beta \overline{z}$ with $|\alpha|^2 - |\beta|^2 = 1$. The Beltrami coefficient of the pullback $T^*(dz + \mu d\overline{z})$ is equal to $(\overline{\alpha}\mu + \beta)/(\overline{\beta}\mu + \alpha)$, which is the standard action of $SL(2, \mathbb{R})$ on \mathbb{D} .

Now we can check that the correspondence $\operatorname{Conf}(V) \approx \mathbb{D}$ is isometry. Since both actions of $\operatorname{SL}(2, \mathbb{R})$ preserve the hyperbolic metrics, it is sufficient to check that $\operatorname{dist}_{\operatorname{hyp}}(\sigma, \mu) = \operatorname{dist}_{\operatorname{hyp}}(0, \mu)$, where the former distance in measured in $\operatorname{Conf}(V)$, while the latter is measured in \mathbb{D} . But this is what the first formula of (11.3) tells us.

2.22 Consider the points $z_n = x + (z - x)/2^n$, $0 \le n \le N$, where $1/2 < |z_{N+1}| \le 1$, and use $|z'_n - z'_{n-1}| \le L|z'_n|$ inductively.

??: Ahlfors-Beurling Extension. See [A], Ch IV, Theorem 2.

Chapter 3

3.6. A quadratic differential $\phi \in \mathcal{Q}$ can be represented as $\phi(z)dz^2$ where $\phi(z)$ is a holomorphic function on $\hat{\mathbb{C}} \smallsetminus \mathcal{P}$. Since $\int |\phi| < \infty$, this function can have at most simple poles at finite points z_i , $i = 1, \ldots, n-1$, and $\phi(z) = O(|z|^{-3})$ near ∞ (which is equivalent to saying that the differential $\phi(z)dz^2$ has a simple pole at ∞). Hence

$$\phi(z) = \sum_{i=1}^{n-1} \frac{\lambda_i}{z - z_i}$$

with $\sum \lambda_i = 0$ and $\sum \lambda_i \sum_{k \neq i} z_k = 0$. These two linear conditions are independent,

and in fact, $(\lambda_1, \ldots, \lambda_{n-3})$ can be selected as global coordinates on the correspondent subspace (as the the right-most minor of the corresponding $2 \times (n-1)$ matrix is equal to $z_{n-1} - z_{n-2} \neq 1$).

3.10. Let $[S_n, \phi_n]$ converge to $[S, \phi]$ in $\mathcal{T}(S_0)$. Then one can select representatives ϕ_n and qc maps $h_n : S_n \to S$ with $\text{Dil}(h_n) \to 0$ such that $h_n \circ \phi_n = \phi$. Lift these maps to \mathbb{H} normalizing the Φ_n at three points. Then use Theorem 2.31 to show that the Φ_n converge to Φ uniformly on \mathbb{H} .

3.14. Let $\{g_{\alpha}\}$ be the projective atlas on V. Let us write f in the local parameter $z = g_{\alpha}(x)$ (i.e., consider the function $f_{\alpha} = f \circ g_{\alpha}^{-1}$), and let us take its Schwarzian $Sf_{\alpha}(z) dz^2$. Let $\zeta = g_{\beta}(x)$ be another local chart (with an overlapping domain), and let $\zeta = A_{\beta\alpha}(z)$ be the transit Möbius map. Then $f_{\beta} \circ A_{\beta\alpha} = f_{\alpha}$, and the Chain Rule (19.4) translates into the property that the quadratic differential $Sf_{\alpha}(z) dz^2$ is the pullback of $Sf_{\beta}(\zeta) d\zeta^2$ under $A_{\beta\alpha}$. This means by definition that these local expressions determine a global quadratic differential on V.

Chapter 4

4.9. It follows from the chain rule: $Df^n(z) = \prod_{k=0}^{n-1} Df(f^k z)$.

4.12. (i) Consider fixed points of f and their preimages.

(ii) It is a generality about full sets: a non-trivial loop γ in int K would break $\hat{\mathbb{C}} \smallsetminus K$ into two pieces.

4.36. For $z \in D_f(\alpha)$, $f^n \to \alpha$ uniformly on a neighborhood of z. Let D be the component of int K(f) containing z. Then by normality of the family $\{f^n | D\}$, the $f^n \to \alpha$ uniformly on compact subsets of D.

4.37 (i) $D^0(\alpha)$ is the component of $\{z : f^{pn}(z) \to \alpha \text{ as } n \to \infty\}$ containing α . (ii) Let $P_{\infty} = \cup P_n$. Then $f^p(\partial P_{\infty}) = \partial P_{\infty}$ since $f^p(\partial P_n) = \partial P_{n-1}$.

4.60. The size of the gap in $\mathbb{D}(z, \rho)$ depends lower semi-continuous on z.

5.12. Note that the foliation by round circles is defined dynamically as the closures of the equivalence classes

$$z \sim \zeta : \exists n : g^n z = g^n \zeta,$$

sometimes called "small orbits". Hence a germ ϕ commuting with g must respect this foliation. It follows that ϕ is linear (even if it mapped just one round circle onto a round circle).

4.70 Since $f: D \to D_1$ is a conformal isomorphism, the push-forward $f_*\mu_0$ is a conformal structure of D_1 with the same dilatation as μ_0 . For the same reason, $f_*^2\mu_0$ is a conformal structure of D_2 with the same dilatation, etc. By pushing it further by all iterated of f, we obtain an invariant measurable conformal structure μ on orb D with the same dilatation as the original structure on D.

Let us now pull this structure back to preimages of the domains D_n . Of course, one of these preimages can contain the critical point, where the pullback is not well defined. However, it does not cause a problem since a measurable conformal structure needs to be defined only almost everywhere. Since f is locally conformal outside the critical point, the pullback preserves the dilatation of the structure. Iterating this procedure, we obtain an invariant conformal structure on the grand

orbit Orb $D = \bigcup_{m=0}^{\infty} f^{-m}(\text{orb}\,D)$ (undefined on the critical set C_f (21.1)) with the

same dilatation as the initial structure.

Let us extend this structure to $\hat{\mathbb{C}} \setminus \operatorname{Orb} D$ as the standard one, σ . As σ is invariant in the first place (and has no dilatation), we obtain an invariant measurable structure with bounded dilatation on the whole Riemann sphere.

4.79. Take a Jordan curve Γ close to ∂U with winding number 1 around the origin and, look at the curve $g: \Gamma \to \mathbb{C}$, and apply the Argument Principle.

??. First consider any diffeomorphism $h_1 : \partial U' \to \mathbb{T}_{r^2}$, then lift it to a diffeomorphism $h_2 : \partial U \to \mathbb{T}_r$ satisfying (49.3), and finally interpolate in between h_1 and h_2 .

4.84. The fixed point is attracting by the Schwarz Lemma.

By Exercise 1.118, g extends continuously to the unit circle T. By the Schwarz Reflection Principle, g extends to the whole sphere making it a degree two Blyaschke

product. To bring it to the normal form (28.2), put the fixed point α to the origin. To show that it is expanding on \mathbb{T} , use the hyperbolic metric in $\overline{\mathbb{C}} \smallsetminus \operatorname{clorb}\{c, 1/\overline{c}\}$, where $c \in \mathbb{D}$ is the critical point of g (compare Theorem 4.53). Finally, $B'_a(0) = -a$.

4.85. For a tangent vector v, let

$$\|v\|_{\rho} = \sup_{n} \max_{1 \le i \le 2^n} \lambda^n |Dg_i^{-n}v|,$$

where g_i^{-n} are the inverse branches of g^n .

4.86. Uniqueness easily follows from the expanding property. To prove existence, lift g to the universal covering. We obtain an orientation preserving diffeomorphism $G : \mathbb{R} \to \mathbb{R}$ with the equivariance property: G(x + 1) = G(x) + 2. Obviously, the equation G(x) = x has a solution.

Or, apply the Lefschetz formula instead: the Lefschetz number of g is equal to 1 and the index of any fixed point is also 1 (since it is repelling).

Chapter 6

6.1. (iii) Recall the proof of Proposition 4.11.

(iv) It follows from the dichotomy: $\phi_n \to \infty$ locally uniformly on $\mathbb{C} \setminus M$, and $|\phi_n(z)| < 2$ on M (as in Proposition 4.29).

6.6. (Compare with Theorem ?? (i) Since the family of functions ϕ_n is not normal near $c_* \in \partial M$, one of the equations $\phi_n(c) = 0$ or $\phi_n(c) = \pm \sqrt{c}$ should have roots arbitrary close to $c_* \in \partial M$.

(ii) Consider, for instance, the β -fixed point as a function of c (it branches only at the main cusp 1/4). Then one of the equations $\phi_n(c) = \beta(c)$ or $\phi_n(c) = \sqrt{\beta(c) - c}$ should have roots arbitrary close to $c_* \in \partial M$.

6.33. For a point $\zeta = z^2 \in \mathbb{A}' = \mathbb{A}[R^2, R^4]$, let $H_c(\zeta) = (H_c(z))^2$. This map is correctly defined (does not depend on the choice of $z = \sqrt{\zeta}$), and is a self-homeomorphism of the annulus \mathbb{A}' identical on $\partial \mathbb{A}'$ and commuting with the group of rotations. Moreover, it commutes with $z \mapsto z^2$ (by definition) and depends holomorphically on c. Now extend it further to $\mathbb{A}[R^4, R^8]$, and so on.

Chapter 9: Quadratic-like maps

9.4: Dilatation of tubing. Since $\partial U'$ is 0-symmetric $\kappa(\delta)$ -quasicircle, there is a $L'(\kappa)$ -qs homeomorphism $B : \partial U' \to \mathbb{T}_{r^2}$. Since $g : \partial U \to \partial U'$ has a $C(\delta)$ -bounded distortion, B lifts to a L(D, L')-qc homeomorphism $B : \partial U \to \mathbb{T}_r$. These two qs homeomorphisms can be interpolated by a K-qc homeomorphism $B : A \to \mathbb{A}[r, r^2]$, with K depending only on L', L and bounds for mod A (see Exercise ??).

9.6: Canonical extension of μ . Pull μ back from the fundamental annulus $A = S_0^2 \cap S_\infty^2$ to its preimages $A_n = F^{-n}A$, $\mu|A_n = (F^n)^*(\mu|A)$. Since F is holomorphic in the local chart ϕ_0 (namely, equal to g), all these structures (in this local chart) have the same dilatation as $\mu|A$. Hence they form a single F-invariant measurable conformal structure with bounded dilatation on $S^2 \setminus \phi_0^{-1}K(g)$. Finally, let $\mu = (\phi_0)^* \sigma$ on $\phi_0^{-1}K(g)$.

Chapter 10: Primitive copies

49. INDEX

??: Solid torus. The map $H : \mathbb{V} \to \partial \Lambda \times \mathbb{D}$, $(\lambda, z) \mapsto (\lambda, h_{\lambda}^{-1}(z))$ straightens the tube \mathbb{V} to the solid torus. In this chart, the homotopy ψ_t can be given by moving the point $H(\phi(\lambda))$ straight to 0.

10.8: Hyperbolic components. Compare §28.4.

49. Index

Abelian differential §1.8 automorphism of f §20.2 adjustment §49.1 almost renormalization (canonical) §45.1.1 Banach slice §?? Basic Dichotomy §21.3 Beltrami differential §1.8, ?? Bers' Lemma §13.3 bifurcation saddle-node/superattracting/period doubling §21.5 Blyaschke product §28.1 Böttcher function/coordinate/equation §?? Böttcher motion §42.2 branched coverings $\S2.1$ Carathéodory convergence §4.8 Carathéodory Theorem (First/Second) §§??, ?? characteristic ray §32.2 chordal metric §2 conectedness locus C §?? conformal Riemannian metric §?? conformal sphere/plane/disk §?? Conformal Schönflis Theorem ?? continuum §2.2 convex hull (hyperbolic) $\S1.5.2$ critical set Crit_f §§21.1, 21.9 critical rays and equipotential §32.2 ∂_z and $\partial_{\bar{z}}$ partial derivatives §1.8.2 ∂ and $\bar{\partial}$ operators §1.8.2 dilatation §§11.1, 11.2 dilatation macroscopic/upper §§4.4, 12.2 dilatation quasisymmetric §12.3 Dirichlet barrier §7.7 Dirichlet Problem §7.3,7.7 distance ration §12.3 elementary Fuchsian group §1.5.3 embedding $\S2$ end, end-region, end-nest $\S1.1.2$ equator §2, §1.6.1 equipotential §§7.9, 32.2, 32.3, tubing equipotential §?? equipped q-l family §42.1 equivariant map $\S{31.2}$ ergodicity §31.6.3

```
expanding circle maps §28.1
    extension of a conformal structure (invariant/canonical) §26.1.4
    external coordinates \S??
    external rays \mathcal{R}_c^{\theta} and equipotentials \mathcal{E}_c^t \equiv \mathcal{E}_c^r §32.2
    external neighborhood \S 2.4
    external conjugacy §??
Fatou's Conjecture §28.4
    Fatou set §21.3
    fiber of a ql family §42.1
    full set, filling in the holes \S2.2 21.3
    fundamental annulus §27.1.1
    fundamental domain §20.8
Green function §§7.9, 32.3
harmonic functions and differentials §7.1
    Harnak Inequality §7.4
    Hausdorff distance dist_{\rm H} (2.1) §2
    holomorphic motion §??
    holomorphic equivalence relation §20.8
    Hubbard tree §23.4
    hull \S2.2
    hybrid conjugacy/equivalence/class §49.2
    hyperbolic maps/sets/parameters §§22.2, ??, 28.4
ideal circle, boundary §§1.1.2,1.5.3
    immersion §2
    invariant set, completely invariant set §20.1
Julia set §21.3
    J-stability \S{31.2}
    Jordan curve and disk §??
kneading sequence
    Koebe Distortion Theorem §4.4
limit set of a Fuchsin group §1.5.3
    local sections of an equivalence relation §20.8
Mandelbrot set
    Mandelbrot set of a ql family §42.1
    Mean Value Property (MVP) §7.2,7.5
    Montel Theorem (Little, Big and Refined) §4.3
    multiplier §22
    Multiplier Theorem §30
    normal families, \S4.3, 4.8, 7.4
one-point-per-end compactification §1.1.2
    orbit (grand/small) §20.1
path metric §2
    perfect set §21.3
    periodic point and cycle §20
    - attracting/repelling §22
```

49. INDEX

- neutral/parabolic §22 - Siegel/Cremer §?? Perron method §7.6 phase-parameter relation $\S42.5$, Eq. (42.4)Poisson Formula §7.3 porousity §23.6 proper map §2 pullback §20 quadratic differential §1.8 quadratic-like (q-l) map/conventional/real §27.1.1 quadratic-like family: equipped/proper/unfolded §42.1 quadrilateral §3.3.5 quasiconformal (qc) map §11.4 qc rigidity §31.7 quasiregular (qr) map §11.4 quasisymmetric (qs) map §?? ray (external) §§7.9, 32.2, 32.3, proper ray §32.2, tubing ray§?? renormalization complex/real, primitive/satellite §§27.3, 49.5 radius inner/outer §4.4 rectangle §3.3.5 removability qc/dynamical §23.7, §?? restricted quadratic family §42.2 Riemann-Hurwitz formula §2.2 rotation number $\S{22}$ saturation §20.8 Schwarz Lemma §4.1 set of discontinuity of a Fuchsian group $\S1.5.3$ semicontinuity (upper/lower) §2 simple closed curve \S ?? spider $\S18.2$ straightening §§49.2, 51, ?? structural stability §31.1 subharmonic and superharmonic functions §7.5 subpotential disk $\S{32.2}$ tensor §1.8 topologically holomorphic maps §2.1 topological entropy tube (of a q-l family) §42.1 tubing §49.3: Eq. 49.3, §42.5: Eq. 42.1 tubing position of the critical value $\S42.5$: Eq. (42.5) unimodal map §21.5 wandering set, weakly wandering set §20.1 wandering domains §26.2 Weyl's Lemma §13.1
Bibliography

ANALYSIS, GEOMETRY, TOPOLOGY

- [A] L. Ahlfors. Lectures on quasi-conformal maps. Van Nostrand Co, 1966.
- [AB] L. Ahlfors & L. Bers. Riemann's mapping theorem for variable metrics. Ann. Math., v. 72 (1960), 385 - 404.
- [Ab] W. Abikoff. The real analytic theory of Teichmüller space. Lecture Notes in Math., v. 820. Springer-Verlag, 1980.
- [B] P.P. Belinskii. General properties of quasiconformal maps (in Russian). "Nauka", Novosibirsk 1974.
- [He] W. Heymann
- [Hö] A. Hörmander. The analysis of linear partial differential operators I. A series o Comprehensive Studies in Mathematics, v. 256. Springer-Verlag 1990.
- [Mu] J. Munkres. Topology, Prentice Hall, 2000.
- [Ru] W. Rudin.

DYNAMICS

- [At] P. Atela. Bifurcations of dynamical rays in complex polynomials of degree two. Erg. Th. & Dyn. Syst., v. 12 (1991), 401-423.
- [Ba] Baker.
- [BR] Baker & Rippon.
- [Be] A. Beardon. Iteration of rational functions. Graduate texts in mathematics. Springer-Verlag, 1991.
- [Bu] X. Buff. Points fixes de renormalisation. Thesis, Orsay 1996.
- [Br] B. Branner. The Mandelbrot set. In: "Chaos and fractals", Proceedings of Symposia in Applied Mathematics, v. 39, 75-106.
- [BH] B. Branner & J.H. Hubbard. The iteration of cubic polynomials, Part II, Acta Math. v. 169 (1992), 229-325.
- [BKNS] H. Bruin, G. Keller, T. Nowicki & S. van Strien. Wild Cantor attractors exist. Ann. Math., v. 143 (1996), 97-130.
 - [1] . Carleson & W. Gamelin. Complex dynamics. Springer-Verlag, 1993.
- [D1]
- [D2] A. Douady. Chirurgie sur les applications holomorphes. In: Proceedings of the International Congress of Mathematicians (Berkeley 1986), 724-738.
- [D3] A. Douady. Description of compact sets in C. In: "Topological Methods in Modern Mathematics, A Symposium in Honor of John Milnor's 60th Birthday", Publish or Perish, 1993.
- [DH1] A. Douady and J.H. Hubbard.
- [DH2] A. Douady & J.H. Hubbard. Étude dynamique des polynômes complexes. Publication Mathematiques d'Orsay, 84-02 and 85-04.
- [DH3] A. Douady & J.H. Hubbard. On the dynamics of polynomial-like maps. Ann. Sc. Éc. Norm. Sup., v. 18 (1985), 287-343.
- [DH4] A. Douady & J.H. Hubbard, A proof of Thurston's topological characterization of rational functions. Acta Math. v. 171 (1993), 263–297.
- [E] H. Epstein. Fixed points of composition operators. Nonlinearity, v. 2 (1989), 305-310.
- [EL] A. Eremenko & M. Lyubich. Iterates of entire functions.
- [F] P. Fatou. Memoirs.
- [GK] L. Goldberg & L. Keen.

BIBLIOGRAPHY

- [GJ] J. Guckenheimer & S. Johnson. Distortion of S-unimodal maps. Ann. Math., v. 132 (1990), 71-130.
- [GS1] J. Graczyk & G. Swiatek. Induced expansion for quadratic polynomials. Stony Brook IMS Preprint # 1993/8.
- [GM] L. Goldberg & J. Milnor. Part II: Fixed point portraits. Ann. Sc. Éc. Norm. Sup.
- [H] J.H. Hubbard. Local connectivity of Julia sets and bifurcation loci: three theorems of J.-C. Yoccoz. In: "Topological Methods in Modern Mathematics, A Symposium in Honor of John Milnor's 60th Birthday", Publish or Perish, 1993.
- [HPS] R.H.Herman, I.F.Putnam, C.F.Skau, Ordered Bratteli diagrams, dimension groups and topological dynamics, International J. of Math., v.3 (1992), 827-864.
- [HJ] J. Hu & Y. Jiang. The Julia set of the Feigenbaum quadratic polynomial is locally connected. Preprint 1993.
- [K] J. Kahn. Holomorphic removability of Julia sets. Thesis, 1995.
- [Kr] S.L. Krushkal. Quasi-conformal mappings and Riemann surfaces. John Wiley 1979.
- [KN] G. Keller & T. Nowicki. Fibonacci maps re(al)visited. Ergod. Th. & Dynam. Syst., v. 15 (1995), 99-120.
- [LS] G. Levin, S. van Strien. Local connectivity of Julia sets of real polynomials, Preprint IMS at Stony Brook, 1995/5.
- [L1] M. Lyubich. The dynamics of rational transforms: the topological picture. Russian Math. Surveys, v. 41 (1986), 43-117.
- [L2] M. Lyubich. On the Lebesgue measure of the Julia set of a quadratic polynomial, Preprint IMS at Stony Brook, #1991/10.
- [L3] M. Lyubich. Milnor's attractors, persistent recurrence and renormalization, "Topological Methods in Modern Mathematics, A Symposium in Honor of John Milnor's 60th Birthday", Publish or Perish, 1993.
- [L4] M. Lyubich. Combinatorics, geometry and attractors of quasi-quadratic maps. Annals of Math., v. 140 (1994), 347-404.
- [L5] M. Lyubich. Teichmüller space of Fibonacci maps. Preprint IMS at Stony Brook, #1993/12.
- [L6] M. Lyubich. On the borderline of real and complex dynamics. Proc. IMS, Zürich 1994. Birkhäuser Verlag, 1995. 1203-1215.
- [L7] M. Lyubich. Renormalization ideas in conformal dynamics. Current Developments in Mathematics, 1995. International Press, Cambridge, MA, 155-184.
- [L8] M. Lyubich. Dynamics of quadratic polynomials, III. Parapuzzle and SBR measures. Preprint IMS at Stony Brook, # 1996/5.
- [L9] M. Lyubich. Feigenbaum-Coullet-Tresser Universality and Milnor's Hairiness Conjecture. Preprint IHES/M/96/61.
- [L10] M. Lyubich. How big is the set of infinitely renormalizable quadratics? Preprint 1996.
- [LM] M. Lyubich & J.Milnor. The unimodal Fibonacci map. Journal of AMS, v. 6 (1993), 425-457.
- [LY] M. Lyubich & M. Yampolsky. Complex bounds for real maps. Preprint MSRI # 034-95. To appear in Ann. Inst. Fourier.
- [LV] Lehto & Virtanen. Quasiconformal mappings in the plane. Springer 1973.
- [Mi1] J. Milnor. Dynamics in one complex variable: Introductory lectures, Stony Brook IMS Preprint # 1990/5.
- [Mi2] J. Milnor. Local connectivity of Julia sets: expository lectures. Preprint IMS at Stony Brook #1992/11.
- [Mi3] J. Milnor. Self-similarity and hairiness of the Mandelbrot set. In: "Computers in geometry and topology. Lect. Notes in Pure and Appl. Math., Dekker, 1989.
- [Mi4] J. Milnor. Periodic orbits, external rays and the Mandelbrot set: An expository account. Preprint 1995.
- [MT] J. Milnor & W. Thurston. On iterated maps of the interval. "Dynamical Systems", Proc. U. Md., 1986-87, ed. J. Alexander, Lect. Notes Math., v. 1342 (1988), 465-563.
- [McM1] C. McMullen. Families of rational maps, Ann. Math., 125 (1987), 467-493.
- [McM2] C. McMullen. Complex dynamics and renormalization. Annals of Math. Studies, v. 135, Princeton University Press, 1994.

BIBLIOGRAPHY

- [McM3] C. McMullen. Renormalization and 3-manifolds which fiber over the circle. Annals of Math. Studies, to appear.
- [McM4] C. McMullen. The classification of conformal dynamical systems. Current Developments in Mathematics, 1995. International Press, Cambridge, MA, 1995, pp. 323-360.
- [Ma] M. Martens. Distortion results and invariant Cantor sets for unimodal maps. Erg. Th. & Dyn. Syst., v. 14 (1994), 331-349.
- [MN] M. Martens & T. Nowicki. Invariant measures for Lebesgue typical quadratic maps. Preprint IMS at Stony Brook, # 1996/6.
- [Man] A. Manning. Logarithmic capacity and renormalizability for landing on the Mandelbrot set. Preprint 1995.
- [MS] W. de Melo & S. van Strien. One dimensional dynamics. Springer-Verlag, 1993.
- [P] C.L. Petersen. Local connectivity of some Julia sets containing a circle with an irrational rotation, Preprint IHES/M/94/26, to appear in Acta Math.
- [Poi] H. Poinvaré. Les methodes nouvelle de la mécanique cêleste.
- [Pra2] Ergodicity of conformal measures for unimodal polynomials. Preprint IMS at Stony Brook, # 1996/8.
- [Prz] F. Przytycki. Holomorphic Collet-Eckmann maps. Preprint 1995.
- [R] M. Rees. A possible approach to a complex renormalization problem. In: "Linear and Complex Analysis Problem Book 3, part II", p. 437-440. Lecture Notes in Math., v. 1574.
 [S1] D. Sullivan. CR note
- [S2] D. Sullivan. Preprint "Seminar on Conformal and Hyperbolic Geometry".
- [S3] D. Sullivan. No wandering domains.
- [S3] D. Sullivan. Quasiconformal homeomorphisms and dynamics, topology and geometry. Proc. ICM-86, Berkeley, v. II, 1216-1228.
- [S4] D. Sullivan. Bounds, quadratic differentials, and renormalization conjectures. AMS Centennial Publications, v. 2: Mathematics into Twenty-first Century (1992).
- [Sc1] D. Schleicher. The structure of the Mandelbrot set. Preprint 1995.
- [Sc2] D. Schleicher. Rational external rays of the Mandelbrot set. Preprint 1996.
- [Sh1] M. Shishikura
- [Sh2] M. Shishikura. The Hausdorff dimension of the boundary of the Mandelbrot set and Julia sets. Preprint IMS at Stony Brook, 1991/7.
- [Sh3] M. Shishikura. Topological, geometric and complex analytic properties of Julia sets. Proceedings of the ICM-94, 886-895. Birkhäuser.
- [SaN] L. Sario & M. Nakai. Classification theory of Riemann surfaces, Springer-Verlag, 1970.
- [Th] W. Thurston. The geometry and topology of 3-manifolds. Princeton University Lecture Notes, 1982.
- [V] L.I. Volkovyski. Quasi-conformal mappings (in Russian). Lvov University, 1954.
- [Ve] A.M. Vershik. Uniform algebraic approximation of shift and multiplication operators. Soviet Math. Dokl., v. 24 (1981), 97-100.
- [W] L. Wenstrom. Parameter scaling for the Fibonacci point. Preprint IMS at Stony Brook, # 1996/4.
- [Y] M. Yampolsky. Complex bounds for critical circle maps. Preprint IMS at Stony Brook, # 1995/12.
- [Yar] B. Yarrington. Local connectivity and Lebesgue measure of polynomial Julia sets. Thesis, Stony Brook 1995.