The Resident's View Revisited

Mary Rees

University of Liverpool

9th W. R. Hamilton Geometry and Topology Workshop,
Geometry and Groups after Thurston, 27-31 August 2013, HMI,
Trinity College Dublin.

Branched coverings

Branched coverings

- The sphere and the torus are the only compact surfaces to admit self-branched coverings of degree greater than one.
- Branched coverings of the sphere which are not homeomorphisms necessarily have branch points, also called critical values.
- Rational maps are examples of branched coverings.
- The dynamics of rational maps is one of the main areas of complex dynamics.

Critically finite branched coverings

Critically finite branched coverings

A branched covering f is said to be critically finite if the postcritical set

$$X(f) = \{f^n(c) : c \text{ critical }, n > 0\}$$

is finite.

Critically finite branched coverings

A branched covering f is said to be critically finite if the postcritical set

$$X(f) = \{f^n(c) : c \text{ critical }, n > 0\}$$

is finite.

Two critically finite branched coverings f_0 and f_1 are usually said to be Thurston equivalent if there is a homotopy f_t $(t \in [0,1])$ through critically finite branched coverings such that $X(f_t)$ varies isotopically for $t \in [0,1]$.

Critically finite branched coverings

A branched covering f is said to be critically finite if the postcritical set

$$X(f) = \{f^n(c) : c \text{ critical }, n > 0\}$$

is finite.

Two critically finite branched coverings f_0 and f_1 are usually said to be Thurston equivalent if there is a homotopy f_t ($t \in [0, 1]$) through critically finite branched coverings such that $X(f_t)$ varies isotopically for $t \in [0, 1]$.

In this talk a slightly stronger notion of Thurston equivalence will be used. We consider branched coverings of the Riemann sphere for which the critical values are numbered, and then f_0 and f_1 are said to be Thurston equivalent if the isotopy of $X(f_t)$ preserves the numbering of critical values.

Thurston's Theorem for critically finite branched coverings (1982)

Thurston's Theorem for critically finite branched coverings (1982)

The quotient by Möbius conjugation of a Thurston equivalence class is contractible to the Möbius conjugacy class of a unique rational map, if and only if a certain orbifold is hyperbolic, and a certain combinatorial condition holds, which can be described as the non-existence of a Thurston obstruction.

Thurston's theorem has been deliberately formulated as a result about the topology of a space of maps. It is a geometrisation result in two ways.

It gives a condition under which a map is holomorphic, modulo the appropriate type of homotopy equivalence (Thurston equivalence).

Thurston's theorem has been deliberately formulated as a result about the topology of a space of maps. It is a geometrisation result in two ways.

It gives a condition under which a map is holomorphic, modulo the appropriate type of homotopy equivalence (Thurston equivalence).

It also shows that the corresponding space of maps is contractible to a space with a geometric structure – although there is little to say about geometries on a space consisting of a single point. But a point is just the simplest case. . .

A generalisation

A generalisation

The elements of a connected topological space $B = B(f_0, Y(f_0))$ are $[f, Y(f)]$, and include $[f_0, Y(f_0)]$, where:

A generalisation

The elements of a connected topological space $B = B(f_0, Y(f_0))$ are $[f, Y(f)]$, and include $[f_0, Y(f_0)]$, where:

- f is a branched covering of the Riemann sphere;

A generalisation

The elements of a connected topological space $B = B(f_0, Y(f_0))$ are $[f, Y(f)]$, and include $[f_0, Y(f_0)]$, where:

- f is a branched covering of the Riemann sphere;
- $Y(f)$ is a finite set which contains all the critical values, which are numbered, and is the union of the set of critical values and $Z(f)$, where $f(Z(f)) \subset Z(f)$;

A generalisation

The elements of a connected topological space $B = B(f_0, Y(f_0))$ are $[f, Y(f)]$, and include $[f_0, Y(f_0)]$, where:

- f is a branched covering of the Riemann sphere;
- $Y(f)$ is a finite set which contains all the critical values, which are numbered, and is the union of the set of critical values and $Z(f)$, where $f(Z(f)) \subset Z(f)$;
- $Y(f)$ and$Z(f)$ vary isotopically with f, for $[f, Y(f)] \in B$;

A generalisation

The elements of a connected topological space $B = B(f_0, Y(f_0))$ are $[f, Y(f)]$, and include $[f_0, Y(f_0)]$, where:

- f is a branched covering of the Riemann sphere;
- $Y(f)$ is a finite set which contains all the critical values, which are numbered, and is the union of the set of critical values and $Z(f)$, where $f(Z(f)) \subset Z(f)$;
- $Y(f)$ and$Z(f)$ vary isotopically with f, for $[f, Y(f)] \in B$;
- $[f, Y(f)]$ denotes the conjugacy class of $(f, Y(f))$ by Möbius transformations, using only Möbius transformations which preserve the numbering of critical values.

A generalisation

The elements of a connected topological space $B = B(f_0, Y(f_0))$ are $[f, Y(f)]$, and include $[f_0, Y(f_0)]$, where:

- f is a branched covering of the Riemann sphere;
- $Y(f)$ is a finite set which contains all the critical values, which are numbered, and is the union of the set of critical values and $Z(f)$, where $f(Z(f)) \subset Z(f)$;
- $Y(f)$ and$Z(f)$ vary isotopically with f, for $[f, Y(f)] \in B$;
- $[f, Y(f)]$ denotes the conjugacy class of $(f, Y(f))$ by Möbius transformations, using only Möbius transformations which preserve the numbering of critical values.

For example, B could be the Thurston equivalence class of a critically finite branched covering f_0, with $Y(f_0) = Z(f_0) = X(f_0)$.

The connection with Teichmüller space

The connection with Teichmüller space

If $[f_t, Y(f_t)]$ is a path in B then $f_t = \varphi_t \circ f_0 \circ \psi_t^{-1}$ where:

The connection with Teichmüller space

If $[f_t, Y(f_t)]$ is a path in B then $f_t = \varphi_t \circ f_0 \circ \psi_t^{-1}$ where:

- φ_0 and ψ_0 are the identity is the identity and $t \mapsto \varphi_t$ and $t \mapsto \psi_t$ are isotopies;

The connection with Teichmüller space

If $[f_t, Y(f_t)]$ is a path in B then $f_t = \varphi_t \circ f_0 \circ \psi_t^{-1}$ where:

- φ_0 and ψ_0 are the identity is the identity and $t \mapsto \varphi_t$ and $t \mapsto \psi_t$ are isotopies;
- $\varphi_t(Y(f_0)) = Y(f_t)$, $\psi_t(Z(f_0) = Z(f_t)$ and φ_t and ψ_t are isotopic homeomorphisms through an isotopy which is constant on $Z(f_0)$;

The connection with Teichmüller space

If $[f_t, Y(f_t)]$ is a path in B then $f_t = \varphi_t \circ f_0 \circ \psi_t^{-1}$ where:

- φ_0 and ψ_0 are the identity is the identity and $t \mapsto \varphi_t$ and $t \mapsto \psi_t$ are isotopies;
- $\varphi_t(Y(f_0)) = Y(f_t)$, $\psi_t(Z(f_0) = Z(f_t)$ and φ_t and ψ_t are isotopic homeomorphisms through an isotopy which is constant on $Z(f_0)$;
- φ_t is determined up to isotopy constant on $Y(f_0)$, and post-composition by a Möbius transformation, by $[f_t, Y(f_t)]$: that is, as an element $[\varphi_t]$ of the Teichmüller space $\mathcal{T}(Y(f_0))$ of the sphere with marked set bijective to $Y(f_0)$.

The connection with Teichmüller space

If $[f_t, Y(f_t)]$ is a path in B then $f_t = \varphi_t \circ f_0 \circ \psi_t^{-1}$ where:

- φ_0 and ψ_0 are the identity is the identity and $t \mapsto \varphi_t$ and $t \mapsto \psi_t$ are isotopies;
- $\varphi_t(Y(f_0)) = Y(f_t)$, $\psi_t(Z(f_0) = Z(f_t)$ and φ_t and ψ_t are isotopic homeomorphisms through an isotopy which is constant on $Z(f_0)$;
- φ_t is determined up to isotopy constant on $Y(f_0)$, and post-composition by a Möbius transformation, by $[f_t, Y(f_t)]$: that is, as an element $[\varphi_t]$ of the Teichmüller space $\mathcal{T}(Y(f_0))$ of the sphere with marked set bijective to $Y(f_0)$.

This correspondence maps the universal cover of B to $\mathcal{T}(Y(f_0))$ with contractible fibres, so that B is a $K(\pi, 1)$. The fundamental group maps to a subgroup of the pure mapping class group $PMG(\overline{\mathbb{C}}, Y(f_0))$. Since B is a $K(\pi, 1)$, the Topographer's View is a result about the structure of its fundamental group.

The Topographer's View

The Topographer's View

I obtained a result, or sequence of results, which I called The Topographer's View – which I do not particularly want to revisit at this juncture, but it is not possible to separate the Topographer and Resident's views completely, because they complement each other.

The Topographer's View

I obtained a result, or sequence of results, which I called The Topographer's View – which I do not particularly want to revisit at this juncture, but it is not possible to separate the Topographer and Resident's views completely, because they complement each other. The Topographer's View of $B = B(f_0, Y(f_0))$ is a homotopy equivalence to an ordered graph of countably manytopological spaces of maps, where these topological spaces have some geometric structure.

The Topographer's View

I obtained a result, or sequence of results, which I called The Topographer's View – which I do not particularly want to revisit at this juncture, but it is not possible to separate the Topographer and Resident's views completely, because they complement each other. The Topographer's View of $B = B(f_0, Y(f_0))$ is a homotopy equivalence to an ordered graph of countably manytopological spaces of maps, where these topological spaces have some geometric structure.
The "base" geometric pieces are the rational maps in B quotiented by Möbius conjugation.

An important part of the result, not easy, is that the inclusion of each component V of rational maps in the larger space B is injective on π_1.

An important part of the result, not easy, is that the inclusion of each component V of rational maps in the larger space B is injective on π_1.
This result was only obtained for B consisting of degree two maps, or maps "of polynomial type", $Z(f_0)$ contained in the full orbit of periodic critical points.

Question: nature of the embedding

Question: nature of the embedding

Having seen that the universal cover $\widetilde{V}$ of V embeds in $\mathcal{T}(Y(f_0))$, one can obviously ask about the nature of the embedding. From the group-theoretic point of view, this is a question about a subgroup of the Pure Mapping Class Group $PMG(\overline{\mathbb{C}}, Y(f_0))$ which identifies with the fundamental group of the space of rational maps. Of course the embedding is Lipschitz, but I know nothing about the inverse map.

A projection of the embedding

A projection of the embedding

In the cases that we consider, f_0 is of degree two, and $Y(f_0) \setminus Z(f_0) = Y \setminus Z$ is a single point, a critical value of f_0, denoted by $v_2 = v_2(f_0)$.

A projection of the embedding

In the cases that we consider, f_0 is of degree two, and $Y(f_0) \setminus Z(f_0) = Y \setminus Z$ is a single point, a critical value of f_0, denoted by $v_2 = v_2(f_0)$.
There is a natural projection from $\mathcal{T}(Y)$ to the universal cover $\widetilde{\overline{\mathbb{C}} \setminus Z}$ of $\overline{\mathbb{C}} \setminus Z$, defined as follows.

A projection of the embedding

In the cases that we consider, f_0 is of degree two, and $Y(f_0) \setminus Z(f_0) = Y \setminus Z$ is a single point, a critical value of f_0, denoted by $v_2 = v_2(f_0)$.
There is a natural projection from $\mathcal{T}(Y)$ to the universal cover $\widetilde{\overline{\mathbb{C}} \setminus Z}$ of $\overline{\mathbb{C}} \setminus Z$, defined as follows.
Let $\pi_Z : \mathcal{T}(Y) \to \mathcal{T}(Z)$ denote the natural projection and let d_Y and d_Z denote the Tiechmüller metrics on $\mathcal{T}(Y)$ and $\mathcal{T}(Z)$

If $[\varphi_t]$ is a path in $\mathcal{T}(Y)$ from a basepoint $[\varphi_0] = x_0$, then let $\chi_t : \overline{\mathbb{C}} \to \overline{\mathbb{C}}$ be the unique homeomorphism minimizing qc-distortion such that

$$\pi_Z([\varphi_t]) = \pi_Z([\chi_t \circ \varphi_0])$$

If $[\varphi_t]$ is a path in $\mathcal{T}(Y)$ from a basepoint $[\varphi_0] = x_0$, then let $\chi_t : \overline{\mathbb{C}} \to \overline{\mathbb{C}}$ be the unique homeomorphism minimizing qc-distortion such that

$$\pi_Z([\varphi_t]) = \pi_Z([\chi_t \circ \varphi_0])$$

Then $t \mapsto \chi_t^{-1} \circ \varphi_t(v_2)$ defines an element of $\widetilde{\overline{\mathbb{C}} \setminus Z}$.

If $[\varphi_t]$ is a path in $\mathcal{T}(Y)$ from a basepoint $[\varphi_0] = x_0$, then let $\chi_t : \overline{\mathbb{C}} \to \overline{\mathbb{C}}$ be the unique homeomorphism minimizing qc-distortion such that

$$\pi_Z([\varphi_t]) = \pi_Z([\chi_t \circ \varphi_0])$$

Then $t \mapsto \chi_t^{-1} \circ \varphi_t(v_2)$ defines an element of $\widetilde{\overline{\mathbb{C}} \setminus Z}$. The composition with the embedding gives a map

$$\rho : \widetilde{V} \to \widetilde{\overline{\mathbb{C}} \setminus Z}.$$

The Resident's View Theorem (or part of it)

The Resident's View Theorem (or part of it)

In the cases considered, V is a finite type Riemann surface of negative Euler characteristic, and so, of course is $\overline{\mathbb{C}} \setminus Z$. So the universal cover of each of these Riemann surfaces is the unit disc up to conformal equivalence, with boundary the unit circle.

The Resident's View Theorem (or part of it)

In the cases considered, V is a finite type Riemann surface of negative Euler characteristic, and so, of course is $\overline{\mathbb{C}} \setminus Z$. So the universal cover of each of these Riemann surfaces is the unit disc up to conformal equivalence, with boundary the unit circle.

$\rho : \widetilde{V} \to \widetilde{\overline{\mathbb{C}} \setminus Z}$ extends continuously and monotonically to map the boundary $\partial\widetilde{V}$ into $\partial\widetilde{\overline{\mathbb{C}} \setminus Z}$ with just countably many discontinuities which can be naturally characterised, and where right and left limits exist.

The Resident's View Theorem (or part of it)

In the cases considered, V is a finite type Riemann surface of negative Euler characteristic, and so, of course is $\overline{\mathbb{C}} \setminus Z$. So the universal cover of each of these Riemann surfaces is the unit disc up to conformal equivalence, with boundary the unit circle.

$\rho : \widetilde{V} \to \widetilde{\overline{\mathbb{C}} \setminus Z}$ *extends continuously and monotonically to map the boundary* $\partial\widetilde{V}$ *into* $\partial\widetilde{\overline{\mathbb{C}} \setminus Z}$ *with just countably many discontinuities which can be naturally characterised, and where right and left limits exist.*
Once the continuity is proved, monotonicity is straightforward.

The limit of ρ along geodesics

We use d_P to denote the Poincaré (or hyperbolic) metric on the unit disc (the universal cover of $\overline{\mathbb{C}} \setminus Z$).

The limit of ρ along geodesics

We use d_P to denote the Poincaré (or hyperbolic) metric on the unit disc (the universal cover of $\overline{\mathbb{C}} \setminus Z$).

Theorem

$\lim_{x \to \infty} \rho(x)$ *exists along any half geodesic segment* ℓ *in* $\mathcal{T}(Y)$ *such that* $\lim_{x \to \infty} d_P(0, \rho(x)) = +\infty$. *In fact, if* ℓ *starts at* x_0 *and* $d_P(0, \rho(z)) \geq n$ *for all* $z \in \ell$ *between* x *and* y *then for a suitable constant* C,

$$|\rho(x) - \rho(y)| \leq Cne^{-n}.$$

The limit of ρ along geodesics

We use d_P to denote the Poincaré (or hyperbolic) metric on the unit disc (the universal cover of $\overline{\mathbb{C}} \setminus Z$).

Theorem

$\lim_{x\to\infty} \rho(x)$ exists along any half geodesic segment ℓ in $\mathcal{T}(Y)$ such that $\lim_{x\to\infty} d_P(0, \rho(x)) = +\infty$. In fact, if ℓ starts at x_0 and $d_P(0, \rho(z)) \geq n$ for all $z \in \ell$ between x and y then for a suitable constant C,

$$|\rho(x) - \rho(y)| \leq Cne^{-n}.$$

It therefore seems natural to consider geodesic segment with endpoints in $\widetilde{V}$ (which is a subset of $\mathcal{T}(Y)$) and to compare ρ on the geodesic with a path with the same endpoints.

A key idea in the proof of this theorem is that geodesic segments on $\overline{\mathbb{C}} \setminus Z$ tend to have many self-intersections.

A key idea in the proof of this theorem is that geodesic segments on $\overline{\mathbb{C}} \setminus Z$ tend to have many self-intersections.

More precisely, given any geodesic segment γ of length Δ there is a constant C_1 such if we consider lifted geodesics in the unit disc starting from 0, every such geodesic segment of length n ends within Euclidean distance $C_1 e^{\Delta - n}$ of a geodesic segment which has endpoints within a bounded Poincaré distance of the endpoints of γ.

A key idea in the proof of this theorem is that geodesic segments on $\overline{\mathbb{C}} \setminus Z$ tend to have many self-intersections.

More precisely, given any geodesic segment γ of length Δ there is a constant C_1 such if we consider lifted geodesics in the unit disc starting from 0, every such geodesic segment of length n ends within Euclidean distance $C_1 e^{\Delta - n}$ of a geodesic segment which has endpoints within a bounded Poincaré distance of the endpoints of γ.

So we can choose γ to do anything we like by choosing Δ sufficiently large, e.g. to cut $\overline{\mathbb{C}} \setminus Z$ into topological discs with at most one puncture.

Continuity at punctures

It is relatively easy to check that $\lim_{x \to \infty} \rho(x)$ exists along half-geodesics in V ending at any puncture corresponding to a rational map f at which $v_2(f) \in Z(f)$.

Continuity at punctures

It is relatively easy to check that $\lim_{x\to\infty} \rho(x)$ exists along half-geodesics in V ending at any puncture corresponding to a rational map f at which $v_2(f) \in Z(f)$.

Note that $f \notin V$ because $Z(g) \cup \{v_2(g)\} = Y(g)$ varies isotopically for $g \in V$.

Continuity at punctures

It is relatively easy to check that $\lim_{x\to\infty}\rho(x)$ exists along half-geodesics in V ending at any puncture corresponding to a rational map f at which $v_2(f)\in Z(f)$.

Note that $f\notin V$ because $Z(g)\cup\{v_2(g)\}=Y(g)$ varies isotopically for $g\in V$.

In these cases, $\lim_{x\to\infty}\rho(x)$ is a "lift" of a point in Z – the endpoint of a geodesic in $\overline{\mathbb{C}}\setminus Z$ which ends at a point of Z.

Continuity at punctures

It is relatively easy to check that $\lim_{x\to\infty}\rho(x)$ exists along half-geodesics in V ending at any puncture corresponding to a rational map f at which $v_2(f)\in Z(f)$.

Note that $f\notin V$ because $Z(g)\cup\{v_2(g)\}=Y(g)$ varies isotopically for $g\in V$.

In these cases, $\lim_{x\to\infty}\rho(x)$ is a "lift" of a point in Z – the endpoint of a geodesic in $\overline{\mathbb{C}}\setminus Z$ which ends at a point of Z.

It is also quite easy to show that, for a lift x of any other puncture of V, either $\lim_{y\to x}\rho(y)$ exists, or left and right limits exist outside a horosphere.

Discontinuities

The discontinuities $x \in \partial\widetilde{V}$ are the points such that $\liminf_{y\to x} d_P(0,\rho(y)) < \infty$. These are quite easily characterised and also it is quite easy to show that right and left limits exist outside a horosphere or Stoltz angle at such a point. However I have never managed to find such a point of *pseudo-Anosov type*. It is possible (although unlikely) that they do not exist.

Continuity along one path is sufficient

Continuity along one path is sufficient

The points at which limits exist are sufficiently dense that the proof of the Resident's View is completed by showing that for each $x \in \widetilde{V}$, and for a choice of basepoint x_0, there is just one path x_t in $\widetilde{V}$ from x_0 to x on which the same uniform continuity for limits holds, that is, if $d_P(\rho(x_u), 0) \geq n$ for all $u \in [s, t]$ then $|\rho(x_s) - \rho(x_t)| \leq C_1 n e^{-n}$.

Continuity along one path is sufficient

The points at which limits exist are sufficiently dense that the proof of the Resident's View is completed by showing that for each $x \in \widetilde{V}$, and for a choice of basepoint x_0, there is just one path x_t in $\widetilde{V}$ from x_0 to x on which the same uniform continuity for limits holds, that is, if $d_P(\rho(x_u), 0) \geq n$ for all $u \in [s, t]$ then $|\rho(x_s) - \rho(x_t)| \leq C_1 n e^{-n}$.

It is natural to start with the geodesic segment in $\mathcal{T}(Y(f_0))$ and to try to modify this to a path in $\widetilde{V}$.

Continuity along one path is sufficient

The points at which limits exist are sufficiently dense that the proof of the Resident's View is completed by showing that for each $x \in \widetilde{V}$, and for a choice of basepoint x_0, there is just one path x_t in $\widetilde{V}$ from x_0 to x on which the same uniform continuity for limits holds, that is, if $d_P(\rho(x_u), 0) \geq n$ for all $u \in [s, t]$ then $|\rho(x_s) - \rho(x_t)| \leq C_1 n e^{-n}$.

It is natural to start with the geodesic segment in $\mathcal{T}(Y(f_0))$ and to try to modify this to a path in $\widetilde{V}$.

It is only this part of the proof that I have been revisiting.

Use of a chain of geodesic segments to get the path in $\widetilde{V}$

Use of a chain of geodesic segments to get the path in $\widetilde{V}$

Suppose that $[x_0, x]$ is a geodesic segment in $\mathcal{T}(Y)$ with endpoints x_0 and x in $\widetilde{V}$.

Use of a chain of geodesic segments to get the path in $\widetilde{V}$

Suppose that $[x_0, x]$ is a geodesic segment in $\mathcal{T}(Y)$ with endpoints x_0 and x in $\widetilde{V}$.

The path in $\widetilde{V}$ with these endpoints is obtained by taking $\lim_{n\to\infty} x_n(x_1) \in \widetilde{V}$ corresponding to $x_1 \in [x_0, x]$, where:

Use of a chain of geodesic segments to get the path in $\widetilde{V}$

Suppose that $[x_0, x]$ is a geodesic segment in $\mathcal{T}(Y)$ with endpoints x_0 and x in $\widetilde{V}$.

The path in $\widetilde{V}$ with these endpoints is obtained by taking $\lim_{n\to\infty} x_n(x_1) \in \widetilde{V}$ corresponding to $x_1 \in [x_0, x]$, where:

$$d_Y(x_n, x_{n+1}) - d_Z(x_n, x_{n+1}) \leq C;$$

Use of a chain of geodesic segments to get the path in $\widetilde{V}$

Suppose that $[x_0, x]$ is a geodesic segment in $\mathcal{T}(Y)$ with endpoints x_0 and x in $\widetilde{V}$.

The path in $\widetilde{V}$ with these endpoints is obtained by taking $\lim_{n\to\infty} x_n(x_1) \in \widetilde{V}$ corresponding to $x_1 \in [x_0, x]$, where:

$$d_Y(x_n, x_{n+1}) - d_Z(x_n, x_{n+1}) \leq C;$$

a chain of *long thick and dominants* (α_i, ℓ_i) with $\ell_i \subset [x_i, x_{i+1}]$ and $\alpha_i \cap \alpha_{i+1} \neq \emptyset$ is non-cancelling, so that ℓ_i is a bounded d_{α_i}-distance from $[x_1, x_n]$ for $1 \leq i < n$.

Use of a chain of geodesic segments to get the path in $\widetilde{V}$

Suppose that $[x_0, x]$ is a geodesic segment in $\mathcal{T}(Y)$ with endpoints x_0 and x in $\widetilde{V}$.

The path in $\widetilde{V}$ with these endpoints is obtained by taking $\lim_{n\to\infty} x_n(x_1) \in \widetilde{V}$ corresponding to $x_1 \in [x_0, x]$, where:

$$d_Y(x_n, x_{n+1}) - d_Z(x_n, x_{n+1}) \leq C;$$

a chain of *long thick and dominants* (α_i, ℓ_i) with $\ell_i \subset [x_i, x_{i+1}]$ and $\alpha_i \cap \alpha_{i+1} \neq \emptyset$ is non-cancelling, so that ℓ_i is a bounded d_{α_i}-distance from $[x_1, x_n]$ for $1 \leq i < n$.

From this we can deduce properties of the difference between $\rho(x_1)$ and $\rho(x_n)$.

In revisiting this proof, my hope was to use a result about chains of long thick and dominants on Teichmüller geodesics which I found rather hard to prove, and useful in another situation.

In revisiting this proof, my hope was to use a result about chains of long thick and dominants on Teichmüller geodesics which I found rather hard to prove, and useful in another situation.

I am not quite there yet, but have found an application to the analysis of the relation between $\tau([x,\tau(x)])$ and $[\tau(x),\tau^2(x)]$ where τ is the appropriate analogue for $B(f_0, Y(f_0))$ of the *Thurston pullback*.

The use of Teichmüller space

The proof of Thurston's Theorem for critically finite branched coverings used a distance-non-increasing map $\tau : \mathcal{T}(X(f_0)) \to \mathcal{T}(X(f_0))$ known as the *Thurston pullback*. For a suitable integer k, τ^k is a uniform contraction on any set $x : d(x, \tau(x) < M\}$. So there is a unique fixed point.

The use of Teichmüller space

The proof of Thurston's Theorem for critically finite branched coverings used a distance-non-increasing map $\tau : \mathcal{T}(X(f_0)) \to \mathcal{T}(X(f_0))$ known as the *Thurston pullback*. For a suitable integer k, τ^k is a uniform contraction on any set $x : d(x, \tau(x) < M\}$. So there is a unique fixed point.

The map τ is given by $\tau([\varphi]) = [\psi]$ where $\varphi \circ f_0 = s \circ \psi$ where s is a holomorphic branched covering and ψ is a homeomorphism.

The use of Teichmüller space

The proof of Thurston's Theorem for critically finite branched coverings used a distance-non-increasing map $\tau : \mathcal{T}(X(f_0)) \to \mathcal{T}(X(f_0))$ known as the *Thurston pullback*. For a suitable integer k, τ^k is a uniform contraction on any set $x : d(x, \tau(x) < M\}$. So there is a unique fixed point.

The map τ is given by $\tau([\varphi]) = [\psi]$ where $\varphi \circ f_0 = s \circ \psi$ where s is a holomorphic branched covering and ψ is a homeomorphism.

This equation determines s and $[\psi]$ uniquely.

The use of Teichmüller space

The proof of Thurston's Theorem for critically finite branched coverings used a distance-non-increasing map $\tau : \mathcal{T}(X(f_0)) \to \mathcal{T}(X(f_0))$ known as the *Thurston pullback*. For a suitable integer k, τ^k is a uniform contraction on any set $x : d(x, \tau(x) < M\}$. So there is a unique fixed point.

The map τ is given by $\tau([\varphi]) = [\psi]$ where $\varphi \circ f_0 = s \circ \psi$ where s is a holomorphic branched covering and ψ is a homeomorphism.

This equation determines s and $[\psi]$ uniquely.
If $\tau([\varphi]) = [\varphi]$ then $\varphi \circ f_0 \circ \varphi^{-1} = s$, and s is critically finite and Thurston equivalent to f_0. Since τ is a contraction, (s, φ) is unique, up to Möbius conjugation of s and post-composition of φ by this same Möbius transformation.

The iteration in the generalised case

The iteration in the generalised case

in this more general case, the Teichm|'uller space used is $\mathcal{T}(Y)$, where Y contains all the critical values of f_0 but may not be forward invariant.

The iteration in the generalised case

in this more general case, the Teichm|'uller space used is $\mathcal{T}(Y)$, where Y contains all the critical values of f_0 but may not be forward invariant. Only $Z \subset Y$ is forward invariant, and Z does not contain all the critical values.

The iteration in the generalised case

in this more general case, the Teichm|'uller space used is $\mathcal{T}(Y)$, where Y contains all the critical values of f_0 but may not be forward invariant. Only $Z \subset Y$ is forward invariant, and Z does not contain all the critical values.
If we use the same formula as before and define $\tau([\varphi]) = [\psi]$ then we can consider $[\psi]$ as an element of $\mathcal{T}(Z)$. But we want an iteration on $\mathcal{T}(Y)$.

The iteration in the generalised case

in this more general case, the Teichm|'uller space used is $\mathcal{T}(Y)$, where Y contains all the critical values of f_0 but may not be forward invariant. Only $Z \subset Y$ is forward invariant, and Z does not contain all the critical values.
If we use the same formula as before and define $\tau([\varphi]) = [\psi]$ then we can consider $[\psi]$ as an element of $\mathcal{T}(Z)$. But we want an iteration on $\mathcal{T}(Y)$.

There is a natural way to do this, simply by defining $\tau([\varphi])$ by the two conditions

$$\pi_Z(\tau([\varphi]) = \pi_Z([\psi])$$

The iteration in the generalised case

in this more general case, the Teichm|'uller space used is $\mathcal{T}(Y)$, where Y contains all the critical values of f_0 but may not be forward invariant. Only $Z \subset Y$ is forward invariant, and Z does not contain all the critical values.
If we use the same formula as before and define $\tau([\varphi]) = [\psi]$ then we can consider $[\psi]$ as an element of $\mathcal{T}(Z)$. But we want an iteration on $\mathcal{T}(Y)$.

There is a natural way to do this, simply by defining $\tau([\varphi])$ by the two conditions

$$\pi_Z(\tau([\varphi]) = \pi_Z([\psi])$$

$$d_Y([\varphi], \tau([\varphi])) = dZ([\varphi], \tau([\varphi])),$$

where d_Z denotes Teichmüller distance in $\mathcal{T}(Z)$.

Properties of τ: decreasing along orbits

Properties of τ: decreasing along orbits

$\tau^k : \mathcal{T}(Y) \to \mathcal{T}(Y)$ is not a global contraction for any k

Properties of τ: decreasing along orbits

$\tau^k : \mathcal{T}(Y) \to \mathcal{T}(Y)$ is not a global contraction for any k However

$$d_Y(\tau([\varphi]), \tau^2([\varphi])) \leq d_Y([\varphi], \tau([\varphi])),$$

Properties of τ: decreasing along orbits

$\tau^k : \mathcal{T}(Y) \to \mathcal{T}(Y)$ is not a global contraction for any k However

$$d_Y(\tau([\varphi]), \tau^2([\varphi])) \leq d_Y([\varphi], \tau([\varphi])),$$

and for a suitable n depending only on $\#(Y)$,

$$d_Y(\tau^n([\varphi]), \tau^{n+1}([\varphi])) < d_Y([\varphi], \tau([\varphi])).$$

Properties of τ: the fixed set

Properties of τ: the fixed set

The most important property is the characterisation of the fixed set of τ.

Properties of τ: the fixed set

The most important property is the characterisation of the fixed set of τ.

If $\tau([\varphi]) = [\varphi]$ then $s(\varphi(Z) = \varphi(Z)$ where s is the holomorphic map such that $\varphi \circ f_0 = s \circ \psi$. Also, φ and ψ are isotopic via an isotopy which is constant on Z.

Properties of τ: the fixed set

The most important property is the characterisation of the fixed set of τ.

If $\tau([\varphi]) = [\varphi]$ then $s(\varphi(Z) = \varphi(Z)$ where s is the holomorphic map such that $\varphi \circ f_0 = s \circ \psi$. Also, φ and ψ are isotopic via an isotopy which is constant on Z.

This means that the fixed set of τ is the union of all the sets $\widetilde{V} \subset \mathcal{T}(Y)$, where V runs over the components of rational maps in B.

Definition of the sequence x_n

The sequence $x_n = x_n(x_1)$ is then obtained by defining x_{n+1} to be a modification of $\tau(x_n)$.