1 Introduction

These notes are from a minicourse by Matt Clay at The Geometry of Outer Space: Investigated through its analogy with Teichmueller space, a conference held at Aix-Marseille Universite June 24,2013- June 30, 2013. They are aimed at beginning graduate students with an interest in outer space and intersection numbers. Some knowledge is assumed: the reader should know the definitions of $\mathbb{R}-$trees, outer space, isometric actions, simple closed curves on surfaces (these can all be seen in the notes from Thierry Coulbois' course at `http://www.latp.univ-mrs.fr/~catherine.pfaff/Notes.html`), as well as hyperbolic metrics, geodesics, and splittings (HNN extensions, amalgamated free products).

The figures are scanned from Matt's handwritten notes for the course.

Goal and Outline

Given a finitely generated group G and two $\mathbb{R}$-trees T_0, T_1 equipped with isometric actions, we want a space where we can "view" both actions. One such space is $T_0 \times T_1$, where G acts diagonally, but this space is "too big". Instead, we define the *Guirardel core*, the smallest subset $\mathcal{C} \subset T_0 \times T_1$ that carries the actions: that is, there are two foliations, such that the vertical is $\mathcal{C} \cap \{x_0\} \times T_1$, and the horizontal is $\mathcal{C} \cap T_0 \times \{x_1\}$, and T_0 is the leaf space of the vertical foliation while the T_1 is the leaf space of the horizontal foliation.

Outline:

- Definition of Guirardel Core
- Properties of the Core
- Examples
- Connections with geometric intersection number
- Applications

2 Definition of Guirardel Core

Motivating example

Let α, β be simple closed curves on a closed surface S. Assume that α, β *fill* S (that is, $S - \{\alpha \cup \beta\}$ is a disjoint union of discs, or equivalently, for all simple closed curves γ, $i(\alpha, \gamma) + i(\beta, \gamma) > 0$).

These give $\mathbb{Z}-$splittings on $\pi_1(S)$. If the curves are separating, then we have an amalgamated free product and $\pi_1(S) \cong \pi_1(S_0) *_{\mathbb{Z}} \pi_1(S_1)$. If the curves are nonseparating, then the group is an HNN extension and $\pi_1(S) \cong \pi_1(S_0)*_{\mathbb{Z}}$.

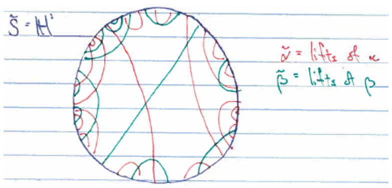

Bass-Serre theory gives us the dual trees T_α, T_β. We can concretely describe these trees using the universal cover $\tilde{S}$, and proceed to do so here. The figure above shows us what we are taking the dual of: the lifts of α, β.

First fix a hyperbolic metric on S, and realize α and β as geodesics (remember geodesics are unique in hyperbolic space). Lift these to $\tilde{S}$. The vertices of T_α are the complementary components of $\tilde{\alpha}$ in $\tilde{S}$, that is, $V(T_\alpha) = \tilde{S} - \tilde{\alpha}$. The edges of T_α correspond to lifts of α, with the endpoints of $\tilde{\alpha_0}$ being the components whose closure contains that lift. We build a similar graph for T_β.

Next, we build equivariant maps $f_\alpha : \tilde{S} \to T_\alpha, f_\beta : \tilde{S} \to T_\beta$. Pick some small $\epsilon > 0$ so that $\epsilon-$neighborhoods of lifts of α do not intersect. Then f_α will map $(-\epsilon, \epsilon) \times \tilde{\alpha_0} \to (-\epsilon, \epsilon)$, while complementary components of these $\epsilon-$neighborhoods will map to their corresponding vertices. Similar for β.

Define $f = (f_\alpha, f_\beta) : \tilde{S} \to T_\alpha \times T_\beta$, by $f(x) = (f_\alpha(x), f_\beta(x))$, and homotope f so that it is a homeomorphism onto its image.

This allows us to define the ***core*** of f: $\mathcal{C} = f(\tilde{S}) \subset T_\alpha \times T_\beta$. Note a few properties of $\mathcal{C}$:

- $\mathcal{C}$ is $\pi_1 S$-invariant.
- $\mathcal{C}$ is closed and connected.

- $\mathcal{C}$ is fiberwise convex: $\mathcal{C} \cap \{x_0\} \times T_\beta$ and $\mathcal{C} \cap T_\alpha \times \{x_1\}$ are connected if and only if they are convex. In fact, $\mathcal{C} \cap \{x_0\} \times T_\beta = \{x_0\} \times T_{\text{Stab}(x_0)}$, the minimal subtree of the action of the stabilizer of x_0 on T_β. By minimal subtree, we mean the smallest subtree of T_β left invariant by the action of $Stab(x_0)$.

We explain the third property in more detail. If x_0 is the midpoint of an edge in T_α, then its preimage under f_α is a particular lift $\tilde{\alpha_0}$ inside $\tilde{S}$. This lift crosses infinitely many lifts of $\tilde{\beta}$. Each of these lifts become edges in T_β, so $f_\beta(f_\alpha^{-1}(x_0))$ is the line formed by these edges (this is a line since there is one component between adjacent lifts, which is a vertex in the line). The third property says that this line is the axis for the subgroup of hyperbolic isometries generated by $Stab(x_0)$, and hence its minimal subtree (since, we recall, hyperbolic isometries act on their axes by translation and hence leave the axes invariant).

OK so we have the line, now why would it be the axis? Since α is a simple closed curve, $[\alpha]$ represents a conjugacy class in $\pi_1 S$. Elements in this conjugacy class correspond to different lifts $\tilde{\alpha}$. So if an element translates along a fixed axis, it fixes the particular element of the conjugacy class $\tilde{\alpha_0}$. So it is in the stabilizer $Stab(x_0)$. And if an element is in the stabilizer, then it fixes the conjugacy class, and so fixes the lift, and so must act by translations on it.

If x_0 is a point on an edge in T_α, not necessarily the midpoint, then its preimage in $\tilde{S}$ is a line at constant distance from a particular lift $\tilde{\alpha_0}$, and the same argument holds.

If x_0 is a vertex in T_α, then its preimage under f_α^{-1} is a neighborhood of $\tilde{S}$. Infinitely many lifts $\tilde{\beta}$ intersect this neighborhood, and so $f_\beta(f_\alpha^{-1}(x_0))$ is

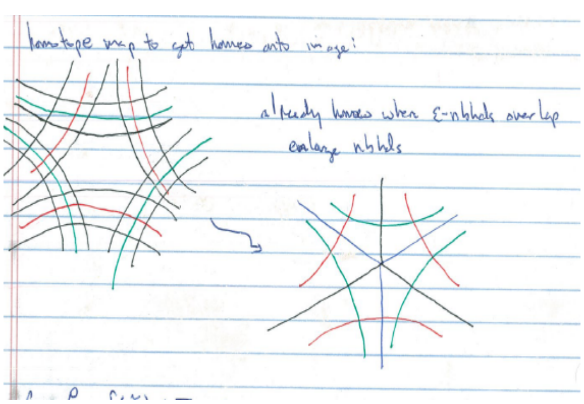

an infinite tree, and the claim is that this tree is the smallest invariant tree under the action of $Stab(x_0)$.

Exercise: Show that $\mathcal{C}$ is the smallest (with respect to inclusion) subset of $T_\alpha \times T_\beta$ that satisfies these properties.

Each square in $\mathcal{C}/\pi_1(S)$ corresponds to an edge $e_\alpha \times e_\beta \subset T_\alpha \times T_\beta$, and the number of squares is the geometric intersection number of α and β, that is $i(\alpha, \beta)$.

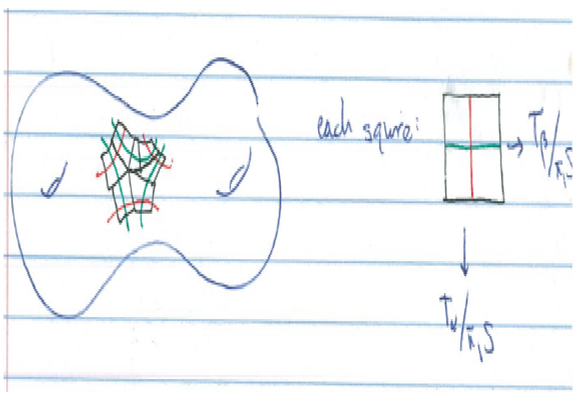

Exercise: $i(\alpha, \gamma) = \|\gamma\|_{T_\alpha}$ for any simple closed curve γ. That is, the intersection number for α and γ is the translation length of γ in the tree T_α. Note that this is well-defined because intersection number is defined up to conjugacy class.

Definitions

Given that T_0, T_1 are $\mathbb{R}-$trees, a *direction* $\delta \subset T_0$ based at $x_0 \in T_0$ is a complementary component of x_0 in T_0. A *quadrant* in $T_0 \times T_1$ is a product of directions. If you fix a basepoint $* = (*_0, *_1) \in T_0 \times T_1$, we say a quadrant $Q \subset T_0 \times T_1$ is *heavy* if there exists a sequence $\{g_k\} \subset G$ such that $g_k.* \in Q\ \forall k$ and $d_{T_i}(*_i, g_k.*_i) \to \infty$ as $k \to \infty$. A quadrant is *light* if it is not heavy. Note that this definition is independent of our choice of basepoint. Intuitively, a quadrant is heavy if there is a sequence of group elements that translates some basepoint arbitrarily deep in both directions.

Finally, the *Guirardel Core* is defined as

$$\mathcal{C} = \mathcal{C}(T_0 \times T_1) = T_0 \times T_1 - \cup_I Q$$

where I is all of the light quadrants. A few remarks:

- The core is G-invariant (since heavyness and lightness are).
- The core is closed
- $\mathcal{C}$ is fiberwise convex, as $\mathcal{C} \times \{x_0\} \cap T_1, \mathcal{C} \cap T_0 \times \{x_1\}$ are connected.

We define the *intersection number* of two trees by $i(T_0, T_1) = vol(\mathcal{C}/G)$. This is not necessarily finite.

If T_0, T_1 are simplicial and we assign 1 to all of the edge lengths, then $i(T_0, T_1)$ is the number of squares in the quotient.

Example Consider the action of $\mathbb{Z}^2$ on $\mathbb{R}$, so we have both $T_0, T_1 = \mathbb{R}$, with the action on T_0 given by $a(x) = x + 1, b(x) = x$, and the swapped action on T_1, so $a(x) = x, b(x) = x+1$. In this case, every quadrant is heavy, and $\mathcal{C} = T_0 \times T_1$. Notice also that $i(T_0, T_1) = 1$.

Example Next, consider the action by $a(x) = x + 1, b(x) = x + r$, where r is some positive real number. Then, again $T_0 \times T_1$ is the core. Here, $i(T_0, T_1) = r$.

3 Properties of the core

Motivating example

Suppose α, β are simple closed curves on S Now suppose that they have lifts $e_{\alpha_0}, e_{\beta_0}$ that correspond to edges of T_α, T_β, and the lifts are disjoint in $\tilde{S}$. The product of these in $T_\alpha \times T_\beta$ is a square. We claim: $e_{\alpha_0} + e_{\beta_0}$ is not in $\mathcal{C}(T_\alpha \times T_\beta)$.

Since the lifts are disjoint, we can choose a direction δ_{α_0} that begins at the initial vertex of the edge e_{α_0} and contains the edge. We can do a similar process for δ_{β_0}, so the directions are pointing “away” from each other. Pick the quadrant $Q = \delta_{\alpha_0} \times \delta_{\beta_0}$, and note that $e_{\alpha_0} + e_{\beta_0} \in Q$. To show that this square is not in the core, we need only show that Q is a light quadrant.

Pick some basepoint corresponding to a component of $\tilde{S}$ that lies between the disjoint $e_{\beta_0}, e_{\alpha_0}$. There can be no sequence of elements sending this basepoint arbitrarily deep in both directions. So Q is light.

Properties

- $\mathcal{C}(T_0 \times T_1) \neq \emptyset$ if one of T_i is *irreducible*. This means that there exist $g, h \in G$ that act hyperbolically and the axes $Ax(g) \cap Ax(h)$ is compact.

That is, neither are dihedral nor abelian.

A nonexample: Suppose $\mathbb{Z}^2$ acts on $\mathbb{R}$ by $a(x) = x+1, b(x) = x$. Then $\mathcal{C}(T_0 \times T_0) = \emptyset$, since every point is in a light quadrant.

Another nonexample: If T is simplicial and minimal, and every vertex has valence ≥ 3, with any finitely generated group acting on it. Then $\mathcal{C}(T \times T) = \{(v,v) : v \in V(T)\}$ is disconnected.

Proof of nonexample: a quadrant $\delta_0 \times \delta_1$ is light if and only if $\delta_0 \cap \delta_1 \subset e$, where e is an edge: if Q is not contained in an edge, then it contains an infinite ray, and so can send something off deeply. And if Q is heavy, then it can't be contained in an edge.

- $\mathcal{C}(T_0 \times T_1)$ is disconnected if and only if T_0 and T_1 are *refinements* of the same tree T. T_i is a refinement of T if there is a map $T_i \to T$ such that the preimages of points are connected in T_i.

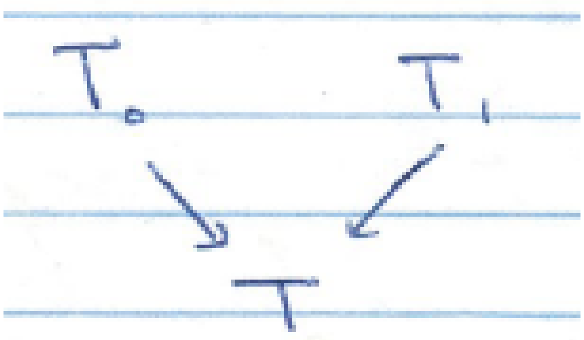

- We can always add "diagonal" edges to make our core connected (we call this the "augmented core"). These will not affect intersection numbers.

- If the lifts α, β do not intersect, then $e_\alpha \times e_\beta \notin \mathcal{C}(T_\alpha \times T_\beta)$, that is, there is a light quadrant containing this square. Let δ_α be the direction containing e_α in $T_\alpha - i(e_\alpha)$, and similar for β. Then $\delta_\alpha \times \delta_\beta$ is light.

Example in Outer Space Consider $T_0, T_1 \in CV_N$. Fix a function $f : T_0 \to T_1$ that sends vertices to vertices, is linear on edges, is F_N-equivariant, and is such that f has more than one *gate* at every vertex: that is, not all the images of edges adjacent to a vertex share a common initial segment. Then f is a quasi-isometry and induces a F_N−equivariant homeomorphism $\partial f : \partial T_0 \to \partial T_1$.

Lemma 1. *The quadrant $\delta_0 \times \delta_1$ is heavy if and only if $\partial f(\partial\delta_0) \cap \partial\delta_1 \neq \emptyset$. Here, $\partial\delta_i \subset \partial T_i$ are rays that intersect the directions δ_i in a ray.*

Proof. ($\Rightarrow$) Suppose $\delta_0 \times \delta_1$ is heavy, so there exists a sequence $\{g_k\} \subset F_N$ such that $g_k(*_0, f(*_0)) \in \delta_0 \times \delta_1$ and $d_{T_0}(*_0, g_k.*_0), d_{T_1}(f(*_0), g_k.f(*_0)) \to \infty$ as $k \to \infty$, for some basepoint $*_0 \in T_0$.

Pass to a subsequence so that $g_k.*_0 \to X \in \partial\delta_0$. Then pass to another subsequence so we have the same thing in the other direction: $g_k.f(*_0) \to Y \in \partial\delta_1$. By continuity, $\delta f(X) = Y$, and so $\partial f(\partial\delta_0) \cap \partial\delta_1 \neq \emptyset$.

($\Leftarrow$) Suppose $\partial f(\partial\delta_0) \cap \partial\delta_1 \neq \emptyset$. Pick a ray $R \in \partial\delta_0$ so that $f(R) \in \partial\delta_1$. Choose a sequence $\{g_k\} \subset F_N$ that translates a point $*_0$ along R. Since f is a quasi-isometry, this sequence also translates $f(*_0)$ along the δ_1 direction (since $R \subset \delta_1$). □

A bit of notation: if $e \subset T$ is an oriented edge, we say that $[e]$ is the set of rays that cross e in the same orientation. That is, $[e]$ is the set of infinite rays that start at $i(e)$, continue through $t(e)$, and go on. This is a *one-sided cylinder*. Note that $[e], [\bar{e}]$ decompose ∂T into two parts.

Lemma 2. *$e_0 \times e_1 \in \mathcal{C}(T_0 \times T_1)$ if and only if each of the following four subsets is nonempty:*

- $\partial f([e_0]) \cap [e_1]$
- $\partial f([\bar{e}_0]) \cap [e_1]$

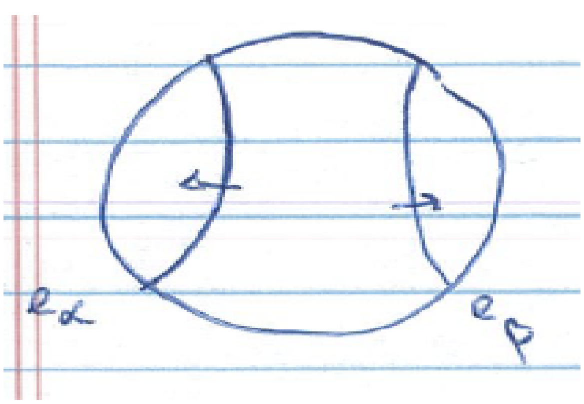

- $\partial f([\bar{e_0}]) \cap [\bar{e_1}]$
- $\partial f([e_0]) \cap [\bar{e_1}]$

Proof. Note that for any direction δ_i that contains e_i, $\partial\delta_i$ contains either $[e]$ or $[\bar{e}]$. Then apply the previous lemma. □

We have an alternative way of saying the lemma. $e_0 \times e_1 \in \mathcal{C}$ if and only there exist geodesics $p_+, p_- : \mathbb{R} \to T_0$ such that in T_0, both geodesics contain e_0 while, after tightening their images to geodesics in T_1, e_1 separates them (so one is contained in $[e_1]$, and the other in $[\bar{e_1}]$ This lemma helps us build the core, allowing us to determine when squares are contained in the core.

Definition. Given an edge $e_1 \in T_1$, the *slice* of the core $\mathcal{C}(T_0 \times T_1)$ above e_1 is $\mathcal{C}_{e_1} = \{e_0 \in T_0 : e_0 \times e_1 \in \mathcal{C}(T_0 \times T_1)\}$. Note that $\mathcal{C}_{e_1}$ is a subtree of T_0.

Since F_N acts freely on edges, $\mathcal{C}_{e_1} \times \{x_1\}$, where x_1 is the midpoint of e_1, embeds in $\mathcal{C}/F_N$.

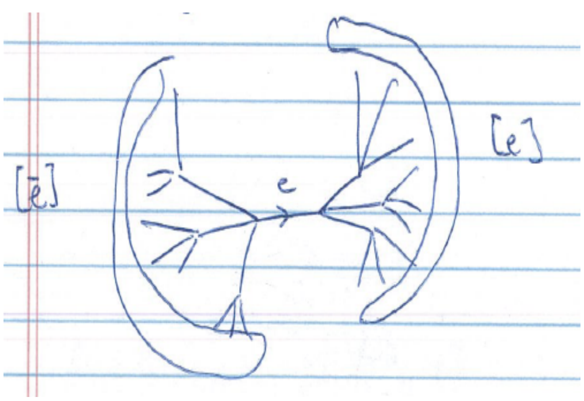

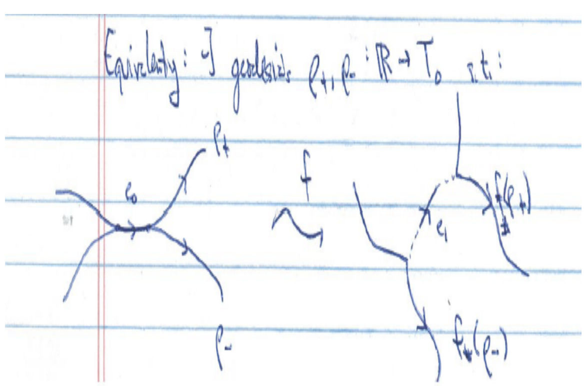

geodesics.png

So $i(T_0, T_1) = Vol(\mathcal{C}/F_N) = \sum_{e \in T_1/F_N} l_{T_1}(e) vol_{T_0}(\mathcal{C}_{e_1})$.

We want to understand the map on the boundaries: $\partial f : \partial T_0 \to \partial T_1$. Fix a basepoint $*_0 \in T_0$. Orient edges of the tree to point away from $*_0$ and $f(*_0)$. Let's describe $(\partial f)^{-1}([e_1])$, where e_1 is an infinite ray in T_1.

Choosing some point $x_1 \in e_1$, we have that $f^{-1}(x_1)$ is a finite set of points in T_0 [f sends vertices to vertices, is F_N invariant]. Say these are $\{p_1, \ldots, p_k\}$, and each point is in a unique edge $q_i \in T_0$. We give these points signs: if $f(q_i)$ crosses e_1, it has a positive sign and is in I^+, otherwise, it crosses $\bar{e}_1$ and has a negative sign and is in I^-..

Then we define $(\partial f)^{-1}([e_1]) = \sum_{i \in I^+} [q_i] - \sum_{j \in I^-} [q_j]$.

Examples

Example

$f : R_2 \to R_2$ by $f(a) = a, f(b) = a^2 b$. Let $T_0 = \tilde{R_2}, T_1 = \tilde{R_2} f$. Looking at T_0 compared to T_1, the map f shears $\tilde{R_2}$ two units to the right: the a "coordinate" is fixed, while the b "coordinate" moves two branches to the right. In the notation defined above, $(\partial f)^{-1}([a]) = [a]+[b]+[a^{-1}b]$. In words, $(\partial f)^{-1}$ maps infinite words that start with a homeomorphically to infinite words that start with a, b or $a^{-1}b$, and no other initial segments result in a word that begins with a. *Example*

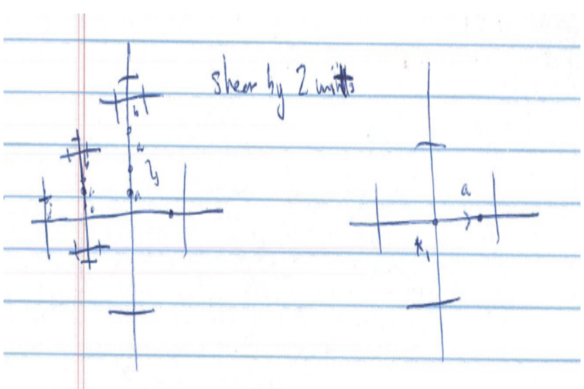

Now suppose $f : R_2 \to R_2$ by $f(a) = ab, f(b) = bab$. Let's try to find $(\partial f)^{-1}([b])$.

Fix a path α from a basepoint in R_2 to the midpoint x_1 of the b edge. $f^{-1}(x_1) = \{p_1, p_2, p_3\}$. Here, p_1, p_2 are points along the original b edge, since

$f(b) = bab$, so assuming constant speed these points land us at x_1, and p_3 is a point on the original a edge, since $p(a) = ab$.

Now we need to find a path from the original basepoint to p_i whose image is homotopic (rel endpoints) to α. The problem being that another path will be mapped (up to homotopy)to a different lift rather than our specified α.

- For p_1, we have b_1, a path along b from the basepoint to p_1, which is positive.
- For p_2, we have $ba^{-1}b^{-1}b_1b_2 = ba^{-1}b_3^{-1}$, which is negative.

  Note here that b_1b_2 does not work, so this is different from the p_1 case. For $f(b_1b_2)$ maps to ba, which is not homotopic to α (but is homotopic to $ba\alpha$, the incorrect lift). So we have to append a word before in order to map to α and get rid of the extra ba- we need something that maps to $a^{-1}b^{-1}$. Inspection gives that $f^{-1}(a^{-1}b^{-1}) = ba^{-1}b^{-1}$, so we toss that on at the beginning and get $ba^{-1}b^{-1}b_1b_2$. Tracing this through the graph, we end with negative orientation b_3^{-1}m and so we end up with a negative term.
- And for p_3, we have $ba^{-1}a_2^{-1}$, which is negative. Our reasoning for p_3 is similar to that of p_2 and is left as an exercise for the reader.

For the actual preimage, we need to consider the full edges inside the span of the three preimages. So we "round up" from our half edges, and have $(\partial f)^{-1}([b]) = [b] - [ba^{-1}b^{-1}] - [ba^{-2}]$.

Now let's find the slices. We have $\mathcal{C}_{e_1} = \{e_0 \in T_0 : e_0 \times e_1 \subset \mathcal{C}\}$. We claim that $\mathcal{C}_b = ba^{-1}$. We use the lemma below, proven by Bestvina, Behrstock, and Clay. We do not prove it.

Lemma 3. *If $T_0, T_1 \in CV_N, e_1$ an edge of T_1, then $\mathcal{C}_{e_1}$ consists of the interior edges of the span of $f^{-1}(x_1)$, where x_1 is the midpoint of e_1.*

So looking at our preimage of x_0, we had $(\partial f)^{-1}([b]) = [b] - [ba^{-1}b^{-1}] - [ba^{-2}]$. If we draw this on a infinite tree, we see that the interior edges are ba^{-1}. So $\mathcal{C}_b = ba^{-1} \in \tilde{R_2}$, as desired.

Now let's build the core $\mathcal{C}(\tilde{R_2} \times \tilde{R_2}f)$ for the edge $b \in \tilde{R_2}f$.

We found the slice $\mathcal{C}_b$ above, so now let's find the slice $\mathcal{C}_a$. We calculate in a similar way to before that $(\partial f)^{-1}([a]) = [a] - [ab^{-1}]$. Then $\mathcal{C}_a$ is just the vertex $a \in \tilde{R_2}$, from our lemma above.

We can draw $\mathcal{C}(\tilde{R}_2 \times \tilde{R}_2 f) \subset \tilde{R}_2 \times \tilde{R}_2 f$: we have a box for $\mathcal{C}_b$, an edge for $\mathcal{C}_a$, and another two edges to connect the fibers. To find the core, we need to mod out by the action of F_N, which acts by multiplication by a, b on the left in the first factor, and by $f(a), f(b)$ on the right in the second factor.

For example, the element $a^2b^{-1} \in F_2$ sends the point $(ba^{-1}, *)$ to $(a^2b^{-1}ba^{-1}, f(a^2)f(b^{-1})) = (a, ababb^{-1}a^{-1}b^{-1}) = (a, a)$. We have another identification of two edges, so our final $\mathcal{C}/F_2$ looks like:

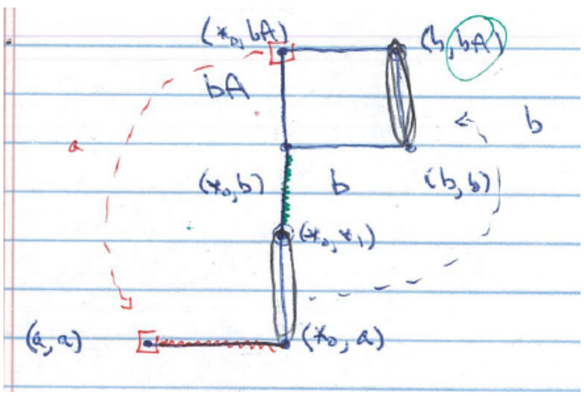

Connection to Geometric Intersection Number

Note that $i : CV_N \times CV_N \to \mathbb{R}_{\geq 0}$ is continuous, however, the natural extension $i : C\bar{V}_N \times C\bar{V}_N \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ is not.

For a nonexample, suppose $f : R_2 \to R_2$ by $a \mapsto a, b \mapsto ab$. The slice $\mathcal{C}_b$ is empty, while the slice $\mathcal{C}_a$ consists of $n-1$ edges, the preimage of a shear by n units to the right. If we consider the intersection number of these universal covers, we have $i(\tilde{R}_2, \tilde{R}_2 f^n) = n - 1$. Note here that the core is not connected. If we want to get to the boundary of outer space, we must scale $\frac{1}{n}\tilde{R}_2 f^n$, and take the $\lim_{n\to\infty}$, which will lead to a $T \in C\bar{V}_N$.

Now consider the quotient T/F_2: it is an HNN extension of $< a >$ with the stable letter b. Why? Consider the length of the images of a, b as $n \to \infty$ under the iterated $\frac{1}{n}f^n$. Since $a \mapsto a$, the length of its image goes to 0, while $b \mapsto ab$ means $f^n(b) = aaa \cdot ab$ and so the length of its image under $\frac{1}{n}f^n$ goes to 1.

Since $\tilde{R}_2 \to T$ is a refinement, $i(\tilde{R}_2, T) = 0$, while $i(\tilde{R}_2, \frac{1}{n}\tilde{R}_2 f^n) \to 1$. So i is not continuous.

Theorem 1. *$i(T_0, T_1) = 0$ if and only if there exists a common refinement $T \to T_i$, where point preimages are connected.*

We can read this theorem as saying that if the core has no area, then it is the common refinement T, and vice-versa.

Whitehead Model, Sphere Systems Definitions

Let's think of a 3- manifold M_N as the connected sum of N copies of $S^1 \times S^2$, or as the double of a handlebody of genus N, glued along their boundaries by the identity map. By Van Kampen's theorem, $\pi_1(M_N) = F_N$. Yet another way to think of M_N is by using *Whitehead's model*, where we can think of it as S^3 with the interior of $2N$ disjoint balls cut out of it. Match each ball with a mate, and then identify the boundaries of these balls (reversing orientation in the identification map).

If we collapse these balls, we get the *Whitehead graph*. Suppose $A = \{a_1, \ldots, a_n\}$ is an alphabet, and g a cyclically reduced word $x_1 \ldots x_k$. Then $Wh_A(g)$ is a graph with $2n$ vertices, labeled $a_i^{\pm}$, and with an edge from x^{-1} to y for each subword xy contained in g, considered cylically. For instance, ab^{-1} would have an edge from a^+ to b^+ and another from b^- to a^-.

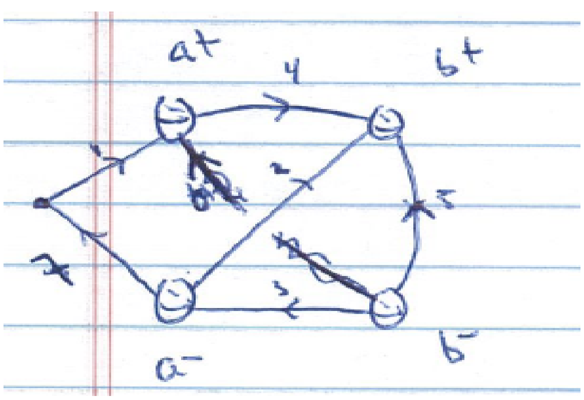

The Whitehead graph has many interesting applications, and we give an example of one here without proof:

Theorem 2. *If g is primitive, then $Wh_A(g)$ is disconnected or has a cut vertex.*

We can define a *Dehn twist* about a 2-sphere (the boundary of these matched balls) by taking a neighborhood $S^2 \times I$ of the sphere, and then rotating the sphere $S^2 \times \{t\}$ by $2\pi t$. Note that any curve that meets the sphere, then twists, can be homotoped back to its original state. So Dehn twists act trivially on $\pi_1(M_N)$. The next theorem says these twists are all the diffeomorphisms that do so.

Theorem 3 (Laudenbach). *The natural map from the mapping class group to $Out(F_N)$, is surjective and has a finite kernel, which is a direct sum of N copies of $\mathbb{Z}_2$, generated by Dehn twists about the 2-pheres in M_N.*

Recall that we can think of $Map(M_N)$ as $\pi_0(\text{Diffeo}^+(M_N))$.

Now a *sphere system* in M_N is a collection of disjoint embedded 2-spheres, no two of which are non-isotopic, and none of which bound a 3-ball. Let $S(M_N)$ denote the isotopy classes of sphere systems. There is a bijection between $S(M_N)$ and automorphisms of F_N acting on simplicial trees. We have a similar lemma.

A *free splitting* of F_N is a minimal simplicial action of F_N on a tree, where edge stabilizers are trivial.

Lemma 4 (Aramoyena-Souto). *There exists a bijection between $S(M_N)$ and free splittings $FS(F_N)$*

Proof. Given a sphere system Σ, consider the lifts $\tilde{\Sigma}$ in $\tilde{M}_N$. Build the dual tree T_Σ as we did earlier, in our example with simple closed curves. So vertices are the components of $\tilde{M}_N - \tilde{\Sigma}$, while edges are the lifts $\tilde{\Sigma}$, with endpoints where you think they are. This gives a free splitting of F_N.

For the other direction, given an action of F_N on T which is minimal, simplicial, and has trivial edge stabilizers, we want to build a sphere system with an equivariant map $f : \tilde{M}_N \to T$. If we consider $f^{-1}(x_0)$, where x_0 is the midpoint of an edge, and then do surgeries on this graph so that these preimages are spheres, we have the desired sphere system. □

Intersection number and Sphere systems

Given two sphere systems Σ_0, Σ_1, define their *intersection number* $i(\Sigma_0, \Sigma_1)$ as the minimum number of components of $\Sigma_0' \cap \Sigma_1'$, where Σ_i' is a representative of the isotopy class Σ_i.

This next proposition is the main goal of this section.

Proposition 1 (Horbey). $i(T_{\Sigma_0}, T_{\Sigma_1}) = i(\Sigma_0, \Sigma_1)$.

To prove this proposition, we need a quick definition. We say a sphere system is in *Hatcher normal form* if, for any component $P \subset M_N - \Sigma_1$, each component of $\Sigma_0 \cap P$ is one of the following:

- A disk separating 2 spheres in ∂P
- A cylinder with boundary consisting of two spheres in ∂P
- A pair of pants with the boundary contained in three spheres

Proof. Combine our two maps $f_{\Sigma_i} : \tilde{M_N} \to T_{\Sigma_i}$ to form $f = f(f_{\Sigma_0}, f_{\Sigma}) : \tilde{M_N} \to T_{\Sigma_0} \times T_{\Sigma_1}$.

Under this map, squares in $f(\tilde{M_N})$ correspond to an intersection circle in $S_0 \cap S_1$, where S_i are spheres in $\tilde{\Sigma_i}$. Since the core is minimal, this set of squares, $f(\tilde{M_N})$m contains the core. Then take the quotient of this by F_N, and we have that $f(\tilde{M_N})/F_N \supset \mathcal{C}/F_N$. So by definition, $i(T_{\Sigma_0}, T_{\Sigma_1}) \leq i(\Sigma_0, \Sigma_1)$.

For the other direction, it is a fact that Σ_0 and Σ_1 intersect minimally when they are in Hatcher normal form. Assume that Σ_i are maximal, so their complementary components are 3-spheres with the interiors of disjoint 3-balls removed. So if $S_i \in \tilde{\Sigma_i}$, $S_0 \cap S_1$ has at most one circle in $\tilde{M_N}$. If this circle exists, then the four regions of $\tilde{M_N} - S_0 \cap S_1$ are all unbounded. Then the square $e_{S_0} \times e_{S_1} \subset \mathcal{C}$, as every quadrant containing it is heavy.

□

Applications

Theorem 4 (Guirardel). *If $\phi \in Out(F_N)$ is an iwip and its attracting and repelling trees $T_\pm$ are geometric, then ϕ has a psseudo-Anosov on a surface representation. In this case, the core $\mathcal{C} \subset T_+ \times T_-$, and the quotient $\mathcal{C}/F_N$ is a surface with marked points for the boundary.*

Theorem 5 (Handel-Mosher, Behrstock-Bestvina-Clay). *If $\phi \in Out(F_N)$ is fully irreducible, T_+ is geometric, and T_- is not, then $\lambda_\rho > \lambda_{\rho-1}$, where these represent the growth rates.*

The proof of this second theorem comes from understanding the relationship between intersection number and growth rates.