Basic Ideas about Laminations

Daniel Mathews

March 6, 2007

Contents

1 Preface **2**

2 Prerequisites **2**

3 Introduction **2**
3.1 So what is the definition, then? 3
3.2 Some examples 3
3.3 A quibble, and unique decomposition 4

4 What do laminations look like? **5**
4.1 Laminations are not jerky... 5
4.2 ... but not very smooth, either. 6
4.3 And hence they have unique decomposition. 7

5 Closed in the set of closed sets **7**
5.1 Where does $\mathcal{L}$ naturally lie? 7
5.2 $\mathcal{L}$ is closed, the world rejoices. 8

6 Isolated and boundary leaves **9**
6.1 Derivation: Destroying isolated leaves 9
6.2 Boundary leaves and why they are everywhere 10

7 Principal regions and What They Look Like **12**
7.1 They are very finite... 12
7.2 ... can be lifted to the hyperbolic plane 13
7.3 ... look nice when there aren't any closed leaves... 13
7.4 ... without any annoying closed leaves, are finite at infinity... . . 14
7.5 ... and sometimes even when there *are* closed (isolated) leaves... . 15
7.6 ... and take up all the space... 16
7.7 ... while the rest is rather weird. 16

8 Derivations and their properties **17**
8.1 Stability, Perfection and Their Consequences 17
8.2 How the proof goes. 17

9 Surface Automorphisms **19**
9.1 Telling them apart 20
9.2 Surface automorphisms meet laminations: The invariant lamination 21
9.3 What the invariant lamination means 22
9.4 Attraction and Repulsion 24
9.5 The Proof I: How good is our invariant lamination? 26
9.6 The Proof II: Alternating attraction and repulsion 27
9.7 The Proof III: The End 28

1 Preface

The story of laminations is quite amazing. It throws around a whole lot of interesting ideas from geometry, topology, and analysis.

I read Casson and Bleiler's book "Automorphisms of Surfaces after Nielsen and Thurston". I liked it. All the story is there in their theorems. But stories can also be told directly. That's the idea here.

2 Prerequisites

I'm assuming you know about basic topology. I assume you know the classification of surfaces. And I'm assuming you're familiar with the basics of hyperbolic geometry: for instance, the properties of the hyperbolic plane; the types of hyperbolic isometries; and so on.

3 Introduction

We are dealing with laminations on surfaces. So what is a lamination? You might have heard something like:

> A laminate is a material constructed by uniting two or more layers of material together. The process of creating a laminate is lamination, which usually refers to sandwiching something between layers of plastic and sealing them with heat and/or pressure.[1]

Well, forget everything you thought you knew about *this* kind of lamination! This is completely different. Actually, if you have an interest in composite materials, you may want to remember this more well-known version of lamination, but be aware that it's entirely unrelated.

No, a lamination on a surface is something completely different.

[1] From the Wikipedia entry on "Laminate", http://en.wikipedia.org/wiki/Laminate.

3.1 So what is the definition, then?

Before we give the definition, however, let's be clear what page we're on. Our surface is not just any surface. First, it's a *closed* surface, that is, it's compact and has no boundary. No punctures, no cusps, not infinite genus, nothing like that. Second, it's an *oriented* surface. No mobius strips, no crosscaps, no klein bottles or anything like that. And third, it's a *geometric* surface, in fact we require it to have a very particular geometry, *hyperbolic* geometry. This means it has a Riemannian metric of constant curvature -1, and implies that the surface has negative Euler characteristic, hence has genus at least 2.

Thus, the topology of the surface we're talking about is completely specified by the genus, which can be any positive integer 2 or greater. (On a particular surface, however, there are many possible hyperbolic metrics — in fact, a $(6g-6)$-dimensional space of them.)

Now, a lamination will just be a bunch of geodesics. First, these will be *complete*, in that they don't stop. They can be closed loops of finite length; or if they don't close up, they will be of infinite length. No endpoints! Second, no intersections will be allowed. So the geodesics must be *simple*, that is, have no self-intersections. And they must be *disjoint*: they can't intersect each other. The only other requirement — which really is the one that induces all the structure — is that together this bunch of geodesics forms a *closed set.* Every limit point of the set is in the set.

That's the definition.

Definition 3.1 *Let S be a closed oriented hyperbolic surface. A* lamination *on S is a closed subset L of S which is a union of disjoint complete simple geodesics.*

3.2 Some examples

There are some easy examples of laminations. For instance, there is the empty lamination. A single simple closed geodesic forms a closed set, hence is a lamination. A finite set of disjoint simple closed geodesics will be closed, and hence taken together form a lamination. A bit more tricky is an example where we take a non-closed simple geodesic which *spirals* towards a simple closed geodesic. The spiralling geodesic does not form a lamination, because it's not a closed set — its limit set includes the simple closed geodesic it's spiralling towards. But if we take the spiralling geodesic, *and* the simple closed geodesic it's spiralling towards, *together*, then we have a lamination.

Can you think of any other examples?

Well, it turns out that the types of phenomena I just mentioned — simple closed curves, and non-closed curves which spiral towards closed curves — are actually quite unusual. Laminations may have properties such as: no closed geodesics; every geodesic is dense; no isolated geodesics. This is something we haven't defined, but it probably has the definition you think it does.

These generic examples are difficult to draw a picture of, unfortunately. But there's another way to think about them, which can be quite useful. This way relies on the fact that *the set of all laminations on S is closed.* What does "closed" mean here? What's the topology on this space? Where does it naturally lie? We will answer these questions shortly. But this fact means that examples of laminations can be given as *limits* of laminations we can draw. For instance, consider a sequence of laminations L_n, each consisting of a single simple closed geodesic, but where these geodesics get more and more complicated, and as a closed set "approaches" some subset L_∞ of S. Once we define what all this means, our fact means that L_∞ will be a lamination; and that may not have been an obvious thing to see. And since it's a limit of increasingly complicated laminations, L_∞ may have more exotic properties than its relatively uncomplicated progenitors.

3.3 A quibble, and unique decomposition

But first: there's a question with the definition. It might sound like a quibble, but it's important and exposes something interesting about hyperbolic geometry.

We'd like to say that a lamination L consists of geodesics γ_i, where i is in some indexing set, possibly infinite, possibly uncountable. We'd like to say that the geodesics γ_i are *the leaves* of L. Note the use of the definite article "the".

But given the lamination L, there might be different ways to decompose this closed set into geodesics. If that were the case, we would have to be careful about the different decompositions of L.

If you think about Euclidean surfaces, you'll see that this can actually happen. For instance, the whole of the Euclidean plane (i.e. $L = S = \mathbb{R}^2$) can be decomposed as a disjoint union of vertical geodesics, or as a disjoint union of horizontal geodesics, in fact lines with any single given slope. The same is true on the standard Euclidean torus $L = S = \mathbb{R}^2/\mathbb{Z}^2$: it can be decomposed as a disjoint union of geodesics of any given slope.

Thankfully, it turns out that we don't have to worry.

Proposition 3.2 *Let S be a closed oriented hyperbolic surface and let L be a lamination on S. Then L can be decomposed as a union of disjoint complete simple geodesics in one and only one way.*

This proposition tells us something about the nature of hyperbolic geometry. In some sense, you can't have geodesics too close together. If you do, the negative curvature doesn't allow them to do so within a compact surface.

The same principle seems to be at work in the following simple picture. If you take a small open ball in the hyperbolic plane, you can try to fill it up with geodesics. It's possible to do so, in fact in many different ways. Think of one way. But now look at where those geodesics go after they leave the ball and continue through the hyperbolic plane. They spread out like crazy through at least some regions — they will fill up a region of infinite area. That's going to be a problem on a compact surface. Well, maybe you might think that it's okay, because you're going to quotient the hyperbolic plane by some group of isometries to get the surface, and maybe it will all work out. But it won't, and you might like to ponder why. We'll get to the proof soon enough.

In any case, pending the proof of proposition 3.2, we will make a

Definition 3.3 *The geodesics γ_i in a lamination L are called* the leaves *of L. At each point $x \in L$, we write γ_x for the leaf of L passing through x.*

4 What do laminations look like?

Since it's hard to draw pictures of laminations, we can calibrate our intuition by deriving some properties about them. These properties have some pictorial value, which we can bear in mind for later reference.

4.1 Laminations are not jerky...

Laminations don't jump around too much. By which I mean, the geodesics involved can't change direction too sharply. That is, if you go from one point x in L to another which is close by — call it y — then the directions of the geodesics γ_x and γ_y don't change too much. We can state this as a proposition, and it's so simple that we might even prove it. Actually, it's *so* simple that you can prove it for yourself!

Proposition 4.1 *Let S be a closed oriented hyperbolic surface. Let L is a lamination which is decomposed as a union of disjoint complete simple geodesics. Let γ_x denote the geodesic passing through x. The direction of γ_x varies continuously with x.*

(Why did I not just say that γ_x is the leaf of L through x? Because I'm proving something, and I haven't proved anything about unique decomposition into geodesics yet! Hence the slightly more complicated phrasing. This fact about direction varying continuously is true for *any* decomposition of a given lamination L.)

We can make this rigorously precise, if we like, though the intuitive picture does not depend on it. At each point $x \in S$, the set of directions from S is a real projective line $\mathbb{R}P^1$, which is homeomorphic to a circle. If you put together these projective lines at each point, we have the projectivized tangent bundle of S, which we write $PT(S)$. At each $x \in S$ the fibre $PT(S)_x$ is a copy of $\mathbb{R}P^1$.

This bundle clearly restricts to $L \subset S$, and we call the restricted bundle $PT(S)|_L$. (It's different from $PT(L)$!) There is still an $\mathbb{R}P^1$ fibre above each point of L. Let $\pi : PT(S)|_L \to L$ be the projection associated with the bundle. The function assigning to each point x of L the direction of γ_x is a section of this bundle, that is a map $s : L \longrightarrow PT(S)|_L$ which takes each $x \in L$ to a point above it, $s(x) \in (PT(S)|_L)_x$; equivalently $\pi \circ s = 1_L$.

Proposition 4.1 says that the section s is continuous.

Did that general nonsense help? Probably not. The picture is simple to draw in any case.

Locally our surface is isometric to the hyperbolic plane, so we can just draw pictures there. If the direction of γ_x doesn't vary continuously with x, then you can find two points which are close but where their directions are not close. Draw the geodesics in these directions. Then they meet, so the geodesics are not disjoint, neither in the hyperbolic plane or in the quotient surface S. If you want to put in some epsilons and some hyperbolic trigonometry to completely convince yourself, go right ahead.

4.2 ... but not very smooth, either.

On the other hand, laminations do jump around a lot. They have lots of holes in them. As we argued rather unconvincingly earlier, a lamination can't fill up much of a surface. In fact it can't even fill up an arbitrarily small ball in the surface. Proving this fact, we will have shown that geodesics are "not too close together", in some sense. We will then be able to prove proposition 3.2, the uniqueness of decomposition into geodesics.

But we will start with a simpler proposition and build up to a climax: L can't be the *whole* manifold. This is an extremely weak result, in the light of what we're about to prove, but still an interesting application of the Poincaré-Hopf index theorem.

Proposition 4.2 *Let S be a closed oriented hyperbolic surface and L a lamination on S. Then L is a proper subset of S, i.e. $L \neq S$.*

(Note that hyperbolic is essential here! We have already seen examples of laminations of the Euclidean plane or Euclidean torus which fill the whole manifold.)

PROOF Take a decomposition of L into geodesics (we still don't know anything about uniqueness) and for $x \in L$ let γ_x be the geodesic through x. If $L = S$ then we consider the direction of γ_x at x. From above, proposition 4.1, this direction varies continuously. These directions thus give a line field on S. If you like, you can obtain from this a nowhere vanishing vector field on S. But in any case, this is prohibited by the Poincaré-Hopf index theorem on a surface of negative Euler characteristic. The indices of singularities must sum to the Euler characteristic, which is nonzero! ∎

And now for something stronger — a *lot* stronger — which piggy-backs on top of this rather weak proposition.

Proposition 4.3 *Let S be a closed oriented hyperbolic surface and L a lamination on it. Then L is nowhere dense in S.*

PROOF Take a decomposition of L into geodesics. Suppose to the contrary that there's an open ball U in L. Take $x_0 \in U$ and take an arc α transverse to the geodesic γ_{x_0} through x_0. Since L contains the neighbourhood U of x_0, if α is sufficiently small then $\alpha \subset U$. By proposition 4.1 the direction of the geodesics γ_x varies continuously, by replacing α with a smaller arc if necessary we can assume α is transverse to the lamination L at every point of α.

Now we have $\alpha \subset L$ which is transverse to L. This spells doom for L. Identify the universal cover $\tilde{S}$ with $\mathbb{H}^2$, and take a lift $\tilde{\alpha} \subset \mathbb{H}^2$. Define a function $\Phi : \alpha \times \mathbb{R} \longrightarrow \mathbb{H}^2$ which takes (y, t) to the point on γ_y, at (signed) distance t

from α. Want to worry about the technicalities of defining this properly in the universal cover? Then go right ahead; it works fine. We see that $\Phi(\alpha \times \mathbb{R})$ is a huge region of $\tilde{S} \cong \mathbb{H}^2$ bounded by two geodesics — namely, the geodesics in L through the endpoints of α. These must enclose a region of infinite area. This massive region of $\tilde{S}$ projects to S and all lies in L. You can fit arbitrarily large balls in this region, including ones larger than the diameter of S; such a large ball will project onto the whole of S. So $L = S$; but this is a fatal blow because it contradicts the previous proposition. ∎

This proves that geodesics really aren't very close together.

4.3 And hence they have unique decomposition.

What really mattered in the previous proof was the arc $\alpha \subset L$ transverse to the decomposition of L into geodesics. This allowed us to take the transverse geodesics — and this smooth family of disjoint geodesics had to fill up way too much of the hyperbolic plane, hence way too much (all) of S.

But unique decomposition follows is proved in exactly the same way.

For suppose we can decompose L into a union of disjoint complete simple geodesics two different ways. Take a point x_0 where the two decompositions have different geodesics through x_0. Call the first γ_{x_0}; call the second α. By shortening α sufficiently, we obtain an arc $\alpha \subset L$ which is transverse to the first decomposition of L. This spells doom for L in the same way that it did four paragraphs ago. The same proof applies, word for word.

Now we will use the phrase "the leaves of L" with a clear conscience. Being careful with it was getting a little annoying!

5 Closed in the set of closed sets

We'll now return to the rather exciting prospect mentioned earlier when discussing examples of laminations. That prospect is that we can define really complicated laminations without being able to draw them: we define them as the *limit* of simpler laminations. To do this properly, however, requires the *set of laminations* to be *closed*.

We will now answer the questions raised in this previous discussion.

As always, S is a closed oriented hyperbolic surface. Let $\mathcal{L}$ be the set of all laminations on S.

5.1 Where does $\mathcal{L}$ naturally lie?

Every lamination is a closed subset of S by definition. Hence $\mathcal{L}$ is a subset of the *set of all closed subsets of* S, which we will denote $\mathcal{P}(S)$.

Turns out $\mathcal{P}(S)$ has a nice topology. In fact, if S has a metric, then so does $\mathcal{P}(S)$. This metric is sometimes called *Hausdorff distance*. Given two closed sets $A, B \subseteq S$, we define their Hausdorff distance $d(A, B)$ by

$$d(A,B) \leq \epsilon \quad \text{if and only if} \quad A \subseteq N_\epsilon(B) \text{ and } B \subseteq N_\epsilon(A),$$

or equivalently, as

$$d(A, B) = \inf \{\epsilon \geq 0 \ : \ A \subseteq N_\epsilon(B) \text{ and } B \subseteq N_\epsilon(A)\}.$$

Here $N_\epsilon(A)$ denotes an open ϵ-neighbourhood of A.

So the distance between closed sets A, B is essentially the least amount required to expand them so that either one engulfs the other.

Clearly there's nothing special about any hyperbolic or geometric properties of S here. In fact we can define Hausdorff distance on $\mathcal{P}(X)$ for any metric space X and get a metric. If X — like S — is compact, then any two metrics on X are equivalent; hence the topology on $\mathcal{P}(X)$ doesn't depend on the metric. This is left as an exercise for the reader; it is standard fare in a course on metric spaces.

So $\mathcal{P}(S)$ is a well-defined metric space, and while the metric depends on the metric on S, the topology does not. The set $\mathcal{P}(S)$ has one more important property: *it is compact.* (In fact, $\mathcal{P}(X)$ is compact for any compact topological space X.) This is also a standard fact in metric spaces and left to the reader. But it is important. If you take any sequence of closed subsets of S, they contain a convergent subsequence, in the Hausdorff metric. And the intuitive picture of Hausdorff convergence is one where the pictures of sets look more and more similar.

5.2 $\mathcal{L}$ is closed, the world rejoices.

Every lamination is a closed subset of S, so $\mathcal{L} \subset \mathcal{P}(S)$. (Clearly, a strict subset, since every lamination is nowhere dense!) We endow it with the subspace topology.

Now all our talk previously of limits of laminations requires that $\mathcal{L}$ be *closed.* That is, we want $\mathcal{L}$ to be a closed subset of $\mathcal{P}(S)$. If so, then as a closed subset of a compact space, it will be compact itself. Any Hausdorff-convergent sequence of laminations will converge not just to any old closed set, but to a lamination. And, any sequence of laminations at all will have a subsequence converging to a lamination.

This satisfies our aesthetic senses, and also means that we can define examples in the way discussed earlier.

Proposition 5.1 *Let S be a closed oriented hyperbolic surface and let $\mathcal{L}$ be the set of all laminations on S. Then $\mathcal{L}$ with Hausdorff distance is closed in $\mathcal{P}(S)$; so $\mathcal{L}$ is compact.*

The idea of the proof is as follows. Take a sequence of laminations L_n and suppose that L_n converges to a closed set $L \in \mathcal{P}(S)$. We must show L is actually a lamination. Take $x \in L$, and $x_n \in L_n$ converging to x, with x_n lying on the leaf γ_n of L_n. Now we want to show that the geodesics γ_n converge (locally, Hausdorff-wise) to a geodesic γ through x. This is not so easy to do directly.

What is easier to do, although more technical, is to take the projectivized tangent bundle (alas!) $PT(S)$, and for each lamination L_n take a lift $\tilde{L_n}$ in

$PT(S)$, which is the image of a section of $PT(S)|_{L_n}$. We consider the $\tilde{L_n}$ as closed subsets of $PT(S)$, i.e. in $\mathcal{P}(PT(S))$. Now $\mathcal{P}(PT(S))$ is compact just like $\mathcal{P}(S)$ is. To make this proof rigorous, you have to figure out how this works and that, whether we consider laminations in $\mathcal{P}(S)$ or their lifts in $\mathcal{P}(PT(S))$, we get the same topology on $\mathcal{L}$. With this fearsome technological structure in place we return to the original argument.

The geodesics γ_n, or rather their lifts $\tilde{\gamma}_n \subset \tilde{L}_n$, converge to a geodesic $\tilde{\gamma}$ *by definition*: the points x_n on them converge to x, and the directions converge because that's what convergence in $PT(S)$ means.

6 Isolated and boundary leaves

We now turn to two particular types of leaves in a lamination, which are useful and important: isolated leaves, and boundary leaves.

6.1 Derivation: Destroying isolated leaves

Sounds like a rather ruthless strategy for victory in a battle with a tree. But it's a very useful thing to do with laminations and has some amazing properties.

As we have seen, laminations can't have leaves which are so close together that they fill up a ball in the surface. But they can have leaves which are relatively close together, as in the spiralling-into-a-closed-geodesic example, for instance. They can also have leaves which are far apart, for instance, if the lamination is merely a finite collection of simple closed geodesics. The bits that are relatively close together are the interesting part, because they can be very very subtle and weird and Cantor-like. We therefore want an operation to remove everything else. That's what derivation is: it removes isolated leaves.

What's an isolated leaf? Think about what it looks like, write the definition you think it has and you're probably correct. In any case I'll write a definition too.

Definition 6.1 *Let S be a closed oriented hyperbolic surface and L be a lamination on S. A leaf γ of S is* isolated *if every point x of γ has a neighbourhood U_x in S such that $L \cap U_x$ only contains one arc, namely that of γ. Equivalently, every $x \in \gamma$ has a neighbourhood U_x such that the pair $(U, U \cap L)$ is homeomorphic to (disk, diameter).*

Note that in this definition, the neighbourhood U_x is allowed to have different sizes for different points on γ. There is not a uniform bound so that you can take a tubular neighbourhood of γ with that width. If the leaf γ is closed, there will be such a tubular neighbourhood. But if γ is not closed, i.e. γ is infinitely long, there need not be any such tubular neighbourhood.

So, in the example of a finite union of disjoint simple closed geodesics, all the leaves are isolated. In the example of a closed leaf and a non-closed leaf spiralling towards it, the closed leaf is not isolated. But the spiralling leaf *is* — and the sizes of the neighbourhoods in the above definition approach zero as we approach the closed leaf. Maybe this isn't quite a consequence we first had in mind, but we will go with it.

A derived lamination is obtained, as the heading suggests, by destroying the closed leaves.

Definition 6.2 *Let S be a closed oriented hyperbolic surface and L a lamination on S. The* derived lamination *L' of L consists of the non-isolated leaves of L.*

Is the derived lamination actually a lamination? Thankfully it is, else the language would have been most unfortunate. We just need to see why it's closed. Points on non-isolated leaves can't converge to isolated leaves. (Of course not, because they're isolated! This definition of "isolated" is good: it adds rhetorical force to arguments that rely not on rigour but on persuasion...) So a sequence in L', if it converges, can only converge to a point in L'; hence L' is closed.

To give you a taste of what this operation can do, here are some theorems about derivation.

As always, let S be a closed oriented hyperbolic surface, and L a lamination on S.

Theorem 6.3 *$L''' = L''$.*

That is, derivation has a sort of stability property: once you do it enough times, you get to a stable sort of lamination to which derivation does nothing. A lamination of this type has a special name.

Definition 6.4 *A lamination is said to be* perfect *if $L' = L$.*

So the theorem above says that L'' is always perfect. In fact, laminations can become perfect even earlier, if there are no closed leaves.

Theorem 6.5 *If L has no closed leaves then $L'' = L'$.*

In this sense, and as we will see later, closed leaves are actually something of an annoyance to a lamination. We originally may like closed leaves, because they are the easiest ones to see! But in fact they just make life difficult. We will come to prefer laminations with no closed leaves. When we come to discuss the correspondence between laminations and the mapping class group, a closed leaf will indicate that a surface automorphism is reducible, and reducible is no fun for anybody. Pseudo-anosov is fun, and that means no closed leaves.

6.2 Boundary leaves and why they are everywhere

In addition to isolated leaves, there are boundary leaves. These are called boundary for a reason!

Soon, we will be considering what happens when we cut up the surface along the lamination. We will get a disconnected bunch of smaller surfaces. Leaves that are "close" to others. Now what is the boundary of these smaller surfaces? You might say, they are the leaves of the lamination! Well, each boundary component will indeed be a leaf of the lamination. But not every leaf will be a boundary component. Can you see why?

If you can't see why immediately, consider that any leaf which forms a boundary component must have a certain property — it must be "isolated on one side". If we have a leaf γ which is not isolated on either side, then any chunk of surface near it has other leaves closer to it than γ. So γ will not be a boundary component. In fact you should be able to see that this is condition of being isolated on one side is equivalent to being a boundary component of the dissected surface. Just like the definition of "isolated", you can probably write a definition of this idea of "boundary" or "isolated on one side" for yourself. But I'll do it as well, again, for the sake of completeness.

Definition 6.6 *Let S be a complete oriented hyperbolic surface and L a lamination on S. A leaf γ of L is a* boundary *leaf if at each point of x of γ there is a half-disk neighbourhood V_x of x, with diameter lying on γ, such that $L \cap V_x$ contains only the diameter of V_x.*

(Equivalently, for all $x \in \gamma$ there exists a disk neighbourhood U_x such that $U_x \cap L$ contains at least one component of $U_x - \sigma$, where σ is the arc of γ in U_x containing x.)

Now, if you cut the surface S along the lamination L, you would think that the regions you obtain should consist of most of the original surface! If the leaves of the lamination were densely packed somewhere, so they covered an entire ball in the surface, then that whole part of the surface would be missing from the dissected surface. But we know that our lamination L is nowhere dense; so that can't happen. Hence we expect that the remaining leaves can't count for much. The boundary leaves account for most of what's going on — we lose non-boundary leaves after decomposing the surface, but they don't count for much, in some sense.

In fact what we are trying to get at here is that the boundary leaves are *dense* in the lamination. Everything is close to a boundary leaf; boundary leaves are everywhere!

Lemma 6.7 *Let S be a compact oriented closed surface and L a lamination on S. The union of the boundary leaves of L is dense in L.*

PROOF If you are a leaf of L, but not a boundary leaf, you can't bound anything. So there have to be other things close by on both side — you are not isolated on either side. Hence there are leaves arbitrarily close on both sides. What we want to show is that there are *boundary leaves* arbitrarily close; this is not difficult.

So, take a point $x \in L$ on a leaf γ; we must show there is a point of a boundary leaf arbitrarily close to L. If γ is a boundary leaf we are done. If γ is not a boundary leaf, then as we just mentioned, there are leaves arbitrarily close to γ. (In fact on either side, but we don't care.) Now take a point $y \in S - L$ close to x; this is possible since L is nowhere dense. Take an arc connecting y to x; proceeding from y, this arc hits a first point of L, say z. Now z lies on a leaf which is not γ, which is a boundary leaf (bounding the region containing y), and which is arbitrarily close to x. ∎

7 Principal regions and What They Look Like

"Principal region" is the name given to the components of the dissected surface we have been discussion — the ones whose boundary components are the boundary leaves of L.

Definition 7.1 *Let S be a closed oriented hyperbolic surface and L a lamination on S. A* principal region *for L is a component of $S - L$.*

Now that we have them, we may ask what these principal regions look like! Let U be a principal region for L. Clearly U is also an oriented hyperbolic surface. But U may not be compact. Its boundary components are the boundary leaves of L. A closed boundary leaf will correspond to a closed geodesic boundary. A non-closed, hence infinite length, boundary leaf will correspond to a non-compact boundary component; the ends of this leaf will be ends of the principal region.

A hyperbolic surface also has a lift to $\tilde{S} \cong \mathbb{H}^2$. So a good way to understand principal regions will be to consider them and how they sit in the hyperbolic plane, relative to the rest of the surface. This is good, because if we try to draw them in the surface itself they may look hideously complicated, with boundary leaves that are a complete mess. In the hyperbolic plane, however, geodesics look straight (or at least, circular in the unit disk model), and hence the principal regions can't look too bad.

To understand what principal regions look like, we will make great use of their lifts to $\mathbb{H}^2$.

7.1 They are very finite...

It might seem at first that our lamination could have extremely many extremely complicated leaves, many of them boundary; and hence, there could be finitely many principal regions. Thankfully, this is not the case. The reason: the Gauss–Bonnet theorem, which has the oft-forgotten consequence that the area of a hyperbolic surface (with geodesic boundary, even if it has punctures) is tightly constrained — in fact, an integer multiple of π

Lemma 7.2 *Let S be a closed oriented hyperbolic surface and L a lamination on S. Then L has only finitely many principal regions, each with only finitely many boundary leaves. (Hence only finitely many boundary leaves.)*

The above observation shows that there are finitely many principal regions. And if a principal region has infinitely many boundary leaves, its Euler characteristic is infinite, implying infinite area by Gauss–Bonnet, a contradiction.

7.2 ... can be lifted to the hyperbolic plane ...

Let's now take a look at a lift of a principal region P of a lamination in $\mathbb{H}^2$. It's going to bound be a region of the hyperbolic plane bounded by geodesics (possibly infinitely many, since its lift may consist of infintiely many copies). The geodesics bounding the plane have to be disjoint, else they would project to intersecting leaves of L in S. So topologically it's a disk, and contractible; in fact, hyperbolically convex.

Lemma 7.3 *Let S be a closed oriented surface and L a lamination on S. Let P be a principal region for L and $\tilde{P}$ a component of the preimage of P in $\mathbb{H}^2$. Then $\tilde{P}$ is a contractible hyperbolic surface with geodesic boundary.*

We can also see that a principal region can't bound a disk: a geodesic in a hyperbolic surface can't bound a disk.

Lemma 7.4 *If P is a principal region for L, then $\pi_1(P) \longrightarrow \pi_1(S)$ is injective.*

7.3 ... look nice when there aren't any closed leaves...

Now we come to the first really interesting result about the structure of principal regions. It turns out that when *there are no closed leaves*, they look particularly simple.

What will the regions look like? If there are no closed leaves, the principal regions will be surfaces having boundary components homeomorphic to $\mathbb{R}$. (Only boundary leaves which were closed leaves could give boundary components homeomorphic to S^1).

One possibility is that a principal region P is a polygon with vertices at infinity. This possibility gives P no interesting topology. Now P can have topology, however; but without any S^1 boundary components, the geometry at the boundary must be a little subtle. If you have some familiarity with the geometry of hyperbolic surfaces, you will know that a non-compact hyperbolic surface can have a *compact core*; that is, a connected compact subsurface (or, in a degenerate case, a single geodesic) with boundary consisting of geodesics, and which contains all the interesting topology of the surface; and such that it separates the non-compact ends of the surface, which "flare" out to infinity. In our case, the compact core will separate ends which have one S^1 boundary component and finitely many $\mathbb{R}$ components, and have zero genus. Thus each end looks like a "crown": namely, an annulus with finitely many punctures on one of its boundary components; the geometry near each puncture goes out to infinity.

It turns out these are the only possibilities. If you know enough about the geometry of hyperbolic surfaces, this may be obvious, once you consider that we start with a surface of finite area. In any case we will say something about why it's true.

Lemma 7.5 *Let S be a closed oriented hyperbolic surface, and let L be a lamination on S without closed leaves. Let U be a principal region for L. Then U is either:*

(i) isometric to a finite sided polygon with vertices at infinity; or

(ii) there exists a unique compact subset of U_0 (either a subsurface, or a geodesic) such that $U - U_0$ is isometric to a finite disjoint union of crowns.

To see why it's true, we consider $\tilde{U} \subset \mathbb{H}^2$. The boundary of this region consists of $\mathbb{R}$ components. Take a boundary leaf β_0; since the region has finite area, there are adjacent boundary leaves adjacent to each end of β_0. So we obtain a sequence β_n. This sequence may cycle around, in which case $\tilde{U}$ is a finite sided polygon with vertices at infinity — and hence, so is U (it is the universal cover of something, something orientable, but that something can't have an infinite fundamental group). Or it may not, in which case the β_n in the disk model get smaller and smaller (in the Euclidean metric!) and converge to two limit points $\beta_{-\infty}$ and β_∞. (Note although U can only have finitely many boundary leaves, its universal cover can have infinitely many; and will, if $\pi_1(U)$ is infinite.) Joining these we obtain a convex region $\tilde{W} \subset \tilde{U}$; it is the convex hull of the β_n.

Now $\tilde{W}$ covers W, which will be a crown. Well clearly none of the boundary leaves β_n project to circles, since we have no closed leaves. But the deck transformation taking $\beta_n \mapsto \beta_{n+1}$ will project the axis joining $\beta_{-\infty}$ to β_∞ to a circle. So this gives a crown. Cutting all these out gives the compact core.

7.4 ... without any annoying closed leaves, are finite at infinity...

One more property. We now see that our principal regions will lift to regions of $\mathbb{H}^2$, with compact cores and crown sets and polygons tessellating the plane. These are all bounded in $\mathbb{H}^2$ by geodesics with endpoints at infinity. We know that there are only finitely many principal regions, each with only finitely many boundary leaves; but in $\mathbb{H}^2$, a compact core can have infinitely many boundary components in its universal cover, and a crown has infinitely many boundary components in its universal cover. So it looks like things are much more infinite in $\mathbb{H}^2$, perhaps as expected.

But not *so* infinite. Turns out each vertex at infinity is very finite! Again, *provided* there are no closed leaves.

Lemma 7.6 *Let S be a closed oriented hyperbolic surface and let L be a lamination on S without closed leaves. Any point on the circle at infinity is the endpoint of only finitely many leaves of $\tilde{L}$.*

Why is this? Suppose there were infinitely many leaves meeting at the same point at infinity x. Since there are only finitely many boundary leaves, by the pigeon hole principle there's a deck translation g fixing x. Now translating boundary leaves ending at x under g, they converge to the axis of g. As a lamination is closed, the axis of g is a leaf of L. But since g is a deck transformation, this projects to a closed leaf of L, contradiction.

It's precisely the absence of closed leaves that makes this impossible. If closed leaves are allowed, all manner of infinities in the degrees of these vertices can and will arise. In fact, if you think about taking the universal cover of any surface with some closed leaves in a lamination, you'll see that there will generally be infinite degree vertices at infinity; these arise whenever there is spiralling in towards a closed leaf.

So, we see that the absence of closed leaves really does simplify matters, perhaps to a bizarre extent, when we consider the picture in $\mathbb{H}^2$.

7.5 ... and sometimes even when there *are* closed (isolated) leaves...

We now have a nice picture of the lift of a lamination to $\mathbb{H}^2$. There are principal regions, which look rather nice when there are no closed leaves. If there are no closed leaves, we either have finite sided ideal polygons, or compact cores and crowns, both of which are quite simple phenomenon. In fact, this picture can be useful even when there are no closed leaves: given a lamination L, take out the (necessarily finitely many) closed leaves, look at the principal regions, and put the closed leaves back in again. However it's not that simple: our closed leaf may be isolated; or it may be isolated only on one side, i.e. a non-isolated boundary leaf; or it may be isolated on neither side, i.e. a non-boundary leaf. Only when the closed leaf is isolated is the picture easy to see.

What does this amount to? Let us consider the picture in $\mathbb{H}^2$. If our closed leaf is isolated, then inserting it into the picture simply subdivides some existing principal regions. So this amounts to inserting a geodesic into the interior of finite sided ideal polygons and compact-cores-with-crowns. But the closed leaves can't cross the non-closed leaves — a lamination consists of *disjoint* geodesics! So each closed leaf either lies inside a finite sided ideal polygon, or lies inside a compact-core-with-crowns.

If an isolated closed leaf lies inside a finite sided ideal polygon, it's easy to see it must just be a diagonal. Actually, as it turns out this picture can't happen. The reason why is the argument of the previous lemma. The deck translation g corresponding to our closed leaf γ must carry some sides of the polygon closer to γ, contradicting the fact that we have a principal region. (And, for that matter, showing that γ is not isolated. Because diagonals in ideal polygons are isolated.)

If an isolated closed leaf lies in the interior of a compact-core-with-crowns, it must lie inside the compact core. If an endpoint in $\mathbb{H}^2$ lifts to a crown, it can't possibly project to a closed leaf, since it has infinite distance going out to a point of a crown. So its endpoints must lie in the compact core; and as the core is convex, so must the entire leaf.

This case includes our simplest example: when L is a finite collection of simple closed geodesics. Then we remove the closed geodesics, obtaining the empty lamination; we take a principal region, which is the entire surface, for which there are no crowns and the entire surface is the compact core; and then we re-insert the closed leaves, which certainly lie within the compact core! Maybe not such an interesting example.

Things are more complicated when we have a non-isolated closed curve. The problem is, once it is removed, we no longer have a lamination; the set is no longer closed. So all our previous arguments for laminations without closed leaves no longer apply.

7.6 ... and take up all the space...

We have already said that geodesics can't be too close; that a lamination is nowhere dense in the surface (proposition 4.3); and that boundary leaves are dense in the lamination (lemma 6.7). It would seem, then, that they should be small. But we have not said anything about the *area* of the lamination. It should be measure zero, surely! Thankfully it is.

Lemma 7.7 *Let S be a closed orientable hyperbolic surface and L a lamination on S. The area of L is zero; the area of $S - L$ is the full area of S.*

This is a vector-field and Gauss–Bonnet bash. Construct some vector fields on the principal regions, with some singularities, in a canonical way, and use that the Euler characterisic.

7.7 ... while the rest is rather weird.

Now for perhaps the most confusing part of the picture of all. We've got an excellent picture of principal regions, at least when there are no closed leaves. We have principal regions all over $\mathbb{H}^2$. You might have got the impression that they in fact *tessellate* $\mathbb{H}^2$. Sadly, if you have this impression, it's not quite correct. It will be correct only if every leaf is isolated. We actually have the plane split into principal regions, with strange "twilight zones" in between of non-boundary leaves — except that there is no zone as such. This sounds weird. What does this mean?

Well, certainly if we have a boundary leaf which is not isolated, beyond it will be non-boundary leaves. But leaves are nowhere dense, and moreover, boundary leaves are dense in the lamination. So there's no region of non-boundary leaves or anything; in fact, there must be a boundary leaf, and hence a principal region, arbitrarily close to any non-boundary leaf. On the other hand, if we have a non-boundary leaf, there are leaves which come arbitrarily close to it; and then it is not difficult to show that there are leaves which come arbitrarily close to any point of the leaf (to be close at a particular point, it suffices to be very very *very* close at another point!).

Reconciling these apparently-contradictory-but-not statements gives a picture of the non-boundary leaves. Well, not so contradictory at all, come to think of it. The non-boundary leaves can't just disappear from the picture in $\mathbb{H}^2$ — and they can't bound any principal regions! It follows that they have to be doing something like this.

8 Derivations and their properties

8.1 Stability, Perfection and Their Consequences

We now look a bit more closely at derivations. Turns out a derived lamination L' has interesting stability properties. As we will see, *the minimal sublaminations of L' are precisely its components.* Let's see what this amounts to.

Well, it implies some pretty strong and amazing stuff. Because L'' is a sublamination of L', hence a union of minimal sublaminations — so L'' just consists of a subset of the components of L'. Some of the components of L', upon being derived, cease to exist; others of them persist — and since they are minimal, their derivation must be themselves. So the components of L' either disappear under derivation, or are perfect; and L'' consists precisely of the perfect components. It follows that L'' is perfect and $L''' = L''$. So once we prove the above statement, we have proved theorem 6.3.

It also implies theorem 6.5: if L has no closed leaves then $L'' = L'$. For when we consider the components of L', the ones that disappear are sublaminations whose derivation is empty — they are sublaminations consisting entirely of isolated leaves. Any lamination consisting entirely of isolated leaves actually consists entirely of closed leaves; an infinite leaf will spiral towards something. But if we rule out closed leaves, then none of these occur, so L' is perfect.

There is another corollary as well:

Theorem 8.1 *The following are equivalent:*

(i) every leaf of L is dense in L;

(ii) L is connected and $L' = L$.

If every leaf is dense, then the lamination is certainly connected, and there are no isolated leaves, so $L' = L$. On the other hand, if L is connected and $L' = L$, then $L' = L$ is a union of perfect components. Since it is connected, it is one perfect component. So the closure of any leaf is that component, hence the whole lamination.

8.2 How the proof goes.

Now, let's see why it might be true that the minimal sublaminations of L' are precisely its components. A good example to bear in mind is a lamination consisting of 3 leaves: a closed leaf, and two non-closed leaves spiralling towards it from either side.

A component of a lamination is always a sublamination; but not necessarily a minimal one. There may be extra leaves in there. In our example there is only one component, and the whole lamination is certainly a sublamination, but it's not minimal. For instance, just take the closed leaf. On the other hand, if you take a minimal sublamination, it is the closure of any leaf in it; but this minimal sublamination may not be a component. In our example, the closed leaf is a minimal sublamination, but certainly not a component.

In our example, however, deriving removes both the spiralling leaves, so we are left only with one leaf, and clearly the minimal sublaminations are the components now.

Note what this means: upon taking L', everything near closed leaves gets removed — leaving closed leaves isolated. This makes sense — closed leaves will now become isolated components of L', and will be among the components of L' that disappear under further derivation. And in L'', there are no closed leaves, and no leaves spiralling towards closed leaves — which is good, for such things are not perfect.

In fact, in our example there is a general phenomenon occurring: any leaf close to a closed leaf is isolated.

Lemma 8.2 *Let S be a closed oriented hyperbolic surface and let L be a lamination on S. Let C be a closed leaf of L. Then any other leaf sufficiently close to C is isolated. That is, there exists a neighbourhood N of C such that any leaf other than C intersecting N is isolated; equivalently, $L' \cap N \subseteq C$.*

To see why, again we look to $\mathbb{H}^2$. The closed leaf C lifts to a line $\tilde{C}$ with a deck transformation g translating along it. Think about what a leaf close to C looks like. We might expect it to share an endpoint at infinity with $\tilde{C}$. This is right. If it shares no endpoint at infinity, and is sufficiently close, then translating it by g will see it intersecting its translate. This is bad, because there is no self-intersection in leaves. So it does share an endpoint at infinity with $\tilde{C}$. But we know that there are only finitely many leaves of $\tilde{L}$ intersecting at a point at infinity, so C is isolated.

Now for the main theorem.

Theorem 8.3 *Let S be a closed oriented hyperbolic surface and L a lamination on S. If L_1 is a sublamination of L' then L_1 is a union of components of L'. Equivalently, the minimal sublaminations of L' are precisely its components.*

Proof The idea is simply to analyse what's in $L_1 \subset L'$ and where they lie in the picture in $\mathbb{H}^2$. In particular, L_1, being a subset of L', contains no isolated leaves of L: it contains only non-isolated leaves of L. Amongst these there are

(i) closed leaves; and

(ii) non-closed leaves — these we call L_2.

Now L_1 will consist of closed leaves, as well as L_2. We have showed that anything near a closed leaf is isolated, so the leaves in category (i) will be isolated in L', hence components of L'. We only need to show that the leaves in category (ii), namely L_2, also form a union of connected components of L'.

Now L_2 is a lamination, being a sublamination of L' where closed leaves in L' are isolated. The point is to show that if we take any other leaf in L', it lies far away from L_2; for then L_2 is a union of connected components of L'. The way we do this is to consider the principal regions of L_2; anything in L' but not L_2, namely anything in $L' - L_2$, we will show lies in the compact cores of the principal regions of L_2. Thus they are far from L_2, which are all near boundary leaves and the boundaries of the principal regions.

We consider the possibilities for a principal region U of L_2. If $\tilde{U}$ is a finite-sided ideal polygon, and γ is a leaf of $L' - L_2$, then $\tilde{\gamma}$ is a diagonal, so is isolated. In fact, it is also in L of course ($L' \subset L$!), and must have been isolated there, a contradiction since then it wouldn't be in L'!

If $\tilde{U}$ consists of a compact core and crown regions, then $\tilde{\gamma}$ either lies in a compact core or has an end going to an end of a crown region. The first case is what we want. In the second case, the leaf is isolated (and was in L), a contradiction again to being in L'. ∎

In fact we can do something which is maybe slightly stronger and definitely much more convoluted! (This is the statement Casson and Bleiler prove, which is perhaps not the most friendly approach from the point of view of understanding.) We have shown that any sublamination of L' consists of components of L', but in fact, *any sublamination of* L, when intersected with L', consists of components of L'.

Corollary 8.4 *Let S be a closed oriented hyperbolic surface and L a lamination on S. If L_1 is a sublamination of L, then $L_1 \cap L'$ is a union of components of L'.*

For $L_1 \cap L'$ is certainly a sublamination of L': the leaves of L_1 which intersect L' are certainly leaves of L'; and the intersection of two closed sets is closed. The corollary now follows immediately.

9 Surface Automorphisms

The most amazing thing about surface automorphisms is their classification: periodic, reducible, or pseudo-Anosov. In essence, reducible ones are annoying because they really belong not on our surface, but on a decomposition of the surface into smaller pieces. Periodic ones also just get in the way; it's pretty unusual to have a periodic automorphism. What's amazing is that *everything else* is pseudo-Anosov, which means having some very interesting dynamical properties (defined in terms of measured foliations). It's not clear at all why "most" nontrivial automorphisms should have such an interesting dynamical description.

These definitions are topological; they do not involve hyperbolic geometry, or any other type of geometry for that matter. It is by introducing geometry that we get so much interesting structure. The definitions are up to homotopy, and we'll be homotoping stuff around quite freely.

Definition 9.1 *Let S be a closed orientable surface. Let $h : S \longrightarrow S$ be an automorphism of S. The automorphism h is called:*

(i) periodic *if there exists a positive integer n such that h^n is homotopic to the identity;*

(ii) reducible *if h is homotopic to an automorphism which leaves invariant an essential closed 1-submanifold of S, i.e. leaves invariant a finite collection of essential closed curves.*

(iii) pseudo-Anosov *if there exist transverse singular foliations $\mathcal{F}^s, \mathcal{F}^u$ (called* stable *and* unstable *foliations) equipped with transverse measures μ^s, μ^u such that for some $\lambda > 1$,*

$$h(\mathcal{F}^s, \mu^s) = (\mathcal{F}^s, \lambda\mu^s), \quad h(\mathcal{F}^u, \mu^u) = (\mathcal{F}^u, \lambda^{-1}\mu^u)$$

The definition of pseudo-Anosov means that the leaves of the stable foliation attract, and the leaves of the unstable foliation repel.

The idea is that reducible automorphisms really belong on the surface obtained by cutting along the invariant 1-submanifold. And periodic automorphisms cannot possibly have the pseudo-Anosov property.

The above-mentioned amazing fact is contained in the following theorem.

Theorem 9.2 *Every non-periodic irreducible automorphism of a closed oriented hyperbolic surface is isotopic to a pseudo-Anosov automorphism.*

We'll get to this eventually.

9.1 Telling them apart

We can tell something about reducibility or periodicity from looking at the induced map of a surface automorphism on its homology. As always, let S be a closed oriented surface — for now it's just topological, not hyperbolic, not geometric. Then the automorphism $h : S \longrightarrow S$ induces $h_* : H_1(S) \longrightarrow H_1(S)$. The first homology over $\mathbb{Z}$ is just a free abelian group (of rank $2g$ where g is the genus of S). So h_* can be written as a $(2g \times 2g)$ matrix A with integer entries and which is invertible. We write $\chi_h(t)$ for the (degree $2g$) characteristic polynomial of this matrix.

Lemma 9.3 *If h is periodic then all the zeroes of $\chi_h(t)$ are roots of unity.*

This is clear; since $h_*^n = 1$, we have $A^n = 1$; so the minimal polynomial of A is a factor of $t^n - 1$, hence has all zeroes being roots of unity; hence also the characteristic polynomial.

Lemma 9.4 *If h is reducible then either $\chi_h(t)$ has a root of unity as a zero, or is reducible, or is a polynomial in t^n for some positive integer $n > 1$.*

This is just a little linear algebra, and requires different considerations depending on how the invariant 1-submanifold separates the surface. If you can get $h_*(C) = C$ for some (homology class of a) closed curve C, you'll get a root of unity in the characteristic polynomial. But h_* may permute the components of the invariant 1-submanifold, and then we have to think about the matrix: the characteristic polynomial then may be reducible; if it is not reducible, we find it is of the form $\det(B - t^n I)$ for some $n > 1$, and hence a polynomial in t^n.

9.2 Surface automorphisms meet laminations: The invariant lamination

So far we have only discussed surface automorphisms in the abstract. No mention of what they have to do with laminations! The definition of pseudo-anosov has to do with *foliations*, which may look a little like laminations, but not really; they're very different, in fact.

The connection is: an automorphism which is not periodic — (which we secretly know will be reducible or pseudo-Anosov) — has an invariant *lamination.*

Theorem 9.5 *Let S be a closed oriented hyperbolic surface and let h be a non-periodic automorphism of S. Then there exists a lamination L on S such that $h(L) = L$.*

In saying that $h(L) = L$, we are employing the standard procedure of taking geodesics to geodesics by a straightening process. (We are homotoping rather blithely, then.) This is done by lifting h to the universal cover $\mathbb{H}^2$, where it acts on $\mathbb{H}^2$ and extends continuously to the circle at infinity S^1_∞. (I have not proved this, but it's a standard fact, and it's proved in Casson and Bleiler.) Then straightening out geodesics by noting where their endpoints map to at infinity.

The proof is a bit tricky. It uses a clever trick to obtain an infinite sequence of distinct geodesics; and then it uses the compactness of the space $\mathcal{L}$ of laminations; and then some more!

Proof We first show that there exists a simple closed geodesic C such that for all $n \geq 1$, $h^n(C) \neq C$. Suppose to the contrary that there is not; then for all simple closed geodesics C, there is a sufficiently high power of h which fixes it. Taking enough geodesics and taking a multiple of all the sufficiently high powers — which is alright since h is not periodic — we fix so much of the surface that h must in fact be periodic. This is a contradiction. So there is such a C.

Now we look at $h^n(C)$. This is a sequence of simple closed geodesics in S, hence a sequence of laminations. But $\mathcal{L}$, the space of laminations, is compact — so this sequence has a convergent subsequence $h^{n_i}(C)$, which converges to a lamination K.

Seems like we might be done. Well we would be, if $h(K) = K$. But sadly, there is no reason why this should true. (We would have been fine, if our sequence $h^n(C)$ were convergent — but we only have a subsequence.) We're going to try to get an invariant lamination out of this.

It is therefore natural to consider the laminations $h^n(K)$. If they were disjoint, we could just take their union, (and close it,) and then we would have an invariant lamination. In fact, even if they were not disjoint, but contained some overlaps of the same geodesics, this would be good enough: we still just take union and closure. But they may not be disjoint! So we ask: how much might they intersect? Now, *transverse* interections are what matter; degenerate intersections of geodesics mean that both laminations contain the same geodesic, which as we just noticed is no problem at all.

We therefore need to consider $h^r(K) \cap h^s(K)$. Well we can cancel some h's here so it is sufficient to consider $K \cap h^r(K)$, and transverse intersections in it. This is a limit of $h^{n_i}(C) \cap h^{n_i+r}(C)$, which has the same number of transverse intersections as $C \cap h^r(C)$. When we take a limit, the number of transverse intersections might change — but it *can only decrease*. For take a small neighbourhood of each transverse intersection point in the limit; these also contain transverse intersections of the converging laminations — and possibly more of them.

Therefore, we let N_r be the number of transverse intersections of $C \cap h^r(C)$. Then $K \cap h^r(K)$ has $\leq N_r$ transverse intersections. In particular, this number *is finite*, and consists of discrete points. So every leaf where a transverse isolation occurs *is isolated*. It follows that K' and $h^r(K')$ have no transverse intersection points; and hence, any $h^r(K')$ and $h^s(K')$ have no transverse intersections.

So we now carry out our strategy, take the union $\cup^\infty h^r(K')$ and its closure $\overline{\cup^\infty}(K')$. This is the lamination we want. ∎

9.3 What the invariant lamination means

We have looked at several features of laminations. Some of these correspond to features of the automorphism.

As a first example, suppose that L has a closed leaf. Since h fixes L, it takes a closed leaf to another closed leaf. There can be only finitely many closed leaves (because they are disjoint simple closed geodesics), and hence after possibly cycling around we have a finite collection of disjoint closed leaves invariant under L; so L is reducible.

Lemma 9.6 *Let S be a closed oriented hyperbolic surface and L a lamination on S which is invariant under the automorphism h. If L has a closed leaf then h is reducible. Equivalently, if h is irreducible then L has no closed leaves.*

This is the real proof that closed leaves are annoying. Closed leaves are annoying because reducible automorphisms are annoying; and closed leaves imply reducible automorphisms.

As another example, consider the principal regions. If there are no closed leaves, then these come in two types: finite sided ideal polygons; and surfaces with compact cores with crowns attached. If the lamination is fixed under an automorphism h, however, then these regions must be carried to themselves. In particular:

(i) boundary leaves map to boundary leaves

(ii) principal regions map to principal regions; and in particular

(iii) finite ideal polygon principal regions map to finite ideal polygon principal regions, and

(iv) compact cores and crowns map to compact cores and crowns.

(v) In particular, boundaries of the compact cores map to boundaries of compact cores.

However, boundaries of compact cores are simple closed curves in our surface. So they give us a 1-submanifold invariant under h.

Lemma 9.7 *Let S be a closed oriented hyperbolic surface and L a lamination on S. Suppose L is invariant under an automorphism $h : S \longrightarrow S$. If L has no closed leaves, and there exists a principal region of L which has a compact core, then h is reducible. Equivalently, (and better!), if h is irreducible then (as we previously saw, there are no closed leaves and) each principal region of L is a finite sided ideal polygon.*

Now, this picture says even more. If we only have finite sided ideal polygons for principal regions, then *the lamination is connected.* Just go from one principal region to another along boundary leaves — which meet at infinity (and hence are arbitrarily close, so in the same connected component). Boundary leaves are dense in the lamination, as we know, so the lamination is connected.

Lemma 9.8 *We still suppose L is invariant under h as above. If h is irreducible, then L is connected.*

We can still say more. Being connected as a lamination has even more relevance if you are of the form L', i.e. you are a derived lamination. If $h(L) = L$, then certainly $h(L') = L'$. If h is irreducible then L has no closed leaves, so certainly not L'. So what does L' look like? As we know, its components are precisely its minimal sublaminations, and there are no closed leaves. If L' is connected then: it's just one component, perfect, and has no proper sublaminations. So $L'' = L'$ and the closure of any leaf of L' is all L'. Said another way, every leaf of L' is dense in L'.

Even more than this: take a leaf γ of L; as h is irreducible, γ is not closed. Then its closure $\bar{\gamma}$ is a sublamination of L, and by corollary 8.4, $\bar{\gamma} \cap L'$ is a union of components of L. (If you don't like annoying corollary 8.4, you can quite easily still use the more approachable theorem 8.3.) It's nonempty, since otherwise $\bar{\gamma}$ would consist entirely of isolated leaves of L, impossible as γ is not closed. Since L' is connected, $\bar{\gamma} \cap L' = L'$, so γ is dense in L. So not only is every leaf of L' dense in L'; better, every leaf of L is dense in L'.

Lemma 9.9 *We still let L be invariant under h. If h is irreducible, then L' is connected and perfect and every leaf of L is dense in L'.*

We are not done yet. Let us consider L' further; in particular, let us consider its principal regions. We know that $h(L') = L'$, so from the argument we applied above to L we see that each principal region of L' is a finite sided ideal polygon. But now, how do these join up? Suppose one edge of such a polygon in $\mathbb{H}^2$ joins directly to an edge of another polygon; then the edge along which they were joined is isolated (in L'). But L' contains no isolated leaves; in fact, it is connected and perfect. What must happen is that non-boundary leaves interfere on the other side of the edge, which are nowhere dense but come arbitrarily close to our edge; and then there is a whole mess of boundary leaves among them (since boundary leaves are dense in the lamination). In any case it follows that the vertices at infinity of the polygon principal region have degree precisely 2. If they had larger degree, the other edges would interfere with nearby non-boundary leaves.

Lemma 9.10 *Still let L be invariant under h. If h is irreducible, then every principal region of L' is a finite sided ideal polygon for which every ideal vertex is the endpoint of precisely two leaves of L'.*

To summarise:

Proposition 9.11 *Let S be an oriented closed hyperbolic surface. Let h be an irreducible automorphism of S and let L be a lamination invariant under S. Then:*

(i) L has no closed leaves;

(ii) every principal region of L is a finite sided ideal polygon, and every principal region of L' is a finite sided ideal polygon for which every ideal vertex is the endpoint of precisely two leaves of L';

(iii) L is connected;

(iv) L' is connected and perfect;

(v) every leaf of L is dense in L'.

That's quite a lot to say about a lamination, from the simple fact of irreducibility — and every automorphism is irreducible once you cut up the surface enough.

9.4 Attraction and Repulsion

In nature, as in life, some things attract, and other things repel. Pseudo-anosov automorphisms do both: the leaves of the stable foliation repel, while the leaves of the unstable foliation attract. We want to see why an irreducible and non-periodic automorphism has these attracting and repelling properties.

Let's consider the stable foliation. What we hope to see is a singular foliation, leaves of which attract. The leaves attract each other, but to conserve area (it's a finite area surface!), each individual leaf should be elongated, in some sense.

As it turns out, the stable and unstable foliation of a pseudo-anosov automorphism will arise from an invariant lamination. How do you go from a lamination to a foliation? This will be demonstrated shortly, in a dazzling burst of analysis that uses a Cantor function. (Yes, as in the Cantor set and the devil's staircase.)

So then, in a pseudo-anosov automorphism and an invariant lamination we should perhaps expect to see some fixed leaves; and nearby ones are attracted to them. (Such a fixed leaf will not be closed, since otherwise it would be a reducible automorphism, by the previous section.) This doesn't seem to make much sense, since we have shown that leaves cannot be too close — for instance, they are nowhere dense; furthermore, boundary leaves are dense, and there are only finitely many of them — and boundary leaves are all that show up in the hyperbolic plane! What we will see in $\mathbb{H}^2$, however, is boundary leaves shifting around to get "closer" to our fixed leaf, even though they are still actually quite far away.

In the universal cover $\mathbb{H}^2$, such a leaf γ becomes to a geodesic $\tilde{\gamma}$; even if it's a "singular" leaf which ends at a singularity, its continuation is a geodesic. On either side of γ, other leaves should be attracted. So upon iteration of our iteration h, they should move closer and closer, and in fact converge to $\tilde{\gamma}$. Since we know that each vertex at infinity has only finitely many geodesics ending there, this means that the vertices at infinity converge to the endpoints of $\tilde{\gamma}$.

So, we require that under iteration of h, the circle at infinity moves in such a way that every point converges to an endpoint of $\tilde{\gamma}$. Well, this is not quite right. For one thing, it is impossible since h acts continuously on the circle at infinity; there must be intervening repelling fixed points. For another, this only applies if the geodesic $\tilde{\gamma}$ is "innermost" in some sense.

How could we define "innermost" here? Consider the interval on the circle at infinity S^1_∞ cut of by $\tilde{\gamma}$. What we *can't* have is a situation, say, where this interval is cut in half by two geodesics, lifts of leaves, which form an ideal triangle with $\tilde{\gamma}$. If this happens, the automorphism (being an automorphism of the interval, and leaving $\tilde{L}$ invariant) must fix those geodesics — so that $\tilde{\gamma}$ is not "innermost". When $\tilde{\gamma}$ is "innermost", we call the interval it cuts off on S^1_∞ *stable.*

Try to find a definition that makes all this rigorous. If you can, it's probably equivalent to the following definition.

Definition 9.12 *Let S be a closed oriented hyperbolic surface and L a lamination on S. Let $\tilde{L}$ be the lift of L to the universal cover $\tilde{S} \cong \mathbb{H}^2$. A* stable interval *for L is a closed interval $I \subset S^1_\infty$ such that for any two points P, Q in the interior of I, there is a leaf of $\tilde{L}$ which separates P and Q from I.*

It follows that any stable interval for L is a (lift of a) leaf of L — just take P and Q arbitrarily close to the endpoints of I. It also follows that there are leaves of the lamination which have endpoints at infinity arbitrarily close to the endpoints of I.

We should obtain in this situation a point of I which is a point at infinity fixed under h, and which is the unique repelling fixed point of h in I. The endpoints of I will be only attracting fixed points; the interior fixed point will attract everything else.

To prove our amazing theorem 9.2, we will show: that this definition of stable interval, in our situation, implies something nice about attracting and repelling fixed points; that when h is irreducible and non-periodic, it has nicely behaved fixed points at infinity, alternately repelling and attracting; it can be taken to fix a lamination which involves stable intervals; and that this implies the pseudo-anosov property.

9.5 The Proof I: How good is our invariant lamination?

We now embark on the first part of the proof of our amazing theorem 9.2. We will consider an irreducible and non-periodic h, and the invariant lamination L we obtained earlier in theorem 9.5. Note this required h to be non-periodic.

Where did L come from? Recall we took a curve C for which $h^n(C) \neq C$ — we showed one exists, but any would do. Then we considered $h^n(C)$ and took a convergent subsequence, converging to a lamination K — such a convergent subsequence exists, but any would do. Then we considered $h^n(K)$, showed that these all contain no transverse intersections, so that $h^n(K')$ were all disjoint; and took their union for L.

It is worth thinking about how C, K and L interact. First let us consider K and L. Can they intersect? They certainly can, in fact all of K' lies in $L = \overline{\cup^\infty h^n(K')}$. But what about transverse intersection? Well, any leaf of K which has a transverse intersection with L has a transverse intersection with some $h^n(K')$; and hence has infinitely many points of transverse intersection with $h^n(K)$, a contradiction to what we found in our proof of theorem 9.5. So K and L have no transverse intersection.

We know quite a lot about L', from section 9.3. In particular, we know that every principal region is a finite sided ideal polygon, and that arbitrarily close to each side, there are leaves on the outside. We can ask where our special curve C lies in this picture. Clearly C can't intersect L' tangentially; this would mean it coincides with a leaf of L', but L' contains no closed leaves. Can it avoid L' altogether? If it doesn't intersect L, then it lies in one principal region; and hence it must be a diagonal in the ideal polygon. But since C is a closed curve, this would mean that leaves get arbitrarily close to C; and then deck transformatoins corresponding to C would contradict our picture of our finite sided ideal polygon. So C intersects L' transversely somewhere.

This is interesting: C intersects L' transversely, but K, which is the limit of $h^{n_i}(C)$ for some sequence n_i, does not intersect L' (in fact, not even L) transversely. What has happened? In $\mathbb{H}^2$, C crosses leaves of L'; but upon iterating with h this has straightened out into a leaf of L'. We will exploit this interesting property in regard to stable intervals.

Now let's suppose we have a stable interval I on S^1_∞ cut off by a geodesic $\tilde{\gamma}$. As we have seen, $\tilde{\gamma} \in L$. It follows from the definition of stable interval that there are other leaves arbitrarily close to $\tilde{\gamma}$ on one side, so γ is not isolated, and γ is in L'.

We want to see what happens if h fixes γ, and hence fixes I.

Well, we consider C, which intersects L'. We lift to $\mathbb{H}^2$ and take a component $\tilde{C}$ which intersects $\tilde{\gamma}$, with endpoints A, B. Without loss of generality, A lies in the interior of I, and B lies outside I. Upon iterating h, we know that $h^{n_i}(C) \longrightarrow K$, so $h^{n_i}(\tilde{C})$ converges to a geodesic $\tilde{C}_\infty$, with endpoints A_∞, B_∞. Clearly $\tilde{C}_\infty$ is a lift of a leaf of K, which does not intersect L' transversely. There are not many ways this can happen. Since B lies outside I, so does B_∞. And since A is in the interior of I, $A_\infty \in I$. The only way to avoid transverse intersection with I is to have $A_\infty \in \partial I$. So we have shown that $h^{n_i}(A) \longrightarrow \partial I$.

It follows that A_∞ is an attracting fixed point of $h|_I$.

Are there any other fixed points? Since h fixes L', and since boundary leaves are dense, we can see that the open subinterval $U = (A, A_\infty)$ contains no fixed points; A moves towards A_∞, and so does everything else in between. What about the other end of the interval? Since one end attracts, and by the stable-interval property there are leaves arbitrarily close by, the other end attracts too. What about in the interior? Well if we a leaf close to $\tilde{\gamma}$, and iterate under h^{-1}, it will converge either to a geodesic or a single point at infinity. If it converges to a geodesic, however, then we have ourselves a nested stable interval, which has attracting endpoints under h; but we just showed it attracts points under h^{-1}. This is a contradiction unless the geodesic is degenerate, consisting of a single point Z and all in the interior of I attract to Z under h^{-1}.

Our invariant lamination, then, is good. How good? Very good. Stable intervals which are invariant under h actually are stable: their endpoints are attracting fixed points, and there is a unique repelling fixed point in the interior.

Lemma 9.13 *Let S be a closed oriented hyperbolic surface, and h an irreducible non-periodic automorphism of S. Let L be the lamination invariant under h constructed in theorem 9.5. If any stable interval I is fixed by h then the two endpoints of I are attracting fixed points and the only other fixed point of I is a repelling fixed point in the interior of I.*

Why did we need non-periodic? We needed this for the proof of the invariant lamination. Why did we need irreducible? To use nice properties of the principal regions of L'.

9.6 The Proof II: Alternating attraction and repulsion

We have shown in our situation of a non-periodic irreducible h, stable intervals actually have the simple dynamical property we hope they have. Now we will show that such a property of alternate attraction and repulsion applies globally.

Lemma 9.14 *Let S be a closed orientable surface and h an irreducible non-periodic automorphism of S. Then h has finitely many fixed points on S^1_∞, alternately attracting and repelling.*

Proof We use theorem 9.5 to get a lamination L such that $h(L) = L$. We now know a lot about L and L'; and we know that stable intervals invariant under h have endpoints which are attracting and unique repelling points in their interior.

We consider three cases, which obviously exhaust all the possibilities.

(i) h fixes the endpoints at infinity of a boundary leaf of L';

(ii) h fixed the endpoints at infinity of a non-boundary leaf of L';

(iii) h does not fix the endpoints of any leaf of L'.

In the first case, the boundary leaf γ in question is part of the boundary of a the lift $\tilde{U}$ of a principal region U. Here $\tilde{U}$ is a finite sided ideal polygon; and as discussed previously, each ideal vertex is the endpoint of precisely two leaves. It follows that h fixes the entire polygon. On the outside of each edge, there are arbitrarily close leaves, which come arbitrarily close to any point of U. It follows that each edge of the polygon cuts off an interval on the circle at infinity which is a stable interval. In this stable interval there is a unique repelling fixed point. So we have alternate repelling and attracting fixed points.

In the second case, a similar argument applies, though the picture is a bit more complicated to see. Nearby to our fixed non-boundary leaf there are again arbitrarily close leaves, and hence that it cuts off a stable interval, on both sides. So the endpoints are attracting fixed points; and there is a unique repelling fixed point on either side.

In the final case, h may or may not fix any points at infinity. If it has no fixed points at infinity, there is nothing to prove; so suppose h fixes a point x on the circle at infinity. For any lift $\tilde{\gamma}$ of a leaf of L', we can consider the open interval $U(\tilde{\gamma})$ on the circle at infinity which has the same endpoints as $\tilde{\gamma}$ and avoids x. All the sets $U(\tilde{\gamma})$ are either nested or disjoint. It follows that under h, either x is an attracting fixed point and there is precisely one other repelling fixed point, or vice versa. ∎

In fact, we can say more. Note that L' is invariant under h; and is perfect; and connected. If h has more than one attracting fixed point at infinity, then the geodesics joining consecutive attracting fixed points are in L'. But this is not just true for h; it is true for any positive power of h as well: if h^m has an attracting fixed point then it's an attracting fixed point of h also. Conversely, (well, almost conversely,) every boundary leaf of L' connects two endpoints at infinity fixed under some positive power of h: since there are only finitely many boundary leaves, there is a power of h which fixes that leaf and its orientation.

And this property actually is sufficient to define L': the boundary leaves are precisely those which have endpoints which are fixed under some positive power of h; and boundary leaves are dense. So L' is in this sense unique.

Lemma 9.15 *Make the same assumptions as above. Any lift of a strictly positive power of h has finitely many fixed points at infinity, alternately attracting and repelling. There is a unique perfect lamination L^s which contains the geodesics joining consecutive attracting fixed points of any positive power of h.*

9.7 The Proof III: The End

We are now almost done. Given a non-periodic irreducible automorphism h of S, we can obtain a stable lamination L^s, and an unstable lamination L^u (stable for h^{-1}).

These two laminations intersect transversely: we can see this from the alternating attracting and repelling fixed points of h. Their intersection points are mapped to intersection points by h. So we get rectangles bounded by arcs of

L^s and L^u, which are mapped to other rectangles under h. This allows us to straighten out h.

We call points on S equivalent if they are in the same rectangle: that is, they are either in the (closure of the) same component of $L^s - L^u$; or in the (closure of the) same component of $L^u - L^s$; or in the (closure of the) same component of $S - L^s - L^u$. The quotient by this equivalence relation is just like the Cantor function, and gives us a homeomorphic surface on which the lamination has become a foliation. Using this idea, we can finish the proof of the amazing theorem 9.2.

References

[1] Andrew J. Casson and Steven A. Bleiler, Automorphisms of Surfaces after Nielsen and Thurston