

TILING THE PROJECTIVE FOLIATION SPACE OF A PUNCTURED SURFACE

LEE MOSHER

ABSTRACT. There is a natural way to associate, to each ideal triangulation of a punctured surface a cell decomposition of the projective foliation space of the punctured surface.

A recent theme in the topology of surfaces has been the special role played by punctured surfaces, both in the study of geometric structures, and in the study of measured foliations. For example, one important result has been the construction of an equivariant “ideal cell-decomposition” of the Teichmüller space of a punctured surface, a result due to many people, including Harer, Strebel, Mumford, Epstein and Bowditch, and Penner (see, for example, [Har], [EB], and [P]). In all of these proofs an important construct has been an ideal cell-division of the punctured surface: indeed, the *cells* of Teichmüller space are in 1-1 correspondence with *isotopy classes of ideal cell-divisions* of the punctured surface.

In [M1] and [M2], ideal cell-divisions were used in the study of measured foliations on a punctured surface, in order to define conjugacy invariants of certain elements of the mapping class group of a punctured surface, and to present an algorithm for computing these invariants. It is the purpose of this work to provide a deeper foundation for the algorithm described in [M1]. In particular, we shall develop the theory of measured foliations on a punctured surface from the ground up, using the plan and several technical tools from [FLP], but utilizing ideal triangulations in place of pairs of pants decompositions for actually constructing measured foliations. Our main result here is that to each isotopy class of ideal triangulations, there is a naturally associated cell-decomposition of \mathcal{PF} , Thurston’s space of projective measured foliations. This result parallels the result of Harer and Penner that on any surface (punctured or not), to each choice of a pair of pants decomposition (plus an additional choice of data sufficient to determine a marking of the surface), there is a naturally associated cell-decomposition of \mathcal{PF} . Many of the techniques of proof used here are not new, being adaptations and reorganizations of the techniques of [FLP]; the notable exception occurs in the use of “elementary moves” on ideal triangulations to understand the overlap maps of different parameterizations of \mathcal{MF} .

Received by the editors May 25, 1985 and, in revised form, April 30, 1986.

1980 *Mathematics Subject Classification* (1985 *Revision*). Primary 57N05; Secondary 58F15, 58F18.

Partially supported by a National Science Foundation grant.

The plan of the paper is as follows. §1 gives the basic definitions. §2 describes an expanded category of measured foliations, the “partial measured foliations”, which are, perhaps, somewhat easier to use than the traditional category. §3 states the Normal Form Theorem, which says that each class in \mathcal{MF} has a unique normal form with respect to a given ideal triangulation, and that the process of constructing normal forms defines a parameterization of \mathcal{MF} which depends naturally on an ideal triangulation. This statement is analogous to the Normal Form Theorem proven in [FLP, exposé 5], which is stated in terms of pairs of pants decompositions. §4 proves the existence of normal forms, starting from the standard technical tools proven in [FLP, exposé 5]. §5 proves that the parameterization of \mathcal{MF} described above is, in fact, a homeomorphism; the tool used here is the “elementary move”, a simple combinatorial operation which generates all ideal triangulations. §6 contains the statement and proof of the Cell-Decomposition Theorem, which is quite short after all the work of the previous sections. §7 develops some additional structure of our cell-decompositions, using some results about train tracks; in particular, we give a proof using ideal triangulations and elementary moves of Thurston’s theorem that \mathcal{PF} is a manifold.

Though our cell-decompositions are different from those of Harer and Penner, one should wonder what the purpose is of reproving the existence of natural cell-decompositions of \mathcal{PF} , especially in the more restricted setting of punctured surfaces. The primary reason, mentioned above, is to use the tools developed here to present a complete classification of conjugacy classes in the mapping class group of a punctured surface, and to describe a structured and efficient algorithm for computing this classification. The results of [M1] and [M2] are a first step in this direction, although there are several shortcomings: the classification of pseudo-Anosovs described in [M1] and [M2] is not complete; and the algorithm described in [M1] is not efficient, suffering from the computational disadvantage of a built-in tree search. The correction of these problems is the subject of the sequel [M3], which uses the foundations laid down here.

In addition, using the cell-decompositions presented here, it is possible to understand the structure at infinity of the decomposition of Teichmüller space: an ideal triangulation δ of a punctured surface has associated to it a cell complex structure $\mathcal{C}(\delta)$ on \mathcal{PF} , and a set $\sigma(\delta) \subset \mathcal{T}$, where \mathcal{T} is the Teichmüller space, and $\sigma(\delta)$ is a top-dimensional simplex minus some subcomplex, embedded properly in \mathcal{T} . Using the Thurston compactification, \mathcal{PF} can be thought of as the sphere at infinity of Teichmüller space; it turns out then that $\sigma(\delta)$ is bounded at infinity by a subcomplex of $\mathcal{C}(\delta)$. This shall hopefully appear in a future work.

1. Let S be a closed, orientable surface, of genus $g \geq 0$, and let P be a finite, nonempty set of points of S . The pair (S, P) is called a *punctured surface*, and the points of P are called *punctures*. The set $S - P$ is often called the *interior* of (S, P) . The only cases we disallow are where (S, P) is a once, twice, or thrice punctured sphere. Maps, homotopies, and isotopies of (S, P) will often be referred to as maps, homotopies, and isotopies of S rel P . In addition, continuous maps $f: X \rightarrow S$ will

often be altered by ambient isotopy rel P ; this shall occasionally be shortened to “isotopy rel P ” or just “isotopy”.

The mapping class group of (S, P) , denoted $\mathcal{M}(S, P)$, is the group $\pi_0(\text{Homeo}_+(S, P)) \approx \text{Homeo}_+(S, P)/\text{Homeo}_0(S, P)$, where $\text{Homeo}_+(S, P)$ denotes orientation preserving homeomorphisms of (S, P) , and $\text{Homeo}_0(S, P)$ denotes homeomorphisms of (S, P) isotopic rel P to the identity.

A *measured foliation* on (S, P) is a codimension 1 singular foliation equipped with a transverse measure. The type of singularity which can occur at a point x is an *n-pronged singularity*: this means that, for some complex coordinate z near x , f is foliated by the horizontal trajectories of the meromorphic quadratic differential $z^{n-2}dz^2$ (holomorphic if $n \geq 2$). An allowable singularity *at a puncture* is an *n-pronged singularity* with $n \geq 1$, and an allowable singularity *in the interior* is an *n-pronged singularity* with $n \geq 3$. (By convention, a puncture is always considered to be a singularity, even if it has 2 prongs.) A *transverse measure* on such a foliation f is an assignment of a Borel measure with positive total weight to every embedded interval transverse to f ; this measure is invariant under isotopy parallel to leaves of f , and it is preserved under restriction to subintervals. If α is an interval embedded transversely to f , we use $\int_\alpha f$ to denote the integral over α of the measure defined on α by f .

A *leaf segment* ℓ of a measured foliation f is the image of an embedding of the closed interval I into a leaf of f , such that $\text{int}(I)$ maps to regular points of f . The endpoints of ℓ may be either regular points, or singularities; if both endpoints are singularities, then ℓ is called a *singular leaf segment*. A *leaf cycle* of f is a 1-complex which is a union of leaf segments of f . We say that a leaf cycle c is *closed* if the only degree one vertices of c occur at punctures. A measured foliation f is *arational* if f has no closed leaf cycles.

There is an equivalence relation on measured foliations, called *Whitehead equivalence*, generated by isotopy and Whitehead moves; a *Whitehead move* on a measured foliation f is the operation of collapsing a singular leaf segment ℓ of f to a point, resulting in a new measured foliation f' ; the only restriction on ℓ is that at least one of its endpoints be in the interior of (S, P) (a more precise formulation of Whitehead moves is given below). The equivalence class of f is denoted $\{f\}$, and the collection of equivalence classes, with the empty foliation included, is denoted $\mathcal{MF}(S, P)$, or more briefly, just \mathcal{MF} (we shall denote the empty foliation as 0). Note that for a measured foliation, the property of being arational is preserved under Whitehead equivalence. Therefore we can properly speak of an arational class in \mathcal{MF} .

A measured foliation can also be multiplied by a positive scalar r , by defining $\int_\alpha r \cdot f = r \cdot \int_\alpha f$, for every measured foliation f , embedded interval α transverse to f , and $r \geq 0$. This scalar multiplication respects Whitehead equivalence, and so descends to a scalar multiplication on $\mathcal{MF} - \{0\}$. This gives a projectivization map $p: \mathcal{MF} - \{0\} \rightarrow \mathcal{PF}$; the projective class of a measured foliation class $\mathcal{F} = \{f\}$ is denoted $p(\mathcal{F})$, or $p\{f\}$. Note that the group $\mathcal{M}(S, P)$ acts on both \mathcal{MF} and \mathcal{PF} , via the action of $\text{Homeo}^+(S, P)$ on measured foliations; these actions commute with p .

We shall have need of the following more precise formulation of Whitehead moves. Given a measured foliation f and a singular leaf segment ℓ of f such that at least one endpoint of ℓ is in $S - P$, there is a map $\varphi: (S, P) \rightarrow (S, P)$, homotopic to the identity, which takes ℓ to a point and acts homeomorphically on the complement of ℓ . It follows that $\varphi(f)$ is a measured foliation on (S, P) , which we say is obtained from f by a *Whitehead move*. The map φ is called the *collapsing map* of the Whitehead move. If the endpoints of ℓ are m and n -pronged singularities of f , then $\varphi(\ell)$ is an $(m + n - 2)$ -pronged singularity of $\varphi(f)$.

Here are several technical results needed for the proof of the Normal Form Theorem, which are proven in [FLP, exposé 5]; they are restated here in the form appropriate for punctured surfaces. The proofs are exactly the same as in the closed surface case.


EULER - POINCARÉ INDEX FORMULA. *Let f be a singular foliation of a closed surface M , where the singularities can be k -pronged singularities for $k \geq 0$ (a 0-pronged singularity has a neighborhood foliated by concentric circles). Then $\chi(M) = \sum(\text{index}(s))$, summed over the singularities s of f , where for an n -pronged singularity s , $\text{index}(s) = 1 - n/2$.*

STABILITY LEMMA. *Let f be a measured foliation on (S, P) . Let ℓ be a leaf cycle of f which is homeomorphic to the closed interval, with $\partial\ell = \{x_0, x_1\}$. Suppose that on one side of ℓ , the side, the foliation is regular (i.e. that side is foliated locally like horizontal lines in the upper half-plane). Let α, β be a pair of transverse arcs to f such that: $x_0 \in \partial\alpha$ and $x_1 \in \partial\beta$; $\text{int}(\alpha), \text{int}(\beta)$ and $\text{int}(\ell)$ are pairwise disjoint; and α, β each emanate from the nonsingular side of ℓ . Then there exists a continuous mapping $H: I \times I \rightarrow S$ such that: $H|_{\text{int}(I) \times \text{int}(I)}$ is an embedding; $H(I \times 0) = \ell$; $H(0 \times I) \subset \alpha$; $H(1 \times I) \subset \beta$; $H(I \times t)$ is a nonsingular leaf segment of f for each $t \in \text{int}(I)$; if $\partial\alpha \cap \partial\beta = \emptyset$, $H(I \times 1)$ is a leaf cycle containing either an endpoint of α or β , or a singularity of f ; if $\partial\alpha \cap \partial\beta \neq \emptyset$, then $H(I \times 1)$ is the point $\partial\alpha \cap \partial\beta$. Moreover, the image of H is uniquely determined by these conditions.*

We shall also need a version of the Poincaré Recurrence Theorem; the statement is given at the end of the next section, after we have introduced some new notions.

2. For many purposes, it turns out that measured foliations are not so convenient to work with. What usually happens is that a certain measured foliation class might assign zero weight to all curves on a certain subsurface of S , and one might want to be able to see this subsurface, or at least a portion of it. This inconvenience can be remedied in several ways. The most common is to put a hyperbolic structure on $S - P$, and to straighten the leaves of a measured foliation f , giving a “geodesic lamination”, which depends only on the class of f . In order to avoid the intricacies of the hyperbolic category, we shall follow a different tack, using a variation on the “partial measured foliations” defined in exposé 9 of [FLP].

A partial measured foliation f on a closed surface S , as defined in [FLP], is (roughly) a measured foliation supported on a subsurface S' , such that each component c of $\partial S'$ is an essential simple closed curve on S , and c is a cycle of



A pinch point

FIGURE 1

leaves of f . On a punctured surface (S, P) , the support S' of a partial measured foliation on (S, P) is allowed to have a slightly more general topology at a puncture—specifically, near a puncture p , S' can look like a union of sectors of a measured foliation. (Note: All our definitions in this section apply equally well to the closed surface case—just allow the puncture set P to be empty.)

Here is the precise definition. A *pinched subsurface* of (S, P) is a nonempty closed subspace $S' \subset S$ such that $S' - P$ is a subsurface of $S - P$ (with boundary, if S' is a proper subspace), and S' contains no isolated punctures. Thus, if $p \in \partial S' \cap P$, then near p , S' looks like a union of finitely many pie wedges. If there are ≥ 2 wedges which are disjoint except at p , we will say that p is a *pinch point* of S' . See Figure 1.

A *relative boundary component* of a pinched subsurface S' is defined to be the closure c of a boundary component of $S - S'$. Thus, c is either a simple closed curve in $S - P$, or c is the image of a continuous map $(I, \partial I, \text{int}(I)) \rightarrow (S, P, S - P)$ which is 1-1 on $\text{int}(I)$. To say that c is *essential* means:


- (1) If $c \approx S^1$ and $c \cap P = \emptyset$, then c does not bound a nonpunctured or a once-punctured disc.
- (2) If $c \approx S^1$ and $c \cap P = \{p\}$, then c does not bound a disc with no interior punctures (this is just the definition of an “essential arc” given in §3).

In particular, if $c \neq S^1$, then c is essential.

If each relative boundary component of a pinched subsurface S' is essential, we will say that S' is *essential*.

A *partial measured foliation* on (S, P) is a measured foliation f defined on some essential pinched subsurface S' of (S, P) , such that each relative boundary component of S' is a cycle of leaves of f . Moreover, if $x \in \partial S'$ and x is not a pinch point, then either f is regular at x , or f has a boundary singularity at x : that is, near x , f looks like the foliation of the upper half-plane by the horizontal trajectories of $z^{2n-2}dz^2$ ($n \geq 1$). Finally, if $p \in \partial S' \cap P$ is a pinch point of S' , then near p , f looks like a union of sectors of an n -pronged singularity at p ($n \geq 1$); we shall call this a *pinch singularity*. See Figure 2.

We call S' the *support* of f , denoted $\text{supp}(f)$; S' may be empty, or all of S , or any arbitrary essential pinched subsurface. Unless the context suggests otherwise, the term “measured foliation” shall be reserved to mean a partial measured foliation



A pinch singularity

FIGURE 2

whose support is all of S ; to clarify certain situations, we shall occasionally refer to a measured foliation as a *total* measured foliation.

In constructing partial measured foliations, we shall make use of the following lemma (notational note: the closure operator is denoted “ cl ”):

PUNCTURE SINGULARITY LEMMA. *Let S' be a pinched subsurface (not necessarily essential); suppose $D \subset \text{cl}(S - S')$ is a once-punctured disc such that $D \cap S'$ is a union of finitely many closed arcs contained in $\partial D \cap \partial S'$ (this includes the case where $D \cap S' = \partial D$). Let $\varphi: (S, P) \rightarrow (S, P)$ be a continuous map which takes D to $p \in D$, and is injective on $S - D$. Let f be a measured foliation on S' such that $\partial S'$ is a leaf cycle of f , and each pinch point on S' is a pinch singularity of f ; moreover, if $D \cap S' = \partial D$, it is required that f has a singularity on ∂D . Then $\varphi(f)$ is a measured foliation on $\varphi(S')$ such that $\partial(\varphi(S'))$ is a leaf cycle of $\varphi(f)$, and each pinch point of $\varphi(S')$ is a pinch singularity of $\varphi(f)$.*


The proof of this lemma is an elementary construction, and is omitted. The point of this lemma is: if $D \cap S' = \partial D$, then $\varphi(f)$ has an n -pronged singularity at p , where $n \geq 1$ is guaranteed by the condition that f has a singularity on ∂D ; if $D \cap S'$ is a single arc, then $\varphi(f)$ has a boundary singularity at p ; and if $D \cap S'$ is two or more arcs, then $\varphi(f)$ has a pinch singularity at p (see Figure 3).

We now need a generalization of the concept of an “enlargement” (élargissement), presented in [FLP, exposé 5]. This is a method of assigning, to each partial measured foliation on (S, P) , a total measured foliation whose class is well defined.

Let f be a partial measured foliation on (S, P) , with support S' , and suppose S' is not all of S . Let Σ be a spine of $\text{cl}(S - S')$ relative to $P' = P \cap \text{cl}(S - S')$. This means that Σ is a finite 1-complex contained in $\text{cl}(S - S')$, such that $P' \subset \Sigma$, and $\text{cl}(S - S')$ has a strong deformation retraction to Σ . We also impose the condition that Σ has no degree 1 vertices, except possibly at some singularities of f on $\partial S'$, or possibly at some punctures in P' . Such a spine can always be found (see Figure 4).


Then there exists a continuous mapping $\varphi: (S, P) \rightarrow (S, P)$, homotopic to the identity, called an *enlarging map* (Figure 4) for f , such that:

- (1) $\varphi|_{\text{cl}(S - S')}$ is a strong deformation retraction onto Σ .
- (2) $\varphi|_{\text{int}(S')}$ is a homeomorphism onto $S - \Sigma$;



The Puncture Singularity Lemma

FIGURE 3



An enlarging map of a partial measured foliation

FIGURE 4

(3) $\varphi|_{\partial S'}$ from $\partial S'$ onto Σ is a topological immersion, except at pinch points of $\partial S'$ and degree one vertices of Σ .

Thus, $\varphi(f)$ is a foliation with singularities of (S, P) , equipped with a transverse measure. It is easy to see that the singularities are n -pronged singularities, with $n \geq 0$. But the condition that S' is essential guarantees that no component of Σ is a point, therefore $\varphi(f)$ has no 0-pronged singularities. Also, the condition on the degree 1 vertices of Σ guarantees that $\varphi(f)$ has no interior 1-pronged singularities. It follows that $\varphi(f)$ is a measured foliation on (S, P) , called an *enlargement* of f .

Although $\varphi(f)$ depends on the choice of Σ and φ , the class of $\varphi(f)$ is well defined. The proof of this fact is a straightforward generalization of the arguments found in [FLP, exposé 5, §3]; the idea is to show by techniques of *PL* topology that

any two spines of a punctured surface with boundary are equivalent up to isotopy and “Whitehead moves”; one could also use the techniques of “elementary moves”, presented in §6.

Whitehead moves on partial measured foliations are defined almost exactly as for total measured foliations, though we must be slightly more careful. Suppose f is a partial measured foliation and ℓ is a singular leaf segment of f . As before, we exclude the case where ℓ connects two punctures. We also put the following condition on ℓ : if $\partial\ell \subset \partial(\text{supp}(f)) - P$, then we require that $\ell \subset \partial(\text{supp}(f))$; this excludes the case where $\ell \not\subset \partial(\text{supp}(f))$, but ℓ connects two boundary singularities of f lying in $S - P$; without excluding that case, it would be possible to create pinch singularities in $S - P$, which we do not want. A collapsing map on ℓ is now defined just as for measured foliations, resulting in a partial measured foliation f' which is said to be obtained from f by a Whitehead move.

Whitehead equivalence is now defined, as in the case of total measured foliations, as the equivalence relation generated by isotopy and Whitehead moves. If f, g are partial measured foliations which are Whitehead equivalent, then clearly there are enlargements f', g' of f, g that are Whitehead equivalent. It follows that there is an equivalence relation on partial measured foliations generated by isotopies, Whitehead moves, and enlargements, and the collection of equivalence classes is in 1-1 correspondence with \mathcal{MF} , under the map induced by the inclusion of the set of measured foliations into the set of partial measured foliations. Also, this bijection is natural with respect to the action of $\mathcal{M}(S, P)$. Therefore, we shall alter our notation and use \mathcal{MF} to refer to equivalence classes of partial measured foliations, under the above equivalence relation (note that this equivalence relation on partial measured foliations is stronger than Whitehead equivalence). The associated projectivization will similarly be denoted \mathcal{PF} .

One important reason for introducing partial measured foliations is that it clarifies the connection between measured foliations and measured geodesic laminations. In particular, given a partial measured foliation f whose components are all either minimal partial foliations or (nonpairwise isotopic) annuli (we call f a “complete dissection”), there is a 1-1 correspondence between the components of f and the components of the measured geodesic lamination corresponding to f . This kind of partial measured foliation is often the most convenient to work with in applications, so we shall develop the general theory here before stating the Normal Form Theorem.

Some notation: given a partial measured foliation f and a union K of components of $\text{supp}(f)$, we use $f|K$ to denote the partial measured foliation with support K which agrees with f on K .

Given $\mathcal{F} \in \mathcal{MF}$ and $f \in \mathcal{F}$, we say that f is a *complete dissection* of \mathcal{F} if, for every component S' of $\text{supp}(f)$, $\partial S'$ consists of simple closed curves in $S - P$ (which are necessarily closed leaf cycles of f), and either $f|S'$ is arational (meaning that the boundary components are the only closed leaf cycles of $f|S'$), or S' is an annulus embedded in $S - P$. In the first case, we say that $f|S'$ is *partial arational*; in the second case, we say that $f|S'$ is *annular*. If $f|S'$ is annular, it follows from the Stability Lemma that $f|S'$ is a foliation by leaves parallel to the boundary. We

also impose the condition that two distinct annular components of $\text{supp}(f)$ cannot be isotopic; one can easily see that this condition is necessary in order for complete dissections to be unique, as stated in the following lemma:

COMPLETE DISSECTION LEMMA. *Every $\mathcal{F} \in \mathcal{MF}$ is represented by a complete dissection, unique up to Whitehead equivalence.*

This fact should be viewed in parallel with the obvious fact that each $\mathcal{F} \in \mathcal{MF}$ is represented by a total measured foliation, unique up to Whitehead equivalence. Thus, complete dissections give us an alternate way of using Whitehead equivalence to decide identity in \mathcal{MF} .

PROOF. First we prove existence of complete dissections. Given a measured foliation $f_0 \in \mathcal{F}$, let Σ be the union of all closed leaf cycles of f_0 containing a singularity. Note that Σ does *not* contain smooth closed leaves. It follows that Σ is itself a closed leaf cycle of f_0 . Moreover, Σ is a maximal closed leaf cycle with respect to the property that every component of Σ contains a singularity.

Let f_1 be the partial measured foliation obtained by cutting f_0 along Σ ; f_1 is characterized (up to isotopy rel P) by the property that there exists a closed regular neighborhood N of Σ , and a mapping $\varphi: (S, P) \rightarrow (S, P)$ homotopic to the identity, such that $\varphi|_N$ is a deformation retraction onto Σ , φ takes $\text{int}(S - N)$ homeomorphically to $S - \Sigma$, and $\varphi(f_1) = f_0$.

$\text{Supp}(f_1)$ is the closure of the complement of N . Thus, $\partial(\text{supp}(f_1))$ is a collection of simple closed curves in $S - P$. Moreover, each of these curves is essential, for otherwise, f_0 has a closed leaf cycle bounding a nonpunctured or once-punctured disc. It follows that f_1 is a partial measured foliation on (S, P) .

Let S' be a component of $\text{supp}(f_1)$. Suppose S' is not an annulus, and $f|_{S'}$ is not arational. We shall arrive at a contradiction. Since $f|_{S'}$ is not arational, it has a closed leaf cycle ℓ , which does not intersect $\partial S'$. We can assume that ℓ is connected.

If ℓ contains a singularity, then $\varphi(\ell)$ is a closed leaf cycle of f_1 containing a singularity, but then $\varphi(\ell) \cap \Sigma = \emptyset$, contradicting maximality of Σ .

If ℓ contains no singularity, then ℓ is a smooth closed leaf of f_1 . It follows from the Stability Lemma that ℓ is a leaf of some maximal foliated annulus A , embedded except at its boundary, where it is immersed. If ∂A is not contained in $\partial S'$, then ∂A has a component which is a closed leaf cycle of $f_1|_{S'}$; by maximality of A , this cycle must contain a singularity, but as above, this leads to a contradiction of the maximality of Σ . If ∂A is contained in $\partial S'$, then A is embedded and $A = S'$; this contradicts the assumption that S' is not an annulus.

We now prove that complete dissections are unique, up to Whitehead equivalence. This is clear in the case that \mathcal{F} is arational, for in that case, any two representatives of \mathcal{F} are measured foliations, and are Whitehead equivalent by virtue of being in the same class in \mathcal{MF} . So for the remainder of the proof, we shall consider the case where \mathcal{F} is not arational.

Suppose f_1 and f_2 are two complete dissections of \mathcal{F} . Then there are enlargements f'_1 and f'_2 of f_1 and f_2 which are Whitehead equivalent. So we can assume, without loss of generality, that there is a map $\varphi': (S, P) \rightarrow (S, P)$, homotopic to the

identity, such that $\varphi'(f'_1) = f'_2$, and either φ' is a homeomorphism, or φ' is the collapsing map of a Whitehead move. We shall prove in this case that there is a map $\varphi: (S, P) \rightarrow (S, P)$, homotopic to the identity, such that $\varphi(f_1) = f_2$, and either φ is a homeomorphism, or φ is the collapsing map of a Whitehead move.

For $i \in \{1, 2\}$, let $\Psi_i: (S, P) \rightarrow (S, P)$ be a collapsing map taking f_i to f'_i , and let Σ_i be the spine onto which $S - \text{supp}(f_i)$ is collapsed by Ψ_i . Thus, Σ_i is a leaf cycle of f'_i . Also, each component of Σ_i contains a singularity of f'_i : for noncircular components this is obvious; for circular components, this follows because of the condition on complete dissections that each complementary non-punctured annulus contain a singularity on at least one boundary component. Finally, note that Σ_i is the largest leaf cycle of f'_i such that each component contains a singularity: this is because the only closed leaf cycles of f_i are either smooth closed leaves in annular components of $\text{supp}(f_i)$, or are components of $\partial(\text{supp}(f_i))$. This property of leaf cycles in measured foliations is evidently preserved under Whitehead equivalence, so $\varphi'(\Sigma_1) = \Sigma_2$.

Now note that $\Psi_i|_{\text{supp}(f_i)}$ is a local homeomorphism onto S (except at the inverse image of a degree one vertex of Σ_i ; this is not a problem, however). Together with the fact that $\varphi'(\Sigma_1) = \Sigma_2$, we see that there is a coherent way to define Ψ_2^{-1} on $\varphi' \circ \Psi_1(U)$, for each U in some finite open cover of $\text{supp}(f_1)$ —coherent in the sense that Ψ_2^{-1} is well defined by restriction for any open set dominated by this cover. This allows us to define $\varphi|_{\text{supp}(f_1)}$ by $\varphi = \Psi_2^{-1} \circ \varphi' \circ \Psi_1$; it is now easy to see how to extend φ to all of S so that φ is the required Whitehead equivalence. Q.E.D.

We close this section with a version of the Poincaré Recurrence Theorem which is useful for what follows. The proof is exactly the same as the proof given in [FLP].

POINCARÉ RECURRENCE THEOREM. *Let f be a measured foliation on (S, P) , and let S' be a pinched subsurface of S such that $\partial S'$ is a union of leaf cycles of f and arcs transverse to f . Then every leaf ℓ of f , contained in S' and originating at $\partial S'$, eventually either runs into a singularity of f in S' or returns to $\partial S'$. Since there are only finitely many singular leaves, it follows that almost every such leaf ℓ returns to $\partial S'$.*

3. On a closed surface, essential simple closed curves are used to parameterize \mathcal{MF} . As presented in [FLP], there are two steps to this process. First one maps \mathcal{MF} into $\mathbf{R}_{\geq 0}^{\mathcal{S}}$, where \mathcal{S} is the set of isotopy classes of simple closed curves; this mapping is accomplished using intersection numbers. Then one picks out a finite set of elements of \mathcal{S} which has some geometrical significance on the surface, and shows that this set parameterizes \mathcal{MF} ; for example, \mathcal{MF} can be parameterized by intersection numbers with the curves of two pairs of pants decompositions which together fill up the surface.

On a punctured surface one can also carry out this program using simple closed curves; this is done in [HP]. But there is another construct which is not available on a closed surface, namely, an “ideal arc”, i.e. an essential simple arc with both ends in the puncture set; and the entire program can be carried out with ideal arcs alone. Thus, letting \mathcal{H} represent the set of isotopy classes of ideal arcs, we shall map \mathcal{MF} into $\mathbf{R}_{\geq 0}^{\mathcal{H}}$ using intersection numbers; then we shall use the arcs of an “ideal

triangulation" of (S, P) to parameterize \mathcal{MF} . We start with the definitions concerning ideal arcs.

An *ideal arc* on (S, P) is the image h of a continuous mapping $(I, \partial I, \text{int}(I)) \rightarrow (S, P, S - P)$, which is 1-1 on $\text{int}(I)$, and satisfying the "essentiality" condition that there does *not* exist an embedded disc $D \subset S$ such that $\partial D = h$ and $\text{int}(D) \cap P = \emptyset$. Two ideal arcs h, h' are *isotopic* if there exists $\varphi \in \text{Homeo}_0(S, P)$ such that $\varphi(h) = h'$. \mathcal{H} shall denote the set of all isotopy classes of ideal arcs on (S, P) .

For a fixed partial measured foliation f and ideal arc h , let $\int_h f$ denote the total mass of h , measured with respect to f . For $\mathcal{F} \in \mathcal{MF}$, let $\langle \mathcal{F}, h \rangle = \inf(\int_h f)$, taken over all $f \in \mathcal{F}$. Note that, if h, h' are isotopic, then $\langle \mathcal{F}, h \rangle = \langle \mathcal{F}, h' \rangle$. Thus, given $h \in \mathcal{H}$, we can unambiguously define $\langle \mathcal{F}, h \rangle$ to be $\langle \mathcal{F}, h \rangle$, for any $h \in h$. For each $\mathcal{F} \in \mathcal{MF}$, define $\nu_{\mathcal{F}} \in \mathbf{R}_{\geq 0}^{\mathcal{H}}$ by the formula $\nu_{\mathcal{F}}(h) = \langle \mathcal{F}, h \rangle$. This gives a mapping $\nu: \mathcal{MF} \rightarrow \mathbf{R}_{\geq 0}^{\mathcal{H}}$. Note that the map ν commutes with the actions of $\mathcal{M}(S, P)$ on \mathcal{MF} and on $\mathbf{R}_{\geq 0}^{\mathcal{H}}$, i.e. for each $\Phi \in \mathcal{M}(S, P)$, $\mathcal{F} \in \mathcal{MF}$, and $h \in \mathcal{H}$, we have $\langle \mathcal{F}, h \rangle = \langle \Phi(\mathcal{F}), \Phi(h) \rangle$.

We now topologize \mathcal{MF} by giving it the smallest topology with respect to which ν is continuous. As we shall see in the Normal Form Theorem, ν is 1-1; thus, the topology on \mathcal{MF} is chosen so that ν is a homeomorphism onto its image.

Let f be a partial measured foliation and h an ideal arc. We now wish to develop a geometric condition on f and h which guarantees that $\int_h f$ assumes its minimal value, over the class of f . To get this condition, we need to know that $\partial(\text{supp}(f))$ and h intersect nicely: given a pinched subsurface S' of (S, P) and an ideal arc h , we say that $\partial S'$ and h *intersect efficiently* if there does *not* exist a closed disc $D \subset S$ that satisfies the following conditions:

- (1) $\partial D = \alpha \cup \beta$, where α, β are closed intervals.
- (2) $\alpha \cap \beta = \partial\alpha = \partial\beta$.
- (3) $\alpha \subset h$.
- (4) $\beta \subset \partial S'$.
- (5) $D \cap P \subset \alpha \cap \beta$.


(6) D is embedded, except possibly when $\partial\alpha \subset P$, in which case $\partial\alpha$ is allowed to be a single point, arising from an identification of two points on the boundary of the disc D .

Such a disc D , if it exists, is called a *compressing disc* for h and $\partial S'$. Thus, h and $\partial S'$ intersect efficiently if and only if there is no compressing disc for h and $\partial S'$.

Note. Condition (5) is stated so that in Figure 5, h and $\partial S'$ do intersect efficiently, but in Figure 6, h and $\partial S'$ do not intersect efficiently. In this latter case, h is isotopic to a relative boundary component of S' ; thus, according to the definition given, if h is isotopic rel P to a relative boundary component of S' , then h and $\partial S'$ intersect efficiently if and only if h actually is a relative boundary component of S' .


MINIMAL INTERSECTION PROPOSITION. *Given $\mathcal{F} \in \mathcal{MF}$ and $f \in \mathcal{F}$, and given an ideal arc h , suppose that h is transverse to f and h intersects $\partial(\text{supp}(f))$ efficiently. Then $\int_h f = \langle \mathcal{F}, h \rangle$.*

PROOF. If $\int_h f = 0$, this is clearly true, so we will assume that $\int_h f \neq 0$. First we show that $\int_h f$ is minimal within the isotopy class of f .



$h \cap \partial S'$ is efficient

FIGURE 5



$h \cap \partial S'$ is not efficient

FIGURE 6

LEMMA. Under the same hypotheses as above, if f' is isotopic to f , then $\int_h f \leq \int_h f'$.

PROOF. This is equivalent to the following: again under the same hypotheses as the Minimal Intersection Proposition, if h' is an ideal arc isotopic to h , then $\int_h f \leq \int_{h'} f$. The proof now follows very closely the proof of Proposition II.3 of exposé 5 of [FLP].

We can assume h and h' intersect transversely; thus, $\text{int}(h)$ and $\text{int}(h')$ intersect in finitely many points. The proof is by induction on the number of intersection points.

If $\text{int}(h) \cap \text{int}(h') = \emptyset$, then h and h' together bound a pinched subsurface S' , and $P \cap S' = \partial h = \partial h'$. Note that S' is either a disc or a disc with two boundary points identified. We show that almost every leaf ℓ of f entering S' through h must leave through h' ; it follows that $\int_h f \leq \int_{h'} f$. By Poincaré Recurrence almost every such leaf ℓ returns to $h \cup h'$. Suppose there is a leaf ℓ of f entering and leaving S' through h . Let α be the subarc of h cut off by ℓ . Let S'' be the disc bounded by ℓ and α . By hypothesis, $\ell \not\subset \partial(\text{supp}(f))$; moreover, no leaf of f intersecting $\text{int}(S'')$ can be in $\partial(\text{supp}(f))$. Since α intersects f transversely, this implies that $S'' \subset \text{int}(\text{supp}(f))$, so S'' is foliated by f , with ℓ as a leaf, and α transverse to the foliation. Since $S'' \cap P = \emptyset$, this foliation has no positive index singularities. But

the double of S'' then gives a foliation of the sphere with two one-pronged singularities and no other positive index singularities, in contradiction with the Euler-Poincaré Index Formula (see the Foliated Disc Lemma in §5).

If $\text{int}(h) \cap \text{int}(h') \neq \emptyset$, there exists an embedded disc $D \subset S$, $\text{int}(D) \cap P = \emptyset$, $\partial D = \alpha \cup \beta$, where α, β are closed intervals, $\partial\alpha = \partial\beta = \alpha \cap \beta$, $\alpha \subset h$, and $\beta \subset h'$. Again by the conditions on f and h , almost every leaf entering D through α must leave through β . Then we can isotope h' across D , resulting in an arc h'' having at least one fewer intersection with h , and such that $\int_{h''} f \leq \int_{h'} f$. Now apply the induction hypothesis. Q.E.D.


We shall use $[f]$ to denote the isotopy class of f , and $\langle [f], h \rangle = \inf \int_h f'$, where the infimum is taken over $f' \in [f]$.

The remainder of the proof of Minimal Intersection is to show that if f and f' are related by a Whitehead move or an enlargement, then $\langle [f], h \rangle = \langle [f'], h \rangle$.

First suppose that f is obtained from f' by a Whitehead move. Then f can be chosen in its isotopy class so that $\int_h f \leq \langle [f], h \rangle + \epsilon$, and so that h misses the singularity resulting from collapsing a leaf of f' . f' can be chosen in its isotopy class so that the collapsing map taking f' to f has support disjoint from h . Thus, $\int_h f' = \int_h f \leq \langle [f], h \rangle + \epsilon$. Since this is true for all ϵ , $\langle [f'], h \rangle \leq \langle [f], h \rangle$.

To get the reverse inequality, choose f' in its isotopy class so that $\int_h f' \leq \langle [f'], h \rangle + \epsilon/2$. Let ℓ be the leaf segment of f' that is collapsed in producing f . If h does not intersect ℓ , leave f' alone. Otherwise, recall that at least one end of ℓ is not at a puncture. Therefore, f' can be altered by isotopy rel P so that h and ℓ no longer intersect; this isotopy can be chosen so that $\int_h f'$ is altered by no more than $\epsilon/2$. Thus, we can assume that h does not intersect ℓ , and $\int_h f' \leq \langle [f'], h \rangle + \epsilon$. The collapsing map taking f' to f can now be chosen to have support disjoint from h , so $\int_h f = \int_h f' \leq \langle [f'], h \rangle + \epsilon$; therefore, $\langle [f], h \rangle \leq \langle [f'], h \rangle$.

Now suppose that f is obtained from f' by enlargement. Let Σ be a spine for $S - \text{supp}(f')$. Thus, under the collapsing map φ , Σ becomes a union of leaf segments in f . Choose f and Σ so that h intersects Σ transversely, and $\int_h f \leq \langle [f], h \rangle + \epsilon$. The partial measured foliation f' can be chosen within its isotopy class so that $\text{cl}(S - \text{supp}(f'))$ is contained in any neighborhood of Σ ; moreover, the collapsing map φ taking f to f' can be chosen to have support in any neighborhood of $\text{cl}(S - \text{supp}(f'))$, and then to map h to itself, collapsing each component of $h \cap \text{cl}(S - \text{supp}(f'))$ to the point where that component intersects Σ . See Figure 7.



φ collapses a component of $h \cap (S - \text{supp}(f'))^c$ to a point

FIGURE 7

With these choices, $\int_h f' = \int_h f$, so $\langle [f'], h \rangle \leq \langle [f], h \rangle$.

To get the reverse inequality, choose f' so that h intersects $\partial(\text{supp}(f'))$ transversely, and so that $\int_h f' \leq \langle [f'], h \rangle + \varepsilon/2$.

If h does not intersect $\partial(\text{supp}(f'))$ efficiently, choose an innermost compressing disc D for h and $\partial(\text{supp}(f'))$, and write $\partial D = \alpha \cup \beta$, where α, β are closed intervals, $\alpha \cap \beta = \partial\alpha = \partial\beta$, $\alpha \subset h$, $\beta \subset \partial(\text{supp}(f'))$. Note that no punctures lie along $\text{int}(\beta)$. Thus we can isotope f' across D , removing one compressible component of $h \cap \text{cl}(S - \text{supp}(f'))$, and increasing $\int_h f'$ by an arbitrarily small amount.

After performing this operation finitely many times, we can assume h and $\partial(\text{supp}(f'))$ intersect efficiently, and $\int_h f' \leq \langle [f'], h \rangle + \varepsilon$.

The collapsing map φ taking f' to f can now be chosen to take h to itself, and on each component of $h \cap \text{cl}(S - \text{supp}(f'))$, φ restricts to a retraction map onto a point. Thus, $\int_h f = \int_h f'$, so $\langle [f], h \rangle \leq \langle [f'], h \rangle$.

To prove the Minimal Intersection Proposition, note that we can get from f to f' within the same class in \mathcal{MF} by a finite sequence of isotopies, Whitehead moves, inverse Whitehead moves, enlargements, and inverse enlargements. As was shown above, if f achieves the minimum of $\int_h f$ within its isotopy class, then $\int_h f \leq \int_h f'$, so $\int_h f = \langle \{f\}, h \rangle$. By the first lemma, if f and h satisfy the conditions of the proposition, $\int_h f$ achieves the minimum in the isotopy class of f , so $\int_h f = \langle \{f\}, h \rangle$. Q.E.D.

Now we describe ideal triangulations, and the way in which they are used to parameterize \mathcal{MF} . An *ideal triangulation* of (S, P) is a cell division δ of S such that:

- (1) The set of 0-cells of δ is P .
- (2) Each 1-cell of δ is an ideal arc.
- (3) Each 2-cell T of δ is a *triangle*, i.e. T is obtained by attaching a Euclidean triangle τ to the 1-skeleton of δ , mapping each vertex of τ to a puncture, and each side of τ to an arc of δ .

We remark that an Euler characteristic argument can be employed to show that: given a 1-cell h of a cell division δ satisfying (1) and (3) only, h necessarily satisfies the essentiality condition for ideal arcs, and so h is an ideal arc; thus, δ is an ideal triangulation.

A triangle T of δ may have its vertices identified in an arbitrary manner. In addition, T may have a pair of sides identified (in an orientation preserving way, of course), resulting in a single ideal arc h . If the latter happens, we say that T is a *puncture piece*, and that h is the *interior arc* of the puncture piece T . Also, the puncture p which is incident to h but not incident to the remaining side of T is called the *interior puncture* of the puncture piece T . See Figure 8.


The existence of puncture pieces leads to certain pathologies in the theory. Of course, in order for there to be any puncture pieces, there must be at least two punctures. Thus, the once-punctured case is the nicest, being free of the pathologies.

Now let δ be an ideal triangulation of (S, P) , and let $H(\delta)$ denote the collection of arcs of δ . Notice that if $h \neq h' \in H(\delta)$, then h and h' are not isotopic. To see why, notice that δ lifts to an equivariant ideal triangulation $\tilde{\delta}$ of $\mathbf{H}^2 = \{\text{universal}$



A puncture piece

FIGURE 8



Local models for a foliation in normal form

FIGURE 9

cover of $S - P$ }, in such a way that no arc of $\tilde{\delta}$ has its ends at the same point of $S_\infty^1 = \partial\mathbf{H}^2$; for otherwise, there would be an arc of δ which violates the essentiality property. It follows that if $h \neq h' \in H(\delta)$, and if \tilde{h}, \tilde{h}' are any lifts of h, h' (resp.), then \tilde{h}, \tilde{h}' do not have the same endpoints on S_∞^1 . From this it follows that h and h' are not isotopic.

It now follows that there is a well-defined projection map $q_\delta: \mathbf{R}_{\geq 0}^{\mathcal{L}} \rightarrow \mathbf{R}_{\geq 0}^{H(\delta)}$. The main objective of the Normal Form Theorem is to show that $q_\delta \circ \nu: \mathcal{MF} \rightarrow \mathbf{R}_{\geq 0}^{H(\delta)}$ is injective, and to identify the image. To accomplish this, we shall give a construction for partial measured foliations in “normal form” with respect to δ , and show that these normal forms can be used to completely understand the map $q_\delta \circ \nu$.

Let δ be an ideal triangulation of (S, P) and f a partial measured foliation on (S, P) . We say that f is in *normal form* with respect to δ if, for each triangle T of δ , the underlying foliation of f restricted to T looks like one of the four pictures in Figure 9.

Here is a formal description for each of the cases.

In case 0, we start with a regular Euclidean triangle τ , and three “subdividing segments” which connect the barycenter of τ to the barycenters of the three sides of τ ; these segments subdivide τ into three quadrilaterals denoted Q_1, Q_2, Q_3 . Now let A be a regular Euclidean triangle with one horizontal side s , and choose homeomorphisms $\Psi_i: A \rightarrow Q_i$, $i = 1, 2, 3$, in such a way that Ψ_i maps the vertex of A opposite s to a vertex of τ , and Ψ_i maps s onto the union of the two subdividing segments lying on ∂Q_i . Now push forward the horizontal foliation on A , to give a foliation of each Q_i . Since s is a leaf of the horizontal foliation of A , then the foliations of Q_i match up to give a singular foliation of τ , with a 3-pronged




FIGURE 10

singularity at the point where all the Q_i intersect; this defines a case 0 foliation of τ . Notice that each corner of a case 0 foliation of τ has a foliated neighborhood modelled on the foliation by vertical lines of the sector $-\pi/4 \leq \arg(z) \leq \pi/4$ of \mathbf{C} ; we say that such a corner is *transversely foliated* (see Figure 10).

The remaining cases are easier to describe. A case 1 foliation of a triangle is the foliation of \mathbf{C} by vertical lines, restricted to a regular triangle with horizontal base. Two corners of a case 1 triangle are transversely foliated; the remaining corner has a foliated neighborhood modelled on the foliation by horizontal lines of the sector $-\pi/4 \leq \arg(z) \leq \pi/4$ of \mathbf{C} ; we say that such a corner is a *case 1 singular corner*. A case 2 foliation of a triangle is the foliation of \mathbf{C} by horizontal lines, restricted to a triangle with horizontal base. One corner of a case 2 triangle is transversely foliated; each of the remaining corners has a foliated neighborhood modelled on the foliation by horizontal lines of the sector $0 \leq \arg(z) \leq \pi/2$ of \mathbf{C} , we say that such a corner is a *case 2 singular corner*. Finally, a case 3 foliation of a triangle is just the empty foliation.

Thus, a partial measured foliation f on (S, P) is in normal form with respect to an ideal triangulation δ of (S, P) if and only if, for each triangle T of δ , the underlying foliation of f restricted to T is a case 0, 1, 2, or 3 foliation of T .

Notice that, if f is in normal form with respect to δ , then for every arc h of δ , $\int_h f = \langle \{f\}, h \rangle$. This is obvious if $\int_h f = 0$; and if $\int_h f \neq 0$, one need only to observe that f and h satisfy the hypotheses of the minimal intersection proposition. Thus, for each $h \in H(\delta)$, the h -coordinate of $q_\delta \circ \nu(\{f\})$ is precisely $\int_h f$.

Let $W(\delta) \subset \mathbf{R}_{\geq 0}^{H(\delta)}$ be the collection of all $\mu \in \mathbf{R}_{\geq 0}^{H(\delta)}$ such that $\mu = q_\delta \circ \nu(\{f\})$, for some partial measured foliation f in normal form with respect to δ .

LEMMA: CLASSIFICATION OF NORMAL FORMS. *For each $\mu \in \mathbf{R}_{\geq 0}^{H(\delta)}$, $\mu \in W(\delta)$ if and only if μ satisfies the following properties:*

(a) *If T is a triangle of δ with sides h_1, h_2, h_3 (ignoring possible identification), then $\mu(h_1), \mu(h_2), \mu(h_3)$ satisfy the nonstrict triangle inequalities.*

(b) *For each $p \in P$, if $\mu(h) \neq 0$ for each $h \in H(\delta)$ incident to p , then there exists a triangle T of δ , incident to p , and an enumeration h_1, h_2, h_3 of the three sides of T , such that h_1 and h_2 are incident to p , and $\mu(h_1) + \mu(h_2) = \mu(h_3)$.*

Moreover, there is a natural map $L_\delta: W(\delta) \rightarrow \mathcal{MF}$ with the property that $q_\delta \circ \nu \circ L_\delta$ is the identity on $W(\delta)$; it follows that L_δ is 1-1.

PROOF. First we prove necessity of (a) and (b). Let f be a partial measured foliation in normal form with respect to δ , and let $\mu = q_\delta \circ \nu(\{f\})$.

The necessity of (a) is obvious. In fact, it is evident that for each $n = 0, 1, 2, 3$, in a case n triangle, exactly n of the three triangle inequalities are equalities, and $3 - n$ are strict inequalities.

To see why (b) is necessary, suppose that $\mu(h) \neq 0$ for each $h \in H(\delta)$ incident to a particular puncture $p \in P$. This has several implications. First of all, it implies that none of the triangles incident to p are case 0 triangles, since μ is zero on each side of a case 0 triangle; thus, $p \in \text{supp}(f) - \partial(\text{supp}(f))$. Also, none of the corners incident to p are case 2 singular corners, for μ is zero on each side incident to such a corner. Thus, each corner incident to p is either transversely foliated, or is a case 1 singular corner.

Now if *every* corner incident to p were transversely foliated, it would follow that p has a neighborhood which is foliated either by concentric circles, or by leaves spiraling around p , eventually limiting on p . Both of these violate the definition of a partial measured foliation on (S, P) .

It follows that at least one corner incident to p is a case 1 singular corner. Let this be a corner of a triangle T ; it follows that T is a case 1 triangle, and that if h_1, h_2, h_3 are the sides of T enumerated in such a way that the corner between h_1 and h_2 is the given case 1 singular corner, then $\mu(h_1) + \mu(h_2) = \mu(h_3)$.

To prove sufficiency of (a) and (b), for each $\mu \in \mathbf{R}_{\geq 0}^{H(\delta)}$ satisfying (a) and (b), we shall construct a partial measured foliation f_μ on (S, P) , in normal form with respect to δ , with the property that $q_\delta \circ \nu(\{f_\mu\}) = \mu$; in other words, for each $h \in H(\delta)$, $\int_h f_\mu = \mu(h)$. f_μ will be well defined up to isotopy, and thus we will have a map $L_\delta: W(\delta) \rightarrow \mathcal{MF}$, defined by $L_\delta(\mu) = \{f_\mu\}$, which evidently satisfies the conditions of the lemma.

The construction of f_μ is simple: in each triangle T of δ , with sides h_1, h_2, h_3 (ignoring identifications), consider the three triangle inequalities associated to $\mu(h_1), \mu(h_2), \mu(h_3)$; if exactly n of these are equalities, put an appropriately oriented case n foliation in T , and define a transverse measure on the foliation in such a way that the total mass of each side is as specified by μ ; it is easy to check that this can always be done. Now put all these foliations together to form the partial measured foliation f_μ , gluing by isometries along each arc $h \in H(\delta)$ with $\mu(h) \neq 0$, and gluing arbitrarily along other arcs.

Before verifying that f_μ actually satisfies the definition of a partial measured foliation, let us check to make sure that f_μ is well defined, at least up to isotopy.

There are several choices made in the definition of f_μ , although it is already clear that no matter how these choices are made, $\int_h f_\mu = \langle \mu, h \rangle$ for each $h \in H(\delta)$. So to show that f_μ is well defined up to isotopy it suffices to prove the following:

LEMMA: UNIQUENESS OF NORMAL FORMS. *Given partial measured foliations f, f' in normal form with respect to δ , if $\int_h f = \int_h f'$ for each $h \in H(\delta)$, then there exists $\varphi \in \text{Homeo}_0(S, P)$ such that φ preserves each cell of δ , φ preserves the orientation of each 1-cell, and $\varphi(f) = f'$; in particular, $\{f\} = \{f'\}$.*

We refer to the property expressed in this lemma by saying that f and f' are isotopic via an isotopy which *strictly preserves* δ .

PROOF. We construct an isotopy from f to f' , strictly preserving δ , in two steps. first, for each arc $h \in H(\delta)$, find an isotopy of h taking f' -measure on h to f -measure on h ; this can be done because both measures are σ -finite Borel measures which are positive on every open set, and having the same total mass. Do this simultaneously for every h . Now for any triangle T of δ , f and f' make T into a triangle of the same case; once $f|_T$ and $f'|_T$ agree on ∂T , then it is elementary to construct an isotopy from $f|_T$ to $f'|_T$ constant on ∂T . Therefore the isotopy defined on the 1-skeleton of δ can be extended to an isotopy of each individual triangle. This gives an isotopy taking f to f' , strictly preserving δ . Q.E.D.

Now we check that f_μ satisfies all the defining conditions for a partial measured foliation.

Note that any union of triangles of δ forms an essential pinched subsurface of (S) ; thus, the condition that $\text{supp}(f_\mu)$ be an essential pinched subsurface is satisfied. Clearly f_μ has the required singularity structure in $S - P$.

We check that f_μ has the required singularity structure at each puncture. One could set up an application of the puncture singularity lemma, but it seems less clumsy to just make a direct argument. Notice that at a puncture p , we have constructed f_μ in a neighborhood of p by gluing together several corners in a circle, and in each corner we have put one of the following kinds of foliations: a transversely foliated corner; a case 1 singular corner; a case 2 singular corner; or an empty foliation. Moreover, if all of the corners incident to p are transversely foliated, then for each arc h incident to p , we would have $\mu(h) \neq 0$, contradicting condition (b). Thus, there is at least one corner which is not transversely foliated. If there are *no* empty corners, then a direct geometric argument shows that near p , f has the structure of an n -pronged singularity, with one prong for every case 1 singular corner and adjacent pair of case 2 singular corners; $n \geq 1$ is guaranteed by the fact that there is at least one corner which is not transversely foliated. If there are empty corners, then a similar argument shows that near p , f has the structure of a pinch singularity, with one wedge of pie for each maximal circular interval of nonempty corners; if all corners are empty, then $p \notin \text{supp}(f)$. Q.E.D.

Now we are ready to state our main technical result:

NORMAL FORM THEOREM. *Let δ be an ideal triangulation of (S, P) . Every class $\mathcal{F} \in \mathcal{MF}$ is represented by a partial measured foliation in normal form with respect to δ . This representation is unique up to an isotopy which strictly preserves δ . Thus, the map $L_\delta: W(\delta) \rightarrow \mathcal{MF}$ is a bijection. Moreover, L_δ is a homeomorphism.*

PROOF. We have already proven uniqueness of normal forms. We have also already proven that L_δ is 1-1. The gist of the theorem is to prove that L_δ is onto, and then to prove that it is a homeomorphism.

The proof that L_δ is onto, i.e. that every class $\mathcal{F} \in \mathcal{MF}$ is represented by a partial measured foliation in normal form with respect to δ , is basically adapted from exposé 5 of [FLP], where a normal form theorem is proven for closed surfaces

in terms of a pair of pants decomposition of the surface. In particular, for several pieces of this proof, we shall give references to similar proofs found in [FLP]. The emphasis in our presentation will be on those elements of the proof that are peculiar to punctured surfaces.

4. The proof of existence of normal forms has several parts. Given an ideal triangulation δ and a class $\mathcal{F} \in \mathcal{MF}$, we first choose a complete dissection $f \in \mathcal{F}$, and then alter f by Whitehead equivalence so that each arc h of δ intersects $\partial(\text{supp}(f))$ efficiently, and is transverse to f . (Definition: If each arc h of δ intersects $\partial(S')$ efficiently, where S' is a pinched subsurface, then we shall say that δ intersects $\partial(S')$ efficiently.) The f we have found will not be in normal form with respect to δ : a triangle T of δ whose interior intersects $\text{supp}(f)$ will not have its interior completely contained in $\text{supp}(f)$ —there will be “strips” and other pieces missing from T , as indicated by the pictures accompanying the Partial Enlargement Lemma (Figures 16–18).

The next part of the proof will be to define a “partial enlargement” of f , by collapsing all “strips”, to get a partial measured foliation which fills up each triangle whose interior it intersects. Finally, we show that the foliation so obtained is in normal form.

PROPOSITION A. *Given $\mathcal{F} \in \mathcal{MF}$, a complete dissection $f \in \mathcal{F}$, and an ideal triangulation δ , f can be altered by isotopies and Whitehead moves so that δ and $\partial(\text{supp}(f))$ intersect efficiently, and each arc h of δ is transverse to f .*

PROOF. The proof is similar in parts to the proof of Proposition II.6, exposé 5 of [FLP]. The introduction of complete dissections circumvents the difficulties encountered there when an arc (in their setting, a simple closed curve) has zero minimal intersection with the class of f . Beyond this, we shall rely on the proof of Proposition II.6 in [FLP], indicating only the differences which arise because of working with punctured surfaces instead of closed surfaces.

EFFICIENT INTERSECTION LEMMA. *Given $\mathcal{F} \in \mathcal{MF}$, a complete dissection $f \in \mathcal{F}$, and an ideal triangulation δ , f can be isotoped so that $\partial(\text{supp}(f))$ and δ intersect efficiently.*

PROOF. Since $\partial(\text{supp}(f))$ is a collection of essential simple closed curves in $S - P$, we can isotope f so that $\partial(\text{supp}(f))$ is transverse to each arc h of δ . There may be compressing discs of $\partial(\text{supp}(f))$ with respect to arcs of δ . f can be isotoped to eliminate these compressing discs one at a time, starting with an innermost disc. Q.E.D.

ANNULUS TRANSVERSALITY LEMMA. *Under the conditions of the Efficient Intersection Lemma, let C be an annular component of $\text{supp}(f)$. Then f can be altered by isotopy, constant outside of C , so that δ intersects $f|_C$ transversely.*

PROOF. By efficiency of $\partial C \cap \delta$, each component of $C \cap \delta$ is an arc connecting opposite components of ∂C . Also, $f|_C$ foliates C by circles parallel to ∂C . Give C a Euclidean structure with geodesic boundary, where leaves of f are parallel geodesic

circles. Straighten the arcs of $C \cap \delta$ to geodesics by an ambient isotopy constant outside of C . The inverse of this isotopy takes f to a partial measured foliation satisfying the lemma. Q.E.D.

Now let S' be a partial arational component of $\text{supp}(f)$. We can perform an isotopy of f , supported near $\partial S'$, so that δ and $\partial S'$ will intersect efficiently, and so that near $\partial S'$, δ and f are transverse. Let $\delta' = \delta \cap S'$. Let (S'', P'') be the punctured surface obtained from $(S', P \cap S')$ by collapsing all boundary components of S' to punctures. Let δ'' be the resulting ideal cell-division of (S'', P'') , induced from δ' under the collapsing. Note that δ'' is not necessarily a triangulation: a 2-cell of δ'' can be a triangle or a bigon—this shall not cause any difficulties.

Let f'' be the measured foliation of (S'', P'') induced by $f|_{S'}$. f'' is arational. If we can alter f'' by Whitehead equivalence on (S'', P'') , so that f'' and δ'' intersect transversely, this will lift to an alteration of f by isotopies and Whitehead moves supported on S' so that δ intersects $f|_{S'}$ transversely. Therefore, we have reduced the proof of Proposition A to the following proposition, which is stated so as to allow us to make the arcs of δ'' transverse to f'' one at a time.

PROPOSITION B. *Let f be an arational measured foliation on (S, P) , let H be a finite collection of ideal arcs whose interiors are pairwise disjoint, and let $h \in H$ be given. Suppose that, for every $h' \in H - \{h\}$, h' is transverse to f . Then f can be altered by Whitehead equivalence, constant on $\cup(H - \{h\})$, so that h is transverse to f .*

PROOF. We follow the proof of Proposition II.6 of exposé 5 of [FLP], which says in part that if f is an arational measured foliation on a closed surface and h is an essential simple closed curve, then f can be altered by isotopies and Whitehead moves so that h is transverse to f . We shall provide enough details only to show how the proof differs as a result of working on punctured surface surfaces rather than closed surfaces. Also, note that there are some simplifications in their arguments, since we have dealt with the case $\langle \{f\}, h \rangle = 0$ in a different way.

As in Proposition II.6, we begin by using isotopies and Whitehead moves to alter f so that h is an alternating concatenation of α subarcs which are transverse to f , and β subarcs which are leaf cycles of f . The proof consists of performing further alterations of f so that the number of α 's and β 's is reduced. h cannot consist of a single β subarc, for that would make h a closed leaf cycle of f , contradicting arationality of f . So by induction, we are reduced to the case where f consists of a single α subarc, finishing the proof.

First note that by performing isotopies and Whitehead moves, we can eliminate any β subarc occurring at an end of h . Thus, each β subarc occurring in h is nested by two α 's. So we can assume h has the form $\alpha_0 * \beta_0 * \alpha_1 * \cdots * \beta_{n-1} * \alpha_n$.

Following the proof of Proposition II.6, we can assume that no β_j contains a singularity of f , by performing isotopies and Whitehead moves on f , without increasing the number of β_j . Then by performing an isotopy of f , we can remove any β_j such that α_j and α_{j+1} leave β_j from opposite sides of β_j .

This reduces to the case where each β_j is a nonsingular leaf segment, such that α_j and α_{j+1} leave β_j from the same side of β_j . So by the Stability Lemma there exists a

“maximal foliated rectangle” R containing β_j as a horizontal side, whose verticle sides are contained in α_j and α_{j+1} , respectively. More precisely, there exists a mapping $R: I \times I \rightarrow S$ satisfying the following properties:

(a) R is an embedding on $I \times I - I \times 1$, with image in $S - P$, taking interior horizontal segments to leaves of f .

(b) $R(I \times 0) = \beta_j$.


(c) $R(0 \times I)$ and $R(1 \times I)$ are contained (respectively) in α_j and α_{j+1} .

(d) on $I \times 1$, R is an immersion into a cycle of leaves of f , except that there may be folds at 1-pronged singularities of f . Moreover, either $R(I \times 1)$ contains a singularity of f , or one of α_j and α_{j+1} is entirely contained in a vertical side of R .

Any given $h' \in H - \{h\}$ is transverse to f , so h' does not intersect R . The isotopies and Whitehead moves to be performed will all be supported near R , so transversality with h' will not be affected. Moreover, β_j can be chosen so that R is an innermost rectangle. This implies that h does not intersect R in the interior.


If R is an embedding, first remove any puncture singularities on $R(I \times 1)$, by performing inverse Whitehead moves (see Figure 11). The proof of Proposition II.6 then shows that by isotopies and Whitehead moves, the number of α 's and β 's in h can be reduced; we include the diagrams from [FLP, exposé 5] (Figure 12) which illustrate this.

The remainder of the proof consists of reducing to the case where R is an embedding. This is where our proof diverges from the proof of Proposition II.6; to reduce to the case where R is an embedding, we must deal with 1-pronged singularities, which do not arise on closed surfaces.



Removing punctures from $R(I \times 1)$

FIGURE 11



Removing an embedded rectangle

FIGURE 12

Let $\lambda: [0, 1] \rightarrow S$ be the composition $[0, 1] = I \rightarrow I \times 1 \subset I \times I \rightarrow S$, the last map being R . If λ is an embedding, then so is R , and we are done. If λ is not an embedding, then $(0, 1)$ must contain singularity preimages. If $(0, 1)$ does not contain any 1-pronged singularity preimages, then by arationality of f , λ is an embedding, and we are done. If $(0, 1)$ contains at least two 1-pronged singularity preimages, then $\text{im}(\lambda)$ contains a closed leaf cycle of f , contradicting arationality.


So we can assume $(0, 1)$ contains exactly one 1-pronged singularity preimage t_0 . In addition, $\lambda[0, t_0]$ and $\lambda[t_0, 1]$ must both be embeddings, or else arationality of f is contradicted.

If $(0, t_0)$ and $(t_0, 1)$ both contain singularity preimages (Figure 13), then there is a Whitehead move which removes the 1-pronged singularity, and we are done.

If one or both of $(0, t_0)$ and $(t_0, 1)$ do not contain singularity preimages, it follows that one of $[0, t_0]$ and $[t_0, 1]$ is mapped into the image of the other, say $\lambda[0, t_0] \subset \lambda[t_0, 1]$. There are two cases:


(I) If $\lambda[0, t_0] = \lambda[t_0, 1]$ (Figure 14), this leads to the contradiction that $h - P$ contains a copy of S^1 .

(II) If $\lambda[0, t_0] \subsetneq \lambda[t_0, 1]$ (Figure 15), consider the subarc α_j . In order to avoid contradicting the fact that h does not intersect the interior of R , α_j must equal $R(0 \times I)$, β_{j-1} must run along $\lambda[t_0, 1]$ away from the 1-pronged singularity, and α_{j-1} must turn out of R . Using the rectangle between α_{j-1} , β_{j-1} and α_j , we can isotop h slightly so that β_{j-1} no longer intersects R . Then we can isotop h so that $\beta_{j-1} * \alpha_j * \beta_j * \alpha_{j+1}$ is replaced by a β subarc followed by an α subarc (see Figure 15). Combining these two isotopies and taking the inverse, we have an isotopy of f which reduces the total number of α and β subarcs of h .



$(0, t_0)$ and $(t_0, 1)$ both contain singularity preimages

FIGURE 13



$\lambda[0, t_0] = \lambda[t_0, 1]$

FIGURE 14




FIGURE 15

This finishes the proof of Proposition B, and therefore also of Proposition A.

Consider now an ideal triangulation δ of (S, P) and a partial measured foliation f on (S, P) such that δ and $\partial(\text{supp}(f))$ intersect efficiently, and δ is transverse to f . Consider a triangle T of δ . If $\text{int}(T) \cap \text{supp}(f) = \emptyset$, we shall say that f is *empty* in T ; if $\text{int}(T) \subset \text{supp}(f)$, we shall say that f is *full* or *strictly full* in T ; otherwise, we say that f is *partially full* in T .

PARTIAL ENLARGEMENT LEMMA. *For any ideal triangulation δ of (S, P) , and any $\mathcal{F} \in \mathcal{MF}$, there exists $f \in \mathcal{F}$ such that:*

- (a) δ and $\partial(\text{supp}(f))$ intersect efficiently.
- (b) Each arc of δ whose interior intersects $\text{supp}(f)$ is either transverse to f , or is a leaf cycle of f .
- (c) For each triangle T of δ , either f is full in T or f is empty in T .

PROOF. Let $f' \in \mathcal{F}$ be a partial measured foliation as given by Proposition A. Thus f' satisfies (a) and each arc h of δ is transverse to f' . However, there may be triangles in which f' is partially full. We shall define a new partial measured foliation f by taking a “partial enlargement” of f' . That is, instead of collapsing all of $(S - \text{supp}(f'))^c$ to a spine, we shall only collapse the subsurface L of $(S - \text{supp}(f'))^c$ consisting of the union of $(T - \text{supp}(f'))^c$, taken over all triangles T of δ in which f' is partially full.

To describe how L is collapsed, we look one at a time at the components of $(T - \text{supp}(f'))^c$, for each T in which f' is partially full.

Let T be a triangle of δ . It is convenient to think of T without the identifications on its boundary (more formally, consider T as being a euclidean triangle which is the domain for the characteristic map of a 2-cell of δ ; we shall be considering the pulled-back foliation on T). Since f' is a complete dissection of \mathcal{F} , $\partial(\text{supp}(f'))$ is a collection of simple closed curves in $S - P$. Since δ and $\partial(\text{supp}(f'))$ intersect efficiently, it follows that every component of $\partial(\text{supp}(f')) \cap T$ is a segment with endpoints on the interiors of two distinct sides of T . Therefore, these components fall into three distinct types, depending on which pair of sides the segment connects. Any number of these types may be empty, though all three are empty if and only if f' is empty in T or full in T .

We now classify a component C of $T - \text{supp}(f')$ by the nature of ∂C .

If ∂C contains a single segment of $\partial(\text{supp}(f'))$ and a single corner of T , we call C a *corner piece*. If ∂C contains a single segment of $\partial(\text{supp}(f'))$ and two corners of T ,




FIGURE 16




FIGURE 17

we call C a *side piece*. If ∂C contains two segments of $\partial(\text{supp}(f'))$ of the same type, we call C a *strip*. If ∂C contains two segments of $\partial(\text{supp}(f'))$ of different types, we call C a *house*. If ∂C contains three segments of $\partial(\text{supp}(f'))$ of three different types, we call C a *hexagon* (see Figure 16). It is easy to see that this exhausts all possibilities for C .

We shall define a mapping $\varphi: (S, P) \rightarrow (S, P)$ which will collapse L to a spine, will be the identity on any triangle in which f' is empty or full, and will take f' to f satisfying the lemma. φ will be defined as a composition $\varphi = \varphi_2 \circ \varphi_1$ where φ_1 collapses all “1-handles” and “interior zero handles” of L , and φ_2 collapses all punctured zero handles of L .

For each C , the definition of φ_1 on C depends on the type of C .

If C is a strip, let $\Sigma(C)$ be a closed interval in C with one end on each component of $\partial C - \partial(\text{supp}(f'))$. Similarly, if C is a hexagon, let $\Sigma(C)$ be a Y -shaped 1-complex in C with one end on each component of $\partial C - \partial(\text{supp}(f'))$. In either case, define $\varphi_1|_C$ to be a deformation retract from C to $\Sigma(C)$, taking a component of $\partial C - \partial(\text{supp}(f'))$ to the associated endpoint of $\Sigma(C)$.

If C is a side piece or a house, $\varphi_1|_C$ will not collapse all of C —the corners will remain intact. If C is a side piece, let $\Sigma(C)$ be a closed interval embedded in the interior of the side of T contained in C . Let R be a rectangle embedded in C , with one horizontal side of R equal to $\Sigma(C)$, the other horizontal side of R contained in $\text{int}(\partial C \cap \partial(\text{supp}(f')))$, and the vertical sides of R having their interiors contained in $\text{int}(C)$. Define $\varphi_1|_R$ to be the map which takes each vertical segment of R to the endpoint where that segment intersects $\Sigma(C)$. Extend φ_1 to all of C so that φ_1 is the identity outside a small neighborhood of T , and also on $(\partial C \cap \partial T) - \Sigma(C)$ (see Figure 17).

If C is a house, define R to be the “first story” of C : that is, $R \approx I \times I$ is a rectangle embedded in C , with the bottom $I \times 0$ equal to the nonpunctured component of $\partial C \cap \partial T$; with sides $0 \times I$ and $1 \times I$ contained in either component of $\partial C \cap \partial(\text{supp}(f'))$; and with top $I \times 1$ contained in $\text{int}(T)$. Let $\Sigma(C)$ be a vertical




FIGURE 18

segment in R connecting the top and bottom of R . Define $\varphi_1|_R$ to be the map which takes each horizontal segment of R to the point where that segment intersects $\Sigma(C)$. Extend φ_1 to all of C so that φ_1 is the identity outside a neighborhood of R (see Figure 18).

Finally, if C is a corner piece, define $\varphi_1|_C$ to be the identity.

Note that φ_1 can be made well defined along all arcs of δ : just make sure that when two pieces C and C' meet along an arc of δ , then $\Sigma(C) \cap \Sigma(C')$ is a point on that arc. This defines $\varphi_1|_L$. Now extend φ_1 to all of S so that φ_1 takes each cell of δ to itself, so that φ_1 is the identity on each triangle in which f' is empty or full, and so that φ_1 is the identity outside a neighborhood of L .

Note that $\varphi_1(f')$ satisfies all the conditions to be a partial measured foliation, except that a component of $\partial(\text{supp}(\varphi_1(f')))$ might be a simple closed curve in $S - P$ which bounds a once-punctured disc made up of corner pieces in $S - \text{supp}(\varphi_1(f'))$. The remaining corner pieces of $S - \text{supp}(\varphi_1(f'))$ will form “pie wedges” around punctures p such that p is incident to both empty and partially full triangles with respect to f' .

Now define $\varphi_2(S, P) \rightarrow (S, P)$ so that: φ_2 maps each corner piece of $S - \text{supp}(\varphi_1(f'))$ to the puncture contained in the corner piece; φ_2 maps each cell of δ to itself; φ_2 is constant on triangles in which $\varphi_1(f')$ is full; φ_2 is constant outside a small neighborhood of the union of all corner pieces of $S - \text{supp}(\varphi_1(f'))$. Note that φ_2 will not necessarily be constant on a triangle T in which $\varphi_1(f')$ is empty, because for every side piece of $S - \text{supp}(f')$ adjacent to T , there are two corner pieces of $S - \text{supp}(\varphi_1(f'))$ adjacent to T .

Let $f = \varphi_2 \circ \varphi_1(f')$. There are several things to check: that f is a partial measured foliation; that f satisfies (a), (b), and (c); and that f is in the same class as f' .

To show that f is a partial measured foliation, first note that $\text{supp}(f)$ is the union of all triangles of δ in which f' is full or partially full, therefore $\text{supp}(f)$ is an essential pinched subsurface of (S, P) . Second, note that for an interior singularity s of f : either s is an interior singularity of f' and φ is constant near s ; or s arises when φ_1 collapses a hexagon, identifying three points of $\partial(\text{supp}(f'))$; or s arises when φ_1 identifies two points of $\partial(\text{supp}(f'))$, at least one of which is a boundary singularity of f' . In the latter two cases, φ_2 is constant near s . In any case, it is easy to see that s is an n -pronged singularity of f , with $n \geq 3$.

Finally, at a puncture p , we apply the Puncture Singularity Lemma to $\varphi_1(f')$. The only thing that needs checking is the following: if c is a component of $\partial(\text{supp}(\varphi_1(f')))$ which bounds a once-punctured disc, then c contains a singularity of $\varphi_1(f')$. For

suppose not: Then c is a smooth closed leaf of $\varphi_1(f')$ on the boundary. By the Stability Lemma this implies the existence of a smooth closed leaf of $\varphi_1(f')$ in the interior which bounds a once-punctured disc. This implies in turn that f' , and any enlargement of f' , contains a smooth closed leaf bounding a once-puncture disc; a standard argument using the Euler-Poincaré Index Formula implies a singularity structure which violates the definition of a measured foliation.

f satisfies (a) and (c), since $\text{supp}(f)$ is a union of triangles of δ . To verify (b), let h be an arc of δ . If h bounds empty triangles of f' on both sides, then h bounds empty triangles of f on both sides. If h bounds an empty triangle of f' and a side piece of $S - \text{supp}(f')$, then h is a boundary leaf cycle of f . If h bounds side pieces of $S - \text{supp}(f')$ on both sides, then h is an interior leaf cycle of f . Otherwise, h is transverse to f' , implying that h is transverse to f .

Finally, we show that f' and f are in the same class in \mathcal{MF} . To do this, let φ' be a collapsing map taking f to a total measured foliation g on (S, P) . It would be nice if $\varphi' \circ \varphi$ were a collapsing map taking f' to g , but unfortunately for our definitions, $\varphi' \circ \varphi$ does not act as a topological immersion on $\partial(\text{supp}(f'))$: $\varphi = \varphi_2 \circ \varphi_1$ collapses whole segments of f' to a point, namely those segments whose image under φ_1 bound corner pieces of $S - \text{supp}(\varphi_1(f'))$. However, it is easy to see that $\varphi' \circ \varphi = \Psi' \circ \Psi$, where Ψ is an enlarging map taking f' to some total measured foliation g' , and Ψ' is a composition of Whitehead moves taking g' to g . Q.E.D.

Here is another way to look at the construction of the partial enlarging map φ , which emphasizes the analogy with enlarging maps as originally defined. Let $\Sigma = \bigcup \varphi(\Sigma(C))$, taken over all components C of $(T - \text{supp}(f'))^c$, for each T in which f' is partially full. Then Σ is a spine of the subsurface L , relative to the punctures in P and the collection of boundary arcs $\partial L \cap (\text{union of arcs of } \delta)$. Moreover, Σ is the largest closed leaf cycle of f which contains no smooth closed leaves. The map φ can then be viewed as a map which collapses L to its spine Σ , and acts homeomorphically on the complement of L . Of course, φ also preserves cells of δ .

To finish the proof of existence of normal forms, we show that the partial measured foliation f constructed in the partial enlargement lemma is in normal form with respect to δ .

We shall use repeatedly the following corollary to the Euler-Poincaré Index Formula.

FOLIATED DISC LEMMA. *Given a singular foliation F of the closed ball B such that the maximal subset of ∂B which is transverse to F is either empty or an open arc, then F has a positive index singularity in $\text{int}(B)$.*

PROOF. Consider the doubled foliation F' on the double S^2 of B ; we think of B as the northern hemisphere of S^2 . The hypotheses on F imply that the index sum of singularities of F' over $S^1 = \partial B \subset S^2$ is at most $+1$. Since $\chi(S^2) = +2$, F' must have positive index singularities in $S^2 - S^1$, by the Euler-Poincaré Index Formula. By symmetry, these must be evenly distributed between the open northern and southern hemispheres. Therefore, F has a positive index singularity in $\text{int}(B)$. Q.E.D.




FIGURE 19

Let T be a triangle in which f is full. Again, we think of T without boundary identifications. Let α, β, γ denote the sides of T . Let $p_\alpha, p_\beta, p_\gamma$ denote the vertices of T , each one opposite the indicated side. Recall that each side of T is either a leaf cycle of f or is transverse to f .

First we note that at most one side of T can be a leaf cycle. For otherwise (by the Foliated Disc Lemma), f has a positive index singularity in $\text{int}(T)$, a contradiction.

Next, if one side of T , say α , is a leaf cycle of f , we shall analyze the structure of f near α . Note that there does not exist a leaf cycle ℓ in T , distinct from α , which is homeomorphic to the closed interval, with one end on α and the other end on ∂T , for this implies again the contradiction that f has a positive index singularity in $\text{int}(T)$ (we can assume $\text{int}(\ell) \cap \text{int}(\alpha) = \emptyset$; then ℓ divides T into two discs, at least one of which satisfies the hypotheses of the Foliated Disc Lemma). It follows that f has no singularities on α , and also that the corners of T adjacent to α are foliated by f as case 2 singular corners.

Now we analyze f near a vertex of T whose adjacent sides are transverse to f : say, β, γ are transverse to f ; we shall analyze f near p_α . Recall that f is a partial measured foliation on (S, P) ; thus, we can analyze $f|_T$ near p_α in terms of the number of singular leaves of f at p_α that lie in T in the angle between β and γ . If there are no singular leaves between β and γ , the corner of T at p_α is transversely foliated by f . If there is one singular leaf between β and γ , then the corner at p_α is foliated as a case 1 singular corner. If there are two or more singular leaves between β and γ , then by an argument similar to the one in the previous paragraph, we can find a positive index singularity of f in $\text{int}(T)$, a contradiction.

We now have a complete picture of how f behaves on ∂T ; in particular, every corner is foliated as a transversely foliated corner, a case 1 singular corner, or a case 2 singular corner.

To analyze f on $\text{int}(T)$, we introduce the double D of T , together with the doubled foliation f' on D . The possibilities for the foliation f' are analyzed using the Euler-Poincaré Index Formula. Notice that D is a sphere, with the three arcs α, β , and γ (a slight abuse of notation) decomposing D into two triangles. Each of these arcs is either transverse to f' , or is a leaf segment of f' , but at most one can be a leaf segment. Each vertex $p_\alpha, p_\beta, p_\gamma$ of D is a singularity of index either $+1$, $+1/2$, or 0 , corresponding (respectively) to a transversely foliated corner of T , a case 2 singular corner of T , or a case 1 singular corner. See Figure 19.

Consider the case where one arc of D , say, α , is a leaf segment; then β and γ are transverse. Moreover, p_β and p_γ are both index $+1/2$ singularities of f' . f' has no positive index singularities away from the vertices, so by Euler-Poincaré, f' has an index $+1$ singularity at p_α , and f' has no other singularities on D . Therefore the foliation f has no singularities in $\text{int}(T)$. By the Stability Lemma T is foliated by f as a case 2 triangle.

Now consider the case where all three arcs of D are transverse to f' , and one vertex, say, p_α , has index 0. Then p_β and p_γ must both have index $+1$, and there can be no other singularities of f' in D . Therefore in T , the foliation f has the corner structure of a case 1 triangle. In particular, there is a leaf ℓ emerging into $\text{int}(T)$ from the vertex p_α . There is no other singularity in T for ℓ to run into, so by Poincaré recurrence, ℓ must leave through some side of T . But if this side is β or γ , then there is a positive index singularity of f in $\text{int}(T)$; thus, ℓ leaves through α . Then by the Stability Lemma f must foliate T as a case 1 triangle.

Finally, consider the case where α , β , and γ are all transverse to f' , and f' has index $+1$ singularities at each of p_α , p_β , and p_γ . Thus, $f|T$ has the corner structure of a case 0 triangle. By Euler-Poincaré, the additional singularities of f' must have an index sum of -1 . There can be no other positive index singularities, and the total amount -1 must be distributed evenly between the front and back of D . It follows that f' has two 3-pronged (index $-1/2$) singularities on D , and therefore f has one 3-pronged singularity in $\text{int}(T)$. The three singular leaves of this singularity must leave T through different sides, to avoid creating positive index singularities. Then by the Stability Lemma f foliates T as a case 0 triangle.

This finishes the proof of existence of normal forms.

As a footnote to the proof of existence and uniqueness of normal forms, it is nice to see that the theorem can be reformulated entirely in terms of complete dissections; we shall need this reformulation in the proof of continuity of L_δ .

NORMAL FORM THEOREM (COMPLETE DISSECTION VERSION). *For each ideal triangulation δ of (S, P) , and each class $\mathcal{F} \in \mathcal{MF}$, there is a complete dissection $f \in \mathcal{F}$ which δ intersects efficiently; moreover, f is unique up to an isotopy which strictly preserves δ .*

PROOF. The existence of f is proved using Proposition A, plus the fact that any partial measured foliation which is Whitehead equivalent to a complete dissection is itself a complete dissection. The proof of uniqueness parallels the proof of the uniqueness half of the Complete Dissection Lemma; in place of the fact used there that two total measured foliations in the same class are Whitehead equivalent, we use here the uniqueness part of the Normal Form Theorem.

Let f_1 and f_2 be complete dissections of \mathcal{F} , each of which δ intersects efficiently. By the proof of the Partial Enlargement Lemma for $i \in \{1, 2\}$, there are maps $\Psi_i(S, P) \rightarrow (S, P)$ preserving cells of δ , such that $f'_i = \Psi_i(f_i)$ is in normal form with respect to δ . By uniqueness of normal forms, there is a map $\varphi' \in \text{Homeo}_0(S, P)$ preserving cells of δ such that $\varphi'(f'_1) = f'_2$.

Let Σ_i be the largest closed leaf cycle of f_i' containing no closed leaves; clearly $\varphi'(\Sigma_1) = \Sigma_2$. Moreover, if we let L_i be the union of $(T - \text{supp}(f_i))^c$, taken over all triangles T in which f_i is partially full, then Ψ_i is a map which collapses L_i onto Σ_i , and acts homeomorphically on $S - L_i$. Unfortunately, Ψ_i is not a local homeomorphism on $(S - L_i)^c$, for as we saw in the Partial Enlargement Lemma, Ψ_i collapses whole arcs of ∂L_i to points, when such an arc lies on the boundary of a corner piece of L . However, this is only a perceptual disadvantage: there is still a coherent way to define a local homeomorphism $\varphi|U$ on each U in some finite open cover of $\text{supp}(f_1)$, so that $\varphi' \circ \Psi_1 = \Psi_2 \circ \varphi$ on U . Pasting all these maps together defines φ on $\text{supp}(f_1)$; this is easily extended first over L and then over the remainder of S , in such a way that φ preserves cells of δ and takes f_1 to f_2 . Q.E.D.

5. We now turn to the proof that $L_\delta: W(\delta) \rightarrow \mathcal{MF}$ is a homeomorphism. Since we already know that $L_\delta^{-1} = \nu \circ q_\delta$ is 1-1, onto, and continuous, it remains to prove that L_δ is continuous. Since the topology on \mathcal{MF} is defined by the embedding $\nu: \mathcal{MF} \rightarrow \mathbf{R}_{\geq 0}^{\mathcal{H}}$, and since $\mathbf{R}_{\geq 0}^{\mathcal{H}}$ is given the product topology, it suffices to show that: for each $h \in \mathcal{H}$, the function $i_h: W(\delta) \rightarrow \mathbf{R}$, given by $i_h(\mu) = \langle L_\delta(\mu), h \rangle$, is continuous. In fact, we shall show how to explicitly derive homogeneous piecewise linear formulas for the map i_h .

First we recouch the problem in different terms. Suppose it is true that, for any two ideal triangulations δ and δ' , the bijective “coordinate change” map $L_{\delta'}^{-1} \circ L_\delta: W(\delta) \rightarrow W(\delta')$ is continuous, given by homogeneous piecewise linear formulas. If we are given an ideal triangulation δ and $h \in \mathcal{H}$, it is a simple exercise, which we leave to the reader, to show that there exists an ideal triangulation δ' and $h' \in H(\delta')$ such that $\{h'\} = h$. Thus, i_h is one of the coordinates of the map $L_{\delta'}^{-1} \circ L_\delta$, and the desired properties of i_h follow.

So to verify continuity of each of the maps L_δ , it suffices to show that all coordinate change maps are continuous.

Now for an arbitrary pair of ideal triangulations δ, δ' , one would expect the formulas for the coordinate change map to be quite complicated. Therefore, we shall start with the case where the relationship between δ and δ' is as simple as possible: where δ and δ' are related by an “elementary move”. This is a combinatorial construction which has become quite familiar recently, in the context of studying geometric structures on punctured surfaces. Briefly, given an ideal triangulation δ and $h \in H(\delta)$, if h bounds two distinct triangles of δ , then these two triangles are joined along h to form a quadrilateral; removing h and inserting the opposite diagonal of this quadrilateral, we obtain a new ideal triangulation δ' which we say is obtained from δ by an “elementary move”.

Once we have made the definitions, we shall prove the Coordinate Change Theorem for Elementary Moves, which implies that the coordinate change map $L_{\delta'}^{-1} \circ L_\delta$ is continuous, when δ' is obtained from δ by an elementary move. Then to complete the proof of the Normal Form Theorem, we shall show that for any two ideal triangulations δ, δ' , there exists a finite sequence of elementary moves connecting δ and δ' , implying that $L_{\delta'}^{-1} \circ L_\delta$ is continuous. Here are the formal definitions.




FIGURE 20

Given an ideal triangulation δ , an arc $h_0 \in H(\delta)$ is *removable* if h_0 bounds two distinct triangles of δ , i.e. if h_0 is not the interior arc of a puncture piece. Let these two triangles be denoted T_W and T_E (for the compass points), and let us choose characteristic maps for T_W and T_E , defined an Euclidean triangles τ_W and τ_E , in such a way that τ_W and τ_E share a single common edge e , and so that the two characteristic maps agree along e , producing a characteristic map for h_0 . We denote the quadrilateral $\tau_W \cup \tau_E$ by Q ; we shall systematically confuse Q with its image in S , under the union of the two characteristic maps $\tau_W \rightarrow T_W$ and $\tau_E \rightarrow T_E$. Note that this map may identify certain vertices of Q to the same puncture, and certain sides of Q may be identified to the same ideal arc of δ ; we shall usually ignore these identifications. With this understanding, we shall also cease to refer to τ_W and τ_E , confusing them systematically with T_W and T_E . Let us denote the vertices of Q as v_N, v_S, v_W, v_E , in such a way that h_0 connects v_N and v_S , and so that v_W (resp. v_E) is a vertex of T_W (resp. T_E) (see Figure 20). Then we define an *opposite diagonal* to h_0 to be any ideal arc of (S, P) which is contained in Q and which connects v_W to v_E ; clearly an opposite diagonal to h_0 is well defined up to isotopy. It is also clear that there is an ideal triangulation δ' , well defined up to isotopy, characterized by the property that $H(\delta') = \{h'_0\} \cup [H(\delta) - \{h_0\}]$, where h'_0 is any opposite diagonal to h_0 . We say that δ' is *obtained from δ by an elementary move performed on h_0* , and we write $\delta \rightarrow \delta'$; and we say that Q is the *support* of the elementary move. It is an exercise to show that elementary moves performed on distinct arcs of δ result in nonisotopic ideal triangulations; thus, the notation $\delta \rightarrow \delta'$ completely determines the arc that the elementary move is performed on, as well as the opposite diagonal.

To state the coordinate change theorem for elementary moves, label the four sides of Q as $h_{NE}, h_{SE}, h_{SW}, h_{NW}$, as indicated in Figure 20.

COORDINATE CHANGE THEOREM FOR ELEMENTARY MOVES. *Given an ideal triangulation δ and a removable arc $h_0 \in H(\delta)$, let h'_0, δ' , and Q be as above, with the vertices and sides of Q as labelled above. Choose $\mu \in W(\delta)$ and $\mu' \in W(\delta)$ so that $\mu' = L_{\delta'}^{-1} \circ L_{\delta}(\mu)$. Then for each $h \in H(\delta) \cap H(\delta')$, we have $\mu'(h) = \mu(h)$; and furthermore,*

$$\mu'(h'_0) = \begin{cases} \mu(h_{NW}) + \mu(h_{SE}) - \mu(h_0) & \text{if } \mu(h_{NW}) + \mu(h_{SE}) \geq \mu(h_{NE}) + \mu(h_{SW}), \\ \mu(h_{NE}) + \mu(h_{SW}) - \mu(h_0) & \text{if } \mu(h_{NW}) + \mu(h_{SE}) \leq \mu(h_{NE}) + \mu(h_{SW}). \end{cases}$$

Thus, $L_{\delta'}^{-1} \circ L_{\delta}$ is a continuous map.

Note the nice symmetry in these formulas: the formula for $\mu(h_0)$ in terms of $\mu'(h'_0)$ is identical to the formula for $\mu'(h'_0)$ in terms of $\mu(h_0)$. This, of course, reflects the combinatorial symmetry of an elementary move, which holds in the following sense: $\delta \rightarrow \delta'$ is an elementary move performed on h_0 supported in Q , and $\delta' \rightarrow \delta$ is an elementary move performed on h'_0 supported in Q .

PROOF. Let $\mathcal{F} = L_\delta(\mu)$. The idea of the proof is as follows: let f be a normal form with respect to δ . We shall give an explicit method for altering f , staying in the equivalence class \mathcal{F} , so as to produce a normal form f' with respect to δ' . This shall be done in such a way that we will be able to read off the value of $\mu'(h'_0) = \langle \mathcal{F}, h'_0 \rangle$ from the intersection number of f' and h'_0 , and verify the above formula.

However, instead of using normal forms, we shall use complete dissections, which are technically easier to work with. Thus, if $f \in \mathcal{F}$ is a complete dissection which intersects δ efficiently, we shall give an explicit method for altering f by isotopies and Whitehead moves, supported entirely in $\text{int}(Q)$, and for choosing an opposite diagonal h'_0 to h_0 , so that f' and h'_0 intersect efficiently. Since f' and f are identical outside of $\text{int}(Q)$, it will follow that f' intersects all the arcs in $H(\delta) \cap H(\delta') = H(\delta') - \{h_0\} = H(\delta) - \{h'_0\}$ efficiently, and thus f' intersects δ' efficiently. It follows immediately that f and f' give the same measure to each arc of $H(\delta) \cap H(\delta')$. As mentioned above, once we have done the work of constructing f' , we will be able to read off the value of $\mu'(h'_0) = \int_{h'_0} f'$.

We shall adopt the notation $\mu_{\text{NW}} = \mu(h_{\text{NW}})$, $\mu_{\text{NE}} = \mu(h_{\text{NE}})$, etc.

So let $f \in L_\delta(\mu)$ be a complete dissection which intersects δ efficiently. We shall begin by stating some numerical properties of f .

Let $\mu_{\text{W}} = (1/2)[\mu_{\text{NW}} + \mu_{\text{SW}} - \mu(h_0)]$. Note that μ_{W} is the total transverse weight of all leaf segments of f which cross Q from h_{NW} to h_{SW} .

Similarly, if $\mu_{\text{E}} = (1/2)[\mu_{\text{NE}} + \mu_{\text{SE}} - \mu(h_0)]$, then μ_{E} is the total transverse weight of all leaf segments of f crossing Q from h_{NE} to h_{SE} .

Now let $\mu_d = (1/2)[\mu_{\text{NE}} + \mu_{\text{SW}} - \mu_{\text{NW}} - \mu_{\text{SE}}]$ (d stands for “diagonal”). Notice that, if $\mu_d > 0$, then there are leaf segments of f crossing Q from h_{NE} to h_{SW} , and the total weight of such leaf segments is μ_d ; we shall say that these leaf segments have “positive slope”. And if $\mu_d < 0$, then there are leaf segments of f crossing Q from h_{NW} to h_{SE} , and their total weight is $-\mu_d$; we say that these leaf segments have “negative slope”. Finally, if $\mu_d = 0$, there are no leaf segments of either positive or negative slope. Clearly there cannot be leaf segments of both positive and negative slope, for that would violate the foliation properties of f .

Now we shall show, by separating into several cases, how to choose an opposite diagonal h'_0 to h_0 , and a partial measured foliation f' Whitehead equivalent to f , so that f' and f are identical outside of $\text{int}(Q)$; in each case, it will be evident by construction that h'_0 intersects each leaf segment of f' contained in Q at most once; later we shall give the proof, which does not depend on the case, that f' and h'_0 intersect efficiently.

Case I. Suppose $\mu_d > 0$. Then there exist smooth leaf segments of f having positive slope. Thus, there exists a horizontally foliated rectangle $R \subset Q$ such that:

$R \subset \text{supp}(f)$; the two vertical sides of R are contained in h_{NE} and h_{SW} , respectively; and each horizontal leaf of R is a leaf of f .

In this case, we choose $f' = f$. It remains only to choose the opposite diagonal h'_0 so that h'_0 intersects each leaf segment of f in Q at most once. Construct h'_0 as follows: starting from v_{W} , run along nearly parallel to h_{SW} until you enter R , then proceed diagonally across R until nearly encountering h_{NE} , then run nearly parallel to h_{NE} to v_{E} . Clearly h'_0 satisfies the required conditions. Also, it is clear that h'_0 intersects precisely the following leaf segments of f in Q (and intersects each exactly once): segments connecting h_{NW} and h_{SW} ; positive slope segments; and segments connecting h_{NE} and h_{SE} . Thus, $\int_{h'_0} f' = \mu_{\text{W}} + \mu_{\text{E}} + |\mu_d| = \mu_{\text{W}} + \mu_{\text{E}} + \mu_d = \mu_{\text{NE}} + \mu_{\text{SW}} - \mu(h_0)$.

Case II. Suppose $\mu_d < 0$. Then there exist smooth leaf segments of f having negative slope. We argue analogously to case I, setting $f' = f$, and finding an opposite diagonal h'_0 to h_0 which intersects precisely the following leaf segments of f in Q ; segments connecting h_{NW} and h_{SW} ; negative slope segments; and segments connecting h_{NE} and h_{SE} . Thus,

$$\int_{h'_0} f' = \mu_{\text{W}} + \mu_{\text{E}} + |\mu_d| = \mu_{\text{W}} + \mu_{\text{E}} - \mu_d = \mu_{\text{NW}} + \mu_{\text{SE}} - \mu(h_0).$$

Case III. Suppose $\mu_d = 0$. There are several subcases to consider.

Case III.A. Suppose there exists an embedded interval $\alpha \subset Q$ such that α connects opposite sides of Q , and $\alpha \subset S - \text{supp}(f)$; suppose that α connects h_{NE} and h_{SW} (a similar argument covers the case where α connects h_{NW} and h_{SE}). Again we choose f' to be f . Construct h'_0 by starting from v_{W} , running along nearly parallel to h_{SW} until encountering α , running along α until nearly encountering h_{NE} , then running along nearly parallel to h_{NE} until encountering v_{E} . Clearly h'_0 intersects precisely: leaf segments of f connecting h_{SW} and h_{NW} , as well as leaf segments connecting h_{SE} and h_{NE} . Thus,

$$\begin{aligned} \int_{h'_0} f' &= \mu_{\text{W}} + \mu_{\text{E}} = (1/2)[\mu_{\text{SW}} + \mu_{\text{NW}} + \mu_{\text{SE}} + \mu_{\text{NE}}] - \mu(h_0) \\ &= \mu_{\text{NE}} + \mu_{\text{SW}} - \mu(h_0) = \mu_{\text{NW}} + \mu_{\text{SE}} - \mu(h_0), \end{aligned}$$

using the fact that $\mu_{\text{NE}} + \mu_{\text{SW}} = \mu_{\text{NW}} + \mu_{\text{SE}}$, which follows from the fact that $\mu_d = 0$.

Case III.B. Suppose that the hypothesis of Case III.A is false. We need some definitions to describe the structure of f in T_{W} and T_{E} .

Define a *tine* in a triangle T to be a 1-complex embedded in T , homeomorphic to the letter y , such that the three order 1 vertices lie on distinct sides of T , and the remainder lies in $\text{int}(T)$. We say that it is an *f-tine* if each edge is a leaf segment of f .

Define a *bisector* of T to be an embedded interval in T with one end on a vertex of T and the other end on the opposite side of T . We say that it is an *f-bisector* if it is a leaf segment of f .

LEMMA. *Suppose that the hypothesis of Case III.A is false; i.e. there does not exist an embedded interval in Q connecting opposite sides of Q . Then there does not exist a tine in either T_W or T_E contained in $S - \text{supp}(f)$.*

PROOF. Suppose otherwise. Let Y be a tine in, say, T_W , with $Y \subset S - \text{supp}(f)$, and let x be the vertex of Y lying on h_0 . Since f intersects h_0 efficiently, there exists an embedded interval β in T_E , $\beta \subset S - \text{supp}(f)$, connecting x to another side of T_E . Then clearly we can find an embedded interval $\alpha \subset Y \cup \beta \subset S - \text{supp}(f)$ connecting opposite sides of Q , a contradiction. Q.E.D.

LEMMA. *Let T be any triangle of δ , and suppose that there does not exist a tine in T contained in $S - \text{supp}(f)$. Then either there exists an f -tine in T or there exists an f -bisector in T .*


PROOF. Let the sides of T be h, h', h'' . The hypothesis of the lemma implies that there exists a point $x \in h \cap \text{int}(\text{supp}(f))$ with the following property: any leaf segment of f entering T through h , closer to h' than the point x , must leave T through h' ; and a similar statement holds for h'' .

It could happen that x lies at an end of h , where, say, h meets the side h' of T . It follows that every leaf segment of f entering T through h leaves through h'' . Since $x \in \text{int}(\text{supp}(f))$, and since f intersects δ efficiently, it follows that there is a neighborhood of x in T which is foliated as either a case 1 singular corner or a case 2 singular corner. In the latter case, it follows that h' is a leaf segment of f ; but h' connects two punctures, contradicting the fact that f is a complete dissection; thus, x must have a neighborhood foliated as a case 1 singular corner. If ℓ is the leaf segment of f in T ending at x , then since f intersects δ efficiently, it follows that ℓ leaves T through the opposite side h'' ; so ℓ is an f -bisector of T . Notice that ℓ actually lies in $\text{int}(\text{supp}(f))$; in fact, since f is a complete dissection, no f -bisector of a triangle can lie in $\partial(\text{supp}(f))$.

Now suppose that $x \in \text{int}(h)$. Since $s \in \text{int}(\text{supp}(f))$, there is a leaf segment ℓ of f entering T at the point x . Now two leaf segments entering T arbitrarily close to x on opposite sides of x leave T through different sides. It follows that ℓ must end at a singularity of f in T . This can occur in one of two ways: either ℓ ends at a 3-pronged singularity of f in $\text{int}(T)$; or ℓ ends at the opposite vertex of T from h . In the former case, ℓ is part of an f -tine in T ; in the latter case, ℓ is an f -bisector of T . Q.E.D.

Therefore, in Case III.B it follows that in both T_W and T_E , there is either an f -tine or an f -bisector. We now consider several subcases, depending on the possible combinatorial combinations of f -tines and f -bisectors.

Case III.B.1. Suppose there exists an f -bisector of T_W connecting v_W to h_0 and an f -bisector of T_E connecting v_E to h_0 . If these bisectors do not end at the same point of h_0 , then we arrive at the contradiction that either $\mu_d \neq 0$ or there is a leaf segment $\alpha \subset Q \cap (S - \text{supp}(f))$ connecting opposite sides of Q . But if the bisectors end at the same point of h_0 , we arrive at the conclusion that there is a leaf segment of f connecting two punctures, contradicting the hypothesis that f is a complete dissection. So this case cannot occur.



Case III.B.4

FIGURE 21

Case III.B.2. Suppose there is an f -bisector ℓ in T_W connecting v_W to h_0 , and an f -bisector in T_E connecting v_N to h_{SE} (there are three other similar cases, which the reader may enumerate). Extending ℓ into T_E , we see that it must leave T_E through h_{SE} . But this implies that there is either a leaf segment of f in Q or an embedded interval $\alpha \subset Q \cap (S - \text{supp}(f))$ connecting h_{NW} to h_{SE} , a contradiction. So this case cannot occur.

Case III.B.3. There exists an f -bisector of T_W connecting v_N to h_{SW} and an f -bisector of T_E connecting v_S to h_{NE} (there is one other similar case). Then we arrive at the same contradiction as in Case III.B.2, so this case cannot occur.

Case III.B.4. There exists an f -bisector of T_W connecting v_N to h_{SW} and an f -bisector of T_E connecting v_N to h_{SE} (there is one other similar case). This case *can* occur, and we must now show how to choose f' and h'_0 . Notice that v_N has neighborhoods in both T_W and T_E which are foliated like case 1 singular corners. Thus, v_N has a neighborhood in Q whose foliation by f is locally modelled on the foliation of the sector $\pi/6 \leq \arg(z) \leq 5\pi/6$ of \mathbb{C} by the horizontal leaves of the quadratic differential $z^4 dz^2$. Clearly we can alter f near v_N by Whitehead equivalence, resulting in a partial measured foliation f' , so that v_N has a neighborhood in Q foliated by f' as a case 1 singular corner, and there a leaf segment ℓ of f' in Q connecting v_N to a 3-pronged singularity of f' . Whitehead equivalence preserves the property of being a complete dissection, so f' is a complete dissection. Also, the Whitehead equivalence can clearly be done in such a way that f and f' coincide outside of $\text{int}(Q)$. See Figure 21.

Now we define h'_0 : starting from v_W , run along nearly parallel to h_{NW} until encountering ℓ , and then run along nearly parallel to h_{NE} until encountering v_E . Clearly h'_0 satisfies the required conditions, and $\int_{h'_0} f' = \mu_W + \mu_E$; now compute as in Case III.A.

Case III.B.5. Suppose that there is an f -bisector of T_W connecting v_W to h_0 , and there is an f -tine in T_E (there is one other similar case). The f -bisector in T_W and the f -tine in T_E must end at the same point in h_0 , or else there will be either a leaf segment of f , or an embedded interval $\alpha \subset Q - \text{supp}(f)$, connecting opposite sides of Q . Thus, there exists a leaf segment ℓ of f in Q connecting v_W to a 3-pronged singularity s in T_E . Since f is a complete dissection, we must have $s \in \text{int}(\text{supp}(f))$. Thus, there is a Whitehead move which collapses ℓ to a point, resulting in a complete dissection f' ; this can be done so that f' coincides with f outside of


$\text{int}(Q)$. To define h'_0 , first take a transversal in T_E which goes straight from v_E to s , intersecting f efficiently; let h'_0 be the image of this transversal under the collapsing map of the Whitehead move. Clearly h'_0 satisfies the required conditions. Also, we see that

$$\int_{h'_0} f' = \mu_E = (1/2)[\mu_{NE} + \mu_{SE} - \mu(h_0)] = (1/2)[\mu_{NW} + \mu_{SW} - \mu(h_0)],$$

using the fact that $\mu_d = 0$ and $\mu_w = 0$. The diagram for this case can be obtained from the diagram for Case III.B.4, reading that diagram backwards.

Case III.B.6. Suppose that there is an f -bisector of T_W connecting v_N to h_{SW} , and there is an f -tine in T_E (there are three other similar cases). It follows that there exist leaf segments of f in Q connecting h_{NE} to h_{SW} , a contradiction; so this case cannot occur.

Case III.B.7. Suppose that there are f -tines in both T_W and T_E . These tines must end at the same point of h_0 , or else there will be either leaf segments of f or intervals $\alpha \subset Q - \text{supp}(f)$ connecting opposite sides of Q . Thus, there is a leaf segment ℓ of f connecting two 3-pronged singularities of f , s_W in T_W and s_E in T_E . Define leaf segments ℓ_W and ℓ_E of f , ℓ_W connecting h_{NW} and h_{SW} and ℓ_E connecting h_{NE} and h_{SE} , so that $s_W \in \ell_W$ and $s_E \in \ell_E$. Since f is a complete dissection, ℓ_W and ℓ_E cannot both be in $\partial(\text{supp}(f))$, for the leaf segment ℓ connects s_W to s_E . It follows that there is a Whitehead move which collapses ℓ to a point, resulting in a partial measured foliation f'' with a 4-pronged singularity; we can pull apart this 4-pronged singularity in the opposite direction, resulting in a partial measured foliation f' with a singular leaf segment ℓ' connecting two 3-pronged singularities s_N and s_S , in such a way that $s_N \subset \ell'_N$ and $s_S \subset \ell'_S$, where ℓ'_N, ℓ'_S are leaf segments of f' , ℓ'_N connecting h_{NE} to h_{NW} , and ℓ'_S connecting h_{SE} to h_{SW} . Clearly this can be done so that f' and f agree outside of $\text{int}(Q)$. Now we define h'_0 : start by choosing transversals α_W and α_E so that α_W connects v_W to s_W , intersecting f efficiently, and similarly for α_E ; let α'' be the image of $\alpha_W \cup \alpha_E$ under the collapsing map of the first Whitehead move. Let h'_0 be any arc intersecting ℓ' transversely such that its image under the second Whitehead move is α'' . Clearly h'_0 satisfies the required conditions. Also, it is clear that $\int_{h'_0} f' = \mu_W + \mu_E$; now compute as in Case III.A. See Figure 22.



Case III.B.7.




FIGURE 23

We have now exhausted all the possible cases; this finishes the definition of f' and h'_0 , as well as the verification that $\int_{h'_0} f'$ satisfies the required equation.

It remains to show that h'_0 and f' intersect efficiently; the same proof works for all cases. Suppose this is false: suppose that there exists a compressing disc D for h'_0 and f' . We shall show that, of the two possibilities $D \subset \text{int}(Q)$ and $D \not\subset \text{int}(Q)$, both lead to contradictions. It is obvious by construction in each case of h'_0 and f' that $D \subset \text{int}(Q)$ is contradictory, for any leaf segment of f' completely contained in Q intersects h'_0 at most once. So suppose that $D \not\subset \text{int}(Q)$. Let c, d be the intervals in ∂D such that $\partial c \cap \partial d = \{\text{two points}\}$, $c \subset h'_0$, and d is a leaf segment of f' . Since f' intersects all arcs of δ' transversely, it follows that ∂D intersects the arcs of δ' transversely. In particular, we can choose points $x, y \in \partial D$ and intervals $a, b \subset \partial D$ such that $a \cap b = \partial a \cap \partial b = \{x, y\} \subset \partial Q$, $c \subset a$ and $b \subset d$; moreover, we can assume that x and y lie on the same side \bar{h} of Q ; this follows from the fact that D is a nonpunctured disc in S , and so ∂D must intersect each arc of δ' in an even number of points.

Let e be the interval on \bar{h} cut off by x and y . It now follows that the subdisc of D bounded by $e \cup b$ is a compressing disc for f' and \bar{h} ; thus, f' and \bar{h} do not intersect efficiently. But this implies that f and \bar{h} do not intersect efficiently; and since $\bar{h} \in H(\delta)$, this contradicts the fact that f and δ intersect efficiently. See Figure 23. Q.E.D.

The coordinate change theorem for elementary moves shows that if two ideal triangulations δ, δ' are related by a finite sequence of elementary moves, say $\delta = \delta(0) \rightarrow \delta(1) \rightarrow \cdots \rightarrow \delta(n-1) \rightarrow \delta(n) = \delta'$, then $L_{\delta'}^{-1} \circ L_{\delta} = (L_{\delta(n)}^{-1} \circ L_{\delta(n-1)}) \circ \cdots \circ (L_{\delta(1)}^{-1} \circ L_{\delta(0)})$ is continuous. Thus, as noted earlier, to show that $L_{\delta}: W(\delta) \rightarrow \mathcal{MF}$ is continuous it suffices to prove:

CONNECTIVITY THEOREM FOR ELEMENTARY MOVES. *Any two ideal triangulations of (S, P) are related by a finite sequence of elementary moves.*

This is a theorem which, by this point in time, is very well known. Indeed, several different proofs have already appeared. In fact, there is a much stronger theorem which has been proven and which we now proceed to describe.

A generalization of an ideal triangulation of (S, P) is an *ideal cell-division* of (S, P) , which is a cell division satisfying all the same properties as an ideal triangulation except that a 2-cell is allowed to be an n -gon for any $n \geq 3$; that is, each 2-cell has a characteristic map defined on an n -sided Euclidean polygon, so that each vertex of the polygon goes to a puncture, and each edge of the polygon goes to an ideal arc. Given an ideal cell-division γ , the *deficiency* of γ is defined to be the number of extra ideal arcs that need to be added in order to make γ an ideal triangulation; more precisely, the deficiency is equal to $\sum k_n(n - 3)$, where k_n is the number of n -gons of γ . A *refinement* of γ is an ideal cell-division γ' of (S, P) such that $H(\gamma') \subset H(\gamma)$ ($H(\cdot)$ denotes the set of arcs). As with ideal triangulations, we are primarily interested in ideal cell-divisions up to isotopy.

CONTRACTIBILITY THEOREM. *There exists a contractible cell-complex X that has one cell $c(\gamma)$ of dimension d for every (isotopy class of an) ideal cell-division γ of (S, P) of deficiency d ; the subcomplex $\bigcup c(\gamma')$ taken over all (isotopy classes of) refinements γ' of γ is a sphere of dimension $d - 1$ along which $c(\gamma)$ is attached.*

The Connectivity Theorem follows from the Contractibility Theorem by noting that the 1-skeleton X^1 of X is connected, and consists of a 0-cell $c(\delta)$ for each isotopy class of an ideal triangulation δ , with a 1-cell connecting $c(\delta)$ to $c(\delta')$ if and only if there is an elementary move $\delta \rightarrow \delta'$.

The Contractibility Theorem has been proven from many different points of view. Each proof constructs an “ideal triangulation” of some high-dimensional open ball, and proceeds by showing that X is isomorphic to the dual cell-division of this ideal triangulation. Harer constructed an ideal triangulation of the space of noncompactly supported geodesic laminations on $S - P$; Mumford [Har] noticed that Strebel differentials could be used to triangulate the Teichmüller space of complex analytic structures on (S, P) ; most recently, Epstein and Bowditch [EB], and also Penner [P] used hyperbolic structures on $S - P$ to triangulate the Teichmüller space. There is also a completely combinatorial proof of the Connectivity Property, given under another guise in the appendix of exposé 5 of [FLP] by Poénaru. It would be very interesting to extend Poénaru’s techniques to give a combinatorial proof of the Contractibility Theorem.¹

There has been much work done to exploit the cell-complex X —in particular, using the fact that $\mathcal{M}(S, P)$ acts by cellular homeomorphisms on X with finite order cell stabilizers, one can conceivably compute the rational cohomology of $\mathcal{M}(S, P)$ from $X/\mathcal{M}(S, P)$, and with more care, even the integral cohomology. Although specific calculations of Betti numbers are still forthcoming, Harer [Har] has used this idea to compute the rational cohomological dimension of $\mathcal{M}(S, P)$, and Harer and Zagier [HZ] have computed the Euler characteristic of $\mathcal{M}(S, P)$.

Despite this plethora of proofs of the Connectivity Theorem, we wish to give a new proof, completely combinatorial in spirit, and quite elementary. Our proof has

¹ Hatcher now has an elementary combinatorial proof of the Contractibility Theorem [Hat].

the advantage that it implicitly gives an algorithm for constructing a sequence of elementary moves connecting δ and δ' , when δ' is given in terms of certain intersection numbers with the arcs of δ . We shall not explicitly describe this algorithm; we leave that to the interested reader. It should be a straightforward matter to construct the algorithm from the information provided in the proof.

We begin by defining, for any two isotopy classes of ideal arcs h, h' in (S, P) , an intersection number $i(h, h')$ to be the minimum cardinality of $\text{int}(h) \cap \text{int}(h')$, taken over all $h \in h$ and $h' \in h'$. We also say that h and h' intersect efficiently if there does not exist a compressing disc for h and h' , where a compressing disc is defined in exactly the same manner as it was defined for an ideal arc h and the boundary of a pinched subsurface of (S, P) .

The following is an elementary exercise in the standard techniques of intersection numbers:

EFFICIENT INTERSECTION LEMMA FOR IDEAL ARCS. *Given two ideal arcs h, h' which are not isotopic, if h and h' intersect efficiently, then $i(\{h\}, \{h'\}) = |\text{int}(h) \cap \text{int}(h')|$.*


Now suppose that δ is an ideal triangulation, and h is an ideal arc on (S, P) . Choose an orientation of h arbitrarily (we shall use the orientation of h to define the “first” point of any compact subset of $\text{int}(h)$). We wish to define an intersection number $i(\delta, \{h\})$; one might try taking the sum of $i(\{h\}, \{h'\})$, taken over all $h' \in H(\delta)$; but for technical reasons, the following definition is easier to work with. If h is isotopic to an arc of δ , we define $i(\delta, \{h\}) = 0$. If h is not isotopic to an arc of δ , then it is easy to see that we can isotop h until it simultaneously intersects each arc of δ efficiently (we say that h and δ intersect efficiently); once this is done, the Efficient Intersection Lemma for Ideal Arcs shows that h assumes its minimum intersection with all arcs of δ simultaneously. Let $x \in h$ be the first point of $\text{int}(h)$ lying on an arc of δ , and let h^0 be the unique arc in $H(\delta)$ satisfying $x \in h^0$ (note that there must exist such a point x , since h is not isotopic to an arc of δ). We define $i(\delta, \{h\})$ to be $\sum_{h' \in H(\delta) - \{h^0\}} i(h, h')$. In other words, add up all intersection numbers of h with arcs of δ , except with h^0 .

To see that $i(\delta, \{h\})$ is well defined: since h intersects all the arcs of δ efficiently, the vector of intersection numbers $v \in \mathbf{Z}_{\geq 0}^{H(\delta)}$ is well defined. Then by mimicking the proof of uniqueness of normal forms, it is easy to see that h is determined by v , up to an isotopy which strictly preserves δ . Thus, the arc $h^0 \in H(\delta)$ is well defined, implying that $i(\delta, \{h\})$ is well defined.

Here is an elementary property of $i(\delta, \{h\})$:

LEMMA. *Given an ideal triangulation δ and an ideal arc h , $i(\delta, \{h\}) = 0$ if and only if either h is isotopic to an arc of δ , or h is isotopic to an opposite diagonal of a removable arc of δ .*

PROOF. If h can be isotoped so that $i(h, h') = 0$ for all $h' \in H(\delta)$, then $\text{int}(h)$ lies in $\text{int}(T)$ for some triangle T of δ , from which it follows that h is isotopic to a side of T , and thus to an arc of δ .



Suppose h^0 is the interior arc of a puncture piece T

FIGURE 24

So suppose h is *not* isotopic to an arc of δ , and suppose without loss of generality that h intersects δ efficiently. It follows from the previous paragraph that h intersects *some* arc of δ . But since $i(\delta, \{h\}) = 0$, it follows that h^0 is the *only* arc of δ that h intersects. We shall show that h^0 is removable, and that h is a diagonal to h^0 .

If h^0 were nonremovable, being the interior arc of some puncture piece T of δ , then $\text{int}(h)$ would have to lie entirely in the set $\text{int}(T) \cup \text{int}(h^0)$, since $\text{int}(h)$ does not intersect the boundary arc of T . This implies that h is isotopic to either the interior arc of T (namely, h^0), or to the boundary arc of T , contradicting the hypothesis that h is not isotopic to an arc of δ .

We have proven that h^0 is removable. Let Q be the quadrilateral of δ that h^0 bisects. Since $\text{int}(h)$ does not intersect ∂Q , then $\text{int}(h) \subset \text{int}(Q)$. Since h is not isotopic to an arc of δ , this implies that h is an opposite diagonal to h^0 . Q.E.D.

Here is the critical property of $i(\delta, \{h\})$:

LEMMA. *Let δ be an ideal triangulation, let h be an oriented ideal arc which is not isotopic to any arc of δ , or to the opposite diagonal of any removable arc of δ . Suppose furthermore that h and δ intersect efficiently. Let x be the first point of $\text{int}(h)$ lying on an arc of δ , and suppose $x \in h^0 \in H(\delta)$. Then h^0 is removable; and if $\delta \rightarrow \delta'$ is the elementary move performed on h^0 , then $i(\delta', \{h\}) < i(\delta, \{h\})$.*

PROOF. First we must verify that h^0 is a removable arc of δ . Suppose otherwise: suppose h^0 is the interior arc of a puncture piece T of δ ; let \bar{h} be the boundary arc of T , and let p be the interior puncture of T . Since h^0 is the first arc of δ that h crosses, and since h and h^0 intersect efficiently, then there must be embedded closed intervals $\alpha \subset h$, $\beta \subset h^0$, satisfying the following properties (see Figure 24):

(i) $\partial\alpha = \partial\beta = \{x, y\}$, where y is the puncture at which \bar{h} has both ends; note that $x \in \text{int}(h) \cap \text{int}(h^0)$;

(ii) $\alpha \cup \beta$ bounds an embedded disc D with $p \in \text{int}(D)$.

It is evident, now, that after first crossing h^0 , h enters $\text{int}(D)$. There are now two possibilities. The first possibility is that h never leaves D ; this implies that h ends at p and is isotopic to h^0 , contradicting the hypothesis that h is not isotopic to an arc of δ . The second possibility is that h leaves D , it is easy to show that this implies the existence of a compressing disc for h and h^0 , contradicting the hypothesis that h intersects δ efficiently.

We have shown that h^0 is removable. Let $\delta \rightarrow \delta'$ be the elementary move performed on h^0 . We need some notation for the proof that $i(\delta', \{h\}) < i(\delta, \{h\})$.

Let Q be the quadrilateral of δ bisected by h^0 ; thus, Q is the support of $\delta \rightarrow \delta'$. We label the boundary arcs, triangles, and vertices of Q just as in the proof of the Coordinate Change Theorem for Elementary Moves, and we assume that h^0 bisects Q horizontally. Moreover, we assume that h starts from v_N and traverses T_N before its first intersection with h^0 . Since h intersects h^0 efficiently, and since h is not isotopic to an opposite diagonal of h^0 , then it follows that after first intersecting h^0 , h intersects one of h_{SW} or h_{SE} ; by symmetry, we can suppose that it is h_{SW} .

It is easy, now, to construct an opposite diagonal h' to h^0 such that there is no compressing disc for h and h' completely contained in Q : construct h' by starting from v_S , following along nearly parallel to h_{SW} until nearly encountering h , then follow along nearly parallel to h until encountering v_N . Mimicking the argument from the coordinate change theorem for elementary moves, it follows that h and h' intersect efficiently. For the remainder of the proof, we shall drop the braces in all our intersection number expressions, since all arcs concerned intersect efficiently.

Note that h_{SW} is the first arc of δ' whose interior intersects $\text{int}(h)$. Thus; we have the equation

$$(*) \quad i(\delta, h) - i(\delta', h) = i(h_{SW}, h) - i(h', h).$$

Let T_W be the triangle of δ' bounded by h_{NW} , h_{SW} , and h' . The lemma follows from equation (*), together with

$$(**) \quad i(h_{SW}, h) = i(h', h) + i(h', h_{NW}) + K, \quad \text{where } K = 1 \text{ or } 2.$$

To prove equation (**), first note that there is an initial segment of h in T_W connecting v_N to h_{SW} ; thus, every segment of h in T_W with one end on h_{NW} or h' must have its opposite end on h_{SW} . This accounts for all possible intersections of h with $\text{int}(h_{SW})$ with the following exceptions: there exists an initial segment of h in T_W connecting v_N to h_{SW} ; and there is also, possibly, a *final* segment of h in T_W connecting v_N to h_{SW} . Thus, the points of $[\text{int}(h) \cap \text{int}(h_{NW})] \cup [\text{int}(h) \cap \text{int}(h')]$ can be put in 1-1 correspondence with the points of $\text{int}(h) \cap \text{int}(h_{SW})$, with the exception of one or two points of $\text{int}(h) \cap \text{int}(h_{SW})$. Q.E.D.

COROLLARY. *If δ is an ideal triangulation and h is any ideal arc, there exists a finite sequence of elementary moves $\delta \rightarrow \cdots \rightarrow \delta'$ such that h is isotopic to an arc of δ' .*

PROOF. Applying induction and the above lemma, we find that there exists a finite sequence of elementary moves $\delta = \delta_0 \rightarrow \cdots \rightarrow \delta_m$ such that $i(\delta_m, \{h\}) = 0$. If h is not already isotopic to an arc of δ_m , then h is isotopic to the opposite diagonal of

some removable arc of δ_m , and it follows that there is an elementary move $\delta_m \rightarrow \delta_{m+1}$ such that h is isotopic to an arc of δ_{m+1} . Q.E.D.

REMARK. Given an ideal triangulation δ_0 and an ideal arc h , h determines an “intersection vector” $\nu_h \in \mathbf{Z}_{\geq 0}^{H(\delta_0)}$, whose coordinates are intersection numbers. ν_h does *not* satisfy all of the (nonstrict) triangle inequalities; in fact, if we compute the “triangle inequality excess” by summing $\max\{0, i(h_1, h) - i(h_2, h) - i(h_3, h)\}$ over all triples (h_1, h_2, h_3) forming the oriented boundary of a triangle of δ_0 , then we get exactly 2; and for any triple (h_1, h_2, h_3) such that $i(h_1, h) - i(h_2, h) - i(h_3, h) > 0$, we can choose h^0 to be h_1 . Thus, a necessary condition for an element of $\mathbf{Z}_{\geq 0}^{H(\delta_0)}$ to arise as the intersection vector of an arc is that the triangle inequality excess equals 2; but this is not a sufficient condition. In fact, we can apply the proof of the corollary to construct an algorithm which takes an element of $\mathbf{Z}_{\geq 0}^{H(\delta_0)}$, decides if it is the intersection vector of some arc h , and if so, computes a sequence of elementary moves $\delta_0 \rightarrow \cdots \rightarrow \delta_n$ such that $h \in H(\delta_n)$; the only missing ingredient is to have a formula for the “coordinate change map” $\mathbf{Z}_{\geq 0}^{H(\delta)} \rightarrow \mathbf{Z}_{\geq 0}^{H(\delta')}$ for intersection vectors of arcs, when $\delta \rightarrow \delta'$ is an elementary move; this formula can be easily worked out. Then at each stage, one computes the new vector, checks that the triangle inequality excess equals two, and if so, chooses the next h^0 and continues. If $\nu \in \mathbf{Z}_{\geq 0}^{H(\delta)}$ is not an intersection vector for some arc, eventually one arrives at a stage where the triangle inequality excess is not 2; if ν is an intersection vector, eventually one arrives at a stage where the transformed ν has a single nonzero coordinate equal to 1; if this is the \bar{h} coordinate of $\mathbf{Z}_{\geq 0}^{H(\delta_n)}$ for some $\bar{h} \in H(\delta_n)$, it follows that the original $\nu \in \mathbf{Z}_{\geq 0}^{H(\delta)}$ represents the intersection vector of an opposite diagonal of \bar{h} with respect to $H(\delta_n)$.

At this point, we already have enough to prove that $L_\delta: W(\delta) \rightarrow \mathcal{MF}$ is continuous, for if δ, δ' are as in the above corollary, then $L_{\delta'}^{-1} \circ L_\delta: W(\delta) \rightarrow W(\delta')$ is continuous, and since $h \in H(\delta')$ (up to isotopy), it follows that the composition of L_δ with $\langle \cdot, \{h\} \rangle$ is continuous. Nonetheless, here is a

PROOF OF CONNECTIVITY THEOREM. Note from the proof of the above corollary that, if h' is any arc of δ such that $i(h', h) = 0$, then the sequence of elementary moves $\delta \rightarrow \cdots \rightarrow \delta'$ can be chosen so that h' is an arc of each ideal triangulation along the way: essentially one needs only observe that the elementary moves are always performed on arcs having nonzero intersection number with h .

Now let δ, δ' be any two ideal triangulations, and let h_1, \dots, h_k be an enumeration of $H(\delta')$; notice that $i(h_i, h_j) = 0$ when $i \neq j$. Thus, there exists a sequence of elementary moves $\delta = \delta_0 \rightarrow \cdots \rightarrow \delta_{n(1)} \rightarrow \cdots \rightarrow \delta_{n(2)} \rightarrow \cdots \rightarrow \cdots \rightarrow \delta_{n(k)}$ such that for each $i = 1, \dots, k$, $h_i \in \bigcap_{j \geq n(i)} H(\delta_j)$. In particular, $h_i \in H(\delta_{n(k)})$ for each $i = 1, \dots, k$. Since the h_i are isotopically distinct, and since $|H(\delta_{n(k)})| = |H(\delta')|$, it follows that $H(\delta_{n(k)}) = H(\delta')$, and so $\delta_{n(k)} = \delta'$. Q.E.D.

6. For each ideal triangulation δ of (S, P) we define a decomposition $\mathcal{C}(\delta)$ of \mathcal{PL} by saying that two projective classes are in the same decomposition element if their respective normal forms have the same shape in each triangle of δ . To make this precise, we need a few definitions.

First recall the definition of a circular ordering on a set X : this is a partition of the collection of distinct ordered triples $\{(x_1, x_2, x_3) \mid x_i \neq x_j \text{ when } i \neq j\}$ into two subsets, the *positively ordered triples* and the *negatively ordered triples* such that the following axioms are satisfied:

- (i) if π is a permutation of $\{1, 2, 3\}$, then (x_1, x_2, x_3) and $(x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)})$ are identically oriented if and only if π is even;
- (ii) if (x_1, x_2, x_3) and (x_3, x_4, x_1) are positively ordered, then (x_2, x_3, x_4) is positively ordered.

For example, S^1 has a circular ordering, as well as any subset of S^1 . Note that a 2-point set and a 1-point set each have a unique circular ordering; they happen to be vacuous, since there do not exist distinct ordered triples.

Given a circularly ordered set X , we say that x' is the *successor* of x if there does not exist $x'' \in X$ such that (x, x'', x') is positively ordered; we also say that x is the *predecessor* of x' . In a 1-point set $\{x\}$, x is its own successor and predecessor; in a 2-point set $\{x, x'\}$, x and x' are each other's successors and predecessors. In general, successors and predecessors always exist in a *finite* circularly ordered set.

Now consider an ideal triangulation δ . Each arc h of δ has two *ends*, defined as the ends, in the usual topological sense, of $\text{int}(h)$. Each end is located at a particular puncture p , defined as the unique limit point of the end. The set of all arc ends of δ located at a particular $p \in P$ is denoted $E(\delta, p)$. $E(\delta, p)$ has a natural circular ordering which is defined as follows. Choose a closed coordinate disc D around p such that the arcs of δ intersect D in radii. Each such radius defines a unique element of $E(\delta, p)$. The endpoints of the radii inherit a circular ordering from the counterclockwise circular ordering on ∂D (this depends on the orientation of S). Clearly this induces a well-defined circular ordering on $E(\delta, p)$, which is natural with respect to the action of $\text{Homeo}_+(S, P)$. Note that since $E(\delta, P)$ is finite, successors and predecessors exist in $E(\delta, p)$.

A *prong* of δ located at p is defined to be an ordered pair (e, e') in $E(\delta, p)$ such that e' is the successor of e . The set of all prongs located at p is denoted $\Pi(\delta, p)$. Note that, for each $e \in E(\delta, p)$ there is a unique $e' \in E(\delta, p)$ and $e'' \in E(\delta, p)$ such that $(e', e) \in \Pi(\delta, p)$ and $(e, e'') \in \Pi(\delta, p)$. Note the following special cases, where $E(\delta, p)$ consists of two ends and one end, respectively: if $E(\delta, p) = \{e, e'\}$, then $\Pi(\delta, p) = \{(e, e'), (e', e)\}$; and if $E(\delta, p) = \{e\}$, then $\Pi(\delta, p) = \{(e, e)\}$; the latter case occurs precisely when p is the interior puncture of a puncture piece. Note that $\Pi(\delta, p)$ has its own natural circular ordering, where $(e_0, e_1), (e_2, e_3), (e_4, e_5)$ is positively ordered in $\Pi(\delta, p)$ if and only if e_0, e_2, e_4 is positively ordered in $E(\delta, p)$.

The set $\bigcup \{\Pi(\delta, p)\}_{p \in P}$ of all groups of δ will be denoted $\Pi(\delta)$, and an individual prong will usually be denoted with the variable π .

$\Pi(\delta)$ is in 1-1 correspondence with the set of corners of triangles of δ in the following sense: given a triangle T of δ , choose an orientation preserving characteristic map $\gamma: \tau \rightarrow S$ for T , where τ is a Euclidean triangle; we *define* the corners and sides of T to be the corners and sides of τ ; this gets around the problem of identifications under the characteristic map γ . Given a corner v of τ , let $p = \gamma(v)$; clearly there exists a unique prong $(e, e') \in \Pi(\delta, p)$ such that e and e' are



 $h^L, h^R, \text{ and } h^{\text{opp}}$

FIGURE 25

represented (respectively) by $\gamma(\tilde{e})$ and $\gamma(\tilde{e}')$, where \tilde{e} and \tilde{e}' represent the ends located at v of the two sides of τ incident to v . We shall use $\Pi(T)$ to denote the three prongs of δ corresponding in this manner to corners of T . Suppose that α is a closed interval embedded τ such that $\partial\alpha = \alpha \cap [\partial\tau - \{\text{corners}\}]$, and α separates v from the other two corners of τ ; suppose furthermore that γ does not identify the two boundary points of α (which is certainly the case if T is not a puncture piece); then $\gamma(\alpha)$ is an embedded closed interval in T , and we say that $\gamma(\alpha)$ is a *traversal* of the prong $\pi \in \Pi(T)$ corresponding to v , or that $\gamma(\alpha)$ *traverses* π .

Given $\mathcal{F} \in \mathcal{MF}$, $f \in \mathcal{F}$ in normal form with respect to δ , and a prong $\pi \in \Pi(\delta)$, we ask: does f have a leaf segment ℓ traversing π ? The *prong set* $\Pi(\delta, \mathcal{F})$ of \mathcal{F} relative to δ is defined to be the set consisting of those $\pi \in \Pi(\delta)$ for which there is no transversal of π by a leaf segment ℓ of f . This is well defined, for the uniqueness part of the Normal Form Theorem guarantees that the existence of such an ℓ is independent of the choice of a normal form representing \mathcal{F} . Also, we shall use $\Pi(T, \mathcal{F})$ to denote the set $\Pi(\delta, \mathcal{F}) \cap \Pi(T)$, for each triangle T of δ .

Here is an equivalent formulation of prong sets. Given $\pi = (e^L, e^R) \in \Pi(\delta)$, let h^L, h^R be the ideal arcs of δ having e^L, e^R (respectively) as ends. Let T be the triangle of δ (with characteristic map $\gamma: \tau \rightarrow S$) having a vertex v corresponding to π . Let h^{opp} be the image under γ of the side of τ opposite v ; h^{opp} is an ideal arc of δ . Thus, $h^L \cup h^R \cup h^{\text{opp}} = \partial T$; these arcs are referred to as the arcs to the left, right, and opposite from π (see Figure 25); as usual, two of these arcs may be identified, when T is a puncture piece.

Given $\mathcal{F} \in \mathcal{MF}$ and an ideal triangulation δ , let $\mu = L_\delta^{-1}(\mathcal{F}) \in W(\delta)$; let $\pi \in \Pi(\delta)$ be given, and let h^L, h^R, h^{opp} be as above. By definition of $W(\delta)$, we have $\mu(h^L) + \mu(h^R) \geq \mu(h^{\text{opp}})$; we abbreviate this by saying that μ satisfies the π -inequality. Similarly, we say that μ satisfies the π -equality if $\mu(h^L) + \mu(h^R) = \mu(h^{\text{opp}})$ and that μ satisfies the π -strict inequality if $\mu(h^L) + \mu(h^R) > \mu(h^{\text{opp}})$.

Given $\mu \in W(\delta)$, a normal form f_μ for $L(\mu)$ with respect to δ , and $\pi \in \Pi(\delta)$, it is clear from the definition of normal forms that μ satisfies the π -strict inequality if and only if there is a traversal of π by a leaf segment of f_μ . Thus:

PRONG SET LEMMA. *For each ideal triangulation δ and each $\mathcal{F} \in \mathcal{MF}$,*

$$\Pi(\delta, \mathcal{F}) = \{ \pi \in \Pi(\delta) \mid L_\delta^{-1}(\mathcal{F}) \text{ satisfies the } \pi\text{-strict inequality} \}.$$

We now define a decomposition $\mathcal{C}(\delta)$ on \mathcal{PF} by saying that $p(\mathcal{F}), p(\mathcal{G}) \in \mathcal{PF}$ are in the same decomposition element if and only if $\Pi(\delta, \mathcal{F}) = \Pi(\delta, \mathcal{G})$ in $\Pi(\delta)$ (note that this definition is independent of the choice of \mathcal{F} and \mathcal{G} in their projective classes). Clearly $\mathcal{C}(\delta)$ depends only on the isotopy class of δ . Also, the correspondence $\delta \rightarrow \mathcal{C}(\delta)$ is natural with respect to the action of $\text{Homeo}_+(S, P)$, for if f is in normal form with respect to δ and $\varphi \in \text{Homeo}_+(S, P)$, then $\varphi(f)$ is in normal form with respect to $\varphi(\delta)$, and φ takes the prong set of f in $\Pi(\delta)$ to the prong set of $\varphi(f)$ in $\Pi(\varphi(\delta))$. Thus, the correspondence $[\delta] \rightarrow \mathcal{C}(\delta)$ is well defined and natural with respect to the action of $\mathcal{M}(S, P)$ ($[\delta]$ denotes isotopy class).

It is not true that every subset of $\Pi(\delta)$ occurs as the prong set of some $\mathcal{F} \in \mathcal{MF}$. If $\Pi = \Pi(\delta, \mathcal{F})$ for some $\mathcal{F} \in \mathcal{MF}$, then we will say that Π is *realizable*. Later we shall give necessary and sufficient conditions for Π to be realizable. In either case, given an arbitrary subset $\Pi \subset \Pi(\delta)$, define $c(\delta, \Pi) = \{p(\mathcal{F}) \in \mathcal{PF} \mid \Pi(\delta, \mathcal{F}) \supset \Pi\}$, and define $c^0(\delta, \Pi) = \{p(\mathcal{F}) \in \mathcal{PF} \mid \Pi(\delta, \mathcal{F}) = \Pi\}$ (Note: “ \supset ” denotes nonproper containment). Thus, either $c^0(\delta, \Pi)$ is a decomposition element of $\mathcal{C}(\delta)$, or it is empty.

Now we examine the topological nature of the decomposition elements $c^0(\delta, \Pi)$.

PROPOSITION: DECOMPOSITION ELEMENTS ARE CELLS. *If Π is realizable, then $c(\delta, \Pi)$ is a topological closed cell, whose interior is $c^0(\delta, \Pi)$, and whose boundary is the union of those decomposition elements $c^0(\delta, \Pi')$ for which Π' is properly contained in Π .*

As a corollary, we have our main theorem:

CELL-DECOMPOSITION THEOREM. *$\mathcal{C}(\delta)$ is a cell-decomposition of \mathcal{PF} which is natural with respect to the action of $\mathcal{M}(S, P)$.*

To prove the proposition, we work in \mathcal{MF} , projecting our results down to \mathcal{PF} via the projection map $p: \mathcal{MF} \rightarrow \mathcal{PF}$. By the Normal Form Theorem, utilizing the homeomorphism $L: W(\delta) \rightarrow \mathcal{MF}$, we can actually work entirely in $W(\delta)$. The sets $L^{-1}(p^{-1}(c(\delta, \Pi)))$ and $L^{-1}(p^{-1}(c^0(\delta, \Pi)))$ will be denoted $W(\delta, \Pi)$ and $W^0(\delta, \Pi)$. Thus, we shall show that $W(\delta, \Pi) \subset W(\delta)$ is a convex, closed cone in $\mathbf{R}_{\geq 0}^{H(\delta)}$, whose manifold interior is $W^0(\delta, \Pi)$. From this it follows that $c(\delta, \Pi)$ is a closed cell, with interior $c^0(\delta, \Pi)$.

We shall assume for the remainder of the proof that $\Pi \subset \Pi(\delta)$ is a realizable prong set.

First we restate the definition of $W(\delta)$:

LEMMA. *$W(\delta)$ equals the set of all $\mu \in \mathbf{R}_{\geq 0}^{H(\delta)}$ such that:*

- (1) *For each $\pi \in \Pi(\delta)$, μ satisfies the π -inequality.*
- (2) *For each $p \in P$, there exists $\pi \in \Pi(\delta, p)$ such that μ satisfies the π -equality.*

PROOF. By definition, $W(\delta)$ is the set of $\mu \in \mathbf{R}_{\geq 0}^{H(\delta)}$ such that: (1) holds; and for each puncture $p \in P$, either there is an arc h incident to p such that $\mu(h) = 0$, or μ satisfies the π -equality for some $\pi \in \Pi(\delta)$ located at p . But if $\mu(h) = 0$ for some h incident to p , let π be any prong located at p adjacent to h ; say, h is the arc h^L to

the left of π . Since $\mu(h^L) = 0$, and since the two inequalities $\mu(h^L) + \mu(h^R) \geq (h^{\text{Opp}})$, $\mu(h^L) + \mu(h^{\text{Opp}}) \geq (h^R)$ hold, it follows that $\mu(h^R) = \mu(h^{\text{Opp}})$, and so $\mu(h^L) + \mu(h^R) = (h^{\text{Opp}})$, which is the π -equality. Q.E.D.

Now consider the set $W^0(\delta, \Pi) \subset W(\delta)$. By the Prong Set Lemma, $W^0(\delta, \Pi)$ is the set of all $\mu \in \mathbf{R}_{\geq 0}^{H(\delta)}$ satisfying the conditions:

- (a) $\mu \in W(\delta)$;
- (b) for all $\pi \in \Pi$, μ satisfies the π -equality;
- (c) for all $\pi \in \Pi(\delta) - \Pi$, μ satisfies the π -strict inequality.

Since Π is realizable, meaning that $W^0(\delta, \Pi) \neq \emptyset$, we can drop condition (a). To see why, notice that since $W^0(\delta, \Pi) \neq \emptyset$, then by the above lemma, Π contains at least one prong at each puncture. Thus, (a) follows from (b) and (c). Therefore, $W^0(\delta, \Pi) = \{\mu \in \mathbf{R}_{\geq 0}^{H(\delta)} \mid \text{(b) and (c) are satisfied}\}$.

Consider also the condition:

- (c') For all $\pi \in \Pi(\delta) - \Pi$, μ satisfies the π -inequality.

Let $W' = \{\mu \in \mathbf{R}_{\geq 0}^{H(\delta)} \mid \text{(b) and (c') are satisfied}\}$; we want to show that $W' = W(\delta, \Pi)$. Note that $W' \subset W(\delta)$: this is so because if $\mu \in W'$, then the subset $\{\pi \in \Pi(\delta) \mid \mu \text{ satisfies the } \pi\text{-equality}\}$ includes Π , and so includes at least one prong at each puncture. Thus, $W' \subset W(\delta) \cap W(\delta, \Pi) = W(\delta, \Pi)$; the opposite inclusion is trivial. Also, it is clear that W' is the closure in $\mathbf{R}_{\geq 0}^{H(\delta)}$ of $W^0(\delta, \Pi) = \{\mu \in \mathbf{R}_{\geq 0}^{H(\delta)} \mid \text{(b) and (c) are satisfied}\}$; since these sets are both contained in $W(\delta)$, it follows that $W' = W(\delta, \Pi)$ is the closure in $W(\delta)$ of $W^0(\delta, \Pi)$.

Thus, $W(\delta, \Pi)$ is an intersection in $\mathbf{R}^{H(\delta)}$ of closed linear half-spaces, linear subspaces, and $\mathbf{R}_{\geq 0}^{H(\delta)}$; therefore $W(\delta, \Pi)$ is a convex, closed cone in a Euclidean space, and it follows that $\epsilon(\delta, \Pi) = p(L(W(\delta, \Pi)))$ is a closed topological cell.

To see that $W^0(\delta, \Pi)$ is an open cone in a Euclidean space, suppose that $\mu(h) = 0$ for some $\mu \in W^0(\delta, \Pi)$ and $h \in H(\delta)$. Then each prong in $\Pi(\delta)$ incident to h is in Π . From this it follows that $\mu(h) = 0$ for all $\mu \in W(\delta, \Pi)$. Set $H' = \{h \in H(\delta) \mid \text{some prong of } \Pi(\delta) \text{ incident to } h \text{ is in } \Pi(\delta) - \Pi\}$; under the standard inclusion $\mathbf{R}^{H'} \subset \mathbf{R}^{H(\delta)}$, where the remaining coordinates are set to zero, it follows that $W^0(\delta, \Pi) \subset \mathbf{R}^{H'}$ is the intersection of open linear half-spaces, linear subspaces, and $\mathbf{R}_+^{H'}$ (not $\mathbf{R}_{\geq 0}^{H'}$). Thus, $W^0(\delta, \Pi)$ is an open cone in some Euclidean space. Since $\mathbf{R}^{H'}$ is closed in $\mathbf{R}^{H(\delta)}$, then $W(\delta, \Pi)$ is the closure of $W^0(\delta, \Pi)$ in \mathbf{R}^H , and it follows that $W^0(\delta, \Pi)$ is the manifold interior of $W(\delta, \Pi)$; so $\epsilon^0(\delta, \Pi)$ is the manifold interior of $\epsilon(\delta, \Pi)$.

Clearly $\partial W(\delta, \Pi)$ is the union of all $W^0(\delta, \Pi')$ such that $\Pi' \supset \Pi$ and $\Pi' \neq \Pi$; this follows from conditions (b) and (c') for Π , and conditions (b) and (c) for Π' . Thus, $\partial \epsilon(\delta, \Pi) = \bigcup \{\epsilon(\delta, \Pi') \mid \Pi' \supset \Pi, \Pi' \neq \Pi\}$. Q.E.D.

7. In this section we shall apply the techniques of triangulations and elementary moves to the proof of Thurston's theorem that \mathcal{PF} is a compact manifold of dimension $6 \cdot g - 7 + 2 \cdot |P|$, where g is the genus of S . Compactness follows from the finiteness of $\mathcal{C}(\delta)$. To prove that \mathcal{PF} is a manifold of the appropriate dimension takes several steps. First we shall compute the dimensions of cells of $\mathcal{C}(\delta)$, for any ideal triangulation δ ; from this it will follow that the top-dimensional

cells have dimension $6 \cdot g - 7 + 2 \cdot |P|$. Then we will show that each point in \mathcal{PF} lies in the top dimensional cell of $\mathcal{C}(\delta)$, for some δ .

It is beyond the scope of this work to prove the additional fact, due again to Thurston, that \mathcal{PF} is a sphere. The reason for this difficulty is that neither the cell-decomposition $\mathcal{C}(\delta)$ nor the coordinate space $W(\delta)$ lend themselves to a direct combinatorial attack; this is in contrast to the “pairs of pants” coordinates on a surface, for which it is immediately clear that the coordinate space is homeomorphic to Euclidean space. It should be possible to revert to Thurston’s original proof, given in [T], where he shows that \mathcal{PF} is covered by just two open sets, and then uses the Alexander trick; for instance, one could try to find two ideal triangulations δ, δ' and a top dimensional cell in each of $\mathcal{C}(\delta), \mathcal{C}(\delta')$ which together cover \mathcal{PF} .


Another purpose of this section is to relate the techniques of ideal triangulations to the more “classical” setting of train track theory; in fact, we shall apply train track theory to get results about dimensions of cells of $\mathcal{C}(\delta)$. So we commence with a short exposition of train tracks, consisting mostly of well-known “folk” theorems that have previously managed to escape print (or, at least, the none too careful eyes of this writer); in particular, we shall need a computation of the dimension of the space of invariant measures on an arbitrary train track.

The definition of a train track that we use is very general and abstract to start with; we shall soon show how to relate this to the standard theory of train tracks on a surface.

An *abstract train track* τ is a finite 1-complex with no isolated 0-cells, whose 1-cells are called *branches* and whose 0-cells are called *switches*, and for each switch s , there is a partition of the set of branch ends located at s into two subsets, called *sides* of s . We *do* allow the possibility that one of the sides of s is empty; when this happens, we shall call s a *terminus* of τ .

Given an orientation for a branch b of τ , a branch end e of b induces an orientation of the switch s at which e is located: a “+” orientation on s if b points toward e , and a “−” orientation on s if b points away from e . An *orientation* of τ is a choice of orientation for each branch of τ with the property that, for each switch s of τ and any two branch ends e, e' located at s , e and e' induce the same orientation on s if and only if e and e' are in the same side of s . τ is *orientable* if there exists an orientation for τ ; otherwise τ is *nonorientable*.

If τ is an abstract train track with branch set B and switch set S , consider the vector space \mathbf{R}^B and its dual space \mathbf{R}^{B*} . Let $\{v_b\}_{b \in B}$ be the standard basis for \mathbf{R}^B , where $\langle v_b, b' \rangle = \delta_{b,b'}$. Let $\{v_b^*\}_{b \in B}$ be the dual basis of \mathbf{R}^{B*} . For each switch $s \in S$, we define a covector $\alpha_s \in \mathbf{R}^{B*}$ called a *switch condition* as follows: let $e(1), \dots, e(m)$ and $e'(1), \dots, e'(n)$ be enumerations of the two sides of s , and let $b(1), \dots, b(m)$ and $b'(1), \dots, b'(n)$ denote the corresponding branches; α_s is defined to be $v_{b(1)}^* + \dots + v_{b(m)}^* - v_{b'(1)}^* - \dots - v_{b'(n)}^*$. This definition depends on an ordering of the two sides of s ; but clearly α_s is well defined up to a factor of -1 . When τ is an orientable train track, we shall adopt the following convention: once an orientation of τ is fixed, the ordering of sides of s is chosen so that the side $e(1), \dots, e(m)$ consists of ends inducing the “+” orientation on s , and the side $e'(1), \dots, e'(n)$ consists of ends inducing the “−” orientation.



Abstract train tracks with one branch

FIGURE 26

We now define an *invariant measure* on τ to be an element of \mathbf{R}^B which is in the null-space of $\{\alpha_s\}_{s \in S}$; and let $W(\tau)$ denote the space of all invariant measures on τ . Thus, $\mu \in \mathbf{R}^B$ is in $W(\tau)$ if and only if, for each switch s , $\mu(b(1)) + \cdots + \mu(b(m)) = \mu(b'(1)) + \cdots + \mu(b'(n))$.

PROPOSITION. *If τ is a connected abstract train track, then $\dim(W(\tau))$ is computed as follows:*

If τ is nonorientable, $\dim(W(\tau)) = |B| - |S|$,

If τ is orientable, $\dim(W(\tau)) = |B| - |S| + 1$.

PROOF. The proof is to show that the switch conditions are linearly independent, with the exception of a single dependency relation when τ is orientable. First we describe that relation: if τ is orientable, notice that for each $b \in B$ the term v_b^* occurs once with a positive sign and once with a negative sign among the list of switch conditions $\{\alpha_s\}$; this follows from our convention for the choice of α_s when τ is orientable. It follows that $\sum \alpha_s = 0$.

Claim. Given an abstract train track τ , if τ is nonorientable then the switch conditions are linearly independent; and if τ is orientable then the only dependency relations on the switch conditions have constant coefficients. From this the proposition clearly follows. We shall prove the claim by induction on the number of branches.

To start the induction, consider the abstract train tracks having exactly one branch. There are three of these, one nonorientable and two orientable; they are pictured in Figure 26.

In the nonorientable example, the only switch condition is $\alpha_s = 2 \cdot v_b^*$, which is nonzero and so forms an independent set. In the first orientable example, the line segment, the two switch conditions are $\alpha_s = v_b^*$ and $\alpha_{s'} = -v_b^*$, and the claim clearly follows. In the second orientable example, the circular track, the only switch condition is $\alpha_s = v_b^* - v_b^* = 0$, and the claim also follows.

Now suppose τ is an arbitrary connected train track with ≥ 2 branches. We shall consider several cases, depending on the combinatorial structure of τ .

Case 1. τ has a single switch s . Thus, every branch has both ends at s , and τ is orientable if and only if each branch has its ends in opposite sides of s .

If τ is orientable, then the unique switch condition is just $\alpha_s = 0$, and the claim follows.

If τ is nonorientable, then there exists at least one branch having both ends in the same side of s . Let $b(1), \dots, b(m)$ denote those branches with both ends on one side of s and $b'(1), \dots, b'(n)$ denote those branches with both ends on the other side of s ; these two sets are disjoint and at least one is nonempty, so $\alpha_s = 2 \cdot v_{b(1)}^* + \cdots + 2 \cdot v_{b(m)}^* - 2 \cdot v_{b'(1)}^* - \cdots - 2 \cdot v_{b'(n)}^*$ is nonzero, and the claim follows.

Case 2. Now suppose τ is connected and has at least two switches. Thus, there exists a branch b connecting two distinct switches s, s' of τ . In the analysis which follows, we shall write the collection of switch conditions of τ as an integer matrix, with rows indexed by S and columns indexed by B . There are again several subcases, depending on the combinatorial properties of τ .

Case 2a. Suppose that one of the two switches s, s' , say s , has a unique branch end located there, necessarily one of the ends of b . Since τ has ≥ 2 branches and is connected, s' has other branch ends located there. Note that τ is orientable if and only if $\tau - b$ is orientable. Applying induction, the claim is true for $\tau - b$; we shall compare the matrix of switch conditions of τ to that of $\tau - b$.

The matrix of switch conditions of τ can be written in the following form:

$$M \quad S \quad \begin{matrix} & & B \\ & & b \quad \cdot \quad \cdot \quad \cdot \\ s & \left(\begin{array}{c|ccc} 1 & & & \\ \hline -1 & & & 0 \\ \hline 0 & & A & \end{array} \right) \\ s' & & & \\ \vdots & & & \\ \vdots & & & \end{matrix}$$

Case 2a. The matrix of switch conditions of τ

It is clear, then, that the matrix of switch conditions of $\tau - b$ is just A ; also, row s' of M must have other nonzero entries (since s' has other branch ends located there), so row s' of A has nonzero entries. Clearly any dependency relation on the rows of M has equal coefficients on row s and row s' .

If τ is nonorientable, then $\tau - b$ is nonorientable, so by induction, the rows of A are linearly independent; thus, any dependency relation on the rows of M must have zero coefficients on all rows except possibly row s ; it follows that all coefficients are zero.

If τ is orientable, then $\tau - b$ is orientable, so by induction, the only dependency relations on rows of A have constant coefficients. Thus, any dependency relation on rows of M must have constant coefficients on all rows except possibly row s ; it follows that all coefficients are equal.

Case 2b. Now suppose that s and s' both have ≥ 2 ends located there. Again, $\tau - b$ has one fewer branch, but now it may happen that $\tau - b$ is disconnected. Nonetheless, each component has fewer total branches, so the induction hypothesis applies to each component of $\tau - b$. Again we have several subcases:

Case 2bi, $\tau - b$ is connected. Thus, the matrix of switch conditions of τ can be written as follows:

$$M \quad S \quad \begin{matrix} & & B \\ & & b \quad \cdot \quad \cdot \quad \cdot \\ s & \left(\begin{array}{c|ccc} x & & & \\ \hline y & & & \\ \hline 0 & & A & \end{array} \right) \\ s' & & & \\ \vdots & & & \\ \vdots & & & \end{matrix}$$

Case 2bi. The matrix of switch conditions of τ

In this diagram, $x, y = \pm 1$. Notice that the matrix of switch conditions of $\tau - b$ is A . Notice also that rows s and s' of A have nonzero entries, since s and s' both have ends other than ends of b located there.

If τ is nonorientable and $\tau - b$ is nonorientable, then by induction, the rows of A are linearly independent, so the rows of M are linearly independent.

If τ is nonorientable and $\tau - b$ is orientable, pick an orientation for $\tau - b$. Let e denote the end of b located at s , and let E denote the side of s in τ containing e ; let E' be similarly defined. Since τ is nonorientable, it follows that the ends $E - \{e\}$ and $E' - \{e'\}$ induce the same orientation on s and s' ; our convention for α_s and $\alpha_{s'}$ now implies that x and y are equal. Now any dependency relation on the rows of M is also a dependency relation on the rows of A , and by induction, it follows that all coefficients of the relation are equal. However, since $x = y = \pm 1$, it follows that the coefficients on rows s and s' are zero. Thus, all coefficients are zero, and the rows of M are linearly independent.

If τ is orientable, then $\tau - b$ is orientable; any dependence relation on rows of M is also a dependency relation on rows of A , so by induction, the relation has constant coefficients.

Case 2bii. $\tau - b$ is disconnected, consisting of two components τ_1 and τ_2 . Thus, the matrix of switch conditions for τ can be written as

$$M \quad \begin{matrix} s \\ \cdot \\ s' \\ \cdot \end{matrix} \quad \begin{array}{c|cc} & \begin{matrix} b \\ 1 \\ 0 \end{matrix} & \begin{matrix} \cdot \\ \cdot \\ \cdot \end{matrix} & \\ \hline & A_1 & & 0 \\ \hline & \begin{matrix} -1 \\ 0 \end{matrix} & 0 & A_2 \end{array}$$

Case 2bii. The matrix of switch conditions of τ


It is clear that A_1 and A_2 are the matrices of switch conditions for τ_1 , τ_2 , respectively.

If τ is oriented, then τ_1 and τ_2 are oriented, and any dependency relation on the rows of M gives rise to a dependency relation on the rows of A_1 and on the rows of A_2 . By induction, the coefficients on the rows of A_1 are equal, and the coefficients on the rows of A_2 are equal. But the coefficients on row s and row s' must also be equal, so the dependency relation has constant coefficients.

If τ is nonorientable, then at least one of τ_1 , τ_2 must also be nonorientable.

If τ_1 , τ_2 are both nonorientable, then any dependency relation on rows of M must, by induction, be zero on rows of A_1 and on rows of A_2 , so the rows of M are linearly independent.

If, say, τ_1 is orientable and τ_2 is nonorientable, then given a dependency relation on rows of M , the rows of A_1 must have constant coefficients, and the rows of A_2 must have zero coefficients. Since row s and row s' must have equal coefficients, it follows that all rows have zero coefficients. Q.E.D.



No smooth simple closed curve contains the middle branch of τ

FIGURE 27

Recall now the definition of “recurrence”: a *smooth closed curve* in a train track τ is a map $f: S^1 \rightarrow \tau$ which is a local injection, and which always crosses a switch from one side to the opposite side; τ is called *recurrent* if, for each branch b of τ , there exists a smooth closed curve in τ whose image contains b . Note that if τ has any terminus, then τ is not recurrent.

The definition that we have given for recurrence is not an effective definition, as it requires checking through infinitely many homotopy classes of smooth closed curves. However, it is easy to parlay this into an effective condition: given a train track τ , the *oriented double cover* $\tilde{\tau}$ of τ is defined just as in manifold theory: the branches of $\tilde{\tau}$ are in 1-1 correspondence with *oriented* branches of τ . Just as in manifold theory, if τ is connected, then τ is orientable if and only if $\tilde{\tau}$ is disconnected.

RECURRENCE LEMMA. τ is recurrent if and only if, for each branch \tilde{b} of the oriented double cover $\tilde{\tau}$ of τ , \tilde{b} is contained in the image of some smooth simple closed curve in $\tilde{\tau}$. Consequently, τ is recurrent if and only if, for each branch b of τ , b is in the image of a smooth oriented closed curve f such that, for each branch b' of τ , f passes over b' at most once in each direction.

Note. The necessity of taking oriented double covers is shown by the accompanying example of a recurrent train track τ , which has the property that the middle branch of τ is not contained in the image of any smooth simple closed curve in τ . Figure 27 also shows the oriented double cover of τ .

PROOF. The proof is very easy. First one observes that τ is recurrent if and only if $\tilde{\tau}$ is recurrent. This follows from the fact that: any smooth closed curve in τ lifts to a smooth closed curve in $\tilde{\tau}$, since S^1 is orientable; and any smooth closed curve in $\tilde{\tau}$ projects down to a smooth closed curve in τ .

Now we show that an *orientable* train track τ is recurrent if and only if, for each branch b of τ , there exists a smooth *simple* closed curve passing through b . One direction is trivial. For the other direction, suppose τ is recurrent; let b be a branch of τ , and let f be a smooth closed curve in τ passing through b , which has the minimum possible number of “branch traversals”, where a *branch transversal* of f :

$S^1 \rightarrow \tau$ is a segment in S^1 whose interior maps bijectively to the interior of a branch of τ . We shall show that s simple. Suppose not, then there is a branch b' of τ that f passes over at least twice; moreover, by orientability of S^1 and of τ , each time f passes over b' it passes over in the same direction. Therefore, if we choose two points $x_1, x_2 \in S^1$ which map to the same point of b' , then we can cut and paste f at x_1 and x_2 to form two smooth closed curves f_1, f_2 in τ . At least one of these passes through b , and each has fewer branch traversals than f . Q.E.D.

Now we relate recurrence to the existence of transverse measures with positive weights. Let $W_{\geq 0}(\tau)$ denote the subspace of $W(\tau)$ consisting of elements with nonnegative coordinates, and let $W_+(\tau)$ denote those elements with all positive coordinates.

PROPOSITION. τ is recurrent if and only if $W_+(\tau) \neq \emptyset$.

PROOF. This is a familiar proof to anyone acquainted with train tracks; we include it here for completeness.

Each smooth closed curve c in τ gives rise to an element $\mu_c \in W_{\geq 0}(\tau)$ by counting, for each branch b of τ , the number of times that c passes over b . If τ is recurrent, let $c(b)$ denote a smooth closed curve passing through b . It follows that $\mu_{c(b)}(b) > 0$; setting $\mu = \sum_b \mu_{c(b)}$, then $\mu \in W_+(\tau)$.

Conversely, suppose $W_+(\tau) \neq \emptyset$. Note that each switch condition has rational coefficients. Thus, the points of $W_+(\tau)$ are all the points in \mathbf{R}^B which are solutions to a certain collection of strict inequalities and rational equations. Since $W_+(\tau) \neq \emptyset$, it follows that there exists a real solution, and therefore a rational solution. Clearing the denominators of all the coefficients gives an integer solution. Thus, there exists $\mu \in W_+(\tau)$ with all integer coefficients.

Now we construct a finite collection of smooth closed curves $c(i)$ in τ such that $\sum_i \mu_{c(i)} = \mu$. Choose a collection of $\sum_{b \in B} \mu(b)$ disjoint intervals, and for each $b \in B$, choose maps of $\mu(b)$ of these intervals onto b each of which is a homeomorphism on the interior. For each switch s of τ , since μ satisfies the switch condition α_s , the set of ends of intervals mapped to one side of s is equal in size to the set of ends of intervals mapped to the other side of s ; so choose an arbitrary bijection between these two sets, and identify corresponding boundary points. This gives a pairwise identification of the set of all ends of intervals, giving rise to finitely many disjoint circles. Each of these circles is mapped smoothly into τ , giving rise to closed curves c_i which evidently satisfy the property that $\sum_i \mu_{c(i)} = \mu$. Since $\mu(b) > 0$ for each $b \in B$, then there exists some $c(i)$ with $\mu_{c(i)}(b) > 0$. Q.E.D.

Now we turn to train tracks on surfaces. Consider a punctured surface (S, P) and a finite 1-complex $\tau \subset S - P$. For each 0-cell v of τ , there exists a well-defined natural circular ordering on the set of ends of 1-cells located at v : this ordering is defined by choosing a closed coordinate disc D around v such that v corresponds to the origin, and $D \cap \tau$ corresponds to a collection of radial segments; the circular ordering is then induced by the natural circular ordering on ∂D . Suppose τ is given the structure of an abstract train track, still denoted by τ (i.e. choose sides, for each switch of τ); we say that τ is *embedded* in S if, for each switch s of τ , each side of s


is a circular interval, in the natural circular ordering on the set branch ends located at s .

Let τ be an abstract train track embedded in (S, P) , and suppose that τ has no terminus. Consider a complementary component C of τ in S . C is an open topological surface, and there is a natural way to complete C to a surface with boundary \bar{C} having the following property: there is a map $\varphi: \bar{C} \rightarrow S$ which is the identity on C and which immerses $\partial\bar{C}$ into τ . We denote $\varphi|_{\partial\bar{C}}$ by $\partial\varphi$. Note that $\partial\varphi$ is not necessarily smooth, for as you travel around $\partial\bar{C}$, there will be isolated points, called *boundary cusps*, where you double back on the same side of a switch of τ . We say that the complementary component C is *hyperbolic* if the surface-with-boundary \bar{C} -{boundary cusps}-{punctures} is hyperbolic, i.e. if the double of this surface across its boundary has negative Euler characteristic. The kinds of complementary components that are eliminated by this condition are as follows: discs with m boundary cusps and n punctures, where $(m, n) = (0, 0), (0, 1), (1, 0)$ or $(2, 0)$; and annuli with no cusps and no punctures. Finally, we say that τ is a *train track on* (S, P) if τ has no terminus, and each complementary component of τ is hyperbolic.

In order to motivate our use of train tracks, here is a short expository account of the theory of train tracks and their relation to measured foliations; this account has no role in the logical flow of this paper, other than to clarify scattered comments in what follows. A more complete study of train tracks is made in [T].

Given a train track τ on (S, P) and an element $\mu \in W_{\geq 0}(\tau)$, there is a construction due to Thurston for a partial measured foliation f_μ whose class $L(\mu) = \{f_\mu\}$ is well defined in \mathcal{MF} : this is the famous highways construction. One replaces each branch b of τ with a “highway” or rectangle foliated by leaves parallel to b , having a transverse measure of total weight $\mu(b)$; then for each switch s of τ , one glues ends of highways together in a measure-preserving manner; that this can be done is guaranteed by the switch condition $\langle \alpha_s, \mu \rangle = 0$. In order to obtain a true partial measured foliation, one must then collapse complementary components of the support which are nonpunctured or once-punctured discs. The primary property of this map $L: W_{\geq 0}(\tau) \rightarrow \mathcal{MF}$ is that if τ is recurrent and “transversely recurrent” (see [T] for a definition), then L is injective. The image of L can be characterized as follows: we say that a partial measured foliation f is *carried* by an abstract train track τ embedded in (S, P) if, after replacing each singularity of f by a disc, there is a homotopy from the identity of $\text{supp}(f)$ to a map $\varphi: \text{supp}(f) \rightarrow \tau$ such that for each smooth leaf segment α of f , $\varphi(\alpha)$ maps smoothly into τ . If τ is a train track on (S, P) , then $\text{image}(L: W_{\geq 0}(\tau) \rightarrow \mathcal{MF})$ consists of the classes of all partial measured foliations carried by τ . In fact, if f and φ are as above, then for each point x on a branch b of τ , $\varphi^{-1}(x)$ is a transversal to f whose f -measure depends only on b ; this defines a transverse measure on τ depending only on the class of f , and it is easy to see that this is the inverse of the highways construction, up to equivalence. Thus, τ is recurrent if and only if there exists a partial measured foliation f carried on τ with positive weights on every branch.

For our purposes, we shall not need all of these concepts, since we have already done the hard work using ideal triangulations; we shall refer to these concepts only to show how certain of our results can be couched in the “classical” terminology.



A piece of train track in a triangle

FIGURE 28

However, we *shall* use train tracks to decide when a subset $\Pi \subset \Pi(\delta)$ is realizable, and to compute dimensions of cells in the cell-decomposition $\mathcal{C}(\delta)$.

First we give two necessary conditions for $\Pi \subset \Pi(\delta)$ to be realizable. The first condition is that $\Pi \cap \Pi(\delta, p) \neq \emptyset$, for $p \in P$. We refer to this as the *puncture condition* with respect to p . We have already seen that for Π to be realizable, it is necessary that Π satisfy all puncture conditions.

Consider now an arc $h \in H(\delta)$. Suppose first that h is a removable arc, bounding distinct triangles T, T' ; then there are four prongs adjacent to h , which come in two pairs, one pair $\{\pi_1, \pi_2\} \subset \Pi(T)$, the other pair $\{\pi'_1, \pi'_2\} \subset \Pi(T')$. Given $\Pi \subset \Pi(\delta)$, we say that Π satisfies the *arc condition* with respect to h if: assuming that Π contains both prongs in one pair, then Π contains both prongs in the other pair. To see that this is necessary, suppose $\mu \in W^0(\delta, \Pi)$ and, say, $\{\pi_1, \pi_2\} \subset \Pi$; then μ satisfies π_1 -equality and π_2 -equality, therefore $\mu(h) = 0$, implying that μ satisfies π'_1 -equality and π'_2 -equality, so $\{\pi'_1, \pi'_2\} \subset \Pi$.

There is also an arc condition for a nonremovable arc $h \in H(\delta)$. h is the interior arc of some puncture piece T , whose interior puncture is p . There is a unique prong $\pi_1 \in \Pi(T) \cap \Pi(\delta, p)$ and $\Pi(T)$ has two other prongs π_2, π_3 located at the boundary puncture of T . The puncture condition on Π implies that $\pi_1 \in \Pi$. We say that Π satisfies the *arc condition* with respect to h if: $\pi_1 \in \Pi$; and either $\{\pi_2, \pi_3\} \subset \Pi$ or $\{\pi_2, \pi_3\} \cap \Pi = \emptyset$. To see that this is necessary, if $\mu \in W^0(\delta, \Pi)$ and if, say, $\pi_2 \in \Pi$, then μ satisfies π_2 -equality; from the fact that μ satisfies π_1 equality it follows that $\mu(h) = 0$, and so μ also satisfies π_3 -equality; thus $\pi_3 \in \Pi$.

Now we shall show how to associate, to each $\Pi \subset \Pi(\delta)$ satisfying all puncture and arc conditions, a train track $\tau(\delta, \Pi)$ on (S, P) ; this will be done in such a way that Π is realizable if and only if $\tau(\delta, \Pi)$ is recurrent.

First of all, given an ideal triangulation δ , embed an abstract train track $\tau(\delta)$ in (S, P) by putting a piece of train track in each triangle of δ , as indicated in Figure 28; the diagrams are drawn as if there were a smooth structure on S , so that all branch ends at a given switch s are tangent, with two branch ends tangent from the same direction if and only if they are in the same side of s ; this, of course, is the classical definition of a train track on a smooth surface. For each arc $h \in H(\delta)$, there is a branch $b(h)$ intersecting h once transversely; and for each prong $\pi \in \Pi(\delta)$, there is a branch $b(\pi)$ with the property that any traversal of π

embedded smoothly in τ must contain $b(\pi)$. Notice that the only thing that keeps $\tau(\delta)$ from being a train track on (S, P) is that each puncture p is contained in a nonhyperbolic complementary component, which is a disc with no boundary cusps and one puncture (namely, p).

Incidentally, the Normal Form Theorem can be couched in terms of the abstract embedded train track $\tau(\delta)$: that theorem says simply that each class in \mathcal{FF} is represented by some partial measured foliation carried on $\tau(\delta)$.

Consider now an ideal triangulation δ and a set of prongs $\Pi \subset \Pi(\delta)$ satisfying all puncture and arc conditions. Define $\tau(\delta, \Pi)$ to be the train track obtained by erasing all branches of the form $b(\pi)$ where $\pi \in \Pi$, and all branches of the form $b(h)$ where each prong adjacent to h is in Π . Notice that the arc conditions on Π guarantee that $\tau(\delta, \Pi)$ has no terminus; and the puncture conditions on Π guarantee that $\tau(\delta, \Pi)$ has all hyperbolic complementary components.

PROPOSITION. *Given $\Pi \subset \Pi(\delta)$ satisfying all puncture and arc conditions, Π is realizable if and only if $\tau(\delta, \Pi)$ is recurrent. When this happens, then $\dim(W^0(\delta, \Pi)) = \dim(W_+(\tau(\delta, \Pi)))$, and $\dim(c^0(\delta, \Pi)) = \dim(W_+(\tau(\delta, \Pi))) - 1$.*

PROOF. Recall that $\tau(\delta, \Pi)$ is recurrent if and only if $W_+(\tau(\delta, \Pi)) \neq \emptyset$. Recall also that $W^0(\delta, \Pi) = L^{-1}(p^{-1}(c^0(\delta, \Pi))) = \{\mu \in \mathbf{R}_{\geq 0}^{H(\delta)} \mid \text{for each } \pi \in \Pi, \mu \text{ satisfies the } \pi\text{-equality, and for each } \pi \in \Pi(\delta) - \Pi, \mu \text{ satisfies the } \pi\text{-strict inequality}\}$; also, Π is realizable if and only if $W^0(\delta, \Pi) \neq \emptyset$. We shall construct a homeomorphism $\chi: W_+(\tau(\delta, \Pi)) \rightarrow W(\delta, \Pi)$; the proposition follows directly from the existence of χ .

In fact, first we define $\chi: W_{\geq 0}(\tau(\delta)) \rightarrow \mathbf{R}_{\geq 0}^{H(\delta)}$, and then show that the restriction to $W_+(\tau(\delta, \Pi))$ is a homeomorphism onto $W(\delta, \Pi)$. The definition is simple: given $\mu \in W_{\geq 0}(\tau(\delta))$, define $\chi(\mu) \in \mathbf{R}_{\geq 0}^{H(\delta)}$ by $\chi(\mu)(h) = \mu(b(h))$. First note that $\chi(\mu)$ automatically satisfies the triangle inequalities, for if $\pi \in \Pi(\delta)$ and h^L, h^R, h^{Opp} are the arcs to the left, right and opposite π , then

(*)

$$\begin{aligned} \chi(\mu)(h^L) + \chi(\mu)(h^R) - \chi(\mu)(h^{\text{Opp}}) &= \mu(b(h^L)) + \mu(b(h^R)) - \mu(b(h^{\text{Opp}})) \\ &= 2 \cdot \mu(b(\pi)) \geq 0. \end{aligned}$$

Evidently χ is 1-1. Moreover, the image of χ is $\{\nu \in \mathbf{R}_{\geq 0}^{H(\delta)} \mid \nu \text{ satisfies } \pi\text{-inequality, for all } \pi \in \Pi(\delta)\}$: for if ν satisfies all π -inequalities, then we can define $\chi^1(\nu) \in W_{\geq 0}(\tau(\delta))$ by

$$\chi^{-1}(\nu)(b(\pi)) = (1/2) \cdot [\nu(h^L) + \nu(h^R) - \nu(h^{\text{Opp}})],$$

where h^L, h^R and h^{Opp} are the arcs to the left, right and opposite π . Evidently χ and χ^1 are continuous.

Now suppose that $\Pi \subset \Pi(\delta)$ satisfies all puncture and arc conditions; thus, $\tau(\delta, \Pi)$ is defined and is a train track on (S, P) . Note that $W_{\geq 0}(\tau(\delta, \Pi))$ is naturally embedded in $W_{\geq 0}(\tau(\delta))$, by defining $\mu \in W_{\geq 0}(\tau(\delta, \Pi))$ to be zero on all branches in $\tau(\delta) - \tau(\delta, \Pi)$. We shall henceforth identify $W_{\geq 0}(\tau(\delta, \Pi))$ with its image in $W_{\geq 0}(\tau(\delta))$. It remains to show that χ takes $W_+(\tau(\delta, \Pi))$ onto $W^0(\delta, \Pi)$.

Equation (*) above show that, for $\mu \in W_{\geq 0}(\tau(\delta))$, $\chi(\mu)$ satisfies the π -strict inequality if and only if $\mu(b(\pi)) > 0$. Thus, $\chi(\mu) \in W^0(\delta, \Pi)$ if and only if $\mu \in W_+(\tau(\delta, \Pi))$. Since $W^0(\delta, \Pi) \subset \{\mu \in \mathbf{R}_{\geq 0}^{H(\delta)} \mid \text{all prong inequalities are satisfied}\} = \text{image}(\chi)$, it follows that $\chi|_{W_+(\tau(\delta, \Pi))}$ is a homeomorphism onto $W(\delta, \Pi)$. Q.E.D.


COROLLARY. $\Pi \subset \Pi(\delta)$ is realizable if and only if the following conditions are satisfied:

- (i) Π satisfies all puncture and arc conditions.
- (ii) For each prong $\pi \in \Pi(\delta) - \Pi$, there exists an essential simple closed curve $c \subset S - P$, intersecting δ efficiently, such that: c traverses π ; c traverses no prong in Π ; c traverses each prong in $\Pi(\delta) - \Pi$ at most once in each direction; and c crosses each arc of δ at most once in each direction.

PROOF. First suppose Π satisfies (i) and (ii). (i) implies that $\tau(\delta, \Pi)$ is defined and is a train track on (S, P) . Now for any branch of the form $b(\pi)$ where $\pi \in \Pi(\delta) - \Pi$, take a curve c as in condition (ii); clearly c can be homotoped to a smooth closed curve γ in $\tau(\delta, \Pi)$ traversing $b(\pi)$. And for any branch of the form $b(h)$, where $h \in H(\delta)$ and some $\pi \in \Pi(\delta) - \Pi$ is adjacent to h , let c be a curve satisfying condition (ii) with respect to π ; clearly c intersects h , and so c is recurrent, and it follows from the previous proposition that Π is realizable. What we have done, in the language of train tracks, is to construct, for each branch b of $\tau(\delta, \Pi)$, an essential simple closed curve in (S, P) carried on $\tau(\delta, \Pi)$ with positive weight on b .

Now suppose Π is realizable. The previous proposition shows that $\tau(\delta, \Pi)$ is defined and is recurrent. Consider $\pi \in \Pi(\delta) - \Pi$. Choose $\mu \in W_+(\tau(\delta, \Pi))$ with integral weights and having the smallest total weight, so that $\mu(b(\pi)) \neq 0$. Using the homeomorphism $\chi: W_+(\tau(\delta, \Pi)) \rightarrow W^0(\delta, \Pi)$ from the previous proposition, let $\mu' = \chi(\mu) \in W^0(\delta, \Pi)$. Since μ' has integral weights, there exists $c \subset S - P$ which is a disjoint union of essential simple closed curves such that c transverses each $\pi' \in \Pi(\delta)$ exactly $\mu'(\pi')$ times. Homotop c to a union c' of smooth closed curves on $\tau(\delta, \Pi)$. Clearly μ is the set of transverse measures of c' . Moreover, using the minimality of μ , the argument of the Recurrence Lemma shows that c' consists of just one closed curve, which traverses each branch of τ at most once in each direction. It follows that c is a simple closed curve which traverses each prong of δ at most once in each direction, and crosses each arc of δ at most once in each direction. The remaining properties of (ii) are easily verified. Q.E.D.

Finally, if $\Pi \subset \Pi(\delta)$ is realizable, then $\dim(W_+(\tau(\delta, \Pi)))$ is computed as follows: say that Π is *orientable* if $\tau(\delta, \Pi)$ is orientable. Define $H(\delta, \Pi) = \{h \in H(\delta) \mid \text{any prong of } \delta \text{ adjacent to } h \text{ is in } \Pi\}$; clearly $|H(\delta) - H(\delta, \Pi)|$ is the number of branches of $\tau(\delta, \Pi)$ of the form $b(h)$. Also, $|\Pi(\delta) - \Pi|$ is the number of branches of $\tau(\delta, \Pi)$ of the form $b(\pi)$. Let $n_i(\delta, \Pi)$ be the number of triangles of δ having i prongs in Π , i.e. the number of case i triangles with respect to Π . It is easy to see that the number of switches of $\tau(\delta, \Pi)$ is $2 \cdot n_2(\delta, \Pi) + 3 \cdot [n_1(\delta, \Pi) + n_0(\delta, \Pi)]$ (note that our definition of train track allows a switch to have only one branch end on either side). Then from the formula for $\dim(W_+(\tau(\delta, \Pi)))$ and the above



Gluing Triangles To Make a Once-Punctured Torus

FIGURE 29


A 1-cell in \mathcal{PF} of a punctured torus

FIGURE 30

proposition, we have

$$(**) \quad \dim(\mathcal{C}^0(\delta, \Pi)) = |H(\delta) - H(\delta, \Pi)| + |\Pi(\delta) - \Pi| - 2 \cdot n_2(\delta, \Pi) \\ - 3 \cdot [n_1(\delta, \Pi) + n_0(\delta, \Pi)] - d$$

where $d = 0$ if Π is orientable, and $d = 1$ if Π is nonorientable.

Now we shall pause to give a complete description of $\mathcal{C}(\delta)$ in the case that (S, P) is a once-punctured torus, (T^2, p) . First note that any two ideal triangulations of (T^2, p) are “combinatorially equivalent”, which means that they are obtained from the same pattern of triangle gluings (this does not mean that they are all isotopic; far from it). In fact, Euler-characteristic considerations show that there are exactly three ideal arcs, two triangles, six arc ends, and six prongs. Moreover, in order that the two triangles be glued together to form an oriented surface with all vertices identified, the gluing pattern must be exactly as diagrammed in Figure 29, up to symmetries of the triangles.

In the diagram, the six prongs of $\Pi(\delta)$ come in three pairs, $\{\pi_i, \pi'_i\}$, $i \in \{1, 2, 3\}$. Each pair has the following property: the arcs to the left of π_i and π'_i are identical; similarly for the arcs to the right and the opposite arcs. Denote these arcs by h_i^L , h_i^R , h_i^{Op} . Then for any realizable $\Pi \subset \Pi(\delta)$, $\pi_i \in \Pi$ if and only if $\pi'_i \in \Pi$, for both of these conditions are equivalent to the condition that $\mu(h_i^L) + \mu(h_i^R) = \mu(h_i^{\text{Op}})$, for any $\mu \in W(\delta, \Pi)$. Thus, there are three possible realizable prong sets consisting of



 A 0-cell in \mathcal{PF} of a punctured torus

FIGURE 31

exactly one pair $\{\pi_i, \pi'_i\}$, and it is easily verified that each of these three is actually realizable; in fact, by symmetry, it suffices to verify this for only one of the three, say, $\{\pi_1, \pi'_1\}$. In Figure 30 we show how to visualize the resulting train track by first gluing together the two triangles along the side h_1^{opp} opposite π_1 and π'_1 ; and then gluing all three sides together and picturing the result on a punctured torus embedded in \mathbf{R}^3 . It is evident that the resulting train track τ is recurrent, and that $\dim(W_+(\tau)) = 2$, so the corresponding cell in \mathcal{PF} is of dimension 1.

Also, there are three possible realizable prong sets consisting of two pairs $\{\pi_i, \pi'_i, \pi_j, \pi'_j\}$, and each possibility does actually result in a realizable prong set; in fact, the corresponding train track is a circle, so the corresponding cell in \mathcal{PF} has dimension 0. We illustrate this in Figure 31 for the prong set $\{\pi_2, \pi'_2, \pi_3, \pi'_3\}$.

It is evident from the rules for gluing cells together, given in the proposition Decomposition Elements Are Cells, that these cells fit together into a circle, with three 1-cells and three 0-cells.

Now we turn to considerations of an ideal triangulation δ on a general punctured surface (S, P) . Our goal, while not to give a complete combinatorial description of $\mathcal{C}(\delta)$, is to at least describe the top-dimensional cells of $\mathcal{C}(\delta)$, and to use this description to prove that \mathcal{PF} is a manifold of dimension $6 \cdot g - 7 + 2 \cdot |P|$. We shall need the fact that for each triangulation δ of (S, P) , δ has $6 \cdot g - 6 + 3 \cdot |P|$ arcs, $4 \cdot g - 4 + 2 \cdot |P|$ triangles and $12 \cdot g - 12 + 6 \cdot |P|$ prongs, where $g = \text{genus}(S)$; this is easily verified by an Euler-characteristic calculation.

LEMMA: TOP-DIMENSIONAL CELLS OF $\mathcal{C}(\delta)$. *Suppose $\Pi \subset \Pi(\delta)$ is realizable, and has exactly one prong located at each puncture. Then $\dim(c(\delta, \Pi)) = 6 \cdot g + 2 \cdot |P| - 7$. Moreover, $c^0(\delta, \Pi)$ is an open subset of \mathcal{PF} .*

PROOF. Note first that $|\Pi| = |P|$; thus, $|\Pi(\delta) - \Pi| = 12 \cdot g - 12 + 5 \cdot |P|$.

Note next that every arc $h \in H(\delta)$ is adjacent to at least one prong of $\Pi(\delta) - \Pi$: for h has an end adjacent to two distinct prongs located at the same puncture, one of which is not in Π , by hypothesis. Thus, $|H(\delta) - H(\delta, \Pi)| = 6 \cdot g - 6 + 3 \cdot |P|$.

Also, every triangle T contains at most one prong of Π : for otherwise, T would have a side h such that both prongs on one side of h are in Π ; the arc condition on Π implies then that each prong adjacent to h is in Π , and as in the previous paragraph, it follows that two prongs located at the same puncture are in Π , a contradiction. Thus, $n_3(\delta, \Pi) = n_2(\delta, \Pi) = 0$, $n_1(\delta, \Pi) = |P|$, and $n_0(\delta, \Pi) = 4 \cdot g - 4 + |P|$.

Finally, we show that Π is nonorientable. Given $p \in P$, since $|\Pi \cap \Pi(\delta, p)| = 1$, then the complementary component C of $\tau(\delta, \Pi)$ containing p is a once punctured disc with one cusp on the boundary. Since there are an odd number of cusps on ∂C , there is no way to assign compatible orientations to the branches composing ∂C , and so $\tau(\delta, \Pi)$ is nonorientable.

So by formula (**)

$$\begin{aligned} \dim(\varepsilon^0(\delta, \Pi)) &= \{6 \cdot g - 6 + 3 \cdot |P|\} + \{12 \cdot g - 12 + 5 \cdot |P|\} \\ &\quad - 3 \cdot \{|P| + 4 \cdot g - 4 + |P|\} - 1 \\ &= 6 \cdot g - 7 + 2 \cdot |P|. \end{aligned}$$

It remains to show that $\varepsilon^0(\delta, \Pi)$ is open in \mathcal{PF} , when Π has exactly one prong located at each puncture (if Π is not realizable, then $\varepsilon^0(\delta, \Pi) = \emptyset$, so it is certainly open). Note that for any $\pi \in \Pi(\delta)$,

$$\{p(\mathcal{F}) \mid \pi \in \Pi(\delta, \mathcal{F})\} = \{p(L(\mu)) \mid \mu \in W(\delta) \text{ satisfies } \pi \text{ equality}\}$$

is a closed subset of \mathcal{PF} . So if Π is an arbitrary subset of $\Pi(\delta)$, then the set $A = \bigcup_{\pi \in \Pi(\delta) - \Pi} \{p(\mathcal{F}) \mid \pi \in \Pi(\delta, \mathcal{F})\}$ is closed, being a finite union of closed sets; the complement $A^c = \mathcal{PF} - A$ of this set is therefore open. When Π has exactly one prong at each puncture, we shall identify A^c with $\varepsilon^0(\delta, \Pi)$, which is therefore open in \mathcal{PF} . Note that

$$\begin{aligned} A^c &= \bigcap_{\pi \in \Pi(\delta) - \Pi} \{p(\mathcal{F}) \mid \pi \notin \Pi(\delta, \mathcal{F})\} \\ &= \{p(\mathcal{F}) \mid \text{for each } \pi \in \Pi(\delta) - \Pi, \pi \notin \Pi(\delta, \mathcal{F})\}. \end{aligned}$$

Note, however, that each $\Pi(\delta, \mathcal{F})$ must contain at least one prong in $\Pi(\delta, p)$, for each $p \in P$; but if $p(\mathcal{F}) \in A^c$, then all prongs in $\Pi(\delta, p)$ but one are excluded, and that one is the unique prong of $\Pi(\delta, p) \cap \Pi$. Thus, $A^c = \{p(\mathcal{F}) \mid \text{for each } \pi \in \Pi(\delta) - \Pi, \pi \notin \Pi(\delta, \mathcal{F})\} = \varepsilon^0(\delta, \Pi)$. Q.E.D.

It is not hard to use (**) to show that the *only* way that $\dim(\varepsilon^0(\delta, \Pi))$ can be $6 \cdot g - 7 + 2 \cdot |P|$ is in one of the two following situations:

- (i) Π is realizable and $|\Pi \cap \Pi(\delta, p)| = 1$ for each $p \in P$; or,
- (ii) Π is realizable and orientable, $|\Pi \cap \Pi(\delta, p)| = 2$ for some $p \in P$, and $|\Pi \cap \Pi(\delta, p')| = 1$ for all $p' \neq p \in P$.

In other words, although adding additional prongs might decrease the number of switches in the associated train track, it would in addition decrease the number of branches by at least the same amount, resulting in a net decrease in dimension. As

we have seen, case (ii) is in fact what happens on a once punctured torus: the top-dimensional cells have two prongs at the puncture (in fact, for the once-punctured torus, the hypotheses of the above lemma are vacuous). However, if (S, P) is not a once-punctured torus, no $\Pi \subset \Pi(\delta)$ can satisfy (ii) above. To see why, suppose that Π satisfies everything except possibly the orientability requirement. If $|P| > 2$, then there will exist some $p' \in P$ for which $|\Pi \cap \Pi(\delta, p')| = 1$, and $\tau(\delta, \Pi)$ will have a complementary region with one boundary cusp, containing p . And if $|P| = 1$, then $g \geq 2$, and there will be at least one triangle T of δ such that $\Pi(T) \cap P = \emptyset$; then $\tau(\delta, \Pi)$ will contain a complementary region, located in T , which has three boundary cusps. In either case, $\tau(\delta, \Pi)$ is nonorientable. Thus, if (S, P) is not a once-punctured torus, and δ is an ideal triangulation of (S, P) , then any realizable $\Pi \subset \Pi(\delta)$ which does not satisfy the hypotheses of Top-Dimensional Cells of $\mathcal{C}(\delta)$ represents a cell of dimension strictly less than $6 \cdot g - 7 + 2 \cdot |P|$.

As the title of the above lemma suggests, it is, in fact, true that every top-dimensional cell of $\mathcal{C}(\delta)$ (i.e. one which is not a face of a higher-dimensional cell) has dimension $6 \cdot g - 7 + 2 \cdot |P|$, though to show this directly would require verifying realizability for lots of subsets $\Pi \subset \Pi(\delta)$ having exactly one prong at each $p \in P$, a difficult task to perform in general. However, we can actually prove a lot more:


THEOREM. \mathcal{PF} is a compact manifold of dimension $6 \cdot g + 2 \cdot |P| - 7$.

PROOF. The proof given here is by no means as slick as the “classical” proof using arbitrary train tracks. In the classical proof, one “pinches” complementary components of a nonmaximal train track so as to make it maximal (see [T] or [HP]). The same idea underlies the proof we give, but the category of ideal triangulations and elementary moves is somewhat more “rigid” than the category of train tracks and splittings, and we have to do some rather complicated combinatorial analysis. This combinatorial analysis is not without interest in its own right, however: there is an interesting counting argument which is closely analogous to the counting argument used to prove the Connectivity Property for Elementary Moves.

We shall assume that (S, P) is not a once-punctured torus. To prove the theorem, for each $p(\mathcal{F}) \in \mathcal{PF}$, we shall show that there exists an ideal triangulation δ such that $\Pi(\delta, \mathcal{F})$ has exactly one prong at each puncture. It follows from the lemma Top-Dimensional Cells of $\mathcal{C}(\delta)$ that $p(\mathcal{F})$ is contained in the open set $\epsilon^0(\delta, \Pi(\delta, \mathcal{F}))$ of \mathcal{PF} , which has the correct dimension, proving the theorem.

The idea is as follows: start from an arbitrary ideal triangulation δ . In general, $\Pi(\delta, \mathcal{F})$ will have more than one prong at each puncture; what is worse, there may even be case 2 or (horrors!) case 3 triangles of δ , with respect to \mathcal{F} . To get rid of these, we will perform elementary moves on δ , resulting in new ideal triangulations which have fewer and fewer bad triangles, and fewer and fewer prongs in the prong set of \mathcal{F} . In order to find the correct sequence of elementary moves, we will have to make a precise combinatorial analysis of how the prong set can change under an elementary move.

To begin with, note the following fact, which is a trivial consequence of the Coordinate Change Theorem for Elementary Moves.



Removing Case 3 Triangles

FIGURE 32

LEMMA. Let $\delta \rightarrow \delta'$ be an elementary move supported on a quadrilateral Q of δ . Then for each $\pi \in \Pi(\delta) \cap \Pi(\delta') = \Pi(\delta) - \Pi(Q)$ and each $p(\mathcal{F}) \in \mathcal{PF}$, $\pi \in \Pi(\delta, \mathcal{F})$ if and only if $\pi \in \Pi(\delta', \mathcal{F})$.

We shall use this fact repeatedly without mention in what follows.

Now we show how to get rid of all case 3 triangles:


LEMMA: REMOVING CASE 3 TRIANGLES. Suppose $\delta \rightarrow \delta'$ is an elementary move performed on $h \in H(\delta)$, supported on the quadrilateral Q of δ . Let T_W, T_E be the triangles of δ in Q , and T_N, T_S be the triangles of δ' in Q , as in the proof of the Coordinate Change Theorem for Elementary Moves. Given $p(\mathcal{F}) \in \mathcal{PF}$, suppose T_E is a case 2 triangle and T_W is a case 3 triangle with respect to \mathcal{F} . Then T_N and T_S are case 2 triangles with respect to \mathcal{F} .

The proof is an elementary consequence of the Coordinate Change Theorem for Elementary Moves. The reader may fill in the details, consulting Figure 32. In the diagram, we also show how the train tracks $\tau(\delta, \Pi(\delta, \mathcal{F}))$ and $\tau(\delta', \Pi(\delta', \mathcal{F}))$ are related; in fact, they are isotopic.

As a consequence of this lemma: if δ has any case 3 triangles with respect to \mathcal{F} , then by connectivity there must be at least one such triangle T adjacent to a triangle T' that is not case 3, and the arc condition on $\Pi(\delta, \mathcal{F})$ guarantees that T' is a case 2 triangle; if $\delta \rightarrow \delta'$ is performed on the common arc of T and T' , the above lemma guarantees that δ' has one fewer case 3 triangle than δ . Iterating this process, we see that for any $\mathcal{F} \in \mathcal{PF}$, there exists a δ having no case 3 triangles with respect to \mathcal{F} .

For our next trick:

LEMMA: REMOVING CASE 2 TRIANGLES. Assume the same notation as in the previous lemma. Given $p(\mathcal{F}) \in \mathcal{PF}$, if T_W and T_E are case 2 triangles with respect to \mathcal{F} , and if the prongs of $\Pi(\delta, \mathcal{F}) \cap \Pi(T_W)$ and of $\Pi(\delta, \mathcal{F}) \cap \Pi(T_E)$ are all adjacent to h , then T_N, T_S are case 1 triangles.




Removing Case 2 triangles

FIGURE 33

Again we only indicate the proof with Figure 33. Note that in this case, $\tau(\delta', \Pi(\delta', \mathcal{F}))$ is obtained by “pinching” $\tau(\delta, \Pi(\delta, \mathcal{F}))$; classically, these “pinching” operations were what Thurston used to prove that \mathcal{PF} was a manifold. In most cases, this operation will increase the dimension of the corresponding cell in \mathcal{PF} ; the only exception is when $\tau(\delta, \Pi(\delta, \mathcal{F}))$ is orientable and $\tau(\delta', \Pi(\delta', \mathcal{F}))$ is nonorientable.


To apply this lemma: suppose δ has no case 3 triangles with respect to \mathcal{F} , but there is a case 2 triangle T_0 with respect to δ . Let h be the unique arc adjacent to T_0 such that both prongs of $\Pi(\delta, \mathcal{F}) \cap \Pi(T_0)$ are adjacent to h . By the arc condition, all prongs adjacent to h are in $\Pi(\delta, \mathcal{F})$. This shows that h is removable, bounding another triangle T_1 of δ , and that the two prongs of $\Pi(T_1)$ adjacent to h are in $\Pi(\delta, \mathcal{F})$. By hypothesis, T_1 cannot be case 3, so it is case 2. The above lemma now applies to the elementary move $\delta \rightarrow \delta'$ performed on h , and it follows that δ' has two fewer case 2 triangles than δ (note that in general, if there are no case 3 triangle, there are an even number of case 2 triangles). Iterating this process, we arrive at an ideal triangulation having only case 0 and 1 triangles with respect to \mathcal{F} .

To reduce the number of case 1 triangles, we use another combinatorial fact about elementary moves, similar in form to those above. To state it, we set up some notation. Let $\delta \rightarrow \delta'$ be an elementary move performed on $h \in H(\delta)$ supported on Q , and let T_W, T_E be the triangles of δ in Q as above, and T_N, T_S the triangles of δ' in Q as above. We need some notation for the prongs of δ and of δ' in Q . This notation follows the “compass points” conventions established in the Coordinate Change Theorem; it is given in Figure 34. In general, if A denotes a compass point, then π_A is the unique prong of δ in Q located at v_A (when this makes sense); and if A, B denote nonantipodal, distinct compass points, then π_{AB} is the unique prong located at v_A which is in $\Pi(T_B)$ (again, when this makes sense). Similar conventions hold for the prongs π'_A and π'_{AB} of δ' .



Labellings of prongs in an elementary move

FIGURE 34



Removing Case 1 triangles


FIGURE 35

LEMMA: REMOVING CASE 1 TRIANGLES. *Given $p(\mathcal{F}) \in \mathcal{PF}$, suppose T_W and T_E are both case 1 triangles with respect to \mathcal{F} , and assume that $\pi_{NW}, \pi_{NE} \in \Pi(\delta, \mathcal{F})$. Then T_N is a case 1 triangle with $\pi'_N \in \Pi(\delta', \mathcal{F})$, and T_S is a case 0-triangle.*

The proof is indicated in Figure 35. Note once again that $\pi(\delta', \Pi(\delta', \mathcal{F}))$ is obtained by pinching $\tau(\delta, \Pi(\delta, \mathcal{F}))$.

Unfortunately, it is not at all clear how to apply this to remove all but one case 1 triangle for each puncture, for it may well happen that $\Pi(\delta, \mathcal{F}) \cap \Pi(\delta, p)$ consists of several, nonsuccessive prongs. Thus, we must do an intermediate sequence of elementary moves to force nonsuccessive prongs closer together. This will be accomplished with the aid of a further result on elementary moves, paralleling the results above. Again, we assume that $\delta \rightarrow \delta'$ is an elementary move, with all triangles and prongs involved in the action labelled as above.

LEMMA: ALTERING CASE 1 TRIANGLES. *Given $p(\mathcal{F}) \in \mathcal{PF}$, suppose that T_W is a case 1 triangle with respect to \mathcal{F} , with $\pi_{NW} \in \Pi(T_W, \mathcal{F})$; suppose also that T_E is*



Altering Case 1 Triangles

FIGURE 36

either case 0 or 1 with respect to \mathcal{F} , but that $\pi_{NE} \notin \Pi(T_E, \mathcal{F})$. Then the following properties hold:

- (i) T_N is a case 1 triangle, with $\pi'_N \in \Pi(T_N; \mathcal{F})$;
- (ii) T_S is case 0 or 1, according to whether T_E is case 0 or 1 (resp.);
- (iii) $\pi'_{ES} \in \Pi(T_S, \mathcal{F})$ if and only if $\pi_E \in \Pi(T_E, \mathcal{F})$;
- (iv) $\pi'_S \in \Pi(T_S, \mathcal{F})$ if and only if $\pi_{SE} \in \Pi(T_E, \mathcal{F})$.

Again, this is a direct consequence of the Coordinate Change Theorem for Elementary Moves. We indicate in Figure 36 the three possible cases that can arise: T_E is case 0; T_E is case 1 with $\pi_E \in \Pi(T_E, \mathcal{F})$; and T_E is case 1 with $\pi_{SE} \in \Pi(T_E, \mathcal{F})$.

It is still not apparent how to apply this to reduce the number of case 1 triangles. We need a dynamic picture of how prongs can “move” under a sequence of elementary moves.

First we introduce some new terminology. Let $\delta \rightarrow \delta'$ be an elementary move performed on an arc $h \in H(\delta)$, let $e \in E(\delta, p)$ be an end of h , let e_- be the predecessor of e in $\Pi(\delta, p)$, and let e_+ be the successor of e in $E(\delta, p)$. Note that $\pi_- = (e_-, e)$ and $\pi_+ = (e, e_+)$ are both in the set $\Pi(\delta, p)$, and $\pi' = (e_-, e_+)$ is in the set $\Pi(\delta', p)$. We shall say that π_- and π_+ coalesce into π' under the elementary

move $\delta \rightarrow \delta'$, and that π' is a *coalescence* of π_- and of π_+ . We shall also say that π' *splits* into π_- and π_+ under the elementary move $\delta' \rightarrow \delta$.

Utilizing the same notation as above, note that $\pi_{ES} \in \Pi(\delta')$ is uniquely characterized by the property that it is not in $\Pi(T_N)$, and that it coalesces into $\pi_E \in \Pi(\delta) - \Pi(T_W)$ under the inverse elementary move $\delta' \rightarrow \delta$. Similarly, $\pi_S \in \Pi(\delta')$ is uniquely characterized by the property that it is not in $\Pi(T_N)$, and that $\pi_{SE} \in \Pi(\delta) - \Pi(T_N)$ coalesces into π_S under the elementary move $\delta \rightarrow \delta'$. These remarks motivate the following definitions:

Given an ideal triangulation δ and $\pi^* \in \Pi(\delta)$, let T^* be the triangle of δ such that $\pi^* \in \Pi(T^*)$, and let $\Pi^*(\delta, \pi^*) = [\Pi(\delta) - \Pi(T^*)] \cup \{\pi^*\}$; thus, we throw away the other two prongs of T^* to obtain $\Pi^*(\delta, \pi^*)$ from $\Pi(\delta)$. Now suppose $\delta \rightarrow \delta'$ is an elementary move performed on an arc $h \in H(\delta)$ incident to π^* ; let Q be the support of $\delta \rightarrow \delta'$, let $\pi'^* \in \Pi(\delta')$ be the coalescence of π^* under $\delta \rightarrow \delta'$, and let T'^* be the triangle of δ' such that $\pi'^* \in \Pi(T'^*)$. We define a map $pf: \Pi^*(\delta, \pi^*) \rightarrow \Pi^*(\delta', \pi'^*)$, called the *prong flow*, as follows: given $\pi \in \Pi^*(\delta, \pi^*)$, $pf(\pi) \in \Pi^*(\delta', \pi'^*)$ is uniquely defined by one of the following cases:

- (i) if $\pi \in \Pi(T)$ where $\text{int}(T) \cap Q = \emptyset$, then $\pi = pf(\pi)$;
- (ii) $pf(\pi^*) = \pi'^*$;
- (iii) if $\pi \in \Pi(T)$ where $T \neq T^*$ and $T \subset Q$, then:
 - (a) if π and π^* are incident to the same end of h , then $pf(\pi) = pf(\pi^*) = \pi'^*$;
 - (b) otherwise, $pf(\pi) \in \Pi(T')$ where $T' \neq T'^*$ and $T' \subset Q$, and either: $pf(\pi)$ is a coalescence to π under $\delta \rightarrow \delta'$; or π is a coalescence of $pf(\pi)$ under $\delta' \rightarrow \delta$.

In case (iii)(b), notice that $pf(\pi)$ is a coalescence of π if and only if π and π^* are incident to opposite ends of h ($\pi^* = \pi_{NW}$ and $\pi = \pi_{SE}$ in Figure 36); and π is a coalescence of $pf(\pi)$ if and only if π corresponds to a corner opposite the arc h ($\pi^* = \pi_{NW}$ and $\pi = \pi_E$ in Figure 36). Thus, pf is 1-1, with the single exception that it maps two prongs to π'^* (i.e. $pf(\pi_{NW}) = pf(\pi_{NE}) = \pi'_N$ in Figure 36). Also, only one prong of $\Pi^*(\delta', \pi'^*)$ is left out of the image of pf (corresponding to π'_{WS} in Figure 36).

The critical property of the prong flow is expressed in the following lemma, which is a direct consequence of the two lemmas Removing Case 1 Triangles and Altering Case 1 Triangles.

LEMMA: NATURALITY OF THE PRONG FLOW. *Let δ be an ideal triangulation, $\pi^* \in \Pi(\delta)$, T^* the triangle of δ such that $\pi^* \in \Pi(T^*)$. Given $p(\mathcal{F}) \in \mathcal{PF}$, suppose that $\Pi(\delta, \mathcal{F}) \cap \Pi(T^*) = \{\pi^*\}$. Let $\delta \rightarrow \delta'$ be an elementary move performed on an arc adjacent to π^* . Then for each $\pi \in \Pi^*(\delta, \pi^*)$: if $\pi \in \Pi(\delta, \mathcal{F})$, then $pf(\pi) \in \Pi(\delta', \mathcal{F})$; and when $pf(\pi) \neq pf(\pi^*)$, if $pf(\pi) \in \Pi(\delta', \mathcal{F})$, then $\pi \in \Pi(\delta, \mathcal{F})$.*

This says that the prong flow is natural with respect to prong sets of measured foliations. This lemma is true as well if the hypothesis is weakened to say merely that $\pi^* \in \Pi(\delta, \mathcal{F})$, as can easily be checked by a closer examination of the proof of the two lemmas Removing Case 3 Triangles and Removing Case 2 Triangles.

We can now state a proposition which, together with Naturality of the Prong Flow, implies that \mathcal{PF} is a manifold.

PROPOSITION: COALESCING NONSUCCESSIVE PRONGS. *Let δ be an ideal triangulation, $\pi^* \in \Pi(\delta, p)$, T^* the triangle of δ such that $\pi^* \in \Pi(T^*)$. Then for any $\pi \in \Pi^*(\delta, \pi^*) \cap \Pi(\delta, p)$, there exists a sequence of elementary moves $\delta = \delta_0 \rightarrow \delta_1 \rightarrow \dots \rightarrow \delta_n$ such that $\delta_i \rightarrow \delta_{i+1}$ is performed on an arc adjacent to $pf^i(\pi^*)$, and $pf^n(\pi^*) = pf^n(\pi)$, in the set $\Pi(\delta_n, p)$.*

Before proving this proposition, let us use it to prove that \mathcal{PF} is a manifold. Given $p(\mathcal{F}) \in \mathcal{PF}$, we can choose an ideal triangulation δ such that δ has only case 0 and 1 triangles with respect to \mathcal{F} . Choose $p \in P$, and if $|\Pi(\delta, p) \cap \Pi(\delta, \mathcal{F})| > 1$, choose $\pi^*, \pi \in \Pi(\delta, p) \cap \Pi(\delta, \mathcal{F})$ with $\pi^* \neq \pi$. Now apply the above proposition, and let $\delta = \delta_0 \rightarrow \dots \rightarrow \delta_n$ be the resulting sequence of elementary moves. From Naturality of the Prong Flow it follows that $|\Pi(\delta_n, p) \cap \Pi(\delta_n, \mathcal{F})| < |\Pi(\delta_0, p) \cap \Pi(\delta_0, \mathcal{F})|$, and that $|\Pi(\delta_n, p') \cap \Pi(\delta_n, \mathcal{F})| = |\Pi(\delta_0, p') \cap \Pi(\delta_0, \mathcal{F})|$ for any $p' \neq p \in P$. Iterating this process, eventually we arrive at an ideal triangulation δ' such that $|\Pi(\delta', \mathcal{F}) \cap \Pi(\delta', p)| = 1$ for any $p \in P$, and we are done.

Now we turn to the proof of Coalescing Nonsuccessive Prongs. Suppose that $\pi^*, \pi \in \Pi(\delta, p)$, $\pi^* \neq \pi$. We define the *positive circular interval* between π^* and π to be

$$I_+(\pi^*, \pi) = \{ \pi' \in \Pi(\delta, p) \cap \Pi^*(\delta, \pi^*) \mid (\pi^*, \pi', \pi) \text{ is positively oriented in } \Pi(\delta, p) \},$$

and we define the *positive circular distance* from π^* to π to be $d_+(\pi^*, \pi) = 1 + |I_+(\pi^*, \pi)|$. Similarly, we define the *negative circular interval* between π^* and π to be

$$I_-(\pi^*, \pi) = \{ \pi' \in \Pi(\delta, p) \cap \Pi^*(\delta, \pi^*) \mid (\pi^*, \pi', \pi) \text{ is negatively oriented in } \Pi(\delta, p) \},$$

and we define the *negative circular distance* from π^* to π to be $d_-(\pi^*, \pi) = 1 + |I_-(\pi^*, \pi)|$. Thus, $d_+(\pi^*, \pi)$ is computed by going around $\Pi(\delta, p)$ in the positive direction from π^* to π , skipping over any prong of T^* ; similarly for $d_-(\pi^*, \pi)$. The reason for skipping over the prongs of T^* in this peculiar manner is made clear in the counting arguments that follow.

Note that for $\pi^* \neq \pi \in \Pi(\delta, p)$, $d_+(\pi^*, \pi) = 1$ if and only if π is the successor of π^* in $\Pi(\delta, p)$. For the only other possibility is that the successor of π^* is a prong in T^* , implying that T^* is a puncture piece with interior puncture p ; but that means that $|\Pi(\delta, p)| = 1$, contradicting the hypothesis that $\pi^* \neq \pi \in \Pi(\delta, p)$. Thus, if we can perform a sequence of elementary moves so that, under the iterated prong flow, two prongs located at p come within circular distance 1, then we are in a position to coalesce the prongs with a single additional elementary move. For this purpose, we now prove:

LEMMA: REDUCING CIRCULAR DISTANCES. *Given an ideal triangulation δ and $\pi^* \in \Pi(\delta, p) \cap \Pi(T^*)$, let $h^{*L}, h^{*R} \in H(\delta)$ be the arcs to the left and right of π^* (resp.), let π^{*L} be the prong in $\Pi(T^*)$ adjacent to h^{*L} , and let π^{*R} be the prong in*




FIGURE 37

$\Pi(T_*)$ adjacent to h^{*R} . Let $\pi \in \Pi(\delta, p) \cap \Pi^*(\delta, \pi^*) - \{\pi^*\}$ be given. If $\delta \rightarrow \delta'$ is performed on h^{*R} , then:

(i) $d_+(\pi'^*, pf(\pi)) = d_+(\pi^*, \pi) - 1$ if and only if (π^*, π^{*L}, π) is not positively oriented in $\Pi(\delta, p)$; and

(ii) $d_+(\pi'^*, pf(\pi)) = d_+(\pi^*, \pi)$ if and only if (π^*, π^{*L}, π) is positively oriented in $\Pi(\delta, p)$.

Also, if $\delta \rightarrow \delta^*$ is performed on h^{*L} , then a similar conclusion holds replacing each “+” with “-”, replacing π^{*L} with π^{*R} and replacing “positively oriented” with “negatively oriented”.

PROOF. We shall prove only the case where $\delta \rightarrow \delta'$ is performed on h^{*R} . Figure 37 accompanies the proof. Thus, we are identifying π^* with π_{NW} and π^* with π'_N , in the earlier notation.

The proof goes as follows. First observe that the prong flow preserves circular ordering, as long as it does not coalesce any two prongs in the ordered triple under consideration.

Thus, $\pi_E \in I_+(\pi^*, \pi)$ if and only if $pf(\pi_E) = \pi'_{ES} \in I_+(\pi'^*, pf(\pi))$ (note that $pf(\pi^*) = \pi'^*$). Similarly, $\pi_{SE} \in I_+(\pi^*, \pi)$ if and only if $pf(\pi_{SE}) = \pi'_S \in I_+(\pi'^*, pf(\pi))$. Also, for any prong $\bar{\pi} \in \Pi(\delta) - \Pi(T_W) - \Pi(T_E)$, $\bar{\pi} \in I_+(\pi^*, \pi)$ if and only if $pf(\bar{\pi}) = \bar{\pi} \in I_+(\pi'^*, pf(\pi))$.

This accounts for all prongs in $I_+(\pi^*, \pi)$ with the exception of π_{NE} . This also accounts for all prongs in $I_+(\pi'^*, pf(\pi))$, with one possible exception: if (π^*, π^{*L}, π) is positively ordered in $\Pi(\delta, p)$, then $\pi'_{WS} \in I_+(\pi'^*, pf(\pi))$, showing that $d_+(\pi'^*, pf(\pi)) = d_+(\pi^*, \pi)$. If (π^*, π^{*L}, π) is not positively ordered in $\Pi(\delta, p)$, then there is no exception, because $\pi'_{WS} \notin I_+(\pi'^*, pf(\pi))$; thus, $d_+(\pi'^*, pf(\pi)) = d_+(\pi^*, \pi) - 1$. Q.E.D.

Thus, in order to reduce circular distances with a single elementary move, we wish to be in a situation where the triple of prongs (π^*, π^{*L}, π) is not positively ordered in $\Pi(\delta, p)$. There are several ways that this condition can occur. We shall consider them case by case.

Case 0. If $\pi^{*L} \notin \Pi(\delta, p)$, then it is certainly true that (π^*, π^{*L}, π) is not positively ordered in $\Pi(\delta, p)$. Moreover, after performing an elementary move $\delta \rightarrow \delta'$ on h^{*R} , it remains true that $pf(\pi^*)^L \notin \Pi(\delta', p)$, since π^{*L} is a coalescence of $pf(\pi^*)^L$. Continuing inductively, we see that for any $\pi \in \Pi^*(\delta, \pi^*)$, setting

$n = d_+(\pi^*, \pi)$, there is a sequence of elementary moves $\delta = \delta_0 \rightarrow \delta_1 \rightarrow \cdots \rightarrow \delta_{n-1}$ such that $d_+(pf^i(\pi^*), pf^i(\pi)) = n - i$, for each $i = 1, \dots, n - 1$. In fact, as long as $|\Pi(\delta_i, p)| > 1$, we can perform an elementary move on the arc to the right of $pf^i(\pi^*)$, which reduces all positive circular distances in $\Pi(\delta_i, p)$; since there is an upper bound on the positive circular distances, this upper bound must decrease as i increases. This cannot continue forever, which shows that for some i , $|\Pi(\delta_i, p)| = 1$. The case where $\pi^{*R} \notin \Pi(\delta, p)$ is handled similarly. This proves Coalescing Non-successive Prongs in the case where the three prongs π^* , π^{*L} , π^{*R} are not all located at the same puncture.

Now suppose that π^* , π^{*L} , π^{*R} are all located at the same puncture $p \in P$. There are two cases to consider, depending on the ordering of the triple π^* , π^{*L} , π^{*R} .

Case 1. $(\pi^*, \pi^{*R}, \pi^{*L})$ is positively ordered in $\Pi(\delta, p)$. It follows from the properties of circularly ordered sets that for any $\pi \in \Pi^*(\delta, \pi^*) \cap \Pi(\delta, p)$, either (π^*, π, π^{*L}) is positively ordered, or (π^*, π, π^{*R}) is negatively ordered; we assume the former. Thus, (π^*, π^{*L}, π) is not positively ordered in $\Pi(\delta, p)$, and we can apply Reducing Circular Distances to find an elementary move $\delta \rightarrow \delta'$ which reduces the positive distance from π^* to π . Moreover, it is easy to check that $(pf(\pi^*), pf(\pi^*)^L, pf(\pi))$ is still not positively ordered in $\Pi(\delta', p)$, using the fact that π^{*L} is the coalescence of $pf(\pi^*)^L$ under the inverse elementary move $\delta' \rightarrow \delta$. Proceeding inductively, we can construct a sequence of elementary moves which eventually coalesces π^* and π , proving the proposition Coalescing Nonsuccessive Prongs in this case.

Case 2. $(\pi^*, \pi^{*R}, \pi^{*L})$ is negatively ordered in $\Pi(\delta, p)$. In this case, it happens that there are prongs $\pi \in \Pi(\delta, p) \cap \Pi^*(\delta, \pi^*)$ such that (π^*, π, π^{*L}) is negatively ordered, and (π^*, π, π^{*R}) is positively ordered. In fact, such “inaccessible” prongs are all contained in the positive circular interval from π^{*L} to π^{*R} , $I_+(\pi^{*L}, \pi^{*R})$.

It is useful here to have a diagram of the triangulation. In such a diagram, we represent $E(\delta, p)$ as points on a circle, with counterclockwise ordering. An arc on the circle between two successive points then represents a prong, and the circular ordering of such arcs corresponds to the circular ordering on $\Pi(\delta, p)$. Furthermore, we connect two points by a chord of the circle if the corresponding ends are the two ends of an ideal arc of δ (in the diagrams that follow, in order to be sufficiently general, we shall not show all possible chords). It is possible that an end located at p has its opposite end located at a different puncture, so not all points are connected up to other points by chords.

In Figure 38 we show the chords representing the arcs of δ adjacent to T^* , when $(\pi^*, \pi^{*L}, \pi^{*R})$ is positively ordered in $\Pi(\delta, p)$; for evident reasons, we say that T^* is a “twisted” triangle. A similar diagram shows that when $(\pi^*, \pi^{*L}, \pi^{*R})$ is negatively ordered, T^* can be thought of as an “untwisted” triangle.

Now we define certain circular intervals of prongs in $\Pi(\delta, p)$, when T^* is twisted. Let $I_+ = I_+(\pi^*, \pi^{*L})$, $I_- = I_-(\pi^*, \pi^{*R})$, $I_0 = I_+(\pi^{*L}, \pi^{*R})$. These are indicated in the diagram of a twisted triangle. Since $(\pi^*, \pi^{*L}, \pi^{*R})$ is positively ordered, it follows that $I_0 \neq \emptyset$, and that’s our whole problem.





FIGURE 38



The diagram of an ideal triangulation of a punctured torus

FIGURE 39

Now we use the fact that (S, P) is not a punctured torus: we claim that $I_+ \cup I_-$ consists of at least three prongs. Suppose otherwise: since π^* , π^{*L} are not adjacent and π^* , π^{*R} are not adjacent, it would follow that I_+ and I_- each consist of a single prong, say π_+ and π_- . Now π_+ , π^{*L} are successive prongs in $\Pi(\delta, p)$, and π^{*R} , π_- are successive prongs. Since π^{*L} , π^{*R} are in a single triangle, namely T^* , it follows that π_+ , π_- are in a single triangle, say T' . Where is the third prong π' of T' ? Since π^* , π_+ are successive prongs, it follows that π' , π^{*R} are successive prongs. And since π_- , π^* are successive, then π^{*L} , π' are successive. Thus, I_0 consists of the single prong π' . But then it follows that the three sides of T^* are identified in pairs with the three sides of T' , so $T \cup T' = S$ by connectivity, and a simple Euler-characteristic argument shows that S is a torus, whose unique puncture is p , again by connectivity. The complete diagram for this ideal triangulation of a punctured torus is given in Figure 39 with all prongs labelled as in this argument.

We shall now suppose, without loss of generality, that I_+ consists of at least two successive prongs. Let $\pi^* = (e_0, e_1)$, and $\pi^{*L} = (e_2, e_3)$. Let E_+ denote the positive circular interval in $E(\delta, p)$ from e_1 to e_2 , noninclusive; thus, $e \in E_+$ if and only if there are successive prongs in I_+ each adjacent to e . Since I_+ contains successive prongs, it follows that $E_+ \neq \emptyset$. Let E_- and E_0 be similarly defined.

Case 2a. Suppose there is an arc $h \in H(\delta)$ connecting two ends in E_+ . Figure 40 is the triangulation diagram for this situation.




FIGURE 40




FIGURE 41

From the proposition Reducing Circular Distances, it follows that there is a sequence of elementary moves $\delta = \delta_0 \rightarrow \cdots \rightarrow \delta_n$ such that $\delta_i \rightarrow \delta_{i+1}$ is performed on the arc to the right of $pf^i(\pi^*)$, h^{*L} is the arc to the left of $pf^n(\pi^*)$, and h is the arc to the right of $pf^n(\pi^*)$. Thus, the triangle containing $pf^n(\pi^*)$ is untwisted, and we are reduced to case 1.

Case 2b. Suppose there is an arc h connecting an end of E_+ to some end not in $E_+ \cup E_-$. If the opposite end of h is not located at p , then using the same argument as above, the sequence of elementary moves $\delta = \delta_0 \rightarrow \cdots \rightarrow \delta_n$ terminates in an ideal triangulation δ_n such that $pf^n(\pi^*)$ is adjacent to the arc h , which has its opposite end at a different puncture. Thus, $pf^n(\pi^*)^L \notin \Pi(\delta, p)$, and we are reduced to case 0.

If the opposite end of h is located in E_0 , then the sequence of elementary moves $\delta = \delta_0 \rightarrow \cdots \rightarrow \delta_n$ terminates in an ideal triangulation δ_n such that $pf^n(\pi^*)$ is still located in a twisted triangle, but the number of inaccessible prongs has been reduced. By induction, we are reduced to the following case:

Case 2c. Suppose that I_0 consists of a single prong, say $I_0 = \{\pi_0\}$. Let T_0 be the triangle with prong π_0 . Since π_0 is adjacent to h^{*L} and to h^{*R} , it follows that T_0 has its two other prongs in I_+ and I_- , both adjacent to π^* . Figure 41 is the triangulation diagram.

As we saw above, since (S, P) is not a punctured torus, there exists at least one other prong in $I_+ \cup I_-$, say, π . Let T be the triangle with $\pi \in \Pi(T)$. Since $T \neq T^*$ and $T \neq T_0$, then the three prongs of T are distributed among the three sets I_+ , I_- , and $\Pi(\delta) - \Pi(\delta, p)$. By the pigeonhole principle, one of the following must happen: T has a prong in $\Pi(\delta) - \Pi(\delta, p)$; T has two prongs in I_+ ; or T has two prongs in I_- . The first possibility reduces to case 0; the latter two reduce to case 2a. Q.E.D.

BIBLIOGRAPHY

- [EB] D. B. A. Epstein and B. H. Bowditch, *Triangulations associated with punctured surfaces*, Topology (to appear).
- [FLP] A. Fathi, et al, *Travaux de Thurston sur les surfaces*, Asterisque **66–67** (1979).
- [HZ] J. Harer and D. Zagier, *The Euler characteristic of the moduli space of curves*, Invent. Math. **85** (1986), 157–185.
- [Har] J. Harer, *The virtual cohomological dimension of the mapping class group of an oriented surface*, Invent. Math. **84** (1986), 157–176.
- [Hat] A. Hatcher, *On triangulations of surfaces*, preprint.
- [M1] L. Mosher, *Pseudo-Anosovs on punctured surfaces*, Dissertation, Princeton Univ., 1983.
- [M2] L. Mosher, *The classification of pseudo-Anosovs*, Low Dimensional Topology and Kleinian Groups, Ed., D. B. A. Epstein, London Math. Soc. Lecture Notes Ser. 112, Cambridge Univ. Press, 1985.
- [M3] L. Mosher, *Topological invariants of measured foliations*, preprint.
- [P] R. C. Penner, *The Teichmüller space of punctured surfaces*, preprint.
- [T] W. Thurston, *The geometry and topology of 3-manifolds*, Princeton Univ. Lecture Notes, Chapters 8–9, 1978.

DEPARTMENT OF MATHEMATICS, HARVARD UNIVERSITY, CAMBRIDGE, MASSACHUSETTS 02138

Current address: Department of Mathematics, The Graduate School and University Center of the City University of New York, 33 West 42nd Street, New York, New York 10036