Commentarii Mathematici Helvetici

Generic leaves

John Cantwell and Lawrence Conlon

Abstract. A remarkable theorem of E. Ghys asserts that, for any harmonic measure μ on a compact, foliated metric space, μ–almost every leaf has 0, 1, 2 or a Cantor set of ends. In this paper, analogous results are proven for topologically almost all (*i.e.,* residual families of) leaves. More precisely, if some leaf is totally recurrent, a residual family of leaves is totally recurrent with 1, 2 or a Cantor set of ends. A "local" version of this theorem asserts that, in general, topologically almost all leaves have 0, 1, 2 or a Cantor set of dense ends.

Mathematics Subject Classification (1991). 57R30.

Keywords. Foliated metric space, generic, residual, meager, endset, totally recurrent leaf.

1. Introduction

Let $(X, \mathcal{F})$ be a locally compact, separable, complete, foliated metric space. The leaves are manifolds of some fixed dimension p, but transversely the foliation is modeled on a locally compact, separable, metric space S. It follows that there is induced a metric along the leaves which is compatible with their manifold structure and in which each leaf is complete. Hereafter, the locally compact and separable properties of our spaces will be assumed without further mention.

If the transverse space S is $\mathbb{R}^q$, then $(X, \mathcal{F})$ is an honest foliated manifold of codimension q, but there are important examples in which S is totally disconnected (*e.g.*, exceptional minimal sets in compact foliated manifolds of codimension one, many essential laminations of 3–manifolds, *etc.*). We study the topology of the *generic leaf* of $\mathcal{F}$.

The term "generic leaf" can be defined from the measure theoretic point of view, where it means "μ–almost every leaf" relative to a suitable measure μ, or from the purely topological viewpoint, where it refers to a residual family of leaves.

In a remarkable paper [6], E. Ghys shows that, given an arbitrary harmonic probability measure μ for a compact, leafwise C^3, foliated space, μ–almost every leaf has 0, 1, 2, or a Cantor set of ends. If, in addition, the leaf dimension is $p = 2$, μ–almost every noncompact leaf either has genus 0 or the leaf is orientable and all ends are nonplanar or all ends are nonorientable. This restricts the pos-

Figure 1. Generic, noncompact, orientable, 2–dimensional leaf types

sible homeomorphism types of the generic, noncompact, orientable leaves to the six depicted in Figure 1. There are also three possible nonorientable types with crosscaps clustering at all ends.

All compact foliated spaces which are leafwise C^3 admit harmonic probability measures μ (the proof in [5], formulated for compact, foliated manifolds, readily generalizes to compact, foliated metric spaces). By an application of [5, Theorem 4], the union of the supports of the ergodic components of μ has full μ–measure, so restricting to the support of an ergodic component reduces Ghys's theorem to the case that μ is ergodic with $\operatorname{supp}\mu = X$. In this case, one sees that μ–almost every leaf has topologically the *same* endset.

Definition 1.1. An end e of a leaf L in a foliated space $(X, \mathcal{F})$ is dense if every neighborhood of e in L is dense in X. The leaf L is totally recurrent if it is noncompact and every end of L is dense.

If μ is an ergodic, harmonic probability measure with $\operatorname{supp}\mu = X$, an easy application of the Fundamental Proposition of Ghys [6, p. 402], using the separability of X, shows that either X is a single compact leaf or μ–almost every leaf is totally recurrent in X. This is not used in [6], but suggests the correct hypothesis for a theorem about the topologically generic leaf. We turn to this.

Since X is locally compact and Hausdorff, it is a Baire space. That is, a countable union of closed, nowhere dense subsets has empty interior. Recall that a subset $Y \subset X$ is *meager* if it is contained in the countable union of closed, nowhere dense subsets of X. This is the topological analogue of a set of measure

zero. The two notions, of course, have no logical relation, as is illustrated by the existence of Cantor sets $Y \subset \mathbb{R}$ of positive Lebesgue measure. The complement of a meager set, called a *residual* set, contains the countable intersection of open, dense subsets of X. In particular, a residual set is dense in X, any countable intersection of residual sets is residual, and relatively residual (respectively, meager) subsets of residual sets are residual (respectively, meager) as subsets of X.

In this paper, we take "generic" in the topological sense and prove an analog of Ghys's theorem for every complete, foliated metric space $(X, \mathcal{F})$ which contains a totally recurrent leaf. It is not assumed that X is compact, no differentiability is needed and the proof is surprisingly elementary. We will generally blur the distinction between a family of leaves and the $\mathcal{F}$–saturated set which is the union of the leaves in the family.

Theorem A. *Let $(X, \mathcal{F})$ be a complete, foliated metric space. If $\mathcal{F}$ has a totally recurrent leaf, then there is a residual family G of totally recurrent leaves without holonomy such that one of the following holds:*
(1) *every leaf in G has a Cantor set of ends*;
(2) *every leaf in G has exactly two ends*;
(3) *every leaf in G has exactly one end.*
If, in addition, the leaf dimension is $p = 2$, either all leaves L in G have genus 0, or all leaves L in G are orientable and have only nonplanar ends, or all leaves L in G have only nonorientable ends.

Whether or not there is a totally recurrent leaf, one can characterize the set $\mathcal{E}_d(L)$ of dense ends of the topologically generic leaf L. Remark that $\mathcal{E}_d(L)$ is a closed, hence compact subset of the endset $\mathcal{E}(L)$.

Theorem B. *If $(X, \mathcal{F})$ is a complete, foliated metric space, a residual family G has the property that its leaves have 0, 1, 2, or a Cantor set of dense ends, the cardinality of $\mathcal{E}_d(L)$ being constant as L varies over G. If the leaf dimension is 2, the dense ends of the leaves in G are either all planar, all orientable but nonplanar, or all nonorientable.*

These results are analogous to the theorem of H. Hopf [10], according to which a regular covering space of a compact, connected manifold has 0, 1, 2, or a Cantor set of ends. Hopf's theorem is proven by determining the endset of the covering group, while our proofs and those of Ghys analyze the endset of the generic holonomy orbit. In their details, however, the proofs in this paper and those in [6] differ substantially.

Finally, one can define more general foliated spaces in which the leaves belong to some suitable class of path connected metric spaces. For instance, the leaves might be connected simplicial complexes. With very little change, our proofs work for such foliated spaces.

The authors wish to thank Renato Feres for several helpful conversations.

2. A Heuristic proof

The main idea in the proof of Theorem A is quite simple. We sketch it here for the case in which $(X, \mathcal{F})$ is a compact, minimal foliated space. In this case, each end $e \in \mathcal{E}(L)$ of each leaf L has nonempty asymptote $A_e \subseteq X$. This is a compact, $\mathcal{F}$–saturated set, so minimality implies that $A_e = X$ and every leaf is totally recurrent.

It is an elementary but surprising theorem that the leaves without holonomy form a residual set G_0 [4, 8], so we restrict our attention to these leaves. Let L be such a leaf having three or more ends and fix a compact, connected submanifold $N \subset L$ with three boundary components, each interfacing a distinct unbounded component of $L \smallsetminus N$. Since L has trivial holonomy, the usual Reeb stability argument localizes to N, providing an open product neighborhood $V \times N \subset X$ such that $V \subset S$ is open and, for each $z \in V$, $N_z = \{z\} \times N$ is contained in a leaf L_z. Let $V_0 \subseteq V$ be the residual set of points z such that $L_z \subset G_0$. If there is a relatively open subset $U \subseteq V_0$ such that, for each $z \in U$, N_z separates L_z into three unbounded components, it follows that every leaf in G_0 has a Cantor set of ends. Indeed, let F be such a leaf, $e \in \mathcal{E}(F)$ an end, and let $W \subset F$ be a neighborhood of that end. By total recurrence and the fact that $F \subset G_0$, W passes through $U \times N$, hence picks up a copy of N which separates F into three unbounded components. It follows that the arbitrary neighborhood W of the arbitrary end e is the neighborhood of at least one other end $e' \in \mathcal{E}(F)$. This is illustrated in Figure 2. As a result, the leaf F has no isolated ends. Since $\mathcal{E}(F)$ is a compact, totally disconnected, separable metric space, it must be a Cantor set.


Figure 2. A Cantor set of ends

Assume, then, that some leaf in G_0 does not have a Cantor set of ends and let $Z_3 \subseteq V_0$ be the set of points z such that N_z splits L_z into three unbounded components. By the above, the relative interior of Z_3 in V_0 is empty and there is $z \in V_0 \smallsetminus Z_3$. Since $L_z \smallsetminus N_z$ has at most two unbounded components, either one of its components is compact or there is a curve in L_z connecting two components of ∂N_z and missing int N_z. Local Reeb stability propagates this to an open neighborhood of z in V, proving that Z_3 is relatively closed in V_0, hence meager. One now needs to show that only countably many neighborhoods of the form $V_i \times N_i$ need to be considered. In order to do this rigorously, we will pass to the "1–skeleton" of each leaf L_z, this being the graph of the orbit $\Gamma(z)$, where Γ is the holonomy pseudogroup. Waving our hands at this detail, we conclude to the following.

Claim 1. *Either every leaf in G_0 has a Cantor set of ends, or the set of leaves with three or more ends is meager.*

Continuing to assume that some leaf in G_0 does not have a Cantor set of ends, we neglect the meager set of leaves with more than two ends. Consider the remaining leaves L_z in G_0 and compact, connected submanifolds $N_z \subset L_z$ with two boundary components. Reasoning as above, replacing Z_3 with the subset $Z_2 \subseteq V_0$ of z such that N_z splits L_z into two unbounded components, we establish the following.

Claim 2. *If the set of leaves with more than two ends is meager, then the set of leaves with two ends is either residual or meager. If it is meager, then the set of leaves with one end is residual.*

Since we have been working in G_0, all of the residual sets we have found consist of (totally recurrent) leaves without holonomy.

Similarly, a crosscap on a 2–dimensional leaf without holonomy propagates to a set of crosscaps on all leaves which (by total recurrence) cluster at all ends of the leaves. If some leaf is orientable, no crosscaps are allowed on any leaf of G_0 and a handle on some leaf without holonomy propagates (by total recurrence) to handles clustering at all ends of all leaves. Finally, if some leaf is orientable with finite genus, no handles are allowed on the leaves without holonomy.

In the next two sections, we set up the combinatorial results needed to make Claim 1 and Claim 2 rigorous. These results are also needed to prove the fundamental proposition (Proposition 5.5) that either no leaf is totally recurrent or total recurrence is topologically generic, as well as the analogous fact (Proposition 5.5) that either no leaf has a dense end or a residual family of leaves have some dense ends. For completeness, we also give the proof that trivial holonomy is topologically generic (Proposition 3.3).

3. Regular covers and holonomy

Most of the material in this section will be familiar to foliators, but a review seems worthwhile in order to establish notation, terminology and conventions.

The foliated metric space $(X, \mathcal{F})$ admits a locally finite cover

$$\mathcal{U} = \{(U_\alpha, x_\alpha, y_\alpha)\}_{\alpha \in \mathfrak{A}}$$

by open, relatively compact $\mathbb{R}^p \times S$ charts which are compatible with $\mathcal{F}$, where (S, ρ) is a fixed, locally compact, separable metric space. Here, $x_\alpha : U_\alpha \to \mathbb{R}^p$ is a continuous map onto an open, relatively compact p–cell $E_\alpha^p \subset \mathbb{R}^p$ and $y_\alpha : U_\alpha \to S$ is a continuous map onto an open, relatively compact metric r_α–ball $T_\alpha \subseteq S$, $\forall \, \alpha \in \mathfrak{A}$. Moreover,

$$z \mapsto (x_\alpha(z), y_\alpha(z))$$

defines a homeomorphism of U_α onto the open subset $E_\alpha^p \times T_\alpha \subset \mathbb{R}^p \times S$. Compatibility with the foliation means that the level sets of y_α are open, relatively compact p–cells in leaves of $\mathcal{F}$, called the *plaques* of $\mathcal{F}$, $\forall \, \alpha \in \mathfrak{A}$. Since X is separable, this locally finite atlas is at most countably infinite.

The metric d on X restricts to a metric on each plaque compatible with its manifold structure. It is easy to produce a metric $d_\mathcal{F}$ along the leaves which agrees with d on each plaque. If L is a leaf and $x, y \in L$, consider all finite sequences $x = x_0, x_1, \ldots, x_n = y$ of points in L such that x_{k-1} and x_k lie in a common plaque, $1 \leq k \leq n$. One sets $d_\mathcal{F}(x, y)$ equal to the infimum of the numbers

$$d(x_0, x_1) + d(x_1, x_2) + \cdots + d(x_{n-1}, x_n)$$

taken over all such sequences. Remark that $d_\mathcal{F}(x, y) \geq d(x, y)$. In fact, as distinct plaques in $L \cap \mathcal{U}_\alpha$ get close in the metric d, they get arbitrarily far apart in the metric $d_\mathcal{F}$.

We can and do require that there be uniform finite, positive lower and upper bounds on the diameters of the plaques, as measured by $d_\mathcal{F}$. Finally, we require that the atlas $\mathcal{U}$ be *regular* in the sense that, on overlaps $U_\alpha \cap U_\beta \neq \emptyset$, the change of coordinates has the form

$$x_\alpha = x_\alpha(x_\beta, y_\beta)$$
$$y_\alpha = y_\alpha(y_\beta).$$

In particular, a plaque of U_α meets at most one plaque of U_β. The fact that a regular foliated atlas can be found uses local Lebesgue numbers to refine a preliminary choice of foliated atlas. This construction guarantees, in particular, that each chart $(U_\alpha, x_\alpha, y_\alpha)$ has compact closure $\overline{U}_\alpha \subset U_\alpha'$, where $(U_\alpha', x_\alpha', y_\alpha')$ is a chart compatible with $\mathcal{F}$, $x_\alpha = x_\alpha'|U_\alpha$ and $y_\alpha = y_\alpha'|U_\alpha$. Thus, $\overline{T}_\alpha$ is a compact metric ball in (S, ρ) sitting in an open, relatively compact metric ball T_α' with

the same center as T_α and radius $r'_\alpha > r_\alpha$. One can choose $r'_\alpha = r_\alpha + \varepsilon$ for arbitrary values of $\varepsilon > 0$ sufficiently small. Of course, since S is not required to be connected, it is possible that $T_\alpha = \overline{T}_\alpha = T'_\alpha$. At any rate, we include the existence of the extensions $(U'_\alpha, x'_\alpha, y'_\alpha)$ as part of the meaning of "regular foliated atlas".

We let T denote the disjoint union of the open subsets $T_\alpha = y_\alpha(U_\alpha)$, $\alpha \in \mathfrak{A}$. This transverse space has two useful interpretations. Since y_α sets up a one–to–one correspondence between the plaques of U_α and the points of T_α, we can view T as the set of plaques of the regular cover $\mathcal{U}$. Also, the local product structure of U_α allows us to imbed $T_\alpha \hookrightarrow U_\alpha$ as a cross section of $\mathcal{F}|U_\alpha$, $\forall \alpha \in \mathfrak{A}$. The local finiteness of $\mathcal{U}$ makes it possible to guarantee that these imbeddings have disjoint images and define an imbedding $T \hookrightarrow X$ of T as a transverse cross section to $\mathcal{F}$. We will use both of these interpretations without further comment.

Note that the change of coordinate formula $y_\alpha = y_\alpha(y_\beta)$ yields a homeomorphism

$$g_{\alpha\beta} : y_\beta(U_\alpha \cap U_\beta) \to y_\alpha(U_\alpha \cap U_\beta)$$

between open subsets of T such that $g_{\alpha\beta} \circ y_\beta = y_\alpha$ on $U_\alpha \cap U_\beta$. The system $\gamma = \{g_{\alpha\beta}\}_{\alpha,\beta \in \mathfrak{A}}$ satisfies the *cocycle* conditions

$$g_{\alpha\alpha} = \mathrm{id}\,|T_\alpha, \quad \forall\, \alpha \in \mathfrak{A}$$
$$g_{\alpha\beta} = g_{\beta\alpha}^{-1}, \quad \forall\, \alpha, \beta \in \mathfrak{A}$$
$$g_{\alpha\beta} = g_{\alpha\lambda} \circ g_{\lambda\beta}, \quad \text{on } y_\beta(U_\alpha \cap U_\lambda \cap U_\beta), \forall\, \alpha, \lambda, \beta \in \mathfrak{A}.$$

We call γ the *holonomy cocycle* and the pseudogroup Γ of local homeomorphisms on T generated by γ is called the *holonomy pseudogroup*.

Points $x, y \in X$ lie in the same leaf of $\mathcal{F}$ if and only if there is a finite chain of plaques $P_0, P_1, \ldots, P_N$ such that $x \in P_0$, $y \in P_N$, and $P_{k-1} \cap P_k \neq \emptyset$, $1 \leq k \leq N$. Under the interpretation of T as the space of plaques, $g_{\alpha\beta}(P) = Q$ means that P is a plaque in U_β, Q is a plaque in U_α, and $P \cap Q \neq \emptyset$. Thus, the plaque chain connecting x to y in a leaf corresponds to a pure composition

$$g = g_{\alpha_N \alpha_{N-1}} \circ g_{\alpha_{N-1}\alpha_{N-2}} \circ \cdots \circ g_{\alpha_1\alpha_0}$$

of elements of γ, with maximal domain an open neighborhood of P_0 in T, where P_k is a plaque in U_{α_k}, $0 \leq k \leq N$. Let Γ' denote the countable set of pure compositions of elements of γ with maximal domains. General elements of Γ agree locally with such pure compositions, hence, for each $x \in T$, the set of all points $g(x)$ with $g \in \Gamma'$ is the full Γ–orbit of x.

There is a canonical one–to–one correspondence between the set of Γ–orbits in T and the set of leaves in X. Indeed, the $\mathcal{F}$–saturated subsets $Y \subseteq X$ correspond exactly to the Γ–invariant subsets of T by $Y \leftrightarrow Y \cap T$.

Lemma 3.1. *Under the above correspondence, the meager (respectively, residual) $\mathcal{F}$–saturated sets in X correspond exactly to the meager (respectively, residual) Γ–invariant sets in T.*

Proof. If $Y \subset X$ is a meager, $\mathcal{F}$–saturated set, it is clear that $Y \cap T$ is a meager subset of T. Conversely, if $Y \cap T$ is meager, local compactness and separability of T imply that $Y \cap T$ is contained in the union of compact, nowhere dense subsets $\{S_n\}_{n=1}^\infty$ of T. Since the plaques are relatively compact, the union of the closures of the plaques corresponding to S_n is a compact subset $Z_n \subseteq X$ and $Y \subseteq \bigcup_{n=1}^\infty Z_n$. The local product structure implies that each Z_n has empty interior, so Y is a meager subset of X. $\qquad\square$

Lemma 3.2. *If $Y \subseteq T$ is meager (not necessarily Γ–invariant), the union $\Gamma(Y)$ of the Γ–orbits of the elements of Y is a meager, Γ–invariant set. Thus, the $\mathcal{F}$–saturation of Y is meager in X.*

Proof. Since $Y \subseteq Z$, where Z is a countable union of closed, nowhere dense subsets of T, it will be enough to prove that $\Gamma(Z)$ is meager. Let $g \in \Gamma'$. Since T is locally compact and separable, so is $\operatorname{dom} g$, hence $Z \cap \operatorname{dom} g$, is a countable union of compact, nowhere dense sets. Since g is a homeomorphism between open sets, $g(Z \cap \operatorname{dom} g)$ is also a countable union of compact, nowhere dense subsets of $\operatorname{im} g$, hence a countable union of closed, nowhere dense subsets of T. Since Γ' is countable

$$\Gamma(Z) = \bigcup_{g \in \Gamma'} g(Z \cap \operatorname{dom} g)$$

is meager. $\qquad\square$

If L is a leaf and $z \in L$. Fix a chart $(U_\alpha, x_\alpha, y_\alpha)$ containing z and let $\zeta = y_\alpha(z)$. The holonomy group of the leaf L at z is defined to be the group $\mathcal{H}_z(L)$ of germs at ζ of elements $g \in \Gamma$ such that $\zeta \in \operatorname{dom} g$ and $g(\zeta) = \zeta$. It is rather easy to see that this group is the homomorphic image of $\pi_1(L, z)$ and, up to isomorphism, is independent of the choice of $z \in L$. A leaf L is said to be *without holonomy* if $\mathcal{H}_z(L)$ is trivial. The following is due to Epstein, Millett and Tischler [4] and Hector [8].

Proposition 3.3. *Let $(X, \mathcal{F})$ be a foliated space. The union G_0 of leaves of $\mathcal{F}$ without holonomy is residual.*

Proof. Let $Z \subset T$ be the set of all points $\zeta \in T$, fixed by at least one $g \in \Gamma'$, but such that g has nontrivial germ at ζ. It should be clear that Z is Γ–invariant, being the intersection of T with the family of leaves with nontrivial holonomy group. For each $g \in \Gamma'$, the set F_g of fixed points is a closed subset of $\operatorname{dom} g$ and the set theoretic boundary $Z_g = \partial F_g$ is closed and nowhere dense in $\operatorname{dom} g$. As in the proof of Lemma 3.2, such a set is meager in T. It should be clear that

$$Z = \bigcup_{g \in \Gamma'} Z_g,$$

so the countability of Γ' implies that Z is meager. $\qquad\square$

In much of what follows, we work in the residual set G_0, neglecting the leaves with holonomy. It will also be necessary to neglect another meager set of leaves.

Lemma 3.4. *The set B of leaves which meet $\partial(\mathrm{dom}\, g_{\alpha\beta}) \subset T_\beta$, for at least one $g_{\alpha\beta} \in \gamma$, is meager.*

Proof. Each T_β is a locally compact metric space and $\partial(\mathrm{dom}\, g_{\alpha\beta})$ is a closed subset of T_β with empty interior. The countable union of these sets is a meager subset $Z \subset T$ and the saturation $B = \mathcal{F}(Z)$ is meager. $\qquad\square$

Remark. Unlike the meager set of leaves with holonomy, B depends on the choice of regular foliated atlas $\mathcal{U}$. If L is a specific leaf in which we are interested, we can choose the regular atlas so that L does not lie in B. Indeed, as in our discussion of regular foliated atlases, there is an extension $(U'_\alpha, x'_\alpha, y'_\alpha)$ of each $(U_\alpha, x_\alpha, y_\alpha)$ to a foliated chart with transverse space the $(r_\alpha + \varepsilon)$–ball $T'_\alpha \supseteq \overline{T}_\alpha$. There are continuously many allowed choices of $\varepsilon > 0$, while L meets T'_α in only countably many points, so a suitable small increase in r_α will make sure that the leaf L will not meet the set theoretic boundary of $\mathrm{dom}\, g_{\alpha\beta}$, for any of the finitely many choices of $\beta \in \mathfrak{A}$ for which $g_{\alpha\beta}$ is defined. Hereafter, we replace G_0 with the residual saturated set $G_0 \smallsetminus B$. Abusing notation, we let G_0 denote this set since no particular leaf without holonomy need ever be excluded.

4. The 1–skeleta of leaves

The plaque cover of each leaf L determines a graph L^*, called the 1–skeleton of L. The vertices of L^* are the plaques in L and two distinct vertices P and Q are joined by an edge if, as plaques, they intersect. Thus, when these vertices are viewed as points of T, we can write $Q = g_{\alpha\beta}(P)$, for a unique $g_{\alpha\beta} \in \gamma$, and the directed edge from P to Q can be labelled by $g_{\alpha\beta}$. Degenerate edges with labels $g_{\alpha\alpha} = \mathrm{id}\,|T_\alpha$ will be suppressed. It is well known and elementary that the space of ends $\mathcal{E}(L^*)$ of this 1–complex is canonically the same as the space of ends $\mathcal{E}(L)$ of the leaf. The isomorphism $\kappa : \mathcal{E}(L^*) \to \mathcal{E}(L)$ is determined by the condition that a sequence $\{P_n\}_{n=1}^\infty$ of vertices of L^* converges to an end $e \in \mathcal{E}(L^*)$ if and only if the corresponding sequence of plaques in L converges to $\kappa(e)$. It is here that the assumption of leafwise completeness becomes essential, along with the existence of a finite upper bound to the diameters of the plaques.

Remark. With some caution, it is possible to view the 1–skeleta as defining a kind of foliated space in which the leaves are graphs. There is some problem

with local product structure since, if a vertex $x \in T_\beta$ lies in $\partial(\mathrm{dom}\, g_{\alpha\beta})$, for some $\alpha \in \mathfrak{A}$, no edge emanating from x is labeled by $g_{\alpha\beta}$, but there are vertices $y \in T_\beta$, arbitrarily near x, out of which there do emanate edges with label $g_{\alpha\beta}$. This difficulty disappears if we disallow 1–skeleta of leaves in the meager family B of Lemma 3.4.

Recall that Γ' denotes the countable set of all pure compositions

$$g = g_{\alpha_N \alpha_{N-1}} \circ g_{\alpha_{N-1} \alpha_{N-2}} \circ \cdots \circ g_{\alpha_1 \alpha_0}$$

of elements of the holonomy cocycle. For each such $g \in \Gamma'$ and $x \in \mathrm{dom}\, g$, g defines an edgepath on the 1–skeleton L_x^* of the leaf L_x through x with successive vertices

$$P_0 = x, P_1 = g_{\alpha_1 \alpha_0}(P_0), \ldots, P_N = g_{\alpha_N \alpha_{N-1}}(P_{N-1}).$$

The union of the vertices and edges of the edgepath is a finite, connected subcomplex $K_g(x) \subseteq L_x^*$, called the *trace* of g at x. As g varies over the countable set Γ' and x varies over $\mathrm{dom}\, g$, $K_g(x)$ varies over all finite, connected subcomplexes of the 1–skeleta of leaves. Of course, the same complex can be described as the trace at any one of its vertices of countably many distinct elements of Γ'. We will routinely use this observation to write a finite, connected supercomplex $K \supseteq K_g(x)$ as $K = K_{f \circ g}(x)$, for suitable $f \in \Gamma'$.

Given $g \in \Gamma'$ as above, set $g_0 = g_{\alpha_0 \alpha_0}$ and

$$g_k = g_{\alpha_k \alpha_{k-1}} \circ \cdots \circ g_{\alpha_1 \alpha_0}, \quad 1 \le k \le N.$$

Each vertex $P \in K_g(x)$ is equal to $g_k(x)$, for at least one value of k and, if P and Q are vertices connected by an edge in $K_g(x)$, there is at least one value of k for which one of these vertices, say P, is $g_k(x)$ and $Q = g_{k+1}(x)$. The directed edge from P to Q is labelled by $g_{\alpha_{k+1} \alpha_k}$, this label being independent of the allowable choices of the integer k.

Lemma 4.1. *Let $g \in \Gamma'$ and $x \in \mathrm{dom}\, g$. Then there is a neighborhood $V_x \subseteq \mathrm{dom}\, g$ of x and a canonical surjection*

$$\pi_y : K_g(y) \to K_g(x)$$

of 1–complexes, defined for each $y \in V_x$, which preserves the labels $g_{\alpha\beta}$ of directed edges.

Proof. By continuity, the conditions $g_k(x) \ne g_\ell(x)$, $0 \le k < \ell \le N$ are *open*. That is, there is an open neighborhood V_x of x in $\mathrm{dom}\, g$ such that

$$g_k(x) \ne g_\ell(x) \Rightarrow g_k(y) \ne g_\ell(y), \quad \forall\, y \in V_x.$$

Thus, π_y is well defined on the vertices of $K_g(y)$ by $\pi_y(g_k(y)) = g_k(x)$, for each $y \in V_x$, $0 \le k \le N$. By the above remarks, this is surjective and extends linearly to a canonical surjection of 1–complexes preserving the labels on directed edges.$\square$

Note that an arbitrary edgepath in L_x^*, based at x, is determined by the unique $h \in \Gamma'$ such that $x \in \operatorname{dom} h$ and the successive vertices of this path are

$$x = h_0(x), h_1(x), \ldots, h_N(x) = h(x).$$

We denote the path by $\sigma_h(x)$. It is a *simple* edgepath if it has no repeated vertices and it is a *loop* if $h(x) = x$. A loop is *simple* if its only repeated vertices are the initial and terminal vertex x and it is *basic* if it has the form $\sigma_{h^{-1} \circ f \circ h}(x)$ where $\sigma_h(x)$ is a simple edgepath and $\sigma_f(h(x))$ is a simple edgeloop. A finite, connected subcomplex $K \subseteq L_x^*$ admits only finitely many basic edgeloops and a standard trick of factoring an edgeloop essentially into a composition of basic ones reduces the following condition to a finite one..

Definition 4.2. A finite, connected subcomplex $K \subseteq L_x^*$ containing the vertex x has trivial holonomy if, for every edgeloop $\sigma_h(x)$ in K, h has trivial germ at x.

The following is a combinatorial version of local Reeb stability which will be essential.

Corollary 4.3. *If $g \in \Gamma'$ and $x \in \operatorname{dom} g$ are such that $L_x \not\subseteq B$ and $K_g(x)$ has trivial holonomy, there is a neighborhood $V_x \subseteq \operatorname{dom} g$ of x such that the projection $\pi_y : K_g(y) \to K_g(x)$ is an isomorphism of 1–complexes, $\forall\, y \in V_x$. Furthermore, $K_g(y) \cap K_g(z) = \emptyset$, whenever $z, y \in V_x$ and $z \ne y$.*

Proof. We must prove that the neighborhood V_x of Lemma 4.1 can be chosen so small that, for $0 \le k < \ell \le N$,

$$g_k(x) = g_\ell(x) \Rightarrow g_k(y) = g_\ell(y), \quad \forall\, y \in V_x.$$

But the condition implies that $x = h_{k\ell}(x)$, where $h_{k\ell} = g_\ell^{-1} \circ g_k$. Clearly $h_{k\ell}$ defines an edgeloop $\sigma_{h_{k\ell}}(x)$ in $K_g(x)$. By the hypothesis of trivial holonomy, $h_{k\ell}$ is defined and fixes every point in some neighborhood $W_{k\ell}$ of x in V_x. The new V_x should be the intersection of these $W_{k\ell}$'s.

For the second assertion, Let $z, y \in V_x$ and suppose that $y \ne z$. The claim that $K_g(y) \cap K_g(z) = \emptyset$ means that these subcomplexes have no vertex in common. Suppose, to the contrary, that $g_\ell(y) = g_k(z)$, for suitable k and ℓ. Since g_k is one–to–one, $k \ne \ell$ and we write $y = h_{k\ell}(z)$, where $h_{k\ell} = g_\ell^{-1} \circ g_k \in \Gamma'$. If such points y and z can be found arbitrarily near x, the fact that $L_x \not\subseteq B$ implies that $x \in \operatorname{dom} h_{k\ell}$, hence $h_{k\ell}(x) = x$. But this implies that $h_{k\ell}$ defines an edgeloop $\sigma_{h_{k\ell}}(x)$ in $K_g(x)$ such that $h_{k\ell}(z) \ne z$. This possibility has been eliminated by the choice of V_x.

Let $T_0 = T \cap G_0$, a Γ–invariant, residual subset of T. Our usual application of Corollary 4.3 is to $K_g(x)$ with $x \in T_0$, the hypothesis being clearly satisfied. But we do have need of the more general case.

Definition 4.4. The star of a vertex P of L_x^* is the union $\operatorname{star}(P)$ of the open edges emanating from P. A vertex P of $K_g(x)$ is an interior point of $K_g(x)$ if $\operatorname{star}(P) \subset K_g(x)$ and, otherwise, the vertex is a boundary point. The subcomplex $\partial(K_g(x))$ spanned by the boundary points is called the boundary of $K_g(x)$ and the subcomplex $\operatorname{int}(K_g(x))$ spanned by the interior points is called the interior.

The following is a fairly obvious consequence of the fact that T_0 meets no leaf of the meager set B (*cf.* Lemma 3.4 and the remark following).

Lemma 4.5. *If $g \in \Gamma'$ and $x \in \operatorname{dom} g$ are as in* Corollary 4.3, *the neighborhood $V_x \subset \operatorname{dom} g$ in that corollary can be chosen so that $\operatorname{star}(P)$ and $\operatorname{star}(\pi_y(P))$ have edges with exactly the same labels, $\forall P \in K_g(y)$, $\forall y \in V_x$. Consequently, $\pi_y(\partial(K_g(y))) = \partial(K_g(x))$ and $\pi_y(\operatorname{int}(K_g(y))) = \operatorname{int}(K_g(x))$.*

5. Dense ends

The following result plays a role in our theory analogous to that of [6, Proposition Fondamentale] in Ghys's theory, but the proof is completely different.

Proposition 5.1. *In a complete foliated metric space $(X, \mathcal{F})$, either no leaf of $\mathcal{F}$ is totally recurrent in X or a residual family $G \subseteq G_0$ consists of leaves that are totally recurrent in X.*

We will prove this fundamental proposition using 1–skeleta of leaves. Some preliminary discussion is needed. To begin with, remark that the definition of "dense end" of a noncompact leaf L can be reformulated as follows: $e \in \mathcal{E}_d(L)$ *if and only if, given an arbitrary open subset $W \subseteq X$, there is a sequence of points $\{x_k\}_{k=1}^{\infty} \subset L \cap W$ which converges to the end e.* Since we assume a finite upper bound on diameters of plaques, we can reformulate this by requiring existence of a sequence $\{P_k\}_{k=1}^{\infty}$ of plaques in L converging to e and satisfying $P_k \cap W \neq \emptyset$. The analogous condition is formulated for 1–skeleta as follows.

Definition 5.2. Let L^* be the 1–skeleton of a leaf L, $e \in \mathcal{E}(L^*)$. We say that e is a dense end if, for each open subset $V \subseteq T$, there is a sequence of vertices $\{P_k\}_{k=1}^{\infty}$ of $L^* \cap V$ converging to e. The set of dense ends will be denoted by $\mathcal{E}_d(L^*)$. We say that L^* is totally recurrent if $\mathcal{E}(L^*) = \mathcal{E}_d(L^*)$.

Lemma 5.3. *The canonical isomorphism $\kappa : \mathcal{E}(L^*) \to \mathcal{E}(L)$ carries $\mathcal{E}_d(L^*)$ exactly*

onto $\mathcal{E}_d(L)$. *In particular, L is totally recurrent if and only if the 1–skeleton L^* is totally recurrent.*

Lemma 5.4. *Let $g \in \Gamma'$ and let $O \subseteq T$ be open. If $\mathcal{F}$ has a totally recurrent leaf, then the set O_g of $x \in \operatorname{dom} g$ for which each unbounded component of $L_x^* \smallsetminus K_g(x)$ meets O contains an open dense subset of $\operatorname{dom} g$.*

Lemma 5.4 is delicate. Before proving it, we show how it implies our fundamental proposition.

Proof of Proposition 5.1. Assume that some leaf $L \subset G_0$ is totally recurrent. We will prove that the Γ–invariant set Z of points $x \in T_0$ such that L_x^* is not totally recurrent is meager. By Lemma 5.3, this will be enough. Let $O \subseteq T$ be open, choose $g \in \Gamma'$, and define $Z_g(O)$ to be the set of points $x \in T_0 \cap \operatorname{dom} g$ such that some unbounded component of $L_x^* \smallsetminus K_g(x)$ does not meet O. By Lemma 5.3, $Z_g(O) = (T_0 \cap \operatorname{dom} g) \smallsetminus O_g$ is meager. Uniting these sets as g ranges over the countable set Γ' gives a meager set $Z(O) \subset T_0$ with the property that the 1–skeleton F^* of a leaf $F \subset G_0$ meets $Z(O)$ if and only if F^* contains a finite, connected subcomplex K such that some unbounded component of $F^* \smallsetminus K$ does not meet O. Taking the union of the sets $Z(O)$ as O ranges over a countable base of the topology of T gives a meager set with Γ–saturation exactly Z. By Lemma 3.2, Z is meager. □

We turn to the proof of Lemma 5.4. Write

$$g = g_{\alpha_N \alpha_{N-1}} \circ \cdots \circ g_{\alpha_1 \alpha_0},$$

let $x \in T_0 \cap \operatorname{dom} g$ and let V_x be any open neighborhood of x in $\operatorname{dom} g$ small enough to satisfy Corollary 4.3 and Lemma 4.5. We will show that $O_g \cap V_x$ contains an open subset. Since T_0 meets $\operatorname{dom} g$ in a dense set of points and V_x can be chosen as small as desired, the assertion will follow.

Claim 1. *For arbitrary $y, z \in V_x$, there is a canonical isomorphism*

$$\pi_y^z : K_g(y) \to K_g(z),$$

defined by $\pi_y^z(g_k(y)) = g_k(z)$, $0 \leq k \leq N$. These isomorphisms preserve the interiors and boundaries of these complexes, as well as the labels of edges emerging from corresponding boundary vertices. In particular, the subcomplex $K_g(y)$ has trivial holonomy, $\forall\, y \in V_x$.

Proof. Indeed, $\pi_y : K_g(y) \to K_g(x)$ satisfies Corollary 4.3 and Lemma 4.5, for each $y \in V_x$, so we define

$$\pi_y^z = \pi_z^{-1} \circ \pi_y$$

and all assertions follow. □

Claim 2. *We lose no generality in assuming that each edge emerging from a boundary vertex of $K_g(z)$ either lies in $K_g(z)$ or has its other vertex in $L_z^* \smallsetminus K_g(z)$, $\forall\, z \in V_x$.*

Proof. If ℓ is an edge of L_x^* violating this, it is a bounded component of the complement and can be adjoined to $K_g(x)$, forming a supercomplex $K_{f \circ g}(x)$. Since $\mathrm{dom}\, f \circ g \subseteq \mathrm{dom}\, g$ and $x \in T_0 \cap \mathrm{dom}\, f \circ g$, a possibly smaller choice of V_x guarantees that this edge propagates via π_x^z to a bounded component of $L_z^* \smallsetminus K_{f \circ g}(z)$, $\forall\, z \in V_x \subset \mathrm{dom}\, f \circ g$. Thus, we can replace $K_g(z)$ with $K_{f \circ g}(z)$. Finite repetition of this procedure proves the claim. $\qquad\square$

For each $z \in V_x$, let ℓ_i range over those edges in $L_z^* \smallsetminus K_g(z)$ emerging from boundary vertices of $K_g(z)$, $1 \leq i \leq m$. Each component of $L_z^* \smallsetminus K_g(z)$ contains at least one of these edges and, by Claim 1, we are allowed to use these same labels ℓ_i for the analogously defined edges issuing from $\partial K_g(w)$, for all $w \in V_x$. Denote by z_i the initial vertex of the edge ℓ_i emerging from $K_g(z)$. Of course, this will generally give multiple indices to the same vertex of $\partial K_g(z)$. By Claim 2 the other vertex of ℓ_i does not lie in $K_g(z)$.

By hypothesis, there is a point $z \in V_x$ such that L_z is totally recurrent. By total recurrence, each unbounded component of $L_z^* \smallsetminus K_g(z)$ meets O. If the component containing ℓ_i is unbounded, there is a simple edgepath $s_i(z)$ having initial segment ℓ_i, meeting $K_g(z)$ only in z_i, and terminating in O. Such a path might also exist even if the component of $L_z^* \smallsetminus K_g(z)$ containing ℓ_i is bounded. At any rate, for as many components as possible, construct one such path in each. By relabelling, suppose that these paths are indexed by $1 \leq i \leq n$, some $n \leq m$. Since these paths are simple and lie in distinct components of the complement of $K_g(z)$, Claim 1 implies the following.

Claim 3. *The finite, connected subcomplex $K_g(z) \cup s_1(z) \cup \cdots \cup s_n(z)$ of the totally recurrent skeleton L_z^* has trivial holonomy.*

Write the subcomplex in Claim 3 as $K_{f \circ g}$, for suitable $f \in \Gamma'$. By this Claim, we can find a small enough neighborhood V_z of z in V_x, so that π_z^y extends canonically to an isomorphism

$$\pi_z^y : K_{f \circ g}(z) \to K_{f \circ g}(y),$$

for each $y \in V_z$. If V_z is small enough, π_z^y carries $s_i(z)$ to a simple path $s_i(y)$ in L_y^* which connects the boundary vertex y_i of $K_g(y)$ to a vertex in O and meets $K_g(y)$ only at y_i. Since L_z^* meets T in a dense set, we can choose $y \in V_z \smallsetminus \{z\}$ so that $L_z^* = L_y^*$. Of course, new components of $L_y^* \smallsetminus K_g(y)$ may appear, but m is an absolute upper bound on the number of components. If, for some such y, there is $i \geq n+1$, say $i = n+1$, such that ℓ_{n+1} enters a new unbounded component of $L_y^* \smallsetminus K_g(y)$, we build a new simple edgepath $s_{n+1}(y)$ in that component, ending in O and meeting $K_g(y)$ only at y_{n+1}. It may be possible to build several such new

edgepaths, $s_{n+1}(y), \ldots, s_{n+r}(y)$, even in some bounded components. We build as many as possible, but at most one per component. We now replace z with y and replay this game. This establishes the following.

Claim 4. *Without loss of generality, we can assume that, for every $y \in V_z \cap L_z^*$, the number of edgepaths $s_i(y)$, $1 \leq i \leq n$, cannot be increased by the above procedure. In particular, either $n = m$ or, for $n < i \leq m$, the edge ℓ_i either enters a bounded component of $L_z^* \smallsetminus K_g(z)$ or an unbounded component containing an $s_j(z)$, some $j \leq n$.*

Denote by $K_i(z)$ the closure of the bounded component entered by ℓ_i and reindex so that the distinct bounded components given by Claim 4 have closures $K_i(z)$, $n < i \leq r$, where $r \leq m$.

Claim 5. *For V_z sufficiently small, $y \in V_z$ arbitrary and $n < i \leq r$, every edgepath $\sigma_h(z_i)$ in $K_i(z)$, with initial and final vertices $z_i, v_i \in K_g(z)$ has lift $\sigma_h(y_i)$ with initial and final vertices $y_i = \pi_z^y(z_i), \pi_z^y(v_i) \in K_g(y)$.*

Proof. We can assume that $\sigma_h(z)$ meets $K_g(z)$ only in its initial and final vertex. Otherwise, break it down into such segments, treating each separately. In standard fashion, arbitrary edgepaths of the given type essentially factor into a product of simple paths and basic loops, all of this type, so we only need to check finitely many such paths. Choose V_z small enough that $h(y_i)$ is defined, $\forall\, y \in V_z$. That is, the lift $\sigma_h(y_i)$ is defined. Of course, the initial vertex y_i is in $K_g(y)$, but suppose the final vertex $h(y_i) \notin K_g(y)$, for suitable choices of $y \in V_z$ arbitrarily near z. Evidently these points form an open subset $U \subset V_z$ which clusters at z. Since L_z^* meets $\operatorname{dom} g$ in a dense set of points, we can choose $y \in U \cap L_z^*$. Since y can be chosen arbitrarily near z, the projection π_y^z extends to a (nonisomorphic) projection $\widetilde{\pi}_y^z$ of $K_g(y) \cup \sigma_h(y_i)$ onto $K_g(z) \cup \sigma_h(z_i)$ (Lemma 4.1). Since this projection carries $K_g(y)$ isomorphically onto $K_g(z)$ and $\sigma_h(y_i)$ onto $\sigma_h(z_i)$, the assumption that this latter edgepath has no intermediate vertices in $K_g(z)$ implies that the edgepath $\sigma_h(y_i)$ meets $K_g(y)$ only in its initial vertex. Its terminal vertex is a boundary point of $K_g(w)$ and $w \neq y$. By Claim 3 and Corollary 4.3, it should be clear that $\sigma_h(y_i)$ can be extended to an edgepath which meets $K_g(y)$ only in its initial vertex y_i and terminates in O. If the edgepath is not simple, we make it so by deleting subloops, hence contradicting Claim 4. Since the restriction of $\widetilde{\pi}_y^z$ to $K_g(y)$ is an isomorphism inverting π_z^y, the initial and final vertices of $\sigma_h(y_i)$ are as asserted. $\square$

Claim 6. *The subcomplex $K_g(z) \cup s_1(z) \cup \cdots \cup s_n(z) \cup K_{n+1}(z) \cdots \cup K_r(z)$ has trivial holonomy.*

Proof. We proceed by finite induction, adding $K_i(z)$, $n+1 \leq i \leq r$, one at a time.

Consider $K_g(z) \cup K_{n+1}(z)$ and loops in this complex based at a common boundary vertex z'. Such a loop factors into a sequence of paths, alternately in $K_g(z)$ and in $K_{n+1}(z)$. By Claim 1 and Claim 5, the loop has trivial holonomy. If it is necessary to attach another $K_{n+2}(z)$, replace $K_g(z)$ with $K_{f \circ g}(z) = K_g(z) \cup K_{n+1}(z)$ and repeat the argument. By finite induction and the fact that the paths $s_j(z)$ are simple, disjoint (except, perhaps, for initial vertices), and meet no $K_i(z)$ (except, perhaps, at initial vertices), the assertion follows.   □

To complete the proof of Lemma 5.4, choose $h \in \Gamma'$ such that

$$K_{h \circ g}(z) = K_g(z) \cup s_1(z) \cup \cdots \cup s_n(z) \cup K_{n+1}(z) \cup \cdots \cup K_m(z)$$

(repetitions allowed, as well as the possibility $n = m$). By Claim 6, Corollary 4.3 and Lemma 4.5, we can choose V_z so that every unbounded component of $L_y^* \smallsetminus K_g(y)$ (and, perhaps, some bounded ones) meets O, $\forall\, y \in V_z$. Then $V_z \subseteq V_x \cap O_g$. This is exactly what we needed to show.

Proposition 5.1 is fundamental to the proof of Theorem A. The following somewhat easier result is fundamental to the proof of Theorem B.

Proposition 5.5. *In a complete, foliated metric space, either no leaf has a dense end or a residual family $G_d \subseteq G_0$ consists of leaves having at least one dense end.*

The following terminology will simplify language in the proof of Proposition 5.5.

Definition 5.6. Let L^* be the 1–skeleton of a leaf, $W \subseteq L^*$. Then W is transversely dense if the vertices of L^* in W form a dense subset of T.

Remark. For L^* to have a dense end, it is necessary but not sufficient that L^* be transversely dense. For example, L^* might be transversely dense, but have a vertex x that is an isolated point of T. Then every end of L^* will have a neighborhood W, the vertices of which do not cluster at x in T.

Lemma 5.7. *Let $g \in \Gamma'$ and $x \in \mathrm{dom}\, g$. Let W be a component of $L_x^* \smallsetminus K_g(x)$. If L_x^* is tranversely dense and if W contains a sequence $\{P_n\}_{n=1}^{\infty}$ of vertices which converge to x in T, then W is a neighborhood of a dense end of L_x^* and, in particular, is transversely dense.*

Proof. Since $x \notin W$, the terms of the sequence $\{P_n\}_{n=1}^{\infty}$ can be assumed to be distinct. Thus, any finite subcomplex of L_x^* can contain at most finitely many of the terms P_n. That is, the sequence "diverges to infinity" in L_x^* and, passing to a subsequence if necessary, we assume that $\{P_n\}_{n=1}^{\infty}$ converges to an end e of L^*. Clearly, W is a neighborhood of e and we will prove that $e \in \mathcal{E}_d(L^*)$.

Let $K_1 \subset K_2 \subset \cdots \subset K_m \subset \cdots$ be an exhaustion of L^* by finite, connected subcomplexes, chosen so that the component W_m of $L^* \smallsetminus K_m$ which is a neighborhood of e is a subset of W. Fix $m \geq 1$. Given an arbitrary open subset $O \subseteq T$,

the fact that L_x^* is transversely dense allows us to choose an edgepath s in L^* connecting x to a point of O. This edgepath corresponds to a composition

$$h = g_{\alpha_N \alpha_{N-1}} \circ \cdots \circ g_{\alpha_1 \alpha_0},$$

hence lifts to pairwise disjoint paths s_n in L_x^* connecting P_n to O, for n sufficiently large. For large enough values of n, the vertex P_n lies in W_m and, for possibly larger values, s_n has no vertex in common with K_m. It follows that a vertex of W_m lies in O. Since O is an arbitrary open subset of T, W_m is transversely dense. Since $m \geq 1$ is arbitrary, $e \in \mathcal{E}_d(L^*)$. $\square$

Corollary 5.8. *The 1–skeleton L^* of a leaf L has a dense end $\iff L^* \smallsetminus K$ is transversely dense, for all finite, connected subcomplexes $K \subset L^*$.*

Proof. The implication "$\Rightarrow$" is clear. For "$\Leftarrow$", choose $g \in \Gamma'$ and $x \in \operatorname{dom} g$ such that $L^* = L_x^*$ and take $K = K_g(x)$ in Lemma 5.7. Since $L_x^* \smallsetminus K_g(x)$ is transversely dense, so is L^*. Furthermore, $L_x^* \smallsetminus K_g(x)$ contains a sequence $\{P_n\}_{n=1}^{\infty}$ of vertices converging in T to x. Since $L_x^* \smallsetminus K_g(x)$ has only finitely many components, a subsequence lies entirely in one of these components W. By Lemma 5.7, W is a neighborhood of a dense end of L^*. $\square$

Lemma 5.9. *Let $g \in \Gamma'$ and let $O \subseteq T$ be open. If some leaf of $\mathcal{F}$ has a dense end, then the set O_g of $x \in \operatorname{dom} g$ for which $L_x^* \smallsetminus K_g(x)$ meets O contains an open dense subset of $\operatorname{dom} g$.*

Proof. Let $y \in T_0 \cap \operatorname{dom} g$ and choose an open neighborhood V_y of y in $\operatorname{dom} g$ small enough to satisfy Corollary 4.3 and Lemma 4.5. If L^* has a dense end, there is a point $x \in L^* \cap V_y$ and a simple edgepath ℓ starting at a boundary vertex of $K_g(x)$, otherwise lying in $L_x^* \smallsetminus K_g(x)$, and having its terminal vertex in O. As in the proof of Lemma 5.4, $K_g(x) \cup \ell$ satisfies Reeb stability, so there is an open neighborhood of x in $V_y \cap O_g$. Since G_0 meets $\operatorname{dom} g$ in a dense subset, the assertion follows. $\square$

Proof of Proposition 5.5. Given $g \in \Gamma'$, Lemma 5.9 implies that the intersection of the sets O_g as O ranges over a countable base of the topology of T is residual in $\operatorname{dom} g$. The meager complement Z_g in $\operatorname{dom} g$ of this set consists of those points $x \in \operatorname{dom} g$ such that, for some open subset $O \subseteq T$, $L_x^* \smallsetminus K_g(x)$ does not meet O. The union of the Z_g's as g ranges over Γ' is meager and its meager Γ–saturation Z consists of all vertices of those skeleta F^* such that $F^* \smallsetminus K$ does not meet T in a dense subset, for some choice of finite, connected subcomplex K. By Corollary 5.8, Z is exactly the family of leaves having no dense end. $\square$

6. Proof of Theorem A

By Proposition 5.1, Theorem A is really a corollary of Theorem B. Nevertheless, a direct proof of Theorem A, along the lines outlined in §2, is significantly easier than the proof of Theorem B, so we give it here.

By Proposition 5.1, we assume that total recurrence is topologically generic. Let $g \in \Gamma'$ and let $G_g \subset T_0 \cap \operatorname{dom} g$ be the residual set of points x such that L_x is totally recurrent.

Lemma 6.1. *If $k \geq 1$ is an integer, the set of $x \in G_g$ such that $L_x^* \smallsetminus K_g(x)$ has at most k components is relatively open in G_g.*

Proof. The condition on x, call it $\mathbf{P}_x$, can be stated equivalently that $K_g(x)$ is contained in $\operatorname{int}(K_{f \circ g}(x))$ and $K_{f \circ g}(x) \smallsetminus K_g(x)$ has at most k components, for a suitable choice of $f \in \Gamma'$. This makes sense for arbitrary $x \in \operatorname{dom} g$. By Corollary 4.3 and Lemma 4.5, if $x \in G_g$, then $\mathbf{P}_x \Rightarrow \mathbf{P}_y$, for all y in a suitable open neighborhood V_x of x in $\operatorname{dom} g$, hence for the points of a relatively open neighborhood of x in G_g. $\qquad\square$

Corollary 6.2. *If $k \geq 1$ is an integer, the set of $x \in G_g$ such that $L_x^* \smallsetminus K_g(x)$ has at most k unbounded components is relatively open in G_g.*

Proof. Let $x \in G_g$ be such that $L_x^* \smallsetminus K_g(x)$ has at most k unbounded components and choose $f \in \Gamma'$ such that $K_{f \circ g}(x)$ is exactly the union of $K_g(x)$ and the bounded components of its complement. Then $L_x^* \smallsetminus K_{f \circ g}(x)$ consists of the unbounded components of $L_x^* \smallsetminus K_g(x)$, the number of these being $n \leq k$. By the proof of Lemma 6.1, there is an open neighborhood V_x of x in $\operatorname{dom} f \circ g \subseteq \operatorname{dom} g$ such that $L_y^* \smallsetminus K_{f \circ g}(y)$ has at most n components, $\forall y \in V_x$, so the number of unbounded components is at most $n \leq k$. $\qquad\square$

Proposition 6.3. *The set Z_g^3 of $x \in G_g$ such that $K_g(x)$ separates L_x^* into three or more unbounded components is relatively closed in G_g. If it has nonempty interior in the relative topology of G_g, then every leaf in the residual set of totally recurrent leaves in G_0 has a Cantor set of ends.*

Proof. The first assertion is an immediate corollary of Corollary 6.2, applied to the case $k = 2$. For the second, assume that $V = \operatorname{int} Z_g^3 \neq \emptyset$ and let $L \subset G_0$ be totally recurrent. Then every neighborhood W of every end e of L^* meets V. That is, W contains a finite, connected subcomplex K_W that separates L^* into at least three unbounded components (*cf.* Figure 2, where the corresponding picture in the leaf L is drawn.). Equivalently, every neighborhood of every end e is a neighborhood of at least one end e' distinct from e. The compact, separable, totally disconnected, metrizable set of ends of L^* has no isolated points, hence is a Cantor set. $\qquad\square$

Corollary 6.4. *If the set of leaves in G having a Cantor set of ends is not residual, the set of leaves having more than two ends is meager, so a residual set of leaves has at most two ends.*

Proof. By Proposition 6.3, Z_g^3 is relatively closed and has empty interior in G_g, $\forall\, g \in \Gamma'$. That is, $Z_g^3 = G_g \cap Z_g$, where Z_g is relatively closed in $\mathrm{dom}\, g$, and Z_g must have empty interior in $\mathrm{dom}\, g$ since G_g is residual, hence dense in $\mathrm{dom}\, g$. As usual, we conclude that Z_g is meager, hence that $Z = \bigcup_{g \in \Gamma'} Z_g$ is meager. The $\mathcal{F}$–saturation of Z is meager (Lemma 3.2) and contains the set of leaves in G having three or more ends. Since G is residual, all assertions follow. $\square$

We assume that only a meager set of leaves in G has more than two ends and replace G_g with the residual subset of points x for which the totally recurrent leaf L_x^* has at least one and at most two ends. Hereafter, G_g denotes this set. We define $Z_g^2 \subseteq G_g$ to be the set of points x such that $K_g(x)$ separates L_x^* into components, exactly two of which are unbounded.

Proposition 6.5. *The set Z_g^2 is relatively closed in G_g. If it has nonempty relative interior in G_g, a residual set of totally recurrent leaves in G_0 consists of leaves with exactly two ends. If the set of leaves having exactly two ends is not residual, it is meager and a residual set of totally recurrent leaves in G_0 has exactly one end.*

Proof. The first assertion is immediate by Corollary 6.2, applied to the case $k = 1$. For the second, assume that $V = \mathrm{int}\, Z_g^2 \neq \emptyset$. By Corollary 6.4, a residual set of totally recurrent leaves $L \subset G_0$ have at most two ends. These leaves meet V, hence have exactly two ends. As in the proof of Corollary 6.4, if the set of 2–ended leaves is not residual, it is meager and the set of 1–ended leaves is residual. $\square$

There remain the assertions in Theorem A about 2–dimensional leaves. Skeleta are not needed. Indeed, we work in the residual subset $G \subseteq G_0$ consisting of totally recurrent leaves without holonomy and argue exactly as in §2.

7. Proof of Theorem B

If no leaf has a dense end, the conclusion of Theorem B holds, so Proposition 5.5 allows us to assume that a residual family $G_d \subseteq G_0$ consists of leaves having at least one dense end. Relatively residual subsets of G_d are residual in X and relatively meager subsets are meager in X. Similar remarks hold for $T_d = T \cap G_d$ versus T, allowing us to replace X with G_d, T with T_d, $\mathcal{F}$ with $\mathcal{F}|G_d$ and Γ with $\Gamma|T_d$. The following, therefore, enables us to simplify notation and exposition.

Lemma 7.1. *Without loss of generality, it can be assumed that every leaf of $\mathcal{F}$ has a dense end and trivial holonomy.*

The one thing we need be careful with is the loss of local compactness. For instance, the proof of Lemmas 3.1 and 3.2 made essential use of the local compactness of X and T, but the conclusions of those lemmas relativise.

Fix $g \in \Gamma'$ and let Z_g^3 denote the set of points $x \in \operatorname{dom} g$ such that at least three components of $L_x^* \smallsetminus K_g(x)$ are transversely dense. The analogously defined set in Proposition 6.3 was proven to be closed, but this is no longer guaranteed and we must take its closure $\overline{Z}_g^3$ in $\operatorname{dom} g$.

Proposition 7.2. *If $\overline{Z}_g^3$ has nonempty interior V, then all leaves in a residual family $G \subseteq X$ have a Cantor set of dense ends.*

We prove this proposition in a series of lemmas.

Since no leaf has holonomy, Corollary 4.3 and Lemma 4.5 can be applied freely at all points $x \in V$.

Lemma 7.3. *If $V \neq \emptyset$ and $O \subseteq T$ is open, the subset $O_g \subseteq V$ of points x such that 3 or more distinct components of $L_x^* \smallsetminus K_g(x)$ meet O contains an open dense subset of V.*

Proof. In the following argument, we avoid clumsy phraseology by saying that an edgepath having initial vertex in $\partial K_g(x)$, but otherwise not meeting $K_g(x)$, "lies in a component of $L_x^* \smallsetminus K_g(x)$". The points of $Z_g^3 \cap V$ are dense in V. Let x be such a point and let V_x be an open neighborhood of x in V as in Corollary 4.3 and Lemma 4.5. We will show that some open subset of V_x lies in O_g. For $y \in V_x$, let $\ell_1, \ldots, \ell_n$ be the distinct edges emanating from vertices of $\partial K_g(y)$ and not lying in $K_g(y)$. As in § 5, Claim 2, we can assume that the terminal vertex of each ℓ_i does not lie in $K_g(y)$. As in Claim 1 of § 5, the labels of these edges do not depend on the choice of $y \in V_x$ and, in particular, n is an upper bound to the number of components of $L_y^* \smallsetminus K_g(y)$, $\forall y \in V_x$. Since $x \in Z_g^s$, we can find edgepaths $s_i(x)$, $1 \leq i \leq 3$, each having initial segment one of the ℓ_j's, each lying in distinct components of $L_x^* \smallsetminus K_g(x)$, and having terminal vertices in O. We renumber so that ℓ_i is the initial edge of $s_i(z)$, $1 \leq i \leq 3$. Since Corollary 4.3 applies to $K_g(x) \cup s_1(x) \cup s_2(x) \cup s_3(x)$, we can make V_x smaller, if necessary, so that, as y varies over V_x, the lifts $s_i(y)$ of these paths have all these same properties except, perhaps, for lying in distinct components of $L_y^* \smallsetminus K_g(y)$. If, for some y, this property does fail, assume that $s_2(y)$ and $s_3(y)$ lie in the same component. Since there is an edgepath in this component joining a vertex of $s_2(y)$ to a vertex of $s_3(y)$, Corollary 4.3 insures that there is an open neighborhood V_y of y in V_x such that $s_2(z)$ and $s_3(z)$ lie in the same component C of $L_z^* \smallsetminus K_g(z)$, $\forall z \in V_y$. In particular, choose $z \in Z_g^3 \cap V_y$. Choose paths $s_4(z)$ and $s_5(z)$ in $L_z^* \smallsetminus \operatorname{int}(K_g(z))$,

at least one of which does not lie in C, having terminal vertices in O, such that $s_1(z)$, $s_4(z)$ and $s_5(z)$ lie in distinct components. At least one of these has an initial edge, call it ℓ_4 that is not one of ℓ_i, $1 \leq i \leq 3$. Whichever of these initial edges have been discarded will never be used again as we repeat this procedure. Since $n - 2$ is an absolute upper bound on the number of repetitions possible, we will stop upon finding a point w, a neighborhood $V_w \subseteq V_x$ of w and edgepaths $\sigma_i(w)$, $1 \leq i \leq 3$, such that, for every $u \in V_w$, the paths $\sigma_i(u)$ lie in distinct components of $L_u^* \smallsetminus K_g(u)$ and have terminal vertices in O. Thus, $V_w \subseteq O_g$. $\square$

Continuing to assume that $V \neq \emptyset$, let $Y \subseteq V$ denote the intersection of the O_g's as O ranges over a countable base $\mathcal{B}$ of the topology of T. Thus, Y is a residual subset of V. Let $X = V \smallsetminus Y$, a meager set.

Lemma 7.4. *The set Y is exactly $V \cap Z_g^3$.*

Proof. It is clear that $V \cap Z_g^3 \subseteq Y$. For the reverse inclusion, select an arbitrary point $y \in Y$ and remark that L_y^* is transversely dense. Let $\{O_k\}_{k=1}^\infty$ be a fundamental system of neighborhoods of y in V, all belonging to $\mathcal{B}$ and let $W_1, \ldots, W_n$ denote the distinct components of $L_y^* \smallsetminus K_g(y)$. These can be numbered so that $W_1, \ldots, W_{n_1}$ are the components meeting O_1. Here, by the definition of Y, $3 \leq n_1 \leq n$. At least three of these will also meet $O_2 \subset O_1$, so another renumbering gives $W_1, \ldots, W_{n_2}$ as the subset of these n_1 components that also meet O_2, $3 \leq n_2 \leq n_1$. After finitely many such steps, we obtain $W_1, \ldots, W_m$, all meeting O_k, $1 \leq k < \infty$, where $m \geq 3$. By Lemma 5.7, W_i is transversely dense, $1 \leq i \leq m$. $\square$

Proof of Proposition 7.2. The Γ–saturation of $X = V \smallsetminus Y$ is meager, so we work in the complement $H \subseteq T$ of this saturation, a residual, Γ–invariant set. If $x \in H$, each neighborhood W of each dense end e of L_x^* meets the open set V, necessarily in points $y \in Y$. For all but finitely many of these points y, $K_g(y) \subset W$. By Lemma 7.4, $K_g(y)$ splits W into two or more transversely dense components, W_1 and W_2 and, by Lemma 5.7, these are neighborhoods of dense ends e_1 and e_2. Thus, W is a neighborhood of at least one dense end distinct from e, proving that $\mathcal{E}_d(L_x^*)$ has no isolated points. Since $\mathcal{E}_d(L_x^*)$ is compact, totally disconnected, separable and metrizable, it is a Cantor set. $\square$

Corollary 7.5. *If the family of leaves having a Cantor set of dense ends is not residual, then the family of leaves with three or more dense ends is meager.*

Proof. By Proposition 7.2, if there is not a residual set of leaves with a Cantor set of dense ends, then $\overline{Z}_g^3$ has empty interior, $\forall\, g \in \Gamma'$. The union of these is then meager, as is its $\mathcal{F}$–saturation Z, and every leaf with at least three dense ends lies in Z. $\square$

We are reduced to the case in which the leaves of a residual set have at most 2 dense ends. We are assuming that all leaves have at least one such end. Let Z_g^2 be the set of points $x \in \text{dom}\, g$ such that at least two components of $L_x^* \smallsetminus K_g(x)$ are transversely dense. In fact, any transversely dense component is a neighborhood of a dense end (Lemma 5.7), so $L_x^* \smallsetminus K_g(x)$ will have exactly two such components, $\forall\, x \in Z_g^2$.

Proposition 7.6. *If the set of leaves having* 2 *dense ends is not residual, it is meager and the set of leaves having one dense end is residual.*

Proof. As in the proof of Proposition 7.2, if $\overline{Z}_g^2$ has nonempty interior V, then $Y = V \cap Z_g^2$ is residual in V. The Γ–saturation of $V \smallsetminus Y$ is meager and its complement is a residual, Γ–invariant set $G \subseteq T$. For each $x \in G$, L_x^* meets V, necessarily in Y, hence has 2 dense ends. Alternatively, $\overline{Z}_g^2$ is meager, as is the $\mathcal{F}$–saturation of the union of these sets as g ranges over Γ', so the remaining leaves have exactly one dense end and form a residual set. $\square$

The assertions about 2–dimensional leaves are proven exactly as in Theorem A, so the proof of Theorem B is complete.

8. Generic 2–ended leaves

In this section, we restrict ourselves to compact, leafwise C^1 laminations of a closed manifold M. That is, the compact, foliated metric space $(X, \mathcal{F})$ is topologically imbedded in M in such a way that each leaf is a C^1–immersed, p–dimensional submanifold of M. We also require that there be a continuous orientation of the tangent bundle $T(X, \mathcal{F})$ of the lamination. These hypotheses allow us to use the theory of structure cycles [13] for the lamination. These are closed de Rham p–currents on M in the cone of p–currents generated by the Dirac currents $v_x \in \Lambda^p(T_x(X, \mathcal{F}))$, $x \in X$. A fundamental result of Sullivan [13] identifies the cone of structure cycles canonically with the cone of holonomy invariant measures for $(X, \mathcal{F})$ which are finite on compact subsets of T.

In the following, the term "generic leaf" refers to an arbitrary leaf in the residual family G of Theorem A.

Theorem 8.1. *If $(X, \mathcal{F})$ has a totally recurrent leaf and the generic leaf is* 2–*ended, then there is a probability measure μ on T, invariant under the holonomy pseudogroup and supported in $\overline{L} \cap T$, where L is a generic,* 2–*ended leaf.*

We sketch the proof. Referring to the proof of Proposition 6.5, we see that, if the generic leaf (totally recurrent and without holonomy) has 2 ends, there is $g \in \Gamma'$ such that Z_g^2 has nonempty interior. Let L be one of these generic leaves

and $x \in \operatorname{dom} g$ a point where L^* meets int Z_g^2. Replacing g with suitable $f \circ g$, we can assume that $K_g(x)$ separates L_x^* into exactly two components W_1 and W_2, neighborhoods of the respective ends e_1 and e_2. Then, for every point $y \in \operatorname{int} Z_g^2$ such that L_y^* is the skeleton of a generic leaf, $K_g(y)$ separates L_y^* into exactly two components, neighborhoods of the respective ends.

In the leaf L, $K_g(x)$ can be realized as a compact, connected submanifold K which is the union of the closures of the plaques which are vertices of $K_g(x)$. These closed plaques can be assumed to have C^1 boundaries intersecting transversely, so K will have piecewise C^1 boundary. By an arbitrarily small modification, the boundary of K can be made C^1. Let N denote the union of those boundary components that interface one of the complementary components, say W_1, and remark that N is a compact submanifold of L of dimension $p-1$ which separates L into exactly two components, neighborhoods of e_1 and e_2, respectively.

By local Reeb stability, there is an open subset $V \subset \operatorname{int} Z_g^2$ and an imbedding $V \times N \hookrightarrow M$ such that $N_y = \{y\} \times N$ lies in the leaf L_y through y. If L_y is one of the generic leaves (in particular, L itself), N_y separates it into two components, neighborhoods of its two ends. We can choose sequences $\{y_k\}_{k=1}^\infty$ and $\{z_k\}_{k=1}^\infty$ in $L \cap V$ such that $y_k \to e_1$ and $z_k \to e_2$ in $L \cup \mathcal{E}(L)$ as $k \to \infty$. Let $A_k \subset L$ be the compact submanifold cobounded by N_{z_k} and N_{y_k}. Since these boundary manifolds remain uniformly bounded as $k \to \infty$, they also converge to the respective ends and $\operatorname{vol}_p(A_k) \to \infty$. The linear functional $\varphi_k : A^p(M) \to \mathbb{R}$, defined by

$$\varphi_k(\omega) = \frac{1}{\operatorname{vol}_p(A_k)} \int_{A_k} \omega,$$

defines a de Rham p–current and the sequence $\{\varphi_k\}_{k=1}^\infty$ is easily seen to be bounded away from ∞ and 0 in the space of p–currents. These are structure currents for $(X, \mathcal{F})$, so a subsequence converges to a nontrivial structure current μ. Stokes's theorem and the fact that $\{\operatorname{vol}_p(A_k)\}_{k=1}^\infty$ is unbounded, while $\{\operatorname{vol}_{p-1}(\partial A_k)\}_{k=1}^\infty$ is bounded, implies that μ is a nontrivial structure cycle, canonically identified as a nontrivial, Γ–invariant Borel measure on T, finite on compact subsets. Finally, this measure is readily seen to be supported in $\overline{L} \cap T$.

The regularity of the foliated atlas allows us to replace T with a slightly smaller transverse space with compact closure, so we can assume that $0 < \mu(T) < \infty$. Normalizing produces the desired probability measure.

We remark that the sequence $\{A_k\}_{k=1}^\infty$ is an example of an *averaging sequence* in the sense of Goodman and Plante [7].

9. Codimension one

We indicate some applications of our results to closed, C^2–foliated manifolds $(M, \mathcal{F})$ of codimension one. It is assumed that M is orientable and that $\mathcal{F}$ is transversely orientable.

Let $X \subset M$ be an exceptional minimal set. A celebrated, but unpublished, theorem of Duminy asserts that the semiproper leaves of $\mathcal{F}|X$ have a Cantor set of ends. If the foliation is real analytic, this result holds for all the leaves $\mathcal{F}|X$. If the holonomy of $\mathcal{F}|X$ is generated by a Markov chain (all of the standard examples of exceptional minimal sets have this property), the authors have shown that all the leaves have a Cantor set of ends [3], but the general case remains open. The following is a small but interesting step in the right direction.

Theorem 9.1. *If $(X, \mathcal{F}|X)$ is an exceptional minimal set of a foliated manifold $(M, \mathcal{F})$ as above, there is a residual family $G \subseteq X$ of leaves such that either every leaf in G has a Cantor set of ends or every leaf in G has exactly one end.*

Proof. Since $(X, \mathcal{F}|X)$ is minimal and compact, every leaf is totally recurrent, so the family G can be chosen as in Theorem A. If the leaves in G are 2–ended, Theorem 8.1 provides a holonomy invariant probability measure on T. By minimality, $\operatorname{supp} \mu = X \cap T$. The C^2 hypothesis allows us to apply a theorem of Sacksteder [12] according to which some leaf L of $\mathcal{F}|X$ supports a 2–sided holonomy contraction. In standard fashion, it follows that each of the infinitely many points x of $L \cap T$ has the same *positive* measure $\mu(\{x\})$, contradicting the fact that $\mu(T) = 1$. $\square$

According to the theory of levels [1], also called the "architecture" of foliations [9], there is a filtration

$$\emptyset = M_{-1} \subset M_0 \subseteq M_1 \subseteq \cdots \subseteq M_k \subseteq \cdots \subseteq M$$

by compact, $\mathcal{F}$–saturated subsets such that $M_k \smallsetminus M_{k-1}$ is the union of the minimal sets of $\mathcal{F}|(M \smallsetminus M_{k-1})$, $0 \le k < \infty$. The fact that $\mathcal{F}|U$ has minimal sets, where $U \subseteq M$ is any open, $\mathcal{F}$–saturated subset, is a consequence of the C^2 hypothesis. The leaves of $\mathcal{F}|(M_k \smallsetminus M_{k-1})$ are said to be at level k. The union $M_* = \bigcup_{k \ge 0} M_k$ is the family of leaves at finite level and $M_\infty = M \smallsetminus M_*$ is the family of leaves at infinite level. An end e of a leaf L is at level $k \le \infty$ if the maximal level of leaves in its asymptote $A_e \subseteq M$ is k. We denote by $\mathcal{E}_k(L)$ the set of ends of L at level $k \le \infty$.

If L is a leaf at finite level k, the family of leaves in $X = \overline{L}$ at level k is an open, dense subset of X. By the methods of [1], one shows that an end e of such a leaf is at level k if and only if e is a dense end for the foliated space $(X, \mathcal{F}|X)$. Of course, if L is a *proper* leaf at level k, the highest level of any of its ends is $k - 1$, but in all other cases, level k leaves have some level k ends.

In light of Theorem B, we have the following consequence of this discussion.

Theorem 9.2. *If L is a leaf at level $k < \infty$, there is a residual family of leaves $G \subseteq \overline{L}$ such that every leaf in G has 0, 1, 2, or a Cantor set of ends at level k, the cardinality of $\mathcal{E}_k(F)$ being constant as F ranges over the leaves in G*

If L is a leaf at infinite level, $X = \overline{L}$ contains uncountably many leaves at

infinite level, all asymptotic to L [1, §5]. This family of leaves is a residual subset of X. Indeed, the sets $X \cap M_k$ are compact and, being in the closure of L, they have empty interior in X, so the union of leaves at finite level in X is meager. Again, the ends at infinite level are exactly the dense ends relative to the foliated space $(X, \mathcal{F}|X)$ and every leaf at infinite level has such an end.

Theorem 9.3. *If L is a leaf at infinite level, there is a residual family $G \subset \overline{L}$ of leaves, every one of which has 1, 2, or a Cantor set of ends at infinite level, the cardinality of $\mathcal{E}_\infty(F)$ being constant as F ranges over the leaves in G.*

10. Examples

There are numerous examples which illustrate Theorem A and, in many cases, the measure theoretic version in [6]. We sketch a few here and refer the reader to [6] for more.

Example 10.1. Consider the Cantor set $K = \{0,1\}^{\mathbb{Z}}$, the bi–infinite sequence of 0's and 1's with the Tychonoff topology. The shift map $\tau : K \to K$ is defined by

$$\tau\left((x_i)_{i \in \mathbb{Z}}\right) = (y_i)_{i \in \mathbb{Z}}, \quad y_{i+1} = x_i, \quad \forall\, i \in \mathbb{Z}.$$

This is a homeomorphism with well understood dynamics. There are countably many periodic orbits and these are dense in K. There are also minimal sets made up of nonperiodic orbits. There are some nonperiodic orbits which are dense in forward time, some in backward time, and some in both forward and backward time. We claim that these latter form a residual subset of K. Indeed, a base for the topology of K is given by the countable family

$$U(\varepsilon_{i_1}, \varepsilon_{i_2}, \dots, \varepsilon_{i_n}) = \{(x_i)_{i \in \mathbb{Z}} \mid x_{i_j} = \varepsilon_{i_j}, 1 \leq j \leq n\},$$

where the choice of integers $i_1 < i_2 < \cdots < i_n$ is fixed, as are the choices of $\varepsilon_{i_j} \in \{0,1\}$, $1 \leq j \leq n$. Fix one such basic open set U. The reader will easily check that the set of $x \in K$ such that the forward τ–orbit $\{\tau^k(x)\}_{k=0}^{\infty}$ meets U at least once is open and dense in K. In fact, given an integer $n \geq 1$, the same considerations show that the set of points whose forward τ–orbit meets U at least n times is open and dense. The residual intersection of these sets is the set of points x such that $\tau^k(x) \in U$, for infinitely many integers $k \geq 0$. Intersecting these residual sets as U ranges over the given base gives a residual set consisting of those points of K with τ–orbit dense in forward time. Similarly, one obtains a residual set with τ–orbit dense in backward time, hence the intersection of these two residual sets is again residual and consists of the points $x \in K$ with τ–orbits dense in both forward and backward time.

If Σ is a compact, orientable surface with infinite fundamental group, there are surjective homomorphisms

$$\varphi : \pi_1(\Sigma) \to \mathbb{Z}.$$

Composing φ with the $\mathbb{Z}$–action generated by τ on K, we obtain an action of $\pi_1(\Sigma)$ on K. The standard suspension construction then produces a foliated bundle over Σ with fiber the Cantor set K. This is a compact, foliated metric space $(X, \mathcal{F})$ in which the leaves are covering spaces of Σ and correspond one–one to the τ–orbits in K. The leaves corresponding to periodic orbits are compact and these are the only leaves with nontrivial holonomy. If Σ is a torus, so are the compact leaves, though they occur with arbitrarily large volumes proportional to the lengths of the corresponding periodic orbits. If genus $\Sigma > 1$, the compact leaves occur with arbitrarily large genera. The remaining leaves are 2–ended, being cylinders if Σ is a torus, ladders if it has genus greater than one. Corresponding nonorientable leaves are produced by taking Σ nonorientable of large enough genus. By the analysis of the dynamics of τ, a residual family consists of 2–ended leaves that are totally recurrent, but there are infinitely many 2–ended leaves that are not totally recurrent. This shows that the residual set G in Theorem A can be a proper subset of G_0. Evidently, this is a non–minimal example.

Remark that the set $K = \{1, 2, \ldots, n\}^{\mathbb{Z}}$ is also a Cantor set and one can define a closed subset H by using an $n \times n$ incidence matrix (a_{ij}) of 0's and 1's. The sequences $(x_k)_{k \in \mathbb{Z}} \in H$ are determined by the condition that i, j can be consecutive terms x_k, x_{k+1} only if $a_{ij} = 1$. It is clear that the shift map τ on K restricts to a homeomorphism of H onto itself, called "a subshift of finite type". In interesting cases, H itself is a Cantor set, the dynamical properties are similar to those described above, and one obtains a foliated space by suspension.

Example 10.2. Subshifts of finite type can also be realized in foliations of honest manifolds. Let $f : T^2 \to T^2$ be the Anosov diffeomorphism which lifts to the linear transformation of $\mathbb{R}^2$ having unimodular matrix

$$A = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}.$$

This f–action admits a Markov partition making it semi–conjugate to a subshift $\tau|H : H \to H$ of finite type. That is, there is a continuous, finite–to–one surjection $\pi : H \to T^2$ such that the diagram

$$
\begin{array}{ccc}
H & \xrightarrow{\ \tau|H\ } & H \\
\pi \downarrow & & \downarrow \pi \\
T^2 & \xrightarrow[\ f\ ]{} & T^2
\end{array}
$$

commutes. Here, H is a Cantor subset of $\{1, 2, 3, 4, 5\}^{\mathbb{Z}}$, the matrix for the subshift being

$$\begin{bmatrix} 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \end{bmatrix}$$

(*cf.* [11, pp. 84–86]). By considerations similar to those in the previous example, $\tau|H$ has a countable dense set of periodic points and a residual family of nonperiodic orbits that are dense in both forward and backward time. The periodic orbits project by π to a dense family of periodic orbits of f. The projection π restricts to a homeomorphism of a certain residual subset of H onto a residual subset of T^2, hence f has a residual family of orbits that are dense in forward and backward time. The suspension construction yields a foliated T^2–bundle $p : M \rightarrow \Sigma$ with a dense family of compact leaves and a residual family G of totally recurrent, 2–ended leaves without holonomy. Again, only the compact leaves have holonomy while many noncompact leaves fail to be totally recurrent, so G is a proper subfamily of G_0.

This construction also exemplifies Ghys's result. Since the matrix A is unimodular, f preserves Lebesgue measure on T^2. It follows that Lebesgue measure on M is a completely invariant harmonic measure for the foliation. Since the countable family of compact leaves has Lebesgue measure zero, the 2–ended leaves form a set of full measure.

Example 10.3. The projective special linear group $\mathrm{PSl}(2, \mathbb{R})$ is the group of orientation preserving isometries of the hyperbolic plane $\mathbb{H}^2$ and has a natural identification with $T_1(\mathbb{H}^2)$, the bundle of unit tangent vectors to $\mathbb{H}^2$. Since $\mathbb{H}^2$ covers the compact, orientable, hyperbolic surface Σ of genus 2, the covering group is a discrete, cocompact subgroup $\Gamma \subset \mathrm{PSl}(2, \mathbb{R})$ and the compact quotient $M = \Gamma\backslash \mathrm{PSl}(2, \mathbb{R})$ is canonically identified with the the unit tangent bundle $T_1(\Sigma)$. Here, as the notation indicates, the quotient is the set of right cosets of Γ. Let $H \subset \mathrm{PSl}(2, \mathbb{R})$ be the subgroup

$$H = \left\{ \begin{bmatrix} a & b \\ 0 & a^{-1} \end{bmatrix} \right\}_{a > 0}$$

and consider the foliation $\mathcal{H}$ of $\mathrm{PSl}(2, \mathbb{R})$ by left cosets of H. This foliation is invariant under left translations in the group, hence passes to a foliation $\mathcal{F}$ of $T_1(\Sigma)$ which is transverse to the circle fibers. It is well known that each leaf of $\mathcal{H}$ is the unit tangent field to a geodesic pencil in $\mathbb{H}^2$ issuing from a point on the circle at infinity. In the quotient, most of these leaves survive as copies of the hyperbolic plane, but a countable infinity are hyperbolic cylinders and have nontrivial holonomy. Each leaf is everywhere dense, so $(M, \mathcal{F})$ is a minimal

foliated manifold and the planar leaves form the residual set G_0. This foliated manifold supports a unique ergodic harmonic measure [5, Proposition 5] and the union of the countably many leaves with holonomy has measure zero. Finally, a simple modification of this foliated manifold produces an example in which the generic leaves are 1–ended with infinite genus. Simply drill out a small tubular neighborhood N of one of the circle fibers of $T_1(\Sigma)$, noting that $\mathcal{F}$ restricts to a foliation of this solid torus by disks transverse to ∂N. Replace N with a copy of $K \times S^1$, foliated by copies of K, where K is either a handle or a crosscap. This foliation matches up with $\mathcal{F}|(M \smallsetminus \text{int } N)$ and the generic leaves (in either sense) are 1–ended with handles (respectively, crosscaps) clustering at that end.

Example 10.4. In [2], it was shown that, given an arbitrary closed 3–manifold M and an arbitrary noncompact, orientable surface L, there is a smooth foliation $\mathcal{F}$ of M having a leaf diffeomorphic to L. This leaf eludes visual intuition by lying at *infinite level* in the foliation. By the theory of levels, $X = \overline{L} \subseteq M$ is a compact foliated subspace containing leaves at every finite level, as well as uncountably many leaves at infinite level, each of the latter being asymptotic to every leaf in X [1, §5]. Generally, the leaf L will not be totally recurrent in X, but if we take $L \cong \mathbb{R}^2$, then the fact that L has one end and is asymptotic to every leaf (including itself) in X does imply total recurrence. This leaf, being simply connected, also has trivial holonomy. By the proof of Proposition 6.3 and Corollary 6.4, the set of leaves with a Cantor set of ends must be meager. Otherwise, the totally recurrent leaf L without holonomy would have a Cantor set of ends. Then, by the proof of Proposition 6.5, the set of leaves with exactly two ends must be meager. Otherwise, the totally recurrent leaf $L \subset G_0$, having at most two ends, would have exactly two ends. Similarly, every leaf without holonomy has genus zero and it follows that the topologically generic leaf in $(X, \mathcal{F}|X)$ is the plane. We do not know whether these leaves lie in the support of a harmonic measure.

Many other topological types can also occur, nongenerically, among the leaves of $\mathcal{F}|X$. If $H_2(M; \mathbb{R}) \neq 0$, the construction in [2] can be carried out so that X contains leaves at every finite depth. The well understood theory of leaves at finite depth (*cf.* [1, §6], where they are called "totally proper" leaves) shows that there are at least countably many distinct topological types of noncompact leaves in X. It is also possible to carry out the construction in arbitrary M so that X contains an exceptional minimal set. By Duminy's (unpublished) theorem, X would then contain some leaves with a Cantor set of ends. Finally, the trick of modifying the foliation along a closed transveral through L produces a generic family of one ended leaves of infinite genus or one with infinitely many crosscaps. This example is not minimal.

Example 10.5. The easiest example in which the generic leaf has a Cantor set of ends is obtained by a construction of M. Hirsch. This is discussed in [6], but we

Figure 3. Forming the solid torus with wormhole drilled out

want to give more detail, making clear the role of the leaves without holonomy. Let P denote the "pair of pants", the compact surface obtained by deleting from S^2 the interiors of three disjoint disks, and consider $P \times [0, 1]$. Let $\varphi : P \to P$ be an orientation preserving diffeomorphism which leaves invariant one boundary component and interchanges the other two. The 3–manifold

$$N = P \times [0, 1]/\{(x, 0) \equiv (\varphi(x), 1)\}$$

is a solid torus with a wormhole drilled out which winds around twice longitudinally while winding once meridianally (Figure 3). This is fibered over $S^1 = [0, 1]/\{0 \equiv 1\}$ with fiber P. We parametrize these fibers as P_t, $0 \le t \le 1$, where $P_0 = P_1$. One can glue the inner toral boundary to the outer one smoothly and in such a way that each of the inner boundary components of P_t is glued to outer boundary components of $P_{h_0(t)}$ and $P_{h_1(t)}$, respectively, where $h_0 : [0, 1] \to [0, 1/2]$ and $h_1 : [0, 1] \to [1/2, 1]$ are orientation preserving diffeomorphisms (Figure 4).

The resulting 3–manifold is closed and the pairs of pants fit together to form the leaves of a smooth foliation. We consider the case in which $h_0(t) = t/2$ and $h_1(t) = (t + 1)/2$. In order to see the topology of the leaves, it will be convenient to write $t \in [0, 1]/\{0 \equiv 1\}$ by its dyadic expansion $t = 0.t_1 t_2 \cdots t_n \cdots$, where $t_i \in \{0, 1\}$, $i \ge 1$. In this notation, the diffeomorphisms take the form

$$h_0(0.t_1 t_2 \cdots) = 0.0t_1 t_2 \cdots$$
$$h_1(0.t_1 t_2 \cdots) = 0.1t_1 t_2 \cdots$$

Thus, if t has a periodic expansion
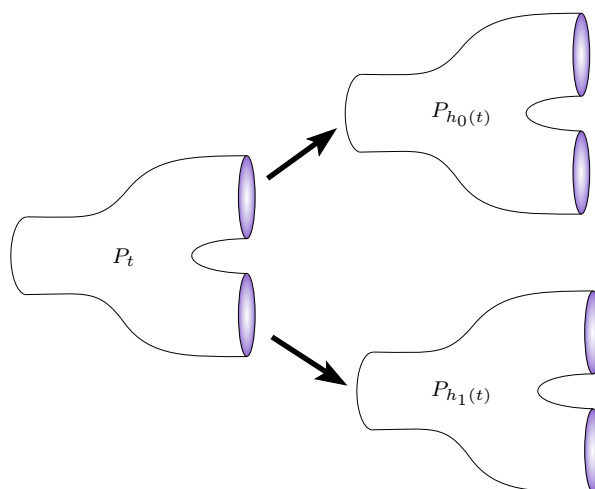
$$t = 0.t_1 \cdots t_r t_1 \cdots t_r \cdots,$$

Figure 4. Gluing together the pairs of pants

$h_{t_1} \circ \cdots \circ h_{t_r}(t) = t$ and the leaf containing P_t has a single cycle of pairs of pants, hence has genus one. Since there are no other cycles, the leaf has a Cantor set of ends. If t is irrational, the leaf containing P_t is a Cantor tree of genus zero, there being no cyclic connections.

One can let the P_t's play the role of plaques. While h_0 and h_1 are not globally well defined on the parameter circle $[0,1]/\{0 \equiv 1\}$, local restrictions are defined and can be used to fashion a pseudogroup on this circle which can serve as a holonomy pseudogroup for $(M, \mathcal{F})$. A little care is needed at the point $0 = 1$, where h_0 on either side matches up with h_1 on the other side. Those leaves containing a plaque P_t with t cyclic of period $t_1 \cdots t_r$ clearly have a holonomy contraction defined by $h_{t_1} \circ \cdots \circ h_{t_r}$. There are countably many of these, all of genus 1. The remaining leaves only contain P_t's with t irrational. They have trivial holonomy and, as noted above, are Cantor trees of genus zero. Every leaf is dense in M, hence $(M, \mathcal{F})$ is minimal and all leaves are totally recurrent. Again, the trick of modifying along a closed transversal produces Cantor trees with handles or Cantor trees with crosscaps.

Finally, we remark that the uniformity of the topology in the generic 2–dimensional leaf has no analogues in higher dimensions. In a truly startling example [6, pp. 396–399], Ghys exhibits a compact, foliated 6–manifold of codimension 2 in which a minimal set X consists of 2–ended leaves no two of which are homeomorphic.

References

[1] J. Cantwell and L. Conlon, Poincaré–Bendixson theory for leaves of codimension one, *Trans. Amer. Math. Soc.* **265** (1981), 181–209.

[2] J. Cantwell and L. Conlon, Every surface is a leaf, *Topology* **26** (1987), 265–285.

[3] J. Cantwell and L. Conlon, Leaves of Markov local minimal sets in foliations of codimension one, *Publicacions Matemàtiques* **33** (1989), 461–484.

[4] D. B. A. Epstein, K. C. Millet, and D. Tischler, Leaves without holonomy, *J. London Math. Soc.* **16** (1977), 548–552.

[5] L. Garnett, Foliations, the ergodic theorem and Brownian motion, *J. Funct. Anal.* **51** (1983), 285–311.

[6] E. Ghys, Topologie des feuilles génériques, *Ann. of Math.* **141** (1995), 387–422.

[7] S. E. Goodman and J. F. Plante, Holonomy and averaging in foliated sets, *J. Diff. Geo.* **14** (1979), 401–407.

[8] G. Hector, *Feuilletages en cylindres*, Lecture Notes in Math. **597**, Springer–Verlag, 1977, pp. 252–270.

[9] G. Hector, Architecture des feuilletages de class C^2, *Astérisque* **107–108** (1983), 243–258.

[10] H. Hopf, Enden offener Räume und unendliche diskontinuerliche Gruppen, *Comm. Math. Helv.* **16** (1944), 81–100.

[11] A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*, Cambridge University Press, Cambridge 1995.

[12] R. Sacksteder, Foliations and pseudogroups, *Amer. J. Math.* **87** (1965), 79–102.

[13] D. Sullivan, Cycles for the dynamical study of foliated manifolds and complex manifolds, *Inv. Math.* **36** (1976), 225–255.

John Cantwell
Department of Mathematics
St. Louis University
St. Louis, MO 63103, USA
e-mail: cantwelljc@slu.edu

Lawrence Conlon
Department of Mathematics
Washington University
St. Louis, MO 63130, USA
e-mail: lc@math.wustl.edu