

SHAPE OF MATCHBOX MANIFOLDS

ALEX CLARK, STEVEN HURDER, AND OLGA LUKINA

ABSTRACT. In this work, we develop shape expansions of minimal matchbox manifolds without holonomy, in terms of branched manifolds formed from their leaves. Our approach is based on the method of coding the holonomy groups for the foliated spaces, to define leafwise regions which are transversely stable and are adapted to the foliation dynamics. Approximations are obtained by collapsing appropriately chosen neighborhoods onto these regions along a “transverse Cantor foliation”. The existence of the “transverse Cantor foliation” allows us to generalize standard techniques known for Euclidean and fibered cases to arbitrary matchbox manifolds with Riemannian leaf geometry and without holonomy. The transverse Cantor foliations used here are constructed by purely intrinsic and topological means, as we do not assume that our matchbox manifolds are embedded into a smooth foliated manifold, or a smooth manifold.

1. INTRODUCTION

In this work, we consider topological spaces \mathfrak{M} which are continua; that is, compact, connected metric spaces. We assume that \mathfrak{M} has the additional structure of a codimension-zero foliated space, so are *matchbox manifolds*. The path-connected components of \mathfrak{M} form the leaves of a foliation \mathcal{F} of dimension $n \geq 1$. The precise definitions are given in Section 2 below. Matchbox manifolds arise naturally in the study of exceptional minimal sets for foliations of compact manifolds, and as the tiling spaces associated to repetitive, aperiodic tilings of Euclidean space \mathbb{R}^n which have finite local complexity. They also arise in some aspects of group representation theory and index theory for leafwise elliptic operators for foliations, as discussed in the books [16, 47].

The class of Williams solenoids and the results about these spaces provide one motivation for this work. Recall that an expanding attractor Λ for an Axiom A diffeomorphism $f: M \rightarrow M$ of a compact manifold is a continuum; that is, a compact, connected metric space. Williams developed a structure theory for these spaces in his seminal works [63, 64, 65]. The hyperbolic splitting of the tangent bundle to TM along Λ yields a foliation of the space Λ by leaves of the expanding foliation for f , and the contracting foliation for f gives a transverse foliation on an open neighborhood $\Lambda \subset U \subset M$. Williams used this additional structure on a neighborhood of Λ to obtain a “presentation” of Λ as an inverse limit of “branched n -manifolds”, $\hat{f}: M_0 \rightarrow M_0$, where n is the dimension of the expanding bundle for f , and the map f induces the map \hat{f} between the approximations. The notion of a 1-dimensional branched manifold is easiest to define, as the branches are required only to meet each other at disjoint vertices. In higher dimensions, the definition of branched manifolds becomes more subtle, and especially for the “transversality condition” imposed on the cell attachment maps. The spaces Λ with this structure are called *Williams solenoids* in the literature. The topological properties of the approximating map $\hat{f}: M_0 \rightarrow M_0$ is used to study the dynamical system defined by f , for example as discussed in the works [24, 40, 57, 58] and others.

The Riemann surface laminations introduced by Sullivan [59] are compact topological spaces locally homeomorphic to a complex disk times a Cantor set, and a similar notion is used by Lyubich and Minsky in [45], and Ghys in [32]. These are also well-known examples of matchbox manifolds.

Key words and phrases. solenoids, matchbox manifold, laminations, Delaunay triangulations.
2000 *Mathematics Subject Classification.* Primary 57R30, 37C55, 37B45; Secondary 53C12 .
AC and OL supported in part by EPSRC grant EP/G006377/1.
Version date: August 15, 2013; revised February 6, 2014.

Associated to the foliation \mathcal{F} of a matchbox manifold \mathfrak{M} is a compactly generated pseudogroup $\mathcal{G}_{\mathcal{F}}$ acting on a “transverse” totally disconnected space \mathfrak{X} , which determines the transverse dynamical properties of \mathcal{F} . The first two authors showed in [17] that if the action of $\mathcal{G}_{\mathcal{F}}$ on \mathfrak{X} is equicontinuous, then \mathfrak{M} is homeomorphic to a weak solenoid (in the sense of [46, 55, 26]). Equivalently, the shape of \mathfrak{M} is defined by a tower of proper covering maps between compact manifolds of dimension n .

The purpose of this work is to study the shape properties of an arbitrary matchbox manifold \mathfrak{M} , without the assumption that the dynamics of the associated action of $\mathcal{G}_{\mathcal{F}}$ is equicontinuous, though with the assumption that \mathfrak{M} is minimal; that is, that every leaf of \mathcal{F} is dense. We also assume that \mathcal{F} is without holonomy. Our main result shows that all such spaces have an analogous structure as that of a Williams solenoid. An important difference between the case of Williams solenoids, and the general case we consider, is that the tower of approximations is not in general defined by a single map, but uses a sequence of maps between compact branched manifolds.

A *presentation* of a space Ω is a collection of continuous maps $\mathcal{P} = \{p_{\ell}: M_{\ell} \rightarrow M_{\ell-1} \mid \ell \geq 1\}$, where each M_{ℓ} is a connected compact branched n -manifold, and each $p_{\ell}: M_{\ell} \rightarrow M_{\ell-1}$ is a proper surjective map of branched manifolds, as defined in Section 10. It is assumed that there is given a homeomorphism h between Ω and the inverse limit space defined by

$$(1) \quad \mathcal{S}_{\mathcal{P}} \equiv \varprojlim \{p_{\ell}: M_{\ell} \rightarrow M_{\ell-1}\}.$$

A *Vietoris solenoid* [61] is a 1-dimensional solenoid, where each M_{ℓ} is a circle, and each $p_{\ell}: \mathbb{S}^1 \rightarrow \mathbb{S}^1$ is a covering map of degree greater than 1. More generally, if each M_{ℓ} is a compact manifold and each p_{ℓ} is a proper covering map, then we say that $\mathcal{S}_{\mathcal{P}}$ is a *weak solenoid*, as discussed in [46, 55, 26].

Here is our main result.

THEOREM 1.1. *Let \mathfrak{M} be a minimal matchbox manifold of dimension n . Assume that the foliation \mathcal{F} of \mathfrak{M} is without holonomy. Then there exists a presentation $\mathcal{P} = \{p_{\ell}: M_{\ell} \rightarrow M_{\ell-1} \mid \ell \geq \ell_0\}$ where each M_{ℓ} is a triangulated branched n -manifold and each p_{ℓ} is a proper simplicial map, so that \mathfrak{M} is homeomorphic to the inverse limit of the system of maps they define,*

$$(2) \quad h: \mathfrak{M} \cong \mathcal{S}_{\mathcal{P}} = \varprojlim \{p_{\ell}: M_{\ell} \rightarrow M_{\ell-1} \mid \ell \geq \ell_0\}.$$

The hypothesis that \mathcal{F} is without holonomy may possibly be removed, using a more general definition of the bonding maps allowed. The work by Benedetti and Gambaudo in [12] gives a model for such a result, in the case of tiling spaces associated with an action of a connected Lie group G . Also, the results of the first two authors in [17] applies to equicontinuous matchbox manifolds which may have non-trivial leafwise holonomy. On the other hand, a generalization of Theorem 1.1 to arbitrary minimal matchbox manifolds would yield solutions to several classical problems in the study of exceptional minimal sets for codimension-one foliations, for example as listed in [37, 38], so an extension of Theorem 1.1 to this generality would be important, but is expected to involve additional subtleties.

In this work, no assumptions are made on the geometry and topology of the leaves in \mathfrak{M} , beyond that they are complete Riemannian manifolds. In order to construct the branched manifold approximations to \mathfrak{M} , we use a technique based on “dynamical codings” for the orbits of the pseudogroup associated to the foliation, a method which generalizes that used by Gambaudo and Martens [31] for flows. This coding method can also be seen as extending a technique used by Thomas in his study of equicontinuous flows in [60]. The first two authors used this coding approach in their study of equicontinuous actions of groupoids on Cantor sets in [17], and it also appears implicitly in the work by Forrest [27] in his study of minimal actions of \mathbb{Z}^n on Cantor sets. The coding method we use is more formal, so applies in complete generality. The other ingredient required is the existence of a transverse Cantor foliation in the generality we consider. We construct these foliations in this work, extending the results of the authors’ paper [18]. The existence of a compatible transverse Cantor foliation \mathcal{H} on \mathfrak{M} is fundamental for defining the branched manifold approximations.

As mentioned above, tiling spaces provide an important class of examples for which Theorem 1.1 applies. Let \mathbf{T} be a tiling of \mathbb{R}^n which is repetitive, aperiodic, and has finite local complexity. The

tiling space $\Omega_{\mathbf{T}}$ is defined as the closure of the set of tilings obtained via the translation action of \mathbb{R}^n , in a suitable metric topology [7, 11, 28, 54]. The assumptions imply that $\Omega_{\mathbf{T}}$ is locally homeomorphic to a disk in \mathbb{R}^n times a Cantor set, so that $\Omega_{\mathbf{T}}$ is a matchbox manifold. (See [29] for a discussion of variants of this result.) The seminal result of Anderson and Putnam in [7] showed that under these assumptions, and under the assumption that the tiling \mathbf{T} is defined by an expanding substitution, then $\Omega_{\mathbf{T}}$ is homeomorphic to the inverse limit of a tower of branched flat n -manifolds. The approach of Anderson and Putnam, and the alternative approaches to generalizing their method and extending the conclusion to all tilings by Gähler and Sadun [53], and in works [4, 11, 12, 31, 43], used variations on a technique called “collaring” or a method called “inflation” of a Voronoi tessellation of \mathbb{R}^n , to obtain the branched manifold approximations for $\Omega_{\mathbf{T}}$. The more recent work [10] gives a version of this method closest to the approach we take in this paper. The work here is an important step in developing the general theory of point patterns and tilings for general, non-Euclidean spaces. As described in Senechal [56], the theory of point patterns and tilings is directly connected to the theory of quasicrystals, and so our work can be seen as a step towards understanding possible quasicrystalline structures in non-Euclidean spaces.

We discuss next the contents of the rest of this work, which culminates in the proof of Theorem 1.1. Section 2 introduces the basic definitions of a matchbox manifold, and Section 3 introduces the definitions and properties of the holonomy pseudogroups for matchbox manifolds, as developed in [17] and [18]. Section 4 discusses the Voronoi tessellations of the leaves associated to transversal clopen sets, then Section 5 develops some fundamental properties of restricted pseudogroups.

A coding of the orbits of the action of the holonomy pseudogroup of the foliation \mathcal{F} on \mathfrak{M} is developed in Section 6, and used to construct “dynamically defined” nested sequences of clopen coverings \mathcal{V}_ℓ for $\ell \geq \ell_0$ of a Cantor transversal to \mathcal{F} , which are “centered” on the transverse orbit defined by a fixed leaf L_0 . The study of the dynamics of the induced action on a descending chain of clopen neighborhoods is a standard technique in almost all approaches to the study of the dynamics of minimal actions on a Cantor set. It corresponds to the initial steps in the construction of a Kakutani tower for the measurable theory of group actions.

In Section 7, we associate to each clopen covering of the transversal space \mathfrak{X} a collection of open sets, called *Reeb neighborhoods* which form a covering of \mathfrak{M} . The results of [18] show that each such Reeb neighborhood projects along a transverse Cantor foliation to its compact base. Section 9 uses the results in [18] to show that these bi-foliated structures on the Reeb neighborhoods can be chosen to be compatible on overlaps. In Section 10, the proof of Theorem 1.1 is completed, by showing that the Reeb neighborhoods collapsed along the Cantor transversal foliation, can be identified (or “glued”) each to the other on their overlaps, to obtain a branched n -manifold. By iterating this process for a chain of clopen covers for \mathfrak{X} , we obtain a tower of maps whose inverse limit is homeomorphic to \mathfrak{M} .

Note that the works [4, 43] follow a similar approach to the above, but using an intuitive extension of the zooming technique for Voronoi cells as in [11, 12]. The geometry and topology of a generic leaf L_0 can be quite complicated, and not at all as intuitively straightforward as for the case of Euclidean spaces. Hence the geometry and topology of the Voronoi cells which decompose L_0 may be similarly non-intuitive. Some of the subtleties for Voronoi partitions at large scale in non-Euclidean manifolds are discussed, for example, in the work [42]. The more formal approach using codings that we use in this work avoids the need to control the geometry of the Voronoi cells introduced. The reader will also find the care taken with the domains of holonomy maps makes the proofs of some results quite tedious at times, but a careful consideration of examples shows that these concerns are required.

This work is part of a program to generalize the results of the thesis of Fokkink [25], started during a visit by the authors to the University of Delft in August 2009. The papers [17, 18] are the initial results of this study, and this work is the preparation for the forthcoming paper [19] which completes the program. The authors would like to thank Robbert Fokkink for the invitation to meet in Delft, and the University of Delft for its generous support for the visit. The authors’ stay in Delft was also supported by a travel grant No. 040.11.132 from the *Nederlandse Wetenschappelijke Organisatie*.

2. FOLIATED SPACES AND MATCHBOX MANIFOLDS

In this section, we give the precise definitions and results for matchbox manifolds, as required for this work. More detailed discussions with examples can be found in the books of Candel and Conlon [16, Chapter 11] and Moore and Schochet [47, Chapter 2]. In the following, details of proofs are omitted whenever possible, as they are presented in detail in these two texts and the paper [17].

DEFINITION 2.1. *A foliated space of dimension n is a continuum \mathfrak{M} , such that there exists a compact metric space \mathfrak{X} , and for each $x \in \mathfrak{M}$ there is a compact subset $\mathfrak{T}_x \subset \mathfrak{X}$, an open subset $U_x \subset \mathfrak{M}$, and a homeomorphism $\varphi_x: \bar{U}_x \rightarrow [-1, 1]^n \times \mathfrak{T}_x$ defined on the closure of U_x in \mathfrak{M} , such that $\varphi_x(x) = (0, w_x)$ where $w_x \in \text{int}(\mathfrak{T}_x)$. Moreover, it is assumed that each φ_x admits an extension to a foliated homeomorphism $\hat{\varphi}_x: \hat{U}_x \rightarrow (-2, 2)^n \times \mathfrak{T}_x$ where $\bar{U}_x \subset \hat{U}_x$ is an open neighborhood.*

The subspace \mathfrak{T}_x of \mathfrak{X} is called the *local transverse model* at x .

Let $\pi_x: \bar{U}_x \rightarrow \mathfrak{T}_x$ denote the composition of φ_x with projection onto the second factor.

For $w \in \mathfrak{T}_x$ the set $\mathcal{P}_x(w) = \pi_x^{-1}(w) \subset \bar{U}_x$ is called a *plaque* for the coordinate chart φ_x . We adopt the notation, for $z \in \bar{U}_x$, that $\mathcal{P}_x(z) = \mathcal{P}_x(\pi_x(z))$, so that $z \in \mathcal{P}_x(z)$. Note that each plaque $\mathcal{P}_x(w)$ for $w \in \mathfrak{T}_x$ is given the topology so that the restriction $\varphi_x: \mathcal{P}_x(w) \rightarrow [-1, 1]^n \times \{w\}$ is a homeomorphism. Then $\text{int}(\mathcal{P}_x(w)) = \varphi_x^{-1}((-1, 1)^n \times \{w\})$.

Let $U_x = \text{int}(\bar{U}_x) = \varphi_x^{-1}((-1, 1)^n \times \text{int}(\mathfrak{T}_x))$. Note that if $z \in U_x \cap U_y$, then $\text{int}(\mathcal{P}_x(z)) \cap \text{int}(\mathcal{P}_y(z))$ is an open subset of both $\mathcal{P}_x(z)$ and $\mathcal{P}_y(z)$. The collection of sets

$$\mathcal{V} = \{\varphi_x^{-1}(V \times \{w\}) \mid x \in \mathfrak{M}, w \in \mathfrak{T}_x, V \subset (-1, 1)^n \text{ open}\}$$

forms the basis for the *fine topology* of \mathfrak{M} . The connected components of the fine topology are called *leaves*, and define the foliation \mathcal{F} of \mathfrak{M} . For $x \in \mathfrak{M}$, let $L_x \subset \mathfrak{M}$ denote the leaf of \mathcal{F} containing x .

Definition 2.1 does not impose any smoothness conditions on the leaves of \mathcal{F} , so they may just be topological manifolds. The next definition imposes a *uniform* smoothness condition on the leaves.

DEFINITION 2.2. *A smooth foliated space is a foliated space \mathfrak{M} as above, such that there exists a choice of local charts $\varphi_x: \bar{U}_x \rightarrow [-1, 1]^n \times \mathfrak{T}_x$ such that for all $x, y \in \mathfrak{M}$ with $z \in U_x \cap U_y$, there exists an open set $z \in V_z \subset U_x \cap U_y$ such that $\mathcal{P}_x(z) \cap V_z$ and $\mathcal{P}_y(z) \cap V_z$ are connected open sets, and the composition*

$$\psi_{x,y;z} \equiv \varphi_y \circ \varphi_x^{-1}: \varphi_x(\mathcal{P}_x(z) \cap V_z) \rightarrow \varphi_y(\mathcal{P}_y(z) \cap V_z)$$

is a smooth map, where $\varphi_x(\mathcal{P}_x(z) \cap V_z) \subset \mathbb{R}^n \times \{w\} \cong \mathbb{R}^n$ and $\varphi_y(\mathcal{P}_y(z) \cap V_z) \subset \mathbb{R}^n \times \{w'\} \cong \mathbb{R}^n$. The leafwise transition maps $\psi_{x,y;z}$ are assumed to depend continuously on z in the C^∞ -topology on maps between subsets of \mathbb{R}^n .

A map $f: \mathfrak{M} \rightarrow \mathbb{R}$ is said to be *smooth* if for each flow box $\varphi_x: \bar{U}_x \rightarrow [-1, 1]^n \times \mathfrak{T}_x$ and $w \in \mathfrak{T}_x$ the composition $y \mapsto f \circ \varphi_x^{-1}(y, w)$ is a smooth function of $y \in (-1, 1)^n$, and depends continuously on w in the C^∞ -topology on maps of the plaque coordinates y . As noted in [47] and [16, Chapter 11], this allows one to define leafwise smooth partitions of unity, vector bundles, and tensors for smooth foliated spaces. In particular, one can define leafwise smooth Riemannian metrics. We recall a standard result, whose proof for foliated spaces can be found in [16, Theorem 11.4.3].

THEOREM 2.3. *Let \mathfrak{M} be a smooth foliated space. Then there exists a leafwise Riemannian metric for \mathcal{F} , such that for each $x \in \mathfrak{M}$, the leaf L_x inherits the structure of a complete Riemannian manifold with bounded geometry, and the Riemannian geometry of L_x depends continuously on x .*

Bounded geometry implies, for example, that for each $x \in \mathfrak{M}$, there is a leafwise exponential map $\exp_x^{\mathcal{F}}: T_x \mathcal{F} \rightarrow L_x$ which is a surjection, and the composition $\exp_x^{\mathcal{F}}: T_x \mathcal{F} \rightarrow L_x \subset \mathfrak{M}$ depends continuously on x in the compact-open topology on maps.

DEFINITION 2.4. *A matchbox manifold is a continuum with the structure of a smooth foliated space \mathfrak{M} , such that the transverse model space \mathfrak{X} is totally disconnected, and for each $x \in \mathfrak{M}$, $\mathfrak{T}_x \subset \mathfrak{X}$ is a clopen (closed and open) subset. A matchbox manifold \mathfrak{M} is minimal if every leaf of \mathcal{F} is dense.*

Intuitively, a 1-dimensional matchbox manifold has local coordinate charts which are homeomorphic to a “box of matches”, which gave rise to this terminology in the works [1, 2, 3].

The maximal path-connected components of \mathfrak{M} define the leaves of a foliation \mathcal{F} of \mathfrak{M} . All matchbox manifolds are assumed to be smooth, with a leafwise Riemannian metric, and metric $d_{\mathfrak{M}}$ on \mathfrak{M} .

2.1. Metric properties and regular covers. We introduce some local metric considerations for a matchbox manifold, and give the definition of a *regular covering* of \mathfrak{M} .

For $x \in \mathfrak{M}$ and $\epsilon > 0$, let $D_{\mathfrak{M}}(x, \epsilon) = \{y \in \mathfrak{M} \mid d_{\mathfrak{M}}(x, y) \leq \epsilon\}$ be the closed ϵ -ball about x in \mathfrak{M} , and $B_{\mathfrak{M}}(x, \epsilon) = \{y \in \mathfrak{M} \mid d_{\mathfrak{M}}(x, y) < \epsilon\}$ the open ϵ -ball about x .

Similarly, for $w \in \mathfrak{X}$ and $\epsilon > 0$, let $D_{\mathfrak{X}}(w, \epsilon) = \{w' \in \mathfrak{X} \mid d_{\mathfrak{X}}(w, w') \leq \epsilon\}$ be the closed ϵ -ball about w in \mathfrak{X} , and $B_{\mathfrak{X}}(w, \epsilon) = \{w' \in \mathfrak{X} \mid d_{\mathfrak{X}}(w, w') < \epsilon\}$ the open ϵ -ball about w .

Each leaf $L \subset \mathfrak{M}$ has a complete path-length metric, induced from the leafwise Riemannian metric:

$$d_{\mathcal{F}}(x, y) = \inf \{ \|\gamma\| \mid \gamma: [0, 1] \rightarrow L \text{ is piecewise } C^1, \gamma(0) = x, \gamma(1) = y, \gamma(t) \in L \quad \forall 0 \leq t \leq 1 \}$$

where $\|\gamma\|$ denotes the path-length of the piecewise C^1 -curve $\gamma(t)$. If $x, y \in \mathfrak{M}$ are not on the same leaf, then set $d_{\mathcal{F}}(x, y) = \infty$.

For each $x \in \mathfrak{M}$ and $r > 0$, let $D_{\mathcal{F}}(x, r) = \{y \in L_x \mid d_{\mathcal{F}}(x, y) \leq r\}$ be the closed leafwise ball.

For each $x \in \mathfrak{M}$, the Gauss Lemma implies that there exists $\lambda_x > 0$ such that $D_{\mathcal{F}}(x, \lambda_x)$ is a *strongly convex* subset for the metric $d_{\mathcal{F}}$. That is, for any pair of points $y, y' \in D_{\mathcal{F}}(x, \lambda_x)$ there is a unique shortest geodesic segment in L_x joining y and y' and contained in $D_{\mathcal{F}}(x, \lambda_x)$ (cf. [20, Chapter 3, Proposition 4.2], or [36, Theorem 9.9]). Then for all $0 < \lambda < \lambda_x$ the disk $D_{\mathcal{F}}(x, \lambda)$ is also strongly convex. As \mathfrak{M} is compact and the leafwise metrics have uniformly bounded geometry, we obtain:

LEMMA 2.5. *There exists $\lambda_{\mathcal{F}} > 0$ such that for all $x \in \mathfrak{M}$, $D_{\mathcal{F}}(x, \lambda_{\mathcal{F}})$ is strongly convex.*

The following proposition summarizes results in [17, Sections 2.1 - 2.2].

PROPOSITION 2.6. *For a smooth foliated space \mathfrak{M} , given $\epsilon_{\mathfrak{M}} > 0$, there exist constants $\lambda_{\mathcal{F}} > 0$ and $0 < \delta_{\mathcal{U}}^{\mathcal{F}} < \lambda_{\mathcal{F}}/5$, and a covering of \mathfrak{M} by foliation charts $\{\varphi_i: \bar{U}_i \rightarrow [-1, 1]^n \times \mathfrak{I}_i \mid 1 \leq i \leq \nu\}$ with the following properties: For each $1 \leq i \leq \nu$, let $\pi_i = \pi_{x_i}: \bar{U}_i \rightarrow \mathfrak{I}_i$ be the projection, then*

- (1) *Interior:* $U_i \equiv \text{int}(\bar{U}_i) = \varphi_i^{-1}((-1, 1)^n \times \mathfrak{I}_i)$
- (2) *Locality:* for $x_i \equiv \varphi_i^{-1}(w_i, 0) \in \mathfrak{M}$, $\bar{U}_i \subset B_{\mathfrak{M}}(x_i, \epsilon_{\mathfrak{M}})$ where $w_i = \pi_i(x_i)$.

For $z \in \bar{U}_i$, the plaque of the chart φ_i through z is denoted by $\mathcal{P}_i(z) = \mathcal{P}_i(\pi_i(z)) \subset \bar{U}_i$.

- (3) *Convexity:* the plaques of φ_i are strongly convex subsets for the leafwise metric.
- (4) *Uniformity:* for $w \in \mathfrak{I}_i$ let $x_w = \varphi_{x_i}^{-1}(0, w)$, then

$$D_{\mathcal{F}}(x_w, \delta_{\mathcal{U}}^{\mathcal{F}}/2) \subset \mathcal{P}_i(w) \subset D_{\mathcal{F}}(x_w, \delta_{\mathcal{U}}^{\mathcal{F}})$$

- (5) *The projection $\pi_i(U_i \cap U_j) = \mathfrak{I}_{i,j} \subset \mathfrak{I}_i$ is a clopen subset for all $1 \leq i, j \leq \nu$.*

A regular foliated covering of \mathfrak{M} is one that satisfies these conditions.

We assume in the following that a regular foliated covering of \mathfrak{M} as in Proposition 2.6 has been chosen. Let $\mathcal{U} = \{U_1, \dots, U_{\nu}\}$ denote the corresponding open covering of \mathfrak{M} . We can assume without loss of generality, that the spaces \mathfrak{I}_i form a *disjoint clopen covering* of \mathfrak{X} , so that $\mathfrak{X} = \mathfrak{I}_1 \dot{\cup} \dots \dot{\cup} \mathfrak{I}_{\nu}$.

Let $\epsilon_{\mathcal{U}} > 0$ be a Lebesgue number for \mathcal{U} . That is, given any $z \in \mathfrak{M}$ there exists some index $1 \leq i_z \leq \nu$ such that the open metric ball $B_{\mathfrak{M}}(z, \epsilon_{\mathcal{U}}) \subset U_{i_z}$. Also, introduce a form of “leafwise Lebesgue number”, defined by

- (3) $\epsilon_{\mathcal{U}}^{\mathcal{F}} = \min \{ \epsilon_{\mathcal{U}}^{\mathcal{F}}(y) \mid \forall y \in \mathfrak{M} \}$, $\epsilon_{\mathcal{U}}^{\mathcal{F}}(y) = \max \{ \epsilon \mid D_{\mathcal{F}}(y, \epsilon) \subset D_{\mathfrak{M}}(y, \epsilon_{\mathcal{U}}/4) \}$.

Thus, for all $y \in \mathfrak{M}$, $D_{\mathcal{F}}(y, \epsilon_{\mathcal{U}}^{\mathcal{F}}) \subset D_{\mathfrak{M}}(y, \epsilon_{\mathcal{U}}/4)$. Note that for all $r > 0$ and $z' \in D_{\mathcal{F}}(z, \epsilon_{\mathcal{U}}^{\mathcal{F}})$, the triangle inequality implies that $D_{\mathfrak{M}}(z', r) \subset D_{\mathfrak{M}}(z, r + \epsilon_{\mathcal{U}}/4)$. Note that for $y \in \mathfrak{M}$, we have

$$D_{\mathcal{F}}(y, \epsilon_{\mathcal{U}}^{\mathcal{F}}) \subset D_{\mathfrak{M}}(y, \epsilon_{\mathcal{U}}/4) \subset D_{\mathfrak{M}}(y, \epsilon_{\mathcal{U}}) \subset U_{i_y}, \quad \mathcal{P}_{i_y}(\pi_{i_y}(y)) \subset D_{\mathcal{F}}(y, \delta_{\mathcal{U}}^{\mathcal{F}})$$

which imply that $\epsilon_{\mathcal{U}}^{\mathcal{F}} < \delta_{\mathcal{U}}^{\mathcal{F}}$, and hence $\delta_{\mathcal{U}}^{\mathcal{F}}/\epsilon_{\mathcal{U}}^{\mathcal{F}} > 1$. This last inequality will be used in Section 5.

For $1 \leq i \leq \nu$, let $\lambda_i: \bar{U}_i \rightarrow [-1, 1]^n$ be the projection, so that for each $z \in U_i$ the restriction $\lambda_i: \mathcal{P}_i(z) \rightarrow [-1, 1]^n$ is a smooth coordinate system on the plaque $\mathcal{P}_i(z)$.

For each $1 \leq i \leq \nu$, the set $\mathcal{T}_i = \varphi_i^{-1}(0, \mathfrak{T}_i)$ is a compact transversal to \mathcal{F} . Without loss of generality, we can assume that the transversals $\{\mathcal{T}_1, \dots, \mathcal{T}_\nu\}$ are pairwise disjoint in \mathfrak{M} , so there exists a constant $0 < \delta_{\mathcal{T}} < \delta_{\mathcal{U}}^{\mathcal{F}}$ such that

$$(4) \quad d_{\mathcal{F}}(x, y) \geq \delta_{\mathcal{T}} \quad \text{for } x \neq y, x \in \mathcal{T}_i, y \in \mathcal{T}_j, 1 \leq i, j \leq \nu.$$

In particular, this implies that the centers of disjoint plaques on the same leaf are separated by distance at least $\delta_{\mathcal{T}}$. Also, define sections

$$(5) \quad \tau_i: \mathfrak{T}_i \rightarrow \bar{U}_i, \quad \text{defined by } \tau_i(\xi) = \varphi_i^{-1}(0, \xi), \quad \text{so that } \pi_i(\tau_i(\xi)) = \xi.$$

Then $\mathcal{T}_i = \mathcal{T}_{x_i}$ is the image of τ_i and we let $\mathcal{T} = \mathcal{T}_1 \cup \dots \cup \mathcal{T}_\nu \subset \mathfrak{M}$ denote their disjoint union, and $\tau: \mathfrak{X} \rightarrow \mathcal{T}$ is defined by the union of the maps τ_i .

Define the metric $d_{\mathfrak{X}}$ on \mathfrak{X} via the restriction of $d_{\mathfrak{M}}$ to \mathcal{T} , and use the map τ to pull it back to \mathfrak{X} .

2.2. Local estimates. The local projections $\pi_i: \bar{U}_i \rightarrow \mathfrak{T}_i$ and sections $\tau_i: \mathfrak{T}_i \rightarrow \bar{U}_i$ are continuous maps of compact spaces, so admit uniform metric estimates as shown in [17].

LEMMA 2.7. *There exists a modulus of continuity function ρ_{π} which is continuous and increasing, such that:*

$$(6) \quad \forall 1 \leq i \leq \nu \text{ and } x, y \in \bar{U}_i, \quad d_{\mathfrak{M}}(x, y) < \rho_{\pi}(\epsilon) \implies d_{\mathfrak{X}}(\pi_i(x), \pi_i(y)) < \epsilon.$$

Proof. Set $\rho_{\pi}(\epsilon) = \min \{ \epsilon, \min \{ d_{\mathfrak{M}}(x, y) \mid 1 \leq i \leq \nu, x, y \in \bar{U}_i, d_{\mathfrak{X}}(\pi_i(x), \pi_i(y)) \geq \epsilon \} \}$. \square

LEMMA 2.8. *There exists a modulus of continuity function ρ_{τ} which is continuous and increasing, such that:*

$$(7) \quad \forall 1 \leq i \leq \nu \text{ and } w, w' \in \mathfrak{T}_i, \quad d_{\mathfrak{X}}(w, w') < \rho_{\tau}(\epsilon) \implies d_{\mathfrak{M}}(\tau_i(w), \tau_i(w')) < \epsilon.$$

Proof. Set $\rho_{\tau}(\epsilon) = \min \{ \epsilon, \min \{ d_{\mathfrak{X}}(w, w') \mid 1 \leq i \leq \nu, w, w' \in \mathfrak{T}_i, d_{\mathfrak{M}}(\tau_i(w), \tau_i(w')) \geq \epsilon \} \}$. \square

Next, for each $1 \leq i \leq \nu$, consider the projection map $\pi_i: \bar{U}_i \rightarrow \mathfrak{T}_i$, and define

$$\epsilon_i^{\mathcal{T}} = \max \{ \epsilon \mid \forall x \in \bar{U}_i \text{ such that } D_{\mathfrak{M}}(x, \epsilon_{\mathcal{U}}/2) \subset \bar{U}_i, \text{ then } D_{\mathfrak{X}}(\pi_i(x), \epsilon) \subset \pi_i(D_{\mathfrak{M}}(x, \epsilon_{\mathcal{U}}/2)) \}.$$

The assumption that $D_{\mathfrak{M}}(x, \epsilon_{\mathcal{U}}/2) \subset \bar{U}_i$ implies that x has distance at least $\epsilon_{\mathcal{U}}/2$ from the exterior of \bar{U}_i . Then the projection of the closed ball $D_{\mathfrak{M}}(x, \epsilon_{\mathcal{U}}/2)$ to the transversal \mathfrak{T}_i contains an open neighborhood of $\pi_i(x)$ by the continuity of projections, and $\epsilon_i^{\mathcal{T}}$ is the distance from this center $\pi_i(x)$ to the exterior of the projected ball. Then introduce $\epsilon_{\mathcal{U}}^{\mathcal{T}} \geq \rho_{\tau}(\epsilon_{\mathcal{U}}/2)$ given by

$$(8) \quad \epsilon_{\mathcal{U}}^{\mathcal{T}} = \min \{ \epsilon_i^{\mathcal{T}} \mid \forall 1 \leq i \leq \nu \}.$$

Finally, we give a plaque-wise estimate on the Riemannian metrics. The assumption that the leafwise Riemannian metric on \mathfrak{M} is continuous, means that for each coordinate chart $\varphi_i: \bar{U}_i \rightarrow [-1, 1]^n \times \mathfrak{T}_i$, the push-forwards of the Riemannian metric to the slices $[-1, 1]^n \times \{w\}$ vary continuously with $w \in \mathfrak{T}_i$. Let $\|\cdot\|_w$ denote the norm defined on the tangent bundle $T([-1, 1]^n \times \{w\})$, and denote a tangent vector by (\vec{v}, ξ, w) , for $\vec{v} \in \mathbb{R}^n$ and $\xi \in (-1, 1)^n$. Let $\|\cdot\|$ denote the Euclidean norm on \mathbb{R}^n . Then by the compactness of \bar{U}_i and the continuity of the metric, for $\epsilon > 0$, there exists $\delta_{\mathcal{U}}(\epsilon) > 0$ such that the following holds, for each $1 \leq i \leq \nu$:

$$(9) \quad \sup \left\{ \left| \max \left\{ \frac{\|\vec{v}\|_w}{\|\vec{v}\|_{w'}}, \frac{\|\vec{v}\|_{w'}}{\|\vec{v}\|_w} \right\} - 1 \right| \mid d_{\mathfrak{X}}(w, w') \leq \delta_{\mathcal{U}}(\epsilon), \vec{v} \in \mathbb{R}^n, \|\vec{v}\| = 1 \right\} < \epsilon.$$

2.3. Foliated maps. A map $f: \mathfrak{M} \rightarrow \mathfrak{M}'$ between foliated spaces is said to be a *foliated map* if the image of each leaf of \mathcal{F} is contained in a leaf of \mathcal{F}' . If \mathfrak{M}' is a matchbox manifold, then each leaf of \mathcal{F} is path connected, so its image is path connected, hence must be contained in a leaf of \mathcal{F}' .

A *leafwise path* is a continuous map $\gamma: [0, 1] \rightarrow \mathfrak{M}$ such that there is a leaf L of \mathcal{F} for which $\gamma(t) \in L$ for all $0 \leq t \leq 1$. If \mathfrak{M} is a matchbox manifold and $\gamma: [0, 1] \rightarrow \mathfrak{M}$ is continuous, then γ is a leafwise path. This yields:

LEMMA 2.9. *Let \mathfrak{M} and \mathfrak{M}' be matchbox manifolds, and $h: \mathfrak{M} \rightarrow \mathfrak{M}'$ a continuous map. Then h maps the leaves of \mathcal{F} to leaves of \mathcal{F}' . In particular, any homeomorphism $h: \mathfrak{M} \rightarrow \mathfrak{M}$ of a matchbox manifold is a foliated map. \square*

3. HOLONOMY OF FOLIATED SPACES

The holonomy pseudogroup of a smooth foliated manifold (M, \mathcal{F}) generalizes the induced discrete dynamical system associated to a section of a flow. A standard construction in foliation theory (see [15], [16, Chapter 2] or [35]) associates to a leafwise path γ a holonomy map h_γ . The collection of all such maps with initial and ending points on a transversal to \mathcal{F} defines the holonomy pseudogroup $\mathcal{G}_\mathcal{F}$ for a matchbox manifold $(\mathfrak{M}, \mathcal{F})$. We recall below the ideas and notations, as required in the proofs of our main theorems, especially the delicate issues of domains which must be considered. See [17, 18] for a complete discussion of the ideas of this section and related technical results.

3.1. Holonomy pseudogroup. Let $\mathcal{U} = \{U_1, \dots, U_\nu\}$ be a regular foliated covering of \mathfrak{M} as in Proposition 2.6. A pair of indices (i, j) , for $1 \leq i, j \leq \nu$, is said to be *admissible* if $U_i \cap U_j \neq \emptyset$.

For (i, j) admissible, set $\mathfrak{X}_{i,j} = \pi_i(U_i \cap U_j) \subset \mathfrak{X}_i$. The regular foliated covering assumption implies that plaques in admissible charts are either disjoint, or have connected intersection. This implies that there is a well-defined transverse change of coordinates homeomorphism $h_{i,j}: \mathfrak{X}_{i,j} \rightarrow \mathfrak{X}_{j,i}$ with domain $\text{Dom}(h_{i,j}) = \mathfrak{X}_{i,j}$ and range $R(h_{i,j}) = \text{Dom}(h_{j,i}) = \mathfrak{X}_{j,i}$. By definition they satisfy $h_{i,i} = \text{Id}$, $h_{i,j}^{-1} = h_{j,i}$, and if $U_i \cap U_j \cap U_k \neq \emptyset$ then $h_{k,j} \circ h_{j,i} = h_{k,i}$ on their common domain of definition. Note that the domain and range of $h_{i,j}$ are clopen subsets of \mathfrak{X} by Proposition 2.6.(5).

Recall that for each $1 \leq i \leq \nu$, $\tau_i: \mathfrak{X}_i \rightarrow \mathcal{T}_i$ denotes the transverse section for the chart U_i , and $\mathcal{T} = \mathcal{T}_1 \cup \dots \cup \mathcal{T}_\nu \subset \mathfrak{M}$ denotes their disjoint union. Then $\pi: \mathcal{T} \rightarrow \mathfrak{X}$ is the coordinate projection restricted to \mathcal{T} , which is a homeomorphism, and $\tau: \mathfrak{X} \rightarrow \mathcal{T}$ denotes its inverse.

The *holonomy pseudogroup* $\mathcal{G}_\mathcal{F}$ of \mathcal{F} is the topological pseudogroup modeled on \mathfrak{X} generated by the elements of $\mathcal{G}_\mathcal{F}^{(1)} = \{h_{j,i} \mid (i, j) \text{ admissible}\}$. A sequence $\mathcal{I} = (i_0, i_1, \dots, i_\alpha)$ is *admissible* if each pair $(i_{\ell-1}, i_\ell)$ is admissible for $1 \leq \ell \leq \alpha$, and the composition

$$(10) \quad h_\mathcal{I} = h_{i_\alpha, i_{\alpha-1}} \circ \dots \circ h_{i_1, i_0}$$

has non-empty domain $\text{Dom}(h_\mathcal{I})$, which is defined to be the maximal *clopen* subset of \mathfrak{X}_{i_0} for which the compositions are defined. Given any open subset $U \subset \text{Dom}(h_\mathcal{I})$, we obtain a new element $h_\mathcal{I}|U \in \mathcal{G}_\mathcal{F}$ by restriction. Introduce

$$(11) \quad \mathcal{G}_\mathcal{F}^* = \{h_\mathcal{I}|U \mid \mathcal{I} \text{ admissible and } U \subset \text{Dom}(h_\mathcal{I})\} \subset \mathcal{G}_\mathcal{F} .$$

The range of $g = h_\mathcal{I}|U$ is the open set $R(g) = h_\mathcal{I}(U) \subset \mathfrak{X}_{i_\alpha} \subset \mathfrak{X}$. Note that each map $g \in \mathcal{G}_\mathcal{F}^*$ admits a continuous extension $\bar{g}: \overline{\text{Dom}(g)} = \bar{U} \rightarrow \mathfrak{X}_{i_\alpha}$ as $\text{Dom}(h_\mathcal{I})$ is a clopen set for each \mathcal{I} .

3.2. Plaque-chains. Let $\mathcal{I} = (i_0, i_1, \dots, i_\alpha)$ be an admissible sequence. For each $1 \leq \ell \leq \alpha$, set $\mathcal{I}_\ell = (i_0, i_1, \dots, i_\ell)$, and let $h_{\mathcal{I}_\ell}$ denote the corresponding holonomy map. For $\ell = 0$, let $\mathcal{I}_0 = (i_0, i_0)$. Note that $h_{\mathcal{I}_\alpha} = h_\mathcal{I}$ and $h_{\mathcal{I}_0} = \text{Id}: \mathfrak{X}_0 \rightarrow \mathfrak{X}_0$.

Given $w \in \text{Dom}(h_\mathcal{I})$, let $x = \tau_{i_0}(w) \in L_w$. For each $0 \leq \ell \leq \alpha$, set $w_\ell = h_{\mathcal{I}_\ell}(w)$ and $x_\ell = \tau_{i_\ell}(w_\ell)$. Recall that $\mathcal{P}_{i_\ell}(x_\ell) = \mathcal{P}_{i_\ell}(w_\ell)$, where each $\mathcal{P}_{i_\ell}(w_\ell)$ is a strongly convex subset of the leaf L_w in the leafwise metric $d_\mathcal{F}$. Introduce the *plaque-chain*

$$(12) \quad \mathcal{P}_\mathcal{I}(w) = \{\mathcal{P}_{i_0}(w_0), \mathcal{P}_{i_1}(w_1), \dots, \mathcal{P}_{i_\alpha}(w_\alpha)\} .$$

Adopt the notation $\mathcal{P}_{\mathcal{I}}(x) \equiv \mathcal{P}_{\mathcal{I}}(w)$. Intuitively, a plaque-chain $\mathcal{P}_{\mathcal{I}}(x)$ is a sequence of successively overlapping convex “tiles” in L_w starting at $x = \tau_{i_0}(w)$, ending at $y = x_\alpha = \tau_{i_\alpha}(w_\alpha)$, where each $\mathcal{P}_{i_\ell}(x_\ell)$ is “centered” on the point $x_\ell = \tau_{i_\ell}(w_\ell)$.

3.3. Leafwise paths to plaque-chains. Let $\gamma: [0, 1] \rightarrow \mathfrak{M}$ be a path. Set $x_0 = \gamma(0) \in U_{i_0}$, $w = \pi(x_0)$ and $x = \tau(w) \in \mathcal{T}_{i_0}$. Let \mathcal{I} be an admissible sequence with $w \in \text{Dom}(h_{\mathcal{I}})$. We say that (\mathcal{I}, w) covers γ , if there is a partition $0 = s_0 < s_1 < \dots < s_\alpha = 1$ such that $\mathcal{P}_{\mathcal{I}}(w)$ satisfies

$$(13) \quad \gamma([s_\ell, s_{\ell+1}]) \subset \mathcal{P}_{i_\ell}(\xi_\ell), \quad 0 \leq \ell < \alpha, \quad \text{and} \quad \gamma(1) \in \mathcal{P}_{i_\alpha}(\xi_\alpha).$$

For a path γ , we construct an admissible sequence $\mathcal{I} = (i_0, i_1, \dots, i_\alpha)$ with $w \in \text{Dom}(h_{\mathcal{I}})$ so that (\mathcal{I}, w) covers γ , and has “uniform domains”. Inductively choose a partition of the interval $[0, 1]$, say $0 = s_0 < s_1 < \dots < s_\alpha = 1$, such that for each $0 \leq \ell \leq \alpha$,

$$\gamma([s_\ell, s_{\ell+1}]) \subset D_{\mathcal{F}}(x_\ell, \epsilon_{\mathcal{U}}^{\mathcal{F}}) \quad , \quad x_\ell = \gamma(s_\ell).$$

As a notational convenience, we have let $s_{\alpha+1} = s_\alpha$, so that $\gamma([s_\alpha, s_{\alpha+1}]) = x_\alpha$. Choose $s_{\ell+1}$ to be the largest value of $s_\ell < s \leq 1$ such that $d_{\mathcal{F}}(\gamma(s_\ell), \gamma(t)) \leq \epsilon_{\mathcal{U}}^{\mathcal{F}}$ for all $s_\ell \leq t \leq s$, then $\alpha \leq \|\gamma\|/\epsilon_{\mathcal{U}}^{\mathcal{F}}$.

For each $0 \leq \ell \leq \alpha$, choose an index $1 \leq i_\ell \leq \nu$ so that $B_{\mathfrak{M}}(x_\ell, \epsilon_{\mathcal{U}}) \subset U_{i_\ell}$. Note that, for all $s_\ell \leq t \leq s_{\ell+1}$, $B_{\mathfrak{M}}(\gamma(t), \epsilon_{\mathcal{U}}/2) \subset U_{i_\ell}$, so that $x_{\ell+1} \in U_{i_\ell} \cap U_{i_{\ell+1}}$. It follows that $\mathcal{I}_\gamma = (i_0, i_1, \dots, i_\alpha)$ is an admissible sequence. Set $h_\gamma = h_{\mathcal{I}_\gamma}$ and note that $h_\gamma(w) = w'$.

The construction of the admissible sequence \mathcal{I}_γ above has the important property, that $h_{\mathcal{I}_\gamma}$ is the composition of generators of $\mathcal{G}_{\mathcal{F}}^*$ which each have a uniform lower bound estimate $\epsilon_{\mathcal{U}}^{\mathcal{F}}$ on the radii of the metric balls centered at the orbit and which are contained in their domains. To see this, let $0 \leq \ell < \alpha$, and note that $x_{\ell+1} \in D_{\mathcal{F}}(x_{\ell+1}, \epsilon_{\mathcal{U}}^{\mathcal{F}})$ implies that for some $s_\ell < s'_{\ell+1} < s_{\ell+1}$, we have that $\gamma([s'_{\ell+1}, s_{\ell+1}]) \subset D_{\mathcal{F}}(x_{\ell+1}, \epsilon_{\mathcal{U}}^{\mathcal{F}})$. Hence,

$$(14) \quad B_{\mathfrak{M}}(\gamma(t), \epsilon_{\mathcal{U}}/2) \subset U_{i_\ell} \cap U_{i_{\ell+1}}, \quad \text{for all } s'_{\ell+1} \leq t \leq s_{\ell+1}.$$

Then for all $s'_{\ell+1} \leq t \leq s_{\ell+1}$, the uniform estimate defining $\epsilon_{\mathcal{U}}^{\mathcal{F}} > 0$ in (8) implies that

$$(15) \quad B_{\mathfrak{X}}(\pi_{i_\ell}(\gamma(t)), \epsilon_{\mathcal{U}}^{\mathcal{F}}) \subset \mathfrak{X}_{i_\ell, i_{\ell+1}} \quad \text{and} \quad B_{\mathfrak{X}}(\pi_{i_{\ell+1}}(\gamma(t)), \epsilon_{\mathcal{U}}^{\mathcal{F}}) \subset \mathfrak{X}_{i_{\ell+1}, i_\ell}.$$

For the admissible sequence $\mathcal{I}_\gamma = (i_0, i_1, \dots, i_\alpha)$, recall that $x_\ell = \gamma(s_\ell)$ and $w_\ell = \pi_{i_\ell}(x_\ell)$. By definition (10) of $h_{\mathcal{I}_\gamma}$, the condition (15) implies that $D_{\mathfrak{X}}(w_\ell, \epsilon_{\mathcal{U}}^{\mathcal{F}}) \subset \text{Dom}(h_\ell)$ as was claimed.

DEFINITION 3.1. *Let $\gamma: [0, 1] \rightarrow \mathfrak{M}$ be a path starting at $x = \tau(w)$ and ending at $y = \tau(w')$. Then a good plaque-chain covering of γ is the plaque-chain, starting at x , associated to an admissible sequence $\mathcal{I}_\gamma = (i_0, i_1, \dots, i_\alpha)$ as constructed above with $\alpha \leq \|\gamma\|/\epsilon_{\mathcal{U}}^{\mathcal{F}}$, and which satisfies (15).*

3.4. Equivalence of leafwise holonomy. We give several criteria for when two holonomy maps must agree. Consider first the case where we are given a path γ with $x = \gamma(0)$ and $y = \gamma(1)$. Let $\mathcal{I} = (i_0, i_1, \dots, i_\alpha)$ and $\mathcal{J} = (j_0, j_1, \dots, j_\beta)$ be admissible sequences such that both (\mathcal{I}, ξ_0) and (\mathcal{J}, ξ'_0) cover the path γ . Then

$$x \in \mathcal{P}_{i_0}(\xi_0) \cap \mathcal{P}_{j_0}(\xi'_0), \quad y \in \mathcal{P}_{i_\alpha}(\xi_\alpha) \cap \mathcal{P}_{j_\beta}(\xi'_\beta).$$

Thus both (i_0, j_0) and (i_α, j_β) are admissible, and $\xi'_0 = h_{j_0, i_0}(\xi_0)$, $\xi_\alpha = h_{i_\alpha, j_\beta}(\xi'_\beta)$. The proof of the following is intuitively clear, with details in [17].

LEMMA 3.2. *The point ξ_0 is contained in the domains of both $h_{\mathcal{I}}$ and $h_{i_\alpha, j_\beta} \circ h_{\mathcal{J}} \circ h_{j_0, i_0}$, and the intersection of their domains is a non-empty clopen subset on which the maps agree.* \square

Next, consider paths $\gamma, \gamma': [0, 1] \rightarrow \mathfrak{M}$ with $x = \gamma(0) = \gamma'(0)$ and $y = \gamma(1) = \gamma'(1)$. Suppose that γ and γ' are homotopic relative endpoints. That is, assume there exists a continuous map $H: [0, 1] \times [0, 1] \rightarrow \mathfrak{M}$ with

$$H(0, t) = \gamma(t), \quad H(1, t) = \gamma'(t), \quad H(s, 0) = x \quad \text{and} \quad H(s, 1) = y \quad \text{for all } 0 \leq s \leq 1$$

Then there exists partitions $0 = s_0 < s_1 < \dots < s_\beta = 1$ and $0 = t_0 < t_1 < \dots < t_\alpha = 1$ such that for each pair of indices $0 \leq j < \beta$ and $0 \leq k < \alpha$, there is an index $1 \leq i(j, k) \leq \nu$ such that

$$H([s_j, s_{j+1}] \times [t_k, t_{k+1}]) \subset D_{\mathcal{F}}(H(s_j, t_k), \epsilon_{\mathcal{U}}^{\mathcal{F}}) \subset U_{i(j,k)}$$

Then proceeding using the methods above and a standard induction argument, we obtain:

LEMMA 3.3. *Let $\gamma, \gamma': [0, 1] \rightarrow \mathfrak{M}$ be paths with $x = \gamma(0) = \gamma'(0)$ and $y = \gamma(1) = \gamma'(1)$, and suppose they are homotopic relative endpoints. Then the induced holonomy maps h_γ and $h_{\gamma'}$ agree on an open neighborhood of $\xi_0 = \pi_{i_0}(x)$.*

The strongly convex property of the plaques and the techniques above also yield the following result.

LEMMA 3.4. *Let $\gamma, \gamma': [0, 1] \rightarrow \mathfrak{M}$ be paths. Suppose that $x = \gamma(0), x' = \gamma'(0) \in U_i$ and $y = \gamma(1), y' = \gamma'(1) \in U_j$. If $d_{\mathfrak{M}}(\gamma(t), \gamma'(t)) \leq \epsilon_{\mathcal{U}}/4$ for all $0 \leq t \leq 1$, then the induced holonomy maps $h_\gamma, h_{\gamma'}$ agree on their common domain $\text{Dom}(h_\gamma) \cap \text{Dom}(h_{\gamma'}) \subset \mathfrak{X}_i$.*

3.5. Plaque-chains to leafwise paths. Recall another basic construction, which associates to an admissible sequence $\mathcal{I} = (i_0, i_1, \dots, i_\alpha)$ and $\xi_0 \in \text{Dom}(h_{\mathcal{I}_0})$ a leafwise path. For each $0 < \ell \leq \alpha$, choose $z_\ell \in \mathcal{P}_{\ell-1}(\xi_{\ell-1}) \cap \mathcal{P}_\ell(\xi_\ell)$. Let $\gamma_\ell: [(\ell-1)/\alpha, \ell/\alpha] \rightarrow L_{x_0}$ be the leafwise piecewise geodesic segment from $x_{\ell-1}$ to z_ℓ to x_ℓ . Define the path $\gamma_{\mathcal{I}}: [0, 1] \rightarrow L_{x_0}$ from x_0 to x_α to be the concatenation of these paths. Note that $\|\gamma_{\mathcal{I}}\| \leq 2\alpha\delta_{\mathcal{U}}^{\mathcal{F}}$ by Proposition 2.6.4.

Let (\mathcal{I}', ξ_0) be a plaque-chain covering of $\gamma_{\mathcal{I}}$, then for $U \subset \text{Dom}(h_{\mathcal{I}}) \cap \text{Dom}(h_{\gamma_{\mathcal{I}}})$ we have $h_{\mathcal{I}}|_U = h_{\gamma_{\mathcal{I}}}|_U$ by Lemma 3.2. Of course, we may take $\mathcal{I}' = \mathcal{I}$ and then the conclusion is tautologous, or may take \mathcal{I}' to define a good plaque-chain covering as in Definition 3.1.

The study of the dynamics of the pseudogroup $\mathcal{G}_{\mathcal{F}}$ acting on \mathfrak{X} differs from the case of a group Γ acting on \mathfrak{X} , in that for a group action each $\gamma \in \Gamma$ defines a homeomorphism $h_\gamma: \mathfrak{X} \rightarrow \mathfrak{X}$. For a pseudogroup action, given $g \in \mathcal{G}_{\mathcal{F}}$ and $w \in \text{Dom}(g)$, there is some clopen neighborhood $w \in U \subset \text{Dom}(g)$ for which $g|_U = h_{\mathcal{I}}|_U$ where \mathcal{I} is admissible sequence with $w \in \text{Dom}(h_{\mathcal{I}})$. By the definition of a pseudogroup, every $g \in \mathcal{G}_{\mathcal{F}}$ is the ‘‘union’’ of such maps in $\mathcal{G}_{\mathcal{F}}^*$. For our applications to the structure of the dynamics of \mathcal{F} , it is very useful to estimate the maximal domain for the holonomy map defined by an admissible sequence. Recall that $\epsilon_{\mathcal{U}}^{\mathcal{F}}$ was defined by (8).

PROPOSITION 3.5. *For each $0 < \epsilon \leq \epsilon_{\mathcal{U}}^{\mathcal{F}}$ and integer $\alpha > 0$, there exists $0 < \delta(\epsilon, \alpha) \leq \epsilon$ so that for each admissible sequence \mathcal{I} with length at most α , $w \in \text{Dom}(h_{\mathcal{I}})$ and $w' = h_{\mathcal{I}}(w)$, there exists an admissible sequence \mathcal{I}' with $D_{\mathfrak{X}}(w_0, \delta(\epsilon, \alpha)) \subset \text{Dom}(h_{\mathcal{I}'})$ such that $h_{\mathcal{I}'}$ coincides with $h_{\mathcal{I}}$ on the intersection of their domains, and*

$$(16) \quad h_{\mathcal{I}'}(D_{\mathfrak{X}}(w_0, \delta(\epsilon, \alpha))) \subset D_{\mathfrak{X}}(w', \epsilon)$$

Proof. Let $\gamma_{\mathcal{I}}$ be the piecewise geodesic from $x = \tau(w) \in \mathcal{T}$ to $y = \tau(w')$ as constructed above. Then by the method of Section 3.3, there is an admissible sequence \mathcal{I}' which defines a good plaque-chain covering (\mathcal{I}', w) of $\gamma_{\mathcal{I}}$ satisfying the condition (15) on domains. The result then follows using induction, as in the proof of [18, Proposition 5.7]. \square

3.6. Dynamics of matchbox manifolds. We recall two definitions from topological dynamics, that of *equicontinuous* and *expansive* dynamics, as adapted to actions of pseudogroups.

DEFINITION 3.6. *Say that the action of the pseudogroup $\mathcal{G}_{\mathcal{F}}$ on \mathfrak{X} is expansive, or more properly that it is ϵ -expansive, if there exists $\epsilon > 0$ such that for all $w, w' \in \mathfrak{X}$, there exists $g \in \mathcal{G}_{\mathcal{F}}^*$ with $w, w' \in D(g)$ such that $d_{\mathfrak{X}}(g(w), g(w')) \geq \epsilon$.*

DEFINITION 3.7. *Say that the action of the pseudogroup $\mathcal{G}_{\mathcal{F}}$ on \mathfrak{X} is equicontinuous if for all $\epsilon > 0$, there exists $\delta > 0$ such that for all $g \in \mathcal{G}_{\mathcal{F}}^*$, if $w, w' \in D(g)$ and $d_{\mathfrak{X}}(w, w') < \delta$, then $d_{\mathfrak{X}}(g(w), g(w')) < \epsilon$. Thus, $\mathcal{G}_{\mathcal{F}}^*$ is equicontinuous as a family of local group actions.*

Equicontinuity for $\mathcal{G}_{\mathcal{F}}$ gives *uniform* control over the domains of arbitrary compositions of generators, so that a much stronger conclusion than Proposition 3.5 is true, as shown in [17]:

PROPOSITION 3.8. *Assume the holonomy pseudogroup $\mathcal{G}_{\mathcal{F}}$ of \mathcal{F} is equicontinuous. Then there exists $\delta_{\mathcal{U}}^{\mathcal{T}} > 0$ such that for every leafwise path $\gamma: [0, 1] \rightarrow \mathfrak{M}$, there is a corresponding admissible sequence $\mathcal{I}_{\gamma} = (i_0, i_1, \dots, i_{\alpha})$ so that $B_{\mathfrak{X}}(w_0, \delta_{\mathcal{U}}^{\mathcal{T}}) \subset D(h_{\mathcal{I}_{\gamma}})$, where $x = \gamma(0)$ and $w_0 = \pi_{i_0}(x)$. Moreover, for all $0 < \epsilon \leq \epsilon_{\mathcal{U}}^{\mathcal{T}}$ there exists $0 < \delta \leq \delta_{\mathcal{U}}^{\mathcal{T}}$ independent of the path γ , such that $h_{\mathcal{I}_{\gamma}}(D_{\mathfrak{X}}(w_0, \delta)) \subset D_{\mathfrak{X}}(w', \epsilon)$ where $w' = \pi_{i_{\alpha}}(\gamma(1))$.*

It is possible that these are the only two possibilities for a minimal action of a pseudogroup on a Cantor space. A proof of the following would generalize results in [8, 9, 23]:

CONJECTURE 3.9. *Let $\mathcal{G}_{\mathcal{F}}$ act minimally on a Cantor space \mathfrak{X} as above. Then either the action is equicontinuous, or it is expansive for some $\epsilon > 0$.*

3.7. Holonomy groupoids. We next consider the groupoid [35] formed by germs of maps in $\mathcal{G}_{\mathcal{F}}$.

Let $U, U', V, V' \subset \mathfrak{X}$ be open subsets with $w \in U \cap U'$. Given homeomorphisms $h: U \rightarrow V$ and $h': U' \rightarrow V'$ with $h(w) = h'(w)$, then h and h' have the same *germ at w* , and write $h \sim_w h'$, if there exists an open neighborhood $W \subset U \cap U'$ such that $h|_W = h'|_W$. Note that \sim_w defines an equivalence relation.

DEFINITION 3.10. *The germ of h at w is the equivalence class $[h]_w$ under the relation \sim_w . The map $h: U \rightarrow V$ is called a representative of $[h]_w$. The point w is called the source of $[h]_w$ and denoted $s([h]_w)$, while $w' = h(w)$ is called the range of $[h]_w$ and denoted $r([h]_w)$.*

The *holonomy groupoid* $\Gamma_{\mathcal{F}}$ is the collection of all germs $[h]_w$ for $h \in \mathcal{G}_{\mathcal{F}}$ and $w \in \text{Dom}(h)$, equipped with the sheaf topology for maps over \mathfrak{X} . Let $\mathcal{R}_{\mathcal{F}} \subset \mathfrak{X} \times \mathfrak{X}$ denote the equivalence relation on \mathfrak{X} induced by \mathcal{F} , where $(w, w') \in \mathcal{R}_{\mathcal{F}}$ if and only if w, w' correspond to points on the same leaf of \mathcal{F} . The product map $s \times r: \Gamma_{\mathcal{F}} \rightarrow \mathcal{R}_{\mathcal{F}}$ is étale; that is, a local homeomorphism with discrete fibers.

We introduce a convenient notation for elements of $\Gamma_{\mathcal{F}}$. Let $(w, w') \in \mathcal{R}_{\mathcal{F}}$, and let γ denote a path from $x = \tau(w)$ to $y = \tau(w')$. Proposition 3.5 implies that we can choose an admissible sequence \mathcal{I}_{γ} for which the domain of the induced map $h_{\mathcal{I}}$ satisfies condition (16). By Lemma 3.3, the germ $[h_{\mathcal{I}_{\gamma}}]_w$ depends only on the endpoint-fixed homotopy class of γ , so that we may assume γ is a geodesic between x and y . Moreover, by Lemma 3.4 the germ $[h_{\mathcal{I}_{\gamma}}]_w$ is independent of the admissible sequence covering γ . Thus, we introduce the notation $\gamma_w = [h_{\mathcal{I}_{\gamma}}]_w$ for this germ, which depends only on the end-point fixed homotopy class of γ . Given $[h]_w \in \Gamma_{\mathcal{F}}$ by Section 3.2 the germ is induced by a plaque-chain covering a piecewise geodesic between x and y , so all elements of $\Gamma_{\mathcal{F}}$ are represented by γ_w for some path γ from x to y .

These remarks imply there is a well-defined surjective homomorphism, the *holonomy map*,

$$(17) \quad h_{\mathcal{F},x}: \pi_1(L_x, x) \rightarrow \Gamma_w^w \equiv \{[g]_w \in \Gamma_{\mathcal{F}} \mid r([g]_w) = w\}.$$

Moreover, if $y, z \in L$ then the homomorphism $h_{\mathcal{F},y}$ is conjugate (by an element of $\mathcal{G}_{\mathcal{F}}$) to the homomorphism $h_{\mathcal{F},z}$. A leaf L is said to have *non-trivial germinal holonomy* if for some $y \in L$, the homomorphism $h_{\mathcal{F},y}$ is non-trivial. If the homomorphism $h_{\mathcal{F},y}$ is trivial, then we say that L_y is a *leaf without holonomy*. This property depends only on L , and not the choice of $y \in L$.

A matchbox manifold \mathfrak{M} is said to be *without holonomy*, if every leaf is without holonomy.

As an application of these remarks, we obtain:

LEMMA 3.11. *Given a path $\gamma: [0, 1] \rightarrow \mathfrak{M}$ with $x = \gamma(0)$ and $y = \gamma(1)$. Suppose that L_x is a leaf without holonomy. Then there exists a leafwise geodesic segment $\gamma': [0, 1] \rightarrow \mathfrak{M}$ with $x = \gamma'(0)$ and $y = \gamma'(1)$, such that $\|\gamma'\| = d_{\mathcal{F}}(x, y)$, and h_{γ} and $h_{\gamma'}$ agree on an open neighborhood of ξ_0 .*

Proof. The leaf L_x containing x is a complete Riemannian manifold, so there exists a geodesic segment γ' which is length minimizing between x and y . Then the holonomy maps h_{γ} and $h_{\gamma'}$ agree on an open neighborhood of $\xi_0 = \pi_{i_0}(x)$ by the definition of germinal holonomy. \square

3.8. Word and path length. We recall the word length function on $\mathcal{G}_{\mathcal{F}}^*$ and $\Gamma_{\mathcal{F}}$, and the path length function on $\Gamma_{\mathcal{F}}$, and give estimates comparing these notions of length.

For $\alpha \geq 1$, let $\mathcal{G}_{\mathcal{F}}^{(\alpha)}$ be the collection of holonomy homeomorphisms $h_{\mathcal{I}}|U \in \mathcal{G}_{\mathcal{F}}^*$ determined by admissible paths $\mathcal{I} = (i_0, \dots, i_k)$ such that $k \leq \alpha$ and $U \subset \text{Dom}(h_{\mathcal{I}})$. For each α , let $C(\alpha)$ denote the number of admissible sequences of length at most α . As there are at most ν^2 admissible pairs (i, j) , we have the basic estimate that $C(\alpha) \leq \nu^{2\alpha}$. This upper bound estimate grows exponentially with α , though the exact growth rate of $C(\alpha)$ may be much less.

For each $g \in \mathcal{G}_{\mathcal{F}}^*$ there is some α such that $g \in \mathcal{G}_{\mathcal{F}}^{(\alpha)}$. Let $\|g\|$ denote the least such α , which is called the *word length* of g . Note that $\mathcal{G}_{\mathcal{F}}^{(1)}$ generates $\mathcal{G}_{\mathcal{F}}^*$.

We use the word length on $\mathcal{G}_{\mathcal{F}}^*$ to define the word length on $\Gamma_{\mathcal{F}}$, where for $\gamma_w \in \Gamma_{\mathcal{F}}$, set

$$(18) \quad \|\gamma_w\| = \min \{ \|g\| \mid [g]_w = \gamma_w \text{ for } g \in \mathcal{G}_{\mathcal{F}}^* \}.$$

Introduce the *path length* of $\gamma_w \in \Gamma_{\mathcal{F}}$, by considering the infimum of the lengths $\|\gamma'\|$ for all piecewise smooth curves γ' for which $\gamma'_w = \gamma_w$. That is,

$$(19) \quad \ell(\gamma_w) = \inf \{ \|\gamma'\| \mid \gamma'_w = \gamma_w \}.$$

Note that if L_w is a leaf without holonomy, set $x = \tau(w)$ and $y = \tau(w')$, then Lemma 3.11 implies that $\ell(\gamma_w) = d_{\mathcal{F}}(x, y)$. This yields a fundamental estimate:

LEMMA 3.12. *Let $[g]_w \in \Gamma_{\mathcal{F}}$ where w corresponds to a leaf without holonomy. Then*

$$(20) \quad d_{\mathcal{F}}(x, y)/2\delta_{\mathcal{U}}^{\mathcal{F}} \leq \|[g]_w\| \leq 1 + d_{\mathcal{F}}(x, y)/\epsilon_{\mathcal{U}}^{\mathcal{F}}$$

Proof. Let $[g]_w$ be represented by $h_{\mathcal{I}}$ where $\mathcal{I} = (i_0, i_1, \dots, i_{\alpha})$ of length $\alpha = \|[g]_w\|$. Then \mathcal{I} defines a plaque-chain $\mathcal{P}_{\mathcal{I}}$ as in (12) where each plaque is contained in a disk $D_{\mathcal{F}}(x_{w_i}, \delta_{\mathcal{U}}^{\mathcal{F}})$, so the piecewise-geodesic $\gamma_{\mathcal{I}}$ which $\mathcal{P}_{\mathcal{I}}$ defines as in Section 3.2 has length at most $2\alpha \delta_{\mathcal{U}}^{\mathcal{F}}$. Thus, $d_{\mathcal{F}}(x, y) \leq 2\alpha \delta_{\mathcal{U}}^{\mathcal{F}}$.

Conversely, as noted in Section 3.3, given a path γ from x to y there is a good plaque-chain \mathcal{I} which covers γ , with $\alpha \leq 1 + \|\gamma\|/\epsilon_{\mathcal{U}}^{\mathcal{F}}$ which yields the right-hand-side estimate in (20). \square

In general, the leafwise distance function $d_{\mathcal{F}}$ of distinct leaves cannot be compared for points which have a large separation, except when it is possible to define a “shadowing” of a path in one leaf by a path in a nearby leaf, as described in the following.

Let $x \in \mathfrak{M}$ with L_x the leaf containing it, and assume that L_x is without holonomy. Given $y \in L_x$ let $\gamma_{x,y}$ be a geodesic from x to y with $\|\gamma_{x,y}\| = d_{\mathcal{F}}(x, y)$. Choose a plaque-chain $\mathcal{P}_{\mathcal{I}}$ covering $\gamma_{x,y}$ as in Section 3.3 so that Proposition 3.5 holds for the holonomy map $h_{\mathcal{I}}$ it induces. Let U_{i_0} be the first chart of the plaque-chain, so that $x \in \mathcal{P}_{i_0}(\xi_0)$, with notation as in (13). Let $x' \in U_{i_0}$ have the same leaf coordinate as x , so $\lambda_{i_0}(x') = \lambda_{i_0}(x)$. Assume that $w' = \pi(x') \in \text{Dom}(h_{\mathcal{I}})$, then we define $y' \in U_{i_{\alpha}}$ corresponding to y in the same way, for the last plaque $\mathcal{P}_{i_{\alpha}}(\xi_{\alpha})$ in the plaque-chain. The curve $\gamma_{x,y}$ is given in (13) as a concatenation of piecewise-geodesic segments, each of which is contained in a plaque of the plaque-chain $\mathcal{P}_{\mathcal{I}}$. Then by shadowing these geodesic segments within each respective foliation chart, we obtain a piecewise geodesic $\tilde{\gamma}_{x',y'}$ from x' to y' . Then using the estimates (9) and Proposition 3.5, we obtain:

LEMMA 3.13. *Let $\epsilon > 0$ and $R > 0$, then there exists $\delta(\epsilon, R) > 0$ such that if $x \in \mathfrak{M}$ is contained in a leaf L_x without holonomy, and $y \in L_x$ satisfies $d_{\mathcal{F}}(x, y) \leq R$, then for x' as above with $w' = \pi(x') \in B_{\mathfrak{X}}(w, \delta(\epsilon, R))$, it follows that $w' \in \text{Dom}(h_{\mathcal{I}})$ and for the corresponding endpoint $y' \in L_{x'}$ we have*

$$(21) \quad d_{\mathcal{F}}(x', y') \leq (1 + \epsilon) d_{\mathcal{F}}(x, y) \quad \text{and} \quad d_{\mathcal{F}}(x, y) \leq (1 + \epsilon) d_{\mathcal{F}}(x', y').$$

Proof. Let $\gamma_{x,y}$ be a leafwise geodesic from x to y with $\|\gamma_{x,y}\| = d_{\mathcal{F}}(x, y)$. Construct a plaque-chain covering $\mathcal{P}_{\mathcal{I}}$ of $\gamma_{x,y}$ as above. Let $\delta_{\mathcal{U}}(\epsilon)$ be the constant appearing in (9), then by Lemma 3.12 and Proposition 3.5, there exists $\delta(\delta_{\mathcal{U}}(\epsilon), R) > 0$ so that $B_{\mathfrak{X}}(w, \delta(\delta_{\mathcal{U}}(\epsilon), R)) \subset \text{Dom}(h_{\mathcal{I}})$. Moreover, for any $w' \in B_{\mathfrak{X}}(w, \delta(\delta_{\mathcal{U}}(\epsilon), R))$ the plaques defined by the shadowed points along the piecewise geodesic

$\tilde{\gamma}_{x',y'}$ from x' to y' have separation at most $\delta_{\mathcal{U}}(\epsilon)$. Then use the estimates in (9) in each plaque to deduce the first estimate in (21). The second estimate follows by considering the path $\gamma_{x,y}$ as shadowing the path $\tilde{\gamma}_{x',y'}$. \square

The technical result Lemma 3.13 illustrates some of the difficulties working with the dynamics of pseudogroups, as the notion of distance along paths in leaves requires care with the initial point.

3.9. Induced groupoids. For an open subset $W \subset \mathfrak{X}$, consider the following subsets of $\Gamma_{\mathcal{F}}$,

$$\begin{aligned}\Gamma_W &= \{[g]_w \in \Gamma_{\mathcal{F}} \mid w \in W\} \\ \Gamma_W^W &= \{[g]_w \in \Gamma_{\mathcal{F}} \mid w \in W, r([g]_w) \in W\} \\ \Gamma_w^W &= \{[g]_w \in \Gamma_{\mathcal{F}} \mid r([g]_w) \in W\}\end{aligned}$$

so we have the inclusions

$$\Gamma_w^w \subset \Gamma_w^W \subset \Gamma_W^W \subset \Gamma_W \subset \Gamma_{\mathcal{F}}.$$

The set Γ_W^W is called the induced groupoid on W . On the other hand, the sets Γ_W and Γ_w^W are not groupoids, as they need not contain the inverses of their elements.

Given $w \in W$ and $R > 0$, define the word and path length filtration of the groupoid Γ_w^W ,

$$(22) \quad \Gamma_w^{W,R} = \{\gamma_w \in \Gamma_w^W \mid \ell(\gamma_w) \leq R\}, \quad \Gamma_w^{W,\alpha} = \left\{ [g]_w \in \Gamma_w^W \mid g \in \mathcal{G}_{\mathcal{F}}^{(\alpha)} \text{ with } w \in \text{Dom}(g) \right\}.$$

By Lemma 3.12, we then have the estimates:

LEMMA 3.14. *Let L_w be a leaf without holonomy. For $\alpha > 0$, set $R^{(\alpha)} = 2\alpha\delta_{\mathcal{U}}^{\mathcal{F}}$ and $R_{(\alpha)} = \alpha\epsilon_{\mathcal{U}}^{\mathcal{F}}$ where we recall from Section 2.1 that $\epsilon_{\mathcal{U}}^{\mathcal{F}} < \delta_{\mathcal{U}}^{\mathcal{F}}$ so that $R^{(\alpha)} > R_{(\alpha)}$, then*

$$(23) \quad \Gamma_w^{W,R_{(\alpha)}} \subset \Gamma_w^{W,\alpha} \subset \Gamma_w^{W,R^{(\alpha)}}.$$

The inclusions in (23) are used in the definition of the coding of the orbits of $\mathcal{G}_{\mathcal{F}}^*$ in Section 6.

4. DELONE SETS AND VORONOI TESSELLATIONS

In this section, we consider the *Delone sets* and their associated *Voronoi tessellations* on the leaves of \mathcal{F} which are associated to a choice of a clopen transversal set $W \subset \mathfrak{X}$. These ideas are discussed in extensive detail in [18]. It is important to note that, as no geometric assumption is made on the leaves of \mathcal{F} , many standard techniques for Voronoi tessellations in the literature for Euclidean spaces do not apply to this general case.

Assume that \mathfrak{M} is a minimal matchbox manifold, so for every $w \in \mathfrak{X}$, its $\mathcal{G}_{\mathcal{F}}$ -orbit $\mathcal{O}(w)$ is dense. This implies that the restriction of $\mathcal{G}_{\mathcal{F}}$ to any clopen subset $U \subset \mathfrak{X}$ yields an induced dynamical system which is “equivalent” to the action of $\mathcal{G}_{\mathcal{F}}$ on \mathfrak{X} , as discussed in [19]. This restriction property is fundamental for the study of minimal Cantor actions. Here is an essential application of minimality:

LEMMA 4.1. *Let $W \subset \mathfrak{X}$ be an open subset. Then there exists an integer α_W such that \mathfrak{X} is covered by the collection $\{h_{\mathcal{I}}(W) \mid h_{\mathcal{I}} \in \mathcal{G}_{\mathcal{F}}^{(\alpha_W)}\}$. Moreover, as $\text{diam}_{\mathfrak{X}}(W) \rightarrow 0$, we have $\alpha_W \rightarrow \infty$.*

Proof. Consider the collection of open sets $\mathcal{Q} = \{h_{\mathcal{I}}(W) \mid h_{\mathcal{I}} \in \mathcal{G}_{\mathcal{F}}\}$ where we adopt the abuse of notation that $h_{\mathcal{I}}(W) = h_{\mathcal{I}}(W \cap \text{Dom}(h_{\mathcal{I}}))$. As the action of $\mathcal{G}_{\mathcal{F}}$ on \mathfrak{X} is minimal, for all $y \in \mathfrak{X}$ there exists $h_{\mathcal{I}}$ with $h_{\mathcal{I}}(y) \in W$. Thus, the collection \mathcal{Q} is an open covering of \mathfrak{X} . Since \mathfrak{X} is compact, there exists a finite subcover $\{h_{\mathcal{I}_1}(W), \dots, h_{\mathcal{I}_m}(W)\}$ of \mathcal{Q} , where $\mathcal{I}_i = (j_0, \dots, j_{\alpha_i})$. Let

$$(24) \quad \alpha_W = \max\{\alpha_i \mid 1 \leq i \leq m\}.$$

The conclusion that $\alpha_W \rightarrow \infty$ as $\text{diam}_{\mathfrak{X}}(W) \rightarrow 0$ follows immediately, as assuming α_W is bounded implies that \mathfrak{X} admits coverings by clopen sets with arbitrarily small diameters, yet bounded in number, which is impossible. \square

4.1. **Delone sets.** We define nets in the leaves of \mathcal{F} related to the dynamics of $\mathcal{G}_{\mathcal{F}}$.

DEFINITION 4.2. *Let (X, d_X) be a complete separable metric space. Given $0 < \lambda_1 < \lambda_2$, a subset $\mathcal{N} \subset X$ is a (λ_1, λ_2) -net (or Delone set) if:*

- (1) \mathcal{N} is λ_1 -separated: for all $y \neq z \in \mathcal{N}$, $d_X(y, z) \geq \lambda_1$;
- (2) \mathcal{N} is λ_2 -dense: for all $x \in X$, there exists some $z \in \mathcal{N}$ such that $d_X(x, z) \leq \lambda_2$.

Given $w \in \mathfrak{X}$, let L_w denote the leaf containing $x = \tau(w)$. Define $\mathcal{N}_w = L_w \cap \mathcal{T}$. Then (4) implies that \mathcal{N}_w is an $\delta_{\mathcal{T}}$ -separated subset of L_w . On the other hand, \mathcal{U} is an open covering of \mathfrak{M} , and therefore by Proposition 2.6.4 each point of L_w has distance at most $\delta_{\mathcal{U}}^{\mathcal{F}}$ from a point of \mathcal{N}_w . That is, \mathcal{N}_w is always an $(\delta_{\mathcal{T}}, \delta_{\mathcal{U}}^{\mathcal{F}})$ -net in L_w .

In the following, we assume without loss of generality that the open set W satisfies $W \subset \mathfrak{T}_1$, and set $\mathcal{T}_W = \tau_1(W) \subset \mathfrak{M}$. For each $w \in W$, define $\mathcal{N}_w^W = L_w \cap \mathcal{T}_W \subset \mathcal{N}_w$.

If w is a fixed point for some non-trivial holonomy map h_{γ} defined by a path γ , then the length of γ gives an approximate upper bound on the constant λ_1 for the net \mathcal{N}_w^W . Define the function

$$(25) \quad \lambda_1(\delta) = \inf \{ \lambda \mid w \in W \subset \mathfrak{X}, \text{diam}_{\mathfrak{X}}(W) \leq \delta, \mathcal{N}_w^W \text{ is } \lambda\text{-separated} \}.$$

We always have $\lambda_1(\delta) \geq \delta_{\mathcal{T}}$ as $\mathcal{N}_w^W \subset \mathcal{N}_w$ for all $w \in \mathfrak{X}$, and \mathcal{N}_w is $\delta_{\mathcal{T}}$ -separated. The following result shows there is no upper bound on the function λ_1 for a foliation without holonomy.

LEMMA 4.3. *Let \mathfrak{M} be a matchbox manifold without holonomy, then $\lambda_1(\delta)$ is increasing, and unbounded as $\delta \rightarrow 0$.*

Proof. If $0 < \delta' < \delta$ then $\lambda_1(\delta') \geq \lambda_1(\delta)$ follows from the definition, so the function λ_1 is increasing.

We show that the function λ_1 is unbounded. Assume there exists $\lambda_* > 0$, such that for each integer $\ell > 0$, there exists distinct points $u_{\ell}, v_{\ell} \in \mathfrak{X}$ such that $d_{\mathfrak{X}}(u_{\ell}, v_{\ell}) \leq 1/\ell$ and there is a path γ_{ℓ} joining $y_{\ell} = \tau(u_{\ell})$ and $z_{\ell} = \tau(v_{\ell})$ with $\|\gamma_{\ell}\| \leq \lambda_*$. As \mathfrak{X} is compact, by passing to subsequences, we can assume that there exists $w \in \mathfrak{X}$ which is the limit of both sequences $\{u_{\ell}\}$ and $\{v_{\ell}\}$.

By Lemma 3.12, we can choose a plaque-chain \mathcal{I}_{ℓ} which covers γ_{ℓ} with length at most $1 + \lambda_*/\epsilon_{\mathcal{U}}^{\mathcal{F}}$. Since there are only a finite number of admissible plaque-chains of length at most $1 + \lambda_*/\epsilon_{\mathcal{U}}^{\mathcal{F}}$, we can pass to a subsequence, and assume without loss of generality that $\mathcal{I} = \mathcal{I}_{\ell}$ is independent of ℓ . Then Proposition 3.5 implies that we can assume that there is some $\delta' > 0$ and $\ell_0 > 0$ with $u_{\ell}, v_{\ell} \in B_{\mathfrak{X}}(w, \delta') \subset \text{Dom}(h_{\mathcal{I}})$ for all $\ell \geq \ell_0$. Then the holonomy map $h_{\mathcal{I}}$ satisfies $h_{\mathcal{I}}(u_{\ell}) = v'_{\ell}$ for all $\ell \geq \ell_0$, and thus we have $h_{\mathcal{I}}(w) = w$. The assumption $u_{\ell} \neq v_{\ell}$ for all $\ell \geq \ell_0$ implies that $h_{\mathcal{I}}$ is not the identity map on any open neighborhood of w . Thus, the germ $[h_{\mathcal{I}}]_w$ is non-trivial, so that the leaf L_w has is germinal holonomy, which contradicts our assumption that \mathcal{F} is without holonomy. \square

REMARK 4.4. The function $\lambda_1(\delta)$ is a fundamental property of the dynamics of the pseudogroup $\mathcal{G}_{\mathcal{F}}$ acting on \mathfrak{X} . From the above proof, it is clear that the function is closely related to the “return times” of the action near to w , though little more seems to be known about its properties in general. Note that Forrest in [27] derives estimates for the constants $\{\lambda_1, \lambda_2\}$ for the net \mathcal{N}_w^W in the case where $\mathcal{G}_{\mathcal{F}}$ is induced from a minimal action of \mathbb{Z}^n on a Cantor set. It is a very interesting problem to derive estimates for these constants in the more general cases of group actions and pseudogroups.

We next show that the density of the net \mathcal{N}_w^W in the leaf L_w has a uniform estimate in terms of the integer α_W introduced in Lemma 4.1.

LEMMA 4.5. *The set \mathcal{N}_w^W is $(2\alpha_W \delta_{\mathcal{U}}^{\mathcal{F}})$ -dense in the set \mathcal{N}_w .*

Proof. Let L_w be the leaf determined by $w \in W$. Let $y \in \mathcal{N}_w$ and $w' = \pi(y) \in \mathfrak{X}$. The collection $\{h_{\mathcal{J}}(W) \mid h_{\mathcal{J}} \in \mathcal{G}_{\mathcal{F}}^{(\alpha_W)}\}$ is an open covering of \mathfrak{X} by Lemma 4.1, so there exists \mathcal{J} with length at most α_W with $w' \in h_{\mathcal{J}}(W)$, which defines a plaque-chain from x to y with at most α_W plaque intersections. Each plaque is contained in a disk $D_{\mathcal{F}}(x_{w_i}, \delta_{\mathcal{U}}^{\mathcal{F}})$, thus there is a piecewise geodesic path in L_w from x to y with length at most $2\alpha_W \delta_{\mathcal{U}}^{\mathcal{F}}$. \square

COROLLARY 4.6. For $e_W = (2\alpha_W + 1)\delta_{\mathcal{U}}^{\mathcal{F}}$ the set \mathcal{N}_w^W is e_W -dense in L_w . That is, for every $x \in \mathfrak{M}$ there is a piecewise geodesic path with length at most e_W joining x to a point of \mathcal{T}_W .

Proof. Given $x \in L_w$ there exists $y \in \mathcal{N}_w$ with $d_{\mathcal{F}}(x, y) \leq \delta_{\mathcal{U}}^{\mathcal{F}}$. By Lemma 4.5, there is $z \in \mathcal{N}_w^W$ with $d_{\mathcal{F}}(y, z) \leq 2\alpha_W \delta_{\mathcal{U}}^{\mathcal{F}}$ and thus $d_{\mathcal{F}}(x, z) \leq (2\alpha_W + 1)\delta_{\mathcal{U}}^{\mathcal{F}}$, as was to be shown. \square

4.2. Voronoi tessellations. We develop some of the basic concepts of *Voronoi tessellations* for complete Riemannian manifolds, and in particular for the leaves of \mathcal{F} . As the leaves of \mathfrak{M} are not necessarily Euclidean, some of the usual results for tessellations of \mathbb{R}^n with the Euclidean metric need not be true in this context (see [42] and [48, Introduction]).

The *Voronoi tessellation* associated to \mathcal{N}_w^W is a partition of L_w into compact regions with disjoint interiors, called the “cells”, where each cell is “centered” at a unique point of \mathcal{N}_w^W . We describe the construction of these cells and some of their properties.

Introduce the “leafwise nearest-neighbor” distance function, where for $y \in L_w$,

$$(26) \quad \kappa_w^W(y) = \inf \{ d_{\mathcal{F}}(y, z) \mid z \in \mathcal{N}_w^W \}.$$

Note that $\kappa_w^W(y) < \infty$ if and only if $y \in L_w$, and $\kappa_w^W(y) = 0$ if and only if $y \in \mathcal{N}_w^W$.

DEFINITION 4.7. For $y \in \mathcal{N}_w^W$, define its Voronoi cell in L_w by

$$(27) \quad \mathcal{C}_w^W(y) = \{ z \in L_w \mid d_{\mathcal{F}}(y, z) = \kappa_w^W(z) \}.$$

That is, for $x \in \mathcal{N}_w^W$ the Voronoi cell $\mathcal{C}_w^W(y)$ consists of the points $z \in L_w$ which are at least as close to y in the leafwise metric as to any other point of \mathcal{N}_w^W .

Note that if L_w is a convex metric space for the leaf metric, for example when the leaf L_w is a Euclidean space, then each cell $\mathcal{C}_w^W(y)$ is a convex subset. However, in the general case, the cells $\mathcal{C}_w^W(y)$ need not be convex, or even simply connected, especially if the leaf L_w has non-trivial topology and the constants λ_1 and λ_2 for the net \mathcal{N}_w^W are “large”. In this work, we assume only that the leaves of \mathfrak{M} are complete metric spaces, and so develop some of the properties of the cells $\mathcal{C}_w^W(y)$ in the decomposition (28) for this generality.

First, note that by the definition, for each $y \in L$, the set $\mathcal{C}_w^W(y)$ is an intersection of closed subsets of L_w , hence is closed. Also, from the definitions every point of L_w belongs to some $\mathcal{C}_w^W(z)$ for $z \in \mathcal{N}_w^W$, so we have:

DEFINITION 4.8. The Voronoi tessellation of L_w associated to \mathcal{N}_w^W is the decomposition

$$(28) \quad L_w = \bigcup_{y \in \mathcal{N}_w^W} \mathcal{C}_w^W(y)$$

In general, the geometry of a cell $\mathcal{C}_w^W(z) \subset L_w$ may be difficult to describe, though we always have the following bounds, which are direct consequences of Lemma 4.3 and Corollary 4.6.

LEMMA 4.9. Suppose that $W \subset B_{\mathfrak{X}}(w, \delta)$, then for each $y \in \mathcal{N}_w^W$,

$$(29) \quad D_{\mathcal{F}}(y, \lambda_1(\delta)/2) \subset \mathcal{C}_w^W(y) \subset B_{\mathcal{F}}(y, e_W)$$

In particular, each cell $\mathcal{C}_w^W(y)$ is closed with diameter at most $2e_W$, and thus is compact.

If \mathcal{F} is a foliation without holonomy, then Lemmas 4.3 and 4.9 imply that if W is chosen to have sufficiently small diameter, then each Voronoi cell in L_w defined by the net $\mathcal{N}_w^W = L_w \cap \tau(W)$ contains a ball of radius larger than any prescribed size.

Next, we introduce the *star-neighborhoods* of Voronoi cells. Given $y \in \mathcal{N}_w^W$, introduce the *vertex-set*

$$(30) \quad \mathcal{V}_w^W(y) = \{ z \in \mathcal{N}_w^W \mid \mathcal{C}_w^W(z) \cap \mathcal{C}_w^W(y) \neq \emptyset \}.$$

DEFINITION 4.10. For $y \in \mathcal{N}_w^W$ the “star-neighborhood” of the Voronoi cell $\mathcal{C}_w^W(y)$ is the set

$$(31) \quad \mathcal{S}_w^W(y) = \bigcup_{z \in \mathcal{V}_w^W(y)} \mathcal{C}_w^W(z).$$

LEMMA 4.11. ([18]) For each $y \in \mathcal{N}_w^W$ and $z \in \mathcal{V}_w^W(y)$, we have that $d_{\mathcal{F}}(y, z) \leq 2e_W$. Hence, $\mathcal{S}_w^W(y) \subset B_{\mathcal{F}}(y, 3e_W)$, and the collection $\mathcal{V}_w^W(y)$ is finite.

For each $x \in \mathcal{N}_w^W$ define the boundary of $\mathcal{C}_w^W(x)$ to be the set $\partial\mathcal{C}_w^W(x) \equiv \mathcal{C}_w^W(x) - \text{int}(\mathcal{C}_w^W(x))$. We consider the relation between the star-neighborhood of a cell and its boundary $\partial\mathcal{C}_w^W(x)$.

LEMMA 4.12. For each $x \in \mathcal{N}_w^W$, we have that $\mathcal{C}_w^W(x) \subset \text{int}(\mathcal{S}_w^W(x))$.

Proof. We always have $\text{int}(\mathcal{C}_w^W(x)) \subset \text{int}(\mathcal{S}_w^W(x))$, so we must show $\partial\mathcal{C}_w^W(x) \subset \text{int}(\mathcal{S}_w^W(x))$.

Let $y \in \partial\mathcal{C}_w^W(x)$ with $y \notin \text{int}(\mathcal{S}_w^W(x))$. Then for every $0 < \epsilon < e_W$, the open ball $B_{\mathcal{F}}(y, \epsilon)$ must intercept the complement of $\mathcal{S}_w^W(x)$ in L_w in a non-empty set; let y'_ϵ be a point in this set. Thus, there is some $z_\epsilon \in \mathcal{N}_w^W$ for which $y'_\epsilon \in \mathcal{C}_w^W(z_\epsilon)$ so that $d_{\mathcal{F}}(x, z_\epsilon) \leq 2e_W + \epsilon < 3e_W$. There are at most a finite number of such net points z_ϵ , so we can choose a subsequence of the points $y'_{1/\ell}$ with $d_{\mathcal{F}}(y, y'_{1/\ell}) < 1/\ell$ and $y'_{1/\ell} \in \mathcal{C}_w^W(z')$ for all ℓ , where $z' \in \mathcal{N}_w^W$ is fixed. As $\mathcal{C}_w^W(z')$ is compact, this implies that $\mathcal{C}_w^W(z') \cap \mathcal{C}_w^W(x) \neq \emptyset$ and so $y'_{1/\ell} \in \mathcal{S}_w^W(x)$ contrary to choice. \square

For $z \in \mathcal{N}_w^W$ with $z \neq y$, introduce the closed subsets of the leaf L_w defined by

$$\begin{aligned} H_w^W(y, z) &= \{\xi \in L_w \mid d_{\mathcal{F}}(y, \xi) \leq d_{\mathcal{F}}(z, \xi)\} \\ L_w^W(y, z) &= \{\xi \in L_w \mid d_{\mathcal{F}}(y, \xi) = d_{\mathcal{F}}(z, \xi)\} \end{aligned}$$

If L_w is isometric to Euclidean space \mathbb{R}^n , then $H_w^W(y, z)$ is a closed half-space, and $L_w^W(y, z)$ is the boundary hyperplane. In general, the closed subspace $L_w^W(y, z)$ need not be a manifold.

The next result is a “local” version of the observation that if L_w is isometric to Euclidean space \mathbb{R}^n , then each Voronoi cell $\mathcal{C}_w^W(y) \subset \mathbb{R}^n$ is the intersection of half-spaces defined by its boundary planes.

LEMMA 4.13. For each $y \in \mathcal{N}_w^W$ we have $\mathcal{C}_w^W(y) = \mathcal{S}_w^W(y) \cap \bigcap_{z \in \mathcal{V}_w^W(y)} H_w^W(y, z)$.

Proof. The identity below and the subsequent inclusion follow from the definitions:

$$(32) \quad \mathcal{C}_w^W(y) = \bigcap_{z \in \mathcal{N}_w^W} H_w^W(y, z) \subset \mathcal{S}_w^W(y) \cap \bigcap_{z \in \mathcal{V}_w^W(y)} H_w^W(y, z).$$

We must show the reverse inclusion in (32) holds as well. Observe that for $z \neq y$, we have $\mathcal{C}_w^W(y) \cap \mathcal{C}_w^W(z) \subset L_w^W(y, z)$ by the definition. We claim that also, $\mathcal{C}_w^W(y) \cap L_w^W(y, z) \subset \mathcal{C}_w^W(z)$. Let $\xi \in \mathcal{C}_w^W(y) \cap L_w^W(y, z)$, so that ξ is closer to y than any other point in \mathcal{N}_w^W and $d_{\mathcal{F}}(y, \xi) = d_{\mathcal{F}}(z, \xi)$. Suppose that $\xi \notin \mathcal{C}_w^W(z)$, then there exists $z' \in \mathcal{N}_w^W$ with $d_{\mathcal{F}}(z', \xi) < d_{\mathcal{F}}(z, \xi) = d_{\mathcal{F}}(y, \xi)$, so $z' \notin \mathcal{C}_w^W(y)$, which contradicts the choice of ξ . Thus, $\xi \in \mathcal{C}_w^W(z)$.

Now consider $z \in \mathcal{N}_w^W$ and $z \notin \mathcal{V}_w^W(y)$. Then $\mathcal{C}_w^W(y) \cap L_w^W(y, z) = \emptyset$. The cell $\mathcal{C}_w^W(y)$ is compact by Lemma 4.9, so there exists $\epsilon > 0$ such that the ϵ -neighborhood of $\mathcal{C}_w^W(y)$ is also disjoint from $L_w^W(y, z)$. Since $\mathcal{C}_w^W(y) \subset H_w^W(y, z)$, this ϵ -neighborhood is also contained in $H_w^W(y, z)$. Note that for all $z' \in \mathcal{N}_w^W$, either $\mathcal{C}_w^W(z') \subset H_w^W(y, z)$ or $\mathcal{C}_w^W(z') \subset H_w^W(z, y)$ so by Lemma 4.12, we conclude that $\mathcal{C}_w^W(z') \subset H_w^W(y, z)$ for all $z' \in \mathcal{V}_w^W(y)$. That is, $\mathcal{S}_w^W(y) \subset H_w^W(y, z)$, which implies the reverse inclusion, as was to be shown. \square

5. INDUCED PSEUDOGRUUPS

Let $W \subset \mathfrak{X}$ be an open subset, and let $\mathcal{G}_{\mathcal{F}}^*(W)$ denote the restriction of $\mathcal{G}_{\mathcal{F}}^*$ to W . That is,

$$(33) \quad \mathcal{G}_{\mathcal{F}}^*(W) = \{h := (h_{\mathcal{I}}, U_h) \in \mathcal{G}_{\mathcal{F}}^* \mid \emptyset \neq U_h \subseteq \text{Dom}(h_{\mathcal{I}}) \cap W \text{ and } h_{\mathcal{I}}(U_h) \subseteq W\}.$$

The goal of this section is to show that if the action of $\mathcal{G}_{\mathcal{F}}$ on \mathfrak{X} is minimal, then $\mathcal{G}_{\mathcal{F}}^*(W)$ admits a *finite* symmetric generating set $\mathcal{G}_{\mathcal{F}}^*(W)^{(1)}$, whose elements have ‘‘maximal’’ domains in W . We begin with some preliminary considerations. Introduce the constants:

$$(34) \quad \begin{aligned} \epsilon_{H,i} &< \max \{ \epsilon \mid \forall x \in \bar{U}_i \text{ such that } D_{\mathfrak{M}}(x, \epsilon_{\mathcal{U}}/4) \subset \bar{U}_i, \text{ then } D_{\mathfrak{X}}(\pi_i(x), \epsilon) \subset \pi_i(D_{\mathfrak{M}}(x, \epsilon_{\mathcal{U}}/4)) \} \\ \epsilon_H &= \min \{ \epsilon_{H,i} \mid \forall 1 \leq i \leq \nu \}. \end{aligned}$$

Note that $\epsilon_H \leq \epsilon_{\mathcal{U}}^{\mathcal{F}}$ follows by comparing the definition (34) of ϵ_H with the definition (8) of $\epsilon_{\mathcal{U}}^{\mathcal{F}}$.

Recall that α_W is an upper bound on the length of plaque-chains \mathcal{I} required for a covering of \mathfrak{X} by open sets of the form $h_{\mathcal{I}}(W)$. and that Lemma 4.5 shows that $h_{\mathcal{I}}$ can be realized by the holonomy along a piecewise geodesic path of length at most $2\alpha_W \delta_{\mathcal{U}}^{\mathcal{F}}$.

Note that while α_W bounds the lengths of plaque-chains required to obtain a covering of \mathfrak{X} by images of W , there is no control over the domains of the holonomy maps defined by the overlaps between the plaques in these chains. Introduce the integer

$$(35) \quad \beta_W = \lceil 2\alpha_W \delta_{\mathcal{U}}^{\mathcal{F}} / \epsilon_{\mathcal{U}}^{\mathcal{F}} \rceil \geq 2\alpha_W$$

where the inequality follows from the inequality $\delta_{\mathcal{U}}^{\mathcal{F}} / \epsilon_{\mathcal{U}}^{\mathcal{F}} > 1$ of Section 2. By allowing plaque-chains of length β_W we can obtain uniform estimates on the domains of the holonomy maps used, as shown by the following remarks. First, introduce the constant defined by Proposition 3.5 for $\epsilon = \epsilon_H$,

$$(36) \quad \delta_W = \delta(\epsilon_H, \beta_W) \leq \epsilon_H.$$

Let $h_{\mathcal{I}}[U_h \in \mathcal{G}_{\mathcal{F}}^*(W)$ where \mathcal{I} has length at most α_W . Let $w \in U_h \subset W$ define a leaf L_w without holonomy. Set $w' = h(w) \in W$ and so $w' \in \mathcal{N}_w^W$. Recall that the leaves of \mathcal{F} are complete Riemannian manifolds, so there exists a geodesic segment $\gamma_{x,y}$ from $x = \tau(w)$ to $y = \tau(w')$ with length $d_{\mathcal{F}}(x, y) \leq 2\alpha_W \delta_{\mathcal{U}}^{\mathcal{F}}$. By the method of Section 3.3 used in the proof of Proposition 3.5, we can choose a plaque-chain \mathcal{I}' covering $\gamma_{x,y}$ with length at most $\alpha \leq d_{\mathcal{F}}(x, y) / \epsilon_{\mathcal{U}}^{\mathcal{F}} \leq \beta_W$, and such that the holonomy map $h_{\mathcal{I}'}$ satisfies $D_{\mathfrak{X}}(w, \delta_W) \subset \text{Dom}(h_{\mathcal{I}'})$, and

$$(37) \quad h_{\mathcal{I}'}(D_{\mathfrak{X}}(w, \delta_W)) \subset D_{\mathfrak{X}}(w', \epsilon_H).$$

As L_w is a leaf without holonomy, by Lemma 3.11 we have $[h]_w = [h_{\mathcal{I}'}]_w$. We thus obtain the main result of this section:

PROPOSITION 5.1. *Let $V \subset W$ be an open subset with $\text{diam}_{\mathfrak{X}}(V) \leq \delta_W$, and let $w \in V$ be such that L_w is without holonomy. Then for all $h \in \mathcal{G}_{\mathcal{F}}^*(W)$ and $w \in V \cap U_h$ with $w' = h(w)$ represented by a path $\gamma_{x,y}$ with length $d_{\mathcal{F}}(x, y) \leq 2\alpha_W \delta_{\mathcal{U}}^{\mathcal{F}}$, then there exists an admissible sequence \mathcal{I}' with length at most β_W such that $V \subset \text{Dom}(h_{\mathcal{I}'})$ and $[h]_w = [h_{\mathcal{I}'}]_w$.*

Introduce the following subset of $\mathcal{G}_{\mathcal{F}}^*(W)$,

$$(38) \quad \mathcal{G}_{\mathcal{F}}^*(W)^{(1)} = \{h_{\mathcal{I}} \mid \mathcal{I} = (i_0, i_1, \dots, i_k), k \leq \beta_W, \text{Dom}(h_{\mathcal{I}}) \cap W \neq \emptyset, h_{\mathcal{I}}(W) \cap W \neq \emptyset\}.$$

The number of admissible sequences \mathcal{I} with length at most β_W is bounded above by $\nu^{(\beta_W+1)}$, so $\mathcal{G}_{\mathcal{F}}^*(W)^{(1)}$ is a finite collection of maps.

6. DYNAMICAL PARTITIONS OF THE TRANSVERSAL

In this section, we construct coding functions for a minimal action of the pseudogroup $\mathcal{G}_{\mathcal{F}}^*$ on \mathfrak{X} , and show some basic properties for such codings. We note that Gromov [33, 34] and Fried [30] described applications of the coding technique in dynamics to the study of actions by finitely-generated groups, as described in the text by Coornaert and Papadopoulos [22] which gives an excellent overview. Our development of codings for the orbits for a pseudogroup action is an extension of this method.

The coding technique was developed for equicontinuous actions of pseudogroups on a Cantor space \mathfrak{X} in the work [17], where it was also shown that the domains of the coding maps have uniform estimates, which fails in the general case considered here. Thus, an important aspect of our development of the coding of the action of the pseudogroup $\mathcal{G}_{\mathcal{F}}^*$ is the careful treatment of the relation between the dynamics of the holonomy maps, and the domains on which the coding function is defined.

After establishing the basic properties of the coding functions, we give an inductive method to construct a nested sequence of partitions of \mathfrak{X} into clopen sets, defined using the level sets of the coding function. The resulting clopen covers of \mathfrak{X} so defined yield “multi-dimensional” Kakutani-Rokhlin towers for the action of $\mathcal{G}_{\mathcal{F}}^*$, analogous to those constructed by Forrest in [27] for minimal \mathbb{Z}^n -actions. Our construction here is also related to the “zooming” procedure in Bellissard, Benedetti and Gambaudo [11, Proposition 2.42] whose proof uses the Euclidean geometry of the leaves. The work of this section and Section 7 is also related to the construction of a “tower system” for a tiling space $\Omega_{\mathbf{T}}$ as sketched in the work of Benedetti and Gambaudo [12, Theorem 3.1], where the leaves of $\Omega_{\mathbf{T}}$ are quotients of a connected Lie group. The method of “inflation” described in Section 3.1 of [4] can also be seen as analogous, though the lack of explicit details in these previous works makes the connections between the approaches to be mostly on an intuitive level.

In this section, we allow \mathcal{F} to have non-trivial holonomy, but work with a leaf without holonomy, which always exists. Fix a point $w_0 \in \mathfrak{X}$ without holonomy, and without loss of generality assume that $w_0 \in \mathfrak{T}_1$. Let L_0 denote the leaf determined by w_0 and let $\mathcal{N}_0 = L_0 \cap \mathcal{T}$ be the induced net.

6.1. Setting constants. The analysis of the domains of compositions of elements in $\mathcal{G}_{\mathcal{F}}^*$ requires careful attention to the metric properties of the action. We first fix some basic constants.

Let $0 < \epsilon_1 \leq \epsilon_H$ where ϵ_H is defined by (34). Choose $V_1 \subset \mathfrak{T}_1 \subset \mathfrak{X}$ a clopen neighborhood of w_0 with $\text{diam}_{\mathfrak{X}}(V_1) < \epsilon_1$. Set $\mathcal{N}_1 = L_0 \cap \tau(V_1)$, which is a subnet of \mathcal{N}_0 .

Let α_1 be the word length of the elements needed to cover \mathfrak{X} by translates of V_1 , as defined by Lemma 4.1. Then the generators of the induced pseudogroup $\mathcal{G}_{\mathcal{F}}(V_1)^{(1)}$ are represented by elements of $\mathcal{G}_{\mathcal{F}}^*$ defined by the holonomy along paths of lengths at most $2\alpha_1\delta_{\mathcal{U}}^{\mathcal{F}}$, and thus by Proposition 5.1 are represented by maps in $\mathcal{G}_{\mathcal{F}}^{(\beta_1)}$, where $\beta_1 = \lceil 2\alpha_1\delta_{\mathcal{U}}^{\mathcal{F}}/\epsilon_{\mathcal{U}}^{\mathcal{F}} \rceil$.

Now let $\theta_1 = (2\alpha_1 + 1)\delta_{\mathcal{U}}^{\mathcal{F}}$, then for all $w \in V_1$, the net $\mathcal{N}_w^{V_1} = L_w \cap \tau(V_1)$ is θ_1 -dense in L_w by Corollary 4.6. Moreover, by Lemma 4.9 for each $y \in \mathcal{N}_w^{V_1}$ the Voronoi cell $\mathcal{C}_w^{V_1}(y) \subset B_{\mathcal{F}}(y, \theta_1)$, and by Lemma 4.11 the star-neighborhood $\mathcal{S}_w^{V_1}(y) \subset B_{\mathcal{F}}(y, 3\theta_1)$. Introduce the constants

$$(39) \quad R'_1 = 2\theta_1 + \lambda_{\mathcal{F}} \quad , \quad R_1 = 2R'_1 = 4\theta_1 + 2\lambda_{\mathcal{F}} .$$

Finally, observe that a leafwise path of length R_1 in \mathfrak{M} can be covered by a good plaque-chain of length at most $\widehat{\beta}_1 + 1$, where $\widehat{\beta}_1 = \lceil R_1/\epsilon_{\mathcal{U}}^{\mathcal{F}} \rceil$.

Proposition 3.5, for the case $\epsilon = \epsilon_H$ and $\alpha = \widehat{\beta}_1$, yields a uniform estimate on the radius of balls contained in the domains of holonomy maps associated to plaque-chains of length at most R_1 , and such that the images of these balls have diameter at most ϵ_H . As L_0 is a leaf without holonomy, the holonomy along a path depends only on the endpoints. As there are only a finite number of such plaque-chains, there are at most a finite number of homotopies required between them as in the proof of Lemma 3.3 which will have length at most $\widehat{\beta}'_1$.

In the proof of Theorem 7.3 below, we construct a transverse Cantor foliation for \mathfrak{M} , based on the techniques used in the proof of Theorem 1.3 in [18]. Applying these techniques requires a (possibly) strong restriction on the diameter of the disks transverse to L_0 . In particular, the proof of [18, Theorem 1.3] shows that for the given bound R_1 on path length distances, there exists $\delta_1^* > 0$ such that a Reeb neighborhood with base given by compact set of diameter less than R_1 and transversal of diameter less than δ_1^* admits a transverse Cantor foliation. Let δ_1^* denote this constant.

Note that $\{V_1, \mathfrak{X} - V_1\}$ defines a clopen partition of \mathfrak{X} , and we introduce a constant which bounds the distance between the sets in it. Let $\varepsilon_1 = \min\{\epsilon_H, d_{\mathcal{F}}(V_1, \mathfrak{X} - V_1), \delta_1^*\}$ if $V_1 \subset \mathfrak{X}$ is a proper

inclusion; otherwise set $\varepsilon_1 = \min\{\varepsilon_H, \delta_1^*\}$. Define the constant $\widehat{\delta}_1$ using Proposition 3.5, where

$$(40) \quad \widehat{\delta}_1 = \delta(\varepsilon_1, \widehat{\beta}'_1) \leq \varepsilon_1.$$

This constant is used to conclude that the partial coding functions introduced next are well-defined, and so assures that the Reeb neighborhoods introduced in Section 7 are well-defined as well.

6.2. Coding functions. We begin with a basic notion.

DEFINITION 6.1. *Let $0 < \varepsilon'_1 < \text{diam}_{\mathfrak{X}}(V_1)/2$. An ε'_1 -coding partition of V_1 is a collection of disjoint clopen sets $\mathcal{W}_1 = \{W(1;1), \dots, W(1;\kappa_1)\}$ with $w_0 \in W(1;1)$ whose union is V_1 , satisfying:*

$$(41) \quad \text{diam}_{\mathfrak{X}}(W(1;i)) < \min\{\varepsilon'_1, \widehat{\delta}_1\} \quad \text{for } 1 \leq i \leq \kappa_1.$$

We then have that $\text{diam}_{\mathfrak{X}}(W(1;i)) < \text{diam}_{\mathfrak{X}}(V_1)/2$ so that $\kappa_1 \geq 2$, thus the partition is not trivial. The assumption $\text{diam}_{\mathfrak{X}}(W(1;i)) < \widehat{\delta}_1$ implies that if $w \in W(1;i)$ and $w \in \text{Dom}(\gamma)$ with $\|\gamma\| \leq R'_1$, then $W(1;i) \subset \text{Dom}(\gamma)$.

The index set $\mathcal{A}_1 = \{0, 1, 2, \dots, \kappa_1\}$ of \mathcal{W}_1 is the ‘‘alphabet’’ used to code the action of the induced pseudogroup $\mathcal{G}_{\mathcal{F}}^*(V_1)$ on V_1 . This coding is used to subdivide the sets of \mathcal{W}_1 into smaller sets with constant codes, in a way compatible with the dynamics of the induced pseudogroup $\mathcal{G}_{\mathcal{F}}^*(V_1)$. Let

$$(42) \quad \eta_1 = \min\{\varepsilon_1, d_{\mathfrak{X}}(W(1;i), W(1;j)) \mid i \neq j\} \quad \text{and set} \quad \zeta_1 = \delta(\eta_1, \widehat{\beta}'_1).$$

Note that $\eta_1 > 0$, as V_1 is clopen in \mathfrak{X} , so the constant $\delta(\eta_1, \widehat{\beta}'_1)$ is well-defined by Proposition 3.5.

With the above preparations, we now give a key construction. Given points $w, w' \in \mathcal{N}_1$, let $\gamma_{w,w'}$ be a length minimizing geodesic from $y = \tau(w)$ to $y' = \tau(w')$. Assume that $d_{\mathcal{F}}(y, y') \leq R_1$ then there is an admissible sequence $\mathcal{I}_{w,w'}$ of length $\alpha \leq \widehat{\beta}'_1$ which defines a good plaque-chain covering for $\gamma_{w,w'}$ as in Definition 3.1. Then $\mathcal{I}_{w,w'}$ yields a holonomy map $h_{\mathcal{I}_{w,w'}}$ satisfying $B_{\mathfrak{X}}(w, \widehat{\delta}_1) \subset \text{Dom}(h_{\mathcal{I}_{w,w'}})$. Note that as L_0 has no holonomy, the map $h_{\mathcal{I}_{w,w'}}$ is independent of the choice of $\gamma_{w,w'}$ on the subset $B_{\mathfrak{X}}(w, \widehat{\delta}_1)$ of its domain.

For $w \in \mathfrak{X}$, recall that $\Gamma_w^{V_1, R}$ was introduced in (22). For simplicity on notation in the following, given $\gamma \in \Gamma_w^{V_1, R}$ we let $\gamma \equiv h_{\mathcal{I}_{w,w'}}$ denote the holonomy map defined by the path.

The *partial coding function* for $\Gamma_w^{V_1, R}$ is defined as follows.

DEFINITION 6.2. *Given $0 < R \leq R_1$ and $w \in W(1;i)$ for $1 \leq i \leq \kappa_1$, the $C_{w,i}^R$ -code of $u \in W(1;i)$ is the function $C_{w,u}^{i,R}: \Gamma_w^{V_1, R} \rightarrow \mathcal{A}_1$ defined as*

$$(43) \quad C_{w,u}^{i,R}(\gamma) = k \text{ where } \gamma(u) \in W(1;k).$$

The value $C_{w,u}^{i,R}(\gamma)$ encodes the index k for the path starting at $\tau(u)$, shadowing the path $\gamma_{w,w'}$ and terminating at $\tau(u')$. By the choice of $\widehat{\delta}_1$ and ε'_1 above, $w \in W(1;i)$ implies $W(1;i) \subset \text{Dom}(\gamma)$ and thus the function $C_{w,u}^{i,R}$ is well-defined.

The next result implies that the coding partition to be defined next will consist of clopen subsets. Recall that ζ_1 is defined above in (42).

LEMMA 6.3. *Let $0 < R \leq R_1$, $w \in V_1$ and $u, v \in W(1;i)$ with $d_{\mathfrak{X}}(u, v) < \zeta_1$. Then $C_{w,u}^{i,R}(\gamma) = C_{w,v}^{i,R}(\gamma)$ for all $\gamma \in \Gamma_w^{V_1, R}$. Hence, the function $C_{w,u}^{i,R}$ defined by $C_{w,u}^{i,R}(u) = C_{w,u}^{i,R}$ is locally constant.*

Proof. Let $\gamma \in \Gamma_w^{V_1, R}$, and suppose that $u, v \in W(1;i)$ with $d_{\mathfrak{X}}(u, v) < \zeta_1$. Set $u' = \gamma(u)$ and $v' = \gamma(v)$. Then $d_{\mathfrak{X}}(u, v) < \zeta_1$ implies $d_{\mathfrak{X}}(u', v') < \eta_1$. If $u' \in W(1;i)$ then $d_{\mathfrak{X}}(u', v') < \eta_1 \leq d_{\mathfrak{X}}(W(1;i), W(1;j))$ for all $j \neq i$ implies that $v' \in W(1;j)$. Thus, $C_{w,u}^{i,R}(\gamma_w) = C_{w,v}^{i,R}(\gamma_w)$. \square

We next introduce clopen partitions of the sets $W(1; i)$ chosen in Definition 6.1, so that the coding function along paths of length at most R'_1 as defined by (39) are constant on each set of the partition. For each $1 \leq i \leq \kappa_1$ define the partition $\mathcal{V}(1; i) = \{V(1; i, 1) \cup \dots \cup V(1; i, \kappa_{1,i})\}$ with

$$(44) \quad W(1; i) = V(1; i, 1) \cup \dots \cup V(1; i, \kappa_{1,i}) .$$

where the coding function C_w^{i, R'_1} is constant on each set $V(1; i, j)$, and if $j \neq j'$ then $V(1; i, j)$ and $V(1; i, j')$ have distinct codes. Assume that $w_0 \in V(1; 1, 1)$.

LEMMA 6.4. *For $1 \leq i \leq \kappa_1$ and each $1 \leq j \leq \kappa_{1,i}$ the set $V(1; i, j)$ is clopen. Thus, $\mathcal{V}(1; i)$ is a clopen partition of $W(1; i)$ and the collection $\mathcal{V}(1; 1) \cup \dots \cup \mathcal{V}(1; \kappa_1)$ is a clopen partition of V_1 .*

Proof. Lemma 6.3 implies that each set $V(1; i, \kappa_{1,i})$ is open, and as they form a finite cover of the clopen set $W(1; i)$, they are also clopen sets. \square

If the action of $\mathcal{G}_{\mathcal{F}}$ on \mathfrak{X} is equicontinuous, then it may happen that points $u, u' \in V(1; i, j)$ have the same code for all paths starting at $w \in V_1$, not just those of length at most R'_1 . In fact, this fact is used in [17] to define clopen partitions which are invariant under the action of $\mathcal{G}_{\mathcal{F}}$. However, if the action of $\mathcal{G}_{\mathcal{F}}^*$ is ϵ -expansive, its action separates points by at least ϵ , so will have distinct codes if the diameters of the sets in the clopen partition \mathcal{W}_1 are less than ϵ .

We introduce a further partition of the clopen sets in $\mathcal{V}(1; i)$ according to their coding for paths of length at most $R_1 = 2R'_1$. This additional decomposition is used for the patching arguments appearing in later sections to construct the transverse Cantor foliation for \mathcal{F} .

For $1 \leq i \leq \kappa_1$ and each $1 \leq j \leq \kappa_{1,i}$ let $\mathcal{V}(1; i, j) = \{V(1; i, j, 1) \cup \dots \cup V(1; i, j, \kappa_{1,i,j})\}$ with

$$(45) \quad V(1; i, j) = V(1; i, j, 1) \cup \dots \cup V(1; i, j, \kappa_{1,i,j})$$

where the coding function C_w^{i, R_1} is constant on each set $V(1; i, j, k)$, and if $k \neq k'$ then $V(1; i, j, k)$ and $V(1; i, j, k')$ have distinct codes. Assume that $w_0 \in V(1; 1, 1, 1)$. As before, each set in the partition $\mathcal{V}(1; i, j)$ is clopen and the sets are disjoint.

Note that $V(1; i, j) \subset W(1; i)$ and so satisfies $\text{diam}_{\mathfrak{X}}(V(1; i, j)) \leq \min\{\epsilon'_1, \widehat{\delta}_1\}$ by (41).

6.3. Inductive step. We next extend the above process, so that the diameters of the transverse partition sets decrease in a systematic manner. The results are used in Sections 8 and 9 for the construction of nested Reeb neighborhoods. We first show how to proceed to step 2, which refines the construction above, and the general inductive step follows analogously.

Let $0 < \epsilon_2 < \epsilon_1/2$ be chosen sufficiently small so that the function $\lambda_1(\delta)$ defined in Lemma 4.3 satisfies $\lambda_{\epsilon_2} = \lambda_1(\epsilon_2) > R_1$.

By choice, we have $w_0 \in V(1; 1, 1, 1)$, and let $w_0 \in V_2 \subset V(1; 1, 1, 1)$ be a clopen neighborhood with $\text{diam}_{\mathfrak{X}}(V_2) < \epsilon_2$. Set $\mathcal{N}_2 = L_0 \cap \tau(V_2)$, which is a subnet of \mathcal{N}_1 .

Let $\alpha_2 \geq \alpha_1$ be an integer such that \mathfrak{X} is covered by translates of V_2 using words of length at most α_2 , as given by Lemma 4.1. Define $\beta_2 = \lceil 2\alpha_2 \delta_{\mathcal{U}}^{\mathcal{F}} / \epsilon_{\mathcal{U}}^{\mathcal{F}} \rceil$ and $\theta_2 = (2\alpha_2 + 1) \delta_{\mathcal{U}}^{\mathcal{F}}$, and as before, set

$$(46) \quad R'_2 = 2\theta_2 + \lambda_{\mathcal{F}} \quad , \quad R_2 = 2R'_2 \quad , \quad \widehat{\beta}_2 = \lceil R_2 / \epsilon_{\mathcal{U}}^{\mathcal{F}} \rceil .$$

Choose $\widehat{\beta}'_2$ so that homotopies of paths of length at most $\widehat{\beta}_2$ contain chains of length at most $\widehat{\beta}'_2$. Let

$$(47) \quad \varepsilon_2 < \min\{ \delta_2^* , d_{\mathfrak{X}}(V(1; i), V(1; j)) , d_{\mathfrak{X}}(V_2, \mathfrak{X} - V_2) \mid i \neq j \} ,$$

where δ_2^* is given by Corollary 7.5, and then set

$$(48) \quad \widehat{\delta}_2 = \delta(\varepsilon_2, \widehat{\beta}'_2) .$$

For $\epsilon'_2 = \text{diam}_{\mathfrak{X}}(V_2)/2 < \epsilon_2/2 < \epsilon_1/4$, choose a clopen partition of V_2 as in Definition 6.1, given by $\mathcal{W}_2 = \{W(2; 1), \dots, W(2; \kappa_2)\}$ with $w_0 \in W(2; 1)$ and which satisfies

$$\text{diam}_{\mathfrak{X}}(W(2; i)) < \min\{\epsilon'_2, \widehat{\delta}_2\} \quad \text{for } 1 \leq i \leq \kappa_2 .$$

This yields a new coding alphabet $\mathcal{A}_2 = \{1, 2, \dots, \kappa_2\}$ corresponding to \mathcal{W}_2 . For each $1 \leq i \leq \kappa_2$, use the coding function for the sets in \mathcal{W}_2 to refine \mathcal{W}_2 into its coding partition for the holonomy along paths of lengths at most R'_2 , where $\mathcal{V}(2; i) = \{V(2; i, 1) \cup \dots \cup V(2; i, \kappa_{2,i})\}$ with

$$(49) \quad W(2; i) = V(2; i, 1) \cup \dots \cup V(2; i, \kappa_{2,i}).$$

We assume that $w_0 \in V(2; 1, 1)$, and observe that the clopen sets satisfy

$$V(2; i, j) \subset W(2; i) \subset V(1; 1, 1, 1) \subset W_1 \subset V_1.$$

In particular, this implies that $\text{diam}_{\mathfrak{X}}(V(2; i, j)) \leq \min\{\epsilon'_2, \widehat{\delta}_2\} \leq \epsilon_H/4$.

We also introduce the subdivision of these partitions for the holonomy of paths with length at most $R_2 = 2R'_2$. For $1 \leq i \leq \kappa_2$ and $1 \leq j \leq \kappa_{2,i}$, let $\mathcal{V}(2; i, j) = \{V(2; i, j, 1) \cup \dots \cup V(2; i, j, \kappa_{2,i,j})\}$ with

$$V(2; i, j) = V(2; i, j, 1) \cup \dots \cup V(2; i, j, \kappa_{2,i,j}), \quad w_0 \in V(2; 1, 1, 1)$$

where the coding function is constant along paths of length at most R_2 on each set $V(2; i, j, k)$, and if $k \neq k'$ then $V(2; i, j, k)$ and $V(2; i, j, k')$ have distinct codes. As before, each set in the partition $\mathcal{V}(2; i, j)$ is clopen and disjoint.

The above process can now be repeated recursively to obtain:

PROPOSITION 6.5. *Let $w_0 \in \mathfrak{T}_1 \subset \mathfrak{X}$ such that the leaf L_0 it determines is without holonomy. Then each $\ell \geq 1$, there exists:*

- (1) clopen neighborhoods $w_0 \in V_\ell \subset V_{\ell-1} \subset \dots \subset V_1 \subset \mathfrak{T}_1$ with $\text{diam}_{\mathfrak{X}}(V_i) < \epsilon_H/2^i$
- (2) constants $\epsilon_\ell, \alpha_\ell, \theta_\ell, R_\ell, \widehat{\beta}_\ell, \widehat{\delta}_\ell$ such that $\lambda_{\epsilon_\ell} = \lambda_1(\epsilon_\ell) > R_{\ell-1}$
- (3) nets $\mathcal{N}_\ell = L_0 \cap \tau(V_\ell)$ which are λ_{ϵ_ℓ} -separated and θ_ℓ dense
- (4) a clopen partition \mathcal{W}_ℓ of V_ℓ with elements $W(\ell; i)$ labeled by $\mathcal{A}_\ell = \{1, 2, \dots, \kappa_\ell\}$
- (5) a clopen partition $\mathcal{V}(\ell; i)$ of each $W(\ell; i)$ with elements $V(\ell; i, j)$ whose elements are clopen sets with $\text{diam}_{\mathfrak{X}}(V(\ell; i, j)) < \min\{\widehat{\delta}_\ell, \epsilon_H/4^\ell\}$ for $1 \leq j \leq \kappa_{\ell,i}$, and the coding function is constant on each $V(\ell; i, j)$ along paths of length at most R'_ℓ
- (6) a clopen partition $\mathcal{V}(\ell; i, j)$ of each clopen set $V(\ell; i, j)$ such that the coding function is constant on each $V(\ell; i, j, k)$ along paths of length at most $R_\ell = 2R'_\ell$ for $1 \leq k \leq \kappa_{\ell,i,j}$.

The clopen sets $V(\ell; i, j)$ of the partition $\mathcal{V}(\ell; i)$ will be used to create Reeb neighborhoods in the following Sections 7 and 8, by translating them along the holonomy of paths of length at most R'_ℓ .

Note that this technique of decomposing a descending chain of clopen transversals V_ℓ for $\ell \geq 1$ into code blocks $\mathcal{V}(\ell; i)$ is *implicit* in the studies of \mathbb{R}^n -actions by Bellissard, Benedetti and Gambaudo [11, Proposition 2.39] and by Forrest in [27], and in the sketch of the proof of Theorem 3.1 in [12] for case when the leaves of \mathcal{F} are defined by a transitive action of a connected Lie group.

7. TRANSVERSE CANTOR FOLIATIONS

A *Cantor foliation* \mathcal{H} on a continuum Ω is a “continuous decomposition” of it into Cantor sets, which are the “leaves” of \mathcal{H} . Cantor foliations arise naturally in many dynamical and geometric contexts. In this work, we require Cantor foliations which are “transverse” to the foliation \mathcal{F} of a matchbox manifold \mathfrak{M} , as made precise in Definition 7.1 below. All constructions in the literature which build up inverse limit representations for a particular class of matchbox manifolds, such as the Williams solenoids, tiling spaces, and suspensions of minimal actions on Cantor sets [5, 7, 10, 11, 12, 17, 18, 43, 53, 60] use such a transverse Cantor foliation either explicitly or implicitly. The branched manifold quotients of a matchbox manifold are obtained via transverse projections along subsets of the leaves of \mathcal{H} that are the “axes of collapsing” in foliated compact subsets of \mathfrak{M} . As such, the existence of such \mathcal{H} is fundamental to the proof of Theorem 1.1.

The existence of a transverse Cantor foliation \mathcal{H} can be part of the given data, such as when a matchbox manifold is a fiber bundle over a closed manifold whose fibers are Cantor sets. In this case, the leaves of \mathcal{H} are given by Cantor fibers of the bundle. For a tiling space $\Omega_{\mathbf{T}}$ formed from a tiling of \mathbb{R}^n with finite local complexity, the foliation by tilings is defined by a free action of \mathbb{R}^n

on $\Omega_{\mathbf{T}}$, and the Cantor foliation is defined by the choice of punctures in each of the proto-tiles, then taking their translates under the \mathbb{R}^n -action. Finally, a Cantor foliation may be determined by dynamical properties of the ambient space, such as in the case of Williams solenoids [63, 65] where the Cantor foliation is given by intersections of the (unstable) attractor with leaves of the stable foliation. This dynamical construction has been generalized to the notion of a ‘‘Smale space’’ as introduced by Ruelle [52], and studied in the works of Putnam [6, 49, 50, 51] and Wieler [62].

For the general matchbox manifold, the existence of a Cantor foliation on the ‘‘Big Boxes’’ in \mathfrak{M} is shown in the authors’ work [18], where this is proven using local constructions based entirely on intrinsic techniques. The results of that paper are used in the rest of this section, and in Sections 8 and 9, where we extend the results to show that a Cantor foliation can always be defined on a minimal matchbox manifold \mathfrak{M} without holonomy. Moreover, the Cantor foliation on \mathfrak{M} can be constructed to be compatible with the constructions in Section 6.

In Section 7.1, we give a rigorous definition of a Cantor foliation, then in Section 7.2 we recall a main result proved in [18] which gives the existence of a Cantor foliation on a neighborhood of a compact *path-connected* subset K_x in \mathfrak{M} , called a *Reeb neighborhood*. Section 7.3 shows the existence of a covering of \mathfrak{M} by ‘‘Big Boxes’’, which are Reeb neighborhoods with a transverse Cantor foliation.

7.1. Definition of transverse Cantor foliations. We assume there is given a regular covering $\{U_i \mid 1 \leq i \leq \nu\}$ of \mathfrak{M} by foliation charts, as in Proposition 2.6, with charts $\varphi_i: \bar{U}_i \rightarrow [-1, 1]^n \times \mathfrak{T}_i$ where $\mathfrak{T}_i \subset \mathfrak{X}$ is a clopen subset. Moreover, by construction, each chart admits a foliated extension $\widehat{\varphi}_i: \widehat{U}_i \rightarrow (-2, 2)^n \times \mathfrak{T}_i$ where $\bar{U}_i \subset \widehat{U}_i \subset \mathfrak{M}$ is an open neighborhood of the closure \bar{U}_i and $\widehat{\varphi}_i|_{\bar{U}_i} = \varphi_i$. For a clopen set $V \subset \mathfrak{T}_i$ set

$$(50) \quad \mathfrak{U}_i^V = \pi_i^{-1}(V) \subset \bar{U}_i \quad ; \quad \widehat{\mathfrak{U}}_i^V = \widehat{\varphi}_i^{-1}(V) \subset \widehat{U}_i .$$

DEFINITION 7.1. *Let \mathfrak{M} be a matchbox manifold, and $\mathfrak{B} \subset \mathfrak{M}$ a closed subset. An equivalence relation \approx on \mathfrak{B} is said to define a transverse Cantor foliation \mathcal{H} of \mathfrak{B} if for each $x \in \mathfrak{B}$, the class $\mathcal{H}_x = \{y \in \mathfrak{B} \mid y \approx x\}$ is a Cantor set. Moreover, assume that for each $x \in \mathfrak{B}$, there exists:*

- (1) $1 \leq i_x \leq \nu$ with $x \in U_{i_x}$,
- (2) a clopen subset $V_x \subset \mathfrak{T}_{i_x}$ with $w_x = \pi_{i_x}(x) \in V_x$ and $\mathfrak{U}_{i_x}^{V_x} \subset \mathfrak{B}$;
- (3) a homeomorphism into $\Phi_x: [-1, 1]^n \times V_x \rightarrow \widehat{U}_{i_x}$ such that for the point w_x we have

$$\Phi_x(\xi, w_x) = \widehat{\varphi}_i^{-1}(\xi, w_x) \text{ for } \xi \in (-1, 1)^n,$$

- (4) for $\xi \in (-1, 1)^n$ and $z = \widehat{\varphi}_i^{-1}(\xi, w_x)$, the image $\Phi_x(\{\xi\} \times V_x) = \mathcal{H}_z \cap \widehat{\mathfrak{U}}_{i_x}^{V_x}$.

The leaves of the ‘‘foliation’’ \mathcal{H} are defined to be the equivalence classes \mathcal{H}_x of \approx in \mathfrak{B} .

We give some additional remarks on the conditions in the definition above. First, the index i_x in Condition 7.1.1 selects a foliation chart containing x in its interior. Then Condition 7.1.2 specifies a clopen subset V_x which defines a ‘‘plaque’’ in the leaf \mathcal{H}_x , so that \mathfrak{B} is a union of such plaques.

Conditions 7.1.3 and 7.1.4 specify the leaf \mathcal{H}_x containing x as the graph of the function $\Phi_x: \{\xi\} \times V_x \rightarrow \widehat{U}_{i_x}$ where $\xi = \lambda_{i_x}(x) \in (-1, 1)^n$ is the horizontal coordinate of x for the coordinates φ_{i_x} . In particular, as V_x is clopen, it is a Cantor set, so the leaf \mathcal{H}_x is a Cantor set.

Note that the function Φ_x depends on the index i_x as the leaves of \mathcal{H} are not assumed to coincide with the local ‘‘vertical’’ foliation of U_{i_x} defined by the transverse coordinate. The function Φ_x gives the necessary ‘‘adjustment’’ to the local vertical foliation. Also note, the image of the graph of Φ_x through $z \in \bar{U}_{i_x}$ is contained in the open neighborhood \widehat{U}_{i_x} of \bar{U}_{i_x} but need not be contained in \bar{U}_{i_x} .

Recall that the constants $0 < \delta_{\mathcal{U}}^{\mathcal{F}} < \lambda_{\mathcal{F}}/5$ were introduced in Proposition 2.6, and there is a constant $\lambda_{\mathcal{F}}^* \leq \lambda_{\mathcal{F}}/5$ introduced in Section 15.4 of [18], where $\lambda_{\mathcal{F}}^*$ is chosen so that the leafwise balls of radius $\lambda_{\mathcal{F}}^*$ are ‘‘approximately Euclidean’’ as required in the constructions there. We can also assume without loss of generality that $\lambda_{\mathcal{F}}^* < \epsilon_{\mathcal{U}}^{\mathcal{F}} < \delta_{\mathcal{U}}^{\mathcal{F}}$.

7.2. Reeb neighborhoods. Cantor foliations as in Definition 7.1 are constructed using the “Big Box” Theorem [18, Theorem 1.3] as follows. First we define *Reeb neighborhoods* as in [18, Section 6].

For $x \in \mathfrak{M}$, let L_x be the leaf containing x , and set $w = \pi_{i_x}(x) \in \mathfrak{X}_{i_x}$ for some $1 \leq i_x \leq \nu$. We assume that L_x is without holonomy.

Let $Q_x \subset L_x$ be a compact connected subset with $x \in \text{int}(Q_x)$ and assume that Q_x is a union of plaques. Then $Q_x \cap \mathcal{T}$ is a net in Q_x , and for each $z \in Q_x \cap \mathcal{T}$, there is a path γ_z in Q_x between x and z . Let $h_{x,z}$ be the holonomy homeomorphism along $\gamma_{x,z}$ as constructed in Section 3.3.

For a clopen neighborhood $V_x \subset \mathfrak{X}_{i_x}$ of w , we assume that $V_x \subset \text{Dom}(h_{x,z})$ for each $z \in Q_x \cap \mathcal{T}$. Then set $V_x^z = h_{x,z}(V_x)$, which is well-defined, as L_x has no holonomy. Let $1 \leq i_z \leq \nu$ denote the index of the foliation chart containing z , which is the last index for the plaque-chain between x and z used to define $h_{x,z}$. The *Reeb neighborhood* of Q_x associated to V_x is a union of saturated neighborhoods, defined as in (50):

$$(51) \quad \mathfrak{N}_{Q_x}^{V_x} = \bigcup_{z \in Q_x \cap \mathcal{T}} \mathfrak{U}_{i_z}^{V_x^z}.$$

Note that each path-connected component in $\mathfrak{N}_{Q_x}^{V_x}$ is a union of plaques. Here is one of the main results of [18].

THEOREM 7.2 (Big Box). *Let \mathfrak{M} be a matchbox manifold, $x \in \mathfrak{M}$ and suppose that L_x is a leaf without holonomy. Let Q_x be a compact connected subset of L_x with $x \in \text{int}(Q_x)$ and assume that Q_x is a union of plaques. Then there exists a clopen set $V_x \subset \mathfrak{X}$, a closed connected subset \widehat{Q}_x which is a union of plaques with $Q_x \subset \text{int}(\widehat{Q}_x)$, and a foliated homeomorphic inclusion $\Phi: \widehat{Q}_x \times V_x \rightarrow \mathfrak{M}$ such that the images $\Phi \left\{ \{y\} \times V_x \mid y \in \widehat{Q}_x \right\}$ form a continuous family of Cantor transversals on a neighborhood of $\mathcal{F} \mathfrak{N}_{Q_x}^{V_x}$.*

Proof. We give a sketch of the proof from [18], so that we can refer to the steps involved in extending this result in the following. Given Q_x , we want to choose a neighborhood V_x small enough so that Theorem 7.2 follows from Theorem 1.3 of [18].

We require some notations and constants from [18]. The reader can just take these as given, as the reasons for and details of their choices is part of the most technical aspects of the work [18]. The constant $\varepsilon_0 > 0$ defined in [18] is chosen so that there is a prescribed bound on the metric distortions in charts. This is used in [18, Section 15.5] to define a “transverse” constant $r^* > 0$, such that if x and y are points in the same plaque, and x' and y' are points in a plaque in the same chart which is at transverse distance at most r^* from the given one, and x' and y' have the same coordinates as x and y respectively, then the metric distances between x and y , and between x' and y' , differ by at most $\varepsilon_0 \lambda_{\mathcal{F}}^*$, where ε_0 is very small and depends on the geometry of leaves. The definition of $\varepsilon_0 > 0$ as given in [18, Section 15.3] is the most delicate part of the estimates in that work. In particular, we have that $\varepsilon_0 < 1/2000$ and $\lambda_{\mathcal{F}}^* < \epsilon_{\mathcal{U}}^{\mathcal{F}}$, so that $\varepsilon_0 \lambda_{\mathcal{F}}^* < \epsilon_{\mathcal{U}}^{\mathcal{F}}/2000$.

We require the following consequence of the choice of r^* and ε_0 . Suppose that $z \in U_i \cap U_j$ for $i \neq j$. Then the divergence $\text{div}(z, i, j, r) \leq \varepsilon_0 \lambda_{\mathcal{F}}^*$ as defined in [18, Section 14.1]. This technical statement has a simple geometric interpretation. The set of points in U_i with the same horizontal coordinate as z defines a “standard section” $\mathfrak{Z}_{z,i} \subset U_i$ (see [18, (17.2)]). Likewise, z defines a standard section $\mathfrak{Z}_{z,j} \subset U_j$. Let $x' \in \mathfrak{Z}_{z,i}$ and $y' \in \mathfrak{Z}_{z,j}$ lie on the same plaque in $U_i \cap U_j$ and suppose this plaque is within r^* from the plaque containing z . Then the leafwise distance $d_{\mathcal{F}}(x', y') \leq \varepsilon_0 \lambda_{\mathcal{F}}^* < \epsilon_{\mathcal{U}}^{\mathcal{F}}/2000$.

For Q_x given, let R denote its diameter so that $Q_x \subset D_{\mathcal{F}}(x, R)$. By Proposition 2.6.4, the radii of the plaques in the atlas \mathcal{U} are uniformly bounded by $\delta_{\mathcal{U}}^{\mathcal{F}}$, so the saturation of Q_x is contained in the ball $D_{\mathcal{F}}(x, R + 4\delta_{\mathcal{U}}^{\mathcal{F}})$. Form the closed plaque-saturation of $D_{\mathcal{F}}(x, R + 4\delta_{\mathcal{U}}^{\mathcal{F}})$ and denote the set obtained by \widehat{Q}'_x . Then the diameter of \widehat{Q}'_x is bounded by $R + 8\delta_{\mathcal{U}}^{\mathcal{F}} < R + 2\lambda_{\mathcal{F}}$.

Apply Proposition 3.5 for $\epsilon = r_*/2$ to conclude that there exists $\delta_* = \delta(r_*/2, R + 2\lambda_{\mathcal{F}})$ such that if $x \in \widehat{V}_x$ is a clopen neighborhood with $\widehat{V}_x \subset B_{\mathfrak{X}}(w_x, \delta_*)$, then for any path γ with initial point x and length at most $R + 2\lambda_{\mathcal{F}}$, the holonomy translate $h_\gamma(\widehat{V}_x)$ of the set \widehat{V}_x has diameter less than $r_*/2$.

For $\widehat{V}_x \subset B_{\mathfrak{X}}(w_x, \delta_*)$ then by the ‘‘Big Box’’ Theorem [18] there exists a Cantor foliation on a subset $\widehat{Q}_x \subset \widehat{Q}'_x$ with $Q_x \subset \text{int}(\widehat{Q}_x)$ and there is a continuous map $\Phi_x : \widehat{Q}_x \times \widehat{V}_x \rightarrow \mathfrak{M}$. Since $\widehat{Q}_x \subset L_x$ is a subset of a leaf without holonomy, the map Φ_x is injective on \widehat{Q}_x , and so there exists $V_x \subset \widehat{V}_x$ such that the restriction $\Phi_x|_{\widehat{Q}_x \times V_x}$ is a homeomorphism onto its image. As $Q_x \subset \widehat{Q}_x$, this completes the proof of Theorem 7.2. \square

7.3. ‘‘Big Box’’ Coverings. We next use Theorem 7.2 to construct coverings of \mathfrak{M} by ‘‘Big Boxes’’ which are adapted to the coding partitions of \mathfrak{X} defined in Section 6. This will use the constant $\delta_{\mathcal{U}}^{\mathcal{F}}$ introduced in Proposition 2.6, the constant ϵ_H as defined by (34), and $\lambda_{\mathcal{F}}^*$ introduced in Section 7.1.

THEOREM 7.3. *Let \mathfrak{M} be a minimal matchbox manifold without holonomy, and there is given a clopen subset $V_1 \subset \mathfrak{X}$ with $\text{diam}(V_1) < \epsilon_H/2$. Choose $w_0 \in V_1$, and let L_0 denote the leaf containing $x_0 = \tau(w_0)$. Then there exists a constant $\delta_1^* > 0$ such that for any clopen partition $\mathcal{W}_1 = \{W(1; 1), \dots, W(1; \kappa_1)\}$ of V_1 with $\text{diam}(W(1; i)) < \delta_1^*$ for $1 \leq i \leq \kappa_1$, there exists:*

- (1) compact connected subsets $\{K_1, K_2, \dots, K_{\kappa_1}\}$ of L_0 ;
- (2) for each $1 \leq i \leq \kappa_1$, a compact connected subset $\widehat{K}_i \subset L_0$ which is the closure of a union of plaques, such that $K_i \subset \text{int}(\widehat{K}_i)$;
- (3) for each $1 \leq i \leq \kappa_1$, a foliated homeomorphic inclusion $\Phi_i : \widehat{K}_i \times W(1; i) \rightarrow \mathfrak{M}$ which form a continuous family of Cantor transversals for $\mathcal{F}|_{\mathfrak{N}_{K_i}^{W(1; i)}}$;

such that the interiors $U(1, i) \equiv \Phi_i(\text{int}(K_i) \times W(1; i))$ form an open covering of \mathfrak{M} .

Proof. The idea of the proof is use the Voronoi cells defined by the net $\mathcal{N}_{w_0}^{V_1} = L_0 \cap \tau(V_1)$ to construct Reeb neighborhoods as in (51) which cover \mathfrak{M} , and then choose an appropriate subcover.

We assume without loss of generality that $V_1 \subset \mathfrak{T}_1 \subset \mathfrak{X}$, so that $\tau(V_1) = \tau_1(V_1) \subset \mathcal{T}_1$.

Let α_1 be the integer defined by Lemma 4.1 for the clopen set V_1 , and let $\theta_1 = (2\alpha_1 + 1)\delta_{\mathcal{U}}^{\mathcal{F}}$ be the density constant for the net $\mathcal{N}_{w_0}^{V_1}$ as established in Corollary 4.6. Set $R'_1 = 2\theta_1 + \lambda_{\mathcal{F}}$ and $R_1 = 2R'_1$.

For $w \in V_1$, let $\mathcal{O}(V_1, w) = \pi(\mathcal{N}_w^{V_1}) \subset V_1$ denote the orbit of w in V_1 which is dense by minimality. Then for $u \in \mathcal{O}(V_1, w)$, set $z = \tau(u) \in L_w$, and recall that $\mathcal{C}_w^{V_1}(z)$ denotes the Voronoi cell in L_w as defined by (27) which contains z in its interior. Lemma 4.11 implies for all $w \in V_1$ and $z \in \mathcal{N}_w^{V_1}$,

$$(52) \quad \mathcal{C}_w^{V_1}(z) \subset \mathcal{S}_w^{V_1}(z) \subset B_{\mathcal{F}}(z, 3\theta_1).$$

By assumption, for each $w \in V_1$ the leaf L_w has no holonomy. Thus, given $y, z \in \mathcal{N}_w^{V_1}$ with $u = \pi(y) \in V_1$ and $v = \pi(z)$, there is a geodesic path $\gamma_{y,z}$ in L_x joining them which defines a holonomy transformation $h_{u,v}$ defined on a sufficiently small open neighborhood of u . More precisely, Lemma 3.13 implies that for $\epsilon' = \delta_{\mathcal{U}}^{\mathcal{F}}/4R_1$, there exists $\delta'_1 > 0$ so that if $d_{\mathfrak{X}}(u, u') < \delta'_1$ then for a length-minimizing geodesic γ contained in $B_{\mathcal{F}}(x, R_1)$, there is a shadowing piecewise geodesic γ' as in the proof of Lemma 3.13 from y' to y' with the length estimate

$$(53) \quad \|\gamma\| - \delta_{\mathcal{U}}^{\mathcal{F}}/4 \leq \|\gamma'\| \leq \|\gamma + \delta_{\mathcal{U}}^{\mathcal{F}}/4\|.$$

For the constant $\epsilon = \min\{\epsilon_H/2, \delta'_1/2\}$, Proposition 3.5 implies there exists $\delta''_1 > 0$ such that for $y, z \in \mathcal{N}_w^{V_1}$, $d_{\mathcal{F}}(y, z) \leq R_1$ implies $B_{\mathfrak{X}}(u, \delta''_1) \subset \text{Dom}(h_{y,z})$, and $h_{y,z}(B_{\mathfrak{X}}(u, \delta''_1)) \subset B_{\mathfrak{X}}(v, \epsilon)$.

Next, observe that $\pi(B_{\mathcal{F}}(x, R'_1) \cap \mathcal{T}) \subset \mathcal{N}_w^{V_1} \subset \mathfrak{X}$ is a finite subset for all $w \in V_1$. As V_1 is compact, its cardinality has a uniform upper bound for $w \in V_1$. It follows that there exists $\delta'''_1 \leq \delta''_1$ so that for all $w \in V_1$, then for a clopen neighborhood of w , $V \subset V_1$ with $\text{diam}_{\mathfrak{X}}(V) \leq \delta'''_1$,

$$(54) \quad \{h_{w,u}(V) \mid u \in \pi(B_{\mathcal{F}}(x, R'_1) \cap \mathcal{T})\}$$

forms a disjoint collection of clopen sets in \mathfrak{X} . Thus, there is a Reeb neighborhood of the plaque-saturation of $B_{\mathcal{F}}(x, \theta_1 + \lambda_{\mathcal{F}})$, defined as in (51).

DEFINITION 7.4. For each $x \in \mathcal{N}_w^{V_1}$, let $K_x \subset L_w$ denote the closure of the plaque-saturation of $B_{\mathcal{F}}(x, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}})$, and let \widehat{K}_x denote the closure of the plaque-saturation of $B_{\mathcal{F}}(x, \theta_1 + \lambda_{\mathcal{F}})$, so that $K_x \subset \text{int}(\widehat{K}_x)$ as $4\delta_{\mathcal{U}}^{\mathcal{F}} < \lambda_{\mathcal{F}}$.

We next define a constant $\delta_1 > 0$ which sets the scale for the coding partition of V_1 . Theorem 7.2 for $Q_x = \widehat{K}_x$ implies that there exists $0 < \delta_1 \leq \delta_1'''$ so that if $W \subset V_1$ is a clopen neighborhood $w \in V_1$ with $\text{diam}_{\mathfrak{X}}(W) < \delta_1$, and \widehat{K}_x is defined as previously for $x = \tau(w)$, then the corresponding Reeb neighborhood $\mathfrak{N}_{\widehat{K}_x}^W$ admits a Cantor foliation. Moreover, from the proof of Theorem 7.2 in [18], the choice of δ_1 is uniform for all $w \in V_1$. We use this uniformity as follows.

For $w_0 \in V_1$ given, consider the collection $\{B_{\mathfrak{X}}(w, \delta_1/4) \cap V_1 \mid w \in \mathcal{O}(V_1, w_0)\}$. This forms a covering of V_1 by open sets, as $\mathcal{O}(V_1, w_0)$ is dense in V_1 . We modify this covering to obtain a covering of V_1 by clopen subsets. Note that the closure $D_{\mathfrak{X}}(w, \delta_1/4) = \overline{B_{\mathfrak{X}}(w, \delta_1/4)}$ is compact with $D_{\mathfrak{X}}(w, \delta_1/4) \subset B_{\mathfrak{X}}(w, \delta_1/2)$, so there exists a finite covering of $D_{\mathfrak{X}}(w, \delta_1/4)$ by clopen subsets of $B_{\mathfrak{X}}(w, \delta_1/2)$. Apply this remark to each set in the covering $\{B_{\mathfrak{X}}(w, \delta_1/4) \cap V_1 \mid w \in \mathcal{O}(V_1, w_0)\}$, and we obtain a countable covering of V_1 by clopen sets with diameters less than δ_1 . Since V_1 is compact, we can choose a finite subcover $\{W_1, \dots, W_m\}$ of this clopen cover. This finite subcovering need not be disjoint, but we can use it to form a disjoint clopen covering as follows:

$$(55) \quad \widehat{W}_1 = W_1, \widehat{W}_2 = W_2 - W_1, \dots, \widehat{W}_i = W_i - \bigcup_{k=1}^{i-1} W_k, \dots$$

If some of \widehat{W}_i are empty, we discard them and renumber the remaining sets, thus obtaining a clopen partition $\widehat{W}_1 = \{\widehat{W}_1, \dots, \widehat{W}_s\}$ of V_1 where $s \leq m$. Then define:

$$\delta_1^* < \min\{ \delta_1, d_{\mathfrak{X}}(\widehat{W}_i, \widehat{W}_j) \mid i, j = 1, \dots, s \}.$$

Next, choose a clopen partition $\mathcal{W}_1 = \{W(1; 1), \dots, W(1; \kappa_1)\}$ of V_1 with $\text{diam}_{\mathfrak{X}}(W(1; i)) \leq \delta_1^*$. Note that every subset of V_1 of diameter less than δ_1^* must be contained in one of the sets \widehat{W}_i , for some $1 \leq i \leq s$, and so for each $1 \leq i \leq \kappa_1$, there is a unique $1 \leq m_i \leq s$ with $W(1; i) \subset \widehat{W}_{m_i}$.

For $1 \leq i \leq \kappa_1$, choose $x_i \in \tau(W(1; i)) \cap L_0$ and set $w_i = \pi(x_i) \in W(1; i)$.

Then set K_i to be the closed plaque-saturation of $B_{\mathcal{F}}(x_i, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}})$, and let $\widehat{K}_i = \widehat{K}_{x_i}$ be defined as in Definition 7.4. By the choice of δ_1^* above we conclude that there exists \widehat{K}'_i with $\widehat{K}_i \subset \text{int}(\widehat{K}'_i)$ and a foliated homeomorphic inclusion $\Phi_i: \widehat{K}'_i \times W(1; i) \rightarrow \mathfrak{M}$ which restricts to a continuous family of Cantor transversals for $\mathcal{F}|_{\mathfrak{N}_{\widehat{K}_i}^{W(1; i)}}$.

To conclude the proof of Theorem 7.3, it remains to show that the sets $\text{int}(\Phi_i\{K_i \times W(1; i)\})$ yield an open covering of \mathfrak{M} . Recall that for any $w \in V_1$ the leaf L_w admits a Voronoi decomposition as in (28), so it suffices to show that for each $y \in \mathcal{N}_w^{V_1}$ the cell $\mathcal{C}_w^{V_1}(y) \subset \text{int}(\Phi_i(K_i \times W(1; i)))$ for some $1 \leq i \leq \kappa_1$.

Now let $y \in \mathcal{N}_w^{V_1}$ with $w = \pi(y) \in V_1$ then there exists $1 \leq i \leq \kappa_1$ for which $w \in W(1; i)$. Then $\mathcal{C}_w^{V_1}(y) \subset B_{\mathcal{F}}(y, \theta_1)$ by the choice of θ_1 , so it suffices to show that $B_{\mathcal{F}}(y, \theta_1) \subset \Phi_i(K_i \times W(1; i))$.

By construction, we have that $B_{\mathcal{F}}(x_i, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}}) \subset \Phi_i(K_i \times W(1; i))$. By the choice of δ_1^* above, given a geodesic path starting at $y \in \tau(W(1; i))$ with length at most θ_1 , there is a piecewise-geodesic path starting at x_i with length at most $\theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}}/4$ which shadows the given geodesic. The endpoint of the shadowing curve is contained in $B_{\mathcal{F}}(x_i, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}})$, which yields the inclusion $B_{\mathcal{F}}(y, \theta_1) \subset \Phi_i(K_i \times W(1; i))$ as was to be shown. \square

A similar argument may of course be repeated for any clopen subset $V_\ell \subset V_1$. We assume the clopen set $V_\ell \subset V_1$ is given with $w_0 \in V_\ell$. Let $\theta_\ell = (2\alpha_\ell + 1)\delta_{\mathcal{U}}^{\mathcal{F}}$ be the density constant for the net $\mathcal{N}_\ell = L_0 \cap \tau(V_\ell)$ as established in Corollary 4.6. Set $R_\ell = 4\theta_\ell + \lambda_{\mathcal{F}}$ and $R'_\ell = 2\theta_\ell + \lambda_{\mathcal{F}}$.

COROLLARY 7.5. *There exists a constant $\delta_\ell^* > 0$ such that for any clopen subset $W \subset V_\ell$ of $\text{diam}_{\mathfrak{X}}(W) < \delta_\ell^*$ and $w \in W$, the collection $\{h_{w,u}(V) \mid u \in \pi(B_{\mathcal{F}}(x, R'_\ell) \cap \mathcal{T})\}$ forms a disjoint collection of clopen sets in \mathfrak{X} .*

REMARK 7.6. Observe that in the proof above, we cover the leaves of \mathcal{F} with plaque-saturations of the balls $B_{\mathcal{F}}(y, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}})$ centered at points $y \in V_1$. The fact that these sets provide a covering of \mathfrak{M} is a consequence of Corollary 4.6 and the definition of R'_1 , which implies that, for each $y \in \mathcal{N}_w^{V_1}$, the ball $B_{\mathcal{F}}(y, \theta_1)$ contains the Voronoi cell $C_w^{V_1}(y)$, and these Voronoi cells form a covering of the leaf L_w . Note that the balls $B_{\mathcal{F}}(y, \theta_1)$ do not necessarily provide an efficient covering, as for $y, z \in \mathcal{N}_w^{V_1}$ the intersection $B_{\mathcal{F}}(y, \theta_1) \cap B_{\mathcal{F}}(z, \theta_1)$ may be “large”. It would be more efficient to use small neighborhoods of the Voronoi cells $C_w^{V_1}(y)$ themselves to form the covering, except that in the construction of the Reeb neighborhoods, there is no guarantee that the Cantor foliation of the neighborhoods of the form $\Phi_i \{K_i \times W(1; i)\}$ has bounded distortion when used to translate the Voronoi cells between leaves.

8. EXISTENCE OF NICE STABLE TRANSVERSALS

Theorem 7.3 shows that a minimal matchbox manifold \mathfrak{M} without holonomy has a finite covering by Big Box neighborhoods, each of which is endowed with a transverse Cantor foliation. However, on the overlaps of the sets in this covering, the Cantor foliations need not agree, so the union of the leaves of these foliations do not define a Cantor foliation on all of \mathfrak{M} . In this section, we give an inductive procedure for modifying the construction of the transverse Cantor foliations on the Big Boxes in a covering, to obtain a Cantor foliation defined on all of \mathfrak{M} , which also adapted to a given initial coding partition on the transversal.

The approach we take here, is to recall the construction of the transverse Cantor foliations on the Reeb neighborhoods from the authors’ work [18], which reduces the construction to the existence of a *nice stable transversal* on each, which is a finite collection $\mathcal{X} = \{\mathfrak{Z}(\xi, i_z, W(1; i)^z) \mid 1 \leq i \leq p\}$ of transversals which intersect each leaf in a net whose separation and spanning constants (d_1, d_2) are sufficiently small. The leaves of the Cantor foliation on each Big Box neighborhood are then defined as collections of points with the same barycentric coordinates in each simplex of the foliated simplicial decomposition defined by the leafwise triangulations.

The constructions of this section require patience to proceed through the multiple steps required. However, the idea is simple in principle, in that we start with a nice stable transversal defined on a subset of an initial Reeb neighborhood, defined using a coding partition of the transversal V_1 , then extend this to its intersections with neighboring Reeb neighborhoods also defined by the coding partition, and observe that the extension so obtained is well-defined. We then recursively repeat this argument with intersections of the previous results of this procedure, and so obtain a recursive procedure for obtaining a nice stable transversal on all of \mathfrak{M} .

8.1. Reeb neighborhoods from partitions. Assume that \mathfrak{M} is a minimal matchbox manifold without holonomy. Given a clopen subset $V_1 \subset \mathfrak{X}$ with $\text{diam}_{\mathfrak{X}}(V_1) < e_H/2$ and $w_0 \in V_1$, let L_0 denote the leaf containing $x_0 = \tau(w_0)$. Let $\delta_1^* > 0$ be the constant defined in the proof of Theorem 7.3.

Let $\mathcal{W}_1 = \{W(1; 1), \dots, W(1; \kappa_1)\}$ be a clopen partition of V_1 with $\text{diam}_{\mathfrak{X}}(W(1; i)) < \delta_1^*$ for all $1 \leq i \leq \kappa_1$ as constructed in the proof of Theorem 7.3, with “basepoints” $x_i \in \tau(W(1; i)) \cap L_0$ for which $w_i = \pi(x_i) \in W(1; i)$. In addition, for each $1 \leq i \leq \kappa_1$, there is a unique $1 \leq m_i \leq s$ with $W(1; i) \subset \widehat{W}_{m_i}$, and K_i is the closed plaque-saturation of $B_{\mathcal{F}}(x_i, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}})$, and $\widehat{K}_i = \widehat{K}_{x_i}$ is defined as in Definition 7.4. Recall the constant $R'_1 = R_1 = 2\theta_1 + \lambda_{\mathcal{F}}$ was defined by (39).

Then there are coding partitions derived from the partition \mathcal{W}_1 as in Section 6.2. For each $1 \leq i \leq \kappa_1$, we have $\mathcal{V}(1; i) = \{V(1; i, j) \mid 1 \leq j \leq \kappa_{1,i}\}$ is a clopen partition of $W(1, i)$ such that the coding function is constant on each $V(1; i, j)$ along paths of length at most R'_1 . Also recall the refinements of these partitions, where the collection $\mathcal{V}(1; i, j)$ is a clopen partition of the clopen set $V(1; i, j)$, where for each $1 \leq k \leq \kappa_{1,i,j}$, the coding function is constant on $V(1; i, j, k)$ along paths of length at most $R_1 = 2R'_1$.

Next choose “basepoints” for these refined partitions. For each $1 \leq i \leq \kappa_1$, $1 \leq j \leq \kappa_{1;i}$, let $v(1; i, j) \in V(1; i, j)$ be a point such that $x(1; i, j) = \tau(v(1; i, j)) \in L_0$, and let $\widehat{K}_{1;i,j}$ be the path-connected subset of $\mathfrak{N}_{\widehat{K}_i}^{W(1;i)}$ containing this point. Denote by $\mathfrak{N}_{1;i,j} \subset \mathfrak{N}_{\widehat{K}_i}^{W(1;i)}$ the Reeb neighborhood formed from the base $\widehat{K}_{1;i,j}$ and transversal $V(1; i, j)$.

Repeat this procedure for the partitions $\mathcal{V}(1; i, j)$. For each $1 \leq i \leq \kappa_1$ and $1 \leq j \leq \kappa_{1;i}$, let $v(1; i, j, k) \in V(1; i, j, k)$ be a point such that $x(1; i, j, k) = \tau(v(1; i, j, k)) \in L_0$. Let $\widehat{K}_{1;i,j,k}$ be a path-connected subset of $\mathfrak{N}_{1;i,j}$ which contains $x(1; i, j, k)$ in its interior. Denote by $\mathfrak{N}_{1;i,j,k}$, the subsets of $\mathfrak{N}_{1;i,j}$ formed from the base $\widehat{K}_{1;i,j,k}$ and transversal $V(1; i, j, k)$. Observe that we then have three collections of Reeb neighborhoods, each of which forms a covering of \mathfrak{M} :

$$(56) \quad \mathfrak{M} = \bigcup \{ \mathfrak{N}_{1;i} \mid 1 \leq i \leq \kappa_1 \}$$

$$(57) \quad = \bigcup \{ \mathfrak{N}_{1;i,j} \mid 1 \leq i \leq \kappa_1; 1 \leq j \leq \kappa_{1;i} \}$$

$$(58) \quad = \bigcup \{ \mathfrak{N}_{1;i,j,k} \mid 1 \leq i \leq \kappa_1; 1 \leq j \leq \kappa_{1;i}; 1 \leq k \leq \kappa_{1;i,j} \} .$$

The decomposition (56) is that given in the proof of Theorem 7.3. The decomposition (57) arises naturally when considering the intersections of Reeb neighborhoods in (56), while the decomposition (58) is used to show that we obtain a well-defined Cantor foliation.

There is no natural notion of “forcing the border” for the decompositions of \mathfrak{M} that we consider, as is the case for tiling spaces as discussed in [7, 10, 28, 53, 54], and as a result we use the additional structure given by the extended coding partitions used to define the spaces in (57) and (58) to show that our decompositions are well-defined.

8.2. Nice stable transversals. For $1 \leq i \leq \kappa_1$ and $z \in \widehat{K}_{1;i} \cap \mathcal{T}$, there exists $1 \leq i_z \leq \nu$ such that $z \in \mathcal{T}_{i_z} \subset \overline{U}_{i_z}$. Choose a path in $\widehat{K}_{1;i}$ joining $x(1; i) = x_i$ and z , which determines a holonomy homeomorphism $h_{x(1;i),z}$ which will be denoted by h_z for simplicity of notation. Note that $h_z(w_i) \in \mathfrak{X}_{i_z}$ and define

$$(59) \quad W(1; i)^z = h_z(W(1; i)) \subset \mathfrak{X}_{i_z} .$$

Then for $\xi \in \mathcal{P}_{i_z}(h_z(w_i))$, a *standard transversal* through ξ is given by

$$(60) \quad \mathfrak{Z}(\xi, i_z, W(1; i)^z) \equiv \varphi_{i_z}^{-1}(\lambda_{i_z}(\xi), W(1; i)^z)$$

which is a vertical coordinate slice for the coordinates on U_{i_z} which passes through the point ξ , and projects to the set $W(1; i)^z$. Note that if $\xi \in U_{i_z'} \cap U_{i_z}$ for some $i_z' \neq i_z$ then the section $\mathfrak{Z}(\xi, i_z, W(1; i)^z)$ need not be a vertical section in the coordinate chart $U_{i_z'}$, and so the choice of the transversal depends on the index i_z and not just the point ξ .

A *nice stable transversal* $\mathcal{X}(1; i)$ on the Reeb neighborhood $\mathfrak{N}_{1;i}$ of $\widehat{K}_{1;i}$ is a disjoint union

$$(61) \quad \mathcal{X}(1; i) = \mathcal{X}_1^{1;i} \cup \dots \cup \mathcal{X}_{p_i}^{1;i}$$

for some ordering of the individual transversals, where each $\mathcal{X}_j^{1;i}$ has the form (60). The choices of the data defining each $\mathcal{X}_j^{1;i}$, as given in the proof of [18, Theorem 1.3], are made so that the intersection $\mathcal{X}(1; i) \cap \widehat{K}_{1;i}$ is a (d_1, d_2) -net on $\widehat{K}_{1;i}$ where the separation and spanning constants $0 < d_1 < d_2$ are chosen sufficiently small so that the leafwise nets they induce yield stable leafwise triangulations, as discussed in [18]. In particular, we note that the constant $d_2 = \lambda_{\mathcal{F}}^*/5$ satisfies $d_2 \leq \lambda_{\mathcal{F}}/25$, although d_2 may be much smaller, as it depends on the sectional curvatures of the leaves of \mathcal{F} , as explained in [18, Sections 15, 17]

The simplicial decomposition on $\widehat{K}_{1;i}$ corresponding to $\mathcal{X}(1; i)$, and therefore the transverse Cantor foliation constructed on $\mathfrak{N}_{1;i}$, depends on the ordering of the local sections comprising a transversal. In the proof of Proposition 8.1 below, we make an initial choice of ordering for the section in $\mathcal{X}(1; 1)$, and then extend this ordering to the other transversals $\mathcal{X}(1; i)$ for $1 < i \leq \kappa_1$, using a recursive procedure. This will yield an ordered nice stable transversal for all of \mathfrak{M} as a result, which is used to construct the transverse Cantor foliation for \mathfrak{M} .

We will consider the restriction of the transversals in $\mathcal{X}(1; i)$ to subsets $\mathfrak{N}_{1; i, j}$, and so obtain a partial transversal $\mathcal{X}(1; i, j)$ for $\mathfrak{N}_{1; i, j}$. To distinguish between these various sections, we will denote by $\mathcal{X}_s^{1; i}$ a section in $\mathfrak{N}_{1; i}$ and by $\mathcal{X}_s^{1; i, j}$ the corresponding section for $\mathfrak{N}_{1; i, j}$ obtained as a restriction of $\mathcal{X}_s^{1; i}$. We note that the section $\mathcal{X}_s^{1; i, j}$ is assigned the same number s in the ordering of transverse sections.

If the neighborhoods $\mathfrak{N}_{1; i, j}$ and $\mathfrak{N}_{1; k, m}$ intersect, in general the transversals $\mathcal{X}(1; i, j)$ and $\mathcal{X}(1; k, m)$ need not match, as illustrated in Figure 1. The recursive procedure developed in the proof of Proposition 8.1 below yields a collection of transversals $\mathcal{X}(1; i, j)$ which match on the intersections of the corresponding Reeb neighborhoods $\mathfrak{N}_{1; i, j}$.

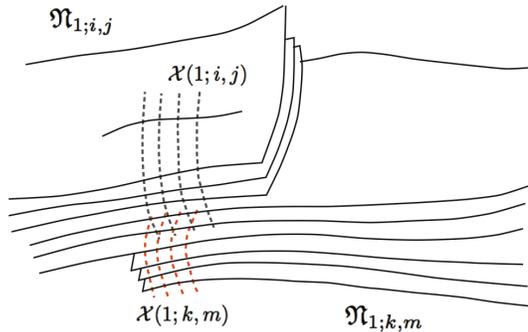


FIGURE 1. Transversals on the overlap of Reeb neighborhoods $\mathfrak{N}_{1; i, j}$ and $\mathfrak{N}_{1; k, m}$.

8.3. Constructing nice stable transversals. We give a recursive procedure for modifying the nice stable transversals on the Reeb covering of \mathfrak{M} as in (57) to obtain a global transversal.

PROPOSITION 8.1. *Nice stable transversals $\mathcal{X}(1; i, j)$ on Reeb neighborhoods $\mathfrak{N}_{1; i, j}$ associated to the partitions $\mathcal{V}(1; i)$ can be chosen so that the following property is satisfied: if $x \in \mathfrak{N}_{1; i, j} \cap \mathfrak{N}_{1; k, m}$, and $x \in \mathcal{X}_s^{1; i, j}$, then $x \in \mathcal{X}_s^{1; k, m}$. In other words, sections of $\mathcal{X}(1; i, j)$ and $\mathcal{X}(1; k, m)$ coincide on the intersection $\mathfrak{N}_{1; i, j} \cap \mathfrak{N}_{1; k, m}$.*

Proof. The recursive construction of the transversals $\mathcal{X}(1; i, j)$ follows a multiple step process. Recall that $R'_1 = 2\theta_1 + \lambda_{\mathcal{F}}$ was defined by (39) and $R_1 = 2R'_1$.

Step 1. Choose nice stable transversals $\mathcal{X}(1; i) = \{\mathcal{X}_1^{1; i}, \dots, \mathcal{X}_{p_i}^{1; i}\}$ on the Reeb neighborhoods $\mathfrak{N}_{1; i}$ for $1 \leq i \leq \kappa_1$ as in (61). By assumption, $w_0 \in V(1; 1, 1)$. Let $\mathcal{X}(1; 1, 1)$ be a nice stable transversal for $\mathfrak{N}_{1; 1, 1}$ obtained by restriction of $\mathcal{X}(1; 1)$.

Step 2. The sets $V(1; i, j)$ for $1 \leq i \leq \kappa_1$ and $1 \leq j \leq \kappa_{1; i}$ are defined using the coding function for paths of length at most R'_1 as above, and their indices are given the lexicographic order. We next extend the transversals in $\mathcal{X}(1; 1, 1)$ to those Reeb neighborhoods which intersect $\mathfrak{N}_{1; 1, 1}$.

Let $(1; i, j)$ be the least index not equal to $(1; 1, 1)$ such that $\mathfrak{N}_{1; 1, 1} \cap \mathfrak{N}_{1; i, j} \neq \emptyset$, if such exists. For any $\mathcal{X}_s^{1; 1, 1} \in \mathcal{X}(1; 1, 1)$ such that $\mathcal{X}_s^{1; 1, 1} \cap \mathfrak{N}_{1; i, j} \neq \emptyset$, we define $\mathcal{X}_s^{1; i, j}$ by

$$(62) \quad \mathcal{X}_s^{1; i, j} = \mathcal{X}_s^{1; 1, 1} \cap \mathfrak{N}_{1; i, j} = \mathfrak{Z}(\xi_s, i_z, W(1; j)^z) \cap \mathfrak{N}_{1; i, j} .$$

That is, we extend the transversal to $\mathfrak{N}_{1; i, j}$ while preserving the ordering of the sections. Figure 2 illustrates the intersections of the Reeb Neighborhood $\mathfrak{N}_{1; 1, 1}$ with $\mathfrak{N}_{1; i, j}$.

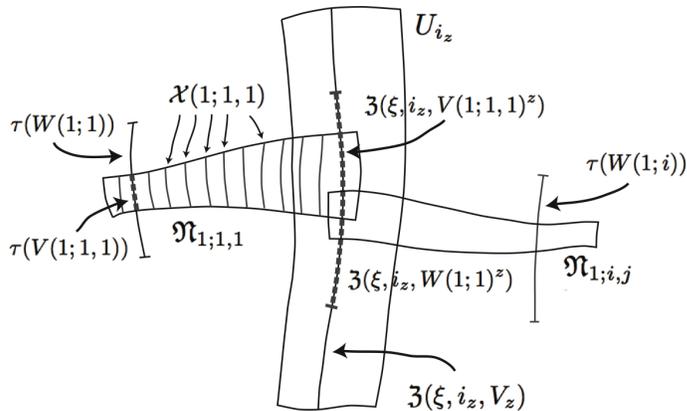


FIGURE 2. A transversal $\mathcal{X}(1; 1, 1)$ constructed on a Reeb neighborhood $\mathfrak{N}_{1;1,1}$. A section $\mathcal{X}_\xi^{1;1,1} = \mathfrak{Z}(\xi, i_z, V(1; 1, 1)^z)$ intersects the Reeb neighborhood $\mathfrak{N}_{1;i,j}$ which can be extended to a section $\mathcal{X}_\xi^{1;i,j}$ since the diameter of $V(1; i, j)$ is so small that $\tau(V(1; i, j)^z) \subset \tau(W(1; 1, 1)^z) \subset \mathfrak{Z}(\xi, i_z, W(1; 1)^z)$.

We repeat this procedure for every $\mathfrak{N}_{1;i,j}$ which intersects $\mathfrak{N}_{1;1,1}$, following the lexicographical ordering on indices. Note that the transversals constructed in the sets $\mathcal{X}(1; i, j)$ do not depend on the order in which the neighborhoods $\mathfrak{N}_{1;i,j}$ are considered, since they are all obtained by restriction of the transversal $\mathcal{X}(1; 1, 1)$. We thus obtain a transversal $\mathcal{X}(1; 1, 1)$ on $\mathfrak{N}_{1;1,1}$, and (partial) transversals $\mathcal{X}(1; i, j)$ defined on the neighborhoods $\mathfrak{N}_{1;i,j}$, $1 \leq i \leq \kappa_1$, $1 \leq j \leq \kappa_{1,i}$ which intersect $\mathfrak{N}_{1;1,1}$.

Next, we consider the case where there is a Reeb neighborhood $\mathfrak{N}_{1;1,m}$ which does not intersect $\mathfrak{N}_{1;1,1}$, but does intersect a subset of $\mathfrak{N}_{1;i,j}$ on which a partial transversal has been defined in the previous step. This situation is illustrated in Figure 3, and involves a fundamental point for the construction of the global Cantor foliation on \mathfrak{M} .

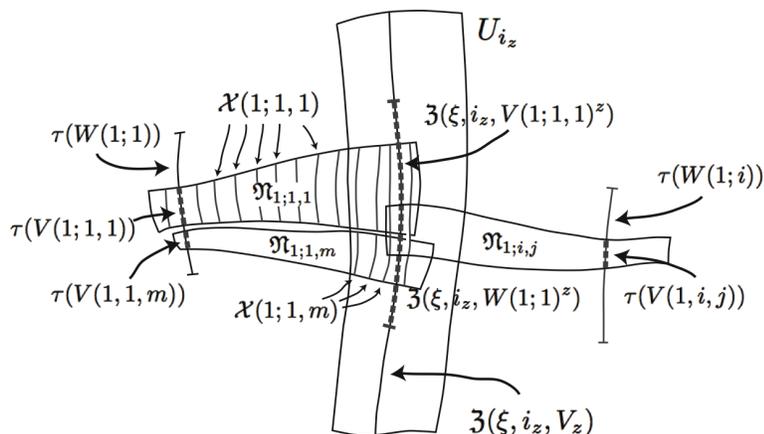


FIGURE 3. A nice stable transversal $\mathcal{X}(1; 1, 1)$ is continued to define partial transversals $\mathcal{X}(1; i, j)$ and $\mathcal{X}(1; 1, m)$.

Let $z \in \mathcal{X}_s^{1;i,j} = \mathcal{X}_s^{1;1} \cap \mathfrak{N}_{1;i,j}$. Then by the construction of the bases for the Reeb neighborhoods in Definition 7.4, and by definitions (59), (60) and (61), there is $x \in \tau(V(1; 1))$ and a path γ of length at most $\theta_1 + 2\delta_{\mathcal{U}}^{\mathcal{F}}$ starting at x and ending at z . Let $h_{x,z}: V(1; 1) \rightarrow \mathcal{X}_s^{1;1}$ be the holonomy map defined by γ .

Suppose there exists $v \in \mathcal{X}_s^{1;i,j} \cap \mathfrak{N}_{1;1,m}$, then there is $u \in \tau(V(1; i, j))$ and a path γ' of length at most $\theta_1 + 2\delta_{\mathcal{U}}^{\mathcal{F}}$ starting at u and ending at v . Let $h_{u,v}: V(1; i) \rightarrow \mathcal{X}_s^{1;1}$ be the holonomy map defined by γ' . Then using the shadowing method in the proof of Lemma 3.13, we can define a path $(\gamma')^{-1} * \gamma$ with length at most $R' = 2\theta_1 + \lambda_{\mathcal{F}}$, which defines the holonomy map $h_{u,v}^{-1} \circ h_{x,y}$.

Then $h_{u,v}^{-1}(\pi(v)) = u \in V(1; i, j)$ and $z \in \mathfrak{N}_{1;i,j}$ implies that $h_{u,v}^{-1}(\pi(z)) \in V(1; i, j)$. This implies that $h_{u,v}^{-1} \circ (\pi(x)) \in V(1; i, j)$. Since the inverse path $\gamma^{-1} * \gamma'$ has length at most R'_1 the definition of the coding partition implies that $h_{x,y}^{-1} \circ h_{u,v}(V(1; i, j)) \subset W(1; 1)$.

It follows that $\mathcal{X}_s^{1;1,m} = \mathcal{X}_s^{1;1} \cap \mathfrak{N}_{1;1,m}$ defines a transversal for $\mathfrak{N}_{1;1,m}$.

We repeat this procedure for every $\mathfrak{N}_{1;1,m}$ which intersects one or more of Reeb neighborhoods with partial transversals along the subset where the transversal is defined. Since the sections $\mathcal{X}_s^{1;i,j}$ are obtained by restriction of the same section $\mathcal{X}_s^{1;1}$, this procedure is well-defined in the case when $\mathfrak{N}_{1;1,m}$ intersects two or more distinct Reeb neighborhoods $\mathfrak{N}_{1;i,j}$ with partial transversal. The partial transversals $\mathcal{X}(1; 1, m)$ do not depend on the order in which the Reeb neighborhoods are considered, and we choose to follow the lexicographic order for definiteness (see Figure 3).

We now have a nice stable transversal $\mathcal{X}(1; 1, 1)$ on $\mathfrak{N}_{1;1,1}$, and transversals $\mathcal{X}(1; i, j)$ on some of the neighborhoods $\mathfrak{N}_{1;i,j}$, $1 \leq i \leq \kappa_1$, $1 \leq j \leq \kappa_{1;i}$, some of them for $i = 1$ and $j \neq 1$.

We must show that the procedure of defining partial transversals eventually stops. For that, let γ, γ' be homeomorphisms associated to paths of length at most $\theta_1 + 2\delta_U^F$. We notice that if for $u \in V(1; 1, 1)$ we have $V(1; 1, 1) \subset \text{Dom}(\gamma)$ and $\gamma(u) \in V(1; i, j)$, then for every γ' with $V(1; i, j) \in \text{Dom}(\gamma')$ the code $C_{w_0}^{1,R_1}(\gamma' \circ \gamma)$ is known. It follows that $\gamma^{-1}(\gamma(V(1; 1, 1)) \cap V(1; i, j))$ is a union of a finite number of sets in the partition $\mathcal{V}(1; 1, 1)$ on which the coding function for paths of length at most R_1 is constant. By a similar argument, $\gamma(\gamma^{-1}(V(1; i, j)) \cap V(1; 1, 1))$ is the union of a finite number of clopen sets in the partition $\mathcal{V}(1; i, j)$. Therefore, an intersection of two slabs $\mathfrak{N}_{1;1,1}$ and $\mathfrak{N}_{1;i,j}$ is always along a subset corresponding to a clopen subset in $V(1; 1)$ which is a union of sets in $\mathcal{V}(1; 1, 1)$. Similarly, $\mathfrak{N}_{1;1,m}$ intersects $\mathfrak{N}_{1;i,j}$ along a subset corresponding to a clopen subset in $V(1; i, j)$ which is a union of subsets in $\mathcal{V}(1; i, j)$. Since the partitions $\mathcal{V}(1; 1)$ and $\mathcal{V}(1; i, j)$ are finite, and intersections always happen over Reeb neighborhoods within the same range of $1 \leq i \leq \kappa_1$, the procedure of defining a partial transversal stops after a finite number of steps.

Step 3 (Inductive step) Suppose for all $(1, 1) \leq (i', j') < (i, j)$ a stable transversal has been defined on $\mathfrak{N}_{1;i',j'}$, and numbers $1, \dots, p$ were assigned to sections in the nice stable transversals so that on each neighborhood, there is at most one section which is assigned a given number s , and two sections are numbered by the same number s in two different neighborhoods if and only if these sections intersect, i.e. $\mathcal{X}_s^{1;i,j} \cap \mathcal{X}_c^{1;k,m} \neq \emptyset$ if and only if $s = c$. Then two cases are possible.

Case 1. There is a partial transversal $\mathcal{X}(1; i, j)$ on $\mathfrak{N}_{1;i,j}$. Then use [18, Proposition 17.4] to complete $\mathcal{X}(1; i, j)$ starting to number the newly defined sections in $\mathcal{X}(1; i, j)$ from $p + 1$.

We now have to repeat the procedure of defining a partial transversal of *Step 2*. We note that it is enough to consider $(k, m) > (i, j)$, since for all $(k, m) < (i, j)$ such that $\mathfrak{N}_{1;k,m} \cap \mathfrak{N}_{1;i,j} \neq \emptyset$, the transversal $\mathcal{X}(1; k, m)$ has already been defined, and, moreover, possibly gave rise to a subset of sections in a partial transversal $\mathcal{X}(1; i, j)$. This ensures that a section $\mathcal{X}_s^{1;i,j}$ in the transversals $\mathcal{X}(1; i, j)$ which intersect the Reeb neighborhood $\mathfrak{N}_{1;k,m}$, $(k, m) < (i, j)$, coincides on this intersection with a section $\mathcal{X}_s^{1;k,m}$, as required.

So suppose there exists $(k, m) > (i, j)$ such that $\mathfrak{N}_{1;k,m} \cap \mathfrak{N}_{1;i,j} \neq \emptyset$. If $\mathcal{X}_s^{1;i,j} \cap \mathfrak{N}_{1;k,m} \neq \emptyset$, then there exists $z \in K_{1;i,j}$ and a chart U_{i_z} such that

$$(63) \quad \mathcal{X}_s^{1;i,j} \subset \mathfrak{Z}(\xi_s, i_z, W(1; i)^z) \quad \text{and} \quad \gamma(V(1; k, m)) \subset W(1, i)$$

for an appropriate homeomorphism γ associated to a path with the same name. Condition (63) makes sure that a section $\mathcal{X}_s^{1;k,m}$ is extended to a maximal subset of $\mathfrak{N}_{1;k,m}$, and in particular to the corresponding subset of $\mathfrak{N}_{1;k,m}$. We define

$$\mathcal{X}_s^{1;k,m} = \mathfrak{Z}(\xi_s, i_z, W(1; i)^z) \cap \mathfrak{N}_{1;k,m},$$

and repeat the procedure for every Reeb neighborhood intersecting $\mathfrak{N}_{1;i,j}$. We may need to repeat the procedure of defining a partial nice stable transversal for $\mathfrak{N}_{1;k',m'}$ such that $\mathfrak{N}_{1;k',m'} \cap \mathfrak{N}_{1;k,m} \neq \emptyset$, and so on. By an argument similar to the one in *Step 2*, the procedure stops after a finite number of repetitions.

Case 2. There is no partial transversal on $\mathfrak{N}_{1;i,j}$. Then implement *Step 1* and *Step 2* with $(1; i, j)$ instead of $(1; 1, 1)$, starting to number sections in $\mathcal{X}(1; i, j)$ from $p + 1$.

Since $\mathcal{V}(1; i)$ is a finite partition, and $1 \leq i \leq \kappa_\ell$, the procedure stops after a finite number of steps. \square

9. CONSTRUCTION OF THE TRANSVERSE CANTOR FOLIATION

In this section, we discuss how to obtain the transverse Cantor foliations \mathcal{H}_ℓ from the nice stable transversal constructed in Section 8. The foliation \mathcal{H}_1 is constructed in Section 9.1. Then in Section 9.2, we show that \mathcal{H}_1 induces Cantor foliations \mathcal{H}_ℓ on families of Reeb neighborhoods associated with smaller coding partitions, which is the key to obtaining the presentation (2) in Theorem 1.1.

9.1. Stable triangulations. The constructions in Section 8.2 yield a nice stable transversal

$$(64) \quad \mathcal{X} = \bigcup \{ \mathcal{X}(1; i, j) \mid 1 \leq i \leq \kappa_1, 1 \leq j \leq \kappa_{1,i} \}$$

on \mathfrak{M} , whose restriction to each Reeb neighborhood $\mathfrak{N}_{1;i,j}$ yields a (d_1, d_2) -net for each path-connected component of $\mathfrak{N}_{1;i,j}$. The net $\mathcal{X} \cap \widehat{K}_{1;i,j}$ determines a triangulation of $\widehat{K}_{1;i,j}$, using the circumscribed sphere method as given in [18, Sections 12,17]. The vertices of the triangulation are contained in the net, and the transverse stability property for \mathcal{X} implies that the triangulations so obtained are isomorphic for each path-connected component of $\mathfrak{N}_{1;i,j}$. As remarked earlier, the leaves of a Cantor foliation on $\mathfrak{N}_{1;i,j}$ are defined by identifying points with the same barycentric coordinates in this simplicial decomposition. This procedure is described in detail in [18, Section 11].

In this section, we take care of a technical issue that arises in this construction of the Cantor foliation. Since taking barycentric coordinates in a simplex only makes sense if all vertices of a simplex are contained in $\mathfrak{N}_{1;i,j}$ we have to restrict to a subset of $\mathfrak{N}_{1;i,j}$. In addition, we have to make sure that after shrinking, the new Reeb neighborhoods still cover \mathfrak{M} . This is achieved if for each $w \in V(1; i, j)$ the restricted neighborhood contains the leafwise ball $B_{\mathcal{F}}(\tau(w), \theta_1)$.

PROPOSITION 9.1. *A Reeb neighborhood $\mathfrak{N}_{1;i,j}$ associated to a clopen set $V(1; i, j)$ contains a foliated neighborhood $\widehat{\mathfrak{N}}_{1;i,j}$ with the following properties.*

- (1) *for every $w \in V(1; i, j)$ we have $B_{\mathcal{F}}(\tau(w), \theta_1) \subset \widehat{\mathfrak{N}}_{1;i,j}$.*
- (2) *$\widehat{\mathfrak{N}}_{1;i,j} \cap K_{1;i,j}$ is the union of simplices of maximal dimension n for the simplicial decomposition of $K_{1;i,j}$ associated to the nice stable transversal $\mathcal{X}(1; i, j)$.*
- (3) *$\widehat{\mathfrak{N}}_{1;i,j}$ is transversely stable; that is, for every path-connected component $K \subset \mathfrak{N}_{1;i,j}$ with $\tau(w) \in K$, the points $\mathcal{X}_{i_0}^{1;i,j} \cap K, \dots, \mathcal{X}_{i_p}^{1;i,j} \cap K$ are vertices of a p -simplex for $1 \leq p \leq n$, if and only if $\mathcal{X}_{i_0}^{1;i,j} \cap K_{1;i,j}, \dots, \mathcal{X}_{i_p}^{1;i,j} \cap K_{1;i,j}$ are vertices of a p -simplex.*

Proof. The intersection $\mathcal{X} \cap \widehat{K}_{1;i,j}$ is a (d_1, d_2) -net for $\widehat{K}_{1;i,j}$, with $d_2 \leq \lambda_{\mathcal{F}}/25$. Thus, for each $y \in \mathcal{X}(1; i, j)$ with $w = \pi(y)$, there is a Voronoi cell $\mathcal{C}_w^{\mathcal{X}}(y) \subset D_{\mathcal{F}}(y, d_2)$. A simplex in the simplicial decomposition associated to the ordered transversal \mathcal{X} which has y as one of the vertices, is contained in the star-neighborhood $\mathcal{S}_w^{\mathcal{X}}(y)$. Then by Lemma 4.11 for this case, we have $\mathcal{S}_w^{\mathcal{X}}(y) \subset D_{\mathcal{F}}(y, 3d_2)$.

Therefore, for each simplex such that one of its vertices is contained in a ball $B_{\mathcal{F}}(y, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}})$, the entire simplex is contained in the ball $B_{\mathcal{F}}(y, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}} + 4d_2)$. From the estimates $\delta_{\mathcal{U}}^{\mathcal{F}} < \lambda_{\mathcal{F}}/5$ and $d_2 \leq \lambda_{\mathcal{F}}/25$ we have $\delta_{\mathcal{U}}^{\mathcal{F}} + 4d_2 < \lambda_{\mathcal{F}}/2$ so that $B_{\mathcal{F}}(y, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}} + 4d_2) \subset B_{\mathcal{F}}(y, \theta_1 + \lambda_{\mathcal{F}}/2)$.

For $v(1; i, j) \in V(1; i, j)$ with $y = x(1; i, j) = \tau(v(1; i, j))$, the ball $B_{\mathcal{F}}(y, \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}}) \subset \widehat{K}_{1;i,j}$ where $\widehat{K}_{1;i,j}$ is the closure of the plaque-saturation of $B_{\mathcal{F}}(y, \theta_1 + \lambda_{\mathcal{F}})$.

Let $K_{1;i,j}^t$ denote the closed triangulated set obtained from the union of all simplices of maximal dimension n associated to the net $\mathcal{X} \cap \widehat{K}_{1;i,j}$ which intersect $K_{1;i,j}$. It then follows that $K_{1;i,j}^t \subset \widehat{K}_{1;i,j}$ so the map Φ_i given in Theorem 7.3 is defined on $K_{1;i,j}^t \times V(1; i, j)$.

Set $\widehat{\mathfrak{N}}_{1;i,j} = \Phi_i \{K_{1;i,j}^t \times V(1;i,j)\}$ for $1 \leq i \leq \kappa_1$ and $1 \leq j \leq \kappa_{1;i}$.

The inclusions $B_{\mathcal{F}}(x(1;i,j), \theta_1 + \delta_{\mathcal{U}}^{\mathcal{F}}) \subset K_{1;i,j}$ implies that these sets form a cover of \mathfrak{M} , exactly as in the proof of Theorem 7.3. The transverse stability of the sets $\widehat{\mathfrak{N}}_{1;i,j}$ follows from the transverse stability of the transversals $\mathcal{X}(1;i,j)$. \square

We have a clopen set V_1 , and a clopen partition $V(1;i,j)$ of V_1 with associated system of Reeb neighborhoods $\widehat{\mathfrak{N}}_{1;i,j}$, $1 \leq i \leq \kappa_1$, $1 \leq j \leq \kappa_{1;i}$ defined in Proposition 9.1. Since each path-connected component of $\widehat{\mathfrak{N}}_{1;i,j}$ is a union of simplices, and the simplicial decomposition is stable in the transverse direction, the following is well-defined.

DEFINITION 9.2. *A Cantor foliation $\mathcal{H}_{1;i,j}$ on $\widehat{\mathfrak{N}}_{1;i,j}$ is defined as follows: two points $x, y \in \widehat{\mathfrak{N}}_{1;i,j}$ are equivalent, $x \approx_{1;i,j} y$ if and only if the following conditions are satisfied:*

- (1) if $x \in \Delta(x_0, x_1, \dots, x_p)$ where $\Delta(x_0, \dots, x_p)$ is a p -simplex with vertices $x_k \in \mathcal{X}_{i_k}^{1;i,j}$, then $y \in \Delta(y_0, \dots, y_p)$ with $y_k \in \mathcal{X}_{i_k}^{1;i,j}$, the same section of $\mathcal{X}(1;i,j)$,
- (2) x and y have the same barycentric coordinates in $\Delta(x_0, \dots, x_p)$ and $\Delta(y_0, \dots, y_p)$ respectively.

For a point $x \in \mathfrak{M}$, $\mathcal{H}_{1;i,j}(x)$ denotes a leaf of $\mathcal{H}_{1;i,j}$ through x .

If two neighborhoods $\widehat{\mathfrak{N}}_{1;i,j}$ and $\widehat{\mathfrak{N}}_{1;k,m}$ intersect, then the foliations $\mathcal{H}_{1;i,j}$ and $\mathcal{H}_{1;k,m}$ match on the intersection, since the sections of transversals $\mathcal{X}(1;i,j)$ and $\mathcal{X}(1;k,m)$ are continuations of each other, and the ordering of sections in nice stable transversals matches as well. We extend $\mathcal{H}_{1;i,j}$ on \mathfrak{M} in the following way.

DEFINITION 9.3. *Let \mathfrak{M} be a matchbox manifold, and $\widehat{\mathfrak{N}}_{1;i,j}$ be a Reeb neighborhood with a Cantor foliation $\approx_{1;i,j}$. We define $\approx_{1;i,j}$ on \mathfrak{M} in the following way: let $x, y \in \mathfrak{M}$. Then $x \approx_{1;i,j} y$ if either $x, y \in \widehat{\mathfrak{N}}_{1;i,j}$ and $x \approx_{1;i,j} y$, or $x, y \notin \widehat{\mathfrak{N}}_{1;i,j}$ and $x = y$.*

Thus the equivalence classes of $\approx_{1;i,j}$ coincide with leaves of $\mathcal{H}_{1;i,j}$ on $\widehat{\mathfrak{N}}_{1;i,j}$, and points outside of $\widehat{\mathfrak{N}}_{1;i,j}$ have only themselves in their equivalence class.

DEFINITION 9.4. *Let \mathfrak{M} be a matchbox manifold, and let $\approx_{1;i,j}$, $1 \leq i \leq \kappa_1$, $1 \leq j \leq \kappa_{1;i}$ be a family of equivalence relations on \mathfrak{M} . For any $x, y \in \mathfrak{M}$ define $x \approx_1 y$ if and only if there exists a finite collection of points z_1, z_2, \dots, z_k such that*

$$x \approx_{1;i,j} z_1 \approx_{1;i_1,j_1} z_2 \approx \dots \approx z_k \approx_{1;i_k,j_k} y$$

A Cantor foliation \mathcal{H}_1 has equivalence classes of \approx_1 as leaves. For a point $x \in \mathfrak{M}$, $\mathcal{H}_1(x)$ denotes a leaf of \mathcal{H}_1 through x .

It is straightforward that \approx_1 is reflexive, symmetric and transitive. Geometrically, this relation identifies $x \in \widehat{\mathfrak{N}}_{1;i,j}$ and $y \in \widehat{\mathfrak{N}}_{1;k,m}$ if and only if the leaves $\mathcal{H}_{1;i,j}(x), \mathcal{H}_{1;i_1,j_1}(z_1), \dots, \mathcal{H}_{1;i_k,j_k}(z_k)$ form an intersecting chain, and $\mathcal{H}_{1;i_k,j_k}(z_k) \ni y$. The equivalence relation \approx_1 determines the Cantor foliation \mathcal{H}_1 .

Thus, by Propositions 8.1 and 9.1, there is a system of Reeb neighborhoods $\widehat{\mathfrak{N}}_{1;i,j}$, $1 \leq i \leq \kappa_1$, $1 \leq j \leq \kappa_{1;i}$, and a Cantor foliation $\mathcal{H}_{1;i,j}$ on each $\widehat{\mathfrak{N}}_{1;i,j}$ such that the Cantor foliations match on the intersections of neighborhoods, and their union equals \mathfrak{M} . These Cantor foliations combine to yield the Cantor foliation \mathcal{H}_1 on \mathfrak{M} .

9.2. Families of nested Cantor foliations. In this section, we give an inductive procedure for restricting the foliation \mathcal{H}_1 to the Reeb neighborhoods $\mathfrak{N}_{\ell;i,j}$ defined using the partitions defined in Proposition 6.5.

PROPOSITION 9.5. *There exists a nested sequence of Cantor foliations \mathcal{H}_ℓ , for all $\ell \geq 1$, with \mathcal{H}_1 given as above, such that each Reeb neighborhood $\mathfrak{N}_{\ell;i,j}$ associated to a clopen set $V(\ell; i, j)$ contains a foliated neighborhood $\widehat{\mathfrak{N}}_{\ell;i,j}$ with the following properties.*

- (1) *for every $w \in V(\ell; i, j)$ we have $B_{\mathcal{F}}(\tau(w), \theta_\ell) \subset \widehat{\mathfrak{N}}_{\ell;i,j}$.*
- (2) *$\widehat{\mathfrak{N}}_{\ell;i,j} \cap K_{\ell;i,j}$ is the union of simplices of maximal dimension n for the simplicial decomposition of $K_{\ell;i,j}$ associated to the nice stable transversal $\mathcal{X}(\ell; i, j)$.*
- (3) *$\widehat{\mathfrak{N}}_{\ell;i,j}$ is transversely stable; that is, for every path-connected component $K \subset \mathfrak{N}_{\ell;i,j}$ with $\tau(w) \in K$, the points $\mathcal{X}_{i_0}^{\ell;i,j} \cap K, \dots, \mathcal{X}_{i_p}^{\ell;i,j} \cap K$ are vertices of a p -simplex for $1 \leq p \leq n$, if and only if $\mathcal{X}_{i_0}^{\ell;i,j} \cap K_{\ell;i,j}, \dots, \mathcal{X}_{i_p}^{\ell;i,j} \cap K_{\ell;i,j}$ are vertices of a p -simplex.*
- (4) *Each $\widehat{\mathfrak{N}}_{\ell;i,j}$ has a Cantor foliation $\mathcal{H}_{\ell;i,j}$ such that they defined a Cantor foliation \mathcal{H}_ℓ on \mathfrak{M} .*
- (5) *The identity map on \mathfrak{M} induces an inclusion of leaves of \mathcal{H}_ℓ into the leaves of $\mathcal{H}_{\ell-1}$.*
- (6) *For each $x \in \mathfrak{M}$ and $\ell \geq 1$, the leaves of \mathcal{H}_ℓ containing x are nested Cantor sets. Moreover, $\text{diam}_{\mathfrak{M}}(\mathcal{H}_\ell(x)) \leq \widehat{\delta}_\ell$.*

Proof. The idea of the inductive construction of \mathcal{H}_ℓ is to restrict the given foliation $\mathcal{H}_{\ell-1}$ on the Reeb neighborhoods $\widehat{\mathfrak{N}}_{\ell-1;i,j}$ to the Reeb neighborhoods $\mathfrak{N}_{\ell;i,j}$, and then modify the neighborhoods $\mathfrak{N}_{\ell;i,j}$ to obtain the neighborhoods $\widehat{\mathfrak{N}}_{\ell;i,j}$, as in the proof of Proposition 9.1. The procedure is straightforward, though care is needed, to show that the restricted Cantor foliations are nested.

Let $\ell > 1$, and assume there is given a Cantor foliation $\mathcal{H}_{\ell-1}$ on \mathfrak{M} satisfying the conditions of Proposition 9.5 for $\ell - 1$.

We also assume there is given a collection of nested clopen sets as given by Proposition 6.5. Recall that $\tau(V_\ell) \cap L$ is a $(\lambda_1(\epsilon_\ell), \theta_\ell)$ -net for each leaf $L \subset \mathfrak{M}$, and in particular for the leaf L_0 . Define the constants $R_\ell = 4\theta_\ell + 2\lambda_{\mathcal{F}}$ and $R'_\ell = 2\theta_\ell + \lambda_{\mathcal{F}}$. Then we have Reeb neighborhoods which cover \mathfrak{M} , defined as in Section 8.2 using this data, so that

$$(65) \quad \mathfrak{M} = \bigcup \{ \mathfrak{N}_{\ell;i,j} \mid 1 \leq i \leq \kappa_\ell; 1 \leq j \leq \kappa_{\ell;i} \}$$

$$(66) \quad = \bigcup \{ \mathfrak{N}_{\ell;i,j,k} \mid 1 \leq i \leq \kappa_\ell; 1 \leq j \leq \kappa_{\ell;i}; 1 \leq k \leq \kappa_{\ell;i,j} \} .$$

For each $1 \leq i \leq \kappa_\ell$ and $1 \leq j \leq \kappa_{\ell;i}$, the inclusion $\mathfrak{N}_{\ell;i,j} \subset \mathfrak{M}$ defines the foliation \mathcal{H}_ℓ by restriction. However, it is necessary that this restriction is of the form in Definition 7.1. This is guaranteed by the following technical result.

LEMMA 9.6. *For each $1 \leq i \leq \kappa_\ell$ and $1 \leq j \leq \kappa_{\ell;i}$, with*

$$(67) \quad \mathfrak{N}_{\ell;i,j} \subset \bigcup \mathfrak{N}_{\ell-1;k,m}$$

where k and m run over a subset of the indexing sets of $\mathcal{W}(\ell - 1)$ and $\mathcal{V}(\ell - 1; k)$, the following condition is satisfied:

(A) *For each $z \in K_{\ell;i,j} \cap \mathcal{T}$, if $\mathfrak{Z}(0, i_z, V(\ell; i, j)^z) \cap \mathfrak{N}_{\ell-1;k,m} \neq \emptyset$, then $\mathfrak{Z}(0, i_z, V(\ell; i, j)^z) \subset \mathfrak{N}_{\ell-1;k,m}$.*

Proof. Since the sets of $\mathcal{V}(\ell; i)$, $1 \leq i \leq \kappa_\ell$, partition $V_\ell = V(\ell - 1, 1, 1)$, the condition (A) is satisfied for each $V(\ell; i, j)$. Also this implies that for every $z \in K_{\ell-1;1,1}$ we have $V(\ell; i, j) \subset V_\ell \subset \text{Dom}(h_z)$, which means that (A) is satisfied for all $z' \in K_{\ell;i,j} \cap \mathfrak{N}_{\ell-1;1,1} \cap \mathcal{T}$.

Suppose $\mathfrak{N}_{\ell;i,j} \cap \mathfrak{N}_{\ell-1;k,m} \neq \emptyset$. We assume for now that $\mathfrak{N}_{\ell;i,j}$ has a single intersection with $\mathfrak{N}_{\ell-1;k,m}$, i.e. for each path-connected component K of $\mathfrak{N}_{\ell;i,j}$ the intersection $K \cap \mathfrak{N}_{\ell-1;k,m}$ is path-connected.

Let $z \in K \cap \mathfrak{N}_{\ell-1;k,m} \subset K'$, where K' is a path-connected component of $\mathfrak{N}_{\ell-1;k,m}$. Recall from conditions (47) and (48) that the restriction $\text{diam}_{\mathcal{X}}(V(\ell; i, j)) < \widehat{\delta}_\ell$ ensures that for every translation of $V(\ell; i, j)$ along a path γ of length at most R_ℓ , points in $V(\ell; i, j)$ do not move apart further than the constant ϵ_ℓ , which was chosen to be smaller than the distance between any two sets in the partition $\mathcal{V}(\ell - 1; k)$. This means that if $\gamma(V(\ell; i, j))$ hits $V_{\ell-1}$, then the image $\gamma(V(\ell; i, j))$ is contained in one of the sets of the clopen partition $\mathcal{V}(\ell - 1; k)$ of $V_{\ell-1}$.

Let γ' be a path between the point $K' \cap \tau(V(\ell - 1; k, m))$ and z . Then

$$\gamma(V(\ell; i, j)) \subset V(\ell - 1; k, m) \subset \text{Dom}(\gamma')$$

and the condition (A) is satisfied for the section $\tau(V(\ell; i, j)^z) = \tau(\gamma'(\gamma(V(\ell; i, j))))$.

Next, suppose that $\mathfrak{N}_{\ell; i, j}$ has multiple intersections with $\mathfrak{N}_{\ell-1; k, m}$, as illustrated in Figure 4, where $\mathfrak{N}_{\ell; i, j}$ intersects Reeb neighborhoods $\mathfrak{N}_{\ell-1; k, m}, \mathfrak{N}_{\ell-1; k', m'}$ and $\mathfrak{N}_{\ell-1; k'', m''}$.

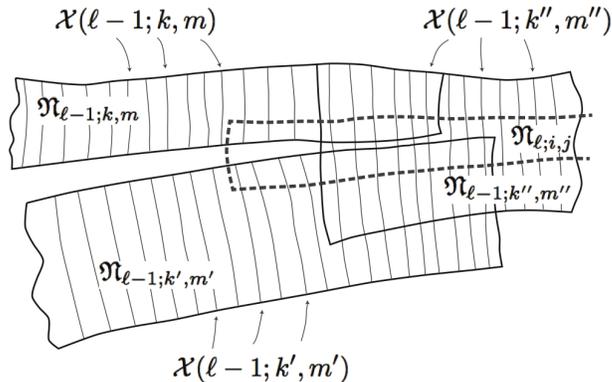


FIGURE 4. $\mathfrak{N}_{\ell; i, j}$ cannot intersect the neighborhoods $\mathfrak{N}_{\ell-1; k, m}$ and $\mathfrak{N}_{\ell-1; k', m'}$ as above.

As the bases of the Reeb neighborhoods are plaque saturated, by compactness the number of possible intersections is finite. That is, for each path-connected component $K \subset \mathfrak{N}_{\ell; i, j}$ the set $K \cap \mathfrak{N}_{\ell-1; k, m}$ has a finite number of path-connected components. Then the proof is completed by repeating the above argument for each intersection. \square

Lemma 9.6 allows us to define a nice stable transversal on each Reeb neighborhood $\mathfrak{N}_{\ell; i, j}$ as follows.

DEFINITION 9.7. Let $\mathcal{X}(\ell - 1; k, m)$ be a family of nice stable transversals defined on Reeb neighborhoods $\mathfrak{N}_{\ell-1; k, m}$ associated to the coding partition $\mathcal{V}(\ell - 1; k)$, and let $\mathfrak{N}_{\ell; i, j}$ be a collection of Reeb neighborhoods associated to the coding partition $\mathcal{V}(\ell; i)$.

Define $\mathcal{X}(\ell; i, j)$ as follows: for each $\mathcal{X}_{\xi}^{\ell-1; k, m}$ such that $\mathcal{X}_{\xi}^{\ell-1; k, m} \cap \mathfrak{N}_{\ell; i, j} \neq \emptyset$, define

$$\mathcal{X}_{\xi}^{\ell; i, j} = \mathcal{X}_{\xi}^{\ell-1; k, m} \cap \mathfrak{N}_{\ell; i, j}.$$

Now use argument similar to the proof of Proposition 9.5 to restrict each Reeb neighborhood $\mathfrak{N}_{\ell; i, j}$ to a Reeb neighborhood $\widehat{\mathfrak{N}}_{\ell; i, j}$ where each connected component is the union of simplices in the simplicial decomposition associated to the transversal \mathcal{X}_{ℓ} . Then use an argument similar to those in Section 9.1 to define a family of equivalence relations $\approx_{\ell; i, j}$ on \mathfrak{M} , and to obtain an equivalence relation \approx_{ℓ} and a Cantor foliation \mathcal{H}_{ℓ} on \mathfrak{M} .

The claims (5) and (6) in Proposition 9.5 follow from the construction of \mathcal{H}_{ℓ} . \square

10. INVERSE LIMIT APPROXIMATIONS OF MATCHBOX MANIFOLDS

In this section we complete the proof of Theorem 1.1. Namely, having established a sequence of nested equivalence relations \approx_{ℓ} on a matchbox manifold \mathfrak{M} , we prove that each quotient $M_{\ell} = \mathfrak{M} / \approx_{\ell}$ is a topological branched manifold, and $\mathfrak{M} \cong \varprojlim \{q_{\ell'} : M_{\ell'} \rightarrow M_{\ell}\}$, where $q_{\ell', \ell}$ is map induced by inclusions of equivalence classes of $\approx_{\ell'}$ into equivalence classes of \approx_{ℓ} .

THEOREM 10.1. Let \mathfrak{M} be a minimal matchbox manifold with a Cantor foliation \mathcal{H}_{ℓ} and associated equivalence relation \approx_{ℓ} . Then the quotient $M_{\ell} = \mathfrak{M} / \approx_{\ell}$ is a compact connected metrizable space. For each $x \in \mathfrak{M}$, an open neighborhood of x is homeomorphic to a finite union of open disks in \mathbb{R}^n modulo identifications. Thus, M_{ℓ} is a topological branched manifold.

Proof. We first show that the quotient $M_\ell = \mathfrak{M}/\approx_\ell$ is a compact connected metrizable space.

Since equivalence classes of \approx_ℓ are closed, for every closed set $U \subset \mathfrak{M}$ its saturation $S(U)$ by equivalence classes of \approx_ℓ is closed. Therefore, \approx_ℓ is a closed equivalence relation [14].

Since \mathfrak{M} is compact and Hausdorff, it is locally compact and so it is completely regular [41, p.42]. Since \approx_ℓ is a closed equivalence relation, its graph is closed in $\mathfrak{M} \times \mathfrak{M}$ [14, p.82]. It follows that M_ℓ is Hausdorff. The quotient M_ℓ is compact and connected since \mathfrak{M} is compact and connected. It is a continuous image of a compact metric space \mathfrak{M} , and so is metrizable.

Since each $x \in \mathfrak{M}$ is contained in at most a finite number of Reeb neighborhoods $\widehat{\mathfrak{N}}_{\ell,i,j}$, a point $[x]_\ell \in M_\ell$ has a neighborhood which is a finite union of open disks modulo identifications. Although leaves of \mathfrak{M} are Riemannian manifolds, there is no canonical way to put a differentiable structure on M_ℓ , and therefore, it is a topological branched manifold (compare [65]). \square

Denote by $q_\ell : \mathfrak{M} \rightarrow \mathfrak{M}/\approx_\ell$ the projection map. Let $\ell' > \ell$, and $M_{\ell'} = \mathfrak{M}/\approx_{\ell'}$ and $M_\ell = \mathfrak{M}/\approx_\ell$ be quotient branched manifolds. We are going to define bonding maps between them.

LEMMA 10.2. *For $\ell' > \ell$, there is a natural continuous map $q_{\ell',\ell} : M_{\ell'} \rightarrow M_\ell$ such that $q_{\ell',\ell} = q_\ell \circ q_{\ell',\ell}$.*

Proof. Note that for $\ell' > \ell$ we have the following diagram

$$\begin{array}{ccc} \mathfrak{M} & \xrightarrow{\text{id}} & \mathfrak{M} \\ q_{\ell'} \downarrow & & \downarrow q_\ell \\ M_{\ell'} & \xrightarrow{q_{\ell',\ell}} & M_\ell \end{array}$$

so we set $q_{\ell',\ell} : \mathfrak{M}/\approx_{\ell'} \rightarrow \mathfrak{M}/\approx_\ell : \bar{x}' = [x]_{\ell'} \rightarrow [x]_\ell$. By Lemma 9.6 the equivalence classes of $\approx_{\ell'}$ are subsets of equivalence classes of \approx_ℓ , and so this map is well-defined. It is straightforward that $q_{\ell',\ell}$ is continuous. \square

Recall that the inverse limit of the directed systems of maps $\{q_{\ell,\ell+1} : M_{\ell+1} \rightarrow M_\ell \mid \ell \geq \ell_0\}$ is the topological space

$$(68) \quad \mathcal{S}\{q_{\ell,\ell'} : M_{\ell'} \rightarrow M_\ell\} \equiv \{\omega = (\omega_{\ell_0}, \omega_{\ell_0+1}, \dots) \in \prod_{\ell \geq \ell_0}^{\infty} M_\ell \mid q_{\ell,\ell+1}(\omega_{\ell+1}) = \omega_\ell\}$$

THEOREM 10.3. *There is a homeomorphism $q : \mathfrak{M} \rightarrow \mathcal{S}\{q_{\ell,\ell'} : M_{\ell'} \rightarrow M_\ell\}$ of foliated spaces.*

Proof. For $x \in \mathfrak{M}$ define

$$q(x) = ([x]_1, [x]_2, \dots) \in \mathcal{S}\{q_{\ell,\ell'} : M_{\ell'} \rightarrow M_\ell\}$$

which is well-defined by Lemma 10.2. The map to each factor, $x \mapsto [x]_\ell = q_\ell(x)$ is continuous, hence q is continuous.

Let $x, y \in \mathfrak{M}$ such that $q(x) = q(y)$. Then $q_\ell(x) = q_\ell(y)$ for all $\ell \geq 1$. That is, $x \approx_\ell y$ for all $\ell \geq 1$.

This means that x and y are always in the same leaf of \mathcal{H}_ℓ . Proposition 9.5.6 shows that the maximal diameter of \mathcal{H}_ℓ -equivalence classes tend to zero as $\ell \rightarrow \infty$. It follows that x and y cannot be separated by an open set, and, therefore, $x = y$.

Surjectivity of q follows again from the fact that equivalence classes of \approx_ℓ are nested, and their maximal length tends to zero. \square

This completes the proof of Theorem 1.1.

REFERENCES

- [1] J.M. Aarts, C.L. Hagopian and L.G. Oversteegen, *The orientability of matchbox manifolds*, **Pacific. J. Math.**, 150:1–12, 1991.
- [2] J. Aarts and L. Oversteegen, *Matchbox manifolds*, In **Continua (Cincinnati, OH, 1994)**, Lecture Notes in Pure and Appl. Math., Vol. 170, Dekker, New York, 1995, pages 3–14..
- [3] J.M. Aarts and M. Martens, *Flows on one-dimensional spaces*, **Fund. Math.**, 131:39–58, 1988.
- [4] F. Alcalde Cuesta, A. Lozano Rojo and M. Macho Stadler, *Transversely Cantor laminations as inverse limits*, **Proc. Amer. Math. Soc.**, 139:2615–2630, 2011.
- [5] J. Aliste-Prieto and D. Coronel, *Tower systems for linearly repetitive Delone sets*, **Ergodic Theory Dyn. Syst.**, 31:1595–1618, 2011.
- [6] M. Amini, I. Putnam and S. Saedi-Gholikandi, *Homology for one-dimensional Solenoids*, arXiv:1307.0977v1.
- [7] J. Anderson and I. Putnam, *Topological invariants for substitution tilings and their associated C^* -algebras*, **Ergodic Theory Dyn. Syst.**, 18:509–537, 1998.
- [8] J. Auslander and J. Yorke, *Interval maps, factors of maps, and chaos*, **Tôhoku Math. J. (2)**, 32:177–188, 1980.
- [9] J. Auslander, E. Glasner and B. Weiss, *On recurrence in zero dimensional flows*, **Forum Math.**, 19:107–114, 2007.
- [10] M. Barge, B. Diamond, J. Hunton and L. Sadun, *Cohomology of substitution tiling spaces*, **Ergodic Theory Dyn. Syst.**, 30:1607–1627, 2010.
- [11] J. Bellissard, R. Benedetti and J.-M. Gambaudo, *Spaces of tilings, finite telescopic approximations and gap-labeling*, **Comm. Math. Phys.**, 261:1–41, 2006.
- [12] R. Benedetti and J.-M. Gambaudo, *On the dynamics of \mathbb{G} -solenoids. Applications to Delone sets*, **Ergodic Theory Dyn. Syst.**, 23:673–691, 2003.
- [13] K. Borsuk, *Concerning homotopy properties of compacta*, **Fund. Math.**, 62:223–254, 1968.
- [14] Bourbaki, Nicolas, **General topology. Chapters 1–4**, Springer-Verlag, Berlin, 1989.
- [15] C. Camacho and A. Lins Neto, **Geometric Theory of Foliations**, Translated from the Portuguese by Sue E. Goodman, Progress in Mathematics, Birkhäuser Boston, MA, 1985.
- [16] A. Candel and L. Conlon, **Foliations I**, Amer. Math. Soc., Providence, RI, 2000.
- [17] A. Clark and S. Hurder, *Homogeneous matchbox manifolds*, **Transactions AMS**, 365:3151–3191, 2013, arXiv:1006.5482v2.
- [18] A. Clark, S. Hurder and O. Lukina, *Voronoi tessellations for matchbox manifolds*, **Topology Proceedings**, 41:167–259, 2013, arXiv:1107.1910v2.
- [19] A. Clark, S. Hurder and O. Lukina, *Classifying matchbox manifolds*, preprint, 2013, arXiv:1311.0226.
- [20] M. do Carmo, **Riemannian geometry**, Translated from the second Portuguese edition by Francis Flaherty, Birkhäuser Boston Inc., Boston, 1992.
- [21] J.-M. Cordier and T. Porter, **Shape theory. Categorical methods of approximation**, Ellis Horwood Ltd., Chichester, 1989.
- [22] M. Coornaert and A. Papadopoulos, **Symbolic dynamics and hyperbolic groups**, Lecture Notes in Mathematics, Vol. 1539, Springer-Verlag, Berlin, Boston, 1993.
- [23] T. Downarowicz and A. Maass *Finite-rank Bratteli-Vershik diagrams are expansive*, **Ergodic Theory Dynam. Systems**, 28:739–747, 2008.
- [24] T. Farrell and L. Jones *Expanding immersions on branched manifolds*, **Amer. J. Math.**, 103:41–101, 1981.
- [25] R. Fokkink, *The structure of trajectories*, Ph. D. Thesis, TU Delft, 1991.
- [26] R. Fokkink and L. Oversteegen, *Homogeneous weak solenoids*, **Trans. Amer. Math. Soc.**, 354:3743–3755, 2002.
- [27] A.H. Forrest, *A Bratteli diagram for commuting homeomorphisms of the Cantor set*, **Internat. J. Math.**, 11:177–200, 2000.
- [28] A. Forrest, J. Hunton and J. Kellendonk, **Topological invariants for projection method patterns**, Mem. Amer. Math. Soc., Vol. 159, 2002.
- [29] N. Priebe Frank and L. Sadun, *Topology of some tiling spaces without finite local complexity*, **Discrete Contin. Dyn. Syst.**, 23:847–865, 2009.
- [30] D. Fried *Finitely presented dynamical systems*, **Ergodic Theory Dynam. Systems**, 7:489–507, 1987.
- [31] J.-M. Gambaudo and M. Martens, *Algebraic topology for minimal Cantor sets*, **Ann. H. Poincaré**, 7:423–446, 2006.
- [32] É Ghys, *Laminations par surfaces de Riemann*, In **Dynamique et Géométrie Complexes**, Panoramas & Synthèses, 8:49–95, 1999.
- [33] M. Gromov, *Hyperbolic manifolds, groups and actions*, In **Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)**, Ann. of Math. Stud., Vol. 97, Princeton Univ. Press, Princeton, NY, 1981, pages 183–213.
- [34] M. Gromov, *Hyperbolic groups*, In **Essays in group theory**, Math. Sci. Res. Inst. Publ., Vol. 8, Springer, New York, 1987, pages 75–263.
- [35] A. Haefliger, *Groupoïdes d’holonomie et classifiants*, In **Transversal structure of foliations (Toulouse, 1982)**, Asterisque, 177–178, Société Mathématique de France, 1984:70–97.
- [36] S. Helgason, **Differential geometry, Lie groups, and symmetric spaces**, Pure and Applied Mathematics, Vol. 80, Academic Press Inc., New York, 1978.

- [37] S. Hurder, *Dynamics and the Godbillon-Vey class: a History and Survey*, In **Foliations: Geometry and Dynamics (Warsaw, 2000)**, World Scientific Publishing Co. Inc., River Edge, N.J., 2002:29–60.
- [38] S. Hurder, *Problem set*, In **Foliations 2005**, World Sci. Publ., Hackensack, NJ, 2006: 441–475.
- [39] S. Hurder, *Lipshitz matchbox manifolds*, preprint, 2013, arXiv:1309.1512.
- [40] L. Jones, *Locally strange hyperbolic sets*, **Trans. Amer. Math. Soc.**, 275:153–162, 1983.
- [41] K. Kuratowski **Topology. Vol. II**, Academic Press, New York, 1968.
- [42] G. Leibon and D. Letscher, *Delaunay triangulations and Voronoi diagrams for Riemannian manifolds*, In **Proceedings of the Sixteenth Annual Symposium on Computational Geometry (Hong Kong, 2000)**, ACM, New York, 2000, 341–349.
- [43] Á. Lozano-Rojo, *Codimension zero laminations are inverse limits*, *Topology Appl.*, 160:341–349, 2013; arXiv:1204.6439.
- [44] O. Lukina, *Hierarchy of graph matchbox manifolds*, **Topology Appl.**, 159:3461–3485, 2012. arXiv:1107.5303.
- [45] M. Lyubich and Y. Minsky, *Laminations in holomorphic dynamics*, **J. Differential Geom.**, 47:17–94, 1997.
- [46] C. McCord, *Inverse limit sequences with covering maps*, **Trans. Amer. Math. Soc.**, 114:197–209, 1965.
- [47] C.C. Moore and C. Schochet, **Analysis on Foliated Spaces**, Math. Sci. Res. Inst. Publ. vol. 9, Second Edition, Cambridge University Press, New York, 2006.
- [48] A. Okabe, B. Boots, K. Sugihara and S.N. Chiu, **Spatial tessellations: concepts and applications of Voronoi diagrams**, second ed., Wiley Series in Probability and Statistics, John Wiley & Sons Ltd., Chichester, 2000.
- [49] I. Putnam, *C^* -algebras from Smale spaces*, **Canad. J. Math.**, 48:175–195, 1996.
- [50] I. Putnam and J. Spielberg, *The structure of C^* -algebras associated with hyperbolic dynamical systems*, **J. Funct. Anal.**, 163:279–299, 1999.
- [51] I. Putnam, *A homology theory for Smale spaces*, **Memoirs A.M.S.**, to appear, 2013.
- [52] D. Ruelle, *Noncommutative algebras for hyperbolic diffeomorphisms*, **Invent. Math.** 93:1–13, 1988.
- [53] L. Sadun *Tiling spaces are inverse limits*, **J. Math. Phys.**, 44:5410–5414, 2003.
- [54] L. Sadun **Topology of tiling spaces**, University Lecture Series, Vol. 46, American Math. Society, 2008.
- [55] R. Schori, *Inverse limits and homogeneity*, **Trans. Amer. Math. Soc.**, 124:533–539, 1966.
- [56] M. Senechal, **Quasicrystals and geometry**, Cambridge University Press, 1995.
- [57] M. Shub and D. Sullivan, *Homology theory and dynamical systems*, **Topology**, 14:109–132, 1975.
- [58] D. Sullivan and R.F. Williams, *On the homology of attractors*, **Topology**, 15:259–262, 1976.
- [59] D. Sullivan, *Bounds, quadratic differentials, and renormalization conjectures*, American Mathematical Society centennial publications, Vol. II (Providence, RI, 1988), pages 417–466, 1992.
- [60] E.S. Thomas, Jr. *One-dimensional minimal sets*, **Topology**, 12:233–242, 1973.
- [61] L. Vietoris, *Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen*, **Math. Ann.**, 97:454–472, 1927.
- [62] S. Wieler, *Smale Spaces via Inverse Limits*, **Ergodic Theory Dynam. Systems**, to appear, 2012. arXiv:1206.0802
- [63] R.F. Williams, *One-dimensional non-wandering sets*, **Topology**, 6:473–487, 1967.
- [64] R.F. Williams, *Classification of one dimensional attractors*, In **Global Analysis (Proc. Sympos. Pure Math., Vol. XIV, Berkeley, Calif., 1968)**, American Mathematical Society, Providence, RI, 1970: 341–361.
- [65] R.F. Williams, *Expanding attractors*, **Inst. Hautes Études Sci. Publ. Math.**, 43:169–203, 1974.

ALEX CLARK, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF LEICESTER, UNIVERSITY ROAD, LEICESTER LE1 7RH, UNITED KINGDOM

E-mail address: adc20@le.ac.uk

STEVEN HURDER, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ILLINOIS AT CHICAGO, 322 SEO (M/C 249), 851 S. MORGAN STREET, CHICAGO, IL 60607-7045

E-mail address: hurder@uic.edu

OLGA LUKINA, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ILLINOIS AT CHICAGO, 322 SEO (M/C 249), 851 S. MORGAN STREET, CHICAGO, IL 60607-7045

E-mail address: lukina@uic.edu