COLLAPSING RIEMANNIAN MANIFOLDS WHILE KEEPING THEIR CURVATURE BOUNDED. II

JEFF CHEEGER & MIKHAEL GROMOV

0. Introduction

This is the second of two papers concerned with the situation in which the injectivity radius at certain points of a riemannian manifold is "small" compared to the curvature.

In Part I [3], we introduced the concept of an *F-structure* of *positive rank*. This generalizes the notion of a torus action, for which all orbits have positive dimension. We showed that if a compact manifold, Y^n, admits an F-structure of positive rank, then it also admits a family of riemannian metrics, g_δ, whose sectional curvatures are uniformly bounded independent of δ and for which the injectivity radius, $i_y(g_\delta)$ goes uniformly to zero at all points $y \in Y^n$, as $\delta \to 0$. Such a sequence is said to *collapse with bounded curvature* (see Part I for variants and refinements of the above result).

In the present paper, we prove a kind of strengthened converse to the collapsing theorem. If $y \in Y^n$, let $|K(y)|$ denote the maximum of the absolute value of the sectional curvature over $\tau \in \Lambda^2(T_y(Y^n))$.

Theorem 0.1. *There exist constants* $c_1(n), c_2(n) > 0$ *such that if* Y^n *is a complete riemannian manifold, then* $Y^n = Y_F^n \cup Y_G^n$, *where*

(1) Y_F^n *is an open set which admits an F-structure of positive rank, whose orbits,* $\mathcal{O}_y$, *have diameter satisfying* $\mathrm{diam}(\mathcal{O}_y) \le c_1(n)i_y$,

(2) *for all* $y \in Y_G^n$, *there exists* w *in the ball* $B_{i_y/c_2(n)}(y)$ *with*

$$(0.2) \qquad\qquad |K(w)|^{1/2}i_y \ge c_2(n).$$

Remark 0.3. For the F-structure we construct, the local actions almost preserve the metric. By applying Lemma 1.3 of [3], we can replace the metric on Y^n by a nearby metric which is invariant for the F-structure on Y_F^n.

Remark 0.4. The set Y_F^n can be taken to be the interior of a submanifold with boundary.

Remark 0.5. The constants $c_1(n)$ and $c_2(n)$ can be estimated explicitly, although we do not do this here. But there is one point in our construction, Proposition 3.4, which is considerably easier to treat by a noneffective argument based on the compactness theorem in riemannian geometry [2], [13], [11], [17]. For completeness, we indicate a second proof of Proposition 3.4, which yields explicit constants.

Remark 0.6. If $|K(y)|$ is uniformly bounded, say $|K(y)| \leq 1$, then the set Y_G^n has *bounded geometry*. In this case, roughly speaking, by the compactness theorem, all geometrical and topological measurements of Y_G^n can be estimated in terms of its size. Thus, the thrust of Theorem 0.1 for the case of bounded curvature is that Y^n admits a decomposition into two pieces, on each of which there is a certain kind of control. Earlier versions of this decomposition were known to Margulis (unpublished) for manifolds of negative curvature, in which case they can be obtained much more directly by special arguments; see [18] for an exposition in the case of 3-manifolds of constant negative curvature.

Remark 0.7. The hypothesis of completeness in Theorem 0.1 is just a convenience since, for an arbitrary manifold, the same decomposition holds sufficiently far from $\overline{Y}^n \backslash Y^n$ (here $\overline{Y}^n$ is the completion of Y^n).

Remark 0.8. Although there is an essentially canonical set of choices for the F-structure on M_F (which are dictated by the local geometry) there is a certain ambiguity in the construction which cannot be entirely removed. In fact, if the F-structure were uniquely determined, it would vary continuously with the local geometry. Then, of necessity, it would always be *pure* (see Part I, §1). But this would contradict the results of Part I (see Theorems 4.1 and A.1).

By combining Theorem 0.1 with the main results of Part I [3] we obtain corollaries such as the following.

Corollary 0.9. (*Critical radius*) *If a compact manifold Y^n admits a metric which is sufficiently collapsed at all points (say $|K(y)| \leq 1$, $i_y < c_2(n)$), then Y^n admits a family of metrics which collapses with bounded curvature.*

The proof of Theorem 0.1 will be given in the remaining sections.

F-structures are discussed in §1.

An *F-structure*, $\mathscr{F}$, on U consists of a *sheaf*, $\not{\ell}$, on U whose stalk, $\not{\ell}_x$, at each point $x \in U$, is isomorphic to some torus and a *local action*, μ, of $\not{\ell}$ on U, for which certain additional conditions are satisfied.

Suppose we are given a finite normal covering $\tilde{U}$ of U and a representation $\rho\colon \Gamma \to \mathrm{Sl}(k, \mathbb{Z})$ of the covering group Γ. Then ρ determines a flat T^k bundle, $\mathcal{f}$, over U. Given an action of the semidirect product, $\Gamma \times_\rho T^k$, on $\tilde{U}$, which extends the action of Γ, we obtain a local action μ of $\mathcal{f}$ on U. The pair $(\mathcal{f}, \mu)$ determines a so-called *elementary F-structure*, $\mathscr{F}$.

Typically, an F-structure is specified by a locally finite collection of open sets, $\{V_\alpha\}$, each of which carries an elementary F-structure, $\mathscr{F}_\alpha$. On nonempty intersections, $V_\alpha \cap V_\alpha$, we require that $\mathcal{f}_\alpha$ agrees with a sub-bundle of $\mathcal{f}_\beta$, or vice versa, that the corresponding local actions agree, and that $V_\alpha \cap V_\beta$ is saturated for the local action of the larger of $\mathcal{f}_\alpha$, $\mathcal{f}_\beta$. In this situation, $\mathcal{f} = \bigcup_\alpha \mathcal{f}_\alpha$.

There is a stability result for elementary structures which follows from a simple generalization of the stability theorem for compact group actions. As a consequence, a collection, $\{(V_\alpha, \mathscr{F}_\alpha)\}$ as above, for which the corresponding local actions on intersections only agree to a high degree of approximation, can be perturbed to one which determines an F-structure. This observation (see Lemma 1.5) provides the framework for the proof of Theorem 0.1. (Actually, Lemma 1.5 will be formulated in terms of the concept of *weak F-structure*, since this turns out to be more convenient for the application to the proof of Theorem 0.1; see §1 for details.)

In proving Theorem 0.1, first we find a covering of the sufficiently collapsed part of Y^n by a collection of sets which are the homeomorphic images of certain subsets of complete flat manifolds. The homeomorphisms are almost isometries. Then, we transfer to Y^n, certain elementary F-structures which are defined over these subsets. Finally, we fit together the transferred elementary F-structures, using Lemma 1.5.

The relevant discussion of elementary F-structures on complete flat manifolds is given in §2. First we describe a class of elementary F-structures of positive rank, which are carried by a noncontractible flat manifold, X^n; for these manifolds $|K(x)|^{1/2} \cdot i_x \equiv 0$. Each such structure is determined by a union of conjugacy classes, $\{\gamma_j\}$, of geodesic loops γ_j. The γ_j lie in the canonical normal abelian subgroup, $A \subset \pi_1(X^n)$, whose existence follows from the Bieberbach Theorem (and the Soul Theorem). In particular, a loop γ lies in A if the rotational angles of its holonomy are not too big.

Next we describe the elementary F-structures which are utilized in the proof of Theorem 0.1. Each of these is specified by a collection of loops at x which lie in A, with the following property. A loop γ is in the

collection if and only if every loop at x with the same length and iso-morphic holonomy transformation is also included. These collections are not necessarily invariant under conjugation by elements of $\pi_1(X^n)$ and, in general, the corresponding elementary F-structures are only defined over proper open subsets, $V \subset X^n$.

In §3, we show that if $|K(w)|^{1/2} \cdot i_y$ is sufficiently small for w near $y \in Y^n$, then we can find an open neighborhood U of y, a complete flat manifold $\mathbf{Y}^n$ and a quasi-isometry $f \colon U \to T_u(S^m)$. Here $T_u(S^m)$ is the u-tubular neighborhood of a soul, $S^m \subset \mathbf{Y}$. The quasi-isometry, f, is almost an isometry if $|K(w)|^{1/2} \cdot i_y$ is sufficiently small.

In §4, this approximation is regularized so that holonomies P_y and P_y of corresponding loops γ and γ in U and $T_u(S^m)$ are close if the loops are not too long.

With the results of §§3 and 4, we can transfer an elementary F-structure from a subset of $T_u(S^m)$ to a subset of U. Moreover, a structure so obtained has an approximate description in terms of geodesic loops of Y^n.

The proof of Theorem 0.1 is carried out in §5, by implementing Lemma 1.5.

If $y \in Y^n$ is a point such that $|K(w)|^{1/2} i_y$ is small for w near y, then there exist various local flat approximations to (Y^n, y) as in §§3 and 4. To each such point y, we assign a flat approximation $f_y \colon U_y \to T_{u_y}(S_y)$, a thin subset V_y, with $y \in V_y \subset U_y$, and an elementary F-structure, $\mathcal{F}_y$, as above, over V_y.

The main point is to make these choices such that on all intersections, $V_{y_1} \cap V_{y_2}$, either $\mathcal{F}_{y_1} \supseteq \mathcal{F}_{y_2}$ or vice versa. This condition is called property (F_1); compare the discussion above, of the contents of §1.

Since the corresponding local actions for both $\mathcal{F}_{y_1}$ and $\mathcal{F}_{y_2}$ have an approximate description in terms of geodesic loops of Y, these actions will be close if the maps f_{y_1} and f_{y_2} are sufficiently close to being isome-tries. In fact, were it not for the fact that $\{V_y\}$ has infinite multiplicity, $\{(V_y, \mathcal{F}_y)\}$ would actually satisfy the hypothesis of Lemma 1.5.

Thus, if we choose a locally finite subcollection, $\{V_{y_\alpha}\}$, with suitably bounded multiplicity, then the full hypothesis of Lemma 1.5 is satisfied for the collection $\{(V_{y_\alpha}, \mathcal{F}_\alpha)\}$ and we obtain a weak F-structure (of positive rank). Our particular method of selecting $\{(V_y, \mathcal{F}_y)\}$ (which guarantees that property (F_1) holds) will also enable us to conclude that our weak F-structure is actually an F-structure.

A more detailed outline of the argument is given at the beginning of §5.

In the Appendix to §2 we give some examples which show that the elementary F-structures discussed in §2 which are defined over all of X^n do not satisfy the hypothesis of §1, since the size of their orbits grows too rapidly at infinity.

Let us mention that by replacing the compactness theorem used in §3 by one proved recently by M. Anderson (see his preprint "Convergence and Rigidity of Manifolds under Ricci Curvature Bounds") the hypothesis of Theorem 0.1 can be replaced by the following assumptions: In (0.2), one can substitute "Ricci curvature" for "sectional curvature," provided one also assumes that for some sufficiently small constant, $c_3(n)$,

$$(0.10) \qquad \int_{B_{i_y/c_3(n)}(y)} |R|^{n/2} < c_3(n).$$

Finally, we point out that K. Fukaya has obtained a number of remarkable results on collapsing in the case of bounded curvature and diameter; see [7]–[10]. His techniques are rather different from those employed here and in [4]. In recent joint work with Fukaya, a common generalization of a portion of his work and ours is obtained by combining the two approaches.

1. F-structures and their stability

Before beginning we recall an elementary fact which is used (sometimes without further mention) in this section and the next.

Let G be a connected topological group which acts on a space Z. Then this action lifts (necessarily uniquely) to the action of a covering group, $\tilde{G}$, on a covering space, $(\tilde{Z}, \tilde{z})$, if and only if

$$(\phi_z)_*(\pi_1(\tilde{G}, \tilde{e})) \subset \pi_1(\tilde{Z}, \tilde{z}) \subset \pi_1(Z, z),$$

where $\phi_z(g) \overset{\text{def}}{=} g(z)$.

Equivalently, let $\tilde{G}$, the universal covering of G, act on $\tilde{Z}$, the universal covering of Z. If $G = \tilde{G}/H$ and $Z = \tilde{Z}/\Gamma$ then the action of $\tilde{G}$ on $\tilde{Z}$ descends to an action of G on Z if and only if the action of G normalizes that of Γ and $H \subset \Gamma$.

For the convenience of the reader, we begin by reviewing some definitions from [3] (to which we refer for further details).

A *partial action*, A, of a topological group, G, on a Hausdorff space, X, is given by

(1) a neighborhood $\mathscr{D} \subset G \times X$ of $e \times X$, where e is the identity of G, and a continuous map $A : \mathscr{D} \to X$, also written $(g, x) \to gx$, such that

(2) $(g_1 g_2) x = g_1 (g_2 x)$ whenever $(g_1 g_2, x)$ and $(g_1, g_2 x)$ lie in $\mathscr{D}$, and such that $ex = x$ for all x.

Two partial actions $(A_1, \mathscr{D}_1)$ and $(A_2, \mathscr{D}_2)$ are called equivalent if there is a neighborhood $\mathscr{D} \subset \mathscr{D}_1$, $\mathscr{D}_2$ containing $e \times X$, such that $A_1 | \mathscr{D} = A_2 | \mathscr{D}$. A *local* action, $\{A\}$, is an equivalence class of partial actions.

Assume G is connected.

A subset $X_0 \subset X$ is called $\{A\}$-invariant if for some (equivalently, any) representative we have $gx \in X_0$ for all $x \in X_0$ with $(g, x) \in \mathscr{D}$. It is easy to see that the X is partitioned into minimal invariant sets called *orbits*. Let $\mathscr{O}_x$ denote the orbit of X.

A local action can be restricted to any open set $U \subset X$ by restricting the domain, $\mathscr{D}$, of some representative to $\mathscr{D}' \supset e \times X$, such that $gx \in U$ for $(g, x) \in \mathscr{D}'$. Similarly a local action can be pulled back under a locally homeomorphic map.

Now consider a sheaf, $\mathscr{g}$, of connected topological groups over X. Let $\mathscr{g}(U)$ denote the group of sections over U. An *action* of $\mathscr{g}$ on X is a local action of $\mathscr{g}(U)$ on U, for every connected open set $U \subset X$, such that the structure homomorphisms $\mathscr{g}(U) \to \mathscr{g}(U')$ (for $U' \subset U$) commute with the restriction of local actions.

A set is *invariant* if its intersection with U is invariant for all U. Again, X is partitioned into minimal invariant subsets called *orbits*. A set is called *saturated* if it is a union of orbits. The *rank* of the action at $x \in X$ is the dimension of the orbit, $\mathscr{O}_x$. The action has *positive rank* if $\dim \mathscr{O}_x > 0$, for all $x \in X$.

An action of $\mathscr{g}$ is called *complete* if for all $x \in X$ there is an open neighborhood, $V(x)$, of x and a locally homeomorphic map, $\tilde{V}(x) \to V(x)$ ($\tilde{V}(x)$ Hausdorff), such that:

(1) If $\pi(\check{x}) = x$, then for any open neighborhood $W \subset \tilde{V}(x)$ of $\check{x}$, the structure homomorphism, $\pi^*(\mathscr{g})(W) \to \mathscr{g}_{\check{x}} \overset{\text{def}}{=} \mathscr{g}_x$ is an isomorphism.

(2) The local action of $\pi^*(\mathscr{g})$ comes from a global action of $\pi^*(\mathscr{g})(\tilde{V}(x)) = \mathscr{g}_{\check{x}}$.

Definition 1.1. A $\tilde{\mathscr{G}}$-*structure* on X is given by the complete action of a sheaf of connected topological groups, $\mathscr{g}$, on X, such that the neighborhood, $\tilde{V}(x)$, can be chosen to satisfy:

(1) $\pi : \tilde{V}(x) \to V(x)$ is a normal covering.

(2) For all x, $V(x)$ is saturated.

(3) For an orbit, $\mathcal{O}$, if $x, y \in \mathcal{O}$, then $V(x) = V(y)$.

Definition 1.2. A $\tilde{\mathcal{G}}$-structure is called an *F-structure* if

(1) For all x, the stalk, $\mathcal{G}_x$, is isomorphic to a torus.

(2) For all x, the normal covering, $\tilde{V}(x) \to V(x)$, can be chosen to be finite.

A structure satisfying (1) and (3) of Definition 1.2 (but not necessarily (2)) is called a *weak $\tilde{\mathcal{G}}$-structure*. A weak $\tilde{\mathcal{G}}$-structure which satisfies the additional conditions of Definition 1.2 is called a *weak F-structure*.

We emphasize that the existence of a weak F-structure of positive rank does *not* guarantee that we can perform the collapsing constructions of [3]. However, we will formulate Lemma 1.5 in terms of this concept, since this turns out to be convenient for the application to the proof of Theorem 0.1.

For the remainder of this section we restrict attention to F-structures (although everything we say generalizes to $\tilde{\mathcal{G}}$-structures).

Definition 1.3. An F-structure is called *elementary* if $\tilde{V}(x) \to V(x)$ can be chosen independent of x.

Note that in Definition 1.3, necessarily, we have $V(x) = X$. Also, as indicated in the introduction, the concept of elementary F-structure can be reformulated as follows.

Suppose we are given

(1) a (possibly disconnected) finite normal covering, $\tilde{X} \to X$, with covering group Γ,

(2) a representation, $\rho: \Gamma \to \operatorname{Aut}(t^k)$, for some torus T^k,

(3) an action of the semidirect product, $\Gamma \times_\rho T^k$, extending the action of $\gamma \subset \Gamma \times_\rho T^k$.

The above data determines an elementary F-structure, $\mathcal{F}$ on X, for which the sheaf, $\mathcal{f}$, is the associated flat bundle on X, with fiber isomorphic to T^k and holonomy representation isomorphic to ρ. The action of $T^k \subset \Gamma \times_\rho T^k$ on $\tilde{X}$ determines an obvious action of $\mathcal{g}$ on X.

For $\mathcal{F}$ as above, let $\mathcal{f}' \subset \mathcal{f}$ be a sub-bundle with fiber $T^{k'} \subset T^k$. Then the action of $\mathcal{f}$ restricts to an action of $\mathcal{f}'$. Moreover, the restriction of $\mathcal{f}'$ to any set U' which is saturated by the orbits of $\mathcal{f}'$ determines an elementary F-structure over U'.

Typically, an F-structure is determined by specifying the following data.

Let $\{V_\alpha\}$ be a locally finite collection of open subsets of X and, for each α, let $\mathcal{F}_\alpha = (\mathcal{f}_\alpha, \mu_\alpha)$ be an elementary F-structure over V_α. Assume

that

(F_1) for all α, β, either $f_\alpha | V_\alpha \cap V_\beta$ is a sub-bundle of $f_\beta | V_\alpha \cap V_\beta$ or vice versa;

(F_2) in the former case, μ_α is obtained restricting μ_β and $V_\alpha \cap V_\beta$ is saturated for μ_β.

Note that in (F_1) above, we allow $f_\alpha | V_\alpha \cap V_\beta$ to coincide with $f_\beta | V_\alpha \cap V_\beta$.

Obviously, a collection, $\{V_\alpha, \mathscr{F}_\alpha\}$, satisfying (F_1) and (F_2) determines an F-structure, $\mathscr{F}$, over $\bigcup_\alpha V_\alpha$, for which the associated sheaf, f, is $\bigcup_\alpha f_\alpha$.

If we replace condition (F_2) by

$(F_2)^w$ in the former case, μ_α is obtained by restricting μ_β and $V_\alpha \cap V_\beta$ is saturated for μ_α,

then a collection satisfying (F_1) and $(F_2)^w$ determines a weak F-structure.

In the proof of Theorem 0.1, we will apply Lemma 1.5 to obtain a collection satisfying (F_1) and $(F_2)^w$. But, it will turn out that two additional conditions $((F_3)$ and $(F_4))$ are satisfied. These guarantee that the weak F-structure is actually an F-structure.

(F_3) If $V_{\alpha_0}, \cdots, V_{\alpha_l}$ is any sequence such that, for $i = 0, \cdots, l-1$, $V_{\alpha_i} \cap V_{\alpha_{i+1}} \neq \varnothing$ and f_{α_i} is *properly* contained in $f_{\alpha_{i+1}}$ on $V_{\alpha_i} \cap V_{\alpha_{i+1}}$, then f_{α_0} extends over $\bigcup_0^l V_{\alpha_i}$.

(F_4) If $V_{\beta_0}, \cdots, V_{\beta_{l'}}$ is a second such sequence and $V_{\alpha_l} \cap V_{\beta_{l'}} \neq \varnothing$ then the extensions of f_{α_0}, f_{β_0} to V_{α_l}, $V_{\beta_{l'}}$ satisfy $f_{\alpha_0} \subseteq f_{\beta_0}$ or vice versa on $V_{\alpha_l} \cap V_{\beta_{l'}}$.

Note that the extension of f_{α_0}, assumed to exist in (F_3), is necessarily unique.

Let $s(\alpha)$ denote those β for which there exists a sequence as in (F_3) with $\alpha = \alpha_0$ and $\beta = \alpha_l$. Put $W_\alpha = \bigcup_{\beta \in s(\alpha)} V_\beta$. Then if (F_3) and (F_4) hold, we claim that $\{(W_\alpha, f_\alpha)\}$ satisfies (F_1) and (F_2). Hence $\{(W_\alpha, f_\alpha)\}$, or, equivalently, $\{(V_\alpha, \mathscr{F}_\alpha)\}$, determines an F-structure.

Observe that the part of condition (F_2) which relates to the actions is automatic. Also, $W_{\alpha_0} \cap W_{\beta_0}$ is a union of sets, $V_{\alpha_l} \cap V_{\beta_{l'}}$ as in (F_4), and we can assume that $f_{\alpha_l} = f_{\beta_{l'}}$. For if, say, f_{α_l} is properly contained in $f_{\beta_{l'}}$, then $V_{\beta_{l'}} \subset W_{\alpha_0}$ and we can replace the sequence $V_{\alpha_0}, \cdots, V_{\alpha_l}$ by $v_{\alpha_0}, \cdots, V_{\alpha_l}, V_{\beta_{l'}}$. Thus, $V_{\alpha_l} \cap V_{\beta_{l'}}$ is saturated for $f_{\alpha_l} = f_{\beta_{l'}}$ and hence for f_{α_0} and f_{β_0}. Therefore, (F_2) holds. (F_1) is obvious from (F_4).

The main result of this section, Lemma 1.5, says essentially that if (F_1) is satisfied and $(F_2)^W$ holds to a high degree of approximation, then the collection can be perturbed to one for which both (F_1) and $(F_2)^W$ hold. This is a consequence of the stability theorem for compact group actions, in the form given in [14] (compare also [16]).

We begin by adapting their theorem to our context.

Let $V_j \subset X$ be open sets, $j = 1, 2$. Let $(V_j, \mathcal{F}_j, \mu_j)$ be an elementary F-structure such that μ_j is induced by an action of $\Gamma_j \times_{\rho_j} T^k$ on a normal covering space, $\tilde{V}_j \overset{\pi_j}{\to} V_j$. We suppose that $\mathcal{F}_1|V_1 \cap V_2$ agrees with a subbundle, $\mathcal{F}_{12}$, of $\mathcal{F}_2|V_1 \cap V_2$.

Let $T^k = S^1 \times \cdots \times S^1$ and let $d(g)$ denote the distance of $g \in T^k$ from the identity element, under the metric obtained by averaging the product metric under the holonomy of $\mathcal{F}_1|V_1 \cap V_2$. Assume that V_j has a metric $\langle\ ,\ \rangle_j$, which is invariant for μ_j and let $V_j^\rho \subset V_j$ denote the set of points at distance $\geq \rho$ from ∂V_j for the metric $\langle\ ,\ \rangle_j$. Assume that the injectivity radius for $\langle\ ,\ \rangle_j$ is bounded below by $\frac{1}{2}$ and that the sectional curvature is bounded by 1 in absolute value. Finally, assume there is a $\frac{1}{2}$-quasi-isometry between $\langle\ ,\ \rangle_1$ and $\langle\ ,\ \rangle_2$ (see (3.3).

Let $x_0 \in V_1 \cap V_2$ and let $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2$ be representative partial actions for μ_1, μ_2 on some contractible neighborhood W of x_0. If $d(g)$ is sufficiently small, we define $\eta(g): W \to X$ by $\eta(g) = \boldsymbol{\mu}_2(g^{-1})\boldsymbol{\mu}_1(g)$. We say that $(\mu_1, \mathcal{F}_1)$ $(\mu_2, \mathcal{F}_{12})$ are δ (C^1-close) on $V_1 \cap V_2$, if for all such x_0, g the map $\eta(g)$ is $d(g)\delta$ (C^1-close) to the inclusion, $W \hookrightarrow S$.

Let $\phi: V_1^\rho \to V_1$ be an imbedding which is ε (C^1-close) (in the sense of [14]) to the inclusion, with $\varepsilon < \frac{1}{2}$. Since the injectivity radius of the metric $\langle\ ,\ \rangle_1$ is $> \frac{1}{2}$, there is a natural identification of $(\phi^{-1})^*(\mathcal{F}_1 V_1^\rho)$ with $\mathcal{F}_1|\phi(V_1^\rho)$. This identification is understood implicitly in (2) and (4) of Lemma 1.4 below.

Lemma 1.4. *For all $1 > \rho > 2\varepsilon > 0$, there exists $\delta = \delta(\rho, \varepsilon, N) > 0$ such that if $(\mu_1, \mathcal{F}_1)$ and $(\mu_2, \mathcal{F}_2)$ are δ (C^1-close), and the coverings $\tilde{V}_j \to V_j$ have order, $N_j \leq N$, then there exists an embedding, $\phi: V_1^\rho \to V_1$, with the following properties:*

(1) *ϕ is ε (C^1-close) to the inclusion $V_1^\rho \hookrightarrow V_1$ and $\phi(x) = x$ for $x \in V_1 \setminus V_2^{\rho/2}$.*

(2) *$(\mathcal{F}_1, \phi\mu_1\phi^{-1})$ agrees with $(\mathcal{F}_{12}, \mu_2)$ on $\phi(V_1) \cap V_2$.*

(3) *If for some $x \in \phi(V_1^\rho)$ and all g with $d(g)$ sufficiently small, we have $\eta(g)(x) = x$, then $x \in V_1^\rho$ and $\phi(x) = x$.*

(4) $(\phi(V_1^\rho), \mathcal{F}_1, \phi\mu_1\phi^{-1})$ and $(V_2^\rho, \mathcal{F}_2, \mu_2)$ determine an F-structure over $\phi(V_1^\rho) \cup V_2^\rho$.

Proof. Consider the subset $\mu_1(\pi_1^{-1}(V_1^{\rho/4} \cap V_2^{\rho/4})) \subset \tilde{V}$, the saturation of $\pi_1^{-1}(V_1^{\rho/4} \cap V_2^{\rho/4})$ by the action of T^k which lifts μ_1. By writing an arbitrary element $g \in T^k$ as $g \in h^m$, where h is sufficiently close to the identity, and then comparing with the local action of the lift of μ_2, we easily find that for δ sufficiently small,

$$\mu_1(\pi_1^{-1}(V_1^{\rho/4}) \cap V_1^{\rho/4}) \subset \pi_1^{-1}(V_1 \cap V_2).$$

We also obtain the corresponding statement with the roles of μ_1 and μ_2 reversed (for the action of $\mu_2(\mathcal{F}_{12})$).

Let $\widetilde{V_1 \cap V_2} \xrightarrow{\pi} V_1 \cap V_2$ be a common covering of $\pi_j^{-1}(V_1 \cap V_2)$. We can assume $\widetilde{V_1 \cap V_2}$ is normal and of order $N < N^2$. Put $N/N_j = l_j$. The action of $T^k = \mathbb{R}^k/\mathbb{Z}^k$ on $\mu_j(\pi_j^{-1}(V_1^{\rho/4} \cap V_2^{\rho/4}))$ lifts to an action of $\mathbb{R}^k/l_j\mathbb{Z}^k$ on the inverse images of $\mu_j(\pi_j^{-1}(V_1^{\rho/4} \cap V_2^{\rho/4}))$ in $\widetilde{V_1 \cap V_2}$. By composing with the homomorphisms $\tilde{T}^k = \mathbb{R}^k/l_1l_2\mathbb{Z}^k \to \mathbb{R}^k/l_j\mathbb{Z}^k$, we obtain actions $\tilde{\mu}_1$ and $\tilde{\mu}_2$ of the same torus on these inverse images (in general, these actions are noneffective). Let Γ_j and Γ denote the covering groups of $\pi_j^{-1}(V_1 \cap V_2)$ and $\widetilde{V_1 \cap V_2}$. By using the homomorphisms $\tilde{\rho}_j: \Gamma \to \Gamma_j \xrightarrow{\rho_j} \mathrm{Sl}(k, \mathbb{Z})$, we extend $\tilde{\mu}_j$ to an action of the semidirect product $\Gamma \times_{\tilde{\rho}_j} \tilde{T}^k$.

Since the order of the covering $\tilde{T}^k \to T^k$ is bounded (by N^2) it is clear that if μ_1 and μ_2 are C^1-close, then $\tilde{\mu}_1$ and $\tilde{\mu}_2$ are C^1-close on the intersection of their domains (write $g = h^m$ as above).

If δ is sufficiently small, we can restrict the domains of the $\tilde{\mu}_j$ to obtain domains $\tilde{W}_j$ for $\tilde{\mu}_j$ such that

$$\pi^{-1}(V_1^\rho \cap V_2^\rho) \subset \tilde{W}_1 \subset \tilde{W}_2 \subset \pi^{-1}(V_1^{\rho/2} \cap V_2^{\rho/2})$$

and the boundaries of these sets are at mutual distance at least $\rho/24$ for $\langle \, , \, \rangle_1$. Again, for δ sufficiently small, the argument of [14] gives an embedding, $\tilde{\psi}: \tilde{W}_1 \to \tilde{W}_2$, as C^1-close as we like to the inclusion, satisfying $\tilde{\psi}\tilde{\mu}_1 = \tilde{\mu}_2\tilde{\psi}$. Moreover, $\tilde{\psi}$ is the identity at points at which $\tilde{\mu}_1$ and $\tilde{\mu}_2$ agree locally.

Put $\pi(\tilde{W}_j) = W_j$. The embedding $\tilde{\psi}$ induces $\psi: W_1 \to W_2$, satisfying $\psi\mu_1 = \mu_2\psi$ with ψ as C^1-close to the inclusion as we like. Let U_1 be

invariant for μ_1 and satisfy $V_1^\rho \cap V_2^\rho \subset U_1 \subset W_1$, with the boundaries of these sets at mutual distance at least $\rho/100$ for the metric $\langle\ ,\ \rangle_1$. By using the Isotopy Extension Theorem, we can find an embedding $\phi\colon W_1 \to W_1$, as C^1-close to ψ as we like, such that $\phi|U_1 = \psi|U_1$, ϕ is the identity near ∂W_1, and $\phi(x) = x$ if $\psi(x) = x$. Then we can extend ϕ to all of V_1 by making it the identity off W_1. Finally, we can assume that ϕ is close enough to the inclusion so that $\phi(V_1^\rho) \cap V_2^\rho \subset \phi(U_1)$. The resulting map satisfies (1)–(4).   q.e.d.

Let $\{V_\alpha\}$ be a covering. Assume there are at most N_1 of those sets whose intersection with any fixed V_{α_0} is nonempty. Let $\mathscr{F}_\alpha = \{\not\!f_\alpha, \mu_\alpha\}$ be a collection of elementary F-structures over the sets $\{V_\alpha\}$ such that condition (F_1) above holds. Assume that the orders of the coverings $\tilde{V}_\alpha \to V_\alpha$ are all $\leq N_2$ and that the fibers of the $\not\!f_\alpha$ all have dimension $\leq N_3$. Finally, assume that each V_α carries an invariant metric for μ_α, with injectivity radius $\geq \frac{1}{2}$ and curvature ≤ 1 in absolute value and that these metrics are $\frac{1}{2}$-quasi-isometric on intersections.

In the following lemma we identify $(\phi_\alpha^{-1})^*(\not\!f_\alpha|V_\alpha^\rho)$ with $\not\!f_\alpha|\phi_\alpha(V_\alpha^\rho)$ as in (2) and (4) of Lemma 1.4.

Lemma 1.5. *For all $1 > \rho > 2\varepsilon > 0$, there exists $\delta = N_1 2^{N_1} \cdot \delta(\rho, \varepsilon, N_2, N_3) > 0$ such that if for all α, β (say) $\not\!f_\alpha|V_\alpha \cap V_\beta$ agrees with $\not\!f_{\alpha,\beta}|V_\alpha \cap V_\beta$ (where $\not\!f_{\alpha,\beta} \subset \not\!f_\beta$), and $(\not\!f_\alpha, \mu_\alpha)$ and $(\not\!f_{\alpha,\beta}, \mu_\beta)$ are δ (C^1-close), then there are embeddings $\phi_\alpha\colon V_\alpha^\rho \to V_\alpha$, with $\rho' \leq \rho$, such that the following holds:*

(1) *For all α, the embedding ϕ_α is ε (C^1-close) to the inclusion $V_\alpha^{\rho'} \to V_\alpha$.*

(2) *The collection $\{(\phi_\alpha(V_\alpha^{\rho'}), \not\!f_\alpha, \phi_\alpha\mu_\alpha\phi_\alpha^{-1})\}$ satisfies (F_1) and $(F_2)^w$, and hence determines a weak F-structure over $\bigcup_\alpha \phi_\alpha(V_\alpha^{\rho'})$.*

Proof. Consider the collections $\alpha = (\alpha_0, \cdots, \alpha_j)$ of indices such that $V_{\alpha_0} \cap \cdots \cap V_{\alpha_j}$ is maximal with respect to the property of having nonempty intersection. Choose an enumeration, $\alpha_1, \alpha_2, \ldots$, of these. For each α_j, we can reorder the subscripts, $\alpha_k \in \alpha_j$ such that on $V_{\alpha_0} \cap \cdots \cap V_{\alpha_j}$, we have $\not\!f_{\alpha_1} \subseteq \not\!f_{\alpha_2} \subseteq \cdots \subseteq \not\!f_{\alpha_i}$.

Now we go through the α_j in order and for each one we do the following. Order the pairs $(\alpha_k, \alpha_{k'})$ with $k < k'$ by $(\alpha_k, \alpha_{k'}) < (\alpha_l, \alpha_{l'})$ if $k' < l'$ or $k' = l'$ and $k < l$. Then run through these pairs in descending order. At each stage apply Lemma 1.4, with $\rho/(N_1 2^{N_1})$, $\varepsilon/(N_1 2^{N_1})$

in place of ρ, ε to the subsets $V_{\alpha_k}^{\cdot}$, $V_{\alpha_{k'}}^{\cdot}$ of V_{α_k}, $V_{\alpha_{k'}}$, which have produced possible previous applications of Lemma 1.4, at earlier stages of the process.

We claim that the above process produces a collection for which (1) and (2) hold.

To see this let $x \in \bigcup_\alpha \phi_\alpha(V_\alpha^{\rho'})$ and let $\alpha(x)$ be the set of those α with $x \in V_\alpha$. Let $\alpha_{j_1} < \alpha_{j_2} < \cdots$ where $j_1 < j_2 < \cdots$, be those α_j which contain $\alpha(x)$ and put $\overline{\alpha}(x) = \alpha_{j_1}$. By referring to (3) of Lemma 1.4 we see that if the actions on those $V_\alpha^{\cdot}$ with $\alpha \in \alpha(x)$ agree at the point x, after the stage of the process corresponding to $\overline{\alpha}(x)$ has been concluded, then they do not change during the remainder of the process.

It suffices to check that after this stage has been concluded, all of these actions agree at x. Recall that Lemma 1.4 is applied for each pair of subscripts α_k, $\alpha_{k'} \in \alpha(x)$, with $k < k'$. Moreover, these pairs are con­sidered in descending order and the action is changed only on a subset of $V_{\alpha_k}^{\cdot}$. Thus, we can assume that for some α_l with $l > k'$, the actions for the pairs (α_k, α_l) and $(\alpha_{k'}, \alpha_l)$ are compatible before the step corresponding to $(\alpha_k, \alpha_{k'})$ but the actions corresponding to (α_k, α_l) are not compatible after this step. However, by (3) of Lemma 1.4 (and induction) this does not happen.   q.e.d.

2. Elementary F-structures on complete flat manifolds

(a) **Preliminaries; short loops.**   Let M^n be a complete riemannian manifold. For c a curve in M^n, let $L[c]$ denote the length of c.

Given curves c_1 and c_2 with the same end points, we say that c_1 and c_2 are *short homotopic*, if they are homotopic keeping end points fixed, through curves of length at most $\max_j L[c_j]$.

Let $m \in M^n$. Let R_m be the largest number such that $\exp_m |B_{R_m}(0) \subset M_m^n$ is nonsingular. If c is closed with $c(0) = m$, $L[c] < R_m$, then c is short homotopic to a unique geodesic loop γ on m. Suppose, in particular, that $c = \gamma$ and that τ is a curve with $\tau(0) = m$. Let τ^s denote $\tau|[0, s]$. As long as the closed curve $\tau^s \cup \gamma \cup -\tau^s$, on $\tau(s)$, is homotopic to a geodesic loop γ_s on $\tau(s)$, with $L[\gamma_s] < R_{\tau(s)}$, then γ_s is unique. We say that γ_s is obtained from $\gamma_0 = \gamma$ by *sliding* along τ. The map, $\gamma_0 \to \gamma_s$, is compatible with the isomorphism between $\pi_1(X^n, \tau(0))$ and $\pi_1(X^n, \tau(s))$ induced by τ^s.

If γ_1 and γ_2 are geodesic loops on m with $L[\gamma_1] + L[\gamma_2] < R_m$, then $\gamma_1 \cup \gamma_2$ is short homotopic to a unique geodesic loop, $\gamma_1 * \gamma_2$. In particular, if $R_m = \infty$, then $\pi_1(M^n, m)$ is isomorphic to the group of geodesic loops on m with the product $*$. In this case, a loop at m gives rise to a collection of loops $\{\gamma\}_{m_1}$ at each point $m_1 \in M^n$, each of which is free homotopic to γ. The collection $\{\gamma\}_{m_1}$ represents a conjugacy class in $\pi_1(M^n, m_1)$.

Let i_m denote the injectivity radius at m.

Lemma 2.1. *There is a constant $c(n)$ such that if the sectional curvature of M^n satisfies $|K| \leq 1$ and $\Lambda i_m < \pi/2$ $(\Lambda > 0)$, then there are at most $c(n)\Lambda^n$ geodesic loops on m of length $\leq \Lambda i_m$.*

Proof. Each loop γ lifts to a segment of a ray, $\hat{\gamma}$, through the origin in M_m^n. Clearly, there exists $c(n)$ such that if there are more than Λ^n geodesics of length at most $\Lambda \cdot i_m$ then endpoints of some pair $\hat{\gamma}_1$, $\hat{\gamma}_2$ are at distance less than $2i_m(\Lambda i_m / \sinh \Lambda i_m)$. It follows that the loop which is short homotopic to $\gamma_1 * \gamma_2^{-1}$ has length $< 2i_m$. This is a contradiction.

(b) **Geometry of complete flat manifolds.** Let X^n be a complete flat manifold. Write $X^n = \overline{X}^l \times R^k$, isometrically, where $\overline{X}^l$ has no Euclidean factor. Then $\overline{X}^l$ contains a unique compact flat totally geodesic submanifold, S^m, the *soul*, such that $\overline{X}^l$ is isometric to the total space of the normal bundle $\nu(S^m)$ (see [3; 19, Theorem 3.3]). There the metric on $\nu(S^m)$ is induced by its natural flat connection.

Note that any tubular neighborhood $T_u(S^m)$ $(u \geq 0)$ is *totally convex*, i.e., any geodesic with endpoints in $T_u(S^m)$ lies in $T_u(S^m)$.

From now on we assume $k < n$, or, equivalently, $m > 0$.

Let S^m be a soul of X^n and let $\tilde{S}^m \xrightarrow{\pi} S^m$ denote the holonomy covering of the compact flat manifold $\tilde{S}^m$. By Bieberbach's theorem, $\tilde{S}^m$ is isometric to a flat torus and $\tilde{S}^m \to S^m$ has order at most $\lambda(n)$, for some constant $\lambda(n)$ depending only on n $(\geq m)$. Since $S^m \hookrightarrow X^n$ is a homotopy equivalence, we can regard $\mathbb{Z}^k \simeq A = \pi_1(\tilde{S}^m)$ as a normal subgroup of $\pi_1(X^n)$. Clearly, A is independent of the particular choice S^m.

Let γ be a geodesic loop with orientation preserving holonomy, having all its rotational angles $< \pi/\lambda(n)$ in absolute value. We write $\mathrm{rot}(P_\gamma) < \pi/\lambda(n)$. In this case $\gamma \in A$. In fact, let τ be a minimal geodesic with $\tau(l) = \gamma(0)$ and $\tau(0)$ the point on S^m closest to $\gamma(0)$. By sliding γ along τ we obtain a geodesic loop $\gamma_s \subset T_s(S^m)$ at $\tau(s)$. In particular, $P_{\gamma_0} \simeq P_\gamma$ (since X^n is flat), $P_{\gamma_0} \subset S^m$, and the claim follows from Bieberbach's theorem.

Note that $L[\gamma_s]$ is given by the increasing function

$$(2.2) \qquad L[\gamma_s] = (L^2[\gamma_0] + (2\sin\theta/2)^2 s^2)^{1/2},$$

where P_{γ_0} rotates $\tau'(0)$ through an angle θ. This follows by an elementary argument after one lifts γ_0 to the universal covering space of X^n.

If $\gamma_0 \in A$, then γ_0 is automatically smooth closed since it lifts to a loop $\tilde{\gamma}_0$ contained in the torus $\tilde{S}^m$.

(c) **Elementary F-structures.** We will explain how a finite subset of A which is invariant under conjugation by elements of $\pi_1(X^n)$ and for which the corresponding holonomy transformations are orientation preserving, gives rise to an elementary F-structure. This construction depends on a suitable set of choices of logarithms for the holonomy transformations.

Let (w, e^B) represent an isometry of R^n, with translational part w. Put $w = w' + w''$, where $e^B(w') = w'$ and w'' is orthogonal to the $+1$-eigenspace of e^B. Let $(1 - e^B)^{-1}w''$ denote the unique inverse image of w'' orthogonal to $\ker(1 - e^B)$. Then the curve

$$(2.3) \qquad t \to (tw' + (1 - e^{Bt})(1 - e^B)^{-1}w'', e^{Bt})$$

is a 1-parameter subgroup passing through (w, e^B) at $t = 1$. The orbit, $\mathscr{O}$, of the origin, is the curve $t \to tw' + (1 - e^{Bt})(1 - e^B)^{-1}w''$. Let L be the length of the restriction of this curve to the interval $0 \leq t \leq 1$. An elementary computation shows that

$$(2.4) \qquad \|w\| \leq L \leq \left[\|w'\|^2 + \left(\frac{\lambda}{2\sin\lambda/2}\right)\|w''\|^2\right]^{1/2},$$

where λ is the largest eigenvalue of B which is not an integral multiple of $2\pi i$.

Let $\{(w_j, e^{B_j})\}$ be a collection of mutually commuting isometries, such that the $\{B_j\}$ are mutually commuting skew symmetric transformations with no eigenvalue of the form $2\pi i k$, for $k \neq 0$. By a trivial calculation, for all j, k, we have

$$(2.5) \qquad (1 = e^{B_k})w_j = (1 - e^{B_j})w_k,$$

$$(2.6) \qquad (1 - e^{B_k})w_j' = (1 - e^{B_j})w_k' = 0.$$

It follows easily that the subgroups given by (2.4) are mutually commuting.

Conversely, let $\{g_j\}$ be mutually commuting elements of $\mathrm{SO}(n)$. Then we can find skew symmetric transformations, $\{B_j\}$, such that $e^{B_j} = g_j$, the $\{B_j\}$ are mutually commuting, and each B_j has no eigenvalue of

the form $2\pi i k$ for $k \neq 0$. In particular, if $\mathrm{rot}(g_j) < \pi$, then the B_j are uniquely determined if we require $\|B_j\| < \pi$. In any case, given a mutually commuting set $\{(w_j, g_j)\}$, we can obtain mutually commuting 1-parameter subgroups as above.

Now assume that the $\{(w_j, e^{B_j})\}$ form a group $\Delta \simeq \mathbb{Z}^k$ of covering transformations of R^n. Given a finite subset $\{(w_j, e^{B_j})\}$, $j = 1, \cdots, N$, we obtain an action of the Cartesian product of the corresponding 1-parameter subgroups on R^n, which descends to a T^N action on R^n/Δ (see the discussion at the beginning of §1). This action need not be effective, but an effective action can be obtained by passing to a quotient of T^N.

Example 2.7. Let $(w, e^{B(\theta)})$ denote the isometry of R^3 such that w is a translation in the direction of a unit vector along the x-axis and $B(\theta)$ is given by the matrix

$$(2.8) \qquad\qquad B(\theta) = \begin{pmatrix} 0 & -\theta \\ \theta & 0 \end{pmatrix}$$

in the y, z-plane. The isometries $(w, e^{B(\theta)})$ and $(2w, e^{B(2\theta)})$ generate a group $\Lambda = \Delta \simeq \mathbb{Z}$ (we assume θ, $2\theta \not\equiv 0 \bmod 2\pi$). The construction above gives a noneffective T^2 action on R^3/Δ, inducing an effective action of T^1. If we use $B(2\theta - 2\pi)$ in place of $B(2\theta)$, we obtain an effective T^2 action. Note for $0 < \theta < \pi/2$, $|2\theta| < \pi$ while for $\pi/2 < \theta < \pi$, $|(2\theta - 2\pi)| < \pi$.

Now suppose that π_1 is a group of covering transformations and that $\Delta \simeq \mathbb{Z}^k$ is a normal subgroup of finite index $\leq \lambda(n)$. Suppose $\{(w_j, e^{B_j})\}$, $j = 1, \cdots, N$, is invariant under conjugation by elements of π_1. Then there is an induced representation $\rho \colon \pi_1/\Delta \to \mathrm{Aut}(T^N)$, which together with the action of T^n on R^n/Δ determines an elementary F-structure on R^n/π_1.

The F-structure just constructed can also be described in terms of geodesic loops on $X^n = R^n/\pi_1$. Identify $X_x^n \simeq \mathbb{R}^n$ with the universal covering space of X^n. Then the group of isometric covering transformations is isomorphic to the group of geodesic loops at x. The element corresponding to a loop, γ, can be recovered as $(V_\gamma, P_{-\gamma})$; where V_γ denotes translation by $L[\gamma] \cdot \gamma'(0)$ and $-\gamma$ denotes γ transversed in the opposite sense.

A collection $\{\gamma_j\}_x$, $j = 1, \cdots, N$, of conjugacy class of loops, $\gamma_j \in A$, determines an elementary F-structure, $\mathscr{F}$, on X^n. In the sequel we are always concerned with the case $\mathrm{rot}(P_{\gamma_j}) < \pi/\lambda(n)$. Note that for any

$\gamma_j \in A$, the conjugacy class $\{\gamma_j\}_x$ contains at most $\lambda(n)$ loops. The fiber, ℓ_{x_1}, at an arbitrary point $x_1 \in X^n$ of the sheaf (flat T^N bundle) ℓ, associated to $\mathscr{F}$, can be identified with the Cartesian product of loops in $\{\gamma_j\}_{x_1}$.

We now describe a class of elementary F-structures which, in general, are defined only over proper subsets of X^n. These will be used in the construction of the F-structure of Theorem 0.1.

Let $[P_\gamma]$ denote the isomorphism class of P_γ.

Let $\gamma_1, \cdots, \gamma_N$ be loops at $x \in X^n$ which lie in A. Fix $\varepsilon > 0$. Assume that if $\gamma \in A$ and $\gamma \neq \gamma_i$ for any i, then for all i, at least one of the following holds:

$$(2.9) \qquad |L[\gamma] - L[\gamma_i]| \geq \varepsilon$$

or

$$(2.10) \qquad [P_\gamma] \neq [P_{\gamma_i}].$$

Let $\mathscr{F}'$ be any elementary F-structure as above on X^n and let $T_\varepsilon(\mathscr{O}'_x)$ denote the open tubular neighborhood of the orbit, $\mathscr{O}'_x$, of radius ε.

Lemma 2.11. (1) *At each* $x_1 \in T_{\varepsilon/4}(\mathscr{O}'_x)$ *there are exactly* N *loops,* $\hat{\gamma}$, *which, for some* i, *satisfy*

$$(2.12) \qquad |L[\hat{\gamma}] - L[\gamma_i]| \leq \varepsilon/2, \qquad [P_{\hat{\gamma}}] = [P_{\gamma_i}].$$

(2) *The collection* $\hat{\gamma}_1, \cdots, \hat{\gamma}_N$ *of such loops is the collection obtained from* $\gamma_1, \cdots, \gamma_N$ *under homotopy in* $T_{\varepsilon/4}(\mathscr{O}'_x)$; *i.e., sliding a loop,* γ_i, *from* x *to* x_1 *along any curve* $c \subset T_{\varepsilon/4}(\mathscr{O}'_x)$ *gives a loop,* $\hat{\gamma}_j$, *for some* j.

Proof. Note first that sliding a loop, γ, does not change $[P_\gamma]$. Then, by an obvious continuity argument, (2) implies (1).

Since $A \subset \pi_1(X^n)$ is normal, the collection of loops at x lying in A can be obtained by sliding the collection of loops lying in A at x along any curve c. If $x_1 \in T_{\varepsilon/4}(\mathscr{O}'_x)$, there is a minimal geodesic σ of length $s < \varepsilon/4$ connecting x_1 to a point on $\mathscr{O}'_x$. Since sliding a loop along σ changes its length by at most $2s < \varepsilon/2$, it suffices to assume $x_1 \in \mathscr{O}'_x$ and to show that for some curve c, from x_1 to x_2, sliding loops of A along c leaves their lengths unchanged.

Let $\tilde{x} \in \tilde{X}$ be a lift of x, let $\tilde{\gamma}$ be a loop at $\tilde{x}$ lifting $\gamma \in A$, and let T^l be the torus corresponding to $\mathscr{F}'$, which acts on $\tilde{X}^n$. We can find a curve $g(t) \subset T^l$ with $g(0)$ the identity element and $g(1)\tilde{x} = \tilde{x}_1$ a lift of x_1. The curve $g(t)$ projects to a curve c from x to x_1 and by an obvious

continuity argument, $g(1)(\tilde{\gamma})$ projects to the loop obtained by sliding γ along c. Since $g(1)$ is an isometry, our claim follows.   q.e.d.

Let $\gamma_1, \cdots, \gamma_N, \gamma_{N+1}, \cdots, \gamma_{N'}$ be a collection of loops at x which lie in A and let $\mathscr{F}'$ be the elementary F-structure determined by the union of conjugacy classes, $\{\gamma_j\}_x$, $1 \leq j \leq N'$. Assume that $\gamma_1, \cdots, \gamma_N$ satisfy (2.9) and (2.10) above. Then Lemma 2.11 implies

Corollary 2.13. *The set* $\gamma_1, \cdots, \gamma_N$ *is invariant under conjugation in* $\pi_1(T_{\varepsilon/4}(\mathscr{O}'_x))$ *and hence defines an elementary F-structure,* $\mathscr{F}$, *over* $T_{\varepsilon/4}(\mathscr{O}'_x)$.

Let $\gamma \in A$ be a loop at x, with lift $\tilde{\gamma}$ at $\tilde{x}$. For the circle action on $\tilde{X}^n$ corresponding to $\tilde{\gamma}$, the orbit of $\tilde{x}$ (counted with multiplicities) is homotopic to $\tilde{\gamma}$ (see (2.3)). Since $L[\tilde{\gamma}] > 0$ is of shortest length in its homotopy class, the orbit of $\tilde{x}$ has positive length. Thus, the elementary F-structures constructed above all have *positive rank.*

Clearly, an orbit of any elementary structure as above lies at constant distance from any soul, S^m. The maximum size of an orbit is controlled by the upper bound in (2.4). If $\|B_j\| < \pi$ for all j, then the orbit in X^n corresponding to the jth circle in $T^N = S^1 \times \cdots \times S^1$, has length at most $\frac{\pi}{2} L[\gamma_j]$.

Remark 2.14. The injectivity radius need not be constant on orbits. However, in view of the obvious relation

$$(2.15) \qquad\qquad i_x \leq i_{\tilde{x}} \leq \lambda(n) i_x,$$

the ratio of the maximum value of the injectivity radius to the minimum value, on an orbit, is bounded by $\lambda(n)$.

Appendix to §2: Growth of the injectivity radius

We claim that it is not possible to assign to each complete flat manifold, X^n, an elementary F-structure, $\mathscr{F}(X^n)$, of the type considered in §2, in such a way the ratio of the diameter of the orbit, $\mathrm{diam}(\mathscr{O}_x)$, to the injectivity radius, i_x, remains uniformly bounded as x and X^n vary.

Suppose first that the rotational angles of P_{γ_0} are all rational multiples of 2π, for some loop γ_0 on $z \in S^m$. Then

$$\gamma_0 \overbrace{* \cdots *}^{N} \gamma_0 \stackrel{\mathrm{def}}{=} N\gamma_0$$

has trivial holonomy, for some smallest integer N. Let τ be a geodesic normal to S^m with $\tau(0) = z$. Let $N\gamma_s$ be the geodesic at $\tau(s)$ obtained by sliding $N\gamma_0$ along τ. Then $L[N\gamma_x] = L[N\sigma_0]$.

On the other hand, if σ_0 is any loop on z with $\sphericalangle(P_\gamma(\tau'(0)),\ \tau'(0)) = \theta > 0$ then $L[\sigma_s]$ grows linearly along τ (see (2.1)).

It follows that those elementary F-structures constructed in §2, for which the diameter of the orbits does not grow linearly in almost all directions, are precisely the ones generated by loops with trivial holonomy.

The following example is typical.

Example A.1. Let X_θ^3 be the total space of the flat 2-plane bundle over S^1 with holonomy θ. For each $\theta = \frac{p}{q}\pi$ (with $\frac{p}{q} < 1$ in lowest terms) there is an elementary F-structure with sublinear (actually constant) asymptotic growth with orbits, $\mathcal{O}_{\tau(s)}$, of length $qL[S^1] = 2i_{\tau(s)}$, for s large. Then however,

$$(A.2) \qquad\qquad L[\mathcal{O}_{\tau(s)}]/i_{\tau(s)} \approx 2q$$

for s small. Here q can be taken arbitrarily large.

If X^n is such that there exists no geodesic loop with rational holonomy, then for all γ_0, the function $L[\gamma_s]$ grows linearly in almost all directions. Hence, the same holds for the orbits of any elementary F-structure arising from the construction of §2. But the injectivity radius itself always satisfies the following estimate (put $i_s = i_{\tau(s)}$).

Lemma A.3. *For say* $s > \frac{1}{3}i_0$,

$$(A.4) \qquad\qquad i_s \le c(n)[\mathrm{Vol}(S^m)]^{1/m+c}s^{c/m+c},$$

where

$$(A.5) \qquad\qquad c = [(n-m)/2].$$

Proof. We can assume $i_{\tau(0)} = 1$. There are at least $c_1(n)r^m/\mathrm{Vol}(S^m)$ geodesic loops in S^m on $\tau(0)$ of length $\le r$, where $r \ge i_{\tau(0)}$. At least one of these, σ, has $\mathrm{rot}(P_\sigma) \le \varepsilon \cdot \pi$, if

$$(A.6) \qquad\qquad c_1(n)\frac{r^m}{\mathrm{Vol}(S^m)} = \varepsilon^{-c}.$$

Then, by (2.2),

$$(A.7) \qquad\qquad L[\sigma_s] \le \left(r^2 + \left(2s \cdot \sin\frac{\varepsilon}{2}\pi\right)^2\right)^{1/2}.$$

Given s, choose r and ε, which satisfy (A.6) and

$$(A.8) \qquad\qquad r = \varepsilon s.$$

Then

$$(A.9) \qquad L[\sigma_s] \le (r^2 + (\varepsilon s)^2)^{1/2} \le \sqrt{2}r$$
$$= c(n)(\mathrm{Vol}(S^m))^{1/m+c}s^{c/m+c}. \qquad \text{q.e.d.}$$

Let $X^n = X^l \times R^{n-l}$. The isometry group of $\tilde{X}^l$ is generated by a collection of circle actions, one for each set of generators for $\pi_1(\tilde{S}^m) \simeq \mathbb{Z}^m$ and the orthogonal transformations of the normal bundle $\nu(S^m)$ (leaving S^m pointwise fixed) which centralize the holonomy group. The function $\tilde{i}_x$ is constant on orbits and the isometry group is transitive on fibers of $\nu(\tilde{S}^m)$.

Lemma A.10. *Let $\sigma(s)$ be a normal geodesic in X^n and put $i_{\sigma(s)} = i_s$. Then for all s*

(A.11) $$i_s \leq i_0 + 2s,$$

and for say $s \geq \frac{1}{3}i_0$,

(A.12) $$i_s \leq c(n)i_0^{1/c+1}s^{c/c+1} = c(n)i_0(s/i_0)^{c/c+1}.$$

Proof. The estimate in (A.12) is clear. The proof of (A.13) is completely analogous to that of (A.4). We just restrict attention to multiples of a fixed loop.

3. Local approximation by complete noncontractible flat manifolds

Let Y^n be a complete riemannian manifold and let $y \in Y^n$. Set

(3.1) $$v(y, R) \stackrel{\text{def}}{=} \sup_{B_{R \cdot i_y}(y)} |K(w)|^{1/2} i_y.$$

By Theorem 4.3 of [5] (see also [6]) it follows that

(3.2) $$i_w \geq i_y \min(\pi/v(y, R), c(n))e^{-(n-1)R \cdot v(y, R)}.$$

If U_1 and U_2 are riemannian manifolds and $f: U_1 \rightarrow U_2$ is a C^1-smooth quasi-isometry, let $M(f)$ denote the infimum of those ε such that if $V(y, \delta^{-1}) \leq \delta$,

(3.3) $$e^{-\varepsilon}g_1 \leq f^*(g_2) \leq e^\varepsilon g_1.$$

The following proposition will allow us to transfer the elementary F-structures on complete noncontractible flat manifolds which were discussed in §2 to more general manifolds.

Proposition 3.4. *Given a continuous decreasing function $h: (0, \infty) \rightarrow (0, \infty)$ and $k > 0$, there exist $\delta = \delta(h, k, n)$, $R(h, k, n)$, such that if $v(y, \delta^{-1}) \leq \delta$, then there exists*

(i) *a complete flat manifold Y^n and a soul $S \subset Y^n$,*

(ii) *a quasi-isometry, $f: U \rightarrow T_u(S^m)$, with $u < R(h, k, n)i_y$, and U an open neighborhood of y, such that*

(iii) $M(f) \leq h(u/i_y)$,

(iv) $\max(i_y, \overline{f(y), S}, \operatorname{diam}(S)) < u/k$,

(v) $i_y = i_{f(y)}$.

Proof. Assume the contrary. Then (after possible rescaling) there are sequences (Y_j^n, y_j) such that $i_{y_j} = 1/v(y_j, j) \leq 1/j$ and either there exists no f as above satisfying (iii) and (iv) or the smallest u for which there exists such an f is $\geq j$. By the compactness theorem in riemannian geometry, there is a pointed C^∞ manifold $(\mathbf{Y}^n, \mathbf{y})$ with a $C^{1,\alpha}$ riemannian metric (for all $\alpha > 1$) such that for some infinite subsequence (Y_{j_s}, y_{j_s}), and any r, the sequence of balls $B_r(y_{j_s})$ converges in the Lipschitz metric to $B_r(\mathbf{y})$. Clearly, $\mathbf{Y}^n$ is complete flat and noncontractible ($i_y = 1$). In particular its metric is C^∞. Since $i_y = 1$, $\overline{\mathbf{y}, S^m} < \infty$, $\operatorname{diam}(S^m) < \infty$ for some soul $S^m \subset \mathbf{Y}^n$, we obtain a contradiction.

Remark 3.5. Although the fact that h can be chosen to be an arbitrary decreasing function of r is of interest in describing the local geometry of the manifolds considered in Proposition 3.4, for the application to the proof of Theorem 0.1 it will suffice to choose h to be a sufficiently small constant.

Remark 3.6. Lipschitz convergence (i.e., (iii) above) is actually not strong enough for our purposes since we will want to compare holonomies around corresponding loops in Y^n, $\overline{Y}^n$ and not just their lengths. In fact the versions of the compactness theorem proved in [11] or [17] show that in harmonic coordinates the convergence of metric tensors actually takes place in the $C^{1,\alpha}$ topology. The compactness theorem as stated in [13] would also suffice. However, in order to emphasize the elementary nature of our result, we show in the next section, by a simple direct argument, that Lipschitz convergence implies C^1 convergence, in case the limit is flat. For this result we do not require a special coordinate system.

Example 3.7. Fix $\theta > 0$ and let E_θ^3 denote the complete flat manifold obtained by dividing R^3 by the group of isometries generated by the isometry $(w, e^{B(\theta)})$ of Example 2.7. Let S be the soul of E_θ^3. We will show directly that Proposition 3.4 holds for the family (E_θ^3, y), where y is a variable point in E_θ^3.

Observe that if γ_y is a shortest geodesic loop at y, then the holonomy, P_{γ_y}, converges to the identity transformation as $\overline{y, S} \to \infty$. This is an immediate consequence of the discussion of the Appendix to §2.

Let S_l^1 denote the circle of length l. Then $0 \times S_l^1$ is a soul of the riemannian product $R^2 \times S_l^1$. Fix $k > 0$. It follows easily from the observation above that for $\overline{y, S}$ sufficiently large there exists a neighborhood U_y of y and a quasi-isometry $f_y: U_y \to T_{2ki_y}(0 \times S_{2i_y}^1)$, with $f(y) \in 0 \times S_{2i_y}^1$. Moreover, $M(f_y) \to 0$ as $\overline{y, S} \to \infty$.

Given a function h as in Proposition 3.4, choose Λ such that $M(f_y) < h(2k)$, if $\overline{y, S} > \Lambda$. For such points, the quasi-isometry, f_y, satisfies the conditions of Proposition 3.4 (with $u = 2ki_y$, $u/k = 2i_y$). Moreover, we can take $R(k, h, 3) = 2k$ for the subfamily consisting of the (E_θ^3, y) with $\overline{y, S} > \Lambda$.

The set of points for which $\overline{y, S} \leq \Lambda$ is compact. Thus, for all these points, we can take f_y to be the identity map on a sufficiently large tubular neighborhood of S. Then we take $R(k, h, 3)$ for the whole family (E_θ^3, y) to be the larger of $2k$ and the radius of this tube.

In order to estimate explicitly the constants $c_1(n)$ and $c_2(n)$ in Theorem 0.1, it is necessary to give a proof of Proposition 3.4 which does not depend on an argument by contradiction. We now briefly outline such an argument; details will appear elsewhere.

(1) Rescale the metric on Y such that $i_y = 1$ and view $B_R(g)$ as the quotient of a ball on the tangent space by an isometric pseudo-group, Γ. In the spirit of [12] (see also [1]), we can imitate the proof of the Soul Theorem for flat manifolds, given in [19, Theorems 3.2.8 and 3.3.3]. In this way we obtain a group, Γ, which acts isometrically in $\mathbb{R}^n$ and freely on a large ball about the origin. Moreover, Γ has an abelian subgroup, $A \simeq \mathbb{Z}^k$, of index $\leq \lambda(n)$. Finally, Γ is isomorphic to a subpseudogroup of Γ.

(2) By deforming the action of Γ slightly if necessary, we can assume that Γ acts freely on $\mathbb{R}^n$.

(3) By a generalization of the argument of Example 3.7, after making a second small deformation of the action of Γ, we can assume that the bounds of (iv) of Proposition 3.4 hold for $\mathbb{R}^n/\Gamma$.

(4) Finally we construct a quasi-isometry f between a slightly smaller ball $B_{R'}(y) \subset B_R(y)$ and a ball in $\mathbb{R}^n/\Gamma$. Here we use the result of [15] to take care of the finite group Γ/A.

4. Regularization of the approximation

Let $y \in Y^n$ and let $f: U \to T_u(S^m)$ be as in Proposition 3.4. Let H_f denote the Hessian of f.

Proposition 4.1. *The constants $\delta(h, k, n)$ and $R(h, k, n)$ can be chosen such that there exists f satisfying* (i)–(iv) *of Proposition 3.4 and the additional estimate*

$$\text{(4.2)} \qquad\qquad \|H_f\| \le h(u/i_y).$$

The idea of the proof is to regularize f by convolving with a suitable smoothing kernel. For an arbitrary map, this would only have the effect of making the Hessian bounded. But by using the fact that f maps U to a flat space with $M(f)$ small, it will follow that the Hessian of the regularized map is actually small.

Proof of Proposition 4.1. We can assume $i_y = 1$.

Let $\bar{\psi}(s): [0, 1] \to [0, 1]$ be a C^∞ function such that $\bar{\psi} \equiv 1$ near $s = 0$ and $\bar{\psi} \equiv 0$ near $s = 1$. Put $\bar{\psi}_\lambda(s) = \bar{\psi}(s/\lambda)$. Let $w_1, w_2 \in Y^n$ and denote the distance from w_1 to w_2 by $\overline{w_1, w_2}$. Finally, let ω denote the volume form on Y^n. Put

$$\text{(4.3)} \qquad\qquad \psi_\lambda(w_1, w_2) = \frac{\bar{\psi}_\lambda(\overline{w_1, w_2})}{\int_{B_2(w_1)} \bar{\psi}(w - 1, w_2)\omega},$$

where the integration is with respect to w_2.

Choose $\delta = \delta(h_1, 2k, n)$ where $h_1 < \frac{1}{10}$ is to be determined later (see Proposition 3.5). If $v(y, \delta^{-1}) \le \delta$, standard estimates give

$$\text{(4.4)} \qquad\qquad \|d\psi_\lambda\| \le c(\delta)\lambda^{-1},$$

$$\text{(4.5)} \qquad\qquad \|H_{\psi_\lambda}\| \le c(\delta)\lambda^{-2},$$

on $B_{\delta^{-1}-\lambda}(y)$.

Let $f: U \to T_u(S^m)$ be the map provided by Proposition 3.4. Lemma A.3 and Remark A.10 give a lower bound, i_0, for i_z on $T_u(S^m)$. If we choose

$$\text{(4.6)} \qquad\qquad \lambda < \tfrac{1}{2}i_0,$$

then for all $y_1 \in U$, the range of $f|B_\lambda(y_1)$ is contained in a convex subset of a flat space. Hence,

$$\text{(4.7)} \qquad\qquad f_\lambda = \int \psi_\lambda(w_1, w_2)f(w_2)\omega$$

is well defined.

Let l be a real valued, affine linear function, with $\|l\| \le 1$ on $B_r(y) \subset T_u(S^m)$ where $r < \frac{1}{2} i_0$. Then

(4.8) $$\|H_{f_\lambda}\| = \|H_{(l \circ f)_\lambda}\|.$$

Up to a constant, any l as above can be written in the form

(4.9) $$l = \frac{\rho_{a_1}^2 - \rho_{a_2}^2 - \mathbf{d}^2}{2\mathbf{d}},$$

where ρ_{a_j} is the distance function from $a_j \in T_u(S^m)$, and $\mathbf{d} = \overline{a_1, a_2} = \frac{1}{2}$. Let $f(y_j) = a_j$ and consider the function

(4.10) $$l = \frac{\rho_{y_1}^2 - \rho_{y_2}^2 - \mathbf{d}^2}{2\mathbf{d}}.$$

Then by (4.4) (and (4.5)) l has differential everywhere close to 1, small Hessian and is uniformly close to $l \circ f$. The explicit bounds depend on h_1. It suffices to estimate $H_{(l \circ f - l)_\lambda}$. Since $l \circ f - l$ is arbitrarily small for suitably small h_1, it is clear that given h, we can choose h_1 such that f_λ will satisfy (3.2).   q.e.d.

Let $f: U \to T_u(S^m)$ be as in Propositions 3.4 and 4.1. Let $\gamma \subset U$ be a geodesic loop on y with $L[\gamma] < R_y$, where $\exp_y B_{R_y}(0) \subset Y_y^n$ is nonsingular. Let $\gamma \subset \mathbf{Y}^n$ be the unique geodesic loop which is short homotopic to $f(\gamma)$.

Corollary 4.11. *Put $h = h(u/i_y)$. Then*

(4.12) $$e^{-h} L[\gamma] = L[\gamma] \le e^h L[\gamma],$$

(4.13) $$\sphericalangle(\gamma'(0), df^{-1}(\gamma'(0)) \le \frac{c(n) \cdot L[\gamma] \cdot h}{i_y},$$

(4.14) $$\|P_\gamma - df^{-1} P_\gamma \, df\| < \frac{c(n) \cdot L[\gamma] \cdot h}{i_y}.$$

Proof. Relation (4.12) follows from the minimizing properties of γ, γ and (iii) Proposition 3.4. By using, in addition, (4.2), relations (4.13) and (4.14) also follow by straightforward arguments.

Suppose that for γ as above, $NL[\gamma] < R_y$. Let $N\gamma$ denote the unique geodesic loop which is short homotopic to the N-fold iterate of γ. Then

we have

Corollary 4.15.

$$(4.16) \qquad \|P_{N\gamma} - (P_\gamma)^N\| \leq \frac{N \cdot c(n) \cdot L[\gamma] \cdot h}{i_y}.$$

Proof. This follows immediately from Corollary 4.11 and the fact that the holonomy of a curve depends only on its homotopy class in the flat case.

5. Construction of the F-structure

(a) **Outline of the construction.** In this section we prove our main result, Theorem 0.1, by using the results of §§2, 3, and 4 to implement Lemma 1.5.

Our basic strategy was sketched in §0. Given a complete riemannian manifold Y^n, let Y_δ^n denote the set of points at which $v(y, \delta^{-1}) < \delta$ (see (3.1)). To each $y \in Y_\delta^n$ (δ sufficiently small) we assign a set, V_y, containing y, and an elementary F-structure, $\mathscr{F}_y$, over V_y. This is done in such a way that $\{(V_y, \mathscr{F}_y)\}$ satisfies all the conditions of Lemma 1.5, apart from the bound, N_1, on the multiplicity. Then we extract a suitable locally finite subcover $\{V_{y_\alpha}\}$. The collection $\{(V_{y_\alpha}, \mathscr{F}_{y_\alpha})\}$ satisfies the hypothesis of Lemma 1.5 and leads to the desired F-structure.

In this subsection, we outline the steps involved in selecting $\{(V_y, \mathscr{F}_y)\}$ and $\{(V_{y_\alpha}, \mathscr{F}_{y_\alpha})\}$. Further details are given in subsections (b)-(g) (which correspond to Steps 1-6 below).

Step 1. To each point $y \in Y_\delta^n$ we assign a set of short geodesic loops $[\gamma_j]_y$, with $\mathrm{rot}(P_{\gamma_j}) < \pi/3\lambda(n)$ ($\lambda(n)$ as in §2). Our choice depends only on the lengths of the short loops at y and on the isomorphism classes of their holonomy transformations. Moreover, the following precursor of property (F_1) holds. If y_1, y_2 are sufficiently close, then $[\gamma_j]_{y_1}$ contains or is contained in $[\gamma_j]_{y_2}$. (As usual we identify loops at y_1 with loops at y_2 by sliding them along the unique minimal geodesic from y_1 to y_2.)

Step 2. Let $f_y: U_y \to T_{u_y}(S_y)$ be any map as provided by Proposition 3.4. The set of loops of $T_{u_y}(S_y)$ corresponding to $[\gamma_j]_y$ determines an elementary F-structure, $\mathscr{F}_y$, over a neighborhood, $\mathbf{V}_y$ of $f_y(y)$, as in Corollary 2.13. The fiber of the corresponding elementary F-structure, $\mathscr{F}_y$, over $V_y = f_y^{-1}(\mathbf{V}_y)$ can be identified with the Cartesian product of the loops in $[\gamma_j]_y$. It follows that the collection $\{(V_y, \mathscr{F}_y)\}$ satisfies a weak version of property (F_1): If y_1 and y_2 are sufficiently close, either $\mathscr{F}_{y_2} \supseteq \mathscr{F}_{y_1}$, or vice versa, on $V_{y_1} \cap V_{y_2}$.

Step 3. Clearly, V_{y_1} and V_{y_2} can have nonempty intersection even if y_1 and y_2 are not close. But by using Lemma 2.9 and Remark 2.14, we find that property (F_1) holds for $\{(V_y, \mathscr{F}_y)\}$.

Step 4. On a set $V_{y_1} \cap V_{y_2}$ the closeness of corresponding local actions for $\measuredangle_{y_1}$ and $\measuredangle_{y_2}$ is determined by the deviation from isometry (in the C^2-topology) of the maps f_{y_1} and f_{y_2}. This is an immediate consequence of the description of elementary F-structures in terms of geodesic loops, for the flat case discussed in §2.

To apply Lemma 1.5 to a subcollection, $\{(V_{y_\alpha}, \mathscr{F}_{y_\alpha})\}$, these deviations must be small relative to the size of the V_{y_α} and the multiplicity, N_1, of $\{V_{y_\alpha}\}$.

Step 5. By a simple variant of a standard packing construction, we select a subcover, $\{V_{y_\alpha}\}$, with $\bigcup_\alpha V_{y_\alpha} \subset Y_\delta^n$, whose multiplicity, N_1, is bounded by $c(n)$.

Step 6. By the results of §4, the deviation from isometry (in the C^2-topology) of a map f_y is controlled by the function h of Proposition 3.4. In view of the bound of Step 5, it suffices to take $h(r) = \varepsilon(n)$, for $\varepsilon(n) > 0$ sufficiently small. Then the covering $\{V_{y_\alpha}\}$ satisfies the hypothesis of Lemma 1.5. The weak F-structure obtained by applying Lemma 1.5 is easily seen to have properties (F_3) and (F_4) of §1. Hence it is an F-structure.

(b) **Assigning short loops to points.** Our procedure for choosing the collections $[\gamma_j]_y$ is based on some trivial observations about sequences.

Let $b_1 \leq b_1 \leq \cdots \leq b_M$ be a nondecreasing sequence such that for some $c_1 < c_2$ and $N \leq M$

(5.1) $b_1 \leq c_1 < c_2 < b_{N+1}.$

Clearly, there exists at least one index, $J \leq N$, such that

(5.2) $b_j + \dfrac{c_2 - c_1}{2N} \leq b_{J+1},$

(5.3) $b_J \leq \dfrac{c_1 + c_2}{2}.$

Remark 5.4. The collection of all such J depends only on the subsequence, $b_1 \leq b_2 \leq \cdots \leq b_N$.

The following lemma is obvious.

Lemma 5.5. *Let $b_1' \leq b_2' \leq \cdots \leq b_M'$ be a second sequence and let π be a permutation of $\{1, \cdots, M\}$ such that for $j \leq M$,*

(5.6) $|b_j - b_{\pi(j)}'| < \dfrac{c_2 - c_1}{4N}.$

Then if J satisfies (5.2), π *preserves the sets* $\{0, \cdots, J\}$ *and* $\{J + 1, \cdots, M\}$.

Choose a nondecreasing function $\phi: [0, \pi] \to [0, \infty)$, with

(5.7) $$\phi|[0, \pi/6\lambda(n)] \equiv 1$$

and

(5.8) $$\phi|[\pi/3\lambda(n), \pi] \equiv 6(6\lambda(n))^{[n/2]}.$$

Define a function $a(\gamma)$ on loops at y by

(5.9) $$a(\gamma) = \phi(\mathrm{rot}(P_\gamma)) \cdot L[\gamma].$$

Clearly, we have $L[\gamma] \leq a(\gamma)$.

   Lemma 5.10.   *For $\delta \leq \delta_0$ sufficiently small,*

(5.11) $$\min_\gamma a(\gamma) \leq 2(6\lambda(n))^{[n/2]} \cdot i_y.$$

The inequality

(5.12) $$L[\gamma] \leq a(\gamma) \leq 6(6\lambda(n))^{[n/2]} \cdot i_y$$

holds for at most $N = N(n)$ loops. For all such loops

(5.13) $$\mathrm{rot}(P_\gamma) \leq \frac{\pi}{3\lambda(n)}.$$

   Proof.   Let γ be a shortest loop at y. Thus, $L[\gamma] = 2i_y$,. By Corollary 4.15 and the standard packing argument, if δ_0 is sufficiently small, there exists k such that

(5.14) $$L[k\gamma] \leq 2ki_y \leq 2(6\lambda(n))^{[n/2]},$$

(5.15) $$\mathrm{rot}(P_{k\gamma}) \leq \frac{\pi}{3\lambda(n)}.$$

Lemma 2.1 implies (5.12), and (5.13) is clear from (5.8) and (5.9).   q.e.d.

   From now on, we assume $\delta \leq \delta_0$ as above.

   Let $y \in Y_\delta^n$ and let $\gamma_1, \gamma_2, \ldots$ be an ordering of the loops at y such that

(5.16) $$a(\gamma_1) \leq a(\gamma_2) \leq \cdots.$$

It follows from Lemma 5.5 that there exists a smallest index, $J \leq N = N(n)$, such that

(5.17) $$a(\gamma_J) + \frac{c_2 - c_1}{2N} \leq a(\gamma_{J+1}),$$

with $c_1 = 2(6\lambda(n))^{[n/2]}$ and $c_2 = 2c_1$.

Define $[\gamma_j]_y$ to be the set $\{\gamma_1, \cdots, \gamma_J\}$. Note that the ordering $\gamma_1, \gamma_2,$ $\ldots$ need not be uniquely determined if the numbers $\{a(\gamma_j)\}$ are not all distinct. However, the set $[\gamma_j]_y$ is independent of the choice of ordering. Also, by (5.13), for $\gamma_j \in [\gamma_j]_y$

$$(5.18) \qquad \qquad \mathrm{rot}(P_{\gamma_j}) \le \frac{\pi}{3\lambda(n)}.$$

Lemma 5.19. *There exists $0 < \varepsilon(n) < 1$ such that if $y_1, y_2 \in Y_\delta^n$ and $\overline{y_1, y_2} \le \varepsilon(n)i_{y_1}$, then either $[\gamma_j]_{y_1} \supseteq [\gamma_j]_{y_2}$ or vice versa.*

Proof. Let $\overline{y_1, y_2} \le \varepsilon i_{y_1}$. Let $\{\gamma_1^k \cdots \gamma_{T_k}^k\} = \mathscr{L}^k$, $k = 1, 2$, be the loops at y_k, with $h(\gamma_1^k) \le \cdots \le h(\gamma_{T_k}^k) \le c_2$. By Remark 5.4, the sets $[\gamma_j]_{y_k}$ are determined by $\{a(\gamma_1^k), \cdots, a(\gamma_{T_k}^k)\}$ or by any larger subsets of $\{a(\gamma_1^k), \ldots\}$.

If $\varepsilon \le \varepsilon(n)$, it is clear that by using (4.14), we can find subsets $\widetilde{\mathscr{L}}^k \supset \mathscr{L}^k$, which are identified with each other under the correspondence between loops at y_1 and y_2, and such that for $\gamma_j^k \in \widetilde{\mathscr{L}}^k$,

$$(5.20) \qquad \qquad a(\gamma_j^k) \le 3c_2 = 6(6\lambda(n))^{[n/2]}.$$

Let $b_1 \le \cdots \le b_M$ $(M \le N)$ be the sequence obtained by arranging the numbers $\{a(\gamma_j^1)\}$, $\gamma_j^1 \in \widetilde{\mathscr{L}}^1$, in ascending order. Let $b_1' \le \cdots \le b_M'$ be obtained similarly from $\widetilde{\mathscr{L}}^2$. Let π be the permutation of $\{1, \cdots, M\}$ induced by the correspondence between $\widetilde{\mathscr{L}}^1$ and $\widetilde{\mathscr{L}}^2$. Our claim now is a direct consequence of Lemma 5.10.

(c) **Assigning elementary F-structures to points.** A map $f: U \to T_u(S^m)$ as in Proposition 3.4 is determined by a number $k > 0$ and a decreasing function $h(r)$. In what follows, it will suffice to choose h to be a sufficiently small constant, and to take

$$(5.21) \qquad \qquad k = 18(6\lambda(n))^{[n/2]}.$$

For each $y \in Y_\delta^n$ $(\delta \le \delta_0$ sufficiently small) we can, by Proposition 3.4, find a map $f_y: U_y \to T_{u_y}(S_y)$. Note that by our choice of k, each loop of $[\gamma_j]_y$ is contained in U. Let $[\gamma_j]_{f_y(y)}$ denote the collection of loops at $f_y(y)$ which are homotopic to the images of $[\gamma_j]_y$. By Corollary 2.13 $[\gamma_j]_y$ determines an elementary F-structure, $\mathscr{F}_y$, over a neighborhood $\mathbf{V}_y \overset{\mathrm{def}}{=} T_{r_y}(\mathscr{O}_{f_y(y)})$. Here we take

$$(5.22) \qquad \qquad r_y = \frac{t}{6}(6\lambda(n))^{-[n/2]} \min_\gamma a(\gamma),$$

where the minimum is over all loops γ at y. The number $t < 1$ will be specified below. Note that $r_y \leq \frac{t}{3} i_y$.

Let $\mathscr{F}_y$ be the elementary F-structure over $V_y = f_y^{-1}(\mathbf{V}_y)$. The fiber $(\mathcal{F}_y)_y$ of $\mathcal{F}_y$ at y can be identified with the Cartesian product of the loops in $[\gamma_j]_y$. Thus, it follows from Lemma 5.19 that if $\overline{y_1, y_2} < 2(r_{y_1} + r_{y_2})$ and $t < \varepsilon(n)$ for $\varepsilon(n) > 0$ sufficiently small, then $\mathcal{F}_{y_1} \subseteq \mathcal{F}_{y_2}$ on $V_{y_1} \cap V_{y_2}$ or vice versa. This is the precursor of property (F_1) of §1.

(d) **Property (F_1) for $\{(V_y, \mathscr{F}_y)\}$.** By Lemma 2.9, the set of values which the function $a(\gamma)$ takes on loops of $A \subset \pi_1(T_{u_y}(S_y))$ is constant on orbits of the elementary F-structure $\mathscr{F}_y$. Let $N = N(n)$ and let t be as in parts (b) and (c). For each point $y_1 \in V_y$, consider the set consisting of the N smallest values (counted with multiplicities) of the function a. Then if $\varepsilon(n)$ is sufficiently small, $t < \varepsilon(n)$, and f_y is sufficiently C^2-close to being an isometry, the above set of values is as close as we like to being independent of the point y_1. Now, the argument of part (c) shows that $\{(V_y, \mathscr{F}_y)\}$ has property (F_1).

(e) **Closeness of local actions.** The fiber of $\mathcal{F}_y$ at y can be identified with the Cartesian product of at most $N(n)$ loops (see (5.17)) of length bounded by (5.12). By (2.3), (2.4), Corollary 4.11, and (4.16) we can insure that the local actions of $\mathcal{F}_{y_1}$ and $\mathcal{F}_{y_2}$ are as C^1-close as we like on $V_{y_1} \cap V_{y_2}$, provided that h of Proposition 3.4 and δ_0 above are sufficiently small. (In measuring the closeness of local actions we rescale the metric so that, say, $i_{y_1} = 1$, to conform to the context of Lemma 1.5.)

The degree of closeness required in Lemma 1.5 depends on the maximum fiber dimension, N_3, on the maximum order, N_2, of a covering space associated to the elementary F-structure and on the multiplicity, N, of the covering. In our situation $N_3 \leq N(n)$ and $N_2 \leq \lambda(n)$. In part (f) below, we will extract a subcovering, $\{V_{y_\alpha}\}$ of $\{V_y\}$, with bounded multiplicity.

(f) **The subcover $\{v_{y_\alpha}\}$.** Let $q(y)$ denote the number of loops in $[\gamma_j]_y$ and let $Y_{\delta, q}^n \subset Y_\delta^n$ be the set of points, y, with $q(y) = q$.

Let q_0 be the largest value of q for which $Y_{\delta, q}^n$ is nonempty. Choose a maximal set of points from $Y_{\delta; q_0}^M$ such that

$$(5.23) \qquad \overline{\mathscr{O}_{y_\alpha}, \mathscr{O}_{y_\beta}} \geq \tfrac{1}{2} \min(r_{y_\alpha}, r_{y_\beta}),$$

where r_{y_α} and r_{y_β} are as in (c) above. Then choose a maximal set of points from $Y_{\delta, q-1}^n$ such that (5.24) continues to hold for all points (in

$Y^n_{\delta,q_0} \cup Y^n_{\delta,q_0-1})$ selected so far. By proceeding in this way, we obtain a set of points $\{y_\alpha\}$. Clearly, for every point $y \in Y^n_\delta$ there exist y_α with $q(y_\alpha) \geq q(y)$ and

$$(5.24) \qquad \overline{\mathscr{O}_y, \mathscr{O}_{y_\alpha}} \leq \tfrac{1}{2}\min(r_y, r_{y_\alpha}).$$

Since $q(y_\alpha) \geq q(y)$, it is clear that for δ sufficiently small, say

$$(5.25) \qquad \overline{y, \mathscr{O}_{y_\alpha}} \leq \tfrac{2}{3}\min(r_y, r_{y_\alpha})$$

(and the same holds for all points of $\mathscr{O}_y$). Thus, $\{V_{y_\alpha}\}$ covers Y^n_δ and in fact, $\{f^{-1}_{y_\alpha}(T_{3r_\alpha/4}(\mathscr{O}_{f_{y_\alpha}(y_\alpha)}))\}$ still covers.

Now, by using the standard packing argument as in [13, Theorem 5.3], the multiplicity of $\{V_{y_\alpha}\}$ can now be bounded by some $N_1(n)$.

(g) **Fitting together local F-structures.** The collection $\{(V_{y_\alpha}, \mathscr{F}_{y_\alpha})\}$ constructed in (f) above satisfies the hypothesis of Lemma 1.5. Thus, we obtain a weak F-structure, $\mathscr{F}$, on a set containing Y^n_δ, for $\delta < \delta_0(n)$ sufficiently small. Since the elementary F-structures, $\mathscr{F}_{y_\alpha}$, have positive rank, so does $\mathscr{F}$. The bound on the diameter of orbits (see (1) of Theorem 0.1) is also satisfied.

To see that the structure we have constructed is actually an F-structure, we observe that property (F_3) of §1 holds if t of (5.22) is sufficiently small. Note that the maximal length of a chain $V_{\alpha_0}, \cdots, V_{\alpha_l}$, as in (F_3) is, of course, bounded by $N(n)$, the maximal dimension of the fiber. Now it is clear from Corollary 2.13 that if t in (5.22) is taken to be $1/4N(n)$ times the value dictated by our previous considerations, then (F_3) and (F_4) hold.

As mentioned in §2, the local actions might be noneffective for the structure just constructed, but this can be remedied by passing to a quotient.

References

[1] P. Buser & H. Karcher, *Gromov's almost flat manifolds*, Asterisque, No. 81, Soc. Math. France, 1981.

[2] J. Cheeger, *Finiteness theorems for Riemannian manifolds*, Amer. J. Math. **92** (1970), 61–74.

[3] J. Cheeger & D. Gromoll, *On the structure of complete manifolds of nonnegative curvature*, Ann. of Math. (2) **96** (1972) 413–443.

[4] J. Cheeger & M. Gromov, *Collapsing riemannian manifolds while keeping their curvature bounded*, I, J. Differential Geometry **23** (1986), 309–346.

[5] J. Cheeger, M. Gromov, & M. Taylor, *Finite propagation speed, kernel estimates for functions of the Laplace operator, and the geometry of complete riemannian manifolds*, J. Differential Geometry **17** (1982) 15–53.

[6] S. Y. Cheng, P. Li, & S. T. Yau, *On the upper estimate of the heat kernel of a complete riemannian manifold*, Amer. J. Math. **103** (1981) 1021–1063.

[7] K. Fukaya, *On a compactification of the set of riemannian manifolds with bounded curvatures and diameters*, Lecture Notes in Math., Vol. 1201, Springer, 1986, 89–107.

[8] ____, *Collapsing riemannian manifolds to ones of lower dimension*, J. Differential Geometry **25** (1987) 139–156.

[9] ____, *A boundary of the set of the riemannian manifolds with bounded curvatures and diameters*, J. Differential Geometry **28** (1988) 1–21.

[10] ____, *Collapsing riemannian manifolds to ones of lower dimension. II* (preprint).

[11] R. Greene & H. Wu, *Lipschitz convergence of riemannian manifolds*, Pacific J. Math. **131** (1988) 119–141.

[12] M. Gromov, *Almost flat manifolds*, J. Differential Geometry **13** (1978) 231–241.

[13] ____, *Structures métriques our les variétés Riem. Red.*, par J. Lafontaine et. P. Pansu, Paris, 1981.

[14] K. Grove & S. Karcher, *How to conjugate C^1-close group actions*, Math. Z. **132** (1973) 11–20.

[15] K. Grove, H. Karcher & E. Ruh, *Group actions and curvature*, Invent. Math. **23** (1974) 31–48.

[16] P. Pansu, *Effondrement des variétés riemanniennes*, d'apres J. Cheeger et M. Gromov, Asterique, No. 121, Soc. Math. France, 1985.

[17] S. Peters, *Convergence of riemannian manifolds*, Compositio Math. **62** (1987) 3–16.

[18] W. Thurston, *The geometry and topology of 3-manifolds* (preprint).

[19] J. Wolf, *Spaces of constant curvature*, McGraw-Hill, New York, 1967.

STATE UNIVERSITY OF NEW YORK, STONY BROOK

INSTITUT DES HAUTES ETUDES SCIENTIFIQUES, FRANCE