Large Scale Detection of Half-flats in CAT(0) *Spaces*

STEFANO FRANCAVIGLIA & JEAN-FRANÇOIS LAFONT

ABSTRACT. Let M be a complete locally compact CAT(0)-space, and X an asymptotic cone of M. For $\gamma \subset M$ a k-dimensional flat, let γ_ω be the k-dimensional flat in X obtained as the ultralimit of γ. In this paper, we identify various conditions on γ_ω that are sufficient to ensure that γ bounds a $(k+1)$-dimensional half-flat.

As applications we obtain:

(1) constraints on the behavior of quasi-isometries between locally compact CAT(0)-spaces;
(2) constraints on the possible non-positively curved Riemannian metrics supported by certain manifolds;
(3) a correspondence between metric splittings of a complete, simply connected non-positively curved Riemannian manifolds, and metric splittings of its asymptotic cones; and
(4) an elementary derivation of Gromov's rigidity theorem from the combination of the Ballmann, Burns-Spatzier rank rigidity theorem and the classic Mostow rigidity theorem.

1. INTRODUCTION

A *k-flat* in a CAT(0)-space X is defined to be an isometrically embedded copy of the standard $\mathbb{R}^k$, $k \geq 1$. In the case where $k = 1$, a k-flat is just a geodesic in X. By a k-dimensional *half-flat*, $k \geq 1$, in a CAT(0) space, we mean an isometric copy of $\mathbb{R}^{k-1} \times \mathbb{R}^+$ (where $\mathbb{R}^+ = [0, \infty)$ is the usual half line). For example, when $k = 1$, a half-flat is just a geodesic ray in X. In the study of CAT(0)-spaces, a key role is played by the presence of flats and half-flats of *higher rank*, i.e., satisfying $k \geq 2$. In the present paper, our goal is to identify some coarse geometric conditions which are sufficient to ensure the existence of half-flats in a CAT(0)-space X. We provide

Indiana University Mathematics Journal ©, Vol. 59, No. 2 (2010)

three results towards this goal, as well as an example showing that our results are close to optimal.

Before stating our main results, let us recall that an *asymptotic cone* of a metric space X is a new metric space, which encodes the large-scale geometry of X, when viewed at an increasing sequence of scales. A precise definition, along with some basic properties of asymptotic cones, is provided in our Section 2. For a k-flat γ_ω inside the asymptotic cone X of a CAT(0)-space, we introduce (see Section 3) the notion of a *flattening sequence* of maps for γ_ω. These are a sequence of maps from a k-disk $\mathbb{D}^k$ into X, whose images are getting further and further away from γ_ω, and whose projections onto γ_ω satisfy certain technical conditions. The main point of such flattening sequences is that we can prove the following result:

Theorem A (Flattening sequence ⇒ half flat)**.** *Let M be a locally compact* CAT(0)*-space and let X be an asymptotic cone of M. Let γ be a k-flat of M (possibly a geodesic) and let γ_ω be its ultralimit in X. Suppose that there exists a flattening sequence of maps for γ_ω. Then the original k-flat γ bounds a $(k+1)$-half-flat in M.*

The reader will readily see that, in the special case where γ_ω itself bounds a half-flat, it is very easy to construct a flattening sequence. So an immediate consequence of Theorem A is the following result:

Theorem B (Half-ultraflat ⇒ half-flat)**.** *Let M be a locally compact* CAT(0)*-space and let X be an asymptotic cone of M. Let γ be a k-flat of M (possibly a geodesic) and let γ_ω be its ultralimit in X. If γ_ω bounds a half-flat, then γ itself must bound a half-flat.*

We also provide an example showing that, in the context of locally compact CAT(0)-spaces, the analogue of Theorem B with "half-flats" replaced by "flats" is *false*.

In Section 4, we weaken the hypothesis of Theorem B, by replacing a $(k+1)$-flat in the ultralimit by a *bi-Lipschitzly embedded* $(k+1)$-flat. We compensate for this by requiring the original flat to satisfy some mild periodicity requirement, and establish the following result:

Theorem C (Bilipschitz half-ultraflat + periodicity ⇒ half-flat)**.** *Let M be a locally compact* CAT(0)*-space and let X be an asymptotic cone of M. Let γ be a k-flat of M (possibly $k = 1$) and let γ_ω its ultralimit in X. Suppose that there exists $G <$ Isom(M) that acts co-compactly on γ.*

If there is a bi-Lipschitz embedding $\Phi : \mathbb{R}^k \times \mathbb{R}^+ \to X$, whose restriction to $\mathbb{R}^k \times \{0\}$ is a homeomorphism onto γ_ω, then γ bounds a $(k+1)$-half-flat in M.

In Section 5, we provide various geometrical applications of our main results. These include:

- Constraints on the possible quasi-isometries between certain locally compact CAT(0)-spaces.
- Restrictions on the possible locally CAT(0)-metrics supported by certain non-positively curved Riemannian manifolds.

- A proof that splittings of simply connected non-positively curved Riemannian manifolds correspond exactly with metric splittings of their asymptotic cones.
- An elementary argument deducing Gromov's rigidity theorem [3] (a closed higher rank locally symmetric space supports a unique metric of non-positive curvature up to homothety) from a combination of the Ballman, Burns-Spatzier rank rigidity theorem and the classic Mostow rigidity theorem.

In January 2008, the authors posted a preliminary version [14] of this work on the arXiv, which contained special cases of Thereoms A, B, C, under the additional hypothesis that the flats be 2-dimensional, and the ambient space M was a Riemannian manifold of non-positive sectional curvature (rather than a CAT(0)-space). Shortly thereafter, Misha Kapovich was kind enough to inform the authors of his paper with B. Leeb [16], in which (amongst other things) they proved a version of Theorem C in the special case of 2-dimensional flats, and where the ambient space M was an arbitrary locally compact CAT(0)-space. Their paper provided the motivation for us to write the present paper, which includes a generalization to higher dimensional flats of the result in [16, Proposition 3.3].

2. Background Material on Asymptotic Cones

In this section, we provide some background on ultralimits and asymptotic cones of metric spaces. Let us start with some basic reminders on ultrafilters.

Definition. A *non-principal ultrafilter* on the natural numbers $\mathbb{N}$ is a collection $\mathcal{U}$ of subsets of $\mathbb{N}$, satisfying the following four axioms:

(1) If $S \in \mathcal{U}$, and $S' \supset S$, then $S' \in \mathcal{U}$.
(2) If $S \subset \mathbb{N}$ is a finite subset, then $S \notin \mathcal{U}$.
(3) If $S, S' \in \mathcal{U}$, then $S \cap S' \in \mathcal{U}$.
(4) Given any finite partition $\mathbb{N} = S_1 \cup \cdots \cup S_k$ into pairwise disjoint sets, there is a unique S_i satisfying $S_i \in \mathcal{U}$.

Zorn's Lemma guarantees the existence of non-principal ultrafilters. Now given a compact Hausdorff space Z and a map $f : \mathbb{N} \to Z$, there is a unique point $f_\omega \in Z$ such that every neighborhood U of f_ω satisfies $f^{-1}(U) \in \mathcal{U}$. This point is called the ω-limit of the sequence $\{f(i)\}$; we will occasionally write $\omega \lim\{f(i)\} := f_\omega$. In particular, if the target space Z is the compact space $[0, \infty]$, we have that f_ω is a well-defined real number (or ∞).

Definition. Let $(M, d, *)$ be a pointed metric space, $M^{\mathbb{N}}$ the collection of M-valued sequences, and $\lambda : \mathbb{N} \to (0, \infty) \subset [0, \infty]$ a sequence of real numbers satisfying $\lambda_\omega = \infty$. Given any pair of points $\{x_i\}$, $\{y_i\}$ in $X^{\mathbb{N}}$, we define the pseudo-distance $d_\omega(\{x_i\}, \{y_i\})$ between them to be f_ω, where $f : \mathbb{N} \to [0, \infty)$ is the function $f(k) = d(x_k, y_k)/\lambda(k)$. Observe that this pseudo-distance takes on values in $[0, \infty]$.

Next, note that $M^{\mathbb{N}}$ has a distinguished point, corresponding to the constant sequence $\{*\}$. Restricting to the subset of $M^{\mathbb{N}}$ consisting of sequences $\{x_i\}$ satisfying $d_\omega(\{x_i\},\{*\}) < \infty$, and identifying sequences whose d_ω distance is zero, one obtains a genuine pointed metric space $(M_\omega, d_\omega, *_\omega)$, which is called an *asymptotic cone* of the pointed metric space $(M, d, *)$.

We will usually denote an asymptotic cone by $\text{Cone}(M)$. The reader should keep in mind that the construction of $\text{Cone}(M)$ involves a number of choices (base-points, sequence λ_i, choice of non-principal ultrafilters) and that different choices could give different (non-homeomorphic) asymptotic cones (see the papers [26], [18], [22]). However, in the special case where $M = \mathbb{R}^k$, *all* asymptotic cones are isometric to $\mathbb{R}^k$ (i.e., we have independence of all choices).

We will require the following basic facts concerning asymptotic cones of non-positively curved spaces (see for instance [17, Propositions 2.4.4, 2.4.6]):

- If (M, d) is a CAT(0)-space, then $\text{Cone}(M)$ is likewise a CAT(0)-space.
- If $\varphi : M \to N$ is a (C, K)-quasi-isometric map, then φ induces a C-bi-Lipschitz map $\varphi_\omega : \text{Cone}(M) \to \text{Cone}(N)$.
- If $\gamma \subset M$ is a k-flat, then $\gamma_\omega := \text{Cone}(\gamma) \subset \text{Cone}(M)$ is likewise a k-flat.
- If $\{a_i\}$, $\{b_i\} \in \text{Cone}(M)$ are an arbitrary pair of points, then the ultralimit of the geodesic segments $\overline{a_i b_i}$ gives a geodesic segment $\overline{\{a_i\}\{b_i\}}$ joining $\{a_i\}$ to $\{b_i\}$.

Concerning the second point above, we remind the reader that a (C, K)-quasi-isometric map $\varphi : (M, d_M) \to (N, d_N)$ between metric spaces is a (not necessarily continuous) map having the property that:

$$\frac{1}{C} \cdot d_M(p, q) - K \le d_N(\varphi(p), \varphi(q)) \le C \cdot d_M(p, q) + K.$$

We also comment that, in the second point above, the asymptotic cones of M, N, have to be taken with the same scaling sequence and the same ultrafilters.

Lemma 2.1 (Translations on asymptotic cone). *Let M be a geodesic space, $\gamma \subset M$ a k-flat, and $\gamma_\omega \subset X$ the corresponding k-flat in an asymptotic cone $X :=$* $\text{Cone}(M)$ *of M. Assume that there exists a subgroup $G <$* $\text{Isom}(M)$ *with the property that G leaves γ invariant, and acts cocompactly on γ. Then for any pair of points p, $q \in \gamma_\omega \subset X$, there is an isometry $\Phi : X \to X$ satisfying $\Phi(p) = q$.*

Proof. Let $\{p_i\}$, $\{q_i\} \subset \gamma \subset M$ be sequences defining the points p, q respectively. Since G leaves γ invariant, and acts cocompactly on γ, there exists elements $g_i \in G$ with the property that for every index i, we have $d(g_i(p_i), q_i) \le R$.

Now observe that the sequence $\{g_i\}$ of isometries of M defines a self-map (defined componentwise) of the space $M^{\mathbb{N}}$ of sequences of points in M. Let us denote by g_ω this self-map, which we now proceed to show induces the desired isometry on $X = \text{Cone}(M)$. First note that it is immediate that g_ω preserves the pseudo-distance d_ω on $M^{\mathbb{N}}$, and has the property that $d_\omega(\{g_i(p_i)\}, \{q_i\}) = 0$. So to see that g_ω descends to an isometry of X, all we have to establish is that for $\{x_i\}$ a sequence satisfying $d_\omega(\{x_i\}, *) < \infty$, the image sequence also satisfies $d_\omega(\{g_i(x_i)\}, *) < \infty$. But we have the series of equivalences:

$$\begin{aligned} d_\omega(\{x_i\}, *) < \infty &\iff d_\omega(\{x_i\}, \{p_i\}) < \infty\,, \\ &\iff d_\omega(\{g_i(x_i)\}, \{g_i(p_i)\}) < \infty\,, \\ &\iff d_\omega(\{g_i(x_i)\}, \{q_i\}) < \infty\,, \\ &\iff d_\omega(\{g_i(x_i)\}, *) < \infty\,, \end{aligned}$$

where the first and last equivalences come from applying the triangle inequality in the pseudo-metric space $(M^{\mathbb{N}}, d_\omega)$, and the second and third equivalences follow from our earlier comments. We conclude that the induced isometry g_ω on the pseudo-metric space $M^{\mathbb{N}}$ of sequences leaves invariant the subset of sequences at finite distance from the distinguished constant sequence, and hence descends to an isometry of X. Finally, it is immediate from the definition of the isometry g_ω that it will leave γ_ω invariant, as each g_i leaves γ invariant. This concludes the proof. □

Let us now specialize the previous lemma to the case of geodesics (i.e., $k = 1$). Observe that any element $g \in \text{Isom}(M)$ as in the previous lemma gives rise to a $\mathbb{Z}$-action on M leaving γ invariant. It is worth pointing out that the lemma does *not* state that the $\mathbb{Z}$-action on M induces an $\mathbb{R}$-action on $X = \text{Cone}(M)$. The issue is that for each $r \in \mathbb{R}$, there is indeed a corresponding isometry of X, but these will not, in general, vary continuously with respect to r.

A simple example illustrating this phenomena is provided by $M = \mathbb{H}^2$, and g a hyperbolic translation along a geodesic γ. The asymptotic cone X is homeomorphic to an $\mathbb{R}$-tree with uncountable branching at every point, and γ_ω is a fixed geodesic within this $\mathbb{R}$-tree. Lemma 2.1 ensures that, for each $r \in \mathbb{R}$, we have a map $\Phi_r \in \text{Isom}(X)$ which acts on $\gamma_\omega \subset X$ via a translation by the real number r. To see that the map $\mathbb{R} \to \text{Isom}(X)$ defined via $r \mapsto \Phi_r$ is *not* continuous, we can consider the orbit of any point $x \in X$ which does not lie on the geodesic γ_ω. The point x lies at distance $\varepsilon > 0$ from the geodesic γ_ω, and has a unique closest point $y \in \gamma_\omega$. Now for any $r \in \mathbb{R} \setminus \{0\}$, Φ_r is an isometry leaving γ_ω invariant, hence $\Phi_r(x)$ also lies at distance ε from γ_ω, and has closest point $\Phi_r(y) \in \gamma_\omega$. We can now easily compute the distance between x and $\Phi_r(x)$:

$$d(x, \Phi_r(x)) = d(x, y) + d(y, \Phi_r(y)) + d(\Phi_r(y), \Phi_r(x)) = r + 2\varepsilon.$$

Since this distance is uniformly bounded away from zero, the image $\Phi_r(x)$ of x does not vary continuously with respect to the parameter r, and hence the maps constructed in Lemma 2.1 do *not* define an $\mathbb{R}$-action.

3. FLATTENING SEQUENCES AND HALF-FLATS

In this section, we will provide a proof of Theorem A. Our goal is to show how certain sequences of maps from the disk to the asymptotic cone of a CAT(0)-space M can be used to construct flats in M. We recall that ω denotes the ultrafilter used to construct $X = \mathrm{Cone}(M)$, that λ_j denotes the sequence of scaling factors, and that $*$ denotes both the base-point of M, and the base-point of X represented by the constant sequence $\{*\}$.

We are given a k-flat $\gamma \subset M$ (possibly a geodesic), and we have the corresponding k-flat $\gamma_\omega \subset X$ in the asymptotic cone X. By abuse of notation, we will use π to denote both the nearest point projection $\pi : M \to \gamma$, as well as the nearest point projection $\pi : X \to \gamma_\omega$ (these are well defined since both M and X are CAT(0) spaces). We can now make the following definition:

Definition (Flattening sequences)**.** We say that γ_ω has a *flattening sequence* provided there exists a sequence of continuous maps $(f_r)_{r\in\mathbb{N}}$ from the k-disk $\mathbb{D}^k$ to X such that

(1) The diameter of $f_r(\mathbb{D}^k)$ is smaller than one.
(2) The image of $\pi \circ f_r$ contains an open neighborhood of the base-point $*$.
(3) The restriction of $\pi \circ f_r$ to $\partial\mathbb{D}^k = S^{k-1}$ does not contain $*$, and represents a non-zero element in the homotopy group $\pi_{k-1}(\gamma_\omega \setminus \{*\}) \cong \mathbb{Z}$.
(4) There exists a constant $D > 0$ with the property that

$$d(f_r(\mathbb{D}^k), \gamma_\omega) = \inf_{x\in\mathbb{D}^k} d(f_r(x), \gamma_\omega) \geq D \cdot r.$$

The sequence of maps $(f_r)_{r\in\mathbb{N}}$ will be called a *flattening sequence* for γ_ω.

We now assume that the k-flat γ_ω has a flattening sequence consisting of maps $f_r : \mathbb{D}^k \to X$. To establish Theorem A, we need to prove that γ bounds a $(k+1)$-dimensional half-flat. In order to do this, we have to construct geodesic rays emanating from various points on γ.

Each such ray will be constructed as a limit of a sequence of longer and longer geodesic segments, originating from a fixed point on γ, and terminating at a sequence of suitably chosen points in the space M. In order to select this suitable sequence of points in M, we start by noting that each of the maps $f_r : \mathbb{D}^k \to X$ in our flattening sequence can be obtained as an ultralimit of a sequence of maps $f_{r,j} : \mathbb{D}^k \to M$ (see for instance Kapovich [15]). The desired collection of points will be carefully chosen to lie on the image of some of the maps $f_{r,j}$. The precise selection process is contained in the following statement:

Assertion 3.1. *Let us be given an arbitrary finite $(m+1)$-tuple of points $\{P^0, \ldots, P^m\} \subset \gamma$. Then for each $r \in \mathbb{N}$, we can choose indices $j_r \in \mathbb{N}$, and $(m+1)$-tuples of points $\{x^0_{r,j_r}, \ldots, x^m_{r,j_r}\} \subset f_{r,j_r}(\mathrm{Int}(\mathbb{D}^k)) \subset M$ with the property that:*

(1) *For each r, and $0 \leq i \leq m$, we have that $\pi(x^i_{r,j_r}) = P^i$.*

(2) *For any* i, i', $$\frac{d(x^i_{r,j_r}, x^{i'}_{r,j_r})}{\lambda_{j_r}} < 2.$$

(3) *For any* i, $$\frac{d(P^i, x^i_{r,j_r})}{r \cdot \lambda_{j_r}} > \frac{D}{2}.$$

We temporarily delay the proof of Assertion 3.1, and focus on explaining how our Theorem A can be deduced from this statement. We first recall some terminology: we say that two half-rays η_1 and η_2 *bound a flat strip*, provided there exists an isometric embedding of $\mathbb{R}^+ \times [0, a]$ (for some $a > 0$) into M, with the property that η_1 coincides with $\mathbb{R}^+ \times \{0\}$ and η_2 coincides with $\mathbb{R}^+ \times \{a\}$. For a fixed ray η, we will denote by $\mathrm{Par}(\eta) \subset M$ the union of all geodesic rays in M which, together with η, jointly bound a flat strip. Our first step is to use Assertion 3.1 to show the following statement:

Claim 3.2. *Given any finite set of points* $\{P^0, \ldots, P^m\} \subset \gamma$, *there exist a collection of geodesic rays* $\{\eta^0, \ldots, \eta^m\}$, *satisfying:*

- *For each* i, *we have that* η^i *originates at* P^i, *and satisfies* $\pi(\eta^i) = P^i$.
- *Each pair of geodesic rays* η^i, $\eta^{i'}$ *jointly bound a flat strip.*

Proof. This can be seen as follows: first, apply Assertion 3.1 to the finite set of points, obtaining a sequence of $(m + 1)$-tuples of points $\{x^0_{r,j_r}, \ldots, x^m_{r,j_r}\} \subset M$. Now for each i, consider the sequence of geodesic segments η^i_r, which joins the point P^i to the point x^i_{r,j_r}. Since M is locally compact, we can extract a subsequence which simultaneously converges for all the $0 \le i \le m$. We define the limiting geodesic rays to be our η^i. So to complete the proof of the claim, we just need to verify that these η^i have the desired property. By construction, we know that each of the η^i_r are geodesic segments originating at P^i, which immediately gives us the corresponding property for η^i. Likewise, we have that each of the geodesic segments η^i_r project to the point P^i, which yields the same statement for η^i. Note that this implies that, for any $t > 0$, η^i is the minimal length path from $\eta^i(t)$ to γ. In particular, the angle (in the CAT(0) sense, see [6]) of η^i with γ must be $\ge \pi/2$.

So we now need to establish the second property: that each pair η^i and $\eta^{i'}$ jointly bound a flat strip. Let us set $d = d(P^i, P^{i'})$; we start by showing that for $t \ge 0$, we have $d(\eta^i(t), \eta^{i'}(t)) = d$. Since the geodesic rays are limits of the geodesic segments, this is equivalent to showing $\lim_{r\to\infty} d(\eta^i_r(t), \eta^{i'}_r(t)) = d$. From condition (1) in Assertion 3.1, and the fact that the projection map is distance non-increasing, we obtain the inequality:

$$\lim_{r\to\infty} d(\eta^i_r(t), \eta^{i'}_r(t)) \ge \lim_{r\to\infty} d(\eta^i_r(0), \eta^{i'}_r(0)) = d(P^i, P^{i'}) = d.$$

For the reverse inequality, we need to estimate from above the distance between $\eta^i_r(t)$ and $\eta^{i'}_r(t)$. To do this, we first truncate the longer of the two geodesic segments η_r^i and $\eta_r^{i'}$ to have length equal to the smaller one. We will denote by $\bar{\eta}_r^i$, $\bar{\eta}_r^{i'}$ the new pair of equal length geodesics, and let L denote their common length. Since we are trying to estimate from above the distance between the points $\bar{\eta}_r^i(t)$ and $\bar{\eta}_r^{i'}(t)$, convexity of the distance function in CAT(0)-spaces yields:

$$d(\bar{\eta}_r^i(t), \bar{\eta}_r^{i'}(t)) \leq \left(1 - \frac{t}{L}\right) \cdot d(\bar{\eta}_r^i(0), \bar{\eta}_r^{i'}(0)) + \frac{t}{L} d(\bar{\eta}_r^i(L), \bar{\eta}_r^{i'}(L)).$$

But from condition (3) in Assertion 3.1, we have that $L > r \cdot \lambda_{j_r} \cdot D/2$. Furthermore, from condition (2) in Assertion 3.1, and an application of the triangle inequality, we obtain that $d(\bar{\eta}_r^i(L), \bar{\eta}_r^{i'}(L)) < 4\lambda_{j_r} + d$. Substituting these expressions (along with $d(\bar{\eta}_r^i(0), \bar{\eta}_r^{i'}(0)) = d$) into our inequality, and simplifying, we get:

$$d(\bar{\eta}_r^i(t), \bar{\eta}_r^{i'}(t)) \leq d + 4t \cdot \frac{\lambda_{j_r}}{L} < d + 4t \cdot \frac{\lambda_{j_r}}{r\lambda_{j_r} \cdot D/2} = d + \frac{8t}{D \cdot r}$$

where we recall that d, t, D are constants. Now taking the limit as $r \to \infty$, we obtain that $\lim_{r\to\infty} d(\bar{\eta}_r^i(t), \bar{\eta}_r^{i'}(t)) \leq d$, as desired.

This verifies that the two geodesic rays η^i, $\eta^{i'}$ remain at a constant distance apart, in the sense that $d(\eta^i(t), \eta^{i'}(t))$ is a constant function of t. Furthermore, from our earlier discussion, we have that the geodesic segment joining $P^i = \eta^i(0)$ to $P^{i'} = \eta^{i'}(0)$ forms an angle $\geq \pi/2$ with each of the geodesic rays η^i, $\eta^{i'}$, and hence both these angles must actually be $= \pi/2$. But in a CAT(0) space, this forces the geodesics η^i and $\eta^{i'}$ to bound a flat strip (see the proof of the flat strip theorem [6, p. 182]). This concludes the proof of Claim 3.2. □

So we now know that any finite set of points $\{P^0, \dots, P^m\} \subset \gamma$ are common endpoints of geodesic rays that pairwise bound a flat strip. But ultimately, we want to show that *every* point on γ is an endpoint of a parallel geodesic ray. Our next step is to establish the following statement:

Claim 3.3. *Given any compact set $K \subset \gamma$, we can find an isometric embedding of $K \times [0, \infty) \hookrightarrow M$ with the property that $K \times \{0\}$ maps to K.*

Proof. Recall that, for a geodesic η, the set Par(η) is the union of all geodesic rays which, together with η, bound a flat strip. From the proof of the product region theorem in CAT(0)-spaces (see for example [6]), one has that Par(η) forms a convex subset of M, which splits as a metric product $B \times [0, \infty)$. Here $B \subset M$ is a convex subset, and consists of the collection of all the base-points of the parallel geodesic rays.

Since $K \subset \gamma$ is compact, we can find a finite set of points $\{P^0, \dots, P^m\} \subset \gamma$ whose convex hull contains K. Applying Claim 3.2, we have that there exist a corresponding collection of geodesic rays $\{\eta^0, \dots, \eta^m\}$, with each η^i originating from P^i, and which pairwise bound a flat strip. Considering the set Par(η^0), we have that $\mathrm{Par}(\eta^0)$ is isometric to $B \times [0, \infty)$, where B is the collection of base-points. But we know that $\{P^0, \dots, P^m\} \subset B$, so their convex hull is likewise contained in B, forcing $K \subset B$. We conclude that there is an isometric copy of $K \times [0, \infty)$ embedded inside the convex subset $\mathrm{Par}(\eta) \subset M$. This completes the proof of Claim 3.3. ◻

Finally, let us take a sequence of compact sets $K_i \subset \gamma$ exhausting the flat γ (for instance, take K_i to be the radius i metric ball centered at $*$). From Claim 3.3, we have a corresponding sequence of isometrically embedded copies of $K_i \times [0, \infty) \hookrightarrow M$, where each $K_i \times \{0\}$ maps to the corresponding compact K_i. From local compactness, we can extract a convergent subsequence, whose limit will be the desired half-flat bounding γ. So to complete the proof of Theorem A, we are left with verifying Assertion 3.1.

Proof of Assertion 3.1. Let us recall the framework: we have a finite collection of points $\{P^0, \dots, P^m\} \subset \gamma$, and for each $r \in \mathbb{N}$, a collection of maps $f_{r,j} : \mathbb{D}^k \to M$ having the property that $\omega \lim f_{r,j} = f_r : \mathbb{D}^k \to X$. We want to find, for each index r, a corresponding index j_r and set of points $\{x^0_{r,j_r}, \dots, x^m_{r,j_r}\} \subset f_{r,j_r}(\mathrm{Int}(\mathbb{D}^k))$. The chosen set of points should have the property that

(1) for each i, $\pi(x^i_{r,j_r}) = P^i$;
(2) for any i, i', $d(x^i_{r,j_r}, x^{i'}_{r,j_r}) < 2\lambda_{j_r}$;
(3) for any i, $d(P^i, x^i_{r,j_r}) > r \cdot \lambda_{j_r} \cdot D/2$.

Our approach is as follows: fixing r, we define three subsets of $\mathbb{N}$ by setting

- J_1 to be the set of indices j for which $\{P^0, \dots, P^m\} \subset \pi \circ f_{r,j}(\mathrm{Int}(\mathbb{D}^k))$;
- J_2 to be the set of indices where $\mathrm{diam}(f_{r,j}(\mathbb{D}^k)) < 2\lambda_j$;
- J_3 to be the set of indices where $d(\gamma, f_{r,j}(\mathbb{D}^k)) > r \cdot \lambda_j \cdot D/2$.

Now assuming we could show that each of these three sets are in ω, property (3) of ultrafilters (closure under finite intersections) implies that $J_1 \cap J_2 \cap J_3 \in \omega$. Finally, every set in ω is infinite (property (2) of ultrafilters), and in particular non-empty, allowing us to find an index $j_r \in J_1 \cap J_2 \cap J_3$. For this index j_r, we can choose arbitrary points $x^i_{r,j_r} \in f_{r,j_r}(\mathrm{Int}(\mathbb{D}^k))$ satisfying $\pi(x^i_{r,j_r}) = P^i$ (such points exist since $j_r \in J_1$). And from $j_r \in J_2 \cap J_3$, it immediately follows that the tuple of points $\{x^0_{r,j_r}, \dots, x^m_{r,j_r}\} \subset f_{r,j_r}(\mathrm{Int}(\mathbb{D}^k))$ has the desired properties.

Step 1: The set J_2 lies in the ultrafilter ω. To see this, we recall that property (1) in the definition of a flattening sequence requires $\mathrm{diam}(f_r(\mathbb{D}^k)) \leq 1$. Since we know that $f_r = \omega \lim f_{r,j}$, the definition of distances in the ultralimit tells us that the set of indices j for which $\mathrm{diam}(f_{r,j}(\mathbb{D}^k))/\lambda_j < 2$ lies in the ultrafilter. This verifies Step 1.

Step 2: The set J_3 lies in the ultrafilter ω. For this, we argue similarly. Recall that property (4) in the definition of a flattening sequence requires the existence of a $D > 0$ so that $d(f_r(\mathbb{D}^k), \gamma_\omega) = \inf_{x\in\mathbb{D}^k} d(f_r(x), \gamma_\omega) \geq D \cdot r$. Since f_r is the ultralimit of the maps $f_{r,j}$, the definition of distances in the ultralimit tells us that the set of indices j for which $d(f_{r,j}(\mathbb{D}^k), \gamma)/\lambda_j \geq rD/2$ lies in the ultrafilter. This verifies Step 2.

Step 3: The set J_1 lies in the ultrafilter ω. This last step is much more involved than the first two. Let us fix one of the points P^i, and consider the restriction $f_{r,j}|_{\partial D^k} : S^{k-1} \to M$, composed with the projection $\pi : M \to \gamma$. We have three distinct possibilities:

(1) P^i lies in the image of $\pi \circ f_{r,j}|_{\partial D^k}$, or
(2) $\pi(f_{r,j}(\partial D^k)) \subset \gamma \setminus \{P^i\}$, and $[\pi \circ f_{r,j}|_{\partial D^k}] = 0$ in $\pi_{k-1}(\gamma \setminus \{P^i\})$, or
(3) $\pi(f_{r,j}(\partial D^k)) \subset \gamma \setminus \{P^i\}$, and $[\pi \circ f_{r,j}|_{\partial D^k}] \neq 0$ in $\pi_{k-1}(\gamma \setminus \{P^i\})$.

This gives us a partition $\mathbb{N} = I_1^i \cup I_2^i \cup I_3^i$ into three disjoint sets, according to which of these three properties holds for the index j. From property (4) of ultrafilters, we have that exactly one of these three sets must lie in ω. If we could show that $I_3^i \in \omega$, then property (3) of ultrafilters would force $I_3^0 \cap \cdots \cap I_3^k \in \omega$. Since we have a containment $I_3^0 \cap \cdots \cap I_3^k \subset J_1$, property (1) of ultrafilters would give us that $J_1 \in \omega$.

So to conclude the proof of Step 3 (and hence, of Assertion 3.1), we are left with showing that, for each fixed choice of i, $I_1^i \notin \omega$ and $I_2^i \notin \omega$. We will argue both of these by contradiction. We also remind the reader that, from this point on in the proof, i should be considered fixed.

Supposing that $I_1^i \in \omega$, we would have that the set of indices for which $d(P^i, \pi \circ f_{r,j}|_{\partial D^k})/\lambda_j = 0$ is contained in ω. So the point represented by the constant sequence $\{P^i\} = \{*\} \in \gamma_\omega$ lies on the set:

$$\omega \lim (\pi \circ f_{r,j}|_{\partial D^k}) = \pi(\omega \lim f_{r,j}|_{\partial D^k}) = \pi \circ f_r|_{\partial D^k}.$$

But this contradicts property (3) in the definition of a flattening sequence.

Similarly, to see that $I_2^i \notin \omega$, we again argue by contradiction. So let us assume that $I_2^i \in \omega$. Note that the indices in I_2^i are those for which the map $\pi \circ f_{r,j}|_{\partial D^k} : S^{k-1} \to \gamma \setminus \{P^i\}$ is homotopically trivial. Let us define real numbers

$$a_j = \inf_{x\in\partial D^k} d(P^i, \pi \circ f_{r,j}(x)),$$
$$b_j = \sup_{x\in\partial D^k} d(P^i, \pi \circ f_{r,j}(x)),$$

and consider the set $A_j := \{x \in \gamma \mid a_j \leq d(x, P_i) \leq b_j\}$. The inclusion $A_j \hookrightarrow \gamma \setminus \{P^i\}$ is a homotopy equivalence, and the map $\pi \circ f_{r,j}|_{\partial D^k}$ has its image lying inside A_j. Hence for indices $j \in I_2^i$, we have that $[\pi \circ f_{r,j}|_{\partial D^k}] = 0 \in \pi_{k-1}(A_j)$, and we can construct a map $F_j : \mathbb{D}^k \to A_j \subset \gamma \setminus \{P^i\}$ with the property that $F_j|_{\partial \mathbb{D}^k} = \pi \circ f_{r,j}|_{\partial D^k}$.

Now for each $\varepsilon > 0$, we can further partition the set $I_2^i = I_2^i(\varepsilon) \cup \bar{I}_2^i(\varepsilon)$ by defining $I_2^i(\varepsilon)$ to be the set of indices where the inequality $a_j < \varepsilon\lambda_j$ holds, and $\bar{I}_2^i(\varepsilon)$ the set of indices where $a_j \geq \varepsilon\lambda_j$. From property (4) of ultrafilters, we have that precisely one of the sets $I_2^i(\varepsilon)$, $\bar{I}_2^i(\varepsilon)$ is contained in ω. We now have two possibilities:

- either we have that $I_2^i(\varepsilon) \in \omega$ for *every* $\varepsilon > 0$,
- or there exists some $\varepsilon > 0$, for which $\bar{I}_2^i(\varepsilon) \in \omega$.

In the first case, we again obtain that the point represented by the constant sequence $\{P^i\} = \{*\} \in \gamma_\omega$ lies on the set $\omega \lim(\pi \circ f_{r,j}|_{\partial D^k}) = \pi \circ f_r|_{\partial D^k}$, which contradicts property (3) in the definition of a flattening sequence.

In the second case, we can try to take the ultralimit of the collection of maps $F_j : \mathbb{D}^k \to \gamma$. The *upper bound* on the distance between F_j and P^i ensures that the F_j escape to infinity slowly enough for the ultralimit to be defined. More precisely, from the fact that the maps $\pi \circ f_{r,j}|_{\partial D^k}$ have as ultralimit $\pi \circ f_r|_{\partial \mathbb{D}^k}$, we must have that $\omega \lim\{b_j/\lambda_j\} < \infty$, which in turn implies that the ultralimit $F_\omega : \mathbb{D}^k \to \gamma_\omega$ exists. On the other hand, the *lower bound* on the distance between F_j and P^i ensures that the ultralimit F_ω does not pass through the constant sequence $\{P^i\}$. More precisely, for the $\varepsilon > 0$ satisfying $\bar{I}_2^i(\varepsilon) \in \omega$, and for any index set $j \in \bar{I}_2^i(\varepsilon)$, we have that

$$d(F_j(\mathbb{D}^k), P^i)/\lambda_j \geq d(A_j, P^i)/\lambda_j = a_j/\lambda_j \geq \varepsilon.$$

Since this holds for a set of indices in the ultrafilter, we immediately deduce the corresponding property for the ultralimit: $d(F_\omega(\mathbb{D}^k), \{P^i\}) \geq \varepsilon$. In particular, $F_\omega : \mathbb{D}^k \to \gamma_\omega$ has its image lying in the complement of the point corresponding to the constant sequence $\{P^i\} = \{*\}$, and restricts to the map $\pi \circ f_r|_{\partial \mathbb{D}^k}$ on the boundary $S^{k-1} = \partial \mathbb{D}^k$. This tells us that $[\pi \circ f_r|_{\partial \mathbb{D}^k}] = 0 \in \pi_{k-1}(\gamma_\omega \setminus \{*\})$, which contradicts property (3) in the definition of a flattening sequence. This concludes the verification of Step 3, and hence completes the proof of Assertion 3.1. ◻

Having established the Assertion 3.1, we have now concluded the proof of Theorem A. From the definition of a flattening sequence, it is obvious that these exist whenever γ_ω bounds a half-flat in X. As a result, we see that Theorem B follows immediately from Theorem A.

Finally, let us conclude this section by providing a family of cautionary examples. These will be locally compact CAT(0)-spaces X_k, each containing a geodesic γ, with the property that for a suitable choice of scales, γ_ω is contained inside a k-dimensional flat, but the individual X_k do not contain any flats of dimension

> 1. In particular, these examples show that the analogue of Theorem B with "half-flats" replaced by "flats" is *false*.

Example. Let us fix a $k \geq 2$, and for $n \in \mathbb{N}$, define the spaces $C_n := [-n^3, n^3]^k \subset \mathbb{R}^k$. Each C_n is isometric to the standard k-dimensional cube with side length $2n^3$; we let $\ell_n \subset C_n$ be the geodesic segment of length $2n^3$ joining the two points $(\pm n^3, 0, \dots, 0)$ inside C_n. Now consider the closed upper half space $\mathbb{R} \times \mathbb{R}_{\geq 0} := \{(x, y) \mid y \geq 0\}$ with the standard flat metric, and for $n \in \mathbb{N}$, let us denote by $\bar{\ell}_n$ the segment of length $2n^3$ joining the pair of points $(\pm n^3, n)$ inside $\mathbb{R} \times \mathbb{R}_{\geq 0}$. We now form the space X_k by gluing together all the C_n to $\mathbb{R} \times \mathbb{R}_{\geq 0}$. More precisely, we isometrically identify each ℓ_n with the corresponding $\bar{\ell}_n$. Observe that this space X_k, with the natural induced metric, is a locally compact CAT(0)-space. Furthermore, it is clear that X_k *does not contain any flats of dimension* > 1. Now consider the ultralimit X obtained by fixing the origin as the sequence of base-points, and setting $\lambda(i) = i^2$ to be the sequence of scales. Let us consider the geodesic $\gamma \subset X_k$ given by the x-axis in the $\mathbb{R} \times \mathbb{R}_{\geq 0}$ portion of M. We claim that the corresponding geodesic $\gamma_\omega \subset \omega \lim X_k$ is contained inside a k-flat in $\omega \lim X_k$. Indeed, this follows readily from the following two observations

- Since the distance from $\bar{\ell}_n$ to the γ grows linearly, while the scaling factor λ grows quadratically, every point $P \in \gamma_\omega$ can be represented by a sequence $\{p_n\}$ with the additional property that $p_n \in C_n$.
- Since the size of the cubes C_n grows cubically, while the scaling factor λ grows quadratically, the subset $C_\omega \subset \omega \lim X_k$ consisting of all points having a representative sequence of the form $\{c_i\}$ (with each $c_i \in C_i$) is isometric to the standard $\mathbb{R}^k$.

This concludes our family of locally CAT(0) examples.

4. FROM BI-LIPSCHITZ HALF-ULTRAFLATS TO HALF-FLATS

In this section we prove Theorem C, allowing us to deduce the presence of half-flats in M from the presence of bi-Lipschitz half-flats in the ultralimit X along with a mild periodicity condition.

The context is the following: we have a locally compact CAT(0)-space M (for instance, the universal cover of a non-positively curved Riemannian manifold) and an asymptotic cone X of M. We have a k-flat γ in M, and its limit γ_ω in X. Moreover, we are supposing that there exists $G < \text{Isom}(M)$ that acts co-compactly on γ. We are assuming that there is a bi-Lipschitz embedding $\varphi : \mathbb{R}^k \times \mathbb{R}_{\geq 0} \to X$, whose restriction to $\mathbb{R}^k \times \{0\}$ maps onto γ_ω, and we want to show that γ bounds a $(k+1)$-dimensional half-flat in M. In view of our Theorem A, it is sufficient to find a flattening sequence for γ_ω.

Let C be the bi-Lipschitz constant of φ, and for $r \in \mathbb{R}$, let us denote by $L_r = \mathbb{R}^k \times \{r\} \subset \mathbb{R}^k \times \mathbb{R}_{\geq 0}$ the horizontal flat at height r. We will use ρ to denote the obvious projection map $\rho : L_r \to L_0$. To make our various expressions more readable, we use d to denote the distance in X (as opposed to d_ω), and the norm notation to denote the distance inside $\mathbb{R}^k \times \mathbb{R}_{\geq 0}$.

We now define, for each $r \in [0, \infty)$, a map

$$\psi_r : L_r \to L_0$$

as follows: given $p \in L_r$, we have $\varphi(p) \in X$. Since $\gamma_\omega \subset X$ is a flat inside the CAT(0) space X, there is a well defined, distance non-increasing, projection map $\pi : X \to \gamma_\omega$, which sends any given point in X to the (unique) closest point on γ_ω. Hence, given $p \in L_r$, we have the composite map $\pi \circ \varphi : L_r \to \gamma_\omega$. But recall that, by hypothesis, φ maps L_0 homeomorphically to γ_ω. We can now set $\psi_r : L_r \to L_0$ to be the composite map

$$\psi_r = \varphi^{-1} \circ \pi \circ \varphi$$

We now show that ψ_r is at finite distance from the projection map $\rho : L_r \to L_0$.

We first observe that for arbitrary $x \in L_r$, the distance from x to L_0 is exactly r, and hence from the bi-Lipschitz estimate, we have

$$d(\varphi(x), \gamma_\omega) = d(\varphi(x), \varphi(L_0)) \leq Cr.$$

Since π is the nearest point projection onto γ_ω, this implies that

$$d(\varphi(x), (\pi \circ \varphi)(x)) \leq Cr.$$

Since $(\pi \circ \varphi)(x) = \varphi(\psi_r(x))$, we can again use the bi-Lipschitz estimate to conclude that:

$$Cr \geq d(\varphi(x), (\pi \circ \varphi)(x)) = d(\varphi(x), \varphi(\psi_r(x))) \geq \frac{1}{C} \cdot \|x - \psi_r(x)\|,$$

which gives us the estimate $\|x - \psi_r(x)\| \leq C^2 r$. This implies that ψ is at bounded distance from the projection map ρ. Since the latter is a homeomorphism onto $L_0 \cong \mathbb{R}^k$, it follows that ψ_r is surjective.

Since ψ_r is a surjective, Lipschitz map, its differential exists almost everywhere and it is almost everywhere non-degenerate (this follows easily from Rademacher's theorem, see for example [27, Chapter 3]). It follows that we can find a k-disk D_r in L_r, of diameter smaller than $1/C$, and a point p_r in $\psi_r(D_r)$ such that $\psi_r(\partial D_r)$ is homotopically non trivial in $L_0 \setminus p_r$. By Lemma 2.1, we can find an isometry g_r of X, that leaves γ_ω invariant, and satisfies $g_r(\varphi(p_r)) = *$.

Now we define the maps $f_r : \mathbb{D}^k \cong D_r \to X$ via the composition

$$f_r = g_r \circ \varphi|_{D_r},$$

and observe that we have $\pi \circ f_r = g_r \circ \varphi \circ \psi_r$. Moreover, since the diameter of D_r is smaller than $1/C$, the diameter of $f_r(D_r)$ is smaller than one. Finally, it is clear that our choices for D_r imply that f_r satisfies all the conditions for being a flattening sequence for γ. Invoking Theorem A completes the proof of Theorem C. ◻

5. SOME APPLICATIONS

Finally, let us discuss some consequences of our main results. As a first application, we obtain some constraints on the behavior of a quasi-isometry between locally compact CAT(0)-spaces.

Corollary 5.1 (Constraints on quasi-isometries). *Let $\tilde{M}_1$, $\tilde{M}_2$ be two locally compact* CAT(0)*-spaces, and assume that $\varphi : \tilde{M}_1 \to \tilde{M}_2$ is a quasi-isometry. Let $\gamma \subset \tilde{M}_1$ be a k-flat, $\gamma_\omega \subset X_1 := \mathrm{Cone}(\tilde{M}_1)$ the corresponding k-flat in the asymptotic cone, and assume that there exists a bi-Lipschitz $(k+1)$-dimensional half-flat $F \subset X_1$ bounding the k-flat $\gamma_\omega \subset X_1$. Then we have the following dichotomy: either*

(1) Non-periodicity: *Every k-flat η at bounded distance from $\varphi(\gamma)$ has the property that $\eta / \mathrm{Stab}_G(\eta)$ is non-compact, where $G = \mathrm{Isom}(\tilde{M}_2)$,* or

(2) Bounding: *Every k-flat η at bounded distance from $\varphi(\gamma)$ bounds a $(k+1)$-dimensional half-flat.*

Proof. This follows immediately from our Theorem C. Assume that the first possibility does not occur, i.e., there exists a k-flat η at bounded distance from $\varphi(\gamma)$ with the property that $\mathrm{Stab}_G(\eta) \subset G = \mathrm{Isom}(\tilde{M}_2)$ acts cocompactly on η. Now recall that the quasi-isometry $\varphi : \tilde{M}_1 \to \tilde{M}_2$ induces a bi-Lipschitz homeomorphism $\varphi_\omega : \mathrm{Cone}(\tilde{M}_1) \to \mathrm{Cone}(\tilde{M}_2)$. Since $\eta \subset \tilde{M}_2$ was a k-flat at finite distance from $\varphi(\gamma)$, we have the containment:

$$\varphi_\omega(\gamma_\omega) \subseteq \eta_\omega \subset \mathrm{Cone}(\tilde{M}_2).$$

Since $\varphi_\omega(\gamma_\omega)$ is a bi-Lipschitz copy of $\mathbb{R}^k$ inside the k-flat η_ω, we conclude that $\varphi(\gamma_\omega) = \eta_\omega$. But recall that we assumed that γ_ω was contained inside a bi-Lipschitz flat $\gamma_\omega \subset F \subset \mathrm{Cone}(\tilde{M}_1)$, and hence we see that $\eta_\omega \subset \varphi_\omega(F)$ is likewise contained inside a bi-Lipschitz flat. Since the hypotheses of Theorem C are satisfied, we conclude that η must bound a $(k+1)$-dimensional half-flat, concluding the proof of Corollary 5.1. ◻

The statement of our first corollary might seem somewhat complicated. We now proceed to isolate a special case of most interest:

Corollary 5.2 (Constraints on perturbations of metrics). *Assume that (M, g_0) is a closed Riemannian manifold of non-positive sectional curvature, and assume that $N^k \hookrightarrow M$ is an isometrically embedded compact flat k-manifold with image γ_0. Let*

$\tilde{\gamma}_0 \subset \tilde{M}$ be the k-flat obtained by taking a connected lift of γ_0, and assume that $\tilde{\gamma}_0$ bounds a $(k+1)$-dimensional half-flat F_0.

Then if (M, g) is any other Riemannian metric on M with non-positive sectional curvature, and $\gamma \subset M$ is an isometrically embedded flat k-manifold (in the g-metric) freely homotopic to γ_0, then the lift $\tilde{\gamma} \subset (\tilde{M}, \tilde{g})$ must also bound a $(k+1)$-dimensional half-flat F.

We can think of Corollary 5.2 as a "non-periodic" version of the Flat Torus theorem. Indeed, in the case where F is $\pi_1(M)$-periodic, the Flat Torus theorem applied to (M, g) implies that $\tilde{\gamma}$ is likewise contained in a periodic flat.

Proof. Since M is compact, the identity map provides a quasi-isometry $\varphi : (\tilde{M}, \tilde{g}_0) \to (\tilde{M}, \tilde{g})$. The half-flat F_0 containing $\tilde{\gamma}_0$ gives rise to a flat $(F_0)_\omega \subset \mathrm{Cone}(\tilde{M}, \tilde{g}_0)$ containing $(\tilde{\gamma}_0)_\omega$. In particular, we can apply the previous Corollary 5.1.

Next note that, since γ_0, γ are freely homotopic to each other, there is a lift $\tilde{\gamma}$ of γ which is at finite distance (in the g-metric) from the given $\tilde{\gamma}_0 \subset (\tilde{M}, \tilde{g})$. Indeed, taking the free homotopy $H : N^k \times [0,1] \to M$ between $H_0 = \gamma_0$ and $H_1 = \gamma$, we can then take a lift $\tilde{H} : \mathbb{R}^k \times [0,1] \to \tilde{M}$ satisfying the initial condition $\tilde{H}_0 = \tilde{\gamma}_0$ (the given lift of γ_0). The time one map $\tilde{H}_1 : \mathbb{R}^k \to \tilde{M}$ will be a lift of $H_1 = \gamma$, hence a k-flat in $(\tilde{M}, \tilde{g})$. Furthermore, the distance (in the g-metric) between $\tilde{\gamma}_0$ and $\tilde{\gamma}$ will clearly be bounded above by the supremum of the g-lengths of the (compact) family of maps $H_p : [0,1] \to (M, g)$, $p \in N^k$, defined by $H_p(t) = H(p, t)$.

Now observe that by construction, the $\tilde{\gamma} \subset (\tilde{M}, \tilde{g})$ from the previous paragraph has $\mathrm{Stab}_G(\tilde{\gamma})$ acting cocompactly on $\tilde{\gamma}$, where $G = \mathrm{Isom}(\tilde{M}, \tilde{g})$. Hence the first possibility in the conclusion of Corollary 5.1 cannot occur, and we conclude that $\tilde{\gamma}$ must bound a $(k+1)$-dimensional half-flat F, as desired. This concludes the proof of Corollary 5.2. ◻

Next we recall some terminology from differential geometry: for M a complete, simply connected, Riemannian manifold of non-positive sectional curvature, the *rank* of a geodesic $\gamma \subset X$ is the dimension $\mathrm{rk}(\gamma)$ of the vector space of parallel Jacobi fields along γ. Note that the unit tangent vector field is always parallel, hence the rank of a geodesic is always ≥ 1; a geodesic is said to have *higher rank* provided $\mathrm{rk}(\gamma) \geq 2$. A geodesic γ that bounds a 2-dimensional half-flat automatically has $\mathrm{rk}(\gamma) \geq 2$, as the unit normal vector field within the half-flat will be a parallel Jacobi field along γ. Finally, the *manifold* M is said to have higher rank provided *every* geodesic $\gamma \subset M$ satisfies $\mathrm{rk}(\gamma) \geq 2$. The celebrated rank-rigidity theorem, established independently by Ballmann [2] and Burns-Spatzier [7], states that if M has higher rank, then:

(1) either M is isometric to an irreducible, higher-rank symmetric space of non-compact type,

(2) or M is reducible, and splits isometrically as a product $M_1 \times M_2$ of lower dimensional manifolds of non-positive sectional curvature.

Our next two applications will exploit the combination of our main results with the rank-rigidity theorem to deduce some information concerning manifolds of non-positive sectional curvature.

Now recall that the classic de Rham theorem [25] states that any simply connected, complete Riemannian manifold admits a decomposition as a metric product $\tilde{M} = \mathbb{R}^k \times M_1 \times \cdots \times M_k$, where $\mathbb{R}^k$ is a Euclidean space equipped with the standard metric, and each M_i is metrically irreducible (and not $\mathbb{R}$ or a point). Furthermore, this decomposition is unique up to permutation of the factors. This result was recently generalized by Foertsch-Lytchak [13] to cover finite dimensional geodesic metric spaces (such as ultralimits of Riemannian manifolds). Our next corollary shows that, in the presence of non-positive Riemannian curvature, there is a strong relationship between splittings of $\tilde{M}$ and splittings of $\mathrm{Cone}(\tilde{M})$.

Corollary 5.3 (Asymptotic cones detect splittings)**.** *Let M be a closed Riemannian manifold of non-positive curvature, $\tilde{M}$ the universal cover of M with induced Riemannian metric, and $X = \mathrm{Cone}(\tilde{M})$ an arbitrary asymptotic cone of $\tilde{M}$. If*

$$\tilde{M} = \mathbb{R}^k \times M_1 \times \cdots \times M_n$$

is the de Rham splitting of $\tilde{M}$ into irreducible factors, and

$$X = \mathbb{R}^\ell \times X_1 \times \cdots \times X_m$$

is the Foertsch-Lytchak splitting of X into irreducible factors, then $k = \ell$, $n = m$, and up to a relabeling of the index set, we have that each $X_i = \mathrm{Cone}(M_i)$.

Proof. Let us first assume that $\tilde{M}$ is irreducible (i.e., k=0, n=1), and show that $X = \mathrm{Cone}(\tilde{M})$ is also irreducible (i.e., $\ell = 0$, $m = 1$). By way of contradiction, let us assume that X splits as a metric product, and observe that this clearly implies that every geodesic $\gamma \subset X$ is contained inside a flat. In particular, from our Theorem A, we see that every geodesic inside $\tilde{M}$ must bound a 2-dimensional half-flat, and hence must have higher rank. Applying the Ballmann, Burns-Spatzier rank rigidity result, and recalling that $\tilde{M}$ was irreducible, we conclude that $\tilde{M}$ is in fact an irreducible higher rank symmetric space. But now Kleiner-Leeb have shown that for such spaces, the asymptotic cone is irreducible (see [17, Section 6]), giving us the desired contradiction.

Let us now proceed to the general case: from the metric splitting of $\tilde{M}$, we get a corresponding metric splitting $\mathrm{Cone}(\tilde{M}) = \mathbb{R}^k \times Y_1 \times \cdots \times Y_n$, where each $Y_i = \mathrm{Cone}(M_i)$. Since each M_i is irreducible, the previous paragraph tells us that each Y_i is likewise irreducible. So we now have two product decompositions of $\mathrm{Cone}(\tilde{M})$ into irreducible factors. So assuming that each Y_i is distinct from a point and is not isometric to $\mathbb{R}$, we could appeal to the uniqueness portion of Foertsch-Lytchak [13, Theorem 1.1] to conclude that, up to relabeling of the index set, each $X_i = Y_i = \text{Cone}(M_i)$, and that the Euclidean factors have to have the same dimension $k = \ell$.

To conclude the proof of our corollary, we need to argue that if $\tilde{M}$ is a simply connected, complete, Riemannian manifold of non-positive sectional curvature, and $\dim(M) \geq 2$, then $\text{Cone}(\tilde{M})$ is distinct from a point or $\mathbb{R}$. Equivalently, one just needs to argue that the asymptotic cone of $\pi_1(M)$ is distinct from a point or $\mathbb{R}$. But Druţu and Sapir [10, Corollary 6.2] have shown that any finitely generated group with an asymptotic cone homeomorphic to a point or $\mathbb{R}$ has to be virtually cyclic, hence must either be finite or have virtual cohomological dimension = 1. Since $\pi_1(M)$ is an infinite group with virtual cohomological dimension = $\dim(M) \geq 2$, we see that neither of these two options can occur. This concludes the proof of Corollary 5.3. □

Before stating our next result, we recall that the celebrated rank rigidity theorem of Ballmann, Burns-Spatzier was motivated by Gromov's well-known rigidity theorem, the proof of which appears in the book [3]. Our next corollary shows how in fact Gromov's rigidity theorem can now be directly deduced from the rank rigidity theorem. This is our last result:

Corollary 5.4 (Gromov's higher rank rigidity [3])**.** *Let M^* be a compact locally symmetric space of $\mathbb{R}$-rank ≥ 2, with universal cover $\tilde{M}^*$ irreducible, and let M be a compact Riemannian manifold with sectional curvature $K \leq 0$. If $\pi_1(M) \cong \pi_1(M^*)$, then M is isometric to M^*, provided* $\text{Vol}(M) = \text{Vol}(M^*)$.

Proof. Since both M and M^* are compact with isomorphic fundamental groups, the Milnor-Švarc theorem gives us quasi-isometries:

$$\tilde{M}^* \simeq \pi_1(M^*) \simeq \pi_1(M) \simeq \tilde{M}$$

which induce a bi-Lipschitz homeomorphism $\varphi : \text{Cone}(\tilde{M}^*) \to \text{Cone}(\tilde{M})$. Now in order to apply the rank rigidity theorem, we need to establish that every geodesic in $\tilde{M}$ has rank ≥ 2.

We first observe that the proof of Corollary 5.2 extends almost verbatim to the present setting. Indeed, in Corollary 5.2, we used the identity map to induce a bi-Lipschitz homeomorphism between the asymptotic cones, and then appealed to Corollary 5.1. The sole difference in our present context is that, rather than using the identity map, we use the quasi-isometry between $\tilde{M}$ and $\tilde{M}^*$ induced by the isomorphism $\pi_1(M) \cong \pi_1(M^*)$. This in turn induces a bi-Lipschitz homeomorphism between asymptotic cones (see Section 2). The reader can easily verify that the rest of the argument in Corollary 5.2 extends to our present setting, establishing that *every lift to $\tilde{M}$ of a periodic geodesic* in M has rank ≥ 2.

So we now move to the general case, and explain why *every* geodesic in $\tilde{M}$ has higher rank. To see this, assume by way of contradiction that there is a geodesic $\eta \subset \tilde{M}$ with $\text{rk}(\eta) = 1$. Note that the geodesic η cannot bound a half-plane. But once we have the existence of such an η, we can appeal to results of Ballmann [1, Theorem 2.13], which imply that η can be approximated (uniformly on compacts) by lifts of periodic geodesics in M; let $\{\tilde{\gamma}_i\} \to \eta$ be such an approximating sequence. Since each $\tilde{\gamma}_i$ has $\mathrm{rk}(\tilde{\gamma}_i) \geq 2$, it supports a parallel Jacobi field J_i, which can be taken to satisfy $\|J_i\| \equiv 1$ and $\langle J_i, \tilde{\gamma}_i' \rangle \equiv 0$. Now we see that:

- The limiting vector field J defined along η exists, due to the control on $\|J_i\|$;
- The vector field J along η is a parallel Jacobi field, since both the "parallel" and "Jacobi" condition can be encoded by differential equations with smooth coefficients, solutions to which will vary continuously with respect to initial conditions; and
- J will have unit length and will be orthogonal to η', from the corresponding condition on the J_i.

But this contradicts our assumption that $\mathrm{rk}(\eta) = 1$. So we conclude that every geodesic $\eta \subset \tilde{M}$ must satisfy $\mathrm{rk}(\eta) \geq 2$, as desired.

From the rank rigidity theorem, we now obtain that $\tilde{M}$ either splits as a product, or is isometric to an irreducible higher rank symmetric space. Since the asymptotic cone of the irreducible higher rank symmetric space is topologically irreducible (see [17, Section 6]), and $\mathrm{Cone}(\tilde{M})$ is homeomorphic to $\mathrm{Cone}(\tilde{M}^*)$, we have that $\tilde{M}$ cannot split as a product. Finally, we see that $\pi_1(M) \cong \pi_1(M^*)$ acts cocompactly, isometrically on two irreducible higher rank symmetric spaces $\tilde{M}$ and $\tilde{M}^*$. By Mostow rigidity [21], we have that the quotient spaces are, after suitably rescaling, isometric. This completes our proof of Gromov's higher rank rigidity theorem. □

Finally, let us conclude our paper with a few comments on this last corollary.

Remark. (1) The actual statement of Gromov's theorem in [3, p. (i)] does not assume $\tilde{M}^*$ to be irreducible, but rather M^* to be irreducible (i.e., there is no finite cover of M^* that splits isometrically as a product). This leaves the possibility that the universal cover $\tilde{M}^*$ splits isometrically as a product, but no finite cover of M^* splits isometrically as a product. However, in this specific case, the desired result was already proved by Eberlein (see [11]). And in fact, in the original proof of Gromov's rigidity theorem, the very first step (see [3, p. 154]) consists of appealing to Eberlein's result to reduce to the case where $\tilde{M}^*$ is irreducible.

(2) In the course of writing this paper, the authors learnt of the existence of another proof of Gromov's rigidity result, which bears some similarity to our reasoning. As the reader has surmised from the proof of Corollary 5.4, the key is to somehow show that M also has to have higher rank. But a sophisticated result of Ballmann-Eberlein [4] establishes that the rank of a non-positively curved Riemannian manifold M can in fact be detected directly from algebraic properties of $\pi_1(M)$, and hence the property of having "higher rank" is in fact algebraic (see

also the recent paper of Bestvina-Fujiwara [5]). The main advantage of our approach is that one can deduce Gromov's rigidity result directly from rank rigidity, and indeed, that one can geometrically "see" that the property of having higher rank is preserved.

(3) We point out that various other mathematicians have obtained results extending Gromov's theorem (and which do not seem tractable using our methods). A variation considered by Davis-Okun-Zheng [9], requires $\tilde{M}^*$ to be *reducible* and M^* to be irreducible (the same hypothesis as in Eberlein's rigidity result). However, Davis-Okun-Zheng allow the metric on M to be locally CAT(0) (rather than Riemannian non-positively curved), and are still able to conclude that M is isometric (after rescaling) to M^*. The optimal result in this direction is due to Leeb [19], giving a characterization of certain higher rank symmetric spaces and Euclidean buildings within the broadest possible class of metric spaces, the Hadamard spaces (complete geodesic spaces for which the distance function between pairs of geodesics is always convex). It is worth mentioning that Leeb's result relies heavily on the viewpoint developed in the Kleiner-Leeb paper [17].

(4) We note that our method of proof can also be used to establish a *non-compact, finite volume* analogue of the previous corollary. Three of the key ingredients going into our proof were

(i) Ballmann's result on the density of periodic geodesics in the tangent bundle;
(ii) Ballmann-Burns-Spatzier's rank rigidity theorem; and
(iii) Mostow's strong rigidity theorem.

A finite volume version of (i) was obtained by Croke-Eberlein-Kleiner (see [8, Appendix]), under the assumption that the fundamental group is finitely generated. A finite volume version of (ii) was obtained by Eberlein-Heber (see [12]). The finite volume versions of Mostow's strong rigidity were obtained by Prasad in the $\mathbb{Q}$-rank one case [23] and Margulis in the $\mathbb{Q}$-rank ≥ 2 case [20] (see also [24]). One technicality in the non-compact case is that isomorphisms of fundamental groups no longer induce quasi-isometries of the universal cover. In particular, it is no longer sufficient to just assume $\pi_1(M) \cong \pi_1(M^*)$, but rather one needs a homotopy equivalence $f : M \to M^*$ with the property that f lifts to a quasi-isometry $\tilde{f} : \tilde{M} \to \tilde{M}^*$. We leave the details to the interested reader.

Acknowledgements. The authors would like to thank Vincent Guirardel, Juha Heinonen, Tadeusz Januszkiewicz, Bruce Kleiner, Ralf Spatzier, and Stefan Wenger for their helpful comments. We are indebted to Michael Kapovich for pointing out the existence of his joint paper with Bernard Leeb [16], which put us on the track of the results contained in this paper. We would also like to thank the anonymous referee for his careful reading of the paper, and several helpful suggestions.

This work was partly carried out during a visit of the first author to the Ohio State University (supported in part by the MRI), and a visit of the second author to the Università di Pisa. The work of the first author was partly supported by the European Research Council—MEIF-CT-2005-010975 and MERG-CT-2007-046557. The work of the second author was partly supported by NSF grants DMS-0606002, DMS-0906483, and by an Alfred P. Sloan research fellowship.

REFERENCES

[1] W. BALLMANN, *Axial isometries of manifolds of nonpositive curvature*, Math. Ann. **259** (1982), 131–144, http://dx.doi.org/10.1007/BF01456836. MR 656659 (83i:53068)

[2] ______, *Nonpositively curved manifolds of higher rank*, Ann. of Math. (2) **122** (1985), 597–609, http://dx.doi.org/10.2307/1971331. MR 819559 (87e:53059)

[3] W. BALLMANN, M. GROMOV, and V. SCHROEDER, *Manifolds of Nonpositive Curvature*, Progress in Mathematics, vol. 61, Birkhäuser Boston Inc., Boston, MA, 1985, ISBN 0-8176-3181-X. MR 823981 (87h:53050)

[4] W. BALLMANN and P. EBERLEIN, *Fundamental groups of manifolds of nonpositive curvature*, J. Differential Geom. **25** (1987), 1–22. MR 873453 (88b:53047)

[5] M. BESTVINA and K. FUJIWARA, *A characterization of higher rank symmetric spaces via bounded cohomology*, Geom. Funct. Anal. **19** (2009), 11–40, http://dx.doi.org/10.1007/s00039-009-0717-8. MR 2507218

[6] M.R. BRIDSON and A. HAEFLIGER, *Metric Spaces of Non-positive Curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999, ISBN 3-540-64324-9. MR 1744486 (2000k:53038)

[7] K. BURNS and R. SPATZIER, *Manifolds of nonpositive curvature and their buildings*, Inst. Hautes Études Sci. Publ. Math. (1987), 35–59, http://dx.doi.org/10.1007/BF02698934. MR 908215 (88g:53050)

[8] C.B. CROKE, P. EBERLEIN, and B. KLEINER, *Conjugacy and rigidity for nonpositively curved manifolds of higher rank*, Topology **35** (1996), 273–286, http://dx.doi.org/10.1016/0040-9383(95)00031-3. MR 1380497 (96m:53047)

[9] M.W. DAVIS, B. OKUN, and F. ZHENG, *Piecewise Euclidean structures and Eberlein's rigidity theorem in the singular case*, Geom. Topol. **3** (1999), 303–330 (electronic), http://dx.doi.org/10.2140/gt.1999.3.303. MR 1714914 (2000j:53054)

[10] C. DRUŢU and M. SAPIR, *Tree-graded spaces and asymptotic cones of groups*, Topology **44** (2005) (With an appendix by D. Osin and M. Sapir), 959–1058, http://dx.doi.org/10.1016/j.top.2005.03.003. MR 2153979 (2006d:20078)

[11] P. EBERLEIN, *Rigidity of lattices of nonpositive curvature*, Ergodic Theory Dynam. Systems **3** (1983), 47–85.

[12] P. EBERLEIN and J. HEBER, *A differential geometric characterization of symmetric spaces of higher rank*, Inst. Hautes Études Sci. Publ. Math. (1990), 33–44, http://dx.doi.org/10.1007/BF02699876. MR 1079642 (91j:53022)

[13] T. FOERTSCH and A. LYTCHAK, *The de Rham decomposition theorem for metric spaces*, Geom. Funct. Anal. **18** (2008), 120–143, http://dx.doi.org/10.1007/s00039-008-0652-0. MR 2399098 (2010c:53061)

[14] S. FRANCAVIGLIA and J.-F. LAFONT, *Asymptotic cones, bi-Lipschitz ultraflats, and the geometric rank of geodesics*, arXiv:0801.3636 (preprint).

[15] M. KAPOVICH, *Triangle inequalities in path metric spaces*, Geom. Topol. **11** (2007), 1653–1680, http://dx.doi.org/10.2140/gt.2007.11.1653. MR 2350463 (2009e:51019)

[16] M. KAPOVICH and B. LEEB, 3-*manifold groups and nonpositive curvature*, Geom. Funct. Anal. **8** (1998), 841–852, http://dx.doi.org/10.1007/s000390050076. MR 1650098 (2000a:57040)

[17] B. KLEINER and B. LEEB, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, Inst. Hautes Études Sci. Publ. Math. (1997), 115–197 (1998), http://dx.doi.org/10.1007/BF02698902. MR 1608566 (98m:53068)

[18] L. KRAMER, S. SHELAH, K. TENT, and S. THOMAS, *Asymptotic cones of finitely presented groups*, Adv. Math. **193** (2005), 142–173, http://dx.doi.org/10.1016/j.aim.2004.04.012. MR 2132762 (2006d:20075)

[19] B. LEEB, *A Characterization of Irreducible Symmetric Spaces and Euclidean Buildings of Higher Rank by Their Asymptotic Geometry*, Bonner Mathematische Schriften [Bonn Mathematical Publications], vol. 326, Universität Bonn Mathematisches Institut, Bonn, 2000. MR 1934160 (2004b:53060)

[20] G.A. MARGULIS, *Non-uniform lattices in semisimple algebraic groups*, Lie Groups and Their Representations (Proc. Summer School on Group Representations of the Bolyai János Math. Soc., Budapest, 1971), Halsted, New York, 1975, pp. 371–553. MR 0422499 (54 #10486)

[21] G.D. MOSTOW, *Strong Rigidity of Locally Symmetric Spaces*, Annals of Mathematics Studies, vol. 78, Princeton University Press, Princeton, N.J., 1973. MR 0385004 (52 #5874)

[22] A.Yu. OL′SHANSKII and M.V. SAPIR, *A finitely presented group with two non-homeomorphic asymptotic cones*, Internat. J. Algebra Comput. **17** (2007), 421–426, http://dx.doi.org/10.1142/S021819670700369X. MR 2310154 (2008c:20079)

[23] G. PRASAD, *Strong rigidity of* $\mathbb{Q}$*-rank* 1 *lattices*, Invent. Math. **21** (1973), 255–286, http://dx.doi.org/10.1007/BF02564308. MR 0385005 (52 #5875)

[24] M. S. RAGHUNATHAN, *Discrete groups and* $\mathbb{Q}$*-structures on semi-simple Lie groups*, Discrete Subgroups of Lie Groups and Applications to Moduli (Internat. Colloq., Bombay, 1973), Oxford Univ. Press, Bombay, 1975, pp. 225–321. MR 0404527 (53 #8327)

[25] G. DE RHAM, *Sur la reductibilité d'un espace de Riemann*, Comment. Math. Helv. **26** (1952), 328–344, http://dx.doi.org/10.1007/BF02564308. MR 0052177 (14,584a) (French)

[26] S. THOMAS and B. VELICKOVIC, *Asymptotic cones of finitely generated groups*, Bull. London Math. Soc. **32** (2000), 203–208, http://dx.doi.org/10.1112/S0024609399006621. MR 1734187 (2000h:20081)

[27] S. WENGER, *Isoperimetric inequalities of Euclidean type and applications*, Ph.D. Thesis, ETH, Zürich, 2004.

STEFANO FRANCAVIGLIA:
Departimento di Matematica
Università di Bologna
Piazza di Porta S. Donato 5
40126, Bologna, Italy

E-MAIL: francavi@dm.unibo.it

JEAN-FRANÇOIS LAFONT:
Department of Mathematics
The Ohio State University
231 West 18th Avenue
Columbus, OH 43210, U.S.A.

E-MAIL: jlafont@math.ohio-state.edu

KEY WORDS AND PHRASES: CAT(0)-space; asymptotic cone; flats; rigidity.

2000 MATHEMATICS SUBJECT CLASSIFICATION: 51F99; 53C24; 53C35.

Received: *March 23rd, 2010; revised: June 15th, 2010.*

Article electronically published on August 3rd, 2010.