BILIPSCHITZ TRIANGULATIONS OF RIEMANNIAN MANIFOLDS

BRIAN H. BOWDITCH

ABSTRACT. We give a proof that a riemannian manifold of bounded geometry admits a uniformly bilipschitz smooth triangulation. The bilipschitz constants depend only on the curvature and injectivity radius bounds. A number of refinements of this statement, and a generalisation to manifolds with boundary, are described.

1. Introduction

A classical result of differential topology tells us that any smooth manifold is triangulable, and thereby admits a PL structure which is canonical up to topological isotopy. A proof is given in [W], inspired by earlier work in [Ca]. A more recent account (including manifolds with boundary) is given in [M].

Since every smooth manifold admits a riemannian metric, it is enough to deal with riemannian manifolds. It can be shown that if one places bounds on curvature and injectivity radius (i.e. "bounded geometry"), then one can arrange for the triangulation to be bilipschitz with respect to the standard euclidean metric on the simplicial complex (that is, where each simplex is regular, with unit side-lengths). Moreover, one can put another euclidean metric on the complex, so that the bilipschitz constants are arbitrarily close to 1. A proof of this has recently been given in [BoDG], making use of some constructions in [DyVW]. Some earlier related results can be found in [S] and [Br].

We will give a self-contained account of such triangulations. The overall strategy uses generic Delaunay triangulations, similarly to that of [BoDG], though the details are somewhat different. We also generalise to the case of manifolds with boundary.

Here are some more precise statements of the results.

Let M be a complete smooth riemannian n-manifold.

Definition. We say M is (κ, χ)-*bounded* if all its sectional curvatures lie between $-\kappa$ and κ, and its convexity radius is at least $\chi > 0$ everywhere.

(Recall that the convexity radius at any given point is at least half of the injectivity radius, so it would be essentially equivalent to place a lower bound on the the latter quantity.)

Date: 11th April 2020.

We write ρ_M for the induced path-metric on M.

Let Θ be a simplicial complex. (All simplicial complexes here will be locally finite.) A "euclidean structure" on Θ is the path-metric, ρ_Θ, obtained by giving each simplex the structure of a euclidean simplex. (Note that the euclidean structure is determined by the edge-lengths.) A "standard (euclidean)" simplex is one where all edge-lengths are equal to 1. If all simplices of Θ are (intrinsically isometric to) standard ones, then we denote the induced path metric by ρ^0_Θ, and refer to it as the "standard structure". Of course, this is equivalent to saying that the edge-lengths of the 1-skeleton are all equal to 1.

By a *triangulation* of M we mean a simplicial complex, Θ, together with a homeomorphism, $\tau : \Theta \longrightarrow M$. It is *smooth* if its restriction to each simplex is smooth (in the sense defined below). We say that it is λ-*bilipschitz* if it is a triangulation, and if it is λ-bilibschitz with respect to the metric ρ^0_θ.

We show:

Theorem 1.1. *Given $\kappa, \chi > 0$ and $n \in \mathbb{N}$, there is some $\lambda > 0$ such that any (κ, χ)-bounded n-manifold admits a λ-bilipschitz smooth triangulation.*

In fact, this is an immediate consequence of a stronger statement which we formulate as follows.

Theorem 1.2. *Suppose that M is a (κ, χ)-bounded m-manifold, and that $\mu > 1$. Then there is some $\eta_0 > 0$, depending only on κ, μ, such that if $\eta \leq \min\{\chi, \eta_0\}$, then there is a euclidean simplicial complex (Θ, ρ_Θ), and a μ-bilipschitz smooth triangulation, $\tau : \Theta \longrightarrow M$, with $\operatorname{diam}(\tau(\Sigma)) \leq \eta$ for all simplices Σ of Θ. Moreover, we can assume that τ maps every 1-simplex of Θ isometrically to a geodesic segment in M. In fact, we can assume that $\rho_\Theta \leq \eta\rho^0_\Theta \leq \nu\rho_\Theta$, where $\nu \geq 1$ depends only on the dimension, m.*

Note that the penultimate statement determines the metric, ρ_Θ, on Θ uniquely.

One has a similar statement for a manifold M with boundary, ∂M. In this case we say that M is (κ, χ)-*bounded* if in addition to the previous requirements, the intrinsic convexity radius of ∂M is at least 2χ, the extrinsic curvature of ∂M is at most κ in norm, and ∂M has a (2χ)-collar. The last statement means that the (2χ)-neighbourhood of ∂M retracts onto ∂M.

Theorem 1.3. *Suppose that M is a (κ, χ)-bounded manifold with boundary. Then the conclusion of Theorem 1.2 holds with the following modification. Any edge of the 1-skeleton of the triangulation is either an intrinsic geodesic segment in ∂M, or a geodesic segment in M which meets ∂M, if at all, in one of its endpoints.*

(See Section 8 for a more precise statement of this result.)

It's not hard to see that ∂M is necessarily a subcomplex of any smooth triangulation. (In fact, this will follow directly from our construction.)

In the case where M has constant curvature, without loss of generality, -1, 0 or 1, we can arrange that each simplex is totally geodesic. In other words, we get:

Theorem 1.4. *Suppose that M is a hyperbolic, euclidean, or spherical n-manifold of convexity radius at least $\chi > 0$. Then M admits a triangulation where each n-simplex is respectively (isometric to) a hyperbolic, euclidean or spherical simplex, respectively, whose inradius is bounded below, and whose circumradius is bounded above, and where these bounds depend only on n and χ.*

In particular, this means that we can find an arbitrarily fine triangulation of M, which is arbitrarily close to being isometric to the euclidean structure on the triangulation with the same edge-lengths. Moreover, after appropriate rescaling, it is uniformly bilipschitz equivalent to the standard euclidean structure. We can allow M to have totally geodesic boundary, provided that it has a uniform collar.

For convenience in this paper, we will interpret "smooth" to mean C^∞, though one could make similar statements with different degrees of regularity. To interpret curvature bounds as stated here one would need to assume M to be at least C^2. We make use of this degree of regularity in Sections 4 and 8. However, one can meaningfully interpret curvature bounds in terms of comparison axioms, without any a-priori differentiability assumption. We will not explore these technical issues here.

To say that a triangulation, τ, is *smooth* we mean that τ restricted to each simplex Σ is smooth. In other words, if we take Σ embedded as a standard simplex in the euclidean space of the same dimension, then we can find a smooth extension to some open set containing containing Σ. Note that the collection of such extensions (or more precisely, their respective inverse maps) gives rise to a smooth atlas for M.

We note that a smooth triangulation, $\tau : \Theta \longrightarrow M$, is a combinatorial triangulation, in the sense that the link of every simplex is homeomorphic to a sphere. Hence, as usual, we get a PL structure on M. See [M] for more discussion of these matters.

We remark that, in Theorem 1.1, we could equivalently replace "triangulation" with "cubulation" — in the latter case, Θ is a cube complex built out of regular unit euclidean cubes. (Note that, by coning over centroids, we can canonically subdivide an n-simplex into $n+1$ n-cubes, and an n-cube into $2^n n!$ n-simplices. These give rise to bilipschitz equivalent metrics. Moreover, it is a simple matter to smooth out these subdivisions.) Cube complexes have the geometric advantage that subdividing into equal smaller regular cubes does not change the induced path-metric — at least up to rescaling.

Every smooth manifold admits a complete riemannian metric of bounded geometry (indeed, one which is asymptotically flat and with the convexity radius tending to infinity. We therefore recover the fact that any smooth manifold admits a smooth combinatorial triangulation, and hence a PL structure. However, there are additional facts about triangulations which we do not address here. For example, one can show that the PL structure is unique up to topological isotopy (see [M]).

We briefly describe how our results and constructions relate to some of those elsewhere. As noted earlier the construction (for manifolds without boundary) is similar to that in [DyVW, BoDG]. Namely, one starts with a discrete net of points in the manifold, and then perturbs it to become sufficiently "generic" (or "regular", as we call it here). This net is identified as the vertex set of a combinatorial "Delaunay" complex. One extends the inclusion of the vertex set to a map of the whole complex to give a smooth triangulation. We use a different method to perturb the net to that described in [BoDG]. Also the method for extending to a triangulation is different. The paper [DyVW] uses barycentric coordinates as described by Karcher [K], whereas we use a different extension method (as can be found in [L] for example). The former is arguably more natural, though the latter adapts more readily to manifolds with boundary.

A related construction using Vorornoi cells was proposed in [ECHLPT] (Theorem 10.3.1 thereof). However, they do not give a detailed proof, and their "coning" construction would not in general give rise to a smooth triangulation.

The paper [S] gives a construction of a "fat" triangulation of a smooth riemannian manifold. The methods are rather different. The hypotheses are more general, but the conclusion is somewhat weaker than ours.

The paper [Br] gives a "thick" triangulation of (constant curvature) hyperbolic manifolds. It also uses Delaunay triangulations. The strategy is broadly similar, but their process for perturbing the net is again a little different. The paper shows that individual simplices are uniformly bilipschitz to standard ones. It does not discuss smooth or bilipschitz triangulations as such. However this could be achieved with a little extra work (cf. Section 9 here).

The outline of this paper is as follows. In Section 2, we describe the Delaunay triangulation in euclidean space. In Section 3, we construct regular nets, again in euclidean space. In Section 4, we give some results relating to smooth maps on simplices. In Section 5, we explain how curvature bounds imply that transition maps can be chosen to C^1-close to euclidean isometries. We apply this to define a notion of "flatness" in Section 6. In Section 7, we construct a triangulation for a manifold without boundary, and prove Theorem 1.2. In Section 8, we generalise to manifolds with boundary, and prove Theorem 1.3. In Section 9, we discuss how the results can be strengthened in the case of constant curvature, and prove Theorem 1.4.

I am grateful to Daryl Cooper for a discussion we had some time ago, regarding methods for showing that smooth manifolds are triangulable. An idea of perturbing nets, broadly along the lines of that discussed in Sections 3 and 7 here, emerged from this discussion.

2. The Delaunay construction

In this section, we recall the construction of the Delaunay triangulation of $\mathbb{R}^m$, with the euclidean metric, ρ_E. This will be used for general manifolds in

Section 6. It will also serve to introduce some notation and terminology used later. In general, the Delaunay construction [De] gives a tessellation by polyhedral cells (dual to the "Voronoi tessellation"). In the "generic" case, this will be a triangulation.

To begin, we recall some linear algebra. We equip $\mathbb{R}^m$ with the euclidean metric, ρ_E, that is, $\rho_E(x, y) = ||x - y||$, where $||.||$ is the euclidean norm.

Suppose that $B = \{x_0, \ldots, x_m\}$ is a set of $m+1$ points in $\mathbb{R}^m$. If B comes with an order defined up to even permutation, we write it as $\vec{B}$. We set

$$\vec{V}_m(\vec{B}) = \begin{vmatrix} x_{01} & x_{02} & \cdots & x_{0m} & 1 \\ x_{11} & x_{12} & \cdots & x_{1m} & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ x_{m1} & x_{m2} & \cdots & x_{mm} & 1 \end{vmatrix}$$

where x_{ij} is the jth coordinate of x_i. We write $V_m(B) = |\vec{V}_m(\vec{B})|$. Note that $V_m(B)$ is invariant under any isometry of $\mathbb{R}^m$. If B is a singleton, then $V_0(B) = 1$. If $B = \{x_0, x_1\}$, then $V_1(B) = \rho_E(x_0, x_1)$.

More generally, if $B = \{x_0, x_1, \ldots, x_k\} \subseteq \mathbb{R}^m$ for some $m > k$, then we define $V_k(B)$ using any k-dimensional subspace containing B. This is well defined by the above observation.

The geometrical interpretation is that V_m is the m-volume of the convex hull, hull(B), of B.

Suppose that $C = \{x_0, \ldots, x_{m+1}\}$ is a set of $m+2$ points in $\mathbb{R}^m$. Again, given an order defined up to even permutation, we set

$$\vec{Q}_m(\vec{C}) = \begin{vmatrix} \sigma_0 & x_{01} & x_{02} & \cdots & x_{0m} & 1 \\ \sigma_1 & x_{11} & x_{12} & \cdots & x_{1m} & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \sigma_{m+1} & x_{m+1,1} & x_{m+1,2} & \cdots & x_{m+1,m} & 1 \end{vmatrix}$$

where $\sigma_i = x_{i1}^2 + x_{i2}^2 + \cdots + x_{im}^2$. We write $Q_m(C) = |\vec{Q}_m(\vec{C})|$.

Note that $Q_m(C) = 0$ if and only if C lies in a codimension-1 sphere in $\mathbb{R}^m$. In general, $Q_m(C)$ can be viewed as a measure of the "non-sphericity" of C.

(To explain this, let V be the top right $(m+1) \times (m+1)$ cofactor in the determinant. In other words, $V = V_m(x_0, \ldots, x_m)$. We can substitute x_{m+1} with $x \in \mathbb{R}^m$, thought of as a variable. If $V \neq 0$, then the determinant has the form $V||x-a||^2 - k$. In other words, it is the equation of a sphere. Substituting $x = x_i$ for any $i \leq m$ gives 0. It is therefore the sphere containing the points $x_0, \ldots, x_m$. Similarly, if $V = 0$, we get a linear equation: that of the plane containing the points $x_0, \ldots, x_m$.)

Let Π be a connected simplicial complex, and let Π^0 be its vertex set. We write $\mathcal{C}^i(\Pi)$ for the set of formal i-simplices (in other words finite subsets of Π^0 of cardinality $i+1$ which bound simplices). Here we are abusing notation slightly by identifying a 0-simplex with a vertex of Π. In other words, we won't bother to distinguish x and $\{x\}$ for $x \in \Pi^0$. We write $\mathcal{C}(\Pi) = \bigcup_{i=0}^{\infty} \mathcal{C}(\Pi)$. Given $C \in \mathcal{C}(\Pi)$, write $\Sigma(C) = \Sigma(\Delta, A) \subseteq \Pi$, for the simplex with vertex set C (i.e. the "realisation" of C). Write $\mathcal{C}^i_{\rm real}(\Pi) = \{\Sigma(C) \mid C \in \mathcal{C}^i(\Pi)\}$, and $\mathcal{C}_{\rm real}(\Pi) = \bigcup_{i=0}^{\infty} \mathcal{C}^i_{\rm real}(\Pi)$. Write $\Pi^i = \bigcup_{j=0}^{i} \bigcup \mathcal{C}^j_{\rm real}$ for the i-skeleton of Π. We write ρ^0_Π for the path-metric on Π induced by giving each simplex the structure of a "standard simplex"; that is, a regular euclidean simplex with all side-lengths equal to 1.

Note that any map $f : \Pi^0 \longrightarrow \mathbb{R}^m$ has a unique "affine extension", $f : \Pi \longrightarrow \mathbb{R}^m$, that is affine on each simplex.

The following definitions make sense in any metric space, (X, ρ_X) (though, for the moment, we can think of X as $\mathbb{R}^m$. We will use $N(.; r)$ to denote an open r-neighbourhood.

Definition. We say that $A \subseteq X$ is *locally finite* (or *discrete*) if every bounded subset is finite.
Given $r > 0$, we say that A is r-*dense* if $X = N(A; r)$.
We say that A is r-*separated* if $\rho(x, y) \geq r$ for all distinct $x, y \in A$. Given $s \geq r > 0$ we say that A is an (r, s)-net if it is r-separated and s-dense.

Note that for any $r > 0$, any maximal r-separated set will be an (r, r)-net. By Zorn's Lemma such always exists.

In this paper, all metric spaces will be proper (complete and locally compact) geodesic spaces. Thus, any r-separated subset will be locally finite.

Given $h > 0$, let $\mathcal{C}(A, h) = \{B \subseteq A \mid \mathrm{diam}(B) \leq h\}$. Let $\Pi = \Pi(A, h)$ be the simplicial complex with vertex set A, and with formal simplices $\mathcal{C}(A, h)$. In other words, $\mathcal{C}(\Pi(A, h)) = \mathcal{C}(A, h)$. If A is locally finite, then $\Pi(A, h)$ is locally finite as a simplical complex.

We now restrict again to the case where $X = \mathbb{R}^m$.

Note that if $A \subseteq \mathbb{R}^m$ is locally finite, then the inclusion of A into $\mathbb{R}^m$ extends to a proper affine map of $\Pi(A, t)$ into $\mathbb{R}^m$. If A is $(t/2)$-dense, then it is easily seen that this map is surjective.

We can now define the Delaunay triangulation.

Let $A \subseteq \mathbb{R}^m$ be locally finite and t-dense. Given $x \in \mathbb{R}^m$, let $A(x) = \{x \in A \mid \rho_E(x, a) = \rho_E(x, A)\}$, i.e. the set of nearest points of A to x. Write $\mathcal{B}(A) = \{A(x) \mid x \in \mathbb{R}^m\}$, and let $\mathcal{C}(A)$ be the set of subsets $C \subseteq A$ such that $C \subseteq B$ for some $B \in \mathcal{B}(A)$. Clearly $\mathcal{C}(A)$ is closed under inclusion, and so we can define $\Delta(A)$ to be the simplicial complex with vertex set $\Delta^0(A) = A$, and with formal simplices $\mathcal{C}(A)$ (so $\mathcal{C}(\Delta(A)) = \mathcal{C}(A)$). Note that $\Delta(A)$ is a subcomplex of $\Pi(A, 2t)$. Moreover, the affine extension, $\tau : \Delta(A) \longrightarrow \mathbb{R}^m$ is surjective.

In general, τ need not be injective. (Suppose, for example, that $|A(x)| = n \geq m + 3$. Then $A(x)$ is the vertex set of an $(n-1)$-simplex in $\Delta(A)$, which maps under τ to $\mathrm{hull}(A(x)) \subseteq \mathbb{R}^m$.) However, generically, this cannot happen.

Theorem 2.1. *Suppose that $A \subseteq \mathbb{R}^m$ is locally finite and t-dense. Suppose that $Q_m(C) \neq 0$ for all $C \in \mathcal{C}^{m+1}(A, 2t)$ (that is, for all $C \subseteq A$ with $|C| = m + 2$ and* diam$(C) \leq 2t$) *then the affine extension,* $\tau : \Delta(A) \longrightarrow \mathbb{R}^m$ *is a triangulation of* $\mathbb{R}^m$.

We refer to $\Delta(A)$ as the *Delaunay triangulation* with vertex set A. It has its origins in [De] where a statement equivalent to Theorem 2.1 can be found. A simple way of understanding this is to note that the m-simplices of the Delaunay triangulation are those whose circumscribing spheres contain no point of A strictly inside.

As well as the standard metric, ρ^0_Δ, $\Delta(A)$ also admits a euclidean metric, ρ_Δ, such that $(\Delta(A), \rho_\Delta)$ is isometric to $(\mathbb{R}^m, \rho_E)$. Here ρ_Δ is a "euclidean structure" on $\Delta(A)$, in the sense defined in the introduction.

If we assume, in addition, that $V_m(B) \geq \omega > 0$ for all $B \in \mathcal{C}^m(\Delta(A))$, (i.e. there is a positive lower bound on the volumes of m-simplices in the Delaunay triangulation) then τ is λ-bilipschitz on each simplex, where $\lambda \geq 1$ depends only on ω, t and m. In other words, the metrics ρ_Δ and ρ^0_Δ are λ-bilipschitz equivalent.

In view of the above, we make the following definition.

Definition. Suppose that $A \subseteq \mathbb{R}^m$, and $t, \omega > 0$. We say that A is ω-*regular on scale* t (or just ω-*regular*, when t is understood) to mean that the following two conditions hold:
(1): whenever $B \subseteq A$ satisfies $|B| = m+1$ and diam$(B) \leq 2t$, we have $V_m(B) \geq \omega$, and
(2): whenever $C \subseteq A$ satisfies $|C| = m+2$ and diam$(C) \leq 2t$, we have $Q_m(C) \geq \omega$.

(We will always assume that $\omega \leq 1$.)

We also note that if we move the points of a regular net a small distance, we get another regular net:

Lemma 2.2. *Given* $\omega, \delta, t > 0$*, there is some* $\omega' > 0$ *with the following property. Suppose that* $A, A' \subseteq \mathbb{R}^m$ *with* $[a \mapsto a']$ *a bijection from* A *to* A' *with* $\rho_E(a, a') \leq \delta$ *for all* $a \in A$*. If* A *is* ω' *regular on scale* $t + \delta$*, then* A' *is* ω*-regular on scale* t*.*

This is a simple consequence of the continuity of the functions, V_m and Q_m defined above. We also note:

Lemma 2.3. *Given* $\omega, t > 0$ *and there is some* $\delta > 0$ *with the following property. Suppose that* A, A' *satisfy the hypotheses of Lemma 2.2 for* ω, δ, t*. Then the map* $[a \mapsto a']$ *induces a simplicial isomorphism of* $\Delta(A)$ *to* $\Delta(A')$*.*

Proof. Let $B \subseteq A$ be a formal simplex of $\Delta(A)$. By ω-regularity, all points of $A \setminus B$ lie a definite distance from the circumscribing sphere centred on A. By definition of $\Delta(A)$, these all lie outside. This remains the case on moving all points a small enough distance. It is a simple exercise to properly quantify the above statements. ∎

We note moreover that if $\mu > 1$, then we can choose $\delta > 0$, depending on ω, t, such that the simplicial isomorphism $\Delta(A) \longrightarrow \Delta(A')$ is μ-bilipschitz with respect to the metrics $\rho_{\Delta(A)}$ and $\rho_{\Delta(A')}$.

We will need the following general observations.

Lemma 2.4. *Given v, δ, there is some $\delta' > 0$ such that the following holds. Suppose $x_1, \ldots, x_p \in \mathbb{R}^m$ with $V_p(x_1, \ldots, x_p) \geq v$. Suppose $l_{ij} > 0$ with $|l_{ij} - \rho_E(x_i, x_j)| \leq \delta'$. Then there exist $y_0, \ldots, y_m \in \mathbb{R}^m$ with $\rho_E(x_i, y_i) \leq \mathbb{R}^m$, with $\rho_E(x_i, x_j) = l_{ij}$ and $\rho_E(x_i, y_i) \leq \delta$ for all i, j.*

Moreover, given $\mu > 1$, we can choose δ' so that, in addition, the affine extension of the map $[x_i \mapsto y_i]$ is μ-bilipschitz. The proof is a simple exercise.

We can apply this to the Delaunay triangulation as follows.

Lemma 2.5. *Given $\omega, t > 0$ there is some $\delta > 0$ such that the following holds. Suppose that $A \subseteq \mathbb{R}^m$ is ω-regular and t-dense and let $\Delta = \Delta(A)$, Suppose that $l' : \mathcal{C}^1(\Delta) \longrightarrow [0, \infty)$ satisfies $|l(e) - l'(e)| \leq \delta$ for all $e \in \mathcal{C}^1(\Delta)$, where $l(e)$ is the length of the edge, $\Sigma(e)$, in the metric ρ_Δ. Then there is a (unique) euclidean metric, ρ'_Δ, on $\Delta(A)$, such that the ρ'_Δ-length of each edge e is $l'(e)$.*

One could easily add that given any $\mu > 1$, we could choose δ so that the metrics ρ_Δ and ρ'_Δ are μ-bilipschitz equivalent (via the identity map). However, this fact won't be used explicitly.

Finally, we note that the Delaunay triangulation only depends locally on A. Indeed, we only need to have A defined on a subset of $\mathbb{R}^m$ for the construction to work. One way to express this is as follows.

Suppose that $A \subseteq N(0; 4t) \subseteq \mathbb{R}^m$ is t-dense in $N(0; 4t)$. Let $\Delta(A)$ be the simplicial complex with vertex set A defined as before. We make the same hypotheses as in Theorem 2.1, namely that $Q_m(C) \neq 0$ for all $C \in \mathcal{C}^{m+1}(A, 2t)$. Let $D_0 \subseteq \Delta(A)$ be the subcomplex whose simplices meet $N(0; 3t)$. Then $\tau | D_0$ is injective and $N(0, 2t) \subseteq \tau(D_0)$. This is a simple observation, noting that nothing outside $N(0; 4t)$ plays any role in the construction in $N(0; 2t)$.

3. Constructing regular nets

We will show that we can perturb any net in $\mathbb{R}^m$ a little bit so that it becomes uniformly regular. The real interest in this comes from its adaptation to sufficiently flat manifolds in Section 7, where we modify some of the arguments given here.

Here is a formal statement:

Lemma 3.1. *Suppose $t_1 > t_0 > 0$ and $\delta > 0$. Then there is some $\omega = \omega(\delta, m, t_0, t_1)$ with the following property. Suppose $A \subseteq \mathbb{R}^m$ is t_0-separated. Then there is an injective map, $[a \mapsto a'] : A \longrightarrow \mathbb{R}^m$, such that $A' = \{a' \mid a \in A\}$ is ω-regular on scale t_1, and $\rho_E(a, a') \leq \delta$ for all $a \in A$.*

Note that if $\delta < t_0/4$, say, then A' is $(t_0/2)$-separated. If it is also t_0-dense, then A' is $(2t_0)$-dense. This means that we can always construct a regular net, by starting with a t_0-net (a maximal t_0-separated subset) and then perturbing it. (Of course one could easily construct uniformly regular nets explicity in $\mathbb{R}^m$.)

Given an upper bound on diameter, then a lower bound on the volume of a simplex implies a lower bound on the volumes of each of its faces. For the proof of Lemma 3.1, it will be convenient to formalise this as follows.

Definition. We say that A is *strongly ω-regular* (at scale t), if the following two conditions hold:
(1): If $1 \leq k \leq m$, and $B \subseteq A$ with $|B| = k+1$ and $\operatorname{diam}(B) \leq t$, then $V_k(B) \geq \omega$.
(2): If $C \subseteq A$ with $|C| = m+2$ and $\operatorname{diam}(C) \leq t$, then $Q_m(B) \geq \omega$.

In other words, we have strengthened (1) so that it holds for all dimensions at most m.

The following is a key step in the process.

Lemma 3.2. *Given $t_1 > t_0 > 0$, $\delta, \omega > 0$, there is some $\omega' = \omega'(m, \delta, \omega, t_1, t_0)$, with the following property. Suppose that $A \subseteq \mathbb{R}^m$ and $a \in \mathbb{R}^m \setminus A$, with $A \cup \{a\}$ t_0-separated. Suppose that A is strongly ω-regular on scale t_1. Then there is some $a' \in \mathbb{R}^m$, with $\rho_E(a, a') \leq \delta$, and with $A \cup \{a'\}$ strongly ω'-regular on scale t_1.*

Proof. We only need to consider $E = A \cap N(a, 2t_1 + \delta)$. Note that $|E|$ is bounded above in terms of m, t_0, t_1, δ.

Suppose that $B \subseteq E$ with $|B| = k$, with $k \leq m-1$. By hypothesis, $V_k(B) \geq \omega$. Let $P_B \subseteq \mathbb{R}^m$ be the (unique) k-plane containing B. If $x \in \mathbb{R}^m$ with $\rho_E(x, P_B) \geq \eta > 0$, say, then $V_k(B \cup \{x\})$ is bounded below by a fixed positive function (depending on m) of ω and η. (In fact, $V_{k+1}(B \cup \{x\}) \geq \omega\eta/m!$.)

Similarly, suppose $C \subseteq E$ with $|C| = m+1$. Note that $V_{m+1}(C) \geq \omega$. Let $S_C \subseteq \mathbb{R}^m$ be the (unique) m-sphere containing C. If $x \in \mathbb{R}^m$ with $\rho_E(x, S_C) \geq \eta > 0$, say, then $Q_m(C \cup \{x\})$ is bounded below by a fixed positive function (depending on m) of ω and η.

Now consider the union, Y, of all sets of the form P_B and S_C, and B, C range over all such subsets. Note that there are boundedly many such B, C, and so Y is a union of boundedly many planes and spheres of codimension at least 1. If η is small enough in relation to these bounds, then $N(Y; \eta)$ cannot contain any δ-ball in $\mathbb{R}^m$. In other words, we can find some $a' \in N(a; \delta)$ with $\rho_E(a', Y) \geq \eta$, where $\eta > 0$ depends only on δ, m, ω and t_2/t_0. This places a lower bound, say $\omega' > 0$, on $V_{k+1}(B \cup \{a'\})$ and $Q_m(C \cup \{a'\})$ for all such B, C.

We can suppose that $\omega' < \omega$. Given that A is assumed strongly ω-regular, this account for all subsets of $A \cup \{a'\}$ of cardinality at most $k+1$. Therefore $A \cup \{a'\}$ is strongly ω'-regular as required. ☐

We will also need:

Lemma 3.3. *Suppose that $A \subseteq \mathbb{R}^m$ is t_0-separated, and $t_2 \geq t_0 > 0$. We can write $A = \bigsqcup_{i=0}^{\nu} A_i$, where each A_i is t_2-separated, and where $\nu \in \mathbb{N}$, depends only on m and t_2/t_0.*

Proof. Note that there is a bound, ν, on the cardinality of any subset of A of diameter at most $2t_2$.

First set $A_0 \subseteq A$ be a maximal t_2-separated subset. Now define A_i inductively, by taking A_{i+1} to be a maximal t_2-separated subset of $A \setminus \bigcup_{j \leq i} A_j$. Note that if $a \in A$, then at least one element of the set $N(a; t_2) \cap A$ gets added at each stage until it is filled. But this set has at most ν elements, so the process must terminate after at most ν steps. $\square$

(The above can be seen as an expression of the fact that any graph of degree at most d is $(d+1)$-colourable.)

Proof of Lemma 3.1. We write $A = \bigsqcup_{i=0}^{\nu} A_i$ as given by Lemma 3.3, with $t_2 = 2t_1$.

We will inductively define $[a \mapsto a']$ on the sets $A_{\leq i} = \bigcup_{j \leq i} A_j$, so that $A'_{\leq i} = \{a' \mid a \in A_{\leq i}\}$ are strongly ω_i-regular on scale t_1, where ω_i depends only on i, t_0, t_1, δ, m.

To begin, set $a' = a$ for all $a \in A_0$. Note that A_0 is vacuously strongly ω-regular for all ω, so we can set $\omega_0 = 1$, say.

Suppose that we have inductively defined, $A'_{\leq i}$, and let $a \in A_{i+1}$. Now $A'_{\leq i}$ is strongly ω_i-regular, so by Lemma 3.3, we can find $a' \in \mathbb{R}^m$ with $\rho_E(a, a') \leq \delta$ and with $A'_{\leq i} \cup \{a'\}$ strongly ω_{i+1}-regular, for some $\omega_{i+1} > 0$. We can find such a' for all $a \in A_{i+1}$. Now A_{i+1} is $(2t_1)$-separated, and A'_{i+1} is t_1-separated provided $\delta < t_1/2$. In particular, any subset of $A_{\leq(i+1)}$ of diameter at most t_1 contains at most one element of A'_{i+1}. Therefore, $A'_{\leq(i+1)}$ is strongly ω_{i+1}-regular.

Setting $\omega = \omega_\nu$, we end up with $A' = A'_{\leq \nu}$, which is (strongly) ω-regular as required. $\square$

We remark that the construction of [Br] (for hyperbolic manifolds) offers an alternative approach to this, which could probably be adapted for the above purposes.

4. Smooth maps on simplices

Given $A \subseteq \mathbb{R}^m$, we say that a map, $f|A \longrightarrow \mathbb{R}^n$ is *smooth* if there is an open set, $U \subseteq \mathbb{R}^m$, with $A \subseteq U$, and a smooth function, $F : U \longrightarrow \mathbb{R}^n$ with $F|A = f$. (One can check that a map is smooth if it is smooth in a neighbourhood of every point.)

In the cases of interest here, where A is an m-simplex, or an open subset of $\mathbb{R}^m$, we can write $T_x A = T_x \mathbb{R}^m \equiv \mathbb{R}^m$, for the tangent space at $x \in M$. We then get a well defined derivative map, $d_x f : \mathbb{R}^m \longrightarrow \mathbb{R}^n$ (independent of the extension, F).

We write $\Lambda_0(f) = \max\{||f(x)|| \mid x \in A\}$ and $\Lambda_1(f) = \max\{||d_x f||_o \mid \in A\}$, where $||.||$ denotes the euclidean norm on $\mathbb{R}^m$, and $||.||_o$ denotes the operator norm.

Definition. We say that f is ϵ-*small* if $\Lambda_1(f) \leq \epsilon$.

In other words, $||(d_x f)(v)|| \leq \epsilon ||v||$ for all $x \in A$ and all tangent vectors, $v \in \mathbb{R}^m$.

Note that if A is convex, $\mathrm{diam}(A) \leq d$ and $f(x) = 0$ for some $x \in a$, then $\Lambda_0(f) \leq d\epsilon$.

Let $\Sigma \subseteq \mathbb{R}^m$ be an m-simplex. Let $\Sigma_0, \Sigma_1, \ldots, \Sigma_m$ be the codimension-1 faces of Σ, and write $\partial\Sigma = \bigcup_{i=0}^m \Sigma_i$.

Now let Σ be the regular m-simplex with unit side-lengths, which we refer to as the *standard* m-simplex. The following construction can be found in [L] (see Lemma 16.8 thereof). We make explicit here the control on derivatives.

Lemma 4.1. *Suppose $f : \partial\Sigma \longrightarrow \mathbb{R}^n$ is a map with $f|\Sigma_i$ smooth and ϵ-small for all i. Then there is a $(K\epsilon)$-small smooth map, $\hat{f} : \Sigma \longrightarrow \mathbb{R}^n$, where $K = K(m)$ depends only on m.*

In particular, it follows that f is smooth (on $\partial\Sigma$).

(We also note that we could equivalently replace $\partial\Sigma$ by any union of faces, in the above statement. The equivalence can be seen by using induction over i-skeletons.)

We remark that the construction we give below is canonical.

Proof. Note that, taking each coordinate separately, we can take $n = 1$; that is, f maps to $\mathbb{R}$. For the proof, we will identify Σ (up to a scale factor of $\sqrt{2}$) with a simplex in $\mathbb{R}^{m+1}$. To this end, we are use coordinates, $t = (t_0, \ldots, t_m)$ for $\mathbb{R}^{m+1}$. Let $I_m = \{0, 1, \ldots, m\}$. Given $I \subseteq I_m$, write $I^C = I_m \setminus I$. Given $i \in I_m$ write $R_i = \{t \in \mathbb{R}^{m+1} \mid t_i = 0\}$. Given $I \subseteq I_m$, write $R_I = \bigcap_{i \in I} R_i$ (so $R_i = R_{\{i\}}$).

Given $t \in \mathbb{R}^{m+1}$, write $\bar{t} = \sum_{i=0}^m t_i$. Let $P = \{t \in \mathbb{R}^{m+1} \mid \bar{t} = 1\}$, and write $P_I = P \cap R_I$. Let $\Sigma = P \cap [0, \infty)^{m+1} = P \cap [0, 1]^{m+1}$, and $\Sigma_I = \Sigma \cap R_I$. Thus, $\Sigma_I \cap \Sigma_J = \Sigma_{I \cup J}$. If $i \in I_m$, write $\Sigma_i = \Sigma_{\{i\}}$ and $\Sigma_{\{i\}^C} = \{v_i\}$. In other words, Σ_i is the codimension-1 face of Σ opposite the vertex v_i.

Fix, for the moment, some $i \in I_m$, and let $\mathcal{I}_i = \{I \subseteq I_m \mid i \notin I\}$. Given $I \in \mathcal{I}_i$, we define a linear map $\pi_{I,i} : \mathbb{R}^{m+1} \longrightarrow \mathbb{R}^{m+1}$ as follows. Given $t \in \mathbb{R}^{m+1}$, set $\pi_{I,i}(t) = u$, where $u_j = 0$ if $j \in I$; $u_j = t_j$, if $j \notin I \cup \{i\}$; and where $u_i = t_i + \sum_{j \in I} t_j$. Thus $\bar{u} = \bar{t}$. Also, $\bar{u} \in R_I$. Therefore, if $t \in P$, then $u \in P \cap R_I = P_i$. In particular, $\pi_{I,i} : P \longrightarrow P_i$ is an affine retraction of P to P_i, which restricts to a retraction of Σ to Σ_i. Note that $\pi_{\varnothing,i}|\Sigma$ is the identity. Also if, $i, j \in I_m$ are distinct, and $I \subseteq \{i, j\}^C$, then $\pi_{I,i}|P_j = \pi_{I \cup \{j\},i}|P_j$.

Suppose $f : P \longrightarrow \mathbb{R}$ is any map. Write $f_{I,i} = f \circ \pi_{I,i} : P \longrightarrow \mathbb{R}$ (so $f_{\varnothing,i} = f$). If $j \in \{i\}^C$ and $I \subseteq \{i, j\}^C$, then $f_{I,i}|P_j = f_{I \cup \{i\},i}|P_j$. It follows that $\sum_{I \in \mathcal{I}_i} (-1)^{|I|} f_{I,i} = 0$ (since the terms in I and $I \cup \{j\}$ cancel, and we can partition $\mathcal{I}_i$ into such pairs). Set $\hat{f}_i = -\sum_{I \in \mathcal{I}_i \setminus \{\varnothing\}} (-1)^{|I|} f_{I,i}$. Then $\hat{f}_i|P_j = f_{\varnothing,i}|P_j = f|P_j$.

Note that the definition of $\hat{f}_i|\Sigma$ makes sense if f is only defined on each Σ_i for $j \neq i$. (Since $\pi_{I,i}$ maps to Σ into some such face, if $I \neq \varnothing$.) In this case, $\hat{f}_i| \bigcup_{j \neq i} \Sigma_j = f$. Moreover, if f is smooth on each Σ_j, then $\hat{f}_i$ is smooth. (Note that $f|\Sigma_j$ can be extended independently to a smooth function in a neighbourhood of each Σ_j in P_j, so this gives $\hat{f}$ defined on a neighbourhood of Σ in P by the same formula.) Moreover, $\Lambda_1(\hat{f}_i)$ is bounded above by some fixed multiple of $\Lambda_1(f, \partial\Sigma)$, where $\Lambda_1(f, \partial\Sigma) = \max_{i=0}^m \Lambda_1(f|\Sigma_i)$.)

We now fix $\phi : \mathbb{R} \longrightarrow [0,1]$ to be any smooth function satisfying $\phi(x) = 0$ if $x \leq 1/2(m+1)$ and $\phi(x) > 0$ if $x > 1/2(m+1)$. Given $i \in I_m$, define $\phi_i : P \longrightarrow [0,1]$, by $\phi_i(t) = \phi(t_i)$. We set $\Phi(t) = \sum_{i=0}^m \phi_i(t)$, and $\sigma_i(t) = \phi_i(t)/\Phi(t)$. (Note that $\Phi(t) > 0$ since $t_i > 1/2(m+1)$ for at least one i.) Thus $\sum_{i=0}^m \sigma_i \equiv 1$. In other words $\{\sigma_i\}_i$ is a smooth partition of unity for Σ. Note $\sigma_i|\Sigma_i \equiv 0$, and that $\Lambda_1(\sigma_i)$ is bounded above in terms of m.

Suppose now that $f : \partial\Sigma \longrightarrow \mathbb{R}$ with $f|\Sigma_i$ smooth for each i. Let $\hat{f}_i : \Sigma \longrightarrow \mathbb{R}$ be defined as above, and set $\hat{f} = \sum_{i=0}^m \sigma_i \hat{f}_i$. Again $\hat{f}$ is smooth Σ and $\Lambda_1(\hat{f})$ is bounded above by some multiple of $\Lambda_1(f, \partial\Sigma)$ depending only on m.

We claim that $\hat{f}|\partial\Sigma = f$. For suppose that $t \in \partial\Sigma$. Then $t \in \Sigma_j$ for some $j \in I_m$. We have noted that $\hat{f}_i|P_j = f|P_j$ when $i \neq j$, so $\hat{f}_i(t) = f(t)$ for all $j \neq i$. Also, $\sigma_j(t) = 0$. Therefore, $\hat{f}(t) = \sum_{i=0}^m \sigma_i(t)\hat{f}_i(t) = \sum_{i\neq j} \sigma_i(t)\hat{f}_i(t) = f(t)\sum_{i\neq j} \sigma_i(t) = f(t)\sum_{i=0}^m \sigma_i(t) = f(t)$. □

(We illustrate the above construction when $m = 1$. In this case, $\Sigma = \{(t_0, t_1) \mid t_0, t_1 \geq 0,\ t_0 + t_1 = 1\}$ is an interval. Then $\mathcal{I}_0 = \{\varnothing, \{1\}\}$ and $\pi_{\varnothing,0}(t_0, t_1) = (t_0, t_1)$ and $\pi_{\{1\},0}(t_0, t_1) = (1, 0)$. Similarly, swapping 0 and 1, we get $\pi_{\varnothing,1}(t_0, t_1) = (t_0, t_1)$ and $\pi_{\{0\},1}(t_0, t_1) = (0, 1)$. So $\hat{f}_0(t_0, t_1) = f(1, 0)$ and $\hat{f}_1(t_0, t_1) = f(0, 1)$. Thus $\hat{f}$ is a diffeomorphism of Σ to the interval between $f(1, 0)$ and $f(0, 1)$. Of course, we could instead have taken a linear map, in this particular case.)

Note that, if $f \geq 0$, then $\hat{f} \geq 0$. Also the construction is linear in f. In other words, if $f = \sum_{k=1}^n \lambda_k f_k$, where the λ_k are constant, then $\hat{f} = \sum_{k=1}^n \lambda_k \hat{f}_k$.

We will make use of the following notion. Suppose that $A \subseteq \mathbb{R}^m$ and $f : A \longrightarrow \mathbb{R}^n$ is smooth.

Definition. We say that f is ϵ-*affine* if there is a linear map, $L : \mathbb{R}^m \longrightarrow \mathbb{R}^n$, such that $||d_x f - L||_o \leq \epsilon$ for all $x \in A$.

This is equivalent to saying that there is an affine map, $T : \mathbb{R}^m \longrightarrow \mathbb{R}^m$, such that $f - T$ is ϵ-small. This follows directly from the definitions: take any affine map with linear part (or derivative) L.

Clearly we always postcompose T with a translation, so that $f(x) = T(x)$ for any fixed $x \in A$. Note that if A is convex and has diameter at most d, then this implies that $\Lambda_0(f - T|A) \leq d\epsilon$.

Now let $\Sigma \subseteq \mathbb{R}^m$ be an m-simplex, with vertex set Σ^0, so that $\mathrm{diam}(\Sigma) = \mathrm{diam}(\Sigma^0)$.

For future reference we note the following easily verified observation:

Lemma 4.2. *There is some $\epsilon_0 = \epsilon_0(m)$ such that if $f : \Sigma \longrightarrow \mathbb{R}^n$ is ϵ_0-affine, then* $\mathrm{diam}(f(\Sigma)) \leq \mathrm{diam}(f(\Sigma^0))$.

In fact, for our application, it would be enough to have $\mathrm{diam}(f(\Sigma))$ bounded above by some fixed multiple of $\mathrm{diam}(f(\Sigma^0))$.

We can give another description of ϵ-affine maps as follows.

Suppose first that $S : \mathbb{R}^m \longrightarrow \mathbb{R}^n$ is any affine map. Then $\Lambda_0(S|\Sigma) = \Lambda_0(S|\Sigma^0)$. Moreover, the following is a simple exercise in linear algebra:

Lemma 4.3. *Let Σ be an m-simplex. Suppose* $\operatorname{diam}(\Sigma) \leq d$ *and* $\operatorname{vol}(\Sigma) \geq v > 0$, *and $S : \mathbb{R}^m \longrightarrow \mathbb{R}^n$ is affine. Then $\Lambda_1(S) \leq K_0\Lambda_0(L|\Sigma^0)$, where $K_0 \geq 0$ depends only on m, n, d, v.*

Now, given any map, $f : \Sigma \longrightarrow \mathbb{R}^m$, let $T_f : \mathbb{R}^m \longrightarrow \mathbb{R}^n$ be the unique affine map with $T_f|\Sigma^0 = f|\Sigma^0$.

Lemma 4.4. *Suppose that f is ϵ-affine. Then $\Lambda_1(f - T_f|\Sigma) \leq K_1\epsilon$, where K_1 depends only on m, n, d, v.*

Proof. By definition, we have $\Lambda_1(f - T|\Sigma) \leq \epsilon$ for some affine $T : \mathbb{R}^m \longrightarrow \mathbb{R}^n$. Choosing any $x_0 \in \Sigma^0$, we can suppose that $f(x_0) = T(x_0)$. Therefore, $\Lambda_0(f - T|\Sigma) \leq d\epsilon$, so $\Lambda_0((T - T_f)|\Sigma^0 \leq d\epsilon$, so $\Lambda_1((T - T_f)|\Sigma) \leq K_0 d\epsilon$. Therefore, $\Lambda_1(f - T_f|\Sigma) \leq \Lambda_1(f - T|\Sigma) + \Lambda_1((T - T_f)|\Sigma) \leq \epsilon + K_0 d\epsilon = K_1\epsilon$, where $K_1 = 1 + K_0 d$. □

We note the following result about extending ϵ-affine maps.

Lemma 4.5. *Let Σ be the standard m-simplex. Suppose that $f : \partial\Sigma \longrightarrow \mathbb{R}^n$ is a map with $f|\Sigma_i$ ϵ-affine for all i. Then there is a $(K\epsilon)$-affine map, $F : \Sigma \longrightarrow \mathbb{R}^n$ with $F|\partial\Sigma = f$. Here $K = K(m)$ depends only on m and n.*

Proof. Let $T = T_f$. Then $T|\Sigma_i = T_{(f|\Sigma_i)}$, so $\Lambda_1((f - T_f|\Sigma_i) \leq K_1\epsilon$. Apply Lemma 4.1 to give $G : \Sigma \longrightarrow \mathbb{R}^n$, with $G|\partial\Sigma = f - T$ and $\Lambda_1(G - (f - T)) \leq kK_1\epsilon$. Set $F = G + T$ and $K = kK_1$. □

Now suppose that $m = n$. Let $X \subseteq \mathbb{R}^m$ and suppose that $f : X \longrightarrow \mathbb{R}^m$ is smooth.

Definition. We say that f is *ϵ-congruent* (or an *ϵ-congruence*) if there is an orthogonal linear map, $L : \mathbb{R}^m \longrightarrow \mathbb{R}^m$, such that $||d_x f - L||_o \leq \epsilon$ for all $x \in X$. We say that f is *ϵ-translating* (or an *ϵ-translation*) if $||d_x f - I||_o \leq \epsilon$ for all $x \in X$, where I is the identity matrix.

In other words, f is ϵ-congruent if there is an isometry $T : \mathbb{R}^m \longrightarrow \mathbb{R}^m$ such that $f - T|X$ is ϵ-small. It is ϵ-translating if $f - \iota$ is ϵ-small, where $\iota : A \longrightarrow \mathbb{R}^m$ is the inclusion map. Clearly any ϵ-translation is an ϵ-congruence. Conversely, we can postcompose any ϵ-congruence with an isometry of $\mathbb{R}^m$ so that it becomes an ϵ-translation.

Now let $\Sigma \subseteq \mathbb{R}^m$ be an m-simplex.

Lemma 4.6. *Given $\mu > 1$, there is some $\epsilon > 0$ such that if $f : \Sigma \longrightarrow \mathbb{R}^m$ is an ϵ-congruence, then f is a bilipschitz homeomorphism to $f(\Sigma)$.*

Proof. After postcomposing f by an isometry of $\mathbb{R}^m$, we can assume that it is ϵ-translating. Let $x, y \in \Sigma$, write $l = ||x - y||$, and let $\alpha : [0, l] \longrightarrow \Sigma$ be the unit-speed geodesic from x to y. Let $\beta = f \circ \alpha : [0, l] \longrightarrow \mathbb{R}^m$. Let $v = (y - x)/l = \alpha'(t)$ for all $t \in [0, l]$. Then $||\beta'(t) - v|| \leq \epsilon$. Then $||y - x - vl|| = || \int_0^l \beta'(t) - v \, dt|| \leq \epsilon l$, so $(1 - \epsilon)l \leq ||f(y) - f(x)|| \leq (1 + \epsilon)l$. It follows that f is $(1/(1 - \epsilon))$-bilipschitz for any $\epsilon < 1$. □

Lemma 4.7. *Suppose that* $\mathrm{diam}(\Sigma) \leq d$ *and* $\mathrm{vol}(\Sigma) \geq v > 0$. *Suppose that* $f : \Sigma \longrightarrow \mathbb{R}^m$ *is ϵ-affine, and that* $\rho_E(x, f(x)) \leq \epsilon$ *for all* $x \in \Sigma^0$. *Then f is* $(K_2\epsilon)$*-translating, where K_2 depends only m, d, v.*

Proof. By assumption, we have $\Lambda_0(f - \iota|\Sigma^0) \leq \epsilon$. By Lemma 4.4, we have $\Lambda_1(f - T_f) \leq K_1\epsilon$. Also, $\Lambda_0((T_f - \iota)|\Sigma^0) \leq \epsilon$, so by Lemma 4.3, $\Lambda_1(T_f - \iota) \leq K_0\epsilon$. Thus $\Lambda_1(f - \iota) \leq \Lambda_1(f - T_f) + \Lambda_1(T_f - \iota) \leq K_1\epsilon + K_0\epsilon = K_2\epsilon$, where $K_2 = K_1 + K_0$. □

We note that being almost affine is preserved under postcomposition with an almost congruent map:

Lemma 4.8. *Suppose that* $f : \Sigma \longrightarrow \mathbb{R}^n$ *with* $||d_x f - L|| \leq \epsilon$ *for all* $x \in \Sigma$, *where L is a fixed matrix with* $||L||_o \leq N$. *Suppose that* $g : f(\Sigma) \longrightarrow \mathbb{R}^n$ *is ϵ-congruent. Then* $g \circ f : \Sigma \longrightarrow \mathbb{R}^n$ *is* $(K_3\epsilon)$*-affine, where K_3 depends only on m, n, N.*

Proof. This is just an application of the chain rule. Let $v \in \mathbb{R}^m$. If $x \in \Sigma$, then $||(d_x f)v - Lv|| \leq \epsilon||v||$, and $||(d_{fx} g)(d_x f)v - (d_x f)v|| \leq \epsilon||(d_x f)v||$. So $||d_x(g \circ f)v - Lv|| \leq (1 + ||d_x f||)\epsilon||v|| \leq (1 + (N + \epsilon))\epsilon||v||$. So if $\epsilon \leq 1$, say, then $||d_x(g \circ f) - L||_o \leq K_3\epsilon$, where $K_3 = 2 + N$. □

Now suppose that Π is a simplicial complex with some euclidean metric ρ_Π.

Let $f : \Pi \longrightarrow \mathbb{R}^n$. We say that f is *smooth* (respectively *ϵ-affine*) if its restriction to each simplex of Π is smooth (respectively ϵ-affine).

The following says that a small perturbation of a Delaunay triangulation is still a triangulation.

Lemma 4.9. *Given* $\omega, t > 0$ *and* $\mu > 1$, *there is some* $\epsilon_1 > 0$ *with the following property. Suppose that* $A \subseteq \mathbb{R}^m$ *is t-dense and ω-regular on scale t. Let $\Delta(A)$ be the Delaunay triangulation with induced metric ρ_Δ (isometric to $\mathbb{R}^m$). Suppose that* $f : \Delta(A) \longrightarrow \mathbb{R}^m$ *is ϵ_1-affine, and that* $\rho_E(a, f(a)) \leq \epsilon_1$ *for all* $a \in A$. *Then f is a μ-bilipschitz triangulation of $\mathbb{R}^m$.*

Proof. Note that each m-simplex, Σ, of $\Delta(A)$ has diameter at most t and volume at least $\omega > 0$. Let $\epsilon > 0$ be the constant given by Lemma 4.6 given μ. Let K_2 be the constant of Lemma 4.7, given $n = m$, $v = \omega$ and $d = t$. Let $\epsilon_1 = \epsilon/K_2$. By assumption, $\Lambda_1(a, f(a)) \leq \epsilon$, for all $a \in \Sigma^0 \subseteq A$, and $f|\Sigma$ is ϵ_1-affine. Therefore, by Lemma 4.7, $f|\Sigma$ is ϵ-translating. Therefore, by Lemma 4.6, $f|\Sigma$ is bilipschitz onto $f(\Sigma)$.

In fact, the same proof as Lemma 4.6 shows directly that $f : \Delta(A) \longrightarrow \mathbb{R}^m$ is μ-bilipschitz. Let $x, y \in \Delta(A) \cong \mathbb{R}^m$, and let α be the euclidean geodesic joining them. The path $\beta = f \circ \alpha$ is continuous, and smooth for all but finitely many parameters, and so the argument goes through as before. □

Of course, Lemma 4.9 is not quite what we need for the general case, but it serves to illustrate the argument (see Section 7).

5. Curvature bounds

We make a digression into a general result of riemannian geometry.

Note that the metric on any riemannian manifold of bounded geometry can be rescaled to make the curvature arbitrarily small and the indectivity radius arbitrarily large. By the Rauch comparison theorem [Ch], this means that the exponential map on from balls of any given fixed radius will be bilipschitz with constant arbitrarily close to 1. We will also need a more subtle fact, namely that the transition functions can be taken to be ζ-congruences, in the sense defined in Section 4 for some fixed (arbitrarily small) ζ. We give the background to this in this section. Accounts of this can be found in [BuK] and Section 3.3 of [DyVW].

First, we introduce the following notation. Given unit vectors, v_1, v_2 in a (finite dimensional) inner-product space, we write $v_1 \sim_\zeta v_2$ to mean that $||v_1 - v_2|| \leq \zeta$. Given two linear maps, L_1, L_2, between two such normed space, we write $L_1 \sim_\zeta L_2$ to mean that $||L_1 - L_2||_o \leq \zeta$. In other words, $L_1 v \sim_\zeta L_2 v$ for all unit vectors, v, in the domain.

Let P be a geodesically convex smooth riemannian manifold, with metric ρ_P. (Typically it will be an open subset of M or of $\mathbb{R}^m$.) Given $x, y \in P$, let $\alpha^P_{xy} : [0, \rho_P(x, y)] \longrightarrow P$ be the unique unit-speed geodesic from x to y. Let $\tau^P_{xy} : T_x P \longrightarrow T_y P$ be the isometry of tangent spaces given by parallel translation along α^P_{xy}. Note that $\tau^P_{xx} = 1$, and $\tau^P_{yx} = (\tau^P_{xy})^{-1}$.

Suppose that $\mathrm{diam}(P) \leq r$, and that all sectional curvatures of P are bounded in norm by some $\kappa > 0$ (which we think of as being small).

Let β be any piecewise smooth path in P from x to y, of length at most l. Let $\tau : T_x P \longrightarrow T_y P$ be parallel transport along β. There is some fixed constant, $h > 0$, such that $\tau \sim_{k\kappa l} \tau^P_{xy}$. In particular it follows that if $x, y, z \in P$, then $\tau^P_{xz} \sim_{\zeta_0} \tau^P_{yz} \circ \tau^P_{xy}$, where $\zeta_0 = 2k\kappa r$ (which we can again take to be small).

Now suppose that Q is another such manifold, and that $\phi : P \longrightarrow Q$ is a diffeomorphism. Given $x \in P$, we have a linear isomorphism $d_x\phi : T_x P \longrightarrow T_{\phi x} Q$.

We will assume:

(C1): There is some $\xi \geq 1$ such that for all $x \in P$ and $y \in Q$, $||d_x\phi||_o \leq \xi$ and $||d_y(\phi^{-1})||_o \leq \xi$.

In other words, the map f is ξ-bilipschitz. Note that this implies that there is an orthogonal matrix, O, such that $d_x f \sim_{\zeta_1} O$ for some fixed ζ_1, depending on ξ and m, and with $qz_1 \to 0$ as $\xi \to 0.$, and with $qz_1 \to 0$ as $\xi \to 0$.

We also assume that the derivative approximately agrees with parallel transport in the following sense:

(C2): There is some $\zeta > 0$ such that if $x, y \in P$, then $\tau^Q_{\phi x, \phi y} \circ d_x\phi \sim_\zeta d_y\phi \circ \tau^P_{xy}$.

Note that swapping P and Q, and replacing ϕ with ϕ^{-1}, all the above maps are inverted. In particular, it follows that ϕ^{-1} also satisfies (C2), but with ζ replaced by $\zeta||\xi||^2$. We will assume henceforth that $||\xi||^2 \leq 2$, say.

The above properties are also respected by composition in the following sense. Suppose P, Q, R are such manifolds, and $\phi : P \longrightarrow Q$ and $\psi : Q \longrightarrow R$ both satisfy (C1) and (C2). Clearly, $\psi \circ \phi$ is ξ^2-bilipschitz. Moreover, suppose $x, y \in P$. Then $\tau^Q_{\phi x, \phi y} \circ d_x\phi \sim_\zeta d_y\phi \circ \tau^P_{xy}$, and $\tau^R_{\psi\phi x, \psi\phi y} \circ d_{\phi x}\psi \sim_\zeta d_{\phi y}\psi \circ \tau^Q_{\phi x, \phi y}$. So

$$\tau^R_{\psi\phi x, \psi\phi y} \circ d_{\phi x}\psi \circ d_x\phi \sim_{2\zeta} d_{\phi y}\psi \circ \tau^Q_{\phi x, \phi y} \circ d_x\phi \sim_{2\zeta} d_{\phi y}\psi \circ d_y\phi \circ \tau^P_{xy},$$

and so $\tau^R_{\psi\phi x, \psi\phi y} \circ d_x(\psi \circ \phi) \sim_{4\zeta} d_y(\psi \circ \phi) \circ \tau^P_{xy}$. In other words, $\psi \circ \phi$ satisfies (C2) with ζ replaced by 4ζ.

If (C2) holds for some fixed x and all $y \in P$, then it holds for all $y, z \in P$, modulo adjusting the constant. More precisely, $\tau^Q_{\phi y, \phi z} \circ d_y\phi \sim_{\zeta + 4\zeta_0} d_z\phi \circ \tau^P_{yz}$.

Now suppose that $P, Q \subseteq \mathbb{R}^m$ are convex and equipped with the euclidean metrics. We can identify all tangent spaces with $\mathbb{R}^m$. Then parallel transport is the identity map. In particular, it follows that if $f : P \longrightarrow Q$ satisfies (C2), then for all $x, y \in P$, we have $d_x\phi \sim_\zeta d_y\phi$.

Now let M be any manifold with sectional curvatures bounded by κ, and convexity radius at least χ at $a \in M$. Let $\Omega = N(o; \chi) \subseteq \mathbb{R}^m$ and $W = N(a; \chi) \subseteq M$. The exponential map, $\theta : \Omega \longrightarrow W$ is a diffeomorphism, and W is geodesically convex. Moreover (by the Rauch comparison theorem) it is ξ-bilipschitz, where ξ depends only on κ. Indeed, we can make ξ arbitrarily close to 1, by assuming κ small enough. We also claim (cf. [BuK]):

Lemma 5.1. *Given any $\zeta > 0$, if κ is sufficiently small in relation to χ, then θ satisfies (C2) for this ζ.*

Proof. Let $x \in \Omega \setminus \{o\}$, and let $b = \theta(x)$. Let $l = \rho_E(o, x) = \rho_M(a, b)$. Let $\alpha = \alpha^\Omega_{ox} : [0, l] \longrightarrow \Omega$ be the geodesic from o to x, and let $\beta = \theta \circ \alpha$. By the definition of the exponential map, β is the geodesic from a to b in $W \subseteq M$. Let $v_0 \in \mathbb{R}^m$ be the fixed unit vector, $\alpha'(0) = (x - o)/l$ for $t \in [0, l]$. Then $d_{\alpha(t)}v_0 = \beta'(t)$ is the tangent vector to β in W. This is parallel transport, and so $(d_b\theta)(v_0) = \tau^W_{ab} \circ (d_o\theta)v_0$.

Now let $v \in \mathbb{R}^m$ be a unit vector orthogonal to v_0. Let $V = (d_o\theta)(v) \in T_aW$. Note that the vector field, $[v \mapsto tv/l]$ where $tv/l \in \mathbb{R}^m \equiv T_{\alpha(t)}\mathbb{R}^m$, is a Jacobi field along α in $\Omega \subseteq \mathbb{R}^m$. Let $X(t) = (d_{\alpha(t)}\theta)(tv/l)$, so that $X(l) = (d_x\theta)(v)$. This is an orthogonal vector field along β. In fact, by the definition of the exponential map, it is again a Jacobi field along β. It therefore satisfies the Jacobi equation:

$\frac{D^2}{\partial t^2}X(t) = R(\beta'(t), X(t))\beta'(t)$, where $\frac{D}{\partial t}$ denotes the covariant derivative in the direction $\beta'(t)$, and where R is the Riemann curvature tensor. From the linearity of R, the norm, $|R|$, is bounded by a fixed multiple of the norms of the sectional curvatures. In fact (from the symmetries of R) one can show that $|R| \leq 4\kappa/3$ (see **??** for example). (The factor of $4/3$ is not important to our discussion.) It follows that $||\frac{D^2}{\partial t^2}X(t)|| \leq \frac{4}{3}\kappa||X(t)||$. Note also that at $t = 0$, we have $\frac{D}{\partial t}X(t) = V/l$.

Now let $Y_0(t)$ be parallel transport of the vector V along β. In other words, $\frac{D}{\partial t}Y_0(t) = 0$. By definition, $Y_0(0) = V$ and $Y_0(l) = \tau^W_{ab}V$. Let $Y(t) = tY_0(t)/l$. In particular, $Y(l) = \tau^W_{ab}V = \tau^W_{ab} \circ d_o\theta(v_0)$. This is another orthogonal vector field along β. Now $\frac{D}{\partial t}Y(t) = Y_0(t)/l$. In particular, at $t = 0$, we have $\frac{D}{\partial t}Y(t) = V/l$. Also, $\frac{D^2}{\partial t^2}Y(t) = 0$ for all t.

Let $Z = X - Y$. We have $Z(0) = 0$ and $\frac{D}{\partial t}Z(t) = 0$ at $t = 0$. Also $||\frac{D^2}{\partial t^2}Z(t)|| \leq \frac{4}{3}\kappa||Z(t)||$ for all t, so $||Z(l)|| \leq \frac{2}{3}\kappa l^2$. Therefore, by choosing κ hence ζ_2 small enough we can arrange that $||Z(l)||$ is arbitrarily small. In other words, $||\tau^W_{ab} \circ d_o\theta(v) - d_b\theta(v)||$ is arbitrarily small.

Recall that $\tau^W_{ab} \circ d_o\theta(v_0) = d_b\theta(v_0)$. Since v is was an arbitrary unit vector orthogonal to v_0, it now follows that $||\tau^W_{ab} \circ d_o\theta - d_b\theta||_o$ is arbitrarily small. Since τ^Ω_{ox} is just the identity map (identifying $T_x\mathbb{R}^m = \mathbb{R}^m$), it follows that θ satisfies (C2) as required. □

Lemma 5.2. *Given $\epsilon > 0$, there is some $\kappa > 0$ with the following property. Suppose that M is (κ, χ)-bounded, and $a, b \in M$. Let $W_a = N(a; \chi)$ and $W_b = N(b; \chi)$. Let $\theta_b^{-1} \circ \theta_a|\theta_a^{-1}(W_a \cap W_b) : \theta_a^{-1}(W_a \cap W_b) \longrightarrow \theta_b^{-1}(W_a \cap W_b)$ be the transition map. Then $\theta_b^{-1} \circ \theta_a$ is an ϵ-congruence. Since parallel translation is the identity map on $\mathbb{R}^m$, this implies that $\theta_b^{-1} \circ \theta_a$ is an ϵ-congruence, where $\epsilon > 0$ is arbitrarily small.*

Proof. First, choose $\kappa > 0$ so that θ_a and θ_b are ξ-bilipschitz for ξ arbitrarily close to 1. Choose any $x \in \theta_a^{-1}(W_a \cap W_b)$ (assuming this is non-empty). Then $d_x(\theta_b^{-1} \circ \theta_a)$ is ξ^2-bilipschitz. By Lemma 5.1, θ_a and θ_b both satisfy (C2), and so therefore does $\theta_b^{-1} \circ \theta_a$. □

A variation on these results also holds for manifolds with boundary, as we discuss in Section 8.

6. Flatness

The results of Section 5 tell us that a bounded geometry manifold is almost euclidean on sufficiently small scales. We summarise this in the notion of "flatness", and describe some of its consequences. To simplify notation, we will rescale the metric so that it is flat on some (arbitrarily) fixed scale.

Let M be a riemannian manifold. Given $\alpha \in M$ we choose an identification of T_xM with $\mathbb{R}^m$, and let $\theta_\alpha : \mathbb{R}^m \longrightarrow M$ be the exponential map. We fix some $R > 0$, and assume that the convexity radius of M is at least R everywhere. (For example, if the injectivity radius is at least $2R$.) Let $\Omega = N(0;R) \subseteq \mathbb{R}^m$ be the euclidean R-ball about the origin. Given $\alpha \in M$, the map $\theta_\alpha|\Omega$ is a diffeomorphism to the geodesically convex set, $W_\alpha = N(\alpha;R) \subseteq M$. Let $\phi_\alpha = \theta_\alpha^{-1} : W_\alpha \longrightarrow \Omega$ be the inverse ("logarithm") map. Then the family of maps $(\phi_\alpha)_{\alpha \in M}$ is a smooth atlas for M.

Definition. We say that M is (ξ,η)-*flat* (at the scale R) if each of the charts, $\phi_\alpha : W_\alpha \longrightarrow \Omega$, is ξ-bilipschitz, and if each of the transition functions, $\phi_\beta \circ \phi_\alpha^{-1}|\phi_\alpha(W_\alpha \cap W_\beta) : \phi_\alpha(W_\alpha \cap W_\beta) \longrightarrow \phi_\beta(W_\alpha \cap W_\beta)$, is an η-congruence.

We will generally write $\phi_\beta \circ \phi_\alpha^{-1}$ for the transition function, where the restriction to $\phi_\alpha(W_\alpha \cap W_\beta)$ is implicitly assumed. Clearly, each of the transition functions is ξ^2-bilipschitz.

Suppose we are given $R > 0$, $\xi > 1$, $\eta > 0$ and $m \in \mathbb{N}$. If M is (r,χ)-bounded, then we can rescale the metric by a constant factor, depending only on R, r, χ, so that it is (ξ,η)-flat (at scale R). This follows by Lemma 5.2. To simplify notation, we fix R (say $R = 1000$).

This means that $\xi > 1$ is sufficiently close to 1, and $\eta > 0$ sufficiently small, in relation to the various parameters to make the arguments go through. We will express this loosely by saying that M is "sufficiently flat".

We also fix some $R_0 \leq R/10$ (say, $R_0 = 100$), and write $W_\alpha^0 = N(\alpha;R_0)$, and $\Omega^0 = N(o;R_0) \subseteq \mathbb{R}^m$. Thus $\phi_\alpha(W_\alpha^0) = \Omega^0$.

Given the convexity radius bound, if $x, y \in M$ with $\rho_M(x,y) \leq R$, then x, y are connected by a unique geodesic, $[x,y] \subseteq M$ of length $\rho_M(x,y)$. In particular, if $\alpha \in M$, and $x, y \in W_\alpha^0$, then $[x,y] \subseteq W_\alpha$.

To simplify terminology, we will write $F = (\xi,\eta)$ where it is assumed that $\xi > 1$ and $\eta > 0$. We then refer to the manifold M as being "F-flat".

Lemma 6.1. *Given $\epsilon > 0$, there is some F such that the following holds. Suppose that M is F-flat, and $f : \Sigma \longrightarrow M$ is smooth. Suppose that $\phi_\alpha \circ f : \Sigma \longrightarrow \mathbb{R}^m$ is ϵ-affine for some $\alpha \in M$ with $f(\Sigma) \subseteq W_\alpha$. Then f is (2ϵ)-affine.*

Proof. In other words, if $\beta \in M$ with $f(\Sigma) \subseteq W_\beta$, we need that $\phi_\beta \circ f$ is (2ϵ)-affine. Given the fact that the transition function, $\phi_\beta \circ \phi_\alpha^{-1}$ is η-congruent, this follows from Lemma 4.8, and assuming that $\eta < \epsilon/K_3$, where K_3 is the constant featuring there. $\square$

Note that, by Lemma 4.8, if η is small enough in relation to ϵ (namely that $\eta \leq \epsilon/2K_3$), then if $\phi_\alpha \circ f : \Sigma \longrightarrow \mathbb{R}^n$ is $(\epsilon/2)$-affine for some such α, then f will be ϵ-affine.

If Π is any simplicial complex, we say that a map $f : \Pi \longrightarrow M$ is ϵ-*affine* if $f|\Sigma$ is ϵ-affine for all $\Sigma \in \mathcal{C}_{\rm real}(\Pi)$. Here, each simplex is given the structure of the standard simplex.

We say that f is *fine* if for all $\Sigma \in \mathcal{C}_{\rm real}(\Pi)$, there is some $\alpha \in M$ with $f(\Sigma) \subseteq W_\alpha$. (Of course, it's enough that $\operatorname{diam}(f(\Sigma)) \leq R$.) By Lemma 4.2, we can take ϵ small enough (depending on m) such that if f is fine and ϵ-affine, then $\text{diam}(f(\Sigma)) \leq \text{diam}(f(\Sigma^0))$ for all $\Sigma \in \mathcal{C}_{\text{real}}(\Pi)$.

From the above, we see that if f is fine and ϵ-affine, and $\alpha \in M$, $\Sigma \in \mathcal{C}_{\text{real}}(\Pi)$ with $f(\Sigma^0) \subseteq W_\alpha^0$, then $f(\Sigma) \subseteq W_\alpha$.

Let $A \subseteq M$ be locally finite. Let $\Pi = \Pi(A, R_0)$ be the complex defined in Section 2 (for an arbitrary metric space). We can extend the inclusion of $A = \Pi^0$ into M to a fine map, $f : \Pi^1 \longrightarrow M$, of the 1-skeleton, such that each 1-simplex gets mapped linearly to the unique shortest geodesic segment connecting its endpoints. If M is sufficiently flat, then this map will be ϵ-affine for arbitrarily small $\epsilon > 0$.

Lemma 6.2. *Given $\epsilon > 0$, there is some $F = F(\epsilon, m)$ such that if M is F-flat, then there is a fine ϵ-affine extension, $f : \Pi^m \longrightarrow M$.*

Proof. Let $K = K(m)$ be the constant of Lemma 4.5. Let $\epsilon_i = \epsilon/(2K)^{m-i}$. We claim inductively on $i = 1, \ldots, m$, that, provided M is sufficiently flat, there is a fine ϵ_i-affine extension $f : \Pi^i \longrightarrow M$.

To begin, we can suppose, by the earlier discussion, that $f : \Pi^1 \longrightarrow M$ as already defined is ϵ_1-affine.

Suppose then that we have defined f on Π^i. Let $\Sigma \in \mathcal{C}^{i+1}_{\text{real}}(\Pi)$. Now (by the definition of Π) we have $\text{diam}(f(\Sigma^0)) \leq R_0$. Thus, there is some α with $f(\Sigma^0) \subseteq W^0_\alpha$. Let Σ_0 be an i-face of Σ. Since $f(\Sigma_0^0) \subseteq W_\alpha^0$, we can assume (if M is sufficiently flat) that $f(\Sigma_0) \subseteq W_\alpha$. By the induction hypothesis, $f|\Sigma_0$ is ϵ_i-affine. By definition, this means that $\phi_\alpha \circ f|\Sigma_0 : \Sigma_0 \longrightarrow \Omega \subseteq \mathbb{R}^i$ is ϵ_i-affine. Since this holds for all i-faces, we have that $\phi_\alpha \circ f|\partial\Sigma = \partial\Sigma \longrightarrow \mathbb{R}^m$ is ϵ_i-affine. By Lemma 4.5, it has a $(K\epsilon_i)$-affine extension $g : \Sigma \longrightarrow \mathbb{R}^m$. Since $g(\Sigma^0) = \phi_\alpha \circ f(\Sigma^0) \subseteq \Omega^0$, we can assume that $g(\Sigma) \subseteq \Omega$. We can therefore extend f over Σ by setting it equal to $\phi_\alpha^{-1} \circ g$. By Lemma 6.1, this extension is $(2K\epsilon_i)$-affine, provided M is sufficiently flat. Doing this for all $\Sigma \in \mathcal{C}^{i+1}_{\text{real}}(\Pi)$, we get an extension $f : \Pi^{i+1} \longrightarrow M$. This is fine, and by an earlier observation, we can assume it to be $(2K\epsilon_i)$-affine. In other words, f is ϵ_{i+1}-affine as required.

Given that $\epsilon_m = \epsilon$, we eventually end up with a fine ϵ-affine map $f : \Pi^m \longrightarrow M$ as required. □

If we assume that A is t-separated for some $t > 0$, then there is a bound, depending on t and m, on the dimension of Π. In this case, the argument would give us a fine ϵ-affine map $f : \Pi \longrightarrow M$, provided we allow ξ to depend also on t.

7. Construction of the triangulation

Let M be F-flat, where $F = (\xi, \eta)$. For each $\alpha \in M$, we have a ξ-bilipschitz chart $\phi_\alpha : W_\alpha \longrightarrow \mathbb{R}^m$. Recall that $W_\alpha = N(\alpha; R)$ and that $W_\alpha^0 = N(\alpha, R_0)$, where $R = 1000$ and $R_0 = R/10 = 100$, say. By definition, the transition functions are η-congruent. We will assume that $\xi \leq 2$.

Lemma 7.1. *Given $\delta > 0$, there is some F such that the following holds. Suppose that M is F-flat. Let $\alpha, \beta \in M$. Then after postcomposting β by an isometry of $\mathbb{R}^m$, we can suppose that $\rho_E(\phi_\alpha x, \phi_\beta x) \leq \delta$, for all $x \in W_\alpha \cap W_\beta$.*

Proof. We can suppose that $W_\alpha \cap W_\beta \neq \emptyset$. After postcomposing ϕ_β with an isometry, we can assume that $\phi_\beta \circ \phi_\alpha^{-1}$ is (δ/R)-translating, and that $\phi_\beta \circ \phi_\alpha^{-1}(o) = o$. Since $\phi_\alpha(W_\alpha \cap W_\beta) \subseteq \Omega = N(o; R)$, it follows that $\phi_\beta \circ \phi_\alpha^{-1}$ moves each point a distance at most δ. □

Given $A \subseteq M$, $i \in \mathbb{N}$ and $\alpha \in M$, write $\mathcal{C}^i_\alpha(A)$ for the set of $B \subseteq A \cap W^0_\alpha$ with $|B| = i + 1$. Write $\mathcal{C}^i(A) = \bigcup_{\alpha \in M} \mathcal{C}^i_\alpha(A)$. Clearly, if $B \subseteq A$ with $|B| = i + 1$ and $\operatorname{diam}(B) \leq R_0$, then $B \in \mathcal{C}(A)$. Also if $B \in \mathcal{C}(A)$ then $\operatorname{diam}(B) \leq 2R_0$.

Definition. Given $\omega > 0$ and $A \subseteq M$, we say that A is *ω-regular* if:
(R1): for all $\alpha \in M$ and $B \in \mathcal{C}^m_\alpha(A)$, $V_m(\phi_\alpha B) \geq \omega$, and
(R2): for all $\alpha \in M$ and $C \in \mathcal{C}^{m+1}_\alpha(A)$, $Q_m(\phi_\alpha C) \geq \omega$.

Lemma 7.2. *Given $\omega > 0$ there is some F such that if M is F-flat, then the following holds. Suppose $B \subseteq \mathcal{C}^k_\alpha(A) \cap \mathcal{C}^k_\beta(A)$, with $V_k(\phi_\alpha B) \geq \omega$ for $k \leq m$. Then $V_k(\phi_\beta B) \geq \omega/2$. Similarly, if $C \subseteq \mathcal{C}^{m+1}_\alpha(A) \cap \mathcal{C}^{m+1}_\beta(A)$ with $Q_m(\phi_\alpha C) \geq \omega$, then with $Q_m(\phi_\beta C) \geq \omega$.*

Proof. This follows easily from Lemma 7.1, given the continuity of the functions V_m and Q_m. □

For the proof below, we note that we can define "strongly ω-regular" allowing for lower dimensional simplices, similarly as in Section 3. A similar discussion applies.

Now fix some $R_1 = R_0/10$, so that $R_1 = 10 = R_0/10$.

Definition. A *net* in M is an R_1-separated $(6R_1)$-dense subset.

(Here the factor "6" could in principle be replaced by any number bigger than 2.)

Lemma 7.3. *There exist $\omega > 0$, F, depending only on m such that if M is F-flat, then it admits an ω-regular net.*

Proof. We follow a similar argument to that of Lemma 3.1.

Let $A \subseteq M$ be any $2R_1$-separated $4R_1$-dense subset. (Take any maximal $2R_1$-separated subset.)

We will construct another set, $A' \subseteq M$, which is ω-regular, together with a bijection $[a \mapsto a'] : A \longrightarrow A'$, for all $a \in A$. Note that this implies that A' is a net.

We first write $A = \bigsqcup_{i=0}^{\nu} A_i$, where $\nu \in \mathbb{N}$ depends only on m, and where each A_i is $(4R_0)$-separated. (This follows as with Lemma 3.3: Given that $\xi \leq 2$, there is a bound, depending only on m, on the cardinality of any subset of A of diameter at most $8R_0$, say.)

We write $A_{\leq i} = \bigcup_{j \leq i} A_j$ and $A'_{\leq i} = \bigcup_{j \leq i} A'_j$.

We claim inductively that we can construct A'_i so that $A'_{\leq i}$ is strongly ω_i-regular, where $\omega_i > 0$ depends only on i.

We begin by setting $a' = a$ for all $a \in A_0$. It follows vacuously that $A'_0 = A_0$ is strongly ω-regular for all $\omega > 0$, so we can set $\omega_0 = 1$, say.

Now suppose that we have constructed A'_i so that $A'_{\leq i}$ is strongly ω_i-regular.

Let $a \in A_{i+1}$. We now apply a similar argument to that of Lemma 3.2 to $\phi_a(W^1_a \cap A)$, where $W^1_a = N(a; 3R_0)$. Suppose that $B \subseteq \mathcal{C}^p_a(A'_{\leq i})$. Then $V_p(\phi_a B) \geq \omega_i$ (by the inductive assumption). We can now find $a' \in M$, with $\rho_M(a, a') \leq r$, so that if $B \in \mathcal{C}^p(A_{\leq i})$ and $B \subseteq W^1_a$, then $V_{p+1}(\phi_a(B \cup \{a'\})) \geq \omega'$, where $\omega' > 0$ depends only on ω_i and i, hence ultimately, only in i. Similarly, if $C \in \mathcal{C}^m(A_{\leq i})$ with $C \subseteq W^1_a$, then $Q_m(\phi_a(C \cup \{a'\})) \geq \omega'$.

We do this for all $a \in A_{i+1}$, and thereby construct A'_{i+1}. Note that $A'_{i=1}$ is $(3R_0)$-separated. We claim that $A'_{\leq(i+1)}$ is ω_{i+1}-regular, where $\omega_{i+1} = \min\{\omega_i, \omega'/2\}$, thereby proving the inductive step.

To see this suppose $B \in \mathcal{C}^p(A'_{\leq(i+1)})$, that is, $B \subseteq A'_{\leq(i+1)} \cap W^0_\alpha$ for some $\alpha \in M$. Now since A'_{i+1} is $(3R_0)$-separated, B can contain at most one point of A'_{i+1}. If $B \cap A'_{i+1} = \emptyset$, then $B \in \mathcal{C}^p(A'_{\leq i})$, and so $V_p(\phi_\alpha B) \geq \omega_i \geq \omega_{i+1}$ by the inductive hypothesis. If $B \cap A_{\leq i} = \{a'\}$ with $a \in A_{i+1}$, then $B \subseteq W^1_a$ (since $\operatorname{diam} B \leq 2R_0$), and so $V_{p+1}(\phi_a(B \cup \{a'\})) \geq \omega'$, by the construction of A'_{i+1}. Therefore, by Lemma 7.2, provided M is sufficiently flat, we see that $V_{p+1}(\phi_\alpha(B \cup \{a'\})) \geq \omega'/2 = \omega_{i+1}$.

The argument works similarly for $C \in \mathcal{C}^{p+1}(A'_{\leq(i+1)})$, thereby proving the claim.

We finally end up with $A' = A'_\nu$, and we set $\omega = \omega_\nu$. □

Remark: it is easily seen that any ω-regular r-dense subset is necessarily s-separated, where s depends only on ω, r and m. However, the fact that our sets are uniformly separated is an immediate consequence of the construction.

Lemma 7.4. *Given $\omega > 0$, there is some F such that if M is F-flat, then the following holds. Suppose that $A \subseteq M$ is an ω-regular net, and $B \subseteq A \cap W^0_\alpha \cap W^0_\beta$ for $\alpha, \beta \in M$. Then $B \in \mathcal{C}(\Delta(\phi_\alpha(A \cap W^0_\alpha)))$ if and only if $B \in \mathcal{C}(\Delta(\phi_\beta(A \cap W^0_\beta)))$.*

Proof. Let $\delta > 0$ be as given by Lemma 2.3 given ω and $t = R$. The construction of the Delaunay triangulation is invariant under isometries of $\mathbb{R}^m$, so by Lemma 7.1, we can assume that $\rho_E(\phi_\alpha x, \phi_\beta x) \leq \delta$ for all $x \in W_\alpha \cap W_\beta$. We now use Lemma 2.3, together with the observation that the Delaunay triangulation is determined locally, as discussed at the end of Section 2. □

Now, let $\Delta(A)$ be the simplicial complex defined by letting $B \in \mathcal{C}(\Delta(A))$ whenever $\phi_\alpha B \in \mathcal{C}(\Delta(\phi_\alpha(A \cap W^0_\alpha)))$ for some, hence any, $\alpha \in M$ with $B \subseteq A \cap W^0_\alpha$.

Note that any such B has diameter at most $2R_0$, therefore $\Delta(A)$ is a subcomplex of $\Pi = \Pi(A, R_0)$, as defined in Section 6. In fact, it lies in the m-skeleton, Π^m, thereof.

Given $\epsilon > 0$, we can assume that M is sufficiently flat so that we have an ϵ-affine map, $f : \Pi \longrightarrow M$, as given by Lemma 6.2. This restricts to a map $f : \Delta(A) \longrightarrow M$.

Lemma 7.5. *$f : \Delta(A) \longrightarrow M$ is a triangulation.*

Proof. By construction and Lemma 7.4, f coincides with the image of the Delaunay triangulation under the exponential map with domains W^0_α. Therefore (by Theorem 2.2, and the discussion at the end of Section 2) f is a local homeomorphism. In fact, since the sets W^0_α have uniform diameter, it is a covering map. Since $f|A$ is injective, it follows that f is a homeomorphism: in other words, a triangulation of M. □

Lemma 7.6. *Given $\omega > 0$, then there is some F such that if M is F-flat, then the metric on the 1-skeleton, $\Delta^1(A)$, extends to a euclidean metric, ρ_Δ, on $\Delta(A)$.*

Proof. Let Σ be an m-simplex of $\Delta(A)$, and let $\alpha \in \Sigma^0$, so that $f(\Sigma) \subseteq W_\alpha$. Since $\phi_\alpha : W_\alpha \longrightarrow \mathbb{R}^m$ is ξ-bilipschitz, we can assume (if ξ is sufficiently close to 1) that each euclidean edge-length agrees with that induced by ρ_M, up to an additive error of at most the constant δ, given by Lemma 2.5 (given ω and $t = R$). Therefore there is a euclidean m-simplex (unique up to isometry) where each edge-length agrees with that given by ρ_M. Gluing these simplices together and taking the induced path-metric gives us the required metric, ρ_Δ, on $\Delta(A)$. □

Lemma 7.7. *Given $\mu > 1$ and $\omega > 0$, there is some F such that if M is F-flat, then $f : (\Delta(A), \rho_\Delta) \longrightarrow (M, \rho_M)$ is μ-bilipschitz.*

Proof. We first take M sufficiently flat so that $\xi \leq \sqrt{\mu}$. In other words, the charts ϕ_α are $\sqrt{\mu}$-bilipschitz. By Lemma 6.2, f is ϵ-affine, where we can make ϵ arbitrarily small, by choosing F appropriately. Similarly, as with Lemma 4.9, we can now assume that the maps $\phi_\alpha \circ f$ are also ξ-bilipschitz. The composition $f = \phi_\alpha^{-1} \circ (\phi \circ f)$ is therefore μ-bilipschitz. □

Proof of Theorem 1.2. Recall that we have chosen a fixed scale $R = 1000$, say. Given κ, χ we choose $\eta_0 > 0$ so that if we rescale the metric, ρ_M by a factor $1/\eta_0$, it becomes F-flat, where F is determined by the following condition. First, we require that F should satisfy the conclusion of Lemma 7.3. Now Lemma 7.3 also gives us a constant $\omega > 0$ (so that M admits an ω-regular net). We assume also that, given this ω, the constants F are also sufficient for Lemmas 7.2, 7.4 and 7.6 to hold, as well as for Lemma 7.7 to hold, given the constant $\mu > 1$ of our hypotheses. Note that, these statements also hold on recaling the metric, ρ_M, by any factor $1/\eta$ for $\eta \leq \eta_0$ (as required for the conclusion of Theorem 1.2).

Let A be the ω-regular net as given by Lemma 7.3. Let $f : \Delta(A) \longrightarrow M$ be as described above (constructed using Lemma 6.2). By Lemma 7.5, this is a triangulation. Let ρ_Δ be the metric on $\Delta(A)$ given by Lemma 7.6. By Lemma 7.7 the triangulation is μ-bilipschitz with respect to the metrics ρ_Δ and ρ_M/η.

Moreover, ρ_Δ is λ-bilipschitz to the standard metric, ρ^0_Δ, where λ depends only on ω and m, and hence ultimately only on κ, χ, m.

Now let $\Theta = \Delta(A)$. We rescale everything back by a factor of η. In other words, $\rho^0_\theta = \rho^0_\Delta$ and $\rho_\theta = \eta\rho^0_\theta$. Write $\tau = f$. Then $\tau : \Theta \longrightarrow M$ is the required triangulation. □

8. Manifolds with boundary

Let M be a riemannian m-manifold with smooth boundary, ∂M.

If $a \in \partial M$ we can identify $T_a(\partial M)$ as an $(m-1)$-dimensional subspace of $T_a(M)$. Given $a \in \partial M$, write $\nu(a) \in T_a(M)$ for the inward-pointing unit normal to ∂M. Let $\beta_a : [0, h] \longrightarrow M$ be the unit-speed geodesic with $\beta_a(0) = a$ and $\beta'_a(0) = \nu(a)$. We say that M is *h-collared* if the map $\psi : \partial M \times [0, h] \longrightarrow M$, defined by $\psi(a, t) = \beta_a(t)$ for all $a \in \partial M$ and $t \in [0, h]$, is an injective immersion. Its image is then the closed h-neighbourhood of ∂M in M. We refer to it as the (closed) *h-collar*.

Thinking of the boundary, ∂M, as a hypersurface in M, it has a second fundamental form, S, given by $S(v, w) = \langle \nu, \nabla_v w\rangle = -\langle \nabla_v \nu, w\rangle$, for $v, w \in T_a(\partial M)$, where ∇ is the Levi-Civita connection on M.

Definition. We will say that M is (κ, χ)-*bounded* if the following conditions hold:
The sectional curvatures of M and the norm of the second fundamental form of ∂M are all at most κ.
The convexity radius of M is everywhere at least χ.
The intrinsic convexity radius of ∂M is everywhere at least χ, and ∂M is (2χ)-collared.

The conditions on ∂M imply that the intrinsic sectional curvatures of ∂M are bounded in norm by some fixed multiple of κ.

We need to define a modified exponential map for $a \in \partial M$. To this end, let $\mathbb{R}^m_+ = \mathbb{R}^{m-1} \times [0, \infty)$ and $\partial\mathbb{R}^m_+ = \mathbb{R}^{m-1} \times \{0\}$. Given $\chi > 0$, write $\Omega_+ = N(o; 2\chi) \cap \mathbb{R}^m_+$ and $\partial\Omega_+ = \Omega \cap \partial\mathbb{R}^m_+$. (We will later rescale, and set χ to be equal to a fixed constant.)

Given $a \in \partial M$, we choose an isometric identification of T_aM with $\mathbb{R}^m$, such that $T_a(\partial M)$ gets identified with $\partial\mathbb{R}^{m-1}$, and inward-pointing vectors with $\mathbb{R}^{m-1} \times (0, \infty)$. Let $\theta^\partial_a : \partial\Omega_+ \longrightarrow \partial M$ be the intrinsically defined exponential map on ∂M. Given $x \in \Omega_+$, write $x = (\hat{x}, t)$, where $\hat{x} \in \partial\Omega_+$ and $t \in [0, 2\chi)$. Let $c = \theta^\partial_a(\hat{x})$ and set $\theta_a(x) = \beta_c(t)$. This defines a map $\theta = \theta_a : \Omega_+ \longrightarrow M$. Since M is (2χ)-collared, θ is injective. Let $W = W_a = \theta(\Omega_+)$. This is an open neighbourhood of a in M. We write $\partial W = \theta(\partial\Omega_+) = W \cap \partial M$.

By the Rauch comparison theorem, $\theta : \Omega_+ \longrightarrow W$ is ξ-bilipschitz, where ξ depends only on κ. In fact, if κ is small enough, then ξ is arbitrarily close to 1.

We also claim that θ satisfies property (C2) of Section 5, where ϵ is arbitrarily small depending on κ. To make sense of this, we first need to define parallel transport on W. In what follows we will write $\sim$ for the relation $\sim_\epsilon$, where $\epsilon > 0$ depends only the stage of the argument and parameters of the hypotheses. It can be made arbitrarily small by taking κ sufficiently small.

Given $b, c \in W$, we define an isometry $\tau^W_{bc} : T_b W \longrightarrow T_c W$, by parallel transport along a path form b to c. For definiteness, we can take the θ-image of the geodesic from $\theta^{-1}b$ to $\theta^{-1}c$ in Ω_+. However the choice doesn't matter much. If τ is parallel transport along another path in W of length at most 4χ say, then $\tau \sim \tau^W_{bc}$.

Also, if $b, c \in \partial W$ and $X \in T_b(\partial W)$ is any unit vector, then $\tau^{\partial W}_{bc} X \sim \tau^W_{bc} X$, where $\tau^{\partial W}_{bc}$ is the intrinsic parallel transport along geodesics in ∂W. This follows since the extrinsic curvature of W in ∂W is arbitrarily small.

Now suppose that $x, y \in \Omega_+$. We claim that $\tau^W_{\theta x, \theta y} \circ d_x\theta \sim d_y\theta$. Here we have identified $T_x\Omega_+ \equiv T_y\Omega_+ \equiv \mathbb{R}^m$, so that parallel transport is just the identity map. We can assume that $x = o$, so that $\theta(x) = a$. The argument, similar to that of Lemma 5.1, goes as follows.

Write $y = (\hat{y}, l)$ and let $z = (\hat{y}, 0)$. Let $c = \theta(z) \in \partial W$. Let α_0, α_1 be respectively the unit-speed geodesics from o to z and from z to y. Let $\beta_i = \theta \circ \alpha_i$. Thus, β_0 is the intrinsic geodesic from a to c in ∂W, and $\beta_1 = \beta_c$ is the geodesic from c to b in M, orthogonal to ∂M. Let τ_i be parallel transport in M along β_i. Thus, $\tau^W_{ab} \sim \tau_1 \circ \tau_0$.

Let $v_0 = (0, \ldots, 0, 1)$ be the unit vector in $\mathbb{R}^m_+$ orthogonal to $\partial\mathbb{R}^m_+$. Let $V_0 = \nu(a) = (d_a\theta)v_0 \in T_a W$. Then $\tau_0(V_0) \sim \nu(c)$ and so $\beta_1'(l) = (d_b\theta)v_0 \sim \tau_1 \circ \tau_0(V_0)$.

Suppose that $v \in \partial\mathbb{R}^{m-1} \times \{0\}$. Let $V = (d_a\theta)v \in T_a(\partial W)$. Let $X_0 = (d_b\theta)v \in T_b(\partial W)$. Let X_0' be the parallel transport of V along β_0 intrinsically in ∂M. By Lemma 5.1, intrinsic to ∂M, we have $X_0' \sim X_0$. Let $X_0'' = \tau_0 X_0''$ (that is, parallel transport in M), so that $X_0'' \sim X_0' \sim X_0$.

We now define two orthogonal vector fields, X, Y, along β_1, with $X(0) = Y(0) = X_0$. First set $X(t) = (d_{\beta_1'(t)}\theta)v$. Note that this is a Jacobi field along β_1. Therefore, as in Section 5, we have $||\frac{D^2}{\partial t^2}X|| \leq 4\kappa/3$. Also at $t = 0$, we have $||\frac{D}{\partial t}X|| = ||\nabla_{X_0}\nu|| \leq \kappa$. We let $Y(t)$ be parallel transport along β_1, and so $\frac{D}{\partial t}Y(t) = 0$. It follows that $||Y(l) - X(l)||$ is arbitrarily small.

Now $Y(l) = (d_b\theta)v$, and $X(l) = \tau_1 \circ \tau_0(X_0) \sim \tau_1(X_0'') = \tau_1 \circ \tau_0(V)$, so $(d_b\theta)v \sim \tau^W_{ab}(d_y\theta)v$. It follows that $\tau^W_{ab} \circ d_x\theta \sim d_y\theta$ as required.

Recall that we have defined maps, $\theta_a : \Omega_+ \longrightarrow M$, for $a \in \partial M$. Given $a \in M \setminus N(\partial M, \chi)$, we define define $\theta_a : \Omega \longrightarrow M$ to be the usual exponential map, where $\Omega = N(o, \chi) \subseteq \mathbb{R}^m$. We now see that the statement of Lemma 5.2 holds, with $a, b \in \partial M \cup (M \setminus N(\partial M; \chi_0))$. The inverses of these maps define a smooth atlas for M. Exactly as in Section 6, we say that the atlas is (ξ, η)-*flat* if the charts are all ξ-bilipschitz and the transition functions are η-congruences. Which after rescaling, can be assumed to be arbitrarily flat i.e. with $\xi > 1$ arbitrarily close to 1 and with $\eta > 0$ arbitrarily small.

The next step is to observe that the Delaunay triangulation works also for the euclidean half-space. Suppose that $A \subseteq \mathbb{R}^m_+$ is locally finite and t-dense in $\mathbb{R}^m_+$ and also that $A \cap \partial\mathbb{R}^m_+$ is t-dense in $\partial\mathbb{R}^m_+$. We construct the Delaunay complex, $\Delta(A) \subseteq \Pi(A, 2t)$, exactly as in Section 2, allowing (for the moment) x to be any point of $\mathbb{R}^m$. (In other words allowing the centres of circumscribing spheres to lie outside $\mathbb{R}^m_+$.) Under the same non-degeneracy assumption as in Theorem 2.1, the map $\tau : \Delta(A) \longrightarrow \mathbb{R}^m_+$ is a triangulation of $\mathbb{R}^m_+$. Moreover, its restriction to $\partial\mathbb{R}^m_+$ is precisely the intrinsically defined Delaunay triangulation, $\Delta(A \cap \partial\mathbb{R}^m_+)$.

We would like the centres of circumscribing spheres to lie inside $\mathbb{R}^m_+$. This is necessarily the case if we suppose, for example, that $\rho_E(A \cap \partial\mathbb{R}^m_+, \partial\mathbb{R}^m_+) \geq t$. Such a set can always be constructed. For example, choose any $s \leq t/2$, and let A be a maximal s-separated set in $\partial\mathbb{R}^m_+ \cup (\mathbb{R}^m_+ \setminus N(\partial\mathbb{R}^m_+; t))$. We will however want a more general construction which works for our manifold, M.

To this end, we fix constants, $R = 1000$ and $R_0 = R/10$, as in Section 7. Let $\Omega = N(o; R)$ and $\Omega^+ = N(o; R) \cap \mathbb{R}^m_+$. and Let $M' = M \setminus N(\partial M; 2R)$. After rescaling the metric appropriately, we can define an altlas, with indexing set $M' \cup \partial M$, as above.

Namely, if $\alpha \in \partial M$, let $W_\alpha = \theta_\alpha(\Omega_+)$, and let $\phi_\alpha = \theta_\alpha^{-1} : W_\alpha \longrightarrow \Omega_+$. $\alpha \in M'$. If $\alpha \in M'$, let $W_\alpha = \theta_\alpha(\Omega) = N(\alpha; R)$, and let $\phi_\alpha = \theta_\alpha : W_\alpha \longrightarrow \mathbb{R}^m$ be the inverse exponential (logarithm) map. Then $\{\phi_\alpha\}_{\alpha \in M' \cup \partial M}$, is an atlas for M. Moreover, by rescaling sufficiently, we can assume that the atlas is as flat as we want.

We also have an atlas of smaller charts. Let $\Omega^0 = N(0, 2R_0)$ and $\Omega^0_+ = N(o; R_0) \cap \mathbb{R}^m_+$. Let $W^0_\alpha = \theta_\alpha(\Omega^0_+)$ for $\alpha \in \partial M$. Let $W^0_\alpha = \theta_\alpha(\Omega^0) = N(\alpha; R_0)$ for $\alpha \in M'$. Then $\{\phi_\alpha | W^0_\alpha\}_{\alpha \in M' \cup \partial M}$ is also a smooth atlas.

We need to modify the definition of "ω-regular".

As before, given $A \subseteq M$, $i \in \mathbb{N}$ and $\alpha \in M' \cup \partial M$, write $\mathcal{C}^i_\alpha(A)$ for the set of $B \subseteq A \cap W^0_\alpha$ with $|B| = i + 1$ (with W^0_α now defined as above). We now write $\mathcal{C}^i(A) = \bigcup_{\alpha \in M' \cup \partial M} \mathcal{C}^i_\alpha(A)$.

Definition. Given $\omega > 0$ and $A \subseteq M$, we say that A is *ω-regular* if:
(R1′): for all $\alpha \in M' \cup \partial M$ and $B \in \mathcal{C}^m_\alpha(A)$, we have $V_m(\phi_\alpha B) \geq \omega$,
(R2′): for all $\alpha \in M' \cup \partial M$ and $C \in \mathcal{C}^{m+1}_\alpha(A)$, we have $Q_m(\phi_\alpha C) \geq \omega$, and
(R3′): for all $\alpha \in M' \cup \partial M$ and $C \in \mathcal{C}^m_\alpha(A)$, we have $Q_{m-1}(\phi_\alpha C) \geq \omega$.

(Note that this implies the intrinsic regularity of $A \cap \partial M$ in ∂M.)

Let $R_1 = R_0/10$.

Definition. A *net* $A \subseteq M$, is a subset $A \subseteq M' \cup \partial M$ which is R_1-separated and $(6R_1)$-dense in $M' \cup \partial M$.

(As before, the "6" could be replaced by number greater than 2.)

We can generalise Lemma 7.3 to manifolds with boundary. (The statement is identical.)

Lemma 8.1. *There exist $\omega > 0$ and F, depending only on m such that if M is F-flat, then it admits an ω-regular net.*

Proof. We start with a $(2R_1)$-separated $(4R_1)$-dense subset of $M' \cup \partial M'$. (Take any maximal $(2R_1)$-separated subset thereof.) We now move it slightly so that it becomes ω-regular. The argument is a slight modification to that of 7.7. First perturb the set on ∂M so that it becomes uniformly regular there. Now fixing that set, we perturb it on M', by partitioning it into a bounded number of subsets, and moving each in turn, so as the satisfy the remaining conditions. This proceeds by essentially the same construction as with the proof of Lemma 7.7. □

The results of Section 6 go through similarly. In particular, Lemma 6.2 applies in this case, moreover with the addendum that if $B \in \mathcal{C}^i(\Pi)$ with $B \subseteq \partial M$, then $f(\Sigma(B)) \subseteq \partial M$. In other words, if the vertex set of a simplex of Π lies in ∂M, then the extension of the map to the simplex also lies in ∂M. This is an immediate consequence of the construction for extending smooth maps on boundaries of simplices.

To prove Theorem 1.3, we now proceed as in Section 7. We rescale so that M is sufficiently flat (as required by Lemma 8.1) and let $A \subseteq M' \cup \partial M$ be an ω-regular net.

We extend the inclusion of the net to a map of the whole complex. First we map in the 1-skeleton, so that edges map to geodesics either in M or to intrinsic geodesics in ∂M. The remainder of the extension process follows as in Section 7. We need to observe that if the boundary of a simplex lies in ∂M, then so does the extension. This is an immediate consequence of the construction.

To see that this gives the required triangulation, we follow the proof of Theorem 1.2 at the end of Section 7. The relevant results all hold in the case of a manifold with boundary. We substitute Lemma 8.1 for Lemma 7.3. The other ingredients, Lemmas 7.2, 7.4, 7.5, 7.6 and 7.7, go through without any essential change.

This proves Theorem 1.3.

9. Constant curvature

In this final section, we say how the results can be strengthened in the case of constant curvature: namely that the simplices of the triangulation can be taken to be totally geodesic (Theorem 1.4).

If we only care about the images of individual simplices, this would be relatively straightforward once we have constructed a regular net. The Delaunay construction (which we have described for euclidean space) works for any constant curvature manifold. This gives a partition into simplices, each of which is uniformly bilipschitz equivalent to regular euclidean simplex. A construction of a regular net for hyperbolic manifolds (slightly different from ours) and the resulting Delaunay triangulation, is described in [Br].

If we want a genuine smooth triangulation, then we need to consistently reparameterise the simplices. In other words, we construct an explicit map from a simplicial complex. This follows from the earlier constructions, though taking care to ensure that the images are totally geodesic.

Up to scale there are three cases, spherical, euclidean and hyperbolic. We first discuss the case of hyperbolic geometry. Let $\mathbb{H}^n$ be hyperbolic n-space.

Given a finite subset $A \subseteq \mathbb{H}^n$, we write $\Sigma(A)$ for its convex hull. This is a polyhedron with vertex set, $\Sigma^0(A)$, contained in A. We say that A is *non-degenerate* if $\Sigma^0(A)$ is a simplex with $\Sigma^0(A) = A$. In this case, $|A| \leq n+1$. Note that if $H \subseteq \mathbb{H}^n$ is a totally geodesic subspace, then H is isometric to $\mathbb{H}^m$ for some $m \leq n$, and we can take the convex hull intrinsically in $\mathbb{H}^m$.

Let Σ_m be the standard euclidean m-simplex. We choose a preferred ordering of the vertex set as $\Sigma^0_m = \{v_0, \dots, v_m\}$. Note that any permutation, π, of the vertex set extends to an isometry, $\bar{\pi}$, of Σ_m. If $i \leq m$, we can identify Σ_i with the face of Σ_m with vertex set, $\{v_0, \dots, v_i\}$.

Given a map $f : \Sigma^0_m \longrightarrow \mathbb{H}^n$, we there is a canonical smooth extension, $\bar{f} : \Sigma_m \longrightarrow \mathbb{H}^n$, such that $[f \mapsto \bar{f}]$, satisfies the following properties:

(E1): $\bar{f}(\Sigma_m) = \Sigma(f(\Sigma^0_m))$.
(E2): If f is injective and $f(\Sigma^0_m)$ is non-degenerate, then $\bar{f}$ is a diffeomorphism from Σ_m to $\Sigma(f(\Sigma^0_m))$.
(E3): If π is any permutation of Σ^0_m, then $\overline{f \circ \pi} = \bar{f} \circ \bar{\pi}$.
(E4): If $i \leq m$, then $\overline{f|\Sigma^0_i} = \bar{f}|\Sigma_i$.
(E5): The map $\bar{f}|\Sigma_1$ is a linear map to the geodesic segment $\Sigma(f(\Sigma^0_1))$.
(E6): If $g : \mathbb{H}^n \longrightarrow \mathbb{H}^n$ is any isometry, then $\overline{g \circ f} = g \circ \bar{f}$.
(E7) Given any $m \in \mathbb{N}$ and $\epsilon > 0$, there is some $r > 0$, so that if $\operatorname{diam} f(\Sigma^0_m) \leq \epsilon$, then $\phi \circ \bar{f}$ is ϵ-affine, where ϕ is the logarithm (inverse exponential) map, $\mathbb{H}^n \longrightarrow \mathbb{R}^n$, based at any point of $\bar{f}(\Sigma_m)$.

In fact, we can also assume that the map $[f \mapsto \bar{f}]$ is smooth, thought of as a function from $(\mathbb{H}^n)^m \times \Sigma_m$ to $\mathbb{H}^n$, though we will not need this.

One way to construct the map would be to use the extension process described by Lemma 4.1. To this end, we assume inductively that we have defined $\bar{f}|\partial\Sigma_m$. We now take the Klein model for $\mathbb{H}^n$, centred on the barycentre of $\bar{f}(\Sigma^0_m)$, and extend $\bar{f}$ as in Lemma 4.1. Since hyperbolic and euclidean convex hulls coincide in the Klein model, the image, $\bar{f}(\Sigma_m)$, will be the hyperbolic convex hull of $f(\Sigma^0_m)$.

A more natural process would be to use the "barycentre" construction of [K]. Given $x \in \Sigma_m$, write $x = \sum_{i=0}^m \lambda_i v_i$, where $\Sigma^m_0 = \{v_0, \dots, v_m\}$. Let $y \in \mathbb{H}^n$ be the unique point which minimises $\sum_{i=0}^m \lambda_i d_{\mathbb{H}^n}(y, f(v_i))^2$, and set $f(x) = y$. Again this has the properties laid out above.

The proof of Theorem 1.4 in the hyperbolic case now follows by previous arguments. We construct a regular net (as in Section 7 or 8 here, or by the method of [Br]). We then take the hyperbolic Delaunay triangulation, and construct a smooth triangulation using the canonical maps as described above. Properties (E3) and (E4) ensure that this is well-defined on intersecting simplices.

In the case of euclidean space, we just take the triangulation to be affine on each simplex.

For spherical geometry, we could use the barycentric construction of [K] as mentioned above.

References

[BoDG] J.-D.Boissonnat, R.Dyer, A.Ghosh, *Delaunay triangulation of manifolds* : Found. Comput. Math. **18** (2018) 399–431.

[Br] W.Breslin, *Thick triangulations on hyperbolic n-manifolds* : Pacific J. Math. **241** (2009) 215–225.

[BuK] P.Buser, H.Karcher, *Gromov's almost flat manifolds* : Astérisque, Vol. 81. Société Mathématique de France (1981).

[Ca] S.S.Cairns, *On the triangulation of regular loci* : Ann. of Math. **35** (1934) 579–587.

[Ch] I.Chavel, *Riemannian geometry: A modern introduction* : Cambridge Tracts in Mathematics, Vol 108. Cambridge University Press (1993).

[De] B.Delaunay, *Sur la sphère vide* : Bulletin de l'Académie des Sciences de l'URSS, Classe des sciences mathématiques et naturelles **6** (1934) 793–800.

[DyVW] R.Dyer, G.Vegter, M.Wintraecken, *Riemannian simplices and triangulations* : Geom. Dedicata **179** (2015) 91–138.

[ECHLPT] D.B.A.Epstein, J.W.Cannon, D.F.Holt, S.V.F.Levy, M.S.Paterson, W.P.Thurston, *Word processing in groups* : Jones and Bartlett (1992).

[K] H.Karcher, *Riemannian center of mass and mollifier smoothing* : Comm. Pure Appl. Math. **30** (1977) 509–541.

[L] J.M.Lee, *Introduction to smooth manifolds* : Graduate Texts in Mathematics, Vol. 218, Springer (2003).

[M] J.R.Munkres, *Elementary Differential Topology* : Annals of Mathematics Studies No. 54, Princeton University Press (1966).

[S] E.Saucan, *Note on a theorem of Munkres* : Mediterr. J. Math. **2** (2005) 215–229.

[W] J.H.C.Whitehead, *On C^1-complexes* : Ann. of Math. **41** (1940) 809–824.

Mathematics Institute, University of Warwick, Coventry, CV4 7AL, Great Britain