COMBINATORIAL SYSTOLIC INEQUALITIES

RYAN KOWALICK, JEAN-FRANÇOIS LAFONT, AND BARRY MINEMYER

Abstract. We establish combinatorial versions of various classical systolic inequalities. For a smooth triangulation of a closed smooth manifold, the minimal number of edges in a homotopically non-trivial loop contained in the 1-skeleton gives an integer called the combinatorial systole. The number of top-dimensional simplices in the triangulation gives another integer called the combinatorial volume. We show that a class of smooth manifolds satisfies a systolic inequality for all Riemannian metrics if and only if it satisfies a corresponding combinatorial systolic inequality for all smooth triangulations. Along the way, we show that any closed Riemannian manifold has a smooth triangulation which "remembers" the geometry of the Riemannian metric, and conversely, that every smooth triangulation gives rise to Riemannian metrics which encode the combinatorics of the triangulation. We give a few applications of these results.

1. Introduction

For a closed Riemannian manifold (M, g), the *systole* is the minimal length of a homotopically non-trivial loop, denoted $\mathrm{Sys}_g(M)$, while the volume of (M, g) is denoted $\mathrm{Vol}_g(M)$. Systolic inequalities are expressions which relate the systole with other geometric quantities, typically the volume. In this paper, we are interested in *combinatorial* versions of the systolic inequalities.

We view smooth triangulations of a manifold M as a combinatorial model for M. For such a triangulation $(M, \mathcal{T})$, we define the combinatorial systole $\mathrm{Sys}_{\mathcal{T}}(M)$ to be the minimal number of edges for a combinatorial loop in the 1-skeleton of $\mathcal{T}$ which is homotopically non-trivial in M. The discrete volume $\mathrm{Vol}_{\mathcal{T}}(M)$ is just the number of top-dimensional simplices in the triangulation $\mathcal{T}$. The main goal of this paper is to establish the following:

Main Theorem. *Let $\mathcal{M}$ be a class of closed smooth n-manifolds. Then the following two statements are equivalent:*

(1) *for every Riemannian metric (M, g) on a manifold $M \in \mathcal{M}$, we have*

$$\mathrm{Sys}_g(M) \leq C \sqrt[n]{\mathrm{Vol}_g(M)},$$

where C is a constant which depends solely on the class $\mathcal{M}$.

(2) *for every smooth triangulation $(M, \mathcal{T})$ of a manifold $M \in \mathcal{M}$, we have*

$$\mathrm{Sys}_{\mathcal{T}}(M) \leq C' \sqrt[n]{\mathrm{Vol}_{\mathcal{T}}(M)},$$

where C' is a constant which depends solely on the class $\mathcal{M}$.

Date: October 24, 2016.

In [Gr1] Gromov proved that the class of closed smooth *essential* Riemannian manifolds satisfies the above Riemannian systolic inequality. So an immediate consequence is the following.

Corollary. *Let $\mathcal{M}$ denote the class of closed smooth essential n-manifolds. Then for every smooth triangulation $(M, \mathcal{T})$ of a manifold $M \in \mathcal{M}$, we have*

$$\mathrm{Sys}_{\mathcal{T}}(M) \leq C \sqrt[n]{\mathrm{Vol}_{\mathcal{T}}(M)},$$

where C is a constant which depends solely on the dimension n.

In the process of proving our **Main Theorem**, we establish a number of auxiliary results which might also be of independent interest. After some preliminaries in Section 2, we show:

Theorem 1 (Encoding a Riemannian metric)**.** *There exists a constant δ_n depending solely on the dimension n, with the property that for any closed Riemannian manifold (M, g), there exists a smooth triangulation $\mathcal{T}$ with the property that*

$$\frac{\sup_{e \subset \mathcal{T}}\{\ell_g(e)\}}{\inf_{\sigma \subset \mathcal{T}}\{\sqrt[n]{\mathrm{Vol}_g(\sigma)}\}} \leq \delta_n,$$

where the volume of the top-dimensional simplices σ, and the lengths of the edges e, are measured in the ambient g-metric.

Roughly speaking, the triangulation $\mathcal{T}$ produced in the theorem has no simplices that are "long and thin" (as measured in the Riemannian metric g). Moreover, Theorem 1 still holds for a possibly different constant $\delta_{n,k}$ when considering the collection of k-dimensional simplices σ (and replacing $\sqrt[n]{\mathrm{Vol}_g(\sigma)}$ with the k^{th} root of the k-dimensional volume of σ). Theorem 1 is established in Section 3. The idea is as follows. In [HW], Whitney proved that every closed smooth manifold M^n supports a triangulation. The method Whitney used was to first smoothly embed M^n into $\mathbb{R}^{2n+1}$, then equip the latter with a sufficiently fine cubulation. Then one perturbs the embedding to be transverse to the $(n+1)$-cubes in the cubulation – the intersection will then give a collection of points. One then uses these points as the vertex set of a certain piecewise affine (polyhedral) complex in $\mathbb{R}^{2n+1}$. If the cubulation is chosen fine enough, this complex lies in a small normal neighborhood of M^n, and one can subdivide to get a simplicial complex, then project down onto M^n. Whitney then argues that this projection gives a smooth triangulation of M.

One quick remark is that in dimensions ≥ 4, the corresponding statement is *false* for topological manifolds (work of Freedman and Casson in dimension $= 4$ [Fr], and of Manolescu [Ma] in dimensions ≥ 5), see Section 7. Also, the existence of triangulations of manifolds was known before Whitney. For example, see Cairns [Cai] and Whitehead [JW].

Now the proof of Theorem 1 also uses Whitney's procedure, but rather than starting from a smooth embedding into $\mathbb{R}^{2n+1}$, we want to start with an embedding that "remembers" the Riemannian structure on (M, g). A natural choice to use is Nash's isometric embedding. We then follow through Whitney's arguments, and check that the resulting triangulation has the desired property. This is done in Section 4.

Theorem 2 (Encoding a triangulation)**.** *There exists a constant κ_n depending solely on the dimension n, with the property that for any smooth triangulation $(M, \mathcal{T})$ of a smooth compact manifold M, and for any $\varepsilon > 0$, there exists a Riemannian metric g on M which satisfies the following:*

(1) $|Vol_g(M) - Vol_{\mathcal{T}}(M)| < \varepsilon$
(2) *If γ is a closed path on M, then there exists a closed edge loop p, freely homotopic to γ, so that*

$$\ell_{\mathcal{T}}(p) \leq \kappa_n \ell_g(\gamma).$$

By an *edge loop p* we mean a closed simplicial path in the 1-skeleton of $\mathcal{T}$, and the notation $\ell_{\mathcal{T}}(p)$ denotes the number of edges contained in the image of p. The idea behind the proof is to put a piecewise Euclidean metric on M, by making each n-dimensional simplex in the triangulation $\mathcal{T}$ isometric to a Euclidean simplex with all edges of equal length, and of volume equal to one. This metric has singularities along the codimension two strata, which can be inductively smoothed out. This gives a metric g satisfying property (1). For property (2), one can easily reduce to the case that γ is a g-geodesic which is not null-homotopic. From there, we remove the sections of γ near the codimension 2 skeleton and, in a Lipschitz manner, replace them with geodesic segments in the singular metric. This results in a loop of roughly comparable length in the singular metric, and property (2) is easy to establish for the singular metrics. The details of this argument can be found in Section 5.

In this paper, a *triangulation* of a manifold M is a simplicial complex $\mathcal{T}$ together with a homeomorphism $\iota : \mathcal{T} \to M$. If the link of every simplex in $\mathcal{T}$ is a piecewise-linear sphere then we call this triangulation a *piecewise-linear triangulation*, and if $\iota|_{|\sigma|}$ is smooth for all simplices $\sigma \in \mathcal{T}$ then we call $\mathcal{T}$ a *smooth triangulation*. In the proceeding sentence, for σ an n-simplex, $|\sigma| \subset \mathbb{R}^n$ denotes the canonical n-simplex in $\mathbb{R}^n$. It is known that all smooth triangulations of a smooth manifold are piecewise-linear, but this containment is strict. And of course the celebrated Cannon-Edwards double suspension theorem ([Ed] and [Can]) shows that not all triangulations of a manifold are piecewise-linear. Nevertheless, all three of these notions agree if the dimension of M is at most three.

In the proof of Theorem 2 we only use the assumption that the triangulation $\mathcal{T}$ is compatible with the smooth structure on M in only one place: when using a smooth partition of unity to patch together locally defined metrics in the construction of the Riemannian metric g. We need the metric g to be smooth in order to use Theorem 2 to prove one implication in our **Main Theorem**. But if all one requires is a C^0-Riemannian metric satisfying the two statements in Theorem 2, then the assumption can be weakened to $\mathcal{T}$ being a *piecewise linear* triangulation of M. Our technique of proof does not extend to continuous triangulations, unfortunately. See Section 7 for a further discussion. As a final remark about Theorem 2, we observe that by simply scaling the metric g, one may obtain equality in property (1) above at the cost of slightly altering the Lipschitz constant κ_n in property (2).

Using these two theorems, the proof of the **Main Theorem** is easy.

Proof of Main Theorem. ($\Rightarrow$) Assume you have a class $\mathcal{M}$ of smooth n-manifolds satisfying condition (1) of the theorem, i.e. satisfying a Riemannian systolic inequality. Let $\mathcal{T}$ be a smooth triangulation of a manifold $M \in \mathcal{M}$ lying within the class, and $\epsilon > 0$ an arbitrary positive constant. Let g be the Riemannian metric on M whose existence is provided by our Theorem 2, γ the closed g-geodesic whose length realizes the Riemannian systole of (M, g), and p the edge path freely homotopic to γ given by Theorem 2. Then we have the sequence of inequalities:

$$\begin{aligned}\mathrm{Sys}_{\mathcal{T}}(M) &\leq \ell_{\mathcal{T}}(p) \leq \kappa_n \ell_g(\gamma) = \kappa_n \cdot \mathrm{Sys}_g(M)\\ &\leq \kappa_n \cdot C \sqrt[n]{\mathrm{Vol}_g(M)}\\ &\leq (\kappa_n \cdot C) \sqrt[n]{\mathrm{Vol}_{\mathcal{T}}(M) + \epsilon}\end{aligned}$$

Letting ϵ tend to zero, we see that the class $\mathcal{M}$ satisfies condition (2) of the theorem (i.e. satisfies a combinatorial systolic inequality), with constant $C' = \kappa_n \cdot C$.

($\Leftarrow$) Conversely, let us assume that you have a class $\mathcal{M}$ of smooth n-manifolds satisfying condition (2) of the theorem, i.e. satisfying a combinatorial systolic inequality. Let g be an arbitrary Riemannian metric on one of the manifolds $M \in \mathcal{M}$ lying within the class. Let $\mathcal{T}$ be the smooth triangulation of M obtained by applying our Theorem 1. We denote by E the supremum of the g-lengths of edges in $\mathcal{T}$, and by v the infimum of the volume of top dimensional simplices in $\mathcal{T}$. So by Theorem 1, we have that $\frac{E}{v^{1/n}} \leq \delta_n$. Let p be an edge path in the triangulation $\mathcal{T}$ which realizes the combinatorial systole. Then we have the series of inequalities:

$$\begin{aligned}\mathrm{Sys}_g(M) &\leq \ell_g(p) \leq E \cdot \ell_{\mathcal{T}}(p) = E \cdot \mathrm{Sys}_{\mathcal{T}}(M)\\ &\leq C' \cdot E \cdot \sqrt[n]{\mathrm{Vol}_{\mathcal{T}}(M)} \leq C' \cdot E \cdot \sqrt[n]{\frac{\mathrm{Vol}_g(M)}{v}} = \delta_n C' \cdot \sqrt[n]{\mathrm{Vol}_g(M)}\end{aligned}$$

Thus, we see that the class $\mathcal{M}$ satisfies condition (1) of the theorem (i.e. satisfies a Riemannian systolic inequality), with constant $C = \delta_n \cdot C'$. This concludes the proof of our **Main Theorem**. $\square$

After the proof of Theorem 2, we discuss some applications of our **Main Theorem** in Section 6. Our paper concludes with a discussion about some open problems in Section 7, and an Appendix listing some general topology results due to Whitney [HW] which are used in Section 4.

Remark. Our **Main Theorem** still holds if we instead consider the homological systole of M. This is clear since Theorem's 1 and 2 only deal with Lipschitz homotopies of paths. But an interesting question is if one can obtain higher dimensional analogues of both the homotopy and homological systolic inequalities, such as those discussed by Brunnbauer in [Br].

Remark. Most of Sections 3, 4, and 6 are contained in the Ph. D. Thesis of Ryan Kowalick [Ko]. Similar results were independently obtained by de Verdière, Hubard, and de Mesmay [VHM]. Their results are focused on the 2-dimensional closed surfaces case (and includes other applications), but they include an Appendix where they discuss analogous results in higher dimensions.

Acknowledgments. The authors would like to thank Dylan Thurston for some helpful comments. The work of the second author was partially supported by the NSF, under grants DMS-1207782 and DMS-1510640.

2. Background material and notation

Suppose X is a metric space, with S, $T \subset X$. We define the *distance between S and T* by

$$\operatorname{dist}(S,T) = \inf\{d(s,t) \,:\, s \in S, t \in T\}$$

and we note that this definition remains unchanged in the event that either S or T consists of a single point. Also, we define the *r-neighborhood of S in X*, denoted $U_r(S)$ to be

$$U_r(S) = \{x \in X \,:\, \operatorname{dist}(x,S) < r\}.$$

In this paper all manifolds, (Riemannian) metrics, and (simplicial) triangulations are assumed to be smooth. A *triangulated manifold* is a tuple $(M, \mathcal{T})$ where M is a manifold and $\mathcal{T}$ is a triangulation. When there is the possibility of confusion, we will denote the triangulation of a manifold M by $\mathcal{T}_M$ instead of $\mathcal{T}$. We also may abuse notation and use either $\mathcal{T}$ or $\mathcal{T}_M$ to denote M when confusion will not arise. A *filling* of a closed triangulated n-dimensional manifold $(M, \mathcal{T}_M)$ is a triangulated $(n+1)$-dimensional manifold $(N, \mathcal{T}_N)$ with $\partial N = M$ and $\mathcal{T}_N|_{\partial N} = \mathcal{T}_M$.

If $\mathcal{T}$ is a simplicial complex, the *k-skeleton of $\mathcal{T}$*, denoted $\mathcal{T}^{(k)}$, will refer to the subcomplex of $\mathcal{T}$ consisting of all simplices of dimension at most k. A *facet* of a triangulation is a simplex of maximal dimension. For any triangulation $\mathcal{T}$, the notation $|\mathcal{T}|$ will refer to the number of facets in the triangulation. In the case of a triangulated manifold, this will be used as a discrete analogue of volume.

The *systole* of a Riemannian manifold (M, g), denoted $\operatorname{Sys}_g(M)$, is the length of the shortest non-contractible loop in M. The *homological systole* of a Riemannian manifold (M, g), denoted $\operatorname{Sys}_g^H(M)$, is the length of the shortest homologically nontrivial loop in M.

If p is an edge path in the triangulated manifold $(M, \mathcal{T})$, the *discrete length* of p, denoted $\ell_{\mathcal{T}}(p)$, will be the number of edges in p. The *discrete systole* of a triangulated manifold $\mathcal{T}$, denoted $\operatorname{Sys}_{\mathcal{T}}(M)$, will refer to the discrete length of the shortest non-contractible edge loop in $\mathcal{T}$. The *discrete homological systole*, denoted $\operatorname{Sys}_{\mathcal{T}}^H(M)$, is defined analogously.

If σ is an n-simplex in $\mathbb{R}^m$, we define its *fullness* to be

$$\Theta(\sigma) = \frac{\operatorname{Vol}_n(\sigma)}{(\operatorname{diam}\sigma)^n},$$

where Vol_n denotes the n-dimensional Hausdorff volume in $\mathbb{R}^m$. We also note that the diameter, $\operatorname{diam}\sigma$, is the length of the longest side in this case.

Tubular neighborhoods and horizontal tangent vectors. Suppose M is an embedded submanifold of $\mathbb{R}^m$, and let U be a tubular neighborhood of M in $\mathbb{R}^m$. Then the projection map $\pi^*\colon U \to M$ is a Riemannian submersion. So $TU \cong TU_h \oplus TU_v$ where TU_h is canonically isomorphic to the tangent bundle TM of M, and similarly for TU_v and the normal bundle TN of M in $\mathbb{R}^m$. We will refer to TU_h as the *horizontal* component of TU and TU_v as the *vertical* component of TU. So for $q \in U$ and $w \in T_qU$, we will write $w = w_h + w_v$ where $w_h \in T_qU_h$ and $w_v \in T_qU_v$. Also note that for any point $q \in U$, the space T_qU_v is equal to the kernel of the derivative of the projection map $D\pi_{\pi^*(q)} = \pi_{\pi^*(q)}$. So if $w \in T_qU$ and $w = w_h + w_v$, then

$$|w_h| = |\pi_{\pi^*(q)}(w)|. \tag{2.1}$$

Now, for any $p \in M$, the map

$$D\pi^*\Big|_{T_pU_h} : T_pU_h \to T_pM$$

is the identity. Thus for any $\epsilon > 0$, there is a smaller tubular neighborhood $U' \subset U$ so that, for any $q \in U'$, the map

$$D\pi^*\Big|_{T_qU'_h} : T_qU'_h \to T_{\pi^*(q)}M$$

has the property that, for any $w \in T_qU'_h$,

$$\frac{1}{\sqrt{3/2}}|w| \leq |D\pi^*(w)| \leq \sqrt{3/2}|w|. \tag{2.2}$$

3. Constructing a triangulation which encodes a Riemannian metric

In this section we prove Theorem 1. Namely, given a smooth Riemannian manifold (M, g), we want to construct a triangulation where the ratio

$$\frac{\sup_{e\subset\mathcal{T}}\{\ell_g(e)\}}{\inf_{\sigma\subset\mathcal{T}}\{\sqrt[n]{\mathrm{Vol}_g(\sigma)}\}}$$

is uniformly bounded above by a constant δ_n which only depends on the dimension of M.

We do this by using the following result, which will be proved in Section 4.

Theorem 3. *Let M be a compact n-dimensional smooth Riemannian submanifold of $\mathbb{R}^m$. Then there is an n-dimensional simplicial complex $T \subset \mathbb{R}^m$ with the following properties:*

(1) *Each simplex of T is a secant simplex of M*
(2) *T is contained in a tubular neighborhood of M. The projection π^* from this neighborhood onto M induces a homeomorphism $\pi^* : T \to M$.*
(3) *If σ is a simplex of T (of any dimension), then its fullness is bounded below by $\Theta_{n,m}$, which depends only on the dimensions of the manifold and the ambient space.*
(4) *For any n-simplex σ of T, point $q \in \sigma$ and tangent vector $v \in T_q\sigma$, we get that*

$$|\pi_{\pi^*(q)}(v)| \geq \frac{1}{2}|v|, \tag{3.1}$$

where $\pi_{\pi^*(q)}$ is the orthogonal projection onto the tangent plane $P_{\pi^*(q)}$.
(5) *If L is the length of an edge in T, then*

$$C_{n,m}\bar{L} \leq L \leq \bar{L} \tag{3.2}$$

for some positive constant $\bar{L}$ which depends on $\mathcal{T}$ but not on L, and positive constant $C_{n,m}$ depending only on n and m.

Let M be an n-dimensional Riemannian manifold. By the Nash Isometric Embedding theorem [Na], M embeds smoothly and isometrically into $\mathbb{R}^m$, where m depends only on n. Thus we may consider the case where M is a smooth Riemannian submanifold of $\mathbb{R}^m$. Then applying Theorem 3, we get an n-dimensional simplicial complex T contained in a tubular neighborhood of M so that the tubular neighborhood projection π^* induces a homeomorphism from T to M.

We will proceed in two parts. The first part will consist of showing that the restriction of π^* to any n-simplex of T is bi-Lipschitz with constants that do not depend on the given simplex. In the second part we will prove Theorem 1 by using this fact to relate the geometry of T with the geometry of $\pi^*(T)$.

π^* is bi-Lipschitz on every n-simplex of T. Let σ be an n-simplex of T and suppose $x_1, x_2 \in \sigma$. Let $p_1 = \pi^*(x_1)$ and $p_2 = \pi^*(x_2)$.

Suppose p is a unit-speed geodesic in σ from x_1 to x_2. For every $t \in [0, \ell(p)]$, $p'(t)$ is a tangent vector in σ, so by Theorem 3 and equation (2.1),

$$1 = |p'(t)| \geq |p'(t)_h| = |\pi_{\pi^* p(t)}(p'(t))| \geq \frac{1}{2}|p'(t)| = \frac{1}{2}. \tag{3.3}$$

Note that, read all the way from left to right, inequality (3.3) is trivial. But various portions of this inequality will be used throughout this Section.

Now $\pi^* \circ p$ is a path on M from p_1 to p_2, and since M is isometrically embedded in $\mathbb{R}^m$,

$$d_M(p_1, p_2) \leq \ell(\pi^* \circ p). \tag{3.4}$$

Combining (2.2) and (3.3), we get that

$$\begin{aligned} \ell(\pi^* \circ p) &= \int_0^{\ell(p)} |D\pi^*_{p(t)}(p'(t))|\, dt \\ &= \int_0^{\ell(p)} |D\pi^*_{p(t)}(p'(t)_h)|\, dt \\ &\leq \sqrt{3/2} \int_0^{\ell(p)} |p'(t)_h|\, dt \\ &\leq \sqrt{3/2} \int_0^{\ell(p)} dt \\ &= \sqrt{3/2} \cdot \ell(p). \end{aligned}$$

Combining the above with (3.4) gives that

$$d_M(p_1, p_2) \leq \sqrt{3/2} \cdot |x_1 - x_2|. \tag{3.5}$$

Now suppose γ is a unit-speed geodesic in M from p_1 to p_2 so that $d_M(p_1, p_2) = \ell(\gamma)$. Then $(\pi^*)^{-1} \circ \gamma$ is a piecewise smooth path in T from x_1 to x_2. We may take a partition of the interval $[0, \ell(\gamma)]$ into

$$0 = a_0 < a_1 < a_2 < \cdots a_N = \ell(\gamma)$$

so that for each i, $(\pi^*)^{-1} \circ \gamma([a_i, a_{i+1}]) \subset \sigma_i$ where σ_i is an n-simplex of T. Let $\gamma|_{[a_i, a_{i+1}]} = \gamma_i$ and let $(\pi^*)^{-1} \circ \gamma(a_i) = b_i$.

Then $(\pi^*)^{-1} \circ \gamma_i$ is a path in σ_i and for every $t \in [a_i, a_{i+1}]$, $(D\pi^*)^{-1}(\gamma_i'(t))$ is a tangent vector in σ_i. So for every $t \in [a_i, a_{i+1}]$, (3.3) gives that

$$|(D\pi^*)^{-1}(\gamma_i'(t))| \leq 2|[(D\pi^*)^{-1}(\gamma_i'(t))]_h|. \tag{3.6}$$

Suppose $D\pi^*(v) = w$ for $v \in T_p\sigma_i$. Then $w = D\pi^*(v_h)$ and (2.2) gives that

$$|v_h| \leq \sqrt{3/2}|w|.$$

Using the above, we get that, for any $t \in [a_i, a_{i+1}]$,

$$|[(D\pi^*)^{-1}(\gamma_i'(t))]_h| \leq \sqrt{3/2}|\gamma_i'(t)|. \tag{3.7}$$

Combining (3.6) and (3.7) gives that

$$\begin{aligned} |b_i - b_{i+1}| &\leq \ell((\pi^*)^{-1} \circ \gamma) \\ &= \int_0^{a_{i+1}-a_i} |(D\pi^*)^{-1}(\gamma_i'(t))|\, dt \\ &\leq 2\int_0^{a_{i+1}-a_i} |[(D\pi^*)^{-1}(\gamma_i'(t))]_h|\, dt \\ &\leq 2\sqrt{3/2}\int_0^{a_{i+1}-a_i} |\gamma_i'(t)|\, dt \\ &\leq 2\sqrt{3/2} \cdot d_M(\gamma(a_i), \gamma(a_{i+1})). \end{aligned}$$

Since γ is a minimizing geodesic, $\sum d_M(\gamma(a_i), \gamma(a_{i+1})) = d_M(p_1, p_2)$. So

$$|x_1 - x_2| \leq \sum_0^N |b_i - b_{i+1}| \leq 2\sqrt{3/2}\sum_0^N d_M(\gamma(a_i), \gamma(a_{i+1})) = 2\sqrt{3/2} \cdot d_M(p_1, p_2). \tag{3.8}$$

Combining (3.5) and (3.8) gives that, for any $x_1, x_2 \in \sigma$,

$$\frac{1}{2\sqrt{3/2}} \cdot |x_1 - x_2| \leq d_M(\pi^*(x_1), \pi^*(x_2)) \leq \sqrt{3/2} \cdot |x_1 - x_2|. \tag{3.9}$$

The geometry of $\pi^*(T)$. By the previous section, if e is an edge of the complex T, then

$$\frac{1}{2\sqrt{3/2}}\ell(e) \leq \ell(\pi^*(e)) \leq \sqrt{3/2}\ell(e). \tag{3.10}$$

Now let σ be an n-simplex of T. It follows from equation (2.2) that

$$\begin{aligned} \mathrm{Vol}_M(\pi^*(\sigma)) &= \int_{\pi^*(\sigma)} dV \\ &\geq \int_\sigma |D\pi^*|^n dV \\ &\geq \left(\frac{1}{\sqrt{3/2}}\right)^n \mathrm{Vol}_n(\sigma) \end{aligned} \tag{3.11}$$

Proof of Theorem 1. Let $\sigma \in \pi^*(T)$ be the simplex in M for which $\mathrm{Vol}_M(\sigma)$ is minimal among all simplices in $\pi^*(T)$. Let E be the length of the longest edge in $\pi^*(T)$, and let L be the length of the longest edge in σ. So there exist edges $e, l \in T$ such that $E = \ell(\pi^*(e))$ and $L = \ell(\pi^*(l))$. By equations (3.10) and (3.2) we have that

$$E = \ell(\pi^*(e)) \leq \sqrt{3/2}\ell(e) \leq \sqrt{3/2}\bar{L} \leq \sqrt{3/2}\frac{L}{C_{n,m}}$$

So,

$$E^n \leq (\sqrt{3/2})^n \frac{L^n}{C_{n,m}^n}. \tag{3.12}$$

By (3.9), we have that

$$L \leq \sqrt{3/2}\,\mathrm{diam}\left((\pi^*)^{-1}\sigma\right). \tag{3.13}$$

Using (3.12), (3.11), (3.13), and Theorem 3 (3) we obtain

$$\begin{aligned}
\frac{\mathrm{Vol}_M(\sigma)}{E^n} &\geq \left(\frac{C_{n,m}}{\sqrt{3/2}}\right)^n \frac{\mathrm{Vol}_M(\sigma)}{L^n} \\
&\geq \frac{C_{n,m}^n}{(\sqrt{3/2})^{2n}} \frac{\mathrm{Vol}_n(\pi^*)^{-1}(\sigma)}{L^n} \\
&\geq \frac{C_{n,m}^n}{(\sqrt{3/2})^{3n}} \frac{\mathrm{Vol}_n(\pi^*)^{-1}(\sigma)}{(\mathrm{diam}(\pi^*)^{-1}\sigma)^n} \\
&\geq \frac{C_{n,m}^n}{(\sqrt{27/8})^n}\Theta_{n,m}.
\end{aligned}$$

Since m depends only on n, we have proved Theorem 1, with the value

$$\delta_n = \frac{(3/2)^{\frac{3}{2}}}{C_{n,m}\Theta_{n,m}^{1/n}}.$$

4. Whitney's triangulation procedure

The goal of this Section is to convince the reader that Theorem 3 is true. This result was mentioned at the beginning of the previous Section 3, where it was used to prove Theorem 1. The proof of this Theorem follows almost directly due to the work of Whitney in ([HW], Ch. IV Part B, pg. 124-135). Unfortunately, Whitney's arguments are very technical and rather difficult to read. So in what follows we give a high-level sketch of Whitney's triangulation procedure, and then we prove Theorem 3 with specific references to all necessary equations in [HW].

We begin by using the smooth Nash isometric embedding theorem [Na] to isometrically embed M^n into $\mathbb{R}^m$, where m is a function of n. Define L_0 to be a cubical subdivision of $\mathbb{R}^m$ with cubes of side length h, and let L be the barycentric subdivision of L_0. Whitney recursively constructs a new triangulation of $\mathbb{R}^m$, L^*, whose $(m-n-1)$-skeleton is sufficiently far away from M.

Whitney then defines the simplicial complex K to be the poset of intersections of simplices of L^* of dimensions $(m-n), \ldots, m$ with M. For h small, this gives us a simplicial complex that sits inside a tubular neighborhood of M. Whitney then proves that the tubular neighborhood projection induces a diffeomorphism of K onto M.

This last remark is for the reader who attempts to tackle Whitney's work in [HW]. On pg. 133 - 134, Whitney defines complexes named K_p, L_p^*, and R_p. These are just small regions in either K or L^* near the point $p \in M$, and their only purpose is in proving that K is diffeomorphic to M.

Proof of Theorem 3. Let us first remind the reader that all Lemmas and equations in this proof reference Whitney's work in [HW].

Whitney proves that $\pi^* : K \to M$ is a diffeomorphism (pg. 134-135). Now, using Lemma 21a on pg. 132, equation (21.2) on pg. 132, and the fact that $4\lambda\xi < \lambda\xi/\alpha$ (since $0 < \alpha << 1$, see equations (17.2) on pg. 128 and (21.2) on pg. 132) gives that the simplicial complex in $\mathbb{R}^m$ whose vertices are $\pi^*(v)$ for each vertex v of K is still homeomorphic to M via π^*. Call this simplicial complex T. Then every simplex of T is a secant simplex of M. We have that every simplex of T has fullness at least $\Theta_1/2 := \Theta_{n,m}$, a number which depends only on n and m (see the proof of Lemma 21a on pg. 132, between equations (21.3) and (21.4). And Θ_1 is defined in equation (17.3) on pg. 128). This proves parts (1), (2), and (3) of Theorem 3.

Let $v \in T_q\sigma$ for $q \in \sigma$ with σ a simplex of T. Then

$$|\pi_{\pi^*(q)}(v)| \geq |v| - |v - \pi_{\pi^*(q)}(v)| \geq |v| - \frac{1}{2}|v| = \frac{1}{2}|v|$$

where the second inequality follows from the last inequality on pg. 132 (beginning with $|u - \pi_p u|$). This proves part (4) of the Theorem. Lastly, via the second-to-last equation on pg. 132 (beginning with $|p_i' - p_0'|$) and equation (21.2) on pg. 132, we have that

$$\frac{\beta\delta}{2} = \frac{b}{2} \leq \text{ length of an edge in } T \leq 2\delta + 8\lambda\xi \leq 3\delta.$$

If $\bar{L} = 3\delta$ and $C_{n,m} = \beta/6$, we have that $C_{n,m}$ depends only on n and m which proves part (5) of Theorem 3.

Let us note that β, δ, and b are defined on pg. 128-129 in equations (17.3), (17.5), and (17.6), respectively. The prameters λ and ξ are likewise defined in equations (17.4) and (17.5). The quantity β depends only on m and h, which in turn both only depend on n. ☐

5. Constructing a Riemannian metric which encodes a triangulation

This section is devoted to the proof of Theorem 2. We are given a smooth triangulated manifold $(M, \mathcal{T})$, and we would like to put a Riemannian metric on M whose geometry captures the combinatorics of the triangulation.

The proof is broken down into four parts. In the first part we will use a metric g_s which will, in general, have singularities, to produce a Riemannian metric g_δ which will depend on a parameter $\delta > 0$. The second step of the proof will be to show that we can choose δ small enough so that $|\mathrm{Vol}_{g_\delta}(M) - \mathrm{Vol}_{\mathcal{T}}(M)| < \varepsilon$. Letting $g = g_\delta$ completes the proof of property (1). We will then give a constructive method to homotope a closed g_s-polygonal path γ in M to a closed edge loop p in such a way that

$$\ell_{\mathcal{T}}(p) \leq \kappa_n \ell_{g_s}(\gamma).$$

Finally, we will argue that, once δ is small enough, the above inequality is preserved when considering closed geodesics in the metric g instead of polygonal paths in the metric g_s, completing the proof of property (2).

Constructing the Riemannian metric g_δ from the singular metric g_s. Let us begin by noting that the following construction of g_δ is pretty basic, but the authors are unaware of any references in the literature to such a construction.

The singular metric g_s is defined by simply requiring that every facet of $\mathcal{T}$ be isometric to an equilateral n-simplex in $\mathbb{E}^n$ with volume 1. Thus it is clear that $\mathrm{Vol}_{g_s}(M) = \mathrm{Vol}_{\mathcal{T}}(M)$. The singular set of g_s will be contained in the codimension two skeleton $\mathcal{T}^{(n-2)}$ of $\mathcal{T}$. A simplex $\sigma \in \mathcal{T}^{(n-2)}$ will be contained in the singular set of g_s if and only if "too many or too few" facets of $\mathcal{T}$ meet at σ. For example, if $n = 2$, a vertex v is contained in the singular set of g_s if and only if the number of 2-simplices of $\mathcal{T}$ which contain v is different from six.

To construct the Riemannian metric g_δ we need to alter g_s within small neighborhoods of simplices of $\mathcal{T}^{(n-2)}$, and then patch these metrics together using a smooth partition of unity. Let us first carefully construct this collection of neighborhoods which we will denote Ω. The construction is recursive with $n-1$ steps, in the l^{th} step we construct a collection of neighborhoods Ω_l with $\Omega_{l-1} \subset \Omega_l$ for $0 \le l \le n-2$. Then $\Omega := \Omega_{n-2}$.

First choose δ so that $0 < \delta < \frac{1}{3}$, and for each vertex v of $\mathcal{T}^{(n-2)}$ insert the open ball $b(v, \delta)$ into Ω_0. Note that, since $\delta < \frac{1}{3}$, each of these balls will be disjoint[1]. Next, let e be an edge of $\mathcal{T}^{(n-2)}$. Let $\bar{e} = e \setminus U$ where U is the union of all of the sets contained in Ω_0. Since e contains exactly two vertices, $\bar{e}$ is simply the interior of the edge e with a segment of length δ removed from each end. Let Ω_1 consist of all of the sets in Ω_0, as well as a set of the form $b(\bar{e}, k_1\delta)$ for each edge $e \in \mathcal{T}^{(1)}$, where k_1 is a small positive constant. For k_1 small enough, $b(\bar{e}, k_1\delta)$ will have nontrivial intersection with exactly two other members of Ω_1, the open δ balls about the vertices of e.

Defining the remaining collection of Ω_l recursively, let $\sigma \in \mathcal{T}^{(n-2)}$ be an l-dimensional simplex. Let $\bar{\sigma} = \sigma \setminus U$ where U is the union of all of the sets contained in Ω_{l-1}. Insert the open neighborhood $b(\bar{\sigma}, k_l\delta)$ into Ω_l where $k_l < k_{l-1}$ is a small positive constant. Also, let $\Omega_{l-1} \subset \Omega_l$. For k_l small enough, $b(\bar{\sigma}, k_l\delta)$ will have nontrivial intersection with exactly the members of Ω_l corresponding to the faces of σ. Letting $\Omega := \Omega_{n-2}$ completes the construction.

Let $\mathcal{O} = \bigcup_{U \in \Omega} U$ and let $\mathcal{U} = b(M \setminus \mathcal{O}, k_{n-1}\delta)$ be the open neighborhood of radius $k_{n-1}\delta$ (for some small positive constant $k_{n-1} < k_{n-2}$) about the closed set $M \setminus \mathcal{O}$. For k_{n-1} small enough, $\mathcal{U}$ will not meet any faces of $\mathcal{T}$ with codimension greater than or equal to two.

The collection $\Omega \cup \{\mathcal{U}\}$ forms an open cover of M. Since we are assuming that the smooth structure on M is compatible with the triangulation $\mathcal{T}$, we may define smooth PL coordinates within each neighborhood of Ω. Within $\mathcal{U}$, we can define smooth PL coordinates interior to each n-simplex of $\mathcal{T}$. Let $\{\phi_i\}$ be a smooth partition of unity subordinate to $\Omega \cup \{\mathcal{U}\}$.

Interior to each open set $U \in \Omega$ we will define a smooth metric g_U. Then the resulting Riemannian metric g_δ will be defined by

$$g_\delta = \phi_{\mathcal{U}} g_s + \sum_{U \in \Omega} \phi_U g_U.$$

Let $U \in \Omega$ be arbitrary. By our construction of Ω, U corresponds to some l-dimensional simplex σ. We define the metric g_U to simply be the pullback of the Euclidean metric under the smooth charts about σ constructed above. We will express this metric in generalized cylindrical coordinates about σ (which, if $l = 0$,

[1]If σ is an equilateral simplex with volume 1, then its edge lengths are an increasing function of its dimension n, hence the edge lengths will always be $\geq \frac{2}{\sqrt[4]{3}} > 1$.

would just be generalized spherical coordinates). More specifically, if we denote the coordinates by $x_1, \ldots, x_l, r, \theta_1, ..., \theta_{n-l-1}$, then

$$g_U = \sum_{i=1}^{l} dx_i^2 + dr^2 + r^2 d\theta_1^2 + \sum_{i=2}^{n-l-1} r^2 \sin^2(\theta_1)... \sin^2(\theta_{i-1}) d\theta_i^2$$

where the last sum is void if $l = n-2$, and where the domains for the coordinates are:

$$-\infty < x_i < \infty \qquad 0 \leq r < \infty \qquad 0 \leq \theta_j \leq \pi \ \ (j \neq n-l-1) \qquad 0 \leq \theta_{n-l-1} < 2\pi$$

Choosing δ so that $|\mathbf{Vol}_{g_\delta}(M) - \mathbf{Vol}_{\mathcal{T}}(M)| < \varepsilon$. To avoid overcomplicating this proof we will proceed as follows. The simplest case where $n = 2$ will be carried out in full detail. We then move on to the case where $n = 3$ and work out in full detail the part that differs from the $n = 2$ case. The $n = 3$ case captures the general behavior of the problem, and so it will be easy to explain from here how the result follows for general n.

Case $n = 2$: In this case, $\mathcal{T}^{(n-2)}$ is just the vertex set of $\mathcal{T}$ and thus Ω is a collection of pairwise disjoint subsets[2] of M. Let $v \in \mathcal{T}^{(n-2)}$ be a vertex with corresponding open set $V \in \Omega$. The metric g_δ differs from g_s only in such neighborhoods V, and in V the metric g_δ has the form

$$g_\delta = \phi_V g_V + \phi_{\mathcal{U}} g_s.$$

Recall that we expressed g_V in polar coordinates. i.e., $g_V = dr^2 + r^2 d\theta^2$ where $0 \leq r < \delta$ and $0 \leq \theta < 2\pi$. We can locally express g_s in Cartesian coordinates by $g_s = dx_1^2 + dx_2^2$. To compute $|\mathrm{Vol}_{g_s}(V) - \mathrm{Vol}_{g_\delta}(V)|$, we need to convert g_s to polar coordinates within V. This is done as in any multivariable calculus course by setting $x_1 = \bar{r}\cos(\bar{\theta})$ and $x_2 = \bar{r}\sin(\bar{\theta})$. But notice that the domains for these parameters are $0 \leq \bar{r} < \delta$ and $0 \leq \bar{\theta} < \frac{\pi}{3} t_v$ where t_v is the number of facets (in this case, triangles) containing v. We then obtain that $0 \leq \frac{6}{t_v}\bar{\theta} < 2\pi$, and so $\bar{r} = r$ and $\bar{\theta} = \frac{t_v}{6}\theta$. Substituting these values and computing the differentials yields

$$g_s = dr^2 + \left(\frac{t_v}{6}\right)^2 r^2 d\theta^2$$

and thus

$$g_\delta = \phi_V g_V + \phi_{\mathcal{U}} g_s = dr^2 + \left(\phi_V + \left(\frac{t_v}{6}\right)^2 \phi_{\mathcal{U}}\right) r^2 d\theta^2.$$

[2]This is one of the two main differences between the $n = 2$ and the higher dimensional cases. The other difference is that the triangulation automatically has high regularity. Vertex links are always $\mathcal{S}^1$ (hence the triangulation is PL), and points are always smoothly embedded in M.

We then compute the change in volume, $|\mathrm{Vol}_{g_\delta}(V) - \mathrm{Vol}_{g_s}(V)|$, to equal

$$\begin{aligned}
&= \left| \int_0^{2\pi} \int_0^{\delta} \sqrt{\left(\phi_V + \left(\frac{t_v}{6}\right)^2 \phi_{\mathcal{U}}\right)} r^2 dr d\theta - \int_0^{2\pi} \int_0^{\delta} \frac{t_v}{6} r dr d\theta \right| \\
&= \left| \int_0^{2\pi} \int_0^{\delta} \left(\sqrt{\left(\phi_V + \left(\frac{t_v}{6}\right)^2 \phi_{\mathcal{U}}\right)} - \frac{t_v}{6} \right) r dr d\theta \right| \\
&\le \left| \sqrt{1 + \left(\frac{t_v}{6}\right)^2} - \frac{t_v}{6} \right| \int_0^{2\pi} \int_0^{\delta} r dr d\theta \\
&= \left| \sqrt{1 + \left(\frac{t_v}{6}\right)^2} - \frac{t_v}{6} \right| \pi\delta^2
\end{aligned}$$

which approaches 0 as δ approaches 0. Since M is compact there are only finitely many vertices, completing the proof when $n = 2$.

Case $n = 3$: When $n = 3$, $\mathcal{T}^{(n-2)}$ is now the 1-skeleton of $\mathcal{T}$. So the neighborhoods in Ω are no longer disjoint. Let v be a vertex of $\mathcal{T}$ and let e be an edge containing v. Denote their corresponding neighborhoods in Ω by U_v and U_e, respectively. The same argument as in the 2-dimensional case applies to the regions of M where only U_v or U_e intersects $\mathcal{U}$. What we need to show is that δ can be chosen small enough so that $|\mathrm{Vol}_{g_\delta}(W) - \mathrm{Vol}_{\mathcal{T}}(W)| < \varepsilon$ with $W = \mathcal{U} \cap U_v \cap U_e$. In what follows we adapt the notation $g_v := g_{U_v}$ and $g_e := g_{U_e}$.

In W, g_v is written in spherical coordinates $g_v = (dr^v)^2 + (r^v)^2(d\theta_1^v)^2 + (r^v)^2 \sin^2(\theta_1^v)(d\theta_2^v)^2$ (where the v is emphasized in the coordinates to distinguish from the coordinates of g_e). The domains of the parameters depend on δ and k_1, but we will not get too caught up in those details here. To compute $\mathrm{Vol}_{g_\delta}(W)$ we convert both $g_{\mathcal{U}}$ and g_e to spherical coordinates. In exactly the same way as when $n = 2$ we have that $g_{\mathcal{U}} = (dr^v)^2 + C_v^2(r^v)^2(d\theta_1^v)^2 + C_e^2(r^v)^2 \sin^2(\theta_1)(d\theta_2^v)^2$ where C_v and C_e are positive constants that depend on the number of facets of $\mathcal{T}$ that contain v and e, respectively.[3]

We need to convert g_e into spherical coordinates within W. Recall that g_e is written in cylindrical coordinates within U_e, i.e. $g_e = dx_1^2 + (dr^e)^2 + (r^e)^2(d\theta_1^e)^2$. Notice that by an orthogonal transformation within U_v, we may align the axis orthogonal to θ_2^v with the edge e (see Figure 1). With this choice of coordinates we see that $\theta_2^v = \theta_1^e$. The domains for these two parameters may differ, but we can get an overestimate for the volume of W by integrating over the larger of the two domains. Also notice that

$$\begin{aligned}
x^1 &= r^v \cos(C_v\theta_1^v) \\
r^e &= r^v \sin(C_v\theta_1^v)
\end{aligned}$$

where the constant C_v is the same constant as in $g_{\mathcal{U}}$ and is due to the change in the domain of θ_1^v between g_δ and g_s, exactly as in the 2-dimensional case where the constant was $\frac{t_v}{6}$. Computing the differential then yields that $g_e = (dr^v)^2 +$

[3]In fact, just as in the 2-dimensional case, $C_e = \frac{t_e}{6}$ where t_e is the number of facets which contain e. To compute C_v, one needs to compute the *solid angle* φ subtended by the three edges emanating from v. Then $C_v = \frac{2\pi}{\varphi t_v}$, where t_v is the number of tetrahedra containing v.

$C_v^2(r^v)^2(d\theta_1^v)^2 + (r^v)^2\sin^2(\theta_1^v)(d\theta_2^v)^2$. Thus the difference in volume, $|\mathrm{Vol}_{g_\delta}(W) - \mathrm{Vol}_{g_s}(W)|$, equals

$$
\begin{aligned}
&= \left| \iiint_W \sqrt{(\phi_v + \phi_e + \phi_{\mathcal{U}} C_v^2)\,(\phi_v + \phi_e C_e^2 + \phi_U C_v^2)}(r^v)^2 \sin(\theta_1^v) dV \right. \\
&\left. - \iiint_W C_v C_e (r^v)^2 \sin(\theta_1^v) dV \right| \\
&\leq \left| \sqrt{(1 + 1 + C_v^2)\,(1 + C_e^2 + C_v^2)} - C_v C_e \right| \iiint_W (r^v)^2 \sin(\theta_1^v) dV \\
&= \left| \sqrt{(1 + 1 + C_v^2)\,(1 + C_e^2 + C_v^2)} - C_v C_e \right| \mathrm{Vol}_{g_v}(W).
\end{aligned}
$$

It is clear that $\mathrm{Vol}_{g_v}(W)$ goes to 0 as either δ or k_1 approaches 0 (recall that k_1 was introduced in the construction of Ω). This completes the proof when $n = 3$.

Case $n > 3$: Once $n > 3$ we must deal with intersections of three or more sets in Ω. But we always change coordinates to those of the simplex of least dimension, just as we did above when changing all coordinates to those of U_v. Then since we are using the same coordinates in each chart, we need only consider the intersection of two neighborhoods (one of which being that lowest dimensional simplex). Denote these simplices by σ and τ with $\dim(\sigma)=j$, $\dim(\tau)=l$, $j < l$. Let U_σ and U_τ denote their corresponding neighborhoods in Ω, and denote their corresponding metrics by g_σ and g_τ. All we need to show is that we can change the coordinates in τ to coordinates in σ in such a way so that each component of the metric g_τ only differs from the corresponding component of g_σ by a constant. Notice that if $j > 0$ then we can project out the coordinates corresponding to the simplex σ in both U_σ and U_τ. So we may assume that $j = 0$, σ is a vertex, and thus the coordinates in U_σ

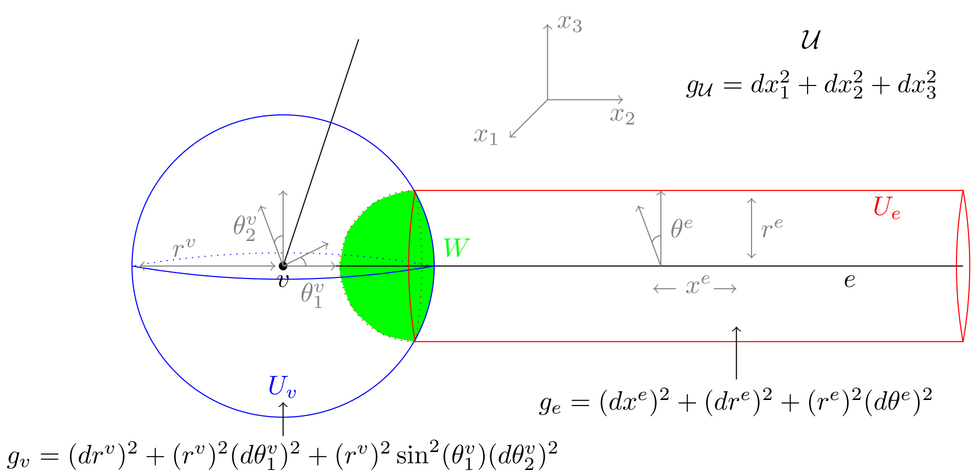

FIGURE 1. The region W is shaded in green, and the three different coordinate charts are written in gray. Notice that, by an orthogonal change of coordinates within U_v, we have aligned the respective axes so that $\theta^e = \theta_2^v$.

are generalized spherical coordinates

$$g_\sigma = d\rho^2 + \rho^2 d\theta_1^2 + \sum_{i=2}^{n-1} \rho^2 \sin^2(\theta_1) ... \sin^2(\theta_{i-1}) d\theta_i^2$$

where $0 \leq \rho < \infty$, $0 \leq \theta_i \leq \pi$ for $1 \leq i \leq n-2$, and $0 \leq \theta_{n-1} < 2\pi$. In U_τ the metric is written in generalized cylindrical coordinates by

$$g_\tau = \sum_{i=1}^{l} dx_i^2 + dr^2 + r^2 d\bar{\theta_1}^2 + \sum_{i=2}^{n-l-1} r^2 \sin^2(\bar{\theta}_1) ... \sin^2(\bar{\theta}_{i-1}) d\bar{\theta_i}^2$$

where $0 \leq r < \infty$, $0 \leq \bar{\theta}_i \leq \pi$ for $1 \leq i \leq n-l-2$, and $0 \leq \bar{\theta}_{n-l-1} < 2\pi$.

Exactly as in the $n = 3$ case, we can rotate the coordinates in U_σ so that $\theta_{i+l} = \bar{\theta}_i$ for $1 \leq i \leq n-l-1$. That still leaves $(n-1)-(n-l-1) = l$ directions in which the coordinates in U_σ can be rotated. Each of the parameters $\theta_1, \ldots, \theta_l$ in U_σ measures an angle from a fixed positive axis. We also rotate the coordinates in U_σ so that the (positive) axis associated with θ_i corresponds to the coordinate x_i in U_τ. We then have that $x_i = \rho \cos(C_i \theta_i)$ where the constant C_i arises from converting the domains of the associated variables in exactly the same way as the 2-dimensional case. By projecting out ρ along each of the dimensions $x_1, ..., x_l$ we can write r as the product $r = \rho \sin(C_1\theta_1) ... \sin(C_l\theta_l)$. Thus, this change in coordinates only differs from the standard (Euclidean) change in coordinates by (possibly) multiplying each variable by a constant factor, and thus the coordinates of the corresponding metric only differ by a constant.

Remark. In the construction of the metrics g_δ, our metrics are changed by stretching or compressing in the radial directions about each simplex – with the amount of distortion determined by the combinatorics of the link of the simplex. It follows that there exists a constant $C_\mathcal{T}$ (depending solely on the triangulation $\mathcal{T}$) with the property that the identity map from (M, g_s) to (M, g_δ) is $C_\mathcal{T}$-Lipschitz.

Lipschitz homotopies of closed paths to closed edge loops. For now on in this Section we refer to g_δ as simply g. First, note that Theorem 2 (2) is trivial for *any* constant if the closed path is null-homotopic. Also, using the *Birkhoff curve shortening process*, a description of which can be found in [Kl], we can homotope any closed path η to a closed geodesic γ in such a way that $\ell_g(\gamma) \leq \ell_g(\eta)$. If we can then find an edge path p satisfying Theorem 2 (2) with respect to γ, then

$$\ell_d(p) \leq \kappa_n \ell_g(\gamma) \leq \kappa_n \ell_g(\eta)$$

where $\ell_d(p)$ denotes the discrete length of p (so, the length of p with respect to $\mathcal{T}$). So we may reduce to the case that γ is a closed g-geodesic which is not null-homotopic.

The following Lemma and Corollary handle the case when γ is a g_s-geodesic.

Lemma 4. *Let Δ be a Euclidean equilateral n-simplex with volume 1, and let α be a straight line segment whose endpoints lie on the boundary of Δ. Then α can be homotoped, rel. endpoints, to a path $\alpha^{(1)} \subseteq \partial\Delta$ such that*

$$\ell(\alpha^{(1)}) \leq C_n \ell(\alpha)$$

where the constant C_n depends only on n.

Proof. First, at the cost of a very slight perturbation, we may assume that α misses the barycenter B of the n-simplex Δ. Now, consider radial projection from B to $\partial\Delta$. Note that this map is *not* Lipschitz – a small segment near the barycenter will get stretched out to a long segment on the boundary. But it *is* C'-Lipschitz if one restricts to segments that are at least some fixed (uniform) distance D away from the barycenter. Finally, if one considers the segments that pass closer than D to the barycenter, their lengths are uniformly bounded below, while the length of their projected images are uniformly bounded above. So again, their is some constant C'' so that $\ell(\alpha^{(1)}) \leq C''\ell(\alpha)$. Letting $C_n := \max\{C', C''\}$ completes the proof. $\square$

Let $\kappa'' := C_n \cdot C_{n-1} \cdot \ldots \cdot C_2$. Define a *$g_s$-polygonal path* to be a path which is a g_s-geodesic (i.e., a straight line) when restricted to any simplex of $\mathcal{T}$. Also, define the *support* of any path to be the collection of facets that it intersects. Recursively applying Lemma 4 proves:

Corollary 5. *Let α be a closed g_s-polygonal path in $(M, \mathcal{T})$. Then there exists an edge path p, freely homotopic to α and with the supp(α) $\subset$ supp(p), such that*

$$e_n \ell_d(p) \leq \kappa'' \ell_{g_s}(\alpha) \tag{5.1}$$

where e_n is the length of an edge of an equilateral n-simplex with volume 1*.*

Proof. Let α be a g_s-polygonal path. We inductively apply Lemma 4 to push our path from the k-skeleton to the $(k-1)$-skeleton. At each stage, we replace a path that is straight in each k-simplex by a path lying in the $(k-1)$-skeleton, and has the property that it is at most C_k times the original length. The path in the $(k-1)$-skeleton may no longer be straight on each $(k-1)$-simplex, but one can straighten it on each of the simplices – this only decreases the length of p which does not effect inequality (5.1) – and then reapply the Lemma. Note that points of α within any simplex stay within the boundary of that simplex throughout this procedure. So the support only grows throughout this process. Finally, we end up with a loop p in the 1-skeleton, homotopic to the original loop, and satisfying inequality (5.1). $\square$

This Corollary proves Theorem 2 (2) for geodesics in the metric g_s instead of g. Intuitively, for $\delta > 0$ very small, geodesics should not differ much in the metrics g and g_s. But this takes a little work to show directly. So in what follows we take an arbitrary g-geodesic and reduce to the case of a g_s-polygonal path. The reader who believes that such a reduction is possible may skip ahead to Section 6.

Note that $e_n \geq 1$ for all $n \geq 1$, and therefore the above Corollary gives that $\ell_d(p) \leq \kappa'' \ell_{g_s}(\alpha)$. Define $\kappa' := 3\kappa''$. Then, using the notation of the above Corollary, we have that

$$\ell_d(p) \leq e_n \ell_d(p) < 3e_n \ell_d(p) \leq \kappa' \ell_{g_s}(\alpha). \tag{5.2}$$

The following Lemma allows us to apply the preceeding Corollary to closed g-geodesics which do not intersect any simplex of $\mathcal{T}$ "too many times".

Lemma 6. *Let γ denote a closed g-polygonal path in $(M, \mathcal{T})$, and let $K > 0$ be some fixed constant. Suppose that:*

(1) *γ is not null-homotopic*

(2) *For all $\sigma \in \mathcal{T}$, $\gamma \cap \sigma$ consists of at most K g-geodesic segments (counted independently in the event that γ has self-intersections).*

Then there exists a closed g_s-polygonal path α, freely homotopic to γ, such that

$$\frac{1}{2}\ell_{g_s}(\alpha) \leq \ell_g(\gamma)$$

In Lemma 6 a *g-polygonal path* is a path which is a g-geodesic when restricted to any simplex of $\mathcal{T}$.

Proof. Let γ be a closed g-polygonal path. If γ does not intersect the δ-neighborhood of $\mathcal{T}^{(n-2)}$ then γ is a polygonal path in the g_s metric and we are done. So assume that γ intersects the δ-neighborhood of $\mathcal{T}^{(n-2)}$, denoted b_δ. Let $\beta_1, \ldots, \beta_k$ denote the connected components of $\gamma \cap b_\delta$.

Consider one of these components β_i. Let $\tau_1, \ldots, \tau_l$ denote the simplices of $\mathcal{T}$ with codimension 2 or greater for which β_i intersects the corresponding warped neighborhood U_{τ_i}.

Recall that, when defining the metric g, we altered the g_s metric about the $k_i\delta$ neighborhood of each i-dimensional face (and $i \leq n-2$). Then, via a compactness argument, we can systematically choose $k_1, \ldots, k_{n-2}$ small enough so that $b_g(x, 2\delta) \setminus b_\delta$ is non-empty for each point x in the $(n-2)$-skeleton of $\mathcal{T}$. The point here is that we can choose the region $\mathcal{O}$ in which we are altering the g_s-metric small enough so that the open 2δ ball about any point of M, measured in the g-metric, contains points outside of $\mathcal{O}$.

Now, we want to remove the interior of β_i and replace it with a g_s-polygonal path, denoted α_i, between its endpoints which stays within a 4δ neighborhood of β_i. We can always find such a path α_i so that

$$\ell_{g_s}(\alpha_i) \leq \ell_g(\beta_i) + 8\delta|\mathcal{T}_i| \tag{5.3}$$

where $\mathcal{T}_i$ denotes the collection of all simplices of $\mathcal{T}$ which contain any of $\tau_1, \ldots, \tau_l$. To see this, just subdivide β_i where it intersects different simplices of $\mathcal{T}$. Sequentially approximate each of these points by a point outside of b_δ at a distance of at most 4δ away (whose existence is guaranteed by the preceding paragraph), and then connect each of these points by a g_s-polygonal path. Note that we could always find a point outside of b_δ at a distance of 2δ away, but we use 4δ above to ensure that we can choose points that can be connected by a g_s-polygonal path outside of b_δ. Please see Figure 2 for a sketch of how α_i is obtained from β_i. Equation (5.3) then follows from repeated application of the triangle inequality. Note that inequality (5.3) is a very crude estimate. In general, one would not need anywhere near $|\mathcal{T}_i|$ polygonal pieces in any such approximation.

Now consider the polygonal path, denoted by α, obtained by replacing each β_i with its corresponding polygonal approximation α_i. Then one sees immediately that

$$\ell_{g_s}(\alpha) \leq \ell_g(\gamma) + 8\delta \sum_{i=1}^{k} |\mathcal{T}_i|. \tag{5.4}$$

Due to assumption (2) we have that

$$\sum_{i=1}^{k} |\mathcal{T}_i| \leq \mu K ||\mathcal{T}|| \tag{5.5}$$

where $||\mathcal{T}||$ denotes the total number of simplices contained in $\mathcal{T}$, and μ denotes the maximal degree of any simplex of $\mathcal{T}$ (i.e., $\mu = \max\{||St(\sigma)|| : \sigma \in \mathcal{T}\}$ where $St(\sigma)$ denotes the *closed star* of the simplex σ).

Now, choose

$$\delta < \frac{1}{16\mu K||\mathcal{T}||}\mathrm{sys}_{g_s}(M) \tag{5.6}$$

where $\mathrm{sys}_{g_s}(M)$ denotes the systole of M with respect to the metric g_s.

Combining inequalities (5.4), (5.5), and (5.6) yields:

$$\begin{aligned}\ell_g(\gamma) &\geq \ell_{g_s}(\alpha) - 8\delta\sum_{i=1}^{k}|\mathcal{T}_i| \\ &\geq \ell_{g_s}(\alpha) - 8\delta\mu K||\mathcal{T}|| \\ &\geq \ell_{g_s}(\alpha) - \frac{1}{2}\mathrm{sys}_{g_s}(M) \\ &\geq \frac{1}{2}\ell_{g_s}(\alpha)\end{aligned}$$

and where, for the last inequality, it is necessary that α is not null-homotopic. $\square$

We can now complete the proof of Theorem 2 with $\kappa_n = 2\kappa'$ for g-geodesics which satisfy the conditions of Lemma 6. Let γ be such a g-geodesic. Let α be the g_s-polygonal path guaranteed by Lemma 6, and let p be the edge path from the Corollary corresponding to α. Then

$$\ell_d(p) \leq \kappa'\ell_{g_s}(\alpha) \leq 2\kappa'\ell_g(\gamma) = \kappa_n\ell_g(\gamma).$$

In order to complete the proof of Theorem 2, we need to fix some $K > 0$ and reduce to the case of g-polygonal paths which intersect each simplex of $\mathcal{T}$ at most K times.

In order to define K, let us first define

$$D := \frac{e_n}{\kappa'} \quad \Longrightarrow \quad e_n = D\kappa' \tag{5.7}$$

where e_n and κ' are as in equation (5.2). Note that D depends only on n.

We now define K as follows. Cover the $(n-2)$-skeleton $\mathcal{T}^{(n-2)}$ of $\mathcal{T}$ with a finite number of open $\frac{1}{8}D$-balls (in the g_s metric). Then extend this cover to an open cover of $\mathcal{T}^{(n-1)}$ by open $\frac{1}{8}D$-balls; denote by $\{p_1, \ldots, p_N\}$ the points where these balls are centered. K is then the maximal number of open sets in this covering required to cover the boundary of any simplex of $\mathcal{T}$. Note that, since the Riemannian manifold (M, g) converges to the geodesic metric space (M, g_s) in the Gromov-Hausdorff sense (as δ approaches 0), we can choose $\delta > 0$ sufficiently small so that the collection of g-metric open $\frac{1}{8}D$-balls, centered at the same collection of points $\{p_1, \ldots, p_N\}$, also forms a cover of $\mathcal{T}^{(n-1)}$ – call this open cover $\mathcal{U}$. Also note that there is no ambiguity with equation (5.6), as K is fixed and then we choose δ.

Remark. Let $U, V \in \mathcal{U}$ be such that $U \cap V \neq \emptyset$. Then by the above construction, $\mathrm{diam}_g(U \cup V) \leq \frac{1}{2}D$. Let $\sigma \in \mathcal{T}$ be a simplex that intersects both U and V. Let $x, y \in \sigma \cap U \cap V$, and let γ be a g-geodesic joining x and y (so, in particular, $\ell_g(\gamma) \leq \frac{1}{2}D$). A priori, γ could weave in and out of σ. But by choosing δ small enough, we can ensure that any such points x and y can be connected by a path of

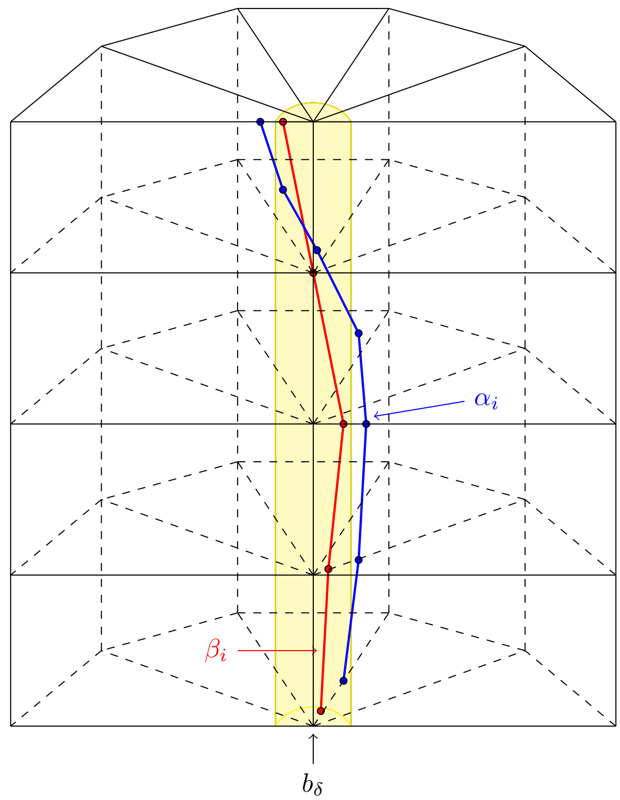

FIGURE 2. Schematic picture for how the g_s-polygonal path α_i is obtained from the g-polygonal path β_i in the proof of Lemma 6. Note that these "wedges" are not actual simplices, but one can add appropriate triangles to the above picture and subdivide the paths accordingly.

length less than D which is a geodesic of $g|_\sigma$. Note that this may not be an actual g-geodesic.

Proof of Theorem 2. Let γ denote a closed g-polygonal path, and let K be as above. Suppose that γ is not null-homotopic, but does not satisfy condition (2) of Lemma 6. Let us assume that $\sigma \in \mathcal{T}$ is the only simplex for which $\gamma \cap \sigma$ consists of more than K connected components, and that $\gamma \cap \sigma$ has exactly $K+1$ components. The procedure described below can be iterated (see Step 4 below) to deal with multiple simplices and/or for a greater intersection number with any simplex.

Fix a base point and orientation of $\mathbb{S}^1$. Using this orientation, each component of $\gamma \cap \sigma$ has an "entrance point" x_i and an "exit point" y_i. Let $x_1, \ldots, x_{K+1}$ denote the $K+1$ entrance points, and let $y_1, \ldots, y_{K+1}$ denote the $K+1$ exit points. By the definition of K, there must exist two entrance points x_i and x_j such that

$$d_g(x_i, x_j) < D \quad \overset{\text{Eqn(5.7)}}{\Longrightarrow} \quad \kappa' d_g(x_i, x_j) < e_n \tag{5.8}$$

Let y_i and y_j denote the corresponding exit points, and assume that $i < j$.

Remove the segments (x_i, y_i) and (x_j, y_j) from γ, and insert g-geodesics (x_i, y_j) and (x_j, y_i) interior to σ. Note that these newly inserted paths are geodesics of $g|_\sigma$, not necessarily of g. By the above Remark, we know that these two segments have length less than D. Denote the closed g-polygonal path containing (x_i, y_j) by γ_1, and the other by γ_2. We may assume that neither path is null-homotopic, for otherwise we could have homotoped the original path interior to σ and reduced the number of components of $\gamma \cap \sigma$. So by Lemma 6, the Corollary, and equation (5.2) there exist closed edge paths p_1 and p_2 freely homotopic to γ_1 and γ_2 such that

$$3e_n\ell_d(p_1) \leq \kappa'\ell_g(\gamma_1) \quad \text{and} \quad 3e_n\ell_d(p_2) \leq \kappa'\ell_g(\gamma_2). \tag{5.9}$$

Now under the free homotopy of γ_1 into the edge path p_1, the points x_i, y_j find themselves lying on the 1-skeleton of σ. Let v_i and w_j be the vertices of σ which are closest to the image of x_i, y_j respectively. Similarly, the free homotopy of γ_2 into the edge path p_2 moves x_j, y_i into the 1-skeleton of σ, and we let v_j, w_i be the vertices of σ closest to these image points. By possibly shortening the paths if necessary, we may assume that the edge $\overline{v_i w_j} \in p_1$ and the edge $\overline{v_j w_i} \in p_2$. Note that it is entirely possible that one (or both) of these edges is degenerate.

We want to simultaneously

- use p_1 and p_2 to construct a closed path p that is freely homotopic to γ.
- Reconstruct γ from γ_1 and γ_2.
- Preserve the key inequality $\ell_d(p) \leq \kappa'\ell_g(\gamma)$.
- Do all of this in a manner which can be iterated.

We will do this in four steps. In steps 2 and 3, we will need to assume that the edges $\overline{v_i w_j}, \overline{v_j w_i}$ are non-degenerate. However, in step 4, we will explain how to allow for degenerate edges.

Step 1: Append edges $\overline{v_i w_i}$ and $\overline{v_j w_j}$ between p_1 and p_2.

Let $p'' := p_1 \cup p_2 \cup \overline{v_i w_i} \cup \overline{v_j w_j}$, and let $\gamma'' := \gamma_1 \cup \gamma_2$ (see Figure 3 below for a schematic illustration). It is important to note that p'' is no longer a path. The image of p'' is (essentially) three closed loops which are glued together along the edges $\overline{v_i w_j}$ and $\overline{v_j w_i}$. We will still talk about the "length" of p'' even though we should probably use something like "one-dimensional Hausdorff measure". And we will still use the notation $\ell_d(p'')$ for this measure. Similar considerations will also

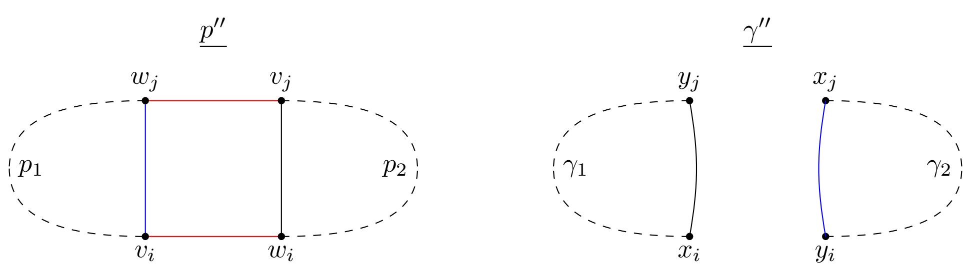

FIGURE 3. Schematic picture for Step 1. In red is what was added in this step, and in blue is what will be deleted in Step 2. Note that this picture is just to keep track of what is going on. It is not geometrically accurate, as x_i is metrically close to x_j in the proof.

apply in Steps 2 and 3. But the point is that at the end of this procedure, we will again have a closed path.

Now let us check that the analogue of the key inequality still holds. This is a consequence of the following series of inequalities:

$$\begin{aligned}
2e_n\ell_d(p'') &\le 2e_n(\ell_d(p_1)+\ell_d(p_2)+2)\\
&= 2e_n(\ell_d(p_1)+1)+2e_n(\ell_d(p_2)+1)\\
&\le 2e_n\left(\ell_d(p_1)+\frac{1}{3}\ell_d(p_1)\right)+2e_n\left(\ell_d(p_2)+\frac{1}{3}\ell_d(p_2)\right)\\
&< 3e_n\ell_d(p_1)+3e_n\ell_d(p_2)\\
&\le \kappa'\ell_g(\gamma_1)+\kappa'\ell_g(\gamma_2)\\
&= \kappa'\ell_g(\gamma'').
\end{aligned}$$

The second inequality holds because both p_1 and p_2 are not null-homotopic, so must consist of at least three edges. The last inequality is due to the "3" present in equation (5.9). Note that, in the event that one of the edges $\overline{v_iw_i}, \overline{v_jw_j}$ is degenerate, this series of inequalities still holds (the only effect is that the first inequality in the chain becomes strict).

Step 2: Insert (x_i, y_i) and remove (x_j, y_i) from γ'', and remove the edge $\overline{v_iw_j}$ from p''.

Let p' and γ' denote the new sets created from the above procedures (see Figure 4 below). Notice that, by the triangle inequality and equation (5.8) we get

$$d_g(x_j, y_i) \le d_g(x_i, y_i) + d_g(x_i, x_j)$$

$$\Longrightarrow \quad -D \le -d_g(x_i,x_j) \le d_g(x_i,y_i) - d_g(x_j,y_i). \tag{5.10}$$

Then by equations (5.10) and (5.7) we obtain

$$2e_n\ell_d(p') = 2e_n\ell_d(p'') - 2e_n \tag{5.11}$$

$$\kappa'\ell(\gamma') = \kappa'[\ell(\gamma'') + d_g(x_i,y_i) - d_g(x_j,y_i)] \ge \kappa'[\ell(\gamma'') - D] = \kappa'\ell(\gamma'') - e_n$$

and so

$$2e_n\ell_d(p') = 2e_n\ell_d(p'') - 2e_n \le \kappa'\ell_g(\gamma'') - e_n \le \kappa'\ell_g(\gamma'). \tag{5.12}$$

Notice that here, it is important that the edge $\overline{v_iw_j}$ is non-degenerate. If it were degenerate, equation (5.11) would not hold, and thus neither would (5.12).

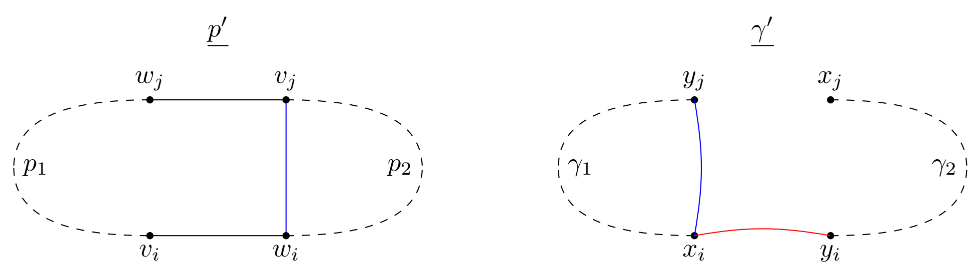

FIGURE 4. Schematic picture for Step 2. Again, what is new is in red, and what will be removed in Step 3 is in blue.

Step 3: Insert (x_j, y_j) and remove (x_i, y_j) from γ', and remove the edge $\overline{v_j w_i}$ from p'.

The first two operations will return our original path γ, and the second will provide a closed edge path p. The path p is freely homotopic to γ since all of our operations occurred within the closed simplex σ. By the triangle inequality, we have that

$$d_g(x_i, y_j) \leq d_g(x_j, y_j) + d_g(x_i, x_j)$$

and the exact same argument as for equation (5.12) proves that

$$\ell_d(p) < 2e_n \ell_d(p) \leq \kappa' \ell(\gamma) \tag{5.13}$$

Again, in this step, one needs the edge $\overline{v_j w_i}$ to be non-degenerate.

Step 4: A few remarks to ensure that this process iterates.

The "3" in equation (5.2) means that, at each vertex of both p_1 and p_2, we can append two edges to obtain new sets p_1' and p_2' which still satisfy that $\ell_d(p_1') \leq \kappa' \ell_g(\gamma_1)$ and $\ell_d(p_2') \leq \kappa' \ell_g(\gamma_2)$

The reason that we need this multiple of three is because any of the vertices v_i, v_j, w_i, and/or w_j could be the same. As already mentioned in Steps 2 and 3, the proofs for inequalities (5.12) and (5.13) do not hold without inequality (5.11). But for this inequality to hold, we must have an edge to delete. This edge may not exist if $v_i = w_j$ and/or $v_j = w_i$, and so we may need an additional edge built in for these steps as sort of an "extraneous edge" that we can delete in order to preserve inequality (5.11).

Since the "3" in equation (5.2) is multiplicative, we can glue in these two additional edges at *every* vertex. Thus, we always have these edges available to us wherever we cut γ into two closed curves γ_1 and γ_2.

□

6. Filling triangulated surfaces

Recall that, given a closed triangulated n-dimensional manifold $(M, \mathcal{T}_M)$, a *filling* of M is a triangulated $(n+1)$-dimensional manifold $(N, \mathcal{T}_N)$ with $\partial N = M$ and $\mathcal{T}_N|_{\partial N} = \mathcal{T}_M$. In this definition, one also requires that the homeomorphism from $|\mathcal{T}_N|$ to N restricts to the homeomorphism from $|\mathcal{T}_M|$ to M along its boundary. Sometimes we refer to a filling of $(M, \mathcal{T}_M)$ as an *extension to a triangulation* of $(N, \mathcal{T}_N)$. A basic question is the following. Given a triangulated manifold $(M, \mathcal{T}_M)$, does such a filling exist and, if so, can you bound $|\mathcal{T}_N|$, the number of facets of such a filling? Theorem 2 leads to the following two solutions to this question in the case when $n = 2$.

Theorem 7. *Let $(M, \mathcal{T}_M)$ be a triangulated surface of genus $\leq g$. Then there exists a filling $(N, \mathcal{T}_N)$ satisfying that*

$$|\mathcal{T}_N| \leq C_g |\mathcal{T}_M|,$$

where C_g depends only on g, and not on the particular surface or triangulation.

Theorem 8. *Let $(M, \mathcal{T}_M)$ be a triangulated surface. Then there exists $(N, \mathcal{T}_N)$, a filling of M, so that*

$$|\mathcal{T}_N| \leq C|\mathcal{T}_M| \left(\log|\mathcal{T}_M|\right)^2 ,$$

where C does not depend on the particular surface or triangulation.

Remark. For all of Section 6, g will refer to the genus of the surface, not a Riemannian metric.

The proofs of both of these theorems are very similar. We first combine Theorem 2 with results of Gromov in [Gr1] and [Gr2] to bound the discrete systole of $(M, \mathcal{T}_M)$ by a factor of the combinatorial volume of $(M, \mathcal{T}_M)$. Then the main idea is to apply a "cut-and-cone" procedure. We begin this procedure by cutting the surface along a short homologically nontrivial edge loop. This will yield a surface of smaller genus with two boundary components. We then cone off the boundary components to get a surface of genus one less than the original surface (See Figure 5). We iterate this procedure until the surface is a 2-sphere, in which case we perform a modified coning-off procedure to get a triangulated 3-ball. By gluing the 3-ball along all of the cuts in the reverse order, we obtain a triangulated 3-manifold with the desired properties.

Remark. The argument for our proofs "builds" the bounding 3-manifold from the triangulation on Σ. One might wonder whether this is really necessary. Indeed, if one takes the genus g handlebody H_g embedded in $\mathbb{R}^3$, any triangulation of the boundary surface Σ_g can be extended in to a triangulation of H_g. The following Lemma shows that H_g is in general not the best filling for Σ_g.

Lemma 9. *One can construct a sequence of triangulations $\mathcal{T}_i$ of the boundary Σ_g with a fixed number of triangles $|\mathcal{T}_i| \leq 24g$, and with the property that any extension to a triangulation $\widehat{\mathcal{T}_i}$ of H_g satisfies $|\widehat{\mathcal{T}_i}| \to \infty$.*

Proof. For simplicity we restrict to $g = 1$. The higher genus cases are completely analogous. When $g = 1$, $H_g \cong \mathbb{D}^2 \times S^1$.

Given an arbitrary triangulation $\phi : L \to \mathbb{D}^2 \times S^1$, where L is a simplicial complex and ϕ a homeomorphism, we associate an invariant $||\phi||$ in the following manner. Consider embedded curves in the 1-skeleton of the boundary ∂L. There are only finitely many such curves, hence these represent finitely many elements $\alpha_1, \ldots, \alpha_k \in H_1(L)$. Under the homeomorphism, they map to finitely many elements in $H_1(\mathbb{D}^2 \times S^1) \cong \mathbb{Z}$, and we can define $||\phi||$ to be $\max\{|\phi_*(\alpha_i)|\}$. Now note that, while this invariant seems to be defined via the map ϕ, it in fact only depends on the simplicial complex L. Indeed, any homeomorphism from L to $\mathbb{D}^2 \times S^1$ induces an isomorphism from $\mathbb{Z} \cong H_1(L) \to H_1(\mathbb{D}^2 \times S^1) \cong \mathbb{Z}$. There are only two such isomorphisms, and the absolute value of the image is independent of which of these isomorphisms one uses.

Now by way of contradiction, assume there is a universal upper bound $|\widehat{\mathcal{T}_i}| \leq K$. Then there are finitely many simplicial complexes one can form with at most K complexes, call them $L_1, \ldots, L_k$. Thus for any corresponding triangulation $\phi : L_i \to \mathbb{D}^2 \times S^1$ one sees that there is an absolute bound on the element in $H_1(\mathbb{D}^2 \times S^1)$ represented by an embedded curve in the 1-skeleton of the boundary.

But starting with a fixed triangulation of $\Sigma_g \cong S^1 \times S^1$, one can compose the homeomorphism with powers of a Dehn twist along a curve $\beta \subset S^1 \times S^1$ which represents the generator in $H_1(\mathbb{D}^2 \times S^1)$. Note that this Dehn twist is a homeomorphism of $S^1 \times S^1$ that does *not* extend to a homeomorphism of $\mathbb{D}^2 \times S^1$. This forms a sequence of triangulations of $S^1 \times S^1$, each of which uses the same number of triangles. It is easy to see that this sequence of triangulations will contain embedded curves in their 1-skeletons representing arbitrarily large elements in $H_1(\mathbb{D}^2 \times S^1)$. It follows that, while every triangulation of $S^1 \times S^1$ in this sequence can be extended to a triangulation of $\mathbb{D}^2 \times S^1$, doing so will require a larger and larger number of tetrahedra, completing the proof. □

Of course, what is underlying the previous example is the fact that the natural homomorphism $MCG(H_g) \to MCG(\Sigma_g)$ has infinite index (where MCG denotes the mapping class group – the group of homotopy classes of homeomorphisms of the manifold). A similar argument can be used to give higher dimensional examples. Lemma 9 shows that the choice of a good filling 3-manifold must depend on the initial triangulation of Σ_g.

Remark. Some variations of our notion of filling function have previously been considered in the literature. For example, Hass, Snoeying, and W. Thurston [HST] have considered unknotted polygonal curves in $\mathbb{R}^3$, and studied the minimal number of triangles in a PL spanning disk for the curve. They give an exponential lower bound for the corresponding filling function, with an upper bound subsequently obtained by Hass, Lagarias, and W. Thurston [HLT]. The corresponding question for knotted polygonal curves bounding PL surfaces was considered by Hass and Lagarias [HL]. In a somewhat different direction, Costantino and D. Thurston [CT] considered a similar question for 3-manifolds – but did not require the optimal triangulation on the filling 4-manifold to restrict to the original triangulation on the 3-manifold.

Discrete analogues of Riemannian systolic inequalities. We first need the following Lemma. This Lemma follows directly from Theorem 2 and uses all of the same notation as this Theorem, except that the Riemannian metric will now be denoted by h:

Lemma 10. *Let $(M, \mathcal{T})$ be a closed triangulated n-dimensional manifold and let $P_1, \ldots, P_N$ be free-homotopy-invariant properties a loop in M can satisfy. Suppose that, for each $\epsilon > 0$, there is a closed geodesic γ_ϵ on the Riemannian manifold (M, h) (where h is the metric from Theorem 2) so that γ_ϵ satisfies properties $P_1, \ldots, P_N$ and*

$$\ell_h(\gamma_\epsilon) \leq C\sqrt{Vol_h(M)}. \tag{6.1}$$

Then there is an edge loop p on M so that p satisfies properties $P_1, \ldots, P_N$ and

$$\ell_{\mathcal{T}}(p) \leq \kappa_n C\sqrt{Vol_{\mathcal{T}}(M)}. \tag{6.2}$$

Corollary 11. *Let $(M, \mathcal{T})$ be a triangulated surface with infinite fundamental group. Then the systole is bounded by*

$$\mathrm{Sys}_{\mathcal{T}}(M) \leq \frac{2}{\sqrt{3}} \kappa_2 \sqrt{Vol_{\mathcal{T}}(M)}.$$

This Corollary follows from Lemma 10 and Corollary 5.2.B [Gr1].

Corollary 12. *Let $(M, \mathcal{T})$ be a triangulated surface of genus $g > 0$. Then the homological systole is bounded by*

$$\mathrm{Sys}^H_{\mathcal{T}}(M) \leq K_g \frac{\log g}{\sqrt{g}} \sqrt{Vol_{\mathcal{T}}(M)}.$$

where K_g depends only on the genus g and not on M or $\mathcal{T}$.

The above Corollary follows from Lemma 10 and Theorem 2.C [Gr2].

The cut-and-cone procedure. Suppose that $(M, \mathcal{T})$ is a triangulated surface with genus $g \geq 2$. The $g = 0, 1$ cases will be dealt with individually later. In order to simplify notation, we will use $|\mathcal{T}|$ to denote $\mathrm{Vol}_{\mathcal{T}}(M)$, the number of triangles in the triangulation $\mathcal{T}$. Set $(M_{(0)}, \mathcal{T}_{(0)}) := (M, \mathcal{T})$. By Corollary 12, there exists a homologically nontrivial edge loop p so that

$$\ell_{\mathcal{T}}(p) \leq K \frac{\log g}{\sqrt{g}} \sqrt{|\mathcal{T}|}. \tag{6.3}$$

By reducing the loop p, if necessary, we may assume that p is simple and still satisfies equation (6.3). Cutting M along p yields a connected surface of genus $g - 1$ with two boundary components. We then cone off the two boundary components to obtain a triangulated surface $(M_{(1)}, \mathcal{T}_{(1)})$ with genus $g - 1$. Note that

$$|\mathcal{T}_{(1)}| \leq |\mathcal{T}| + 2\ell_{\mathcal{T}}(p) \leq |\mathcal{T}| + 2K \frac{\log g}{\sqrt{g}} \sqrt{|\mathcal{T}|} \leq \left(\sqrt{|T|} + K \frac{\log g}{\sqrt{g}} \right)^2.$$

Suppose, inductively, that we have triangulated surfaces[4]

$$\mathcal{T} = \mathcal{T}_{(0)}, \mathcal{T}_{(1)}, \dots, \mathcal{T}_{(n)}$$

where $n \leq g - 1$, $\mathcal{T}_{(i)}$ is obtained from $\mathcal{T}_{(i-1)}$ by the above cut-and-cone procedure, and we have

$$|\mathcal{T}_{(i)}| \leq \left(\sqrt{|\mathcal{T}|} + K \sum_{k=g-(i-1)}^{g} \frac{\log k}{\sqrt{k}} \right)^2.$$

If $n < g - 1$, then $\mathcal{T}_{(n)}$ has genus $g - n \geq 2$, so by Corollary 12, there exists a homologically nontrivial edge loop $p_{(n)}$ so that

$$\ell_{\mathcal{T}_{(n)}}\left(p_{(n)}\right) \leq K \frac{\log(g-n)}{\sqrt{g-n}} \sqrt{|\mathcal{T}_{(n)}|}.$$

[4]Note that, in what follows, we are abusing notation and using $\mathcal{T}_{(i)}$ to denote both the surface and the triangulation.

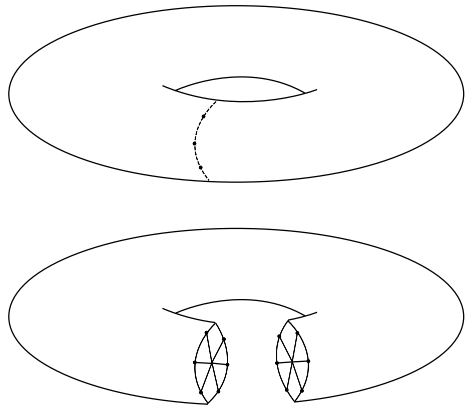

FIGURE 5. An example of the cut-and-cone procedure.

We may cut $\mathcal{T}_{(n)}$ along this path and cone off the boundaries to get a triangulated surface $\mathcal{T}_{(n+1)}$ with genus one less than the genus of $\mathcal{T}_{(n)}$ so that

$$\begin{aligned}
|\mathcal{T}_{(n+1)}| &\leq |\mathcal{T}_{(n)}| + 2K\frac{\log(g-n)}{\sqrt{g-n}}\sqrt{|\mathcal{T}_{(n)}|} \\
&\leq \left(\sqrt{|\mathcal{T}|} + K\sum_{k=g-(n-1)}^{g}\frac{\log k}{\sqrt{k}}\right)^2 \\
&\quad + 2K\frac{\log(g-n)}{\sqrt{g-n}}\left(\sqrt{|\mathcal{T}|} + K\sum_{k=g-(n-1)}^{g}\frac{\log k}{\sqrt{k}}\right) \\
&\leq \left(\sqrt{|\mathcal{T}|} + K\sum_{k=g-n}^{g}\frac{\log k}{\sqrt{k}}\right)^2.
\end{aligned}$$

If $n = g-1$, then $\mathcal{T}_{(n)} = \mathcal{T}_{(g-1)}$ is a torus and we may apply Corollary 11 to get a noncontractible edge loop p so that

$$\ell_{\mathcal{T}_{g-1}}(p) \leq \frac{2}{\sqrt{3}}\kappa_2\sqrt{|\mathcal{T}_{(g-1)}|}. \tag{6.4}$$

Cutting and coning along p gives us a triangulated 2-sphere $\mathcal{T}_{(g)}$ such that

$$\begin{aligned}|\mathcal{T}_{(g)}| &= |\mathcal{T}_{(g-1)}| + 2\ell_{\mathcal{T}_{g-1}}(p)\\ &\leq |\mathcal{T}_{(g-1)}| + 2\left(\frac{2}{\sqrt{3}}\kappa_2\right)\sqrt{|\mathcal{T}_{(g-1)}|}\\ &\leq \left(\sqrt{|\mathcal{T}|} + K\sum_{k=2}^{g}\frac{\log k}{\sqrt{k}}\right)^2 + \frac{4}{\sqrt{3}}\kappa_2\left(\sqrt{|\mathcal{T}|} + K\sum_{k=2}^{g}\frac{\log k}{\sqrt{k}}\right)\\ &\leq 5\kappa_2\left(\sqrt{|\mathcal{T}|} + K\sum_{k=2}^{g}\frac{\log k}{\sqrt{k}}\right)^2.\end{aligned}$$

If $n = g$, then $\mathcal{T}_{(n)} = \mathcal{T}_{(g)}$ is a 2-sphere. We need to perform a special coning off of $\mathcal{T}_{(g)}$. The reason for this is to ensure that, when we glue the surface back together to get our 3-dimensional filling $(N, \mathcal{T}_N)$, we obtain a legitimate simplicial complex decomposition for N. If we would just cone off $\mathcal{T}_{(g)}$, then various tetrahedra could intersect at both the cone point and in their opposite face.

The procedure for the modified coning of the 2-sphere is as follows. For each simplex σ of $\mathcal{T}_{(g)}$, we will triangulate the prism $\sigma \times I$, where I is the unit interval, in the same manner as used by Hatcher in ([Ha], pg 112-113). Suppose the vertices of $\sigma \times \{1\}$ are $\{v_0, v_1, v_2\}$, where the indices represent some fixed ordering of the vertices. Let the corresponding vertices of $\sigma \times \{0\}$ be $\{w_0, w_1, w_2\}$. Then the simplices $\langle v_0v_1v_2w_2\rangle$, $\langle v_0v_1w_1w_2\rangle$, and $\langle v_0w_0w_1w_2\rangle$ triangulate $\sigma \times I$, and if we do this for each simplex of $\mathcal{T}_{(g)}$, adjacent simplices will have consistent triangulations. Finally we cone off $\mathcal{T}_{(g)} \times \{0\}$ to get a triangulated 3-ball B_3 which has two layers: the center, which is a coned off copy of $\mathcal{T}_{(g)}$, and the exterior shell, which is our triangulated $\mathcal{T}_{(g)} \times I$. Note that

$$|B_3| = 4|\mathcal{T}_{(g)}| \leq 20\kappa_2\left(\sqrt{|\mathcal{T}|} + K\sum_{k=2}^{g}\frac{\log k}{\sqrt{k}}\right)^2.$$

By gluing together B_3 along the cuts in the reverse order, we obtain a triangulated 3-manifold $(N, \mathcal{T}')$ which is a filling of $(M, \mathcal{T})$ and

$$\begin{aligned}|\mathcal{T}'| &\leq 20\kappa_2\left(\sqrt{|\mathcal{T}|} + K\sum_{k=2}^{g}\frac{\log k}{\sqrt{k}}\right)^2\\ &= 20\kappa_2\left(\sqrt{|\mathcal{T}|} + K\sum_{k=2}^{7}\frac{\log k}{\sqrt{k}} + K\sum_{k=8}^{g}\frac{\log k}{\sqrt{k}}\right)^2\\ &\leq 20\kappa_2\left(\sqrt{|\mathcal{T}|} + C' + \int_7^g \frac{\log x}{\sqrt{x}}\,dx\right)^2\\ &\leq 20\kappa_2\left(\sqrt{|\mathcal{T}|} + C' + 2\sqrt{g}\log g\right)^2.\end{aligned}$$

Proofs of theorems.

Proof of Theorem 7. Suppose $(M, \mathcal{T}_M)$ is a triangulated surface of genus at most g. After performing the above cut-and-cone procedure we obtain $(N, \mathcal{T}_N)$, a filling of M, so that

$$\begin{aligned}|\mathcal{T}_N| &\leq 20\kappa_2\left(\sqrt{|\mathcal{T}_M|}+C'+2\sqrt{g}\log g\right)^2\\ &\leq 20\kappa_2\left(\sqrt{|\mathcal{T}_M|}+C'_g\right)^2\\ &\leq C_g|\mathcal{T}_M|\end{aligned}$$

for a suitable constant C_g. □

Proof of Theroem 8. Suppose $(M,\mathcal{T}_M)$ is a triangulated surface of genus g. After performing the above cut-and-cone procedure we obtain $(N,\mathcal{T}_N)$, a filling of M, so that

$$|\mathcal{T}_N| \leq 20\kappa_2\left(\sqrt{|\mathcal{T}_M|}+C'+2\sqrt{g}\log g\right)^2. \tag{6.5}$$

Since M is closed, the number of edges in $\mathcal{T}_M$ is $(3/2)|\mathcal{T}_M|$. Thus if $|v(\mathcal{T}_M)|$ is the number of vertices of $\mathcal{T}_M$, we know that the Euler characteristic $\chi(\mathcal{T}_M)$ satisfies

$$2-2g=\chi(\mathcal{T}_M)=|v(\mathcal{T}_M)|-\frac{|\mathcal{T}_M|}{2}.$$

Solving for g then gives that

$$g=\frac{-|v(\mathcal{T}_M)|}{2}+\frac{|\mathcal{T}_M|}{4}+1\leq\frac{|\mathcal{T}_M|}{4}. \tag{6.6}$$

Combining (6.5) and (6.6), we can conclude that

$$\begin{aligned}|\mathcal{T}_N| &\leq 20\kappa_2\left(\sqrt{|\mathcal{T}_M|}+C'+2\sqrt{g}\log g\right)^2\\ &\leq 20\kappa_2\left(\sqrt{|\mathcal{T}_M|}\,(1+\log|\mathcal{T}_M|)+C''\right)^2\\ &= 20\kappa_2\left(|\mathcal{T}_M|\,(1+\log|\mathcal{T}_M|)^2+2C''\sqrt{|\mathcal{T}_M|}(1+\log|\mathcal{T}_M|)+(C'')^2\right)\\ &\leq C|\mathcal{T}_M|\,(\log|\mathcal{T}_M|)^2\end{aligned}$$

for some suitable C. □

7. Concluding remarks

Our results suggest a variety of directions for further work. Firstly, note that throughout our paper we restrict ourselves to *smooth* triangulations. This restriction appears (and is used) in both implications of our Main Theorem. In our Theorem 1, we make use of Whitney's triangulation process, which produces smooth triangulations. In the proof of our Theorem 2, smoothness of the triangulation is used to produce nice local coordinates near the various faces. Note however that, even restricting to smooth manifolds, one can find many non-smooth triangulations. Indeed, for a smooth triangulation, it follows that the image of every simplex is smoothly embedded, and hence has link homeomorphic to a sphere of the appropriate codimension. On the other hand, the celebrated Cannon-Edwards double suspension theorem (see [Ed] and [Can]) states that, if one starts with an arbitrary n-dimensional homology sphere H (i.e. a connected n-manifold whose integral homology vanishes in all degrees $\neq 0,n$), the double suspension $\Sigma^2H=H*S^1$ is homeomorphic to S^{n+2}. Triangulating both H and the S^1, we get an induced triangulation of the join $\Sigma^2H=H*S^1$, and hence a triangulation of S^{n+2}. But in this triangulation, the edges in the S^1 have links homeomorphic to H. As a result, any homeomorphism $\Sigma^2 H \to S^{n+2}$ must take the suspension curve S^1 to a non-smooth curve in S^5. This yields a triangulation of S^{n+2} which is not smooth (in fact, not even PL).

Question: If we have a class of smooth manifolds for which the systolic inequality holds for all smooth triangulations, does the systolic inequality still hold (possibly with a different constant) for *all* triangulations? How about for PL-triangulations?

For a Riemannian manifold, we construct smooth triangulations whose simplices are "metrically nice" (as seen by the Riemannian metric). In the realm of Riemannian geometry, perhaps the most important notion is that of curvature. It is reasonable to ask whether one can produce smooth triangulations which also respect the curvature of the underlying metric.

Question: If M is a closed negatively curved manifold, does M support a piecewise Euclidean, locally CAT(0) metric?

The CAT(0) condition is a metric version of non-positive curvature. In the special case where M is (real) hyperbolic, this question has an affirmative answer, by work of Charney, Davis, and Moussong [CDM]. Note that, if one replaces the "negatively curved" by "non-positively curved", then there are counterexamples (due to Davis, Okun, and Zheng [DOZ]).

In Corollary 1 we used the direction (1) $\Rightarrow$ (2) of our Main Theorem to prove that a specific class of triangulated manifolds satisfied the combinatorial systolic inequality. So the following question is very natural:

Question: Can one directly establish the combinatorial systolic inequality for some classes of manifolds?

Via the implication (2) $\Rightarrow$ (1) in the Main Theorem, this would imply corresponding Riemannian systolic inequalities. Finally, we can ask for improvements on the filling function for triangulated surfaces:

Question: Does the filling function for triangulated surfaces satisfy a linear bound, with constant independent of the genus?

In our Theorem 7, we showed that for each fixed genus g, one has a linear filling function (but with a constant that depends on the genus). If we try to get a genus independent estimate, our Theorem 8 gives a slightly worse bound, with an additional log squared factor. It is unclear whether or not we should expect an affirmative answer to the last question. Of course, the question of finding a good filling is also of interest in higher dimension.

Question: If M is a manifold which bounds, what can one say about the filling function for M? For instance, for closed 3-manifolds, can one compute the (optimal) filling function? Could these filling functions be used to distinguish the topology of the 3-manifold?

If one just tries to minimize the numbers of simplices over all possible triangulations, then Costantino and D. Thurston [CT] have some estimates on the corresponding filling function (note that they do not require the optimal triangulation on the 4-manifold to restrict to the given triangulation on the 3-manifold).

References

[Br] M. Brunnbauer, *On manifolds satisfying stable systolic inequalities*, Math. Ann. **342** (2008), no. 4, 951-968.

[Cai] S. Cairns, *On the triangulation of regular loci*, Ann. of Math. (2) **35** (1934), no. 3, 579-587.

[Can] J. Cannon, *Shrinking cell-like decompositions of manifolds. Codimension three.* Annals Math. **110** (1979), 83–112.

[CDM] R. Charney, M. Davis, and G. Moussong, *Nonpositively curved, piecewise Euclidean structures on hyperbolic manifolds.* Michigan Math. J. **44** (1997), 201–208.

[CT] F. Costantino and D. Thurston, *3-manifolds efficiently bound 4-manifolds.* J. Topol. **1** (2008), 703–745.

[DOZ] M. Davis, B. Okun, and F. Zheng, *Piecewise Euclidean Structures and Eberlein?s Rigidity Theorem in the Singular Case.* Geom. Topol. **3** (1999), 303–330.

[Ed] R. D. Edwards, *Suspensions of homology spheres.* Preprint available on the arXiv:math/0610573

[Fr] M. Freedman, *The topology of four-dimensional manifolds.* J. Differential Geom. **17** (1982), 357–453.

[Gr1] M. Gromov, *Filling Riemannian manifolds.* J. Differential Geom. **18** (1983) no. 1, 1-147.

[Gr2] M. Gromov, *Systoles and intersystolic inequalities.* Actes de la Table Ronde de Géométrie Différentielle (Luminy, 1992), Sémin. Congr., vol. 1, Soc. Math. France, Paris, 1996, pp.291-362 (English, with English and French summaries).

[HL] J. Hass and J. C. Lagarias, *The minimal number of triangles needed to span a polygon embedded in* $\mathbb{R}^d$, in "Discrete and computational geometry", 509526. Algorithms Combin. **25**, Springer, Berlin, 2003.

[HLT] J. Hass, J. C. Lagarias and W. P. Thurston, *Area inequalities for embedded disks spanning unknotted curves.* J. Differential Geom. **68** (2004), 1–29.

[HST] J. Hass, J. Snoeyink and W. P. Thurston, *The size of spanning disks for polygonal curves.* Discrete Comput. Geom. **29** (2003), 1–17.

[Ha] A. Hatcher, *Algebraic topology.* Cambridge University Press, Cambridge (2002).

[Kl] W. Klingenberg, *Riemannian geometry.* de Gruyter Studies in Mathematics, New York (1995), 340–341.

[Ko] R. Kowalick, *Discrete systolic inequalities.* Ph. D. Thesis (2013), The Ohio State University.

[Ma] C. Manolescu, *Pin*(2)*-equivariant Seiberg-Witten Floer homology and the triangulation conjecture.* To appear in J. Amer. Math. Soc.

[Na] J. Nash, *The imbedding problem for Riemannian manifolds.* Ann. of Math. (2) **63** (1956), 20–63.

[VHM] É. Colin de Verdière, A. Hubard, A. de Mesmay, *Discrete systolic inequalities and decompositions of triangulated surfaces.* Disc. and Comp. Geom., *to appear*, arXiv:1408.4036.

[JW] J.H.C. Whitehead, *On* C^1*-complexes*, Ann. of Math. (2) **41** (1940), 809-824.

[HW] H. Whitney. *Geometric integration theory.* Princeton University Press, 1957.

422 E 17th Ave., Columbus, Ohio 43201
E-mail address: rkowalick@gmail.com

Department of Mathematics, Ohio State University, Columbus, Ohio 43210
E-mail address: jlafont@math.ohio-state.edu

Department of Mathematics, Ohio State University, Columbus, Ohio 43210
E-mail address: minemyer.1@osu.edu