Geometriae Dedicata (2019) 202:373–386
https://doi.org/10.1007/s10711-018-00419-9

ORIGINAL PAPER

Filling triangulated surfaces

Ryan Kowalick[1] · **Jean-François Lafont**[2] · **Barry Minemyer**[3]

Received: 17 April 2018 / Accepted: 22 December 2018 / Published online: 28 January 2019

Abstract
Given a triangulated closed oriented surface $(M, \mathcal{T}_M)$, we provide upper bounds on the number of tetrahedra needed to construct a triangulated 3-manifold $(N, \mathcal{T}_N)$ which bounds $(M, \mathcal{T}_M)$. Along the way, we develop a technique to translate (in all dimensions) between the famous Riemannian systolic inequalities of Gromov and combinatorial analogues of these inequalities.

Keywords Systolic geometry · Combinatorial systolic inequality · Fat triangulation · Nash embedding theorem · Whitney triangulation · Bounding manifold · Efficient filling

1 Introduction

Given a closed triangulated n-dimensional manifold $(M, \mathcal{T}_M)$, a *combinatorial filling* of M is a triangulated $(n+1)$-dimensional manifold $(N, \mathcal{T}_N)$ with $\partial N = M$ and $\mathcal{T}_N|_{\partial N} = \mathcal{T}_M$. 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? The main results of this paper are the following two solutions to this question in the case when $n = 2$.

Theorem 1 *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 triangulation.

✉ Jean-François Lafont
jlafont@math.ohio-state.edu

Ryan Kowalick
kowalick.1@osu.edu

Barry Minemyer
bminemyer@bloomu.edu

[1] Columbus, USA

[2] Department of Mathematics, Ohio State University, Columbus, OH 43210, USA

[3] Department of Mathematical and Digital Sciences, Bloomsburg University, Bloomsburg, PA 17815, USA

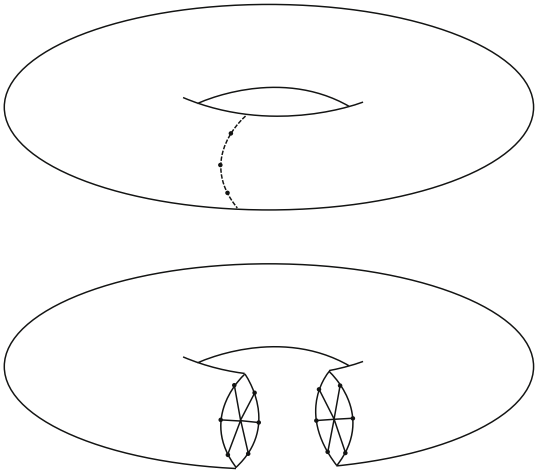

Fig. 1 An example of the cut-and-cone procedure

The next Theorem removes the dependence of the coefficient on the genus, but at the cost of inserting a higher power of $|\mathcal{T}_M|$.

Theorem 2 *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.

We do not know whether these results are optimal. In particular, we do not know whether the constant in Theorem 1 can be chosen independent of g.

The proofs of both of these theorems are very similar. We first develop related *combinatorial systolic inequalities* which bound the combinatorial 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 (compare with Gromov [11, Proof of Theorem 2C, pp. 302–305]). 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 Fig. 1). We iterate this procedure until the surface is a 2-sphere, in which case we can cone off the triangulated 2-sphere to get a triangulated 3-ball. By gluing the 3-ball along all of the cuts in the reverse order, and passing to a barycentric subdivision, we obtain a triangulated 3-manifold with the desired properties.

As we mentioned, our approach to proving these results relies on developing relevant combinatorial systolic inequalities. 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 (curvature free) expressions which relate the systole with other geometric quantities, typically the volume.

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 combinatorial volume $\mathrm{Vol}_{\mathcal{T}}(M)$ is just the number of top-dimensional simplices in the triangulation $\mathcal{T}$. The mechanism which allows us to develop the combinatorial systolic inequalities necessary to prove the above results is the following.

Theorem 3 *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*

$$Sys_g(M) \le C \sqrt[n]{Vol_g(M)},$$

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*

$$Sys_{\mathcal{T}}(M) \le C' \sqrt[n]{Vol_{\mathcal{T}}(M)},$$

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

Theorem 3 allows us to convert Riemannian systolic inequalities developed in [11] into combinatorial systolic inequalities, which are then used to prove Theorems 1 and 2 .

The argument for our proofs "build" the bounding 3-manifold from the triangulation on the surface 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 to a triangulation of H_g. Our next theorem establishes a general criterion, showing such an extension is in general not the most efficient filling for Σ_g.

Theorem 4 *Let W^n be a smooth compact n-manifold with boundary $\partial W = M$. The inclusion induces a homomorphism $\Psi : MCG(W) \to MCG(M)$ between the topological mapping class groups. If the image of Ψ has infinite index in $MCG(M)$, then there exists a sequence of triangulations $\mathcal{T}_i$ of the boundary M with a fixed number of simplices $|\mathcal{T}_i|$, and with the property that any extension to a triangulation $\widehat{\mathcal{T}_i}$ of W satisfies $|\widehat{\mathcal{T}_i}| \to \infty$.*

Notice that in the case where $W = H_g$ and $\partial W = \Sigma_g$, the morphism Φ is well known to have infinite index (indeed, the image is known to be distorted in the mapping class group when $g \ge 2$, see Hamenstädt and Hensel [12]). Thus the previous theorem applies to the case of surfaces.

Let us briefly describe the layout of our paper. In Sect. 2, we establish some basic definitions, and prove Theorem 4. In Sect. 3 we discuss how to covert between Riemannian metrics and smooth triangulations in a manner where the geometry of the Riemannian metric is "related" to the combinatorics of the triangulation (and vice versa). Rigorous statements are formally recorded as Propositions 5 and 6 , but the proofs mostly follow from results already in the literature. These results are then used to prove Theorem 3 in Sect. 4, where we also develop the combinatorial systolic inequalities needed to prove the main results. Finally, we establish Theorems 1 and 2 in Sect. 5. We conclude with Sect. 6 where we give a brief but technical sketch of Whitney's triangulation procedure in [25]—this is needed in the proof of Lemma 7 from Sect. 3.

Remark Many of the results in this paper were developed in the Ph.D. Thesis of Kowalick [17]. Results similar to Theorem 3 were independently obtained by de Verdière et al. [9]. 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.

2 Filling triangulated manifolds

Let us first establish some terminology for this paper. A *triangulated manifold* is a tuple $(M, \mathcal{T})$ where M is a manifold and $\mathcal{T} : K_{\mathcal{T}} \to M$ is a homeomorphism from a simplicial complex $K_{\mathcal{T}}$ to M. If the link of every simplex in $K_{\mathcal{T}}$ is a piecewise-linear sphere then we call this triangulation a *PL-triangulation*. If M is a smooth manifold, and the restriction of $\mathcal{T}$ to each simplex of $K_{\mathcal{T}}$ is smooth, then we call the triangulation $(M, \mathcal{T})$ smooth. A *facet* of a triangulation is a simplex of maximal dimension. For any triangulation $(M, \mathcal{T})$, the notation $|\mathcal{T}|$ will refer to the number of facets in the simplicial complex $K_{\mathcal{T}}$. For a triangulated manifold, this will be used as a combinatorial analogue of volume.

Given a triangulated n-dimensional manifold $(M, \mathcal{T})$, a combinatorial filling is a triangulated $(n+1)$-manifold $(W, \widehat{\mathcal{T}})$ with non-empty boundary, along with commutative diagram:

$$\begin{tikzcd} K_M \arrow[r, "\mathcal{T}"] \arrow[d] & M \arrow[d] \\ K_W \arrow[r, "\widehat{\mathcal{T}}"] & W \end{tikzcd} \tag{2.1}$$

where the vertical map $M \to W$ is a homeomorphism onto ∂W, and the vertical map $K_M \to K_W$ is a linear isomorphism onto a simplicial subcomplex of K_W. Of course, a combinatorial filling will only exist if the manifold M bounds. We are interested in the following.

Question: In each dimension n, estimate the function F with the property that, if $(M, \mathcal{T})$ is a triangulated n-manifold which bounds, then there exists a combinatorial filling $(W, \widehat{\mathcal{T}})$ satisfying $|\widehat{\mathcal{T}}| \leq F(|\mathcal{T}|)$. We call this F the *filling function* in dimension n.

If one knows that M bounds a manifold W, then it is reasonable to ask whether optimal fillings can always be obtained that extend a given triangulation of M. In other words, in the commutative diagram (2.1), we are fixing the right hand vertical arrow, varying the triangulation $(M, \mathcal{T})$, and asking whether the bottom left corner of the diagram can be completed with good control on $|\widehat{\mathcal{T}}|$. Theorem 4, which we now prove, shows that this is in general not the case.

Proof of Theorem 4 Recall that we are given a fixed smooth compact manifold W with $\partial W = M$, where the map $\Psi : \mathrm{MCG}(W) \to \mathrm{MCG}(M)$ induced by the inclusion has image of infinite index. We want to construct a sequence of smooth triangulations of M with a fixed number of simplices, but whose extensions to W require an increasing number of simplices.

Let us argue by way of contradiction. Start with a fixed smooth triangulation $\mathcal{T}$ of M, which we view as a homeomorphism $\mathcal{T} : L \to M$ from a simplicial complex L, which is smooth when restricted to each simplex. Now for each element $\phi \in \mathrm{MCG}(M)$, consider the triangulation $\mathcal{T}_\phi := \phi \circ \mathcal{T} : L \to M$, and observe that this produces infinitely many triangulations with the same fixed number $|\mathcal{T}|$ of top-dimensional simplices.

Now to argue by contradiction, assume that each of these triangulations of M extends to a triangulation $\widehat{\mathcal{T}_\phi} : K_\phi \to W$, with the simplicial complex K_ϕ having $\leq s$ top dimensional simplices. Since there are only finitely many simplicial complexes that have $\leq s$ simplices, each K_ϕ is isomorphic to a simplicial complex from a finite list $K_1, \ldots, K_N$. Moreover, the group of combinatorial automorphisms of the complex L is finite, so there are only finitely many identifications of L with the boundary of each K_i.

Now consider any two elements $\phi, \psi \in \mathrm{MCG}(M)$, and assume that the corresponding extensions to W arise from the same pair (K_i, L). Then we obtain the following commutative diagram

$$\begin{tikzcd} M \arrow[d] & L \arrow[l, "\mathcal{T}_\psi"'] \arrow[r, "\mathcal{T}_\phi"] \arrow[d] & M \arrow[d] \\ W & K_i \arrow[l, "\widehat{\mathcal{T}}_\psi"'] \arrow[r, "\widehat{\mathcal{T}}_\phi"] & W \end{tikzcd}$$

Observe that the composite map $\widehat{\mathcal{T}}_\psi \circ \widehat{\mathcal{T}}_\phi^{-1} \in \text{Homeo}(W)$ gives a self-homeomorphism of W, which extends the self-homeomorphism $\mathcal{T}_\psi \circ \mathcal{T}_\phi^{-1} = \psi \circ \phi^{-1} \in \text{Homeo}(M)$. It follows that the elements $\psi, \phi \in \text{MCG}(M)$ lie in the same left coset of $\Psi\,(\text{MCG}(W))$. Thus the number of cosets is bounded above by N, the cardinality of the finite set of possible pairs (K_i, L). This contradiction completes the proof. □

Since the topological mapping class group of the genus g handlebody H_g has infinite index in the mapping class group of the surface Σ_g, we see that some care must be taken when constructing efficient fillings of triangulated surfaces. Indeed, how the filling 3-manifold attaches to the surface will have to depend on the chosen triangulation of the surface.

Remark Some variations of our notion of filling function have previously been considered in the literature. For example, Hass et al. [15] 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 et al. [14]. The corresponding question for knotted polygonal curves bounding PL surfaces was considered by Hass and Lagarias [13]. In a somewhat different direction, Costantino and Thurston [8] 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.

3 Translating between Riemannian metrics and smooth triangulations

In this section, we establish some results allowing us to translate between Riemannian metrics and smooth triangulations, with a view towards obtaining combinatorial analogues of Riemannian systolic inequalities.

Proposition 5 (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 smooth 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 traversed by the path p. Also, recall from the Introduction that the combinatorial volume $\text{Vol}_{\mathcal{T}}(M)$ is the number of top-dimensional simplices in the triangulation $\mathcal{T}$. As was pointed out to the authors by an anonymous referee, this result was essentially proved by Babenko in [1, Sections 2 and 8, specifically see the proofs of Proposition 2.2 and Lemma 2.3]. So what follows is just a short outline of the proof, but the interested reader may also consult the original version of this paper on the arXiv [18] for complete details.

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.

Proposition 5 will be used to translate Riemannian systolic inequalities into discrete systolic inequalities. For the reverse direction, we will need the following result.

Proposition 6 (Encoding a Riemannian metric) *There exists a constant δ_n depending solely on the dimension n, with the property that for any closed Riemannian n-manifold (M, g), there exists a smooth triangulation $\mathcal{T}$ with the property that*

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

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

Remark Roughly speaking, the triangulation $\mathcal{T}$ produced in the proposition has no simplices that are "long and thin" (as measured in the Riemannian metric g). Such fat triangulations have been considered before, as they can be used to produce quasi-meromorphic mappings (see for instance Peltonen [20], as well as Saucan [22–24]). In those papers, the authors showed that the triangulations obtained via Cairns' method [5–7] could be arranged to be fat. We give a new proof of this result, by appealing instead to Whitney's triangulation method [25]. Our proof also leads to Lemma 7 and Proposition 8 below.

Another approach to constructing fat triangulations comes via Delaunay triangulations. The Delaunay triangulations associated to sets of points in $\mathbb{R}^n$ avoid long skinny simplices—though of course there is no guarantee that the resulting simplicial complex is a manifold. There has been a substantial amount of recent work on obtaining Delaunay triangulations on Riemannian manifolds, see for instance [2,21], though in higher dimensions whether or not a Delaunay complex triangulates a manifold is a subtle question, see [3]. For the case of hyperbolic manifolds, Delaunay triangulations have been constructed with control on the geometry of the simplices, see [4].

Let us now work toward proving Proposition 6. Let (M, g) be an n-dimensional Riemannian manifold. By the Nash Isometric Embedding theorem [19], 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 $\mathbb{R}^m$. We then apply the following Lemma to M, which follows almost directly from work of Whitney in ([25], Chapter IV Part B, pp. 124–135).

Lemma 7 *Let M be a compact n-dimensional smooth Riemannian submanifold of $\mathbb{R}^m$. Then for any tubular neighborhood U of M there exists 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 U, and the projection $\pi^* : U \to M$ induces a homeomorphism $\pi^* : T \to M$.*
(3) *If σ is a simplex of T (of any dimension), then its fatness is bounded below by $\Theta_{n,m}$, which depends only on the dimensions of the submanifold 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 have 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 $T_{\pi^*(q)}M$.
(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 only on T (but not on L), and some positive constant $C_{n,m}$ depending only on n and m.

In part (1) of the Lemma, a *secant simplex* of M is just a simplex σ in the ambient $\mathbb{R}^m$ such that all of the vertices of σ lie on M, and such that the interior of σ is the convex hull of these vertices (in $\mathbb{R}^m$). For σ an n-simplex in $\mathbb{R}^m$, the *fatness* (referred to in part (3) of the Lemma) is defined to be

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

where Vol_n denotes the n-dimensional Hausdorff volume in $\mathbb{R}^m$. This scale invariant quantity distinguishes long and skinny Euclidean simplices from others. Finally, recall that the dimension m of the ambient space is a function of n. So all of the constants in the above Lemma really only depend on n.

Lemma 7 follows directly from Whitney's proof in [25] that every Riemannian submanifold of $\mathbb{R}^m$ admits a smooth triangulation. Unfortunately, Whitney's arguments are very technical and rather difficult to read. In order to not distract from the main argument, we postpone the proof of Lemma 7 to the end of our paper (see Sect. 6).

Proof of Proposition 6 Using the notation of Lemma 7, the projection map $\pi^* : U \to M$ is a Riemannian submersion (for a sufficiently small neighborhood U). So $TU \cong TU_h \oplus TU_v$ where TU_h is canonically isomorphic to the tangent bundle TM of M, and TU_v is canonically isomorphic to the normal bundle of M in $\mathbb{R}^m$. 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 $w = w_h + w_v$, we have that

$$|w_h| = |\pi_{\pi^*(q)}(w)|. \tag{3.3}$$

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

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

is the identity. Thus, if the tubular neighborhood U is chosen sufficiently small, then for any $q \in U$, the map $D\pi^*\big|_{T_qU_h}$ has the property that for any $w \in T_qU_h$,

$$\frac{1}{\lambda}|w| \leq |D\pi^*(w)| \leq \lambda|w| \tag{3.4}$$

where $\lambda > 1$ is an arbitrarily chosen constant.

From here, we show that $\pi^* : T \to M$ is a 2λ bi-Lipschitz homeomorphism. Then since T satisfies Proposition 6 due to (3) of Lemma 7, we have that M also satisfies the Proposition.

Let $q \in \sigma$ and $\sigma \in T$. Then for all $v \in T_q\sigma$, inequalities (3.3) and (3.4) give us that

$$|D\pi^*(v)| = |D\pi^*(v_h)| \leq \lambda |v_h| \leq \lambda |v|.$$

On the other hand, combining the same two inequalities with inequality (3.1) yields

$$|v| \leq 2|\pi_{\pi^*(q)}(v)| = 2|v_h| \leq 2\lambda |D\pi^*(v)|.$$

□

Notice that we have proven that T is bi-Lipschitz homeomorphic to M, which by Lemma 7 implies Proposition 6. We record this as the following Proposition, which may be of independent interest.

Proposition 8 *Let (M, g) be an n-dimensional Riemannian manifold. Then there exists a piecewise-flat metric on M which is C_n bi-Lipschitz to (M, g), where the constant C_n depends only on the dimension of M.*

4 Combinatorial analogues of Riemannian systolic inequalities

Let us briefly recall some definitions. The *systole* of a Riemannian manifold (M, g), denoted $\mathrm{Sys}_g(M)$, is the length of the shortest non-contractible loop in M. The *homological systole* of a Riemannian manifold (M, g), denoted $\mathrm{Sys}_g^H(M)$, 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 *combinatorial length* of p, denoted $\ell_{\mathcal{T}}(p)$, will be the number of edges in p. The *combinatorial systole* of a triangulated manifold $\mathcal{T}$, denoted $\mathrm{Sys}_{\mathcal{T}}(M)$, will refer to the combinatorial length of the shortest non-contractible edge loop in $\mathcal{T}$. The *combinatorial homological systole*, denoted $\mathrm{Sys}_{\mathcal{T}}^H(M)$, is defined analogously.

With this terminology, we are now ready to prove Theorem 3.

Proof of Theorem 3 ($\Rightarrow$) 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 given by Proposition 5, γ the closed g-geodesic whose length realizes the Riemannian systole of (M, g), and p the edge path freely homotopic to γ provided by Proposition 5. 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 Proposition 6. 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 Proposition 6, 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 Corollary. □

In [10] 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 9 *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*

$$Sys_{\mathcal{T}}(M) \leq C\sqrt[n]{Vol_{\mathcal{T}}(M)},$$

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

In [10, Appendix B] Gromov develops analogous systolic inequalities for closed smooth essential manifolds endowed with a *polyhedral metric*. The above Corollary 9 could also be easily deduced from the polyhedral metric version of the systolic inequality.

The anonymous referee pointed out to us that Corollaries 12 and 11 below were proven by Hutchinson in [16]. We give an alternative proof of these results. The key ingredient in our argument is the following Lemma, which follows directly from Proposition 5 (and uses all the same notation).

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, g) (where g is the metric from Theorem 5) so that γ_ϵ satisfies properties $P_1, \ldots, P_N$ and*

$$\ell_h(\gamma_\epsilon) \leq C\sqrt{Vol_g(M)}. \tag{4.1}$$

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

$$\ell_{\mathcal{T}}(p) \leq \kappa_n C\sqrt{Vol_{\mathcal{T}}(M)}. \tag{4.2}$$

Combining Lemma 10 and [10, Corollary 5.2.B] immediately gives us:

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

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

Instead of focusing on homotopically non-trivial loops, one can look instead at homologically non-trivial loops. In a Riemannian manifold, the minimal length of such a loop is the *homological systole*. In the case of a smoothly triangulated manifold $(M, \mathcal{T})$, the minimal number of edges traversed by a homologically non-trivial loop contained in the 1-skeleton similarly defines the combinatorial homological systole $\mathrm{Sys}^H_{\mathcal{T}}(M)$. Combining Lemma 10 and [11, Theorem 2.C] yields:

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

$$Sys_{\mathcal{T}}^{H}(M) \leq K \frac{\log g}{\sqrt{g}} \sqrt{Vol_{\mathcal{T}}(M)}.$$

where K *is a universal constant.*

5 Filling triangulated surfaces

We now focus on the special case of triangulated surfaces. We discuss the cut and cone procedure, and establish a proof of Theorems 1 and 2 .

5.1 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

$$\ell_{\mathcal{T}}(p) \leq K \frac{\log g}{\sqrt{g}} \sqrt{|\mathcal{T}|}. \tag{5.1}$$

By reducing the loop p, if necessary, we may assume that p is simple and still satisfies Eq. (5.1). 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

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

Suppose, inductively, that we have triangulated surfaces

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

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

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

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

$$\ell_{\mathcal{T}_{(g-1)}}(p) \leq \frac{2}{\sqrt{3}}\kappa_2\sqrt{|\mathcal{T}_{(g-1)}|}. \tag{5.2}$$

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 cone off $T_{(g)}$ to obtain a triangulated 3-ball B_3. 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 $(M, \mathcal{T})$. This "triangulation" $\mathcal{T}'$ will generally not be simplicial: various tetrahedra could intersect at both the cone point and in their opposite face. To fix this issue, we pass to the first barycentric subdivision of $\mathcal{T}'$, which always yields a legitimate simplicial complex. By abuse of notation we continue to call this triangulation $\mathcal{T}'$. Performing the barycentric subdivision multiplies the number of tetrahedra by 24. Therefore, we have that

$$\begin{aligned}
|\mathcal{T}'| = 24|\mathcal{T}_{(g)}| &\leq 5\cdot 24\kappa_2\left(\sqrt{|\mathcal{T}|} + K\sum_{k=2}^{g}\frac{\log k}{\sqrt{k}}\right)^2 \\
&= \kappa\left(\sqrt{|\mathcal{T}|} + K\sum_{k=2}^{7}\frac{\log k}{\sqrt{k}} + K\sum_{k=8}^{g}\frac{\log k}{\sqrt{k}}\right)^2 \\
&\leq \kappa\left(\sqrt{|\mathcal{T}|} + C' + \int_7^g\frac{\log x}{\sqrt{x}}\,dx\right)^2 \\
&\leq \kappa\left(\sqrt{|\mathcal{T}|} + C' + 2\sqrt{g}\log g\right)^2
\end{aligned}$$

where $\kappa = 120\kappa_2$.

5.2 Proofs of theorems

Proof of Theorem 1 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

$$|\mathcal{T}_N| \leq \kappa \left(\sqrt{|\mathcal{T}_M|} + C' + 2\sqrt{g}\log g\right)^2 \leq \kappa \left(\sqrt{|\mathcal{T}_M|} + C'_g\right)^2 \leq C_g |\mathcal{T}_M|$$

for a suitable constant C_g. □

Proof of Theorem 2 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

$$|\mathcal{T}_N| \leq \kappa \left(\sqrt{|\mathcal{T}_M|} + C' + 2\sqrt{g}\log g\right)^2. \tag{5.3}$$

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

$$g = \frac{-|v(\mathcal{T}_M)|}{2} + \frac{|\mathcal{T}_M|}{4} + 1 \leq \frac{|\mathcal{T}_M|}{4}. \tag{5.4}$$

Combining (5.3) and (5.4), we can conclude that

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

for some suitable C. □

6 Whitney's triangulation procedure

In this Section we give a short, high-level sketch of Whitney's triangulation procedure from [25] in order to justify Lemma 7 from Sect. 3.

We begin with a Riemannian submanifold M^n of $\mathbb{R}^m$. By the smooth Nash isometric embedding theorem [19] we know that the dimension m of the ambient space can be reduced, if necessary, to a quantity which is a function of n (on the order of n^2). 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 [25]. On pp. 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 Lemma 7 Let us first remind the reader that all Lemmas and equations in this proof reference Whitney's paper [25].

Whitney proves that $\pi^* : K \to M$ is a diffeomorphism (pp. 134–135). Now, using Lemma 21a on p. 132, Eq. (21.2) on p. 132, and the fact that $4\lambda\xi < \lambda\xi/\alpha$ (since $0 < \alpha << 1$, see Eqs. (17.2) on p. 128 and (21.2) on p. 132), we see 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. Every simplex of T is a secant simplex of M. We have that every simplex of T has fatness at least $\Theta_{n,m} := \Theta_1/2$ (see Eq. (17.3) on p. 128 for the definition of Θ_1), a number which depends only on n and m [see the proof of Lemma 21a on p. 132, between Eqs. (21.3) and (21.4)]. This proves parts (1), (2), and (3) of Lemma 7.

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 p. 132 (beginning with $|u - \pi_p u|$). This proves part (4) of the Lemma. Lastly, via the second-to-last equation on p. 132 (beginning with $|p_i' - p_0'|$) and Eq. (21.2) on p. 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 Lemma 7.

Let us note that β, δ, and b are defined on pp. 128–129 in Eqs. (17.3), (17.5), and (17.6), respectively. The parameters λ and ξ are likewise defined in Eqs. (17.4) and (17.5). The quantity β depends only on m and h, which in turn both only depend on n. This concludes the proof of Lemma 7. □

Acknowledgements The authors would like to thank Dylan Thurston for some helpful comments. We would also like to thank the various anonymous referees for remarks which greatly aided in the exposition of this paper, specifically with substantially shortening the proof of Proposition 6, suggesting the addition of Proposition 8, and pointing us towards references [1,16]. The work of the second author was partially supported by the NSF, under Grants DMS-1207782, DMS-1510640, and DMS-1812028. The research of the third author was partially supported by an AMS-Simons travel Grant.

References

1. Babenko, I.K.: Asymptotic invariants of smooth manifolds. Izv. Ross. Akad. Nauk Ser. Mat. **56**(4), 707–751 (1992)
2. Boissonnat, J.-D., Dyer, R., Ghosh, A.: Delaunay triangulation of manifolds. Found. Comput. Math. **18**(2), 399–431 (2018)
3. Boissonnat, J.-D., Dyer, R., Ghosh, A., Martynchuk, N.: An obstruction to Delaunay triangulations in Riemannian manifolds. Discrete Comput. Geom. **59**(1), 226–237 (2018)
4. Breslin, W.: Thick triangulations of hyperbolic n-manifolds. Pacific J. Math. **241**(2), 215–225 (2009)
5. Cairns, S.: On the triangulation of regular loci. Ann. Math. (2) **35**(3), 579–587 (1934)
6. Cairns, S.: Polyhedral approximations to regular loci. Ann. Math. (2) **37**(2), 409–415 (1936)
7. Cairns, S.: A simple triangulation method for smooth manifolds. Bull. Am. Math. Soc. **67**, 389–390 (1961)
8. Costantino, F., Thurston, D.: 3-manifolds efficiently bound 4-manifolds. J. Topol. **1**, 703–745 (2008)
9. de Verdière, E.C., Hubard, A., de Mesmay, A.: Discrete systolic inequalities and decompositions of triangulated surfaces. Discrete Comput. Geom. **53**(3), 587–620 (2015)
10. Gromov, M.: Filling Riemannian manifolds. J. Differ. Geom. **18**(1), 1–147 (1983)

11. Gromov, M.: Systoles and intersystolic inequalities. Actes de la Table Ronde de Géométrie Différentielle (Luminy, 1992), Sémin. Congr., vol. 1, Soc. Math. France, Paris, pp. 291–362 (1996) (English, with English and French summaries)
12. Hamenstädt, U., Hensel, S.: The geometry of handlebody groups I: distortion. J. Topol. Anal. **4**, 71–97 (2012)
13. Hass, J., Lagarias, J.C.: 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)
14. Hass, J., Lagarias, J.C., Thurston, W.P.: Area inequalities for embedded disks spanning unknotted curves. J. Differ. Geom. **68**, 1–29 (2004)
15. Hass, J., Snoeyink, J., Thurston, W.P.: The size of spanning disks for polygonal curves. Discrete Comput. Geom. **29**, 1–17 (2003)
16. Hutchinson, J.P.: On short noncontractible cycles in embedded graphs. SIAM J. Discrete Math. **1**, 185–192 (1988)
17. Kowalick, R.: Discrete systolic inequalities. Ph.D. Thesis, The Ohio State University (2013)
18. Kowalick, R., Lafont, J.F., Minemyer, B.: Combinatorial systolic inequalities. Preprint. arXiv:1506.07121
19. Nash, J.: The imbedding problem for Riemannian manifolds. Ann. Math. (2) **63**, 20–63 (1956)
20. Peltonen, K. : On the existence of quasiregular mappings. Ann. Acad. Sci. Fenn. Ser. A I Math. Dissertationes **85** (1992)
21. Rouxel-Labbé, M., Wintraecken, M., Boissonnat, J.-D.: Discretized Riemannian Delaunay triangulations. IMR25 Proc. Eng. **163**, 97–109 (2016)
22. Saucan, E.: Note on a theorem of Munkres. Mediterr. J. Math. **2**, 215–229 (2005)
23. Saucan, E.: The existence of quasimeromorphic mappings in dimension 3. Conform. Geom. Dyn. **10**, 21–40 (2006)
24. Saucan, E.: The existence of quasimeromorphic mappings. Ann. Acad. Sci. Fenn. Math. **31**, 131–142 (2006)
25. Whitney, H.: Geometric Integration Theory. Princeton University Press, Princeton (1957)