Discrete Systolic Inequalities and Decompositions of Triangulated Surfaces

Éric Colin de Verdière · Alfredo Hubard · Arnaud de Mesmay

Abstract How much cutting is needed to simplify the topology of a surface? We provide bounds for several instances of this question, for the minimum length of topologically non-trivial closed curves, pants decompositions, and cut graphs with a given combinatorial map in triangulated combinatorial surfaces (or their dual cross-metric counterpart).

Our work builds upon Riemannian systolic inequalities, which bound the minimum length of non-trivial closed curves in terms of the genus and the area of the surface. We first describe a systematic way to translate Riemannian systolic inequalities to a discrete setting, and vice-versa. This implies a conjecture by Przytycka and Przytycki from 1993, a number of new systolic inequalities in the discrete setting, and the fact that a theorem of Hutchinson on the edge-width of triangulated surfaces and Gromov's systolic inequality for surfaces are essentially equivalent. We also discuss how these proofs generalize to higher dimensions.

Supported by the French ANR Blanc project ANR-12-BS02-005 (RDAM). Preliminary version in *Proceedings of the 30th Annual Symposium on Computational Geometry*, 2014.

É. Colin de Verdière
CNRS, Département d'informatique, École normale supérieure, Paris, France. E-mail: eric.colin.de.verdiere@ens.fr.

A. Hubard
Laboratoire de l'Institut Gaspard Monge, Université Paris-Est Marne-la-Vallée. E-mail: alfredo.hubard@inria.fr. Portions of this work were done during a post-doctoral visit at the Département d'informatique of École normale supérieure, funded by the Fondation Sciences Mathématiques de Paris.

A. de Mesmay
IST Austria, Vienna. E-mail: arnaud.de.mesmay@ist.ac.at. Portions of this work were done as a Ph.D. student at the Département d'informatique of École normale supérieure.

This paper appeared in *Discrete & Computational Geometry*, 2015, 53(3):587–620. The final publication is available at www.springerlink.com.

Then we focus on topological decompositions of surfaces. Relying on ideas of Buser, we prove the existence of pants decompositions of length $O(g^{3/2}n^{1/2})$ for any triangulated combinatorial surface of genus g with n triangles, and describe an $O(gn)$-time algorithm to compute such a decomposition.

Finally, we consider the problem of embedding a cut graph (or more generally a cellular graph) with a given combinatorial map on a given surface. Using random triangulations, we prove (essentially) that, for any choice of a combinatorial map, there are some surfaces on which any cellular embedding with that combinatorial map has length superlinear in the number of triangles of the triangulated combinatorial surface. There is also a similar result for graphs embedded on polyhedral triangulations.

Keywords systole · edge-width · surface · closed curve · cut graph · pants decomposition

Mathematics Subject Classification (2000) 05C10 · 68U05 · 53C23 · 57M15 · 68R10

CR Subject Classification F.2.2 · G.2.2 · I.3.5

1 Introduction

Shortest curves and graphs with given properties on surfaces have been much studied in the recent computational topology literature; a lot of effort has been devoted towards efficient algorithms for finding shortest curves that simplify the topology of the surface, or shortest topological decompositions of surfaces [7,8,19–23,39] (refer also to the recent surveys [12,18]). These objects provide "canonical" simplifications or decompositions of surfaces, which turn out to be crucial for algorithm design in the case of surface-embedded graphs, where making the graph planar is needed [6, 9, 11, 41], as well as for many purposes in computer graphics and mesh processing [29, 43, 44, 48, 60].

In this article, we study inequalities that relate the size of a triangulated surface with the length of such shortest curves and graphs embedded thereon. The model parameter that we study is the notion of *edge-width* of an (unweighted) graph embedded on a surface [7,56], that is, the length of a shortest closed walk in the graph that is non-contractible on the surface (i.e., cannot be deformed to a single point on the surface). In particular we are interested in the following question: What is the largest possible edge-width, over all triangulations with n triangles of an orientable surface of genus g without boundary? It was known [33] that $O(\sqrt{n/g}\log g)$ is an upper bound for the edge-width, and we prove that this bound is asymptotically tight, namely, that some combinatorial surfaces of arbitrarily large genus achieve this bound. We also study similar questions for other types of curves (non-separating closed curves, null-homologous but non-contractible closed curves) and for decompositions (pants decompositions, and cut graphs with a prescribed combinatorial map), and give an algorithm to compute short pants decompositions.

Most of our results build upon or extend to a discrete setting some known theorems in *Riemannian systolic geometry*, the archetype of which is an upper bound on the systole (the length of a shortest non-contractible closed curve—a continuous version of the edge-width) in terms of the square root of the area of a Riemannian surface without boundary (or more generally the dth root of the volume of an essential Riemannian d-manifold). Riemannian systolic geometry [28,34] was pioneered by Loewner and Pu [54], reaching its maturity with the deep work of Gromov [27]. In Thurston's words, topology is naked and it *dresses* with geometric structures; systolic geometry regards the lengths and areas of all those possible outfits. Similarly, endowing a topological surface with a triangulation is a way to "dress" it and much of this paper leverages on comparing these two types of outfits.

We always assume that the surface has *no boundary*, that the underlying graph of the combinatorial surface is a *triangulation*, and that its edges are *unweighted*; the curves and graphs we seek remain on the edges of the triangulation. Lifting any of these restrictions invalidates or significantly worsens our bounds. In many natural situations, such requirements hold, such as in geometric modeling and computer graphics, where triangular meshes of surfaces without boundary are typical and, in many cases, the triangles have bounded aspect ratio (which immediately implies that our bounds apply, the constant in the $O(\cdot)$ notation depending on the aspect ratio).

After the preliminaries (Section 2), we prove three independent results (Sections 3–5), which are described and related to other works below. This paper is organized so as to showcase the more conceptual results before the more technical ones. Indeed, the results of Section 3 exemplify the strength of the connection with Riemannian geometry, while the results in Sections 4 and 5 are perhaps a bit more specific, but feature deeper algorithmic and combinatorial tools.

Systolic inequalities for closed curves on triangulations. Our first result (Section 3) gives a systematic way of translating a systolic inequality in the Riemannian case to the case of triangulations, and vice-versa. This general result, combined with known results from systolic geometry, immediately implies bounds on the length of shortest curves with given topological properties: On a triangulation of genus g with n triangles, some non-contractible (resp., non-separating, resp., null-homologous but non-contractible) closed curve has length $O(\sqrt{n/g}\log g)$, and, moreover, this bound is best possible.

These upper bounds are new, except for the non-contractible case, which was proved by Hutchinson [33] with a worse constant in the $O(\cdot)$ notation. The optimality of these inequalities is also new. Actually, Hutchinson [33] had conjectured that the correct upper bound was $O(\sqrt{n/g})$; Przytycka and Przytycki refuted her conjecture, building, in a series of papers [51–53], examples that show a lower bound of $\Omega(\sqrt{n \log g/g})$. They conjectured in 1993 [52] that the correct bound was $O(\sqrt{n/g}\log g)$; here, we confirm this conjecture.

In Appendix A, we observe that the proofs of the results mentioned above extend to higher dimensions. However, the situation is not quite as symmetrical as in the two-dimensional case: It turns out that discrete systolic inequalities in terms of the number of vertices or facets imply continuous systolic inequalities, and that continuous systolic inequalities imply discrete systolic inequalities only in terms of the number of facets. This allows us to derive that a systolic inequality in terms of the number of facets holds for every triangulation of an essential manifold.

As pointed out to us by a referee, slight variations of the results of Section 3 and Appendix A were simultaneously and independently discovered by Ryan Kowalick in his Ph.D. thesis [38]. Our approach in Section 3.1 is similar to his. In contrast, we use Voronoi diagrams in Section 3.2, while he uses a different construction inspired by Whitney. We will make some further technical comments on his work at the end of Appendix A.

Short pants decompositions. A pants decomposition is a set of disjoint simple closed curves that split the surface into *pairs of pants*, namely, spheres with three boundary components. In Section 4, we focus on the length of the shortest pants decomposition of a triangulation. As in all previous works, we allow several curves of the pants decomposition to run along a given edge of the triangulation. (Formally, we work in the cross-metric surface that is dual to the triangulation.)

The problem of computing a shortest pants decomposition has been considered by several authors [17, 50], and has found satisfactory solutions (approximation algorithms) only in very special cases, such as the punctured Euclidean or hyperbolic plane [17]. Strikingly, no hardness result is known; the strong condition that curves have to be disjoint, and the lack of corresponding algebraic structure, makes the study of short pants decompositions hard [30, Introduction]. In light of this difficulty, it seems interesting to look for algorithms that compute short pants decompositions, even without guarantee compared to the optimum solution.

Inspired by a result by Buser [5, Th. 5.1.4] on short pants decompositions on Riemannian surfaces, we prove that every triangulation of genus g with n triangles admits a pants decomposition of length $O(g^{3/2}n^{1/2})$, and we give an $O(gn)$-time algorithm to compute one. While it is known that pants decompositions of length $O(gn)$ can be computed for arbitrary combinatorial surfaces [14, Prop. 7.1], the assumption that the surface is unweighted and triangulated allows for a strictly better bound in the case where $g = o(n)$. (It is always true that $g = O(n)$.) We remark that the greedy approach coupled with Hutchinson's bound only gives a bound on the length of the pants decomposition of the form $f(g).\sqrt{n}$ where f is superpolynomial [1, Introduction].

On the lower bound side, some surfaces have no pants decompositions with length $O(n^{7/6-\varepsilon})$, as proved recently by Guth et al. [30] using the probabilistic method. Guth et al. show that polyhedral surfaces obtained by gluing triangles at random have this property.

Shortest embeddings of combinatorial maps. Finally, in Section 5, we consider the problem of decomposing a surface using a short cut graph with a prescribed combinatorial map. A natural approach to build a homeomorphism between two surfaces is to cut both of them along a cut graph, and to put the remaining disks in correspondence. However, for this approach to work, cut graphs defining the same combinatorial map are needed.

In this direction, Lazarus et al. [40] proved that every surface has a *canonical system of loops* (a specific combinatorial map of a cut graph with one vertex) with length $O(gn)$, which is worst-case optimal, and gave an $O(gn)$-time algorithm to compute one.

However, there is no strong reason to focus on canonical systems of loops. It is fairly natural to expect that other combinatorial maps will always have shorter embeddings (in particular, by allowing several vertices on the cut graph instead of just one). Still, we prove (essentially) that for any choice of combinatorial map of a cut graph, there exist triangulations with n triangles on which all embeddings of that combinatorial map have a *superlinear* length, actually $\Omega(n^{7/6-\varepsilon})$. (Since n may be $O(g)$, there is no contradiction with the result by Lazarus et al. [40].) In particular, some edges of the triangulation are traversed $\Omega(n^{1/6-\varepsilon})$ times.

Our proof uses the probabilistic method in the same spirit as the aforementioned article of Guth et al. [30]: We show that combinatorial surfaces obtained by gluing triangles randomly satisfy this property asymptotically almost surely, i.e., that the probability of satisfying this property by a random surface tends to one as the number of triangles tend to infinity. We remark that beyond the extremal qualities that concern us, random surfaces and their geometry have been heavily studied recently [24,45] in connection to quantum gravity [49] and Belyi surfaces [3].

Another view of our result is via the following problem: Given two graphs G_1 and G_2 cellularly embedded on a surface S, is there a homeomorphism $\varphi : S \to S$ such that G_1 does not cross the image of G_2 too many times? Our result essentially says that, if G_1 is fixed, for most choices of trivalent graphs G_2 with n vertices, for any φ, there will be $\Omega(n^{7/6-\varepsilon})$ crossings between G_1 and $\varphi(G_2)$. This is related to recent preprints [25,46], where upper bounds are proved for the number of crossings for the same problem, but with sets of disjoint curves instead of graphs. During their proof, Matoušek et al. [46] also encountered the following problem (rephrased here in the language of this paper): For a given genus g, does there exist a *universal* combinatorial map cutting the surface of genus g into a genus zero surface (possibly with several boundaries), and with a linear-length embedding on every such surface? We answer this question in the negative for cut graphs. In Appendix B, we prove a related result for families of closed curves cutting the surface into a genus zero surface.

2 Preliminaries

2.1 Topology for Graphs on Surfaces

We only recall the most important notions of topology that we will use, and refer to Stillwell [59] or Hatcher [32] for details. We denote by $S_{g,b}$ the (orientable) surface of ***genus*** g with b ***boundaries***, which is unique up to homeomorphism. The surfaces $S_{0,0}$, $S_{0,1}$, $S_{0,2}$, and $S_{0,3}$ are respectively called the ***sphere***, the ***disk***, the ***annulus***, and the ***pair of pants***. Surfaces are assumed to be connected, compact, and orientable unless specified otherwise. The notation ∂S denotes the boundary of S.

A ***path***, respectively a ***closed curve***, on a surface S is a continuous map $p : [0, 1] \to S$, respectively $\gamma : \mathbb{S}^1 \to S$. Paths and closed curves are ***simple*** if they are one-to-one. A ***curve*** denotes a path or a closed curve. We refer to Hatcher [32] for the usual notions of homotopy (continuous deformation) and homology. A closed curve is ***contractible*** if it is null-homotopic, i.e., it cannot be continuously deformed to a point. A simple closed curve is contractible if and only if it bounds a disk.

All the graphs that we consider in this paper are multigraphs, i.e., loops are allowed and vertices can be joined by multiple edges. An ***embedding*** of a graph G on a surface S is, informally, a crossing-free drawing of G on S. A graph embedding is ***cellular*** if its faces are homeomorphic to open disks. Euler's formula states that $v - e + f = 2 - 2g - b$ for any graph with v vertices, e edges, and f faces cellularly embedded on a surface S with genus g with b boundaries. A ***triangulation*** of a surface S is a cellular graph embedding such that every face is a triangle. A graph G cellularly embedded on a surface S yields naturally a ***combinatorial map*** M, which stores the combinatorial information of the embedding G, namely, the cyclic ordering of the edges around each vertex; we also say that G is an ***embedding*** of M on S. Two graphs cellularly embedded on S have the same combinatorial map if and only if there exists a self-homeomorphism of S mapping one (pointwise) to the other.

A graph G embedded on a surface S is a ***cut graph*** if the surface obtained by cutting S along G is a disk. A ***pants decomposition*** of S is a family of disjoint simple closed curves Γ such that cutting S along all curves in Γ gives a disjoint union of pairs of pants. Every surface $S_{g,b}$ except the sphere, the disk, the annulus, and the torus admits a pants decomposition, with $3g + b - 3$ closed curves and $2g + b - 2$ pairs of pants.

2.2 Combinatorial and Cross-Metric Surfaces

We now briefly recall the notions of combinatorial and cross-metric surfaces, which define a discrete metric on a surface; see Colin de Verdière and Erickson [13] for more details. In this paper, all edges of the combinatorial and cross-metric surfaces are unweighted.

A ***combinatorial surface*** is a surface S together with an embedded graph G, which will always be a triangulation in this article. In this model, the only allowed curves are walks in G, and the length of a curve c, denoted by $|c|_G$, is the number of edges of G traversed by c, counted with multiplicity.

However, it is often convenient (Sections 4 and 5) to allow several curves to traverse a same edge of G, while viewing them as being disjoint (implicitly, by "spreading them apart" infinitesimally on the surface). This is formalized using the dual concept of ***cross-metric surface***: Instead of curves in G, we consider curves in ***regular*** position with respect to the dual graph G^*, namely, that intersect the edges of G^* transversely and away from the vertices; the length of a curve c, denoted by $|c|_{G^*}$, is the number of edges of G^* that c crosses, counted with multiplicity. Since, in this article, G is always a triangulation, G^* is always *trivalent*, i.e., all its vertices have degree three. Thus, a cross-metric surface is a surface S equipped with a cellular, trivalent graph (usually denoted by G^*).

We note that the previous definition of cross-metric surface is valid also in the case where the surface has non-empty boundary (see Colin de Verdière and Erickson [13, Section 1.2] for more details). Curves and graph embedded on cross-metric surfaces can be manipulated efficiently [13]. The different notions of systoles are easily translated for both combinatorial and cross-metric surfaces.

Once again, we emphasize that, in this paper, unless otherwise noted, **all combinatorial surfaces are triangulated (each face is a disk with three sides) and unweighted (each edge has weight one)**. Dually, **all cross-metric surfaces are trivalent (each vertex has degree three) and unweighted (each edge has crossing weight one)**.

2.3 Riemannian Surfaces and Systolic Geometry

We will use some notions of Riemannian geometry, referring the interested reader to standard textbooks [15, 37]. A ***Riemannian surface*** (S, m) is a surface S equipped with a metric m, defined by a scalar product on the tangent space of every point. For example, smooth surfaces embedded in some Euclidean space $\mathbb{R}^d$ are naturally Riemannian surfaces—conversely, every Riemannian surface can be isometrically embedded in some $\mathbb{R}^d$ [31] but we will not need this fact. The length of a (rectifiable) curve c is denoted by $|c|_m$. The ***Gaussian curvature*** $\boldsymbol{\kappa_p}$ of S at a point p is the product of the eigenvalues of the scalar product at p. By the Bertrand–Diquet–Puiseux theorem [58, Chapter 3, Prop. 11], the area of the ball $\boldsymbol{B(p, r)}$ of radius r centered at p equals $\pi r^2 - \kappa_p \pi r^4 + o(r^4)$. We now collect the results from systolic geometry that we will use; for a general presentation of the field, see, e.g., Gromov [28] or Katz [34].

Theorem 2.1 ([4, 27, 28, 35, 57]) *There are constants $c, c', c'', c''' > 0$ such that, on any Riemannian surface without boundary, with genus g and area A:*

1. *some non-contractible closed curve has length at most* $c\sqrt{A/g}\log g$*;*
2. *some non-separating closed curve has length at most* $c'\sqrt{A/g}\log g$*;*
3. *some null-homologous non-contractible closed curve has length at most* $c''\sqrt{A/g}\log g$.

Furthermore,

4. *for an infinite number of values of* g*, there exist Riemannian surfaces of constant curvature* -1 *(hence area* $A = 4\pi(g-1)$*) and systole larger than* $\frac{2}{3\sqrt{\pi}}\sqrt{A/g}\log g - c'''$*. In particular, the three previous inequalities are tight up to constant factors.*

In this theorem, (1) and (2) are due to Gromov [27,28], (3) is due to Sabourau [57], and (4) is due to Buser and Sarnak [4, p. 45]. Furthermore, Gromov's proof yields $c = 2/\sqrt{3}$ in (1), which has been improved asymptotically by Katz and Sabourau [35]: They show that for every $c > 1/\sqrt{\pi}$ there exists some integer g_c so that (1) is valid for every $g \geq g_c$.

3 A Two-Way Street

In this section, we prove that any systolic inequality regarding closed curves in the continuous (Riemannian) setting can be converted to the discrete (triangulated) setting, and vice-versa.

3.1 From Continuous to Discrete Systolic Inequalities

Theorem 3.1 *Let* (S, G) *be a triangulated combinatorial surface of genus* g*, without boundary, with* n *triangles. Let* $\delta > 0$ *be arbitrarily small. There exists a Riemannian metric* m *on* S *with area* n *such that for every closed curve* γ *in* (S, m) *there exists a homotopic closed curve* γ' *on* (S, G) *with* $|\gamma'|_G \leq (1+\delta)\sqrt[4]{3}\ |\gamma|_m$.

This theorem, combined with known theorems from systolic geometry, immediately implies:

Corollary 3.1 *Let* (S, G) *be a triangulated combinatorial surface with genus* g *and* n *triangles, without boundary. Then, for some absolute constants* c*,* c'*, and* c''*:*

1. *some non-contractible closed curve has length at most* $c\sqrt{n/g}\log g$*;*
2. *some non-separating closed curve has length at most* $c'\sqrt{n/g}\log g$*;*
3. *some homologically trivial non-contractible closed curve has length at most* $c''\sqrt{n/g}\log g$.

Proof of Corollary 3.1. The proof consists in applying Theorem 3.1 to (S, G), obtaining a Riemannian metric m. For each of the different cases, the appropriate Riemannian systolic inequality is known, which means that a short curve γ

of the given type exists on (S,m) (Theorem 2.1(1–3)); by Theorem 3.1, there exists a homotopic curve γ' in (S,G) such that $|\gamma'|_G \leq (1+\delta)\sqrt[4]{3}\ |\gamma|_m$, for any $\delta > 0$. $\square$

Plugging in the best known constants for Theorem 2.1 (1) allows us to take $c = 2/\sqrt[4]{3}$, or any $c > \sqrt[4]{3/\pi^2}$ asymptotically using the refinement of Katz and Sabourau.

Furthermore, we note that, by Euler's formula and double-counting, we have $n = 2v + 4g - 4$, where v is the number of vertices of G. Thus, on a triangulated combinatorial surface with $v \geq g$ vertices, the length of a shortest non-contractible closed curve is at most $2\sqrt{2}\sqrt[4]{3}\cdot\sqrt{v/g}\log g < 3.73\sqrt{v/g}\log g$. This reproves a theorem of Hutchinson [33], except that her proof technique leads to the weaker constant 25.27. This constant can be improved asymptotically to $\sqrt[4]{108/\pi^2} < 1.82$ with the aforementioned refinement.

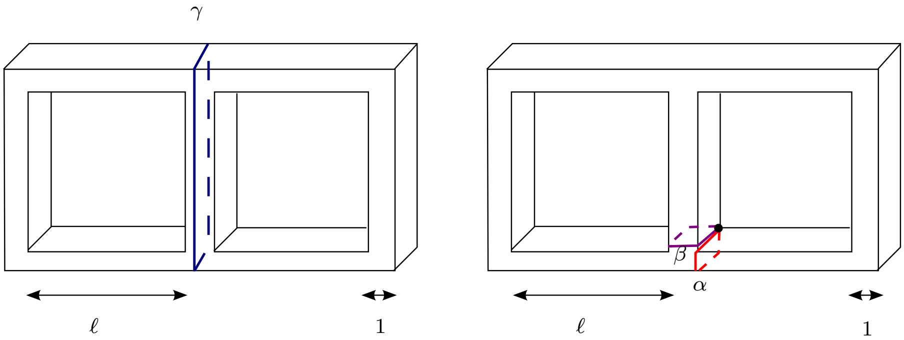

Figure 3.1 A piecewise linear double torus with area A such that the length of a shortest splitting closed curve is $\Omega(A)$ (left), but the length of a shortest homologically trivial non-contractible curve, concatenation of $\alpha\beta\alpha^{-1}\beta^{-1}$, has length $\Theta(1)$.

We also remark that, in (3), we cannot obtain a similar bound if we require the curve to be simple (and therefore to be *splitting* [10]). Indeed, Figure 3.1 shows that the minimum length of a shortest homologically trivial, non-contractible closed curve can become much larger if we additionally request the curve to be simple.

Proof of Theorem 3.1. We first recall that every surface has a unique structure as a smooth manifold, up to diffeomorphism, and we can therefore assume in the following that S is a smooth surface.

The first part of the proof is similar to Guth et al. [30, Lemma 5]. Define m_G to be the singular Riemannian metric given by endowing each triangle of G with the geometry of a Euclidean equilateral triangle of area 1 (and thus side length $2/\sqrt[4]{3}$): This is a genuine Riemannian metric except at a finite number of points, the set of vertices of G. The graph G is embedded on (S, m_G). Let γ be a closed curve $\gamma\colon \mathbb{S}^1 \to S$. Up to making it longer by a factor at most $\sqrt{1+\delta}$, we may assume that γ is piecewise linear and transversal to G. Now, for each triangle T and for every maximal part p of γ that corresponds to a connected component of $\gamma^{-1}(T)$, we do the following. Let x_0 and x_1 be the endpoints of p on the boundary of T. (If γ does not cross any of the edges of G, then it is contractible and the statement of the theorem is trivial.) There are two paths on the boundary of T with endpoints x_0 and x_1; we replace p with the shorter of these two paths. Since T is Euclidean and equilateral, elementary geometry shows that these replacements at most doubled the lengths of the curve. Now, the new curve lies on the graph G. We transform it with a homotopy into a no longer curve that is an actual closed walk in G, by simplifying it each time it backtracks. Finally, from a closed curve γ, we obtained a homotopic curve γ' that is a walk in G, satisfying $|\gamma'|_G = \sqrt[4]{3}/2\ |\gamma'|_{m_G} \leq \sqrt{1+\delta}\sqrt[4]{3}\ |\gamma|_{m_G}$.

The metric m_G satisfies our conclusion, except that it has isolated singularities. For the sake of concision we defer the smoothing procedure to Lemma 3.1. This lemma allows us to smooth and scale m_G to obtain a metric m, also with area n, that multiplies the length of all curves by at least $1/\sqrt{1+\delta}$ compared to m_G. This metric satisfies the desired properties. □

There remains to explain how to smooth the metric, which is done using partitions of unity.

Lemma 3.1 *With the notations of the proof of Theorem 3.1, there exists a smooth Riemannian metric m on S, also with area n, such that any closed curve γ in S satisfies $|\gamma|_m \geq |\gamma|_{m_G}/\sqrt{1+\delta}$.*

Proof. The idea is to smooth out each vertex v of G to make m_G Riemannian, as follows. Recall that $\delta > 0$ is fixed; $\varepsilon > 0$ will be determined later.

On the open ball $B(v, 2\varepsilon)$, consider a Riemannian metric m_v such that (i) m_v has area at most $\delta/3$, and (ii) any path in that ball is longer under m_v than under m_G. This is certainly possible provided ε is small enough: For example, take any diffeomorphism from $B(v, 1/2)$ onto the open unit disk D in the plane; define a metric on $B(v, 1/2)$ by taking the pullback metric of a multiple λ of the Euclidean metric on D, where λ is chosen large enough so that this pullback metric is larger than m_G (and thus (ii) is satisfied). If we take $\varepsilon > 0$ small enough, the restriction of this pullback metric to $B(v, 2\varepsilon)$ also satisfies (i).

We now use a partition of unity to define a smooth metric $\hat{m}$ that interpolates between m_G and the metrics m_v. By choosing an appropriate open cover, and therefore an appropriate partition of unity ρ, we obtain a metric $\hat{m} = \rho_G m_G + \sum_{v \in V} \rho_v m_v$ such that:

- outside the balls centered at a vertex v of radius 2ε, we have $\hat{m} = m_G$;
- inside a ball $B(v, \varepsilon)$, we have $\hat{m} = m_v$;
- in $B(v, 2\varepsilon) \setminus B(v, \varepsilon)$, the metric $\hat{m}$ is a convex combination of m_G and m_v.

The area of $\hat{m}$ is at most the sum of the areas of m_G and the m_v's, which is at most $n(1+\delta)$. Moreover, for any curve γ, we have $|\gamma|_{\hat{m}} \geq |\gamma|_{m_G}$.

Finally, we scale $\hat{m}$ to obtain the desired metric m with area n; for any curve γ, we indeed have $|\gamma|_m \geq |\gamma|_{\hat{m}}/\sqrt{1+\delta}$. □

3.2 From Discrete to Continuous Systolic Inequalities

Here we prove that, conversely, discrete systolic inequalities imply their Riemannian analogs. The idea is to approximate a Riemannian surface by the Delaunay triangulation of a dense set of points, and to use some recent results on intrinsic Voronoi diagrams on surfaces [16].

Theorem 3.2 *Let (S, m) be a Riemannian surface of genus g without boundary, of area A. Let $\delta > 0$. For infinitely many values of n, there exists a triangulated combinatorial surface (S, G) embedded on S with n triangles, such that every closed curve γ in (S, G) satisfies $|\gamma|_m \le (1+\delta)\sqrt{\frac{32}{\pi}}\sqrt{A/n}\ |\gamma|_G$.*

We have stated this result in terms of the number n of triangles; in fact, in the proof we will derive it from a version in terms of the number of vertices; Euler's formula and double counting imply that, for surfaces, the two versions are equivalent. Together with Hutchinson's theorem [33], this result immediately yields a new proof of Gromov's classical systolic inequality:

Corollary 3.2 *For every Riemannian surface (S, m) of genus g, without boundary, and area A, there exists a non-contractible curve with length at most $\frac{101.1}{\sqrt{\pi}}\sqrt{A/g}\log g$.*

Proof. Let $\delta > 0$, and let (S, G) be the triangulated combinatorial surface implied by Theorem 3.2 with $n \ge 6g - 4$ triangles. Euler's formula implies that the number v of vertices of G is at least g, hence we can apply Hutchinson's result [33], which yields a non-contractible curve γ on G with $|\gamma|_G \le 25.27\sqrt{(\frac{n}{2} + 2 - 2g)/g}\log g$. By Theorem 3.2, $|\gamma|_m \le \frac{101.08(1+\delta)}{\sqrt{\pi}}\sqrt{A/g}\log g$. □

On the other hand, using this theorem in the contrapositive together with the Buser–Sarnak examples (Theorem 2.1(4)) confirms the conjecture by Przytycka and Przytycki [52, Introduction]:

Corollary 3.3 *For any $\varepsilon > 0$, there exist arbitrarily large g and v such that the following holds: There exists a triangulated combinatorial surface of genus g, without boundary, with v vertices, on which the length of every non-contractible closed curve is at least $\frac{1-\varepsilon}{6}\sqrt{v/g}\log g$.*

Proof. Let $\varepsilon > 0$, let (S, m) be a Buser–Sarnak surface from Theorem 2.1(4), and let G be the graph obtained from Theorem 3.2 from (S, m), for some $\delta > 0$ to be determined later. Combining these two theorems, we obtain that every non-contractible closed curve γ in G satisfies

$$(1+\delta)\sqrt{\frac{32}{\pi}}\sqrt{\frac{A}{n}}\ |\gamma|_G \ge \frac{2}{3\sqrt{\pi}}\sqrt{\frac{A}{g}}\log g - c''',$$

where $A = 4\pi(g-1)$. If δ was chosen small enough (say, such that $1/(1+\delta) \ge 1-\varepsilon/2$), and g was chosen large enough, we have $|\gamma|_G \ge \frac{1-\varepsilon}{3\sqrt{8}}\sqrt{\frac{n}{g}}\log g$. Finally, we have $n \ge 2v$ by Euler's formula. □

Before delving into the proof of Theorem 3.2, we introduce a refinement of the well-known injectivity radius. The *strong convexity radius* at a point x in a Riemannian surface (S,m) is the largest radius ρ_x such that for every $r < \rho_x$ the ball of radius r centered at x is strongly convex, that is, for any $p,q \in B(x,r)$ there is a unique shortest path in (S,m) connecting p and q, this shortest path lies entirely within $B(x,r)$, and moreover no other geodesic connecting p and q lies within $B(x,r)$; we refer to Klingenberg [37, Def. 1.9.9] for more details. The strong convexity radius is positive at every point, and its value on the surface is continuous (see also Dyer, Zhang, and Möller [16, Sect. 3.2.1]). It follows that for every compact Riemannian surface (S,m), the *strong convexity radius* of (S,m), namely, the infimum of the strong convexity radii of the points in (S,m), is strictly positive. We will need the following lemma, which is a result of of Dyer, Zhang, and Möller [16, Corollary 2] (see also Leibon [42, Theorem 1] for a very related theorem):

Lemma 3.2 *Let (S,m) be a Riemannian surface without boundary, let $\rho' > 0$ be less than half the strong convexity radius of (S,m), and let P a point set of S in general position such that for every x on S, there exists a point p of P such that $d_m(x,p) \le \rho'$. Then the Delaunay graph of P is a triangulation of S, and its edges are shortest paths.*

Proof of Theorem 3.2. Let η, $0 < \eta < 1/2$ be fixed, and $\varepsilon > 0$ to be defined later (depending on η). Let P be an ε-separated net on (S,m), that is, P is a point set such that any two points in P are at distance at least ε, and every point in (S,m) is at distance smaller than ε from a point in P. For example, if we let P be the centers of an inclusionwise maximal family of disjoint open balls of radius $\varepsilon/2$, then P is an ε-separated net. In the following we put P in general position by moving the points in P by at most $\eta\varepsilon$; in particular, no point in the surface is equidistant with more than three points in P.

Let $P = \{p_1, \dots, p_v\}$, and let

$$V_i := \{x \in (S,m) \mid \forall j \neq i, d(x,p_i) \le d(x,p_j)\}$$

be the Voronoi region of p_i. Since every point of (S,m) is at distance at most $(1+\eta)\varepsilon$ from a point in P, each Voronoi region V_i is included in a ball of radius $(1+\eta)\varepsilon$ centered at p_i. Define the Delaunay graph of P to be the intersection graph of the Voronoi regions, and note that if $V_i \cap V_j \neq \emptyset$, then the corresponding neighboring points of the Delaunay graph are at distance at most $2(1+\eta)\varepsilon$.

Assume that ε is small enough so that $(1+\eta)\varepsilon$ is less than half the strong convexity radius. Lemma 3.2 implies that the Delaunay graph, which we denote by G, can be embedded as a triangulation of S with shortest paths representing the edges.

Consider a closed curve γ on G. Since neighboring points in G are at distance no greater than $2(1+\eta)\varepsilon$ on (S,m), we have $|\gamma|_m \le 2(1+\eta)\varepsilon|\gamma|_G$. To obtain the claimed bound, there remains to estimate the number v of points in P. By compactness, the Gaussian curvature of (S,m) is bounded from above by a constant K. By the Bertrand–Diquet–Puiseux theorem, the area of each ball of radius $\frac{1-2\eta}{2}\varepsilon$ is $\pi(1-2\eta)^2\frac{\varepsilon^2}{4} - K\pi(1-2\eta)^4\frac{\varepsilon^4}{16} + o(\varepsilon^4) \geq \pi(1-2\eta)^3\frac{\varepsilon^2}{4}$ if $\varepsilon > 0$ is small enough. Since the balls of radius $(1-2\eta)\frac{\varepsilon}{2}$ centered at P are disjoint, their number v is at most $A/(\pi(1-2\eta)^3\frac{\varepsilon^2}{4})$. In other words, $\varepsilon \leq \frac{2}{\sqrt{\pi(1-2\eta)^3}}\sqrt{A/v}$. Putting together our estimates, we obtain that

$$|\gamma|_m \leq \frac{4(1+\eta)}{\sqrt{\pi(1-2\eta)^3}}\sqrt{\frac{A}{n/2-2g+2}}\ |\gamma|_G,$$

where n is the number of triangles of G. Thus, if $\varepsilon > 0$ is small enough, n can be made arbitrarily large, and the previous estimate implies, if η was chosen small enough (where the dependency is only on δ) that $|\gamma|_m \leq (1+\delta)\sqrt{\frac{32}{\pi}}\sqrt{\frac{A}{n}}\ |\gamma|_G$.
□

Remark on orientability. Notice that Theorems 3.1 and 3.2 hold for non orientable surfaces with the same proofs. We stated the continuous systolic inequality for orientable surfaces. As observed by Gromov [28, p. 306] a double cover argument shows that the same results hold (up to a multiplicative constant factor) for the systole of non-orientable surfaces other than the projective plane. For the projective plane, a systolic inequality also holds, for which the exact constant is known and corresponds to metrics of constant positive curvature [54]. Therefore, since our results do not rely on orientability, the discrete systolic inequalities hold for all surfaces, with similar dependence on the Euler genus, up to a multiplicative factor. Notice that when we talked about homology no coefficients were specified. It is customary to assume $\mathbb{Z}$ coefficients for orientable manifolds and $\mathbb{F}_2$ for non orientable ones.

4 Computing Short Pants Decompositions

Recall that the problem of computing a shortest pants decomposition for a given surface is open, even in very special cases. In this section, we describe an efficient algorithm that computes a short pants decomposition on a triangulation. Technically, we allow several curves to run along a given edge of the triangulation, which is best formalized in the dual cross-metric setting. If g is fixed, the length of the pants decomposition that we compute is of the order of the square root of the number of vertices:

Theorem 4.1 *Let (S, G^*) be a (trivalent, unweighted) cross-metric surface of genus $g \geq 2$, with n vertices, without boundary. In $O(gn)$ time, we can compute a pants decomposition $(\gamma_1, \ldots, \gamma_{3g-3})$ of S such that, for each i, the length of γ_i is at most $C\sqrt{gn}$ (where C is some universal constant).*

With a little more effort, we can obtain that the length of γ_i is at most $C\sqrt{in}$ but we focus on the weaker bound for the sake of clarity.

The inspiration for this theorem is a result by Buser [5], stating that in the Riemannian case, there exists a pants decomposition with curves of length bounded by $3\sqrt{gA}$. The proof of Theorem 4.1 consists mostly of translating Buser's construction to the discrete setting and making it algorithmic. The key difference is that for the sake of efficiency, unlike Buser, we cannot afford to shorten the curves in their homotopy classes, and we have to use contractibility tests in a careful manner.

Given simple, disjoint closed curves Γ in general position on a (possibly disconnected) cross-metric surface (S, G^*), cutting S along Γ, and/or restricting to some connected components, gives another surface S', and restricting G^* to S' naturally yields a cross-metric surface that we denote by $(S', G^*_{|S'})$. To simplify notation we denote by $|c|$ (instead of $|c|_{G^*}$) the length of a curve c on a cross-metric surface (S, G^*).

A key step towards the proof of Theorem 4.1 is the following proposition, which allows us to effectively cut a surface with boundary along closed curves of controlled length.

Proposition 4.1 *Let (S, G^*) be a possibly disconnected cross-metric surface, such that every connected component has non-empty boundary and admits a pants decomposition. Let n be the number of vertices of G^* in the interior of S. Assume moreover that $|\partial S| \leq \ell$, where ℓ is an arbitrary positive integer.*

We can compute a family Δ of disjoint simple closed curves of (S, G^*) that splits S into one pair of pants, zero, one, or more annuli, and another possibly disconnected surface S' containing no disk component, such that $|\partial S'| \leq \ell + 4n/\ell + 2$. The algorithm takes as input (S, G^*), outputs Δ and $(S', G^*_{|S'})$, and takes linear time in the complexity of (S, G^*).

We first show how Theorem 4.1 can be deduced from this proposition. It relies on computing a good approximation of the shortest non-contractible closed curve, cutting along it, and applying Proposition 4.1 inductively:

Proof of Theorem 4.1. To prove Theorem 4.1, we consider our cross-metric surface without boundary (S, G^*), and we start by computing a simple non-contractible curve γ whose length is at most twice the length of the shortest non-contractible closed curve. Such a curve can be computed in $O(gn)$ time [7, Prop. 9] (see also Erickson and Har-Peled [20, Corollary 5.8]) and has length at most $C\sqrt{n}$, where C is a universal constant, see Section 3. This gives a surface $S^{(1)}$ with two boundary components.

Let us define the sequence $\ell_k = C\sqrt{kn}$ for some constant C. We then proceed inductively, applying Proposition 4.1 with $\ell = \ell_k$ to $S^{(k)}$, in order to obtain another surface $S^{(k)\prime}$, from which we remove all the pairs of pants and annuli. We denote the resulting surface by $S^{(k+1)}$ and repeat until we obtain a surface $S^{(m)}$ that is empty. Note that, for every k, $S^{(k)}$ contains no disk, annulus, or pair of pants, and that every application of Proposition 4.1 gives another pair of pants. Therefore, we obtain a pants decomposition of S by taking the initial curve γ together with the union of the collections of curves Δ given by successive applications of Proposition 4.1 and removing, for any subfamily of Δ of several homotopic curves, all but the shortest one of them. The number of applications of Proposition 4.1 is bounded by the number of pants in a pants decomposition, which is $2g - 2$.

There remains to bound the length of the closed curves in the pants decomposition. A small computation shows that $\ell_k + 4n/\ell_k + 2 \leq \ell_{k+1}$ for C large enough and $k \leq 3n$, which holds since $k \leq 3g - 3 \leq 3n$. Now, $|\partial S^{(1)}| \leq C\sqrt{n} = \ell_1$, and applying Proposition 4.2 inductively on $S^{(k-1)}$ with $\ell = \ell_{k-1}$ shows that $|\partial S^{(k)}| \leq \ell_k = C\sqrt{kn}$. Therefore, the length of the kth closed curve of the pants decomposition is at most $C\sqrt{kn}$. The total complexity of this algorithm is $O(gn)$ since we applied $O(g)$ times Proposition 4.1. □

Now, onwards to the proof of the main proposition.

Proof of Proposition 4.1. We will only describe how Δ is computed, since one directly obtains S' by cutting along Δ and discarding the annuli and one pair of pants.

The idea is to *shift* the boundary components simultaneously until one boundary component *splits*, or two boundary components *merge*. This is analogous to Morse theory on the surface with the function that is the distance to the boundary. In this way, we choose the homotopy classes of the curves in Δ, but in order to control their length we actually do some backtracking before splitting or merging.

Initially, let $\Gamma = (\gamma^1, \ldots, \gamma^k)$ be (curves infinitesimally close to) the boundaries of S. We will shift these curves to the right while preserving their simplicity, disjointness, and homotopy classes. We orient each γ^i so that it has the surface to its right at the start. In particular, at any time of the algorithm, any two curves are to the right of each other.

Shifting phase: The idea of shifting a closed curve γ^i one step to the right is to push it so that every point of the resulting curve is exactly at distance one from the original curve. The shifting phase consists of shifting every curve in Γ one step to the right, and to reiterate. During this process, curves will collide, which will allow us to build the new curves of the pants decomposition.

A *piece* of a curve in Γ is a maximal subpath inside a face of G^*. We say that two distinct pieces of curves in Γ are *tangent* if (i) they are not consecutive pieces along the same curve and (ii) there is a path on the surface that starts to the right of one piece, arrives to the right of the other, crosses no piece, and crosses at most one edge of G^*, see Figures 4.1(a, c, e).

Basically, tangencies are the obstacles to shifting the curve to the right. On the other hand, in a tangency, we can *rewire* the curves as shown on Figure 4.1(b, d, f), by locally exchanging the connections between the pieces without changing the orientations of the pieces. Our algorithm needs first to remove all tangencies in Γ, by repeating the following steps while there exists a tangency:

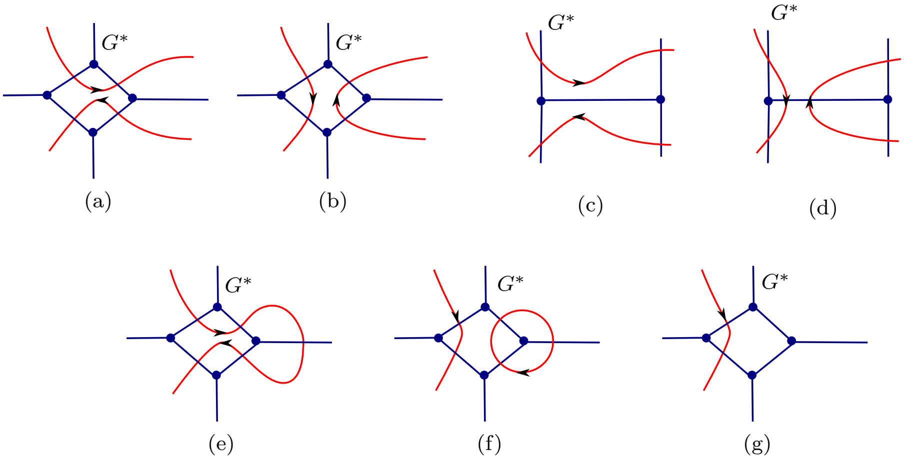

Figure 4.1 (a) Two tangent pieces of curves lying in the same face. (b) The rewiring of these curves. (c) Two tangent pieces of curves lying in adjacent faces. (d) The rewiring of these curves. (e) A curve that is tangent with itself. (f) Its rewiring. (g) The result after discarding the contractible subcurve.

– If the pieces involved in the tangency belong to the same closed curve, then, by the chosen orientation, the rewiring necessarily transforms the initial curve into exactly two curves, which we test for contractibility. If one of them is contractible, we discard it (Figure 4.1(g)) and continue with the other one. Otherwise, both are non-contractible; the shifting phase is over, and we go to the splitting phase below.
– If the tangency involves pieces belonging to different closed curves in Γ, the rewiring transforms the two curves into a single curve; the shifting phase is over, and we go to the merging phase below.

At this step, we removed all tangencies without entering the splitting or the merging phase. Since G^* is trivalent, if γ^i were to cross consecutively two edges that are incident to the same vertex v to the right of the curve, it would form a tangency with the third edge incident to v, a contradiction. Thus, the local picture is as on Figure 4.2(a): The edges of G^* to the right of γ^i, incident to the faces traversed by γ^i, form a cycle (the horizontal line in Figure 4.2); each edge incident to the cycle is either to its left or to its right, and these edges are attached to the cycle by distinct vertices; and γ^i crosses exactly those edges of G^* that are to its left. We transform γ^i so that it now crosses exactly those edges that are to the right of the cycle, as shown on Figure 4.2(b). The absence of tangencies ensures that this still gives disjoint simple curves, with the same homotopy classes; of course, this operation may create one or several tangencies (in particular, a face of G^* may now be traversed by several pieces).

When this is done, we repeat the entire shifting phase (again starting with the tangency detection). Thus, the shifting phase is repeated over and over, until we enter the splitting phase or the merging phase below. Before describing

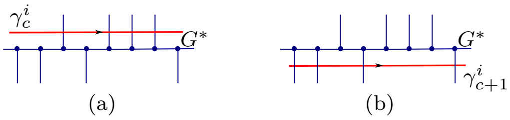

Figure 4.2 Shifting a curve one step to its right.

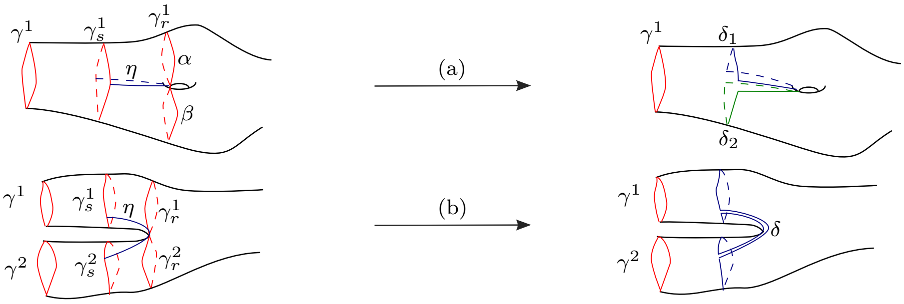

Figure 4.3 (a) Splitting phase. (b) Merging phase.

these phases, let us describe some properties that are satisfied when we exited the shifting phase. Let r be the integer such that each curve has been pushed r steps to the right. For each i, $1 \leq i \leq k$, and each c, $0 \leq c \leq r$, let γ_c^i be the curve γ^i pushed by c steps. Note that by construction, the distance between any point of γ_c^i and the curve γ_{c-1}^i is exactly one. Let s denote the largest $c \leq r$ such that $\sum_{i=1}^k |\gamma_c^i| \leq \ell$. (Remember that this is the case for $c = 0$ by hypothesis.)

Splitting phase: We arrived to the splitting phase because two pieces of the same curve became tangent, and after rewiring, both of the new subcurves are non-contractible, as is pictured on the top of Figure 4.3. The purpose of the splitting phase is to choose geometric representatives of curves in these homotopy classes. For simplicity, let γ^1 denote the curve that became tangent with itself during the shifting phase. First, for every $i \neq 1$, we add γ_s^i to the family Δ. By assumption, γ^1 splits into two non-contractible closed curves α and β. Let η be the shortest path with endpoints on γ_s^1 going through the splitting tangency between α and β. This path can be computed in linear time (in the complexity of the portion of the surface swept during the shifting phase) by backtracking from γ_r^1 to γ_s^1, and adding pieces of η at every step. The path η cuts γ_s^1 into two subpaths μ and ν. We denote by δ_1 the concatenation of μ and η, and by δ_2 the concatenation of ν and η. To finish the splitting phase, we add δ_1 and δ_2 to the family Δ.

Merging phase: We arrived to the merging phase because two distinct shifted curves became tangent in the shifting phase (Figure 4.3, bottom); and we rewired them, obtaining a curve homotopic to their concatenation. The purpose of the merging phase is to choose a geometric representative in this homotopy class. For simplicity, let us denote by γ^1 and γ^2 two curves that became tangent during the shifting phase. First, for every $i \neq 1, 2$, we add γ_s^i to the family Δ. Let η be the shortest path from γ_s^1 and γ_s^2 (as above, we can compute it in linear time). The curve δ is defined by the concatenation $\eta^{-1} \cdot \gamma_s^1 \cdot \eta \cdot \gamma_s^2$. To finish the merging phase, we add δ to Δ.

Analysis: After splitting or merging, we added curves to Δ that cut the surface into an additional pair of pants, (possibly) some annuli, and the remaining surface S'. Observe that we did not add any contractible closed curve to Δ; thus, S' has no connected component that is a disk. There remains to prove that the length of the boundary S' satisfies $|\partial S'| \leq \ell + 4n/\ell + 2$. The key quantitative idea is the way in which the value of s was chosen: If s was equal to r (perhaps the most natural strategy), the boundary of S' would contain (at least) one curve γ_r^i, and we would have no control on its length. On the opposite, if we had chosen $s = 0$, we would have no control on the lengths of the arcs η involved in the merging or the splitting. The choice of s gives the right trade-off in-between: the lengths of the curves γ_i^s are controlled by this threshold, while the lengths of the arcs are controlled by the area of the annulus between γ_s^i and γ_r^i. We now make this explanation precise.

Lengths after the splitting phase. After a splitting phase with the curve γ^1, the boundary $\partial S'$ of S' consists of all the other curves γ_s^i in Γ and of the two new curves, whose sum of the lengths is bounded by $|\gamma_s^1| + 2|\eta|$. Hence $|\partial S'| \leq |\gamma_s^1| + 2|\eta| + \sum_{i=2}^k |\gamma_s^i|$, which is at most $\ell + 2|\eta|$ by the choice of s. Furthermore, by construction, $|\eta| \leq 2(r-s)+1$, as every step of shifting adds at most 2 to the length of η, and it may cost an additional 1 to cross the last tangency edge.

Lengths after the merging phase. After a merging phase with the curves γ^1 and γ^2, the boundary $\partial S'$ of S' consists of all the other curves γ_s^i of Γ, and of the new curve, whose length is bounded by $|\gamma_s^1|+|\gamma_s^2|+2|\eta|$. Hence similarly, $|\partial S'| \leq \ell + 2|\eta|$. Furthermore, by construction, we also have $|\eta| \leq 2(r-s)+1$.

Final analysis. Thus, after either the splitting or the merging phase, we proved that $|\partial S'| \leq \ell + 4(r-s) + 2$. To conclude the analysis, there only remains to prove that $r - s \leq \frac{n}{\ell}$.

Let $c \in \{s, \ldots, r-1\}$. The curves γ_c^i and γ_{c+1}^i bound an annulus K_c^i. We claim that the number $A(K_c^i)$ of vertices in the interior of this annulus, its *area*, is at least $|\gamma_{c+1}^i|$. This follows from the shifting procedure (refer back to Figure 4.2—remember that G^* is trivalent) and from the fact that the contractible closed curves possibly stemming from γ_c^i only make the area larger, by definition of a tangency.

For $c \in \{s, \ldots, r-1\}$ and $i \in \{1, \ldots, k\}$, the annuli K_c^i have disjoint interiors, so the sum of their areas is at most n. By the above formula, this sum is at least $\sum_{j=s}^{r-1} U_{c+1}$, where $U_c = \sum_{i=1}^k |\gamma_c^i|$. On the other hand, we have $U_{c+1} \geq \ell$ if $s \leq c \leq r-1$, by definition of s. Putting all together, we obtain $n \geq (r-s)\ell$, so $r - s \leq \frac{n}{\ell}$.

Complexity: At the start, the complexity of the set of curves is bounded by the complexity of (S, G^*), and by construction, during the algorithm, the complexity of the set of curves is always linear in n. The complexity of the splitting phase or the merging phase is thus also linear in n. The complexity of outputting the new surface $(S', G^*_{|S'})$ is linear in the complexity $\partial S'$, which is, by construction, also linear in n. To conclude, it suffices to prove that the whole shifting phase takes linear time. We study separately the tangency detection step and the contractibility tests.

Tangency detection. Remember that our curves are stored on the cross-metric surface: At each time, we maintain the arrangement A of the overlay of the curves in Γ with G^*. On each face f of A, we store a list $L(f)$ of pointers to the pieces incident to that face and having that face to their right. Thus, f contains a tangency if and only if $|L(f)| \geq 2$. Similarly, if g is a face of A incident to f via an edge of G^*, the union of f and g contains a tangency if and only if $|L(f) \cup L(g)| \geq 3$, or $|L(f) \cup L(g)| = 2$ and the two corresponding pieces are not consecutive. These properties can be tested in constant time.

As we push the curves, we update the corresponding lists $L(f)$. At the start of the shifting, or once the curves have been pushed by one step, we first detect the tangencies within the same face f, and deal with them, updating the lists $L(f)$. At this step, there is at most one piece per face of G^*. For every piece of Γ, we mark the edges incident to the face to the right of that piece; as soon as one edge is marked from both sides, and the two corresponding pieces are not consecutive, there is a tangency, which we handle immediately. The running time for one tangency detection step is the total complexity of the faces that are incident to the curves, and to their right; the sum of these complexities is linear in n. (Note that we only care about the part of the surface that is to the right of the curves; the data structures involving faces of the remaining part of the surface are irrelevant.)

Contractibility tests. Finally, to perform a contractibility test on two subcurves α and β, we perform a tandem search on the surfaces bounded by α and β, and stop as soon as we find a disk. If we find one, the complexity in the tandem search is at most twice the complexity of this disk, which is immediately discarded and never visited again. If we do not find a disk, the complexity is linear in n, but the shifting phase is over. Therefore, the total complexity of the contractibility tests is linear in the number of vertices swept by the shifting phase or in the disks, until the very last contractibility test, which takes time linear in n. In the end, the shifting phase takes time linear in n, which concludes the complexity analysis. □

5 Shortest Cellular Graphs with Prescribed Combinatorial Maps

Guth, Parlier, and Young proved the following result:

Theorem 5.1 ([30, Theorem 2]) *For any $\varepsilon > 0$, the following holds with probability tending to one as n tends to ∞: A random (trivalent, unweighted)* *cross-metric surface without boundary with n vertices has no pants decomposition of length at most $n^{7/6-\varepsilon}$.*

In this statement, two cross-metric surfaces are regarded as equal if some self-homeomorphism of the surface maps one to the other. (Note that vertices, edges, and faces are unlabeled.) As a side remark, by a simple argument, we are actually able to strengthen this result, by replacing, in the statement above, "pants decomposition" by "genus zero decomposition". We defer the proof of this side result, independent of the following considerations, to Appendix B.

The main purpose of this section is to provide an analogous statement, not for pants decompositions or genus zero decompositions, but for cut graphs (or, actually, for arbitrary cellular graphs) with a prescribed combinatorial map. We essentially prove that, for any combinatorial map M of any cellular graph embedding (in particular, of any cut graph) of genus g, there exists a (trivalent, unweighted) cross-metric surface S with n vertices such that any embedding of M on S has length $\Omega(n^{7/6})$. We are not able to get this result in full generality, but are able to prove that it holds for infinitely many values of g. On the other hand, the result is stronger since, as in Theorem 5.1, it holds "asymptotically almost surely" with respect to the uniform distribution on unweighted trivalent cross-metric surfaces with given genus and number of vertices.

Let (S, G^*) be a cross metric surface without boundary, and M a combinatorial map on S. The *M-systole* of (S, G^*) is the minimum among the lengths of all graphs embedded in (S, G^*) with combinatorial map M. Given g and n, we consider the set $\mathcal{S}(g, n)$ of trivalent unweighted cross-metric surfaces of genus g, without boundary, and with n vertices, where we regard two cross-metric surfaces as equal if some self-homeomorphism of the surface maps one to the other. This refines the model introduced by Gamburd and Makover [24]. Here is our precise result:

Theorem 5.2 *Given strictly positive real numbers p and ε, and integers n_0 and g_0, there exist $n \geq n_0$ and $g \geq g_0$ such that, for any combinatorial map M of a cellular graph embedding with genus g, with probability at least $1 - p$, a cross-metric surface chosen uniformly at random from $\mathcal{S}(g, n)$ has M-systole at least $n^{7/6-\varepsilon}$.*

We can obtain a similar result in the case of polyhedral triangulations, namely, metric spaces obtained by gluing n equilateral Euclidean triangles with sides of unit length. We first note that an element of $\mathcal{S}(g, n)$ naturally corresponds to a polyhedral triangulation by gluing equilateral triangles of unit side length on the vertices. The notion of M-systole is defined similarly in this setting, and we now prove that Theorem 5.2 implies an analogous result for polyhedral triangulations:

Theorem 5.3 *Given strictly positive real numbers p and ε, and integers n_0 and g_0, there exist $n \geq n_0$ and $g \geq g_0$ such that, for any combinatorial map M of a cellular graph embedding with genus g, with probability at least $1 - p$,*

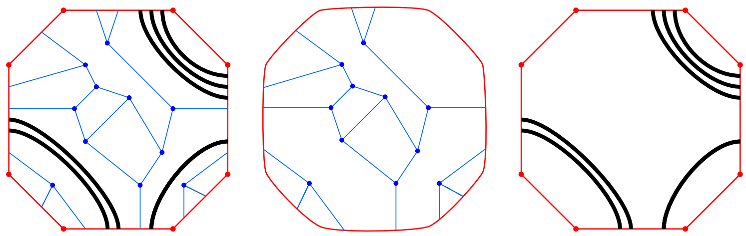

a. b. c.

Figure 5.1 a. The graph H, obtained after cutting S open along C. The vertices in B (on the outer face) and the vertices of G^* (not on the outer face) are shown. The chords are in thick (black) lines. b. The graph H_1. c. The graph H_2.

a polyhedral triangulation chosen uniformly at random from $\mathcal{S}(g,n)$ has M-systole at least $n^{7/6-\varepsilon}$.

5.1 Proof of Theorem 5.2

The general strategy of the proof of Theorem 5.2 is inspired by Guth, Parlier and Young [30], who proved a related bound for pants decompositions; however, the details of the method are rather different. Our main tool is the following proposition.

Proposition 5.1 *Given integers g, n, and L, and a combinatorial map M of a cellular graph embedding of genus g, at most*

$$f(g,n,L) = 2^{O(n)} L\,(L/g+1)^{12g-9}$$

cross-metric surfaces in $\mathcal{S}(g,n)$ have M-systole at most L.

Proof. First, note that it suffices to prove the result for cut graphs with minimum degree at least three. Indeed, one can transform any cellular graph embedding into such a cut graph by removing edges, removing degree-one vertices with their incident edges, and *dissolving* degree-two vertices, namely, removing them and replacing the two incident edges with a single one. For a combinatorial map M with minimum degree at least three, Euler's formula and double-counting immediately imply that M has at most $4g-2$ vertices and $6g-3$ edges. Given a cross-metric surface (S,G^*) in $\mathcal{S}(g,n)$, let C be a cut graph of genus g with combinatorial map M and length at most L.

Let H' be the graph that is the overlay of G^* and C. Cutting S along C yields a topological disk D, and transforms H' into a connected graph H (Figure 5.1(a)) embedded in the plane, where the outer face corresponds to the copies of the vertices and edges of the cut graph C. The set B of vertices of degree two on the outer face of H exactly consists of the copies of the vertices of C; there are at most $12g-6$ of these. A *side* of H is a path on the boundary of D that joins two consecutive points in B.

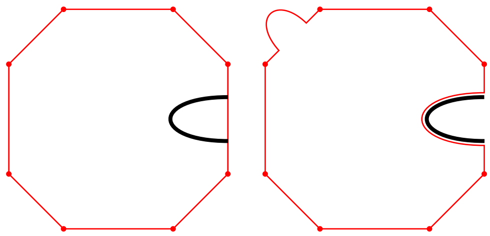

Figure 5.2 The exchange argument to prove (i).

Given the combinatorial map of H in the plane, we can (almost) recover the combinatorial maps corresponding to H' and to (S, G^*). Indeed, the set B of vertices of degree two on the outer face of H determines the sides of D. The correspondence between each side of D and each edge of the combinatorial map M is completely determined once we are given the correspondence between a single half-edge on the outer face of H and a half-edge of M; in turn, this determines the whole gluing of the sides of H and completely reconstructs H' with C distinguished. Finally, to obtain G^*, we just "erase" C. Therefore, one can reconstruct the combinatorial map corresponding to the overlay H' of G^* and C, just by distinguishing one of the $O(L)$ half-edges on the outer face of H.

A *chord* of H is an edge of H that is not incident to the outer face but connects to vertices incident to the outer face. Two chords are *parallel* if their endpoints lie on the same pair of sides of D. We claim that we can assume the following:

(i) no chord has its endpoints on the same side of H (Figure 5.2(a) shows an example not satisfying this property);

and that (at least) one of the two following conditions holds:

(ii) the subgraph of H between any two parallel chords only consists of other parallel chords (Figure 5.3(a) shows an example not satisfying this property), or

(ii') there are two parallel chords such that the subgraph of H between them contains all the interior vertices of H.

Indeed, without loss of generality, we can assume that our cut graph C has minimum length among all cut graphs of (S, G^*) with combinatorial map M. If a chord violates (i), one could shorten the cut graph by sliding a part of the cut graph over the chord (Figure 5.2), which is a contradiction.

For (ii) and (ii'), the basic idea is to use a similar exchange argument as to prove (i), but we need a perturbation argument as well. Specifically, let us temporarily perturb the crossing weights of the edges of G^* as follows: The weight of each edge e of G^* becomes $1 + w_e$, where the w_e's are real numbers that are linearly independent over $\mathbb{Q}$ (e.g., independent and identically distributed random) and strictly between 0 and $1/L$. Let C be a shortest embedded graph with combinatorial map M under this perturbed metric.

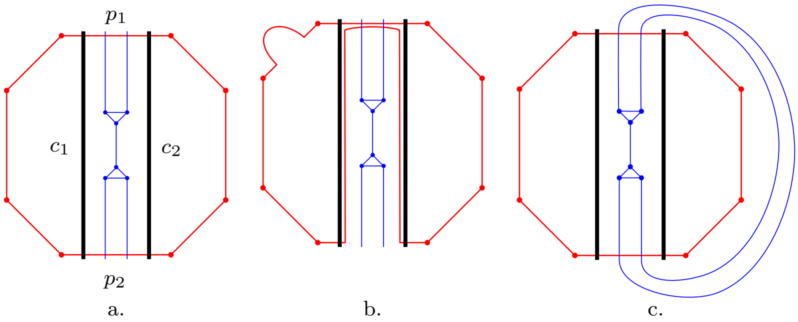

Figure 5.3 a.: Two chords violating (ii). b.: The exchange argument, in case p_1 and p_2 have different perturbed lengths. c.: A schematic view of the situation, in case p_1 and p_2 have the same perturbed length.

It is easy to see that C is also a shortest embedding with combinatorial map M under the unweighted metric: Indeed, two cut graphs C_1 and C_2 with respective (integer) lengths $\ell_1 < \ell_2 \leq L$ in the unweighted metric have respective lengths $\ell'_1 < \ell'_2$ in the perturbed metric, since the perturbation increases the length of each edge by less than $1/L$.

We claim that either (ii) or (ii') holds for this choice of C. Assume that (ii) does not hold; we prove that (ii') holds. So the region R of D between two parallel chords c_1 and c_2 of D contains internal vertices; without loss of generality (by (i)), assume that the region R contains no other chord in its interior. Let p_1 and p_2 be the two subpaths of the cut graph on the boundary of R. If p_1 and p_2 have different lengths under the perturbed metric, e.g., p_1 is shorter, then we can push the part of p_2 to let it run along p_1 and shorten the cut graph (Figure 5.3(b)), which is a contradiction. Therefore, p_1 and p_2 have the same length under the perturbed metric, which implies that they cross exactly the same set E of edges of G^*, since the weights are linearly independent over $\mathbb{Q}$. (We exclude from E the edges on the endpoints of p_1 and p_2.) Since none of the edges in E are chords, all the endpoints of the edges in E belong to R (Figure 5.3(c)), which implies (ii') by connectivity of G^*. This concludes the proof of the claim.

We now estimate the number of possible combinatorial maps for H, by "splitting" it into two connected plane graphs H_1 and H_2, estimating all possibilities of choosing each of these graphs, and estimating the number of ways to combine them.

Let H_1 be the graph (see Figure 5.1(b)) obtained from H by removing all chords and dissolving all degree-two vertices (which are either in B or endpoints of a chord). H_1 is connected, trivalent, and has at most n vertices not incident to the outer face, so $O(n)$ vertices in total. By a classical calculation (see for example [30, Lemma 4]), there are thus $2^{O(n)}$ possible choices for the combinatorial map of this planar trivalent graph H_1.

On the other hand, let H_2 be the graph (see Figure 5.1(c)) obtained from H by removing internal vertices together with their incident edges and dissolving all degree-two vertices not in B. Since the chords are non-crossing and connect distinct sides of D, the pairs of sides connected by at least one chord form a subset of a triangulation of the polygon having one vertex per side of D. To describe H_2, it therefore suffices to describe a triangulation of this polygon with at most $12g-6$ edges, which makes $2^{O(g)} = 2^{O(n)}$ possibilities, and to describe, for each of the $12g-9$ edges of the triangulation, the number of parallel chords connecting the corresponding pair of sides. Since there are at most L chords, the number of possibilities for these numbers equals $\{(x_1, \ldots, x_{12g-9}) \mid x_i \geq 0, \sum_i x_i \leq L\}$, which is the number of weak compositions of L into $12g-8$ parts, namely

$$\binom{L+12g-9}{12g-9} \leq \left(\frac{e(L+12g-9)}{12g-9}\right)^{12g-9} = O\left((L/g+1)^{12g-9}\right) \times 2^{O(n)},$$

the inequality being standard (or following from Stirling's formula).

Finally, in how many ways can we combine given H_1 and H_2 to form H? Let us first assume that (ii) holds; the parallel chords joining the same pair of sides are consecutive, so choosing the position of a single chord fixes the position of the other chords parallel to it. Therefore, given H_1, we need to count in how many ways we can insert the $O(g)$ vertices of B on H_2 into H_1, and similarly the $O(g)$ intervals where endpoints of chords can occur, respecting the cyclic ordering. After choosing the position of a distinguished vertex of H_2, we have to choose $O(g)$ positions on the edges of the boundary of H_1, possibly with repetitions, which leaves us with $\binom{O(n+g)}{O(g)} \leq 2^{O(n+g)} = 2^{O(n)}$ possibilities. In case (ii') holds, a very similar argument gives the same result.

The claimed bound follows by multiplying the number of all possible choices above: there are $O(L)$ choices for the distinguished half-edge of the outer face of H, $2^{O(n)}$ choices for H_1, $O\left((L/g+1)^{12g-9}\right) \times 2^{O(n)}$ choices for H_2, and $2^{O(n)}$ possibilities for combining H_1 and H_2. □

Proof of Theorem 5.2. Let g_0, n_0, p, ε be as indicated. Euler's formula implies that a cross-metric surface with n vertices has genus $g \leq (n+2)/4$. We now show that, if n is large enough,

$$\sum_{g=g_0}^{(n+2)/4} f(g, n, n^{7/6-\varepsilon}) \leq n^{(1-\varepsilon)n/2}. \quad (*)$$

Indeed, by Proposition 5.1 we have

$$f(g, n, n^{7/6-\varepsilon}) \leq 2^{C_0 n} \left(n^{7/6-\varepsilon}/g + 1\right)^{12g-9}$$

for some constant C_0. We need to sum up these terms from $g = g_0$ to $(n+2)/4$. For n large enough, the largest term in this sum is for $g = (n+2)/4$. Thus the desired sum is bounded from above by

$$n 2^{C_0 n} \left(4n^{1/6-\varepsilon} + 1\right)^{12(n+2)/4-9},$$

which is at most $2^{C_1 n} n^{(1/6-\varepsilon)3n}$ (for n large enough, for some constant C_1), which in turn is at most $n^{(1-\varepsilon)n/2}$ for n large enough.

Furthermore, let $h(g,n) = |\mathcal{S}(g,n)|$ be the number of (connected) cross-metric surfaces with genus g and n vertices. We have $\sum_{g=0}^{(n+2)/4} h(g,n) \geq e^{Cn} n^{n/2}$ if n is large enough and even, for some absolute constant C; this is probably folklore, and we provide a proof, deferred to Lemma 5.1. But, if g is fixed, $h(g,n) = O(e^{C'n})$ for some constant C' [30, Lemma 4]. Thus, since g_0 is fixed, there is a constant C'' such that, for n large enough and even, $\sum_{g=g_0}^{(n+2)/4} h(g,n) \geq e^{C''n} n^{n/2}$ (**).

Choose any (even) $n \geq n_0$ such that $n^{-\varepsilon n/2} e^{-C''n} \leq p$ and such that (*) and (**) hold. Thus, we have

$$\sum_{g=g_0}^{(n+2)/4} f(g,n,n^{7/6-\varepsilon}) \leq p \sum_{g=g_0}^{(n+2)/4} h(g,n),$$

which implies that for some $g \geq g_0$,

$$f(g,n,n^{7/6-\varepsilon})/h(g,n) \leq p$$

and the denominator is non-zero. In other words, among all $h(g,n)$ cross-metric surfaces with genus g and n vertices, for any combinatorial map M of a cellular graph embedding of genus g, a fraction at most p of these surfaces have an embedding of M with length at most $n^{7/6-\varepsilon}$. $\square$

We remark that a tighter estimate on the number $h(g,n)$ of triangulations with n triangles of a surface of genus g could lead to the same result *for any large enough g*, instead of *for infinitely many values of g*.

To conclude the proof, there remains to prove the bound on the number of connected surfaces.

Lemma 5.1 *The number of (trivalent, unweighted)* connected *cross-metric surfaces with n vertices without boundary is, for n even large enough, at least $e^{Cn} n^{n/2}$ for some absolute constant C.*

Proof. Let G_n be the set of simple unlabelled trivalent graphs with n vertices. Let G'_n be the set of graphs in G_n that are connected. Let S'_n be the number of connected cross-metric surfaces with n vertices; we want a lower bound on $|S'_n|$. Below we implicitly assume n to be even, for otherwise these sets are empty.

We have $|S'_n| \geq |G'_n|$, because every graph in G'_n leads to a connected cross-metric surface, by cellularly embedding the graph arbitrarily, and these cross-metric surfaces are distinct, because they have distinct vertex-edge graphs.

Moreover, $|G'_n|/|G_n|$ tends to one as n goes to infinity, because the proportion of 3-connected graphs in the set of simple unlabelled trivalent graphs with n vertices goes to one as n goes to infinity [47, p. 338]. (Actually, except for this argument, our proof is heavily inspired by Guth et al. [30, Lemmas 1 and 3].)

The number of simple *labelled* trivalent graphs with $n = 2k$ vertices is, as n goes to infinity, equivalent to $\frac{(6k)!}{(3k)!}288^k e^2$ [55]. The expected number of automorphisms of these graphs tends to one as n goes to infinity [47, Corollary 3.8], which implies that $|G_n|$ is equivalent to $\frac{(6k)!}{(3k)!(2k)!}288^k e^2$, which is at least $e^{Cn}n^{n/2}$ for some absolute constant C. The previous paragraphs imply that $|S'_n|$ is asymptotically at least as large, as desired. $\square$

5.2 Proof of Theorem 5.3

We now show that the result just proved, Theorem 5.2, implies the polyhedral variant, Theorem 5.3:

Proof. As in the proof of Theorem 5.2, it suffices to prove the result for maps M that are cut graphs with minimum degree three, which have at most $4g-2$ vertices and $6g-3$ edges. Let G be the vertex-edge graph of a polyhedral triangulation on a surface S with genus g. Assume that C is a graph with combinatorial map M and of length at most $n^{7/6-\varepsilon}$ on that polyhedral surface. We prove that some cut graph with combinatorial map M has length $O(n^{7/6-\varepsilon})$ in the dual cross-metric surface (S, G^*). Since, by Theorem 5.2, the proportion of such surfaces is arbitrarily small, this implies the theorem.

Without loss of generality, we assume that C is piecewise-linear, and in general position with respect to G. We consider a tubular neighborhood of G (Figure 5.4(a)), obtained by first building a small *disk* around each vertex of G, and then building a rectangular *strip* containing each part of edge not covered by a disk. The disks are pairwise disjoint, the strips are pairwise disjoint, and each strip intersects only the disks corresponding to the incident vertices of the corresponding edge, along paths. We push C into the disks and strips as follows. A *piece* of C in a triangle T is a maximal connected part of C that lies in T; the *side number* of a piece is the number of sides of T it touches.

First, consider all the pieces with side number one. By an ambient isotopy, we can push these pieces across the side of the triangle they touch without increasing their length. So we can assume that no piece has side number one in any triangle.

Next, we deal with the pieces with side number two. By an ambient isotopy of the triangle fixing its boundary, we push all such pieces into the strips of the two sides of the triangle, putting the vertices in the disk touching the two strips (Figure 5.4(b–c)). Elementary geometry implies that this at most doubles the length of the pieces containing no vertex of C, and it increases the length of the pieces with a vertex by an additional term that is linear in the number of edges incident to the vertices of the piece. Since C has $O(g) = O(n)$ edges, the length of the modified cut graph is still $O(n^{7/6-\varepsilon})$.

Finally, there exists at most one piece with side number three lying in each triangle. We can push that piece as well to the three strips of the sides of the triangle, pushing all vertices of that piece to one of the disks, chosen arbitrarily (Figure 5.4(d)); this operation increases the length of C by an additional term

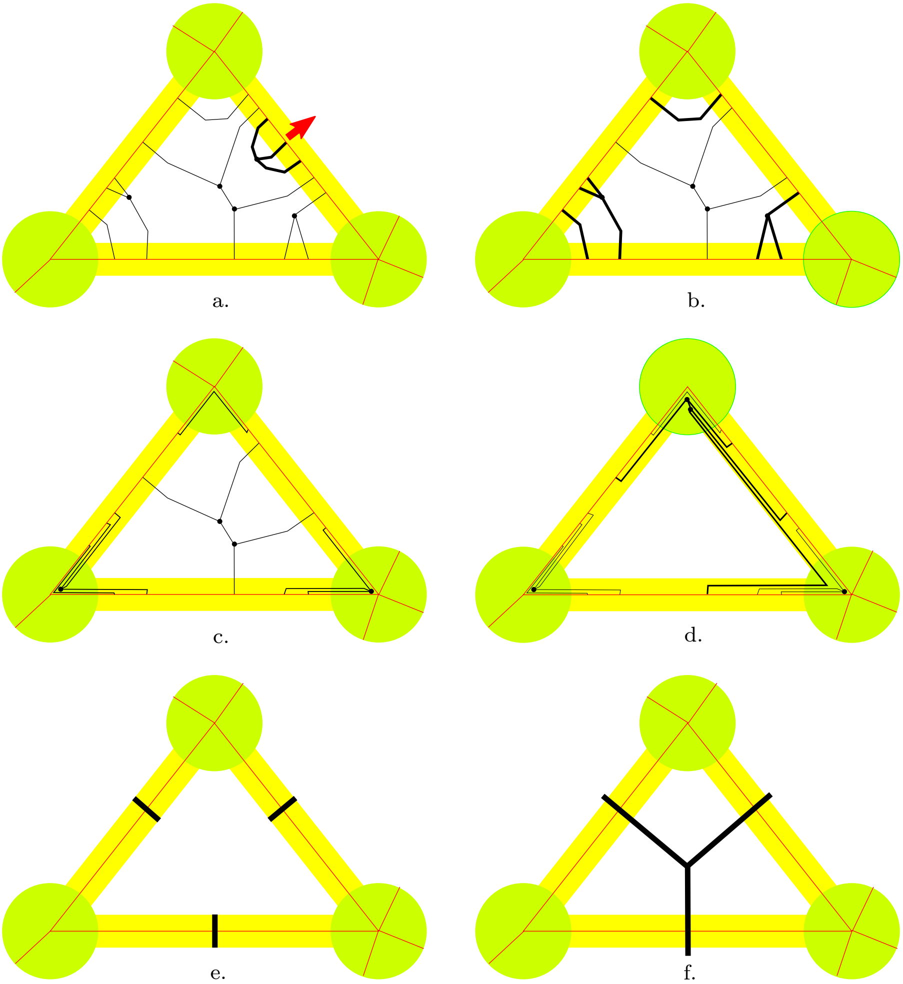

Figure 5.4 Illustration of the proof of Theorem 5.3. a.: The disks and strips inside one triangle of G, and the part of the cut graph C inside the triangle. b.: A piece with side number one is pushed across the side of the triangle. c.: The pieces with side number two are pushed to the disks and strips. d.: The piece with side number three is pushed to the disks and strips. e.: The paths P_s. f.: The cross-metric surface.

that is at most the number of edges of the piece. As before, this additional increase in length is $O(g) = O(n)$.

So, we have obtained an isotopic cut graph C', whose length is still $O(n^{7/6-\varepsilon})$, with the property that the vertices of C' lie in the disks and the edges of C' lie in the union of the disks and the strips. For each strip s, draw a shortest path P_s, with endpoints on its boundary, separating the two incident disks (Figure 5.4(e)). If a portion of C' inside s crosses P_s more than once, it forms a bigon with P_s; by flipping innermost bigons, without increasing the length of C', we can assume that each portion of C' inside s crosses P_s at most once.

Now we extend the paths P_s to form the graph G^* (Figure 5.4(f)). By the paragraph above, each crossing of a path P_s corresponds to a portion of a path of C' that crosses the strip containing P_s, and thus has length at least $1-\delta$, for $\delta > 0$ arbitrarily close to zero (the size of the disks and strips are chosen according to δ). Therefore, the length of C' on the cross-metric surface (S, G^*) is at most $(1-\delta)$ times that of the length of C' on the polyhedral triangulated surface, and thus $O(n^{7/6-\varepsilon})$. □

An interesting question would be to determine whether there exists an analog of Theorem 5.2 when we are not given the embedding of M, but only its abstract graph. More generally, let S and M be two graphs with n vertices that are cellularly embeddable on a surface of genus g; are there cellular embeddings of S and M on this surface such that the graphs cross only $O(n)$ times?

Acknowledgements We would like to thank Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh for pointing out and discussing with us their results on Voronoi diagrams of Riemannian surfaces [16] and manifolds. We are grateful to the anonymous referees for their careful reading of the manuscript, which allowed to correct several problems and to improve the presentation significantly, and for pointing out Kowalick's thesis [38].

References

1. Florent Balacheff, Hugo Parlier, and Stéphane Sabourau. Short loop decompositions of surfaces and the geometry of Jacobians. *Geom. Funct. Anal.*, 22(1):37–73, 2012.
2. Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. Delaunay triangulations of manifolds. arXiv:1311.0117, 2013.
3. Robert Brooks and Eran Makover. Random construction of Riemann surfaces. *Journal of Differential Geometry*, 68(1):121–157, 2004.
4. P. Buser and P. Sarnak. On the period matrix of a Riemann surface of large genus (with an appendix by J.H. Conway and N.J.A. Sloane). *Inventiones Mathematicae*, 117:27–56, 1994.
5. Peter Buser. *Geometry and spectra of compact Riemann surfaces*, volume 106 of *Progress in Mathematics*. Birkhäuser, 1992.
6. Sergio Cabello, Erin W. Chambers, and Jeff Erickson. Multiple-source shortest paths in embedded graphs. *SIAM Journal on Computing*, 42(4):1542–1571, 2013.
7. Sergio Cabello, Éric Colin de Verdière, and Francis Lazarus. Algorithms for the edge-width of an embedded graph. *Computational Geometry: Theory and Applications*, 45:215–224, 2012.
8. Sergio Cabello and Bojan Mohar. Finding shortest non-separating and non-contractible cycles for topologically embedded graphs. *Discrete & Computational Geometry*, 37(2):213–235, 2007.
9. Erin Chambers, Jeff Erickson, and Amir Nayyeri. Minimum cuts and shortest homologous cycles. In *Proceedings of the 25th Annual Symposium on Computational Geometry (SOCG)*, pages 377–385. ACM, 2009.
10. Erin W. Chambers, Éric Colin de Verdière, Jeff Erickson, Francis Lazarus, and Kim Whittlesey. Splitting (complicated) surfaces is hard. *Computational Geometry: Theory and Applications*, 41(1–2):94–110, 2008.
11. Erin W. Chambers, Jeff Erickson, and Amir Nayyeri. Homology flows, cohomology cuts. *SIAM Journal on Computing*, 41(6):1605–1634, 2012.
12. Éric Colin de Verdière. *Topological algorithms for graphs on surfaces*. PhD thesis, École normale supérieure, 2012. Habilitation thesis, available at `http://www.di.ens.fr/~colin/`.

13. Éric Colin de Verdière and Jeff Erickson. Tightening nonsimple paths and cycles on surfaces. *SIAM Journal on Computing*, 39(8):3784–3813, 2010.
14. Éric Colin de Verdière and Francis Lazarus. Optimal pants decompositions and shortest homotopic cycles on an orientable surface. *Journal of the ACM*, 54(4):Article 18, 2007.
15. Manfredo P. do Carmo. *Riemannian geometry*. Birkhäuser, 1992.
16. Ramsay Dyer, Hao Zhang, and Torsten Möller. Surface sampling and the intrinsic Voronoi diagram. *Computer Graphics Forum*, 27(5):1393–1402, 2008.
17. David Eppstein. Squarepants in a tree: sum of subtree clustering and hyperbolic pants decomposition. *ACM Transactions on Algorithms*, 5(3), 2009.
18. Jeff Erickson. Combinatorial optimization of cycles and bases. In Afra Zomorodian, editor, *Computational topology*, Proceedings of Symposia in Applied Mathematics. AMS, 2012.
19. Jeff Erickson, Kyle Fox, and Amir Nayyeri. Global minimum cuts in surface embedded graphs. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1309–1318, 2012.
20. Jeff Erickson and Sariel Har-Peled. Optimally cutting a surface into a disk. *Discrete & Computational Geometry*, 31(1):37–59, 2004.
21. Jeff Erickson and Amir Nayyeri. Computing replacement paths in surface-embedded graphs. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1347–1354, 2011.
22. Jeff Erickson and Kim Whittlesey. Greedy optimal homotopy and homology generators. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1038–1046, 2005.
23. Jeff Erickson and Pratik Worah. Computing the shortest essential cycle. *Discrete & Computational Geometry*, 44(4):912–930, 2010.
24. Alexander Gamburd and Eran Makover. On the genus of a random riemann surface. In *Complex Manifolds and Hyperbolic Geometry*, number 311 in Contemporary Mathematics, pages 133–140. AMS, 2002.
25. J. Geelen, T. Huynh, and R. B. Richter. Explicit bounds for graph minors. arxiv:1305.1451, 2013.
26. M. Gromov. Sign and geometric meaning of curvature. *Rend. Sem. Mat. Fis. Milano*, 61:9–123 (1994), 1991.
27. Mikhael Gromov. Filling Riemannian manifolds. *Journal of Differential Geometry*, 18:1–147, 1983.
28. Mikhael Gromov. Systoles and intersystolic inequalities. In *Actes de la table ronde de géométrie différentielle*, pages 291–362, 1992.
29. Igor Guskov and Zoë J. Wood. Topological noise removal. In *Proceedings of Graphics Interface*, pages 19–26, 2001.
30. Larry Guth, Hugo Parlier, and Robert Young. Pants decompositions of random surfaces. *Geometric and Functional Analysis*, 21:1069–1090, 2011.
31. Qing Han and Jia-Xing Hong. *Isometric embedding of Riemannian manifolds in Euclidean spaces*, volume 130 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2006.
32. Allen Hatcher. *Algebraic topology*. Cambridge University Press, 2002. Available at `http://www.math.cornell.edu/~hatcher/`.
33. Joan P. Hutchinson. On short noncontractible cycles in embedded graphs. *SIAM Journal on Discrete Mathematics*, 1(2):185–192, 1988.
34. Mikhail Katz. *Systolic geometry and topology*, volume 137 of *Mathematical Surveys and Monographs*. AMS, 2007. With an appendix by J. Solomon.
35. Mikhail G. Katz and Stéphane Sabourau. Entropy of systolically extremal surfaces and asymptotic bounds. *Ergodic Theory Dynam. Systems*, 25(4):1209–1220, 2005.
36. Michel A. Kervaire. A manifold which does not admit any differentiable structure. *Comment. Math. Helv.*, 34:257–270, 1960.
37. Wilhelm Klingenberg. *Riemannian geometry*. Philosophie und Wissenschaft. de Gruyter, 1995.
38. Ryan Kowalick. *Discrete systolic inequalities*. PhD thesis, Ohio State University, 2013.
39. Martin Kutz. Computing shortest non-trivial cycles on orientable surfaces of bounded genus in almost linear time. In *Proceedings of the 22nd Annual Symposium on Computational Geometry (SOCG)*, pages 430–438. ACM, 2006.

40. Francis Lazarus, Michel Pocchiola, Gert Vegter, and Anne Verroust. Computing a canonical polygonal schema of an orientable triangulated surface. In *Proceedings of the 17th Annual Symposium on Computational Geometry (SOCG)*, pages 80–89. ACM, 2001.
41. James R. Lee and Anastasios Sidiropoulos. Genus and the geometry of the cut graph. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 193–201, 2010.
42. Gregory Leibon. *Random Delaunay Triangulations, the Thurston-Andreev Theorem and Metric Uniformization*. PhD thesis, University of California at San Diego, 1999. Available on arXiv:math/0011016v1.
43. Bruno Lévy and Jean-Laurent Mallet. Non-distorted texture mapping for sheared triangulated meshes. In *Proceedings of the 25th Annual Conference on Computer Graphics (SIGGRAPH)*, pages 343–352, 1998.
44. Xin Li, Xianfeng Gu, and Hong Qin. Surface mapping using consistent pants decomposition. *IEEE Transactions on Visualization and Computer Graphics*, 15:558–571, 2009.
45. Eran Makover and Jeffrey McGowan. The length of closed geodesics on random Riemann surfaces. *Geometriae Dedicata*, 151:207–220, 2011.
46. Jiří Matoušek, Eric Sedgwick, Martin Tancer, and Uli Wagner. Untangling two systems of noncrossing curves. In Stephen Wismath and Alexander Wolff, editors, *Graph Drawing*, volume 8242 of *Lecture Notes in Computer Science*, pages 472–483. Springer International Publishing, 2013.
47. Brendan D. McKay and Nicholas C. Wormald. Automorphisms of random graphs with specified vertices. *Combinatorica*, 4(4):325–338, 1984.
48. Dan Piponi and George Borshukov. Seamless texture mapping of subdivision surfaces by model pelting and texture blending. In *Proceedings of the 27th Annual Conference on Computer Graphics (SIGGRAPH)*, pages 471–478, 2000.
49. Nicholas Pippenger and Kristin Schleich. Topological characteristics of random triangulated surfaces. *Random Structures & Algorithms*, 28(3):247–288, 2006.
50. Sheung-Hung Poon and Shripad Thite. Pants decomposition of the punctured plane. arXiv:cs.CG/0602080. Preliminary version in *Abstracts of the European Workshop on Computational Geometry*, 2006., 2006.
51. Teresa M. Przytycka and Józef H. Przytycki. On a lower bound for short noncontractible cycles in embedded graphs. *SIAM Journal on Discrete Mathematics*, 3(2):281–293, 1990.
52. Teresa M. Przytycka and Józef H. Przytycki. Surface triangulations without short noncontractible cycles. In Neil Robertson and Paul Seymour, editors, *Graph structure theory*, number 147 in Contemporary Mathematics, pages 303–340. AMS, 1993.
53. Teresa M. Przytycka and Józef H. Przytycki. A simple construction of high representativity triangulations. *Discrete Mathematics*, 173:209–228, 1997.
54. Pao Ming Pu. Some inequalities in certain nonorientable riemannian manifolds. *Pacific Journal of Mathematics*, 2:55–71, 1952.
55. Ronald C. Read. *Some enumeration problems in graph theory*. PhD thesis, University of London, 1958.
56. Neil Robertson and Paul D. Seymour. Graph minors. VII. Disjoint paths on a surface. *Journal of Combinatorial Theory, Series B*, 45:212–254, 1988.
57. Stéphane Sabourau. Asymptotic bounds for separating systoles on surfaces. *Commentarii Mathematici Helvetici*, 83:35–54, 2008.
58. Michael Spivak. *A comprehensive introduction to differential geometry*, volume II. Publish or Perish Press, 1999.
59. John Stillwell. *Classical topology and combinatorial group theory*. Springer-Verlag, New York, 1980.
60. Zoë Wood, Hugues Hoppe, Mathieu Desbrun, and Peter Schröder. Removing excess topology from isosurfaces. *ACM Transactions on Graphics*, 23(2):190–208, 2004.

A Discrete Systolic Inequalities in Higher Dimensions

In this appendix, we show that the proofs from Section 3 extend almost verbatim to higher dimensions. In the following discussion (M, T) will be a triangulated d-manifold.[1] We will denote by $f_d(T)$ the number of d-dimensional simplices of T, and by $f_0(T)$ the number of vertices. The main difference with the two-dimensional case is that while for surfaces, discrete systolic inequalities in terms of f_0 and in terms of f_d are easily seen to be equivalent (by Euler's formula and double counting), in higher dimensions the situation is more complicated.

We consider the supremal values of the functionals $\frac{\mathrm{sys}^d}{f_d}$ and $\frac{\mathrm{sys}^d}{f_0}$, where sys denotes the length of a shortest closed curve in the 1-skeleton of (M, T) that is non-contractible on the manifold M. In particular we focus on when these quantities are bounded from above. As we surveyed in the introduction, the two-dimensional case of this problem has been studied by topological graph theorists and computational topologists; however, as far as we know, it has never been considered in dimension higher than two in the past. We report the results and open problems that we can derive by generalizing our techniques for surfaces.

A.1 From Continuous to Discrete Systolic Inequalities

To infer discrete systolic inequalities from the Riemannian ones, the obvious approach is, as before, to start with a triangulated manifold (M, T) and to endow M with a metric m_T by deciding that each simplex of T is a regular Euclidean simplex of volume one. (Since the simplices are regular, we glue them by facewise isometries.) Hence, length and volume are naturally defined via the restriction to each Euclidean simplex. Following Gromov [27], we will call such a metric a *piecewise Riemannian metric*. Unlike the 2-dimensional case, however, foundational work of Kervaire [36] shows that in higher dimensions such a triangulated manifold is not always smoothable. (We will show how to circumvent this difficulty below.)

Theorem A.1 *There exists a constant C_d, such that for every triangulated compact manifold (M, T) without boundary of dimension d, there exists a piecewise Riemannian metric m on M with volume $f_d(T)$ such that for every closed curve γ in M, there exists a homotopic closed curve γ' on the 1-skeleton G of T with*

$$|\gamma'|_G \leq C_d |\gamma|_m.$$

The proof works inductively, pushing curves from the i-dimensional skeleton to the $(i-1)$-dimensional one. We start with the following lemma.

[1] E.g., (M, T) is a simplicial complex whose underlying space is a d-manifold. However, we can allow more general triangulations obtained from gluing d-simplices, in which, after gluing, some faces (e.g., vertices) of the same d-simplex are identified.

Lemma A.1 *Let Δ be an i-dimensional regular simplex, endowed with the Euclidean metric. There exists an absolute constant C'_i such that, for each arc γ properly embedded in Δ with endpoints in $\partial\Delta$, there exists an arc γ' embedded on $\partial\Delta$, with the same endpoints as γ, such that $|\gamma'| \leq C'_i|\gamma|$.*

Proof of Lemma A.1. Since the statement of the lemma is invariant by scaling all the distances, we can assume that Δ is the regular i-simplex whose circumscribing sphere bounds the unit ball B in $\mathbb{R}^i$. Let us first consider the bijection φ that maps Δ to B by radial projection (such that the restriction of φ to any ray from the origin is a linear function). It is not hard to see that there is a constant C''_i such that, for any arc γ in Δ, we have $|\gamma|/C''_i \leq |\varphi(\gamma)| \leq C''_i|\gamma|$ (one can compute the optimal C''_i by writing the map in hyperspherical coordinates and computing the differential).

Therefore it suffices to prove the lemma for the unit ball B instead of the regular simplex Δ. Let β be an arc embedded in B. Let β' be a shortest geodesic arc on ∂B with the same endpoints as β. Then we have $|\beta'| \leq \frac{\pi}{2}|\beta|$, which proves the result. □

Proof of Theorem A.1. As we mentioned before, we endow M with the piecewise Riemannian metric obtained by endowing each simplex of dimension d with the geometry of the regular Euclidean simplex of volume 1. Then, using Lemma A.1, for every arc A of γ in every d-simplex, we push A to the $(d-1)$-skeleton of (M,T), and we repeat this procedure inductively until γ is embedded in the 1-skeleton. In the end, the length of γ' has increased by at most a multiplicative factor that depends only on d. □

The Riemannian systolic inequality in higher dimensions is now stated in the following theorem.

Theorem A.2 (Gromov [27]) *For every d, there is a constant C_d such that, for any Riemannian metric m on any essential compact d-manifold M without boundary, there exists a non-contractible closed curve of length at most $C_d \operatorname{vol}(m)^{1/d}$.*

For a definition of essential manifold, see [27]. The prime examples of essential manifolds are the so-called aspherical manifolds, which are the manifolds whose universal cover is contractible. These include for example the d-dimensional torus for every d, or manifolds that accept a hyperbolic metric and, more generally, manifolds that are locally CAT(0). In particular, all surfaces except the 2-sphere and the projective plane are aspherical. On the other hand, real projective spaces and lens spaces are examples of essential manifolds that are non aspherical.

Theorem A.2 also holds for piecewise Riemannian metrics. Indeed, its proof revolves around two key inequalities: the filling radius-volume inequality and a systole-filling radius inequality. The former relies on a coarea formula which holds for piecewise Riemannian metrics (see [27, Lemma 4.2b]), and the proof of the latter uses no smoothness property either, see [27, p. 9 and 10]. As a corollary of this refinement to piecewise Riemannian metrics and of our

Theorem A.1, we obtain the following result relating the length of systoles and the number of facets.

Corollary A.1 *Let (M,T) be a triangulated essential compact d-manifold without boundary. Then, for some constant c_d depending only on d, some non-contractible closed curve in the 1-skeleton of T has length at most $c_d f_d(T)^{1/d}$.*

A.2 From Discrete to Continuous Systolic Inequalities

We now turn our attention to the other direction, namely, transforming a discrete systolic inequality into a continuous one.

Theorem A.3 *Let M be a compact Riemannian manifold of dimension d and volume V without boundary, and let $\delta > 0$. For infinitely many values of f_0, there exists a triangulation (M,T) of M with f_0 vertices, such that every closed curve γ in the 1-skeleton G of M satisfies*

$$|\gamma|_m \leq (1+\delta)\frac{10}{\sqrt{\pi}}\Gamma(d/2+1)^{1/d}\left(\frac{V}{f_0}\right)^{1/d}|\gamma|_G.$$

(Here, Γ is the usual Gamma function.) The proof follows the same idea as the proof of Theorem 3.2. We start with the centers of a maximal set of disjoint balls of radius $\varepsilon/2$ in M and want to compute the Delaunay triangulation associated to it, with the hope that if ε is small enough, we will obtain a triangulation of M. However, Delaunay complexes behave differently in higher dimensions, and this hope turns out to be false in many cases. We rely instead on a recent work reported by Boissonnat, Dyer and Ghosh [2] who devised the correct perturbation scheme to triangulate a manifold using a Delaunay complex. We will use the following theorem.

Theorem A.4 *Let M be a compact Riemannian manifold. For a small enough ε, there exists a point set $P \subseteq M$ such that*

(i) For every $x \in M$, there exists $p \in P$ such that $|x-p|_m \leq \varepsilon$.
(ii) For every pair $p \neq p' \in P$, $|p-p'|_m \geq 2\varepsilon/5$.
(iii) The Delaunay complex of P is a triangulation of M.

For completeness, we sketch how to infer this theorem from the paper [2].

Proof of Theorem A.4. We say that a set of points $P \subseteq M$ is ε-dense if $d(x,P) < \varepsilon$ for $x \in M$, $\mu_0\varepsilon$-separated if $d(p,q) \geq \mu_0\varepsilon$ for all distinct $p,q \in P$, and is a (μ_0,ε)-net if it is ε-dense and $\mu_0\varepsilon$-separated.

Taking $\mu'_0 = 1$ and ε' small enough, we start with a (μ'_0,ε')-net in M, which can be obtained for example by taking the centers of a maximal set of disjoint balls of radius $\varepsilon'/2$. Now, the extended algorithm of Boissonnat et al. [2] outputs a (μ_0,ε)-net with $\varepsilon \leq 5\varepsilon'/4$ and $\mu_0 \geq 2\mu'_0/5$, which will be our point set P. These conditions correspond to items (i) and (ii) in our theorem. For ε small enough, all the hypotheses of their main theorem are fulfilled, and therefore we obtain that the Delaunay complex of the (μ_0,ε)-net is a triangulation of M, which is our item (iii). □

The proof of Theorem A.3 now follows the same lines as in the 2-dimensional case.

Proof of Theorem A.3. Let $\varepsilon > 0$ be a constant. Following Theorem A.4, if ε is small enough, there exists a point set P whose Delaunay complex triangulates M. Let G be the 1-skeleton of this complex, and γ be a closed curve embedded in G.

By property (i), neighboring points in G are at distance no more than 2ε, therefore we have $|\gamma|_m \leq 2\varepsilon|\gamma|_G$. There just remains to estimate the value of ε, which we do by estimating the number of balls. By compactness, the scalar curvature of M is bounded from above by some constant K. Now, if ε is small enough, for any $p \in P$ we have:

$$\mathrm{vol}(B(p, \varepsilon/5)) \geq \frac{\varepsilon^d}{5^d}\left(1 - \frac{\varepsilon^2}{6d}K + o(\varepsilon^4)\right)\mathrm{vol}(B^d),$$

where B^d is the unit Euclidean ball of dimension d. This follows from standard estimates on the volume of a ball in a Riemannian manifold, see for example Gromov [26, p. 89]. We recall that $\mathrm{vol}(B^d) = \frac{\pi^{d/2}}{\Gamma(d/2+1)}$.

By property (ii), the balls $B(p, \varepsilon/5)$ are disjoint, therefore their number f_0 is at most $\frac{\Gamma(d/2+1)5^d V}{\pi^{d/2}\varepsilon^d(1-\varepsilon)}$ if ε is small enough. Finally, putting together our estimates, we obtain that

$$|\gamma|_m \leq (1+\delta)\frac{10}{\sqrt{\pi}}\left(\frac{\Gamma(d/2+1)}{f_0}V\right)^{1/d}|\gamma|_G. \quad \square$$

However, this theorem leads to no immediate corollaries, since unlike the two-dimensional case, we do not know of any discrete systolic inequalities involving f_0 in dimensions larger than two. This leads to the following question.

Question A.1 Are there manifolds M of dimension $d \geq 3$ for which there exists a constant c_M such that, for every triangulation (M, T), there is a non-contractible closed curve in the 1-skeleton of T of length at most $c_M f_0(T)^{1/d}$?

Notice that a positive answer to this question for essential compact manifolds without boundary would yield a new proof of Gromov's systolic inequality.

Remark: In his thesis [38], Kowalick states a theorem that is closely related to our Theorem A.3, and thus to this question. Essentially, his result is ours substituting f_0 with f_d. Precisely, he shows that if for a manifold M, there exists $c'_M > 0$ such that, for every triangulation (M, T), there is a non-contractible closed curve in the 1-skeleton of T of length at most $c'_M f_d(T)^{1/d}$, then there exists a constant s_M such that the systole of every Riemannian metric on M is bounded above by $s_M\,\mathrm{vol}(M)^{1/d}$. This statement can be derived from our proof without much extra difficulty. It is enough to show that in the triangulations constructed in the proof of Theorem A.3, for large enough

$f_0(T)$, the number $f_d(T)$ is bounded above by $f_0(T)$ up to a multiplicative constant that depends only on the dimension. This follows again from a packing argument and from the bounds on volume growth of small balls. For ε small enough with respect to the strong convexity radius and the minimal sectional curvature of the manifold, the quotient $\operatorname{vol}(B(x, 2\varepsilon))/\operatorname{vol}(B(x, \varepsilon)$ is bounded from above by an absolute constant k_d depending only on the dimension. Since two points in the same facet of our Delaunay complex are at distance at most 2ε, we have $f_d(T) \leq \frac{{k_d}^d}{d+1} f_0(T)$.

B Lengths of Genus Zero Decompositions

A genus zero decomposition of a surface is a family of disjoint simple closed curves that cut the surface into a (connected) genus zero surface with boundary. Every genus zero decomposition (of a surface with genus at least two) can be extended to a pants decomposition. In this section, we prove the following strengthening of Theorem 5.1:

Theorem B.1 *For any $\varepsilon > 0$, the following holds with probability tending to one as n tends to ∞: A random (trivalent, unweighted) cross-metric surface with n vertices has no genus zero decomposition of length at most $n^{7/6-\varepsilon}$.*

The argument is very similar to the one by Poon and Thite [50, Sect. 2]. As we shall see, this theorem is an immediate consequence of the following proposition:

Proposition B.1 *Let (S, G^*) be a (trivalent, unweighted) cross-metric surface with genus zero and $b \geq 3$ boundary components. Then there exists some pants decomposition Γ of S such that each edge of G^* has $O(\log b)$ crossings with each edge of Γ.*

Proof. Define the *multiplicity* of a set of curves on (S, G^*) to be the maximum number of crossings between an edge of G^* and the set of curves.

Let T be a spanning tree of the boundary components of (S, G^*), that is, a tree of multiplicity one in (S, G^*) so that each boundary component of S is intersected by exactly one leaf of the tree, (Figure B.1(a)). Draw a path p following the tree T, touching it only at the leaves (Figure B.1(b–c)); such a path p has multiplicity two, and touches each boundary component exactly once. Let $B_1, B_2, \ldots, B_b$ be the boundary components in order along p (oriented arbitrarily).

Now, we build the pants decomposition (Figure B.1(d)). First we group the boundary components by pairs, $\{B_1, B_2\}$, $\{B_3, B_4\}$, and so on. Then we cut S into a collection of $\lfloor b/2 \rfloor$ pairs of pants and a genus zero surface with $\lceil b/2 \rceil$ boundary components, and we reiterate the process on the latter surface. After $O(\log b)$ iterations, the remaining surface has at most three boundary components, so we have built a pants decomposition Γ.

We claim that Γ has multiplicity $O(\log b)$. Indeed, each closed curve of Γ is made of (1) pieces that go around a boundary component, and (2) pieces

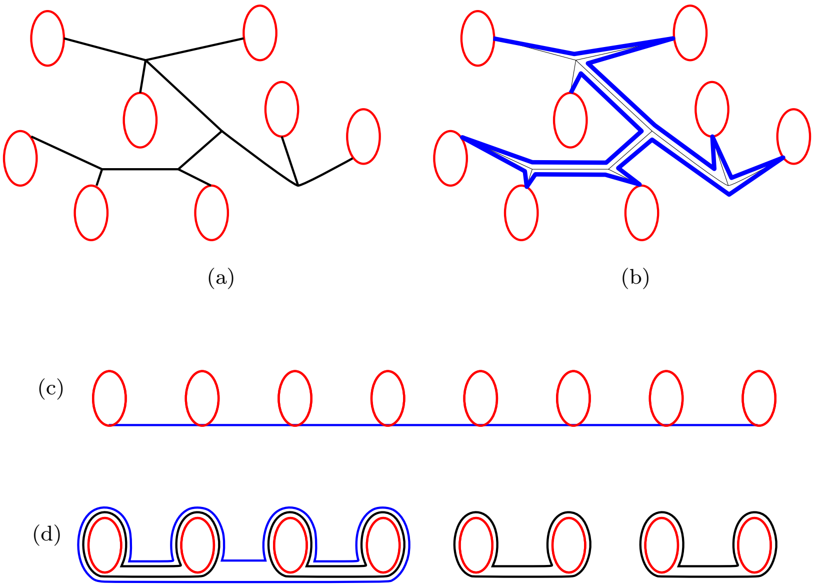

Figure B.1 The construction of the pants decomposition in Proposition B.1. (a) The tree T. (b) The path p. (c) An isomorphic drawing of p. (d) The pants decomposition.

that follow a subpath of p. The pieces of type (1) have overall multiplicity $O(\log b)$, because $O(\log b)$ pieces go around a given boundary component and each edge of G^* is incident to at most two boundary components. The pieces of type (2) have overall multiplicity $O(\log b)$, since $O(\log b)$ pieces run along a given subpath of p, and because p has multiplicity two in (S, G^*). The result follows. □

Proof of Theorem B.1. Consider a random cross-metric surface (S, G^*) with n vertices; let g be its genus.

- It may be that (S, G^*) has genus zero or one; but this happens with probability arbitrarily close to zero, provided n is large enough (this follows by combining Lemma 5.1 with Guth et al. [30, Lemma 4]);
- otherwise, if (S, G^*) admits a genus zero decomposition Γ' of length at most $n^{7/6-\varepsilon}$, we cut (S, G^*) along Γ', obtaining a cross-metric surface with genus zero with $2g \geq 3$ boundary components and $O(n^{7/6-\varepsilon})$ edges. Proposition B.1 implies that this new cross-metric surface has a pants decomposition Γ with length $O(n^{7/6-\varepsilon} \log g) = O(n^{7/6-\varepsilon} \log n)$. The union of Γ and Γ' is a pants decomposition of (S, G^*) of length at most $O(n^{7/6-\varepsilon} \log n + n^{7/6-\varepsilon}) = O(n^{7/6-\varepsilon'})$ for some $\varepsilon' < \varepsilon$ if n is large enough. By Theorem 5.1 above, we conclude that this happens with arbitrarily small probability as $n \to \infty$. □