UNIVERSIDAD DE BUENOS AIRES
Facultad de Ciencias Exactas y Naturales
Departamento de Matemática

Quantitative and geometric invariants for the complexity of spaces and groups

Tesis presentada para optar al título de Doctor de la Universidad de Buenos Aires en el área Ciencias Matemáticas

Eugenio Borghini

Advisor: Elías Gabriel Minian.
Study advisor: Jonathan Ariel Barmak.

Buenos Aires
June 5th 2020

Quantitative and geometric invariants for the complexity of spaces and groups

Abstract

This Thesis is devoted to the study of numerical invariants that measure the complexity of the topology and the homotopy type of a space. A natural way to quantify the complexity of a space is by computing the minimum number of *simple* pieces needed to assemble it, since intuitively, the spaces that exhibit highly non-trivial topology should be hard to build. This is the essential idea behind the definition of some classic invariants that detect non-trivial topology, such as the Lusternik-Schnirelmann (L-S) category and relatives. Minimal triangulations of spaces have also been studied for this reason. In this direction, we show that the minimal triangulations of closed surfaces optimize the number of vertices in triangulations of spaces of their homotopy type, with the only exception of the torus with two handles. This result settles a problem posed by Karoubi and Weibel. We also prove that minimal triangulations of a closed surface S minimize the number of 2-simplices among those simplicial complexes with fundamental group isomorphic to $\pi_1(S)$. This partially answers a question raised by Babenko, Balacheff and Bulteau. Despite the similarity with the first result, the motivation for this problem comes from the close connection between the *simplicial complexity* and the *systolic area* of groups. The systolic area of a group G is a notion due to Gromov that, roughly, consists of the minimum amount of area needed to build a complex K with fundamental group G out of riemannian simplices, normalized by the condition that the length of the shortest non-trivial loop (i.e. the systole) in K equals 1. Meanwhile, the simplicial complexity of a group G is defined as the minimum number of 2-simplices in a simplicial complex K with $\pi_1(K) = G$. This notion, that may be regarded as a discrete version of the area for groups, was recently introduced by Babenko, Balacheff and Bulteau. The authors showed that the simplicial complexity approximates (asymptotically) the systolic area of groups. Thus, it is natural to attack problems about the systolic area in this new combinatorial context. For instance, a basic open problem about the systolic geometry of groups, which goes back to Gromov, is whether the systolic area $\sigma(G * \mathbb{Z})$ of a free product $G * \mathbb{Z}$ equals $\sigma(G)$. The real question here is whether the canonical (wedge sum) model for the groups of the form $G * \mathbb{Z}$ is the most effective one. This motivates the study of the analogue problem for the simplicial complexity of a free product of the form $G * \mathbb{Z}$, which is also open. In this direction, we prove that the simplicial complexity of fundamental groups of surfaces is stable under free product with free groups. We also describe a construction, based on Stallings' topological proof of Grushko theorem, that we think might be a first step towards the computation of both the systolic area and the simplicial complexity of free product of groups.

Lastly, we provide new estimates for the systolic area of some specific groups through the study of the systolic geometry of polyhedra of dimension 2. More concretely, we find a (possibly singular) surface of relatively large area embedded in piecewise riemannian polyhedra of dimension 2 that satisfy a certain cohomological condition. Using this, we prove a new systolic inequality that extends to piecewise riemannian complexes of dimension 2 an inequality of Guth for riemannian manifolds. As a consequence, we conclude that for a large class of groups the systolic area is greater than or equal to $\frac{1}{2}$, which improves the best previously known general lower bound and brings it closer to the conjectured $\frac{2}{\pi}$.

Part of the results of the Thesis appear in the works [16, 15, 17, 18], while some others will

be the subject of an article in preparation.

Keywords. Systolic geometry, Lusternik-Schnirelmann category, minimal triangulations, simplicial complexity of groups.

Introduction

One of the objectives in the field of quantitative topology and geometry is to turn qualitative and existence results in algebraic topology into quantitative ones by studying some measures (often continuous) of the size and complexity of the maps and spaces involved. This process is perhaps best illustrated by some concrete examples.

- It is a basic fact from algebraic topology that a null-homotopic loop $\gamma : S^1 \to X$ is filled by a "disk" $\bar{\gamma} : D \to X$. The space X often comes with some natural inherent geometry (like a distance function or a riemannian metric) which allows to measure how difficult is to realize the trivialization of the loop by looking at the "size" of the filling disks D, for instance, the Hausdorff measure or parameterized area. This leads to the well-known notion of filling function of riemannian manifolds. That the filling function of a riemannian manifold M presents the same qualitative behavior as the Dehn filling of its fundamental group $\pi_1(M)$ is a foundational result in geometric group theory (see [45, 22]).

- One of the most basic invariants of a group G (say, finitely generated) is its cardinality. By fixing a finite (symmetric) generating set $\mathcal{A}$ of G, it is possible to refine the study of the cardinality of G by counting for each natural n the number of different elements of G expressible as words of length at most n in the alphabet $\mathcal{A}$. This is the growth function of the group G with respect to the generating set $\mathcal{A}$, which is a central object of study in geometric group theory. Remarkably, the growth function of a group contains algebraic information about the group, as it is shown by Gromov's characterization of the finitely generated groups with polynomial growth function as the groups having nilpotent subgroups of finite index [41].

The unifying theme of this Thesis is the study of quantitative invariants that measure in different senses the complexity of the topology or the homotopy type of a space, as well as the topological models at which the optimal values for these invariants are attained. As opposed to the classical approach of algebraic topology, which focuses on developing algebraic tools and invariants to study a space, such as homology and homotopy groups, the objective here is to obtain a single non-negative number that accounts for the inherent complexity of a space, or at least, some aspect of it. A natural way to accomplish this task is by first identifying the *simplest* possible spaces (for instance, disks, balls, contractible sets, etc.) and then setting the complexity of a space X to be the minimum number of simple pieces required to assemble X. The classical Lusternik-Schnirelmann (L-S) category of a space, introduced in [67], is a prototypical invariant tailored according to this pattern. Since from the point of view of homotopy theory the simplest sets are the contractible ones, here for a space X the rôle of simple pieces is played by the open sets contractible to a point in X. The L-S category of X is then defined as the minimum number of open sets contractible in the space that cover X. This turns out to be a homotopy invariant of a space which has been applied in a wide range of fields, from estimating the number of critical points of a smooth real function on a manifold to nilpotency aspects of homotopy types (see [29] for more details). Over the last years, there has been a renewed interest in the area of L-S category and related invariants. Besides the introduction of new L-S category type invariants (among them, Farber's *topological complexity* [31] is probably the one that has attracted the most attention), several works have concentrated on giving the *right* analogues of L-S category and related invariants in the context of discrete objects, such as simplicial complexes, polyhedra or finite topological spaces (see for example [1, 34, 33]). The common principle underlying these works is the replacement of contractible sets as the simple building pieces of spaces by other notions of homotopically simple sets that exploit the more rigid nature of the category under study (for example, collapsible sets in the classical Whitehead's simple homotopy theory in [1] or strong collapsible sets in the relatively new strong homotopy types of Barmak and Minian [9] in [34, 33]). Another, perhaps rougher way to measure complexity when one deals with spaces admitting a triangulation is through minimal triangulations. The problem of determining small triangulations of PL manifolds was investigated by a number of authors (see for example [20, 21, 68, 69, 58, 76] and references therein) and it was related explicitly to the complexity of the topology of 3-manifolds by Matveev [70] (more precisely, his notion of *combinatorial complexity* uses minimal pseudotriangulations instead of triangulations).

Interestingly, the point of view of minimal triangulations is relied to L-S category type invariants by Karoubi and Weibel through the introduction of the *covering type* of a space [61]. Recall that an open cover of a space is a *good cover* if every nonempty intersection of its members is contractible. The covering type of a space X is defined as the minimum size of a good cover of a space Y of the homotopy type of X. This invariant is closely related to minimal triangulations of homotopy types by the following argument. By the Nerve Theorem, a topological space X admitting a good cover $\mathcal{U}$ of cardinality n is homotopy equivalent to the nerve of $\mathcal{N}(\mathcal{U})$ of the cover, which is a simplicial complex on n vertices. Hence, for those spaces X of the homotopy type of a finite simplicial complex (for example, all compact CW-complexes fall in this class) the covering type is the minimum number of vertices in a simplicial complex K homotopy equivalent to X. In the cited work [61], Karoubi and Weibel computed the covering type of 1-dimensional simplicial complexes and posed the problem of finding the exact value of the covering type of closed surfaces. Using the reformulation via the Nerve Theorem, this is equivalent to the following natural question:

What is the minimum number of vertices in a simplicial complex homotopy equivalent to a given closed surface?

We provide a complete solution to this problem by proving that the expected result *almost* holds.

Theorem 2.4.5. *Let S be a closed surface, either orientable or non-orientable. Then, the covering type of S coincides with the minimum number of vertices in an optimal triangulation of S, with the only exception of the torus with two handles in which case both numbers differ by 1.*

We remark that an explicit formula for the minimum number of vertices in a triangulation of any closed surface is well known. It was obtained by Ringel [76] in the non-orientable case and Jungerman and Ringel [58] in the orientable case.

A similar (at least, at first sight) question was recently raised by Babenko, Balacheff and

Bulteau ([4]):
What is the minimum number of 2-simplices in a simplicial complex with the fundamental group of a given closed surface S?

This second problem is related to a combinatorial invariant of groups introduced by the authors in [4] called *simplicial complexity.* Given a (finitely presentable) group G, let us say that a connected simplicial complex K is a *triangulation* of G or that K *triangulates* G if $\pi_1(K) = G$. The simplicial complexity of the group G is defined as the minimum number of 2-simplices in a triangulation of G. As opposed to what happens for the covering type, we show that there are no exceptional cases for the simplicial complexity of fundamental groups of surfaces. In what follows, we will refer by *surface groups* to this class of groups, as usual.

Theorem 4.4.7. *The simplicial complexity $\kappa(\pi_1(S))$ of the fundamental group of a non-simply connected closed surface S coincides with the minimum number of 2-simplices in an optimal triangulation of S. Furthermore, the simplicial complexity of surface groups is stable under free product with free groups, that is, $\kappa(\pi_1(S) * \mathbb{Z}) = \kappa(\pi_1(S))$.*

Previously, the exact value of the simplicial complexity was known only for a few specific groups ($\mathbb{Z}_2$, $\mathbb{Z}_3$, $\mathbb{Z} \oplus \mathbb{Z}$ and the fundamental group of the Klein bottle), obtained mainly by a case-by-case analysis in [25]. With regard to the second statement, it was conjectured by the authors in [4] that the simplicial complexity is stable under free product of groups, i.e. that $\kappa(G) = \kappa(G * \mathbb{Z})$ for finitely presented groups G. The conjecture is motivated by the fact that, from a triangulation of a group G, it is always possible to obtain a triangulation of $G * \mathbb{Z}$ with no additional 2-simplices by forming a wedge sum with (a triangulation of) S^1. Beyond the notation, the question is whether this natural construction is the optimal way to triangulate a group of the form $G * \mathbb{Z}$. The verification of the conjecture in the case of surface groups constitutes the first (even if partial) evidence in favor of an affirmative answer to this question.

Both Theorems 2.4.5 and 4.4.7 indicate that in a quantifiable way the closed surfaces are the most effective topological models within their homotopy type. To prove these results, we employed the same core techniques, which we proceed to describe. The first crucial step is the identification of a relevant property in the cohomology ring of spaces and groups, which we call *property (A).*

Definition 2.3.2. *Let X be a topological space or a group. We say that the cohomology ring $H^*(X; \mathbb{Z}_2)$ satisfies property (A) if for every non-trivial $\alpha \in H^1(X, \mathbb{Z}_2)$ there exists $\beta \in H^1(X, \mathbb{Z}_2)$ such that $\alpha \cup \beta$ is non-trivial in $H^2(X, \mathbb{Z}_2)$.*

Notice that if X is a surface (either orientable or non-orientable), its cohomology ring satisfies property (A) by Poincaré Duality. Intuitively, if the cohomology ring of a simplicial complex K satisfies property (A), the 2-skeleton of K should be dense, in the sense that it should contain several 2-simplices (for instance, K cannot be homotopy equivalent to a wedge sum of the form $X \vee S^1$). For complexes of dimension 2, property (A) implies an explicit quantitative lower bound for the number of vertices and 2-simplices through a straightforward Euler characteristic computation. In addition, except for a few exceptional cases this estimate is optimal if the involved 2-dimensional complex has $\mathbb{Z}_2$-(co)homology isomorphic to that of a closed surface S, which means that it coincides with the number of vertices or 2-simplices in a minimal triangulation of S.

The second key step is a homological simplification method controlled by the property (A). The precise meaning of such a process is best illustrated while estimating the covering type of surfaces. Suppose given a simplicial complex K of the homotopy type of a closed surface S. If the dimension of K is greater than 2, its 2-skeleton $K^{(2)}$ has the same (co)homology as S up to degree 1 but in general the induced map $H_2(K^{(2)}, \mathbb{Z}_2) \to H_2(S, \mathbb{Z}_2)$ is only surjective. The homological simplification method allows to find a subcomplex $Z \leq K^{(2)}$ with $\mathbb{Z}_2$-(co)homology isomorphic to that of S and hence, satisfying property (A). Now the desired estimate for the number of vertices and 2-simplices follows from the mentioned Euler characteristic computation.

We have omitted so far to discuss the motivation behind the definition of the simplicial complexity, which brings us to another of the central subjects of this Thesis. Namely, Babenko, Balacheff and Bulteau introduced the simplicial complexity of a group G to approximate an invariant called the *systolic area* of G. To give a meaningful account of this notion, we need to briefly introduce some fundamental concepts and results in systolic geometry and topology.

The systole $\operatorname{sys}(X)$ of a metric space X is the length of a shortest non-contractible loop in the space. If positive, it is a metric measure of the non-triviality of the topology of X. Historically, the first result involving this invariant dates from 1949 (although not under that name, which was coined in 1980), when Loewner proved the following kind of inverse isoperimetric inequality for riemannian 2-tori.

Theorem 3.1.1. *Let $(\mathbb{T}^2, g)$ be a riemannian 2-torus. Then, $\operatorname{sys}(\mathbb{T}^2, g)^2 \leq \frac{2}{\sqrt{3}} \operatorname{Area}(\mathbb{T}^2, g)$. Moreover, the constant $\frac{2}{\sqrt{3}}$ is optimal and it is realized by a flat metric.*

This was later generalized by Accola [2] and Blatter [14], who proved independently the existence of a non-optimal constant $C = C(\gamma)$ such that for any orientable riemannian surface (S, g) of genus γ the inequality

$$\operatorname{sys}(S, g)^2 \leq C(\gamma) \operatorname{Area}(S, g)$$

holds. The subject was popularized by Berger in the seventies (see [13, 12]), but it was not until Gromov's *Filling Riemannian manifolds* [44] that it reached full maturity. One of the main results in that article is a far-reaching generalization of the inequalities by Loewner, Accola and Blatter.

Theorem 3.1.2. *Let (M, g) be an aspherical riemannian manifold of dimension n. Then, there exists a universal constant $C = C_n$ depending only on the dimension n such that*

$$\operatorname{sys}(M, g)^n \leq C_n \operatorname{Vol}(M, g).$$

The best way of understanding the significance of this estimate is seeing it as a part of the constellation of isoperimetric inequalities. To prove this systolic inequality, Gromov introduced a new metric invariant of riemannian manifolds, called *filling radius*, which informally speaking measures the "roundness" of the manifold and related it to both the systole and the volume. More precisely, for a submanifold M of an euclidean space $\mathbb{R}^N$, the filling radius can be viewed as the minimum radius of a neighbourhood around M in which M bounds (that is, such that there exists an $(n+1)$-chain F with $\partial F = M$). Thus the link between the filling radius and the volume is given by the Federer-Fleming isoperimetric inequality [32], which may be interpreted in this new language to provide an upper estimate for the filling radius of a manifold sitting in an Euclidean space in terms only of a dimensional constant and its volume. In order to extend this for a general closed riemannian manifold M, Gromov employed the Kuratowski isometric embedding to view M as a submanifold of the Banach space $L^\infty(M)$. The crucial technical step in Gromov's proof of Theorem 3.1.2 is an extension of the Federer-Fleming isoperimetric inequality [32] to submanifolds of Banach spaces, obtaining in this way an upper bound for the filling radius of a manifold only in terms of its volume and a dimensional constant. As for the estimate involving the systole and the filling radius, it is not difficult to prove that the systole of a riemannian manifold is dominated up to a universal constant by its Filing Radius.

Another perspective on the systolic inequality, more in line with the topological complexity viewpoint, is offered in the field of hyperbolic manifolds. The volume of a hyperbolic manifold (M, h) was long regarded as a good measure of the complexity of the topology of M, especially in the context of 3-manifolds (see for example [82, Chapter 6]). The Thurston-Milnor triangulation estimate [71] is a concrete manifestation of this principle.

Theorem (Thurston-Milnor)**.** *Let (M, h) be a closed hyperbolic n-manifold. Then, any triangulation of M requires at least $c_n \operatorname{Vol}(M, h)$ n-simplices, where c_n is a constant that depends only on the dimension.*

To see the connection with Gromov's systolic inequality, notice that since a hyperbolic manifold M is aspherical, the non-trivial topology of M occurs at a scale larger than $\operatorname{sys}(M, h)$. More precisely, balls of radii less than half the systole are contractible in M. Thus, intuitively it should take many balls of small radius (which play the rôle of the n-simplices in this analogy) to cover a hyperbolic manifold M with complex topology, which therefore cannot have arbitrarily small volume (see also Guth's survey [49] for more details).

With this in mind, the volume of a space normalized by its systole may be considered as a measure of the complexity of the underlying topology. This idea is formalized by the definition of a numerical invariant called *systolic volume*, introduced by Gromov, which is meaningful for the category of *piecewise riemannian polyhedra*, more general and flexible than that of riemannian manifolds. A piecewise riemannian polyhedron is a polyhedron in which every simplex is endowed with a smooth riemannian metric (as a manifold with boundary) in a compatible way. The systolic volume $\sigma(X)$ of a polyhedron X of dimension n (or also *systolic area*, in case that the dimension $n = 2$) is defined as

$$\sigma(X) := \inf_g \frac{\operatorname{Vol}(X, g)}{\operatorname{sys}(X, g)^n},$$

where the infimum is taken over the piecewise riemannian metrics g on X. The study of this invariant is one of the central problems in systolic geometry. In particular, one would like to know under which topological conditions on X this invariant is strictly positive, and once that is established, what is the precise value of σ within each topological type (or at least good estimates for it).

As for the first question, Gromov's proof of the systolic inequality is robust enough to apply to the class of *essential* polyhedra, although with worse constants. Recall that an n-dimensional polyhedron is called essential if there exists an aspherical polyhedron K together with a continuous map $X \to K$ that does not contract to the $(n-1)$-skeleton of K. In particular, by taking this map to be the identity, we see that all aspherical manifolds are essential. Using this new language, Gromov's systolic inequality takes the following form.

Theorem 3.2.6. *(Gromov) Let X be an essential polyhedron of dimension n. Then there exists a positive constant C_n such that $\sigma(X) \geq C_n$.*

With respect to the second problem, as it may be expected, it is an extremely difficult one. To the best of our knowledge, the only closed manifolds for which the systolic volume is known are the torus (due to Loewner), the projective plane [74] and the Klein bottle [11]. Moreover, it is expected that, in general, the metrics realizing the systolic volume of manifolds contain singularities (this happens for example in the case of the Klein bottle), so it could be very difficult to even guess how the optimal metrics look like. However, for the specific case of surfaces several good estimates are available. Most relevantly, the precise asymptotic behavior of the systolic area of surfaces with respect to the genus was determined by Gromov [44] (lower bound) and Buser and Sarnak [27] (upper bound).

Theorem 3.2.5. *There exist positive constants C, C' such that given a surface S of genus g,*

$$C\frac{g}{\log(g)^2} \leq \sigma(S) \leq C'\frac{g}{\log(g)^2}.$$

The minimum value that the systolic area of a surfaces can take is also known, due to yet another Gromov's inequality [44, §5.2.B]. Specifically, it can be deduced that $\sigma(S) \geq \frac{2}{\pi}$ for any closed (either orientable or non-orientable) non-simply connected surface S, with equality attained only at the real projective plane endowed with its standard round metric.

Besides surfaces, we are mainly concerned with the study of the systolic volume of polyhedra of dimension 2 for its connection with finitely presentable groups. Following Gromov [46], it is possible to attach a geometrically defined invariant to each finitely presentable group known as *systolic area*. The systolic area $\sigma(G)$ of a finitely presentable group G is defined as

$$\sigma(G) := \inf_{\pi_1(X)=G} \sigma(X),$$

where the infimum is over the compact piecewise riemannian polyhedra X with fundamental group isomorphic to G. Since plenty is known about the systolic area of closed surfaces, a natural question would be to understand how the systolic area of a closed surface S compares to the systolic area of its fundamental group, beyond the obvious inequality $\sigma(\pi_1(S)) \leq \sigma(S)$. A quite satisfactory answer was provided in [5], where the authors showed that the systolic area of fundamental groups of surfaces grows asymptotically with the genus g as $\frac{g}{\log(g)^2}$ and thus exhibits the same asymptotic behavior as the systolic area of surfaces. Moreover, our Theorem 4.4.7 could be read as an indication that the most efficient geometric-topological model of a surface group is the surface itself. Hence, the most optimistic conjecture would be that $\sigma(\pi_1(S)) = \sigma(S)$. Unfortunately, this problem seems very hard and in fact, the precise value of the systolic area is not known for any non-free group. As an alternative way to compare these quantities, we focus on improving the universal lower bound for the systolic area of surface groups. In contrast to the case of the systolic area of surfaces, for which it is known that $\frac{2}{\pi}$ is the optimum universal lower bound, there were no results in the literature addressing this issue for surface groups. Rather, the best lower bound available is valid for all non-free groups: by the combined works of Katz, Rudyak and Sabourau [64], Rudyak and Sabourau [77], and Katz, Katz, Sabourau, Shnider and Weinberger [63], it is known that $\sigma(G) \geq \frac{1}{4}$ whenever G is a non-free group. This constitutes a considerable improvement over the lower bound $\frac{1}{10^4}$, which is the estimate that can be extracted from the original techniques of Gromov for these groups [44, Theorem 6.7.A]. By analogy to the case of surfaces, the constant $\frac{2}{\pi}$ is a candidate to be optimal also for the systolic area of surface groups (or even for non-free groups, cf. [77, Question 1.5]). In this direction, we refine the universal lower bound for the systolic area of a large class of groups which includes the surface groups, bringing it closer to $\frac{2}{\pi}$. This class of groups, which we call *surface-like* groups is defined in terms of a cohomological condition.

Definition 3.3.9. *Let G be a group. We say that G is surface-like if there exist classes α, β in $H^1(G, \mathbb{Z}_2)$ such that $\alpha \cup \beta \neq 0$ in $H^2(G, \mathbb{Z}_2)$.*

Notice that as the name suggests, surface groups are surface-like as follows from Poincaré Duality.

Theorem 3.3.10. *Let K be a 2-dimensional complex such that there exist classes α, $\beta \in H^1(K, \mathbb{Z}_2)$ with $\alpha \cup \beta \neq 0$. Then, $\sigma(K) \geq \frac{1}{2}$. In particular, if G is a surface-like group, $\sigma(G) \geq \frac{1}{2}$.*

We remark here that the use of a covering theory argument borrowed from [63] (certainly also known to Gromov) allows to extend the inequality to groups containing a surface-like subgroup, a class that includes among others, free abelian groups, most of irreducible 3-manifold groups, non-free Artin groups and Coxeter groups or, more generally, groups containing an element of order 2.

The inequality from Theorem 3.3.10 can be regarded both as an extension of a systolic inequality due to Guth [50] in dimension 2 to complexes of maximal $\mathbb{Z}_2$ cup-length, and as a generalization of Burago and Hebda's inequality for closed essential surfaces [26, 53]. The proof, vaguely inspired in the (co)homological methods previously described, involves realizing a certain $\mathbb{Z}_2$-homology class of the 2-complex K by a simplicial map $f : S \to K$ from a (triangulated) closed surface S in a controlled way. The image of this map is a *surface* embedded in K (strictly speaking it is in general not a surface since its 1-skeleton may contain singularities), which turns out to have relatively large area by a systolic inequality of Nakamura [73] which refines Guth's inequality.

The Thesis is divided in four chapters. The first chapter is opened by a discussion about the Lusternik-Schnirelmann category and some related invariants, with emphasis in the recent discrete versions. The main subject is the definition and analysis of a new invariant of type L-S for compact polyhedra called the *PL geometric category*, designed to exploit the combinatorial structure of spaces admitting triangulations while at the same time capturing some of its inherent topology and geometry. In Chapter 2, we analyze the covering type of spaces and present the resolution of the quoted problem raised by Karoubi and Weibel [61] about minimal triangulations of the homotopy type of surfaces. The third chapter is devoted to systolic geometry, with a focus in the study of the systolic area of groups. Here we give the proof of our extension of Guth's systolic inequality. Also, we describe a construction for complexes mapping to a wedge sum based in Stalling's proof of Grushko theorem [80]. We think that this construction may shed new light in the study of optimal geometrical-topological models for free products of groups. In the final chapter we discuss the simplicial complexity of groups, some of its properties and its relation to the systolic area of groups, to finally present the answer to the question posed by Babenko, Balacheff and Bulteau [4] about optimal triangulations of surface groups.

Some of the original results in this Thesis appear in our works [17, 18, 16] and in the preprint [15]. Other results will be part of a work in preparation.

Contents

Chapter 1

Lusternik-Schnirelmann category and related invariants

The Lusternik-Schnirelmann (L-S) category is a numerical invariant associated to topological spaces, originally introduced in [67] for smooth manifolds in the context of variational problems. One of the goals of Lusternik and Schnirelmann was to relate the number of critical points of a smooth real function defined on a manifold to topological properties of the manifold. The classical Morse theory [72], developed shortly after with the same motivations, provides a deep connection between the structure of the set of critical points of a smooth map and the homotopy type of the manifold. Although the conclusions of Lusternik-Schnirelmann category theory are weaker, it complements Morse theory in the sense that it applies to manifolds of infinite dimension and to maps with degenerate critical points.

Some years later, Fox reformulated in [36] the definition of L-S category, extending it to general topological spaces. In that work, Fox analyzed among other things the interaction between the L-S category of a space and classical topological invariants, such as homology and homotopy groups and homotopy type, thus installing the invariant as a subject of study in the field of algebraic topology. In this new context, the L-S category may be interpreted as a measure of the complexity of the topology (or more precisely, of the homotopy type) of a space. Concretely, the L-S category is defined for a topological space X as the minimum number of open sets that are contractible in the space and cover X. Since contractible sets are the simplest possible sets from the homotopy theory point of view, the L-S category measures how difficult is to decompose a space into *simple* pieces.

In the first section of this chapter we briefly survey some fundamental properties and results about the L-S category. We make no attempt to be comprehensive here. Rather, we highlight some structural aspects, such as the relation of the L-S category with the cohomology and the dimension of a space, that will reappear later in our own work. Next we present some discretizations of the L-S category and relatives for simplicial complexes. In the recent years, besides the introduction of new L-S type invariants, there has been a considerable work aimed to obtain discrete versions of the L-S category and related results. Again, completeness is not the objective of the exposition. We limit ourselves to describe two such invariants, defined with very different motivations, which partially inspired our own contribution to the field of discrete L-S type invariants. Specifically, for the *discrete category* from [1] the authors recover a discrete version of the Lusternik-Schnirelmann theorem, which gives an estimate for the number of critical points of a smooth function on a manifold. On the other hand, with the introduction of the *simplicial L-S category* [34], the authors intend to give the right analogue of the classical L-S category for the category of finite simplicial complexes. In order to do that, they work within the theory of strong homotopy types of Barmak and Minian [9] in substitution of classical homotopy theory, obtaining a notion that depends strongly on the combinatorial structure of the simplicial complex. In the final sections we discuss a new L-S type discrete invariant, called *PL geometric category*, which uses the geometrically flavored notion of *PL collapsible* in replacement of the contractible sets as building blocks of spaces.

1.1 Lusternik-Schnirelmann category

This section is intended as a brief introduction to the Lusternik-Schnirelmann category theory. We want to make clear from the beginning that we regard the Lusternik-Schnirelmann category as one of the numerical invariants that measure the complexity of the topology of a space via optimal special covers. From this point of view, estimating the L-S category in terms of other invariants, both classical and new, is specially relevant. This leads us to privilege in the exposition aspects such as the homotopy invariance of the L-S category, as well as the relation of the invariant with the cohomology ring and the dimension of a space. For a more comprehensive account of the L-S category theory, we refer the reader to the book [29].

As we have already mentioned, the Lusternik-Schnirelmann (L-S) category of a topological space X, denoted $\operatorname{cat}(X)$, is the minimum number of open sets in a cover of X in which each open set is contractible in the space. Here, a set $U \subseteq X$ is contractible in X (or also, *categorical*) if the inclusion $U \hookrightarrow X$ is null-homotopic. At first, this requirement may look less natural than asking that the members of the cover are contractible (in themselves), but it is the key point that makes the L-S category a homotopy invariant. We remark here that several authors define the L-S category of a space with a different normalization, so that it differs by one from our definition (this is the case in the cited [29]). In particular, the L-S category of a contractible space equals 1 according to the definition here, while it is 0 for other authors.

Proposition 1.1.1. *The Lusternik-Schnirelmann category* $\operatorname{cat}(X)$ *of a topological space* X *is a homotopy invariant.*

Proof. Let $f : X \to Y$, $g : Y \to X$ be mutually inverse homotopy equivalences. By symmetry, it is enough to show that $\operatorname{cat}(Y) \leq \operatorname{cat}(X)$. Let $U_1, \ldots, U_n$ be a cover of X by open categorical sets. For each j, consider the open subset $V_j \subseteq Y$ defined as $V_j := g^{-1}(U_j)$. Let us see that the inclusion $i_{V_j} : V_j \to Y$ is null-homotopic. Indeed, since $f \circ g$ is homotopic to the identity map $Y \to Y$, the composition $f \circ g \circ i_{V_j} : V_j \to Y$ is homotopic to the inclusion i_{V_j}. On the other hand, it is straightforward to verify that $f \circ g \circ i_{V_j} = f \circ i_{U_j} \circ g|_{V_j}$, where $g|_{V_j}$ denotes the restriction of the map g to V_j and i_{U_j} is the inclusion $U_j \hookrightarrow X$. Since this last inclusion is null-homotopic, the map $f \circ i_{U_j} \circ g|_{V_j}$ is homotopic to a constant and hence, the inclusion i_{V_j} is null-homotopic as well. Summing up, the open cover of X by categorical sets $\{U_1, \ldots, U_n\}$ pulls back to an open cover of Y by categorical sets via g. Since the cover was arbitrary, the conclusion follows. $\square$

In view of the homotopy invariance of the L-S category, it is natural to try to compare it to the standard homotopy and homology invariants of a topological space. One of the most useful elementary estimates for the L-S category is derived from the ring structure of the cohomology ring of a space.

Definition 1.1.2. Let X be a topological space and fix a coefficient ring R. The *cup-length* $\mathrm{cup}_R(X)$ of X (with coefficients in R) is the least number k such that for any $(k+1)$ cohomology classes $\alpha_1, \dots \alpha_{k+1}$, the cup product $\alpha_1 \cup \cdots \cup \alpha_{k+1}$ is trivial in the reduced cohomology ring $H^*(X, R)$.

In other words, the cup-length of a space X with coefficients in a ring R is the index of nilpotency of the ring $H^*(X, R)$. While it may be infinite, for our principal class of interest which is the class of compact CW-complexes, the index of nilpotency is always finite and in fact, bounded from above by the dimension of the space. The following result relates the cup-length of a space to its L-S category. It constitutes one of the many connections of L-S category theory with notions of nilpotency (see for example [85]).

Lemma 1.1.3. *Let X be a topological space and R a ring. Then,* $\mathrm{cup}_R(X) < \mathrm{cat}(X)$.

Proof. Suppose that the value of the L-S category of X is a finite integer n (otherwise, there is nothing to prove) and take an open cover by categorical sets $\{U_1, \dots, U_n\}$. For every i, since the inclusion $U_i \hookrightarrow X$ is null-homotopic, the induced morphism in cohomology $H^*(X, R) \to H^*(U_i, R)$ is trivial. From the exactness of the sequence

$$\cdots \to H^k(X, U_i; R) \to H^k(X, R) \to H^k(U_i, R) \to \cdots,$$

it follows that the maps $H^k(X, U_i; R) \to H^k(X, R)$ are surjective. Hence, if $\alpha_1, \dots, \alpha_n$ are cohomology classes in $H^*(X, R)$, we can pull back each α_i to the group $H^*(X, U_i; R)$. But then, via the canonical maps $H^*(X, \bigcup U_i; R) \to H^*(X, U_i; R)$ it turns out that the cup product $\alpha_1 \cup \cdots \cup \alpha_n \in H^*(X, R)$ is the image of a corresponding product in $H^*(X, \bigcup U_i; R)$. Since $\bigcup U_i = X$, this ring is trivial and hence $\alpha_1 \cup \cdots \cup \alpha_n = 0$, as we wanted to prove. □

The most classical estimate from above is given by the *covering dimension* of a space. The covering dimension of a (Hausdorff, paracompact) space is defined, similarly as the L-S category, in terms of optimizing a certain quantity related to open covers of the space. Since the covering dimension of a CW-complex coincides with the usual notion of dimension, we opt to prove the result only for this case. Actually, for compact CW-complexes there is a stronger estimate involving the L-S category and the dimension. We take the opportunity to introduce the notion of *geometric category* of a space, which was defined by Fox in [36] to approximate the L-S category. The geometric category $\mathrm{gcat}(X)$ of a topological space X is simply the minimum number of contractible open sets that cover X. The more natural condition of working with contractible sets rather than contractible relative to the space as in the definition of the L-S category comes at the cost of the geometric category not being a homotopy invariant, as the following example due to Fox shows [36, §39] (see also [29, Proposition 3.11]).

Lemma 1.1.4. *The geometric category is not a homotopy invariant.*

Proof. Consider the topological space X formed by adding two edges to the sphere S^2 joining a point p_1 with two different points p_2 and p_3. By identifying to a point two simple paths contained in the sphere that connect p_1 to p_2 and p_3, we see that X is homotopy equivalent to a wedge sum $S^2 \vee S^1 \vee S^1$ (since the union of the two paths is contractible). On the other hand, the quotient Y of X by the two extra edges is homeomorphic to S^2 with the three points p_1, p_2 and p_3 identified. Hence, Y is homotopy equivalent to the wedge sum $S^2 \vee S^1 \vee S^1$.

Let us see that the geometric category of these two spaces differs. It is not difficult to check that $\mathrm{gcat}(S^2 \vee S^1 \vee S^1) = 2$. Indeed, it is possible to cover the space with two contractible open sets, each obtained by choosing one (slightly fattened) hemisphere of each sphere. On the other hand, it is not possible to cover Y by two contractible open sets. Suppose on the contrary that $Y = U_1 \cup U_2$ for contractible open sets U_1 and U_2. Notice that the connected components of the preimage of U_1 and U_2 under the quotient map $q : S^2 \to Y$ are contractible. Moreover, if we set $V_i := q^{-1}(U_i)$ for $i = 1, 2$, by Alexander duality the last statement implies that $S^2 \setminus V_i$ is connected. We split the proof in two cases, according to whether the point $x = q(p_1)$ is in the intersection of U_1 and U_2 or not. In the second case, without loss of generality $x \in U_1 \setminus U_2$. Since the complement $S^2 \setminus V_2$ is connected, it is completely contained in one connected component of V_1. Then, the other connected components of V_1 are contained in V_2. Since no two of the points p_j may belong to the same connected component of V_1 (otherwise, U_1 would not be acyclic), this implies that at least one of the p_j is in U_2, a contradiction. If x is in the intersection of the sets U_1 and U_2, each V_i should have three different connected components V_i^j, one containing each of the points p_j. Since $S^2 \setminus V_1$ is connected, it is contained in one of the connected components of V_2, say, V_2^1. Then, the components V_2^2, V_2^3 are contained in the corresponding components V_1^2 and V_1^3 of V_1. Applying the same argument to the complement $S^2 \setminus V_2$, we see that at least one of the components V_1^2, V_1^3 is contained in the corresponding component V_2^2, V_2^3. But then, one of the components of V_1 equals a component of V_2, which implies that $V_1 \cup V_2$ is the union of mutually disjoint open sets. This is a contradiction, since V_1 and V_2 cover S^2. □

Although not a homotopy invariant, the geometric category clearly constitutes an upper bound for the L-S category and it is often easier to estimate than the L-S category. It is via this invariant that we deduce, for compact CW-complexes, an inequality relating the L-S category and the dimension.

Lemma 1.1.5. *Let X be a compact connected CW-complex. Then* $\mathrm{gcat}(X) \leq \dim(X) + 1$. *In particular,* $\mathrm{cat}(X) \leq \dim(X) + 1$.

Proof. Let U_i be the union of the interior of all the i-cells of X. Since X is connected, hence path-connected, and each connected component of U_i is contractible, these sets can be enlarged to contractible open sets by joining the components appropriately through slightly inflated paths. A word is in order about the 1-dimensional case. Here, the contractible set that can be formed out of the 0-cells is a spanning tree of the complex X, while the other one is a spanning tree of X plus the interior of all the 1-cells not belonging to the spanning tree. We see then that it is always possible to cover X by at most $\dim(X) + 1$ contractible open sets. □

Using these elementary estimations, it is possible to compute the L-S category of some relevant spaces (cf. [29, Example 1.8]). For example, we can check that the L-S category of both the n-dimensional torus $\mathbb{T}^n$ and the n-dimensional projective plane $\mathbb{R}P^n$ is exactly $n+1$. Indeed, that $\mathrm{cat}(\mathbb{T}^n) \leq n+1$ and $\mathrm{cat}(\mathbb{R}P^n) \leq n+1$ follows from Lemma 1.1. For the lower bound, we compute the cup-length of both spaces and apply Lemma 1.1.3. The cohomology ring of the torus $\mathbb{T}^n$ with coefficients in $\mathbb{Q}$ is isomorphic to an exterior algebra on n generators and hence $\mathrm{cup}_{\mathbb{Q}}(\mathbb{T}^n) = n$. With respect to the projective space $\mathbb{R}P^n$, the cohomology ring with coefficients in $\mathbb{Z}$ is isomorphic to $\mathbb{Z}[\omega]/\langle \omega^{n+1} \rangle$, where ω has degree 1, and hence the cup-length of $\mathbb{R}P^n$ with coefficients in $\mathbb{Z}$ equals n. By Lemma 1.1.3, $\mathrm{cat}(\mathbb{T}^n) \geq n+1$, $\mathrm{cat}(\mathbb{R}P^n) \geq n+1$ and the computation is complete.

We close this introductory section by presenting the original motivation for the Lusternik-Schnirelmann theory. Lusternik and Schnirelmann were interested in finding critical points of functionals defined over manifolds (generally, infinite-dimensional) in the context of variational problems. More precisely, they were interested in critical points of the energy functional over the free loop space of a smooth closed riemannian manifold. Recall that the free loop space $\mathcal{L}M$ of a riemannian manifold M is the infinite-dimensional manifold formed by the smooth loops $S^1 \to M$, and that the energy functional $E : \mathcal{L}M \to \mathbb{R}$ sends a loop γ to its *energy* $\frac{1}{2}\int_{S^1} \|\dot{\gamma}(t)\|^2$. The analysis of the critical points for E is motivated by the fact that they clearly contain the closed geodesics on M. In this direction, the Lusternik-Schnirelmann theorem asserts that the number of critical points of a smooth function $f : M \to \mathbb{R}$ defined over a smooth manifold (not necessarily of finite dimension) that satisfies a certain compactness condition is bounded from below by $\mathrm{cat}(M)$ (see [29, Chapter 1] for the precise statement and the proof). Probably, this result reminds the reader of the classical Morse theory, which gives a close connection between the topology of a manifold M and the critical points of a smooth real map $f : M \to \mathbb{R}$, provided that f is *Morse*, that is, those critical points are non-degenerate. The important point to notice is that, while Morse theory has stronger implications than the Lusternik-Schnirelmann theory when both apply, the latter allows for some control in the degenerate case. In the case under consideration the traditional Morse theory does not apply, since among other things constant loops are degenerate critical points of the energy functional E. (We remark however, that for *generic* riemannian metrics on M the Bott's extension [19] of Morse theory allows to obtain the classical conclusions for this functional).

Lusternik and Schnirelmann employed the estimate on the number of critical points in the proof of the three closed geodesics theorem, arguably the most celebrated application of Lusternik-Schnirelmann theory.

Theorem 1.1.6. *Let (S, g) be a riemannian closed surface of genus 0 (that is, S is topologically a sphere). Then, there exist at least three different embedded closed geodesics in S.*

For a proof, see for example [6, 39], where a gap in the original proof of Lusternik and Schnirelmann is corrected.

1.2 Discrete versions of L-S type invariants

Over the past few years, several works were devoted to introduce what we loosely call discrete versions of the L-S category (and related invariants). Such versions usually apply to topological spaces with a rigid combinatorial structure, like simplicial complexes or finite topological spaces. Moreover, the simple building blocks are defined in terms of a homotopy theory that exploits this rigid structure, such as Whitehead's simple homotopy theory. In this section we concentrate in two of these notions, which in part motivated us to introduce another natural discrete L-S category type invariant.

We start by describing the *discrete (geometric) category* from [1], which is inspired in the geometric category but instead of covers by contractible sets employs *collapsible* sets as building blocks. Before the actual definition, we need to recall the basic concepts of Whitehead's simple homotopy theory. A simplex σ of K is a *free face* of a finite simplicial complex K if there is a unique simplex $\tau \in K$ containing σ. In that case, we say that there is an *elementary collapse* from K to $L = K \setminus \{\sigma, \tau\}$, denoted $K \searrow^{e} L$. Notice that the inclusion $L \hookrightarrow K$ is a strong deformation retract and hence, K and L are homotopy equivalent. More generally, K *collapses* to L, denoted by $K \searrow L$, if there is a sequence $K_1 = K, K_2, \ldots, K_r = L$ such that $K_i \searrow^{e} K_{i+1}$ for every i; in particular, K and L are homotopy equivalent in such situation. We also say that L expands to K and denote $L \nearrow K$. The complex K is called *collapsible* if it collapses to a complex with only one vertex. Since collapses are in particular homotopy equivalences, collapsible complexes may be considered as an analogue to contractible sets in this context.

We now give the definition of discrete geometric category, warning the reader that we employ a different normalization than the original definition, so that both differ by 1.

Definition 1.2.1. (cf. [1, Definition 8]) Let K be a simplicial complex and $L \leq K$ a subcomplex. The *discrete geometric pre-category* of L in K is the minimum number (the minimum minus 1 in the original definition) of collapsible subcomplexes of K that cover L. The *discrete geometric category* $\mathrm{dgcat}(K)$ of K is the minimum among the discrete geometric pre-categories of those subcomplexes L to which K collapses.

The use of subcomplexes in the definition instead of open subsets is justified by the well-known (kind of) correspondence between them for a simplicial complex. Namely, every subcomplex L of a simplicial complex K may be inflated to a slightly larger open set which deformation retracts to L. In the other direction, an open subset of K deformation retracts to a subcomplex of an appropriate subdivision of K. Thus, ignoring for one moment the intermediate pre-category, the discrete geometric category may be viewed as a version of the usual geometric category in the context of simplicial complexes and simple homotopy theory. As the authors observe in [1, Remark 9], the definition guarantees that an elementary collapse does not decrease the geometric category of a complex avoiding the need to subdivide. In a sense, we explore in the next section the properties of the invariant that results from taking the other path, that is, allowing subdivisions.

Despite this, the discrete geometric category is not a simple homotopy invariant, just as the geometric category is not a homotopy invariant. Recall that a pair of simplicial complexes K and L are *simple homotopy equivalent* if there exists a finite sequence of complexes $K_1 = K, K_2, \ldots, K_r = L$ such that for every i either $K_i \nearrow^{e} K_{i+1}$ or $K_i \searrow^{e} K_{i+1}$. To show that dgcat is not a simple homotopy invariant, we prove that, in dimension 2, this invariant generalizes the notion of collapsible complex. We state a simple lemma first.

Lemma 1.2.2. *Let K be a collapsible simplicial complex of dimension 2 and let L be a subcomplex of K. If $\dim L = 2$, L collapses to a graph, i.e. a complex of dimension 1.*

Proof. Fix an ordering $\sigma_1, \sigma_2, \ldots, \sigma_r$ of the 2-simplices of K that induces a valid sequence of collapses. It is clear then that the first 2-simplex of L appearing in that list must have a free face in L and hence L collapses to a subcomplex with one fewer 2-simplex. By induction on the number of 2-simplices of L, it follows that L collapses to a graph. □

Proposition 1.2.3. *Let K be a finite, connected simplicial complex of dimension 2. Then,* $\mathrm{dgcat}(K) = 0$ *if and only if K is collapsible.*

Proof. One direction is obvious. For the other one, suppose that $\mathrm{dgcat}(K) = 0$, so that there exists a subcomplex L to which K collapses contained in another collapsible subcomplex $M \leq K$. By Lemma 1.2.2, L collapses to a 1-dimensional complex. Hence, K collapses to a 1-dimensional subcomplex, so it suffices to check that K is acyclic to show that K is collapsible. Now, since the inclusion $L \hookrightarrow K$ is a homotopy equivalence, it induces an isomorphism $H_*(L) \to H_*(K)$ (with integer coefficients). On the other hand, the isomorphism factors through $H_*(M)$ which is trivial because M is collapsible. The conclusion follows. □

As a consequence, we see that dgcat is not a simple homotopy invariant by considering some fixed triangulation of the dunce hat D. As it is well known, the dunce hat is not collapsible (hence $\mathrm{dgcat}(D) > 0$) but it is simple homotopy equivalent to a point.

Another salient feature of the discrete geometric category is that it depends strongly on the simplicial structure and not only on the topology. To illustrate this point, we consider the complete graph K_n on n vertices, which has $\binom{n}{2}$ edges. Since trees (i.e., acyclic connected graphs) are the only collapsible graphs and a subtree of K_n may have at most $n-1$ edges, at least $\frac{1}{n-1}\binom{n}{2} \geq \frac{n}{2}$ collapsible subcomplexes are required to cover K_n, whence $\mathrm{dgcat}(K_n) \geq \frac{n}{2}$. On the other hand, it is not difficult to show that a subdivision of K_n may be decomposed as the union of two collapsible subcomplexes (see Proposition 1.3.5 below for more details).

Nonetheless, there is a point which discrete geometric category shares with the continuous L-S category: it provides a lower bound for the number of critical points of certain functions, called *discrete Morse functions*, defined over finite simplicial complexes. It is this result which motivated the introduction of the discrete geometric category as an analogue of the L-S category for simplicial complexes.

Theorem 1.2.4. *([1, Theorem 26]) Let $f : K \to \mathbb{R}$ be a discrete Morse function with m critical points. Then* $\mathrm{dgcat}(K) \leq m$.

The so-called discrete Morse theory was introduced by Forman in [35]. We content here with referring the reader to that work for the basic definitions of discrete Morse theory and to [1] for the proof of the theorem.

For its part, with the introduction of the *simplicial L-S category* in [34], the authors intend to develop the Lusternik-Schnirelmann theory for simplicial complexes using the notion of *contiguity* as a discrete version of homotopy. We recall next the definition of contiguity classes of simplicial maps, together with the description of this homotopy theory in terms of simple combinatorial moves by Barmak and Minian [9]. Let $\phi, \psi : K \to L$ be two simplicial maps. We say that ϕ and ψ are *contiguous* and denote $\phi \sim_c \psi$ if for every simplex $\sigma \in K$, $\phi(\sigma) \cup \psi(\sigma)$ is a simplex in L. Notice that two contiguous maps are homotopic via the linear homotopy (in each simplex). Two simplicial maps are in the same *contiguity class*, noted $\phi \sim \psi$ whenever they are joined by a finite sequence of contiguous simplicial maps $\phi = \phi_0 \sim_c \cdots \sim_c \phi_n = \psi$. By analogy with the notion of homotopy equivalence, a simplicial map $\phi : K \to L$ is called a *strong equivalence* if there exists another simplicial map $\psi : L \to K$ with $\phi \circ \psi \sim \mathrm{id}_L$ and $\psi \circ \phi \sim \mathrm{id}_K$. In the last case, we note $K \sim L$. This equivalence relation between simplicial complexes was rewritten in terms of *strong homotopy types* in [9].

Definition 1.2.5. [9] Let K be a simplicial complex and $v \in K$ a vertex. We say that v is dominated by a vertex $v' \neq v$ if every maximal simplex that contains v also contains v'. If v is dominated by some vertex v', we say that there is an *elementary strong collapse* from K to $K \setminus v$ and denote $K \searrow\!\!\searrow^{e} K \setminus v$. In that situation we also say that there is an *elementary strong expansion* from $L = K \setminus v$ to K and denote it by $L\, {}^{e}\!\nearrow\!\!\nearrow K$. If there is a sequence of elementary strong collapses that starts in K and ends in L, we say that there is a *strong collapse* from K to L and denote $K \searrow\!\!\searrow L$. The inverse of a strong collapse is called a *strong expansion* and denoted by $L \nearrow\!\!\nearrow K$. Finally, K and L have the same *strong homotopy type* if there is a sequence of strong collapses and expansions that transforms K in L.

Remark 1.2.6. [9, Remark 2.4] $K \searrow\!\!\searrow L$ implies that $K \searrow L$.

As Barmak and Minian show, the strong collapses and expansions completely determine the discrete homotopy theory based in contiguity.

Theorem 1.2.7. *[9, Corollary 2.12] The simplicial complexes K and L have the same strong homotopy type if and only if $K \sim L$.*

The notion of *simplicial L-S category* is obtained almost verbatim from the definition of L-S category, replacing null-homotopic by contiguous to a constant map. Concretely, the *simplicial L-S category* $\operatorname{scat}(K)$ of a simplicial complex K is the minimum number of subcomplexes that cover K such that the inclusion of each one is contiguous to a constant map (by a slight abuse, we will also call *categorical* such subcomplexes). This invariant relies even more strongly than dgcat on the simplicial structure. For example, in [34, Example 3.3] the authors exhibit a certain subdivision K of the 2-simplex which is not strong collapsible, i.e., it does not strong collapse to a vertex, and show that $\operatorname{scat}(K) = 1$ (this complex originally appeared in [9]). Although this example shows that scat is far from being a homotopy invariant, it turns out to be strong homotopy invariant, resembling the classical L-S category in this aspect.

Proposition 1.2.8. *([34, Proposition 3.7] The simplicial L-S category is a strong homotopy invariant.*

Proof. Let $f : K \to L$, $g : L \to K$ be mutually inverse strong simplicial equivalences. The proof is completely analogous to that of Proposition 1.1.1. Take a cover $M_1, \ldots, M_n$ of K by categorical subcomplexes. We check that the subcomplex $N_j := g^{-1}(M_j) \leq L$ is categorical in L for every j. Indeed, it is enough to observe that the simplicial map $f \circ g \circ i_{N_j}$ is contiguous to i_{N_j} on one hand and equals $f \circ i_{M_j} \circ g|_{N_j}$ on the other, which is contiguous to a constant map since i_{M_j} is by hypothesis. ☐

1.3 Covers of polyhedra by PL collapsible subpolyhedra

In this section we define and explore a natural variant of the geometric category in the context of compact connected polyhedra, which we call *PL geometric category* (here *PL* stands for *piecewise linear*). Although this is formally a discrete version of the geometric category and we recover some results that hold for the classical case, our motivation is more naïve: what topological consequences, both local and global, may be derived for a polyhedron that is covered by few collapsible subpolyhedra?

Let us clarify in first place what we understand by the terms *polyhedron* and *(PL) collapsible*. By a *polyhedron* we mean a topological space which admits triangulations, that is, the underlying space of some simplicial complex. More simply, one can think of a polyhedron as a simplicial complex but in which there is no preferred simplicial structure. Thus, a subspace Q of a polyhedron P is called a *subpolyhedron* if it is the underlying space of a subcomplex of some triangulation of P. The basic concepts from simple homotopy theory for polyhedra are defined analogously. We say that a polyhedron P *PL collapses* to a subpolyhedron Q (still denoted by $P \searrow Q$) if there exist coherent triangulations K, L of P and Q respectively such that $K \searrow L$ (see [55, Ch.2]). A polyhedron P is called *PL collapsible* if it PL collapses to a point, i.e. some simplicial complex that triangulates P collapses to a vertex.

Definition 1.3.1. Let P be a polyhedron. The *PL geometric category* $\mathrm{plgcat}(P)$ of P is the minimum number of PL collapsible subpolyhedra that cover P.

Since subpolyhedra and open sets in a polyhedron are, homotopically speaking, equivalent, the geometric category of a polyhedron coincides with the minimum number of contractible subpolyhedra that cover it. This formulation makes clear that the PL geometric category may be regarded as a version of the classical geometric category.

A feature shared by most of the L-S category type invariants (in the continuous context) is that they are bounded by the dimension of the space. The geometric category is no exception: as we showed above, the geometric category of a (compact, connected) polyhedron P of dimension n is at most $n+1$ (see Lemma 1.1). Our first result concerning the PL geometric category shows that it also verifies this nice structural property, namely, that a polyhedron of dimension n is covered by at most $n+1$ PL collapsible subpolyhedra. To prove this result, we follow essentially the same strategy as in the proof of the geometric category version: first, inductively cover the $(n-1)$-skeleton of an appropriate triangulation of the polyhedron by n PL collapsible subpolyhedra, slightly inflate this cover and then form the last PL collapsible subpolyhedron by joining a triangulation of the interior of each n-simplex. However, since being PL collapsible imposes a greater rigidity than just being contractible, we are forced to take a slight technical detour to adapt this strategy to our context. Specifically, we carry out part of the argument by employing the strong homotopy types of [9], which we briefly introduced in Definition 1.2.5.

We take care in first place of the part of the argument that involves extending a cover by PL collapsible subpolyhedra to a slightly inflated polyhedron. To start, we need to recall the notion of *star* and *link* of a vertex v in a simplicial complex, which are the simplicial analogues of ball and sphere around the vertex, respectively. The *star* of a vertex v in a simplicial complex K is the subcomplex $\mathrm{st}_K(v) \subseteq K$ formed by the union of the simplices $\sigma \in K$ such that $\sigma \cup v \in K$. The *link* of v is the subcomplex $\mathrm{lk}_K(v) \subseteq \mathrm{st}_K(v)$ of the simplices that do not contain v. For a given simplex σ, its *boundary* $\partial\sigma$ is the subcomplex formed by the simplices τ strictly contained in σ. We add here an alternative description of domination of vertices that uses these concepts.

Remark 1.3.2. A vertex v in a simplicial complex K is dominated by v' if and only if the link $\mathrm{lk}_K(v)$ is a simplicial cone with apex v', i.e. $\mathrm{lk}_K(v) = v'M$ for certain subcomplex M.

The next two lemmas allow to inflate the cover locally, that is, simplex by simplex. In the first one, we prove the intuitively clear fact that if we remove the central part of a sufficiently fine subdivision of a simplex, the resulting subcomplex will collapse to the boundary.

Lemma 1.3.3. *Let σ_n be the standard n-simplex. Consider the subcomplex of the second barycentric subdivision of σ_n defined as $K_n := \sigma_n'' \setminus \operatorname{st}_{\sigma_n''}(\{v\})$, where v is the barycenter of σ_n. Then K_n strong collapses to $(\partial\sigma_n)''$.*

Proof. For formal reasons, it will be convenient to adopt the set theoretic definition of simplicial complex along the proof. Thus, we consider the simplices of the second subdivision of σ_n as chains of simplices of σ_n' ordered by inclusion. That means that a vertex of σ_n'' is given by a 0-chain $\{\tau\}$ where τ is some simplex of σ_n', an edge of σ_n'' as a 1-chain $\{\tau \subseteq \eta\}$ for some simplices τ, η of σ_n' and so on.

Let $w \in \operatorname{lk}_{\sigma_n''}(\{v\})$ be a vertex. It follows from the definition of link that there is a 1-simplex in σ_n'' with vertices $\{v\}$ and w and so there is a chain of inclusion of simplices of σ_n' $\{\{v\} \subseteq w\}$. Suppose $w = \{e\}$, where $e = \{v, a\}$ is a 1-simplex of σ_n'. Then, any maximal simplex of σ_n'' containing w either contains $\{v\}$ or $\{a\}$. Since $\{v\} \notin K_n$, this shows that w is dominated by $\{a\}$ in K_n.

Consider now the complex $\tilde{K}_n$ obtained from K_n by removing all the vertices of the form $\{e\}$ for e a 1-simplex of σ_n' containing v. Take a vertex $u \in \operatorname{lk}_{\sigma_n''}(\{v\}) \cap \tilde{K}_n$. Suppose that $u = \{\tau\}$ for some 2-simplex $\tau = \{v, a, b\}$ of σ_n'. Since $\tilde{K}_n$ does not contain vertices of the form $\{\{v, x\}\}$, nor the vertex $\{v\}$, any maximal simplex of $\tilde{K}_n$ containing u also contains $\{a, b\}$. Hence u is dominated by $\{a, b\}$ in $\tilde{K}_n$. By removing the vertices of $\operatorname{lk}_{\sigma_n''}(\{v\})$ in non-decreasing order of the dimension of the simplex of σ_n' that they represent, we see that $K_n \searrow\!\!\searrow (\partial\sigma_n)''$. □

Lemma 1.3.4. *Let K, L be simplicial complexes such that $L \nearrow\!\!\nearrow K$. If L can be covered by n strong collapsible subcomplexes, so does K.*

Proof. Let $\{L_1, \ldots, L_n\}$ be a cover of L by n strong collapsible subcomplexes and assume that there is an elementary strong expansion from L to K, say $L = K \setminus v$ for certain $v \in K$. Let $v' \in K$ be a vertex that dominates v, so that $\operatorname{lk}_K(v) = v'M$ for some subcomplex M of L. For each $1 \leq i \leq n$, define the subcomplex K_i of K as

$$K_i = \begin{cases} L_i \cup v(v'M \cap L_i) \text{ if } v'M \cap L_i \neq \emptyset, \\ L_i \text{ otherwise.} \end{cases}$$

If $v'M \cap L_i$ is nonempty, then $v \in K_i$ and is clearly dominated by v' because $\operatorname{lk}_{K_i}(v) = v'(M \cap L_i)$. In any case, K_i strong collapses to L_i and is therefore strong collapsible. This shows that K is covered by n strong collapsible subcomplexes. The conclusion follows by induction on the number of elementary strong expansions from L to K. □

We are now able to prove that the PL geometric category of a polyhedron of dimension n is bounded from above by $n + 1$. We will prove the following slightly stronger result.

Proposition 1.3.5. *Let K be a complex of dimension n. Then, the second barycentric subdivision K'' of K can be covered by $n + 1$ strong collapsible subcomplexes.*

Proof. Proceed by induction on n, the dimension of K. When $n = 1$, K is a simplicial graph. We show first that in this case K' admits a cover by two strong collapsible subcomplexes. In order to produce the strong collapsible cover, let T be a spanning tree of the graph K and note that each edge $e \in K \setminus T$ becomes the union of two edges in K', say $e = e_1 \cup e_2$. Consider the following subcomplexes of K':

$$K_1 = T' \cup \bigcup_{e \in K \setminus T} e_1 \ , \ K_2 = T' \cup \bigcup_{e \in K \setminus T} e_2.$$

As K_1, K_2 both strong collapse to T', they are strong collapsible and they clearly cover K'. Since their barycentric subdivisions are also strong collapsible, the base case is complete.

Let now K be a simplicial complex of dimension n. By inductive hypothesis, the second barycentric subdivision of the $(n-1)$-skeleton $\left(K^{(n-1)}\right)''$ of K can be covered by n strong collapsible subcomplexes $K_1, \dots, K_n$. Let $v_1, \dots, v_r$ be the barycenters of the simplices of dimension n of K. By Lemma 1.3.3, we see that $\left(K^{(n-1)}\right)'' \nearrow\!\!\!\nearrow K'' \setminus \bigcup_{i=1}^{r} \mathrm{st}_{K''}(\{v_i\})$ and so Lemma 1.3.4 implies that this last complex is covered by n strong collapsible subcomplexes. Since K'' is connected and $\mathrm{st}_{K''}(\{v_i\})$ is strong collapsible for every i, we can include their union in a strong collapsible subcomplex of K''. $\square$

For dimension 1, this implies that the PL geometric category only distinguishes trees (contractible graphs) from the rest of graphs, just like the geometric category. In contrast, the PL geometric category of 2-dimensional polyhedra has some subtler properties that we explore in the next section.

1.4 PL geometric category in dimension 2

We know by Proposition 1.3.5 that the PL geometric category of a 2-dimensional polyhedron is bounded by 3. Since PL collapsible polyhedra are relatively well understood, the interest in dimension 2 is centered in distinguishing polyhedra of PL geometric category 2 from those of PL geometric category 3. Without reaching a full characterization, in this section we show among other things that the simple homotopy type of a polyhedron P with $\mathrm{plgcat}(P) = 2$ is severely constrained and derive a criterion, which involves both global and local conditions, to decide that certain polyhedra do not admit covers by two PL collapsible subpolyhedra.

We analyze in first place the simple homotopy type of polyhedra of plgcat 2. If we look at the analogous situation for the geometric category, we find that the homotopy type of the 2-dimensional polyhedra which are the union of two contractible subpolyhedra is completely determined. Indeed, by the Seifert-van Kampen theorem the fundamental group of such a polyhedron is free of rank equal to the number of connected components minus 1 of the intersection of the elements of the contractible cover. Then, by a result of C.T.C. Wall [83] a 2-dimensional polyhedron with free fundamental group is homotopy equivalent to a wedge sum of spheres of dimension 1 and 2. Thus, it is reasonable to expect that an even stronger restriction holds for polyhedra P with $\mathrm{plgcat}(P) = 2$, which is in fact the case. Recall that for a pair of simple homotopy equivalent simplicial complexes K and L we say that K *n-deforms* to L, or that there is an n-deformation from K to L, if there exists a finite sequence of complexes $K_1 = K, K_2, \dots, K_r = L$ such that for every i either $K_i \overset{e}{\nearrow} K_{i+1}$ or $K_i \overset{e}{\searrow} K_{i+1}$ and the dimension of K_i is at most n.

Lemma 1.4.1. *Let K be a simplicial complex of dimension 2 which is covered by collapsible subcomplexes K_1, K_2. Then there is a 3-deformation from K to the suspension $\Sigma(K_1 \cap K_2)$ of $K_1 \cap K_2$.*

Proof. Cone off K_1, K_2 with vertices v_1, v_2. This gives an expansion $K \nearrow v_1K_1 \cup v_2K_2$. Collapse every new simplex based on a simplex contained in K_1 or K_2 but not in both. Hence, $K \nearrow v_1K_1 \cup v_2K_2 \searrow v_1(K_1 \cap K_2) \cup v_2(K_1 \cap K_2)$, which is the desired 3-deformation. $\square$

The combination of this result with Lemma 1.2.2, which provides a certain control over the intersection of collapsible complexes in dimension 2, yields the following restriction on the simple homotopy type of polyhedra of plgcat 2.

Proposition 1.4.2. *Let P be a polyhedron of dimension* 2 *such that* $\mathrm{plgcat}(P) = 2$*. Then P* 3*-deforms to the suspension of a graph.*

Proof. Take a triangulation K of P covered by collapsible subcomplexes K_1, K_2. By Lemma 1.4.1, K 3-deforms to $\Sigma(K_1 \cap K_2) = v_1(K_1 \cap K_2) \cup v_2(K_1 \cap K_2)$ and by Lemma 1.2.2 $K_1 \cap K_2$ collapses to a 1-dimensional subcomplex G. It follows that $v_i(K_1 \cap K_2) \searrow v_iG$ for $i = 1, 2$, and hence K 3-deforms to the suspension of G. $\square$

Remark 1.4.3 (*On the Andrews-Curtis conjecture*)**.** The Andrews-Curtis conjecture [3] states that (compact) contractible polyhedra of dimension 2 3-deform to a point. As a consequence of Proposition 1.4.2, the Andrews-Curtis conjecture is satisfied by contractible polyhedra which admit a cover by two PL collapsible subpolyhedra. Indeed, let P be a contractible polyhedron covered by collapsible subpolyhedra P_1, P_2. From the Mayer Vietoris sequence, the intersection $P_1 \cap P_2$ has trivial homology and by Lemma 1.2.2, $P_1 \cap P_2$ collapses to a tree. Then, by Proposition 1.4.2, P 3-deforms to a point.

Even though Proposition 1.4.2 provides a valuable piece of information, it is far from characterizing those polyhedra that admit covers by two collapsible subpolyhedra. Indeed, as it may be suspected by analogy with the geometric category, the PL geometric category is not a (simple) homotopy invariant of a polyhedron. Even more, this may be deduced from exactly the same example used by Fox to show that the geometric category is not a homotopy invariant, described in Lemma 1.1.4. Let X_1 be the wedge sum of S^2 and two circles and let X_2 be the space obtained from S^2 by identifying three distinct points. Notice that X_1 and X_2 are simply homotopy equivalent (in fact, there is a 3-deformation from X_1 to X_2, which follows from the argument employed in Lemma 1.1.4 to show that X_1 and X_2 are homotopy equivalent). By splitting every sphere in X_1 in two, we see that X_1 admits a cover by two PL collapsible subpolyhedra and hence $\mathrm{plgcat}(X_1) = 2$. On the other hand, since X_2 does not admit covers by two contractible subpolyhedra by the proof of Lemma 1.1.4, $\mathrm{plgcat}(P_2) = 3$.

For the rest of the section, we will try to understand under what conditions it is possible to determine the PL geometric category of a polyhedron once we already know that its simple homotopy type is the correct one. If a given polyhedron P admits a cover by two collapsible subpolyhedra P_1 and P_2, most of the relevant information about the topology of P is concentrated in the way the intersection $P_1 \cap P_2$ is embedded in P. A particularly favorable situation would be that this intersection is 1-dimensional (a graph). With this idea in mind, we describe a class of polyhedra that we call *inner connected* for which it is always the case that the intersection may be deformed to be 1-dimensional.

Definition 1.4.4. Let K be a simplicial complex of dimension 2. We say that an edge of K is *inner* if it is a face of exactly two 2-simplices of K.

Recall that a simplicial complex K of dimension n is *homogeneous* or *pure* if all of its maximal simplices have dimension n.

Definition 1.4.5. Let K be a homogeneous 2-dimensional simplicial complex. We say that K is *inner-connected* if any pair of 2-simplices σ, τ of K is connected by a sequence of 2-simplices $\sigma = \eta_1, \eta_2, \dots, \eta_r = \tau$ such that $\eta_i \cap \eta_{i+1}$ is an inner edge of K for each $1 \leq i < r$. We call such a sequence an *inner sequence*. We say that a polyhedron P is inner-connected if one, (or equivalently, any) of its triangulations is inner-connected.

Example 1.4.6. Surfaces or more generally pseudosurfaces are inner-connected. The presentation complex associated to finite one-relator presentation in which every generator appears at least once in the relator is also inner-connected.

The intuition here is that inner edges of a polyhedron determine locally in a concrete manner two sides and that, in order to cover the polyhedron by two collapsible subpolyhedra, it should be enough to pick one side for each member of the cover around the intersection edges. The global condition required in the definition of inner connected polyhedra guarantees that the elections may be performed coherently. We formalize this in the next result.

Lemma 1.4.7. *Let K be an inner-connected and non collapsible simplicial complex of dimension 2. Suppose that K is the union of collapsible subcomplexes K_1, K_2. Then there exist collapsible subcomplexes L_1, L_2 such that $K = L_1 \cup L_2$ and $L_1 \cap L_2$ is 1-dimensional.*

Proof. Suppose that $K_1 \cap K_2$ has at least one 2-simplex η. Since $K_1 \cap K_2$ is a proper subcomplex of K, we can find a 2-simplex not in $K_1 \cap K_2$ and an inner sequence joining it to η. Then there are 2-simplices σ, τ together with an inner edge $e = \sigma \cap \tau$ such that $\tau \in K_1 \cap K_2$ but $\sigma \notin K_1 \cap K_2$. Without loss of generality, suppose that $\sigma \in K_1$. Then e is a free face of the complex K_2, which implies that we can remove τ from K_2. That is, the complexes K_1 and $\tilde{K}_2 = K_2 \setminus \tau$ form again a collapsible cover of K and $K_1 \cap \tilde{K}_2$ has one fewer 2-simplex than $K_1 \cap K_2$. It follows by induction that it is possible to find collapsible subcomplexes L_1, L_2 that cover K and intersect in a graph. $\square$

Now that we know that the intersection of a PL collapsible cover of size 2 is 1-dimensional for a considerable class of polyhedra, the next object of interest is the structure of the graph itself. We focus on the leaves of this intersection graph, since they are the natural candidates to start simplifying the graph. More specifically, we find that very often the local topology of the polyhedron around a leaf of the intersection graph exhibits a certain rigidity. We introduce a definition from [51] to make this statement precise.

Definition 1.4.8. [51] Let K be a simplicial complex. A vertex v of K is a *bridge* if $K \setminus v$ has more connected components than K. We say that v is *splittable* if the link $\mathrm{lk}_K(v)$ has bridges. Note that it makes sense to say that a point in a polyhedron is splittable because this property depends only on the homeomorphism type of a small closed neighborhood around the point and not on a specific triangulation of the space.

Both the content and the proof of the following lemma are inspired on results from [8, 51].

Lemma 1.4.9. *Let K be a homogeneous complex of dimension 2 which admits a collapsible cover of size two. Suppose additionally that the link of every non splittable vertex of K is connected. Then, there exist collapsible subcomplexes L_1, L_2 that cover K and such that every leaf of the 1-skeleton $(L_1 \cap L_2)^{(1)}$ of $L_1 \cap L_2$ is a splittable vertex of K.*

Proof. Let K_1 and K_2 be subcomplexes of K that form a collapsible cover of K. Take $\eta = vw \in (K_1 \cap K_2)^{(1)}$ an edge such that w is a leaf, i.e. $\mathrm{lk}_{(K_1 \cap K_2)^{(1)}}(w) = v$, but not a splittable vertex. Suppose in first place that η is not maximal in either of the subcomplexes K_1, K_2, so that there exist vertices $v_i \in K_i$ with $vwv_i \in K_i$ for $i = 1, 2$. As w is not a splittable vertex, we can find a path joining v_1 and v_2 in $\mathrm{lk}_K(w) \setminus v$. But then there must be at least another edge in $\mathrm{lk}_{K_1 \cap K_2}(w)$ contradicting the hypothesis that η is a leaf of $(K_1 \cap K_2)^{(1)}$. Suppose now η is maximal in K_1 and take $\tau = v_2\eta$ a 2-simplex of K_2 containing η (we can find one by homogeneity of K). We show that in this case K_1 collapses to $K_1 \setminus w$. If it was not the case, there should be another edge $\eta' \in K_1$ hanging from w. By the homogeneity of K, η' is the face of some 2-simplex $\sigma = v_1\eta'$ which per force is in K_1 but not in K_2. Since by hypothesis w is not splittable and has connected link, there is a path in $\mathrm{lk}_K(w) \setminus v$ joining v_1 to v_2 and so w cannot be a leaf of $(K_1 \cap K_2)^{(1)}$, a contradiction. By performing the collapses that correspond to edges in the second case, we may assume that the leaves of $(K_1 \cap K_2)^{(1)}$ are splittable vertices. □

Suppose now that a 2-dimensional polyhedron P is covered by collapsible subpolyhedra P_1, P_2. In the case that the intersection is a graph, we already know something about its local behavior at leaves. With respect to the global topology of $P_1 \cap P_2$, since by Lemma 1.2.2 it collapses to a graph, its homotopy type is completely determined by the homology. As a straightforward computation using the (reduced) Mayer-Vietoris long sequence reveals,

$$\tilde{H}_0(P_1 \cap P_2) \equiv H_1(P),\ H_1(P_1 \cap P_2) \equiv H_2(P),$$

where the homology groups are taken with coefficients in $\mathbb{Z}$. From Proposition 1.4.2, we know that $H_1(P)$ and $H_2(P)$ are finitely generated free abelian groups. Now, if $\mathrm{rk}\, H_2(P) < \mathrm{rk}\, H_1(P)$ at least two connected components of $P_1 \cap P_2$ are acyclic. Since the polyhedron $P_1 \cap P_2$ collapses to a graph, this implies that at least two connected components of $P_1 \cap P_2$ are collapsible. If moreover these components are graphs, $P_1 \cap P_2$ should have at least two leaves which in turn implies by Lemma 1.4.9 that P has at least two vertices that are splittable or with non connected links, provided P is homogeneous. The conclusion reached in this paragraph is roughly that an inner-connected polyhedron which is regular both in a local and a global sense does not admit PL collapsible covers of size two.

Theorem 1.4.10. *Let P be an inner-connected polyhedron of dimension* 2 *such that* $H_2(P) \equiv 0$ *or* $\mathrm{rk}\, H_2(P) < \mathrm{rk}\, H_1(P)$. *Suppose additionally that P is not PL collapsible, has at most one splittable vertex and that the link of every non splittable vertex is connected. Then* $\mathrm{plgcat}(P) = 3$.

Proof. The case $\mathrm{rk}\, H_2(P) < \mathrm{rk}\, H_1(P)$ was already treated in the paragraph above. Suppose then $H_2(P) \equiv H_1(P) \equiv 0$ and that P is the union of PL collapsible subpolyhedra P_1, P_2 that intersect in a graph. Hence, $P_1 \cap P_2$ is a tree and since we may assume by Lemma 1.4.9 that its leaves are located in splittable vertices, $P_1 \cap P_2$ should be a point. It follows that P is a wedge sum of PL collapsible polyhedra, which contradicts the hypothesis that P be inner-connected. □

Example 1.4.11. The dunce hat D is an inner-connected contractible polyhedron with only one splittable vertex and such that every other vertex has connected link. Hence, by Theorem 1.4.10 no triangulation of D admits a cover by two collapsible subcomplexes. In fact, we can say a little more. The dunce hat D can be viewed as the presentation complex associated to the one-relator presentation $\langle\, a \,|\, aaa^{-1} \,\rangle$ (see the first paragraph of Section 1.4.1). More generally, by Theorem 1.4.10 none of the presentation complexes associated to a presentation of the form $\langle\, a \,|\, a^n a^{-(n-1)} \,\rangle$ ($n \geq 2$) admits a cover by two PL collapsible subpolyhedra.

We remark that this example generalizes a result from [40], where the authors show that a specific triangulation of the dunce hat cannot be written as the union of two collapsible subcomplexes.

Example 1.4.12. The standard Bing's house with two rooms admits a PL collapsible cover of size two (to see this, split the complex in two halves, each one containing the walls which support the vertical tunnels). Note that unlike the dunce hat, the Bing's house with two rooms is not inner-connected. Moreover, as a consequence of the proof of Theorem 1.4.10 it is impossible to cover this polyhedron by two PL collapsible subpolyhedra intersecting in a graph.

1.4.1 The geometry of one-relator presentations

In this brief section we compute completely the PL geometric category of the one-relator presentation complexes, by showing that the value of this invariant can be read off directly from the presentation.

We mentioned in passing in Example 1.4.6 that the complexes associated to one-relator presentations are inner-connected. We recall here the precise construction of complexes associated to presentations. Let $\mathcal{P} = \langle\, X \,|\, R \,\rangle$ be a finite presentation. We associate a topological model to $\mathcal{P}$ in the following way. First we form $K = \vee_{x \in X} S^1_x$, the wedge sum of 1-spheres indexed by X. Every word $r \in R$ spells out a combinatorial loop on the space K based on the wedge point, which is used to attach a 2-cell on K. The resulting 2-dimensional CW-complex is called the *presentation complex* of $\mathcal{P}$ and is denoted by $K_{\mathcal{P}}$. Since the attaching maps are combinatorial, the presentation complex $K_{\mathcal{P}}$ is a polyhedron (see [54, Chapter 2] for more details). When the set R consists of only one word r the presentation $\langle\, X \,|\, r \,\rangle$ is called a *one-relator presentation*.

Homogeneous presentation complexes associated to one-relator presentations are inner connected: since they have only one 2-cell, any two points are connected by a path through the interior of the cell. To avoid unnecessary extra case-by-case analysis we work exclusively with homogeneous one-relator presentation complexes, which amounts to ask that every generator appears in the relator. If this was not the case, such a presentation complex K would decompose as the wedge sum of a bouquet of 1-spheres K_1 and a homogeneous complex K_2, which turns out to be also a presentation complex (associated to the presentation obtained by removing the generators not appearing in the relator). Since K_1 is always the union of two PL collapsible subpolyhedra (simply split in two halves each 1-sphere), the PL geometric category of K coincides with $\mathrm{plgcat}(K_2)$, unless K is PL collapsible.

The computation of the PL geometric category for homogeneous one-relator complexes is divided in two cases, according to whether the presentation admits an *algebraic collapse* or not. Here, we say that a one-relator presentation $\mathcal{P} = \langle\, x_1, \ldots, x_k \,|\, r \,\rangle$ admits an algebraic collapse if one of the generators x_i occurs only once in r, with exponent ± 1. The case of presentations admitting algebraic collapses is the easier of the two.

Lemma 1.4.13. *Let K be a connected simplicial complex and let $L_1, \dots, L_n$ be disjoint collapsible subcomplexes of K. Then there exists a collapsible subcomplex of K containing $\bigcup_{i=1}^n L_i$.*

Proof. Since K is connected and L_1 is disjoint with $\bigcup_{i=2}^n L_i$, there exists a simple path p in the 1-skeleton of K joining a vertex of L_1 with a vertex of some L_i ($i \neq 1$) with no edges in $\bigcup_{i=1}^n L_i$. Consider the subcomplex M of K defined as $M := L_1 \cup p \cup L_i$. Since L_1 and L_i are collapsible, they collapse to any of its vertices and hence M collapses to p which is in turn collapsible. The result now follows from induction. □

Proposition 1.4.14. *Let $\mathcal{P} = \langle x_1, \dots, x_k \,|\, r \rangle$ be a finite one-relator presentation and suppose that r admits an algebraic collapse. Then $K_{\mathcal{P}}$ admits a cover by two PL collapsible subpolyhedra, that is,* $\operatorname{plgcat}(K_{\mathcal{P}}) \leq 2$.

Proof. Without loss of generality, x_1 occurs only once in r and the relator is of the form $r = x_1^{\pm 1} a_1 \dots a_{m-1}$, where each a_i is equal to some $x_j^{\pm 1}$, $j \neq 1$. Picture the complex $K_{\mathcal{P}}$ as a disk with the boundary subdivided in m edges labeled in counterclockwise order according to r. Subdivide the edge labeled x_1 in $2(m-1)+1$ edges and subdivide the rest of the edges in three edges. Join the $2i$-th edge of the subdivided x_1 to the central edge of (the edge labeled as) a_i by a 2-dimensional strip inside the disk in such a way that the strips are pairwise disjoint (see Figure 1.1).

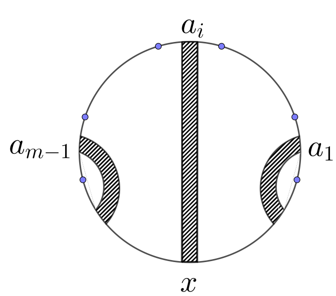

Figure 1.1: The strips (shaded) PL collapse to a tree through the edge which intersects the edge labeled x_1.

Both the subpolyhedron P_1 formed by the union of these strips and its complement P_2 consist of a disjoint union of PL collapsible subpolyhedra of $K_{\mathcal{P}}$. Hence, by Lemma 1.4.13, P_1 and P_2 may be included in PL collapsible polyhedra Q_1 and Q_2 that cover $K_{\mathcal{P}}$. □

As for the other case, notice that in a one-relator presentation complex with no algebraic collapses every point has a connected link, except possibly the wedge point. Since homogeneous one-relator complexes are inner-connected, as a consequence of Theorem 1.4.10 most such complexes do not admit PL collapsible covers of size two.

Proposition 1.4.15. *Let $\mathcal{P} = \langle x_1, \ldots, x_k \,|\, r\rangle$ be a finite one-relator presentation such that r does not admit algebraic collapses. Then, if the number of generators k is greater than 1 or $H_2(K_{\mathcal{P}})$ is trivial,* $\operatorname{plgcat}(K_{\mathcal{P}}) = 3$.

Proof. By cellular homology, the group $H_2(K_{\mathcal{P}})$ is free abelian of rank at most 1. Moreover, by a straightforward Euler characteristic computation we know that

$$\operatorname{rk} H_2(K_{\mathcal{P}}) - \operatorname{rk} H_1(K_{\mathcal{P}}) = 1 - k.$$

Hence, if $\mathcal{P}$ has $k > 1$ generators, we have $\operatorname{rk} H_2(K_{\mathcal{P}}) < \operatorname{rk} H_1(K_{\mathcal{P}})$ and the conclusion follows from Theorem 1.4.10. The case $H_2(P) \equiv 0$ is also covered by Theorem 1.4.10. $\square$

To complete the picture, it remains to handle the specific subcase of the one-relator presentations with only one generator and non trivial second homology group, which are exactly the presentations of the form $\langle\, x \,|\, r\,\rangle$, where r is a word on letters x, x^{-1} with total exponent 0. The computation of the PL geometric category for these complexes requires a careful ad-hoc analysis.

Proposition 1.4.16. *Let $\mathcal{P} = \langle\, x \,|\, r\,\rangle$ be a one-relator presentation such that $H_2(K_{\mathcal{P}})$ is not trivial. Then,* $\operatorname{plgcat}(K_{\mathcal{P}}) = 2$ *if and only $\mathcal{P}$ is of the form $\langle\, x \,|\, (xx^{-1})^{\pm 1}\,\rangle$.*

Proof. Suppose that a triangulation of $K_{\mathcal{P}}$ admits a cover by collapsible subcomplexes K_1, K_2, which we may suppose to intersect in a graph with at most one leaf by Lemmas 1.4.7 and 1.4.9. Since $H_0(K_1 \cap K_2) \equiv \mathbb{Z}^2$ and $H_1(K_1 \cap K_2) \equiv \mathbb{Z}$, one of the connected components of $K_1 \cap K_2$ is acyclic and therefore consists of only one point. For this to be possible, the link of the wedge point v must have more than one connected component.

Recall that the Whitehead graph of r (see [54, Ch.6]) allows to determine the (homeomorphism type of the) link of the wedge point in presentation complexes directly from the relations. Let us briefly indicate how the Whitehead graph is constructed for the specific case at hand. Apart from v, mark two points x_+ and x_- in the loop x, one walking away from the base point following the orientation of x, the other on the contrary sense. Now, if for example xx is a subword of the relator r, there is two-dimensional material connecting the final part of the first x cycle to the initial part of the second one, so that there should be an edge connecting x_- to x_+ in the link of v. The Whitehead graph has as vertices x_+ and x_- and the edges are derived from the word r (viewed cyclically) generalizing the last example in the obvious way. Hence, it is not difficult to see that the link of v is not connected only for presentations of the form $\langle\, x \,|\, (xx^{-1})^{\pm n}\rangle$, $n \in \mathbb{N}$, since the presence of subwords xx or $x^{-1}x^{-1}$ connects the points x_+ and x_-.

Given one such presentation, let us call C the connected component of $K_1 \cap K_2$ which is not the wedge point. Since it is a connected graph with one cycle and no leaves, C is homeomorphic to S^1. Moreover, by perturbing slightly K_1 and K_2 near the loop x if necessary, the intersection of this component with x may by assumed to be transversal, that is, a finite set of points. Also, notice that the intersection $C \cap x$ cannot be empty. Indeed, if the loop x was entirely contained in K_1 (the argument for K_2 is identical), this complex would have non trivial first homology group because the homology class determined by x generates $H_1(K_{\mathcal{P}})$, leading to a contradiction. Let then w be a point in $C \cap x$ and let us label by a and b the edges of the subdivision of x (one in each side) that contain w. Since C intersects the loop x only at points, the edges a and b are completely contained in different members of the collapsible cover, say $a \in K_1 \setminus K_2$ and

$b \in K_2 \setminus K_1$. Notice that, in any triangulation of $K_{\mathcal{P}}$, the edges a, b are the face of $2n$ 2-simplices. Furthermore, the (open) star of w is homeomorphic to a union of $2n$ half euclidean planes with the x axis identified. It follows that vertex w has valency $2n$ in the graph C. This is impossible unless $n = 1$. Finally, observe that the complex associated to a presentation of the form $\langle\, x \,|\, (xx^{-1})^{\pm 1} \,\rangle$ is homeomorphic to a 2-sphere with its poles identified and so admits a cover by two PL collapsible subpolyhedra. □

The complete characterization of the behavior of the PL geometric category for one-relator presentation complexes follows as a corollary to Propositions 1.4.14, 1.4.15 and 1.4.16, and may be summarized as follows.

Theorem 1.4.17. *Let $\mathcal{P} = \langle x_1, \ldots, x_k \,|\, r\rangle$ be a finite one-relator presentation. Then $K_{\mathcal{P}}$ can be covered by two PL collapsible subpolyhedra if and only if r admits an algebraic collapse or $\mathcal{P}$ is of the form $\langle\, x \,|\, (xx^{-1})^{\pm 1} \,\rangle$.*

1.4.2 Inner-connected polyhedra

The purpose of this section is to prove a structural result about inner connected polyhedra. Specifically, we show that all inner connected polyhedra can be obtained as a quotient of a disk, much in the same way that closed surfaces are formed by making identifications on pairs of boundary edges of an appropriate polygon. We follow the treatment and notation of [66, Ch.6].

Definition 1.4.18. Given a finite alphabet S, a *word* of length k in S is an ordered list of k symbols of $S \cup S^{-1}$. A *polygonal presentation* $\mathfrak{P}$ is a finite alphabet S together with a finite set of words $W_1, \ldots, W_r$ in S of length at least three such that every element of S (or its formal inverse) appears in some word. We denote such a presentation by $\mathfrak{P} = \langle S \,|\, W_1, \ldots, W_r\rangle$.

A polygonal presentation $\mathfrak{P}$ determines a topological space (called the *geometric realization* of $\mathfrak{P}$) in the following fashion. For each word W_i in $\mathfrak{P}$ of length k form the regular convex polygon of k sides P_i and label its edges in counterclockwise order according to W_i, starting by an arbitrary vertex. Now identify edges with the same label in $\coprod_i P_i$ by the simplicial homeomorphism that matches the vertices of the edges, inverting orientation when necessary.

There is a number of combinatorial movements on polygonal presentations, called *elementary transformations*, that preserve the (PL) homeomorphism type of the corresponding geometric realizations. We describe here only the transformations we will use and refer to [66, Ch.6] for the complete list.

- Reflection: $\langle S \,|\, a_1 \ldots a_m, W_2, \ldots, W_r\rangle \mapsto \langle S \,|\, a_m^{-1} \ldots a_1^{-1}, W_2, \ldots, W_r\rangle$.
- Rotation: $\langle S \,|\, a_1 a_2 \ldots a_m, W_2, \ldots, W_r\rangle \mapsto \langle S \,|\, a_2 \ldots a_m a_1, W_2, \ldots, W_r\rangle$.
- Pasting: $\langle S, e \,|\, W_1 e, e^{-1} W_2, \ldots, W_r\rangle \mapsto \langle S \,|\, W_1 W_2, \ldots, W_r\rangle$. Note that e does not belong to S so that none of the words $W_1, \ldots, W_r$ should contain e for this transformation to be valid.

The main result of this section states that every inner-connected polyhedron P has a polygonal presentation with one word. The strategy of the proof consists of repeatedly pasting pairs of 2-simplices of a triangulation of P joined by an inner edge until we are left with only one polygon. However, the vertices of an inner edge may be *singular*, i.e. they may have a neighborhood that is not homeomorphic to an open disk. By pasting a pair of 2-simplices through an inner edge with singular vertices, we may create identifications in the interior of the polygon we are building. We show in the next lemma that we can avoid this situation by considering a sufficiently fine triangulation of P.

Lemma 1.4.19. *Let K be an inner-connected simplicial complex of dimension 2. Then, each pair of simplices of the second barycentric subdivision K'' of K may be joined by an inner sequence such that the vertices of the inner edges in the sequence are not singular.*

Proof. Let σ_1, σ_2 be a pair of 2-simplices of K'' and let τ_1, τ_2 be respectively the 2-simplices of K containing them. Since K is inner-connected, there is an inner sequence S in K joining τ_1 to τ_2. It is not difficult to find an inner sequence in K'' formed by 2-simplices contained in 2-simplices of S that avoids the vertices of K and joins σ_1 to σ_2. $\square$

Theorem 1.4.20. *Let P be an inner-connected polyhedron. Then P admits a polygonal presentation with only one word.*

Proof. Take a simplicial complex K that triangulates P. By Lemma 1.4.19, we may assume that every pair of 2-simplices of K is joined by an inner sequence such that the inner edges involved do not have singular vertices. Choose a different label for each edge of K and fix an orientation for its simplices. Consider the polygonal presentation $\mathfrak{P}$ that has as alphabet the set of labels of edges of K and a word for each 2-simplex, determined by the edges of its boundary in the order given by the prescribed orientations. Since the geometric realization of $\mathfrak{P}$ is homeomorphic to P, it suffices to reduce $\mathfrak{P}$ to a presentation with one word by applying elementary transformations. The 2-simplex that corresponds to W_1 has at least one inner edge a with no singular vertices. Without loss of generality, assume that W_2 is the only other word in which a or a^{-1} appears. By applying rotations and reflections we may assume that $W_1 = \tilde{W}_1 a$, $W_2 = a^{-1}\tilde{W}_2$ and paste them to reduce the number of words in $\mathfrak{P}$. Inductively, suppose that there is more than one word in the presentation. We claim that there is one inner edge with no singular vertices of K that appears exactly once in W_1. Indeed, if it was not the case, it would be impossible to connect a 2-simplex of (the subcomplex determined by) W_1 and a 2-simplex not in W_1 by an inner sequence with no singular vertices in its edges. As before, rearrange the words and perform rotations and reflections in such a way that it is possible to paste words W_1 and W_2. $\square$

Remark 1.4.21. Let P be an inner-connected polyhedron and let $\mathfrak{P} = \langle S \,|\, W \rangle$ be a polygonal presentation of P with one word obtained as in Theorem 1.4.20. Consider the subgraph G of P formed by the edges determined by S. The word W defines a surjective combinatorial map $\varphi : S^1 \to G$ for a suitable triangulation of S^1. This provides an alternative description of inner-connected polyhedra. Concretely, given a simplicial graph G and a surjective combinatorial map $\varphi : S^1 \to G$ we obtain an inner-connected polyhedron as the space underlying the CW-complex that consists of one 2-cell attached to G according to φ. In the case that G is homeomorphic to a bouquet of spheres of dimension 1, the resulting space is a (homogeneous) one-relator presentation complex.

Chapter 2

Minimal triangulations of homotopy types of surfaces

Closed smooth manifolds have long known to be triangulable and studied from that perspective. Several classical results are devoted to understanding the restrictions imposed by the topology of a manifold to its triangulations, such as the Dehn-Sommerville relations and the Lower Bound Theorem (see [10, 60]). A particularly active area in the field is the study of minimal triangulations or triangulations with few simplices of manifolds. Besides its intrinsic interest, the availability of optimal or almost optimal triangulations of manifolds would allow to compute (with the assistance of a computer) some invariants such as the first Pontrjagin class [38, 37].

On the other hand, the use of optimal triangulations to measure the complexity of the topology of manifolds has also a long tradition, especially in the context of 3-manifolds (see [70]) and hyperbolic manifolds (see [71] and Chapter 3 for more details). In the recent article [61] Karoubi and Weibel related the topological complexity to minimal triangulations of the homotopy type of spaces through the introduction of the *covering type*. The covering type of a space is a L-S category type invariant designed to measure the complexity of a space through its *good* covers. The connection to minimal triangulations comes from the fact that for a space X of the homotopy type of compact simplicial complexes, the covering type may be reformulated as the number of vertices in a minimal triangulation of a simplicial complex homotopy equivalent to X.

The main objective of this chapter is the exposition of the complete computation of the covering types of closed surfaces, which settles a problem formulated by Karoubi and Weibel. This computation reveals that the optimal triangulations of closed surfaces turn out to be, with only one exception, the most economical models within their homotopy type. For this reason, we open the chapter with a discussion of Jungerman and Ringel theorem [58, 76], which characterizes the optimal triangulations of closed surfaces, both orientable and non-orientable. We continue with a brief introduction to the covering type and some of its basic properties after Karoubi and Weibel. The remaining two sections concern the computation of the covering type of closed surfaces. The tools involved are fairly classical and elementary: the ring structure of the cohomology ring of surfaces and the Euler characteristic formula.

2.1 Minimal triangulations of surfaces

We say as usual that a simplicial complex K is a triangulation of a topological space X if the geometric realization of K is homeomorphic to X. As it is well known, all smooth manifolds admit triangulations. The problem of finding small or minimal triangulations of smooth manifolds has been extensively investigated (see for example [20, 21, 68, 69, 58, 76]), while the interest in optimal triangulations of closed surfaces goes back at least to 1950 [58]. The objective of this section is presenting the complete solution to the problem of minimal triangulations of closed surfaces obtained by Ringel [76], in the non-orientable case, and Jungerman and Ringel [58] in the orientable case.

Theorem 2.1.1. *Let S be a closed surface different from the orientable surface of genus 2 (M_2), the Klein Bottle (N_2) and the non-orientable surface of genus 3 (N_3). There exists a triangulation of S with n vertices if and only if*

$$n \geq \frac{7+\sqrt{49-24\chi(S)}}{2}.$$

For the exceptional cases M_2, N_2 and N_3 it is necessary to replace n by $n-1$ in the formula above.

The proof of this result consists of two separate parts, and, as it may be expected, the hard part is to produce triangulations of a surface S with n vertices whenever n exceeds the number in the formula. It involves extensive use of the theory of current graphs and a complete exposition would deviate us from our objectives, so we refer the interested reader to the original works [76, 58]. We content here with explaining the argument to prove the lower bound on the number of vertices, which amounts to a relatively easy computation with the Euler characteristic formula and will be relevant later. We introduce some convenient notation before the proof.

Notation. Let $k \leq 2$ be an integer number. We denote by $\rho(k)$ the integer defined as

$$\rho(k) := \left\lceil \frac{7+\sqrt{49-24k}}{2} \right\rceil.$$

By abuse of notation, for a simplicial complex K of dimension 2 with $\chi(K) \leq 2$, we will write $\rho(K)$ to mean $\rho(\chi(K))$. Also, $\alpha_i(K)$ will denote the number of i-simplices in the complex K.

Proof of Theorem 2.1.1 (lower bound). Let K be a triangulation of a closed surface S. Then, by the Euler characteristic formula,

$$\chi(S) = \chi(K) = \alpha_0(K) - \alpha_1(K) + \alpha_2(K).$$

Since every edge of K is incident to exactly two 2-simplices, we see that $3\alpha_2(K) = 2\alpha_1(K)$ by double counting. On the other hand, since K is a simplicial complex it has at most $\binom{\alpha_0(K)}{2}$ edges. Then

$$6\chi(S) \geq 6\alpha_0(K) - \alpha_0(K)(\alpha_0(K)-1).$$

If $\chi(S) \leq 0$, the minimum strictly positive integer that satisfies this inequality is precisely $\rho(S) = \rho(\chi(S))$ and therefore $\alpha_0(S) \geq \rho(S)$. An easy analysis shows that $\alpha_0(S) \geq \rho(S)$ also when $\chi(S) = 1, 2$. □

We have implicitly so far talked about minimal triangulations of surfaces as triangulations that minimize the number of vertices. For general smooth manifolds, a minimal triangulation often means one that minimizes the number of maximal dimension faces. In the case of surfaces, there is a linear relation between the number of simplices in each dimension and so the different possible meanings of minimal triangulations coincide in dimension 2.

Notation. Let S be a closed surface. We will denote by $\delta(S)$ the minimum number of 2-simplices in a triangulation of S and by $\lambda(S)$, the minimum number of vertices (0-simplices).

Thus, Theorem 2.1.1 provides an explicit formula for both $\delta(S)$ and $\lambda(S)$, whenever S is a closed surface. We restate it in form of a lemma for future reference.

Lemma 2.1.2. *Let S be a closed surface. Then,*

$$\lambda(S) = \rho(S) = \left\lceil \frac{7 + \sqrt{49 - 24\chi(S)}}{2} \right\rceil,$$

except for $S = M_2, N_2,$ and N_3, in which cases it is $\lambda(S) = \rho(S) + 1$. On the other hand,

$$\delta(S) = 2\lambda(S) - 2\chi(S).$$

Proof. The formula for $\lambda(S)$ is simply a restatement of Theorem 2.1.1. For the second statement, recall that if K triangulates the surface S, $3\alpha_2(K) = 2\alpha_1(K)$ because every edge is the face of exactly two 2-simplices. Hence, from the Euler characteristic formula we have

$$2\chi(S) = 2\alpha_0(K) - 2\alpha_1(K) + 2\alpha_2(K) = 2\alpha_0(K) - \alpha_2(K),$$

and it follows that $\alpha_2(K) = 2\alpha_0(K) - 2\chi(S)$. □

2.2 The covering type of spaces

The purpose of this section is presenting an invariant known as *covering type*, introduced by Karoubi and Weibel in [61] as a way to measure the complexity of topological spaces. The covering type of a space X admitting a triangulation is intimately related to the minimal triangulations of X (in the sense of vertex minimizing) and it was our main motivation to study triangulations of the homotopy type of closed surfaces.

Before giving the definition of the covering type of a space, we need to recall the notion of *good cover*. An open cover $\mathcal{U} = \{U_i\}_{i\in I}$ of a topological space X is called a good cover if every nonempty intersection $U_{i_1} \cap \cdots \cap U_{i_n}$ is contractible. Good covers have been used by Weil to prove some of de Rham's theorems about the cohomology of manifolds (see [84, 61]). They appear naturally in the context of riemannian manifolds: any point in a riemannian manifold has a geodesically convex neighborhood (which, in particular, is contractible) and the intersections of such neighborhoods are again geodesically convex. A refinement of the classical notion of Lusternik-Schnirelmann category in terms of such covers was recently formulated by Karoubi and Weibel [61].

Definition 2.2.1. The *strict covering type* of X is defined to be the minimum number of elements in a good cover of X. The *covering type* of X, denoted by $\mathrm{ct}(X)$, is the minimum of the strict covering types taken over all spaces Y homotopy equivalent to X.

One of the most salient features of this invariant is the connection to minimal (vertex) triangulations of the homotopy types of spaces. Indeed, reinterpreting Proposition 2.1 and Theorem 2.5 from [61] (see also [17, Lemma 2.1]), we obtain the following equivalent definition for the covering type of spaces of the homotopy type of (finite) CW-complexes.

Lemma 2.2.2. *Let X be a topological space homotopy equivalent to a finite CW-complex. Then $\mathrm{ct}(X)$ coincides with the minimum possible number of vertices of a simplicial complex K homotopy equivalent to X.*

Proof. Suppose that $\mathrm{ct}(X) = r$ and let Y be a space homotopy equivalent to X which admits a good cover $\mathcal{U}$ of size r. By the Nerve Theorem, Y is homotopy equivalent to the nerve $\mathcal{N}(\mathcal{U})$, which is a simplicial complex with r vertices. Reciprocally, if K is a simplicial complex homotopy equivalent to X, the open stars of its vertices form a good cover of K. □

To explore the behavior of the covering type of spaces, one may start by analyzing connected graphs, i.e. connected 1-dimensional complexes. The homotopy type of such spaces is completely determined by a positive integer number h, the rank of the first homology group. Concretely, a connected graph Y with $\mathrm{rk}(H_1(Y, \mathbb{Z})) = h$ is homotopy equivalent to a bouquet (wedge sum) of h circles. An explicit formula for the covering type was computed for this family by Karoubi and Weibel.

Proposition 2.2.3. *([61, Proposition 4.1]) Let X_h be a bouquet of h circles. Then, the covering type of X_h is*

$$\mathrm{ct}(X_h) = \left\lceil \frac{3 + \sqrt{1 + 8h}}{2} \right\rceil .$$

That is, $\mathrm{ct}(X_h)$ equals the unique integer n satisfying

$$\binom{n-2}{2} < h \leq \binom{n-1}{2}.$$

To contextualize this result, recall that the L-S category of connected graphs is at most 2. Hence, it indicates that the covering type retains more information about the homotopy type of the underlying space. On the other hand, the formula becomes more transparent when one thinks about the alternative formulation for the covering type of a graph from Lemma 2.2.2. Intuitively, one should try to construct a simplicial complex K homotopy equivalent to X_h with as little vertices as possible. It is clear that, among simplicial graphs on r vertices, the complete graph K_r has the maximum rank $\binom{r-1}{2}$ in H_1. Thus, it is impossible to form a simplicial graph on n vertices homotopy equivalent to X_h if $\binom{n-1}{2} < h$, while by removing some edges (if necessary) from the complete graph K_n one may construct such a graph in the opposite case.

We close this section by relating the covering type of a space to the Betti numbers, a more classical manifestation of the presence of non-trivial topology. Recall that for a space X with finitely generated homology (say, $H_k(x) = 0$ for $k > m$) the *Poincaré polynomial* $P_X(t)$ is defined as

$$P_X(t) = \beta_0 + \beta_1 t + \cdots + \beta_m t^m,$$

where β_i is the rank of the homology group $H_i(X)$ with coefficients taken in a field. Karoubi and Weibel showed the following result.

Theorem 2.2.4. *([61, Theorem 3.3]) Let $P_X(t)$ be the Poincaré polynomial of X and let n be its covering type. If X is not empty then:*

$$P_X(t) \leq \frac{(1+t)^{n-1}-1}{t} + 1 = n + \binom{n-1}{2}t + \binom{n-1}{3}t^2 + \cdots + t^{(n-2)},$$

meaning that $\beta_0 \leq n$, $\beta_1 \leq \binom{n-1}{2}, \ldots, \beta_{n-2} \leq 1$ and $\beta_i = 0$ for $i \geq n-1$.

Proof. The proof is by induction on n, the cases $n = 1, 2$ being trivial since the polynomial $P_X(t)$ equals 1 and 2, respectively, for those cases.

For the inductive step, if $\mathrm{ct}(X) = n$, without loss of generality X admits a good cover of size n and hence $X = \cup_{i=1}^{n} U_i$ for some contractible open subspaces U_i. Set $Y = \cup_{i \neq 1} U_i$ the subspace of X formed by the union of the last $n-1$ members of the good cover. Clearly, both $\mathrm{ct}(Y)$ and $\mathrm{ct}(U_1 \cap Y)$ are at most $n-1$. From the Mayer-Vietoris sequence applied to $X = U_1 \cup Y$, we obtain for each k exact sequences

$$H_k(U_1) \oplus H_k(Y) \to H_k(X) \to H_{k-1}(U_1 \cap Y).$$

Hence, by the inductive hypothesis we have $\beta_0(X) \leq n$ and for $k > 0$,

$$\beta_k(X) \leq \beta_k(Y) + \beta_{k-1}(U_1 \cap Y) \leq \binom{n-2}{k} + \binom{n-2}{k-1} = \binom{n-1}{k}.$$

□

2.3 The covering type of surfaces

As it is well known, the topological type of a closed surface is completely determined by a single integer number (the genus) and whether the surface is orientable or not. Hence, they constitute in this sense the simplest family of 2-dimensional spaces and it is natural to ask about their covering types after computing them for graphs. This problem, posed by Karoubi and Weibel in [61] admits via Lemma 2.2.2 the following equivalent formulation.

Problem. Given a closed surface S, compute the minimum number of vertices in a simplicial complex homotopy equivalent to S.

Here, by closed surface we mean as usual a compact 2-dimensional manifold without boundary, not necessarily orientable. The purpose of this section is to provide a complete solution to this problem.

Although it is only *almost* true, it may serve the reader to take throughout the section as a working hypothesis that minimal triangulations of closed surface are also minimal triangulations of the corresponding homotopy type. If one accepts that the optimal triangulation of the homotopy type of a closed surface is realized by a 2-dimensional complex (which we think is not obvious but it seems reasonable), the conjecture is plausible. Indeed, intuitively a 2-dimensional complex homotopy equivalent but not homeomorphic to a surface appears to contain "extra" 2-simplices. Also, Karoubi and Weibel verified this conjecture for the projective plane and the torus by exploiting the non-triviality of the cup product in the first cohomology groups of those spaces.

Thus, in view of Theorem 2.1.1, it is enough to prove (at least for non-exceptional cases) that a complex K homotopy equivalent to a given closed surface S has at least $\rho(S) = \left\lceil \frac{7+\sqrt{49-24\chi(S)}}{2} \right\rceil$ vertices. The departing point for achieving this is the remark that the proof of the lower bound part of Theorem 2.1.1 is valid under weaker hypotheses than the original.

Lemma 2.3.1. *Let K be a simplicial complex of dimension 2 such that every edge of K belongs to at least two 2-simplices and that $\chi(K) \leq 2$. Then, K has at least $\rho(K)$ vertices.*

Proof. As in the proof of the lower bound of Theorem 2.1.1, we start from the Euler characteristic formula for K,

$$\chi(K) = \alpha_0(K) - \alpha_1(K) + \alpha_2(K).$$

By hypothesis, the number of simplices in K satisfies $3\alpha_2(K) \geq 2\alpha_1(K)$. It is a matter of a simple computation to check that this leads to the inequality

$$6\chi(K) \geq 6\alpha_0(K) - \alpha_0(K)(\alpha_0(K) - 1).$$

If $\chi(K) \leq 0$, the minimum strictly positive integer that satisfies this inequality is precisely $\rho(K)$ and therefore $\alpha_0(K) \geq \rho(K)$. It remains to show that $\alpha_0(K) \geq \rho(K)$ when $\chi(K) = 1$ or $\chi(K) = 2$. If $\chi(K) = 1$, the minimum positive solutions for the above inequality are $a = 1$ and $a = 6 = \rho(K)$, but $a \neq 1$ since $\dim(K) = 2$. It follows that $\alpha_0(K) \geq \rho(K)$. If $\chi(K) = 2$, the inequality is satisfied by every positive integer. However, since every edge of K is the face of at least two 2-simplices, K cannot have 3 or less vertices. Hence $\alpha_0(K) \geq 4 = \rho(K)$. □

Roughly, the previous lemma states that if K is a complex of dimension 2 homology equivalent to a surface S with a triangulation at least as *dense* as a surface, then $\alpha_0(K) \geq \rho(K) = \rho(S)$. We systematize the use of the cohomology ring of surfaces by Karoubi and Weibel to obtain lower bounds for the covering type (notice that the cohomology ring of a space is also exploited in the context of L-S category, see Lemma 1.1.3). Namely, we identify a purely cohomological counterpart to the combinatorial property of a simplicial complex of being *densely* triangulated.

Definition 2.3.2. Let X be a topological space. We say that the cohomology ring $H^*(X, \mathbb{Z}_2)$ satisfies *property (A)* if for every non-trivial α in $H^1(K, \mathbb{Z}_2)$, there exists $\beta \in H^1(K, \mathbb{Z}_2)$ such that $\alpha \cup \beta$ is non-trivial in $H^2(K, \mathbb{Z}_2)$.

In what follows, we will work with reduced (co)homology and the coefficient ring for the (co)homology groups will be $\mathbb{Z}_2$. The reason for this is that it allows to handle the orientable and the non-orientable case at the same time.

Remark 2.3.3. Note that, by Poincaré Duality, the cohomology ring $H^*(S, \mathbb{Z}_2)$ of any closed surface S (orientable or non-orientable) satisfies property (A).

In the following result, we make explicit our assertion that property (A) constitutes a cohomological counterpart to being densely triangulated.

Lemma 2.3.4. *Let K be a simplicial complex of dimension 2 and suppose that $H^*(K, \mathbb{Z}_2)$ satisfies property (A) of Definition 2.3.2. If $\chi(K) \leq 2$, then K has at least $\rho(K)$ vertices.*

Proof. Collapse every free face of K to get a subcomplex L with no free faces. If every edge of L is incident to some 2-simplex, Lemma 2.3.1 applies and we are done.

We treat then the case in which L has maximal edges (i.e. there exists some 1-simplex that is not a face of any 2-simplex of L). We will show that it is always possible to replace L by a homotopy equivalent simplicial complex with less vertices and less maximal edges. Let $e = \{a, b\}$ be a maximal edge. Suppose that there is a simple path P joining a to b in $L \setminus e$ (the simplicial complex obtained from L by removing the edge e). Consider the quotient L/P, which is homotopy equivalent to a wedge sum of the form $T \vee S^1$, where the sphere S^1 is the image under the quotient of the edge e. Since on the other hand L/P is homotopy equivalent to L, we have $L \simeq T \vee S^1$. The cohomology ring of $T \vee S^1$ is isomorphic to the product $H^*(T, \mathbb{Z}_2) \times H^*(S^1, \mathbb{Z}_2)$ (the product in the category of graded algebras, where the operations are defined coordinate-wise), and therefore it does not satisfy property (A). Concretely, for the nonzero element $\gamma \in H^1(S^1, \mathbb{Z}_2)$, the class $\alpha = (0, \gamma) \in H^1(T \vee S^1, \mathbb{Z}_2)$ does not satisfy the requirements of property (A). Since the cohomology ring $H^*(L, \mathbb{Z}_2)$ satisfies property (A), this is a contradiction and hence, there is no path between a and b in $L \setminus e$.

Therefore the quotient L/e has a natural simplicial structure and is homotopy equivalent to L. Thus, if we replace L by L/e we obtain a simplicial complex homotopy equivalent to K with one fewer vertex and less maximal edges than L. By iterating this procedure we end up with a simplicial complex $\hat{K}$ of dimension 2 homotopy equivalent to K with no more vertices than K and such that each edge of $\hat{K}$ belongs to at least two 2-simplices. By Lemma 2.3.1 applied to $\hat{K}$, we deduce that $\hat{K}$ (and in consequence K) has at least $\rho(K)$ vertices. □

In particular, as a corollary to Lemma 2.3.4 we see that minimal triangulations of a (non-exceptional) closed surface S optimize the number of vertices within complexes of dimension 2 of the homotopy type of S (or, more generally, within those with $\mathbb{Z}_2$-homology isomorphic to that of S). We will need a different argument in order to extend this to the general case.

Suppose that K is a simplicial complex, not necessarily of dimension 2, of the homotopy type of a closed surface S. We propose to compare K to its 2-skeleton $K^{(2)}$ to somehow reduce the general problem to the 2-dimensional case. Obviously, the inclusion $K^{(2)} \hookrightarrow K$ induces an isomorphism $i_* : H_n\left(K^{(2)}, \mathbb{Z}_2\right) \to H_n(K, \mathbb{Z}_2)$ for $n < 2$ but it is in general only an epimorphism for $n = 2$. Hence, informally speaking, in order to obtain a complex with homology isomorphic to that of S we need to remove the "extra" homology classes present in $K^{(2)}$. To accomplish this, we will develop a homological simplification method controlled by property (A) to find a subcomplex $Z \leq K^{(2)}$ with the (co)homology of S, consisting of removing some carefully chosen 2-simplices of $K^{(2)}$. Here, it is important to remark that it is not enough to simply remove some 2-simplices from $K^{(2)}$ to kill homology classes until there remains only one. We also need to ensure that the triangulation of the resulting subcomplex is sufficiently dense, that is to say, that at every step of the simplification the involved subcomplexes satisfy property (A). The next lemma describes how property (A) descends to subcomplexes and it is the key to control the homological simplification process.

Lemma 2.3.5. *Let K be a simplicial complex and let $L \leq K$ be a subcomplex such that the inclusion $i : L \hookrightarrow K$ induces isomorphisms $i_* : H_n(L, \mathbb{Z}_2) \to H_n(K, \mathbb{Z}_2)$ for $n < 2$ and an epimorphism for $n = 2$. If the cohomology ring $H^*(K, \mathbb{Z}_2)$ satisfies property (A), then the cohomology ring $H^*(L, \mathbb{Z}_2)$ also satisfies property (A).*

Proof. Identify $H^n(K,\mathbb{Z}_2)$ with $\mathrm{Hom}(H_n(K),\mathbb{Z}_2)$ by the universal coefficient theorem for cohomology. By assumption $H^*(K,\mathbb{Z}_2)$ satisfies property (A) and $i^*: H^n(K,\mathbb{Z}_2)\to H^n(L,\mathbb{Z}_2)$ is an isomorphism for $n<2$ and a monomorphism for $n=2$. The claim follows from the naturality of the cup product. □

Given a simplicial complex K, we will denote by $(C_*(K,\mathbb{Z}_2),d_*)$ its simplicial chain complex with coefficients in $\mathbb{Z}_2$. If $L\le K$ is a subcomplex, we denote by $\dim H_2(L,\mathbb{Z}_2)$ the dimension of $H_2(L,\mathbb{Z}_2)$ as a vector space over $\mathbb{Z}_2$.

Proposition 2.3.6. *Let K be a simplicial complex and let $L\le K^{(2)}$ be a subcomplex. Suppose that $H^*(K,\mathbb{Z}_2)$ has property (A) and that the inclusion $L\hookrightarrow K$ induces isomorphisms $i_*: H_n(L,\mathbb{Z}_2)\to H_n(K,\mathbb{Z}_2)$ for $n<2$ and an epimorphism for $n=2$. If $\dim H_2(L,\mathbb{Z}_2)>\dim H_2(K,\mathbb{Z}_2)$, there is a 2-simplex $\sigma\in L$ such that the inclusion $j: L\setminus\sigma\hookrightarrow K$ also induces isomorphisms $j_*: H_n(L\setminus\sigma,\mathbb{Z}_2)\to H_n(K,\mathbb{Z}_2)$ for $n<2$ and an epimorphism for $n=2$. Moreover, $\dim H_2(L\setminus\sigma,\mathbb{Z}_2)=\dim H_2(L,\mathbb{Z}_2)-1$.*

Proof. Since by hypothesis $\dim H_2(L,\mathbb{Z}_2)>\dim H_2(K,\mathbb{Z}_2)$ there is a non-trivial class B in the kernel of the inclusion induced map $i_*: H_2(L,\mathbb{Z}_2)\to H_2(K,\mathbb{Z}_2)$. Let σ be a 2-simplex of L in the support of B. The topological boundary $\partial\sigma$ viewed as a chain in $C_1(L\setminus\sigma,\mathbb{Z}_2)$ is the boundary of the 2-chain $B-\sigma$. Hence the inclusion induces the zero morphism $H_1(\partial\sigma,\mathbb{Z}_2)\to H_1(L\setminus\sigma,\mathbb{Z}_2)$. It follows that the inclusion $L\setminus\sigma\hookrightarrow L$ induces isomorphisms $H_n(L\setminus\sigma,\mathbb{Z}_2)\to H_n(L,\mathbb{Z}_2)$ for $n<2$. It remains to verify the surjectivity of the map $j_*: H_2(L\setminus\sigma,\mathbb{Z}_2)\to H_2(K,\mathbb{Z}_2)$, where j is the inclusion $j: L\setminus\sigma\hookrightarrow K$. Let $[Z]$ be a class in $H_2(K,\mathbb{Z}_2)$. By hypothesis, there is some class $C\in H_2(L,\mathbb{Z}_2)$ such that $i_*[C]=[Z]$. If σ does not belong to the support of C, when viewed as a class in $H_2(M\setminus\sigma,\mathbb{Z}_2)$ we have $j_*[C]=[Z]$. In the other case, consider the 2-chain $C+B$. Since the coefficients are taken in $\mathbb{Z}_2$, this chain is a well defined 2-cycle in $L\setminus\sigma$ and $j_*[C+B]=i_*[C]+i_*[B]=i_*[C]=[Z]$. Hence, in any case $j_*: H_2(L\setminus\sigma,\mathbb{Z}_2)\to H_2(K,\mathbb{Z}_2)$ is an epimorphism. The fact that $\dim H_2(L\setminus\sigma,\mathbb{Z}_2)=\dim H_2(L,\mathbb{Z}_2)-1$ follows immediately from the Euler characteristic formula, since $\chi(L\setminus\sigma)=\chi(L)-1$. □

Now, given a simplicial complex K homotopy equivalent to a closed surface S its 2-skeleton $K^{(2)}$ satisfies property (A) by Lemma 2.3.5. Hence, by applying iteratively Proposition 2.3.6 to $L=K^{(2)}$, we end up with a subcomplex $T\le K^{(2)}$ such that the inclusion $i: T\hookrightarrow K$ induces an isomorphism in all homology groups (note that here we implicitly use that K is finite). Thus, the reduction from the general case to dimension 2 is complete.

Theorem 2.3.7. *Let K be a simplicial complex homotopy equivalent to a surface S. Then K has at least $\rho(S)$ vertices. In particular, if $S\neq M_2,N_2,N_3$ then $\mathrm{ct}(S)=\lambda(S)$.*

Proof. By Proposition 2.3.6 and the subsequent paragraph, there exists a subcomplex $T\le K^{(2)}$ with $\chi(T)=\chi(S)$ and such that its cohomology ring satisfies property (A). Since the number of vertices of T is less than or equal to the number of vertices of K, by Lemma 2.3.4 K has at least $\rho(S)$ vertices. □

We perform the computation of the covering type of the exceptional surfaces in the next section, since it requires a different set of tools.

2.4 The covering type of surfaces: the exceptional cases

In this section we analyze the covering type of the exceptional surfaces M_2, N_2 and N_3. By Theorem 2.3.7, for $S = M_2, N_2, N_3$ the covering type of S lies between $\lambda(S) - 1$ and $\lambda(S)$. We complete the computation of their covering types by showing that $\mathrm{ct}(N_2) = \lambda(N_2)$, $\mathrm{ct}(N_3) = \lambda(N_3)$ and exhibiting a simplicial complex homotopy equivalent to M_2 with $\lambda(M_2) - 1$ vertices.

To gain some understanding of the difficulties that arise in this case, let S be one of the three exceptional surfaces and suppose that K is a simplicial complex homotopy equivalent to S realizing the covering type, that is, such that K has $\mathrm{ct}(S)$ vertices. By applying Theorem 2.3.7 and Proposition 2.3.6, we obtain a subcomplex L of K of dimension 2 with the (co)homology of S (and thus, its cohomology ring satisfies property (A)) with at least $\rho(S) = \lambda(S) - 1$ vertices. We would like to find out whether it is possible for L to have less than $\lambda(S)$ vertices. To accomplish this, we analyze the interaction of the number of vertices of L with its number of edges and 2-simplices. In first place, by collapsing the free faces of L and proceeding as in Corollary 2.3.4 we can assume further that every edge of L is the face of at least two 2-simplices. As before, this implies that $3\alpha_2(L) \geq 2\alpha_1(L)$ and by the Euler characteristic formula for L we have the inequality

$$3(\alpha_0(L) - \chi(S)) \leq \alpha_1(L).$$

Let us plug in some numbers in the above equation. For example, if $S = N_2$ and L has $7 = \lambda(S) - 1$ vertices we obtain the inequality $\alpha_1(L) \geq 21$ using that $\chi(N_2) = 0$. Since on the other hand $\alpha_1(L) \leq \binom{\alpha_0(L)}{2} = 21$, it turns out that the number of edges and 2-simplices is completely determined, resulting in $\alpha_1(L) = 21$ and $\alpha_2(L) = 14$ respectively. Observe that $3\alpha_2(L) = 2\alpha_1(L)$ in this case, which implies that every edge of L is the face of *exactly* two 2-simplices. Repeating the experiment for $S = N_3$, we obtain that $\alpha_1(L)$ is between 27 and 28. If we proceed to analyze the case when $\alpha_1(L) = 27$, we find analogously as in the previous example that $3\alpha_2(L) = 2\alpha_1(L)$ and thus, that every edge of L belongs to exactly two 2-simplices. Intuitively, since this is one of the principal conditions satisfied by triangulations of (closed) surfaces, such complexes should be close to being a surface. This is formalized in the next result and helps to deal with some of the cases we need to analyze for the exceptional surfaces.

Before the result, we recall some definitions. A simplicial complex K of dimension 2 is *strongly connected* if for every pair of 2-simplices $\sigma, \tau \in K$ (with $\sigma \neq \tau$) there exists a (finite) sequence $\sigma = \sigma_1, \sigma_2, \ldots, \sigma_r = \tau$ of 2-simplices such that for each i, σ_i and σ_{i+1} intersect in an edge. We say that K is a *pseudosurface* (without boundary) if K is strongly connected (in particular, it is connected) and each of its edges is the face of exactly two 2-simplices.

Proposition 2.4.1. *Let K be a simplicial complex of dimension 2 such that each edge of K is the face of exactly two 2-simplices and let S be a closed surface. Suppose that there is a continuous map $K \to S$ inducing isomorphisms in all homology groups. Then K is homeomorphic to S.*

Proof. Each strongly connected component C of K is a pseudosurface without boundary, so that $H^2(C, \mathbb{Z}_2) = \mathbb{Z}_2$. Since two different strongly connected components of K may intersect only at vertices, the dimension of the vector space $H^2(K, \mathbb{Z}_2)$ coincides with the number of its strongly connected components. Hence, K is strongly connected and therefore a pseudosurface. We show next that K is indeed a surface. Since K is a pseudosurface, only the vertices of K may be singular points (i.e. points that do not have a neighborhood in K homeomorphic to a disk).

Since every edge of K is the face of exactly two 2-simplices, the link of every vertex is a 2-regular graph and hence, homeomorphic to a disjoint union of some copies of S^1.

Suppose that there is a vertex v of K such that its link has $m > 1$ connected components. We claim that in this case K is homotopy equivalent to the wedge sum of a suitable 2-complex L with $m-1$ spheres of dimension 1. To see this, consider the star $\mathrm{st}_K(v)$ of v. Note that $\mathrm{st}_K(v)$ is the wedge sum (at v) of m subcomplexes $K_0, \ldots, K_{m-1} \leq \mathrm{st}_K(v)$. Take a star-shaped graph T with $m-1$ leaves, that is, a 1-dimensional simplicial complex formed by a root vertex w_0 and $m-1$ vertices $w_1, \ldots w_{m-1}$ with an edge connecting w_0 to each w_i ($i \geq 1$). Note that $\mathrm{st}_K(v)$ is homotopy equivalent to a complex N obtained from the graph T by attaching each K_i to T at the vertex w_i, since by collapsing the tree T to a point we obtain $\mathrm{st}_K(v)$. Construct a complex K' by replacing in K the subcomplex $\mathrm{st}_K(v)$ by N in the natural way. Now, since K is strongly connected, there is a path in $K' \setminus T$ connecting w_0 to each w_i. It follows that for each edge between w_0 and w_i the attaching map $\phi_i : S^0 \to K'$ of the edge (which takes the values w_0 and w_i) is homotopic to the constant map $c_{w_0} : S^0 \to K'$, and therefore $K' \simeq L \vee_{i=1}^{m-1} S^1$ where L is the subcomplex $K' \setminus T$ of K'. Since the map $f : K \to S$ induces an isomorphism in cohomology, $H^*(K, \mathbb{Z}_2)$ satisfies property (A). Similarly as in the proof of Lemma 2.3.4, the cohomology ring of $H^*(L \vee_{i=1}^{m-1} S^1, \mathbb{Z}_2)$ is isomorphic to the product ring $H^*(L, \mathbb{Z}_2) \times \prod_{i=1}^{m-1} H^*(S^1, \mathbb{Z}_2)$ and therefore it does not satisfy property (A). It follows that the link of every vertex of K is homeomorphic to only one copy of S^1, which in turn implies that K is a surface. By the classification of closed surfaces, K is homeomorphic to S. □

We are now ready to compute the covering type of N_2 and N_3.

Proposition 2.4.2. *Let $S = N_2$ or N_3. Then $\mathrm{ct}(S) = \lambda(S)$.*

Proof. Let K be a simplicial complex on $\mathrm{ct}(S)$ vertices such that there is a homotopy equivalence $f : K \to S$. As described in the discussion opening this section, there is a subcomplex $L \leq K^{(2)}$ such that the map $L \hookrightarrow K \xrightarrow{f} S$ induces an isomorphism in homology and such that every edge of L is the face of at least two 2-simplices. Recall that this implies that $3\alpha_2(L) \geq 2\alpha_1(L)$ and by the Euler characteristic formula for L we have

$$3(\alpha_0(L) - \chi(S)) \leq \alpha_1(L) \leq \binom{\alpha_0(L)}{2}.$$

For $S = N_2$, if it were $\alpha_0(L) = 7 = \lambda(N_2) - 1$, as we have seen it would imply $\alpha_1(L) = 21$ and $\alpha_2(L) = 14$. Hence, $3\alpha_2(L)$ would equal $2\alpha_1(L)$, which means that every edge of L is the face of exactly two 2-simplices. Then, by Proposition 2.4.1 L is homeomorphic to N_2, contradicting Theorem 2.1.1. Hence $\mathrm{ct}(N_2) = 8$.

Consider now the case $S = N_3$. If L has $\lambda(N_3) - 1 = 8$ vertices, using that $\chi(N_3) = -1$ combined with the inequality above leads to $\alpha_1(L) = 27$ or $\alpha_1(L) = 28$. As before, if $\alpha_1(L) = 27$ then L should be homeomorphic to N_3 and this is impossible by Theorem 2.1.1. If $\alpha_1(L) = 28$, since $\chi(N_3) = -1$, it is $\alpha_2(L) = 19$. Since $3\alpha_2(L) = 2\alpha_1(L) + 1$ in this case, every edge of L is the face of two 2-simplices except for one edge that is contained in three 2-simplices. Let $v \in L$ be a vertex of this edge. The link of v is a graph in which exactly one vertex has degree 3 and every other vertex is of degree 2. Since the sum of the degrees of all vertices should be even, this is a contradiction. Hence $\mathrm{ct}(N_3) = \lambda(N_3) = 9$ and the proof is complete. □

When repeating this analysis for the double torus $S = M_2$, more subcases appear. Keeping the notations from Proposition 2.4.2, if $\alpha_0(L) = \lambda_(M_2)-1 = 9$, we would have that L has between 33 and 36 edges and, correspondingly, between 22 and 25 2-simplices. Some of these cases may be handled in the same way as in the proof of Proposition 2.4.2. For example, if $\alpha_1(L) = 33$, L has 22 2-simplices and hence every edge of L is in exactly two 2-simplices which contradicts the Jungerman and Ringel Theorem 2.1.1 in view of Proposition 2.4.1, while the subcase $\alpha_1(L) = 34$ is discarded by a parity argument. However, after devoting a considerable effort to trying to show the impossibility of constructing a simplicial complex that satisfies the hypotheses in every other subcase, we started considering the possibility that M_2 was also exceptional with respect to the covering type. Much to our surprise, it was not too difficult to construct a simplicial complex on $\lambda(M_2) - 1 = 9$ vertices homotopy equivalent to M_2, departing from a minimal triangulation of M_2 described in [58]. In what follows, we give a rigorous construction of such an example, which is illustrated in Figure 2.1 and developed in the next result.

Lemma 2.4.3. *Let K be a simplicial complex of dimension 2. Suppose there are two vertices $v, v' \in K$ not connected by an edge and with disjoint links. Suppose further that there are vertices $w \in \mathrm{lk}_K(v)$, $w' \in \mathrm{lk}_K(v')$ connected by an edge in K. Form the quotient complex L by identifying v and v' into a vertex $[v]$, and attach a 2-simplex $\sigma = \{[v], w, w'\}$ to L to get a complex K'. Then K' is a simplicial complex homotopy equivalent to K.*

Proof. First notice that the quotient complex $L = K/(v \sim v')$ inherits a simplicial structure. Since v and v' do not form an edge in K, the edges of the 1-skeleton of L connect different vertices.

Moreover, since the links of v and v' are disjoint in K, there is at most one edge between two given vertices in the 1-skeleton of L. Finally, since L is obtained from K by identifying a pair of vertices with disjoint links, the 2-simplices of L may intersect only on a vertex or an edge. Consider next the CW-complex M obtained from K by attaching an edge between vertices v and v' and a square (2-cell) with vertices v, v', w and w'. Note that since K is a deformation retract of M and the inclusion $M \hookrightarrow K'$ is an edge collapse, M is homotopy equivalent to both K and K'. This shows that K is homotopy equivalent to K'. □

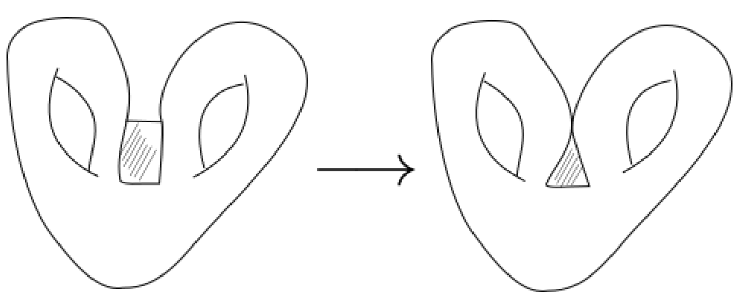

Figure 2.1: Illustration of the construction in Lemma 2.4.3.

Proposition 2.4.4. *There is a simplicial complex of dimension 2 with 9 vertices homotopy equivalent to M_2. In particular,* $\mathrm{ct}(M_2) = \rho(M_2) = 9$.

Proof. Consider the simplicial complex K which realizes the minimal triangulation of M_2 described in [58] (see also [61]). The 1-skeleton of the simplicial complex K is a graph on 10 vertices such that:

- There are two vertices v, v' of degree 4 that do not form an edge and with disjoint links.
- The subgraph induced by the remaining 8 vertices is the complete graph K_8.

Choose vertices $w \in \mathrm{lk}_K(v)$, $w' \in \mathrm{lk}_K(v')$. By Lemma 2.4.3, the simplicial complex $K/(v \sim v') \cup \{[v], w, w'\}$ is homotopy equivalent to K and hence to M_2. □

We collect the results obtained so far in the following theorem.

Theorem 2.4.5. *The covering type of any closed surface S coincides with the number of vertices in a minimal triangulation of S with the exception of M_2, in which case it is one less than the number of vertices in a minimal triangulation.*

Proof. It follows from Theorem 2.3.7 (case $S \neq M_2, N_2, N_3$), Proposition 2.4.2 (cases $S = N_2$ and $S = N_3$) and Proposition 2.4.4 (case $S = M_2$). Concretely $\mathrm{ct}(S) = \lambda(S)$ for the non-exceptional cases, $\mathrm{ct}(M_2) = 9$, $\mathrm{ct}(N_2) = 8$ and $\mathrm{ct}(N_3) = 9$. □

Note that, in fact, $\mathrm{ct}(S) = \rho(S)$ in all cases including M_2 (where we have $\mathrm{ct}(M_2) = \rho(M_2) = 9$ but $\lambda(M_2) = 10$), except for N_2 ($\mathrm{ct}(N_2) = \lambda(N_2) = 8$ but $\rho(N_2) = 7$) and N_3 ($\mathrm{ct}(N_3) = \lambda(N_3) = 9$ but $\rho(N_3) = 8$).

Chapter 3

Systolic geometry

The *systole* of a metric space X is the length of a shortest non contractible loop in X. Systolic geometric is concerned with the study of the systole and its multiple surprising interactions with the global geometry and topology of the underlying space. Such connections often come in the form of inequalities that link the systole to geometric measures of size (notably volume and diameter) or even purely homological invariants such as the Betti numbers. The work on this area could be seen as a part of a broader circle of results, concepts and ideas that loosely aim to relate the complexity of the topology of a space to relatively coarse geometrical measures of its size, such as volume, diameter or filling radius, including curvature-free inequalities. As a paradigmatic example, we cite the Cheeger's finiteness theorem [28], which states that for positive real numbers v, D, κ, there are at most finitely many diffeomorphism types of closed riemannian manifolds of dimension n with volume bounded from below by v, diameter bounded from above by D and sectional curvature between $-\kappa$ and κ. Thus, the systole of a space may be thought of as an invariant of a mixed nature, in the sense that it geometrically detects a manifestation of the non-trivial underlying topology.

The concept of systole made its first apparition in riemannian geometry in an unpublished result of Loewner dating from 1949. Loewner showed that the square of the systole of a riemannian torus (for *any* smooth riemannian metric) bounds from below the area modulo an optimal constant. Through the use of essentially the same techniques, the inequality was extended to the projective plane by his student Pu [74] and subsequently, to all orientable, non-simply connected closed surfaces in the early sixties, independently, by Accola [2] and Blatter [14] (we include below a detailed proof of Pu's inequality in Theorem 3.2.1). These results revealed a deep connection between the geometrical size of a surface (the area) and its systole, which as we remarked, depends both on the geometry and the topology. For many years, systolic inequalities remained confined within the realm of surfaces. The proofs for all the mentioned results relied ultimately on the uniformization theorem, and thus were not directly extendable to higher dimensions. In the 1983 article *Filling Riemannian manifolds* Gromov overcame the difficulties and established a systolic inequality for closed manifolds of arbitrary dimension that satisfy a purely topological condition. The proofs and the results from that article greatly clarified the panorama of systolic phenomena. Most relevantly, Gromov identified the topological substrate underlying systolic inequalities and related these inequalities to generalizations of the classical Federer-Fleming isoperimetric inequality. Besides, in the course of the proof Gromov invented new interesting metric invariants for riemannian manifolds, which he used as intermediate quantities to compare the systole with the volume of a manifold.

After Gromov's paper, the *systolic volume* of a manifold, or more generally of a polyhedron, became one of the central objects of study in systolic geometry. The systolic volume $\sigma(X)$ of a polyhedron X of dimension n is defined by Gromov as

$$\sigma(X) := \inf_g \frac{\mathrm{Vol}(X,g)}{\mathrm{sys}(X,g)^n},$$

where the infimum is over the piecewise riemannian smooth metrics over X. This homeomorphism invariant of a space X identifies optimal riemannian metrics over X, in the sense that it measures the amount of volume needed to form X, normalized by the condition that the systole equals 1. The exponent on the systole guarantees that the systolic volume is invariant under rescalings of the metric. In this language, the most general form of Gromov's inequality reads as the lower bound

$$\sigma(X) \geq C_n$$

for an *essential* polyhedron X of dimension n, where C_n is a positive constant that depends only on the dimension. Thus, the systolic volume is not trivial (i.e. strictly positive) for a considerable class of spaces. For such spaces, a natural problem is to find the value of the systolic volume, together with the optimal metrics. This would allow to somehow sort the spaces according to a numerical measure of its complexity, but more importantly it would help identify optimal geometrical models for certain topological spaces. We remark however that this problem is very hard, to the point that the precise value of the systolic volume is known only for three essential manifolds, and all of them are 2-dimensional. A more achievable goal is the estimation of the systolic volume for some classes of spaces. Advances in this direction also have interesting implications. For example, resuming the case of surfaces, the precise growth of the functional σ with respect to the genus of the surface has been determined by the works of Gromov [44] and Buser and Sarnak [27]. As a consequence, certain arithmetic surfaces (which are in particular hyperbolic from the riemannian geometry point of view) constructed by Buser and Sarnak in the cited work are coarsely optimal with respect to the systolic volume.

Another class of spaces which is of particular interest to us is that of 2-dimensional polyhedra. The reason for this is the classical correspondence between homotopy types of (compact) 2-polyhedra and (finite) presentations of groups. Via this correspondence, given by the fundamental group functor in one direction and the presentation complex (see §1.4.1) on the other, it is possible to extend the systolic invariants to groups. Specifically, for a finitely presentable group G, the systolic area $\sigma(G)$ of G, introduced by Gromov in [46], is defined as

$$\sigma(G) := \inf_{\pi_1(X)=G} \sigma(X),$$

where the infimum is over the *triangulations* X of G, that is, 2-polyhedra with fundamental group isomorphic to G. A new whole range of questions opens up around this invariant, that mainly inquire about the optimal topological models (triangulations) for certain groups. To cite a few that partially motivated our work: What is the exact relation between the systolic area of surface groups and the systolic volume of the surfaces (in other words, are surfaces the most efficient model for their fundamental groups)? What is the behavior of the systolic area under free product of groups? That is to say, is the canonical wedge sum the optimal way to build a topological model for a free product of groups?

In this chapter, we start with a historical survey of some fundamental results in systolic geometry, which is closed by a discussion of the arguments involved in the proof of Gromov's systolic inequality. The second section is devoted to systolic geometry in dimension 2. We present there some of the best known estimates for both the systolic area of surfaces and groups. In section 3 we prove a universal lower bound for the systolic area of a wide class of groups that includes the surface groups, which grew out of our effort to understand in more depth the relation between the systolic area of surfaces and the systolic area of its fundamental groups. This estimate is obtained by extending a systolic inequality by Guth; the new inequality also generalizes a systolic inequality of Burago and Hebda for surfaces to certain 2-dimensional polyhedra.

In the final section of the chapter we study the systolic area of a free product of groups. We give a construction, based on a similar one due to Stallings, to somehow split a polyhedron with fundamental group isomorphic to a free product of two groups in two subpolyhedra, each one responsible, roughly, for one of the free factors. Even though we are not able to establish new estimates for the systolic area of a free product of groups in terms of the systolic area of its free factors, we discuss a possible strategy involving the mentioned construction that could be used for that purpose.

3.1 An overview of systolic geometry

In this section we propose to survey some of the fundamental concepts and results in systolic geometry.

As we have mentioned in the introduction of the chapter, the first result in the area was discovered by Loewner, who proved in 1949 the following result.

Theorem 3.1.1. *Let $(\mathbb{T}^2, g)$ be a riemannian 2-torus. Then,*

$$\operatorname{sys}(\mathbb{T}^2, g)^2 \leq \frac{2}{\sqrt{3}} \operatorname{Area}(\mathbb{T}^2, g).$$

Moreover, the constant $\frac{2}{\sqrt{3}}$ is optimal and it is realized by a flat metric.

Loewner's inequality is interesting at least for three different reasons. In first place, it is valid for any riemannian metric on the torus, and thus it reveals a non-trivial restriction over the metrizations of the topology of a torus. In second place, it is optimal, and hence it identifies a critical point in the space of riemannian metrics on the torus with respect to a certain functional. Lastly, it may be understood as a reverse isoperimetric inequality, in the sense that a function of the perimeter (length) of a non-contractible loop is shown to bound from below the volume of the space.

The proof of Loewner's inequality goes roughly as follows. By the uniformization theorem, a given riemannian torus $(\mathbb{T}^2, g)$ is *conformally* homeomorphic to a flat torus $(\mathbb{T}^2, g_0)$, that is, there exists a positive function $f : \mathbb{T}^2 \to \mathbb{R}$ such that $g = f g_0$. The first step consists on forming the average $\bar{f}$ of f over the isometry group of the flat torus $(\mathbb{T}^2, g_0)$ to obtain a new riemannian metric $\bar{f} g_0$. The key of the argument resides in that, as it is not difficult to check, this averaging process improves (decreases) the ratio

$$\frac{\operatorname{Area}(\mathbb{T}^2, \bar{f} g_0)}{\operatorname{sys}(\mathbb{T}^2, \bar{f} g_0)^2}$$

with respect to the ratio computed with the original riemannian metric $g = fg_0$. On the other hand, since the isometry group of a flat torus is transitive, the average $\bar{f}$ is a constant function and hence Loewner's inequality reduces to the flat case, which can be solved by hand. We give a formal proof of Pu's optimal systolic inequality [74] for the projective plane along these lines below (see Theorem 3.2.1).

About a decade a later, Accola [2] and Blatter [14] independently generalized these inequalities to all orientable closed surfaces different from the sphere, by showing the existence of a non-optimal constant $C = C(\gamma) > 0$ such that

$$\mathrm{sys}(S, g)^2 \leq C(\gamma)\,\mathrm{Area}(S, g),$$

for each surface S of genus γ. In the early seventies, Berger started a more systematic study of the subject [13, 12], posing among other things two natural problems about higher dimensional analogues of these inequalities, both involving higher dimensional versions of the systole and asking for extensions of the inequalities to any dimension. In the 1983 article *Filling Riemannian manifolds* [44] Gromov established a systolic inequality for (certain) manifolds of any dimension through new and radically different methods from the conformal techniques that had been employed in dimension 2 so far.

Theorem 3.1.2. *(Gromov's systolic inequality) Let (M, g) be an essential riemannian manifold of dimension n. Then, there exists a universal constant $C = C_n > 0$ depending only on the dimension n such that*

$$\mathrm{sys}(M, g)^n \leq C_n\,\mathrm{Vol}(M, g).$$

We discuss at some length the hypothesis of *essentialness.* One of the key insights of Gromov to prove his systolic inequality was realizing the rôle of this purely topological condition in systolic phenomena. A closed manifold M of dimension n is *essential* if the image of the fundamental class $[M]$ is non-trivial under the canonical map $H_n(M) \to H_n(K(\pi_1(M), 1))$, where $K(\pi_1(M), 1)$ is an Eilenberg-MacLane space for $\pi_1(M)$. Here and elsewhere, we understand that the coefficient ring for homology groups is $\mathbb{Z}$ if the manifold is orientable and $\mathbb{Z}_2$ in the opposite case. The essential manifolds can be regarded as a generalization of aspherical manifolds which includes, among others, the projective real spaces $\mathbb{R}P^n$ of any dimension and the lens spaces. Already some years before proving his systolic inequality, Gromov was aware that for such inequality to hold for a manifold M, its 1-dimensional topology should detect in some way the fundamental class of M. [1] Let us explain heuristically the apparition of the fundamental class of the manifold by using an analogy from hyperbolic geometry.

The *simplicial volume* (also called *Gromov norm*) of a closed, orientable manifold M is a homotopy invariant introduced in [43] that measures how difficult is to triangulate the fundamental class $[M]$. Concretely, the simplicial volume $\|M\|$ of M is the infimum of

$$\|C\|_1 := \sum_i |a_i|,$$

taken over all those (singular) cycles $C = \sum_i a_i\sigma_i$ in $C_n(M, \mathbb{Z})$ that represent $[M]$. Gromov proved that, for a hyperbolic closed manifold M, the simplicial volume $\|M\|$ is proportional to

[1]"... il faut, on le sent, que, comme dans le cas de T^2, ou $\mathbb{R}P^2$, l'homologie de dimension 1 de M engendre au sens anneau, toute la topologie de M, ou à tout le moins la classe fondamentale de M." Excerpt taken from [42, p. 49].

the volume of the manifold (this is an important step in the Gromov-Thurston proof of Mostow rigidity theorem). So, to put it crudely, the size of the fundamental class of a riemannian manifold acts as an avatar for the volume. In turn, if the 1-dimensional topology of a manifold M generates its fundamental class, since at a scale smaller than the systole the 1-dimensional topology of M is trivial, one may speculate that the systole is controlled by the size of $[M]$.

The proof of Theorem 3.1.2 involves an extensive use of techniques drawn from Geometric Measure Theory, particularly from the area of isoperimetric inequalities. More precisely, Gromov created a new metric invariant of riemannian manifolds called *filling radius* that links the systolic inequality to a generalization of the Federer-Fleming isoperimetric inequality. The filling radius is used by Gromov as an intermediate quantity between the systole and the volume of a manifold. The proof of Gromov's systolic inequality is obtained as a consequence of two estimates, one that involves the systole and the filling radius of a closed essential riemannian manifold and an upper bound for the filling radius of a manifold in terms of its volume.

To motivate the definition of filling radius, consider a codimension 1 closed submanifold M of an euclidean space $\mathbb{R}^n$ (for example, we can take $M = S^{n-1}$). Then, M determines an *interior hole*, that is, the bounded connected component of $\mathbb{R}^n \setminus M$. Provisionally, define the filling radius as the greatest radius of an euclidean ball completely contained in the interior hole of M. For example, in the case that M is a sphere of radius R, this filling radius coincides with R. To extend this notion to a submanifold M of general codimension, a different way to measure the radius of the hole generated by M is needed. It turns out that the technically correct tool to achieve that is homology. Namely, for a submanifold $M \subseteq \mathbb{R}^N$ of dimension $n < N$, the (relative) filling radius $\operatorname{FillRad}(M \subseteq \mathbb{R}^N)$ is defined as the smallest radius $r > 0$ for which the inclusion M is a *boundary* in its r-neighborhood $U_r(M)$, that is, such that the inclusion induced morphism $H_n(M) \to H_n(U_r(M))$ is trivial. Gromov transformed this a priori extrinsically defined invariant of a riemannian manifold into an intrinsic one by considering the isometric (distance-preserving) Kuratowski embedding into an infinite dimensional Banach space. A closed riemannian manifold M (or more generally, a compact metric space) admits a canonical isometric embedding K in the space $L^\infty(M)$ of bounded Borel maps $M \to \mathbb{R}$ endowed with the supremum norm, which sends a point $x \in M$ to the distance function $d_x : M \to \mathbb{R}$ from x. Now, the filling radius of a closed riemannian manifold M is defined relative to this embedding, that is,

$$\operatorname{FillRad}(M) := \inf\{r > 0 \,:\, H_n(M) \to H_n(U_r(K(M))) \text{ is trivial}\},$$

As before, the coefficient ring for homology groups is understood to be $\mathbb{Z}$ if the manifold is orientable and $\mathbb{Z}_2$ otherwise.

The following result links the filling radius to the systole of an essential manifold.

Lemma 3.1.3. *([44, Lemma 1.2.B]) Let (M, g) be a closed essential riemannian n-manifold. Then, $\frac{1}{6}\operatorname{sys}(M, g) \leq \operatorname{FillRad}(M)$.*

Proof. Let us assume first that M is aspherical instead of only essential. Suppose by contradiction that $\operatorname{FillRad}(M) < \frac{1}{6}\operatorname{sys}(M, g)$, so that there is a chain C with support in $U_R(K(M)) \subseteq L^\infty(M)$ which bounds $K(M)$ for some $R < \frac{1}{6}\operatorname{sys}(M, g)$. Triangulate C in such a way that $K(M) = \partial C$ is a subpolyhedron and every simplex of the subdivision has diameter at most δ, for some small $\delta > 0$ to be fixed later. The strategy of the proof consists of constructing a map $f : C \to K(M)$ extending the identity $K(M) \to K(M)$, which leads to a contradiction because $[M] \neq 0$ in

$H_n(M)$. Since M is aspherical, we only need to show that the identity can be extended to the 2-skeleton of C. We proceed, as usual, skeleton by skeleton. For a vertex v of C not in $K(M)$, define $f(v)$ to be any point $x \in K(M)$ such that $d_\infty(v,x) < R$. If v and w form an edge e, by the previous step and the triangular inequality, $d_\infty(f(v), f(w)) < 2R + \delta$. Hence, there exists a geodesic segment in $K(M)$ of length less than $2R + \delta$ joining $f(v)$ to $f(w)$. Extend f to e by sending the edge homeomorphically to this segment. Now, if σ is a 2-simplex in C, the loop $f(\partial\sigma)$ has length less than $6R + 3\delta$ in $K(M)$. Since $6R < \operatorname{sys}(M, g)$, we can choose $\delta > 0$ sufficiently small so that $6R + 3\delta < \operatorname{sys}(M, g)$. In that case, the loop $f(\partial\sigma)$ is null-homotopic in $K(M)$ and so we can extend f to σ by using any homotopy $\sigma \to K(M)$ that contracts $f(\partial\sigma)$ to a point. The extensions of f to the k-skeletons of C for $k > 2$ can be performed analogously, observing that since M is aspherical, the restrictions of f to the boundary of a k-simplex are null-homotopic.

For the general case, notice that if M is an essential manifold, the natural map g from M to an Eilenberg-MacLane space $X = K(\pi_1(M), 1)$ induces a non-trivial map in homology $H_n(M) \to H_n(X)$. The strategy of the proof is completely analogous to the aspherical case, with the difference that now we have to extend the map $K(M) \to X$ instead of the identity $K(M) \to K(M)$. More concretely, if $K(M) \subseteq L^\infty(M)$ is the boundary of an $n+1$-chain C in a R-neighborhood for $R < \frac{1}{6}\operatorname{sys}(M, g)$, it is possible to extend the map $g : K(M) \to X$ to $f : C \to X$ skeleton by skeleton employing the technique described above. This contradicts the fact that M is essential, since the non-trivial map $H_n(M) \to H_n(X)$ would factor through the zero map $H_n(M) \to H_n(C)$. □

The second estimate needed is an upper bound for the filling radius of a riemannian manifold by a function of its volume. When interpreted using Gromov's language, an inequality of that kind for submanifolds of euclidean spaces is a consequence of the classical Federer-Fleming isoperimetric inequality [32].

Theorem 3.1.4 (Federer and Fleming)**.** *Let M be an n-dimensional submanifold (or more generally, an n-cycle) of an euclidean space $\mathbb{R}^N$. Then,* $\operatorname{FillRad}(M \subseteq \mathbb{R}^N) \leq C_N \operatorname{Vol}(M)^{\frac{1}{n}}$*, where $C_N > 0$ depends only on the ambient dimension N.*

The first obstacle to extend this estimate to a general closed riemannian manifold is that the statement holds only for submanifolds of a finite dimensional euclidean space. This can be surmounted by approximating the Kuratowski isometric embedding for a manifold M by finite dimensional embeddings of M in $\ell^\infty_N = (\mathbb{R}^N, \|\cdot\|_\infty)$. More concretely, one may take a finite δ-net $x_1, \ldots, x_N$ for some $\delta > 0$ in M and consider the map $F : M \to \ell^\infty_N$ which in each coordinate is defined as the distance function from the corresponding point of the net. As sketched by Guth (see [48, Lemma 2]) and proved by Katz and Katz (see [62, Theorem 3.1]), provided $\delta > 0$ is chosen sufficiently small, F approximates the Kuratowski embedding up to an arbitrary error.

Theorem 3.1.5. *Let (M, g) be a closed riemannian manifold. Given $\varepsilon > 0$, there exists an embedding $F : M \to \ell^\infty_N$ for a sufficiently large N such that*

$$(1-\varepsilon)d(x, y) \leq \|F(x) - F(y)\|_\infty \leq d(x, y)$$

for every x, y in M, where d denotes the riemannian distance in M.

In order to work with these finite dimensional embeddings instead of the Kuratowski embedding for a manifold, one would also need to extend the bi-Lipschitz constants from Theorem 3.1.5 to the chains bounding the manifold. This can be accomplished by the well-known Lipschitz extension property of L^∞ spaces.

Lemma 3.1.6. *Let X be a metric space, $Y \subseteq X$ a subspace and let $f : Y \to \ell^\infty(S)$ a Lipschitz function (here, S is a set). Then, there is an extension $\tilde{f} : X \to \ell^\infty(S)$ of f with the same Lipschitz constant as f.*

See [48, Lemma 3] for a proof.

As a consequence, given an arbitrary $\varepsilon > 0$, for a closed riemannian manifold M there is a bi-Lipschitz embedding $F : M \to \ell^\infty_N$ for which the relative filling radius $\operatorname{FillRad}(F(M) \subseteq \ell^\infty_N)$ lies between $(1-\varepsilon)\operatorname{FillRad}(M)$ and $\operatorname{FillRad}(M)$. Although the Federer-Fleming estimate from Theorem 3.1.4 applies in principle only to submanifolds of euclidean spaces, after a slight adjustment it may be used to bound the filling radius of $F(M) \subseteq \ell^\infty_N$. Indeed, recall that for a riemannian manifold M of dimension n, the n-dimensional Hausdorff measure, which depends only on the riemannian distance function, coincides with $\operatorname{Vol}(M)$. Hence, since the supremum norm $\|\cdot\|_\infty$ is bi-Lipschitz equivalent to the usual euclidean norm $\|\cdot\|_2$ on $\mathbb{R}^N$, up to changing the dimensional constant C_N the Federer-Fleming estimate implies that

$$\operatorname{FillRad}(F(M) \subseteq \ell^\infty_N) \leq C_N \operatorname{Vol}(F(M))^{\frac{1}{n}},$$

where $\operatorname{Vol}(F(M))$ is to be interpreted as the n-dimensional Hausdorff measure of $F(M)$.

However, there is still a considerable gap between this estimate and the required

$$\operatorname{FillRad}(M) \leq C_n \operatorname{Vol}(M)^{\frac{1}{n}},$$

since we do not control how large the N needs to be taken and the constants C_N go to infinity as $N \to \infty$. The main technical result in [44] is precisely a remarkable generalization of the Federer-Fleming Theorem 3.1.4 to Banach spaces, with a constant that depends only on the dimension of the manifold (or cycle) involved.

Theorem 3.1.7. *([44, §4.3.B]) Let M be an n-dimensional submanifold (or more generally, an n-cycle) of an L^∞ space L. Then,*

$$\operatorname{FillRad}(M \subseteq L) \leq C_n \operatorname{Vol}(M)^{\frac{1}{n}},$$

for some constant $C_n < (n+1)n^n\sqrt{(n+1)!}$.

The proof, which exploits the Federer and Fleming estimate as well as the techniques for settling it, is generally regarded as the most complicated part of the article [44]. We refer the reader to the Guth's notes [48] and the original work [44, §3.2 and §4.3] for the proof.

3.2 Systolic area

In this section we focus on the study of the systolic volume (or better, systolic *area*) of 2-dimensional spaces. Recall that for an individual manifold M, or more generally, a piecewise riemannian polyhedron of dimension n, the systolic volume $\sigma(M)$ is the best constant $C = C(M)$ for which the inequality

$$C(M)\,\mathrm{sys}(M,g)^n \leq \mathrm{Vol}(M,g)$$

holds for any riemannian metric on g on M. As it is easy to check, all closed surfaces different from the sphere (either orientable or not) are essential, so that Gromov's systolic inequality guarantees that the systolic area of such surfaces is strictly positive. However, it is possible to obtain better estimates through arguments different from Gromov's filling techniques.

As we mentioned at the start of the chapter, there are three surfaces for which the exact value of the systolic area is known. These are the torus $\mathbb{T}^2$, the projective plane $\mathbb{R}P^2$ and the Klein bottle K. The systolic area of the torus is $\sigma(\mathbb{T}^2) = \frac{\sqrt{3}}{2}$ by the Loewner's inequality, whose proof was sketched in the previous section. We show below that $\sigma(\mathbb{R}P^2) = \frac{2}{\pi}$, a result due to Pu [74]. The argument can be easily adapted to give a proof of Loewner's inequality. Finally, Bavard proved that $\sigma(K) = \frac{2\sqrt{2}}{\pi}$ in [8] with the optimum attained at a non-smooth riemannian metric.

Theorem 3.2.1. *For every riemannian metric g on $\mathbb{R}P^2$,*

$$\mathrm{sys}(\mathbb{R}P^2, g)^2 \leq \frac{\pi}{2}\,\mathrm{Area}(\mathbb{R}P^2, g).$$

Moreover, equality holds only for metrics of constant curvature 1.

Proof. By the uniformization theorem, there is a positive function $f : \mathbb{R}P^2 \to \mathbb{R}$ (a conformal factor) such that $g = fg_0$, where g_0 is the metric of constant curvature 1 on $\mathbb{R}P^2$ determined by the standard round metric $d\theta^2 + \sin(\phi)d\phi^2$ on the unit sphere. Consider the function $\bar{f}$ obtained by averaging the conformal factor over the isometry group G of $(\mathbb{R}P^2, g_0)$:

$$\bar{f}^{\frac{1}{2}}(x) = \int_G f(\xi x)^{\frac{1}{2}}\, d\xi,$$

where $d\xi$ is the normalized Haar measure of the group G, that is, the unique left invariant measure that integrates 1.

Let us compare the systole and the area of the metric $\bar{f}g_0$ to that of g. For the systole, we take a non-trivial loop γ and compute:

$$\begin{aligned}
\mathrm{length}(\gamma, \bar{f}g_0) &= \int_\gamma \int_G f(\xi\gamma(t))^{\frac{1}{2}}\, d\xi\, g_0(\dot\gamma(t), \dot\gamma(t))\, dt \\
&= \int_G \int_\gamma f(\xi\gamma(t))^{\frac{1}{2}} g_0(\dot\gamma(t), \dot\gamma(t))\, dt\, d\xi \\
&= \int_G \int_\gamma f(\xi\gamma(t))^{\frac{1}{2}} g_0(\dot{\xi\gamma}(t), \dot{\xi\gamma}(t))\, dt\, d\xi \\
&= \int_G \mathrm{length}(\xi\gamma, g)\, d\xi \\
&\geq \int_G \mathrm{sys}(\mathbb{R}P^2, g)\, d\xi = \mathrm{sys}(\mathbb{R}P^2, g),
\end{aligned}$$

where we have used Fubini's theorem in the first step, that the ξ are isometries in the second and that γ is non-trivial in the last step. By taking γ as a systolic loop for the metric $\bar{f}g_0$, we see that $\operatorname{sys}(\mathbb{R}P^2, \bar{f}g_0) \geq \operatorname{sys}(\mathbb{R}P^2, g)$.

With respect to the area, we perform a similar computation, letting ω_0 be the area form of g_0:

$$\begin{aligned}
\operatorname{Area}(\mathbb{R}P^2, \bar{f}g_0) &= \int_{\mathbb{R}P^2} \bar{f}^2 \, d\omega_0 \\
&= \int_{\mathbb{R}P^2} \left(\int_G f(\xi x)^{\frac{1}{2}} \, d\xi \right)^2 d\omega_0 \\
&\leq \int_{\mathbb{R}P^2} \left(\int_G f(\xi x) \, d\xi \right) d\omega_0 \\
&= \int_G \int_{\mathbb{R}P^2} f(\xi x) \, d\omega_0 \, d\xi \\
&= \int_G \operatorname{Area}(\mathbb{R}P^2, g) \, d\xi = \operatorname{Area}(\mathbb{R}P^2, g),
\end{aligned}$$

where the inequality in the third line is a consequence of the Cauchy-Schwarz inequality. Thus,

$$\frac{\operatorname{Area}(\mathbb{R}P^2, g)}{\operatorname{sys}(\mathbb{R}P^2, g)^2} \geq \frac{\operatorname{Area}(\mathbb{R}^2, \bar{f}g_0)}{\operatorname{sys}(\mathbb{R}P^2, \bar{f}g_0)^2}.$$

It remains to notice that, since the group of isometries G is transitive, the function $\bar{f}$ is constant, say $\bar{f} \equiv c$. Hence, $\operatorname{sys}(\mathbb{R}P^2, \bar{f}g_0) = c \operatorname{sys}(\mathbb{R}P^2, g_0) = \pi$ and $\operatorname{Area}(\mathbb{R}P^2, \bar{f}g_0) = c^2 \operatorname{Area}(\mathbb{R}P^2, g_0) = 2\pi c^2$, half the area of the standard sphere of radius c. The desired estimate follows.

To analyze the equality case, notice that the inequality in the area estimation

$$\left(\int_G f(\xi x)^{\frac{1}{2}} \, d\xi \right)^2 \leq \int_G f(\xi x) \, d\xi$$

should be an equality. But this can only happen if f is a constant function, i.e. g is a scalar multiple of the metric of constant curvature 1. $\square$

Actually, the inequality $\sigma(S) \geq \frac{2}{\pi}$ is valid for all essential closed surfaces S by an inequality of Gromov. To prove it, Gromov exploited the topology of the universal coverings of aspherical surfaces, which are, as is well known, homeomorphic to the plane $\mathbb{R}^2$.

Let us make a preliminary observation about the metric on a covering $\tilde{S}$ (not necessarily universal) of a riemannian surface (S, g). We shall always consider a covering space $\tilde{S}$ endowed with the unique riemannian metric $\tilde{g}$ that makes the covering map $p : (\tilde{S}, \tilde{g}) \to (S, g)$ a local isometry. Using this metric it is possible to prove quantitatively that the covering projections are local homeomorphisms, a fact that was explicitly observed in [63].

Lemma 3.2.2. *Let $p : (\tilde{M}, \tilde{g}) \to (M, g)$ be a covering projection of a riemannian manifold (M, g). Then, p maps injectively balls of radius less than $\frac{1}{2}\operatorname{sys}(M, g)$.*

Proof. Let B be a ball of radius $r < \frac{1}{2}\operatorname{sys}(M, g)$ in the covering space $\tilde{M}$. Suppose that $p(y) = p(z)$ for certain y, z in B. In particular, $d(y, z) < 2r$, so that there is a curve λ that joins y to z

and has length at most $2r$. Then, the length of the loop $p \circ \lambda$ is less or equal than $2r < \operatorname{sys}(M, g)$, from where it follows that it is null-homotopic. However, the loop does not lift to a loop in $\tilde{S}$, a contradiction. Hence, p is injective on B. ☐

This fact was exploited by Gromov in his sharp systolic inequality for aspherical surfaces. Recall that given a point x in a riemannian manifold (M, g), $B(x, r)$ denotes the open ball of radius $r > 0$ around x with respect to the distance determined by the riemannian tensor.

Theorem 3.2.3. *([44, Theorem 5.2.A]) Let (S, g) be a closed aspherical riemannian surface and $\gamma \subseteq S$ a systolic loop. Then, given $r < \frac{\operatorname{sys}(S,g)}{2}$ there exists a point $x \in \gamma$ such that*

$$\operatorname{Area} B(x, r) \geq 3r^2.$$

In particular, $\sigma(S) \geq \frac{3}{4}$.

Proof. Consider the average area of the balls of radius r centered at points of γ

$$\frac{1}{\operatorname{sys}(S, g)} \int_\gamma \operatorname{Area} B(v, r)\, d\gamma.$$

In order to show that at least one of those balls has area greater or equal to $3r^2$, in view of the coarea inequality applied to balls:

$$\operatorname{Area} B(v, r) \geq \int_0^s \operatorname{length}(\partial B(v, s))\, ds,$$

it is enough to check that the average length of the spheres $\partial B(v, s)$ over γ is at least $6s$ for each $0 < s < r$. To do this, notice first that the lift $\tilde{\gamma} \subseteq \tilde{S}$ is an infinite geodesic, homeomorphic to a straight line in $\tilde{S}$. Fix $0 < s < r$ and mark an infinite sequence of points $(w_k)_{k \in \mathbb{Z}}$ on $\tilde{\gamma}$ such that $d_{\tilde{S}}(w_i, w_j) = s|i - j|$. We show that the average length of the spheres $\partial B_{\tilde{S}}(w_i, s)$ is at least $6s$. The sphere $\partial B_{\tilde{S}}(w_i, s)$ intersects on segments of length at least $2s$ with each of the balls $B_{\tilde{S}}(w_{i-1}, s)$, $B_{\tilde{S}}(w_{i+1}, s)$. The remaining two segments λ_i^+, λ_i^- that conform $\partial B_{\tilde{S}}(w_i, s)$ lie each on a different side of the separating geodesic $\tilde{\gamma}$. For indices $k < l$, consider the path between w_l and w_k formed by joining all the segments λ_j^+ in the upper side of the spheres $\partial B_{\tilde{S}}(w_j, s)$ for $k \leq j \leq l$ together with segments inside $B_{\tilde{S}}(w_k, s)$ and $B_{\tilde{S}}(w_l, s)$ that join its centers to an endpoint of λ_k^+ and of λ_l^+, respectively (see Figure 3.1). As the distance between w_k and w_l is $s|k - l|$ by hypothesis, we have for the length of that path the estimate

$$\sum_{j=k}^{l} \operatorname{length}(\lambda_j^+) + 2s \geq s|k - l|.$$

Hence, if we let $n = |k - l|$ tend to infinity, we see that

$$\liminf_n \frac{1}{n} \sum_j \operatorname{length}(\lambda_j^+) \geq s,$$

from where it follows that the average length of the spheres $\partial B_{\tilde{S}}(w_i, s)$ is bounded from below by $6s$. Since p maps all these balls injectively to S by Lemma 3.2.2, the average length of the spheres $\partial B(v, s)$ in S is greater than or equal to $6s$, which finishes the proof. ☐

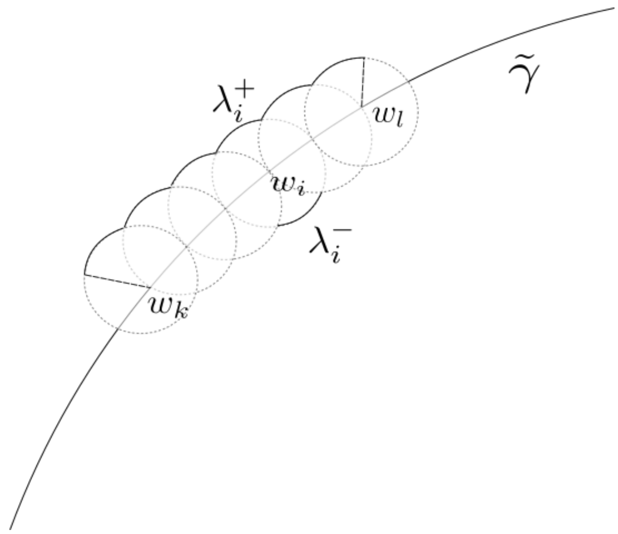

Figure 3.1: In the upper half of the plain, the path that joins w_l to w_k is marked in bold face.

Since all closed surfaces (orientable or not) apart from the sphere and the projective plane are aspherical, we conclude the following universal estimate for the systolic area of surfaces.

Corollary 3.2.4. *Let S be a closed surface different from the sphere. Then,*

$$\sigma(S) \geq \frac{2}{\pi}.$$

Moreover, the equality is attained only at the projective plane with its standard round metric.

Another feature of the systolic area of surfaces which is completely understood is its asymptotic behavior as the genus γ goes to infinite. The best current result in this direction is the following.

Theorem 3.2.5. *(cf. [65, Remark 4.2]) Let S_γ denote a closed surface of genus γ. Then,*

$$\pi \leq \lim_{\gamma \to \infty} \frac{\gamma}{\log(\gamma)^2} \sigma(S_\gamma) \leq \frac{9\pi}{4}.$$

The upper bound is due to Buser and Sarnak [27], who constructed hyperbolic surfaces with the required (asymptotically) systolic area as coverings of an arithmetic Riemann surface. The asymptotically correct lower bound was first established by Gromov in [44, §6.4] through a kind of averaging process over measurable chains known as *diffusion of chains* developed in [43]. The constant π, better than the $\frac{1}{4}$ obtained by Gromov's techniques, was settled by Katz and Sabourau in [65, Theorem 4.1] using Katok's estimate on the *entropy volume* (or *asymptotic volume*) for hyperbolic surfaces.

We close this brief survey about the systolic area of surfaces mentioning some open problems. By analogy with the inequality from Corollary 3.2.4, one may ask if all surfaces of genus ≥ 1 satisfy Loewner's inequality, that is, if $\sigma(S) \geq \frac{\sqrt{3}}{2}$ for S of genus at least 1. By a result of Katz and Sabourau [65, Theorem 5.1], the inequality holds for all surfaces of genus at least 20. Even more ambitiously: is the systolic area σ an increasing function of the genus of the surfaces? The motivation for this question comes from the intuition that the systolic area of a surface S gives a geometric measure of how difficult is to construct a riemannian manifold with the topology of S and the fact that the topology of surfaces becomes more complex as the genus increases.

We shift now our attention to the systolic geometry of 2-dimensional polyhedra. A *piecewise riemannian polyhedron* (X, g) consists of a polyhedron X together with a triangulation in which every simplex is endowed coherently with a smooth riemannian metric. This means that whenever two simplices σ and τ intersect, the corresponding riemannian metrics g_σ and g_τ coincide in the intersection: $g_\sigma|_{\sigma\cap\tau} = g_\tau|_{\sigma\cap\tau}$. The piecewise riemannian structure over X determines in the natural way a length structure for X, while the volume of X is defined, as expected, as the sum of the volumes over the simplices of maximum dimension. Hence, it is clear how to define the systolic volume for compact polyhedra. As Gromov proved in [44, Appendix 2] his systolic inequality (Theorem 3.1.2) applies also to *essential polyhedra* (although with a worse constant), and hence the systolic volume is positive for a considerable class of spaces.

Theorem 3.2.6. *Let X be an n-dimensional compact essential polyhedron. Then, there exists a constant $C_n > 0$ only depending on the dimension n such that*

$$\mathrm{sys}(X, g)^n \leq C_n \,\mathrm{Vol}(X, g)$$

for every piecewise riemannian metric g on X.

Here, following Gromov we call an n-dimensional polyhedron X *essential* if there exists an aspherical polyhedron K and a continuous map $X \to K$ that does not contract to the $(n-1)$-skeleton of K. The intuition is that in an essential polyhedron X, analogously as in essential manifolds, the 1-dimensional topology detects the non-triviality of the n-dimensional topology of X. The difference is that in a general polyhedron there is no fundamental class to accomplish that and hence one uses the n-skeleton of the polyhedron for that purpose.

The strategy to prove Theorem 3.2.6 is similar to the manifold case. Gromov defined a metric invariant for piecewise riemannian polyhedra called *contractibility radius*, which acts as the filling radius in this context. Namely, the $n-1$ contractibility radius of a piecewise riemannian polyhedron X measures how large needs to be taken a neighborhood U of X in $L^\infty(X)$ to guarantee that the inclusion map $X \hookrightarrow U$ factors through an $(n-1)$-dimensional polyhedron. Both Lemma 3.1.3 and Theorem 3.1.7 generalize to this context, replacing the filling radius by the $n-1$ contractibility radius. See [44, Appendix 2] for more details.

The main reason we are interested in the systolic geometry of 2-dimensional polyhedra is for its connection with groups. It is not difficult to check that the essential polyhedra in dimension 2 are exactly those with non-free fundamental group. For such polyhedra X, Gromov's Theorem 3.2.6 implies that the systolic area $\sigma(X)$ is bounded from below by $\frac{1}{10^4}$. In [77], Rudyak and Sabourau significantly improved this estimate for a class of groups that the authors call of zero *Grushko free index*. Recall that according to Grushko's Theorem [80], a finitely presentable group G admits a decomposition

$$G = F_k * H_1 * \cdots * H_n,$$

where F_k is the free group of rank k and the groups H_i are non-free and freely indecomposable groups. Moreover, such decomposition is unique up order and taking conjugates. The number k is referred to as the *Grushko free index* of G in [77, §1].

Theorem 3.2.7. *([77, Theorem 3.5]) Let (X, g) be a piecewise riemannian polyhedron of dimension 2. Suppose that its fundamental group $\pi_1(X)$ is of zero Grushko free index. Then, if x belongs to a systolic loop of x,*

$$\operatorname{Area} B(x, r) \geq r^2,$$

for every $0 < r < \frac{1}{2} \operatorname{sys}(X, g)$. In particular, $\sigma(X) \geq \frac{1}{4}$.

The argument of the proof relies in a careful analysis of the topology of the *spheres* $\partial B(x, r)$ (that is, the level curves of the distance function from x), from which a bound on the area of the balls is obtained via the coarea inequality. Through the use of another topological argument, this time concerning covering theory, the estimate was extended to all essential polyhedra in [63, Theorem 3.1].

Theorem 3.2.8. *([63, Theorem 3.1]) Let (X, g) be an essential piecewise riemannian polyhedron of dimension 2. Then, there is a point x in X such that*

$$\operatorname{Area} B(x, r) \geq r^2,$$

for every $0 < r < \frac{1}{2} \operatorname{sys}(X, g)$. In particular, $\sigma(X) \geq \frac{1}{4}$.

Proof. Since X is essential, its fundamental group is not free and hence X admits a cover $\tilde{X}$ with zero Grushko free index fundamental group. Notice that the covering map $p : (\tilde{X}, \tilde{g}) \to (X, g)$ is a local isometry and induces a monomorphism in fundamental groups. Therefore, $\operatorname{sys}(\tilde{X}, \tilde{g}) \geq \operatorname{sys}(X, g)$. By Theorem 3.2.8 there exists a point $\tilde{x} \in \tilde{X}$ such that

$$\operatorname{Area} B_{\tilde{X}}(\tilde{x}, r) \geq r^2,$$

for every $0 < r < \frac{1}{2} \operatorname{sys}(X, g)$. This implies the claim by Lemma 3.2.2. □

As we explained in the introduction of the chapter, the systolic geometry of 2-dimensional polyhedra extends to finitely presentable groups via the well known correspondence between the two categories. Recall that the systolic area $\sigma(G)$ of a finitely presentable group G is defined as the best systolic area $\sigma(X)$ among its triangulations X. It is easy to see that the systolic area of free groups is trivial, since such groups admit triangulations with no 2-dimensional material. On the other hand, the inequality from Theorem 3.2.8 implies that the systolic area of any finitely presentable non-free group G is greater than or equal to $\frac{1}{4}$ (compare to Corollary 3.2.4). In the next section we improve this lower bound by a factor of 2 for a certain class of groups which includes the surface groups. The intuition behind this refinement is that, while certainly $\sigma(\pi_1(S)) \leq \sigma(S)$ for closed surfaces S, there should be some stronger relations between those two quantities. In particular, the universal lower bounds for the systolic area of surfaces and surface groups should be closer (if not equal).

Another elementary observation about the systolic area of groups is that $\sigma(G * H) \leq \sigma(G) + \sigma(H)$ for finitely presentable groups G, H. Indeed, take near optimal triangulations (X_1, g_1) and (X_2, g_2) for G and H respectively, normalized so that $\operatorname{sys}(X_1, g_1) = \operatorname{sys}(X_2, g_2) = 1$. Clearly, the area of the wedge sum $X_1 \vee X_2$ endowed with the obvious piecewise riemannian metric is the sum of the areas of X_1 and X_2 and the systole equals 1. On the other hand, it seems difficult to imagine more economic models for a free product of groups, especially when one of the two is free (cf. [77, Question 1.2]). However, at the time of writing, the question of whether $\sigma(G * \mathbb{Z}) = \sigma(G)$ remains wide open. We discuss this problem in the last section of the chapter.

3.3 An extension of Guth's systolic inequality

In the present section we generalize to piecewise riemannian polyhedra of dimension 2 a systolic inequality of Guth [49] for riemannian manifolds. Our principal interest in the inequality is that it allows to refine the lower bound for the systolic area of a considerable class of groups which contains in particular the surface groups.

We start with the statement of Guth's systolic inequality.

Theorem 3.3.1. *Let M be a closed smooth n-manifold. Suppose that there are cohomology classes $\alpha_1, \ldots, \alpha_n \in H^1(M, \mathbb{Z}_2)$ with $\alpha_1 \cup \cdots \cup \alpha_n \neq 0$ in $H^n(M, \mathbb{Z}_2)$. Then, given a riemannian metric g on M, there exists a point $x \in M$ such that for every $r < \frac{\operatorname{sys}(M,g)}{2}$*

$$\operatorname{Vol} B(x, r) \geq C_n (2r)^n,$$

where C_n is a positive constant that depends only on the dimension. In particular, it follows that $\sigma(M) \geq C_n$.

Explicitly, according to a refinement due to Nakamura [73] the constant may be taken as $C_n = \frac{1}{n!}$, which is a considerable improvement over the constant

$$C_n = \frac{1}{(6(n+1) \cdot n^n \cdot \sqrt{n+1!}))^n}$$

that may be deduced from the proof of Gromov's systolic inequality.

Let us compare both systolic inequalities further. The topological condition regarding the cup-length with coefficients in $\mathbb{Z}_2$ in Guth's inequality is more restrictive than Gromov's essentialness, but, in the same spirit, it applies to manifolds for which the 1-dimensional topology (in this case, the cohomology) generates in some sense the fundamental class. On the other hand, besides providing tighter constants, Guth's inequality has the advantage of being local, in the sense that it shows that the volume of balls of radii comparable to the systole are relatively large instead of estimating the volume of the whole manifold. As for the methods of the proof, while definitely not trivial, the inequality of Guth avoids most of the technical difficulties of Gromov's argument, particularly the generalization of Federer-Fleming isoperimetric inequality to infinite dimensional Banach spaces. For the same reason, the proof of Guth's inequality does not provide the additional new information about the filling radius and isoperimetric inequalities in infinite-dimensional Banach spaces produced by Gromov's techniques.

The main novelty of the argument in Guth's proof is the use of (near) minimal hypersurfaces. To explain this point, consider a closed manifold M in the conditions of Theorem 3.3.1. By Thom's theorem [81], the $\mathbb{Z}_2$-homology class Poincaré dual to α_n may be represented by a smooth hypersurface (i.e., a codimension 1 submanifold) of M, which is easily seen to be of maximal cup-length with coefficients in $\mathbb{Z}_2$ as well. Suppose for the time being that there exists a hypersurface Z of minimum volume representing that homology class. The technical heart of the proof is a stability estimate, similar to estimates for minimal surfaces in [78], that controls the volume of balls of small radii in Z. More concretely, to prove this estimate, Guth exploited the following relatively intuitive fact: if one takes a ball $B(z, r)$ of an adequately small radius r centered at a point $z \in Z$ and cuts from Z the intersection $Z \cap B(z, r)$ to fill the hole by one of the hemispheres of the sphere $\partial B(z, r)$, one obtains another hypersurface Z' homologous to the original (see Figure 3.2). Hence, by the minimality of Z, the $n-1$ dimensional volume of $Z \cap B(z,r)$ is comparable to the $n-1$ dimensional volume of the sphere $\partial B(z,r)$. Inductively, the $n-1$ dimensional volume of $Z \cap B(z,r)$ (and hence, of $\partial B(z,r)$) is relatively large for all r in a small interval. The final estimate on the volume of balls $B(z,r)$ is obtained by integrating the bounds for the $n-1$ dimensional volume of each sphere $\partial B(z,s)$, $s < r$, that is, using the coarea inequality.

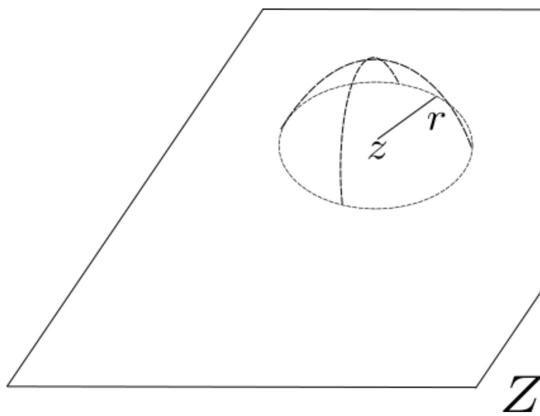

Figure 3.2: The hypersurface Z with a *hemisphere* of the sphere $\partial B(z,r)$.

To give the actual proof, we will need to state some preliminary lemmas and fix some notations. If x is a point in a riemannian manifold (or piecewise riemannian polyhedron) (X,g), the closed (metric) ball of radius r around x will be denoted, as expected, by $\bar{B}(x,r)$. For a non-trivial homology class $[\gamma]$ in $H_1(X,\mathbb{Z}_2)$, the length of $[\gamma]$, denoted length$([\gamma])$, is the infimum of the lengths over 1-cycles representing $[\gamma]$. Following Guth [50], the definition of length is extended to a non-trivial cohomology class α by duality:

$$\text{length}(\alpha) := \inf\{\text{length}([\gamma]) : [\gamma] \in H_1(X,\mathbb{Z}_2),\ \alpha([\gamma]) \neq 0\}.$$

Observe that $\text{length}(\alpha) \geq \text{sys}(X,g)$ for any non-trivial $\alpha \in H^1(X,\mathbb{Z}_2)$. Also, if X is a closed manifold and Z a hypersurface representing a $\mathbb{Z}_2$ homology class in $H_{n-1}(X,\mathbb{Z}_2)$, Z will be called δ-minimizing if the $n-1$ dimensional volume of any other smooth submanifold Z' in the same homology class verifies the inequality $\text{Vol}_{n-1}(Z') \geq \text{Vol}_{n-1}(Z) - \delta$.

We will also use the following homological version of a lemma by Gromov [47, p. 290], due to Guth [50, Curve-factoring Lemma].

Lemma 3.3.2. *Let (M,g) be a complete riemannian manifold. Let γ be a 1-cycle contained in a ball $B(x,r)$ for $x \in M$, $r > 0$. Then, for any given $\varepsilon > 0$, γ is homologous to a finite sum $\sum_i \gamma_i$, where each γ_i is a 1-cycle of length at most $2r + \varepsilon$.*

Proof. It is enough to prove the lemma for the case when γ is homeomorphic to a circle. Subdivide γ in a finite number of intervals σ_i, each of length less than ε. Now, join the i-th vertex of the subdivision to x by a shortest geodesic segment λ_i, where the numbering is cyclic. Then γ is homologous to $\sum_i \gamma_i$ where we set $\gamma_i = \lambda_i + \sigma_i - \lambda_{i+1}$. This concludes the proof, because the length of each segment λ_i is at most r. $\square$

Instead of Guth's original version, we reproduce here an improvement of the stability lemma by Nakamura [73, Lemma 2.1].

Lemma 3.3.3. *Let (M, g) be a closed riemannian manifold. Let $\alpha \in H^1(M, \mathbb{Z}_2)$ be a non-trivial class of length $2R$. For $\delta > 0$, let Z be a δ-minimizing embedded hypersurface in M representing the Poincaré dual homology class to α. Then, if $r < R$,*

$$\mathrm{Vol}_{n-1}(\partial \bar{B}(z, r)) \geq 2 \,\mathrm{Vol}_{n-1}(Z \cap \bar{B}(z, r)) - 2\delta,$$

for every $z \in Z$.

Proof. As we explained above, the key to prove this result is showing that the hypersurface formed by removing $Z \cap \bar{B}(z, r)$ from Z and adding an appropriate "hemisphere" from $\partial \bar{B}(z, r)$ is homologous to Z. To accomplish this, notice in first place that the relative $(n-1)$-cycle $Z \cap \bar{B}(z, r)$ in the pair $(\bar{B}(z, r), \partial \bar{B}(z, r))$ is null-homologous. Indeed, if it was not the case, by Poincaré-Lefschetz duality the cap product between $Z \cap \bar{B}(z, r)$ and some (absolute) 1-cycle $\gamma \subseteq \bar{B}(z, r)$ would be non-trivial. Since $r < R$, by Lemma 3.3.2 γ is homologous to a sum $\sum_i \gamma_i$ where we can make the length of each γ_i to be less than $2R = \mathrm{length}(\alpha)$ by choosing a conveniently small $\varepsilon > 0$. Thus $\alpha([\gamma]) = \sum_i \alpha([\gamma_i]) = 0$ by the definition of the length of α, which contradicts the fact that the cap product between $Z \cap \bar{B}(z, r)$ and γ is non-trivial.

Hence, the relative cycle $Z \cap \bar{B}(z, r)$ bounds a relative chain Q_1. Since the cycle is embedded, this chain must be the sum of some components of $\bar{B}(z, r) \setminus Z$. Moreover, since the cycle is null-homologous and the coefficient field is $\mathbb{Z}_2$, it also bounds the complementary chain Q_0. It follows that $\partial \bar{B}(z, r)$ is decomposed along $Z \cap \partial \bar{B}(z, r)$ into the chains $Q_0 \cap \partial \bar{B}(z, r)$ and $Q_1 \cap \partial \bar{B}(z, r)$. Thus, Z is homologous to the cycle Z' formed by cutting out $Z \cap \bar{B}(z, r)$ from Z and gluing in the smallest "hemisphere" between $Q_0 \cap \partial \bar{B}(z, r)$ and $Q_1 \cap \partial \bar{B}(z, r)$. Since it is possible to smooth Z' near the intersection to become a smooth embedded hypersurface without increasing its volume and Z is δ-minimizing, we deduce that $\mathrm{Vol}_{n-1}(Z') \geq \mathrm{Vol}_{n-1}(Z) - \delta$. Now Z and Z' coincide outside $\bar{B}(z, r)$, so that this amounts to the inequality

$$\begin{aligned}\mathrm{Vol}_{n-1}(\partial \bar{B}(z, r) &\geq 2 \min\{\mathrm{Vol}_{n-1}(Q_0 \cap \partial \bar{B}(z, r)), \mathrm{Vol}_{n-1}(Q_1 \cap \partial \bar{B}(z, r))\} \\ &\geq 2 \,\mathrm{Vol}_{n-1}(Z \cap \bar{B}(y, r)) - 2\delta,\end{aligned}$$

as desired. □

We are now ready to prove the Nakamura's refinement of Guth's systolic inequality.

Theorem 3.3.4. *Let (M, g) be a closed riemannian manifold of dimension n. Let $\alpha_1, \ldots, \alpha_n \in H^1(S, \mathbb{Z}_2)$ be not necessarily distinct classes such that $\alpha_1 \cup \cdots \cup \alpha_n \neq 0$ in $H^n(M, \mathbb{Z}_2)$ and let $2R := \min_i\{\mathrm{length}(\alpha_i)\} > 0$. Then, there exists $x \in M$ such that for any $r \in (0, R)$,*

$$\mathrm{Vol}\, B(x, r) \geq \frac{(2r)^n}{n!}.$$

In particular, $\sigma(M) \geq \frac{1}{n!}$.

Proof. We proceed by induction in the dimension n, the case $n = 1$ being trivial. For $n \geq 2$, take a δ-minimizing hypersurface Z representing the Poincaré dual homology class to α_n. The relation between the cup and the cap product implies that

$$\alpha_1 \cup \cdots \cup \alpha_{n-1}[Z] = \alpha_1 \cup \cdots \cup \alpha_n[M] \neq 0.$$

Hence, we obtain, by restriction to Z, cohomology classes $\alpha'_1, \ldots, \alpha'_{n-1} \in H^1(Z, \mathbb{Z}_2)$ with non-trivial cup product and whose length is no less than the length of the corresponding classes in $H^1(M, \mathbb{Z}_2)$. By induction, there is a point $z \in Z$ such that

$$\mathrm{Vol}_{n-1}(B_Z(z,r)) \geq \frac{(2r)^{(n-1)}}{(n-1)!}$$

whenever $r < R$, where the subscript Z in $B_Z(z,r)$ indicates that the ball is with respect to the induced riemannian metric in Z. Since clearly $B_Z(z,r) \subseteq Z \cap \bar{B}(z,r)$, by Lemma 3.3.3 and the coarea inequality,

$$\begin{aligned}\mathrm{Vol}\, B(z,r) &\geq \int_0^r \mathrm{Vol}_{n-1}(\partial \bar{B}(z,s))\, ds \geq \int_0^r 2\, \mathrm{Vol}_{n-1}(Z \cap \bar{B}(z,r)) - 2\delta\, ds \\ &\geq \int_0^r 2\, \mathrm{Vol}_{n-1}(B_Z(z,r)) - 2\delta\, ds \geq \int_0^r 2\frac{(2s)^{n-1}}{(n-1)!} - 2\delta\, ds = \frac{(2r)^n}{n!} - 2r\delta.\end{aligned}$$

Notice that although it is possible to get this estimate for every $\delta > 0$, the point z belongs to a δ-minimizing hypersurface embedded in M, which therefore depends on δ. To finish the proof, take a sequence of positive numbers (δ_k) converging to 0 and form a corresponding sequence of points $(z_k) \subseteq M$, picking z_k from a δ_k-minimizing hypersurface representing the class $[Z]$. By continuity, the limit $x \in M$ of a convergent subsequence of (z_k) verifies the estimate

$$\mathrm{Vol}\, B(x,r) \geq \frac{(2r)^n}{n!}$$

for $r \in (0, R)$. The bound for the systolic volume of M follows immediately from the fact that $2R \geq \mathrm{sys}(M, g)$. □

This theorem applies, among others, to tori and projective spaces of any dimension, as well as to closed essential surfaces. For manifolds with maximal cup-length with coefficients in $\mathbb{Z}_2$ of dimension $n \geq 3$, it gives the best known lower bound to the systolic volume as Nakamura remarks in [73, §1.4]. In contrast, there are better systolic inequalities in dimension 2, since Theorem 3.3.4 implies only that $\sigma(S) \geq \frac{1}{2}$ for closed surfaces S. However, its conclusion is stronger in one respect: the lower bound for the area of a surface follows only from the existence of two "long" cohomology classes with non-trivial cup product regardless of the size of the systole. This is the crucial observation that allows to extend the inequality to polyhedra of dimension 2. Take a piecewise riemannian 2-polyhedron (X, g) with maximal cup-length. This X has a distinguished $\mathbb{Z}_2$-homology class $[C]$ of dimension 2, namely, the one that detects the non-trivial cup product of 1-dimensional cohomology classes. It can be proved that the class $[C]$ is represented by a (possibly singular) surface embedded in X. Although the systole of S may be arbitrarily small (i.e. it is in general not controlled by the systole of (X, g)), since the class $[C]$ is generated by cohomology classes of length greater than the systole of (X, g), Theorem 3.3.4 applies to give a lower bound for the area of S, hence of X.

Let us formalize the argument sketched above. The first point is describing how to realize 2-dimensional homology classes as singular surfaces, or more precisely, as the continuous image of surfaces in a controlled way. Again by Thom's results, a $\mathbb{Z}_2$-homology class of any dimension in a polyhedron can be represented as the image of the fundamental class of a closed manifold through a continuous map. We record here an explicit construction for the case of 2-dimensional $\mathbb{Z}_2$-homology classes of simplicial complexes similar to the one described in ([52, pp. 108-109]). Recall that a simplicial map is *non-degenerate* if it preserves the dimension of the simplices.

Lemma 3.3.5. *Let X be a simplicial complex and let $[C] \in H_2(X, \mathbb{Z}_2)$ be a non-trivial homology class. Then, there exists a triangulated closed surface S (possibly non-orientable and not connected) together with a non-degenerate simplicial map $h : S \to X$ such that $h_*[S] = [C]$, where $[S] \in H_2(S, \mathbb{Z}_2)$ denotes the fundamental class. Moreover, h does not identify 2-simplices, meaning that $h(\sigma) \neq h(\eta)$ for different 2-simplices σ, η of S.*

Notice in particular that if we equip the 2-complex X with a piecewise riemannian structure g, this construction allows to readily relate the length and the area of subspaces of X to those of subspaces of S endowed with the piecewise riemannian pullback metric through the map $h : S \to X$. More concretely, in this situation the map h preserves lengths and areas.

Proof. Take a 2-cycle $Z = \sum_i \sigma_i$ in $C_2(X, \mathbb{Z}_2)$ representing the homology class $[C]$ and form a disjoint union of 2-simplices $\tilde{\sigma}_i$, one for each σ_i in the support of Z. Since the algebraic boundary of Z is trivial, the edges of the simplices σ_i cancel in pairs. Choose a maximal set of such canceling pairs and identify the edges of $\tilde{\sigma_i}$ accordingly. It is clear that the quotient space obtained from $\coprod_i \tilde{\sigma_i}$ by performing these identifications is a closed surface S and that it gives rise to a simplicial map $h : S \to X$ with the desired properties. □

Remark 3.3.6. In general, for a $\mathbb{Z}_2$-homology class of dimension $n \geq 2$ the construction from the proof gives a realization by a pseudo-manifold whose singularities are of codimension at least 3 (cf. [52, p. 109]). Thus, a different argument is required to extend the inequality in dimensions $n \geq 3$.

Theorem 3.3.7. *Let (X, g) be a connected piecewise riemannian polyhedron of dimension 2. Suppose that there exist classes α, β in $H^1(X, \mathbb{Z}_2)$ such that $\alpha \cup \beta \neq 0$ in $H^2(X, \mathbb{Z}_2)$ and let $2R = \min\{\text{length}(\alpha), \text{length}(\beta)\} > 0$. Then, there is a point $x \in X$ such that for every $r \in (0, R)$,*

$$\text{Area}\, B(x, r) \geq \frac{(2r)^2}{2}.$$

Proof. Since by hypothesis there are classes α, $\beta \in H^1(X, \mathbb{Z}_2)$ such that $\alpha \cup \beta \neq 0$, there exists a homology class $[C] \in H_2(X, \mathbb{Z}_2)$ with $\alpha \cup \beta[C] \neq 0$. By applying Lemma 3.3.5 to the class $[C]$, we obtain a triangulated closed surface S together with a simplicial map $h : S \to X$ representing $[C]$. Endow S with the pullback metric $h^*(g)$, where g is the piecewise riemannian metric on X. By the naturality of the cup product, we have

$$h^*(\alpha) \cup h^*(\beta)[S] = (\alpha \cup \beta)h_*[S] = \alpha \cup \beta[C] \neq 0.$$

If S is not connected, the previous computation implies that for some connected component of S the corresponding components of $h^*(\alpha)$ and $h^*(\beta)$ have non-trivial cup product. By a slight abuse of notation, we will still call S such a component. We claim that both the length of $h^*(\alpha)$ and $h^*(\beta)$ are at least $\text{sys}(X, g)$. Indeed, if γ is a 1-cycle in S such that $h^*(\alpha)\gamma \neq 0$, since $h^*(\alpha)\gamma = \alpha(h_*\gamma)$ and h is length preserving it follows that $\text{length}(h^*(\alpha)) \geq \text{length}(\alpha)$, and, analogously, that $\text{length}(h^*(\beta)) \geq \text{length}(\beta)$. Although in principle Theorem 3.3.4 does not formally apply to $(S, h^*(g))$ because the riemannian metric $h^*(g)$ is only piecewise smooth, one may approximate up to an arbitrarily small error such a metric by a smooth one. Hence, by Theorem 3.3.4 for these smooth approximations and continuity, there is a point y in S such that for every $r \in (0, R)$

$$\operatorname{Area} B_S(y,r) \geq \frac{(2r)^2}{2},$$

where $B_S(y,r)$ stands for the ball of radius r centered at the point $y \in S$. Since h preserves areas we have for $x = h(y)$ that

$$\operatorname{Area} B(x,r) \geq \frac{(2r)^2}{2}.$$

$\square$

Notice that the inequality applies also to non-compact complexes, as long as the length of the cohomology classes involved is positive. This remark is especially relevant when dealing with covering spaces. As we proved in Lemma 3.2.2 after [63], covering projections of piecewise riemannian polyhedra map injectively balls of radius less than half the systole of the base space. Since on the other hand the systole of a non-simply connected covering space is clearly greater than or equal to the systole of the base space, we obtain the following corollary.

Corollary 3.3.8. *Let (X, g) be a compact connected polyhedron of dimension 2 equipped with a piecewise riemannian metric g and let $(\hat{X}, \hat{g})$ be a covering of X. Suppose that there exist classes α, β in $H^1(\hat{X}, \mathbb{Z}_2)$ such that $\alpha \cup \beta \neq 0$ in $H^2(\hat{X}, \mathbb{Z}_2)$ and let $2R := \min\{\operatorname{length}(\alpha), \operatorname{length}(\beta)\}$. Then, there exists $x \in \hat{X}$ such that $\operatorname{Area} B(x,r) \geq \frac{(2r)^2}{2}$ for all $r \in (0,R)$. In particular $\operatorname{Area}(X,g) \geq \frac{1}{2} \operatorname{sys}(X,g)^2$.*

3.3.1 Systolic area of groups

In this section we show how to derive an inequality for the systolic area for the class of surface-like groups from the generalization of Guth's systolic inequality. We recall the definition of surface-like groups from the Introduction.

Definition 3.3.9. Let G be a group. We say that G is *surface-like* if there exist classes α, β in $H^1(G, \mathbb{Z}_2)$ such that $\alpha \cup \beta \neq 0$ in $H^2(G, \mathbb{Z}_2)$.

The cohomological condition in the definition of surface-like groups is of course the translation to the context of groups of the property of having maximal cup-length with coefficients in $\mathbb{Z}_2$ for 2-polyhedra. Likewise, the property that a group contains a surface-like subgroups translates in the context of 2-polyhedra as admitting a covering space of maximal cup-length with coefficients in $\mathbb{Z}_2$. Thus, we obtain a systolic inequality for groups containing a surface-like subgroup through Corollary 3.3.8.

Theorem 3.3.10. *Let G be group which contains a surface-like subgroup T. Then, $\sigma(G) \geq \frac{1}{2}$.*

Proof. Let (X, g) be a piecewise riemannian 2-complex with fundamental group G and let $(\hat{X}, \hat{g})$ be the covering of X with fundamental group T endowed with the pullback metric $\hat{g}$. By Corollary 3.3.8, the proof reduces to verifying that under the hypothesis on its fundamental group, $\hat{X}$ must have maximal cup-length with coefficients in $\mathbb{Z}_2$. Notice that the complex $\hat{X}$ includes as the 2-skeleton of an aspherical CW-complex K (possibly infinite dimensional). Thus K is an Eilenberg-MacLane space $K(T,1)$ and so its $\mathbb{Z}_2$-cohomology ring $H^*(K,\mathbb{Z}_2)$ is isomorphic to the cohomology ring $H^*(T,\mathbb{Z}_2)$ of T. By construction, the inclusion $\hat{X} \hookrightarrow K$ induces an isomorphism $H^1(K,\mathbb{Z}_2) = H^1(T,\mathbb{Z}_2) \to H^1(\hat{X},\mathbb{Z}_2)$ and a monomorphism $H^2(K,\mathbb{Z}_2) = H^2(T,\mathbb{Z}_2) \to H^2(\hat{X},\mathbb{Z}_2)$. Since T is surface-like, this implies that $\hat{X}$ has maximal cup-length with coefficients in $\mathbb{Z}_2$. □

There are many classes of groups to which the bound from Theorem 3.3.10 applies. To begin with, all surface groups are surface-like by Poincaré Duality, as are free abelian groups of rank at least 2, elementary abelian 2-groups and also groups isomorphic to direct and free products of a group with a surface-like group. Examples of groups containing a surface-like subgroups include among others, non-free Artin groups, groups containing an element of order 2 (in particular, to Coxeter groups) and infinite fundamental groups of closed irreducible 3-manifolds. Indeed, Artin groups contain a copy of $\mathbb{Z} \oplus \mathbb{Z}$ unless they are free, while groups with an element of order 2 contain a copy of $\mathbb{Z}_2$. As for the irreducible 3-manifold groups, recall that the surface subgroup conjecture states that every closed, irreducible 3-manifold with infinite fundamental group contains an immersed π_1-injective closed surface. This conjecture was settled in the affirmative by Kahn and Markovic in [59] and thus, by Theorem 3.3.10 the systolic area of infinite fundamental groups of closed irreducible 3-manifolds also admits $\frac{1}{2}$ as a lower bound.

Although the estimate from Theorem 3.3.10 constitutes an improvement upon the general bound $\frac{1}{4}$ for the systolic area of non-free groups, there is still a gap between it and the conjectured optimum $\frac{2}{\pi}$. We end this section by showing that this gap can be closed for fundamental groups of non-orientable surfaces if we assume a strengthened version of Gromov's Filling Area Conjecture. A riemannian manifold (M^{n+1}, g_M) is an *isometric filling* of a riemannian manifold (N^n, g_N) if $\partial M = N$ and the restriction to $N \times N$ of the distance function d_{g_M} determined by g_M coincides with d_{g_N}. As it will become more clear in the next section, estimates on the volume of such fillings could play an important rôle in systolic geometry. The Filling Area Conjecture states that the minimum area among orientable isometric fillings of the circle equipped with its standard riemannian metric of length 2ℓ is attained at the standard hemisphere of area $\frac{2}{\pi}\ell^2$.

Let (X,g) be a piecewise riemannian 2-polyhedron with fundamental group isomorphic to the fundamental group of a non-orientable surface and suppose that $h : S \to X$ represents the relevant homology class of $H_2(X,\mathbb{Z}_2)$ with S a connected surface. By the naturality of the exact sequence from the universal coefficient theorem, S must be non-orientable. Now we proceed as in the proof of [57, Proposition 3.1]. Take the shortest loop $\gamma \in S$ such that $h_*[\gamma]$ is non-trivial in $\pi_1(X)$ and open up S along γ. The resulting riemannian surface $(\Sigma,\tilde{g})$ is an isometric filling (possibly non-orientable) of a circle of twice the length of γ and of the same area as S. If we assume that the conclusion of Filling Area Conjecture holds (also for non-orientable isometric fillings), we have $\mathrm{Area}(\Sigma,\tilde{g}) \geq \frac{2}{\pi}(\mathrm{length}(\gamma))^2 \geq \frac{2}{\pi}(\mathrm{sys}(X,g))^2$. Hence, since $\mathrm{Area}(X,g) \geq \mathrm{Area}(\Sigma,\tilde{g})$ we conclude that $\sigma(X,g) \geq \frac{2}{\pi}$. In particular, this would imply that $\sigma(\mathbb{Z}_2) = \frac{2}{\pi}$.

3.4 Systolic area of a free product of groups

One of the fundamental open problems about the systolic area of groups is understanding the way the systolic area transforms a free product of groups in terms of its free factors. We have already seen that, given finitely presented groups G and H, the inequality $\sigma(G * H) \leq \sigma(G) + \sigma(H)$ holds. It would be interesting to know if there is a stronger relation between the systolic area of $G * H$ and the systolic area of its free factors. More concretely: does the inequality $\sigma(G * H) \geq C(\sigma(G) + \sigma(H))$ hold for some universal constant $C > 0$? While we are not able to answer that question, we intend to describe in this section a possible approach to relate the systolic area of a free product of groups to the systolic area of its free factors.

A first natural attempt to compare the systolic area of a group of the form $G * H$ to that of G and H would go as follows. Take a piecewise riemannian polyhedron (Y, g) that almost realizes the systolic area $\sigma(G * H)$, that is, such that the ratio $\frac{\mathrm{Area}(Y,g)}{\mathrm{sys}(Y,g)^2}$ does not exceed $\sigma(G * H)$ by more than ε for a fixed, arbitrary $\varepsilon > 0$. Fix also triangulations X_G, X_H of the groups G and H, respectively. Since all the polyhedra involved are 2-dimensional and the fundamental group of Y is isomorphic to that of $X_G \vee X_H$, there is a piecewise linear map $g : X_G \vee X_H \to Y$ inducing an isomorphism in fundamental groups (see for example [79, Lemma 1.5]). Then, one expects to be able to estimate the systolic area of $X_G \vee X_H$ endowed with the pullback piecewise riemannian metric via the map g (strictly speaking, such metric could be degenerated but may be approximated by a non-degenerate one) in terms of the systolic area of Y. Unfortunately, there is no way to reasonably control the number of preimages of each 2-simplex in Y via g, at least to the best of our knowledge. Thus, the area of $X_G \vee X_H$ with the metric induced from g could potentially be arbitrarily larger than that of Y.

In view of this, we start with the same setup but instead of g, we use a map in the reverse direction. That is, we take a piecewise linear map $f : Y \to X_G \vee X_H$ which induces an isomorphism on fundamental groups. Essentially, the idea now is to split Y in $f^{-1}(X_G)$ and $f^{-1}(X_H)$ and try to estimate the systolic area of these two subpolyhedra of Y. Let us make this construction more explicit. Fix triangulations of Y, X_G and X_H (which we still denote Y, X_G and X_H) such that the map $f : Y \to X_G \vee X_H$ is simplicial. Let Y_G be the subcomplex of Y defined as $Y_G := f^{-1}(X_G)$ and Y_H the simplicial closure of the complement of Y_G, that is, the subcomplex of Y generated by those simplices s such that $f(s) \not\subseteq X_G$. Notice that the union of Y_G and Y_H covers Y and that the intersection $Y_G \cap Y_H$ is a (possibly not connected) 1-dimensional complex. This last feature is the main reason of the lack of symmetry in the definition of Y_G and Y_H; later on we will attach cones on $Y_G \cap Y_H$ and we will need to keep our constructions in the 2-dimensional world. The first objective is obtaining complexes with fundamental group G and H from Y_G and Y_H. In order to do this, we start by recording two topological properties of the decomposition $Y = Y_G \cup Y_H$ in the next lemma.

Lemma 3.4.1. *Let Y_G and Y_H be as above. Then,*

- *for every connected component C of $Y_G \cap Y_H$ the inclusion $\iota_C : C \hookrightarrow Y$ induces the trivial morphism $\pi_1(C) \to \pi_1(Y)$, and*
- *every connected component of Y_G (symmetrically, of Y_H) contains at least one connected component of $Y_G \cap Y_H$.*

Proof. For the first point, denote v the wedge point of $X_G \vee X_H$ and notice that $f(C) = \{v\}$ for every component C of $Y_G \cap Y_H$. Hence, $f \circ \iota_C$ induces the trivial morphism $\pi_1(C) \to \pi_1(X)$. Since f induces an isomorphism $\pi_1(Y) \to \pi_1(X_G \vee X_H)$, it follows that $\iota_* : \pi_1(C) \to \pi_1(Y)$ is trivial. As for the second statement, let a and b be vertices of Y_G and Y_H respectively. Since Y

is connected, there is a path in the 1-skeleton of Y joining a to b. If there is a first edge in that path not belonging to Y_G, one of its endpoints is in $Y_G \cap Y_H$. In the opposite case, b belongs to the intersection $Y_G \cap Y_H$ and we are done. □

As a consequence, in the case that $Y_G \cap Y_H$ is connected, it is not difficult to construct complexes with fundamental group G and H departing from Y_G and Y_H.

Lemma 3.4.2. *Suppose that $Y_G \cap Y_H$ is connected. Then, the groups $\pi_1(Y_G/(Y_G \cap Y_H))$ and $\pi_1(Y_H/(Y_G \cap Y_H))$ are isomorphic to G and H respectively.*

Proof. Notice that Y_G and Y_H are connected by Lemma 3.4.1 and that $Y/(Y_G \cap Y_H)$ is homeomorphic to $Y_G/(Y_G \cap Y_H) \vee Y_H/(Y_G \cap Y_H)$, where the wedge point is the class of $Y_G \cap Y_H$ in the quotient space $Y/(Y_G \cap Y_H)$. Since $f(Y_G \cap Y_H) = v$, the map f factors through a continuous map $g : Y_G/(Y_G \cap Y_H) \vee Y_H/(Y_G \cap Y_H) \to X_G \vee X_H$. By the first statement of Lemma 3.4.1, the quotient map $q : Y \to Y/(Y_G \cap Y_H)$ induces an isomorphism on fundamental groups. Now, since f_* is an isomorphism, the induced map on fundamental groups $g_* : \pi_1(Y_G/(Y_G \cap Y_H)) * \pi_1(Y_H/(Y_G \cap Y_H)) \to G * H$ is an isomorphism. Since g_* restricts to morphisms $\pi_1(Y_G/(Y_G \cap Y_H)) \to G$ and $\pi_1(Y_H/(Y_G \cap Y_H)) \to H$, it follows that $\pi_1(Y_G/(Y_G \cap Y_H))$ is isomorphic to G and likewise, that $\pi_1(Y_H/(Y_G \cap Y_H))$ is isomorphic to H. □

Although it is possible to coherently metrize the CW-complex $Y_G/(Y_G \cap Y_H)$, the systole of such space is potentially arbitrarily small, since points that are far in Y_G may become very close in the quotient space. Instead, we can attach for example a cone on the intersection $Y_G \cap Y_H$ and work with the complex $Y_G \cup C(Y_G \cap Y_H)$, which is homotopy equivalent to $Y_G/(Y_G \cap Y_H)$ by standard algebraic topology facts. The next result deals with the general case, in which $Y_G \cap Y_H$ is not necessarily connected.

Lemma 3.4.3. *Let Y_G and Y_H be the simplicial complexes defined above. Then, there exists a 2-dimensional CW-complex $\tilde{Y}$ containing Y, together with an extension $\tilde{f} : \tilde{Y} \to X_G \vee X_H$ of f that also induces isomorphism on fundamental groups such that $\tilde{Y}$ decomposes as the union of subcomplexes $\tilde{Y}_G$, $\tilde{Y}_H$ that intersect in a connected graph. Moreover, the inclusion $\tilde{Y}_G \cap \tilde{Y}_H \hookrightarrow \tilde{Y}$ induces the trivial morphism on fundamental groups.*

Proof. This was essentially proved by Stallings in [80, Theorem 3.2]. We show how to adapt his arguments to our context. Suppose that $Y_G \cap Y_H$ is not connected and take a simplicial path γ in Y joining two different connected components of $Y_G \cap Y_H$. Note that $f \circ \gamma$ determines a loop based at v in $X_G \vee X_H$ (v being the wedge point). Since the map f_* induced on fundamental groups by f is surjective, there is a (simplicial) loop α in Y based at $\gamma(0)$ so that $[f \circ \gamma] = [f \circ \alpha]$ in $\pi_1(X_G \vee X_H, v)$. Hence, if we let $\lambda = \alpha^{-1} * \gamma$, the loop $f \circ \lambda$ is null-homotopic in $X_G \vee X_H$. Clearly, we can write λ as a concatenation $\lambda = \lambda_1 * \lambda_2 * \cdots * \lambda_r$, where each λ_i is a path completely contained in Y_G or Y_H and no two consecutive paths are contained in the same subcomplex. In particular, f maps the endpoints of these paths to v and so each $f \circ \lambda_i$ defines a loop in $X_G \vee X_H$. Since $[f \circ \lambda] = [f \circ \lambda_1] * [f \circ \lambda_2] * \cdots * [f \circ \lambda_r]$ in $\pi_1(X_G \vee X_H, v) = G * H$, there exists an index i such that $f \circ \lambda_i$ is null-homotopic in X_G or X_H according to whether λ_i is contained in Y_G or Y_H. If both endpoints of λ_i happen to be in the same connected component of $Y_G \cap Y_H$ we may replace it by another path λ'_i contained in $Y_G \cap Y_H$ and append it to λ_{i+1} to get a decomposition $\lambda = \lambda_1 * \cdots * \lambda_{i-1} * (\lambda'_i * \lambda_{i+1}) * \cdots * \lambda_r$. By induction on r, there exists some path λ_j joining two different connected components of $Y_G \cap Y_H$ such that $f \circ \lambda_j$ is a null-homotopic loop in $X_G \vee X_H$. Consider the mapping cylinder M_λ of the map $\lambda : I \to Y$. Since $f \circ \lambda$ is homotopic to the constant map v in $X_G \vee X_H$, we may extend f to M_λ using the null-homotopy of λ in the added cylinder 2-cell $e = I \times I$. In particular, the image of e under the extension is completely contained in X_G or X_H and is constantly v in the free part of the boundary ℓ, that is, $\ell = \partial(I \times I) \setminus (I \times 0)$ (see Figure 3.3). Set $\tilde{Y}_G$ as $Y_G \cup e$ or $Y_G \cup \ell$, according to whether $\lambda_i \subseteq Y_G$ or not and define $\tilde{Y}_H$

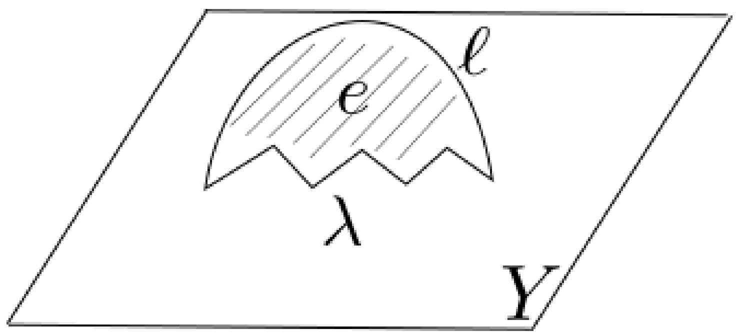

Figure 3.3: Schematic illustration of the mapping cylinder of λ.

analogously. By construction, $\tilde{Y}_G \cap \tilde{Y}_H = (Y_G \cap Y_H) \cup \ell$. After a finite number of steps, we obtain a CW-complex $\tilde{Y}$ of dimension 2, a continuous extension $\tilde{f} : \tilde{Y} \to X_G \vee X_H$ of the map f and subcomplexes $\tilde{Y}_G$, $\tilde{Y}_H$ which cover $\tilde{Y}$ and intersect in a connected 1-dimensional subcomplex. Since the inclusion $Y \hookrightarrow \tilde{Y}$ is a homotopy equivalence, we see that $\tilde{f}$ induces an isomorphism on fundamental groups. Finally, for the last part of the statement notice that $\tilde{f}(\tilde{Y}_G \cap \tilde{Y}_H) = \{v\}$ and since $\tilde{f}_*$ is an isomorphism and the inclusion induced morphism $\pi_1(\tilde{Y}_G \cap \tilde{Y}_H) \to \pi_1(\tilde{Y})$ is trivial. The conclusion follows by the argument in the proof of Lemma 3.4.1. $\square$

Of course, we also need to take care of the area and the systole of the resulting complex $\tilde{Y}$. Let us define the systolic area $\sigma(X, g)$ of a piecewise riemannian polyhedron (X, g) of dimension 2 as the quotient

$$\sigma(X, g) := \frac{\mathrm{Area}(X, g)}{\mathrm{sys}(X, g)^2}.$$

We show next how to endow of a piecewise riemannian metric the mapping cylinders that appear in the proof of the last lemma in a way that the systolic area is almost preserved.

Lemma 3.4.4. *Let (Y, g) be a compact piecewise riemannian polyhedron of dimension 2. Let $\lambda : I \to Y$ be a piecewise linear, non-degenerate map (that is, a map that preserves the dimension of the simplices). Then, given $\varepsilon > 0$, there is a piecewise riemannian metric g' on the mapping cylinder M_λ such that $\sigma(M_\lambda, g') \leq \sigma(Y, g) + \varepsilon$.*

Proof. We need to define a convenient piecewise riemannian metric on the square $I \times I$ which is glued at $I \times 0$ to Y via the map $\lambda : I \to Y$. To do this, start by fixing a sufficiently fine triangulation of Y so that the inclusion of the simplices in Y is an isometric embedding (see [23, Lemma 7.9]). Triangulate $I = I \times 0$ accordingly and endow its 1-simplices with the riemannian metric of their images in Y under λ. Extend this triangulation to $I \times I$ and consider the product metric determined by declaring the interval corresponding to the second coordinate to have the standard metric of an euclidean interval of length δ for some $\delta < \frac{\varepsilon}{\text{length}(\lambda)}$. Thus we obtain a piecewise riemannian metric on $I \times I$ in such a way that $\text{Area}(I \times I) < \varepsilon$. Consider now the equivalence relation in $I \times I$ generated by identifying $(x, 0) \sim (y, 0)$ whenever $\lambda(x) = \lambda(y)$. Since λ is non-degenerate by hypothesis, the quotient map $q : I \times I \to I \times I/\sim$ preserves the dimension of the simplices as well. We declare each simplex $q(\sigma) \in I \times I/\sim$ to be isometric to $\sigma \in I \times I$, so that q is a local isometry. Now, endow M_λ with the piecewise riemannian metric obtained from $I \times I/\sim$ and Y. It is clear that for that metric, $\text{Area}(M_\lambda) < \text{Area}(Y) + \varepsilon$. It remains to estimate the systole of M_λ. First, notice that M_λ is obtained from identifying by a bijective local isometry the subspaces $I \times \{0\}/\sim$ of $I \times I/\sim$ and $\lambda(I)$ of Y. Hence, the resulting piecewise riemannian metric g' on M_λ restricts to the original metric in the interior of each space, that is, in the complement of the subspace where the identification is made (see [23, Chapter I.5] for more on quotient metrics). Let γ be a non-trivial loop in M_λ. Let $p : M_\lambda \to Y$ be the canonical strong deformation retract. Then $p \circ \gamma$ is a non-trivial loop in Y and hence $\text{length}(p \circ \gamma, g) \geq \text{sys}(Y, g)$. Suppose that η is a portion of λ contained in $(I \times I/\sim) \setminus (I \times \{0\}/\sim)$. Then there is a continuous lift $\tilde{\eta}$ of η to $I \times I$. Since the map $q : I \times I \to I \times I/\sim$ is a local isometry, $\text{length}(\tilde{\eta}) = \text{length}(q \circ \tilde{\eta}) = \text{length}(\eta)$. But $\text{length}(\tilde{\eta}) \geq \text{length}(pr(\tilde{\eta}))$, where pr is the projection map to the base of the cylinder. It follows that $\text{length}(\gamma, g') \geq \text{length}(p \circ \gamma, g)$ and hence $\text{sys}(M_\lambda, g') \geq \text{sys}(Y, g)$. This shows that $\sigma(M_\lambda, g') \leq \sigma(Y, g) + \varepsilon$, as desired. □

In conclusion, given an $\varepsilon > 0$, by Lemmas 3.4.3 and 3.4.4 there exists a piecewise riemannian polyhedron (Y, g) that triangulates $G * H$ with $\sigma(Y, g) < \sigma(G * H) + \varepsilon$ and such that the subpolyhedra Y_G and Y_H defined as above intersect in a connected graph. As we remarked in the paragraph following Lemma 3.4.2, under such conditions the complex obtained from Y_G by attaching a cone on $Y_G \cap Y_H$ has fundamental group isomorphic to G. The difficulty resides in endowing the cone $C(Y_G \cap Y_H)$ of a piecewise riemannian metric in such a way that both the area and the systole of the resulting piecewise riemannian polyhedron $Y_G \cup C(Y_G \cap Y_H)$ are controllable. We describe next a condition over the piecewise riemannian metric on the cone that allows to estimate the systole of the whole space.

Call d_{Y_H} the distance induced by the piecewise riemannian metric g on Y_H and consider $d_{Y_H}|_{Y_G \cap Y_H}$. This distance function on the graph $Y_G \cap Y_H$ is typically non-riemannian, meaning that is not induced by any piecewise riemannian metric (to see this, consider for example the distance induced in the 1-sphere by the standard flat riemannian metric on the disk). Let us say that a piecewise riemannian metric h on the cone $C(Y_G \cap Y_H)$ is *homotopy filling* for $(Y_G \cap Y_H, d_{Y_H}|_{Y_G \cap Y_H})$ if the induced distance d_h satisfies $d_h|_{Y_G \cap Y_H} \geq d_{Y_H}|_{Y_G \cap Y_H}$. Notice that such metrics exist. Indeed, notice in first place that the intrinsic distance determined by the piecewise riemannian metric on $Y_G \cap Y_H$ is greater or equal to the induced $d_{Y_H}|_{Y_G \cap Y_H}$. This is simply because there may be paths in Y_H not entirely contained in $Y_G \cap Y_H$ that join two points in the intersection and are shorter than the shortest path within $Y_G \cap Y_H$. Now, to construct a homotopy filling metric on the cone it is enough to endow its 2-simplices with a *round* metric that does not introduce shortcuts. For example, if $Y_G \cap Y_H$ were a topological circle, the cone with this metric would be isometric to the standard round hemisphere. It is not difficult to extend this idea to any finite graph, see [77, Appendix A].

Lemma 3.4.5. *Let h be a homotopy filling piecewise riemannian metric on $C(Y_G \cap Y_H)$ for*

$(Y_G \cap Y_H, d_{Y_H}|_{Y_G \cap Y_H})$. *Then, for the resulting piecewise riemannian metric* g',

$$\mathrm{sys}(Y_G \cup C(Y_G \cap Y_H), g') \geq \mathrm{sys}(Y, g).$$

Proof. Let γ be a non-trivial loop in $Y_G \cup C(Y_G \cap Y_H)$. After a small perturbation, we may suppose that γ intersects transversally $Y_G \cap Y_H$. Let η be a portion of γ contained in $C(Y_G \cap Y_H)$, with distinct endpoints (it they were the same, we could ignore this part and obtain a shorter non-trivial loop γ') in $Y_G \cap Y_H$. Since by hypothesis the distance function on the cone is no less than $d_{Y_H}|_{Y_G \cap Y_H}$, there exists a path $\tilde{\eta}$ in Y_H of length $\leq \mathrm{length}(\eta, h)$. After a finite number of steps, we obtain in this manner a loop $\tilde{\gamma}$ in $Y = Y_G \cup Y_H$ of length at most $\mathrm{length}(\gamma, g')$ such that the restrictions of γ and γ' to Y_G coincide. We have to show that γ' is a non-trivial loop in $\pi_1(Y)$. Notice that the quotient spaces Y/Y_H and $Y_G \cup C(Y_G \cap Y_H)/C(Y_G \cap Y_H)$ are canonically homeomorphic. Moreover the loops $q \circ \gamma$ in $Y_G \cup C(Y_G \cap Y_H)/C(Y_G \cap Y_H)$ and $p \circ \tilde{\gamma}$ in Y/Y_H coincide after applying the canonical homemomorphism $Y/Y_H \to Y_G \cup C(Y_G \cap Y_H)/C(Y_G \cap Y_H)$, where q and p are, respectively, the quotient maps $q : Y_G \cup C(Y_G \cap Y_H) \to Y_G \cup C(Y_G \cap Y_H)/C(Y_G \cap Y_H)$ and $p : Y \to Y/Y_H$. Since q induces an isomorphism on fundamental groups because $C(Y_G \cap Y_H)$ is simply connected, $q \circ \gamma$ and in consequence, $p \circ \tilde{\gamma}$ are non-trivial loops. It follows that the class of $\tilde{\gamma}$ is non-trivial in $\pi_1(Y)$ and therefore, $\mathrm{sys}(Y, g) \leq \mathrm{length}(\tilde{\gamma}, g) \leq \mathrm{length}(\gamma, g')$. The conclusion follows. □

Notice that the only property about the cone $C(Y_G \cap Y_H)$ we required in the proof of the lemma is that it is simply connected. Hence, we could consider different piecewise riemannian polyhedra Z of dimension 2 to "fill" the graph $Y_G \cap Y_H$, as long as these Z contain a homeomorphic copy of the graph, are simply connected and their metric satisfy the property of homotopy filling metrics. We formalize this idea by stating a kind of *homotopy Plateau problem*.

Let P be a finite 1-dimensional complex with a distance d, not necessarily piecewise riemannian. We say that a piecewise riemannian polyhedron (X, g) of dimension 2 is a *homotopy filling* of (P, d) if there is an embedding $\iota : P \to X$ inducing the trivial morphism on fundamental groups such that

$$d(x, y) \leq d_g(\iota(x), \iota(y))$$

for every x, $y \in P$, where d_g is of course the distance induced on X by the piecewise riemannian structure g. By using for instance the mentioned round metric on the cone of P, we see that the minimum area of a homotopy filling (or even of a simply connected homotopy filling) of (P, d) is finite.

In this language, in order to usefully apply Lemma 3.4.5 to estimate the systolic area of $G * H$, we would need to estimate the area of the minimum simply connected homotopy filling of $(Y_G \cap Y_H, d_{Y_H}|_{Y_G \cap Y_H})$. Since the goal is to relate the area of the resulting space to that of $Y = Y_G \cup Y_H$, the round metric on the cone is in general prohibitively large, since its area is proportional to the length of the graph. We do not have an effective bound for the area of (simply connected) homotopy fillings; we speculate below how such an estimate could be established.

The homotopy filling property is related to the *filling volume* of manifolds, a notion defined by Gromov in [44, §2.2]. Let V be a compact null-cobordant manifold with a distance function d (not necessarily riemannian). The *filling volume* $\mathrm{FillVol}(V)$ of V is the minimal volume among the riemannian manifolds (W, g) such that $\partial W = V$ and $d_g|_V \geq d$, that is to say,

$$d_g(x, y) \geq d(x, y)$$

for all x, $y \in V$. Let us call a manifold W a *filling* of V if its boundary is homeomorphic to V. The fillings in the definition should be orientable if the manifold V is orientable. It is shown in [44, Appendix 2] that the filling volume of a null-cobordant manifold of dimension $n \geq 2$ is independent of the topology of the filling. This means that the infimum is realized within the topological type of each filling of V. In contrast, in dimension 1 the filling volume may depend on the topology of the filling (see [44, §2.2.B Counterexamples]).

The value of the filling volume is not known for any riemannian manifold. The long-standing Filling Area Conjecture of Gromov that we stated towards the end of §3.3.1, asserts that the round hemisphere has the minimum area among the fillings of the circle with its standard riemannian metric. It was proved by Gromov that the round hemisphere is the optimal filling of the standard riemannian circle if one restricts to simply connected fillings (this is a consequence of Pu's inequality, Theorem 3.2.1, for the projective plane). This was later extended to fillings of genus at most 1 in [7], but the techniques there do not apply to general surfaces of genus ≥ 2. However, interestingly, the round hemisphere is known to be *coarsely* the optimal filling for the standard riemannian circle, as the following result by Ivanov and Katz shows.

Proposition 3.4.6. *([57, Proposition 3.1]) Let (Σ, g) be a (not necessarily orientable) filling of S^1 such that its distance function d_g restricted to the boundary $\partial\Sigma = S^1$ is greater or equal to the distance function d_{g_0}, where g_0 is the standard riemannian metric on S^1. Then,*

$$\mathrm{Area}(\Sigma, g) \geq \frac{\pi}{4}\,\mathrm{Area}(H),$$

where H is the round hemisphere.

Proof. We may assume that the filling (Σ, g) is isometric, that is, that the distance function $d_g|_{S^1}$ coincides with d_{g_0}. Indeed, we can proceed as in [30, Remark 6.4] to show this. For an arbitrary $\varepsilon > 0$, consider the following riemannian metric g_ε in the collar $S^1 \times I$:

$$g_\varepsilon := tg_0 + (1-t)g|_{S^1} + \varepsilon^2\,dt^2.$$

Thus, g_ε restricts to g_0 on $S^1 \times \{1\}$ and to $g|_{S^1}$ on $S^1 \times \{0\}$, and the area of the collar $(S^1 \times I, g_\varepsilon)$ clearly tends to 0 as $\varepsilon \to 0$. Moreover, since $g|_{S^1} \geq g_0$ by hypothesis, we have that $g_\varepsilon \geq g_0 + \varepsilon^2\,dt^2$, from where it follows that the curves in the collar between points x and y in $S^1 \times \{1\}$ are of length greater than or equal to $d_{g_0}(x, y)$. By gluing the collar to Σ, we obtain an isometric filing $(\Sigma \cup S^1 \times I, \bar{g}_\varepsilon)$ of area arbitrarily close to $\mathrm{Area}(\Sigma, g)$.

Now, fix two points x_0, x_1 in S^1 with $d_{g_0}(x, y) = \frac{\pi}{2}$. Consider the map $f : \Sigma \to \ell_2^\infty = (R^2, \|\cdot\|_{\ell^\infty})$ defined as $f(x) = (d_g(x, x_0), d_g(x, x_1))$. The idea of the proof is that, since the map f is distance non-increasing, the *area* of the image of f should be bounded by the area of Σ. There is however a subtlety here, because $(\mathbb{R}^2, \|\cdot\|_{\ell^\infty})$ is a Finsler manifold and there are several natural but non-equivalent ways to define volumes for such manifolds [56, §4]. Here, we will understand that the area of a Finsler surface is its Ivanov's inscribed riemannian volume (see [56, Example 4.4]), which coincides with the usual notion of volume for riemannian manifolds and subsets of euclidean spaces. We compute now the image of the map f. Notice that, since d_g is isometric to the standard riemannian distance on S^1, the image of $\partial\Sigma = S^1$ is the boundary of an ℓ^1 ball centered at $\left(\frac{\pi}{2}, \frac{\pi}{2}\right)$ of radius $\frac{\pi}{2}$. Since the 1-chain $\partial\Sigma$ is a boundary in Σ and f is continuous, also the 1-chain $f(\partial\Sigma)$ is a boundary in $f(\Sigma) \subseteq \mathbb{R}^2$. Suppose that there is a point y of the ℓ^1 ball of center $\left(\frac{\pi}{2}, \frac{\pi}{2}\right)$ and radius $\frac{\pi}{2}$ not contained in $f(\Sigma)$. Then, the map $f : \Sigma \to \mathbb{R}^2 \setminus \{y\}$ sends the 1-chain $\partial\Sigma$ to a generator of $H_1(\mathbb{R}^2 \setminus \{y\}) = \mathbb{Z}$, which is a contradiction. It follows that $f(\Sigma)$ contains the interior of $f(\partial\Sigma)$ and so

$$\operatorname{Area}(\Sigma, g) \geq \operatorname{Area} f(\Sigma) \geq \frac{\pi^2}{2} = \frac{\pi}{4} \operatorname{Area}(H).$$

$\square$

Interpreted in terms of homotopy fillings, Lemma 3.4.6 states that, if we restrict the fillings to be manifolds, there exists a universal constant $C > 0$ (more precisely, $C \geq \frac{\pi}{4}$) such that for every (manifold) homotopy filling (X, g) of S^1 with its standard riemannian metric,

$$\operatorname{Area}(X, g) \geq C \inf \operatorname{Area}(Z, h),$$

where the infimum is taken over the simply connected (manifold) homotopy fillings (Z, h) of S^1.

By analogy, one may ask if there exists a positive constant relating the area of a minimum simply connected homotopy filling to that of general homotopy fillings for other 1-dimensional complexes and distances. If this was valid in greater generality, for instance, since the inclusion $Y_G \cap Y_H \to Y$ is trivial on fundamental groups, the area of Y would bound up to a constant the area of a minimum simply connected homotopy filling of $(Y_G \cap Y_H, d_Y|_{Y_G \cap Y_H})$. Suppose for a moment that the same conclusion could be obtained for the metric space $(Y_G \cap Y_H, d_{Y_H}|_{Y_G \cap Y_H})$. Then, for a certain simply connected homotopy filling of $(Y_G \cap Y_H, d_{Y_H}|_{Y_G \cap Y_H})$, say, (Z, h), we would have

$$\operatorname{Area}(Y, g) \geq C \operatorname{Area}(Z, h),$$

for a universal constant $C > 0$. By the argument in Lemma 3.4.5, the systolic area of the space $Y_G \cup Z$ with the resulting piecewise riemannian metric g' would satisfy the inequality

$$\sigma(Y_G \cup Z, g') \leq (1 + C)\sigma(Y, g), \tag{3.1}$$

and hence, $\sigma(G) \leq C'\sigma(G * H)$ for $C' = 1 + C$.

However, even if the coarse comparison result between the area of homotopy fillings and simply connected homotopy fillings of the same 1-dimensional space could be established, the proof of inequality (3.1) would still require some work. Namely, the bound from Lemma 3.4.5 applies in principle to $Y_G \cap Y_H$ with the distance induced from Y_H instead of the one induced from Y, which is in general smaller, that is, $d_Y|_{Y_G \cap Y_H} \leq d_{Y_H}|_{Y_G \cap Y_H}$. A possible way to solve this is finding a subgraph $L \leq Y_G \cap Y_H$ such that the inclusion $L \to Y$ is trivial on fundamental groups and the fundamental group of Y_G/L is isomorphic to G. This would intuitively correspond to killing only some of the homotopy classes of loops in $\pi_1(Y_G)$ coming from the inclusion $Y_G \cap Y_H \hookrightarrow Y_G$, the ones needed to quotient out from $\pi_1(Y_G)$ to obtain G.

Chapter 4

A discrete approximation for the systolic area

The *simplicial complexity* is a combinatorial invariant for finitely presentable groups that was recently introduced by Babenko, Balacheff and Bulteau in [4]. Its definition is motivated by the study of the systolic area of groups. As it is shown by the authors in the cited work, the simplicial complexity $\kappa(G)$ of a group G constitutes a good approximation of the systolic area $\sigma(G)$ for large values of $\kappa(G)$. In view of this result, it makes sense to attack problems about the systolic area by analyzing the analogue questions for the simplicial complexity. For instance, one may estimate the systolic area of groups by computing their simplicial complexity or even try to obtain information about structural aspects of the systolic area, such as its behavior under free product of groups, by studying the problem in the context of simplicial complexity.

In the first section we give an introduction to the simplicial complexity of groups and some of its basic properties. The bulk of the chapter is devoted to the partial solution of two problems raised by Babenko, Balacheff and Bulteau, which takes the next three sections. In first place, we compute the simplicial complexity of all surface groups, both in the orientable and in the non-orientable case. Then, using the same techniques we show that $\kappa(G * \mathbb{Z}) = \kappa(G)$ for any surface group G. Philosophically, this result says that an optimal topological model for a group of the form $\pi_1(S) * \mathbb{Z}$, where S is a closed surface, is given by the wedge sum of S with S^1. It also provides the first partial evidence in favor of the conjecture of the stability of the simplicial complexity under free product with free groups. The general stability problem, however, remains wide open. In the last section of the chapter we discuss a potential strategy to approach another instances of the problem, departing from a construction similar to the one described in the final section of Chapter 3.

4.1 The simplicial complexity of groups

Recall that given a group G, we call a simplicial complex a *triangulation* of G if its fundamental group is G. The *simplicial complexity* $\kappa(G)$ of a finitely presented group G is the minimum number of 2-simplices in a triangulation of G (see [4, Definition 2.1]). Here, $\kappa(G)$ should be thought of as a discrete version of area for the group G. Indeed, one may regard a minimal triangulation of G as a geometrical object by declaring each 2-simplex to be the euclidean equilateral triangle of area 1. In this section, we cover the basic properties of the simplicial complexity of groups and give estimations of the invariant for some classes of groups.

The first interesting aspect to explore is the way in which the simplicial complexity measures the complexity of groups. In this sense, the free groups are the simplest groups with respect to this invariant. Indeed, it is clear that $\kappa(F) = 0$ for every finite rank free group F since a bouquet of the appropriate number of 1-spheres triangulates F without using 2-simplices. On the other hand, the simplicial complexity of non-free groups is strictly positive, since the fundamental groups of 1-dimensional complexes are always free. In this aspect, the simplicial complexity is completely analogous to the systolic area. Another salient feature is the subadditivity with respect to the free product of groups: given finitely presentable groups G_1 and G_2, we have

$$\kappa(G_1 * G_2) \leq \kappa(G_1) + \kappa(G_2).$$

For, if K_1, K_2 are triangulations of G_1 and G_2 respectively, the wedge sum $K_1 \vee K_2$ triangulates $G_1 * G_2$. In contrast to the case of the systolic area, it is known that this wedge sum is not, in general, the most effective way to produce a model for $G_1 * G_2$ with respect to the simplicial complexity. To see it, suppose that G_1 and G_2 are not free and consider the model for $G_1 * G_2$ obtained by gluing K_1 and K_2 in a 2-simplex (or more generally, in any simply connected subcomplex of both); this complex has strictly less than $\kappa(G_1) + \kappa(G_2)$ 2-simplices. However, if one of the groups is free it is an open question whether the inequality is strict. In other words: does the identity $\kappa(G * \mathbb{Z}) = \kappa(G)$ hold for any finitely presentable group G? A positive answer for this question would imply that the natural way of extending an optimal triangulation of a group G to a triangulation of $G * \mathbb{Z}$ without adding 2-simplices is the optimal one. In §4.4 we show that this is indeed the case when G is the fundamental group of a closed surface.

The computation of the precise value of the simplicial complexity, as it may be suspected from the definition, is hard even for specific groups. In this respect, there are a few groups for which the simplicial complexity was known, due to Bulteau [25].

Theorem 4.1.1. *The following formulae hold:*

$$\kappa(\mathbb{Z}_2) = 10,\ \kappa(\mathbb{Z} \oplus \mathbb{Z}) = 14,\ \kappa(K_2) = 16,\ \kappa(\mathbb{Z}_3) = 17,$$

where K_2 denotes the fundamental group of the Klein bottle. Moreover, the simplicial complexity of the surface groups $\mathbb{Z}_2$, $\mathbb{Z} \oplus \mathbb{Z}$ and K_2 is attained at optimal triangulations of the corresponding surfaces.

This result is obtained by a combination of some remarks on the local structure of optimal triangulations of groups and an exhaustive analysis of the 2-complexes with at most 17 2-simplices. Bulteau conjectured in [25] that, as it happens in the cases of the projective plane $\mathbb{R}P^2$, the torus $\mathbb{T}^2$ and the Klein bottle, a minimal triangulation of a closed surface S is also an optimal triangulation for its fundamental group $\pi_1(S)$. We verify this conjecture in §4.4.

Although there are, to the best of our knowledge, no other groups for which the simplicial complexity is computed, some estimates for certain classes of groups are available. For example, the precise order of growth of the simplicial complexity of free abelian groups $\kappa(\mathbb{Z}^n)$ with respect to the rank n is known to be n^2 [4, Example 2].

Lemma 4.1.2. *For the simplicial complexity of the group $\mathbb{Z}^n$, the following estimate holds:*

$$\frac{1}{2}n(n-1) \leq \kappa(\mathbb{Z}^n) \leq 7n(n-1).$$

Proof. Notice that the n-dimensional torus $\mathbb{T}^n$ is a $K(\mathbb{Z}^n, 1)$ space. For the lower bound, let K be a 2-dimensional triangulation of $\mathbb{Z}^n$. Since K is the 2-skeleton of an aspherical CW-complex, we see that the free abelian group $H_2(K, \mathbb{Z})$ has rank at least $\binom{n}{2}$, which is the rank of $H_2(\mathbb{T}^n, \mathbb{Z})$. It follows that K must have at least $\binom{n}{2}$ 2-simplices. To prove the upper bound, it is enough to observe that the 2-skeleton of $\mathbb{T}^n$ is formed by the union of $\binom{n}{2}$ 2-dimensional tori $\mathbb{T}^2$. Since it is possible to triangulate $\mathbb{T}^2$ using 14 2-simplices, the result follows. □

The fundamental feature of the simplicial complexity is that it asymptotically approximates the systolic area for non-free groups.

Theorem 4.1.3. *([4, Theorem 1.2]) Let G be a non-free group. Then,*

$$2\pi\sigma(G) \leq \kappa(G) \leq 625(502 \cdot \sigma(G))^{1+\frac{2(1+\ln 5)}{\sqrt{\ln(502\cdot\sigma(G))}}}.$$

We will show the argument for the lower bound, which is elementary. The proof of the upper bound is more involved and it exploits Gromov's covering argument [44, §5.3] (from which the constants in the right-hand side are derived) and some results from [77]. We refer the reader to the original work for that estimate [4, §3].

Proof of Theorem 4.1.3 (lower bound). Let K be a triangulation of G that realizes the simplicial complexity. Endow K with a piecewise riemannian metric h making each edge a segment of length $\frac{2\pi}{3}$ and each 2-simplex a hemisphere of radius 1 with the standard round metric. Thus, the area of (K, h) equals $2\pi\kappa(G)$, since the area of the hemisphere is 2π. Since the arcs of length less than π in the equator of the standard hemisphere are geodesic segments, it is possible to homotope every path in the interior of a 2-simplex to its boundary without increasing the length. Hence, there exists a systolic loop in (K, h) contained in the 1-skeleton and therefore, $\mathrm{sys}(K, h) \geq 2\pi$. Now,

$$\sigma(G) \leq \sigma(K, h) \leq \frac{\kappa(G)}{2\pi},$$

as desired. We remark that the choice of the round metric for the 2-simplices of K is not really essential. It would be possible to obtain a similar inequality for instance endowing each 2-simplex with the metric of the euclidean equilateral triangle of area 1, at the cost however of developing a slightly more difficult argument to estimate the length of paths in K in terms of paths contained in the 1-skeleton. □

4.2 A homological simplification method

The purpose of this section is the proof of a technical lemma which will be used to establish a lower bound of the simplicial complexity for groups of the form $G * T$, where G, T are finitely presentable groups and the cohomology of G satisfies property (A) (see below). The ultimate goal is the computation of the simplicial complexity when G is a surface group and T is a (not necessarily non-trivial) free group. The strategy to achieve this is essentially the same as the one employed in Chapter 2 to find optimal triangulations of the homotopy type of surfaces. The main components of the proof are a version of the Euler characteristic computation in Lemma 2.1.2 (see also the proof of the lower bound in Theorem 2.1.1) and a homological technique that allows to apply this estimate. With this in mind, we develop in this section a homological simplification method for triangulations of groups of the form $G * T$ along the lines of Proposition 2.3.6.

We start by specifying what it means for a group to satisfy property (A). By the cohomology ring of a group G we will refer to its cohomology as a discrete group, i.e. the cohomology of an Eilenberg-MacLane space $K(G,1)$, while $H^*(G,\mathbb{Z}_2)$ denotes the (reduced) cohomology of G with coefficients in $\mathbb{Z}_2$, as usual. We say that the cohomology ring of a group satisfies property (A) whenever the cohomology ring of a $K(G,1)$ space does (see Definition 2.3.2).

Example 4.2.1. Analogously to the case of surfaces, surface groups (orientable or non-orientable) satisfy property (A) by Poincaré Duality. More generally, any one-relator group G with $H_2(G,\mathbb{Z}_2) = \mathbb{Z}_2$ and with a non-degenerate cup product form $H^1(G,\mathbb{Z}_2) \times H^1(G,\mathbb{Z}_2) \to H^2(G,\mathbb{Z}_2) = \mathbb{Z}_2$, satisfies property (A). Using the computation of the cohomology ring of one-relator groups of [75], one may obtain many additional examples of such groups. As concrete examples, the Baumslag-Solitar groups $BS(m,n)$ satisfy property (A) whenever m and n are odd.

Suppose that K is a 2-dimensional triangulation of $G*T$, that is, $\pi_1(K) = G*T$. By standard algebraic topology arguments, K includes as the 2-skeleton of an aspherical CW-complex X, which is possibly infinite dimensional. Since the fundamental group of X is isomorphic to $G*T$, X is an Eilenberg-MacLane space $K(G*T,1)$. By a theorem of Whitehead (see for example [24, Theorem 7.3]), the wedge sum of a $K(G,1)$ space K_G and a $K(T,1)$ space K_T is aspherical, which implies that X is homotopy equivalent to $K_G \vee K_T$. As stated before, we aim to estimate the number of 2-simplices of triangulations of $G*T$ for G a surface group and T a free group. Hence, since informally speaking G should be responsible for the complexity of the triangulation (that is, the number of 2-simplices), the first objective is to isolate a subcomplex of K containing only the 2-dimensional homology classes of K that correspond to classes in $H_2(G,\mathbb{Z}_2)$. We give a definition to formalize this idea.

Definition 4.2.2. Let X be a CW-complex of dimension at least 2, together with a homotopy equivalence $h : X \to K_G \vee K_T$, where K_G, K_T are defined as above. Suppose that its 2-skeleton $X^{(2)}$ has the structure of a finite simplicial complex. Let $M \leq X^{(2)}$ be a (simplicial) subcomplex satisfying the following properties:

1. The inclusion $i : M \hookrightarrow X$ induces isomorphisms $i_* : H_n(M,\mathbb{Z}_2) \to H_n(X,\mathbb{Z}_2)$ for $n < 2$.

2. The composition $H_2(M,\mathbb{Z}_2) \xrightarrow{i_*} H_2(X,\mathbb{Z}_2) \equiv H_2(K_G,\mathbb{Z}_2) \oplus H_2(K_T,\mathbb{Z}_2) \xrightarrow{p} H_2(K_G,\mathbb{Z}_2)$ is an epimorphism, where p is the projection and the isomorphism $H_2(X,\mathbb{Z}_2) \equiv H_2(K_G,\mathbb{Z}_2) \oplus H_2(K_T,\mathbb{Z}_2)$ is induced by h.

We will say that such a subcomplex M is *homologically G-full* with respect to h, or simply *homologically G-full* if the homotopy equivalence h is clear from the context.

The next result states, roughly, that we can kill the "extra" homology classes in $H_2(X,\mathbb{Z}_2)$ one at a time (compare to Proposition 2.3.6).

Lemma 4.2.3. *Let X be a CW-complex of dimension at least 2 homotopy equivalent to a space of the form $K_G \vee K_T$ and such that its 2-skeleton $X^{(2)}$ is a finite simplicial complex. Let $M \leq X^{(2)}$ be a homologically G-full subcomplex. If* $\dim H_2(M,\mathbb{Z}_2) > \dim H_2(G,\mathbb{Z}_2)$*, there exists a 2-simplex $\sigma \in M$ such that $M \backslash \sigma$ is homologically G-full. Moreover,* $\dim H_2(M\backslash\sigma,\mathbb{Z}_2) = \dim H_2(M,\mathbb{Z}_2) - 1$.

Proof. Since by hypothesis $\dim H_2(M,\mathbb{Z}_2) > \dim H_2(K_G,\mathbb{Z}_2)$ there is a non-trivial class B in the kernel of the linear map $p \circ i_* : H_2(M,\mathbb{Z}_2) \to H_2(K_G,\mathbb{Z}_2)$. Let σ be a 2-simplex of M in the support of B. The topological boundary $\partial\sigma$ viewed as a chain in $C_1(M \setminus \sigma,\mathbb{Z}_2)$ is the boundary of the 2-chain $B-\sigma$. Hence the inclusion induces the zero morphism $H_1(\partial\sigma,\mathbb{Z}_2) \to H_1(M\setminus\sigma,\mathbb{Z}_2)$. It follows that the inclusion $M\setminus\sigma \hookrightarrow M$ induces isomorphisms $H_n(M\setminus\sigma,\mathbb{Z}_2) \to H_n(M,\mathbb{Z}_2)$ for $n<2$. It remains to verify the surjectivity of $p\circ j_* : H_2(M\setminus\sigma,\mathbb{Z}_2) \to H_2(K_G,\mathbb{Z}_2)$, where j is the inclusion $j : M\setminus\sigma \hookrightarrow X$. Let $[Z]$ be a class in $H_2(K_G,\mathbb{Z}_2)$. By hypothesis, there is some class $C \in H_2(M,\mathbb{Z}_2)$ such that $p\circ i_*[C] = [Z]$. If σ does not belong to the support of C, when viewed as a class in $H_2(M\setminus\sigma,\mathbb{Z}_2)$ we have $p\circ j_*[C] = [Z]$. In the other case, consider the 2-chain $C+B$. Since the coefficients are taken in $\mathbb{Z}_2$, this chain is a well defined 2-cycle in $M\setminus\sigma$ and $p\circ j_*[C+B] = p\circ i_*[C] + p\circ i_*[B] = p\circ i_*[C] = [Z]$. Hence, in any case $p\circ j_* : H_2(M\setminus\sigma,\mathbb{Z}_2) \to H_2(K_G,\mathbb{Z}_2)$ is an epimorphism. The fact that $\dim H_2(M\setminus\sigma,\mathbb{Z}_2) = \dim H_2(M,\mathbb{Z}_2) - 1$ follows immediately from the Euler characteristic, since $\chi(M\setminus\sigma) = \chi(M) - 1$. □

We introduce some convenient notation before moving on.

Notation. Given a finitely presentable group G, we will denote by $\overline{\chi}(G)$ the *2-truncated* Euler characteristic of G, that is $\overline{\chi}(G) := \dim H_2(G,\mathbb{Z}_2) - \dim H_1(G,\mathbb{Z}_2) + \dim H_0(G,\mathbb{Z}_2)$.

It remains now to combine the homological simplification method with the control provided by the property (A) to obtain, given a triangulation of a group of the form $G * T$, a subcomplex to which the Euler characteristic estimate applies. The next result accomplishes this task.

Lemma 4.2.4. *Let K be a (finite) connected simplicial complex of dimension 2 with fundamental group isomorphic to $G*T$, and suppose that the cohomology ring of G satisfies property (A). Then, there is another simplicial complex L of dimension at most 2 with no more 2-simplices than K such that $\chi(L) \le \overline{\chi}(G)$, $\dim H_2(L,\mathbb{Z}_2) = \dim H_2(G,\mathbb{Z}_2)$ and every edge of L is the face of at least two 2-simplices.*

Proof. Let X be an Eilenberg-MacLane space $K(G*T,1)$ such that $X^{(2)} = K$. Then there is a map $i : K \to K_G \vee K_T$ inducing an isomorphism in H_n for $n=0,1$ and an epimorphism in H_2, where K_G and K_T are respectively a $K(G,1)$ and a $K(T,1)$ space as before. Since the projection $H_2(X,\mathbb{Z}_2) \equiv H_2(K_G,\mathbb{Z}_2) \oplus H_2(K_T,\mathbb{Z}_2) \to H_2(K_G,\mathbb{Z}_2)$ is surjective, K is a homologically G-full subcomplex of X. By applying inductively Lemma 4.2.3, we obtain a subcomplex M of K that is homologically G-full and such that $\dim H_2(M,\mathbb{Z}_2) = \dim H_2(G,\mathbb{Z}_2)$. After collapsing the free faces of M, we may assume that M has no edge that is the face of a unique 2-simplex. Suppose there is a maximal edge $e = \{a,b\}$ in M (otherwise we are done, since we may take the desired complex L as M). If there is no path between a and b in $M\setminus e$, the quotient M/e has a natural structure of simplicial complex with one less maximal edge than M. If, on the contrary, a and b are joined by some path in $M\setminus e$, M is homotopy equivalent to a CW-complex of the form $Z \vee S^1$, where Z is the complex $M\setminus e$ and the S^1 results from attaching a 1-cell by a map that sends both vertices to $a \in Z$. After applying, if needed, finitely many of these moves, we get a CW-complex of the form $L \vee \bigvee_{i=1}^m S^1$ homotopy equivalent to M, where L is a simplicial complex formed by the 2-simplices of M (and hence, with no more 2-simplices than K) in which every edge is the face of at least two 2-simplices. It remains to verify the bound on the Euler characteristic of L. Since $L \vee \bigvee_{i=1}^{m} S^1$ is homotopy equivalent to M, clearly

$$\chi(M) = \chi(L \vee \bigvee_{i=1}^{m} S^1) = \chi(L) - m.$$

On the other hand, by construction $\chi(M) = \overline{\chi}(G) - \dim H_1(T, \mathbb{Z}_2)$, since $\dim H_2(M, \mathbb{Z}_2) = \dim H_2(G, \mathbb{Z}_2)$ and the first homology group of M is isomorphic to $H_1(K_G \vee K_T, \mathbb{Z}_2) = H_1(G, \mathbb{Z}_2) \oplus H_1(T, \mathbb{Z}_2)$. Now, by composing with the homotopy equivalence $L \vee \bigvee_{i=1}^{m} S^1 \simeq M$ we obtain a map $f : L \vee \bigvee_{i=1}^{m} S^1 \to K_G \vee K_T$ which induces an isomorphism in H_n for $n = 0, 1$, and such that $p \circ f_* : H_2(L \vee \bigvee_{i=1}^{m} S^1) = H_2(L, \mathbb{Z}_2) \to H_2(K_G, \mathbb{Z}_2)$ is an epimorphism. In particular, dualizing we get an isomorphism

$$H^1(K_G \vee K_T, \mathbb{Z}_2) = H^1(G, \mathbb{Z}_2) \times H^1(T, \mathbb{Z}_2) \to H^1(L \vee \bigvee_{i=1}^{m} S^1, \mathbb{Z}_2) = H^1(L, \mathbb{Z}_2) \times H^1(\bigvee_{i=1}^{m} S^1, \mathbb{Z}_2).$$

Let $(0, a) \in H^1(L, \mathbb{Z}_2) \times H^1(\bigvee_{i=1}^{m} S^1, \mathbb{Z}_2)$ be a non-trivial class and suppose that $(\alpha, \delta) \in H^1(G, \mathbb{Z}_2) \times H^1(T, \mathbb{Z}_2)$ is the unique class such that $f^*(\alpha, \delta) = (0, a)$. We claim that $\alpha = 0$. Indeed, suppose that it was not the case. Then, since the cohomology ring of G satisfies property (A) there is a class $\beta \in H^1(G, \mathbb{Z}_2)$ with $\alpha \cup \beta \neq 0$. Consider the class $f^*((\alpha, \delta) \cup (\beta, 0)) = f^*(\alpha \cup \beta, 0) \in H^2(L, \mathbb{Z}_2) = H^2(L, \mathbb{Z}_2) \times H^2(\bigvee_{i=1}^{m} S^1, \mathbb{Z}_2)$. It is non-trivial: take a class $\lambda \in H_2(G, \mathbb{Z}_2)$ such that $(\alpha \cup \beta)\lambda \neq 0$ (here we use the identification $H^2(G, \mathbb{Z}_2) = \mathrm{Hom}(H_2(G), \mathbb{Z}_2)$). Since M is homologically G-full, there is some class $\gamma \in H_2(L, \mathbb{Z}_2) \equiv H_2(M, \mathbb{Z}_2)$ such that $f_*(\gamma) = (\lambda, \eta)$, for some $\eta \in H_2(T)$. Then,

$$f^*(\alpha \cup \beta, 0)\gamma = (\alpha \cup \beta, 0) f_*(\gamma) = (\alpha \cup \beta, 0)(\lambda, \eta) \neq 0.$$

On the other hand, from the identity

$$f^*((\alpha, \delta) \cup (\beta, 0)) = (0, a) \cup f^*(\beta, 0) = 0$$

we obtain a contradiction, proving the claim. We conclude that the inverse of the map

$$f^* : H^1(G, \mathbb{Z}_2) \times H^1(T, \mathbb{Z}_2) \to H^1(L, \mathbb{Z}_2) \times H^1(\bigvee_{i=1}^{m} S^1, \mathbb{Z}_2)$$

restricts to a monomorphism $H^1(\bigvee_{i=1}^{m} S^1, \mathbb{Z}_2) \to H^1(T, \mathbb{Z}_2)$. Hence, $m \leq \dim H_1(T, \mathbb{Z}_2)$ and, since $\chi(L) = \overline{\chi}(G) - (\dim H_1(T, \mathbb{Z}_2) - m)$, the result follows. $\square$

4.3 A lower bound for simplicial complexity

In this short section we employ the previous results to prove the announced lower bound on the simplicial complexity for groups of the form $G * T$, where G, T are finitely presentable groups and G satisfies property (A).

The same argument that shows that 2-complexes with a dense triangulation admit a lower bound on the number of vertices in terms of its Euler characteristic may be applied to obtain a bound on the number of 2-simplices. In what follows we will understand that a simplicial complex of dimension 2 is of strict dimension 2, i.e. it has at least one 2-simplex.

Lemma 4.3.1. *(cf. Lemma 2.3.1). Let L be a connected simplicial complex of dimension 2 in which every edge is the face of at least two 2-simplices. Then, if $\chi(L) \leq 2$, the complex L has at least $\rho(L)$ vertices and at least $2\rho(L) - 2\chi(L)$ 2-simplices.*

Recall that here $\rho(L)$ stands for $\rho(\chi(L)) = \left\lceil \frac{7+\sqrt{49-24\chi(L)}}{2} \right\rceil$ as defined in Notation 2.1.

Proof. The lower bound $\rho(L) = \rho(\chi(L))$ for the number of vertices is precisely the conclusion of Lemma 2.3.1. For the estimate on the number of 2-simplices, notice that the Euler characteristic formula for L together with the inequality $3\alpha_2(L) \geq 2\alpha_1(L)$ imply that

$$2\chi(L) \geq 2\alpha_0(L) - 2\alpha_2(L),$$

from where the claimed bound $\alpha_2(L) \geq 2\rho(L) - 2\chi(L)$ follows immediately. □

An explicit formula for the lower bound on the simplicial complexity of groups of the form $G * T$ with G satisfying property (A) follows almost directly from Lemmas 4.2.4 and 4.3.1.

Theorem 4.3.2. *Let G, T be finitely presentable groups. If G satisfies property (A), $\overline{\chi}(G) \leq 2$, and* $\dim H_2(G) > 0$*, then $\kappa(G * T) \geq 2\rho(\overline{\chi}(G)) - 2\overline{\chi}(G)$.*

Proof. Let K be a simplicial complex of dimension 2 with fundamental group isomorphic to $G * T$. Since G satisfies property (A), from Lemma 4.2.4 we obtain a simplicial complex L with $\alpha_2(L) \leq \alpha_2(K)$, $\chi(L) \leq \overline{\chi}(G)$ and such that every edge of L is in at least two 2-simplices. Furthermore, there is an epimorphism $H_2(L) \to H_2(G)$, so that $\dim H_2(L) > 0$ and hence L is of dimension 2. By Lemma 4.3.1, L has at least $2\rho(L) - 2\chi(L)$ 2-simplices. Now, since $\chi(L) \leq \overline{\chi}(G)$ and ρ is a non-increasing function, we conclude that $\alpha_2(L) \geq 2\rho(\overline{\chi}(G)) - 2\overline{\chi}(G)$, as desired. □

We may apply Theorem 4.3.2 to the one-relator groups from Example 4.2.1. For instance, the theorem gives the lower bound $\kappa(BS(m,n)) \geq 14$ for Baumslag-Solitar groups with m, n odd since $\chi(BS(m,n)) = 0$. We know that this bound is not sharp except for the fundamental group of the torus $\mathbb{Z} \oplus \mathbb{Z} = BS(1,1)$. But one would expect stronger lower bounds for one-relator groups with a large number of generators (and hence, small Euler characteristic). As we will see in the next section, the lower bound from Theorem 4.3.2 is sharp for the fundamental groups of surfaces.

4.4 The simplicial complexity of surface groups

This section contains the main result of the chapter: the computation of the simplicial complexity $\kappa(\pi_1(S) * F)$, where S is a closed surface and F is a finite rank free group (possibly trivial). Recall that $\delta(S)$ denotes the number of 2-simplices in a minimal triangulation of a closed surface S and that an explicit formula for these numbers was given by Jungerman and Ringel (see Theorem 2.1.1). Concretely, for non exceptional surfaces, the following identity holds

$$\delta(S) = 2\rho(\chi(S)) - 2\chi(S).$$

For the exceptional cases M_2, N_2 and N_3, it is necessary to replace $\delta(S)$ by $\delta(S) - 2$ in the formula.

As it was observed in Example 4.2.1, the fundamental group of a non-simply connected closed surface S satisfies property (A). Hence, we may apply Theorem 4.3.2 to groups of the form $\pi_1(S) * T$ obtaining the following corollary.

Proposition 4.4.1. *Let S be a non-simply connected closed surface. Then $\kappa(\pi_1(S) * T) \geq 2\rho(\chi(S)) - 2\chi(S)$. In particular, if S is non-exceptional, then $\kappa(\pi_1(S) * T) \geq \delta(S)$.*

Proof. By Theorem 4.3.2, we have that $\kappa(\pi_1(S) * T) \geq 2\rho(\overline{\chi}(\pi_1(S))) - \overline{\chi}(\pi_1(S))$. Hence, it is enough to see that $\overline{\chi}(\pi_1(S)) = \chi(S)$ for all non-simply connected closed surfaces S. The identity is clear for surfaces S different from the real projective plane $\mathbb{R}P^2$ since these surfaces are aspherical. For $\mathbb{R}P^2$, notice that the infinite real projective space $\mathbb{R}P^\infty$ is an Eilenberg-MacLane space for $\pi_1(\mathbb{R}P^2) = \mathbb{Z}_2$, so we have $\overline{\chi}(\pi_1(S)) = \chi(S)$ also for $S = \mathbb{R}P^2$. □

Remark 4.4.2. It can actually be proved for a group G in the conditions of Theorem 4.3.2 that the equality

$$\kappa(G) = 2\rho(\overline{\chi}(G)) - 2\overline{\chi}(G)$$

holds only for G the fundamental group of a non-exceptional surface. Indeed, take an optimal triangulation K of G, that is, one with exactly $\kappa(G) = 2\rho(\overline{\chi}(G)) - 2\overline{\chi}(G)$ 2-simplices. We know that $3\alpha_2(K) \geq 2\alpha_1(K)$ and $\chi(K) \geq \overline{\chi}(G)$. Hence,

$$\kappa(G) = \alpha_2(K) \leq 2\alpha_0(K) - 2\chi(K) \leq 2\rho(\overline{\chi}(G)) - 2\overline{\chi}(G) = \kappa(G).$$

It follows that all the intermediate inequalities are actually identities and in particular, $3\alpha_2(K) = 2\alpha_1(K)$. This implies that K is a pseudosurface and so $G = \pi_1(K)$ is isomorphic to the free product of $\pi_1(S)$ with a free group F for some surface S. Since G satisfies property (A), the free part F is trivial and so $G = \pi_1(S)$ as we wanted to show.

It remains to handle the exceptional cases. Observe that for an exceptional surface S (i.e. $S = N_2, N_3$ or M_2), Proposition 4.4.1 provides the lower bound $\kappa(\pi_1(S)*T) \geq 2\rho(\chi(S)) - 2\chi(S)$, which is slightly weaker than required because $\delta(S) = 2\rho(\chi(S)) - 2\chi(S) + 2$ in these cases. So, for the exceptional surfaces, we will need to refine the proof of the lower bound of Theorem 4.3.2.

Lemma 4.4.3. *Let S be an aspherical closed surface (either exceptional or non-exceptional) and let K be a connected simplicial complex of dimension 2 with fundamental group isomorphic to $\pi_1(S) * T$. Let L be the simplicial complex obtained from K by applying Lemma 4.2.4. If $\chi(L) = \chi(S)$, then there is a continuous map $L \to S$ that induces an isomorphism in (co)homology.*

Proof. From the proof of Lemma 4.2.4 applied to K, we obtain a continuous map $f : L \vee \bigvee_{i=1}^m S^1 \to K_{\pi_1(S)} \vee K_T \simeq S \vee K_T$ (since S is aspherical) which induces an isomorphism in H_n for $n = 0, 1$ and an epimorphism $p \circ f_* : H_2(L \vee \bigvee_{i=1}^m S^1, \mathbb{Z}_2) = H_2(L, \mathbb{Z}_2) \to H_2(S, \mathbb{Z}_2)$, and such that $\dim H_2(L, \mathbb{Z}_2) = \dim H_2(S, \mathbb{Z}_2)$. Consider the natural map $g : L \to S$ defined as the composition $L \hookrightarrow L \vee \bigvee_{i=1}^m S^1 \xrightarrow{f} S \vee K_T \to S$, where the first map is the inclusion and the last one is the projection to the quotient. Since the quotient map $S \vee K_T \to S$ induces the projection $H_*(S, \mathbb{Z}_2) \oplus H_*(K_T, \mathbb{Z}_2) \to H_*(S, \mathbb{Z}_2)$ in homology, g induces an isomorphism in H_2 and so it suffices to show $g_* : H_1(L, \mathbb{Z}_2) \to H_1(S, \mathbb{Z}_2)$ is an isomorphism. Note that from the proof of Lemma 4.2.4 it follows that the inverse of $f^* : H^1(K_{\pi_1(S)}, \mathbb{Z}_2) \times H^1(K_T, \mathbb{Z}_2) \to H^1(L, \mathbb{Z}_2) \times H^1(\bigvee_{i=1}^m S^1, \mathbb{Z}_2)$ restricts to a monomorphism $h : H^1(\bigvee_{i=1}^m S^1, \mathbb{Z}_2) \to H^1(K_T, \mathbb{Z}_2)$. More concretely, the monomorphism

h sends a class $a \in H^1(\bigvee_{i=1}^m S^1, \mathbb{Z}_2)$ to the unique class $\alpha \in H^1(K_T, \mathbb{Z}_2)$ such that $f^*(0, \alpha) = (0, a)$. On the other hand, $\chi(L) = \chi(S) - (\dim H_1(K_T, \mathbb{Z}_2) - m)$ and since $\chi(L) = \chi(S)$ by assumption, $m = \dim H_1(K_T, \mathbb{Z}_2)$ and so h is an isomorphism. It is clear then that also f^* restricts in an analogous way to an isomorphism $H^1(K_T, \mathbb{Z}_2) \to H^1(\bigvee_{i=1}^m S^1, \mathbb{Z}_2)$. Now, the morphism induced by g between the first cohomology groups coincides with the composition $H^1(S, \mathbb{Z}_2) \hookrightarrow H^1(S, \mathbb{Z}_2) \times H^1(K_T, \mathbb{Z}_2) \xrightarrow{f^*} H^1(L, \mathbb{Z}_2) \times H^1(\bigvee_{i=1}^m S^1, \mathbb{Z}_2) \to H^1(L, \mathbb{Z}_2)$, the first arrow being the inclusion and the last one the projection on the first coordinate. Using the fact that f^* restricts to an isomorphism between the second factors, it is not difficult to check that the described map between the first factors $H^1(S, \mathbb{Z}_2) \to H^1(L, \mathbb{Z}_2)$ is an isomorphism. Since $H^1(L, \mathbb{Z}_2)$ and $H^1(S, \mathbb{Z}_2)$ are vector spaces of finite dimension over $\mathbb{Z}_2$, we conclude that g induces an isomorphism in (co)homology as desired. □

From this point, the arguments are very similar to the ones employed to deal with the covering type of the exceptional surfaces. The biggest difference is that there is not enough room to have an exceptional case for the simplicial complexity.

Proposition 4.4.4. *Let $S = N_2, N_3$ or M_2. Then $\kappa(\pi_1(S) * T) \geq \delta(S)$.*

Proof. Let K be a triangulation of $\pi_1(S) * T$. From Lemma 4.2.4, and keeping the notations of the proof of Theorem 4.3.2, we obtain a complex L with $\alpha_2(L) \leq \alpha_2(K)$, $\chi(L) \leq \overline{\chi}(\pi_1(S)) = \chi(S)$ and such that every edge of L is in at least two 2-simplices. By Lemma 4.3.1, this implies that L has at least $\rho(L) \geq \rho(S)$ vertices and at least $2\rho(L) - 2\chi(L)$ 2-simplices. Note that if any of the strict inequalities $\alpha_0(L) > \rho(S)$, $\chi(L) < \chi(S)$ holds, we have

$$\alpha_2(L) \geq 2\rho(S) - 2\chi(S) + 2 = \delta(S)$$

and there is nothing to prove. In view of this, in what follows we will suppose that $\alpha_0(L) = \rho(S)$ and $\chi(L) = \chi(S)$. Observe that the homology of L is isomorphic to the homology of S via a continuous map $L \to S$ by Lemma 4.4.3. Also, since $3\alpha_2(L) \geq 2\alpha_1(L)$, by the Euler characteristic formula for L we have

$$3(\alpha_0(L) - \chi(L)) \leq \alpha_1(L) \leq \binom{\alpha_0(L)}{2}.$$

We solve first the case $S = N_2$. By our assumption, we have that $\chi(L) = \chi(N_2) = 0$ and $\alpha_0(L) = \rho(N_2) = 7$. Hence, from the above inequality we learn that $\alpha_1(L) = 21$ and, since $\chi(L) = 0$, $\alpha_2(L) = 14$. Thus $3\alpha_2(L) = 2\alpha_1(L) = 42$, from where it follows that every edge of L is the face of exactly two 2-simplices. Since there is a map $L \to S$ inducing an isomorphism in homology, by Proposition 2.4.1 L would be homeomorphic to N_2 contradicting Theorem 2.1.1. Hence, $\alpha_0(L) > \rho(N_2)$ or $\chi(L) < \chi(S)$ and consequently $\alpha_2(L) \geq \delta(N_2)$.

For the surface $S = N_3$, we know that $\chi(L) = \chi(N_3) = -1$ and $\alpha_0(L) = \rho(N_3) = 8$. Hence,

$$3(\alpha_0(L) - \chi(L)) = 27 \leq \alpha_1(L) \leq 28 = \binom{\alpha_0(L)}{2}.$$

Suppose first that $\alpha_1(L) = 27$, so that $\alpha_2(L) = 18$. Hence every edge of L is the face of exactly two 2-simplices and from Proposition 2.4.1, L is homeomorphic to N_3 in contradiction to Theorem 2.1.1. Then, $\alpha_1(L) = 28$. In that case, $\alpha_2(L) = 19$ and since $57 = 3\alpha_2(L) = 2\alpha_1(L) + 1$, every edge of L is in two 2-simplices except for one that is the face of three 2-simplices of L. The link of a vertex of this edge is a graph in which every vertex has degree two except for one that has degree three. This is impossible because the sum of the degrees of an undirected graph is even. Therefore, $\alpha_0(L) > \rho(N_3)$ or $\chi(L) < \chi(N_3)$ and hence $\alpha_2(L) \geq \delta(N_3)$ as claimed.

Finally, when $S = M_2$ we have that $\chi(L) = \chi(M_2) = -2$ and $\alpha_0(L) = \rho(M_2) = 9$. In this case, we know that $\alpha_2(L) \geq 22 = \delta(M_2) - 2$ and we want to show that L has at least $\delta(M_2) = 24$ 2-simplices. We will see that the cases $\alpha_2(L) = 22$, $\alpha_2(L) = 23$ are not possible. Suppose first that $\alpha_2(L) = 22$. Then, by the Euler characteristic formula, $\alpha_1(L) = 33$. Therefore, every edge of L is the face of exactly two 2-simplices of L and so, by Proposition 2.4.1 L should be homeomorphic to M_2, which contradicts Theorem 2.1.1. If $\alpha_2(L) = 23$, it is $\alpha_1(L) = 34$, whence $69 = 3\alpha_2(L) = 2\alpha_1(L) + 1$. It follows that every edge of L is the face of exactly two 2-simplices except for one which is the face of three 2-simplices. The same argument as before shows that this is impossible. We conclude that $\alpha_2(L) \geq \delta(M_2)$. □

The main results of the section follow as corollaries to this last proposition.

Corollary 4.4.5. *Let S be a non-simply connected closed surface. Then $\kappa(S) = \delta(S)$.*

Proof. The upper bound $\kappa(\pi_1(S)) \leq \delta(S)$ is clear, while the lower bound follows from Propositions 4.4.1 and 4.4.4. □

Note that, as a consequence of this result, the simplicial complexity of surface groups grows linearly on the genus. This was observed, in the orientable case, in [4, Example 2].

Corollary 4.4.6. *Let S be a non-simply connected closed surface and let T be a finitely presentable group. Then, $\kappa(\pi_1(S) * T) \geq \kappa(\pi_1(S))$. In particular, $\kappa(\pi_1(S) * \mathbb{Z}) = \kappa(\pi_1(S))$.*

Proof. Let T be a finitely presentable group. By Propositions 4.4.1 and 4.4.4, $\kappa(\pi_1(S) * T) \geq \delta(S)$ and since $\kappa(\pi_1(S)) = \delta(S)$ by Corollary 4.4.6, the first claim holds. For the second one, it is enough to observe that the upper bound $\kappa(G * \mathbb{Z}) \leq \kappa(G)$ holds trivially for every finitely presentable group G. □

To sum up, we compile the last results about the simplicial complexity of surface groups in the next theorem.

Theorem 4.4.7. *The simplicial complexity $\kappa(\pi_1(S))$ of the fundamental group of a non-simply connected closed surface S coincides with the minimum number of 2-simplices in an optimal triangulation of S. Furthermore, the simplicial complexity of surface groups is stable under free product with free groups, that is, $\kappa(\pi_1(S) * \mathbb{Z}) = \kappa(\pi_1(S))$.*

4.5 The simplicial complexity of a free product of groups

We have just seen that the simplicial complexity is stable under taking free product with free groups for surface groups. In this section we take another path to study the simplicial complexity of a free product of general groups, using as a starting point the construction described in §3.4. Similarly as in the case of the systolic area, it would be desirable to obtain a lower estimate on the simplicial complexity $G * H$ in terms of $\kappa(G)$ and $\kappa(H)$. To be more specific, it seems plausible that the inequality $\kappa(G * H) \geq \max\{\kappa(G), \kappa(H)\}$ holds. Unfortunately, in the general case we have no definitive results about the simplicial complexity of a free product of groups. Rather, we intend to present a possible strategy for proving that inequality, indicating the missing parts.

The initial setup is almost identical to the setup in §3.4. Let G and H be finitely presentable groups and take a triangulation K of $G * H$ that realizes the simplicial complexity $\kappa(G * H)$. As explained in §3.4, if X_G and X_H are triangulations of G and H, respectively, there is a simplicial map $f : K' \to X_G \vee X_H$ that induces an isomorphism on fundamental groups for an appropriate subdivision K' of K. Define the subcomplex K_G of K' as $K_G := f^{-1}(X_G)$ and let K_H be the simplicial closure of the complement of K_G in K', so that $K_G \cup K_H = K'$ and $K_G \cap K_H$ is 1-dimensional. In Lemma 3.4.3 we showed a way to obtain a subcomplex of K' with fundamental group isomorphic to G formed by K_G, the cone on $K_G \cap K_H$ and certain additional 2-cells. In the present context we modify the output of that result to avoid the need of attaching extra 2-cells.

Lemma 4.5.1. *Let K_G and K_H be the simplicial complexes defined above. Then, there exists a 1-dimensional simplicial complex M such that the simplicial complex $K_G \cup C(K_G \cap K_H) \cup M$ is connected and has fundamental group isomorphic to G. Here, by $C(K_G \cap K_H)$ we mean the complex formed by a disjoint union of cones, exactly one over each connected component of $K_G \cap K_H$.*

Proof. If the intersection of K_G and K_H is connected, we may apply Lemma 3.4.2 and we are done. In the opposite case, analogously as in the proof of Lemma 3.4.3, there is a simplicial path λ in K' joining two different connected components of $K_G \cap K_H$ whose image $f \circ \lambda$ is a null-homotopic loop based in v and completely contained in X_G or X_H. Just as in Lemma 3.4.3, form the mapping cylinder M_λ of $\lambda : I \to K'$ and extend f to M_λ using the homotopy that contracts λ to v. If e is the cylindrical new 2-cell and ℓ is the part of the boundary of e which is not identified with λ, define $\tilde{K}_G$ as $K_G \cup e$ or $K_G \cup \ell$, depending on whether $\lambda_i \subseteq K_G$ or not. Define $\tilde{K}_H$ in the same way and notice that $\tilde{K}_G \cap \tilde{K}_H = (K_G \cap K_H) \cup \ell$.

After a finite number of steps, we obtain a CW-complex $\tilde{K}$ of dimension 2, an extension $\tilde{f} : \tilde{K} \to X_G \vee X_H$ of the map f which induces isomorphism on fundamental groups and subcomplexes $\tilde{K}_G$, $\tilde{K}_H$ which cover $\tilde{K}$ and intersect in a connected 1-dimensional subcomplex. It is clear from the construction that $\tilde{K}_G$ (respectively $\tilde{K}_H$) is connected and collapses to the union of K_G (respectively K_H) with some graph (simply collapse each new 2-cell e in $\tilde{K}_G$ or $\tilde{K}_H$ through the free face ℓ).

Moreover, since $\tilde{f}_*$ is an isomorphism and $\tilde{f}(\tilde{K}_G \cap \tilde{K}_H) = \{v\}$, the argument in the proof of Lemma 3.4.2 reveals that $\pi_1(\tilde{K}_G/(\tilde{K}_G \cap \tilde{K}_H))$ (and hence, also $\pi_1(\tilde{K}_G \cup C(\tilde{K}_G \cap \tilde{K}_H))$ is isomorphic to G. Let us analyze the complex $\tilde{K}_G \cup C(\tilde{K}_G \cap \tilde{K}_H)$ more closely. We may think of the process of attaching a cone on $\tilde{K}_G \cap \tilde{K}_H$ as the concatenation of three steps: first we attach a cylinder $(\tilde{K}_G \cap \tilde{K}_H) \times I$ at height 0, then we identify all points at height 1 over the same connected component of $K_G \cap K_H$ and finally we identify all points at height 1 over an edge in $\tilde{K}_G \cap \tilde{K}_H \setminus K_G \cap K_H$. Let L be the CW-complex obtained after applying the first two steps. Consider the complex obtained from L by collapsing each cylindrical 2-cell over an edge in $\tilde{K}_G \cap \tilde{K}_H \setminus K_G \cap K_H$ to its base (see Figure 4.1). After such collapses, the new 2-cells in $\tilde{K}_G$, that is, those 2-cells that are not in K' have free faces and can be collapsed. Thus, L collapses to the union of $K_G \cup C(K_G \cap K_H)$ with a certain graph M. Since the identifications performed in the third step are homotopy equivalences, we conclude that $\tilde{K}_G \cup C(\tilde{K}_G \cap \tilde{K}_H)$ is homotopy equivalent to $K_G \cup C(K_G \cap K_H) \cup M$ and the claim follows. □

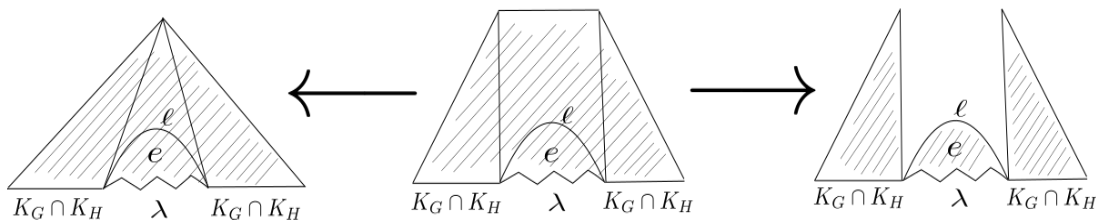

Figure 4.1: The left arrow is the identification of the upper part of the cylinder to a point, while the right arrow is the collapse through this edge.

In general, one may not expect the complex $K_G \cup C(X_G \cap X_H) \cup M$ to have at most $\kappa(G * H)$ 2-simplices. The objective of the rest of the section is showing how to estimate the number of 2-simplices of (some parts) of the complex.

In what follows, we think of the complex K as a rigid geometric structure (for example, as its geometric realization in some euclidean space $\mathbb{R}^n$) in which the subdivision K' has been marked. For a 2-simplex σ of K we write $\sigma' \leq K'$ to mean the subcomplex of the subdivided complex K' given by

$$\sigma' = \bigcup_{\tau \subseteq \sigma, \tau \in K'} \tau.$$

We distinguish two subcomplexes of $K_G \leq K'$ depending on the way they are distributed along the 2-simplices of K. Let A be the subcomplex of K_G generated by the union of those subcomplexes σ' such that $\sigma' \cap K_G = \sigma'$, where σ is a 2-simplex of K. That is, A is the union of the subdivided 2-simplices of K which are completely covered by 2-simplices of K_G. Let B be the complementary subcomplex of K_G. Notice that B collapses to a graph: indeed, for each 2-simplex σ of K, the complex $B \cap \sigma'$ is a proper subcomplex of σ' relative to the subdivision of the boundary $(\partial\sigma)'$, that is, the collapse does not use edges belonging to the subcomplex $(\partial\sigma)'$. Hence, the collapses inside each original 2-simplex $\sigma \in K$ do not interfere and we see that B collapses to a graph.

We speculate for that reason that it is maybe possible to prove that the fundamental group of the union of B with the cone over the edges of $K_G \cap K_H$ that belong to B is free (it should be noted that without some additional hypotheses this need not be true: for example, triangulations of the Möbius band collapse to a graph but the cone on the boundary of the band does not have free fundamental group). Suppose that the fundamental group of $B \cup C(K_G \cap K_H \cap B)$ is free. If in addition the connected components of the intersection of this complex with $A \cup C(K_G \cap K_H \cap A)$ are simply connected, by the Seifert-van Kampen theorem the fundamental group of

$$K_G \cup C(K_G \cap K_H) \cup M = (A \cup C(K_G \cap K_H \cap A)) \cup (B \cup C(K_G \cap K_H \cap B)) \cup M$$

is isomorphic to the free product of $\pi_1(A \cup C(K_G \cap K_H \cap A))$ with a free group. In such situation, since $\pi_1(K_G \cup C(K_G \cap K_H) \cup M)$ is isomorphic to G, the number of 2-simplices in $A \cup C(K_G \cap K_H \cap A)$ bounds the simplicial complexity of G from above.

We will assume in what follows that both the condition over the fundamental group of $B \cup C(K_G \cap K_H \cap B)$ and its intersection with $A \cup C(K_G \cap K_H \cap A)$ are true and concentrate in estimating the number of 2-simplices in the latter complex. We will be able to establish, modulo those two conditions, a bound from below for the simplicial complexity of $G * H$ in terms of $\kappa(G)$.

Of course, to complete the proof of the bound in full generality one would need to guarantee that the two mentioned conditions hold, possibly by carrying out some modifications in the spirit of Lemma 4.5.2 below to gain greater control over the complexes K_G and K_H and its intersection.

At first sight, it would seem that the number of 2-simplices in $A \cup C(K_G \cap K_H \cap A)$ could be arbitrarily larger than the number of simplices in K, since we do not control how fine the subdivision K' needs to be taken. However, the first thing to notice is that after erasing the subdivisions, the complex A may be considered as a subcomplex of the original complex K. To legitimately clear the subdivisions, we need to ensure that such process does not destroy the simplicial structure in the cone $C(K_G \cap K_H \cap A)$. Take an edge e in $K_G \cap K_H \cap A$. Since e belongs in particular to the intersection of K_G and K_H, it cannot be in the interior of an original 2-simplex σ of K. Indeed, if this would be the case, e should be the face of some 2-simplex of the subdivision K' whose image under f is contained in X_H. But since $e \in A$, $f(\sigma)$ is contained in X_G. Hence, e is an edge of the subdivision of some edge E of the 1-skeleton of K.

Now, consider the connected component Z of $E \cap (K_G \cap K_H \cap A)$ containing e. If Z is a proper subcomplex of the subdivision of E, it means that the edge over some of the two endpoints of Z is free (that is, the face of only one 2-simplex) in the complex $A \cup C(K_G \cap K_H \cap A)$. Thus, the cone over Z may be collapsed through this edge to a 1-dimensional complex. Thus, after performing all such collapses and erasing the subdivisions, we arrive at a 2-dimensional complex L in which the 2-simplices are only of one of two types. On one hand, we have the original 2-simplices of K which come from forgetting the subdivisions in A and on the other, the 2-simplices of the cone which are based on 1-simplices of K. To be able to control the number of 2-simplices in L of the second type, we need to perform a modification to the map f and the subdivision of K at the start of the construction (in particular, before the definition of the complexes K_G and K_H).

Lemma 4.5.2. *Let $\sigma = \{a, b, c\}$ be a 2-simplex of K (before the subdivision) and suppose that $f(\sigma) \neq \{v\}$. Then, the map f and the subdivision K' of K can be modified to send to v at most one of the edges of $\partial\sigma := \{a, b\} \cup \{a, c\} \cup \{b, c\}$, while still inducing an isomorphism on fundamental groups.*

Proof. We deal first with the case in which f sends the whole boundary of σ to $\{v\}$. Define $g : K \to X_G \vee X_H$ to be identically v in σ and equal to f outside σ. Since g is simplicial and coincides with f in the 1-skeleton $K^{(1)}$, it induces the same morphism on fundamental groups as f and we are done.

Suppose now that $f(\{a, b\}) = f(\{a, c\}) = \{v\}$. Then $\alpha = f|_{\{b,c\}} : \{b, c\} \to X_G \vee X_H$ is a loop in v. Notice that this loop is trivial in the fundamental group $\pi_1(X_G \vee X_H, v)$. Indeed, α is path-homotopic to a loop β that parameterizes $f(\partial\sigma)$. Since $[\partial\sigma]$ is trivial in $\pi_1(K)$, the class of α in $\pi_1(X_G \vee X_H, v)$ is trivial. From a path-homotopy between α and the constant path v in $X_G \vee X_H$, we obtain by the relative simplicial approximation theorem a simplicial map $F : D \to X_G \vee X_H$ such that D is a triangulated disk with boundary $\ell_1 \cup \ell_2$ where ℓ_1, ℓ_2 are intervals triangulated as (the subdivision in K' of) $\{b, c\}$ and $F|_{\ell_1} = \alpha$, $F|_{\ell_2} = const_v$, the map which is constantly v. Now, for a 2-simplex η of K containing the edge $\{b, c\}$ consider $\eta \cup D_\eta$ where D_η is a simplicially isomorphic copy of D and glue D_η to η by the interval ℓ_1. We can get a new subdivision of K by replacing the simplices η in the star (in K) of $\{b, c\}$ by $\eta \cup D_\eta$, where we glue $\eta \cup D_\eta$ and $\tau \cup D_\tau$ by the isomorphic copy of ℓ_2 in D_η and D_τ. The map defined as f outside the star of $\{b, c\}$ in K and as $f \cup F$ on every $\eta \cup D_\eta$ for η in $\mathrm{st}_K(\{b, c\})$ clearly defines the same map as f on fundamental groups. $\square$

As a consequence, we may assume that every 2-simplex of K has at most one edge that gets mapped entirely to $\{v\}$ under f. The relevance of this relies in the fact that the edges of K that are the base of 2-simplices in the cone part of L are mapped to $\{v\}$ by f since they belong to the intersection $K_G \cap K_H$. It follows that for each 2-simplex of K not in A there is at most one 2-simplex in the cone part of L and hence, the number of 2-simplices in L is at most $\kappa(G * H)$, the number of 2-simplices in K. Summing up, this argument would show that $\kappa(G*H)$ is greater than or equal to $\kappa(G)$ if we were able to account for the effect on the fundamental group of the complex $K_G \cup C(K_G \cap K_H)$ corresponding to the subcomplex B without spending additional 2-simplices.

Bibliography

[1] S. Aaronson and N. A. Scoville. Lusternik-Schnirelmann category for simplicial complexes. *Illinois J. Math.*, 57(3):743–753, 2013.

[2] R. D. M. Accola. Differentials and extremal length on Riemann surfaces. *Proc. Nat. Acad. Sci. U.S.A.*, 46:540–543, 1960.

[3] J. J. Andrews and M. L. Curtis. Free groups and handlebodies. *Proc. Amer. Math. Soc.*, 16:192–195, 1965.

[4] I. Babenko, F. Balacheff, and G. Bulteau. Systolic geometry and simplicial complexity for groups. *J. Reine Angew. Math.*, 757:247–277, 2019.

[5] F. Balacheff, H. Parlier, and S. Sabourau. Short loop decompositions of surfaces and the geometry of Jacobians. *Geom. Funct. Anal.*, 22(1):37–73, 2012.

[6] W. Ballmann. Der Satz von Lusternik und Schnirelmann. In *Beiträge zur Differentialgeometrie, Heft 1*, volume 102 of *Bonner Math. Schriften*, pages 1–25. Univ. Bonn, Bonn, 1978.

[7] V. Bangert, C. Croke, S. Ivanov, and M. Katz. Filling area conjecture and ovalless real hyperelliptic surfaces. *Geom. Funct. Anal.*, 15(3):577–597, 2005.

[8] J. A. Barmak. *Algebraic topology of finite topological spaces and applications*, volume 2032 of *Lecture Notes in Mathematics.* Springer, Heidelberg, 2011.

[9] J. A. Barmak and E. G. Minian. Strong homotopy types, nerves and collapses. *Discrete Comput. Geom.*, 47(2):301–328, 2012.

[10] D. Barnette. A proof of the lower bound conjecture for convex polytopes. *Pacific J. Math.*, 46:349–354, 1973.

[11] C. Bavard. Inégalité isosystolique pour la bouteille de Klein. *Math. Ann.*, 274(3):439–441, 1986.

[12] M. Berger. à l'ombre de Loewner. *Ann. Sci. École Norm. Sup. (4)*, 5:241–260, 1972.

[13] M. Berger. Du côté de chez Pu. *Ann. Sci. École Norm. Sup. (4)*, 5:1–44, 1972.

[14] C. Blatter. Über Extremallängen auf geschlossenen Flächen. *Comment. Math. Helv.*, 35:153–168, 1961.

[15] E. Borghini. A systolic inequality for 2-complexes of maximal cup-length and systolic area of groups. *arXiv e-prints*, page arXiv:1909.12720, Sep 2019.

[16] E. Borghini. Some remarks on PL collapsible covers of 2-dimensional polyhedra. *Osaka J. Math.*, 57(1):39–49, 2020.

[17] E. Borghini and E. G. Minian. The covering type of closed surfaces and minimal triangulations. *J. Combin. Theory Ser. A*, 166:1–10, 2019.

[18] E. Borghini and E. G. Minian. Simplicial complexity of surface groups. *Proc. Roy. Soc. Edinburgh Sect. A*, 2019. (in press).

[19] R. Bott. Nondegenerate critical manifolds. *Ann. of Math. (2)*, 60:248–261, 1954.

[20] U. Brehm and W. Kühnel. Combinatorial manifolds with few vertices. *Topology*, 26(4):465–473, 1987.

[21] U. Brehm and W. Kühnel. 15-vertex triangulations of an 8-manifold. *Math. Ann.*, 294(1):167–193, 1992.

[22] M. R. Bridson. The geometry of the word problem. In *Invitations to geometry and topology*, volume 7 of *Oxf. Grad. Texts Math.*, pages 29–91. Oxford Univ. Press, Oxford, 2002.

[23] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[24] K. S. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.

[25] G. Bulteau. Les groupes de petite complexité simpliciale. Preprint, June 2015.

[26] Y. D. Burago and V. A. Zalgaller. *Geometric inequalities*, volume 285 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1988. Translated from the Russian by A. B. Sosinskiĭ, Springer Series in Soviet Mathematics.

[27] P. Buser and P. Sarnak. On the period matrix of a Riemann surface of large genus. *Invent. Math.*, 117(1):27–56, 1994. With an appendix by J. H. Conway and N. J. A. Sloane.

[28] J. Cheeger. Finiteness theorems for Riemannian manifolds. *Amer. J. Math.*, 92:61–74, 1970.

[29] O. Cornea, G. Lupton, J. Oprea, and D. Tanré. *Lusternik-Schnirelmann category*, volume 103 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2003.

[30] C. B. Croke. Rigidity theorems in Riemannian geometry. In *Geometric methods in inverse problems and PDE control*, volume 137 of *IMA Vol. Math. Appl.*, pages 47–72. Springer, New York, 2004.

[31] M. Farber. Topological complexity of motion planning. *Discrete Comput. Geom.*, 29(2):211–221, 2003.

[32] H. Federer and W. H. Fleming. Normal and integral currents. *Ann. of Math. (2)*, 72:458–520, 1960.

[33] D. Fernández-Ternero, E. Macías-Virgós, E. Minuz, and J. A. Vilches. Discrete topological complexity. *Proc. Amer. Math. Soc.*, 146(10):4535–4548, 2018.

[34] D. Fernández-Ternero, E. Macías-Virgós, and J. A. Vilches. Lusternik-Schnirelmann category of simplicial complexes and finite spaces. *Topology Appl.*, 194:37–50, 2015.

[35] R. Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998.

[36] R. H. Fox. On the Lusternik-Schnirelmann category. *Ann. of Math. (2)*, 42:333–370, 1941.

[37] A. A. Gaĭfullin. Configuration spaces, bistellar moves, and combinatorial formulas for the first Pontryagin class. *Tr. Mat. Inst. Steklova*, 268(Differentsial'nye Uravneniya i Topologiya. I):76–93, 2010.

[38] I. M. Gel'fand and R. D. MacPherson. A combinatorial formula for the Pontrjagin classes. *Bull. Amer. Math. Soc. (N.S.)*, 26(2):304–309, 1992.

[39] M. A. Grayson. Shortening embedded curves. *Ann. of Math. (2)*, 129(1):71–111, 1989.

[40] B. Green, N. A. Scoville, and M. Tsuruga. Estimating the discrete geometric Lusternik-Schnirelmann category. *Topol. Methods Nonlinear Anal.*, 45(1):103–116, 2015.

[41] M. Gromov. Groups of polynomial growth and expanding maps. *Inst. Hautes Études Sci. Publ. Math.*, 53(1):53–73, 1981.

[42] M. Gromov. *Structures métriques pour les variétés riemanniennes*, volume 1 of *Textes Mathématiques [Mathematical Texts]*. CEDIC, Paris, 1981. Edited by J. Lafontaine and P. Pansu.

[43] M. Gromov. Volume and bounded cohomology. *Inst. Hautes Études Sci. Publ. Math.*, 56:5–99 (1983), 1982.

[44] M. Gromov. Filling Riemannian manifolds. *J. Differential Geom.*, 18(1):1–147, 1983.

[45] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[46] M. Gromov. Systoles and intersystolic inequalities. In *Actes de la Table Ronde de Géométrie Différentielle (Luminy, 1992)*, volume 1 of *Sémin. Congr.*, pages 291–362. Soc. Math. France, Paris, 1996.

[47] M. Gromov. *Metric structures for Riemannian and non-Riemannian spaces.* Modern Birkhäuser Classics. Birkhäuser Boston, Inc., Boston, MA, english edition, 2007. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates.

[48] L. Guth. Notes on Gromov's systolic estimate. *Geom. Dedicata*, 123:113–129, 2006.

[49] L. Guth. Metaphors in systolic geometry. In *Proceedings of the International Congress of Mathematicians. Volume II*, pages 745–768. Hindustan Book Agency, New Delhi, 2010.

[50] L. Guth. Systolic inequalities and minimal hypersurfaces. *Geom. Funct. Anal.*, 19(6):1688–1692, 2010.

[51] M. Hachimori. Decompositions of two-dimensional simplicial complexes. *Discrete Math.*, 308(11):2307–2312, 2008.

[52] A. Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[53] J. J. Hebda. Some lower bounds for the area of surfaces. *Invent. Math.*, 65(3):485–490, 1981/82.

[54] C. Hog-Angeloni and W. Metzler, editors. *Two-dimensional homotopy and combinatorial group theory*, volume 197 of *London Mathematical Society Lecture Note Series.* Cambridge University Press, Cambridge, 1993.

[55] J. F. P. Hudson. *Piecewise linear topology.* University of Chicago Lecture Notes prepared with the assistance of J. L. Shaneson and J. Lees. W. A. Benjamin, Inc., New York-Amsterdam, 1969.

[56] S. V. Ivanov. Volumes and areas of Lipschitz metrics. *Algebra i Analiz*, 20(3):74–111, 2008.

[57] S. V. Ivanov and M. G. Katz. Generalized degree and optimal Loewner-type inequalities. *Israel J. Math.*, 141:221–233, 2004.

[58] M. Jungerman and G. Ringel. Minimal triangulations on orientable surfaces. *Acta Math.*, 145(1-2):121–154, 1980.

[59] J. Kahn and V. Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.

[60] G. Kalai. Rigidity and the lower bound theorem. I. *Invent. Math.*, 88(1):125–151, 1987.

[61] M. Karoubi and C. Weibel. On the covering type of a space. *Enseign. Math.*, 62(3-4):457–474, 2016.

[62] K. U. Katz and M. G. Katz. Bi-Lipschitz approximation by finite-dimensional imbeddings. *Geom. Dedicata*, 150:131–136, 2011.

[63] K. U. Katz, M. G. Katz, S. Sabourau, S. Shnider, and S. Weinberger. Relative systoles of relative-essential 2-complexes. *Algebr. Geom. Topol.*, 11(1):197–217, 2011.

[64] M. G. Katz, Y. B. Rudyak, and S. Sabourau. Systoles of 2-complexes, Reeb graph, and Grushko decomposition. *Int. Math. Res. Not.*, pages Art. ID 54936, 30, 2006.

[65] M. G. Katz and S. Sabourau. Entropy of systolically extremal surfaces and asymptotic bounds. *Ergodic Theory Dynam. Systems*, 25(4):1209–1220, 2005.

[66] J. M. Lee. *Introduction to topological manifolds*, volume 202 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2011.

[67] L. A. Lusternik and L. G. Schnirelmann. *Méthodes topologiques dans les problèmes variationnels*. Hermann, Paris, 1934.

[68] F. H. Lutz. *Triangulated manifolds with few vertices and vertex-transitive group actions*. Berichte aus der Mathematik. [Reports from Mathematics]. Verlag Shaker, Aachen, 1999. Dissertation, Technischen Universität Berlin, Berlin, 1999.

[69] F. H. Lutz. Combinatorial 3-manifolds with 10 vertices. *Beiträge Algebra Geom.*, 49(1):97–106, 2008.

[70] S. V. Matveev. Complexity theory of three-dimensional manifolds. *Acta Appl. Math.*, 19(2):101–130, 1990.

[71] J. Milnor and W. Thurston. Characteristic numbers of 3-manifolds. *Enseign. Math. (2)*, 23(3-4):249–254, 1977.

[72] M. Morse. *The calculus of variations in the large*, volume 18 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 1996. Reprint of the 1932 original.

[73] K. Nakamura. On Isosystolic Inequalities for T^n, RP^n, and M^3. *arXiv e-prints*, page arXiv:1306.1617, Jun 2013.

[74] P. M. Pu. Some inequalities in certain nonorientable Riemannian manifolds. *Pacific J. Math.*, 2:55–71, 1952.

[75] J. G. Ratcliffe. The cohomology ring of a one-relator group. In *Contributions to group theory*, volume 33 of *Contemp. Math.*, pages 455–466. Amer. Math. Soc., Providence, RI, 1984.

[76] G. Ringel. Wie man die geschlossenen nichtorientierbaren Flächen in möglichst wenig Dreiecke zerlegen kann. *Math. Ann.*, 130:317–326, 1955.

[77] Y. B. Rudyak and S. Sabourau. Systolic invariants of groups and 2-complexes via Grushko decomposition. *Ann. Inst. Fourier (Grenoble)*, 58(3):777–800, 2008.

[78] R. Schoen. Estimates for stable minimal surfaces in three-dimensional manifolds. In *Seminar on minimal submanifolds*, volume 103 of *Ann. of Math. Stud.*, pages 111–126. Princeton Univ. Press, Princeton, NJ, 1983.

[79] P. Scott and C. T. C. Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge-New York, 1979.

[80] J. R. Stallings. A topological proof of Grushko's theorem on free products. *Math. Z.*, 90:1–8, 1965.

[81] R. Thom. Quelques propriétés globales des variétés différentiables. *Comment. Math. Helv.*, 28:17–86, 1954.

[82] W. Thurston. Geometry and topology of three-manifolds, 1978.

[83] C. T. C. Wall. Finiteness conditions for CW-complexes. *Ann. of Math. (2)*, 81:56–69, 1965.

[84] A. Weil. Sur les théorèmes de de Rham. *Comment. Math. Helv.*, 26:119–145, 1952.

[85] G. W. Whitehead. On mappings into group-like spaces. *Comment. Math. Helv.*, 28:320–328, 1954.