Combinatorial Width Parameters for 3-Dimensional Manifolds

by

Kristóf Huszár

June 2020

A thesis presented to the
Graduate School
of the
Institute of Science and Technology Austria, Klosterneuburg, Austria
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

Institute of Science and Technology

The thesis of Kristóf Huszár, titled *Combinatorial Width Parameters for 3-Dimensional Manifolds*, is approved by:

Supervisor: Uli Wagner, IST Austria, Klosterneuburg, Austria

Signature: _____

Co-supervisor: Jonathan Spreer, The University of Sydney, NSW 2006, Australia

Signature: _____

Committee Member: Éric Colin de Verdière, LIGM, CNRS, Marne-la-Vallée, France

Signature: _____

Committee Member: Herbert Edelsbrunner, IST Austria, Klosterneuburg, Austria

Signature: _____

Defense Chair: Jiří Friml, IST Austria, Klosterneuburg, Austria

Signature: _____

signed page is on file

IST Austria Thesis, ISSN: 2663-337X

ISBN: 978-3-99078-006-0

I hereby declare that this thesis is my own work and that it does not contain other people's work without this being so stated; this thesis does not contain my previous work without this being stated, and the bibliography contains all the literature that I used in writing the dissertation.

I declare that this is a true copy of my thesis, including any final revisions, as approved by my thesis committee, and that this thesis has not been submitted for a higher degree to any other university or institution.

I certify that any republication of materials presented in this thesis has been approved by the relevant publishers and co-authors.

Signature: _____

Kristóf Huszár

June 2020

signed page is on file

Table of Contents

Abstract

Algorithms in computational 3-manifold topology typically take a triangulation as an input and return topological information about the underlying 3-manifold. However, extracting the desired information from a triangulation (e.g., evaluating an invariant) is often computationally very expensive. In recent years this complexity barrier has been successfully tackled in some cases by importing ideas from the theory of parameterized algorithms into the realm of 3-manifolds. Various computationally hard problems were shown to be efficiently solvable for input triangulations that are sufficiently "tree-like."

In this thesis we focus on the key combinatorial parameter in the above context: we consider the *treewidth* of a compact, orientable 3-manifold, i.e., the smallest treewidth of the dual graph of any triangulation thereof. By building on the work of Scharlemann–Thompson and Scharlemann–Schultens–Saito on generalized Heegaard splittings, and on the work of Jaco–Rubinstein on layered triangulations, we establish quantitative relations between the treewidth and classical topological invariants of a 3-manifold. In particular, among other results, we show that the treewidth of a closed, orientable, irreducible, non-Haken 3-manifold is always within a constant factor of its Heegaard genus.

Acknowledgments

There are many people who have stood by me throughout the long and daunting journey of getting a doctorate, without whom I would not have been able to reach this point. It is a great pleasure for me to acknowledge their support.

First and foremost, I would like to thank my supervisor **Uli Wagner** for taking me as a PhD student, for being so generous and supportive during all these years, and for helping me to repeatedly leave my comfort zone, especially when I really doubted myself.

I am very grateful to my co-supervisor **Jonathan Spreer** for our fruitful collaboration and friendship, for his hospitality at FU Berlin, and for the countless long-distance video meetings, that allowed me to enter the world of computational 3-manifold topology.

I thank **Éric Colin de Verdière** and **Herbert Edelsbrunner** for their very careful review of my thesis, for acting as co-examiners at my thesis defense, and for providing me with feedback and guidance at various stages of my PhD; **Roy Meshulam** for acting as a co-examiner at my qualifying exam, for his steady encouragement, and for the inspiring conversations at various occasions; and **Jiří Friml** for chairing my thesis defense.

I am indebted to **Arnaud de Mesmay** and **Saul Schleimer** for noticing a gap in an earlier proof of Theorem 1.4 and for pointing us towards key ingredients of the current one. I also thank Arnaud for inspiring me ever since we shared an office at IST Austria.

I greatly appreciate the friendship of **Anna Gundert**, my "academic older sister," her support at the beginning of my PhD, and her hospitality at the University of Cologne.

I am very grateful to **Ben Burton**, **Clément Maria**, **Jennifer Schultens**, **Abigail Thompson**, **Stefan Friedl** and **Raphael Zentner** for their interest in my work, for their kindness and generosity, and for the stimulating discussions and opportunities; and to the *SoCG* community for allowing me to present our results in a friendly environment. I am indebted to **Ágnes Tuska** for her warm hospitality in Fresno, California during the summer of 2019 and for the thought-provoking conversations on mathematics education.

For nearly seven years, IST Austria has been my academic home, where I was fortunate to work alongside wonderful colleagues. I thank my fellow PhD students in the cohort of 2013 for making the beginning of this journey so enjoyable, despite all the difficulties; the members of the *Wagner Group*—**Alan Arroyo**, **Sergey Avvakumov**, **Marek Filakovský**, **Peter Franek**, **Rado Fulek**, **Marek Krčál**, **Isaac Mabillard**, **Zuzka Masárová**, **Pavel Paták**, **Zuzka Patáková**, **Pascal Wild**, and **Stephan Zhechev**—for sharing the daily experience of being a young researcher and for teaching me a great deal of mathematics; and the scientific community of the Institute for allowing me to gain insights into a wide range of ideas and methods. I am particularly grateful to Pascal and Stephan for the numerous "after hours" conversations about math and much more.

In the course of my PhD, I have enjoyed administrative support of the highest quality, thanks to **Astrid Bonventre-Darthé**, **Vlad Cozac**, **Christine Francois-Rennhofer**, **Elisabeth Hacker**, **Christine Ostermann**, and the team of the *Graduate School Office*. I would also like to thank *IT Support* and the *Library Team* for the reliable assistance on many occasions, and *Communications & Events* for the engaging collaborations.

I am grateful to **Bogi Bálint**, **Carina Baskett**, **Zsuzsi Darida**, **Kata Demeter**, **Arinya Eller**, **Balázs Érdi**, **Dóri Gábriel**, **Iordan Ganev**, **Balázs Garamszegi**, **Máté Gerencsér**, **Eva Heyen**, **Fabienne Jesse**, **Farid Karimipour**, **Réka Kelemen**, **Martin Lukačišin**, **Marta Lukačišinová**, **Lenka Matejovičová**, **Sasha Minets**, **Karin Mitosch**, **Erika Nessel**, **Sebastian Nozzi**, **Georg Osang**, **Márton Palatinszky**, **Georg Rieckh**, **Mišo Rybár**, **Hana Semerádová**, **Magdalena Steinrück**, **Marcin and Mathilde Suskiewicz**, **Peter Synak**, **Barbi Szirányi**, **Pepa Tkadlec**, **Mina Vasileva**, and **Balázs Varga** for enriching my life in so many ways; and to the welcoming community of the parish *Döbling–St. Paul* for helping me to grow in faith.

I am forever grateful to **Vera Lányi**, my high-school math teacher, for setting me on the road, and I am beholden to my mentors **Ádám Besenyei**, **Paul Campbell**, **Dave Dwyer**, **Dave Ellis**, **György Oláh** (†), **János Pach**, **Lajos Pósa**, **János Ruff**, and **András Stipsicz** for cultivating my mathematical interests and for supporting my endeavors. I am grateful to the Hungarian math community and the *Eötvös Collegium* for many friendships and interdisciplinary perspectives across sciences and humanities.

Finally, I would like to thank my family, in particular my parents **Renáta** and **Zoltán**, and my brother **Gergő**, for their constant support and unconditional love.

About the Author

Kristóf Huszár obtained his bachelor's degree in mathematics in June 2013 at the Eötvös Loránd University in Budapest, Hungary. During his undergraduate studies, he spent one semester each at Beloit College (Wisconsin, USA) and at the University of Heidelberg (Germany) as an exchange student.

Kristóf started his PhD at IST Austria in September 2013 and joined the research group of Uli Wagner one year later. His main research interest lies in discrete and computational topology, mostly—but not exclusively—in a three-dimensional setting. He also has a great interest in mathematics education and public outreach.

List of Publications

This thesis is partially based on the following publications.

1. Kristóf Huszár, Jonathan Spreer, and Uli Wagner. On the treewidth of triangulated 3-manifolds. *Journal of Computational Geometry*, 10(2):70–98, 2019. doi:10.20382/jogc.v10i2a5, IST-REx-ID:7093, MR:4039886, Zbl:07150581.

   Extended abstract: *Proceedings of the 34th International Symposium on Computational Geometry (SoCG 2018), June 11–14, 2018, Budapest, Hungary.*

2. Kristóf Huszár and Jonathan Spreer. 3-Manifold triangulations with small treewidth. In *Proceedings of the 35th International Symposium on Computational Geometry (SoCG 2019), June 18–21, 2019, Portland, OR, USA*, volume 129 of LIPIcs, pages 44:1–44:20. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, 2019. doi:10.4230/LIPIcs.SoCG.2019.44, IST-REx-ID:6556, MR:3968630.

List of Tables

List of Figures

1 Introduction

1.1 Motivation and Guiding Questions

In 3-dimensional topology many fundamental problems can be solved algorithmically. Famous examples include deciding whether a given knot is trivial [54], deciding whether a given 3-manifold is homeomorphic to the 3-sphere [125, 144], and, more generally, deciding whether two given 3-manifolds are homeomorphic, see, e.g., [9, 87, 136]. The algorithm for solving the homeomorphism problem is still purely theoretical, and its complexity remains largely unknown [87, 89, 90]. In contrast, the first two problems are known to lie in the intersection of the complexity classes **NP** and **co-NP** [56, 68, 86, 88, 132, 151].[1]

Moreover, implementations of, for instance, algorithms to recognize the 3-sphere exist out-of-the-box (e.g., by using the computational 3-manifold topology software *Regina* [28]) and exhibit practical running times for virtually all known inputs.

In fact, many topological problems have implemented algorithmic solutions that can efficiently handle instances of considerable size. This is despite the fact that most of these implementations have prohibitive worst-case running times, or the underlying problems are even known to be computationally hard in general. In recent years, there have been several attempts to explain this gap using the concepts of parameterized complexity and algorithms for fixed-parameter tractable (FPT) problems. This effort has proven to be highly effective and, today, there exist numerous FPT algorithms in the field [29, 31, 32, 33, 100]. More specifically, given a triangulation $\mathcal{T}$ of a 3-manifold $\mathcal{M}$ with n tetrahedra whose dual graph $\Gamma(\mathcal{T})$ has treewidth[2] at most k, there exist algorithms to compute

[1] The proof that 3-sphere recognition lies in **co-NP** assumes the Generalized Riemann Hypothesis.

[2] We often simply speak of the treewidth of a triangulation, meaning the treewidth of its dual graph.

- taut angle structures[3] of what is called ideal triangulations with torus boundary components in running time $O(7^k \cdot n)$ [33];
- optimal Morse matchings[4] in the Hasse diagram of $\mathcal{T}$ in $O(4^{k^2+k} \cdot k^3 \cdot \log k \cdot n)$ [31];
- the Turaev–Viro invariants[5] for parameter $r \geq 3$ in $O((r-1)^{6(k+1)} \cdot k^2 \cdot \log r \cdot n)$ [32];
- every problem which can be expressed in monadic second-order logic in $O(f(k) \cdot n)$, where f often is a tower of exponentials [29].[6]

Some of these results are not purely theoretical, but are implemented and outperform previous state-of-the-art implementations for typical input. Consequently, they have a significant practical impact. This is in particular the case for the algorithm to compute Turaev–Viro invariants [32] (cf. [100]).

Note that treewidth—the dominating parameter k in the running times given above—is a combinatorial quantity linked to a triangulation, and not a topological invariant of the underlying 3-manifold.[7] This gives rise to the following approach to efficiently solve topological problems on a 3-manifold $\mathcal{M}$: given a triangulation $\mathcal{T}$ of $\mathcal{M}$, search for another triangulation $\mathcal{T}'$ of the same manifold with smaller treewidth.

This approach faces severe difficulties. By a theorem due to Kirby and Melvin [80], the Turaev–Viro invariant for parameter $r = 4$ is #**P**-hard to compute. Thus, if there were a polynomial time procedure to turn an n-tetrahedron triangulation $\mathcal{T}$ into a poly(n)-tetrahedron triangulation $\mathcal{T}'$ with dual graph of treewidth at most k, for some universal constant k, then this procedure, combined with the algorithm from [32], would constitute a polynomial time solution for a #**P**-hard problem. Furthermore, known facts imply that *most* triangulations of *most* 3-manifolds must have large treewidth (see Appendix B.1). However, while these arguments indicate that triangulations of small treewidth may be rare and computationally hard to find, they do not rule out that every manifold has some (potentially very large) triangulation of bounded treewidth. Thus it is intriguing to ask

[3] Taut angle structures are combinatorial versions of semi-simplicial metrics which have implications on the geometric properties of the underlying manifold.

[4] Optimal Morse matchings translate to discrete Morse functions with minimal number of critical points with respect to the combinatorics of the triangulation and the topology of the underlying 3-manifold.

[5] Turaev–Viro invariants are powerful tools to distinguish between 3-manifolds. They are the method of choice when, for instance, creating large censuses of manifolds.

[6] This result is analogous to Courcelle's celebrated theorem in graph theory [38].

[7] The necessary background on graph parameters and 3-manifolds is introduced in Chapter 2.

Question 1. Does there exist a universal constant C, such that every 3-manifold $\mathcal{M}$ has a triangulation $\mathcal{T}$ with dual graph of treewidth at most C?

Indeed, Question 1 and variations thereof had been repeatedly raised by researchers in computational 3-manifold topology at several meetings and open problem sessions including an Oberwolfach workshop in 2015 [30, Problem 8] (formulated in the context of knot theory). It has also served as the main guiding question for the present thesis.

In order to put Question 1 into context, it will be helpful to introduce the *treewidth of a compact, connected* 3*-manifold* $\mathcal{M}$ defined by

$$\mathrm{tw}\,(\mathcal{M}) = \min\{\mathrm{tw}\,(\Gamma(\mathcal{T})) : \mathcal{T} \text{ is a triangulation of } \mathcal{M}\}, \tag{1.1}$$

where $\Gamma(\mathcal{T})$ denotes the dual graph of $\mathcal{T}$. Using (1.1), we can reformulate Question 1 as

Question 1′. Does there exist a universal constant C, such that, for every 3-manifold $\mathcal{M}$, we have $\mathrm{tw}\,(\mathcal{M}) \leq C$?

In this setting it is also natural to consider two closely related open-ended questions.

Question 2. What is the quantitative relation between treewidth $\mathrm{tw}\,(\mathcal{M})$ (and other similar invariants such as the cutwidth $\mathrm{cw}(\mathcal{M})$ and pathwidth $\mathrm{pw}(\mathcal{M})$, cf. Section 5.1.1) and other well-known topological invariants associated with 3-manifolds?

Question 3. What can we say about 3-manifolds $\mathcal{M}$ for which $\mathrm{tw}\,(\mathcal{M})$ is small?

1.2 The Main Results

As the first result in this thesis, we settle Question 1 by answering it in the negative.[8] More specifically, in Section 5.2 we prove the following two statements.

Theorem 1.1. *There exist* 3*-manifolds* $\mathcal{M}$ *with arbitrary large pathwidth* $\mathrm{pw}(\mathcal{M})$.

Theorem 1.2. *There exist* 3*-manifolds* $\mathcal{M}$ *with arbitrary large treewidth* $\mathrm{tw}\,(\mathcal{M})$.

We establish these results through the following theorems, that forge a quantitative link between the pathwidth (resp. treewidth) and the Heegaard genus of a 3-manifold, a classical invariant which has been subject of interest since the end of the 19th century.

[8] See [41] for related work concerning the respective question about knots and their diagrams.

Theorem 1.3. *Let $\mathcal{M}$ be a closed, orientable, irreducible and non-Haken 3-manifold. Then the pathwidth* $\mathrm{pw}(\mathcal{M})$ *and Heegaard genus* $\mathfrak{g}(\mathcal{M})$ *of the 3-manifold $\mathcal{M}$ satisfy*

$$\mathfrak{g}(\mathcal{M}) \leq 4(3\,\mathrm{pw}(\mathcal{M}) + 1).$$

Theorem 1.4. *Let $\mathcal{M}$ be a closed, orientable, irreducible and non-Haken 3-manifold. Then the treewidth* $\mathrm{tw}(\mathcal{M})$ *and Heegaard genus* $\mathfrak{g}(\mathcal{M})$ *of the 3-manifold $\mathcal{M}$ satisfy*

$$\mathfrak{g}(\mathcal{M}) \leq 18(\mathrm{tw}(\mathcal{M}) + 1).$$

The key instrument in proving Theorems 1.3 and 1.4 is the machinery of generalized Heegaard splittings, pioneered by Scharlemann–Thompson [131] and further developed by Scharlemann–Schultens–Saito [130]. We provide a detailed and illustrated exposition to this framework in Section 4.1.

By a result of Agol [1] (Theorem 5.4 in this thesis), there exist closed, orientable, irreducible, non-Haken 3-manifolds of arbitrarily large Heegaard genus. Combining this result with Theorems 1.3 and 1.4 thus immediately implies Theorems 1.1 and 1.2.

Remark 1.5. Note that Theorem 1.1 can be directly deduced from Theorem 1.2 since the pathwidth of a graph is always at least as large as its treewidth.[9] Nonetheless, we provide separate proofs for each of the two statements. The motivation is that while the proof of Theorem 1.3 is considerably simpler than that of Theorem 1.4, it already illustrates several key concepts and ideas which we are building upon in the proof of Theorem 1.4.

In the second half of Chapter 5, we construct small-treewidth triangulations informed by the topology of 3-manifolds. In Section 5.3.1, building on the theory of layered triangulations by Jaco–Rubinstein [71], we show that the Heegaard genus dominates the cutwidth (and so the pathwidth and treewidth as well) by virtue of the following statement.

Theorem 1.6. *Let $\mathcal{M}$ be a closed, orientable 3-manifold, and let* $\mathrm{cw}(\mathcal{M})$ *and* $\mathfrak{g}(\mathcal{M})$ *respectively denote the cutwidth and the Heegaard genus of $\mathcal{M}$. Then we have*

$$\mathrm{cw}(\mathcal{M}) \leq 4\mathfrak{g}(\mathcal{M}) - 2.$$

Theorem 1.6, in combination with Theorems 1.3 and 1.4, implies that for the class of non-Haken 3-manifolds, the Heegaard genus is in fact within a constant factor of both the

[9] This is immediate from the definitions of pathwidth and treewidth, see Section 2.1.1.

cutwidth, pathwidth and treewidth of a 3-manifold, providing an interesting connection between a classical topological invariant and topological properties directly related to the triangulations of a manifold (Corollary 5.19).

In Section 5.3.2, by constructing layered triangulations of compression bodies, we prove a strengthening of Theorem 1.6, where we show that the cutwidth is upper-bounded even by the *linear width* defined by Scharlemann–Thompson [131], which, for Haken-manifolds, can be much smaller than the Heegaard genus. (The difference can be arbitrary large.)

Theorem 1.7. *For every compact, orientable* 3*-manifold* $\mathcal{M}$*, possibly with nonempty boundary, the cutwidth* $\mathrm{cw}(\mathcal{M})$ *and the linear width* $\mathscr{L}(\mathcal{M})$ *satisfy*

$$\mathrm{cw}(\mathcal{M}) \leq 24\mathscr{L}(\mathcal{M}).$$

Similarly as before, Theorem 1.7, in combination with the techniques used to prove Theorem 1.3 (cf. Theorem 5.13) yields that the cutwidth, pathwidth and linear width of a closed, orientable 3-manifold are always within a constant factor (Corollary 5.22).

In Section 5.4 we turn our attention to the profuse family of hyperbolic 3-manifolds. Improving upon recent work of Maria and Purcell [99], we show that the volume of a hyperbolic 3-manifold provides a linear upper bound on its pathwidth.

Theorem 1.8. *There exists a universal constant* $C > 0$ *such that, for any closed hyperbolic 3-manifold* $\mathcal{M}$ *with pathwidth* $\mathrm{pw}(\mathcal{M})$ *and volume* $\mathrm{vol}(\mathcal{M})$*, we have*

$$\mathrm{pw}(\mathcal{M}) \leq C \cdot \mathrm{vol}(\mathcal{M}).$$

In Chapter 6, we inspect 3-manifolds having treewidth at most one. It turns out that this class is essentially the same as that of 3-manifolds with Heegaard genus at most one.

Theorem 1.9. *The class of* 3*-manifolds of treewidth at most one coincides with the class of* 3*-manifolds of Heegaard genus at most one together with the Seifert fibered space* $\mathrm{SFS}[\mathbb{S}^2 : (2,1),(2,1),(2,-1)]$ *of Heegaard genus two.*

In contrast, in Chapter 7 we show, by exhibiting triangulations of treewidth two for all orientable Seifert fibered spaces over the 2-sphere (Theorem 7.2) or a non-orientable surface (Theorem 7.4), that linking Heegaard genus to treewidth fails to hold in general in a very strong sense: there are infinite families of 3-manifolds of unbounded Heegaard

genus which are all of treewidth two (Corollary 7.9). Combining this construction with Theorem 1.9, we deduce that the treewidth of all 3-manifolds with spherical or $\mathbb{S}^2 \times \mathbb{R}$ geometry equals two (Corollary 7.6). As a result, we can determine the treewidth of 4889 out of the 4979 manifolds in the (≤ 10)-tetrahedra census (Table 7.1). Specifically, only 90 of them have treewidth possibly higher than two. These computations also confirm that not all minimal triangulations are of minimum treewidth (Corollary 7.8).

Altogether, our results suggest that *combinatorial width parameters for* 3*-manifolds* (see Section 5.1.1) are interesting notions at the interface of topology and combinatorics, and are well-suited to indicate the power of fixed-parameter tractable algorithms in the field of computational 3-manifold topology.

1.3 Outline of the Thesis

This thesis is organized as follows. We start with basic definitions regarding graphs and parameterized complexity in Chapter 2. Then, in Chapter 3, we provide a quick introduction to 3-manifolds and discuss various ways of constructing them from smaller pieces. This is followed, in Chapter 4, by an exposition about generalized Heegaard splittings and layered triangulations. In Chapters 5, 6 and 7 we present the proofs of our results outlined in Section 1.2. Along the way, we also provide bibliographic remarks and explanations of the additional techniques that are too specific to be included in the introductory chapters. The document is concluded with an Appendix of related constructions and results, that support the main line of discussion.

Graphs versus triangulations. Following a convention adopted by several authors in the field of computational low-dimensional topology, throughout this thesis we use the terms *edge* and *vertex* to refer to an edge or vertex in a 3-manifold triangulation, whereas the terms *arc* and *node* denote an edge or vertex in a graph, respectively.

2 Preliminaries on Graphs and Parameterized Complexity

In this chapter we recall basic terminology about graphs and parameterized complexity, and introduce key concepts that provide the foundation for the subsequent chapters.

2.1 Graphs

A *graph* (more precisely, a *multigraph*) $G = (V, E)$ is an ordered pair consisting of a finite set V of *nodes* and of a multiset E of unordered pairs of nodes, called *arcs*.[1] A *loop* is an arc $\ell \in E$ that is a multiset itself, e.g., $\ell = \{v, v\}$ for some $v \in V$. An element of E with multiplicity higher than one is called a *multiple arc*. A *simple graph* is a graph without loops or multiple arcs. The *degree* $\deg(v)$ of a node $v \in V$ equals the number of arcs containing it, counted with multiplicity. If all of its nodes have the same degree $k \in \mathbb{N}$, a graph is called *k-regular*. A *tree* is a connected graph with n nodes and $n - 1$ arcs, and a *path* is a tree in which every node has degree at most two. A *leaf* is a node of degree one. For general background in graph theory we refer to [46].

Sometimes it will be useful to consider the *simplification* of a multigraph in which we forget the loops and reduce each multiple arc to a single one (cf. Figure 2.1).

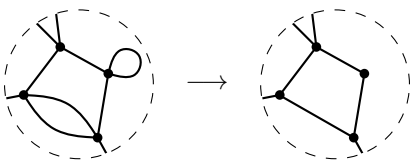

Figure 2.1: The local effect of simplification in a multigraph

[1] If we refer to the nodes (resp. arcs) of a specific graph G, we also use the notation $V(G)$ (resp. $E(G)$).

2.1.1 Selected Width Parameters for Graphs

The theory of parameterized complexity has its sources in graph theory, where many problems which are **NP**-hard in general become tractable in polynomial time if one assumes structural restrictions about the possible input graphs [19]. For instance, several graph theoretical questions have a simple answer if one asks them about trees, or graphs that are similar to trees in some sense. Width parameters capture this similarity in a quantitative way [61]. We are particularly interested in the behavior of these parameters and their relationship with each other when considering bounded-degree graphs or, more specifically, dual graphs of 3-manifold triangulations (cf. Section 3.1).

Treewidth and pathwidth. The concepts of treewidth and pathwidth were introduced by Robertson and Seymour in their early papers on graph minors [123, 124] and have become cornerstones of structural graph theory [14, 16, 77, 93]. The treewidth and pathwidth of a given graph respectively measure how tree-like or path-like the graph is.

Definition 2.1 (Tree decomposition, treewidth)**.** A *tree decomposition* of $G = (V, E)$ is a pair $(\{B_i : i \in I\}, T = (I, F))$ with *bag*s $B_i \subseteq V$, $i \in I$, and a tree $T = (I, F)$, such that

1) $\bigcup_{i \in I} B_i = V$,

2) for every arc $\{u, v\} \in E$, there exists $i \in I$ with $\{u, v\} \subseteq B_i$, and

3) for every node $v \in V$, $T_v = \{i \in I : v \in B_i\}$ spans a connected subtree of T.

The *width* of a tree decomposition equals $\max_{i \in I} |B_i| - 1$, and the *treewidth* $\operatorname{tw}(G)$ is the smallest width of any tree decomposition of G. See Figure 2.2 for some examples.

Definition 2.2 (Path decomposition, pathwidth)**.** A *path decomposition* of a graph G is a tree decomposition for which the tree T is required to be a *path*. The *pathwidth* $\operatorname{pw}(G)$ of a graph G is the minimum width of any path decomposition of G.

Cutwidth and congestion. While treewidth and pathwidth are generally useful in the design and analysis of graph algorithms, congestion and cutwidth have turned out to be helpful, respectively, to relate treewidth and pathwidth to classical topological invariants in a quantitative way [41, 66, 67, 98, 99]. In Chapter 5 we will see this in action.

First we discuss cutwidth. See [45] for an overview of its algorithmic aspects.

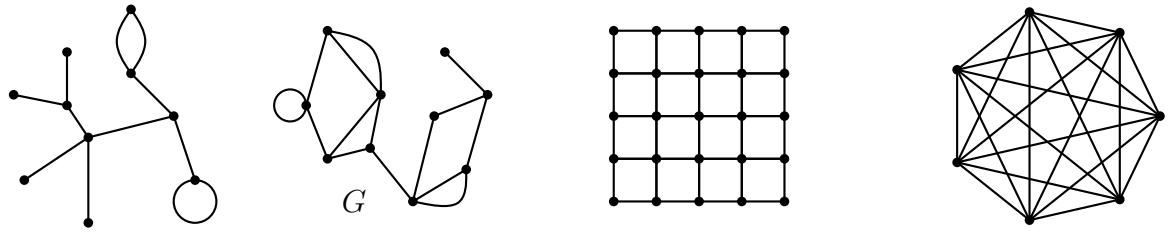

(i) $\mathrm{tw}(\text{tree}) = 1$ (ii) $\mathrm{tw}(G) = 2$ (iii) $\mathrm{tw}(k \times k\text{-grid}) = k$ (iv) $\mathrm{tw}(K_n) = n - 1$

Figure 2.2: (i) Graphs with treewidth one are precisely the trees, and those multigraphs which become trees after simplification. (ii) A graph of treewidth two. (iii) The $k \times k$-grid (in the above example $k = 5$) shows that planar graphs can have arbitrary large treewidth. (iv) The complete graph K_n on n nodes has treewidth $n - 1$

Definition 2.3 (cutset, cutwidth)**.** Let $G = (V, E)$ be a graph with $V = \{v_1, \ldots, v_n\}$. Given a permutation, i.e., a bijection $\sigma \colon [n] \to [n]$ and an integer $\ell \in [n-1]$, we can consider the *cutset* $C(\sigma, \ell) = \{\{v_i, v_j\} \in E : \sigma(i) \leq \ell < \sigma(j)\}$. The *cutwidth* $\mathrm{cw}(G)$ is the size of the largest cutset minimized over all permutations of V. More formally,

$$\mathrm{cw}(G) = \min_{\sigma} \max_{\ell} |C(\sigma, \ell)|.$$

Cutwidth and pathwidth are closely related: for bounded-degree graphs they are within a constant factor. Let $\mathrm{D}(G)$ denote the maximum degree of a node in G.

Theorem 2.4 (Bodlaender, Theorems 47 and 49 from [15][2])**.** *Given a graph G, we have*

$$\mathrm{pw}(G) \leq \mathrm{cw}(G) \leq \mathrm{D}(G)\,\mathrm{pw}(G).$$

Motivated by the study of communication networks, Bienstock introduced *congestion* [12], a generalization of cutwidth, which is a quantity related to treewidth in a similar way as cutwidth to pathwidth (compare Theorems 2.4 and 2.6).

Let us consider two graphs G and H, called the *guest* and the *host*, respectively. An *embedding* $\mathcal{E} = (\iota, \rho)$ of G into H consists of an injective mapping $\iota \colon V(G) \to V(H)$ together with a routing ρ that assigns to each arc $\{u, v\} \in E(G)$ a path in H with endpoints $\iota(u)$ and $\iota(v)$. If $e \in E(G)$ and $h \in E(H)$ is on the path $\rho(e)$, then we say that

[2]The inequality $\mathrm{cw}(G) \leq \mathrm{D}(G)\,\mathrm{pw}(G)$ seems to be already present in the earlier work of Chung and Seymour [37] on the relation of cutwidth to another parameter called topological bandwith (see Theorem 2 in [37]). However, the inequality is phrased and proved explicitly by Bodlaender in [15].

"e is running parallel to h." The congestion of G with respect to an embedding $\mathcal{E}$ of G into a host graph H, denoted as $\mathrm{cng}_{H,\mathcal{E}}(G)$, is defined to be the maximal number of times an arc of H is used in the routing of arcs of G. We also say that H is *realizing* congestion $\mathrm{cng}_{H,\mathcal{E}}(G)$. Several notions of congestion can be obtained by minimizing $\mathrm{cng}_{H,\mathcal{E}}(G)$ over various families of host graphs and embeddings (see, e.g., [115]). In this thesis we work with the following variant originally considered by Bienstock.

Definition 2.5 (Congestion[3])**.** Let $T_{\{1,3\}}$ be the set of unrooted binary trees.[4] The *congestion* $\mathrm{cng}(G)$ of a graph G is defined as

$$\mathrm{cng}(G) = \min\{\mathrm{cng}_{H,\mathcal{E}}(G) : H \in T_{\{1,3\}},\ \mathcal{E} = (\iota, \rho) \text{ with } \iota\colon V(G) \to L(H) \text{ bijection}\},$$

where $L(H)$ denotes the set of leaves of H.

In other words, we minimize $\mathrm{cng}_{H,\mathcal{E}}(G)$ when the host graph H is an unrooted binary tree and the mapping ι maps the nodes of G bijectively onto the leaves of H. The routing ρ is uniquely determined as the host graph is a tree. See Figure 2.3.

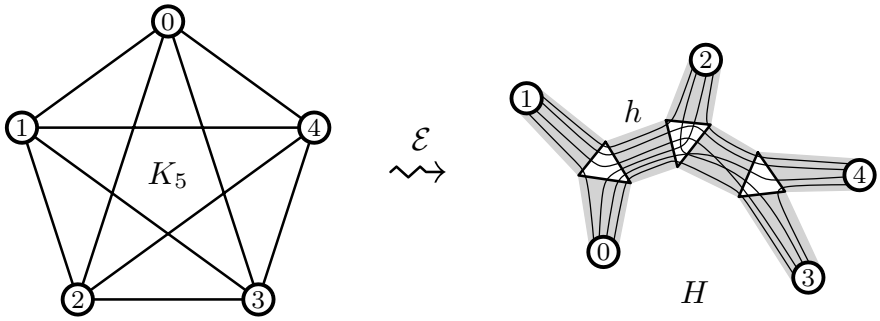

Figure 2.3: The complete graph K_5 (guest) routed along an unrooted binary tree H (host). We have $\mathrm{cng}_{H,\mathcal{E}}(K_5) = 6$ which is witnessed by h: six arcs of K_5 are running parallel to h

Theorem 2.6 (Bienstock [11, p. 108–111][5])**.** *For a graph G with largest degree* $\mathrm{D}(G)$*,*

$$\max\left\{\tfrac{2}{3}(\mathrm{tw}\,(G)+1), \mathrm{D}(G)\right\} \le \mathrm{cng}(G) \le \mathrm{D}(G)(\mathrm{tw}\,(G)+1).$$

[3]It is important to note that congestion in the sense of Definition 2.5 is also known as *carving width*, a term which was coined by Robertson and Seymour in [138]. However, the usual abbreviation for carving width is 'cw' which clashes with that of the cutwidth. Therefore we stick to the name 'congestion' and the abbreviation 'cng' to avoid this potential confusion in notation.

[4]An *unrooted binary tree* is a tree in which each node is incident to either one or three arcs.

[5]Only the right-hand side inequality of Theorem 2.6, $\mathrm{cng}(G) \le (\mathrm{tw}\,(G)+1)\,\mathrm{D}(G)$, is formulated explicitly in [11] as Theorem 1 on p. 111, whereas the left-hand side inequality is stated "inline" in the preceding paragraphs on the same page.

Example 2.7 (The Petersen graph)**.** One of the most widely used examples in graph theory is the *Petersen graph*, denoted P, see Figure 2.4(i). Although it is not a dual graph of a 3-manifold triangulation (since it is not 4-regular), it turns out to be helpful for comparing the graph parameters considered in this article.

- $\mathbf{cw(P) = 6}$. Notice that for any $S \subset V = V(P)$ of cardinality four there are at least six arcs running between S and $V \setminus S$. That is, on the one hand, $\mathrm{cw}(P) \geq 6$. On the other hand, it is easily verified that in the linear ordering $0 < 1 < 2 < \ldots < 9$ the maximal cutset has size six.

- $\mathbf{pw(P) = 5}$. A minimal-width path decomposition (which we computed using the module 'Vertex separation' of *SageMath* [44]) is the following.

$$\{0\} - \{0,1\} - \{0,1,2\} - \{0,1,2,4\} - \{0,1,2,4,5\} - \{1,2,3,4,5\}$$
$$\{4,6,7,9\} - \{3,4,5,6,7,8\} - \{2,3,4,5,6,7\} - \{1,2,3,4,5,6\}$$

- $\mathbf{tw\,(P) = 4}$. An optimal tree decomposition (computed using *SageMath* [44]) is shown in Figure 2.4(ii).

- $\mathbf{cng(P) = 5}$. Every arc $e \in E(H)$ of a host tree H specifies a cut in P, i.e., by deleting e, the leaves of the two components of $H \setminus e$ correspond to a partition $V(P) = S_e \cup R_e$. It is easy to see that, for any host tree H there is an edge $e \in E(H)$ with $\{\#S_e, \#R_e\} \in \{\{3,7\},\{4,6\},\{5,5\}\}$, and that every such cut contains at least five arcs of P, hence $\mathrm{cng}(P) \geq 5$. The reverse inequality is proven through the example shown in Figure 2.4(iii).

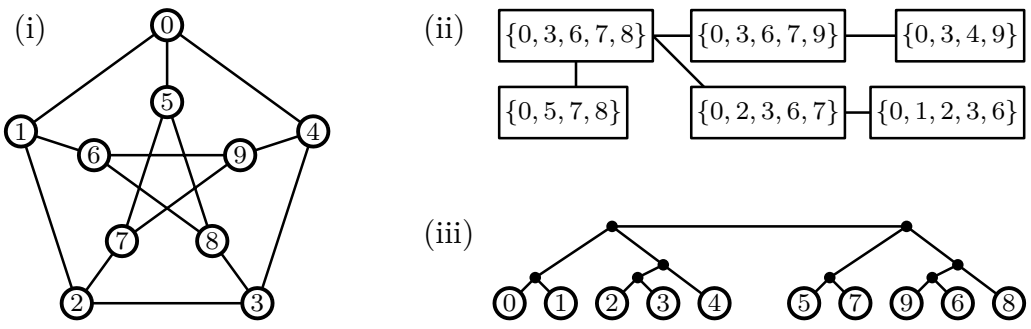

Figure 2.4: (i) The Petersen graph, (ii) a tree decomposition with minimal treewidth, and (iii) an unrooted binary tree realizing minimal congestion

2.2 Parameterized Complexity

There exist various concepts and notions that provide a refined complexity analysis for theoretically difficult problems (e.g., average-case complexity or, more recently, generic-case complexity [75]). In the field of *parameterized complexity theory*, systematized by Downey and Fellows [47] (also see the textbooks [40, 48] for an introduction), a parameter is identified on the input set, which is responsible for the hardness of a given problem.

More precisely, for a problem $\mathcal{P}$ with input set $\mathcal{I}$, a parameter is a computable function $p\colon \mathcal{I} \to \mathbb{N}$. If the parameter p is the output of $\mathcal{P}$, then p is called the *natural parameter*. The problem $\mathcal{P}$ is said to be *fixed-parameter tractable in the parameter* p (or FPT in p, for short) if there exists an algorithm which solves $\mathcal{P}$ for every instance $I \in \mathcal{I}$ with running time $O(f(p(I)) \cdot \mathrm{poly}(n))$, where n is the size of the input I, and $f\colon \mathbb{N} \to \mathbb{N}$ is a computable function. By definition, such an algorithm then runs in polynomial time on the set of inputs with bounded p. Hence, this identifies, in some sense, p as a potential "source of hardness" for $\mathcal{P}$.

The FPT-results cited in Chapter 1 use the treewidth of the dual graph of the input triangulation as a parameter [29, 31, 32, 33]. Similarly, for computational problems in knot theory, it has been very fruitful to work with width parameters of the given knot diagram (considered as a 4-regular multigraph), see, e.g., [18, 95, 96, 98]. Note that these parameters can have very different values for different triangulations (resp. diagrams) of the same 3-manifold (resp. knot). In particular, every 3-manifold (resp. knot) admits a triangulation (resp. diagram) with arbitrarily high width parameters—for all parameters considered in this article. See, e.g., Appendix B.2.

For this reason, in computational 3-manifold topology, a very important family of parameters is the one of *topological invariants*, i.e., properties which only depend on the topology of a given manifold and are independent of the choice of triangulation. At present, the algorithmic scope of these parameters remains largely unexplored. See [100] for such a result, using the first Betti number as parameter.

For computational aspects and comparison of the various graph parameters considered in this thesis we refer to Appendix A.

3 A Primer on 3-Manifolds

The main objects of study in this thesis are 3-dimensional manifolds, or 3-manifolds for short. As we will also work with 2-manifolds, also known as *surfaces*, we give the general definition. A *d-dimensional manifold with boundary* is a topological space[1] $\mathcal{M}$ such that each point $x \in \mathcal{M}$ has a neighborhood which looks like (i.e., is homeomorphic to) the Euclidean d-space $\mathbb{R}^d$ or the closed upper half-space $\{(x_1, \ldots, x_d) \in \mathbb{R}^d : x_d \geq 0\}$. The points of $\mathcal{M}$ that do not have a neighborhood homeomorphic to $\mathbb{R}^d$ constitute the *boundary* $\partial\mathcal{M}$ of $\mathcal{M}$. A compact manifold is said to be *closed* if it has an empty boundary.

Two manifolds $\mathcal{M}_1$ and $\mathcal{M}_2$ are considered equivalent, simply denoted $\mathcal{M}_1 = \mathcal{M}_2$, if they are *homeomorphic*, i.e., if there exists a continuous bijection $f\colon \mathcal{M}_1 \to \mathcal{M}_2$ with f^{-1} being continuous as well. So when talking about a manifold, we are ultimately not interested in its particular shape, but rather in those properties thereof, which are preserved under homeomorphisms. These are called *topological invariants*.

We refer to [134] for an introduction to 3-manifolds (cf. [57, 60, 69, 128, 147]).

All 3-manifolds considered in this thesis are assumed to be compact and orientable.

3.1 Triangulations

A classical result of Moise [108] (cf. [13]) asserts, that every compact 3-manifold admits a triangulation (in fact, infinitely many). To build a triangulation, take a disjoint union $\widetilde{\Delta} = \Delta_1 \cup \ldots \cup \Delta_n$ of finitely many tetrahedra with $4n$ triangular faces altogether. Let $\Phi = \{\varphi_1, \ldots, \varphi_m\}$ be a set of at most $2n$ *face gluings*, i.e., simplicial homeomorphisms, each of which identifies a distinct pair of these triangular faces. Performing all the gluings,

[1] More precisely, we only consider topological spaces which are *second countable* and *Hausdorff*.

the resulting quotient space $\mathcal{T} = \widetilde{\Delta}/\Phi$ is called a *triangulation*, and the pairs of identified triangular faces are referred to as *triangles* of $\mathcal{T}$. Note that these face gluings also identify several tetrahedral edges (or vertices) of $\widetilde{\Delta}$ resulting in a single *edge* (or *vertex*) of $\mathcal{T}$.

The face gluings, however, cannot be arbitrary. A triangulation $\mathcal{T}$ is homeomorphic to a closed 3-manifold $\mathcal{M}$, if and only if the boundary of a small neighborhood around each vertex is $\mathbb{S}^2$, and no edge is identified with itself in reverse.[2] If some of the vertices have small neighborhoods with boundaries being disks, then $\mathcal{T}$ describes a 3-manifold with boundary. In computational topology a 3-manifold is very often presented this way.[3]

To explicitly describe a face gluing $\varphi \in \Phi$, we label the four vertices of the tetrahedron $\Delta_i \in \widetilde{\Delta}$ as $\Delta_i(0)$, $\Delta_i(1)$, $\Delta_i(2)$, and $\Delta_i(3)$. The expression $\Delta_i(123) \xleftrightarrow{\varphi} \Delta_j(103)$ then means that φ identifies $\Delta_i(1)$ with $\Delta_j(1)$, $\Delta_i(2)$ with $\Delta_j(0)$, and $\Delta_i(3)$ with $\Delta_j(3)$.

In the study of triangulations, their dual graphs play an instrumental role.[4] Given a triangulation $\mathcal{T} = \widetilde{\Delta}/\Phi$, its *dual graph* $\Gamma(\mathcal{T}) = (V, E)$ is a multigraph where the nodes in V correspond to the tetrahedra in $\widetilde{\Delta}$, and for each face gluing $\varphi \in \Phi$ that identifies two triangular faces of Δ_i and Δ_j, we add an arc between the corresponding nodes in V, cf. Figure 3.1. (Note that i and j could be equal.) By construction, $\Gamma(\mathcal{T})$ has maximum degree ≤ 4. Moreover, when $\mathcal{T}$ triangulates a closed 3-manifold, then $\Gamma(\mathcal{T})$ is 4-regular.

Remark 3.1. A given pair of triangular faces can be identified six different ways, thus by passing to the dual graph $\Gamma(\mathcal{T})$ we forget the exact structure of $\mathcal{T}$. Still, some information is retained even about the topology of the underlying 3-manifold, cf. Chapter 6.

For a triangulation $\mathcal{T}$, we denote its *size*, i.e., the number of tetrahedra in $\mathcal{T}$, with $|\mathcal{T}|$. Finally, the *triangulation complexity* $\mathrm{tc}(\mathcal{M})$ of a given 3-manifold $\mathcal{M}$ is defined to be the minimum number of tetrahedra in any triangulation of $\mathcal{M}$.[5]

---

[2] Otherwise the midpoint of such an edge would have a small neighborhood with $\mathbb{R}P^2$ boundary.

[3] Triangulations (also called *generalised*, *semi-simplicial*, or *singular* triangulations in the literature) provide a convenient and efficient way to encode 3-manifolds for computational purposes. There exist 13 399 distinct closed, orientable 3-manifolds which admit a triangulation with at most 11 tetrahedra [25]. Note that a triangulation, in the present sense, is not necessarily a simplicial complex. However, it can always be turned into one by taking at most two barycentric subdivisions. See [126, Section 2] for more detailed introduction to 3-manifold triangulations.

[4] The dual graph of a triangulation is also known as the *face pairing graph*, or as the *dual* 1*-skeleton*.

[5] Except for $\mathbb{S}^3$, $\mathbb{R}P^3$ and the lens space $L(3,1)$, the triangulation complexity of a 3-manifold is the same as its 'complexity' introduced and popularized by Matveev, cf. Chapter 2 and Remark 2.1.7 in [103]. We adopted the term 'triangulation complexity' from the recent preprint [91].

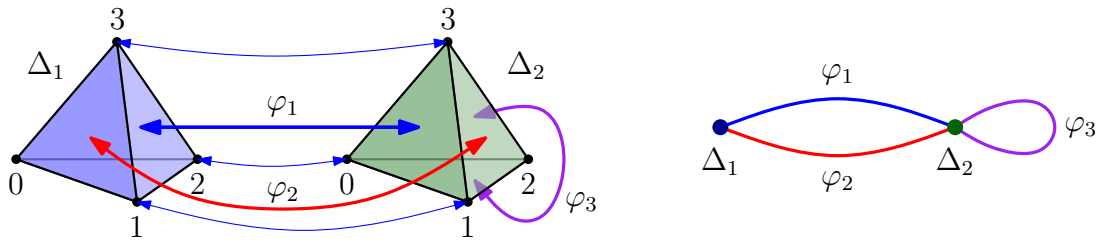

(i) $\mathcal{T}$ (ii) $\Gamma(\mathcal{T})$

Figure 3.1: (i) A triangulation $\mathcal{T} = \widetilde{\Delta}/\Phi$ with two tetrahedra $\widetilde{\Delta} = \{\Delta_1, \Delta_2\}$ and three face gluing maps $\Phi = \{\varphi_1, \varphi_2, \varphi_3\}$. The map φ_1 is specified to be $\Delta_1(123) \xleftrightarrow{\varphi_1} \Delta_2(103)$. (ii) The dual graph $\Gamma(\mathcal{T})$ of the triangulation $\mathcal{T}$

3.2 Handle Decompositions

It follows from Morse theory that every compact 3-manifold can be built from finitely many solid building blocks called 0-, 1-, 2-, and 3-*handles* [134, Appendix B]. In such a *handle decomposition* all handles are homeomorphic to 3-balls, and are only distinguished in how they are glued to the existing decomposition. For instance, to construct a closed 3-manifold from handles, we may start with a disjoint union of 3-balls, or 0-*handles*, where further 3-balls are glued to the boundary of the existing decomposition along pairs of 2-dimensional disks, the so-called 1-*handles*, or along annuli, the so-called 2-*handles*. This process is iterated until the boundary of the decomposition consists of a union of 2-spheres. These are then eliminated by gluing in one additional 3-ball per boundary component, the 3-*handles* of the decomposition.

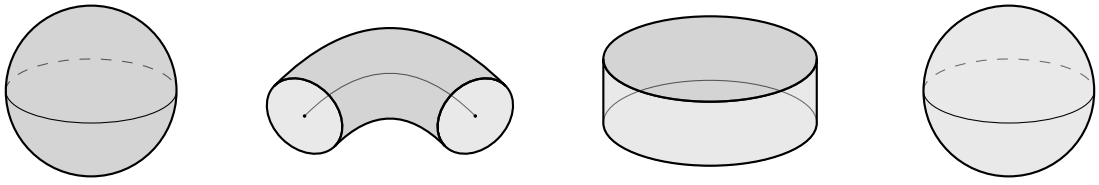

(i) (ii) (iii) (iv)

Figure 3.2: (i) A 0-handle, (ii) a 1-handle, (iii) a 2-handle, and (iv) a 3-handle. The attaching sites are indicated with light gray. For a 1-handle, this is a disjoint union of two disks, for a 2-handle an annulus, and for a 3-handle the entire boundary

In every step of building up a (closed) 3-manifold $\mathcal{M}$ from handles, the existing decomposition is a submanifold whose boundary—called a *bounding surface*—separates $\mathcal{M}$ into

two pieces: the part that is already present, and its complement (each of them possibly disconnected). Bounding surfaces and, more generally, all kinds of surfaces embedded in a 3-manifold, play an important role in the study of 3-manifolds (similar to that of simple closed curves in the study of surfaces). When chosen carefully, an embedded surface reveals valuable information about the topology of the ambient 3-manifold.

Any triangulation gives rise to a handle decomposition in a very natural way.

Definition 3.2. Let $\mathcal{T}$ be a triangulation of a closed 3-manifold $\mathcal{M}$. The *canonical handle decomposition* $\mathrm{chd}(\mathcal{T})$ *of* $\mathcal{M}$ *associated with* $\mathcal{T}$ consists of

- one 0-handle for the interior of each tetrahedron of $\mathcal{T}$,
- one 1-handle for a thickened version of the interior of each triangle of $\mathcal{T}$,
- one 2-handle for a thickened version of the interior of each edge of $\mathcal{T}$, and
- one 3-handle for a neighborhood of each vertex of $\mathcal{T}$.

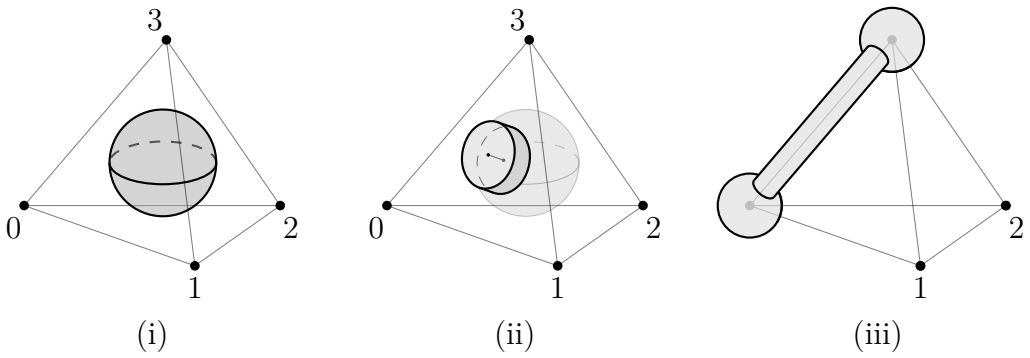

Figure 3.3: (i) A 0-handle in the interior of a tetrahedron. (ii) The 1-handle corresponding to the triangle $\{0, 1, 3\}$. (iii) A 2-handle corresponding to the edge $\{0, 3\}$, and the two 3-handles adjacent to it

Remark 3.3. In Definition 3.2 we associate 0-handles with tetrahedra and 3-handles with vertices of the triangulation. This is motivated by the fact that we model this decomposition on the dual graph rather than on the triangulation itself. The reason for this choice, in turn, is that it is the dual graph of a triangulation which acts as an intermediary between the topology of a 3-manifold and the framework of structural graph theory, which we will exploit in Section 5.2.

Remark 3.4. The notion of "thickened version" in Definition 3.2 can be made precise via barycentric subdivisions. Given a triangulation $\mathcal{T}$, let $\mathrm{sd}_1(\mathcal{T})$ and $\mathrm{sd}_2(\mathcal{T})$ denote its first

and second barycentric subdivision, respectively. Note that vertices of $\mathrm{sd}_1(\mathcal{T})$ are in one-to-one correspondence with the simplices of $\mathcal{T}$. Let $v \in \mathrm{sd}_1(\mathcal{T})$ be a vertex corresponding to an i-simplex σ of $\mathcal{T}$. The $(3-i)$-handle of $\mathrm{chd}(\mathcal{T})$ corresponding to σ is then the union of all simplices in $\mathrm{sd}_2(\mathcal{T})$ that are incident to v. See Figure 3.4 for a 2-dimensional illustration. The 3-dimensional case is analogous.

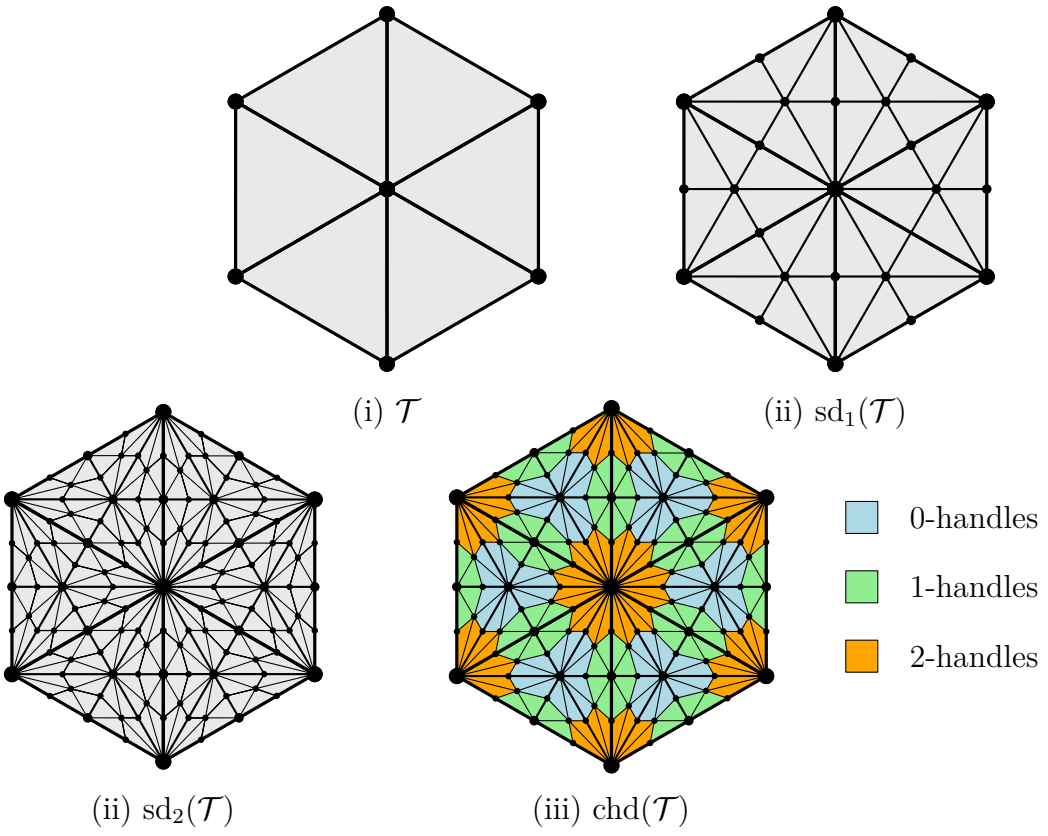

Figure 3.4: The canonical handle decomposition of a 2-dimensional triangulation

3.3 Surfaces in 3-Manifolds

Given a 3-manifold $\mathcal{M}$, a surface $\mathcal{S} \subset \mathcal{M}$ is said to be *properly embedded*, if it is embedded in $\mathcal{M}$, and for the boundary we have $\partial\mathcal{S} = \mathcal{S} \cap \partial\mathcal{M}$. Let $\mathcal{S} \subset \mathcal{M}$ be a properly embedded surface distinct from the 2-sphere, and let D be a disk embedded into $\mathcal{M}$ such that its boundary satisfies $\partial D = D \cap \mathcal{S}$. D is said to be a *compressing disk for* $\mathcal{S}$ if ∂D does not bound a disk on $\mathcal{S}$. If such a compressing disk exists, then $\mathcal{S}$ is called *compressible*, otherwise it is called *incompressible*. An embedded 2-sphere $\mathcal{S} \subset \mathcal{M}$ is called *incompressible* if $\mathcal{S}$ does not bound a 3-ball in $\mathcal{M}$.

Example 3.5. A basic example of a compressible surface is a torus embedded in the 3-sphere $\mathbb{S}^3$, and of an incompressible surface is the 2-sphere $\mathbb{S}^2 \times \{x\} \subset \mathbb{S}^2 \times \mathbb{S}^1$.

A connected 3-manifold $\mathcal{M}$ is called *irreducible*, if every embedded 2-sphere bounds a 3-ball in $\mathcal{M}$. Moreover, it is called P^2*-irreducible*, if it does not contain an embedded 2-sided[6] real projective plane $\mathbb{R}P^2$. This notion is only significant for non-orientable manifolds, since orientable 3-manifolds cannot contain any 2-sided non-orientable surfaces, and are readily P^2-irreducible. If a P^2-irreducible, irreducible 3-manifold $\mathcal{M}$ contains a 2-sided incompressible surface, then it is called *Haken*, otherwise it is called *non-Haken*.

Finally, let $\mathcal{S}$ be a compact (not necessarily connected or orientable) surface. The *genus* of $\mathcal{S}$, denoted by $g(\mathcal{S})$, equals the maximum number of pairwise disjoint simple closed curves one can remove from $\mathcal{S}$ without increasing the number of connected components. The genus and the orientability determine a closed surface up to homeomorphism. We denote the closed orientable (resp. non-orientable) surface of genus g by $\mathcal{F}_g$ (resp. $\mathcal{N}_g$).

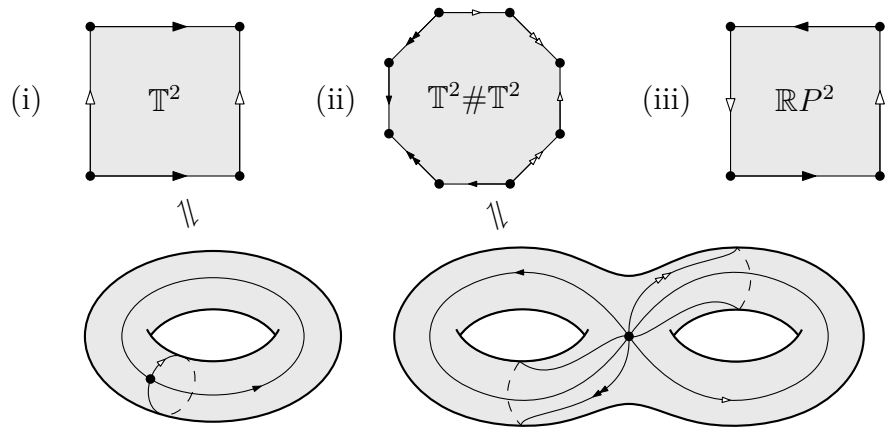

Figure 3.5: (i) The 2-torus $\mathbb{T}^2 = \mathcal{F}_1$, i.e., the closed orientable surface of genus one, can be obtained by identifying the opposite sides of a square. (ii) The closed orientable surface of genus two is the connected sum of two $\mathbb{T}^2$'s. (iii) The real projective plane $\mathbb{R}P^2 = \mathcal{N}_1$

3.4 Handlebodies and Compression Bodies

Similarly to tetrahedra, handles (seen in Section 3.2) can be thought of as fundamental building blocks of 3-manifolds. Here we discuss handlebodies and compression bodies,

[6] A properly embedded surface $\mathcal{S} \subset \mathcal{M}$ is 2-sided in $\mathcal{M}$, if the codimension zero submanifold in $\mathcal{M}$ obtained by thickening $\mathcal{S}$ has two boundary components, i.e., $\mathcal{S}$ locally separates $\mathcal{M}$ into two pieces.

built from several handles, which provide other ways to construct 3-manifolds.

A *handlebody* $\mathcal{H}$ is a connected 3-manifold with boundary that can be described as a single 0-handle with a number of 1-handles attached to it, or, equivalently, as a thickened graph. Up to homeomorphism, $\mathcal{H}$ is determined by the genus $g(\partial\mathcal{H})$ of its boundary.

Let $\mathcal{S}$ be a compact, orientable (not necessarily connected) surface. A *compression body* is a 3-manifold $\mathcal{C}$ obtained from $\mathcal{S}\times[0,1]$ by (optionally) attaching some 1-handles to $\mathcal{S}\times\{1\}$, and (optionally) filling in some of the 2-sphere components of $\mathcal{S}\times\{0\}$ with 3-balls. $\mathcal{C}$ has two sets of boundary components: $\partial_-\mathcal{C} = \mathcal{S}\times\{0\}\setminus\{\text{filled-in 2-sphere components}\}$ and $\partial_+\mathcal{C} = \partial\mathcal{C}\setminus\partial_-\mathcal{C}$. We call $\partial_+\mathcal{C}$ the *upper boundary*, and $\partial_-\mathcal{C}$ the *lower boundary* of $\mathcal{C}$.

Dual to this construction, a compression body $\mathcal{C}$ can also be built by starting with a closed, orientable surface $\mathcal{F}$, thickening it to $\mathcal{F}\times[0,1]$, (optionally) attaching some 2-handles along $\mathcal{F}\times\{0\}$, and (optionally) filling in some of the resulting 2-spheres with 3-balls. The upper and lower boundary are given by $\partial_+\mathcal{C} = \mathcal{F}\times\{1\}$ and $\partial_-\mathcal{C} = \partial\mathcal{C}\setminus\partial_+\mathcal{C}$.

See Figures 3.6 and 3.7 for schematics and illustrations, respectively.

Remark 3.6. Let $\mathcal{C}$ be a compression body constructed as above.

(1) Note that for the genera of the boundaries we always have $g(\partial_+\mathcal{C}) \geq g(\partial_-\mathcal{C})$.

(2) If $\partial_-\mathcal{C} = \emptyset$, then the compression body $\mathcal{C}$ is actually a handlebody. This can be achieved if $\mathcal{S} = \mathbb{S}^2$ and we fill in the single 2-sphere boundary $\mathcal{S}\times\{0\}$ with a 3-ball.

(3) If $\mathcal{C}$ is homeomorphic to $\mathcal{S}\times[0,1]$ then $\mathcal{C}$ is called a *trivial compression body*. This happens if we omit all optional steps in the construction.

(4) If $\mathcal{C}$ is connected, it is reducible if and only if $\partial_-\mathcal{C}$ contains a 2-sphere component.

(5) Each component of the lower boundary $\partial_-\mathcal{C}$ is incompressible in $\mathcal{C}$.

(6) The upper boundary $\partial_+\mathcal{C}$ is incompressible in $\mathcal{C}$ if and only if $\mathcal{C}$ is trivial.

(i) (ii) (iii) Primal (iv) Dual

Figure 3.6: (i) Schematic of a handlebody $\mathcal{H}$. (ii)–(iv) Three schematics of the same compression body $\mathcal{C}$ emphasizing different viewpoints (cf. Figure 3.7)

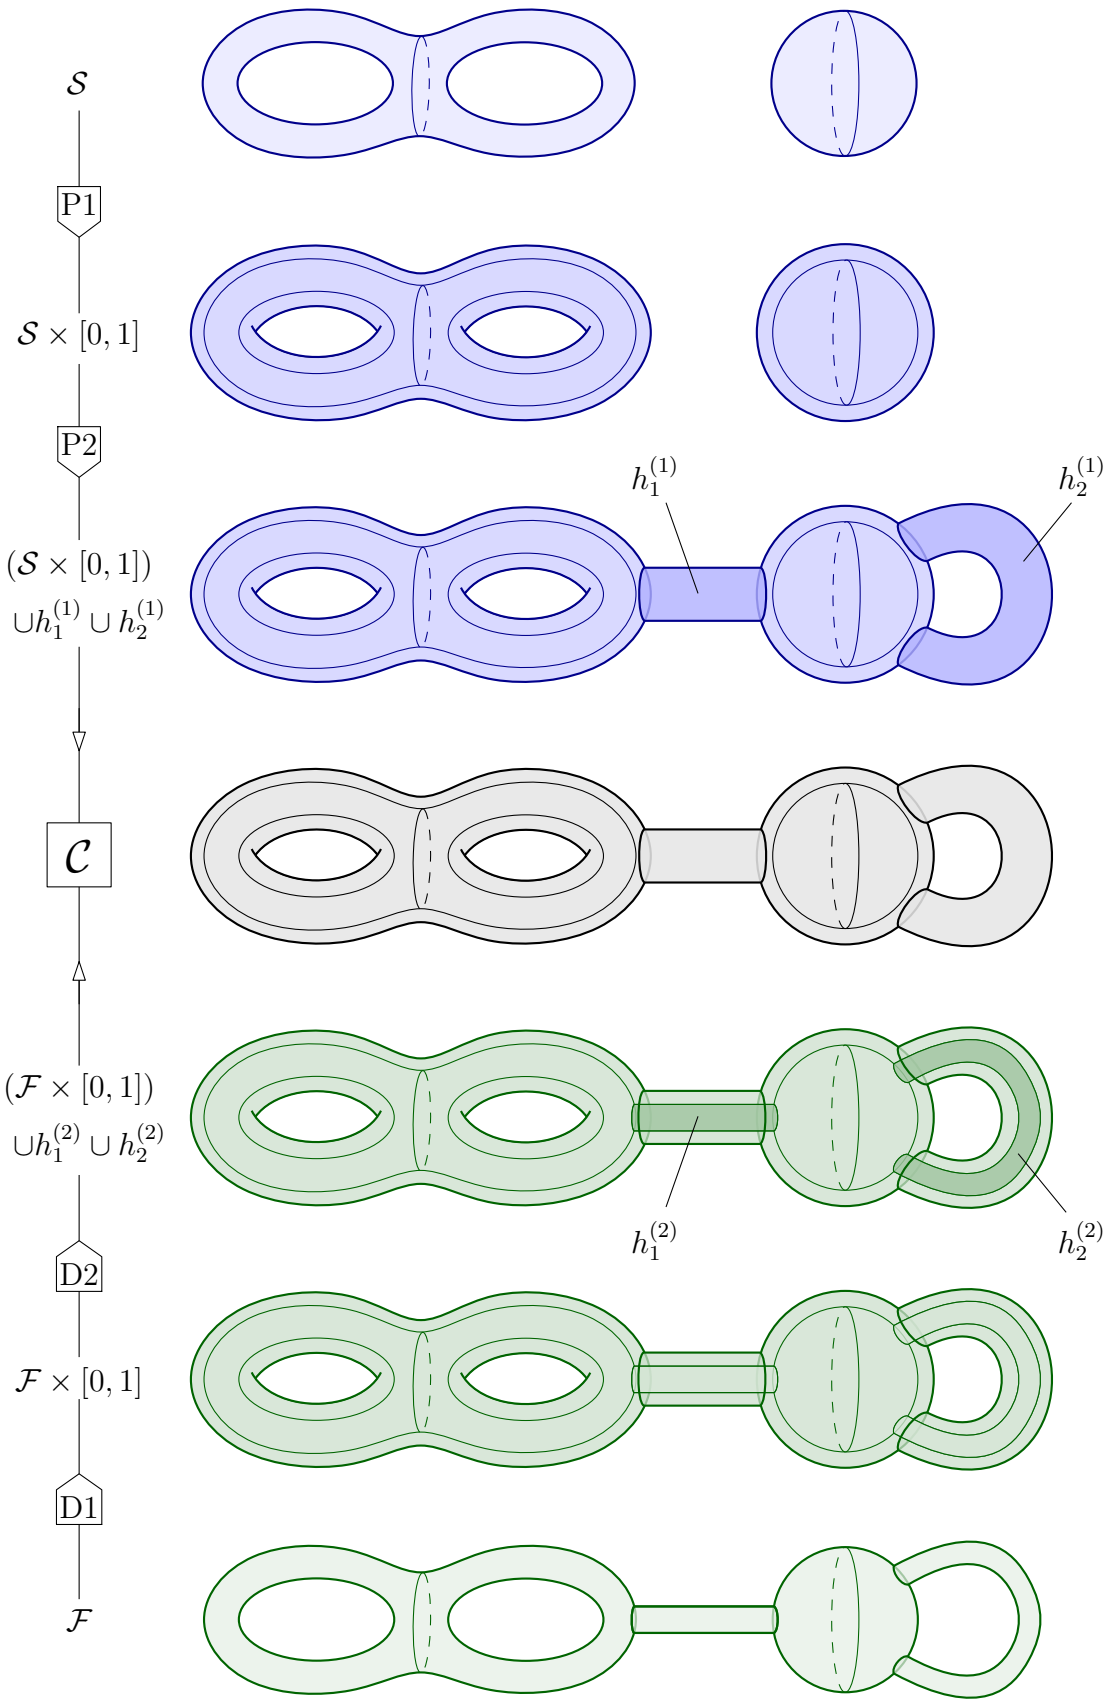

Figure 3.7: Two ways (primal and dual) of constructing a compression body $\mathcal{C}$

3.5 Heegaard Splittings

Introduced in [59], Heegaard splittings have been central to the study of 3-manifolds for over a century. Given a closed, orientable 3-manifold $\mathcal{M}$, a *Heegaard splitting* is a decomposition $\mathcal{M} = \mathcal{H} \cup_{\mathcal{S}} \mathcal{H}'$ where $\mathcal{H}$ and $\mathcal{H}'$ are homeomorphic handlebodies with $\mathcal{H} \cup \mathcal{H}' = \mathcal{M}$ and $\mathcal{H} \cap \mathcal{H}' = \partial\mathcal{H} = \partial\mathcal{H}' = \mathcal{S}$. The *Heegaard genus* $\mathfrak{g}(\mathcal{M})$ of $\mathcal{M}$ equals the smallest genus $g(\mathcal{S})$ over all Heegaard splittings of $\mathcal{M}$. Sometimes we think of a Heegaard splitting in a procedural way: Start with the disjoint union of two identical handlebodies, $\mathcal{H}$ and $\mathcal{H}'$, and identify their boundary surfaces via a homeomorphism $f\colon \partial\mathcal{H} \to \partial\mathcal{H}'$ referred to as the *attaching map*. We use the notation $\mathcal{M} = \mathcal{H} \cup_f \mathcal{H}'$ if we want to stress this viewpoint. Heegaard splittings with isotopic attaching maps yield homeomorphic 3-manifolds, hence are considered equivalent. See [129] for a comprehensive survey.

Example 3.7 (Heegaard splittings from triangulations, I)**.** Given a triangulation $\mathcal{T}$ of a closed, orientable 3-manifold $\mathcal{M}$, let $\mathcal{T}^{(1)}$ denote its 1-skeleton consisting of the vertices and edges of $\mathcal{T}$. Thickening up $\mathcal{T}^{(1)}$, i.e., taking its regular neighborhood, in $\mathcal{M}$ yields a handlebody $\mathcal{H}_1$. The closure $\mathcal{H}_2$ of the complement $\mathcal{M} \setminus \mathcal{H}_1$ is also a handlebody homeomorphic to a regular neighborhood of $\Gamma(\mathcal{T})$, and $\mathcal{M} = \mathcal{H}_1 \cup \mathcal{H}_2$ is a Heegaard splitting of $\mathcal{M}$.

Note that $\mathcal{H}_1$ is the union of 2- and 3-handles of the canonical handle decomposition $\mathrm{chd}(\mathcal{T})$, and $\mathcal{H}_2$ is the union of the 0- and 1-handles thereof, cf. Definition 3.2.

The following proposition is folklore.

Proposition 3.8. *Let $\mathcal{T}$ be a triangulation of a closed 3-manifold $\mathcal{M}$ consisting of n tetrahedra. Then the Heegaard splitting considered in Example 3.7 is of genus $n + 1$.*

Proof. This is a straightforward Euler characteristic computation. First of all, $\mathcal{H}_2$, by construction, is homotopy equivalent to $\Gamma(\mathcal{T})$, thus their Euler characteristic is the same. Second, the nodes and arcs of the dual graph $\Gamma(\mathcal{T})$ correspond to the tetrahedra and triangles of $\mathcal{T}$, respectively. Third, since the underlying 3-manifold $\mathcal{M}$ is assumed to be closed, the triangulation $\mathcal{T}$ with n tetrahedra necessarily has $2n$ triangles. Therefore for the Euler characteristic of $\mathcal{H}_2$ we have $\chi(\mathcal{H}_2) = \chi(\Gamma(\mathcal{T})) = n - 2n = -n$, which in turn yields $g(\partial\mathcal{H}_2) = 1 - \chi(\mathcal{H}_2) = 1 + n$. □

Remark 3.9. Proposition 3.8 implies that the Heegaard genus $\mathfrak{g}(\mathcal{M})$ and the triangulation complexity $\mathrm{tc}(\mathcal{M})$ of any 3-manifold $\mathcal{M}$ satisfy $\mathfrak{g}(\mathcal{M}) + 1 \leq \mathrm{tc}(\mathcal{M})$. (Moreover, this bound is sharp for infinitely many 3-manifolds [50].) This gives a flavor of the endeavors central to our work: minimize a combinatorial quantity over all triangulations of a given 3-manifold $\mathcal{M}$ and relate it to classical topological invariants of $\mathcal{M}$.

Heegaard splittings of 3-manifolds with boundary. Using compression bodies, one can generalize Heegaard splittings to 3-manifolds with nonempty boundary. Let $\mathcal{M}$ be a 3-manifold and $\partial_1\mathcal{M} \cup \partial_2\mathcal{M} = \partial\mathcal{M}$ be an arbitrary partition of its boundary components. Then there exist compression bodies $\mathcal{C}_1$ and $\mathcal{C}_2$ with $\mathcal{C}_1 \cup \mathcal{C}_2 = \mathcal{M}$, $\partial_-\mathcal{C}_1 = \partial_1\mathcal{M}$, $\partial_-\mathcal{C}_2 = \partial_2\mathcal{M}$, and $\mathcal{C}_1 \cap \mathcal{C}_2 = \partial_+\mathcal{C}_1 = \partial_+\mathcal{C}_2$. (See Example 3.10, based on [130, Theorem 2.1.11], for a construction similar to that in Example 3.7. Also see [129, Section 2.2].) The decomposition $\mathcal{M} = \mathcal{C}_1 \cup_{\mathcal{S}} \mathcal{C}_2$ is called a Heegaard splitting of $\mathcal{M}$ *compatible with the partition* $\partial_1\mathcal{M} \cup \partial_2\mathcal{M}$. Its splitting surface, or *Heegaard surface*, is $\mathcal{S} = \mathcal{C}_1 \cap \mathcal{C}_2$. The Heegaard genus $\mathfrak{g}(\mathcal{M})$ is again the minimum genus $g(\mathcal{S})$ over all such decompositions.

Example 3.10 (Heegaard splittings from triangulations, II)**.** Let $\mathcal{T}$ be a triangulation of $\mathcal{M}$ with partition $\partial_1\mathcal{M} \cup \partial_2\mathcal{M}$ of its boundary components. Suppose that no simplex in $\mathcal{T}$ is incident to more than one component of $\partial\mathcal{M}$.[7] Take the first barycentric subdivision $\mathrm{sd}_1(\mathcal{T})$ of $\mathcal{T}$. Recall that $\mathcal{T}^{(1)}$ and $\Gamma(\mathcal{T})$ denote the 1-skeleton and the dual graph of $\mathcal{T}$. Their first barycentric subdivisons $\mathcal{T}^{(1)}_{\mathrm{sd}}$ and $\Gamma(\mathcal{T})_{\mathrm{sd}}$ are both naturally contained in $\mathrm{sd}_1(\mathcal{T})$. Consider the subcomplex $N(\partial_2\mathcal{M}) \subset \mathrm{sd}_1(\mathcal{T})$ consisting of all simplices incident to $\partial_2\mathcal{M}$. We define two further subcomplexes of $\mathrm{sd}_1(\mathcal{T})$, namely

$$\Gamma_1 = \partial_1\mathcal{M} \cup \{\text{vertices and edges of } \mathcal{T}^{(1)}_{\mathrm{sd}} \textit{ not} \text{ incident to } \partial_2\mathcal{M}\}, \text{ and}$$
$$\Gamma_2 = N(\partial_2\mathcal{M}) \cup \Gamma(\mathcal{T})_{\mathrm{sd}}.$$

Now pass to the second barycentric subdivision $\mathrm{sd}_2(\mathcal{T})$ and let $(\Gamma_i)_{\mathrm{sd}}$ denote the image of Γ_i under this operation ($i = 1, 2$). Let $\eta(\Gamma_i)$ be the "thickening" of Γ_i, i.e., the subcomplex of $\mathrm{sd}_2(\mathcal{T})$ formed by all simplices incident to $(\Gamma_i)_{\mathrm{sd}}$. One can readily verify that $\eta(\Gamma_1)$ and $\eta(\Gamma_2)$ are compression bodies whose union is $\mathcal{M}$, their upper boundaries satisfy $\partial_+\eta(\Gamma_1) = \partial_+\eta(\Gamma_2) = \eta(\Gamma_1) \cap \eta(\Gamma_2)$, and for their lower boundaries $\partial_-\eta(\Gamma_1) = \partial_1\mathcal{M}$ and $\partial_-\eta(\Gamma_2) = \partial_2\mathcal{M}$. Hence $\eta(\Gamma_1)$ and $\eta(\Gamma_2)$ form a Heegaard splitting of $\mathcal{M}$ compatible with the given partition of its boundary components. See Figure 3.8 for an example.

[7] This can be achieved, e.g., by passing to the first subdivision of $\mathcal{T}$ if necessary.

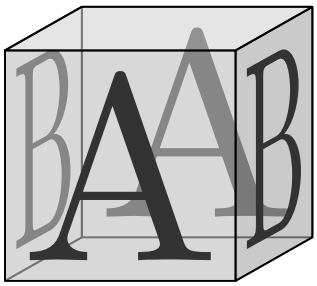

(i) $\mathcal{M} = \mathbb{T}^2 \times [0,1]$

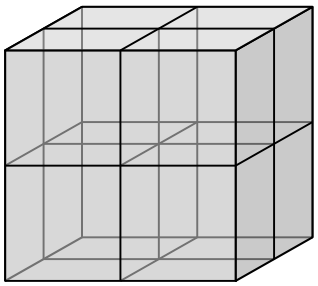

(ii) A quadrangulation $\mathcal{Q}$ of $\mathcal{M}$ with eight cubes

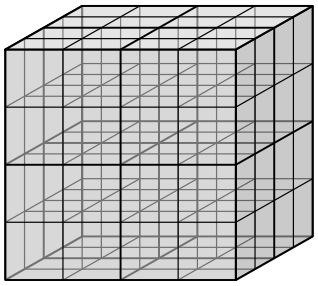

(iii) The first barycentric subdivision $\mathrm{sd}_1(\mathcal{Q})$ of $\mathcal{Q}$

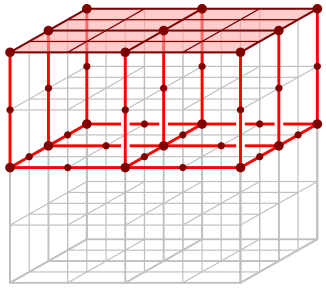

(iv) $\boldsymbol{\Gamma_1}$

$\boldsymbol{\Gamma_1} = \partial_1\mathcal{M} \cup \{\text{vertices \& edges of } \mathcal{T}^{(1)}_{\mathrm{sd}} \text{ avoiding } N(\partial_2\mathcal{M})\}$

$\boldsymbol{\Gamma_2} = N(\partial_2\mathcal{M}) \cup \Gamma(\mathcal{T})_{\mathrm{sd}}$

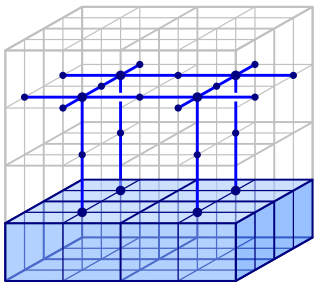

(v) $\boldsymbol{\Gamma_2}$

Figure 3.8: Constructing a Heegaard splitting of $\mathbb{T}^2 \times [0,1]$ from a quadrangulation

3.6 Seifert Fibered Spaces

Seifert fibered spaces, first described and classified in [137], comprise an important class of 3-manifolds. Here we describe those that are closed and orientable following [128].

Let us consider the surface $\mathcal{F}_{g,r} = \mathcal{F}_g \setminus (\operatorname{int} D_1 \cup \cdots \cup \operatorname{int} D_r)$ obtained from the closed, connected, orientable genus g surface by removing the interiors of r pairwise disjoint disks. Taking the product with the circle $\mathbb{S}^1$ yields an orientable 3-manifold $\mathcal{F}_{g,r} \times \mathbb{S}^1$ whose boundary consists of r tori: $\partial(\mathcal{F}_{g,r} \times \mathbb{S}^1) = (\partial D_1) \times \mathbb{S}^1 \cup \cdots \cup (\partial D_r) \times \mathbb{S}^1$. For each $(\partial D_i) \times \mathbb{S}^1$, $1 \leq i \leq r$, we glue in a solid torus so that its meridian wraps a_i times around the meridian $(\partial D_i) \times \{y_i\}$ and b_i times around the longitude $\{x_i\} \times \mathbb{S}^1$ of $(\partial D_i) \times \mathbb{S}^1$. Here a_i and b_i are assumed to be coprime integers with $a_i \geq 2$, and the point $(x_i, y_i) \in (\partial D_i) \times \mathbb{S}^1$ is chosen arbitrarily. This way we obtain a closed orientable 3-manifold $\mathcal{M} = \mathrm{SFS}[\mathcal{F}_g : (a_1, b_1), \ldots, (a_r, b_r)]$ which is called the *Seifert fibered space over $\mathcal{F}_g$ with r exceptional (or singular) fibers*. In relation to $\mathcal{M}$, the surface $\mathcal{F}_g$ is referred to as the *base space* (or *base orbifold*, or *orbit surface*).

Example 3.11. *Lens spaces*, the 3-manifolds of Heegaard genus one, agree with Seifert fibered spaces over $\mathbb{S}^2$ having at most two singular fibers [128, p. 27].[8]

Non-orientable base spaces. With a slight modification of the above construction, one can obtain additional orientable Seifert fibered spaces having non-orientable base spaces. Beginning with $\mathcal{N}_g$, the non-orientable genus g surface, we pass to $\mathcal{N}_{g,r}$ by adding r punctures (i.e., by removing r pairwise disjoint open disks). At this point, however, instead of taking the product $\mathcal{N}_{g,r} \times \mathbb{S}^1$ (which yields a non-orientable 3-manifold) we consider the "orientable $\mathbb{S}^1$-bundle" over $\mathcal{N}_{g,r}$, denoted $\mathcal{N}_{g,r}\widetilde{\times}\mathbb{S}^1$, which has again r torus boundary components. As before, we conclude by gluing in r solid tori, specified by pairs of coprime integers (a_i, b_i) with $a_i \geq 2$, where $1 \leq i \leq r$. The notation for the resulting 3-manifold remains the same. See [94, Section 2] for a concrete and detailed description of orientable Seifert fibered spaces (cf. the classes $\{Oo, g\}$ and $\{On, g\}$ therein).

Remark 3.12. By leaving some of the boundary tori of $\mathcal{F}_{g,r} \times \mathbb{S}^1$ or $\mathcal{N}_{g,r}\widetilde{\times}\mathbb{S}^1$ intact, i.e., not filling them in with solid tori, we obtain Seifert fibered spaces with boundary.

3.7 On the Classification of 3-Manifolds

> *"Perhaps by the year 2000 our understanding of 3-manifolds and Kleinian groups will be solid, and the phenomena we now expect will be proven."*
>
> — William P. Thurston, 1982 [146]

A vast amount of research about 3-manifolds in the 20th century has been directly or indirectly motivated by the famous Poincaré Conjecture from 1904, according to which any simply connected,[9] closed 3-manifold is homeomeomorphic to the 3-dimensional sphere. In a celebrated effort, the conjecture was affirmatively resolved in 2003 by Perelman [116, 117, 118], along with Thurston's more general Geometrization Conjecture [146], by completing the "curvature flow" approach envisioned by Hamilton [55]. (We refer to [9, 81, 110] for complete accounts of the proof of the Geometrization Conjecture, to [4, Chapter 1] for a survey, and to [105, 107, 120, 140] for expository and historical overviews.)

[8] In particular, we regard $\mathbb{S}^2 \times \mathbb{S}^1$ (the SFS over $\mathbb{S}^2$ without exceptional fibers) to be a lens space, too.

[9] A topological space $\mathcal{X}$ is *simply connected* if its fundamental group $\pi_1(\mathcal{X})$ vanishes, or, informally, if any closed loop in $\mathcal{X}$ can be continuously tightened to a single point.

The Geometrization Conjecture can be seen a major step towards the classification of 3-manifolds. Very informally, it states that any compact 3-manifold can be decomposed into "geometric" pieces in a "canonical" way. An in-depth discussion of the Geometrization Conjecture is way beyond the scope of this thesis. In what follows, we provide a high-level overview in order to introduce some important notions we will encounter later.

The classification of surfaces. We start with a brief review of the classification of surfaces, i.e., 2-dimensional manifolds. This is a classical result, often part of advanced undergraduate or introductory graduate courses on topology. We refer to Chapter 1 of [102] for a concise treatment, and to [52] for an extensive one, including historical remarks.

The following operation is central to the classification of both 2- and 3-manifolds.

Definition 3.13 (connected sum)**.** Let $\mathcal{M}_1$ and $\mathcal{M}_2$ be compact, connected d-manifolds. Consider $\mathcal{M}_i' = \mathcal{M}_i \setminus \mathrm{int}(\mathcal{B}_i)$ obtained by removing the interior of a d-dimensional ball $\mathcal{B}_i \subset \mathcal{M}_i$ from $\mathcal{M}_i$. Let $\mathcal{M} = \mathcal{M}_1' \cup_f \mathcal{M}_2'$ be a d-manifold obtained from the disjoint union of $\mathcal{M}_1'$ and $\mathcal{M}_2'$ by identifying $\partial\mathcal{B}_1$ and $\partial\mathcal{B}_2$ via a homeomorphism $f\colon \partial\mathcal{B}_1 \to \partial\mathcal{B}_2$. (If $\mathcal{M}_1'$ and $\mathcal{M}_2'$ are oriented manifolds, we require that f is orientation-reserving.) We say that $\mathcal{M}$ is the *connected sum* of $\mathcal{M}_1$ and $\mathcal{M}_2$, which we denote as $\mathcal{M} = \mathcal{M}_1 \# \mathcal{M}_2$.

The connected sum of two surfaces is always uniquely defined. In dimension $d = 3$, the situation is slightly more complicated: orientable (but not oriented) 3-manifolds can have two non-homeomorphic connected sums. Nevertheless, if we consider *oriented* 3-manifolds, the construction in Definition 3.13 delivers a well-defined result. Moreover, taking the connected sum is an associative and commutative operation, with the d-sphere $\mathbb{S}^d$ being its (unique) neutral element, i.e., $\mathcal{M} \# \mathbb{S}^d = \mathcal{M}$ for any d-manifold $\mathcal{M}$. See [134, Section 1.6] and [60, Chapter 3] for details, and Figure 3.5(ii) for a basic example.

A closed, connected d-manifold $\mathcal{M} \neq \mathbb{S}^d$ is *prime* if it cannot be obtained as a non-trivial connected sum. In other words, whenever $\mathcal{M} = \mathcal{M}_1 \# \mathcal{M}_2$ then one of the summands is $\mathbb{S}^d$. The only prime surfaces are the torus $\mathbb{T}^2$ and the real projective plane $\mathbb{R}P^2$ (Figure 3.5). By Theorem 3.14, any closed surface other than $\mathbb{S}^2$ can be built from them.

Theorem 3.14 (The Classification of Surfaces)**.** *A closed, connected surface $\mathcal{S}$ is either homeomorphic to $\mathbb{S}^2$, or to a connected sum of tori, or to a connected sum of real projective planes. Moreover, in either case the number of summands is uniquely determined and is equal to the genus $g(\mathcal{S})$ of the surface $\mathcal{S}$* (cf. the last paragraph of Section 3.3)*.*

Theorem 3.14 can be viewed as a "prime decomposition theorem" for surfaces, which, together with the fact that there are only two "prime elements" ($\mathbb{T}^2$ and $\mathbb{R}P^2$), readily provides a classification. From Theorem 3.14 one can also deduce the geometrization theorem for orientable surfaces, which has been known since the 1870s.

Theorem 3.15 (The Geometrization of Surfaces)**.** *Any closed, connected and orientable surface $\mathcal{S}$ can be presented as a quotient of $\mathbb{S}^2$ (regarded as the unit sphere in $\mathbb{R}^3$ equipped with its standard round metric), the Euclidean plane $\mathbb{E}^2$, or the hyperbolic plane $\mathbb{H}^2$ by a discrete group of isometries of the respective metric spaces.*

We note that there are only two closed, orientable surfaces that are *not* hyperbolic: $\mathbb{S}^2$ itself and the torus $\mathbb{T}^2$ which is Euclidean. See Section 2 of [105] for details.

Decomposition theorems for 3-manifolds. Triangulations and Heegaard splittings represent two extremes in making up a given 3-manifold from *identical* building blocks: a triangulation consist of several "small" pieces (namely, tetrahedra), whereas a Heegaard splitting is constructed from only two but "large" chunks (namely, handlebodies).

The Geometrization Conjecture takes a different, top-down approach, in which a given 3-manifold is first decomposed into smaller parts (in two stages), after which the individual pieces (still of very high variability) are classified in terms of their geometric structure. Our exposition of the following material is inspired by Chapter 1 of [4] and by [120].

As we generally do throughout this thesis, we only consider orientable 3-manifolds.

The first stage is the same as for surfaces: a decomposition into prime manifolds.

Theorem 3.16 (Prime Decomposition Theorem; Kneser [83], Milnor [106])**.** *Every closed, connected and oriented 3-manifold $\mathcal{M}$ can be decomposed as a connected sum $\mathcal{M} = \mathcal{M}_1 \# \cdots \# \mathcal{M}_k$ of prime 3-manifolds $\mathcal{M}_i$ $(1 \leq i \leq k)$. Moreover, if $\mathcal{M} = \mathcal{N}_1 \# \cdots \# \mathcal{N}_\ell$ is another decomposition into prime 3-manifolds $\mathcal{N}_j$ $(1 \leq j \leq \ell)$, then $\ell = k$ and there exists a permutation $\pi\colon \{1, \ldots, k\} \to \{1, \ldots, k\}$, such that $\mathcal{M}_{\pi(i)} = \mathcal{N}_i$ for every $i \in \{1, \ldots, k\}$.*

Remark 3.17. The set of *decomposing spheres*, along which the connected sums are formed in Theorem 3.16, is *not* unique up to isotopy. However, two such sets of spheres are always related by so-called "slide homeomorphisms" of the ambient manifold, cf. [104, Section 3] for details, and [152, Section 2.2] for a visual description.

It is not hard to show that an orientable 3-manifold is prime if and only if it is irreducible (Section 3.3) or homeomorphic to $\mathbb{S}^2 \times \mathbb{S}^1$, see, e.g., [134, Theorem 3.3.4]. Thus, for the rest of this section, we may restrict our attention to irreducible 3-manifolds.

A 3-manifold $\mathcal{M}$ is (homotopically) *atoroidal* if, for any continuous map $f\colon \mathbb{T}^2 \to \mathcal{M}$ which induces an injective map $f_*\colon \pi_1(\mathbb{T}^2) \to \pi_1(\mathcal{M})$ between the fundamental groups, there is a homotopy $H\colon \mathbb{T}^2 \times [0,1] \to \mathcal{M}$, $H_t(x) = H(x,t)$ with $H_0 = f$ and $\mathrm{im}(H_1) \subseteq \partial\mathcal{M}$. In particular, every incompressible torus embedded in $\mathcal{M}$ is isotopic to a boundary component (and if $\partial\mathcal{M} = \emptyset$, then $\mathcal{M}$ does not contain incompressible tori at all).[10]

In the second stage, irreducible 3-manifolds are further decomposed along tori.

Theorem 3.18 (Torus Decomposition Theorem; Jaco–Shalen [72], Johannson [73])**.** *Let $\mathcal{M}$ be a closed, irreducible, orientable 3-manifold. There exists a family $\mathscr{T} = \{\mathbb{T}^2_1, \ldots, \mathbb{T}^2_s\}$ of pairwise disjoint incompressible tori embedded in $\mathcal{M}$, such that each component of the 3-manifold $\mathcal{M}'$ obtained by cutting $\mathcal{M}$ along $\mathscr{T}$ is atoroidal or Seifert fibered. Moreover, a minimal such family of tori is canonical, i.e., uniquely determined up to isotopy.*

The decomposition presented in Theorem 3.18 is called the *JSJ decomposition* of $\mathcal{M}$, and the connected components of $\mathcal{M}'$ are referred to as its *pieces*. A 3-manifold is a *graph manifold* [149] if all the pieces in its JSJ decomposition are Seifert fibered.

Example 3.19 (3-torus)**.** Consider the 3-dimensional torus, or 3-*torus*, $\mathbb{T}^3 = \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1$. Note, that the tori $\mathbb{T}^2_1 = \mathbb{S}^1 \times \mathbb{S}^1 \times \{x_0\} \subset \mathbb{T}^3$ and $\mathbb{T}^2_2 = \mathbb{S}^1 \times \{y_0\} \times \mathbb{S}^1 \subset \mathbb{T}^3$ are both incompressible and non-isotopic in $\mathbb{T}^3$. However, this does not contradict the uniqueness of JSJ decompositions, simply because $\mathbb{T}^3$ is a Seifert fibered space over the 2-torus without exceptional fibers. Hence its minimal family $\mathscr{T}$ of decomposing tori is empty.

Example 3.20 (torus bundles over $\mathbb{S}^1$)**.** Let $h\colon \mathbb{T}^2 \to \mathbb{T}^2$ be an orientation-preserving homeomorphism. The *mapping torus* $M(h)$ is a closed orientable 3-manifold obtained from $\mathbb{T}^2 \times [0,1]$ by identifying its boundary components via $(x,0) \mapsto (h(x),1)$ where $x \in \mathbb{T}^2$. Such a manifold is also called a *torus bundle over* $\mathbb{S}^1$. If $h = \mathrm{id}$, then $M(h) = \mathbb{T}^3$. However, for h "complicated enough," $M(h)$ is not Seifert fibered, but a so-called *solvmanifold*.[11]

[10] This condition is sometimes called "geometrically atoroidal" and is strictly weaker than being atoroidal in the homotopical sense, see, e.g., [97, p. 383].

[11] In particular, its JSJ decomposition cuts it along a torus fiber resulting in a single piece homeomorphic to $\mathbb{T}^2 \times [0,1]$. For the classification of torus bundles over $\mathbb{S}^1$, cf. [57, §2.2].

Geometrization of 3-manifolds. Any closed, connected and orientable surface can be obtained as a quotient of $\mathbb{S}^2$, $\mathbb{E}^2$, or $\mathbb{H}^2$ by a discrete group of isometries (Theorem 3.15). The analogous statement for 3-manifolds is generally not true. However, Thurston boldly conjectured [146] (and established for 3-manifolds which are Haken) that the pieces in the JSJ decomposition can already be "geometrized." More precisely, he postulated that the interior of each of these pieces can be obtained as a quotient of one of eight simply connected, homogeneous[12] Riemannian 3-manifolds (called *model geometries*) by a discrete group of isometries acting freely. We refer to the surveys [22, 135] and [101, Chapter 12] for details about geometrization and the eight model geometries. (See [110, Chapter 1] for a quick overview.) Below we present a "hybrid" version of the Geometrization Theorem. To that end we only need to deal with two kind of geometries.

A 3-manifold $\mathcal{M}$ is *spherical* (resp. *hyperbolic*) if it (resp. its interior) is the quotient of the standard 3-sphere $\mathbb{S}^3$ (resp. hyperbolic space $\mathbb{H}^3$) by a discrete group of isometries acting freely on $\mathbb{S}^3$ (resp. $\mathbb{H}^3$). The geometry of $\mathcal{M}$ is said to be *modeled* on $\mathbb{S}^3$ (resp. $\mathbb{H}^3$).

The following two theorems, conjectured by Thurston and proved by Perelman, show that the geometry of a compact 3-manifold is closely governed by its fundamental group.

Theorem 3.21 (Elliptization Theorem)**.** *Every closed 3-manifold with finite fundamental group is spherical.* [In particular, the Poincaré Conjecture is true.]

Theorem 3.22 (Hyperbolization Theorem)**.** *Let $\mathcal{M}$ be a compact, irreducible, orientable 3-manifold with $\partial\mathcal{M}$ being empty or the disjoint union of tori. Suppose that $\mathcal{M}$ is atoroidal and not homeomorphic to $\mathbb{S}^1 \times \mathbb{D}^2$, $\mathbb{T}^2 \times [0,1]$ or $\mathbb{K}^2\tilde{\times}[0,1]$ (the orientable I-bundle over the Klein bottle $\mathbb{K}^2$). If the fundamental group of $\mathcal{M}$ is infinite, then $\mathcal{M}$ is hyperbolic.*

It has been known since the 1930s that every spherical 3-manifold is Seifert fibered [145, §7, Hauptsatz] (cf. [114, Chapter 6, Theorem 5]), as well as the 3-manifolds $\mathbb{S}^1 \times \mathbb{D}^2$, $\mathbb{T}^2 \times [0,1]$ and $\mathbb{K}^2\tilde{\times}[0,1]$. These facts together with Theorems 3.18, 3.21 and 3.22 imply

Theorem 3.23 (Geometrization Theorem)**.** *Let $\mathcal{M}$ be a closed, irreducible, orientable 3-manifold. There exists a family $\mathscr{T} = \{\mathbb{T}^2_1, \ldots, \mathbb{T}^2_s\}$ of pairwise disjoint incompressible tori embedded in $\mathcal{M}$ such that each component of the 3-manifold $\mathcal{M}'$ obtained by cutting $\mathcal{M}$ along $\mathscr{T}$ is either hyperbolic or Seifert fibered. Moreover, a minimal such family of tori is canonical, i.e., uniquely determined up to isotopy.*

[12] Here 'homogeneous' means that the isometry group of the respective 3-manifold is acting transitively.

Remark 3.24. As mentioned earlier, Theorem 3.23 may be regarded as a "hybrid" version of the Geometrization Theorem, because some pieces of the decomposition are described in geometric (viz., 'hyperbolic') and others in non-geometric (viz., 'Seifert fibered') terms. With a slight adjustment, however, one can recover from it the "fully geometric" version. In order to do that, one has to geometrize the Seifert fibered pieces as well. For most cases, this can readily be done using the aforementioned spherical ($\mathbb{S}^3$) geometry or one of five other geometries: $\mathbb{E}^3$, $\mathbb{S}^2 \times \mathbb{R}$, $\mathbb{H}^2 \times \mathbb{R}$, $\widetilde{\mathrm{SL}_2(\mathbb{R})}$, and Nil.

Some subtleties have to be taken care of, because the "fully geometric" version of the Geometrization Theorem in fact requires the pieces to have *finite* volume. For example, the interior of $\mathbb{T}^2 \times [0,1]$ is modeled on the Euclidean 3-space $\mathbb{E}^3$ [146, p. 359], it is, however, not of finite volume. At the same time, if some piece $\mathcal{P}'$ in the JSJ decomposition of $\mathcal{M}$ is homeomorphic to $\mathbb{T}^2 \times [0,1]$, then, due to the minimality of $\mathscr{T}$, this piece has to stem from a torus bundle $\mathcal{P}$ over $\mathbb{S}^1$, which is recovered from $\mathcal{P}' = \mathbb{T}^2 \times [0,1]$ by gluing its two boundary components back together via a homeomorphism. Now $\mathcal{P}$ can be endowed with a finite-volume geometry modeled on $\mathbb{E}^3$, Nil, or Sol, depending on the gluing homeomorphism. (Due to the connectedness of $\mathcal{M}$, we actually have $\mathcal{M} = \mathcal{P}$ in this case.) For similar reasons, we also have to allow decompositions along Klein bottles. This is sufficient to state the "fully geometric" version of the Geometrization Theorem:

Every closed, irreducible, orientable 3-manifold $\mathcal{M}$ admits a canonical decomposition[13] *along tori or Klein bottles, such that the resulting pieces can be geometrized using the following eight model geometries:* $\mathbb{E}^3$, $\mathbb{S}^3$, $\mathbb{S}^2 \times \mathbb{R}$, $\mathbb{H}^2 \times \mathbb{R}$, $\widetilde{\mathrm{SL}_2(\mathbb{R})}$, Nil, Sol, *and* $\mathbb{H}^3$.

Again, we refer to [22], [135], [101, Chapter 12] and [110, Chapter 1] for details.

Towards the classification of 3-manifolds. By Theorem 3.14, closed surfaces are determined by their genus and their orientability. Moreover, both of these properties can easily be computed from a triangulation.[14] As the previous paragraphs suggest, characterizing 3-manifolds is a much more complicated task, and one should not expect to tell them apart by means of a "simple" computable invariant. Nevertheless, there are various ways to approach their taxonomy, each of which has been quite productive.

[13] While this canonical decomposition is similar to the JSJ decomposition (and sometimes it coincides with it), the above examples show that they are generally not the same. Cf. [101, Section 11.5.3].

[14] Recall that the Euler characteristic $\chi(\mathcal{S})$ of a closed surface $\mathcal{S}$ equals $\chi(\mathcal{S}) = 2 - 2g(\mathcal{S})$ if $\mathcal{S}$ is orientable, and $\chi(\mathcal{S}) = 1 - g(\mathcal{S})$ if $\mathcal{S}$ is non-orientable.

In view of Theorems 3.16 and 3.23, the classification of 3-manifolds boils down to understanding Seifert fibered manifolds and those that are hyperbolic. There is a striking contrast between the two classes. While Seifert fibered spaces were already classified by Seifert [137], the family of hyperbolic 3-manifolds remains elusive.[15] In addition, most 3-manifolds are hyperbolic (cf. Theorem 2.6 in [146] and the discussion thereafter). Thus, in the past forty years tremendous effort has been dedicated to research on hyperbolic 3-manifolds, cf. [4, Chapter 4], [10] and [97], and the references therein.

A different take on the classification problem aims at building a catalogue, or census, of 3-manifolds listed by increasing "complexity." A natural measure of complexity in this context is the *triangulation complexity*, i.e., the minimum number of tetrahedra required to triangulate a given 3-manifold (Section 3.1). Such a census may be created two steps:

1. generate all triangulations up to a given size, and
2. eliminate the duplicates (up to homeomorphism).

The real challenge lies in the second step, which calls for a solution to the Homeomorphism Problem, a fundamental decision problem in 3-dimensional topology:

Given two triangulations, decide whether they encode homeomorphic 3-manifolds.

Recently it was rigorously shown, that the Geometrization Theorem (in synthesis with other results) yields algorithmic solutions to the Homeomorphism Problem [87, 136]. At present, however, the practical impact of these algorithms is limited (viz., they have not been implemented), and their computational complexity is far from being understood, see Section 3.1 in [90] for the current state of the art.

Therefore, in practice, the Homeomorphism Problem is approached via topological invariants. Computational challenges are still present: algorithms for evaluating invariants that are fine enough to tell large classes of 3-manifolds apart may need exponential time to terminate in the worst case (measured in terms of the number of tetrahedra in the triangulation). At the same time, as mentioned in Chapter 1, there exist powerful invariants (e.g., the Turaev–Viro invariants), which have been successfully employed in creating large censuses of triangulated 3-manifolds, largely due to being fixed-parameter tractable in the treewidth of the input triangulation [32].

[15] Hyperbolic 3-manifolds with finitely generated fundamental groups (regardless the finiteness of their volume) have been classified in terms of the Tameness [2, 34] and the Ending Lamination [23] Theorems.

4 Interfaces between Combinatorics and Topology

3-Manifolds have the pleasant feature that they can be equipped with different structures providing various viewpoints of study. In this chapter we discuss interfaces that will allow us to navigate between the combinatorial and the topological viewpoints in Chapter 5.

4.1 Generalized Heegaard Splittings

The notion of a Heegaard splitting, where a 3-manifold is built by gluing two handlebodies together, was refined by Scharlemann and Thompson in a seminal paper [131] motivated by the notion of *thin position* in knot theory [51].

In a *generalized Heegaard splitting* a 3-manifold is constructed from several pairs of compression bodies. This construction arises naturally from more complicated sequences of handle attachments, e.g., when a 3-manifold can be assembled by first attaching only some of the 0- and 1-handles before attaching any 2- and 3-handles.

4.1.1 Linear Splittings

A *linear generalized Heegaard splitting* (or *linear splitting*, for short) of a 3-manifold $\mathcal{M}$ is a decomposition

$$\mathcal{M} = (\mathcal{N}_1 \cup_{\mathcal{S}_1} \mathcal{K}_1) \cup_{\mathcal{R}_1} \cdots \cup_{\mathcal{R}_{r-1}} (\mathcal{N}_r \cup_{\mathcal{S}_r} \mathcal{K}_r) \tag{4.1}$$

where $(\mathcal{N}_1, \mathcal{K}_1, \ldots, \mathcal{N}_r, \mathcal{K}_r)$ is a sequence of possibly disconnected compression bodies in $\mathcal{M}$. They are pairwise disjoint except for subsequent pairs, which are "glued together" along (pairwise disjoint) closed surfaces $\mathcal{S}_1, \mathcal{R}_1, \ldots, \mathcal{S}_{r-1}, \mathcal{R}_{r-1}, \mathcal{S}_r$. More precisely,

$$\mathcal{S}_i = \partial_+\mathcal{N}_i = \partial_+\mathcal{K}_i \ \ (1 \leq i \leq r) \quad \text{and} \quad \mathcal{R}_j = \partial_-\mathcal{K}_j = \partial_-\mathcal{N}_{j+1} \ \ (1 \leq j \leq r-1).$$

The lower boundaries of $\mathcal{N}_1$ and $\mathcal{K}_r$ constitute the boundary of $\mathcal{M}$, i.e.,

$$\partial\mathcal{M} = \partial_-\mathcal{N}_1 \cup \partial_-\mathcal{K}_r.$$

In the light of Example 3.10, a linear splitting is a decomposition $\mathcal{M} = \mathcal{M}_1 \cup \ldots \cup \mathcal{M}_r$ into 3-manifolds with boundary, such that $\mathcal{M}_i \cap \mathcal{M}_j \neq \emptyset \Rightarrow |i-j| \leq 1$, together with a choice of a Heegaard splitting $\mathcal{M}_i = \mathcal{N}_i \cup_{\mathcal{S}_i} \mathcal{K}_i$ for each constituent compatible with the partition $\partial\mathcal{M}_i = \partial_-\mathcal{N}_i \cup \partial_-\mathcal{K}_i$ of the boundary components of $\mathcal{M}_i$. See Figure 4.1.

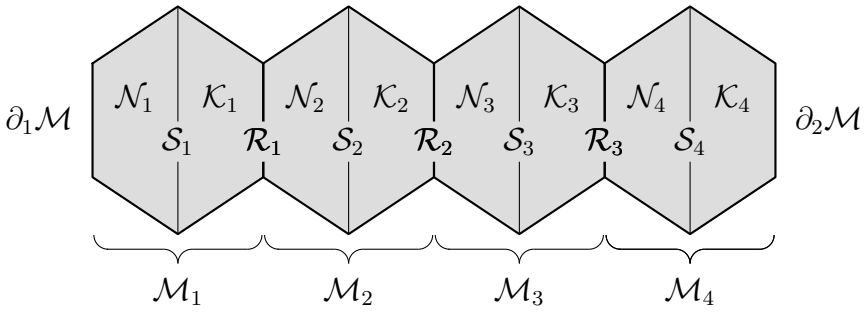

Figure 4.1: Schematic for a linear generalized Heegaard splitting of a 3-manifold $\mathcal{M}$. The picture indicates the fact that $g(\mathcal{S}_i) \geq \max\{g(\mathcal{R}_{i-1}), g(\mathcal{R}_{i+1})\}$. However, it does not contain information about the connectedness of these surfaces or the manifolds $\mathcal{M}_i$

Example 4.1 (Linear splittings from handle decompositions)**.** Assume $\mathcal{M}$ is a closed 3-manifold given via a handle decomposition, i.e., a sequence of handle attachments to build up $\mathcal{M}$. Consider the first terms of the sequence up to (but not including) the first 2- or 3-handle attachment. Let $\mathcal{N}_1$ be the union of all handles in this subsequence. In the second step look at all 2- and 3-handles following the initial sequence of 0- and 1-handles until we reach 0- or 1-handles again, and follow the dual construction to obtain another compression body $\mathcal{K}_1$. More precisely, define $\partial_+\mathcal{K}_1 = \partial_+\mathcal{N}_1$, thicken the top boundary into $\partial_+\mathcal{K}_1 \times [0,1]$, and then attach the given 2- and 3-handles along $\partial_+\mathcal{K}_1 \times \{0\}$. Iterating this procedure results in a linear splitting of $\mathcal{M}$ into compression bodies $(\mathcal{N}_1, \mathcal{K}_1, \ldots, \mathcal{N}_r, \mathcal{K}_r)$.

4.1.2 Graph Splittings

Similarly to linear splittings, a *graph generalized Heegaard splitting* (or *graph splitting*) of a 3-manifold $\mathcal{M}$ is a decomposition

$$\mathscr{D} = \{\mathcal{M}_i : i \in I,\ \textstyle\bigcup_{i\in I} \mathcal{M}_i = \mathcal{M},\ \text{and}\ \operatorname{int}(\mathcal{M}_i) \cap \operatorname{int}(\mathcal{M}_j) = \emptyset \text{ for } i \neq j\} \tag{4.2}$$

into finitely many 3-manifolds with pairwise disjoint interiors that intersect along surfaces, together with an appropriate Heegaard splitting for each $\mathcal{M}_i$. However, here the $\mathcal{M}_i$ are not necessarily glued together in linear order, but along an arbitrary graph.

More precisely, given a decomposition $\mathscr{D}$ as above, consider its *dual graph*,[1] which is a multigraph $\Gamma(\mathscr{D}) = (I, E)$ with nodes corresponding to the $\mathcal{M}_i$ and arcs between i and j to the connected components of $\mathcal{M}_i \cap \mathcal{M}_j$, cf. Figure 4.2. Pick an ordering of I, i.e., a bijection $\ell\colon I \to \{1, \dots, |I|\}$. For any $i \in I$, let $\partial_1\mathcal{M}_i \cup \partial_2\mathcal{M}_i$ be a partition of the connected components of $\partial\mathcal{M}_i$, so that $\partial_1\mathcal{M}_i$ (resp. $\partial_2\mathcal{M}_i$) contains the components glued to those of any $\mathcal{M}_j$ with $\ell(j) < \ell(i)$ (resp. $\ell(j) > \ell(i)$). The "unpaired" components of $\partial\mathcal{M}_i$, i.e., those which contribute to the boundary $\partial\mathcal{M}$ of the ambient manifold $\mathcal{M}$, are partitioned among $\partial_1\mathcal{M}_i$ and $\partial_2\mathcal{M}_i$ arbitrarily. For each $i \in I$, choose a Heegaard splitting $\mathcal{M}_i = \mathcal{N}_i \cup_{\mathcal{S}_i} \mathcal{K}_i$ of $\mathcal{M}_i$ compatible with the partition $\partial_1\mathcal{M}_i \cup \partial_2\mathcal{M}_i$ of the boundary components (Example 3.10). This way we obtain a *graph splitting* of $\mathcal{M}$ (Figure 4.3).

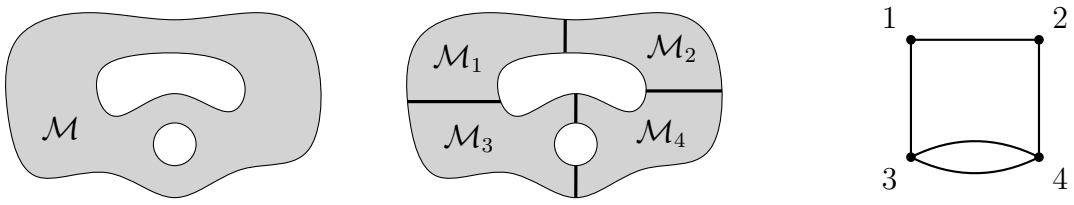

Figure 4.2: (i) Schematic of a closed 3-manifold $\mathcal{M}$ with nontrivial first homology, (ii) a decomposition $\mathscr{D}$ of $\mathcal{M}$ into four submanifolds, and (iii) the dual graph $\Gamma(\mathscr{D})$ of $\mathscr{D}$

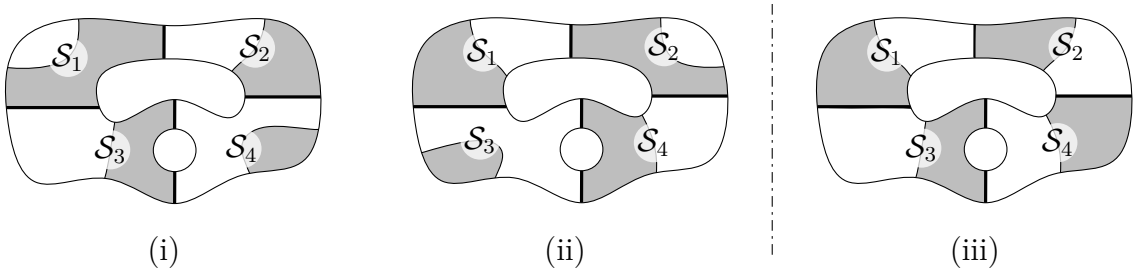

Figure 4.3: (i)–(ii) Two graph splittings of $\mathcal{M}$ stemming from the decomposition shown on Figure 4.2, which respectively correspond to the orderings $\ell_1(i) = i$ ($i \in I = \{1, 2, 3, 4\}$), and $\ell_2(1) = 2$, $\ell_2(2) = 4$, $\ell_2(3) = 1$, $\ell_2(4) = 3$. (iii) A non-example, cf. Remark 4.2(1). The pieces of the splitting are colored with gray and white for increased visibility

[1]Not to be confused with the dual graph of a triangulation.

In some situations we merely need to talk about the constituents of a decomposition or a graph splitting thereof without having to deal with the exact way the parts are glued together. Then we will use the shorthand notation

$$\mathcal{M} = \bigcup_{i\in I} \mathcal{M}_i \quad \text{and} \quad \mathcal{M} = \bigcup_{i\in I} (\mathcal{N}_i \cup_{\mathcal{S}_i} \mathcal{K}_i), \quad \text{where} \quad \mathcal{N}_i \cup_{\mathcal{S}_i} \mathcal{K}_i = \mathcal{M}_i. \tag{4.3}$$

In other situations, e.g., when connectivity properties of the graph $\Gamma(\mathscr{D})$ underlying a given splitting are relevant, we will also work with so-called fork complexes. Here we give a brief overview of this language. For more details, see [130, Chapter 5].

A *fork complex* is essentially a decorated version of $\Gamma(\mathscr{D})$. It is a labeled graph in which the compression bodies of a given decomposition are modeled by *fork*s. More precisely, an *n-fork* is a tree F with $n+2$ nodes $V(F) = \{g, p, t_1, \dots, t_n\}$ with p being of degree $n+1$ and all other nodes being leaves. The nodes g, p, and the t_i are called the *grip*, the *root*, and the *tine*s of F, respectively (Figure 4.4(i) shows a 0- and a 3-fork). We think of a fork $F = F_{\mathcal{C}}$ as an abstraction of a compression body $\mathcal{C}$, such that the grip of F corresponds to $\partial_+\mathcal{C}$, whereas the tines correspond to the connected components of $\partial_-\mathcal{C}$.

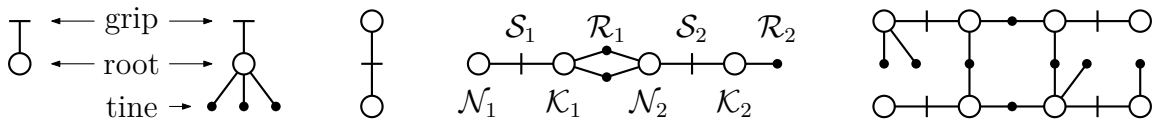

Figure 4.4: Fork complexes of (ii) Heegaard, (iii) linear, and (iv) graph splittings. Note that the labels on (iii) near the roots and tines refer to the underlying compression bodies and surfaces, respectively, and not to the roots and tines themselves

Informally, a *fork complex* $\boldsymbol{F}$ (representing a given graph splitting of a 3-manifold $\mathcal{M}$) is obtained by taking several forks (corresponding to the compression bodies which constitute $\mathcal{M}$), and identifying grips with grips, and tines with tines (following the way the boundaries of these compression bodies are glued together). The set of grips and tines which remain unpaired is denoted by $\partial\boldsymbol{F}$ (as they correspond to surfaces which constitute the boundary $\partial\mathcal{M}$). The formal relationship between $\boldsymbol{F}$ and the underlying 3-manifold $\mathcal{M}$ (possibly with boundary) is described by a map $\rho : (\mathcal{M}; \partial\mathcal{M}) \to (\boldsymbol{F}; \partial\boldsymbol{F})$ which has to satisfy certain natural criteria [130, Definition 5.1.7]. The pair $(\boldsymbol{F}, \rho)$ is called a *graph splitting*. See Figure 4.4 for illustrations, and [130, Section 5.1] for further details.

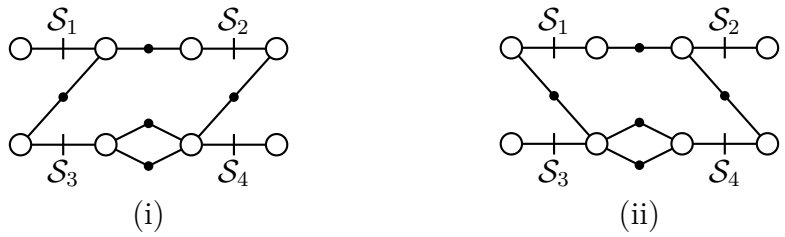

Figure 4.5: Fork complexes representing the graph splittings shown on Figure 4.3.

Remark 4.2. We conclude the section with some remarks.

(1) Ordering the vertices of $\Gamma(\mathscr{D})$ and choosing the Heegaard splittings of the $\mathcal{M}_i$ in a compatible way might seem to be an ad-hoc requirement in the construction of graph splittings. It becomes more natural if we consider graph splittings to be refinements of linear splittings. The above requirement ensures that any graph splitting can be *linearized*, i.e., turned into a linear splitting. See, e.g, Figure 4.6. Note that the fake graph splitting in Figure 4.3(iii) cannot be linearized.

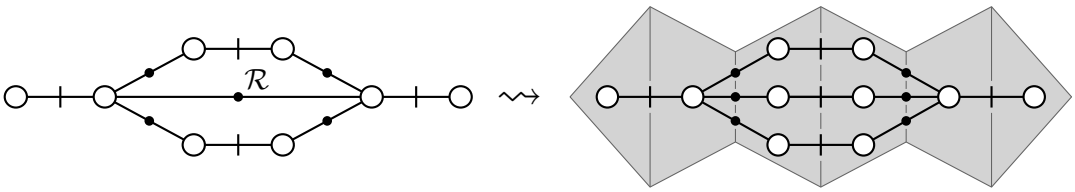

Figure 4.6: Turning a graph splitting into a linear one. The decomposition might need to be (trivially) refined at some of the lower boundaries, e.g., at $\mathcal{R}$

(2) Generalized Heegaard splittings are described in various ways in the literature. Our treatment of graph splittings is inspired by [5] and [130]. The terms "linear splitting" and "graph splitting" were adopted from [62, 63].

(3) Along with the notion of *strong irreducibility*, pioneered by Casson and Gordon [35], generalized Heegaard splittings have been central to the study of classical Heegaard splittings of 3-manifolds, cf. Section 5 of the survey [64].

(4) Any generalized Heegaard splitting can be turned into a classical one via a procedure called *amalgamation* [133]. (Here we also rely on the requirement mentioned in (1) above, cf. [5, Example 2.6].) Amalgamations lie at the heart of many applications, including the main result of [5] according to which the problem of computing the Heegaard genus is **NP**-hard. We also make use of amalgamations in Section 5.4.

4.2 Layered Triangulations

In this section we provide a brief and selective introduction to the theory of *layered triangulations* by Jaco and Rubinstein [71], which captures the inherently topological notion of a Heegaard splitting (see Section 3.5) in a combinatorial way. We focus on the terminology important for our purposes. Despite all the technicalities, the nomenclature is very expressive and encapsulates much of the intuition.

4.2.1 Spines and Layered Triangulations of Handlebodies

Let $\mathcal{N}_{g,r}$ denote the non-orientable surface of genus g with r *punctures*.[2] A *g-spine* is a 1-vertex triangulation of $\mathcal{N}_{g,1}$. It has one vertex, $3g-1$ edges (out of which $3g-2$ are interior and one is on the boundary), and $2g-1$ triangles. In particular, the Euler characteristic of any g-spine equals $1-g$.

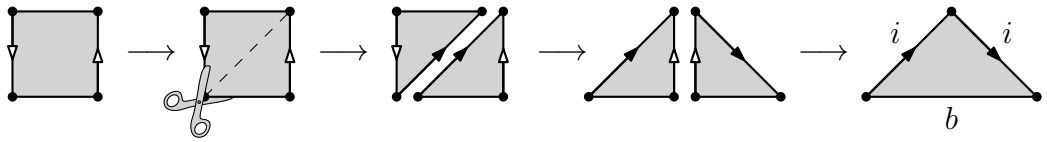

Figure 4.7: Left to right: Transforming the well-known depiction of the Möbius band (the non-orientable surface of genus one with one puncture) into a 1-spine with interior edge i

Now consider a triangulation $\mathcal{S}$ of a surface—usually seen as a g-spine or as the boundary of a triangulated 3-manifold—and let e be an interior edge of $\mathcal{S}$ with t_1 and t_2 being the two triangles of $\mathcal{S}$ containing e. Gluing a tetrahedron Δ along t_1 and t_2 without a twist is called a *layering* onto the edge e of the surface $\mathcal{S}$, cf. Figure 4.8(i). Importantly, we allow t_1 and t_2 to coincide, e.g., when layering on the interior edge of a 1-spine (Figure 4.7, right).

It is a pleasant fact that by layering a tetrahedron onto each of the $3g-2$ interior edges of a g-spine we obtain a triangulation of the genus g handlebody $\mathcal{H}_g$, called a *minimal layered triangulation* thereof (see Figure 4.9). Moreover, we call any triangulation obtained by additional layerings a *layered triangulation* of $\mathcal{H}_g$.

[2]For example, $\mathcal{N}_{1,0} = \mathbb{R}P^2$, $\mathcal{N}_{1,1}$ = Möbius band, $\mathcal{N}_{2,0}$ = Klein bottle.

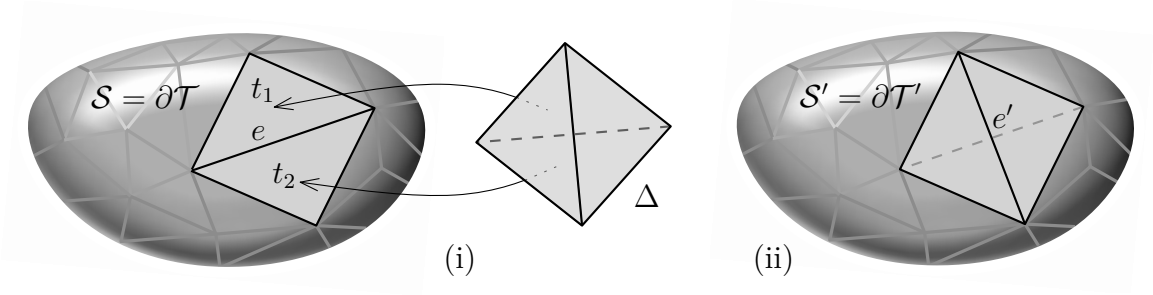

Figure 4.8: (i) Layering onto the edge e of $\mathcal{S} = \partial\mathcal{T}$, (ii) has the effect of flipping e into e'

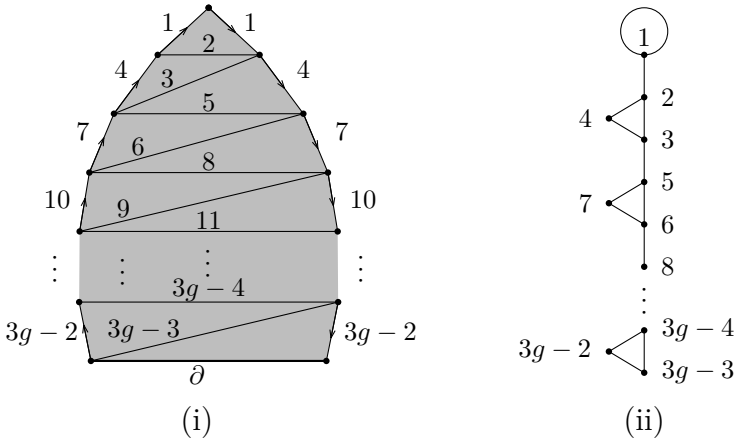

Figure 4.9: (i) A g-spine together with the order in which we layer onto its interior edges. (ii) The dual graph of the resulting minimal layered triangulation of $\mathcal{H}_g$

4.2.2 Layered Triangulations of Heegaard Splittings

Let $\mathcal{M} = \mathcal{H} \cup_f \mathcal{H}'$ be a closed, orientable 3-manifold given by a Heegaard splitting. If $\mathcal{H}$ and $\mathcal{H}'$ can be endowed with layered triangulations $\mathcal{T}$ and $\mathcal{T}'$, respectively, such that the attaching map f is a simplicial isomorphism (i.e., respects the triangulations of $\partial\mathcal{T}$ and $\partial\mathcal{T}'$), then $\mathcal{T} \cup_f \mathcal{T}'$ triangulates $\mathcal{M}$ and is called a *layered triangulation* of $\mathcal{M}$. The next theorem is fundamental and asserts that this is always possible.

Theorem 4.3 (Jaco–Rubinstein, Theorem 10.1 of [71])**.** *Any closed, orientable* 3*-manifold admits a layered triangulation (which is a one-vertex triangulation by construction).*

4.2.3 Layered Solid Tori

The case $g = 1$ is of particular importance. Starting with a 1-spine (Figure 4.7) and layering on its interior edge i produces a 1-tetrahedron triangulation $\mathcal{T}$ of the solid torus $\mathcal{H}_1$ (Appendix C). Its boundary $\mathcal{S} = \partial\mathcal{T}$ is the unique 2-triangle triangulation of the

torus with one vertex and three edges, and layering onto any edge of $\mathcal{S}$ yields another triangulation of $\mathcal{H}_1$. We may iterate this procedure to obtain further triangulations, any of which we call a *layered solid torus*, cf. [24, Section 1.2] for a detailed exposition. By construction, its dual graph consists of a single loop at the end of a path of double arcs; see, e.g, Figure 4.10(v).

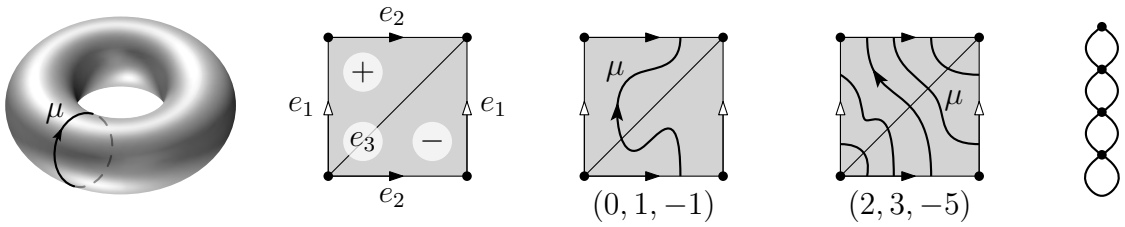

(i) (ii) (iii) (iv) (v)

Figure 4.10: (i) An oriented meridian μ on the torus $\mathbb{T}^2$. (ii) The minimal triangulation of $\mathbb{T}^2$ together with a labeling of its edges and triangles. (iii)–(iv) Two non-isotopic triangulations of $\mathbb{T}^2$ with the relative positions of the meridian μ. (v) The dual graph of a layered solid torus with four tetrahedra

While combinatorially the same, boundary triangulations of layered solid tori generally are not isotopic; they can be described as follows. Consider a "reference torus" $\mathbb{T}$ with a fixed meridian μ, Figure 4.10(i). Given a layered solid torus, its boundary induces a triangulation of $\mathbb{T}$. Label the two triangles with $+$ and $-$, and the three edges with e_1, e_2, and e_3, Figure 4.10(ii); and fix an orientation of μ. Let p, q and r denote the geometric intersection number of μ with e_1, e_2 and e_3, respectively. We say that the corresponding layered solid torus is of type (p, q, r), or $\mathrm{LST}(p, q, r)$ for short, cf. Figure 4.10(iii)–(iv). It can be shown that p, q, r are always coprime with $p + q + r = 0$. Conversely, for any such triplet, one can construct a layered solid torus of type (p, q, r), cf. [24, Algorithm 1.2.17].

5 From Combinatorics to Topology and Back – In a Quantitative Way

In this chapter we forge quantitative relationships between several numerical invariants, which, in one way or another, capture the "width" of a compact 3-manifold.

First we give an overview of selected width parameters for 3-manifolds (Section 5.1). In particular, we discuss a general scheme of turning graph parameters into topological invariants of 3-manifolds, which we call *combinatorial width parameters*. Then we review *topological width parameters* associated with classical and generalized Heegaard splittings. We also briefly consider *geometric width parameters* defined for hyperbolic 3-manifolds.

The next two sections, based on joint work with Spreer and Wagner [67], and Spreer [66], respectively, showcase some of the main results of this thesis. In Section 5.2 we show that combinatorial width parameters dominate their topological counterparts, whereas in Section 5.3 we explore possibilities of reversing these inequalities.

In Section 5.4 we restrict our focus to hyperbolic 3-manifolds. By making use of generalized Heegaard splittings, we improve upon a result of Maria and Purcell [99]; we show that the volume of a hyperbolic 3-manifold gives an upper bound on its pathwidth.

5.1 Selected Width Parameters for 3-Manifolds

5.1.1 Combinatorial Width Parameters

Invariants of "topological objects" are often defined via "combinatorial models" thereof. Well-known examples include the *Euler characteristic* of a topological space presented as a simplicial complex; *knot polynomials*, when defined through planar diagrams; or the *Turaev–Viro invariants* of 3-manifolds, computed from triangulations. Since a given

topological object generally admits infinitely many distinct "models," these definitions always come with a respective theorem, which states that the outcome of the computation does not depend on the chosen model, but only on the underlying topological object. An advantage of this perspective is, that it offers an explicit way of computing these invariants.

In this section we follow another common way of obtaining topological invariants of combinatorial flavor: consider a numerical parameter and minimize its value over all models that represent a given topological object. Basic examples from low-dimensional topology entail the *crossing number* of a knot K, i.e., the minimum number of crossings of any diagram of K, or the *triangulation complexity* (Section 3.1) of a 3-manifold. In contrary to the examples of the previous paragraph, such quantities are readily seen to be topological invariants. At the same time, their definition does not offer a direct pathway to compute them. Indeed, given a "complicated" triangulation $\mathcal{T}$ of a 3-manifold $\mathcal{M}$, it is generally a challenging question to decide whether, and up to what extent, $\mathcal{T}$ can be reduced to a smaller triangulation $\mathcal{T}'$ of the same 3-manifold, even for $\mathcal{M} = \mathbb{S}^3$ [26]. Therefore, these invariants are often understood in terms of their quantitative relationship with other (preferably computable) invariants.

Now recall the graph parameters discussed in Section 2.1.1.

Treewidth & Co. of a 3-manifold. Let $\mathcal{T}$ be a triangulation of a compact 3-manifold $\mathcal{M}$. By the *treewidth* of $\mathcal{T}$ we mean $\mathrm{tw}(\Gamma(\mathcal{T}))$, i.e., the treewidth of its dual graph. The *treewidth* $\mathrm{tw}(\mathcal{M})$ of $\mathcal{M}$ is the smallest treewidth of any triangulation of $\mathcal{M}$, i.e.,

$$\mathrm{tw}(\mathcal{M}) = \min\{\mathrm{tw}(\Gamma(\mathcal{T})) : \mathcal{T} \text{ is a triangulation of } \mathcal{M}\}. \tag{5.1}$$

The definitions of *pathwidth* $\mathrm{pw}(\mathcal{M})$, *cutwidth* $\mathrm{cw}(\mathcal{M})$, and *congestion* $\mathrm{cng}(\mathcal{M})$ are analogous. We collectively refer to these quantities as *combinatorial width parameters*.

Since in the dual graph of any 3-manifold triangulation each node has degree at most four, Theorem 2.4 readily implies that

$$\mathrm{tw}(\mathcal{M}) \leq \mathrm{pw}(\mathcal{M}) \leq \mathrm{cw}(\mathcal{M}) \leq 4\,\mathrm{pw}(\mathcal{M}) \tag{5.2}$$

for any 3-manifold $\mathcal{M}$. Similarly, from Theorem 2.6 we immediately deduce

$$\max\left\{\tfrac{2}{3}(\mathrm{tw}(\mathcal{M})+1), 4\right\} \leq \mathrm{cng}(\mathcal{M}) \leq 4(\mathrm{tw}(\mathcal{M})+1). \tag{5.3}$$

We note that there are simple ways to prove that any 3-manifold admits triangulations of arbitrarily high treewidth, hence pathwidth, cutwidth and congestion (Appendix B.2).

5.1.2 Topological Width Parameters

The Heegaard genus (defined in Section 3.5) is a first example for a topological width parameter of a 3-manifold: the larger the Heegaard genus, the "wider" the manifold. Here we consider two subsequent refinements, the linear width and the graph width, whose properties are essential for proving our results in Sections 5.2.1 and 5.2.2, respectively.

Linear width. In [131] Scharlemann and Thompson extend the concept of thin position from knot theory [51] to 3-manifolds and define the linear width of a manifold.[1] For this they look at linear splittings, i.e., linear decompositions of a 3-manifold into compression bodies. This setup is explained in Section 4.1.1.

Given a linear splitting of a 3-manifold $\mathcal{M}$ into $2r$ compression bodies with upper boundary surfaces $\{\mathcal{S}_i : 1 \leq i \leq r\}$, consider the multiset $\{g(\mathcal{S}_i) : 1 \leq i \leq r\}$ of the genera of the surfaces $\mathcal{S}_i$. This multiset $\{g(\mathcal{S}_i) : 1 \leq i \leq r\}$, when arranged in non-increasing order, is called the *width* of the (linear) splitting.[2] We here define the *linear width of a manifold* $\mathcal{M}$, denoted by $\mathscr{L}(\mathcal{M})$, to be the maximum entry in a lexicographically smallest width ranging over all linear splittings of $\mathcal{M}$.[3] A manifold $\mathcal{M}$ together with a linear spitting of lexicographically smallest width is said to be in *thin position*.

A guiding idea behind thin position is to attach 2-handles as early as possible and 1-handles as late as possible in order to obtain a decomposition for which the "topological complexity" of the upper boundary surfaces is minimized.

One of the key properties of thin position is expressed by the following result.

Theorem 5.1 (Scharlemann–Thompson, Rule 5 from [131]; cf. Casson–Gordon [35])**.** *Let $\mathcal{M}$ be a 3-manifold together with a linear splitting $(\mathcal{N}_1, \mathcal{K}_1, \ldots, \mathcal{N}_r, \mathcal{K}_r)$ in thin position, and $\{\mathcal{R}_i \subset \mathcal{M} : 1 \leq i \leq r-1\}$ be the set of lower boundary surfaces $\mathcal{R}_i = \partial_- \mathcal{K}_i = \partial_- \mathcal{N}_{i+1}$. Then every connected component of every surface $\mathcal{R}_i$ is incompressible.*

---

[1] Also see [64] and the textbook [130] for an introduction to generalized Heegaard splittings and to thin position, and for a survey of recent results.

[2] In [131] the width is defined via the multiset $\{c(\mathcal{S}_i) : 1 \leq i \leq r\}$, where $c(\mathcal{S}) = \max\{0, 2g(\mathcal{S}) - 1\}$. For us it is more natural to work with the genus $g(\mathcal{S})$, which is also in agreement with the way this notion is discussed in [130].

[3] For our purposes it is most convenient to define the linear width to be a single integer rather than a multiset of integers. We thus deviate also at this point from the definition of linear width in [131].

Corollary 5.2 (Scharlemann–Thompson [131]). *Let $\mathcal{M}$ be an irreducible, non-Haken 3-manifold. Then the smallest width linear splitting of $\mathcal{M}$ into compression bodies is a Heegaard splitting of minimal genus $\mathfrak{g}(\mathcal{M})$. In particular, we have $\mathscr{L}(\mathcal{M}) = \mathfrak{g}(\mathcal{M})$.*

Proof (sketch). Let $(\mathcal{N}_1, \mathcal{K}_1, \mathcal{N}_2, \mathcal{K}_2, \ldots, \mathcal{N}_r, \mathcal{K}_r)$ be a thin decomposition of $\mathcal{M}$. By Theorem 5.1, all surface components of all bounding surfaces $\mathcal{R}_i$, $1 \leq i < r$, must be incompressible. Moreover, all $\mathcal{R}_i$ are separating and thus they are 2-sided. However, irreducible, non-Haken 3-manifolds do not have incompressible 2-sided surfaces. Hence $r = 1$ and therefore the decomposition $(\mathcal{N}_1, \mathcal{K}_1)$ must be a Heegaard splitting of $\mathcal{M}$. □

Graph width. In [130] Scharlemann, Schultens and Saito further refine the concept of thin position to graph splittings of 3-manifolds, see Section 4.1.2. In particular, given a manifold $\mathcal{M}$ together with a graph splitting defined by a fork complex $\boldsymbol{F}$, let $\{\mathcal{S}_j : j \text{ grip of } \boldsymbol{F}\}$ be the set of upper boundary surfaces of the decomposition. Then the *width* of the graph splitting coming from $\boldsymbol{F}$ is defined as the multiset $\{g(\mathcal{S}_j) : j \text{ grip of } \boldsymbol{F}\}$ with non-increasing order. Similar to the construction of linear width, the *graph width* $\mathscr{G}(\mathcal{M})$ of $\mathcal{M}$ is defined to be the largest entry of the lexicographically smallest width ranging over all graph splittings of $\mathcal{M}$. A graph splitting of $\mathcal{M}$ which has lexicographically smallest width is said to be *thin*.

Theorem 5.3 (Scharlemann–Schultens–Saito, [130, Lemma 5.2.4]). *Leta $\boldsymbol{F}$ be a fork complex that determines a thin graph splitting of a 3-manifold $\mathcal{M}$, and let $\{\mathcal{R}_i \subset \mathcal{M} : i \text{ tine of } \boldsymbol{F}\}$ be the set of lower boundary surfaces as defined in Section 4.1.2. Then every connected component of every lower boundary surface $\mathcal{R}_i$ is incompressible.*

Similarly to the linear width case, Theorem 5.3 implies that a thin graph splitting of an irreducible, non-Haken 3-manifold must be a Heegaard splitting. In particular, $\mathscr{G}(\mathcal{M}) = \mathscr{L}(\mathcal{M}) = \mathfrak{g}(\mathcal{M})$ for any given irreducible, non-Haken 3-manifold $\mathcal{M}$.

Non-Haken 3-manifolds of large genus. The next theorem provides an infinite family of 3-manifolds for which we can apply our results established in the subsequent sections.

Theorem 5.4 (Agol, Theorem 3.2 in [1]). *There exist orientable, closed, irreducible, and non-Haken 3-manifolds of arbitrarily large Heegaard genus.*

Remark 5.5. The construction used to prove Theorem 5.4 starts with n-component link complements (having Heegaard genus at least $n/2$, cf. Example 3.10), and performs Dehn fillings which neither create incompressible surfaces, nor decrease the (unbounded) Heegaard genera of the complements. The existence of such Dehn fillings is guaranteed by work due to Hatcher [58] and Moriah–Rubinstein [111]. As can be deduced from the construction, the manifolds in question are closed and orientable and non-Haken.

5.1.3 Geometric Width Parameters for Hyperbolic 3-Manifolds

Recall that a 3-manifold $\mathcal{M}$ is hyperbolic, if its interior can be obtained as a quotient of the hyperbolic 3-space $\mathbb{H}^3$ by a discrete group of isometries acting freely on $\mathbb{H}^3$. This is equivalent to saying that the interior of $\mathcal{M}$ admits a complete Riemannian metric of constant sectional curvature -1. Throughout this section, $\mathcal{M}$ is assumed to be an orientable Riemannian 3-manifold.

After fixing an orientation on $\mathcal{M}$, its metric tensor induces a "volume form" ω. This in turn leads to the notion of *volume* defined via the integral

$$\mathrm{vol}(\mathcal{D}) = \int_{\mathcal{D}} \omega$$

for any open set $\mathcal{D} \subseteq \mathcal{M}$. Also, any submanifold of $\mathcal{M}$ (e.g., surfaces) has a Riemannian metric induced by the metric tensor of $\mathcal{M}$. Thus we may measure lengths of paths and areas of surfaces in $\mathcal{M}$ as well. We refer to [101, Section 1.2] for details.

If $\mathcal{M}$ is compact, then $\mathrm{vol}(\mathcal{M})$ is finite. The next striking result has been of paramount importance in geometric topology, as it implies that "metric properties" of finite-volume hyperbolic 3-manifolds are actually topological invariants.

Theorem 5.6 (Mostow Rigidity Theorem [112], cf. [4, Theorem 1.7.1], [101, Chapter 13])**.** *Let $\mathcal{M}$, $\mathcal{N}$ be finite-volume hyperbolic 3-manifolds. Every isomorphism $\pi_1(\mathcal{M}) \to \pi_1(\mathcal{N})$ is induced by a unique isometry $\mathcal{M} \to \mathcal{N}$.*

Corollary 5.7. *If two hyperbolic 3-manifolds have different volume, then they cannot be homotopy equivalent.*

Therefore the volume can be considered a *geometric width parameter*—a topological invariant of a hyperbolic 3-manifold derived from the geometry thereof.

In a series of two papers, Hoffoss and Maher have introduced additional geometric width parameters based on Morse functions [62, 63] (cf. [134, Appendix B] for a quick refresher on Morse theory). For any Morse function $f\colon \mathcal{M} \to \mathbb{R}$, we may define its *area*, denoted $\mathrm{Area}(f)$, as the supremal area of any *level set* $f^{-1}(t) \subset \mathcal{M}$ ranging over the regular values $t \in \mathbb{R}$ of f. (Note that for any regular value $t \in \mathbb{R}$ of f, the *level set* $f^{-1}(t)$ is a codimension-one submanifold of $\mathcal{M}$, i.e., a surface.) Now the *Morse area* $\mathscr{A}_{\mathrm{M}}(\mathcal{M})$ of $\mathcal{M}$ is the infimal area of any Morse function on $\mathcal{M}$ [62]. In short,

$$\mathscr{A}_{\mathrm{M}}(\mathcal{M}) = \inf\left\{\mathrm{Area}(f) \mid f\colon \mathcal{M} \to \mathbb{R} \text{ is a Morse function}\right\}. \tag{5.4}$$

In [63] the *Gromov area* of a Riemannian 3-manifold is defined. This is an adaptation of Morse area based on *generalized Morse functions* $\mu\colon \mathcal{M} \to Y$, where Y is a trivalent (i.e., 3-regular) graph. To define such functions, first recall that the critical points $\mathrm{Crit}(f)$ of a Morse function $f\colon \mathcal{M} \to \mathbb{R}$ form a finite subset of $\mathcal{M}$, they are classified by their index (0, 1, 2, or 3), and $f|_{\mathrm{Crit}(f)}$ is an injection. For the local structure of the level sets around a critical point, there are two possibilities:

A. If $p \in \mathrm{Crit}(f)$ is of index 0 or 3, then there is a level set preserving homeomorphism $h : N(p) \to \mathcal{B}_1(0)$ between a small open ball $N(p)$ around p in $\mathcal{M}$, and the unit open ball $\mathcal{B}_1(0) \subset \mathbb{R}^3$ centered at the origin $0 \in \mathbb{R}^3$ and regarded as the union $\mathcal{B}_1(0) = \bigcup_{t\in[0,1)} S_t$ of the level sets $S_t = \{(x, y, z) \in \mathbb{R}^3 : x^2 + y^2 + z^2 = t\}$.

B. For $p \in \mathrm{Crit}(f)$ of index 1 or 2, we can almost repeat the previous paragraph, but now we equip $\mathcal{B}_1(0)$ with the level set structure $H_t = \{(x, y, z) \in \mathcal{B}_1(0) : x^2 + y^2 - z^2 = t\}$, where $t \in (-1, 1)$. Figure 5.1 shows what these level sets typically look like.

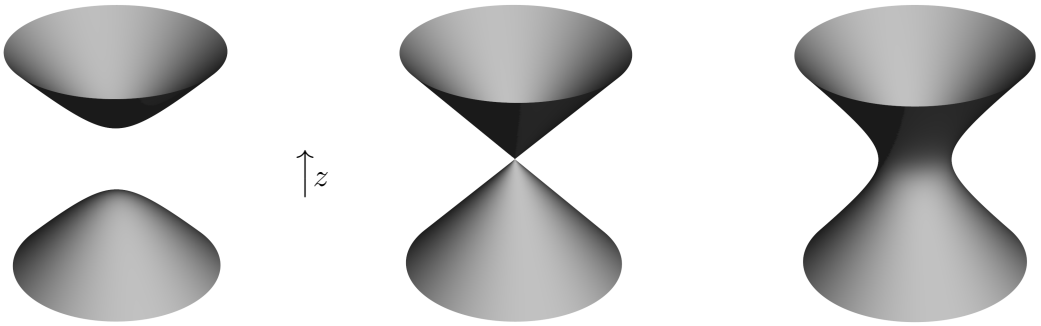

(i) $t < 0$ (ii) $t = 0$ (iii) $t > 0$

Figure 5.1: Solution sets of the equation $x^2 + y^2 - z^2 = t$ near the origin in $\mathbb{R}^3$

Observe that $\mathcal{B}_1(0) \setminus H_0$ has three connected components: $H_+ = \{(x, y, z) \in \mathcal{B}_1(0) : x^2 + y^2 - z^2 > 0\}$ and the two components of $H_- = \{(x, y, z) \in \mathcal{B}_1(0) : x^2 + y^2 - z^2 < 0\}$. Thus, if $p \in \mathcal{M}$ is a critical point of index 1 or 2, and $h\colon N(p) \to \mathcal{B}_1(0)$ is a level set preserving homeomorphism, then $h^{-1}(H_0)$ divides $N(p)$ into three connected components as well (Figure 5.2(i)). Let $\mathscr{C}$ be the quotient map that collapses each component of the level sets $h^{-1}(H_t)$ into a single point. Then $\mathscr{C}(N(p))$ is a "partial graph" with one node $\mathscr{C}(h^{-1}(H_0))$ and three "open-ended" arcs emanating from it (Figure 5.2(ii)). We call $\mathscr{C}$ a *collapsing map* and p a *trivalent singularity*.

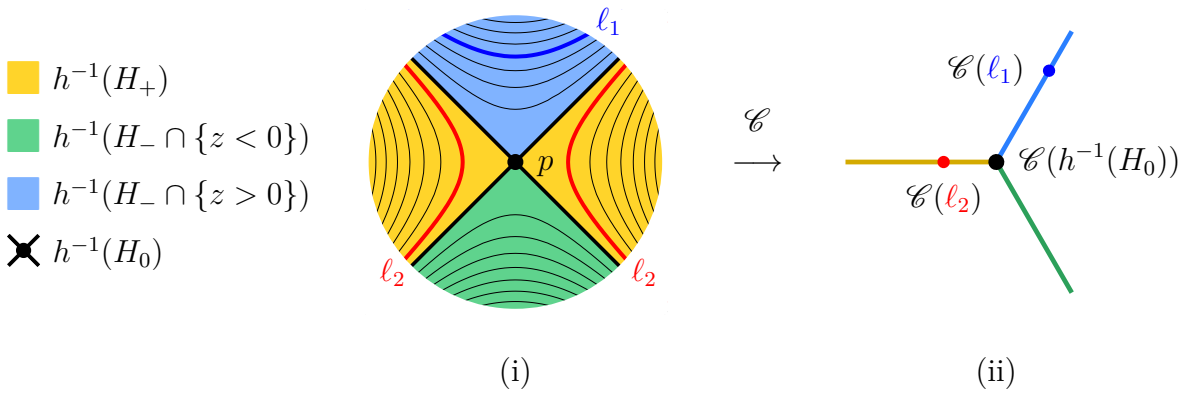

Figure 5.2: (i) Two-dimensional slice of a small open ball $N(p)$ around a critical point $p \in \mathcal{M}$ of index 1 or 2, with contour lines indicating some of the level sets. (ii) The image of $N(p)$ under the collapsing map $\mathscr{C}$

Now let $Y = (V, E)$ be a 3-regular graph (also regarded as a topological space). A map $\mu\colon \mathcal{M} \to Y$ is a *generalized Morse function*, or simply *Morse*, if the restriction $\mu|_{\mu^{-1}(\mathrm{int}(e))}$ to the preimage of the interior of any $e \in E$ is Morse (in the classical sense), and for every $v \in V$, the preimage $\mu^{-1}(v)$ is either empty or contains a unique trivalent singularity around which μ is locally isotopic to a collapsing map.

Similarly as before, the $\mathrm{Area}(\mu)$ of a generalized Morse function $\mu\colon \mathcal{M} \to Y$ is defined to be the maximum area of any preimage $\mu^{-1}(y)$ for $y \in Y$, and the *Gromov area* $\mathscr{A}_{\mathrm{G}}(\mathcal{M})$ of $\mathcal{M}$ is the infimal area of any Morse function on $\mathcal{M}$ to any 3-regular graph [63], i.e.,

$$\mathscr{A}_{\mathrm{G}}(\mathcal{M}) = \inf \{\mathrm{Area}(\mu) \mid Y \text{ is a 3-regular graph, } \mu\colon \mathcal{M} \to Y \text{ is Morse}\}. \tag{5.5}$$

Remark 5.8. Any classical Morse function can be regarded as a generalized one. Indeed, given a Morse function $f\colon \mathcal{M} \to \mathbb{R}$, take a sufficiently large open interval $\mathcal{I} \subset \mathbb{R}$ containing $f(\mathcal{M})$. Consider an inclusion $\iota\colon \mathcal{I} \to Y$ into a 3-regular graph Y, such that $\iota(\mathcal{I}) \subset e$ for some arc e of Y. Then $\iota \circ f\colon \mathcal{M} \to Y$ is a generalized Morse function.

From Remark 5.8 it immediately follows that for any Riemannian 3-manifold

$$\mathscr{A}_{\mathrm{G}}(\mathcal{M}) \leq \mathscr{A}_{\mathrm{M}}(\mathcal{M}). \tag{5.6}$$

Remark 5.9. Using the language of *foliations*, one can considerably streamline the discussion of generalized Morse functions. Consult [134, Section 7.5] and the references therein for an introduction to this framework.

Remark 5.10. One advantage of generalized Morse functions over classical ones is their ability to differentiate between the connected components of a given level set (of a classical Morse function). This is illustrated in Figure 5.3 via a 2-dimensional example.

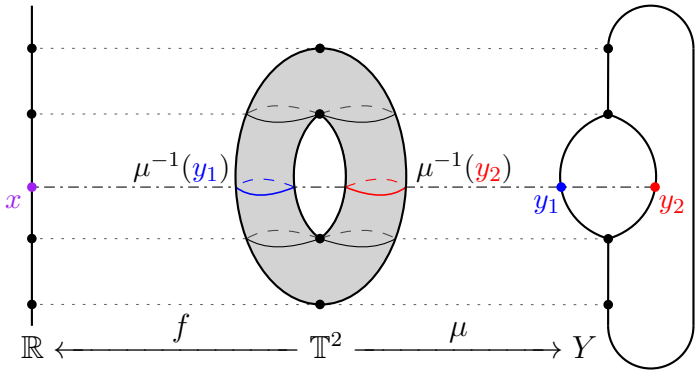

Figure 5.3: Comparison of a classical (f) and a generalized (μ) Morse function on $\mathbb{T}^2$. The latter gives a finer picture of the level sets, e.g., $f^{-1}(x) = \mu^{-1}(y_1) \cup \mu^{-1}(y_2)$

Theorem 5.6 (Mostow Rigidity) implies that, for hyperbolic 3-manifolds the Morse area and the Gromov area are topological invariants. Moreover, by the following result, they are within a constant factor to the linear width and the graph width, respectively.

Theorem 5.11 (Hoffoss–Maher, Theorem 1.2 of [62] and Theorem 1.3 of [63])**.** *There exist universal constants $K_{\mathrm{M}}, K_{\mathrm{G}} > 0$, such that for any closed hyperbolic 3-manifold $\mathcal{M}$ with linear width $\mathscr{L}(\mathcal{M})$, graph width $\mathscr{G}(\mathcal{M})$, Morse area $\mathscr{A}_{\mathrm{M}}(\mathcal{M})$, and Gromov area $\mathscr{A}_{\mathrm{G}}(\mathcal{M})$ the following inequalities hold:*

(1) $K_{\mathrm{M}}\mathscr{L}(\mathcal{M}) \leq \mathscr{A}_{\mathrm{M}}(\mathcal{M}) \leq 4\pi\mathscr{L}(\mathcal{M})$, *and*

(2) $K_{\mathrm{G}}\mathscr{G}(\mathcal{M}) \leq \mathscr{A}_{\mathrm{G}}(\mathcal{M}) \leq 4\pi\mathscr{G}(\mathcal{M})$.

Remark 5.12. Both left-hand side inequalities in Theorem 5.11 rely on a well-known conjecture by Pitts and Rubinstein [119] asserting that, under some natural assumptions, a Heegaard surface in a Riemannian 3-manifold can be isotoped to be a minimal surface. A proof of this conjecture was recently announced by Ketover–Liokumovich–Song [78].

5.2 From Combinatorics to Topology...

5.2.1 An Obstruction to Bounded Cutwidth and Pathwidth

In this section we prove Theorem 1.1, that is, we show that there exists an infinite family of 3-manifolds not admitting triangulations of uniformly bounded pathwidth. The proof relies on the following bound on the linear width of a 3-manifold in terms of its cutwidth.

Theorem 5.13. *Let $\mathcal{M}$ be a closed, orientable 3-manifold. Then, for its linear width $\mathscr{L}(\mathcal{M})$ and cutwidth $\mathrm{cw}(\mathcal{M})$, we have $\mathscr{L}(\mathcal{M}) \leq 3\,\mathrm{cw}(\mathcal{M}) + 4$.*

For non-Haken 3-manifolds, Theorem 5.13, in combination with previously mentioned results, implies the following.

Theorem 1.3 (rephrased using the notation from Section 5.1)**.** *Let $\mathcal{M}$ be a closed, orientable, irreducible and non-Haken 3-manifold. Then, for its Heegaard genus $\mathfrak{g}(\mathcal{M})$ and pathwidth $\mathrm{pw}(\mathcal{M})$, we have $\mathfrak{g}(\mathcal{M}) \leq 4\,(3\,\mathrm{pw}(\mathcal{M}) + 1)$.*

Proof of Theorem 1.3 assuming Theorem 5.13. By Theorem 2.4, $\mathrm{cw}(\mathcal{M}) \leq 4\,\mathrm{pw}(\mathcal{M})$ since dual graphs of 3-manifold triangulations are 4-regular. By Corollary 5.2, $\mathscr{L}(\mathcal{M}) = \mathfrak{g}(\mathcal{M})$ whenever $\mathcal{M}$ is irreducible and non-Haken. Combining these relations with the inequality provided by Theorem 5.13 yields the desired result. □

Theorem 1.1 is now obtained from Theorem 1.3 and Agol's Theorem 5.4. It remains to prove Theorem 5.13. To this end, let $\mathcal{M}$ be a closed, orientable 3-manifold and $\mathcal{T}$ be a triangulation of $\mathcal{M}$. Consider the canonical handle decomposition $\mathrm{chd}(\mathcal{T})$ associated with this triangulation (Definition 3.2). In what follows, we construct a linear generalized Heegaard splitting of $\mathcal{M}$ based on $\mathrm{chd}(\mathcal{T})$, such that the genera of the boundary surfaces of the compression bodies in the splitting is upper-bounded by $3\,\mathrm{cw}(\Gamma(\mathcal{T})) + 4$.

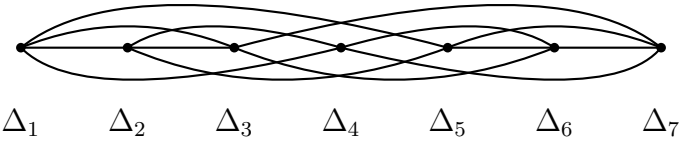

Figure 5.4: A linear layout of the dual graph $\Gamma(\mathcal{T})$ of the smallest triangulation $\mathcal{T}$ of $\mathcal{M} = \mathrm{SFS}[\mathbb{T}^2 : (1,2)]$. The largest cutset has six arcs, hence $\mathrm{cw}(\mathcal{M}) \leq 6$. Source: [28]

Note that, if $\mathcal{T}$ consists of a single tetrahedron, then $\mathcal{M} \in \{\mathbb{S}^3, \mathrm{L}(4,1), \mathrm{L}(5,2)\}$, see [28, Closed Census]. Now these 3-manifolds all have Heegaard genus, hence linear width, at most one, so they satisfy Theorem 5.13. Thus, for the remainder, we assume that $\mathcal{M}$ is a connected 3-manifold and $\mathcal{T}$ consists of at least two tetrahedra.

Construction of the linear splitting. Fix an ordering $\Delta_1 < \Delta_2 < \ldots < \Delta_n$ of the tetrahedra in $\mathcal{T}$, so that the corresponding linear layout of the dual graph $\Gamma(\mathcal{T})$ realizes minimal cutwidth (see, e.g., Figure 5.4). Let $\emptyset = \mathcal{T}_0 \subset \mathcal{T}_1 \subset \mathcal{T}_2 \subset \ldots \subset \mathcal{T}_n = \mathcal{M}$ be a filtration of $\mathcal{M}$ by submanifolds, where $\mathcal{T}_j$ is the union of all those handles of $\mathrm{chd}(\mathcal{T})$ that are disjoint from the union of tetrahedra $\Delta_i \subset \mathcal{M}$ with $i > j$. We now turn this filtration into a linear splitting $\mathcal{M} = (\mathcal{N}_1 \cup_{\mathcal{S}_1} \mathcal{K}_1) \cup_{\mathcal{R}_1} \cdots \cup_{\mathcal{R}_{n-1}} (\mathcal{N}_n \cup_{\mathcal{S}_n} \mathcal{K}_n)$. The construction is the same as in Example 4.1, but here we spell out the details. To this end, recall the primal and dual constructions of compression bodies (Figures 3.6 and 3.7), and, for each $i \in \{1, \ldots, n\}$, let $h(\Delta_i)$ denote the 0-handle in $\mathrm{chd}(\mathcal{T})$ corresponding to Δ_i.

Set $\mathcal{R}_0 = \emptyset$. We think of $\mathcal{R}_0$ as the empty surface. Inductively, for $i \in \{1, \ldots, n\}$, two compression bodies $\mathcal{N}_i$ and $\mathcal{K}_i$ are constructed, in two steps each, as follows.

Construction of $\mathcal{N}_i$

P1$_{(i)}$: Take the disjoint union of $\mathcal{R}_{i-1} \times [0,1]$ and $h(\Delta_i)$, and

P2$_{(i)}$: attach all the 1-handles of $\mathcal{T}_i \setminus \mathcal{T}_{i-1}$ along $\mathcal{R}_{i-1} \times \{1\} \cup \partial h(\Delta_i)$.

The resulting compression body is denoted by $\mathcal{N}_i$.

We set $\partial_- \mathcal{N}_i = \mathcal{R}_{i-1} \times \{0\}$ and $\mathcal{S}_i = \partial_+ \mathcal{N}_i = \partial \mathcal{N}_i \setminus \partial_- \mathcal{N}_i$.

Construction of $\mathcal{K}_i$

D1$_{(i)}$: Start with the thickened surface $\mathcal{S}_i \times [0,1]$, and

D2$_{(i)}$: attach all the 2- and 3-handles of $\mathcal{T}_i \setminus \mathcal{T}_{i-1}$ along $\mathcal{S}_i \times \{1\}$.

The resulting compression body is denoted by $\mathcal{K}_i$.

We set $\partial_+ \mathcal{K}_i = \mathcal{S}_i \times \{0\}$ and $\mathcal{R}_i = \partial_- \mathcal{K}_i = \partial \mathcal{K}_i \setminus \partial_+ \mathcal{K}_i$.

The desired linear splitting of $\mathcal{M}$ is obtained by identifying $\partial_+ \mathcal{N}_i$ with $\partial_+ \mathcal{K}_i$ $(1 \leq i \leq n)$, and $\partial_- \mathcal{K}_i$ with $\partial_- \mathcal{N}_{i+1}$ $(1 \leq i \leq n-1)$. Note that $\mathcal{N}_1$ and $\mathcal{K}_n$ are handlebodies, thus $\partial_- \mathcal{N}_1 = \partial_- \mathcal{K}_n = \emptyset$. See Figure 5.5 for an illustration of this construction.

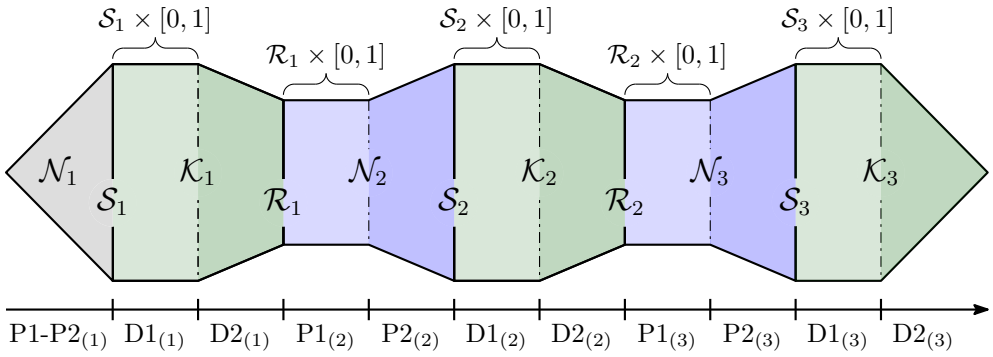

Figure 5.5: The linear splitting of $\mathcal{M}$ constructed in the proof of Theorem 5.13

Bounding the width. Recall that, by the definition of linear width, we have

$$\mathscr{L}(\mathcal{M}) \leq \max_{1 \leq i \leq n} g(\mathcal{S}_i). \tag{5.7}$$

In what follows we give an upper bound on the maximum genus of any surface $\mathcal{S}_i$ that appears in the linear splitting above. The first step is the following general observation.

Lemma 5.14. *Let $\mathcal{T}$ be a triangulation of a 3-manifold $\mathcal{M}$ and $\Delta_1 < \Delta_2 < \ldots < \Delta_n$ be a linear ordering of its tetrahedra. Let $\emptyset = \mathcal{T}_0 \subset \mathcal{T}_1 \subset \mathcal{T}_2 \subset \ldots \subset \mathcal{T}_n = \mathcal{M}$ be a filtration of $\mathcal{M}$ by codimension zero submanifolds as before, where $\mathcal{T}_j$ is the union of all those handles of* $\mathrm{chd}(\mathcal{T})$ *that are disjoint from the union of tetrahedra $\Delta_i \subset \mathcal{M}$ with $i > j$. Then passing from $\mathcal{T}_j$ to $\mathcal{T}_{j+1}$ corresponds to adding at most* 15 *handles. Moreover, if $\widetilde{\mathcal{T}}_j$ denotes a codimension zero submanifold of $\mathcal{M}$ constructed from $\mathcal{T}_j$ by adding an arbitrary subset of these handles, then, for the genus of its boundary, we have*

$$g(\partial \widetilde{\mathcal{T}}_j) \leq g(\partial \mathcal{T}_j) + 4. \tag{5.8}$$

Proof of Lemma 5.14. This is apparent from the fact that every tetrahedron consists of 15 (non-empty) faces and thus at most 15 handles of $\mathrm{chd}(\mathcal{T})$ are disjoint from $\mathcal{T}_j$ but not disjoint from $\mathcal{T}_{j+1}$. In particular, at most 15 handles are added at the j-th level of the filtration. Moreover, note that at most four of the handles added in every step are 1-handles (corresponding to the four triangular faces of the tetrahedron), which are the only handles that can increase the genus of $\partial \mathcal{T}_j$. □

Corollary 5.15. *For the surfaces $\mathcal{R}_{i-1}$ and $\mathcal{S}_i$ $(1 \leq i \leq n)$ constructed before, we have*

$$g(\mathcal{S}_i) \leq g(\mathcal{R}_{i-1}) + 4. \tag{5.9}$$

Proof of Corollary 5.15. By construction, $\mathcal{R}_{i-1} = \partial\mathcal{T}_{i-1}$, and $\mathcal{S}_i$ is obtained by attaching at most four 1-handles to a parallel copy of $\mathcal{R}_{i-1}$, hence (5.9) follows from Lemma 5.14. □

We can now conclude the proof of Theorem 5.13.

Proof of Theorem 5.13. Let v_j $(1 \leq j \leq n)$ be the nodes of $\Gamma(\mathcal{T})$ with corresponding tetrahedra $\Delta_j \in \mathcal{T}$ $(1 \leq j \leq n)$. We may assume, without loss of generality, that the largest cutset of the linear ordering $v_1 < v_2 < \ldots < v_n$ has cardinality $\mathrm{cw}(\Gamma(\mathcal{T})) = k$.

Let $\mathcal{T}_j \subset \mathrm{chd}(\mathcal{T})$, $1 \leq j \leq n$, be the filtration from Lemma 5.14. Moreover, let C_j, $1 \leq j < n$, be the cutsets of the linear ordering above (cf. Definition 2.3). Naturally, the cutset C_j can be associated with at most k triangles of $\mathcal{T}$ with, together, at most $3k$ edges and at most $3k$ vertices of $\mathcal{T}$. Let $\mathcal{M}(C_j) \subset \mathrm{chd}(\mathcal{T})$ be the corresponding submanifold formed from the at most k 1-handles and at most $3k$ 2- and 3-handles each of $\mathrm{chd}(\mathcal{T})$ associated with these faces of $\mathcal{T}$.

By construction, the boundary of $\mathcal{M}(C_j)$ decomposes into two parts, one of which coincides with $\partial\mathcal{T}_j = \mathcal{R}_j$. Since $\mathcal{M}(C_j)$ is the "neighborhood of k triangles in $\mathcal{T}$," and since the 2- and 3-handles of $\mathrm{chd}(\mathcal{T})$ form a handlebody, the 2- and 3-handles of $\mathcal{M}(C_j)$ form a union of handlebodies with sum of genera at most $3k$. To complete the construction of $\mathcal{M}(C_j)$, at most k 1-handles are attached to this union of handlebodies as 2-handles. These either increase the number of boundary surface components, or decrease the overall sum of genera of the boundary components. Altogether, the sum of genera of $\mathcal{R}_j \subset \partial\mathcal{M}(C_j)$ is bounded above by $3k$.

Hence, following Corollary 5.15, we have $g(\mathcal{S}_i) \leq 3k + 4 = 3\,\mathrm{cw}(\Gamma(\mathcal{T})) + 4$, which, together with (5.7), implies $\mathscr{L}(\mathcal{M}) \leq 3\,\mathrm{cw}(\mathcal{M}) + 4$. □

5.2.2 An Obstruction to Bounded Congestion and Treewidth

In this section is to prove Theorems 1.2 and 1.4, the counterparts of Theorem 1.1 and 1.3 for treewidth. At the core of the proof is the following explicit connection between the congestion of the dual graph of any triangulation of a 3-manifold $\mathcal{M}$ and its graph width.

Theorem 5.16. *Let $\mathcal{M}$ be a closed, connected, orientable 3-manifold of graph width $\mathscr{G}(\mathcal{M})$, and $\mathcal{T}$ be a triangulation of $\mathcal{M}$ with dual graph $\Gamma(\mathcal{T})$ of congestion* $\mathrm{cng}(\Gamma(\mathcal{T}))$. *Then either $\mathscr{G}(\mathcal{M}) \leq \frac{9}{2}\,\mathrm{cng}(\Gamma(\mathcal{T}))$, or $\mathcal{T}$ only contains one tetrahedron.*

Proof of Theorem 1.4 assuming Theorem 5.16. First note, that the only closed orientable 3-manifolds which can be triangulated with a single tetrahedron are the 3-sphere of Heegaard genus zero, and the lens spaces of type $\mathrm{L}(4,1)$ and $\mathrm{L}(5,2)$ of Heegaard genus one (see Example 3.11) for which Theorem 1.4 holds. Otherwise, since dual graphs of 3-manifold triangulations are 4-regular, Theorem 2.6 yields $\mathrm{cng}(\Gamma(\mathcal{T})) \leq 4(\mathrm{tw}\,(\Gamma(\mathcal{T})) + 1)$. In addition, Theorem 5.3 implies that $\mathscr{G}(\mathcal{M}) = \mathfrak{g}\,(\mathcal{M})$ whenever $\mathcal{M}$ is irreducible and non-Haken. Combining these relations with the inequality provided by Theorem 5.16 we obtain Theorem 1.4. □

Similar as in Section 5.2.1, Theorem 1.2 immediately follows from Theorems 1.4 and 5.4. Hence, the remainder of the section is dedicated to the proof of Theorem 5.16.

Proof of Theorem 5.16. Let $\mathcal{M}$ be a closed, connected, orientable 3-manifold, and $\mathcal{T}$ be a triangulation of $\mathcal{M}$ whose dual graph $\Gamma(\mathcal{T})$ has congestion $\mathrm{cng}(\Gamma(\mathcal{T})) = k$. Furthermore, let H be an unrooted binary tree realizing $\mathrm{cng}(\Gamma(\mathcal{T})) = k$ (cf. Definition 2.5).

If $k = 0$, then $\mathcal{T}$ must consist of a single tetrahedron, and the theorem holds. Thus we may assume that $k \geq 1$.

The idea of the proof is to first construct a graph splitting of $\mathcal{M}$ from a fork complex $\boldsymbol{F}$ modeled on H (cf. Section 4.1.2), and then to analyze the genera of the upper boundary surfaces appearing in the splitting to see that they are all bounded above by $\frac{9}{2}\,\mathrm{cng}(\Gamma(\mathcal{T}))$.

Construction of the graph splitting. Consider the canonical handle decomposition $\mathrm{chd}(\mathcal{T})$ of $\mathcal{M}$ associated with $\mathcal{T}$ as defined in Definition 3.2. Every compression body in the graph splitting described below is either a union of handles in $\mathrm{chd}(\mathcal{T})$, a thickened surface parallel to the boundary surface of some union of handles from $\mathrm{chd}(\mathcal{T})$, or a combination of both. In particular, the graph splitting maintains the handle structure coming from $\mathrm{chd}(\mathcal{T})$. Note that we do *not* require the following compression bodies to be connected, but rather to be the union of connected compression bodies.

For every leaf $w \in V(H)$, a handlebody $\mathcal{H}_w$ is constructed as follows. Consider the abstract tetrahedron $\Delta_w \in \mathcal{T}$ associated to w. If Δ_w has no self-identifications in $\mathcal{T}$, then $\mathcal{H}_w$ is just the 0-handle of $\mathrm{chd}(\mathcal{T})$ corresponding to Δ_w. If Δ_w exhibits self-identifications, then first note that at most one pair of triangular faces of Δ_w are identified, otherwise Δ_w would be disjoint from the rest of $\mathcal{T}$. Up to symmetry there are two possibilities:

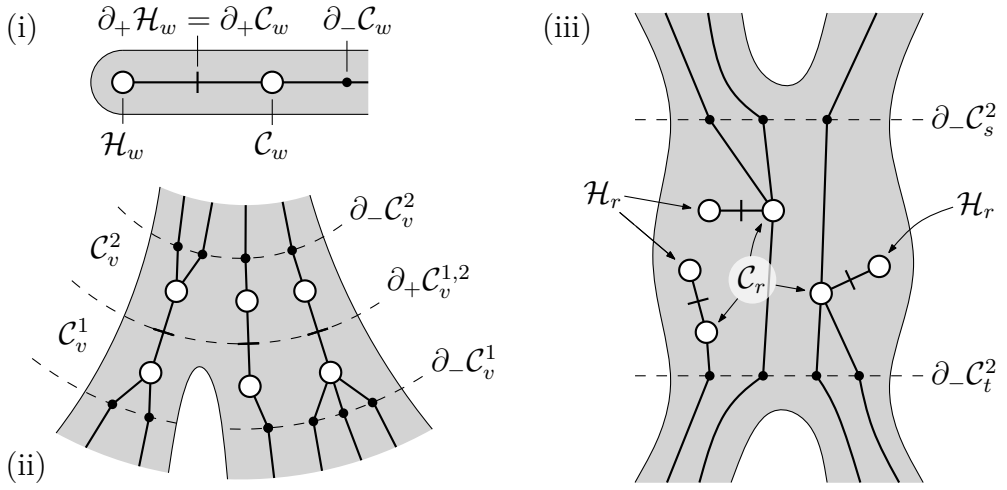

Figure 5.6: Local pictures of the fork complex $\boldsymbol{F}$ built in the proof of Theorem 5.16

Either, Δ_w forms a "snapped" 3-ball in $\mathcal{T}$, see Figure 5.7(i), in which case $\mathcal{H}_w$ is built from the 0-handle, the 1-handle, and the 2-handle of $\mathrm{chd}(\mathcal{T})$ corresponding to Δ_w, to the face gluing, and to the edge $\{1,2\}$ of Δ_w, respectively. Or, Δ_w forms a solid torus in $\mathcal{T}$, see Figure 5.7(ii), and then $\mathcal{H}_w$ consists of the 0-handle and of the 1-handle of $\mathrm{chd}(\mathcal{T})$, corresponding to Δ_w and to the face gluing, respectively. Moreover, for every leaf $w \in V(H)$, a compression body $\mathcal{C}_w = \partial\mathcal{H}_w \times [0,1]$ is attached to $\partial\mathcal{H}_w$ along $\partial_+\mathcal{C}_w = \partial\mathcal{H}_w \times \{0\}$. See Figure 5.6(i).

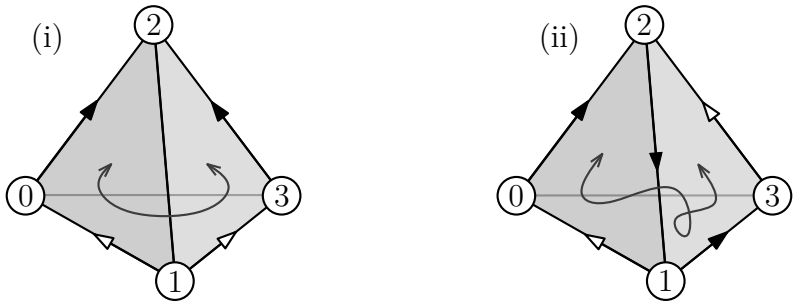

Figure 5.7: (i) A snapped 3-ball, and(ii) a one-tetrahedron solid torus

Before proceeding, let us fix a "root arc" $r \in E(H)$. This choice induces a partial order on $V(H)$: for $x, y \in V(H)$, $x \prec y$ if and only if y is contained by the path connecting x with r. We also say "x is below y". In particular, $x \prec x$ for all $x \in V(H)$. Given $x \in V(H)$, let $\mathcal{T}_x$ denote the submanifold of $\mathcal{M}$ consisting of

- any 0-handle of $\mathrm{chd}(\mathcal{T})$ corresponding to a leaf of H below x,
- any 1-handle of $\mathrm{chd}(\mathcal{T})$ where both adjacent 0-handles are in $\mathcal{T}_x$

- any 2-handle of $\mathrm{chd}(\mathcal{T})$ where all adjacent 0-handles are in $\mathcal{T}_x$, and
- any 3-handle of $\mathrm{chd}(\mathcal{T})$ where all adjacent 0-handles are in $\mathcal{T}_x$.

In other words, $\mathcal{T}_x$ is the submanifold of $\mathcal{M}$ *spanned by* the 0-handles of $\mathrm{chd}(\mathcal{T})$ below x.

Claim 1. If $x \prec y$ then $\mathcal{T}_x \subseteq \mathcal{T}_y$. If $x \not\prec y$ and $y \not\prec x$ then $\mathcal{T}_x \cap \mathcal{T}_y = \emptyset$ and $\partial\mathcal{T}_x \cap \partial\mathcal{T}_y = \emptyset$.

Proof of Claim 1. The first part of the claim is obvious. For the second part, let $x, y \in V(H)$ with $x \not\prec y$ and $y \not\prec x$. The way we construct $\mathcal{T}_x$ and $\mathcal{T}_y$ ensures that $\mathcal{T}_x$ and $\mathcal{T}_y$ do not only have disjoint interiors, but are also separated from each other such that their boundaries are disjoint as well. Indeed, if $\mathcal{T}_x$ and $\mathcal{T}_y$ are both collections of 0-handles this is certainly true as each of these 0-handles can be thought of as living inside a single tetrahedron of $\mathcal{T}$ away from its boundary. As soon as we have two such 0-handles living in two adjacent tetrahedra in, say, $\mathcal{T}_x$, the 1-handle(s) corresponding to their common triangular face(s) can be glued to the two 0-handles and the resulting boundary is still disjoint from any $\mathcal{T}_y$ being such a collection of 0- and 1-handles itself. Now fix an edge f of $\mathcal{T}$ and suppose that $\mathcal{T}_x$ contains all 0- and 1-handles associated to tetrahedra and triangles around f, then we are safe to add the 2-handle corresponding to f to $\mathcal{T}_x$ and still be disjoint from $\mathcal{T}_y$ even if $\mathcal{T}_y$ itself is such a collection of 0-, 1-, and 2-handles. The case of adding 3-handles to $\mathcal{T}_x$ and $\mathcal{T}_y$ corresponding to vertices of $\mathcal{T}$ is completely analogous to the previous case. This proves the second part of the claim.

We can now describe the remaining parts of the graph splitting of $\mathcal{M}$. For this, let $v \in V(H)$ be a degree three node, $u, u' \in V(H)$ be the two nodes below and incident to v, and let $e, e' \in E(H)$ denote the arcs with endpoints v and u, u', respectively. For every such degree three node $v \in V(H)$, a pair of compression bodies $(\mathcal{C}_v^1, \mathcal{C}_v^2)$ is constructed in two steps. (See Figure 5.6(ii) for an example of a local schematic picture of $\boldsymbol{F}$ around v.)

1. To construct the first compression body $\mathcal{C}_v^1$, we start with $(\partial(\mathcal{T}_u \cup \mathcal{T}_{u'})) \times [0,1]$ and attach to $(\partial(\mathcal{T}_u \cup \mathcal{T}_{u'})) \times \{1\}$ all 1-handles of $\mathrm{chd}(\mathcal{T})$ corresponding to triangles of $\mathcal{T}$ associated to arcs of $\Gamma(\mathcal{T})$ running parallel to e *and* e'. (These are the remaining 1-handles of $\mathcal{T}_v$ not attached earlier.) We then define its lower and upper boundary as $\partial_- \mathcal{C}_v^1 = (\partial(\mathcal{T}_u \cup \mathcal{T}_{u'})) \times \{0\}$ and $\partial_+ \mathcal{C}_v^1 = \partial\mathcal{C}_v^1 \setminus \partial_- \mathcal{C}_v^1$, respectively.

2. For the second compression body $\mathcal{C}_v^2$, we start with $\partial_+ \mathcal{C}_v^1 \times [0,1]$ (with the top boundary being defined as $\partial_+ \mathcal{C}_v^2 = \partial_+ \mathcal{C}_v^1 \times \{1\}$). The compression body is then

completed by attaching along $\partial_+\mathcal{C}_v^1 \times \{0\}$ all 2- and 3-handles of $\mathrm{chd}(\mathcal{T})$ which are contained in $\mathcal{T}_v$ but not in $\mathcal{T}_u \cup \mathcal{T}_{u'}$. We set $\partial_-\mathcal{C}_v^2 = \partial\mathcal{C}_v^2 \setminus \partial_+\mathcal{C}_v^1$.

For the root arc $r = \{s, t\}$, we construct a pair of compression bodies $(\mathcal{C}_r, \mathcal{H}_r)$ as follows. ($\mathcal{H}_r$ is a union of handlebodies, hence the notation.)

1. To build $\mathcal{C}_r$, start with $(\partial_-\mathcal{C}_s^2 \cup \partial_-\mathcal{C}_t^2) \times [0, 1]$, define the lower boundary as $\partial_-\mathcal{C}_r = (\partial_-\mathcal{C}_s^2 \cup \partial_-\mathcal{C}_t^2) \times \{0\}$ and attach to $(\partial_-\mathcal{C}_s^2 \cup \partial_-\mathcal{C}_t^2) \times \{1\}$ all 1-handles of $\mathrm{chd}(\mathcal{T})$ corresponding to arcs of $\Gamma(\mathcal{T})$ routed *through* r. As usual, $\partial_+\mathcal{C}_r = \partial\mathcal{C}_r \setminus \partial_-\mathcal{C}_r$.

2. Finally, to obtain $\mathcal{H}_r$, take $\partial_+\mathcal{C}_r \times [0, 1]$, set $\partial_+\mathcal{H}_r = \partial_+\mathcal{C}_r \times \{1\}$ and identify it with $\partial_+\mathcal{C}_r$, and attach all remaining 2- and 3-handles of $\mathrm{chd}(\mathcal{T})$ to $\partial_+\mathcal{C}_r \times \{0\}$.

Figure 5.6(iii) shows a possible scenario near the root arc. This finishes the construction.

Claim 2. The compression bodies $\mathcal{H}_w$, $\mathcal{C}_w$, $\mathcal{C}_v^1$, $\mathcal{C}_v^2$, $\mathcal{C}_r$, and $\mathcal{H}_r$ (where $w, v \in V(H)$, $\deg(w) = 1$, $\deg(v) = 3$, and $r \in E(H)$ is the root arc), glued together along their appropriate boundary components, form a graph splitting of $\mathcal{M}$.

Proof of Claim 2. It follows from Claim 1 and the construction that all compression bodies above have pairwise disjoint interiors. We check that their lower and upper bonudary components match up whenever they are identified (cf. Figure 5.6). For the identifications between $\mathcal{H}_w$ and $\mathcal{C}_w$, between $\mathcal{C}_v^1$ and $\mathcal{C}_v^2$, and between $\mathcal{C}_r$ and $\mathcal{H}_r$ this is immediate. Now let $v \in V(H)$ be of degree three with adjacent nodes $u, u' \in V(H)$ below. Note that, by construction, $\partial_-\mathcal{C}_u^2 = \partial\mathcal{T}_u$ and $\partial_-\mathcal{C}_{u'}^2 = \partial\mathcal{T}_{u'}$ and they are disjoint by Claim 1. Hence $\partial_-\mathcal{C}_v^1 = (\partial(\mathcal{T}_u \cup \mathcal{T}_{u'})) \times \{0\}$ can indeed be identified with the disjoint union of $\partial_-\mathcal{C}_u^2$ and $\partial_-\mathcal{C}_{u'}^2$. For the gluings between $\partial_-\mathcal{C}_r$ and $\partial_-\mathcal{C}_s^2 \cup \partial_-\mathcal{C}_t^2$, where $r = \{s, t\} \in E(H)$ is the root arc, the reasoning is analogous. Finally, as it is modeled on the tree, the fork complex $\boldsymbol{F}$ is *exact* (see [130, Definition 5.1.4]), yielding a graph splitting of $\mathcal{M}$.

Bounding the width. Following [130, Section 5.1], cf. Section 5.1.2, the width of the graph splitting of $\mathcal{M}$ exhibited above is given by the largest genus of a top boundary of a connected compression body in the graph splitting. However, this splitting, by construction, consists of unions of compression bodies. In particular, $\mathcal{C}_v^1, \mathcal{C}_v^2, \mathcal{C}_r$, and $\mathcal{H}_r$ may be disconnected. Note that this is not a problem since the sum of genera of top

boundaries for every such union cannot be smaller than the genus of the largest genus compression body in the union. Hence, with this adjustment, we are left with the multiset

$$\left\{ g(\partial\mathcal{H}_w), g(\partial_+\mathcal{C}^1_v), g(\partial\mathcal{H}_r) \;\middle|\; w, v \in V(H), \deg(w) = 1, \deg(v) = 3, r \in E(H) \text{ root arc} \right\}$$

to determine an upper bound on the graph width of $\mathcal{M}$, where $g(\mathcal{S})$ denotes the sum of the genera of all connected components of $\mathcal{S}$.

The handlebody $\mathcal{H}_w$ has genus at most one: $g(\partial\mathcal{H}_w) = 0$ if Δ_w is a 0-handle, or forms a "snapped" 3-ball in $\mathcal{T}$, and $g(\partial\mathcal{H}_w) = 1$ if Δ_w forms a solid torus in $\mathcal{T}$.

Let us fix a node $v \in V(H)$ of degree three. Our goal is to upper-bound $g(\partial_+\mathcal{C}^1_v)$. Note that, since $\mathrm{cng}(\Gamma(\mathcal{T})) \leq k$, at most k arcs of $\Gamma(\mathcal{T})$ run parallel to each arc of H. Moreover, counting those arcs of $\Gamma(\mathcal{T})$ along the three arcs of H incident to v, we encounter each of them twice, therefore at most $\frac{3}{2}k$ arcs of $\Gamma(\mathcal{T})$ meet v. Based on this fact, we show that $g(\partial_+\mathcal{C}^1_v) \leq \frac{9k}{2}$. The proof relies on the next key observation.

Claim 3. Let $x \in V(H)$ be a node of H and $a \in E(H)$ be the unique arc of H above and incident to x. Then any handle $h \in \mathrm{chd}(\mathcal{T}) \setminus \mathcal{T}_x$ that touches $\partial\mathcal{T}_x$ is adjacent or equal to a 1-handle of $\mathrm{chd}(\mathcal{T})$ corresponding to an arc of $\Gamma(\mathcal{T})$ routed through a.

Proof of Claim 3. Recall that $\mathcal{T}_x$ is spanned by those 0-handles of $\mathrm{chd}(\mathcal{T})$ that correspond to the leaves of H below x. Turning this around, every handle in $\mathrm{chd}(\mathcal{T}) \setminus \mathcal{T}_x$ is either a 0-handle whose corresponding leaf is *not* below x, or is adjacent to at least one such 0-handle. Now let $h \in \mathrm{chd}(\mathcal{T}) \setminus \mathcal{T}_x$ be a handle that touches $\partial\mathcal{T}_x$.

First, observe that h cannot be a 0-handle. Indeed, if h is a 0-handle not in $\mathcal{T}_x$ then its corresponding leaf is not below x, and therefore all $h' \in \mathrm{chd}(\mathcal{T})$ adjacent to h are not part of $\mathcal{T}_x$ either. As the union of these handles h' comprise a neighborhood of h, it follows that $\partial h \cap \partial\mathcal{T}_x = \emptyset$, contradicting the assumption that h touches $\partial\mathcal{T}_x$.

Second, notice that if h is an i-handle ($i \in \{1, 2, 3\}$) and no 0-handles adjacent to h are below x, then h is separated from $\mathcal{T}_x$ by the union of the $h' \in \mathrm{chd}(\mathcal{T}) \setminus \{h\}$ that are adjacent to at least one of these 0-handles. Thus there exists a 0-handle $h_1 \in \mathrm{chd}(\mathcal{T})$ adjacent to h with corresponding leaf below x. Moreover, some other 0-handle adjacent to h, say h_2, must be in $\mathrm{chd}(\mathcal{T}) \setminus \mathcal{T}_x$, since otherwise h must be part of $\mathcal{T}_x$.

If h is a 1-handle, then h_1 and h_2 are precisely the two 0-handles adjacent to h, which thus corresponds to an arc of $\Gamma(\mathcal{T})$ routed through $a \in E(H)$ and we are done. If h is

a 2- or 3-handle, then there is an alternating sequence $h_1 = h^{(0)}, h^{(1)}, h^{(2)}, \ldots, h^{(p)} = h_2$ of 0- and 1-handles adjacent to h, and $h^{(j)}$ being adjacent to $h^{(j+1)}$ $(0 \leq j < p)$. Since $h_1 \in \mathcal{T}_x$ and $h_2 \in \mathrm{chd}(\mathcal{T}) \setminus \mathcal{T}_x$, for some even $q \in \{0, 2, \ldots, p-2\}$ we have $h^{(q)} \in \mathcal{T}_x$ and $h^{(q+2)} \in \mathrm{chd}(\mathcal{T}) \setminus \mathcal{T}_x$. But then $h^{(q+1)}$ is a 1-handle of $\mathrm{chd}(\mathcal{T})$ adjacent to h that corresponds to an arc of $\Gamma(\mathcal{T})$ routed through a. This concludes the proof of Claim 3.

By construction, $\partial_+ \mathcal{C}_v^1 \approx \partial(\mathcal{T}_u \cup \mathcal{T}_{u'} \cup \mathscr{H}_{e,e'})$, where $\mathscr{H}_{e,e'}$ consists of the 1-handles in $\mathrm{chd}(\mathcal{T})$ corresponding to arcs of $\Gamma(\mathcal{T})$ running parallel to both $e = \{u, v\}$ and $e' = \{u', v\}$. (Here $\mathcal{S}_1 \approx \mathcal{S}_2$ denotes that the surfaces $\mathcal{S}_1$ and $\mathcal{S}_2$ are parallel, and hence, in particular, of the same genus.)

Let $\mathcal{X}$ be the submanifold of $\mathcal{M}$ built from the handles in $\mathrm{chd}(\mathcal{T}) \setminus (\mathcal{T}_u \cup \mathcal{T}_{u'})$ touching $\partial(\mathcal{T}_u \cup \mathcal{T}_{u'})$. It follows from Claim 3, that each handle in $\mathcal{X}$ is either a 1-handle routed through e or e', or a 2- or 3-handle adjacent to such a 1-handle. In particular, $\mathcal{X}$ consists of at most $\frac{3k}{2}$ 1-handles, at most $\frac{9k}{2}$ 2-handles, and at most $\frac{9k}{2}$ 3-handles, cf. the paragraph before Claim 3. Since the 2- and 3-handles of $\mathrm{chd}(\mathcal{T})$ form a handlebody, the 2- and 3-handles of $\mathcal{X}$ form a union $\mathcal{X}_{2,3}$ of handlebodies with sum of genera at most $\frac{9k}{2}$.

Consider the submanifold $\mathcal{Y} \subseteq \mathcal{X}$ obtained from $\mathcal{X}_{2,3}$ by attaching to it all 1-handles of $\mathcal{X}$ *not* in $\mathscr{H}_{e,e'}$. (These are precisely the 1-handles of $\mathrm{chd}(\mathcal{T})$ that correspond to arcs of $\Gamma(\mathcal{T})$ running parallel either to e or to e' but not to both.) Note that these 1-handles are attached to $\mathcal{X}_{2,3}$ as 2-handles. Each of these attachments either increases the number of boundary surface components, or decreases the overall sum of genera of the boundary components by one. Consequently, $g(\partial\mathcal{Y}) \leq g(\partial\mathcal{X}_{2,3}) \leq \frac{9k}{2}$. Finally, by construction, $\partial_+ \mathcal{C}_v^1$ is parallel to the union of some components of $\partial\mathcal{Y}$, and therefore $g(\partial_+ \mathcal{C}_v^1) \leq g(\partial\mathcal{Y}) \leq \frac{9k}{2}$.

Bounding above the genus of $\partial\mathcal{H}_r$ is analogous. The only difference is that $\partial\mathcal{H}_r \approx \partial(\mathcal{T}_s \cup \mathcal{T}_t \cup \mathscr{H}_r)$, where $\mathscr{H}_r$ now consists of the at most k 1-handles in $\mathrm{chd}(\mathcal{T})$ which correspond to arcs of $\Gamma(\mathcal{T})$ routed through the root arc $r = \{s, t\} \in E(H)$. Here an even stronger bound holds, i.e., $g(\partial\mathcal{H}_r) \leq 3k < \frac{9k}{2}$.

From the definition of graph width it immediately follows that $\mathscr{G}(\mathcal{M}) \leq \frac{9k}{2}$. □

5.3 ...and Back

5.3.1 Combinatorial Width and Heegaard Genus

In Section 5.2.2 it was shown that for a closed, orientable, irreducible and non-Haken 3-manifold $\mathcal{M}$, the Heegaard genus $\mathfrak{g}(\mathcal{M})$ and the treewidth $\mathrm{tw}(\mathcal{M})$ satisfy

$$\mathfrak{g}(\mathcal{M}) \leq 18(\mathrm{tw}(\mathcal{M}) + 1). \tag{5.10}$$

In this section we further explore the connection between these two parameters, guided by two questions: 1. Does a reverse inequality hold? 2. Can one refine the assumptions?

For the first one, we give an affirmative answer.

Theorem 1.6. Let $\mathcal{M}$ be a closed, orientable 3-manifold, and let $\mathrm{cw}(\mathcal{M})$ and $\mathfrak{g}(\mathcal{M})$ respectively denote the cutwidth and the Heegaard genus of $\mathcal{M}$. Then we have

$$\mathrm{cw}(\mathcal{M}) \leq 4\mathfrak{g}(\mathcal{M}) - 2. \tag{5.11}$$

The result is almost immediate if one takes a layered triangulation of a closed, orientable 3-manifold. Due to Jaco and Rubinstein (cf. Theorem 4.3), this is always possible.

The second question is more open-ended. As a first step, we prove the following.

Proposition 5.17. *There exists an infinite family of 3-manifolds of bounded cutwidth—hence of bounded treewidth—with unbounded Heegaard genus.*

Proof. Consider $\{\mathcal{M}_k = \mathcal{M}^{\#k} : k \in \mathbb{N}\}$, where $\mathcal{M}_k$ is the k-fold connected sum of a 3-manifold $\mathcal{M}$ of Heegaard genus one (e.g., $\mathcal{M} = \mathbb{S}^2 \times \mathbb{S}^1$). Since the Heegaard genus is additive under taking connected sums [69, Corollary II.10.], we have that $\mathfrak{g}(\mathcal{M}_k) = k$. However, $\mathcal{M}_k$ admits a triangulation of bounded cutwidth. Indeed, start with a fixed triangulation $\mathcal{T}$ of $\mathcal{M}$ that contains two tetrahedra Δ_1 and Δ_2 which a) do not share any vertices in $\mathcal{T}$, and b) do not have any self-identifications in $\mathcal{T}$. Let w denote the width of an ordering of $V(\Gamma(\mathcal{T}))$—the nodes of the dual graph $\Gamma(\mathcal{T})$—in which Δ_1 and Δ_2 correspond to the first and the last node, respectively. Moreover, let $\mathcal{T}^i$ $(1 \leq i \leq k)$ be k copies of $\mathcal{T}$. Forming connected sums along Δ_2^i and Δ_1^{i+1} $(1 \leq i \leq k-1)$ yields a triangulation $\mathcal{T}_k$ of $\mathcal{M}_k$ together with an ordering of $V(\Gamma(\mathcal{T}_k))$ of width w, see Figure 5.8. Therefore $\mathrm{cw}(\mathcal{M}_k) \leq \mathrm{cw}(\mathcal{T}_k) \leq w$ for every $k \in \mathbb{N}$. □

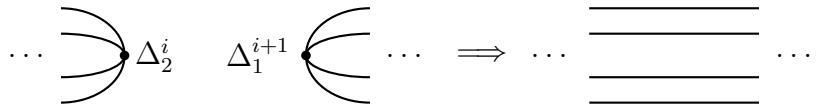

Figure 5.8: The effect of forming $\mathcal{T}^i \# \mathcal{T}^{i+1}$ at the level of the dual graphs

Remark 5.18. Proposition 5.17 shows that among reducible 3-manifolds one can easily find infinite families that violate (5.10). Nevertheless, irreducibility alone is insufficient for (5.10) to hold. In particular, in Chapter 7 we prove that orientable Seifert fibered spaces over $\mathbb{S}^2$ have treewidth at most two (Theorem 7.2). However, all but two of them are irreducible [134, Theorem 3.7.17] and they can have arbitrarily large Heegaard genus [21, Theorem 1.1].

Nevertheless, as mentioned before, the reverse inequality always holds.

Proof of Theorem 1.6. Let $g = \mathfrak{g}(\mathcal{M})$. Consider the g-spine $\mathcal{S}$ in Figure 4.9(i) together with the indicated order in which we layer onto the $3g-2$ interior edges of $\mathcal{S}$ to build two copies $\mathcal{T}'$ and $\mathcal{T}''$ of a minimal layered triangulation of the genus g handlebody. See Figure 4.9(ii) for the dual graph of $\mathcal{T}'$ (and of $\mathcal{T}''$). Note that $\partial\mathcal{T}'$ and $\partial\mathcal{T}''$ consist of $4g-2$ triangles each.

By Theorem 4.3, we may extend $\mathcal{T}'$ to a layered triangulation $\mathcal{T}'''$ which can be glued to $\mathcal{T}''$ along a simplicial map $f : \partial\mathcal{T}''' \to \partial\mathcal{T}''$ to yield a triangulation $\mathcal{T} = \mathcal{T}''' \cup_f \mathcal{T}''$ of $\mathcal{M}$. This construction imposes a natural ordering on the tetrahedra of $\mathcal{T}$: 1. Start by ordering the tetrahedra of $\mathcal{T}'$ according to the labels of the edges of $\mathcal{S}$ onto which they are initially layered. 2. Continue with all tetrahedra between $\mathcal{T}'$ and $\mathcal{T}''$ in the order they are attached to $\mathcal{T}'$ in order to build up $\mathcal{T}'''$. 3. Finish with the tetrahedra of $\mathcal{T}''$ again in the order of the labels of the edges of $\mathcal{S}$ onto which they are layered. This way we obtain a linear layout of the nodes of $\Gamma(\mathcal{T})$ which realizes width $4g-2$ (Figure 5.9). Therefore $\mathrm{cw}(\mathcal{M}) \leq 4g-2$. □

Combining Theorem 1.6 with $\mathrm{tw}(\mathcal{M}) \leq \mathrm{cw}(\mathcal{M})$, we directly deduce the following.

Corollary 5.19. *For any closed, orientable, irreducible, non-Haken 3-manifold $\mathcal{M}$ the Heegaard genus $\mathfrak{g}(\mathcal{M})$ and the treewidth $\mathrm{tw}(\mathcal{M})$ satisfy*

$$\frac{1}{4}(\mathrm{tw}(\mathcal{M})+2) \leq \mathfrak{g}(\mathcal{M}) \leq 18(\mathrm{tw}(\mathcal{M})+1). \tag{5.12}$$

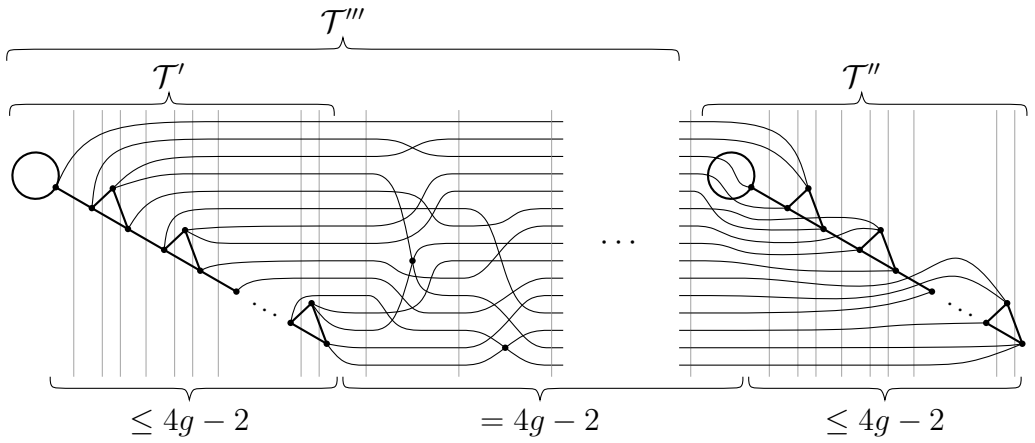

Figure 5.9: A linear layout showing that $\mathrm{cw}(\mathcal{M})$ is bounded above by $4\mathfrak{g}(\mathcal{M}) - 2$

Remark 5.20. Bachman et al. [5] ask (cf. Question 5.3 therein) whether computing the Heegaard genus of a 3-manifold is still hard when restricting to the set of non-Haken 3-manifolds. Corollary 5.19 implies that the answer to this question also has implications on the hardness of approximating the treewidth of non-Haken manifolds.

For an algorithmic consequence of Theorem 1.6 see Appendix D.

5.3.2 Layered Triangulations of Generalized Heegaard Splittings

In this section we prove the following qualitative strengthening of Theorem 1.6.

Theorem 1.7. *There exists a universal constant $C > 0$, such that, for any 3-manifold $\mathcal{M}$ (possibly with $\partial\mathcal{M} \neq \emptyset$), the cutwidth $\mathrm{cw}(\mathcal{M})$ and the linear width $\mathscr{L}(\mathcal{M})$ satisfy*

$$\mathrm{cw}(\mathcal{M}) \leq 24\mathscr{L}(\mathcal{M}). \tag{5.13}$$

At the heart of the proof lies the following lemma.

Lemma 5.21. *Let $\mathcal{C}$ be a compression body, where $\partial_-\mathcal{C}$ has no $\mathbb{S}^2$ components. Then $\mathcal{C}$ admits a triangulation $\mathcal{T}$ with at most $12g(\partial_+\mathcal{C}) - 6$ tetrahedra that induces a minimal (one-vertex) triangulation on each boundary surface.*

Proof. We may assume $\mathcal{C}$ to be connected. Then the disconnected case follows by arguing component-wise. We will construct the desired triangulation $\mathcal{T}$ by performing an adapted version of the *primal construction* (cf. Section 3.4 and the upper part of Figure 3.7) and working with triangulated building blocks. Recall that, via the primal construction, the compression body $\mathcal{C}$ with lower boundary $\partial_-\mathcal{C} = \mathcal{S}$ is simply obtained as follows:

P1. Take the thickened surface $\mathcal{S} \times [0,1]$. **P2.** Attach some 1-handles to $\mathcal{S} \times \{1\}$.

Note that if $\mathcal{S} = \mathcal{S}_1 \cup \ldots \cup \mathcal{S}_m$, where the $\mathcal{S}_i$ are the connected components of $\mathcal{S}$, and $\gamma = g(\partial_+\mathcal{C}) - g(\partial_-\mathcal{C})$, then, by slightly modifying the primal construction, $\mathcal{C}$ can be built by first thickening each $\mathcal{S}_i$ $(1 \leq i \leq m)$, then connecting these thickened surfaces with $m-1$ one-handles called *connectors*, and finally appending a handlebody $\mathcal{H}_\gamma$ of genus γ, again via a connector. See Figure 5.10 for an example, where $m = 2$ and $\gamma = 2$.

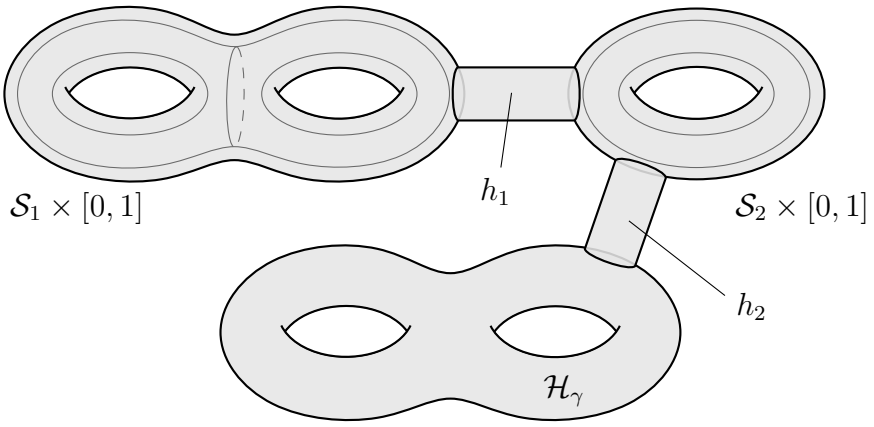

Figure 5.10: A compression body $\mathcal{C} = (\mathcal{S}_1 \times [0,1]) \cup h_1 \cup (\mathcal{S}_2 \times [0,1]) \cup h_2 \cup \mathcal{H}_\gamma$

We now construct triangulations of these three building blocks (i.e., thickened surfaces, connectors, and handlebodies), which can then be assembled into a triangulation $\mathcal{T}$ of $\mathcal{C}$.

Thickened surfaces. This construction is illustrated in Figure 5.11 for the case $g = 2$. Let $\mathcal{F}_g$ denote the closed orientable surface of genus g. It is well-known that $\mathcal{F}_g$ can be obtained from a regular polygon with $4g$ vertices by identifying its edges in pairs. Specifically, let P be a $4g$-gon with vertices $v_1^{(1)}, v_2^{(1)}, v_3^{(1)}, v_4^{(1)}, \ldots, v_1^{(g)}, v_2^{(g)}, v_3^{(g)}, v_4^{(g)}$ in this circular order. Triangulate P by adding the diagonals incident to $v_1^{(1)}$. Performing the identifications $(v_1^{(i)}, v_2^{(i)}) \leftrightarrow (v_4^{(i)}, v_3^{(i)})$, and $(v_2^{(i)}, v_3^{(i)}) \leftrightarrow (v_1^{(i+1)}, v_4^{(i)})$, for $i \in \{1, \ldots, g\}$, where $v_1^{(g+1)} = v_1^{(1)}$, gives a minimal (one-vertex) triangulation of $\mathcal{F}_g$ with $4g-2$ triangles. We refer to Figure 5.11(i) for an example.

This triangulation of $\mathcal{F}_g$ with $4g-2$ triangles naturally lends itself to a decomposition of $\mathcal{F}_g \times [0,1]$ into $4g-2$ triangular prisms glued together along pairs of quadrilaterals of the form $e \times [0,1]$, where e is an edge in the triangulation of $\mathcal{F}_g$. Imagine, for a moment, that we cut up this shape along those quadrilaterals that correspond to the boundary edges of P identified in the previous paragraph. This way we recover a thickened version $\boldsymbol{P}$ of the $4g$-gon P. Let v_1 denote a distinguished vertex $v_1^{(1)} \times \{1\}$, cf. Figure 5.11(ii).

Next we subdivide every triangular prism in $\boldsymbol{P}$ into three tetrahedra each, to get a triangulation of $\boldsymbol{P}$. We do this in two steps. First, we split each quadrilateral in $\partial\boldsymbol{P}$ into two triangles by adding a diagonal, in a way that is compatible with the identifications above. Moreover, in the two quadrilaterals at v_1 we add those diagonals that are incident to v_1. This gives a triangulation of $\partial\boldsymbol{P}$, see Figure 5.11(iii). Now a triangulation of $\boldsymbol{P}$ is obtained by coning over the vertex v_1, cf. Figure 5.11(iv).

By redoing the identifications at the triangulated quadrilaterals, we get a triangulation of $\mathcal{F}_g \times [0,1]$ with exactly $3 \cdot (4g-2) = 12g-6$ tetrahedra.[4]

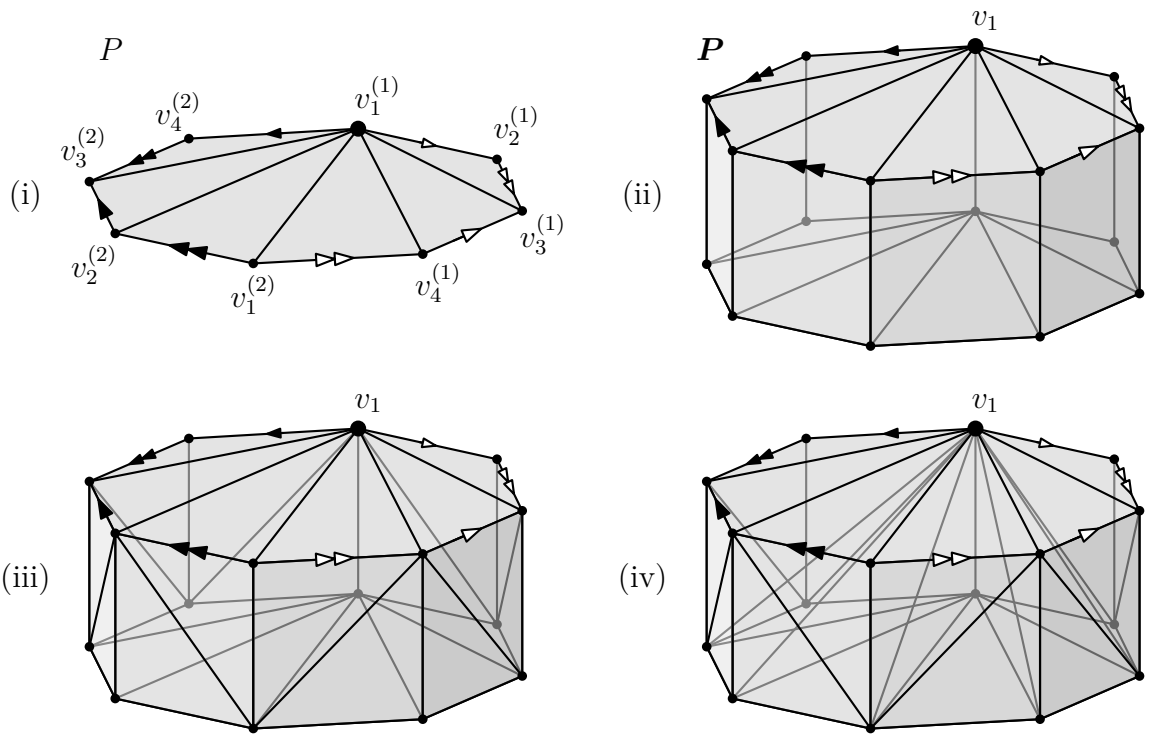

Figure 5.11: Triangulating the compression body $\mathcal{F}_g \times [0,1]$ with $12g-6$ tetrahedra

Connectors. We triangulate a connector as a *triangular bipyramid*, which is the union of two tetrahedra sharing one triangular face (Figure 5.12). Its five vertices are labeled with $0, 1, 2, 4$ and 5. In our example, $\{0,1,3\}$ is the common face of the tetrahedra $\{0,1,2,3\}$ and $\{0,1,3,4\}$. We distinguish two boundary triangles $t = \{0,1,2\}$ and $t' = \{0,3,4\}$, which will act as attaching sites.

[4]It is worth noting that these triangulations are not minimal. In [70] it was shown that $\mathcal{F}_g \times [0,1]$ can always be triangulated with $10g-4$ tetrahedra. Moreover, this bound is sharp, cf. [70, Theorem 25].

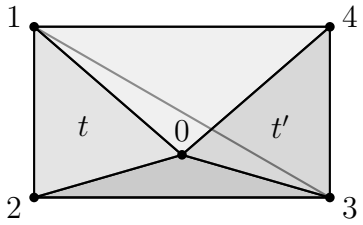

Figure 5.12: A connector is a triangular bipyramid with two attaching sites t and t'

Handlebodies. Here we will work with minimal layered triangulations (Section 4.2.1). For the handlebody $\mathcal{H}_g$ of genus g such a triangulation $\mathcal{T}(\mathcal{H}_g)$ consists of $3g-2$ tetrahedra and induces a minimal (1-vertex) triangulation on the boundary $\partial\mathcal{H}_g = \mathcal{F}_g$.

Having discussed the building blocks, we can now conclude the proof of Lemma 5.21. Let $\mathcal{C}$ be a connected compression body with $\mathcal{S}_1, \ldots, \mathcal{S}_m$ being the connected components of its lower boundary $\partial_-\mathcal{C}$. Furthermore, let $\gamma = g(\partial_+\mathcal{C}) - g(\partial_-\mathcal{C})$.

For $1 \le i \le m$, consider the thickened surfaces $\mathcal{C}_i = \mathcal{S}_i \times [0,1]$ and endow each of them with a triangulation $\mathcal{T}(\mathcal{C}_i)$ as described above. Then $|\mathcal{T}(\mathcal{C}_i)| = 12g(\mathcal{S}_i) - 6$, and at both boundary components $\partial_-\mathcal{C}_i = \mathcal{S}_i \times \{0\}$ and $\partial_+\mathcal{C}_i = \mathcal{S}_i \times \{1\}$ the induced triangulations, respectively denoted as $\partial_-\mathcal{T}(\mathcal{C}_i)$ and $\partial_+\mathcal{T}(\mathcal{C}_i)$, are minimal triangulations with $4g(\mathcal{S}_i) - 2$ triangles and one vertex each. We now formalize the construction shown in Figure 5.10.

Set $\mathcal{T}^{(1)} = \mathcal{T}(\mathcal{C}_1)$, $\partial_+\mathcal{T}^{(1)} = \partial_+\mathcal{T}(\mathcal{C}_1)$ and $\partial_-\mathcal{T}^{(1)} = \partial_-\mathcal{T}(\mathcal{C}_1)$. For $i \in \{2, \ldots, m\}$, iteratively define $\mathcal{T}^{(i)} = \mathcal{T}^{(i-1)} \cup_\varphi h \cup_\psi \mathcal{T}(\mathcal{C}_i)$, where h is a connector with distinguished boundary triangles t and t' sharing a single vertex (cf. Figure 5.12), φ identifies t with a triangle on $\partial_+\mathcal{T}^{(i-1)}$, and ψ identifies t' with a triangle on $\partial_+\mathcal{T}(\mathcal{C}_i)$.

Note that, at each step, $\partial_+\mathcal{T}^{(i)}$ is maintained to be a 1-vertex triangulation (of a closed orientable surface of genus $\sum_{j=1}^{i} g(\mathcal{S}_j)$), and so is each component of $\partial_-\mathcal{T}^{(i)}$.

After m iterations, we have connected the triangulations $\mathcal{T}(\mathcal{C}_i)$ of all thickened surfaces $\mathcal{C}_i = \mathcal{S}_i \times [0,1]$ $(1 \le i \le m)$. For the last step, we take $\mathcal{T} = \mathcal{T}^{(m)} \cup_\varphi h \cup_\psi \mathcal{H}_\gamma$, where φ identifies t with a triangle on $\partial_+\mathcal{T}^{(m)}$, and ψ identifies t' with a triangle on $\partial\mathcal{T}(\mathcal{H}_\gamma)$. Thus we have constructed a triangulation $\mathcal{T}$ of $\mathcal{C}$. For the size $|\mathcal{T}|$ we have

$$\begin{aligned}
|\mathcal{T}| &\le \left(\sum_{i=1}^{m} |\mathcal{T}(\mathcal{C}_i)|\right) + 2m + (3\gamma - 2) \\
&= \left(\sum_{i=1}^{m} 12g(\mathcal{S}_i) - 6\right) + 2m + (3\gamma - 2) \\
&= 12g(\partial_-\mathcal{C}) - 4m + 3\left(g(\partial_+\mathcal{C}) - g(\partial_-\mathcal{C})\right) - 2 \le 12g(\partial_+\mathcal{C}) - 6.
\end{aligned}$$

□

Proof of Theorem 1.7. Suppose that $\mathcal{M} \notin \{\mathbb{B}^3, \mathbb{S}^2 \times \mathbb{S}^1\}$. Assume first that $\mathcal{M}$ is prime. Consider a linear generalized Heegaard splitting $\mathcal{M} = (\mathcal{N}_1 \cup_{\mathcal{S}_1} \mathcal{K}_1) \cup_{\mathcal{R}_1} \cdots \cup_{\mathcal{R}_{r-1}} (\mathcal{N}_r \cup_{\mathcal{S}_r} \mathcal{K}_r)$ in thin position.[5] Being in thin position has two consequences for the splitting that are important for us:

1. Every 2-sphere component of every lower boundary $\mathcal{R}_i$ is *essential*, i.e., it does not bound a 3-ball, see [130, Observation 5.2.1].[6] Therefore, since $\mathcal{M}$ is irreducible,[7] none of the lower boundary surfaces $\mathcal{R}_i$ can have a 2-sphere component.

2. $\max\{g(\mathcal{S}_i) : 1 \leq i \leq r\} = \mathscr{L}(\mathcal{M})$.

The first property precisely means that the compression bodies $\mathcal{N}_i$ and $\mathcal{K}_i$ $(1 \leq i \leq r)$ satisfy the condition of Lemma 5.21. The second one then implies that each of them can be constructed using at most $12\mathscr{L}(\mathcal{M}) - 6$ tetrahedra, so that we have 1-vertex triangulations at their boundary components.

Now the gluings $\partial_+\mathcal{N}_i \leftrightarrow \partial_+\mathcal{K}_i$ $(1 \leq i \leq r)$ and $\partial_-\mathcal{K}_i \leftrightarrow \partial_-\mathcal{N}_{i+1}$ $(1 \leq i \leq r-1)$ can be realized via layerings (Section 4.2), analogously to the procedure in the proof of Theorem 1.6. In particular, the dual graph of the resulting triangulation admits a linear layout as shown in Figure 5.13 (cf. Figure 5.9).

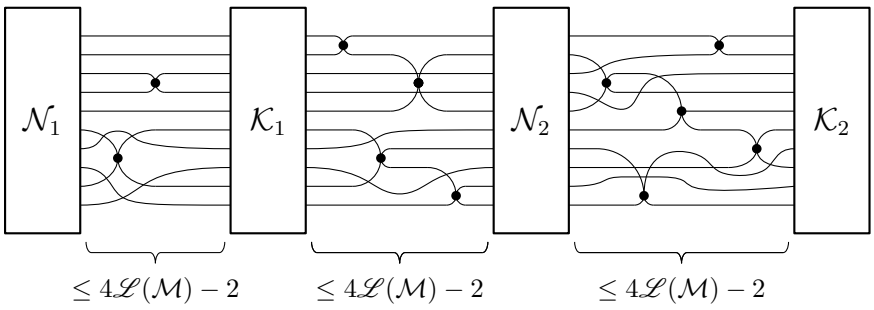

Figure 5.13: Layout of the dual graph of a layered triangulation of a linear splitting. The size of any cutset "between" the compression bodies is at most $4\mathscr{L}(\mathcal{M}) - 2$

Note that, for every $i \in \{1, \ldots, r\}$, the dual graphs of the triangulations of $\mathcal{N}_i$ and $\mathcal{K}_i$ contain at most $24\mathscr{L}(\mathcal{M}) - 12$ arcs each, since they have at most $12\mathscr{L}(\mathcal{M}) - 6$ nodes with maximum degree at most 4. This readily implies that Theorem 1.7 holds in case the 3-manifold $\mathcal{M}$ is prime.

---

[5] See page 41 for details regarding thin position.

[6] This is not true for the genus zero Heegaard splitting of $\mathbb{B}^3$, but we have excluded this case.

[7] The only orientable, reducible prime 3-manifold is $\mathbb{S}^2 \times \mathbb{S}^1$, but this case has been excluded as well.

In the general case, consider the prime decomposition $\mathcal{M} = \mathcal{M}_1 \# \cdots \# \mathcal{M}_r$. It can be shown that the cutwidth $\mathrm{cw}(\mathcal{M})$ and the linear width $\mathscr{L}(\mathcal{M})$ satisfy

$$\mathrm{cw}(\mathcal{M}) = \max\{\mathrm{cw}(\mathcal{M}_i) : 1 \leq i \leq r\}, \quad \text{and} \tag{5.14}$$

$$\mathscr{L}(\mathcal{M}) = \max\{\mathscr{L}(\mathcal{M}_i) : 1 \leq i \leq r\}. \tag{5.15}$$

The proof of (5.14) is the same as the proof of Proposition 5.17 (cf. Figure 5.8), while (5.15) follows from the "additive" property of the original definition of linear width by Scharlemann and Thompson [131, p. 233]. Having these formulas at hand, the general inequality $\mathrm{cw}(\mathcal{M}) \leq 24\mathscr{L}(\mathcal{M})$ immediately follows. □

Theorem 1.7, in combination with other results cited or developed in this thesis, has several consequences. For example, Theorems 5.13 and 1.7 readily imply the following.

Corollary 5.22. *For every closed, orientable* 3*-manifold* $\mathcal{M}$ *with cutwidth* $\mathrm{cw}(\mathcal{M})$ *and linear width* $\mathscr{L}(\mathcal{M})$ *we have*

$$\frac{1}{3}(\mathscr{L}(\mathcal{M}) - 4) \leq \mathrm{cw}(\mathcal{M}) \leq 24\mathscr{L}(\mathcal{M}). \tag{5.16}$$

By (5.2) the pathwidth $\mathrm{pw}(\mathcal{M})$ is also within a constant factor of $\mathrm{cw}(\mathcal{M})$ and $\mathscr{L}(\mathcal{M})$. Furthermore, blending in Theorem 5.11(1) by Hoffoss–Maher, we deduce

Corollary 5.23. *For closed, orientable,* ***hyperbolic*** 3*-manifolds, the cutwidth* $\mathrm{cw}(\mathcal{M})$*, pathwidth* $\mathrm{pw}(\mathcal{M})$*, linear width* $\mathscr{L}(\mathcal{M})$*, and the Morse area* $\mathscr{A}_{\mathrm{M}}(\mathcal{M})$ *are all within a constant factor of each other.*

Remark 5.24. The technique used in the proof of Theorem 1.7 can readily be used to triangulate graph splittings as well. We strongly believe that, for 3-manifolds with bounded first Betti number (in particular for rational homology 3-spheres), the congestion $\mathrm{cng}(\mathcal{M})$, the treewidth $\mathrm{tw}(\mathcal{M})$, the graph width $\mathscr{G}(\mathcal{M})$ and—in case $\mathcal{M}$ is hyperbolic—the Gromov area $\mathscr{A}_{\mathrm{G}}(\mathcal{M})$ are within a constant factor of each other.

5.4 On the Width of Hyperbolic 3-Manifolds

In Sections 5.2 and 5.3 we studied the quantitative relationship of combinatorial and topological width parameters of (closed) 3-manifolds in general. Here we restrict our attention to hyperbolic 3-manifolds that possess a rigid geometry (Theorem 5.6). We have already seen that, for such manifolds, some geometric and topological width parameters are linearly related (Theorem 5.11). In this section we improve upon a recent result by Maria and Purcell that relates the treewidth and the volume of hyperbolic 3-manifolds.

Theorem 5.25 (Maria–Purcell [99])**.** *There exists a universal constant $C > 0$ such that, for any closed hyperbolic 3-manifold $\mathcal{M}$ with treewidth* $\mathrm{tw}(\mathcal{M})$ *and volume* $\mathrm{vol}(\mathcal{M})$*,*

$$\mathrm{tw}(\mathcal{M}) \leq C \cdot \mathrm{vol}(\mathcal{M}). \tag{5.17}$$

Equivalently, using the 'big O notation,' this can be written as $\mathrm{tw}(\mathcal{M}) = O(\mathrm{vol}(\mathcal{M}))$*.*

The main steps of the proof of Theorem 5.25 can informally be summarized as follows:

1. Start with a *thick-thin decomposition* $\mathcal{M} = \mathcal{M}_{\text{thick}} \cup \mathcal{M}_{\text{thin}}$. The fact that hyperbolic manifolds of finite volume (in any dimension, even if non-compact[8]) admit such a decomposition is one the most important structure theorems in geometric topology. Here $\mathcal{M}_{\text{thick}} \cap \mathcal{M}_{\text{thin}}$ is a disjoint union of tori, and $\mathcal{M}_{\text{thin}}$ is that of solid tori.[9]
2. Choose a triangulation $\mathcal{T}_{\text{thick}}$ of the thick part $\mathcal{M}_{\text{thick}}$ with $O(\mathrm{vol}(\mathcal{M}))$ tetrahedra. Due to Jørgensen and Thurston (Theorem 5.30), this is always possible. Trivially, the dual graph of $\mathcal{T}_{\text{thick}}$ satisfies $\mathrm{tw}(\Gamma(\mathcal{T}_{\text{thick}})) = O(\mathrm{vol}(\mathcal{M}))$.
3. Perform an operation called *crushing* on $\mathcal{T}_{\text{thick}}$ to get another, "simpler" triangulation $\widehat{\mathcal{T}}_{\text{thick}}$ of $\mathcal{M}_{\text{thick}}$, where each component of $\partial\widehat{\mathcal{T}}_{\text{thick}}$ is a one-vertex triangulation of $\mathbb{T}^2$. Maria and Purcell show that $\mathrm{tw}\left(\Gamma(\widehat{\mathcal{T}}_{\text{thick}})\right) = O(\mathrm{vol}(\mathcal{M}))$ is still true.
4. Finally, extend $\widehat{\mathcal{T}}_{\text{thick}}$ to a triangulation $\mathcal{T}$ of $\mathcal{M}$ by attaching a layered solid torus to each component of $\partial\widehat{\mathcal{T}}_{\text{thick}}$. As layered solid tori have very simple triangulations, cf. Figure 4.10(v), this final step also maintains the treewidth to be $O(\mathrm{vol}(\mathcal{M}))$.

Here we strengthen Theorem 5.25 by showing that the volume of a closed hyperbolic 3-manifolds provides a linear upper bound even on its pathwidth.[10]

[8] Typical examples of finite-volume non-compact hyperbolic 3-manifolds are those with *cusps*.

[9] This is so, because $\mathcal{M}$ is compact. In general $\mathcal{M}_{\text{thin}}$ also contains the cusps of $\mathcal{M}$, cf. Theorem 5.28.

[10] It would be interesting to understand the extent of the improvement, in particular, whether the difference $\mathrm{pw}(\mathcal{M}) - \mathrm{tw}(\mathcal{M})$ can be arbitrary large.

Theorem 1.8. There exists a universal constant $C' > 0$ such that, for any closed hyperbolic 3-manifold $\mathcal{M}$ with pathwidth $\mathrm{pw}(\mathcal{M})$ and volume $\mathrm{vol}(\mathcal{M})$, we have

$$\mathrm{pw}(\mathcal{M}) \leq C' \cdot \mathrm{vol}(\mathcal{M}). \tag{5.18}$$

The proof of Theorem 1.8 rests on the following result, known to experts in the field (see, e.g., [139, p. 336–337] or [122]), according to which the Heegaard genus of a closed hyperbolic 3-manifold can be linearly upper-bounded in terms of its volume.

Theorem 5.26. *There exists a universal constant $C'' > 0$ such that, for any closed hyperbolic 3-manifold $\mathcal{M}$ with Heegaard genus $\mathfrak{g}(\mathcal{M})$ and volume $\mathrm{vol}(\mathcal{M})$, we have*

$$\mathfrak{g}(\mathcal{M}) \leq C'' \cdot \mathrm{vol}(\mathcal{M}). \tag{5.19}$$

Proof of Theorem 1.8 assuming Theorem 5.26. By combining Theorem 5.26 with the estimate of Theorem 1.6 (i.e., $\mathrm{cw}(\mathcal{M}) \leq 4\mathfrak{g}(\mathcal{M}) - 2$) and the left-hand-side inequality in Theorem 2.4 (i.e., $\mathrm{pw}(\mathcal{M}) \leq \mathrm{cw}(\mathcal{M})$), the claim of Theorem 1.8 is readily deduced. □

Next, we supply a proof of Theorem 5.26 that is likely to be a roundabout way to establish this result, nevertheless reflects the author's understanding of it. The proof is inspired by that of Theorem 5.25. The improvement comes from the observation that a triangulation of the thick part with $O(\mathrm{vol}(\mathcal{M}))$ tetrahedra lends itself to a generalized Heegaard splitting of $\mathcal{M}$, where the sum of genera of the Heegaard surfaces is $O(\mathrm{vol}(\mathcal{M}))$. This generalized Heegaard splitting can then be "amalgamated" into a classical Heegaard splitting, which is still of genus $O(\mathrm{vol}(\mathcal{M}))$. We now elaborate on the proof ingredients.

Thick-thin decompositions. The first ingredient is the same as in the proof Theorem 5.25 by Maria–Purcell: the thick-thin decomposition theorem, a fundamentally important structural result for hyperbolic manifolds. In order to formulate it, we need to introduce some notions about Riemannian manifolds.

Definition 5.27 (injectivity radius)**.** Let $\mathcal{M}$ be a Riemannian manifold and $x \in \mathcal{M}$. The *injectivity radius of $\mathcal{M}$ at x*, denoted $\mathrm{inj}_x(\mathcal{M})$, is the supremal value $r > 0$ such that the metric ball of radius r around x is embedded in $\mathcal{M}$. The *injectivity radius of $\mathcal{M}$* is defined as the infimal value of $\mathrm{inj}_x(\mathcal{M})$, i.e., $\mathrm{inj}(\mathcal{M}) = \inf\{\mathrm{inj}_x(\mathcal{M}) : x \in \mathcal{M}\}$.

After fixing some threshold $\varepsilon > 0$, a Riemannian manifold $\mathcal{M}$ naturally decomposes into an ε-*thick* and an ε-*thin* part based on the injectivity radius of its points:

$$\mathcal{M}_{[\varepsilon,\infty)} = \{x \in \mathcal{M} : \mathrm{inj}_x(\mathcal{M}) \geq \varepsilon/2\} \quad \text{and} \quad \mathcal{M}_{(0,\varepsilon]} = \{x \in \mathcal{M} : \mathrm{inj}_x(\mathcal{M}) \leq \varepsilon/2\}. \tag{5.20}$$

We can now state the Thick-Thin Decomposition Theorem according to which, for a sufficiently small constant $\varepsilon > 0$ only depending on the dimension d, the ε-thin part of any orientable hyperbolic d-manifold has a well-understood structure.[11]

Theorem 5.28 (Thick-Thin Decomposition; cf. [101, Chapter 4], [121, Section 5.3])**.** *There exists a universal constant $\varepsilon_d > 0$, depending only on the dimension d, such that for any $\varepsilon \in (0, \varepsilon_d]$, the ε-thin part of any orientable hyperbolic d-manifold $\mathcal{M}$ consists of tubes around short geodesics diffeomorphic to $\mathbb{S}^1 \times \mathbb{D}^{d-1}$, or cusps.*[12]

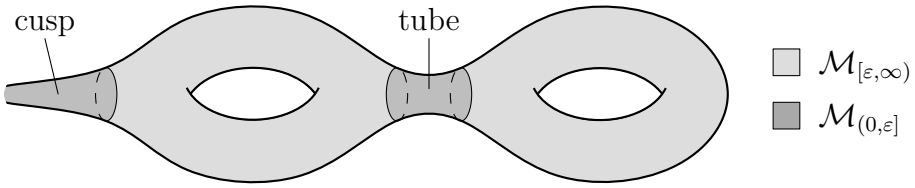

Figure 5.14: Thick-thin decomposition of a non-compact hyperbolic surface

Remark 5.29.

(1) In case of compact 3-manifolds, there are no cusps in the thick-thin decomposition, but only tubes. In dimension three, they are homeomorphic to solid tori. This is important, as Theorem 5.26 is concerned with closed (hence compact) 3-manifolds.

(2) The supremum of all ε_d for which the conclusion of Theorem 5.28 holds is called the d-dimensional *Margulis constant.* As of now, the precise value of ε_d remains unknown. For $d = 3$, it is known that $0.104 \leq \varepsilon_3 \leq 0.616$, cf. [121, p. 92].

(3) Theorem 5.28 is a corollary of a more general result about discrete subgroups of Lie groups, called the Margulis Lemma, which appeared in [76], cf. [101, Section 4.2] and [121, Theorem 5.22].

[11] The manifolds in consideration are also required to be *complete* (as metric spaces). However, the way we define hyperbolic d-manifolds (i.e., quotients of $\mathbb{H}^d$ under discrete groups of isometries acting freely) automatically ensures their completeness.

[12] A d-dimensional cusp is a d-manifold with boundary that is diffeomorphic to $\mathcal{N} \times [0, \infty)$, where $\mathcal{N}$ is a $(d-1)$-dimensional flat, i.e., Euclidean, manifold. See [101, Section 4.1] for a precise definition.

Triangulating the thick part. In Chapter 4 we discussed generalized Heegaard splittings and layered triangulations as "interfaces between combinatorics and topology." Here we invoke the aforementioned theorem by Jørgensen and Thurston (carefully proved by Kobayashi and Rieck), which can be regarded as an interface between combinatorics and geometry in the realm of hyperbolic 3-manifolds. To precisely state this result, let us define the *(closed) δ-neighborhood* $\overline{N_\delta}(\mathcal{X})$ of a subset $\mathcal{X}$ of a Riemannian 3-manifold $\mathcal{M}$ to be the set of those points in $\mathcal{M}$ that have distance at most δ from some point in $\mathcal{X}$.

Theorem 5.30 (Jørgensen–Thurston [148, §5.11], Kobayashi–Rieck [84])**.** *Let $\varepsilon \in (0, \varepsilon_3]$, where ε_3 is the Margulis constant in dimension three.*

(1) *For any $\delta > 0$ there exists a constant $K > 0$, depending on ε and δ, so that for any finite-volume hyperbolic 3-manifold $\mathcal{M}$, the δ-neighborhood $\overline{N_\delta}\left(\mathcal{M}_{[\varepsilon,\infty)}\right) \subset \mathcal{M}$ of the thick part $\mathcal{M}_{[\varepsilon,\infty)}$ admits a triangulation with at most $K \cdot \mathrm{vol}(\mathcal{M})$ tetrahedra.*

(2) *Moreover, $\overline{N_\delta}\left(\mathcal{M}_{[\varepsilon,\infty)}\right)$ is obtained from $\mathcal{M}$ by removing open tubular neighborhoods around short geodesics, and truncating cusps* [84, Proposition 1.2].

We now fix $\varepsilon \in (0, \varepsilon_3]$ and $\delta > 0$. In the proof of Theorem 5.25, Theorem 5.30 plays the crucial role in ensuring that the treewidth of $\overline{N_\delta}\left(\mathcal{M}_{[\varepsilon,\infty)}\right)$ is upper-bounded by a linear function of the volume of $\mathcal{M}$, and that $\partial\overline{N_\delta}\left(\mathcal{M}_{[\varepsilon,\infty)}\right)$ can be filled with solid tori.

For proving Theorem 5.26, we utilize Theorem 5.30 in a slightly different way:

Corollary 5.31. *Let $\mathcal{Y} = \overline{N_\delta}\left(\mathcal{M}_{[\varepsilon,\infty)}\right)$ as in Theorem 5.30. The following are true.*

(1) *For the Heegaard genus of $\mathcal{Y}$ we have $\mathfrak{g}(\mathcal{Y}) = O(\mathrm{vol}(\mathcal{M}))$.*

(2) *$\mathcal{Y}$ has $O(\mathrm{vol}(\mathcal{M}))$ boundary components, each of which are tori.*

Proof. To establish (1), first consider a triangulation $\mathcal{T}$ of $\mathcal{Y}$ with $O(\mathrm{vol}(\mathcal{M}))$ tetrahedra. Such a triangulation is guaranteed to exist by Theorem 5.30(1). Fix an arbitrary partition $\mathscr{P} = \{\partial_1\mathcal{Y}, \partial_2\mathcal{Y}\}$ of the boundary components of $\mathcal{Y}$ (the trivial partition, i.e., $\partial_1\mathcal{Y} = \emptyset$, $\partial_2\mathcal{Y} = \partial\mathcal{Y}$, is also allowed). Following the procedure described in Exampe 3.10, we obtain a Heegaard splitting of $\mathcal{Y}$ compatible with $\mathscr{P}$. By construction, the genus of this splitting is $O(\mathrm{vol}(\mathcal{M}))$, hence, for the Heegaard genus of $\mathcal{Y}$, we have $\mathfrak{g}(\mathcal{Y}) = O(\mathrm{vol}(\mathcal{M}))$.

For the first part of (2) observe, that a tetrahedron can contribute triangles to at most one boundary component. The second part of (2) follows from Theorem 5.30(2). □

Generalized Heegaard splittings. As we saw in Section 4.1.2, any decomposition $\mathcal{M} = \bigcup_{i\in I} \mathcal{M}_i$ of a 3-manifold into codimension zero submanifolds with pairwise disjoint interiors gives rise to generalized Heegaard splittings of $\mathcal{M}$. Here we construct a graph splitting from a thick-thin decomposition of the hyperbolic 3-manifold $\mathcal{M}$. Let

$$\mathscr{D} = \{\mathcal{M}_i : i \in [m],\ \textstyle\bigcup_{i=1}^m \mathcal{M}_i = \mathcal{M},\ \text{and}\ \operatorname{int}(\mathcal{M}_i) \cap \operatorname{int}(\mathcal{M}_j) = \emptyset \text{ for } i \neq j\} \tag{5.21}$$

be a thick-thin decomposition of $\mathcal{M}$, where $\mathcal{M}_1 = \mathcal{Y} = \overline{N_\delta}\left(\mathcal{M}_{[\varepsilon,\infty)}\right)$ is the thick part and $\mathcal{M}_2 \ldots, \mathcal{M}_m$ are the thin ones, see, e.g., Figure 5.15(i).

Note that, by Theorem 5.30(2), each $\mathcal{M}_i$ $(2 \leq i \leq m)$ is homeomorphic to a solid torus $\mathbb{S}^1 \times \mathbb{D}^2$, and, by Corollary 5.31(2), $m = O(\mathrm{vol}(\mathcal{M}))$. Let us label the nodes of $\Gamma(\mathscr{D})$ via the identity map (Figure 5.15(ii)). For each $i \in [m]$, choose a Heegaard splitting $\mathcal{M}_i = \mathcal{N}_i \cup_{\mathcal{S}_i} \mathcal{K}_i$ of minimal genus compatible with this labeling (Figure 5.15(iii)).

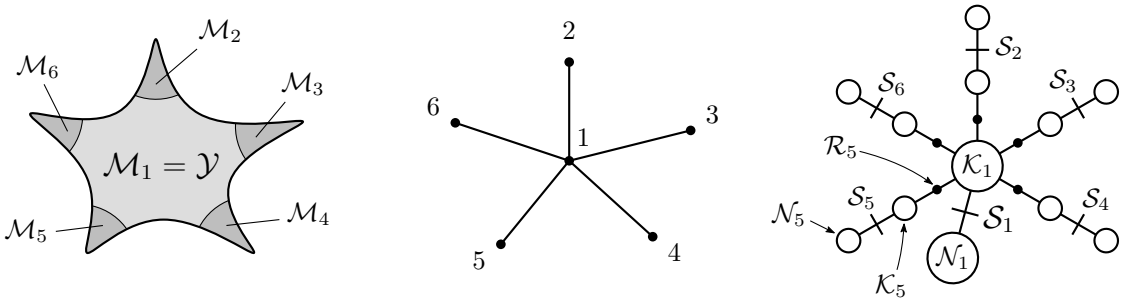

(i) $\mathscr{D}$ (ii) $\Gamma(\mathscr{D})$ (iii) $\boldsymbol{F}$

Figure 5.15: (i) Schematic of a thick-thin decomposition $\mathscr{D}$ of a hyperbolic 3-manifold $\mathcal{M}$. (ii) The dual graph $\Gamma(\mathscr{D})$ of $\mathscr{D}$ with its nodes labeled. (iii) The fork complex $\boldsymbol{F}$ of a generalized Heegaard splitting associated with $\mathscr{D}$ and the given labeling of $V(\Gamma(\mathscr{D}))$

Proposition 5.32. *It follows directly from the construction that the Heegaard splittings $\mathcal{M}_i = \mathcal{N}_i \cup_{\mathcal{S}_i} \mathcal{K}_i$ above possess the following properties:*

(1) *All the $\mathcal{N}_i$ are handlebodies. For $\mathcal{N}_1$ we have $g(\partial\mathcal{N}_1) = g(\mathcal{S}_1) = O(\mathrm{vol}(\mathcal{M}))$.*

(2) *If $2 \leq i \leq m$, then $\mathcal{N}_i$ is a solid torus, therefore $g(\partial\mathcal{N}_i) = g(\mathcal{S}_i) = 1$.*

(3) *$\mathcal{K}_1$ is a compression body with $\partial_+\mathcal{K}_1 = \mathcal{S}_1$ and $\partial_-\mathcal{K}_1 =$ "disjoint union of m tori."*

(4) *If $2 \leq i \leq m$, then $\mathcal{K}_i$ is a trivial compression body homeomorphic to $\mathbb{T}^2 \times [0,1]$. For its boundary components we have $\partial_+\mathcal{K}_i = \mathcal{S}_i$ and $\partial_-\mathcal{K}_i = \mathcal{R}_i = \mathcal{M}_i \cap \mathcal{M}_1$.*

(5) *We note that $\sum_{i=1}^m g(\mathcal{S}_i) = O(\mathrm{vol}(\mathcal{M}))$.*

Amalgamations. Introduced by Schultens in [133], *amalgamation* is a useful procedure that turns a generalized Heegaard splitting into a classical one. There are several excellent references where amalgamations are discussed in detail (cf. [5, Section 2], [43, Section 2.3], [130, Section 5.4]), therefore we rely on a simple example to illustrate this operation.

Let $\mathcal{M} = (\mathcal{N}_1 \cup_{\mathcal{S}_1} \mathcal{K}_1) \cup_{\mathcal{R}} (\mathcal{N}_2 \cup_{\mathcal{S}_2} \mathcal{K}_2)$ be a linear splitting of $\mathcal{M}$, which we would like to amalgamate to form a classical Heegaard splitting $\mathcal{M} = \mathcal{N} \cup_{\mathcal{S}} \mathcal{K}$, see Figure 5.16. Recall that every compression body $\mathcal{C}$ can be obtained by first taking the thickened version $\partial_- \mathcal{C} \times [0,1]$ of its lower boundary $\partial_- \mathcal{C}$ and then attaching some 1-handles to $\partial_- \mathcal{C} \times \{1\}$ (see steps P1 and P2 in Figure 3.7). In our example $\partial_- \mathcal{K}_1 = \mathcal{R} = \partial_- \mathcal{N}_2$, so $\mathcal{K}_1$ can be built from $\mathcal{R} \times [-1,0]$ by attaching two 1-handles $h_1^{(1)}$ and $h_2^{(1)}$ along $\mathcal{R} \times \{-1\}$. Similarly, $\mathcal{N}_2$ is constructed by taking $\mathcal{R} \times [0,1]$ and attaching the 1-handles $h_1^{(2)}$ and $h_2^{(2)}$ to $\mathcal{R} \times \{1\}$.

The *amalgamation* process consists of two steps: **1.** Collapse $\mathcal{R} \times [-1,1]$ to $\mathcal{R} \times \{0\}$, such that the attaching sites of the 1-handles $h_1^{(1)}$, $h_2^{(1)}$, $h_1^{(2)}$ and $h_2^{(2)}$ remain pairwise disjoint. (This can be achieved by slightly deforming the attaching maps, if necessary.) **2.** Set $\mathcal{N} = \mathcal{N}_1 \cup h_1^{(2)} \cup h_2^{(2)}$ and $\mathcal{K} = \mathcal{K}_2 \cup h_1^{(1)} \cup h_2^{(1)}$, see Figure 5.16(ii).

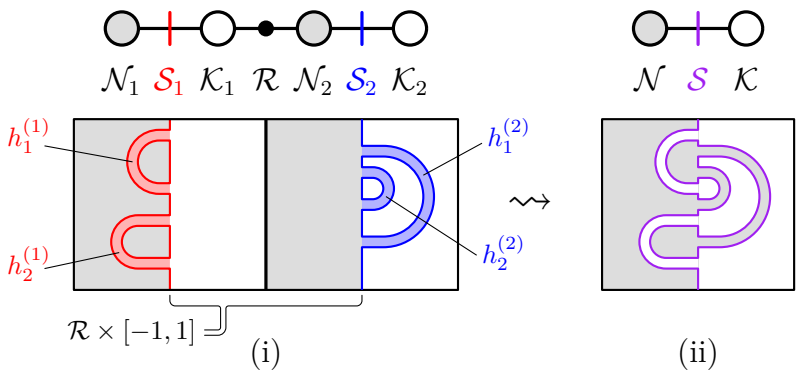

Figure 5.16: Amalgamating a generalized Heegaard splitting into a Heegaard splitting

If $\mathcal{R}$ is connected, then for the genus of the amalgamated Heegaard surface $\mathcal{S}$ we have

$$g(\mathcal{S}) = g(\mathcal{S}_1) + g(\mathcal{S}_2) - g(\mathcal{R}). \tag{5.22}$$

However, in case $\mathcal{R}$ has multiple connected components, then (5.22) does not hold anymore. The procedure of amalgamation nevertheless works for arbitrary generalized Heegaard splittings (cf. Remark 4.2(4) and [5, Section 2], [43, Section 2.3], [130, Section 5.4]), and the formula (5.22) can be adapted to the general setting by taking into account the Euler characteristic of the dual graph of the decomposition.

Theorem 5.33 (Quantitative Amalgamation; cf. Theorems 2.8 and 2.9 in [5])**.**

(1) *Any generalized Heegaard splitting* $\mathcal{M} = \bigcup_{i\in I}(\mathcal{N}_i \cup_{\mathcal{S}_i} \mathcal{K}_i)$ *of a given 3-manifold* $\mathcal{M}$ *can be amalgamated to a (classical) Heegaard splitting* $\mathcal{M} = \mathcal{N} \cup_{\mathcal{S}} \mathcal{K}$ *thereof.*

(2) *Let* $\mathscr{D}$ *be the decomposition* $\mathcal{M} = \bigcup_{i\in I} \mathcal{M}_i$ *underlying the generalized Heegaard splitting above, and* $\Gamma(\mathscr{D}) = (I, E)$ *be its dual graph with Euler characteristic* $\chi(\Gamma(\mathscr{D}))$*. For any* $e = \{u, v\} \in E$*, let* $\mathcal{R}_e$ *be the connected component of* $\mathcal{M}_u \cap \mathcal{M}_v$ *dual to* e*.*[13] *Then the genus* $g(\mathcal{S})$ *of the amalgamated Heegaard surface* $\mathcal{S}$ *satisfies*

$$g(\mathcal{S}) = \sum_{i\in I} g(\mathcal{S}_i) - \sum_{e\in E} g(\mathcal{R}_e) + 1 - \chi(\Gamma(\mathscr{D})). \tag{5.23}$$

Now let $\mathcal{M}$ be a hyperbolic 3-manifold and consider a generalized Heegaard splitting $\mathcal{M} = \bigcup_{i\in I}(\mathcal{N}_i \cup_{\mathcal{S}_i} \mathcal{K}_i)$ with minimum genus Heegaard surfaces, induced by a thick-thin decomposition (Figure 5.15). By Theorem 5.33, we may amalgamate this to a classical Heegaard splitting $\mathcal{M} = \mathcal{N} \cup_{\mathcal{S}} \mathcal{K}$. Finally, by combining the formula (5.23) with the data from Proposition 5.32, we immediately get

$$\begin{aligned} g(\mathcal{S}) &= \sum_{i\in I} g(\mathcal{S}_i) - \sum_{e\in E} g(\mathcal{R}_e) + 1 - \chi(\Gamma(\mathscr{D})) && (5.24)\\ &= g(\mathcal{S}_1) + \sum_{i=2}^{m} (g(\mathcal{S}_i) - g(\mathcal{R}_i)) + 1 - \chi(\Gamma(\mathscr{D})) && (5.25)\\ &= O(\mathrm{vol}(\mathcal{M})) + \sum_{i=2}^{m} (1-1) + 1 - 1 && (5.26)\\ &= O(\mathrm{vol}(\mathcal{M})). && (5.27) \end{aligned}$$

This concludes the proof of Theorem 5.26. □

[13] Note that there might be multiple arcs between the nodes u and v in $\Gamma(\mathscr{D})$. We account for all.

6 The Classification of 3-Manifolds with Treewidth One

This chapter is dedicated to the proof of Theorem 1.9.

Theorem 1.9. *The class of* 3*-manifolds of treewidth at most one coincides with the class of* 3*-manifolds of Heegaard genus at most one together with the Seifert fibered space* $\mathrm{SFS}[\mathbb{S}^2 : (2,1),(2,1),(2,-1)]$ *of Heegaard genus two.*

One direction in Theorem 1.9 readily follows from work of Jaco and Rubinstein.

Theorem 6.1 (cf. [71], Theorem 6.1)**.** *Every lens space has a layered triangulation* $\mathcal{T}$ *with the simplification of* $\Gamma(\mathcal{T})$ *being a path. In particular, all* 3*-manifolds of Heegaard genus at most one have treewidth at most one.*

For the proof of the other direction, the starting point is the following observation.

Lemma 6.2. *If the simplification of a* 4*-regular multigraph is a tree, then it is a path.*

Proof. Let G be a 4-regular multigraph whose simplification $S(G)$ is a tree. Call an arc of $S(G)$ *even* (resp. *odd*) if its corresponding multiple arc in G consist of an even (resp. odd) number of arcs. Let $Odd(G)$ be the subgraph of $S(G)$ consisting of all odd arcs. It follows from a straightforward parity argument that all nodes in $Odd(G)$ have an even degree. In particular, if the set $E(Odd(G))$ of arcs is nonempty, then it necessarily contains a cycle. However, this cannot happen as $S(G)$ is a tree by assumption. Consequently, all arcs of $S(G)$ must be even. This implies that every node of $S(G)$ has degree at most 2 (otherwise there would be a node in G with degree > 4), hence $S(G)$ is a path. □

Consequently, if $\mathrm{tw}\,(\Gamma(\mathcal{T})) \leq 1$ for a triangulation $\mathcal{T}$ of a closed 3-manifold, then $\Gamma(\mathcal{T})$ is a "thick" path. If $\mathrm{tw}\,(\Gamma(\mathcal{T})) = 0$, then $\Gamma(\mathcal{T})$ is a single node with two loops (Figure

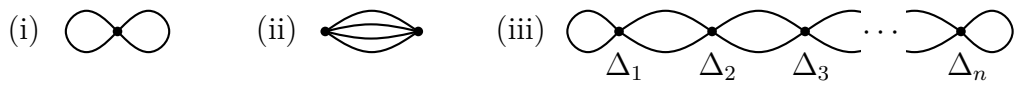

Figure 6.1: The only possible dual graphs (corresponding to closed 3-manifolds) of treewidth at most one

6.1(i)). By looking at the *Closed Census* [28], the only orientable 3-manifolds admitting a dual graph of this kind are $\mathbb{S}^3$ and two lens spaces. If $\Gamma(\mathcal{T})$ has a quadruple arc, then it must be a path of length two (Figure 6.1(ii)), and the only 3-manifold not a lens space appearing here is SFS$[\mathbb{S}^2 : (2,1),(2,1),(2,-1)]$ which has Heegaard genus two, cf. [128, p. 27]. Otherwise, order the tetrahedra $\Delta_1, \ldots, \Delta_n$ of $\mathcal{T}$ as shown in Figure 6.1(iii), and define $\mathcal{T}_i \subset \mathcal{T}$ to be the i^{th} *subcomplex* of $\mathcal{T}$ consisting of $\Delta_1, \ldots, \Delta_i$.

$\mathcal{T}_1$ is obtained by identifying two triangles of Δ_1. Without loss of generality, we may assume that these are the triangles $\Delta_1(013)$ and $\Delta_1(023)$. There are six possible face gluings between them (corresponding to the six bijections $\{0,1,2\} \to \{0,2,3\}$).

The gluing $\Delta_1(013) \mapsto \Delta_1(023)$ yields a 3-vertex triangulation of the 3-ball, called a *snapped* 3*-ball*, and is an admissible choice for $\mathcal{T}_1$, Figure 6.2(i). $\Delta_1(013) \mapsto \Delta_1(032)$ and $\Delta_1(013) \mapsto \Delta_1(203)$ both create Möbius bands as vertex links of the vertices (0) and (2), respectively, and thus these 1-tetrahedron complexes cannot be subcomplexes of a 3-manifold triangulation. $\Delta_1(013) \mapsto \Delta_1(230)$ and $\Delta_1(013) \mapsto \Delta_1(302)$ both produce valid but isomorphic choices for $\mathcal{T}_1$: the minimal layered solid torus of type LST$(1,2,-3)$, Figure 6.2(ii). Lastly, $\Delta_1(013) \mapsto \Delta_1(320)$ identifies the edge (03) with itself in reverse, it is hence invalid. We discuss the two valid cases separately, starting with the latter one.

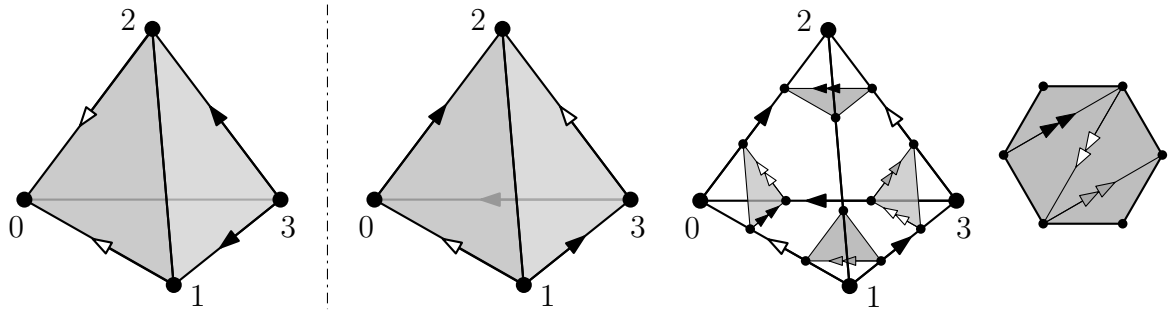

(i) $\Delta_1(013) \mapsto \Delta_1(023)$ (ii) $\Delta_1(013) \mapsto \Delta_1(230)$ (iii) (iv)

Figure 6.2: (i) The snapped 3-ball. (ii) A layered solid torus, (iii) with four normal triangles comprising the single vertex link, (iv) which is a triangulated hexagonal disk

Lemma 6.3. *Let $\mathcal{T}$ be a triangulation of a closed, orientable 3-manifold $\mathcal{M}$. Assume that* $\mathrm{tw}\,(\Gamma(\mathcal{T})) = 1$ *and $\mathcal{T}_1$ is a solid torus. Then $\mathcal{M}$ has Heegaard genus $\mathfrak{g}\,(\mathcal{M}) \leq 1$.*

Proof. The proof consists of the following three parts.

1) We systematize all subcomplexes $\mathcal{T}_2 \subset \mathcal{T}$ which arise from gluing a tetrahedron Δ_2 to $\mathcal{T}_1$ along two triangular faces, and discard all complexes which cannot be part of a 3-manifold triangulation.

2) For the remaining cases we discuss the combinatorial types of complexes $\mathcal{T}_i$, $i > 2$, and triangulations of 3-manifolds arising from them.

3) We show for all resulting triangulations, that the fundamental group of the underlying manifold is cyclic, and that it thus is of Heegaard genus at most one.

To enumerate all possibilities for $\mathcal{T}_2$, we may assume that $\mathcal{T}_1$ is obtained by $\Delta_1(013) \mapsto \Delta_1(230)$. The boundary $\partial\mathcal{T}_1$ is built from two triangles $(012)_\partial$ and $(123)_\partial$, sharing an edge (12), via the identifications (01) = (23) and (02) = (13), see Figure 6.2(ii). The vertex link of $\mathcal{T}_1$ is a triangulated hexagon as shown in Figure 6.2(iii)–(iv).

The second subcomplex $\mathcal{T}_2$ is obtained from $\mathcal{T}_1$ by gluing Δ_2 to the boundary of $\mathcal{T}_2$ along two of its triangles. By symmetry, we are free to choose the first gluing. Hence, without loss of generality, let $\mathcal{T}_2'$ be the complex obtained from $\mathcal{T}_1$ by gluing Δ_2 to $\mathcal{T}_1$ with gluing $\Delta_2(012) \mapsto (012)_\partial$. The result is a 2-vertex triangulated solid torus with four boundary triangles $\Delta_2(013)$, $\Delta_2(023)$, $\Delta_2(123)$ and $(123)_\partial$, see Figure 6.3(ii). Since adjacent edges in the boundary of the unique vertex link of $\mathcal{T}_1$ are always normal arcs in distinct triangles of $\partial\mathcal{T}_1$, the vertex links of $\mathcal{T}_2'$ must be a triangulated 9-gon and a single triangle, shown in Figure 6.3(iii). Note that both vertex links of $\mathcal{T}_2'$ are symmetric with respect to the normal arcs coming from boundary triangles $\Delta_2(013)$, $\Delta_2(023)$ and $\Delta_2(123)$. By this symmetry, we are free to choose whether to glue $\Delta_2(013)$, $\Delta_2(023)$ or $\Delta_2(123)$ to $(123)_\partial$, in order to obtain $\mathcal{T}_2$. Therefore we have the following six possibilities to consider (Figure 6.4).

$\Delta_2(123) \mapsto (123)_\partial$ is a layering onto (12). It yields another layered solid torus with vertex link a triangulated hexagon with edges adjacent in the boundary of the link being normal arcs in distinct faces in $\partial\mathcal{T}_2$. Hence we have the same options for $\mathcal{T}_3$ as the ones in this list. Any complex obtained by iterating this case is of this type.

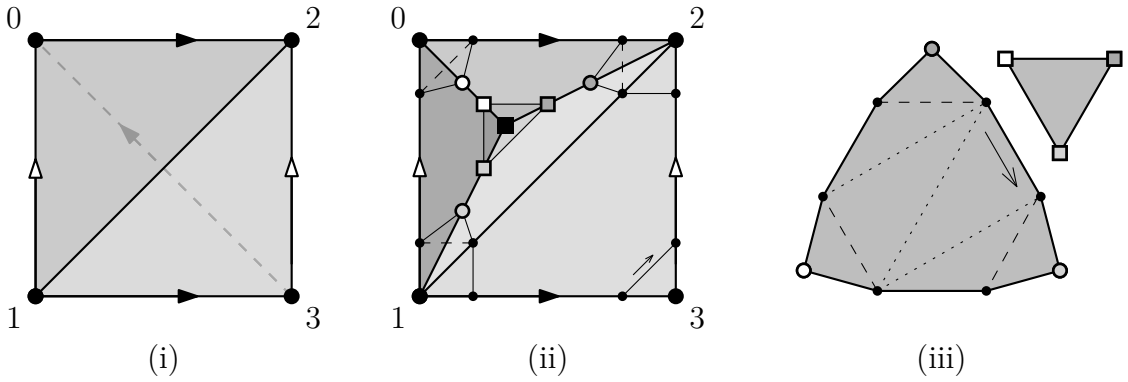

Figure 6.3: (i) The solid torus $\mathcal{T}_1$, (ii) the complex $\mathcal{T}_2'$, and (iii) the vertex links of $\mathcal{T}_2'$

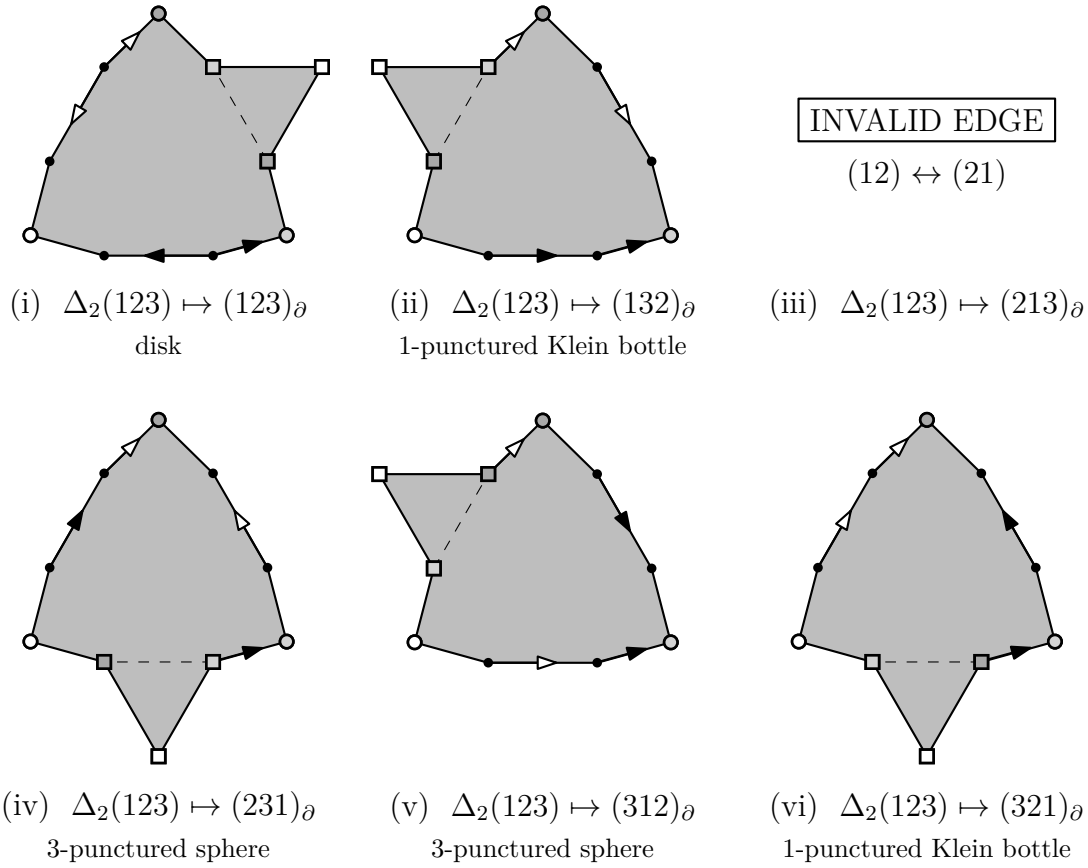

Figure 6.4: The six possibilities for the link of the unique vertex in $\mathcal{T}_2$

Definition 6.4 (type)**.** Here, as well as for the remainder of the proof, whenever we obtain a subcomplex with all cases for the next subcomplex equal to a case already considered (i.e., isomorphic boundary complexes compatible with isomorphic boundaries of vertex links), we talk about these cases to be of the same *type*.

We denote the type obtained via $\Delta_2(123) \mapsto (123)_\partial$ by $\mathscr{T}_\mathrm{I}$.

$\Delta_2(123) \mapsto (132)_\partial$ is invalid, as it creates a punctured Klein bottle as vertex link.

$\Delta_2(123) \mapsto (213)_\partial$ is invalid, as it identifies (12) on the boundary with itself in reverse.

$\Delta_2(123) \mapsto (231)_\partial$ gives a 1-vertex complex with triangles $(013)_\partial$ and $(023)_\partial$ comprising its boundary, which is isomorphic to the boundary of the snapped 3-ball with all of its three vertices identified. The vertex link is a 3-punctured $\mathbb{S}^2$ with two boundary components being normal loop arcs and one consisting of the remaining four normal arcs. This complex is discussed in detail below; we denote its type by $\mathscr{T}_{\mathrm{II}}$.

$\Delta_2(123) \mapsto (312)_\partial$ results in a 1-vertex complex of type $\mathscr{T}_{\mathrm{II}}$ as in the previous case.

$\Delta_2(123) \mapsto (321)_\partial$ is invalid: it produces a punctured Klein bottle in the vertex link.

Now we discuss complexes of type $\mathscr{T}_{\mathrm{II}}$. To this end, let $\mathcal{T}_2$ be the complex in Figure 6.5(ii) defining this type. By gluing Δ_3 to $\mathcal{T}_2$ along a boundary triangle, say $\Delta_3(013) \mapsto (013)_\partial$, we obtain a complex $\mathcal{T}_3'$ (see Figure 6.5(iii)). Note that no boundary edge of the 3-punctured sphere vertex link $\mathcal{L}$ can be identified with an edge in another boundary component of $\mathcal{L}$, for that would create genus in $\mathcal{L}$ (an obstruction to being a subcomplex of a 3-manifold triangulation in which all vertex links must be $\mathbb{S}^2$). As shown in Figure 6.6, there is a unique gluing to avoid this, namely $\Delta_3(023) \mapsto (023)_\partial$, which yields a 1-vertex complex $\mathcal{T}_3$ with vertex link still being a 3-punctured sphere, but now with three boundary components consisting of two edges each, as indicated in Figure 6.6(i). Let $\mathscr{T}_{\mathrm{III}}$ denote its type. Repeating the same argument for $\mathcal{T}_3$ implies that a valid $\mathcal{T}_4$ must be again of type $\mathscr{T}_{\mathrm{II}}$.

Altogether, the type of each intermediate complex $\mathcal{T}_i$ $(i < n)$ is either $\mathscr{T}_{\mathrm{I}}$ (a layered solid torus), or one of the two types $\mathscr{T}_{\mathrm{II}}$ and $\mathscr{T}_{\mathrm{III}}$ of 1-vertex complexes with a 3-punctured sphere as vertex link. If $\mathcal{T}_{n-1}$ is of type $\mathscr{T}_{\mathrm{I}}$, then it can always be completed to a triangulation of a closed 3-manifold by adding a one-tetrahedron solid torus or a snapped 3-ball. If $\mathcal{T}_{n-1}$ is of type $\mathscr{T}_{\mathrm{II}}$, it may be completed by adding a snapped 3-ball. If $\mathcal{T}_{n-1}$ is of type $\mathscr{T}_{\mathrm{III}}$ it cannot be completed to a triangulation of a 3-manifold.

To conclude that any resulting $\mathcal{T}$ triangulates a 3-manifold $\mathcal{M}$ of Heegaard genus at most one, we observe that the fundamental group of $\mathcal{M}$ is generated by one element. Indeed, $\pi_1(\mathcal{T}_1)$ is isomorphic to $\mathbb{Z}$ and is generated by a boundary edge. Furthermore, since $\mathcal{T}_1$ only has one vertex, all edges in $\mathcal{T}_1$ must be loop edges, and no edge which is trivial in $\pi_1(\mathcal{T}_1)$ can become non-trivial in the process of building up the triangulation

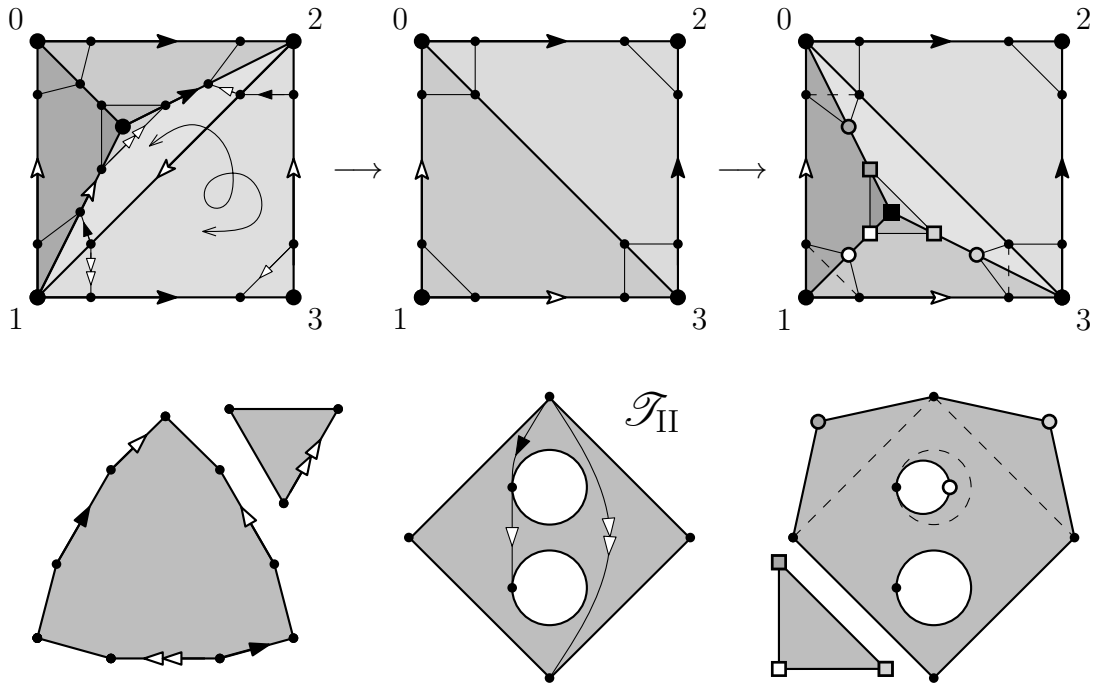

(i) → $\boxed{\Delta_2(123) \mapsto (231)_\partial}$ → (ii) → $\boxed{\Delta_3(013) \mapsto (013)_\partial}$ → (iii)

Figure 6.5: Discussion of the complex of type $\mathscr{T}_{\mathrm{II}}$

of the closed 3-manifold. When we extend $\mathcal{T}_1$ by attaching further tetrahedra along two triangles each, then most edges of the newly added tetrahedron are identified with edges of the previous complex and the unique new boundary edge can be expressed in terms of the existing generator. Therefore the fundamental group of the new complex still admits a presentation with one generator. Moreover, since we have a one-vertex triangulation, no new generator can arise in the last step of gluing the two boundary triangles together.

So either $\pi_1(\mathcal{M})$ is infinite cyclic, i.e., isomorphic to $\mathbb{Z}$, in which case $\mathcal{M}$ must be $\mathbb{S}^2 \times \mathbb{S}^1$ [57]; or $\pi_1(\mathcal{M})$ is finite, but then it is spherical by the Elliptization Theorem (see Theorem 3.21), hence it is a lens space [146, Theorem 4.4.14.(a)]. □

Lemma 6.5. *Let $\mathcal{T}$ be a triangulation of a closed, orientable 3-manifold $\mathcal{M}$ with n tetrahedra, such that $\mathrm{tw}\,(\Gamma(\mathcal{T})) = 1$ and both $\mathcal{T}_1$ and $\mathcal{T} \setminus \mathcal{T}_{n-1}$ is a snapped 3-ball. Then $\mathcal{M}$ has Heegaard genus $\mathfrak{g}\,(\mathcal{M}) \leq 1$.*

Proof. The proof follows the same structure as the one of Lemma 6.3. Let $\mathcal{T}_1$ be a snapped 3-ball, say, obtained by the gluing $\Delta_1(013) \mapsto \Delta_1(023)$. Its boundary is a three-vertex two-triangle triangulation of $\mathbb{S}^2$ with triangles $(012)_\partial$ and $(123)_\partial$ glued along common edge (12) with edge identifications (01) = (02) and (13) = (23), see Figure 6.7(i). One

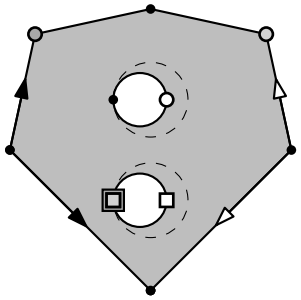

(i) $\Delta_3(023) \mapsto (023)_\partial$
3-punctured sphere

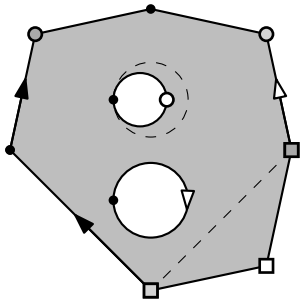

(ii) $\Delta_3(023) \mapsto (032)_\partial$
surface with genus

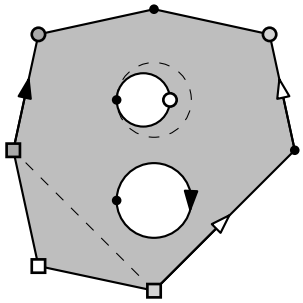

(iii) $\Delta_3(023) \mapsto (203)_\partial$
surface with genus

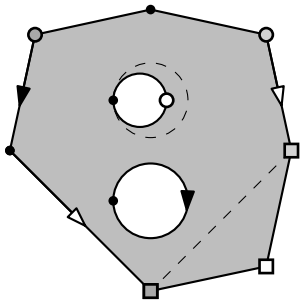

(iv) $\Delta_3(023) \mapsto (230)_\partial$
3-punctured torus

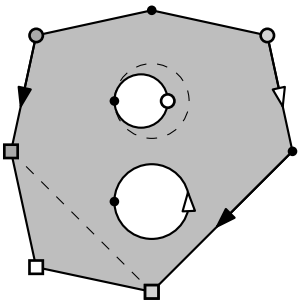

(v) $\Delta_3(023) \mapsto (302)_\partial$
3-punctured torus

INVALID EDGE

$(03) \leftrightarrow (30)$

(vi) $\Delta_3(023) \mapsto (320)_\partial$

Figure 6.6

vertex link of $\mathcal{T}_1$ consists of two triangles identified along a common edge, and two vertex links are single triangles with two of their edges identified (Figure 6.7(iii)).

$\mathcal{T}_2$ is obtained from $\mathcal{T}_1$ by gluing Δ_2 to the boundary of $\mathcal{T}_2$ along two of its triangles. Again, by symmetry, we are free to choose the first gluing. Hence, let $\mathcal{T}_2'$ be the complex obtained from $\mathcal{T}_1$ by gluing Δ_2 to $\mathcal{T}_1$ via $\Delta_2(012) \mapsto (012)_\partial$. The result is a 4-vertex triangulated 3-ball with four boundary triangles $\Delta_2(013)$, $\Delta_2(023)$, $\Delta_2(123)$, and $(123)_\partial$ (Figure 6.8(i)). The vertex links of $\mathcal{T}_2'$ are a two-triangle triangulation of a bigon with an interior vertex of degree one, a triangulation of a hexagon, and two 1-triangle triangulations of a disk, one with a single boundary edge, and one with three boundary edges, cf. Figure 6.8(ii). While all four vertex links of $\mathcal{T}_2'$ are symmetric with respect to the normal arcs of the boundary triangles $\Delta_2(013)$ and $\Delta_2(023)$, the normal arcs of $\Delta_2(123)$ are different. It follows that there are twelve cases to consider (cf. Figures 6.9 and 6.10).

$\Delta_2(013) \mapsto (123)_\partial$ In this case we can see that $\mathcal{T}_2$ has as boundary two triangles identified along their boundaries with two vertices identified. The two vertex links are an annulus and a disk, and all three boundary components of the links are bigons.

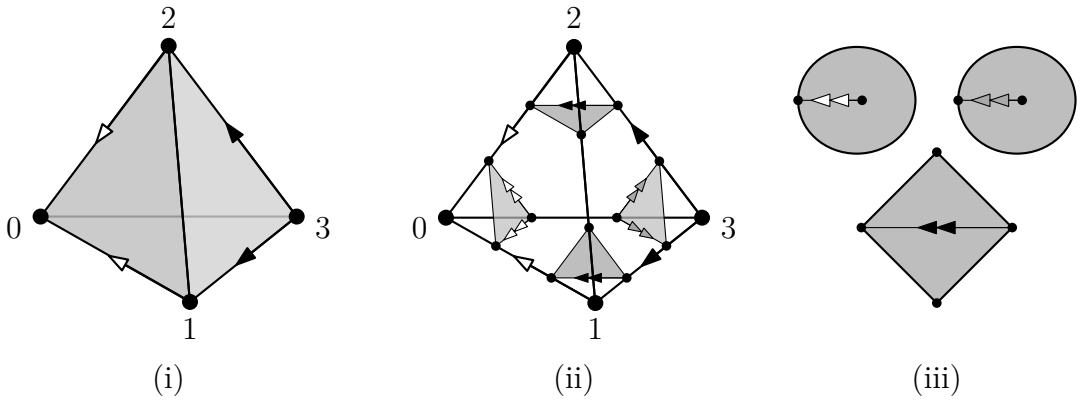

Figure 6.7: (i) A snapped 3-ball, (ii) its normal triangles, and (iii) its vertex links

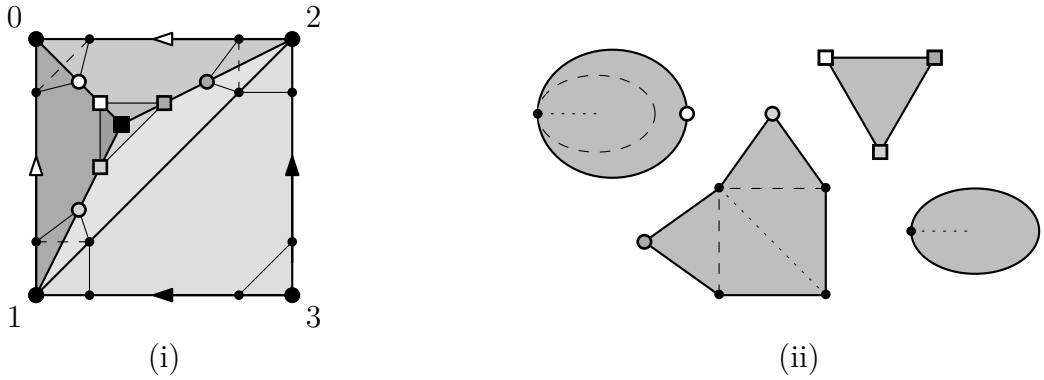

Figure 6.8: (i) The complex $\mathcal{T}_2'$, and (ii) its four vertex links

Extending this case to a valid complex $\mathcal{T}_3$ is only possible in the trivial way, i.e., gluing Δ_3 to $\mathcal{T}_2$ along $\Delta_3(123) \mapsto (123)_\partial$ and $\Delta_3(023) \mapsto (123)_\partial$. This yields a 2-vertex complex with boundary isomorphic to that of the snapped 3-ball but with one apex identified with the vertex of the loop edge. Again, we have one annulus and one disk as vertex links. In accordance with the structure of $\partial\mathcal{T}_3$, the disk is now bounded by a single normal arc while the annulus has a loop normal arc as one boundary component, and four normal arcs in the other.

We can glue Δ_4 to the unique valid complex $\mathcal{T}_3$ described in the previous paragraph in twelve distinct ways: We start by $\Delta_4(012) \mapsto (012)_\partial$ and proceed by gluing $\Delta_4(013)$ and $\Delta_4(023)$ to $(012)_\partial$ in all possible six ways each.

Apart from the trivial gluing, which results in the same type as $\mathcal{T}_2$ above, we obtain three 1-vertex complexes with boundary being $\mathbb{S}^2$ with vertices identified and vertex link a 3-punctured sphere. Two of them have a vertex link with three boundary components consisting of two edges each (i.e., $\partial\mathcal{T}_4$ is of type two triangles identified

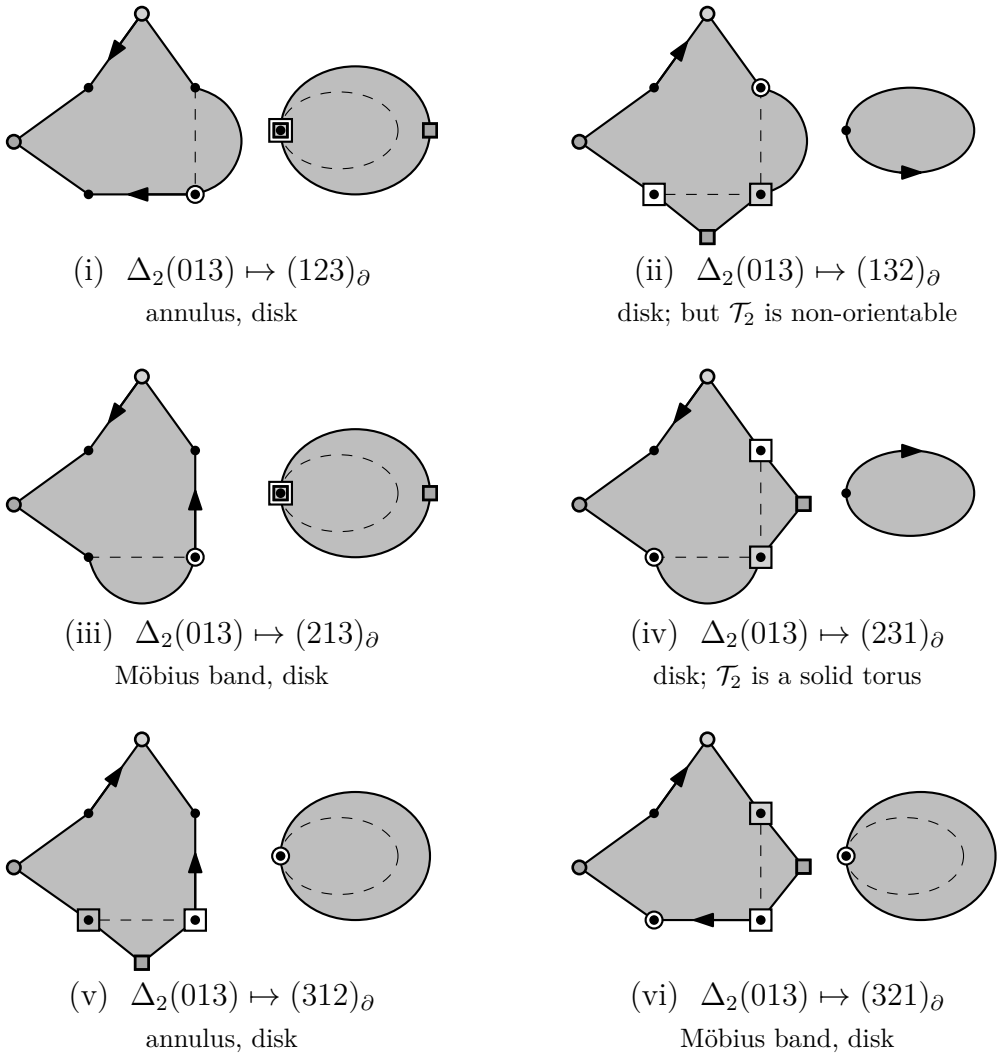

Figure 6.9

along their boundaries and all vertices identified), one of them has a vertex link with two boundary components with one and one boundary component with four edges (i.e., the boundary $\partial\mathcal{T}_4$ is isomorphic to that of the snapped 3-ball with all vertices identified). These cases correspond to types $\mathscr{T}_{\text{III}}$ and $\mathscr{T}_{\text{II}}$ respectively, from the proof of Lemma 6.3.

$\Delta_2(013) \mapsto (132)_\partial$ yields a non-orientable 1-handle with Klein bottle boundary, which cannot be completed to an orientable 3-manifold. Hence we are done with this case.

$\Delta_2(013) \mapsto (213)_\partial$ is invalid, as it produces a Möbius band in one of the vertex links.

$\Delta_2(013) \mapsto (231)_\partial$ gives a 1-vertex triangulation of the solid torus. In particular, the vertex link is a triangulated hexagon with neighboring edges in the boundary of the link being normal arcs in triangles of $\partial\mathcal{T}_2$. We can thus proceed as in the proof of

Lemma 6.3 to conclude that $\mathcal{T}$ must be of Heegaard genus one.

$\Delta_2(013) \mapsto (312)_\partial$ produces a 2-sphere of type "boundary of the snapped 3-ball" in the boundary, and two vertex links. One of them an annulus, the other one a disk. Here, the disk is bounded by a single normal arc while the annulus has a loop normal arc as one boundary component, and four normal arcs in this other. This case is of the same type as $\mathcal{T}_3$ in case $\Delta_2(013) \mapsto (123)_\partial$ above.

$\Delta_2(013) \mapsto (321)_\partial$ creates a Möbius band in the vertex link and can thus be discarded.

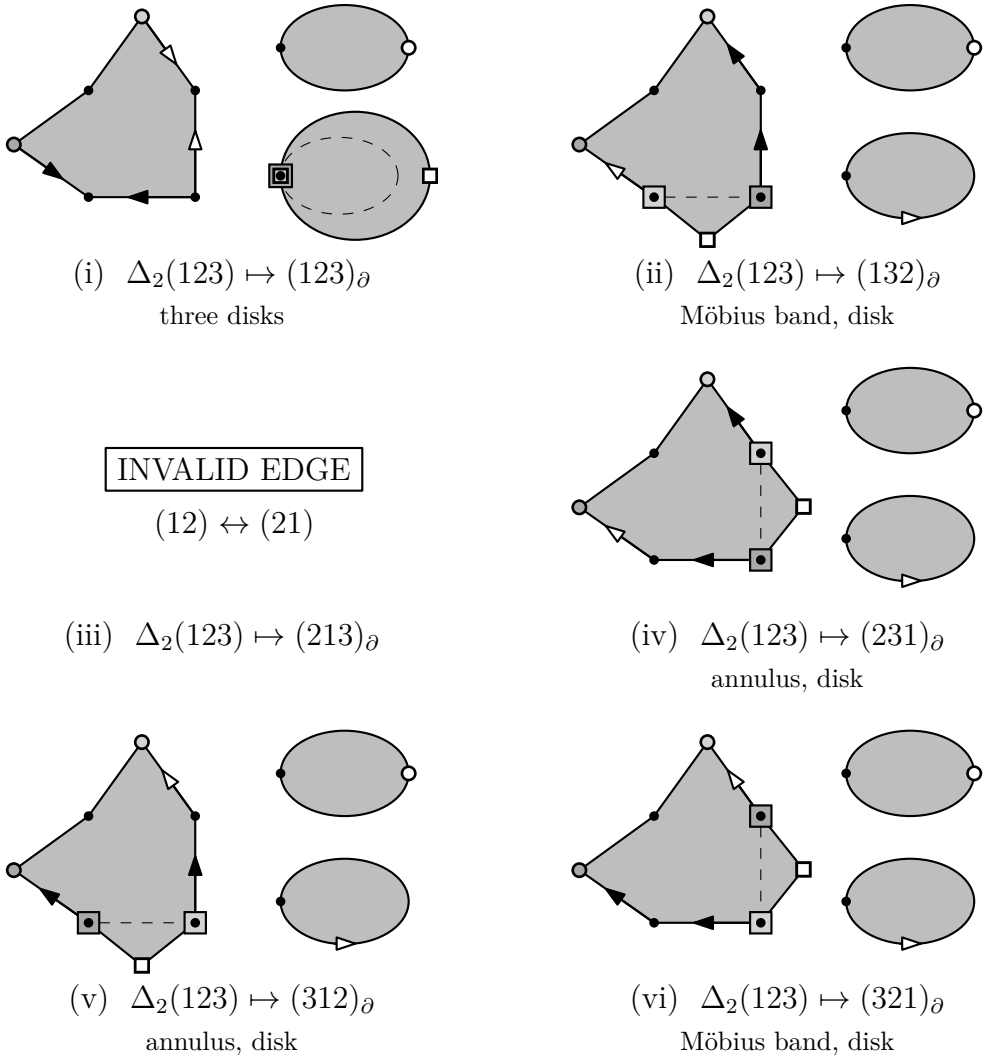

Figure 6.10

$\Delta_2(123) \mapsto (123)_\partial$ is a layering that creates an interior degree two edge. Consequently we obtain two triangles glued along their boundaries for $\partial\mathcal{T}_3$. The three vertex links are all disks with two boundary edges.

Extending this complex by attaching $\mathcal{T}_3$ yields three valid complexes. The first is obtained via the trivial gluing and of the same type as the snapped 3-ball $\mathcal{T}_1$. The other two are 1-vertex triangulations of the solid torus discussed in Lemma 6.3.

$\Delta_2(123) \mapsto (132)_\partial$ yields a Möbius band in the vertex link and can be discarded.

$\Delta_2(123) \mapsto (213)_\partial$ identifies (12) on the boundary with itself in reverse, thus is invalid.

$\Delta_2(123) \mapsto (231)_\partial$ This gluing, again, produces two vertex links. One of them an annulus, the other one a disk. The boundary of $\mathcal{T}_2$ and the boundary components of the vertex links are of the same type as in the case $\Delta_2(013) \mapsto (123)_\partial$.

$\Delta_2(123) \mapsto (312)_\partial$ gives an annulus and a disk as vertex links. $\partial\mathcal{T}_2$ and the boundary components of the vertex links are of the same type as in the case $\Delta_2(013) \mapsto (123)_\partial$.

$\Delta_2(123) \mapsto (321)_\partial$ produces a Möbius band in the vertex link and can be discarded.

It remains to show—along the same lines as in Lemma 6.3—that none of the complexes described above can be completed to a triangulation of Heegaard genus greater than one. It suffices to look at the complexes which can be completed to a triangulation of a manifold. The 1-vertex complexes with torus boundary ($\Delta_2(013) \mapsto (231)_\partial$ and $\Delta_2(123) \mapsto (123)_\partial$) are solid tori and thus admit a fundamental group with one generator. Following the proof of Lemma 6.3, the Heegaard genus of a triangulation of a closed 3-manifold obtained from these subcomplexes are of Heegaard genus at most one. The three 1-vertex complexes with vertex links a 3-punctured sphere have as fundamental group the free group with two generators. However, note that these complexes can only be extended by trivial gluings and completed by inserting a 1-tetrahedron snapped 3-ball or, in the case of three boundary components of size two, closed off by trivially identifying the two boundary triangles. In all of these cases we obtain a triangulation of the 3-sphere and in particular a closed 3-manifold of Heegaard genus at most one. □

7 Some 3-Manifolds with Treewidth Two

In what follows, we use the classification of 3-manifolds of treewidth one (Theorem 1.9) to show that a large class of orientable Seifert fibered spaces and some *graph manifolds* have treewidth two. This is done by exhibiting appropriate triangulations, which have all the hallmarks of a space station. First, we review the building blocks (cf. Appendix E).

7.1 What Makes a Space Station?

Robotic arms. These are the layered solid tori with 2-triangle boundaries introduced in Section 4.2.3 and encountered in the proof of Theorem 1.9. Their dual graphs are thick paths (Figure 4.10(v)). A layered solid torus is of type $\mathrm{LST}(p, q, r)$ if its meridional disk intersects its boundary edges p, q and r times. For any coprime p, q, r with $p + q + r = 0$, a triangulation of type $\mathrm{LST}(p, q, r)$ can be realized by [24, Algorithm 1.2.17].

Example 7.1. A special class of robotic arms are the ones of type $\mathrm{LST}(0, 1, 1)$, as they can be used to trivially fill-in superfluous torus boundary components without inserting an unwanted exceptional fiber into a Seifert fibered space (cf. descriptions of $\mathbb{A}_2$ and $\mathbb{A}_1$ below). One of the standard triangulations of robotic arms of type $\mathrm{LST}(0, 1, 1)$ has three tetrahedra Δ_i, $0 \leq i \leq 2$, and is given by the gluing relations

$$\begin{aligned} &\Delta_0(023) \mapsto \Delta_1(013), \quad \Delta_0(123) \mapsto \Delta_1(120), \quad \Delta_1(023) \mapsto \Delta_2(201), \\ &\qquad\qquad\quad \Delta_1(123) \mapsto \Delta_2(301), \quad \Delta_2(023) \mapsto \Delta_2(312). \end{aligned} \tag{7.1}$$

Core unit with three docking sites. Start with a triangle t, take the product $t \times [0, 1]$, triangulate it using three tetrahedra, Figure 7.1(i)–(ii), and glue $t \times \{0\}$ to $t \times \{1\}$ without a

twist, Figure 7.1(iii). The dual graph of the resulting complex $\mathbb{A}_3$ (which is topologically a solid torus) is K_3, hence of treewidth two. Its boundary (a 6-triangle triangulation of the torus) can be split into three 2-triangle annuli, corresponding to the edges of t, each of which we call a *docking site*. Edges running along a fiber and thus of type $\{\text{vertex of } t\} \times [0,1]$ are termed *vertical edges*. Edges orthogonal to the fibers, i.e., the edges of $t \times \{0\} = t \times \{1\}$, are termed *horizontal edges*. The remaining edges are referred to as *diagonal edges*. The triangulation of $\mathbb{A}_3$ has gluing relations

$$\Delta_0(012) \mapsto \Delta_1(012), \quad \Delta_1(013) \mapsto \Delta_2(013), \quad \Delta_2(023) \mapsto \Delta_0(312). \tag{7.2}$$

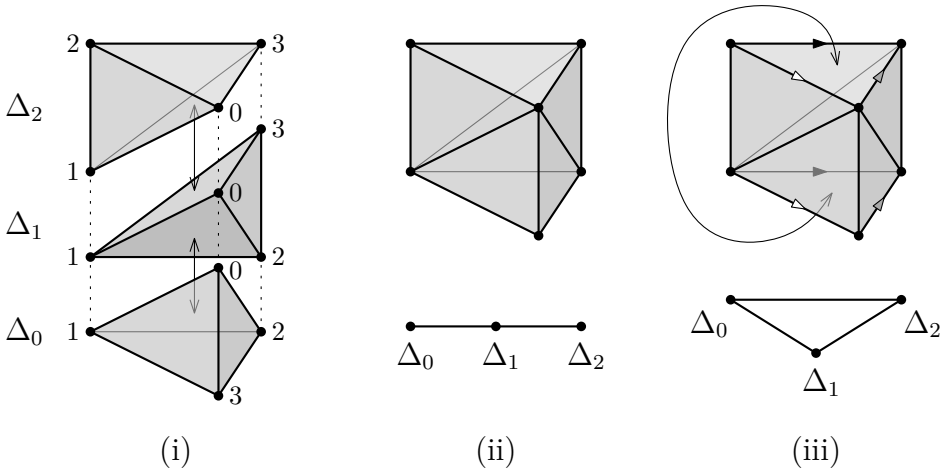

Figure 7.1: Construction of the core unit $\mathbb{A}_3$ with three docking sites

Core assembly with r docking sites. For $r = 2$ (resp. $r = 1$), take a core unit $\mathbb{A}_3$ and glue a robotic arm of type $\mathrm{LST}(0,1,1)$ onto one (resp. two) of its docking sites such that the *unique boundary edge* of the robotic arm (i.e., the boundary edge which is only contained in one tetrahedron of the layered solid torus) is glued to a horizontal boundary edge of $\mathbb{A}_3$. See Example 7.1 for a detailed description of a particular triangulation of a layered solid torus of type $\mathrm{LST}(0,1,1)$ with unique boundary edge being $\Delta_0(01)$. The resulting complex is denoted by $\mathbb{A}_2$ (resp. $\mathbb{A}_1$) and their dual graphs are shown in Figure 7.1. Observe that they have treewidth two.

For $\mathbb{A}_r$ ($r \geq 3$) take $r-2$ copies of $\mathbb{A}_3$, denote them by $\mathbb{A}_3^i$ ($1 \leq i \leq r-2$) with tetrahedra Δ_0^i, Δ_1^i and Δ_2^i ($1 \leq i \leq r-2$). Glue them together by mirroring them across one of their docking sites as shown by (7.3) for $1 \leq i \leq r-3$ odd, and by (7.4) for

$2 \leq i \leq r-3$ even. The resulting complex, denoted by $\mathbb{A}_r$ (see Figure 7.3 for $r = 5$), has $2r$ boundary triangles which become r docking sites.

$$i \text{ odd: } \Delta_1^i(123) \mapsto \Delta_1^{i+1}(123), \quad \Delta_2^i(123) \mapsto \Delta_2^{i+1}(123). \tag{7.3}$$

$$i \text{ even: } \Delta_0^i(023) \mapsto \Delta_0^{i+1}(023), \quad \Delta_1^i(023) \mapsto \Delta_1^{i+1}(023). \tag{7.4}$$

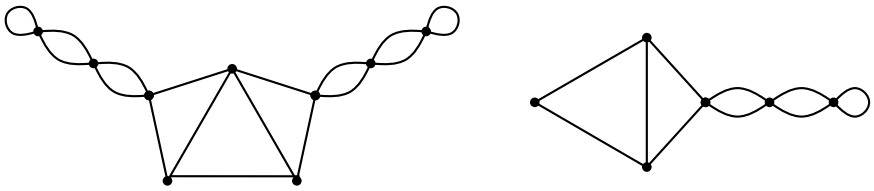

(i) $\Gamma(\mathbb{A}_1)$ (ii) $\Gamma(\mathbb{A}_2)$

Figure 7.2: The dual graphs of the core assemblies $\mathbb{A}_1$ and $\mathbb{A}_2$ with one and two docking sites, respectively

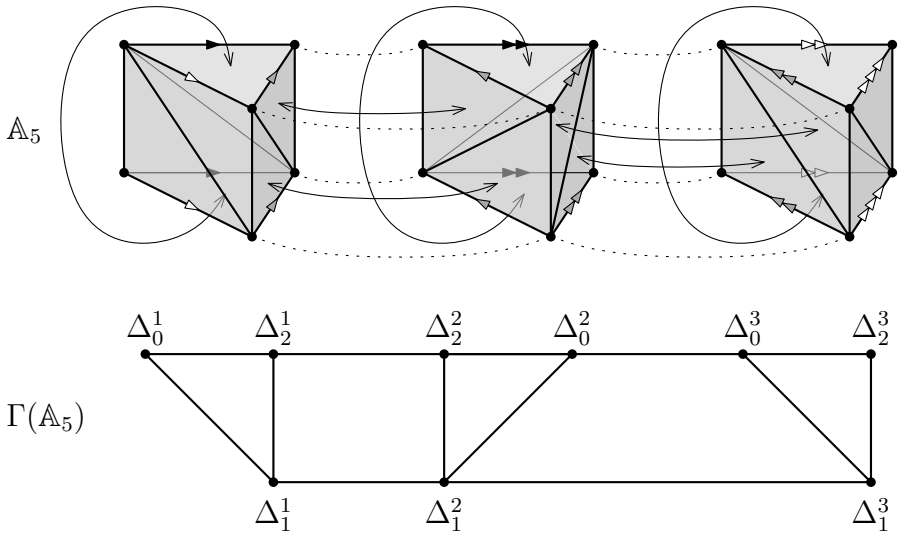

Figure 7.3: The core assembly $\mathbb{A}_5$ with five docking sites and its dual graph $\Gamma(\mathbb{A}_5)$

Möbius laboratory module. This complex, denoted by $\mathbb{M}$, is given by

$$\begin{aligned} T_0(123) \mapsto T_1(123), \quad T_0(023) \mapsto T_1(031), \quad T_1(012) \mapsto T_2(201), \\ T_1(023) \mapsto T_2(023), \quad T_0(013) \mapsto T_2(132). \end{aligned} \tag{7.5}$$

Its dual graph is a triangle with two double edges, and hence of treewidth two (see, for instance, Figure 7.5). $\mathbb{M}$ has one torus boundary component, or docking site, given by the two triangles $T_0(012)$ and $T_2(013)$ with edges $T_0(01) = T_2(13)$, $T_0(02) = T_2(03)$, and $T_0(12) = T_2(01)$. $\mathbb{M}$ triangulates the orientable $\mathbb{S}^1$-bundle over $\mathcal{N}_{1,1}$.

7.2 Constructing Space Stations

Theorem 7.2. *Orientable Seifert fibered spaces over $\mathbb{S}^2$ have treewidth at most two.*

Proof. To obtain a treewidth two triangulation of $\mathrm{SFS}[\mathbb{S}^2 : (a_1, b_1), \ldots, (a_r, b_r)]$, start with the core assembly $\mathbb{A}_r$ and a collection $\{\mathrm{LST}(a_i, \pm|b_i|, -a \mp |b_i|) : 1 \leq i \leq r\}$ of robotic arms. The robotic arms are then glued to the r docking sites (2-triangle torus boundary components, separated by the vertical boundary edges) of $\mathbb{A}_r$, such that boundary edges of type a_i are glued to vertical boundary edges, and edges of type b_i are glued to horizontal boundary edges. The sign in $\mathrm{LST}(a_i, \pm|b_i|, -a \mp |b_i|)$ is then determined by the type of diagonal edge in the i^{th} docking site of $\mathbb{A}_r$ and by the sign of b_i. (See, e.g., Figure 7.4.) □

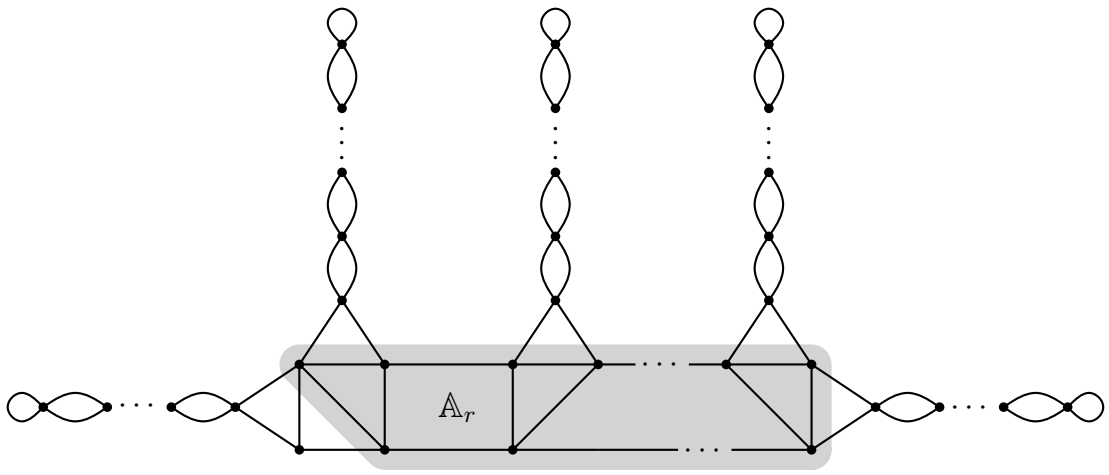

Figure 7.4: Dual graph of a treewidth two triangulation of an orientable SFS over $\mathbb{S}^2$

Remark 7.3. Note that, in some cases, a fiber of type $(2, 1)$ can be realized by identifying the two triangles of a docking site of $\mathbb{A}_r$ with a twist. In the dual graph this appears as a double arc rather than the attachment of a thick path. See Figure 7.6 on the right for an example in the treewidth two triangulation of the *Poincaré homology sphere*.

Theorem 7.4. *An orientable Seifert fibered surface over a non-orientable surface is of treewidth at most two.*

Proof. In order to obtain a treewidth two triangulation of the orientable Seifert fibered space $\mathrm{SFS}[\mathcal{N}_g : (a_1, b_1), \ldots, (a_r, b_r)]$ over the non-orientable surface $\mathcal{N}_g$ of genus g, start

with a core assembly $\mathbb{A}_{r+g}$ and attach g copies $\mathbb{M}_j$, $1 \leq j \leq g$, of the Möbius laboratory module via

$$T_0^j(012) \mapsto \Delta_2^j(201) \quad \text{and} \quad T_2^j(013) \mapsto \Delta_0^j(013), \tag{7.6}$$

where T_0^j, T_1^j and T_2^j are the tetrahedra comprising $\mathbb{M}_j$, and Δ_0^j, Δ_1^j and Δ_2^j denote the tetrahedra making up the first g core units (each being a copy of $\mathbb{A}_3$) in $\mathbb{A}_{r+g}$. By construction, this produces a triangulation of the orientable $\mathbb{S}^1$-bundle over $\mathcal{N}_{g,1}$. Proceed by attaching a robotic arm of type $\mathrm{LST}(a_i, \pm|b_i|, -a_i \mp |b_i|)$, $1 \leq i \leq r$, to each of the remaining r docking sites. Again, for the gluings between the robotic arms and the core assembly $\mathbb{A}_{r+g}$, the edges of type a_i must be glued to the vertical boundary edges, the edges of type b_i must be glued to the horizontal boundary edges, and attention has to be paid to the signs of the b_i and to how exactly diagonal edges run. See Figure 7.5 for the dual graph of the resulting complex, which is of treewidth two by inspection. □

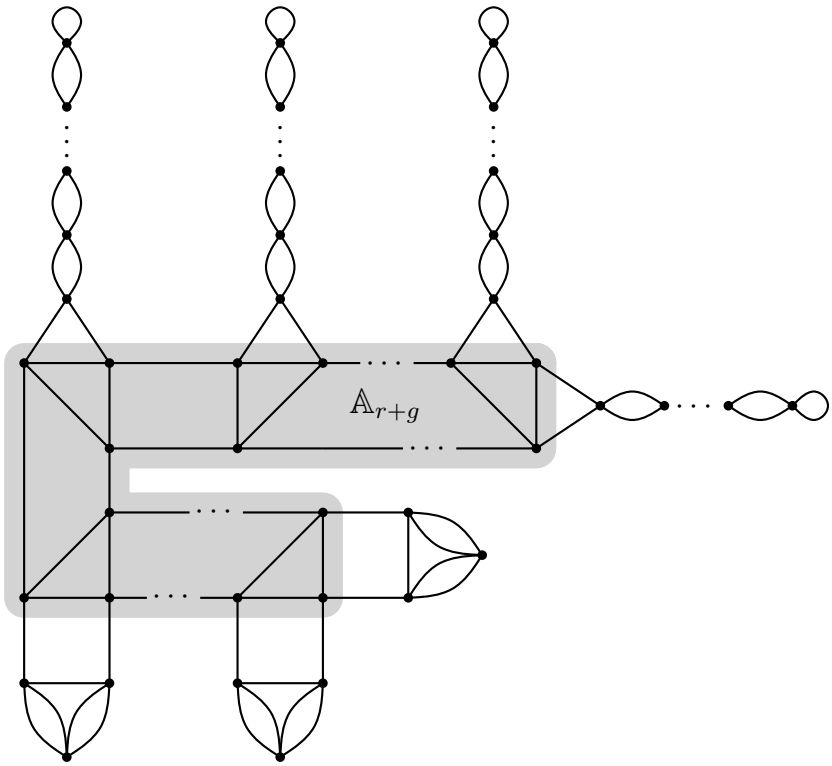

Figure 7.5: Dual graph of a treewidth two triangulation of an orientable SFS over $\mathcal{N}_g$

Corollary 7.5. *An orientable Seifert fibered space $\mathcal{M}$ over $\mathbb{S}^2$ or a non-orientable surface is of treewidth one if $\mathcal{M}$ is a lens space or* $\mathrm{SFS}[\mathbb{S}^2 : (2,1),(2,1),(2,-1)]$*, and two otherwise.*

Corollary 7.5 directly follows from the combination of Theorems 1.9, 7.2 and 7.4.

Corollary 7.6. *Orientable* 3*-manifolds with spherical or* $\mathbb{S}^2 \times \mathbb{R}$ *geometry have treewidth at most two.*

Proof. Every orientable 3-manifold with spherical or $\mathbb{S}^2 \times \mathbb{R}$ geometry can be represented either as a Seifert fibered space over $\mathbb{S}^2$ with at most three exceptional fibers, or as a Seifert fibered space over the projective plane (i.e., $\mathcal{N}_1$) with at most one exceptional fiber [135]. Hence the result follows from Theorems 7.2 and 7.4. □

Corollary 7.7. *Graph manifolds* $\mathcal{M}_T$ *modeled on a tree* T *with nodes being orientable Seifert fibered spaces over* $\mathbb{S}^2$ *or* $\mathcal{N}_g$*,* $g > 0$*, have treewidth at most two.*

Proof of Corollary 7.7. Let $\mathcal{M}$ be such a graph manifold. A treewidth two triangulation of $\mathcal{M}$ can be constructed in the following way: For every node of T of degree k insert a treewidth two triangulation of the respective Seifert fibered space from Theorem 7.2 or Theorem 7.4 with k additional docking sites. As can be deduced from the construction of $\mathbb{A}_r$, this can be done without increasing the treewidth. Proceed by gluing the Seifert fibered spaces along the arcs of T using the additional docking sites (torus boundary components) added in the previous step: Every such gluing is determined by a diffeomorphism on the torus and every such diffeomorphism can be modelled by a sequence of layerings onto the boundary components with dual graph a path of double edges of treewidth one. This fact can be deduced from, for instance, [71, Lemma 4.1]. Altogether, the triangulation constructed this way is of treewidth at most two, cf. Figure 7.7. □

Corollary 7.8. *Minimal triangulations are not always of minimum treewidth.*

Proof. The *Poincaré homology sphere* $\mathrm{SFS}[\mathbb{S}^2 : (2,1),(3,1),(5,-4)]$ has treewidth two but its minimal triangulation has treewidth four, see Figure 7.6. □

Corollary 7.9. *There exist irreducible* 3*-manifolds with treewidth two, but arbitrarily high Heegaard genus.*

Proof. Due to work of Boileau and Zieschang [21, Theorem 1.1(i)], the 3-manifold

$$\mathcal{M}_m = \mathrm{SFS}[\mathbb{S}^2 : (2,1),\ldots,(2,1),(a_m,b_m)\ ],$$

for a_m odd, has Heegaard genus $\mathfrak{g}(\mathcal{M}_m) \geq m-2$. At the same time, $\mathcal{M}_m$ is irreducible [134, Theorem 3.7.17] and $\mathrm{tw}(\mathcal{M}_m) = 2$ due to Theorem 7.4. □

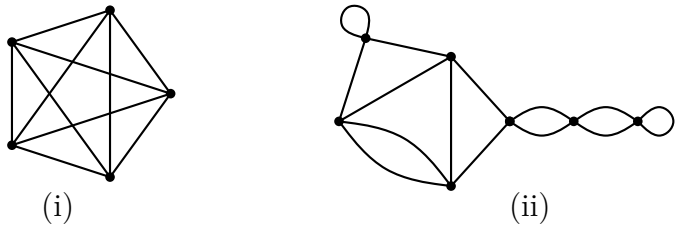

Figure 7.6: Dual graph of the minimal (i), and of a treewidth two (ii) triangulation of the Poincaré homology sphere (cf. Corollary 7.8 and Remark 7.3)

By combining Theorems 1.9, 7.2 and 7.4, we can determine the treewidth of most closed, orientable 3-manifolds triangulable with at most 10 tetrahedra (Table 7.1).

n	# mfds. $\mathcal{M}$	$\mathrm{tw}(\mathcal{M}) = 0$	$\mathrm{tw}(\mathcal{M}) = 1$	$\mathrm{tw}(\mathcal{M}) = 2$	unknown
1	3	3	0	0	0
2	7	0	7	0	0
3	7	0	6	1	0
4	14	0	10	4	0
5	31	0	20	11	0
6	74	0	36	36	2
7	175	0	72	100	3
8	436	0	136	297	3
9	1154	0	272	861	21
10	3078	0	528	2489	61
Σ	4979	3	1087	3799	90

Table 7.1: The 3-manifolds triangulable with ≤ 10 tetrahedra and their treewidths

Using similar constructions it can be shown that orientable Seifert fibered spaces over orientable surfaces have treewidth at most four. Naturally, this only provides an upper bound rather than the actual treewidth of this family of 3-manifolds. Computing the maximum treewidth of an orientable Seifert fibered space is left as future work.

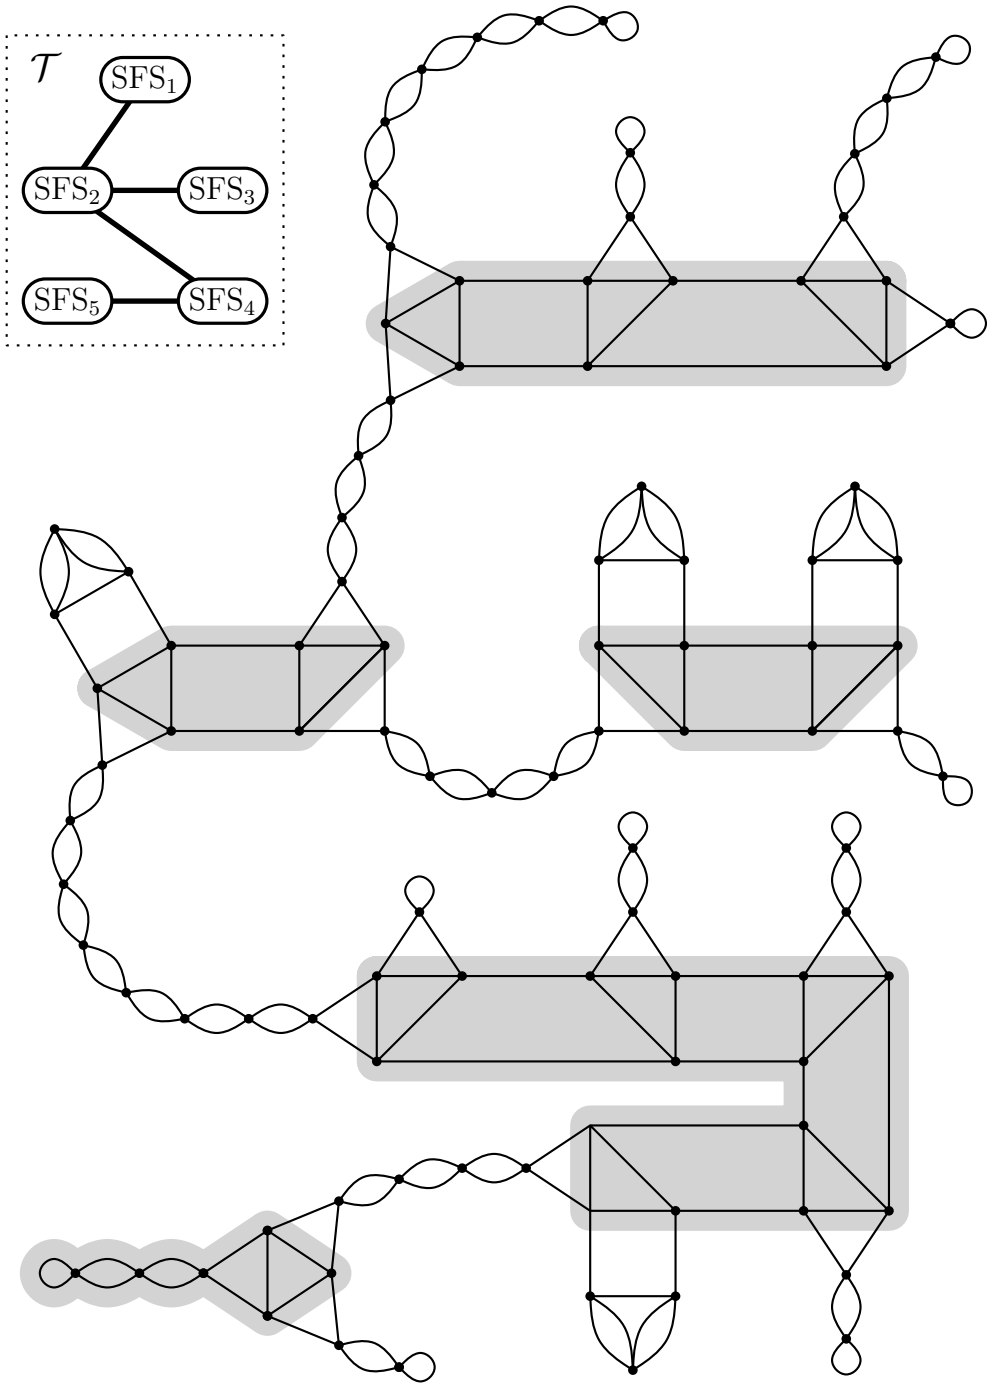

Figure 7.7: Dual graph of a treewidth two triangulation of a graph manifold modeled on the tree T. In order to increase visibility, the core of each constituent Seifert fibered space is highlighted in gray

A Computational Aspects

A.1 Complexity and Fixed Parameter Tractability

In Table A.1 we collect complexity and fixed parameter tractability properties of computing the considered graph parameters. First, we explain the columns of the table.

- **Complexity.** The computational complexity of the question "Is $p(G) \leq k$?". Here k is a variable given as part of the input.
- **FPT.** Fixed parameter tractability in the natural parameter. The check mark (✓) indicates the following: if k is fixed (as opposed to being a variable part of the input) and G is an n-vertex graph, then the answer to the question "Is $p(G) \leq k$?" can be found in $O(\text{poly}(n))$ time.
- **Bounded-degree graphs.** What is known if we restrict our attention to a family of bounded-degree graphs.

p	**Complexity**	**FPT**	**Bounded-degree graphs**
tw	**NP**-complete [3]	✓ [17]	remains **NP**-complete [20]
pw	**NP**-complete [3]	✓ [17]	remains **NP**-complete [109][1]
cw	**NP**-complete [53]	✓ [143]	polynomial if tw bounded [142]
cng	**NP**-complete [138]	✓ [141][2]	

Table A.1: Complexity and fixed parameter tractability of selected graph parameters

[1] **NP**-completeness is shown for the vertex separation number which is equivalent to pathwidth [79].

[2] See the discussion in the introduction of [141].

We point out that there is a more detailed table in [14], showing the complexity of computing pathwidth and treewidth on several different classes of graphs.

A.2 On Working with Different Width Parameters

A small treewidth $k \geq 0$ of the dual graph $\Gamma(\mathcal{T})$ of an n-tetrahedron triangulation $\mathcal{T}$ of a 3-manifold can be exploited by applying standard dynamic programming techniques to the tetrahedra of the triangulation.

In a tree decomposition T of $\Gamma(\mathcal{T})$ with $O(n)$ bags realizing width k, every bag $B \in V(T)$ corresponds to a subcomplex $X_B \subset \mathcal{T}$ of at most $k+1$ tetrahedra of $\mathcal{T}$. Going up from the leaves of T, for each bag $B \in V(T)$, compute a list of candidate solutions of the given problem on $X_B \subset \mathcal{T}$. When processing a new bag $B' \in V(T)$, for all child bags $B_i \in V(T)$, $1 \leq i \leq r$, their lists of candidate solutions (which are already computed) are used to validate or disqualify candidate solutions for $X_{B'}$. Due to property 2) of a tree decomposition (Definition 2.1), every time a tetrahedron disappears from a bag while we go up from the leaves to the root of T, it never reappears. This means that constraints for a global solution coming from such a "forgotten" tetrahedron are fully incorporated in the candidate solutions of the current bags. If for each bag the running time, as well as the length of the list of candidate solutions is a function in k, the procedure must have running time $O(n)$ for triangulations with dual graphs of constant treewidth.

Moreover, for every tree decomposition T with $O(n)$ bags, there exists a linear time procedure to preprocess T into a tree decomposition T' of $\Gamma(\mathcal{T})$, also with $O(n)$ bags, such that every bag is of one of three types: *introduce*, *forget*, or *join bag* [82]. Such a *nice tree decomposition*[3] has the advantage that only three distinct procedures are needed to process all the bags—causing such FPT-algorithms to be much simpler in structure: Introduce and forget bags are bags in T with only one child bag where a node is either added to or removed from the child bag to obtain the parent bag. Procedures to deal with these situations are often comparatively simple to implement and running times are often comparatively feasible. A join bag is a bag with two child bags such that parent bag and both child bags are identical. Depending on the problem to be solved, the procedure of

[3]See the 3-manifold software *Regina* [28] for a visualization of a tree decomposition, and a nice tree decomposition of the dual graph of any given triangulation.

a join bag can be more intricate, and running times are often orders of magnitude slower than in the other two cases.

Pathwidth vs. treewidth. Since every (preprocessed) nice path decomposition is a nice tree decomposition without join bags, every FPT-algorithm in the treewidth of the input is also FPT in the pathwidth with the (often dominant) running time of the join bag removed. Thus, at least in certain circumstances, it can be beneficial to work with nice path decompositions—and thus with pathwidth as a parameter—instead of treewidth and its more complicated join bags in their nice tree decompositions.

For small cutwidth and for small congestion, similar dynamic programming techniques can be applied to the cutsets of the respective linear or binary tree layouts of the nodes and arcs of the dual graph of a triangulation. Thus, in the following paragraph we compare these parameters to treewidth (and pathwidth), and point out some potential benefits from using them as alternative parameters.

Congestion vs. treewidth and cutwidth vs. pathwidth. Parameterized algorithms using pathwidth or treewidth operate on *bags* containing elements corresponding to tetrahedra of the input triangulations. In contrast to this, parameterized algorithms using cutwidth or congestion operate on *cutsets* containing elements corresponding to triangles of the input triangulations. It follows that an algorithm operating on a tree decomposition of width k must handle a 3-dimensional subcomplex of the input triangulation made of up to $15(k+1)$ faces in one step. An algorithm operating on a tree layout of congestion k, however, only needs to consider a 2-dimensional subcomplex of up to $7k$ faces of the input triangulation per step. Moreover, cutwidth and congestion are equivalent to pathwidth and treewidth respectively (for bounded degree graphs, up to a small constant factor, see Theorems 2.4 and 2.6), and parameterized algorithms for 3-manifolds are not just theoretical statements, but may give rise to practical tools outperforming current state-of-the-art algorithms (see, for example, [32]). These observations suggest that, at least for some problems, parameterized algorithms using cutwidth or congestion of the dual graph have a chance to outperform similar algorithms operating on pathwidth or treewidth, respectively.

B High-Treewidth Triangulations

B.1 Most Triangulations Have Large Treewidth

As mentioned in the Introduction, most triangulations of most 3-manifolds must have large treewidth. Here we briefly review two well-known and simple observations that, together, imply this fact. Since they are hard to find in the literature, we also sketch their proofs.

Proposition B.1. *There exists a constant $C > 1$, such that there are at least $C^{n \log(n)}$ 3-manifolds which can be triangulated with $\leq n$ tetrahedra.*

Sketch of the proof. Note that the number of isomorphism classes of graphs is superexponential in the number of nodes. Consider the family of graph manifolds where the nodes are Seifert fibered spaces of constant size.[1] Conclude by the observation that graph manifolds modeled on non-isomorphic graphs are non-homeomorphic, cf. [89, Section 3]. □

Proposition B.2. *Given $k \geq 0$, there exists a constant $C_k > 1$ such that there are at most C_k^n triangulations of 3-manifolds with dual graph of treewidth $\leq k$ and $\leq n$ tetrahedra.*

Sketch of the proof. The property of a graph to be of bounded treewidth is closed under minors. Hence, it follows from a theorem of Norine, Seymour, Thomas and Wollan [113] that the number of isomorphism classes of graphs with treewidth $\leq k$ is at most exponential in the number of nodes n of the graph. Furthermore, any given graph can at most produce a number of combinatorially distinct triangulations exponential in n. □

[1] See [114] for an overview on graph manifolds and Seifert fiber spaces.

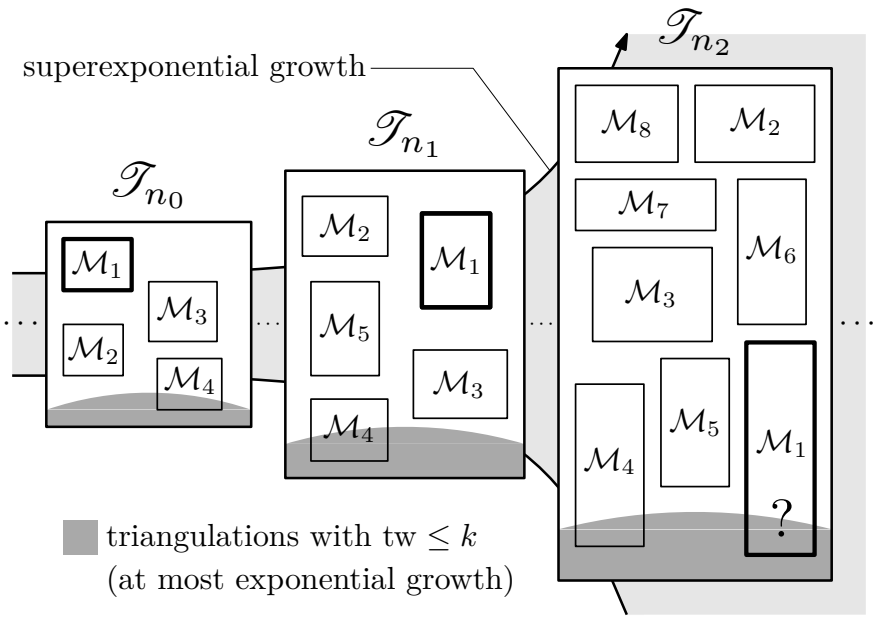

Figure B.1: As n grows larger, the triangulations of treewidth at most k represent a rapidly decreasing fraction of the set $\mathscr{T}_n$ of all ($\leq n$)-tetrahedra triangulations of 3-manifolds

Let $\boxed{\mathcal{M}_i} \subset \mathscr{T}_{n_j}$ denote the set of triangulations of the manifold $\mathcal{M}_i$ with at most n_j tetrahedra. The main question we are investigating in this article is the following:

Given a 3-manifold $\mathcal{M}$, does there always exist some $n_{\mathcal{M}} \in \mathbb{N}$, such that the set of triangulations $\boxed{\mathcal{M}}$ in $\mathcal{T}_{n_{\mathcal{M}}}$ overlaps with the region of treewidth $\leq k$ triangulations? In other words, does the set of triangulations of every 3-manifold eventually behave like the one of $\mathcal{M}_1$ on Figure B.1? Theorem 1.2 answers this question in the negative in general.

B.2 Constructing High-Treewidth Triangulations

Proposition B.3. *Every 3-manifold admits a triangulation of arbitrarily high treewidth.*

Sketch of the proof. Treewidth is monotone with respect to taking subgraphs [15, Lemma 11 (Scheffler)], hence it is sufficient to exhibit high-treewidth triangulations for the 3-ball, which then can be connected (via the 'connected sum' operation) to any triangulation.

For every $k \in \mathbb{N}$, however, it is easy to construct a triangulation of the 3-ball, whose dual graph contains the $k \times k$-grid as a minor (Figures B.2 and B.3). As treewidth is minor-monotone [15, Lemma 16] and $\mathrm{tw}\,(k \times k\text{-grid}) = k$ [15, Section 13.2], the result follows. □

Remark B.4. There is another approach [6], making use of the existence of arbitrarily large *simplicial* 2*-neighborly triangulations* of 3-manifolds (cf. [127], and [150, Section 7]).

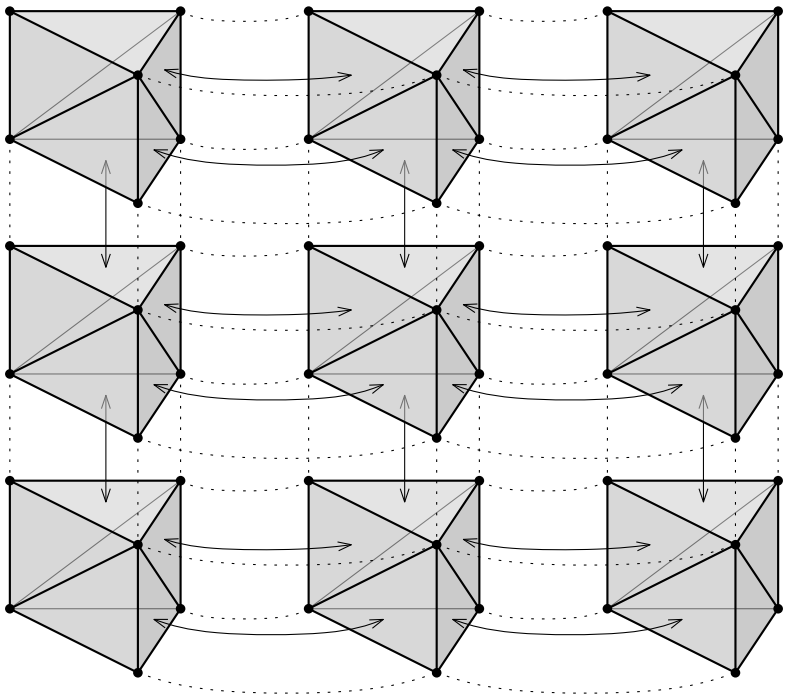

Figure B.2: A "grid-like" triangulation $\mathcal{T}_{3\times 3}$ of the 3-ball

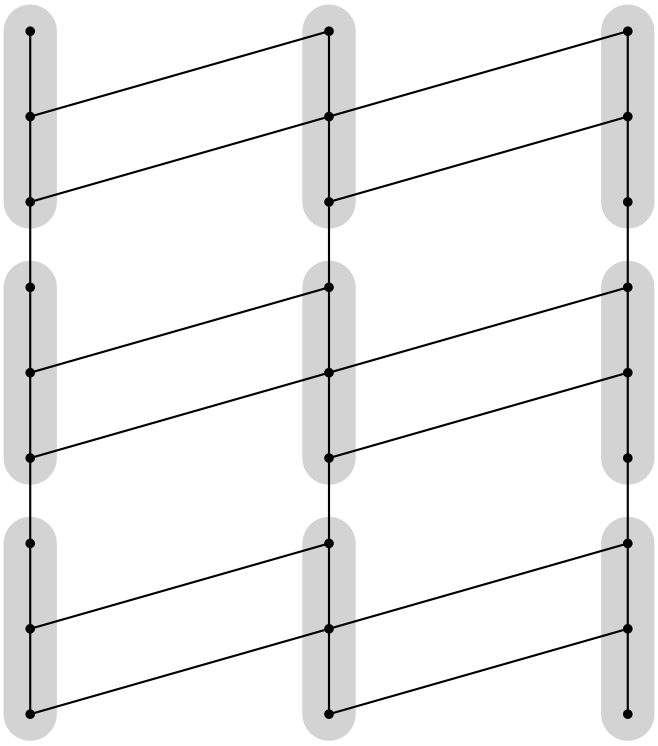

Figure B.3: The dual graph of $\mathcal{T}_{3\times 3}$ from Figure B.2 contains the 3×3 grid as a minor

C The 1-Tetrahedron Layered Solid Torus

Given a tetrahedron Δ, the single face-gluing $\varphi\colon \Delta(012) \to \Delta(231)$ (cf. Figure 5.7) yields a 3-manifold with boundary homeomorphic to a solid torus. Here we aim to shed some light on this simple construction, which, at first, can be quite perplexing.

The idea is to break up the face-identification φ into two smaller steps. To this end, consider the "thickening" $\boldsymbol{M}$ of a Möbius band M, i.e., the orientable I-bundle with base space M. Note that the fiber-wise boundary A of $\boldsymbol{M}$ is homeomorphic to an annulus and covers M twice under the projection $p\colon \boldsymbol{M} \to M$ onto the base space (Figure C.1).

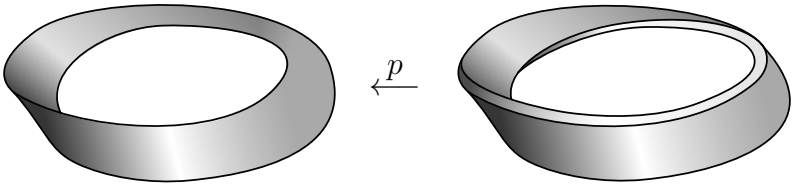

Figure C.1: (i) A Möbius band M, and (ii) the orientable I-bundle $\boldsymbol{M}$ over M

Consider a 2-triangle triangulation of A (Figure C.2). Layering the tetrahedron Δ on edge 1 of A, and composing this layering with the projection p yields a solid torus. Observe that this composition agrees with the face-gluing φ above.

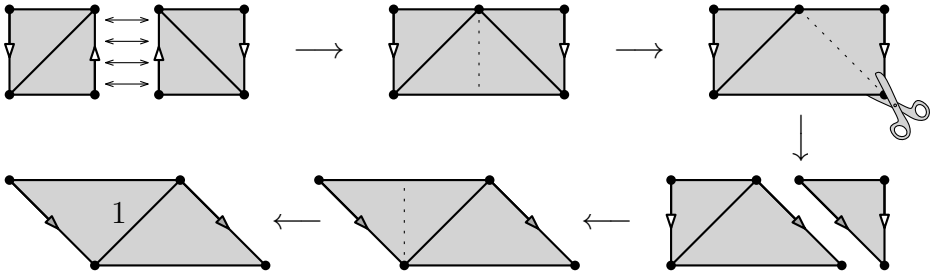

Figure C.2: Triangulating the orientable 2-cover A of the Möbius band M

D An Algorithmic Aspect of Layered Triangulations

As an application of Theorem 1.6, we describe how the machinery of layered triangulations together with work of Bell [7] can be employed to construct a "convenient" triangulation of a 3-manifold when it is presented via a mapping class $w \in \mathrm{MCG}(\mathcal{F}_g^\star)$.[1] This triangulation can then be used—as an input of existing FPT-algorithms, e.g., [29, 31, 32, 33, 100]—to compute difficult properties of $\mathcal{M}$ in running time singly exponential in g and linear in the complexity of the presentation—for some reasonable definition of complexity. First, we introduce the additional background for the statement and proof of Theorem D.2.

The mapping class group. Recall the definition of a Heegaard splitting from Section 3.5. For the study of such splittings of a given genus g, it is crucial to get a grasp on isotopy classes of their attaching maps. To this end, let $\mathcal{F}_g^\star$ be the closed orientable genus g surface $\mathcal{F}_g$ with one *marked point* $\star \in \mathcal{F}_g$ (for reasons provided later), and let $\mathrm{Homeo}^+(\mathcal{F}_g^\star)$ denote the group of orientation-preserving homeomorphisms $f\colon \mathcal{F}_g \to \mathcal{F}_g$ with $f(\star) = \star$. The *mapping class group* $\mathrm{MCG}(\mathcal{F}_g^\star)$ consists of the isotopy classes (also called *mapping classes*) of $\mathrm{Homeo}^+(\mathcal{F}_g^\star)$, where isotopies are required to fix $\star$.

The group $\mathrm{MCG}(\mathcal{F}_g^\star)$ can be generated by some "elementary" homeomorphisms: Let $c \subset \mathcal{F}_g$ be a *non-separating simple closed curve* (i.e., an embedding of the circle which does not split the surface into two connected components). Informally, a *Dehn twist along* c is a homeomorphism $\tau_c\colon \mathcal{F}_g \to \mathcal{F}_g$ where we first cut $\mathcal{F}_g$ along c, twist one of the ends by 2π, and then glue them back together [42]. A commonly used—although non-minimal [65], cf. [74, 85]—set of Dehn twists to generate $\mathrm{MCG}(\mathcal{F}_g^\star)$ is given by Theorem D.1. For more background, we refer to [49, Chapters 2 & 4].

[1] $\mathcal{M} = \mathcal{H} \cup_f \mathcal{H}'$, where $\mathcal{H}$ and $\mathcal{H}'$ are genus g handlebodies and f belongs to the isotopy class w.

Theorem D.1 (Lickorish [92])**.** *The group* $\mathrm{MCG}(\mathcal{F}_g^\star)$ *is generated by the Dehn twists along the simple closed curves* α_i, β_j $(1 \le i, j \le g)$ *and* γ_k $(1 \le k \le g-1)$*, as shown in Figure D.1.*

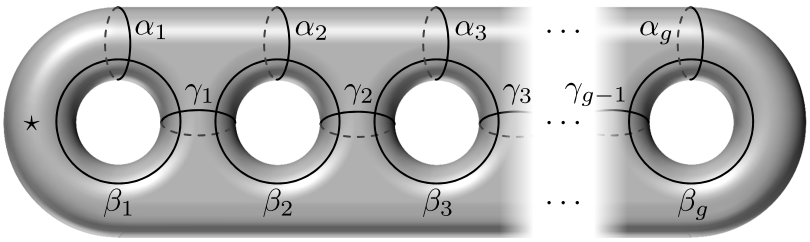

Figure D.1: The marked surface $\mathcal{F}_g^\star$ with $3g-1$ Dehn twists generating $\mathrm{MCG}(\mathcal{F}_g^\star)$

Flips. Recall the definition of a layered triangulation from Section 4.2. Layered triangulations provide the following combinatorial view on the mapping class group. Let $\mathcal{T}$ be a layered triangulation of a handlebody $\mathcal{H}$. By layering a tetrahedron Δ onto an edge $e \in \mathcal{S} = \partial\mathcal{T}$, we obtain another triangulation $\mathcal{T}' = \mathcal{T} \cup \Delta$ of $\mathcal{H}$, where, combinatorially, $\mathcal{S}' = \partial\mathcal{T}'$ is related to $\mathcal{S}$ via an *(edge) flip* $e \leftrightarrow e'$, see Figure 4.8.

Theorem 4.3 implies that for any homeomorphism $f\colon \partial\mathcal{H} \to \partial\mathcal{H}$, there is a sequence $\mathcal{T}^{(0)} \to \cdots \to \mathcal{T}^{(i+1)} = \mathcal{T}^{(i)} \cup \Delta \to \cdots \to \mathcal{T}^{(m)}$ of layered triangulations of $\mathcal{H}$, descending to a *flip sequence* $\mathcal{S}^{(0)} \to \cdots \to \mathcal{S}^{(i)} = \partial\mathcal{T}^{(i)} \to \cdots \to \mathcal{S}^{(m)}$ of one-vertex triangulations of $\partial\mathcal{H}$, such that f, when considered as a map $\mathcal{S}^{(0)} \to \mathcal{S}^{(m)}$, is a simplicial isomorphism.

We can now state and prove the main theorem of this section.

Theorem D.2. *Let* X *be a set of Dehn twists generating* $\mathrm{MCG}(\mathcal{F}_g^\star)$*, and* $\mathcal{M}$ *be a* 3*-manifold given by a word* $w \in \langle X \rangle$ *(representing a mapping class). Then we can construct a layered triangulation of* $\mathcal{M}$ *of size* $O(g + K|w|)$ *and cutwidth* $\le 4g-2$ *in time* $O\left(K(|X| + |w|)\right)$*, where* $|w|$ *denotes the length of* w*, and* K *is a constant only depending on* g *and* X*.*

Proof. Take a minimal layered triangulation $\mathcal{T}'$ of the genus g handlebody (e.g., the one shown in Figure 4.9). Then $\partial\mathcal{T}'$ is a one-vertex triangulation of $\mathcal{F}_g^\star$ (and its vertex is the "marked point" $\star$). For every Dehn twist $\tau_c \in X$, let k_c be the geometric intersection number of the edges of $\partial\mathcal{T}'$ with the curve $c \subset \partial\mathcal{T}'$ defining τ_c. Computing k_c is immediate due to the so-called *bigon criterion*, see [7, Lemma 2.4.1]. Moreover, according to [7,

Theorem 2.4.6], a flip sequence of length $O(k_c)$ realizing τ_c simplicially (cf. Section 4.2 for definitions) can be found in $O(k_c)$ steps. Let $K = \max\{k_c : \tau_c \in X\}$. It follows that we can compute the flip sequences for all generators in X in time $O(K|X|)$.

With this setup it is straightforward to construct a layered triangulation $\mathcal{T}$ of $\mathcal{M}$: start with $\mathcal{T}'$ and for every letter ℓ in w (i.e., $\ell \in X$), perform at most K layerings specified by the respective precomputed flip sequence; denote the resulting triangulation by $\mathcal{T}''$. This can be done in $O(K|w|)$ steps altogether. $\mathcal{T}$ is then completed by gluing a copy of $\mathcal{T}'$ to $\mathcal{T}''$.

By construction, the number of tetrahedra in $\mathcal{T}$ is at most $6g - 4 + K|w|$, and, following the proof of Theorem 1.6, $\mathrm{cw}(\Gamma(\mathcal{T})) \leq 4g - 2$. The overall running time is $O(K(|X| + |w|))$. □

Remark D.3. In this setting, it is natural to ask the following. Given a 3-manifold $\mathcal{M}$ of Heegaard genus $\mathfrak{g}(\mathcal{M}) = g$, what is the minimum length a word $w \in \mathrm{MCG}(\mathcal{F}_g^\star)$ (with respect to some generating set X) that realizes $\mathcal{M}$? If we choose X to be the set given by Theorem D.1, we obtain the so-called *Heegaard–Lickorish complexity* of $\mathcal{M}$ [36].

Remark D.4. Theorem D.2 gives an upper bound on the size of a triangulated Heegaard splitting in terms of the genus of the splitting and the length of the word corresponding to the gluing mapping class. It is worth noting that this construction has been implemented by Bell, Hall and Schleimer in the software *Twister* [8], which has also been included in the widely-used computational topology software bundle *SnapPy* [39].

E Generating Treewidth Two Triangulations Using *Regina*

The following Python functions can be executed using *Regina*'s Python interface or the shell environment within the *Regina* GUI. See [27, 28] for more information.

First note that layered solid tori (of treewidth one) are readily available from within *Regina*. Given an n-tetrahedra triangulation `t`, we can insert a layered solid torus of type $\mathrm{LST}(a, b, a+b)$ at the end of `t` using the command `t.insertLayeredSolidTorus(a,b)` with tetrahedra $\Delta_n, \ldots, \Delta_{n+k}$. Its boundary is always given by the triangles $\Delta_n(012)$ and $\Delta_n(013)$ and the unique edge only contained in Δ_n is thus (01).

Moreover, there is also the possibility of generating triangulations of the orientable Seifert fibered spaces over the sphere (with treewidth two), as well as layered lens spaces (with treewidth one), out-of-the-box using *Regina*.

```
# 3-punctured sphere x S^1
def prism1():
  t = Triangulation3()
  t.newSimplex()
  t.newSimplex()
  t.newSimplex()
  t.simplex(0).join(3,t.simplex(1),NPerm4())
  t.simplex(1).join(2,t.simplex(2),NPerm4())
  t.simplex(2).join(1,t.simplex(0),NPerm4(3,0,1,2))
  return t

# Moebius strip x~ S^1
```

```
def moebius():
  t = Triangulation3()
  t.newSimplex()
  t.newSimplex()
  t.newSimplex()
  t.simplex(0).join(0,t.simplex(1),NPerm4())
  t.simplex(1).join(1,t.simplex(2),NPerm4())
  t.simplex(0).join(1,t.simplex(1),NPerm4(0,2,3,1))
  t.simplex(1).join(3,t.simplex(2),NPerm4(2,0,1,3))
  t.simplex(0).join(2,t.simplex(2),NPerm4(1,3,0,2))
  return t

# Moebius strip union 3-punctured sphere x~ S^1
def ext_moebius():
  t = Triangulation3()
  t.insertTriangulation(moebius())
  t.insertTriangulation(prism1())
  t.simplex(0).join(3,t.simplex(5),NPerm4(2,0,1,3))
  t.simplex(2).join(2,t.simplex(3),NPerm4())
  return t

# Non-orientable genus g surface x~ S^1
def nonor_bundle(g):
  t = Triangulation3()
  for i in range(g):
    t.insertTriangulation(ext_moebius())
  for i in range(g-1):
    if i%2==0:
      t.simplex(6*i+4).join(0,t.simplex(6*i+10),NPerm4())
      t.simplex(6*i+5).join(0,t.simplex(6*i+11),NPerm4())
    if i%2==1:
      t.simplex(6*i+3).join(1,t.simplex(6*i+9),NPerm4())
```

```
      t.simplex(6*i+4).join(1,t.simplex(6*i+10),NPerm4())
  return t

# r-punctured sphere x S^1
def disk(r):
  if r < 0:
    return None
  if r == 0:
    t = prism1()
    t.insertLayeredSolidTorus(0,1)
    t.simplex(3).join(3,t.simplex(2),NPerm4(2,0,1,3))
    t.simplex(3).join(2,t.simplex(0),NPerm4(1,3,2,0))
    t.insertLayeredSolidTorus(0,1)
    t.simplex(6).join(3,t.simplex(0),NPerm4(3,2,0,1))
    t.simplex(6).join(2,t.simplex(1),NPerm4(0,3,1,2))
    t.insertLayeredSolidTorus(0,1)
    t.simplex(9).join(3,t.simplex(2),NPerm4(3,2,1,0))
    t.simplex(9).join(2,t.simplex(1),NPerm4(2,1,0,3))
    return t
  if r == 1:
    t = prism1()
    t.insertLayeredSolidTorus(0,1)
    t.simplex(3).join(3,t.simplex(2),NPerm4(2,0,1,3))
    t.simplex(3).join(2,t.simplex(0),NPerm4(1,3,2,0))
    t.insertLayeredSolidTorus(0,1)
    t.simplex(6).join(3,t.simplex(0),NPerm4(3,2,0,1))
    t.simplex(6).join(2,t.simplex(1),NPerm4(0,3,1,2))
    return t
  if r == 2:
    t = prism1()
    t.insertLayeredSolidTorus(0,1)
    t.simplex(3).join(3,t.simplex(2),NPerm4(2,0,1,3))
```

```
    t.simplex(3).join(2,t.simplex(0),NPerm4(1,3,2,0))
    return t
  if r >= 3:
    t = Triangulation3()
    for i in range(r-2):
      t.insertTriangulation(prism1())
    for i in range(r-3):
      if i%2 == 0:
        t.simplex(3*i+1).join(0,t.simplex(3*i+4),NPerm4())
        t.simplex(3*i+2).join(0,t.simplex(3*i+5),NPerm4())
      if i%2 == 1:
        t.simplex(3*i).join(1,t.simplex(3*i+3),NPerm4())
        t.simplex(3*i+1).join(1,t.simplex(3*i+4),NPerm4())
    return t

# Non-orientable genus g surface x~ S^1 with r punctures
def nonor_bundle_punct(g,r):
  t = Triangulation3()
  t.insertTriangulation(nonor_bundle(g))
  t.insertTriangulation(disk(r))
  t.simplex(3).join(1,t.simplex(6*g),NPerm4())
  t.simplex(4).join(1,t.simplex(6*g+1),NPerm4())
  return t

# Example
t = nonor_bundle_punct(6,12)
print t.detail()
```

Bibliography

[1] I. Agol. Small 3-manifolds of large genus. *Geom. Dedicata*, 102:53–64, 2003. doi:10.1023/B:GEOM.0000006584.85248.c5, MR:2026837, Zbl:1039.57008.

[2] I. Agol. Tameness of hyperbolic 3-manifolds, 2004. 22 pages, 1 figure. arXiv:math/0405568.

[3] S. Arnborg, D. G. Corneil, and A. Proskurowski. Complexity of finding embeddings in a k-tree. *SIAM J. Algebr. Discrete Methods*, 8(2):277–284, 1987. doi:10.1137/0608024, MR:881187, Zbl:0611.05022.

[4] M. Aschenbrenner, S. Friedl, and H. Wilton. *3-Manifold Groups*, volume 20 of *EMS Ser. Lect. Math.* Eur. Math. Soc. (EMS), Zürich, 2015. doi:10.4171/154, MR:3444187, Zbl:1326.57001.

[5] D. Bachman, R. Derby-Talbot, and E. Sedgwick. Computing Heegaard genus is NP-hard. In *A Journey Through Discrete Mathematics: A Tribute to Jiří Matoušek*, pages 59–87. Springer, Cham, 2017. doi:10.1007/978-3-319-44479-6, MR:3726594, Zbl:1388.57020.

[6] B. Bagchi, B. A. Burton, B. Datta, N. Singh, and J. Spreer. Efficient algorithms to decide tightness. In *32nd Int. Symp. Comput. Geom. (SoCG 2016)*, volume 51 of *LIPIcs. Leibniz Int. Proc. Inf.*, pages 12:1–12:15. Schloss Dagstuhl–Leibniz-Zent. Inf., 2016. doi:10.4230/LIPIcs.SoCG.2016.12, MR:3540855, Zbl:1387.52008.

[7] M. C. Bell. *Recognising Mapping Classes*. PhD thesis, Univ. Warwick, June 2015. URL: https://wrap.warwick.ac.uk/77123.

[8] M. C. Bell, T. Hall, and S. Schleimer. Twister (computer software), 2008–2014. Version 2.4.1. URL: https://bitbucket.org/Mark_Bell/twister.

[9] L. Bessières, G. Besson, S. Maillot, M. Boileau, and J. Porti. *Geometrisation of 3-Manifolds*, volume 13 of *EMS Tracts Math.* Eur. Math. Soc. (EMS), Zürich, 2010. doi:10.4171/082, MR:2683385, Zbl:1244.57003.

[10] M. Bestvina. Geometric group theory and 3-manifolds hand in hand: the fulfillment of Thurston's vision. *Bull. Amer. Math. Soc. (N.S.)*, 51(1):53–70, 2014. doi:10.1090/S0273-0979-2013-01434-4, MR:3119822, Zbl:1286.57013S.

[11] D. Bienstock. On embedding graphs in trees. *J. Comb. Theory, Ser. B*, 49(1):103–136, 1990. doi:10.1016/0095-8956(90)90066-9, MR:1056822, Zbl:0646.05025.

[12] D. Bienstock. Graph searching, path-width, tree-width and related problems. In *Reliability Of Computer And Communication Networks*, volume 5 of *DIMACS Ser. Discrete Math. Theor. Comput. Sci.*, pages 33–50. Amer. Math. Soc., Providence, RI, 1991. doi:10.1090/dimacs/005/02, MR:1119138, Zbl:0777.05090.

[13] R. H. Bing. An alternative proof that 3-manifolds can be triangulated. *Ann. Math. (2)*, 69:37–65, 1959. doi:10.2307/1970092, MR:0100841, Zbl:0106.16604.

[14] H. L. Bodlaender. A tourist guide through treewidth. *Acta Cybern.*, 11(1-2):1–21, 1993. URL: https://www.inf.u-szeged.hu/actacybernetica/edb/vol11n1_2/Bodlaender_1993_ActaCybernetica.xml, MR:1268488, Zbl:0804.68101.

[15] H. L. Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theor. Comput. Sci.*, 209(1-2):1–45, 1998. doi:10.1016/S0304-3975(97)00228-4, MR:1647486, Zbl:0912.68148.

[16] H. L. Bodlaender. Discovering treewidth. In *Proc. 31st Conf. Curr. Trends Theory Pract. Comput. Sci. (SOFSEM 2005), Liptovský Ján, Slovakia, January 22–28, 2005*, pages 1–16, 2005. doi:10.1007/978-3-540-30577-4_1, Zbl:1117.68451.

[17] H. L. Bodlaender. Fixed-parameter tractability of treewidth and pathwidth. In *The Multivariate Algorithmic Revolution and Beyond - Essays Dedicated to Michael R. Fellows on the Occasion of His 60th Birthday*, volume 7370 of *Lect. Notes Comput. Sci.*, pages 196–227. Springer, Berlin, Heidelberg, 2012. doi:10.1007/978-3-642-30891-8_12, Zbl:1358.68119.

[18] H. L. Bodlaender, B. Burton, F. V. Fomin, and A. Grigoriev. Knot diagrams of treewidth two, 2019. 19 pages, 16 figures. arXiv:1904.03117.

[19] H. L. Bodlaender and A. M. C. A. Koster. Combinatorial optimization on graphs of bounded treewidth. *Comput. J.*, 51(3):255–269, 2008. doi:10.1093/comjnl/bxm037.

[20] H. L. Bodlaender and D. M. Thilikos. Treewidth for graphs with small chordality. *Discrete Appl. Math.*, 79(1–3):45–61, 1997. doi:10.1016/S0166-218X(97)00031-0, MR:1478240, Zbl:0895.68113.

[21] M. Boileau and H. Zieschang. Heegaard genus of closed orientable Seifert 3-manifolds. *Invent. Math.*, 76(3):455–468, 1984. doi:10.1007/BF01388469, MR:746538, Zbl:0538.57004.

[22] F. Bonahon. Geometric structures on 3-manifolds. In *Handbook of Geometric Topology*, pages 93–164. North-Holland, Amsterdam, 2001. doi:10.1016/B978-044482432-5/50004-4, MR:1886669, Zbl:0997.57032.

[23] J. F. Brock, R. D. Canary, and Y. N. Minsky. The classification of Kleinian surface groups, II: The ending lamination conjecture. *Ann. of Math. (2)*, 176(1):1–149, 2012. doi:10.4007/annals.2012.176.1.1, MR:2925381, Zbl:1253.57009.

[24] B. A. Burton. *Minimal Triangulations and Normal Surfaces.* PhD thesis, Univ. Melbourne, September 2003. URL: https://people.smp.uq.edu.au/BenjaminBurton/papers/2003-thesis.html.

[25] B. A. Burton. Detecting genus in vertex links for the fast enumeration of 3-manifold triangulations. In *Proc. Int. Symp. Symb. Alg. Comput. (ISSAC 2011), San Jose, CA, USA, June 7–11, 2011*, pages 59–66, 2011. doi:10.1145/1993886.1993901, MR:2895195, Zbl:1323.68537.

[26] B. A. Burton. The Pachner graph and the simplification of 3-sphere triangulations. In *Proc. 27th ACM Symp. Comput. Geom. (SoCG 2011), Paris, France, June 13–15, 2011*, pages 153–162. ACM, New York, 2011. doi:10.1145/1998196.1998220, MR:2919606, Zbl:1283.05065.

[27] B. A. Burton. Computational topology with Regina: algorithms, heuristics and implementations. In *Geometry and Topology Down Under*, volume 597 of *Contemp. Math.*, pages 195–224. Amer. Math. Soc., Providence, RI, 2013. doi:10.1090/conm/597/11877, MR:3186674, Zbl:1279.57004.

[28] B. A. Burton, R. Budney, W. Pettersson, et al. Regina: Software for low-dimensional topology, 1999–2019. Version 5.1. URL: https://regina-normal.github.io.

[29] B. A. Burton and R. G. Downey. Courcelle's theorem for triangulations. *J. Comb. Theory, Ser. A*, 146:264–294, 2017. doi:10.1016/j.jcta.2016.10.001, MR:3574232, Zbl:1353.05122.

[30] B. A. Burton, H. Edelsbrunner, J. Erickson, and S. Tillmann, editors. *Computational Geometric and Algebraic Topology*, volume 12 of *Oberwolfach Rep.* EMS Publ. House, 2015. doi:10.4171/OWR/2015/45, Zbl:1380.00045.

[31] B. A. Burton, T. Lewiner, J. Paixão, and J. Spreer. Parameterized complexity of discrete Morse theory. *ACM Trans. Math. Softw.*, 42(1):6:1–6:24, 2016. doi:10.1145/2738034, MR:3472422, Zbl:1347.68165.

[32] B. A. Burton, C. Maria, and J. Spreer. Algorithms and complexity for Turaev–Viro invariants. *J. Appl. Comput. Topol.*, 2(1–2):33–53, 2018. doi:10.1007/s41468-018-0016-2, MR:3873178, Zbl:07089248.

[33] B. A. Burton and J. Spreer. The complexity of detecting taut angle structures on triangulations. In *Proc. 24th Annu. ACM-SIAM Symp. Discrete Algorithms (SODA 2013), New Orleans, LA, USA, January 6–8, 2013*, pages 168–183, 2013. doi:10.1137/1.9781611973105.13, MR:3185388, Zbl:1421.68161.

[34] D. Calegari and D. Gabai. Shrinkwrapping and the taming of hyperbolic 3-manifolds. *J. Amer. Math. Soc.*, 19(2):385–446, 2006. doi:10.1090/S0894-0347-05-00513-8, MR:2188131, Zbl:1090.57010.

[35] A. J. Casson and C. McA. Gordon. Reducing Heegaard splittings. *Topology Appl.*, 27(3):275–283, 1987. doi:10.1016/0166-8641(87)90092-7, MR:918537, Zbl:0632.57010.

[36] J. C. Cha. Complexities of 3-manifolds from triangulations, Heegaard splittings and surgery presentations. *Q. J. Math.*, 69(2):425–442, 2018. doi:10.1093/qmath/hax041, MR:3811586, Zbl:1398.57031.

[37] F. R. K. Chung and P. D. Seymour. Graphs with small bandwidth and cutwidth. *Discrete Math.*, 75(1–3):113–119, 1989. doi:10.1016/0012-365X(89)90083-6, MR:1001391, Zbl:0668.05039.

[38] B. Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H, MR:1042649, Zbl:0722.03008.

[39] M. Culler, N. M. Dunfield, M. Goerner, J. R. Weeks, et al. SnapPy, a computer program for studying the geometry and topology of 3-manifolds, 2009–2019. Version 2.7. URL: http://snappy.computop.org.

[40] M. Cygan, F. V. Fomin, Ł. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, Cham, 2015. doi:10.1007/978-3-319-21275-3, MR:3380745, Zbl:1334.90001.

[41] A. de Mesmay, J. Purcell, S. Schleimer, and E. Sedgwick. On the tree-width of knot diagrams. *J. Comput. Geom.*, 10(1):164–180, 2019. doi:10.20382/jocg.v10i1a6, MR:3957223, Zbl:07084708.

[42] M. Dehn. Die Gruppe der Abbildungsklassen: Das arithmetische Feld auf Flächen. *Acta Math.*, 69(1):135–206, 1938. doi:10.1007/BF02547712, MR:1555438, Zbl:0019.25301.

[43] R. Derby-Talbot. Stabilization, amalgamation and curves of intersection of Heegaard splittings. *Algebr. Geom. Topol.*, 9(2):811–832, 2009. doi:10.2140/agt.2009.9.811, MR:2505126, Zbl:1176.57021.

[44] The Sage Developers. *Sage Mathematics Software System (Version 7.6)*, 2017. URL: http://www.sagemath.org, doi:10.5281/zenodo.820864.

[45] J. Díaz, J. Petit, and M. J. Serna. A survey of graph layout problems. *ACM Comput. Surv.*, 34(3):313–356, 2002. doi:10.1145/568522.568523.

[46] R. Diestel. *Graph Theory*, volume 173 of *Grad. Texts Math.* Springer, Berlin, 5th edition, 2017. doi:10.1007/978-3-662-53622-3, MR:3644391, Zbl:1375.05002.

[47] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Monogr. Comput. Sci. Springer-Verlag New York, Inc., 1999. doi:10.1007/978-1-4612-0515-9, MR:1656112, Zbl:0914.68076.

[48] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Texts Comput. Sci. Springer, London, 2013. doi:10.1007/978-1-4471-5559-1, MR:3154461, Zbl:1358.68006.

[49] B. Farb and D. Margalit. *A Primer on Mapping Class Groups*, volume 49 of *Princeton Math. Ser.* Princeton Univ. Press, Princeton, NJ, 2012. doi:10.1515/9781400839049, MR:2850125, Zbl:1245.57002.

[50] R. Frigerio, B. Martelli, and C. Petronio. Complexity and Heegaard genus of an infinite class of compact 3-manifolds. *Pacific J. Math.*, 210(2):283–297, 2003. doi:10.2140/pjm.2003.210.283, MR:1988535, Zbl:1061.57017.

[51] D. Gabai. Foliations and the topology of 3-manifolds. III. *J. Differ. Geom.*, 26(3):479–536, 1987. doi:10.4310/jdg/1214441488, MR:910018, Zbl:0639.57008.

[52] J. Gallier and D. Xu. *A Guide to the Classification Theorem for Compact Surfaces*, volume 9 of *Geom. Comput.* Springer-Verlag Berlin Heidelberg, 2013. doi:10.1007/978-3-642-34364-3, MR:3026641, Zbl:1270.57001.

[53] F. Gavril. Some NP-complete problems on graphs. In *Proc. 1977 Conf. on Inf. Sci. Syst.*, pages 91–95. Johns Hopkins Univ., 1977.

[54] W. Haken. Theorie der Normalflächen. *Acta Math.*, 105:245–375, 1961. doi:10.1007/BF02559591, MR:0141106, Zbl:0100.19402.

[55] R. S. Hamilton. Three-manifolds with positive Ricci curvature. *J. Differ. Geom.*, 17(2):255–306, 1982. doi:10.4310/jdg/1214436922, MR:664497, Zbl:0504.53034.

[56] J. Hass, J. C. Lagarias, and N. Pippenger. The computational complexity of knot and link problems. *J. ACM*, 46(2):185–211, 1999. doi:10.1145/301970.301971, MR:1693203, Zbl:1065.68667.

[57] A. Hatcher. Notes on basic 3-manifold topology. 2007. URL: https://pi.math.cornell.edu/~hatcher/3M/3Mfds.pdf.

[58] A. E. Hatcher. On the boundary curves of incompressible surfaces. *Pacific J. Math.*, 99(2):373–377, 1982. doi:10.2140/pjm.1982.99.373, MR:658066, Zbl:0502.57005.

[59] P. Heegaard. Sur l'"Analysis situs". *Bull. Soc. Math. France*, 44:161–242, 1916. doi:10.24033/bsmf.968, MR:1504754.

[60] J. Hempel. *3-Manifolds*. AMS Chelsea Publ., Providence, RI, 2004. Reprint of the 1976 original. doi:10.1090/chel/349, MR:2098385, Zbl:1058.57001.

[61] P. Hlinený, S. Oum, D. Seese, and G. Gottlob. Width parameters beyond tree-width and their applications. *Comput. J.*, 51(3):326–362, 2008. doi:10.1093/comjnl/bxm052.

[62] D. Hoffoss and J. Maher. Morse area and Scharlemann–Thompson width for hyperbolic 3-manifolds. *Pacific J. Math.*, 281(1):83–102, 2016. doi:10.2140/pjm.2016.281.83, MR:3459967, Zbl:1335.57027.

[63] D. Hoffoss and J. Maher. Morse functions to graphs and topological complexity for hyperbolic 3-manifolds, 2017. 21 pages, 1 figure. arXiv:1708.04140.

[64] H. Howards, Y. Rieck, and J. Schultens. Thin position for knots and 3-manifolds: a unified approach. In *Workshop on Heegaard Splittings*, volume 12 of *Geom. Topol. Monogr.*, pages 89–120. Geom. Topol. Publ., Coventry, 2007. doi:10.2140/gtm.2007.12.89, MR:2408244, Zbl:1152.57012.

[65] S. P. Humphries. Generators for the mapping class group. In *Topology of Low-Dimensional Manifolds (Proc. 2nd Sussex Conf., Chelwood Gate, 1977)*, volume 722 of *Lect. Notes Math.*, pages 44–47. Springer, Berlin, 1979. doi:10.1007/BFb0063188, MR:547453, Zbl:0732.57004.

[66] K. Huszár and J. Spreer. 3-Manifold triangulations with small treewidth. In *35th Int. Symp. Comput. Geom. (SoCG 2019)*, volume 129 of *LIPIcs. Leibniz Int. Proc. Inf.*, pages 44:1–44:20. Schloss Dagstuhl–Leibniz-Zent. Inf., 2019. doi:10.4230/LIPIcs.SoCG.2019.44, MR:3968630.

[67] K. Huszár, J. Spreer, and U. Wagner. On the treewidth of triangulated 3-manifolds. *J. Comput. Geom.*, 10(2):70–98, 2019. doi:10.20382/jogc.v10i2a5, MR:4039886, Zbl:07150581.

[68] S. V. Ivanov. The computational complexity of basic decision problems in 3-dimensional topology. *Geom. Dedicata*, 131:1–26, 2008. doi:10.1007/s10711-007-9210-4, MR:2369189, Zbl:1146.57025.

[69] W. Jaco. *Lectures on Three-Manifold Topology*, volume 43 of *CBMS Reg. Conf. Ser. Math.* Amer. Math. Soc., Providence, R.I., 1980. doi:10.1090/cbms/043, MR:565450, Zbl:0433.57001.

[70] W. Jaco, J. Johnson, J. Spreer, and S. Tillmann. Bounds for the genus of a normal surface. *Geom. Topol.*, 20(3):1625–1671, 2016. doi:10.2140/gt.2016.20.1625, MR:3523065, Zbl:1350.57024.

[71] W. Jaco and J. H. Rubinstein. Layered-triangulations of 3-manifolds, 2006. 97 pages, 32 figures. arXiv:math/0603601.

[72] W. H. Jaco and P. B. Shalen. Seifert fibered spaces in 3-manifolds. *Mem. Amer. Math. Soc.*, 21(220):viii+192, 1979. doi:10.1090/memo/0220, MR:539411, Zbl:0415.57005.

[73] K. Johannson. *Homotopy equivalences of* 3*-manifolds with boundaries*, volume 761 of *Lect. Notes Math.* Springer, Berlin, 1979. doi:10.1007/BFb0085406, MR:551744, Zbl:0412.57007.

[74] D. Johnson. The structure of the Torelli group. I. A finite set of generators for $\mathcal{I}$. *Ann. Math. (2)*, 118(3):423–442, 1983. doi:10.2307/2006977, MR:727699, Zbl:0549.57006.

[75] I. Kapovich, A. Myasnikov, P. Schupp, and V. Shpilrain. Generic-case complexity, decision problems in group theory, and random walks. *J. Algebra*, 264(2):665–694, 2003. doi:10.1016/S0021-8693(03)00167-4, MR:1981427, Zbl:1041.20021.

[76] D. A. Každan and G. A. Margulis. A proof of Selberg's hypothesis. *Math. USSR–Sbornik*, 4(1):147–152, 1968. Translation from *Mat. Sb. (N.S.)*, 75(117):163–168, 1968. Translated by: Z. Skalsky. doi:10.1070/SM1968v004n01ABEH002782, MR:0223487, Zbl:0241.22024.

[77] K. Kawarabayashi and B. Mohar. Some recent progress and applications in graph minor theory. *Graphs Combin.*, 23(1):1–46, 2007. doi:10.1007/s00373-006-0684-x, MR:2292102, Zbl:1114.05096.

[78] D. Ketover, Y. Liokumovich, and A. Song. On the existence of minimal Heegaard surfaces, 2019. 51 pages, 4 figures. arXiv:1911.07161.

[79] N. G. Kinnersley. The vertex separation number of a graph equals its path-width. *Inf. Process. Lett.*, 42(6):345–350, 1992. doi:10.1016/0020-0190(92)90234-M, MR:1178214, Zbl:0764.68121.

[80] R. Kirby and P. Melvin. Local surgery formulas for quantum invariants and the Arf invariant. In *Proceedings of the Casson Fest*, volume 7 of *Geom. Topol. Monogr.*, pages 213–233. Geom. Topol. Publ., Coventry, 2004. doi:10.2140/gtm.2004.7.213, MR:2172485, Zbl:1087.57006.

[81] B. Kleiner and J. Lott. Notes on Perelman's papers. *Geom. Topol.*, 12(5):2587–2855, 2008. doi:10.2140/gt.2008.12.2587, MR:2460872, Zbl:1204.53033.

[82] T. Kloks. *Treewidth: Computations and Approximations*, volume 842 of *Lect. Notes Comput. Sci.* Springer, 1994. doi:10.1007/BFb0045375, MR:1312164, Zbl:0825.68144.

[83] H. Kneser. Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten. *Jahresber. Dtsch. Math.-Ver.*, 38:248–259, 1929. URL: https://eudml.org/doc/145838, Zbl:55.0311.03.

[84] T. Kobayashi and Y. Rieck. A linear bound on the tetrahedral number of manifolds of bounded volume (after Jørgensen and Thurston). In *Topology and Geometry in Dimension Three*, volume 560 of *Contemp. Math.*, pages 27–42. Amer. Math. Soc., Providence, RI, 2011. doi:10.1090/conm/560/11089, MR:2866921, Zbl:1335.57028.

[85] M. Korkmaz. Generating the surface mapping class group by two elements. *Trans. Amer. Math. Soc.*, 357(8):3299–3310, 2005. doi:10.1090/S0002-9947-04-03605-0, MR:2135748, Zbl:1079.57018.

[86] G. Kuperberg. Knottedness is in NP, modulo GRH. *Adv. Math.*, 256:493–506, 2014. doi:10.1016/j.aim.2014.01.007, MR:3177300, Zbl:1358.68138.

[87] G. Kuperberg. Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization. *Pacific J. Math.*, 301(1):189–241, 2019. doi:10.2140/pjm.2019.301.189, MR:4007377.

[88] M. Lackenby. The efficient certification of knottedness and Thurston norm, 2016. 98 pages, 23 figures. arXiv:1604.00290.

[89] M. Lackenby. Some conditionally hard problems on links and 3-manifolds. *Discrete Comput. Geom.*, 58(3):580–595, 2017. doi:10.1007/s00454-017-9905-8, MR:3690662, Zbl:1384.57010.

[90] M. Lackenby. Algorithms in 3-manifold theory, 2020. 44 pages, 9 figures. arXiv:2002.02179.

[91] M. Lackenby and J. Purcell. Triangulation complexity of fibred 3-manifolds, 2019. 70 pages, 31 figures. arXiv:1910.10914.

[92] W. B. R. Lickorish. A representation of orientable combinatorial 3-manifolds. *Ann. Math. (2)*, 76:531–540, 1962. doi:10.2307/1970373, MR:0151948, Zbl:0106.37102.

[93] L. Lovász. Graph minor theory. *Bull. Amer. Math. Soc. (N.S.)*, 43(1):75–86, 2006. doi:10.1090/S0273-0979-05-01088-8, MR:2188176, Zbl:1082.05082.

[94] F. H. Lutz. Triangulated manifolds with few vertices: Geometric 3-manifolds. Tech. Rep. 03-36, ZIB, Takustr. 7, 14195 Berlin, 2003. urn:nbn:de:0297-zib-7583.

[95] J. A. Makowsky. Coloured Tutte polynomials and Kauffman brackets for graphs of bounded tree width. *Discrete Appl. Math.*, 145(2):276–290, 2005. doi:10.1016/j.dam.2004.01.016, MR:2113147, Zbl:1084.05505.

[96] J. A. Makowsky and J. P. Mariño. The parametrized complexity of knot polynomials. volume 67, pages 742–756. 2003. Special issue on parameterized computation and complexity. doi:10.1016/S0022-0000(03)00080-1, MR:2036511, Zbl:1093.68043.

[97] A. Marden. *Hyperbolic Manifolds: An Introduction in 2 and 3 Dimensions.* Cambridge Univ. Press, Cambridge, 2016. doi:10.1017/CBO9781316337776, MR:3586015, Zbl:1355.57002.

[98] C. Maria. Parameterized complexity of quantum invariants, 2019. arXiv:1910.00477.

[99] C. Maria and J. Purcell. Treewidth, crushing and hyperbolic volume. *Algebr. Geom. Topol.*, 19(5):2625–2652, 2019. doi:10.2140/agt.2019.19.2625, MR:4023324, Zbl:07142614.

[100] C. Maria and J. Spreer. A polynomial-time algorithm to compute Turaev–Viro invariants $\mathrm{TV}_{4,q}$ of 3-manifolds with bounded first Betti number. *Found. Comput. Math.*, pages 1–22, 2019. Online: 11.11.2019. doi:10.1007/s10208-019-09438-8.

[101] B. Martelli. *An Introduction to Geometric Topology.* CreateSpace Indep. Publ. Platf., 2016. URL: http://people.dm.unipi.it/martelli/geometric_topology.html, arXiv:1610.02592.

[102] W. S. Massey. *A Basic Course in Algebraic Topology*, volume 127 of *Grad. Texts Math.* Springer-Verlag, New York, 1991. doi:10.1007/978-1-4939-9063-4, MR:1095046, Zbl:0725.55001.

[103] S. Matveev. *Algorithmic Topology and Classification of 3-Manifolds*, volume 9 of *Algorithms Comput. Math.* Springer, Berlin, 2nd edition, 2007. doi:10.1007/978-3-540-45899-9, MR:2341532, Zbl:1128.57001.

[104] D. McCullough. Mappings of reducible 3-manifolds. In *Geometric and Algebraic Topology*, volume 18 of *Banach Cent. Publ.*, pages 61–76. PWN, Warsaw, 1986. URL: https://eudml.org/doc/267610, MR:925856, Zbl:0637.57009.

[105] C. T. McMullen. The evolution of geometric structures on 3-manifolds. *Bull. Amer. Math. Soc. (N.S.)*, 48(2):259–274, 2011. doi:10.1090/S0273-0979-2011-01329-5, MR:2774092, Zbl:1214.57017.

[106] J. Milnor. A unique decomposition theorem for 3-manifolds. *Amer. J. Math.*, 84:1–7, 1962. doi:10.2307/2372800, MR:142125, Zbl:0108.36501.

[107] J. Milnor. Towards the Poincaré conjecture and the classification of 3-manifolds. *Notices Amer. Math. Soc.*, 50(10):1226–1233, 2003. MR:2009455, Zbl:1168.57303.

[108] E. E. Moise. Affine structures in 3-manifolds. V. The triangulation theorem and Hauptvermutung. *Ann. Math. (2)*, 56:96–114, 1952. doi:10.2307/1969769, MR:0048805, Zbl:0048.17102.

[109] B. Monien and I. H. Sudborough. Min cut is NP-complete for edge weighted trees. *Theoret. Comput. Sci.*, 58(1-3):209–229, 1988. doi:10.1016/0304-3975(88)90028-X, MR:963264, Zbl:0657.68034.

[110] J. Morgan and G. Tian. *The Geometrization Conjecture*, volume 5 of *Clay Math. Monogr.* Amer. Math. Soc., Providence, RI; Clay Math. Inst., Cambridge, MA, 2014. URL: https://bookstore.ams.org/cmim-5, MR:3186136, Zbl:1302.53001.

[111] Y. Moriah and H. Rubinstein. Heegaard structures of negatively curved 3-manifolds. *Comm. Anal. Geom.*, 5(3):375–412, 1997. doi:10.4310/CAG.1997.v5.n3.a1, MR:1487722, Zbl:0890.57025.

[112] G. D. Mostow. Quasi-conformal mappings in n-space and the rigidity of hyperbolic space forms. *Publ. Math. IHÉS*, 34:53–104, 1968. URL: http://www.numdam.org/item?id=PMIHES_1968__34__53_0, MR:236383, Zbl:0189.09402.

[113] S. Norine, P. D. Seymour, R. Thomas, and P. Wollan. Proper minor-closed families are small. *J. Comb. Theory, Ser. B*, 96(5):754–757, 2006. doi:10.1016/j.jctb.2006.01.006, MR:2236510, Zbl:1093.05065.

[114] P. Orlik. *Seifert Manifolds*, volume 291 of *Lect. Notes Math.* Springer-Verlag, Berlin-New York, 1972. doi:10.1007/BFb0060329, MR:0426001, Zbl:0263.57001.

[115] M. I. Ostrovskii. Minimal congestion trees. *Discrete Math.*, 285(1–3):219–226, 2004. doi:10.1016/j.disc.2004.02.009, MR:2062845, Zbl:1051.05032.

[116] G. Perelman. The entropy formula for the Ricci flow and its geometric applications, 2002. 39 pages. arXiv:math/0211159.

[117] G. Perelman. Finite extinction time for the solutions to the ricci flow on certain three-manifolds, 2003. 7 pages. arXiv:math/0307245.

[118] G. Perelman. Ricci flow with surgery on three-manifolds, 2003. 22 pages. arXiv:math/0303109.

[119] J. T. Pitts and J. H. Rubinstein. Existence of minimal surfaces of bounded topological type in three-manifolds. In *Miniconference on geometry and partial differential equations (Canberra, 1985)*, volume 10 of *Proc. Centre Math. Anal. Austral. Nat. Univ.*, pages 163–176. Austral. Nat. Univ., Canberra, 1986. URL: https://projecteuclid.org/euclid.pcma/1416336688, MR:857665, Zbl:0602.49028.

[120] J. Porti. Geometrization of three manifolds and Perelman's proof. *Rev. R. Acad. Cienc. Exactas Fís. Nat. Ser. A Math. RACSAM*, 102(1):101–125, 2008. doi:10.1007/BF03191814, MR:2416241, Zbl:1170.57016.

[121] J. S. Purcell. Hyperbolic Knot Theory (preprint). February 2020. URL: http://users.monash.edu/~jpurcell/book/HypKnotTheory.pdf, arXiv:2002.12652.

[122] J. S. Purcell and A. Zupan. Independence of volume and genus g bridge numbers. *Proc. Amer. Math. Soc.*, 145(4):1805–1818, 2017. doi:10.1090/proc/13327, MR:3601570, Zbl:1364.57009.

[123] N. Robertson and P. D. Seymour. Graph minors. I. Excluding a forest. *J. Comb. Theory, Ser. B*, 35(1):39–61, 1983. doi:10.1016/0095-8956(83)90079-5, MR:723569, Zbl:0521.05062.

[124] N. Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986. doi:10.1016/0196-6774(86)90023-4, MR:855559.

[125] J. H. Rubinstein. An algorithm to recognize the 3-sphere. In *Proc. Int. Congr. Math., Zürich, Switzerland, August 3–11, 1994*, volume 1, pages 601–611. Birkhäuser, 1995. doi:10.1007/978-3-0348-9078-6_54, MR:1403961, Zbl:0864.57009.

[126] J. H. Rubinstein, H. Segerman, and S. Tillmann. Traversing three-manifold triangulations and spines, 2018. 41 pages, 42 figures. To appear in *L'Enseign. Math.* arXiv:1812.02806.

[127] K. S. Sarkaria. On neighbourly triangulations. *Trans. Amer. Math. Soc.*, 277(1):213–239, 1983. doi:10.2307/1999353, MR:690049, Zbl:0522.57010.

[128] N. Saveliev. *Lectures on the Topology of 3-Manifolds: An Introduction to the Casson Invariant*. De Gruyter Textbook. Walter de Gruyter & Co., Berlin, 2nd edition, 2012. doi:10.1515/9783110250367, MR:2893651, Zbl:1246.57003.

[129] M. Scharlemann. Heegaard splittings of compact 3-manifolds. In *Handbook of Geometric Topology*, pages 921–953. North-Holland, Amsterdam, 2001. doi:10.1016/B978-044482432-5/50019-6, MR:1886684, Zbl:0985.57005.

[130] M. Scharlemann, J. Schultens, and T. Saito. *Lecture Notes on Generalized Heegaard Splittings*. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2016. doi:10.1142/10019, MR:3585907, Zbl:1356.57004.

[131] M. Scharlemann and A. Thompson. Thin position for 3-manifolds. In *Geometric topology (Haifa, 1992)*, volume 164 of *Contemp. Math.*, pages 231–238. Amer. Math. Soc., Providence, RI, 1994. doi:10.1090/conm/164/01596, MR:1282766, Zbl:0818.57013.

[132] S. Schleimer. Sphere recognition lies in NP. In *Low-dimensional and symplectic topology*, volume 82 of *Proc. Sympos. Pure Math.*, pages 183–213. Amer. Math. Soc., Providence, RI, 2011. doi:10.1090/pspum/082/2768660, MR:2768660, Zbl:1250.57024.

[133] J. Schultens. The classification of Heegaard splittings for (compact orientable surface) $\times S^1$. *Proc. London Math. Soc. (3)*, 67(2):425–448, 1993. doi:10.1112/plms/s3-67.2.425, MR:1226608, Zbl:0789.57012.

[134] J. Schultens. *Introduction to 3-Manifolds*, volume 151 of *Grad. Stud. Math.* Amer. Math. Soc., Providence, RI, 2014. doi:10.1090/gsm/151, MR:3203728, Zbl:1295.57001.

[135] P. Scott. The geometries of 3-manifolds. *Bull. London Math. Soc.*, 15(5):401–487, 1983. doi:10.1112/blms/15.5.401, MR:705527, Zbl:0561.57001.

[136] P. Scott and H. Short. The homeomorphism problem for closed 3-manifolds. *Algebr. Geom. Topol.*, 14(4):2431–2444, 2014. doi:10.2140/agt.2014.14.2431, MR:3331689, Zbl:1311.57025.

[137] H. Seifert. Topologie Dreidimensionaler Gefaserter Räume. *Acta Math.*, 60(1):147–238, 1933. doi:10.1007/BF02398271, MR:1555366, Zbl:0006.08304.

[138] P. D. Seymour and R. Thomas. Call routing and the ratcatcher. *Combinatorica*, 14(2):217–241, 1994. doi:10.1007/BF01215352, MR:1289074, Zbl:0799.05057.

[139] P. B. Shalen. Hyperbolic volume, Heegaard genus and ranks of groups. In *Workshop on Heegaard Splittings*, volume 12 of *Geom. Topol. Monogr.*, pages 335–349. Geom. Topol. Publ., Coventry, 2007. doi:10.2140/gtm.2007.12.335, MR:2408254, Zbl:1140.57009.

[140] J. Stillwell. Poincaré and the early history of 3-manifolds. *Bull. Amer. Math. Soc. (N.S.)*, 49(4):555–576, 2012. doi:10.1090/S0273-0979-2012-01385-X, MR:2958930, Zbl:1255.57001.

[141] D. M. Thilikos, M. J. Serna, and H. L. Bodlaender. Constructive linear time algorithms for small cutwidth and carving-width. In *Proc. 11th Int. Conf. Algorithms Comput. (ISAAC 2000), Taipei, Taiwan, December 18–20, 2000*, pages 192–203, 2000. doi:10.1007/3-540-40996-3_17, MR:1858364, Zbl:1044.68709.

[142] D. M. Thilikos, M. J. Serna, and H. L. Bodlaender. A polynomial time algorithm for the cutwidth of bounded degree graphs with small treewidth. In *Proc. 9th Annu. Eur. Symp. Algorithms (ESA 2001), Aarhus, Denmark, August 28–31, 2001*, pages 380–390, 2001. doi:10.1007/3-540-44676-1_32, MR:1914050, Zbl:1007.05091.

[143] D. M. Thilikos, M. J. Serna, and H. L. Bodlaender. Cutwidth I: A linear time fixed parameter algorithm. *J. Algorithms*, 56(1):1–24, 2005. doi:10.1016/j.jalgor.2004.12.001, MR:2146375, Zbl:1161.68856.

[144] A. Thompson. Thin position and the recognition problem for S^3. *Math. Res. Lett.*, 1(5):613–630, 1994. doi:10.4310/MRL.1994.v1.n5.a9, MR:1295555, Zbl:0849.57009.

[145] W. Threlfall and H. Seifert. Topologische Untersuchung der Diskontinuitätsbereiche endlicher Bewegungsgruppen des dreidimensionalen sphärischen Raumes (Schluß). *Math. Ann.*, 107(1):543–586, 1933. URL: https://eudml.org/doc/159610, doi:10.1007/BF01448910, MR:1512817, Zbl:58.1203.01.

[146] W. P. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982. doi:10.1090/S0273-0979-1982-15003-0, MR:648524, Zbl:0496.57005.

[147] W. P. Thurston. *Three-Dimensional Geometry and Topology. Vol. 1*, volume 35 of *Princeton Math. Ser.* Princeton Univ. Press, Princeton, NJ, 1997. Edited by S. Levy. doi:10.1515/9781400865321, MR:1435975, Zbl:0873.57001.

[148] W. P. Thurston. The geometry and topology of three-manifolds. Electronic version 1.1, March, 2002. URL: http://library.msri.org/books/gt3m/.

[149] F. Waldhausen. Eine Klasse von 3-dimensionalen Mannigfaltigkeiten. I, II. *Invent. Math.*, 3:308–333; ibid. 4:87–117, 1967. URL: https://eudml.org/doc/141878 and https://eudml.org/doc/141884. doi:10.1007/BF01402956, MR:235576, Zbl:0168.44503.

[150] D. W. Walkup. The lower bound conjecture for 3- and 4-manifolds. *Acta Math.*, 125:75–107, 1970. doi:10.1007/BF02392331, MR:0275281, Zbl:0204.56301.

[151] R. Zentner. Integer homology 3-spheres admit irreducible representations in $\mathrm{SL}(2,\mathbb{C})$. *Duke Math. J.*, 167(9):1643–1712, 2018. doi:10.1215/00127094-2018-0004, MR:3813594, Zbl:06904637.

[152] X. Zhao. On the Nielsen numbers of slide homeomorphisms on 3-manifolds. *Topol. Appl.*, 136(1-3):169–188, 2004. doi:10.1016/S0166-8641(03)00218-9, MR:2023416, Zbl:1042.55002.