



HAL
open science

On the width of complicated JSJ decompositions

Kristóf Huszár, Jonathan Spreer

► **To cite this version:**

Kristóf Huszár, Jonathan Spreer. On the width of complicated JSJ decompositions. 39th International Symposium on Computational Geometry (SoCG 2023), Jun 2023, Dallas, United States. pp.42:1–42:18, 10.4230/LIPIcs.SoCG.2023.42. hal-04055617

HAL Id: hal-04055617

<https://hal.science/hal-04055617>

Submitted on 3 Apr 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

On the Width of Complicated JSJ Decompositions

Kristóf Huszár   

Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Jonathan Spreer   

School of Mathematics and Statistics F07, The University of Sydney, NSW 2006 Australia

Abstract

Motivated by the algorithmic study of 3-dimensional manifolds, we explore the structural relationship between the JSJ decomposition of a given 3-manifold and its triangulations. Building on work of Bachman, Derby-Talbot and Sedgwick, we show that a “sufficiently complicated” JSJ decomposition of a 3-manifold enforces a “complicated structure” for all of its triangulations. More concretely, we show that, under certain conditions, the treewidth (resp. pathwidth) of the graph that captures the incidences between the pieces of the JSJ decomposition of an irreducible, closed, orientable 3-manifold \mathcal{M} yields a linear lower bound on its treewidth $\text{tw}(\mathcal{M})$ (resp. pathwidth $\text{pw}(\mathcal{M})$), defined as the smallest treewidth (resp. pathwidth) of the dual graph of any triangulation of \mathcal{M} .

We present several applications of this result. We give the first example of an infinite family of bounded-treewidth 3-manifolds with unbounded pathwidth. We construct Haken 3-manifolds with arbitrarily large treewidth—previously the existence of such 3-manifolds was only known in the non-Haken case. We also show that the problem of providing a constant-factor approximation for the treewidth (resp. pathwidth) of bounded-degree graphs efficiently reduces to computing a constant-factor approximation for the treewidth (resp. pathwidth) of 3-manifolds.

2012 ACM Subject Classification Mathematics of computing → Geometric topology; Theory of computation → Problems, reductions and completeness; Theory of computation → Fixed parameter tractability

Keywords and phrases computational 3-manifold topology, fixed-parameter tractability, generalized Heegaard splittings, JSJ decompositions, pathwidth, treewidth, triangulations

Funding *Kristóf Huszár*: Supported by the French National Research Agency through the 3IA Côte d’Azur (ANR-19-P3IA-0002), AlgoKnot (ANR-20-CE48-0007), GrR (ANR-18-CE40-0032), and TWIN-WIDTH (ANR-21-CE48-0014) projects. A previous version of this manuscript was written during the author’s visit at Institut Henri Poincaré (IHP). The author acknowledges the hospitality of IHP (UAR 839 CNRS–Sorbonne Université) and LabEx CARMIN (ANR-10-LABX-59-01).

Jonathan Spreer: Supported by the Australian Research Council’s Discovery funding scheme (project no. DP220102588).

Acknowledgements We are grateful to Arnaud de Mesmay and to Clément Maria for their interest in our work and for inspiring discussions. We thank the anonymous referees for several suggestions to improve this article.

1 Introduction

Manifolds in geometric topology are often studied through the following two-step process. Given a piecewise linear d -dimensional manifold \mathcal{M} , first find a “suitable” triangulation of it, i.e., a decomposition of \mathcal{M} into d -simplices with “good” combinatorial properties. Then apply algorithms on this triangulation to reveal topological information about \mathcal{M} .

The work presented in this article is motivated by this process in dimension $d = 3$. Here every manifold can be triangulated [43] and questions about them typically admit algorithmic solutions [34, 38, 53].¹ At the same time, the feasibility of a particular computation can

¹ In higher dimensions none of these statements is true in general. See, e.g., [40], [42] or [45, Section 7].

greatly depend on structural properties of the triangulation in use. Over the past decade, this phenomenon was recognized and exploited in various settings, leading to *fixed-parameter tractable* (FPT) algorithms for several problems in low-dimensional topology, some of which are even known to be **NP**-hard [14, 15, 16, 17, 18].² Although these algorithms have exponential running time in the worst case, for input triangulations with *dual graph* of bounded *treewidth* they always terminate in polynomial (in most cases, linear) time.³ Moreover, some of them have implementations that are highly effective in practice, providing useful tools for researchers in low-dimensional topology [12, 13].

The theoretical efficiency of the aforementioned FPT algorithms crucially depends on the assumption that the dual graph of the input triangulation has small treewidth. To understand their scope, it is thus instructive to consider the *treewidth* $\text{tw}(\mathcal{M})$ of a *compact 3-manifold* \mathcal{M} , defined as the smallest treewidth of the dual graph of any triangulation of \mathcal{M} . Indeed, the relation between the treewidth and other quantities associated with 3-manifolds has recently been investigated in various contexts [25, 26, 27, 28, 41]. For instance, in [28] together with Wagner we have shown that the treewidth of a *non-Haken* 3-manifold is always bounded below in terms of its *Heegaard genus*. Combined with earlier work of Agol [1]—who constructed non-Haken 3-manifolds with arbitrary large Heegaard genus—this implies the existence of 3-manifolds with arbitrary large treewidth. Despite the fact that, asymptotically, most triangulations of most 3-manifolds must have dual graph of large treewidth [28, Appendix A], this collection described by Agol has remained, to this date, the only known family of 3-manifolds with arbitrary large treewidth.

The main result In this work we unravel new structural links between the triangulations of a given 3-manifold and its *JSJ decomposition* [30, 31, 32]. Employing the machinery of *generalized Heegaard splittings* [49], the results developed in [28], and building on the work of Bachman, Derby-Talbot and Sedgwick [5, 6], we show that, under suitable conditions, the dual graph of any triangulation of a given 3-manifold \mathcal{M} inherits structural properties from the *decomposition graph* that encodes the incidences between the pieces of the JSJ decomposition of \mathcal{M} . More precisely, in Section 4 we prove the following theorem.

► **Theorem 1** (Width inheritance). *For any closed, orientable and irreducible 3-manifold \mathcal{M} with sufficiently complicated⁴ torus gluings in its JSJ decomposition \mathcal{D} , the treewidth and pathwidth of \mathcal{M} and that of the decomposition graph $\Gamma(\mathcal{D})$ of \mathcal{D} satisfy*

$$\text{tw}(\Gamma(\mathcal{D})) \leq 18 \cdot (\text{tw}(\mathcal{M}) + 1) \quad (1) \quad \text{and} \quad \text{pw}(\Gamma(\mathcal{D})) \leq 4 \cdot (3 \text{pw}(\mathcal{M}) + 1). \quad (2)$$

An algorithmic construction Much work in 3-dimensional topology has been devoted to the study of 3-manifolds constructed by pasting together simpler pieces along their boundary surfaces via “sufficiently complicated” gluing maps, and to understand how different decompositions of the same 3-manifold interact under various conditions, see, e.g., [4, 5, 6, 7, 35, 39, 47, 52]. Theorem 1 allows us to leverage these results to construct 3-manifolds, where we have tight control over the treewidth and pathwidth of their triangulations [25].

By combining Theorem 1 and [4, Theorem 5.4] (cf. [6, Appendix]), in Section 5 we prove the following result.

² See [8] for an FPT algorithm checking tightness of (weak) pseudomanifolds in arbitrary dimensions.

³ The running times are given in terms of the *size* of the input triangulation, i.e., its number of tetrahedra.

⁴ The notion of “sufficiently complicated” under which we establish Theorem 1 is discussed in Section 4.

► **Theorem 2.** *There is a polynomial-time algorithm that, given an n -node graph G with maximum node-degree Δ , produces a triangulation \mathcal{T}_G of a closed 3-manifold \mathcal{M}_G , such that*

1. *the triangulation \mathcal{T}_G contains $O_\Delta(\text{pw}(G) \cdot n)$ tetrahedra,⁵*
2. *the JSJ decomposition \mathcal{D} of \mathcal{M}_G satisfies $\Gamma(\mathcal{D}) = G$, and*
3. *there exist universal constants $c, c' > 0$, such that*
 - a. *$(c/\Delta) \text{tw}(\mathcal{M}_G) \leq \text{tw}(G) \leq 18 \cdot (\text{tw}(\mathcal{M}_G) + 1)$, and*
 - b. *$(c'/\Delta) \text{pw}(\mathcal{M}_G) \leq \text{pw}(G) \leq 4 \cdot (3 \text{pw}(\mathcal{M}_G) + 1)$.*

Applications In Section 6 we present several applications of Theorem 2. First, we construct a family of bounded-treewidth 3-manifolds with unbounded pathwidth (Corollary 12). Second, we exhibit Haken 3-manifolds with arbitrary large treewidth (Corollary 13). To our knowledge, no such families of 3-manifolds had been known before. Third, we show that the problem of providing a constant-factor approximation for the treewidth (resp. pathwidth) of bounded-degree graphs reduces in polynomial time to computing a constant-factor approximation for the treewidth (resp. pathwidth) of 3-manifolds (Corollary 14). This reduction, together with previous results [46, 55, 56], suggests that this problem may be computationally hard.

Outline of the proof of Theorem 1. We now give a preview of the proof of our main result. As the arguments for showing (1) and (2) are analogous, we only sketch the proof of (1). To show that $\text{tw}(\Gamma(\mathcal{D})) \leq 18(\text{tw}(\mathcal{M}) + 1)$, we start with a triangulation \mathcal{T} of \mathcal{M} whose dual graph has minimal treewidth, i.e., $\text{tw}(\Gamma(\mathcal{T})) = \text{tw}(\mathcal{M})$. Following [28, Section 6], we construct from \mathcal{T} a generalized Heegaard splitting \mathcal{H} of \mathcal{M} , together with a *sweep-out* $\Sigma = \{\Sigma_x : x \in H\}$ along a tree H , such that the genus of any level surface Σ_x is at most $18 \cdot (\text{tw}(\Gamma(\mathcal{T})) + 1)$. If \mathcal{H} is not already *strongly irreducible*, we repeatedly perform *weak reductions* until we get a strongly irreducible generalized Heegaard splitting \mathcal{H}' with associated *sweep-out* $\Sigma' = \{\Sigma'_x : x \in H\}$ along the same tree H . Crucially, weak reductions do not increase the genera of level surfaces [49, Section 5.2], thus $18 \cdot (\text{tw}(\Gamma(\mathcal{T})) + 1)$ is still an upper bound on those in Σ' . Now, by the assumption of the JSJ decomposition of \mathcal{M} being “sufficiently complicated,” each *JSJ torus* can be isotoped in \mathcal{M} to coincide with a connected component of some *thin level* of \mathcal{H}' . This implies that, after isotopy, each level set Σ'_x is incident to at most $18 \cdot (\text{tw}(\Gamma(\mathcal{T})) + 1) + 1$ JSJ pieces of \mathcal{M} . Sweeping along H , we can construct a *tree decomposition* of $\Gamma(\mathcal{D})$ where each *bag* contains at most $18 \cdot (\text{tw}(\Gamma(\mathcal{T})) + 1) + 1$ nodes of $\Gamma(\mathcal{D})$. Hence $\text{tw}(\Gamma(\mathcal{D})) \leq 18 \cdot (\text{tw}(\Gamma(\mathcal{T})) + 1) = 18 \cdot (\text{tw}(\mathcal{M}) + 1)$. ◀

Organization of the paper In Section 2 we review the necessary background on graphs and 3-manifolds. Section 3 contains a primer on generalized Heegaard splittings, which provides us with the indispensable machinery for proving our main result (Theorem 1) in Section 4. In Section 5 we describe the algorithmic construction of 3-manifolds that “inherit” their combinatorial width from that of their JSJ decomposition graph (Theorem 2). Then, in Section 6 we present the aforementioned applications of this construction (Corollaries 12–14). The paper is concluded with a discussion and some open questions in Section 7. Selected results from 3-manifold topology we rely on are collected in the Appendix.

⁵ Similar to the standard *big-O notation*, $O_\Delta(x)$ means “a quantity bounded above by x times a constant depending on Δ .” To ensure that 3a is satisfied, but not necessarily 3b, $O_\Delta(\text{tw}(G) \cdot n)$ tetrahedra suffice.

2 Preliminaries

2.1 Graphs

A *(multi)graph* $G = (V, E)$ is a finite set $V = V(G)$ of *nodes*⁶ together with a multiset $E = E(G)$ of unordered pairs of not necessarily distinct nodes, called *arcs*. The *degree* d_v of a node $v \in V$ equals the number of arcs containing it, where loop arcs are counted twice. G is *k-regular*, if $d_v = k$ for all $v \in V$. A *tree* is a connected graph with n nodes and $n - 1$ arcs.

A *tree decomposition* of a graph $G = (V, E)$ is a pair $(\mathcal{X} = \{B_i : i \in I\}, T = (I, F))$ with *bags* $B_i \subseteq V$ and a tree $T = (I, F)$, such that **1.** $\bigcup_{i \in I} B_i = V$ (node coverage), **2.** for all arcs $\{u, v\} \in E$, there exists $i \in I$ such that $\{u, v\} \subseteq B_i$ (arc coverage), and **3.** for all $v \in V$, $T_v = \{i \in I : v \in B_i\}$ spans a connected sub-tree of T (sub-tree property). The *width* of a tree decomposition equals $\max_{i \in I} |B_i| - 1$, and the *treewidth* $\text{tw}(G)$ is the smallest width of any tree decomposition of G . Replacing all occurrences of “tree” with “path” in the definition of treewidth yields the notion of *pathwidth* $\text{pw}(G)$. We have $\text{tw}(G) \leq \text{pw}(G)$.

2.2 Manifolds

A *d-dimensional manifold* is a topological space \mathcal{M} , where each point $x \in \mathcal{M}$ has a neighborhood homeomorphic to \mathbb{R}^d or to the closed upper half-space $\{(x_1, \dots, x_d) \in \mathbb{R}^d : x_d \geq 0\}$. The latter type of points of \mathcal{M} constitute the *boundary* $\partial\mathcal{M}$ of \mathcal{M} . A compact manifold is said to be *closed* if it has an empty boundary. We consider manifolds up to homeomorphism (“continuous deformations”) and write $\mathcal{M}_1 \cong \mathcal{M}_2$ for homeomorphic manifolds \mathcal{M}_1 and \mathcal{M}_2 .

3-Manifolds and surfaces


The main objects of study in this paper are 3-dimensional manifolds, or *3-manifolds* for short. Here we give a brief introduction to 3-manifolds tailored to our purposes. We refer the reader to [51] (and the references therein) for more details. All 3-manifolds and surfaces encountered in this article are compact and orientable. We let \mathbb{S}_g denote the closed, connected, orientable surface of genus g . We also refer to the d -dimensional torus and sphere as \mathbb{T}^d and \mathbb{S}^d , respectively (hence $\mathbb{S}_0 = \mathbb{S}^2$ and $\mathbb{S}_1 = \mathbb{T}^2$). The genus $g(\mathcal{S})$ of a (not necessarily connected) surface \mathcal{S} is defined to be the sum of the genera of its connected components.

Triangulations and the treewidth of 3-manifolds A *triangulation* \mathcal{T} of a given 3-manifold \mathcal{M} is a finite collection of abstract tetrahedra glued together along pairs of their triangular faces, such that the resulting space is homeomorphic to \mathcal{M} . Unpaired triangles comprise a triangulation of the boundary of \mathcal{M} . Note that the face gluings may also identify several tetrahedral edges (or vertices) in a single *edge* (or *vertex*) of \mathcal{T} . Every compact 3-manifold admits a triangulation [43] (cf. [10]). Given a triangulation \mathcal{T} , its *dual graph* $\Gamma(\mathcal{T})$ is the multigraph whose nodes correspond to the tetrahedra in \mathcal{T} , and arcs to face gluings (Figure 1).

For a compact 3-manifold \mathcal{M} , its *treewidth* $\text{tw}(\mathcal{M})$ (resp. *pathwidth* $\text{pw}(\mathcal{M})$) is defined as the smallest treewidth (resp. pathwidth) of the dual graph of any triangulation of \mathcal{M} .

Incompressible surfaces and essential disks Given a 3-manifold \mathcal{M} , a surface $\mathcal{S} \subset \mathcal{M}$ is said to be *properly embedded* in \mathcal{M} if $\partial\mathcal{S} \subset \partial\mathcal{M}$ and $(\mathcal{S} \setminus \partial\mathcal{S}) \subset (\mathcal{M} \setminus \partial\mathcal{M})$. Given a properly

⁶ Throughout this paper 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.



■ **Figure 1** (i) Example of a triangulation \mathcal{T} with two tetrahedra Δ_1 and Δ_2 , and three face gluing maps φ_1 , φ_2 and φ_3 . The map φ_1 is specified to be $\Delta_1(123) \longleftrightarrow \Delta_2(103)$. (ii) The dual graph $\Gamma(\mathcal{T})$ of the triangulation \mathcal{T} . Reproduced from [26, Figure 1].

embedded surface $\mathcal{S} \subset \mathcal{M}$, an embedded disk $D \subset \mathcal{M}$ with $\text{int}(D) \cap \mathcal{S} = \emptyset$ and $\partial D \subset \mathcal{S}$ a curve not bounding a disk on \mathcal{S} is called a *compressing disk*. If such a disk exists, \mathcal{S} is said to be *compressible* in \mathcal{M} , otherwise – and if \mathcal{S} is not a 2-sphere – it is called *incompressible*. A 3-manifold \mathcal{M} is said to be *irreducible*, if every embedded 2-sphere bounds a 3-ball in \mathcal{M} . A disk $D \subset \mathcal{M}$ properly embedded in a 3-manifold \mathcal{M} is called *inessential* if it cuts off a 3-ball from \mathcal{M} , otherwise D is called *essential*. A compact, orientable, irreducible 3-manifold is called *Haken* if it contains an orientable, properly embedded, incompressible surface, and otherwise is referred to as *non-Haken*.

Heegaard splittings of closed 3-manifolds A *handlebody* \mathcal{H} is a connected 3-manifold homeomorphic to a thickened graph. The *genus* $g(\mathcal{H})$ of \mathcal{H} is defined as the genus of its boundary surface $\partial\mathcal{H}$. A *Heegaard splitting* of a closed, orientable 3-manifold \mathcal{M} 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 $\mathcal{H} \cap \mathcal{H}' = \partial\mathcal{H} = \partial\mathcal{H}' = \mathcal{S}$ called the *splitting surface*. Introduced in [24], the *Heegaard genus* $\mathfrak{g}(\mathcal{M})$ of \mathcal{M} is the smallest genus $g(\mathcal{S})$ over all Heegaard splittings of \mathcal{M} .

The JSJ decomposition A central result by Jaco–Shalen [30, 31] and Johannson [32] asserts that every closed, irreducible and orientable 3-manifold \mathcal{M} admits a collection \mathbf{T} of pairwise disjoint embedded, incompressible tori, where each piece of the complement $\mathcal{M} \setminus \mathbf{T}$ is either Seifert fibered⁷ or atoroidal⁸. A minimal such collection of tori is unique up to isotopy and gives rise to the so-called *JSJ decomposition* (or *torus decomposition*) of \mathcal{M} [23, Theorem 1.9]. We refer to this collection of incompressible tori as the *JSJ tori* of \mathcal{M} . The graph with nodes the JSJ pieces, and an arc for each JSJ torus (with endpoints the two nodes corresponding to its two adjacent pieces) is called the *dual graph* $\Gamma(\mathcal{D})$ of the *JSJ decomposition* \mathcal{D} of \mathcal{M} .

3 Generalized Heegaard Splittings

A Heegaard splitting of a closed 3-manifold is a decomposition into two identical handlebodies along an embedded surface. Introduced by Scharle and Thompson [50], a generalized Heegaard splitting of a compact 3-manifold \mathcal{M} (possibly with boundary) is a decomposition of \mathcal{M} into several pairs of compression bodies along a family of embedded surfaces, subject to certain rules. Following [49, Chapters 2 and 5],⁹ here we give an overview of this framework.

⁷ See [51, Section 3.7] or [23, p. 18] for an introduction to Seifert fibered spaces (cf. Appendix A.2).

⁸ An irreducible 3-manifold \mathcal{M} is called *atoroidal* if every incompressible torus in \mathcal{M} is boundary-parallel.

⁹ For an open-access version, see [48, Sections 3.1 and 4].


3.1 Compression bodies and forks

A *compression body* \mathcal{C} is a 3-manifold with boundary obtained by the following procedure.

1. Consider the thickening $\mathcal{S} \times [0, 1]$ of a closed, orientable, possibly disconnected surface \mathcal{S} ,
2. optionally attach some 1-*handles*, each being of the form $D \times [0, 1]$ (thickened edge, where the disk D is the cross-section), to $\mathcal{S} \times \{1\}$ along $D \times \{0\} \cup D \times \{1\}$, and
3. optionally fill in some collection of 2-sphere components of $\mathcal{S} \times \{0\}$ with 3-balls.


We call $\partial_- \mathcal{C} = \mathcal{S} \times \{0\} \setminus \{\text{filled-in 2-sphere components}\}$ the *lower boundary* of \mathcal{C} and $\partial_+ \mathcal{C} = \partial \mathcal{C} \setminus \partial_- \mathcal{C}$ its *upper boundary*. By construction, $g(\partial_+ \mathcal{C}) \geq g(\partial_- \mathcal{C})$. Note that, if $\partial_- \mathcal{C}$ is empty, then \mathcal{C} is a handlebody. We allow compression bodies to be disconnected.

A *fork*, more precisely an n -*fork*, F is a tree with $n + 2$ nodes $V(F) = \{\rho, \gamma, \tau_1, \dots, \tau_n\}$, where ρ , called the *root*, is of degree $n + 1$, and the other nodes are leaves. One of them, denoted γ , is called the *grip*, and the remaining leaves τ_1, \dots, τ_n are called *tines*. A fork can be regarded as an abstraction of a connected compression body \mathcal{C} , where the grip corresponds to $\partial_+ \mathcal{C}$ and each tine corresponds to a connected component of $\partial_- \mathcal{C}$, see, e.g., Figure 2.




■ **Figure 2** The compression body \mathcal{C} is obtained by first thickening the disconnected surface $\mathbb{T}^2 \cup \mathbb{S}_2$ to $(\mathbb{T}^2 \cup \mathbb{S}_2) \times [0, 1]$, then attaching two 1-handles (h_1 and h_2) between $\mathbb{T}^2 \times \{1\}$ and $\mathbb{S}_2 \times \{1\}$. For the lower boundary of \mathcal{C} we have $\partial_- \mathcal{C} = (\mathbb{T}^2 \cup \mathbb{S}_2) \times \{0\}$, and for its upper boundary $\partial_+ \mathcal{C} = \partial \mathcal{C} \setminus \partial_- \mathcal{C} \cong \mathbb{S}_4$.

Non-faithful forks In certain situations (notably, in the proof of Theorem 1, cf. Section 4) it is useful to also take a simplified view on a generalized Heegaard splitting. To that end, one may represent several compression bodies by a single *non-faithful* fork, where the grip and the tines may correspond to collections of boundary components. To distinguish faithful forks from non-faithful ones, we color the roots of the latter with gray, see Figure 3.



■ **Figure 3** Two faithful forks bundled into a non-faithful fork. The colors show the grouping of the tines.




■ **Figure 4** Spines of a compression body \mathcal{C} and of a handlebody \mathcal{H} .

Spines and sweep-outs of compression bodies A graph Γ embedded in a compression body \mathcal{C} is called a *spine* of \mathcal{C} , if every node of Γ that is incident to $\partial_- \mathcal{C}$ is of degree one, and $\mathcal{C} \setminus (\Gamma \cup \partial_- \mathcal{C}) \cong \partial_+ \mathcal{C} \times (0, 1]$, see Figure 4. Assume first that $\partial_- \mathcal{C} \neq \emptyset$. A *sweep-out of \mathcal{C} along an interval*, say $[-1, 1]$, is a continuous map $f: \mathcal{C} \rightarrow [-1, 1]$, such that $f^{-1}(\pm 1) = \partial_{\pm} \mathcal{C}$,


$$f^{-1}(t) \cong \begin{cases} \partial_+ \mathcal{C}, & \text{if } t \in (0, 1), \\ \Gamma \cup \partial_- \mathcal{C}, & \text{if } t = 0, \text{ and} \\ \partial_- \mathcal{C}, & \text{if } t \in (-1, 0). \end{cases} \quad (3)$$

Writing $\Sigma_t = f^{-1}(t)$, we get a 1-parameter family of surfaces (except for Σ_0) “sweeping through” \mathcal{C} . For handlebodies, the definition of a sweep-out is similar, but it “ends at 0” with the spine: $f^{-1}(1) = \partial \mathcal{H}$, $f^{-1}(t) \cong \partial \mathcal{H}$ if $t \in (0, 1)$, and $f^{-1}(0) = \Gamma$, see Figure 5.



■ **Figure 5** Sweep-outs of the compression body \mathcal{C} and of the handlebody \mathcal{H} shown in Figure 4.

► **Remark 3 (Sweep-out along a fork).** It is straightforward to adapt the definition of sweep-out in a way that, instead of an interval, we sweep a compression body \mathcal{C} along a (faithful or non-faithful) fork F . Such a sweep-out is a continuous map $f: \mathcal{C} \rightarrow \|F\|$ (where $\|F\|$ denotes a geometric realization of F) that satisfies very similar requirements to those in (3), however, the components of the lower boundary $\partial_- \mathcal{C}$ may appear in different level sets, see Figure 6.




■ **Figure 6** Sweep-out of a compression body along $[-1, 1]$ (left) and along its faithful fork (right).

3.2 Generalized Heegaard splittings and fork complexes


Heegaard splittings revisited Let \mathcal{M} be a 3-manifold, and $\{\partial_1 \mathcal{M}, \partial_2 \mathcal{M}\}$ be a partition of the components of $\partial \mathcal{M}$. A *Heegaard splitting* of $(\mathcal{M}, \partial_1 \mathcal{M}, \partial_2 \mathcal{M})$ is a triplet $(\mathcal{C}_1, \mathcal{C}_2, \mathcal{S})$, where \mathcal{C}_1 and \mathcal{C}_2 are compression bodies with $\mathcal{C}_1 \cup \mathcal{C}_2 = \mathcal{M}$, $\mathcal{C}_1 \cap \mathcal{C}_2 = \partial_+ \mathcal{C}_1 = \partial_+ \mathcal{C}_2 = \mathcal{S}$, $\partial_- \mathcal{C}_1 = \partial_1 \mathcal{M}$, and $\partial_- \mathcal{C}_2 = \partial_2 \mathcal{M}$. The *genus* of the Heegaard splitting $(\mathcal{C}_1, \mathcal{C}_2, \mathcal{S})$ is the genus $g(\mathcal{S})$ of the *splitting surface* \mathcal{S} . The *Heegaard genus* $\mathfrak{g}(\mathcal{M})$ of \mathcal{M} is the smallest genus of any Heegaard splitting of $(\mathcal{M}, \partial_1 \mathcal{M}, \partial_2 \mathcal{M})$, taken over all partitions $\{\partial_1 \mathcal{M}, \partial_2 \mathcal{M}\}$ of $\partial \mathcal{M}$.

► **Proposition 4** ([49, Theorem 2.1.11], cf. [26, Appendix B]). *For any partition $\{\partial_1 \mathcal{M}, \partial_2 \mathcal{M}\}$ of the boundary components of \mathcal{M} , there exists a Heegaard splitting of $(\mathcal{M}, \partial_1 \mathcal{M}, \partial_2 \mathcal{M})$.*



■ **Figure 7** Schematic of a 3-manifold \mathcal{M} with a partition of its boundary components (left). Faithful (center) and non-faithful (right) *fork complexes* representing a Heegaard splitting of $(\mathcal{M}, \partial_1\mathcal{M}, \partial_2\mathcal{M})$.

Generalized Heegaard splittings A *generalized Heegaard splitting* \mathcal{H} of a 3-manifold \mathcal{M} consists of **1.** a decomposition $\mathcal{M} = \bigcup_{i \in I} \mathcal{M}_i$ into submanifolds $\mathcal{M}_i \subseteq \mathcal{M}$ intersecting along closed surfaces (Figure 8a), **2.** for each $i \in I$ a partition $\{\partial_1\mathcal{M}_i, \partial_2\mathcal{M}_i\}$ of the components of $\partial\mathcal{M}_i$ that together satisfy an *acyclicity condition*: there is an ordering, i.e., a bijection $\ell: I \rightarrow \{1, \dots, |I|\}$, such that, for each $i \in I$ the components of $\partial_1\mathcal{M}_i$ (resp. $\partial_2\mathcal{M}_i$) not belonging to $\partial\mathcal{M}$ are only incident to submanifolds \mathcal{M}_j with $\ell(j) < \ell(i)$ (resp. $\ell(j) > \ell(i)$), and **3.** a choice of a Heegaard splitting $(\mathcal{C}_1^{(i)}, \mathcal{C}_2^{(i)}, \mathcal{S}_i)$ for each $(\mathcal{M}_i, \partial_1\mathcal{M}_i, \partial_2\mathcal{M}_i)$. Such a choice of *splitting surfaces* \mathcal{S}_i ($i \in I$) is said to be *compatible* with ℓ , see, e.g., Figure 8b.



(a) A decomposition of \mathcal{M} into four submanifolds $\mathcal{M}_1, \dots, \mathcal{M}_4$ intersecting along (possibly disconnected) closed surfaces \mathcal{R}_i .

(b) An admissible choice of splitting surfaces \mathcal{S}_i for the \mathcal{M}_i that is compatible with the trivial ordering $\ell(i) \mapsto i$.

(c) The faithful fork complex that represents the generalized Heegaard splitting shown in the center (Figure 8b).

■ **Figure 8** Schematics of a generalized Heegaard splitting (based on figures from [26, Section 2.4]).

Just as compression bodies can be represented by forks, (generalized) Heegaard splittings can be visualized via *fork complexes*, see Figures 7 and 8c (cf. [49, Section 5.1] for details).

Sweep-outs of 3-manifolds A generalized Heegaard splitting \mathcal{H} of a 3-manifold \mathcal{M} induces a *sweep-out* $f: \mathcal{M} \rightarrow \|\mathcal{F}\|$ of \mathcal{M} along any fork complex \mathcal{F} that represents \mathcal{H} (here $\|\mathcal{F}\|$ denotes a drawing, i.e., a *geometric realization* of the abstract fork complex \mathcal{F}) by concatenating the corresponding sweep-outs of the compression bodies that comprise \mathcal{H} (cf. Remark 3). We also refer to a sweep-out $f: \mathcal{M} \rightarrow \|\mathcal{F}\|$ by the ensemble $\Sigma = \{\Sigma_x : x \in \|\mathcal{F}\|\}$ of its *level sets*, where $\Sigma_x = f^{-1}(x)$.

The width of a generalized Heegaard splitting For a generalized Heegaard splitting \mathcal{H} , the surfaces \mathcal{S}_i ($i \in I$) are also called the *thick levels*, and the lower boundaries $\partial_-\mathcal{C}_1^{(i)}, \partial_-\mathcal{C}_2^{(i)}$ are called the *thin levels* of \mathcal{H} . The *width* $w(\mathcal{H})$ of \mathcal{H} is the sequence obtained by taking a non-increasing ordering of the multiset $\{g(\mathcal{S}_i) : i \in I\}$ of the genera of the thick levels.

A generalized Heegaard splitting \mathcal{H} of a 3-manifold \mathcal{M} for which $w(\mathcal{H})$ is minimal with respect to the lexicographic order ($<$) among all splittings of \mathcal{M} is said to be in *thin position*.


3.3 Weak reductions

A Heegaard splitting $(\mathcal{C}_1, \mathcal{C}_2, \mathcal{S})$ of a connected 3-manifold is said to be *weakly reducible* [19], if there are essential disks $D_i \subset \mathcal{C}_i$ ($i = 1, 2$)¹⁰ with $\partial D_1 \cap \partial D_2 = \emptyset$, see Figure 9. In this case we also say that the splitting surface \mathcal{S} is weakly reducible. A generalized Heegaard splitting \mathcal{H} is *weakly reducible*, if at least one of its splitting surfaces is weakly reducible; otherwise \mathcal{H} is called *strongly irreducible*. Every 3-manifold possesses a strongly irreducible generalized Heegaard splitting and this fact can be exploited in various contexts (e.g., in the proof of Theorem 1). The usefulness of such splittings is mainly due to the following seminal result.

► **Theorem 5** ([49, Lemma 5.2.4], [50, Rule 5]). *Let \mathcal{H} be a strongly irreducible generalized Heegaard splitting. Then every connected component of every thin level of \mathcal{H} is incompressible.*

Given a weakly reducible generalized Heegaard splitting \mathcal{H} with a weakly reducible splitting surface¹¹ \mathcal{S} and essential disks D_1 and D_2 as above, one can execute a *weak reduction* at \mathcal{S} . This modification amounts to performing particular cut-and-paste operations on \mathcal{S} guided by D_1 and D_2 , and decomposes each of the two compression bodies adjacent to \mathcal{S} into a pair of compression bodies. Importantly, this operation results in another generalized Heegaard splitting \mathcal{H}' of the same 3-manifold with $w(\mathcal{H}') < w(\mathcal{H})$.

We illustrate weak reductions via Example 16 and refer to [49, Proposition 5.2.3] for further details (notably, Figures 5.8–5.13 therein, but also Lemma 5.2.2, Figures 5.6 and 5.7, and Proposition 5.2.4), including an exhaustive list of instances of weak reductions.¹²



■ **Figure 9** Local picture of a portion of a weakly reducible splitting surface \mathcal{S} .

4 The Main Result

In this section we prove Theorem 1. The inequalities (1) and (2) are deduced in the same way, thus we only show the proof of (1) in detail, and then explain how it can be adapted to that of (2). First, we specify what we mean by a “sufficiently complicated” JSJ decomposition.

► **Definition 6.** *Given $\delta > 0$, the JSJ decomposition of an irreducible 3-manifold \mathcal{M} is δ -complicated, if any incompressible or strongly irreducible Heegaard surface $\mathcal{S} \subset \mathcal{M}$ with genus $g(\mathcal{S}) \leq \delta$ can be isotoped to be simultaneously disjoint from all the JSJ tori of \mathcal{M} .*


¹⁰The assumption that D_i is essential in the compression body \mathcal{C}_i implies that $\partial D_i \subset \partial_+ \mathcal{C}_i = \mathcal{S}$ ($i = 1, 2$).

¹¹For future reference we remind the reader that splitting surfaces, also called thick levels, correspond to grips in the faithful fork complex that represents the generalized Heegaard splitting \mathcal{H} .

¹²In the open-access version [48] these are Proposition 4.2.3 and Figures 87–92 (as well as Lemma 4.2.2, Figures 85 and 86, and Lemma 4.2.4).

Proof of inequality (1). Our goal is to prove that $\text{tw}(\Gamma(\mathcal{D})) \leq 18(\text{tw}(\mathcal{M}) + 1)$, where $\Gamma(\mathcal{D})$ is the dual graph of the δ -complicated JSJ decomposition of the irreducible 3-manifold \mathcal{M} . To this end, we set $\delta = 18(\text{tw}(\mathcal{M}) + 1)$ and fix a triangulation \mathcal{T} of \mathcal{M} , whose dual graph $\Gamma(\mathcal{T})$ has minimal treewidth, i.e., $\text{tw}(\Gamma(\mathcal{T})) = \text{tw}(\mathcal{M})$. Now, (1) is established in four stages.

1. Setup. By invoking the construction in [28, Section 6], from \mathcal{T} we obtain a generalized Heegaard splitting \mathcal{H} of \mathcal{M} together with a sweep-out $f: \mathcal{M} \rightarrow \|\mathcal{F}_\bullet\|$ along a non-faithful fork complex \mathcal{F}_\bullet representing \mathcal{H} . By construction, \mathcal{F}_\bullet is a tree with all of its nodes having degree one or three (Figure 10), moreover all non-degenerate level surfaces $\Sigma_x = f^{-1}(x)$ have genus bounded above by $18(\text{tw}(\mathcal{M}) + 1)$. Let \mathcal{F}_\circ be the faithful fork complex representing \mathcal{H} . Note that \mathcal{F}_\circ is obtained from \mathcal{F}_\bullet by replacing every non-faithful fork $F \in \mathcal{F}_\bullet$ with the collection \mathcal{C}_F of faithful forks that accurately represents the (possibly disconnected) compression body \mathcal{C} corresponding to F . The inverse operation, i.e., for each $F \in \mathcal{F}_\bullet$ bundling all faithful forks in \mathcal{C}_F into F (see Figure 3), induces a projection map $\pi: \|\mathcal{F}_\circ\| \rightarrow \|\mathcal{F}_\bullet\|$ between the drawings (cf. Figures 11(i)–(ii)). Note that every fork, tine, or grip in \mathcal{F}_\bullet corresponds to a collection of the corresponding items in \mathcal{F}_\circ and so the projection map is well-defined. \triangleleft

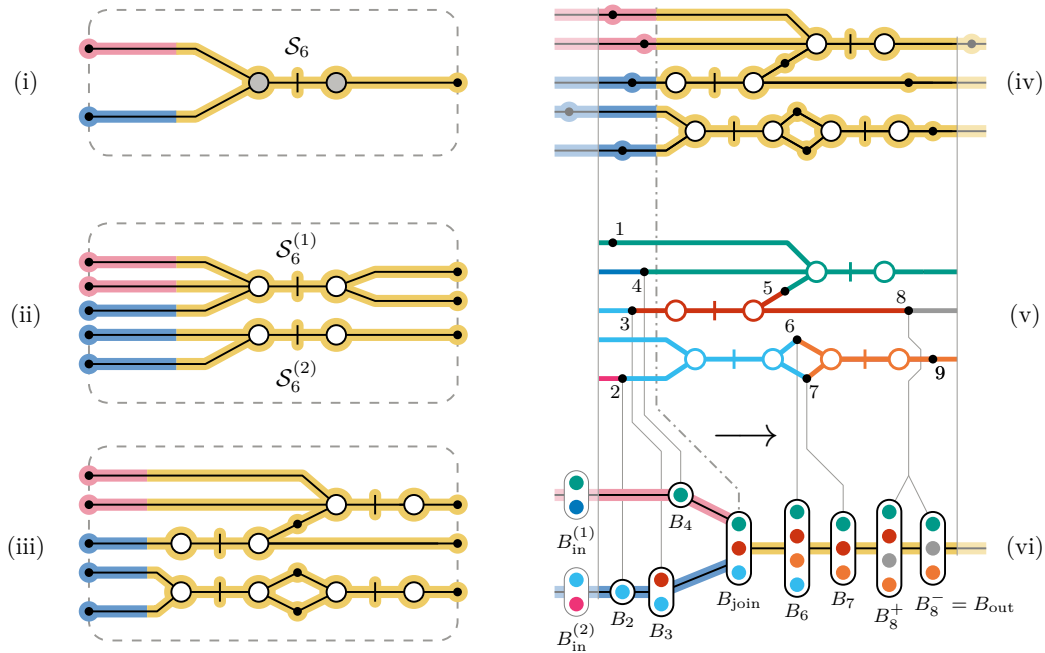


■ **Figure 10** (i) **A.** The dual graph $\Gamma(\mathcal{T})$ of some triangulation \mathcal{T} of a 3-manifold \mathcal{M} . **B.** Low-congestion routing of $\Gamma(\mathcal{T})$ along a host tree with marked root arc r . **C.** Drawing of a non-faithful fork complex \mathcal{F}_\bullet that represents the generalized Heegaard splitting of \mathcal{M} induced by the routing of $\Gamma(\mathcal{T})$. The genera of all (possibly disconnected) thick levels \mathcal{S}_i is bounded above by $18(\text{tw}(\mathcal{M}) + 1)$. (ii) Color-coded segmentation of $\|\mathcal{F}_\bullet\|$ in preparation for the next stage of the proof (cf. Figure 11).

2. Weak reductions. In case \mathcal{H} is weakly reducible, we repeatedly perform weak reductions until we obtain a strongly irreducible generalized Heegaard splitting \mathcal{H}' of \mathcal{M} . Since weak reductions always decrease the width of a generalized Heegaard splitting, this process terminates after finitely many iterations. Throughout, we maintain that the drawings of the associated faithful fork complexes follow $\|\mathcal{F}_\bullet\|$. Let \mathcal{F}'_\circ be the faithful fork complex representing the final splitting \mathcal{H}' , $f': \mathcal{M} \rightarrow \|\mathcal{F}'_\circ\|$ be the sweep-out of \mathcal{M} induced by \mathcal{H}' , and $\pi': \|\mathcal{F}'_\circ\| \rightarrow \|\mathcal{F}_\bullet\|$ be the associated projection map (Figure 11(iii)). Due to the nature of weak reductions, $18(\text{tw}(\mathcal{M}) + 1)$ is still an upper bound on the genus of any (possibly

disconnected) level surface of $\pi' \circ f': \mathcal{M} \rightarrow \|\mathcal{F}_\bullet\|$. As \mathcal{M} is irreducible, we may assume that no component of such a level surface is homeomorphic to a sphere (cf. [36, p. 337]). It follows that the number of components of any level set of $\pi' \circ f'$ is at most $18(\text{tw}(\mathcal{M}) + 1)$. \triangleleft

3. Perturbation and isotopy. We now apply a level-preserving perturbation $\pi' \rightsquigarrow \pi''$ on π' , after which each thin level of \mathcal{H}' lies in different level sets of $\pi'' \circ f': \mathcal{M} \rightarrow \|\mathcal{F}_\bullet\|$ (Figure 11(iv)). Since \mathcal{H}' is strongly irreducible, its thick levels are strongly irreducible (by definition), and its thin levels are incompressible (by Theorem 5). Moreover, all these surfaces have genera at most $18(\text{tw}(\mathcal{M}) + 1)$. Hence, as the JSJ decomposition of \mathcal{M} is assumed to be δ -complicated with $\delta = 18(\text{tw}(\mathcal{M}) + 1)$, we may isotope all JSJ tori to be disjoint from all the thick and thin levels of \mathcal{H}' . Then, by invoking [5, Corollary 4.5] we can isotope each JSJ torus of \mathcal{M} to coincide with a component of some thin level of \mathcal{H}' (Figure 11(v)). As a consequence, every compression body of the splitting \mathcal{H}' is contained in a unique JSJ piece of \mathcal{M} . \triangleleft



■ **Figure 11** Overview of the proof of inequality (1). The figure shows the circled area in Figure 10(ii). Stage 1: Construction of initial fork complex (i) and split into faithful fork complex (ii). Stage 2: Weak reductions (iii), see [49, Proposition 5.2.3, Figures 5.8–5.13] for a complete list and their effects on the underlying fork complexes. Stage 3: perturbations and isotopy (iv). Stage 4: Construction of tree decomposition (v) and (vi).

4. The tree decomposition of $\Gamma(\mathcal{D})$. First note that every level set $(\pi'' \circ f')^{-1}(x)$ is incident to at most $18(\text{tw}(\mathcal{M}) + 1) + 1 = 18\text{tw}(\mathcal{M}) + 19$ JSJ pieces of \mathcal{M} . The ‘plus one’ appears, because if $(\pi'' \circ f')^{-1}(x)$ contains a JSJ torus, then this torus is incident to two JSJ pieces of \mathcal{M} . Also note that, because of the perturbation performed in the previous stage, each level set $(\pi'' \circ f')^{-1}(x)$ can contain at most one JSJ torus of \mathcal{M} .

We now construct a tree decomposition (\mathcal{X}, T) of $\Gamma(\mathcal{D})$ of width $18(\text{tw}(\mathcal{M}) + 1)$. Eventually, T will be a subdivision of \mathcal{F}_\bullet (which is a tree) with nodes corresponding to the bags in \mathcal{X} , which we now describe. By [28, Section 6 (p. 86)], each leaf l of \mathcal{F}_\bullet corresponds to a spine of a handlebody \mathcal{H}_l . We define a bag that contains the unique node of $\Gamma(\mathcal{D})$ associated with

the JSJ piece containing \mathcal{H}_l . As we sweep through $\|\mathcal{F}_\bullet\|$ (cf. the arrows on Figure 10(2)), whenever we pass through a point $x \in \|\mathcal{F}_\bullet\|$ such that the level set $(\pi'')^{-1}(x)$ contains a tine of $\|\mathcal{F}'_\bullet\|$, one of four possible events may occur (illustrated in Figure 11(v)–(vi)): **1.** A new JSJ piece appears. In this case we take a copy of the previous bag, and add the corresponding node of $\Gamma(\mathcal{D})$ into the bag. **2.** A JSJ piece disappears. Then we delete the corresponding node of $\Gamma(\mathcal{D})$ from a copy of the previous bag. **3.** Both previous kinds of events happen simultaneously. In this case we introduce two new bags. The first to introduce the new JSJ piece, the second to delete the old one. **4.** If neither a new JSJ piece is introduced, nor an old one is left behind, we do nothing. Whenever we arrive at a merging point in the sweep-out (i.e., a degree-three node of $\|\mathcal{F}_\bullet\|$), we introduce a new bag, which is the union of the two previous bags. Note that the two previous bags do not necessarily need to be disjoint. \triangleleft

It remains to verify that (\mathcal{X}, T) is, indeed, a tree decomposition of $\Gamma(\mathcal{D})$ of width at most $18 \operatorname{tw}(\mathcal{M}) + 18$. *Node coverage:* Every node of $\Gamma(\mathcal{D})$ must be considered at least once, since we sweep through the entirety of $\|\mathcal{F}_\bullet\|$. *Arc coverage:* we must ensure that all pairs of nodes of $\Gamma(\mathcal{D})$ with their JSJ pieces meeting at a tine are contained in some bag. This is always the case for **1**, because a new JSJ piece appears while all other JSJ pieces are still in the bag. It is also the case for **2** because a JSJ piece disappears but the previous bag contained all the pieces. In **3** a new JSJ piece appears and at the same time another JSJ piece disappears. However, in this case we first introduce the new piece, thus making sure that adjacent JSJ pieces always occur in at least one bag. *Sub-tree property:* This follows from the fact that a JSJ piece, once removed from all bags, must be contained in the part of $\|\mathcal{F}_\bullet\|$ that was already swept. Now, every JSJ piece incident to a given level set $(\pi'' \circ f')^{-1}(x)$ of the sweep-out must contribute a positive number to its genus. Hence, it follows that every bag can contain at most $18 \operatorname{tw}(\mathcal{M}) + 19$ elements (with equality only possible where a tine simultaneously introduces and forgets a JSJ piece). This proves inequality (1). \blacktriangleleft

Proof of inequality (2). We start with the results from [28, Section 5] yielding a fork-complex \mathcal{F}_\bullet whose underlying space $\|\mathcal{F}_\bullet\|$ is a path, and the genus of the level sets of the associated sweep-out is bounded above by $4(3 \operatorname{pw}(\mathcal{M}) + 1)$. Setting $\delta = 4(3 \operatorname{pw}(\mathcal{M}) + 1)$, the remainder of the proof is analogous with the proof of inequality (1). \blacktriangleleft

5 An Algorithmic Construction



Here we establish Theorem 2 that paves the way to the applications in Section 6. In what follows, Δ denotes an arbitrary, but fixed, positive integer. Let $G = (V, E)$ be a graph with $|V| = n$ and maximum degree Δ . Theorem 2 asserts that, in $\operatorname{poly}(n)$ time one can construct a triangulation \mathcal{T}_G of a closed, irreducible 3-manifold \mathcal{M}_G , such that the dual graph $\Gamma(\mathcal{D})$ of its JSJ decomposition \mathcal{D} equals G , moreover, the pathwidth (resp. treewidth) of G determines the pathwidth (resp. treewidth) of \mathcal{M}_G up to a constant factor.

Our proof of Theorem 2 rests on a synthesis of work by Lackenby [37] and by Bachman, Derby-Talbot and Sedgwick [6]. In [37, Section 3] it is shown that the homeomorphism problem for closed 3-manifolds is at least as hard as the graph isomorphism problem. The proof relies on a simple construction that, given a graph G , produces a closed, orientable, triangulated 3-manifold whose JSJ decomposition \mathcal{D} satisfies $\Gamma(\mathcal{D}) = G$. This gives the blueprint for our construction as well. In particular, we use the same building blocks that are described in [37, p. 591]. However, as opposed to Lackenby, we paste together these building blocks via *high-distance torus gluings* akin to the construction presented in [6, Section 4]. This ensures that we can apply Theorem 1 for the resulting 3-manifold \mathcal{M}_G and deduce

the right-hand-side inequalities of 3a and 3b. The left-hand-side inequalities are shown by inspecting \mathcal{T}_G . We now elaborate on the ingredients of the proof of Theorem 2.

The building blocks We first recall the definition of the building blocks from [37, p. 591]. Let $k \in \mathbb{N}$ be a positive integer. Consider the 2-dimensional torus $\mathbb{T}^2 = \mathbb{S}^1 \times \mathbb{S}^1$, and let \mathbb{T}_k^2 denote the compact surface obtained from \mathbb{T}^2 by the removal of k pairwise disjoint open disks. We define $\mathcal{N}(k) = \mathbb{T}_k^2 \times \mathbb{S}^1$. Note that the boundary $\partial\mathcal{N}(k)$ of $\mathcal{N}(k)$ consists of k tori, which will be the gluing sites. See Figure 12 for the example of $\mathcal{N}(3)$.

We choose a triangulation $\mathcal{T}(k)$ for $\mathcal{N}(k)$ that induces the minimal 2-triangle triangulation of the torus (cf. Figure 13) at each boundary component. Note that $\mathcal{T}(k)$ can be constructed from $O(k)$ tetrahedra. An explicit description of $\mathcal{T}(k)$ is given in Appendix B.




■ **Figure 12** Illustration of $\mathcal{N}(3)$.

■ **Figure 13** Minimal triangulation of the torus \mathbb{T}^2 .

Overview of the construction of \mathcal{M}_G We now give a high-level overview of constructing \mathcal{M}_G from the above building blocks. Given a graph $G = (V, E)$, for each node $v \in V$ we pick a block $\mathcal{N}_v \cong \mathcal{N}(d_v)$, where d_v denotes the degree of v . Next, for each arc $e = \{u, v\} \in E$, we pick a homeomorphism $\phi_e: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ of sufficiently high *distance* (we discuss this notion below) and use it to glue together a boundary torus of \mathcal{N}_u with one of \mathcal{N}_v . After performing all of these gluings, we readily obtain the 3-manifold \mathcal{M}_G (cf. Figure 14(i)–(ii)).

▷ **Claim 7** (Based on [37, p. 591]). The JSJ decomposition \mathcal{D} of \mathcal{M}_G satisfies $\Gamma(\mathcal{D}) = G$.

For a proof of Claim 7 we refer to Appendix A.2.



■ **Figure 14** Schematic overview of the construction underlying Theorem 2.

High-distance torus gluings As already mentioned, to ensure that we can apply Theorem 1 to \mathcal{M}_G , we use torus homeomorphisms of “sufficiently high distance” to glue the building blocks together. This notion of *distance*, which is defined through the *Farey distance*, is somewhat technical. Thus, we refer to [6, Section 4.1 and Appendix] for details, and rather recall a crucial result that makes the usefulness of distance in the current context apparent.

► **Theorem 8** ([6, Appendix], [4, Theorem 5.4]¹³). *There exists a computable constant K , depending only on the homeomorphism types of the blocks, so that if any set of blocks are glued with maps of distance at least $K\delta$ along their torus boundary components to form a closed 3-manifold \mathcal{M} , then the JSJ decomposition of \mathcal{M} is δ -complicated (cf. Definition 6).*

Triangulating the gluing maps We have already discussed that the block $\mathcal{N}(k)$ admits a triangulation $\mathcal{T}(k)$ with $O(k)$ tetrahedra, where $\mathcal{T}(k)$ induces a minimal, 1-vertex triangulation at each torus boundary of $\mathcal{N}(k)$. It is shown in [6, Section 4.2] that the gluings beading these blocks together can be realized as *layered triangulations* [29]. These triangulations manifest as “daisy chains” in the dual graph $\Gamma(\mathcal{T}_G)$ of the final triangulation \mathcal{T}_G , see Figure 14(iii).

► **Lemma 9** ([6, Lemma 4.6]). *There exist torus gluings with distance at least D , that can be realized as layered triangulations using 2D tetrahedra.*

▷ **Claim 10.** There exist universal constants $c, c' > 0$ such that

$$(c/\Delta) \text{tw}(\mathcal{M}_G) \leq \text{tw}(G) \quad \text{and} \quad (c'/\Delta) \text{pw}(\mathcal{M}_G) \leq \text{pw}(G).$$

Proof. Since every node of G has degree at most Δ , the construction of \mathcal{M}_G only uses building blocks homeomorphic to $\mathcal{N}(1), \dots, \mathcal{N}(\Delta)$. Hence, for each $v \in V$ the triangulation $\mathcal{T}(d_v)$ of the block \mathcal{N}_v contains $O(\Delta)$ tetrahedra. This, together with the above discussion on triangulating the gluing maps implies that (upon ignoring multi-arcs and loop arcs, which are anyway not “sensed” by treewidth or pathwidth), the dual graph $\Gamma(\mathcal{T}_G)$ is obtained from G by **1.** replacing each node $v \in V$ with a copy of the graph $\Gamma(\mathcal{T}(d_v))$ that contains $O(\Delta)$ nodes, and by **2.** possibly subdividing each arc $e \in E$ several times.

Now, the first operation increases the treewidth (resp. pathwidth) at most by a factor of $O(\Delta)$, while the arc-subdivisions keep these parameters basically the same, cf. Lemma 11. Hence $\text{tw}(\Gamma(\mathcal{T}_G)) \leq O(\Delta \text{tw}(G))$ and $\text{pw}(\Gamma(\mathcal{T}_G)) \leq O(\Delta \text{pw}(G))$, and the claim follows. ◀

► **Lemma 11** (Folklore, cf. [9, Lemma A. 1]). *Let $G = (V, E)$ be a graph. If G' is a graph obtained from G by subdividing a set $F \subseteq E$ of arcs an arbitrary number of times. Then*

$$\text{pw}(G') \leq \text{pw}(G) + 2 \quad \text{and} \quad \text{tw}(G') \leq \max\{\text{tw}(G), 3\}.$$

Finishing the proof of Theorem 2. We have already shown **2** (Claim 7) and the left-hand-sides of the inequalities **3a** and **3b** (Claim 10). To prove the remaining parts of Theorem 2, let $\delta = \max\{18(\text{tw}(G) + 1), 4(3 \text{pw}(G) + 1)\} = O(\text{pw}(G))$. By Theorem 8, there is a computable constant K_Δ depending only on $\mathcal{N}(1), \dots, \mathcal{N}(\Delta)$ and hence only on Δ , so that if we glue together the blocks via maps of distance at least $K_\Delta\delta$, then the JSJ decomposition of \mathcal{M}_G is δ -complicated. By Lemma 9, each such gluing map can be realized as a layered triangulation consisting of $2K_\Delta\delta$ tetrahedra. Since G has at most $\Delta n/2$ arcs, these layered triangulations contain at most $2K_\Delta\delta\Delta n/2 = \Delta K_\Delta\delta n = O(\Delta K_\Delta \text{pw}(G) \cdot n) = O_\Delta(\text{pw}(G) \cdot n)$ tetrahedra altogether. Since the triangulated blocks $\mathcal{T}(d_v)$ contain $O(\Delta \cdot n)$ tetrahedra in total, the triangulation \mathcal{T}_G of the manifold \mathcal{M}_G can be built from $O_\Delta(\text{pw}(G) \cdot n) \leq O_\Delta(n^2)$ tetrahedra. Last, as it follows from [6, Section 4], the construction can be executed in quadratic time. ◀


¹³The notation and the statement of Theorem 8 have been adapted to match the present context.

6 Applications

► **Corollary 12.** *There exist 3-manifolds $(\mathcal{M}_h)_{h \in \mathbb{N}}$ with $\text{tw}(\mathcal{M}_h) \leq 2$ and $\text{pw}(\mathcal{M}_h) \xrightarrow{h \rightarrow \infty} \infty$.*

► **Corollary 13.** *There exist Haken 3-manifolds $(\mathcal{N}_k)_{k \in \mathbb{N}}$ with $\text{tw}(\mathcal{N}_k) \xrightarrow{k \rightarrow \infty} \infty$.*

Proof of Corollaries 12 and 13. Using Theorem 2, the construction of bounded-treewidth (Haken) 3-manifolds with arbitrarily large pathwidth follows by taking $\mathcal{M}_h = \mathcal{M}_{T_h}$ (with the notation of Theorem 2), where T_h is the *complete binary tree of height h* . The construction of Haken 3-manifolds of arbitrary large treewidth is deduced by setting $\mathcal{N}_k = \mathcal{M}_{\text{Grid}(k)}$, where $\text{Grid}(k)$ denotes the $k \times k$ grid graph. See Figure 15. Obtained as JSJ decompositions, where the JSJ tori are two-sided, incompressible surfaces, all of these manifolds are Haken. ◀



■ **Figure 15** (i) The complete binary tree T_h of height h has pathwidth $\lceil h/2 \rceil$, cf. [11, Theorem 67]. (ii) The $k \times k$ grid graph is known to have pathwidth and treewidth both equal to k .

► **Corollary 14.** *Approximating the treewidth (resp. pathwidth) of closed, orientable 3-manifolds up to a constant factor is at least as hard as giving a constant-factor approximation of the treewidth (resp. pathwidth) of bounded-degree graphs.*

Proof. The argument for treewidth and pathwidth is the same. Given a graph G with maximum vertex degree Δ , we use the polynomial-time procedure from Theorem 2 to build a 3-manifold \mathcal{M} with $\text{tw}(\mathcal{M})$ within a constant factor of $\text{tw}(G)$. An oracle for a constant-factor approximation of $\text{tw}(\mathcal{M})$ hence gives us a constant-factor approximation of $\text{tw}(G)$ as well. ◀

► **Remark 15.** Computing a constant-factor approximation of treewidth (resp. pathwidth) for arbitrary graphs is known to be conditionally **NP**-hard under the Small Set Expansion Hypothesis [46, 55, 56]. For proving Corollary 14, however, we rely on the assumption that the graph has bounded degree. Thus the conditional hardness of approximating the treewidth (resp. pathwidth) of a 3-manifold does not directly follow. Establishing such a hardness result would add to the growing, but still relatively short list of algorithmic problems in low-dimensional topology that are known to be (conditionally) hard, cf. [2, 6, 21, 33, 37].

7 Discussion and Open Problems

We have demonstrated that 3-manifolds with JSJ decompositions with dual graphs of large treewidth (resp. pathwidth) and “sufficiently complicated” gluing maps cannot admit triangulations of low treewidth (resp. pathwidth). This provides a technique to construct a wealth of families of 3-manifolds with unbounded tree- or pathwidth that hopefully will prove to be useful for future research in the field.

One obvious limitation of our construction is a seemingly heavy restriction on the JSJ gluing maps in order to deduce a connection between the treewidth of a 3-manifold and that

of the dual graph of its JSJ decomposition. Hence, a natural question to ask is, how much this restriction on gluing maps may be relaxed while still allowing meaningful structural results about treewidth. In particular, we have the following question.

► **Question 1.** Given a 3-manifold \mathcal{M} with JSJ decomposition \mathcal{D} and no restrictions on its JSJ gluings. Is there a lower bound for $\text{tw}(\mathcal{M})$ in terms of $\text{tw}(\Gamma(\mathcal{D}))$?

Note that the assumption that we are considering the JSJ decomposition of \mathcal{M} is necessary: Consider a graph $G = (V, E)$ and a collection of 3-manifolds $\{\mathcal{M}_v\}_{v \in V}$, where \mathcal{M}_v has $\deg(v)$ torus boundary components. Assume that we glue the manifolds \mathcal{M}_v along the arcs of G to obtain a closed 3-manifold \mathcal{M} . Without restrictions on how these pieces are glued together, this cannot result in a lower bound $\text{tw}(\mathcal{M})$ in terms of $\text{tw}(G)$: we can construct Seifert fibered spaces \mathcal{M} in this way, even if $G = (V, E)$ is the complete graph with $|V|$ arbitrarily large. At the same time, Seifert fibered spaces have constant treewidth, see [27].

► **Question 2.** What is the complexity of computing the treewidth of a 3-manifold?

We believe that this should be at least as hard as computing the treewidth of a graph.

References

- 1 I. Agol. Small 3-manifolds of large genus. *Geom. Dedicata*, 102:53–64, 2003. doi:10.1023/B:GEOM.0000006584.85248.c5.
- 2 I. Agol, J. Hass, and W. Thurston. The computational complexity of knot genus and spanning area. *Trans. Am. Math. Soc.*, 358(9):3821–3850, 2006. doi:10.1090/S0002-9947-05-03919-X.
- 3 M. A. Armstrong. *Basic topology*. Undergrad. Texts Math. Springer-Verlag, New York-Berlin, 1983. Corrected reprint of the 1979 original. doi:10.1007/978-1-4757-1793-8.
- 4 D. Bachman. Stabilizing and destabilizing Heegaard splittings of sufficiently complicated 3-manifolds. *Math. Ann.*, 355(2):697–728, 2013. doi:10.1007/s00208-012-0802-4.
- 5 D. Bachman, R. Derby-Talbot, and E. Sedgwick. Heegaard structure respects complicated JSJ decompositions. *Math. Ann.*, 365(3-4):1137–1154, 2016. doi:10.1007/s00208-015-1314-9.
- 6 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_3.
- 7 D. Bachman, S. Schleimer, and E. Sedgwick. Sweepouts of amalgamated 3-manifolds. *Algebr. Geom. Topol.*, 6:171–194, 2006. doi:10.2140/agt.2006.6.171.
- 8 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.
- 9 R. Belmonte, T. Hanaka, M. Kanzaki, M. Kiyomi, Y. Kobayashi, Y. Kobayashi, M. Lampis, H. Ono, and Y. Otachi. Parameterized complexity of (A, ℓ) -path packing. *Algorithmica*, 84(4):871–895, 2022. doi:10.1007/s00453-021-00875-y.
- 10 R. H. Bing. An alternative proof that 3-manifolds can be triangulated. *Ann. Math. (2)*, 69:37–65, 1959. doi:10.2307/1970092.
- 11 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.
- 12 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. Am. Math. Soc., Providence, RI, 2013. doi:10.1090/comm/597/11877.
- 13 B. A. Burton, R. Budney, W. Pettersson, et al. Regina: Software for low-dimensional topology, 1999–2022. Version 7.2. URL: <https://regina-normal.github.io>.
- 14 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.

- 15 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.
- 16 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.
- 17 B. A. Burton and W. Pettersson. Fixed parameter tractable algorithms in combinatorial topology. In *Proc. 20th Int. Conf. Comput. Comb. (COCOON 2014)*, pages 300–311, 2014. doi:10.1007/978-3-319-08783-2_26.
- 18 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)*, pages 168–183, 2013. doi:10.1137/1.9781611973105.13.
- 19 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.
- 20 H. S. M. Coxeter. Regular honeycombs in hyperbolic space. In *Proc. Int. Congr. Math., Amsterdam, The Netherlands, September 2–9, 1954*, volume 3, pages 155–169. Erven P. Noordhoff; North-Holland, 1956. Available at <https://www.mathunion.org/fileadmin/ICM/Proceedings/ICM1954.3/ICM1954.3.ocr.pdf> (accessed: March 12, 2023).
- 21 A. de Mesmay, Y. Rieck, E. Sedgwick, and M. Tancer. The unbearable hardness of unknotting. *Adv. Math.*, 381:Paper No. 107648, 36, 2021. doi:10.1016/j.aim.2021.107648.
- 22 Henri Paul de Saint-Gervais. *Uniformization of Riemann surfaces. Revisiting a hundred-year-old theorem. Translated from the French by R. G. Burns*. Herit. Eur. Math. Eur. Math. Soc. (EMS), Zürich, 2016. doi:10.4171/145.
- 23 A. Hatcher. Notes on basic 3-manifold topology. Available at <https://pi.math.cornell.edu/~hatcher/3M/3Mfds.pdf> (accessed: February 18, 2023).
- 24 P. Heegaard. Sur l' "Analysis situs". *Bull. Soc. Math. France*, 44:161–242, 1916. doi:10.24033/bsmf.968.
- 25 K. Huszár. *Combinatorial width parameters for 3-dimensional manifolds*. PhD thesis, IST Austria, June 2020. doi:10.15479/AT:ISTA:8032.
- 26 K. Huszár. On the pathwidth of hyperbolic 3-manifolds. *Comput. Geom. Topol.*, 1(1):1–19, 2022. doi:10.57717/cgt.v1i1.4.
- 27 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.
- 28 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.
- 29 W. Jaco and J. H. Rubinstein. Layered-triangulations of 3-manifolds, 2006. 97 pages, 32 figures. arXiv:math/0603601.
- 30 W. Jaco and P. B. Shalen. A new decomposition theorem for irreducible sufficiently-large 3-manifolds. In *Algebraic and Geometric Topology*, volume 32, part 2 of *Proc. Sympos. Pure Math.*, pages 71–84. Am. Math. Soc., Providence, RI, 1978. doi:10.1090/pspum/032.2/520524.
- 31 W. H. Jaco and P. B. Shalen. Seifert fibered spaces in 3-manifolds. *Mem. Am. Math. Soc.*, 21(220):viii+192, 1979. doi:10.1090/memo/0220.
- 32 K. Johannson. *Homotopy equivalences of 3-manifolds with boundaries*, volume 761 of *Lect. Notes Math.* Springer, Berlin, 1979. doi:10.1007/BFb0085406.
- 33 D. Koenig and A. Tsvietkova. NP-hard problems naturally arising in knot theory. *Trans. Am. Math. Soc. Ser. B*, 8:420–441, 2021. doi:10.1090/btran/71.
- 34 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.
- 35 M. Lackenby. The Heegaard genus of amalgamated 3-manifolds. *Geom. Dedicata*, 109:139–145, 2004. doi:10.1007/s10711-004-6553-y.
- 36 M. Lackenby. Heegaard splittings, the virtually Haken conjecture and property (τ) . *Invent. Math.*, 164(2):317–359, 2006. doi:10.1007/s00222-005-0480-x.


- 37 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.
- 38 M. Lackenby. Algorithms in 3-manifold theory. In I. Agol and D. Gabai, editors, *Surveys in 3-manifold topology and geometry*, volume 25 of *Surv. Differ. Geom.*, pages 163–213. Int. Press Boston, 2020. doi:10.4310/SDG.2020.v25.n1.a5.
- 39 Tao Li. Heegaard surfaces and the distance of amalgamation. *Geom. Topol.*, 14(4):1871–1919, 2010. doi:10.2140/gt.2010.14.1871.
- 40 C. Manolescu. Lectures on the triangulation conjecture. In *Proc. 22nd Gökova Geom.-Topol. Conf. (GGT 2015)*, pages 1–38. Int. Press Boston, 2016. URL: <https://gokovagt.org/proceedings/2015/manolescu.html>.
- 41 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.
- 42 A. Markov. The insolubility of the problem of homeomorphy (in Russian). *Dokl. Akad. Nauk SSSR*, 121:218–220, 1958.
- 43 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.
- 44 P. Orlik, E. Vogt, and H. Zieschang. Zur Topologie gefaserner dreidimensionaler Mannigfaltigkeiten. *Topology*, 6:49–64, 1967. doi:10.1016/0040-9383(67)90013-4.
- 45 B. Poonen. Undecidable problems: a sampler. In *Interpreting Gödel*, pages 211–241. Cambridge Univ. Press, 2014. doi:10.1017/CB09780511756306.015.
- 46 P. Raghavendra and D. Steurer. Graph expansion and the unique games conjecture. In *Proc. 2010 ACM Int. Symp. Theor. Comput. (STOC’10)*, pages 755–764. ACM, New York, 2010.
- 47 M. Scharlemann and J. Schultens. Comparing Heegaard and JSJ structures of orientable 3-manifolds. *Trans. Am. Math. Soc.*, 353(2):557–584, 2001. doi:10.1090/S0002-9947-00-02654-4.
- 48 M. Scharlemann, J. Schultens, and T. Saito. Lecture notes on generalized Heegaard splittings, 2005. Open-access version of [49] with slightly different structure. [arXiv:math/0504167](https://arxiv.org/abs/math/0504167).
- 49 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.
- 50 M. Scharlemann and A. Thompson. Thin position for 3-manifolds. In *Geometric topology (Haifa, 1992)*, volume 164 of *Contemp. Math.*, pages 231–238. Am. Math. Soc., Providence, RI, 1994. doi:10.1090/conm/164/01596.
- 51 J. Schultens. *Introduction to 3-Manifolds*, volume 151 of *Grad. Stud. Math.* Am. Math. Soc., Providence, RI, 2014. doi:10.1090/gsm/151.
- 52 J. Schultens and R. Weidmann. Destabilizing amalgamated Heegaard splittings. In *Workshop on Heegaard Splittings*, volume 12 of *Geom. Topol. Monogr.*, pages 319–334. Geom. Topol. Publ., Coventry, 2007. doi:10.2140/gtm.2007.12.319.
- 53 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.
- 54 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.
- 55 Y. Wu, P. Austrin, T. Pitassi, and D. Liu. Inapproximability of treewidth, one-shot pebbling, and related layout problems. *J. Artificial Intelligence Res.*, 49:569–600, 2014. doi:10.1613/jair.4030.
- 56 K. Yamazaki. Inapproximability of rank, clique, Boolean, and maximum induced matching-widths under small set expansion hypothesis. *Algorithms (Basel)*, 11(11):Paper No. 173, 10, 2018. doi:10.3390/a11110173.

A Selected Results from 3-Manifold Topology

In this appendix we collect important facts about the topology of 3-manifolds we rely on. We do not aim to reproduce their complete proofs; indeed, those are often quite technical or include lengthy case analyses. However, we do aim to give precise statements and references.

A.1 Weak reductions

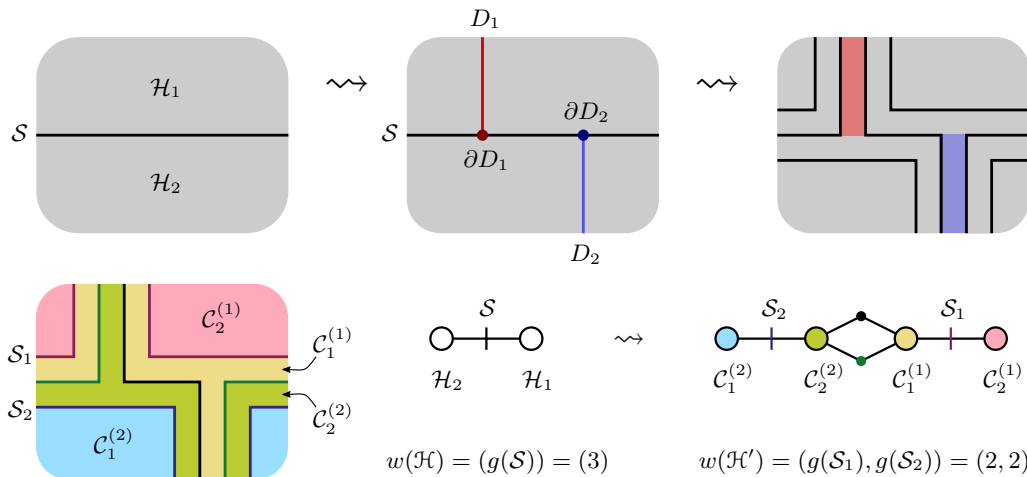
► **Example 16.** The 3-torus $\mathbb{T}^3 = \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1$ is a closed 3-manifold built by identifying opposite faces of a solid cube K (Figure 16a). It is well-known that \mathbb{T}^3 has Heegaard genus $g(\mathbb{T}^3) = 3$. A genus-three Heegaard splitting $\mathcal{H} = (\mathcal{H}_1, \mathcal{H}_2, \mathcal{S})$ of \mathbb{T}^3 can be obtained as follows: Upon the face-identifications, the 1-skeleton of K becomes a bouquet \bigcirc of three circles. We let $\mathcal{H}_1 = \overline{N_{\mathbb{T}^3}(\bigcirc)}$ be a closed regular neighborhood of \bigcirc in \mathbb{T}^3 , \mathcal{H}_2 be the closure of $\mathbb{T}^3 \setminus \mathcal{H}_1$, and $\mathcal{S} = \mathcal{H}_1 \cap \mathcal{H}_2$ be the resulting splitting surface (Figure 16b).



(a) The 3-torus $\mathbb{T}^3 = \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1$ can be obtained by identifying opposite faces of a solid cube. (b) A surface $\mathcal{S} \subset \mathbb{T}^3$ defining a Heegaard splitting $(\mathcal{H}_1, \mathcal{H}_2, \mathcal{S})$ of \mathbb{T}^3 of minimum genus three. (c) Two essential disks $D_i \subset \mathcal{H}_i$ ($i = 1, 2$) witnessing the weak reducibility of $(\mathcal{H}_1, \mathcal{H}_2, \mathcal{S})$.

■ **Figure 16** A weakly reducible minimum-genus Heegaard splitting of the 3-torus $\mathbb{T}^3 = \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1$.

Note that D_1 and D_2 in Figure 16c are essential disks in the handlebodies \mathcal{H}_1 and \mathcal{H}_2 , respectively, and $\partial D_1 \cap \partial D_2 = \emptyset$, hence the splitting $\mathcal{H} = (\mathcal{H}_1, \mathcal{H}_2, \mathcal{S})$ is weakly reducible. Now, a weak reduction at \mathcal{S} may be carried out as presented on Figure 17. We refer to [49, p. 108–111] for details (cf. [48, p. 61–62]).



■ **Figure 17** Schematic “movie” of a weak reduction of the genus-three Heegaard splitting of \mathbb{T}^3 .

A.2 Seifert fibered spaces and JSJ decompositions

A key assertion about the manifold \mathcal{M}_G constructed in Section 5 is Claim 7, according to which the dual graph $\Gamma(\mathcal{D})$ of the JSJ decomposition \mathcal{D} of \mathcal{M}_G satisfies $\Gamma(\mathcal{D}) = G$. This can be shown through the same argument that is outlined by Lackenby in [37, p. 591]. However, while we use the same building blocks as Lackenby (namely the manifolds $\mathcal{N}(k)$), we glue them together via different (*high-distance*) maps. Therefore, some explanations are in order. To this end, closely following [23], we recall some basic definitions and important results.

Seifert fiberings A *model Seifert fibering* of the solid torus $\mathbb{S}^1 \times D$ is a decomposition of $\mathbb{S}^1 \times D$ into disjoint circles, called *fibers*, obtained as follows. Consider the solid cylinder $[0, 1] \times D$ and identify its boundary disks $\{0\} \times D$ and $\{1\} \times D$ through a rotation by $2\pi p/q$, where p and q are relatively prime integers with $q > 0$. After this identification the segment $[0, 1] \times \{0\}$ becomes a fiber $\mathbb{S}^1 \times \{0\}$, whereas every other fiber consists of q segments of the form $[0, 1] \times \{x\}$. Next, a *Seifert fibering* of 3-manifold \mathcal{M} is a decomposition of \mathcal{M} into disjoint circles (also called *fibers*), such that each fiber has a neighborhood homeomorphic to that of a fiber in some model Seifert fibering of $\mathbb{S}^1 \times D$ via a fiber-preserving homeomorphism. A 3-manifold that admits a Seifert fibering is called a *Seifert fibered space*. The following classical result says that most Seifert fibered spaces admit a unique Seifert fibering.

► **Theorem 17** ([44, 54], adapted from [23, Theorem 2.3]). *Every orientable Seifert fibered space admits a unique Seifert fibering up to fiber-preserving homeomorphism, except for:*

- $\mathbb{S}^1 \times D$, the solid torus (various model fiberings; see above),
- $\mathbb{S}^1 \tilde{\times} \mathbb{S}^1 \tilde{\times} I$, the orientable \mathbb{S}^1 -bundle over the Möbius band (two fiberings),
- Seifert fibered spaces over $\mathbb{R}P^2$ with one exceptional fiber (two fiberings),
- lens spaces $L(p, q)$ including \mathbb{S}^3 and $\mathbb{S}^1 \times \mathbb{S}^2$ (various fiberings), and
- $\mathbb{S}^1 \tilde{\times} \mathbb{S}^1 \tilde{\times} \mathbb{S}^1$, the orientable \mathbb{S}^1 -bundle over the Klein bottle (two fiberings).

► **Corollary 18.** *The manifolds $\mathcal{N}(k)$ defined in Section 5 admit unique Seifert fiberings.*

Proof. The statement follows by observing that $\mathcal{N}(k)$ is a Seifert fibered space over the torus, and so it does not belong to any of the exceptional manifolds mentioned in Theorem 17. ◀

► **Proposition 19** (folklore, see [23, Proposition 1.6]). *If an irreducible 3-manifold $\tilde{\mathcal{M}}$ is a covering space of the interior of another 3-manifold \mathcal{M} , then \mathcal{M} is irreducible as well.*

▷ **Claim 20** (folklore, see [23, p. 11]). For any compact, connected and orientable surface F (possibly with boundary) other than \mathbb{S}^2 , the interior of $F \times \mathbb{S}^1$ has \mathbb{R}^3 as its universal cover.¹⁴

Proof. If $F \not\cong \mathbb{S}^2$, it is well-known that the universal cover of F is \mathbb{R}^2 . Indeed, for $F \cong \mathbb{T}^2$ a covering map $\mathbb{R}^2 \rightarrow F$ can be defined through the regular tiling of the Euclidean plane via squares, while for $F \cong \mathbb{S}_g$ with genus $g > 1$ such a map can be obtained through the regular tiling of the hyperbolic plane via regular $4g$ -gons, with $4g$ such polygons meeting at each vertex [20]. If F has non-empty boundary, then by the theory of covering spaces [3, Theorem 10.19] the interior of F has some universal cover \mathcal{X} , and by the *uniformization theorem* [22, Theorem XII.0.1], this space \mathcal{X} – being a non-compact, open, simply-connected surface – must be homeomorphic to \mathbb{R}^2 . As the universal cover of \mathbb{S}^1 is \mathbb{R} , it follows that the universal cover of $F \times \mathbb{S}^1$ is \mathbb{R}^3 . ◀

¹⁴The claim holds for non-orientable surfaces as well with the exception of $\mathbb{R}P^2$.

▷ Claim 21 (folklore, see [23, p. 14]). A 2-sided torus $T \subset \mathcal{M}$ in an irreducible 3-manifold \mathcal{M} is compressible if and only if T either bounds a solid torus $\mathbb{S}^1 \times D \subset \mathcal{M}$ or lies in a ball in \mathcal{M} .

► **Corollary 22.** *For any $k \in \mathbb{N}$, $\mathcal{N}(k)$ is irreducible and has incompressible boundary.*

Proof. The irreducibility of $\mathcal{N}(k)$ is a consequence of Proposition 19 and Claim 20. Then, the incompressibility of $\partial\mathcal{N}(k)$ follows from Claim 21 by noting that $\partial\mathcal{N}(k)$ is a union of tori and, by construction of $\mathcal{N}(k)$, none of them bounds a solid torus or lies in a ball in $\mathcal{N}(k)$. ◀

► **Proposition 23.** *The manifold \mathcal{M}_G constructed in Section 5 is irreducible.*

Proof. Let $G = (V, E)$ be a multigraph on n nodes with m arcs, and $\mathbf{T} = \{T_1, \dots, T_m\}$ be the collection of pairwise disjoint tori $T_i \subset \mathcal{M}_G$ that are the images of the boundary tori of the manifolds $\mathcal{N}_v \cong \mathcal{N}(d_v)$ after performing the torus-gluing in the construction of \mathcal{M}_G .


Let us assume that \mathcal{M}_G is *not* irreducible, i.e., there is some reducing 2-sphere $S \subset \mathcal{M}_G$. By definition, S does not bound a ball in \mathcal{M}_G . Note that S cannot be entirely contained in any of the building blocks $\mathcal{N}_v \subset \mathcal{M}_G$, otherwise it would imply the reducibility of $\mathcal{N}_v \cong \mathcal{N}(d_v)$, contradicting Corollary 22. It follows that $S \cap (\cup_{i=1}^m T_i) \neq \emptyset$. Thus, after an isotopy if necessary, we may assume that S intersects $\cup_{i=1}^m T_i$ transversally, and the intersection $S \cap (\cup_{i=1}^m T_i)$ is a union of pairwise disjoint closed curves $\alpha = \{\alpha_1, \dots, \alpha_\ell\}$ with ℓ being minimal.

Now let $\alpha_j \in \alpha$ be a curve that is *innermost* in S , i.e., one of the two disks bounded by α_j in S , say $D \subset S$, does not contain any other curves from α . After relabeling if necessary, let $T_j \in \mathbf{T}$ denote the torus for which $\alpha_j = S \cap T_j$. Note that, if α_j also bounds a disk $D' \subset T_j$, then $D \cup D'$ is a sphere entirely contained in one of the building blocks \mathcal{N}_w incident to T_j . Since \mathcal{N}_w is irreducible (Corollary 22), $D \cup D'$ bounds a ball $B \subset \mathcal{N}_w$, which can be used to eliminate α_j by first isotoping D inside B to coincide with D' and then pushing D' off the torus T_j , contradicting the minimality of α . Hence the disk D is essential in \mathcal{N}_w . However, this contradicts the incompressibility of T_j in \mathcal{N}_w (Corollary 22). ◀


Proof of Claim 7. We follow the notation introduced in the proof of Proposition 23. By Corollary 18, each building block \mathcal{N}_v admits a unique Seifert fibering. In particular, each boundary component of \mathcal{N}_v has a uniquely determined fibering. Therefore, when two building blocks \mathcal{N}_u and $\mathcal{N}_{u'}$ are glued together along boundary tori $T \subset \partial\mathcal{N}_u$ and $T' \subset \partial\mathcal{N}_{u'}$, the resulting 3-manifold is not Seifert fibered, unless the gluing map aligns the (identical) Seifert fiberings of T and T' . Now, it follows from the definition of distance [6, Section 4.1] that gluing maps of distance at least one do *not* align these Seifert fiberings. (In fact, for any two fixed Seifert fiberings of the torus, there are precisely two maps that align them.) Hence, upon the manifold \mathcal{M}_G is constructed, the Seifert fibrations of adjacent building blocks induce different fibrations on each torus in $\mathbf{T} = \{T_1, \dots, T_m\}$. Now, by construction of the building blocks, none of the tori in \mathbf{T} bounds a solid torus (cf. [37, p. 591]), therefore they are all incompressible (cf. Claim 21) and pairwise non-parallel. This implies that the building blocks \mathcal{N}_v ($v \in V$) are indeed the JSJ pieces of \mathcal{M}_G and the tori in \mathbf{T} comprise its JSJ tori. ◀

B Triangulating the building blocks

Here we describe the triangulation $\mathcal{T}(k)$ of the building block $\mathcal{N}(k)$ introduced in Section 5. Topologically, $\mathcal{N}(k) = \mathbb{T}_k^2 \times \mathbb{S}^1$, where \mathbb{T}_k^2 is obtained by removing the interiors of k pairwise disjoint disks from the torus \mathbb{T}^2 . To construct $\mathcal{T}(k)$, first we triangulate \mathbb{T}_k^2 with $3k+2$ triangles as shown in Figure 18. This triangulation of \mathbb{T}_k^2 naturally lifts to a decomposition of $\mathbb{T}_k^2 \times [0, 1]$ into $3k+2$ *triangular prisms*, each of which can be triangulated with three tetrahedra, see Figure 19. The construction of $\mathcal{T}(k)$ is concluded by identifying the triangulations of $\mathbb{T}_k^2 \times \{0\}$ and $\mathbb{T}_k^2 \times \{1\}$ via the identity map. It immediately follows that $\mathcal{T}(k)$ consists of $9k+6 = O(k)$ tetrahedra and induces a two-triangle triangulation at each boundary torus.



■ **Figure 18** A minimal triangulation of \mathbb{T}_k^2 with $3k+2$ triangles, drawn for $k=3$. The extension of the drawing for larger k is straightforward.



■ **Figure 19** (i) The *triangular prism* (the shape described as $t \times [0, 1]$, where t denotes a triangle) can be triangulated with three tetrahedra. (ii) A triangulation of the triangular prism and its corresponding dual graph. (iii) A triangular prism with its top and bottom triangles identified. Reproduced from [27, Figure 13].