ON THE TREEWIDTH OF TRIANGULATED 3-MANIFOLDS*

Kristóf Huszár,[†] Jonathan Spreer,[‡] and Uli Wagner[§]

ABSTRACT. In graph theory, as well as in 3-manifold topology, there exist several width-type parameters to describe how "simple" or "thin" a given graph or 3-manifold is. These parameters, such as pathwidth or treewidth for graphs, or the concept of thin position for 3-manifolds, play an important role when studying algorithmic problems; in particular, there is a variety of problems in computational 3-manifold topology—some of them known to be computationally hard in general—that become solvable in polynomial time as soon as the dual graph of the input triangulation has bounded treewidth.

In view of these algorithmic results, it is natural to ask whether every 3-manifold admits a triangulation of bounded treewidth. We show that this is not the case, i.e., that there exists an infinite family of closed 3-manifolds not admitting triangulations of bounded pathwidth or treewidth (the latter implies the former, but we present two separate proofs).

We derive these results from work of Agol, of Scharlemann and Thompson, and of Scharlemann, Schultens and Saito by exhibiting explicit connections between the topology of a 3-manifold $\mathcal{M}$ on the one hand and width-type parameters of the dual graphs of triangulations of $\mathcal{M}$ on the other hand, answering a question that had been raised repeatedly by researchers in computational 3-manifold topology. In particular, we show that if a closed, orientable, irreducible, non-Haken 3-manifold $\mathcal{M}$ has a triangulation of treewidth (resp. pathwidth) k then the Heegaard genus of $\mathcal{M}$ is at most $18(k+1)$ (resp. $4(3k+1)$).

1 Introduction

In the field of 3-manifold topology many fundamental problems can be solved algorithmically. Famous examples include deciding whether a given knot is trivial [27], deciding whether a given 3-manifold is homeomorphic to the 3-sphere [54, 65], and, more generally (based on Perelman's proof of Thurston's geometrization conjecture [38]), deciding whether two given 3-manifolds are homeomorphic, see, e.g., [3, 41, 60]. The algorithm for solving the homeomorphism problem is still purely theoretical, and its complexity remains largely unknown [41, 43]. In contrast, the first two problems are known to lie in the intersection of the complexity classes **NP** and **co-NP** [28, 35, 40, 42, 58, 67].[1]

---

†*Institute of Science and Technology Austria*, `kristof.huszar@ist.ac.at`

‡*School of Mathematics and Statistics, The University of Sydney*, `jonathan.spreer@sydney.edu.au`

§*Institute of Science and Technology Austria*, `uli@ist.ac.at`

[1]The proof of **co-NP** membership for 3-sphere recognition assumes the Generalized Riemann Hypothesis.

Moreover, implementations of, for instance, algorithms to recognize the 3-sphere exist out-of-the-box (e.g., using the computational 3-manifold software *Regina* [14]) and exhibit practical running times for virtually all known inputs.

In fact, many topological problems with implemented algorithmic solutions solve problem 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 [23, 24]. This effort has proven to be highly effective and, today, there exist numerous FPT algorithms in the field [15, 16, 17, 18, 44]. 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

- taut angle structures[3] of what is called ideal triangulations with torus boundary components in running time $O(7^k \cdot n)$ [18];

- optimal Morse matchings[4] in the Hasse diagram of $\mathcal{T}$ in $O(4^{k^2+k} \cdot k^3 \cdot \log k \cdot n)$ [16];

- the Turaev–Viro invariants[5] for parameter $r \geq 3$ in $O((r-1)^{6(k+1)} \cdot k^2 \cdot \log r \cdot n)$ [17];

- 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 [15].[6]

Some of these results are not purely theoretical—as is sometimes the case with FPT algorithms—but are implemented and outperform previous state-of-the-art implementations for typical input. As a result, they have a significant practical impact. This is in particular the case for the algorithm to compute Turaev–Viro invariants [17, 44].

Note that treewidth—the dominating factor in the running times given above—is a combinatorial quantity linked to a triangulation, not a topological invariant of the underlying manifold. 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 a triangulation $\mathcal{T}'$ of the same manifold with smaller treewidth.

This approach faces severe difficulties. By a theorem due to Kirby and Melvin [37], 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 [17], would constitute a polynomial time solution for a #**P**-hard problem. Furthermore, known facts imply that

---

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

[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 the minimum 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 [20].

most triangulations of *most* 3-manifolds must have large treewidth[7] (see Propositions 22 and 23 in Appendix A). 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.

In this article we show that this is actually not the case, answering a question that had been raised repeatedly by researchers in computational 3-manifold topology.[8] More specifically, we prove the following two statements.

Theorem 1. *There exists an infinite family of* 3-*manifolds which does not admit triangulations with dual graphs of uniformly bounded pathwidth.*

Theorem 2. *There exists an infinite family of* 3-*manifolds which does not admit triangulations with dual graphs of uniformly bounded treewidth.*

We establish the above results through the following theorems, which are the main contributions of the present paper. The necessary terminology is introduced in Section 2.

Theorem 3. *Let $\mathcal{M}$ be a closed, orientable, irreducible, non-Haken* 3-*manifold and let $\mathcal{T}$ be a triangulation of $\mathcal{M}$ with dual graph $\Gamma(\mathcal{T})$ of pathwidth $\mathrm{pw}(\Gamma(\mathcal{T})) \leq k$. Then $\mathcal{M}$ has Heegaard genus $\mathfrak{g}(\mathcal{M}) \leq 4(3k+1)$.*

Theorem 4. *Let $\mathcal{M}$ be a closed, orientable, irreducible, non-Haken* 3-*manifold and let $\mathcal{T}$ be a triangulation of $\mathcal{M}$ with dual graph $\Gamma(\mathcal{T})$ of treewidth $\mathrm{tw}(\Gamma(\mathcal{T})) \leq k$. Then $\mathcal{M}$ has Heegaard genus $\mathfrak{g}(\mathcal{M}) \leq 18(k+1)$.*

By a result of Agol [1] (Theorem 17 in this paper), there exist closed, orientable, irreducible, non-Haken 3-manifolds of arbitrarily large Heegaard genus. Combining this result with Theorems 3 and 4 thus immediately implies Theorems 1 and 2.

Remark. Note that Theorem 1 can be directly deduced from Theorem 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 3 is considerably simpler than that of Theorem 4, it already illustrates several key concepts and ideas which we are building upon in the proof of Theorem 4.

The paper is organized as follows. After going over the preliminaries in Section 2, we give an overview of selected width-type graph parameters in Section 3. Most notably, we propose the *congestion* of a graph (also known as *carving width*) as an alternative choice of a parameter for FPT algorithms in 3-manifold topology (cf. Appendix C). Section 4 is devoted to results from 3-manifold topology which we build upon. In Section 5 we then prove Theorem 1, and in Section 6 we prove Theorem 2.

---

[7]It is known that, given $k \in \mathbb{N}$, there exist constants $C, C_k > 1$ such that there are at least $C^{n \log(n)}$ 3-manifolds which can be triangulated with $\leq n$ tetrahedra, whereas there are at most C_k^n triangulations with treewidth $\leq k$ and $\leq n$ tetrahedra.

[8]The question whether every 3-manifold admits a triangulation of bounded treewidth, and variations thereof have been asked at several meetings and open problem sessions including an Oberwolfach meeting in 2015 [12, Problem 8] (formulated in the context of knot theory).

[9]This is immediate from the definitions of treewidth and pathwidth, see Section 3.

2   Preliminaries

In this section we recall some basic concepts and terminology of graph theory, 3-manifolds, triangulations, and parameterized complexity theory.

Graphs vs. triangulations.   Following several authors in the field, 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.

2.1   Graphs

For general background on graph theory we refer to [22].

A *graph* (more specifically, a *multigraph*) $G = (V, E)$ is an ordered pair consisting of a finite set $V = V(G)$ of *nodes* and a multiset $E = E(G)$ of unordered pairs of nodes, called *arcs*. We allow *loops*, i.e., an arc $e \in E$ might itself be a multiset, e.g., $e = \{v, v\}$ for some $v \in V$. The *degree* of a node $v \in V$, denoted by $\deg(v)$, equals the number of arcs containing it, counted with multiplicity. In particular, a loop $\{v, v\}$ contributes two to the degree of v. For every node $v \in V$ of a graph G, its *star* $\mathrm{st}_G(v)$ denotes the set of edges incident to v. A graph is called *k-regular* if all of its nodes have the same degree $k \in \mathbb{N}$. A *tree* is a connected graph with n nodes and $n-1$ arcs. A node of degree one is called a *leaf*.

2.2   3-Manifolds and their triangulations

For an introduction to the topology and geometry of 3-manifolds and to their triangulations we refer to the textbook [59] and to the seminal monograph [66].

A 3-*manifold with boundary* is a topological space[10] $\mathcal{M}$ such that each point $x \in \mathcal{M}$ has a neighborhood which either looks like (i.e., is homeomorphic to) the Euclidean 3-space $\mathbb{R}^3$ or the closed upper half-space $\{(x, y, z) \in \mathbb{R}^3 : z \geq 0\}$. The points of $\mathcal{M}$ that do not have a neighborhood homeomorphic to $\mathbb{R}^3$ constitute the *boundary $\partial \mathcal{M}$* of $\mathcal{M}$. A compact 3-manifold is *closed* if it has an empty boundary.

---

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

Informally, two 3-manifolds are equivalent (or *homeomorphic*) if one can be turned into the other by a continuous, reversible deformation. In other words, when talking about a 3-manifold, we are not interested in its particular shape, but only in its qualitative properties (topological invariants), such as "number of boundary components", or "connectedness".

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

Handle decompositions.   Every compact 3-manifold can be built from finitely many building blocks called 0-, 1-, 2-, and 3-handles. 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 build 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.

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.

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 $\mathcal{M}$.[11]

A 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[12] 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 closed and orientable surface (possibly disconnected). 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.

---

[11]A standard example of a compressible surface is a torus (or any other orientable surface) embedded in the 3-sphere $\mathbb{S}^3$, and of an incompressible surface is the 2-sphere $\mathbb{S}^2 \times \{x\} \subset \mathbb{S}^2 \times \mathbb{S}^1$.

[12]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.

Handlebodies and compression bodies. We have already discussed handle decompositions of 3-manifolds. Closely related are the notions of handlebody and compression body.

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

Dual to this construction, a compression body can be built by starting with a closed, orientable surface $\mathcal{F}$, thickening it to $\mathcal{F} \times [0,1]$, attaching 2-handles along $\mathcal{F} \times \{0\}$, and filling in some of the resulting 2-spheres with 3-balls. This 3-manifold $\mathcal{N}$ has again two sets of boundary components given by $\partial_+\mathcal{N} = \mathcal{F} \times \{1\}$ and $\partial_-\mathcal{N} = \partial\mathcal{N} \setminus \partial_+\mathcal{N}$.

In accordance with [31, 32], we call $\partial_+\mathcal{N}$ the *top boundary*, and $\partial_-\mathcal{N}$ the *lower boundary* of $\mathcal{N}$. Note that, by construction, we always have $g(\partial_+\mathcal{N}) \geq g(\partial_-\mathcal{N})$. Moreover, if $\partial_-\mathcal{N} = \emptyset$, then the compression body $\mathcal{N}$ is actually a handlebody.

Heegaard splittings. Let $\mathcal{M}$ be a 3-manifold, possibly with boundary. A *Heegaard splitting* of $\mathcal{M}$ is a decomposition $\mathcal{M} = \mathcal{N} \cup_\mathcal{S} \mathcal{K}$ (i.e., $\mathcal{N} \cup \mathcal{K} = \mathcal{M}$ and $\mathcal{N} \cap \mathcal{K} = \mathcal{S}$) into compression bodies $\mathcal{N}$ and $\mathcal{K}$ with $\mathcal{S} = \partial_+\mathcal{N} = \partial_+\mathcal{K}$ and $\partial\mathcal{M} = \partial_-\mathcal{N} \cup \partial_-\mathcal{K}$. The *Heegaard genus* of $\mathcal{M}$, denoted $\mathfrak{g}(\mathcal{M})$, is the minimum possible genus $g(\mathcal{S})$ over all Heegaard splittings of $\mathcal{M}$.

A fundamental result of Moise [46] implies that every compact orientable 3-manifold admits a Heegaard splitting (also see the survey [55]).

Example 5 (Heegaard splittings from handle decompositions)**.** When building up a closed connected 3-manifold $\mathcal{M}$ from handles, one may assume that (possibly after deforming the attaching maps) all 0- and 1-handles are attached before any 2- or 3-handles. Defining $\mathcal{H}_1$ to be the union of the 0- and 1-handles, and $\mathcal{H}_2$ to be the union of the 2- and 3-handles yields a Heegaard splitting $\mathcal{M} = \mathcal{H}_1 \cup_\mathcal{S} \mathcal{H}_2$, $\mathcal{S} = \partial\mathcal{H}_1 = \partial\mathcal{H}_2$, into handlebodies $\mathcal{H}_1$ and $\mathcal{H}_2$.

Example 6 (Lens spaces)**.** Among the best known 3-manifolds are the closed orientable 3-manifolds of Heegaard genus one. These manifolds are also known under the name of *lens spaces*. To construct them, let p, q be two positive co-prime integers. The lens space $\mathrm{L}(p,q)$ is then obtained by taking two solid tori $\mathbb{T}_i = \mathbb{S}^1 \times D$, $i = 1, 2$, and gluing them together along their boundaries in a way such that the meridian $m_1 = \{x_1\} \times \partial D \subset \mathbb{T}_1$ of $\mathbb{T}_1$ is mapped onto the curve of $\mathbb{T}_2$ which wraps p times around the longitude $l_2 = \mathbb{S}^1 \times \{y_2\} \subset \mathbb{T}_2$ and q times around the meridian $m_2 = \{x_2\} \times \partial D \subset \mathbb{T}_2$.

Linear splittings. In their work on thin position [57], discussed in Section 4, Scharlemann and Thompson consider a generalization of Heegaard splittings, we call a *linear splitting*, which arises naturally from more complicated sequences of handle attachments, e.g., when building up a manifold by first only attaching some of the 0- and 1-handles before attaching

2- and 3-handles (cf. Example 7).[13]

A *linear splitting* of a 3-manifold $\mathcal{M}$ is a decomposition

$$\mathcal{M} = (\mathcal{N}_1 \cup_{\mathcal{R}_1} \mathcal{K}_1) \cup_{\mathcal{F}_1} (\mathcal{N}_2 \cup_{\mathcal{R}_2} \mathcal{K}_2) \cup_{\mathcal{F}_2} \cdots \cup_{\mathcal{F}_{s-1}} (\mathcal{N}_s \cup_{\mathcal{R}_s} \mathcal{K}_s),$$

where $(\mathcal{N}_1, \mathcal{K}_1, \ldots, \mathcal{N}_s, \mathcal{K}_s)$ 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 $\mathcal{R}_1, \mathcal{F}_1, \ldots, \mathcal{R}_{s-1}, \mathcal{F}_{s-1}, \mathcal{R}_s$ in $\mathcal{M}$. More precisely,

$$\mathcal{R}_i = \mathcal{N}_i \cap \mathcal{K}_i = \partial_+ \mathcal{N}_i = \partial_+ \mathcal{K}_i \quad \text{and} \quad \mathcal{F}_j = \mathcal{K}_j \cap \mathcal{N}_{j+1} = \partial_- \mathcal{K}_j = \partial_- \mathcal{N}_{j+1}.$$

The lower boundaries of $\mathcal{N}_1$ and $\mathcal{K}_s$ constitute the boundary of $\mathcal{M}$, i.e., $\partial \mathcal{M} = \partial_- \mathcal{N}_1 \cup \partial_- \mathcal{K}_s$.

Example 7 (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 $(\mathcal{N}_1, \mathcal{K}_1, \ldots, \mathcal{N}_s, \mathcal{K}_s)$.

Graph splittings and fork complexes. The decomposition described above exhibits a linear structure. Here we introduce a more general approach of decomposing a 3-manifold into compression bodies following a graph structure [56].

The main difference is that now we allow the lower boundary components of a compression body to be glued to lower boundary components of distinct compression bodies. The top boundary of a compression body, however, is still identified with the top boundary of a single other compression body. This structure can be represented by a so-called *fork complex* (which is essentially a labeled graph) in which the compression bodies of the decomposition are modeled by *fork*s. More precisely, an *n-fork* is a tree F with $n+2$ nodes $V(F) = \{g, p, t_1, \ldots, t_n\}$ 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 *tines* of F, respectively (Figure 1(i) shows a 0- and a 3-fork). We think of a fork $F = F_{\mathcal{N}}$ as an abstraction of a compression body $\mathcal{N}$, such that the grip of F corresponds to $\partial_+ \mathcal{N}$, whereas the tines correspond to the connected components of $\partial_- \mathcal{N}$.

Informally, a *fork complex* $\boldsymbol{F}$ (representing a given decomposition 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

---

[13]While this construction is sometimes called a *generalized Heegaard splitting*, we prefer the more expressive term of a *linear splitting* [31] to make a distinction from the even more general *graph splittings* [32].
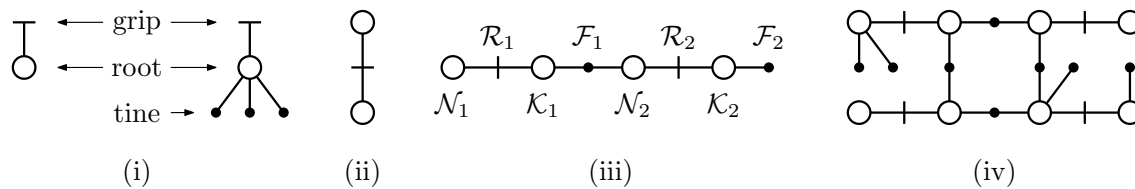
Figure 1: Fork complexes describing Heegaard (ii), linear (iii), and graph splittings (iv)

$\mathcal{M}$ (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 [56, Definition 5.1.7]. The pair $(\boldsymbol{F}, \rho)$ is called a *graph splitting*. See Figure 1 for illustrations, and [56, Section 5.1] for further details.

Triangulations. In this article we typically describe 3-manifolds by *triangulations* (also referred to as *generalised*, *semi-simplicial*, or *singular* triangulations in the literature).[14] That is, a finite collection of abstract tetrahedra, glued together in pairs along their triangular faces (called *triangles*). As a result of these *face gluings*, many tetrahedral edges (or vertices) are glued together and we refer to the result as a single *edge (or vertex) of the triangulation*. A triangulation $\mathcal{T}$ describes a closed 3-manifold if no tetrahedral edge is identified with itself in reverse, and the boundary of a small neighborhood around each vertex is a 2-sphere. Similarly, $\mathcal{T}$ describes a 3-manifold with boundary if, in addition, the boundary of small neighborhoods around some of the vertices are disks.

Given a triangulation $\mathcal{T}$ of a closed 3-manifold, its *dual graph* $\Gamma(\mathcal{T})$ (also called the *face pairing graph*) is the graph with one node per tetrahedron of $\mathcal{T}$, and with an arc between two nodes for each face gluing between the corresponding pair of tetrahedra. By construction, the dual graph is a 4-regular multigraph. Since every triangulation $\mathcal{T}$ can be linked to its dual graph $\Gamma(\mathcal{T})$ this way, we often attribute properties of $\Gamma(\mathcal{T})$ directly to $\mathcal{T}$.

2.3 Parameterized complexity and fixed parameter tractability

There exist various notions and concepts of a refined complexity analysis for theoretically difficult problems. Parameterized complexity, due to Downey and Fellows [23, 24], identifies a parameter on the set of inputs, 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 for 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

---

[14]Triangulations, in the present sense, provide a very efficient way to describe 3-manifolds: more than $11{,}000$ topologically distinct 3-manifolds can be triangulated with 11 tetrahedra or less [13, 45]. This efficiency comes at a price: we allow self-identifications (e.g., gluings of pairs of triangular faces of the same tetrahedron), and thus a triangulation is generally non-regular when seen as a (simplicial) cell-complex. However, this deficiency can be overcome by passing to the first barycentric subdivision. A second barycentric subdivision then yields a simplicial complex. In particular, every triangulation can be turned into a simplicial complex of size at most $24^2 = 576$ times larger than the original triangulation.

The aforementioned theorem of Moise [46] tells us that every 3-manifold has a triangulation.

running time $O(f(p(I)) \cdot \text{poly}(n))$, where n is the size of the input I, and $f \colon \mathbb{N} \to \mathbb{N}$ is arbitrary (computable). 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 $\mathcal{P}$, cf. the results listed in the Introduction.

In computational 3-manifold topology, a very important set 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 (see [44] for such a result, using the first Betti number as parameter). However, most FPT-results in the field use parameters of the dual graph of a triangulation which greatly depend on the choice of the triangulation: every 3-manifold admits a triangulation with arbitrarily high graph parameters—for all parameters considered in this article. The aim of this work is to link these parameters to topological invariants in the only remaining sense: given a 3-manifold $\mathcal{M}$, find lower bounds for graph parameters of dual graphs of triangulations ranging over all triangulations of $\mathcal{M}$.

3   Width-type graph parameters

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. For instance, several graph theoretical questions have a simple answer if one asks them about trees, or graphs which are similar to trees in some sense. Width-type parameters make this sense of similarity precise [30]. 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. (See Appendix C for computational aspects.)

Treewidth and pathwidth.   The concepts of treewidth and pathwidth were introduced by Robertson and Seymour in their early papers on graph minors [52, 53], also see the surveys [6, 8, 10]. Given a graph G, its *treewidth* $\text{tw}(G)$ measures how tree-like the graph is.

Definition 8 (Tree decomposition, treewidth)**.** A *tree decomposition* of a graph $G = (V, E)$ is a tree T with nodes $B_1, \ldots, B_m \subseteq V$, also called *bags*, such that

1. $B_1 \cup \ldots \cup B_m = V$,

2. if $v \in B_i \cap B_j$ then $v \in B_k$ for all bags B_k of T in the path between B_i and B_j, in other words, the bags containing v span a (connected) subtree of T,

3. for every arc $\{u, v\} \in E$, there exists a node B_i such that $\{u, v\} \subseteq B_i$.

The *width* of a tree decomposition equals the size of the largest bag minus one. The *treewidth* $\text{tw}(G)$ is the minimum width among all possible tree decompositions of G.

Definition 9 (Path decomposition, pathwidth)**.** A *path decomposition* of a graph $G = (V, E)$ is a tree decomposition for which the tree T is required to be a *path*. The *pathwidth* $\text{pw}(G)$ of a graph G is the minimum width of any path decomposition of G.

Cutwidth.   The *cutwidth* $\mathrm{cw}(G)$ of a graph G is the graph-analogue of the linear width of a manifold (to be discussed in Section 4). If we order the nodes $\{v_1, \ldots, v_n\} = V(G)$ of G on a line, the set of arcs running from a node v_i, $i \leq \ell$, to a node v_j, $j > \ell$, is called a cutset C_ℓ of the ordering. The cutwidth $\mathrm{cw}(G)$ is defined to be the cardinality of the largest cutset, minimized over all linear orderings of $V(G)$.

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

Theorem 10 (Bodlaender, Theorems 47 and 49 from [7][15]). *Given a graph G, we have*

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

Congestion.   In [5] Bienstock introduced *congestion*, a generalization of cutwidth, which is a quantity related to treewidth in a similar way as cutwidth to pathwidth (compare Theorems 10 and 12).

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 "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., [51]). Here we work with the following.

Definition 11 (Congestion[16]). Let $T_{\{1,3\}}$ be the set of unrooted binary trees.[17] 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.

---

[15]The inequality $\mathrm{cw}(G) \leq \Delta(G)\,\mathrm{pw}(G)$ seems to be already present in the earlier work of Chung and Seymour [19] on the relation of cutwidth to another parameter called topological bandwith (see Theorem 2 in [19]). Pathwidth plays an intermediate, connecting role there. However, the inequality is phrased and proved explicitly by Bodlaender in [7].

[16]It is important to note that congestion in the sense of Definition 11 is also known as *carving width*, a term which was coined by Robertson and Seymour in [61]. 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.

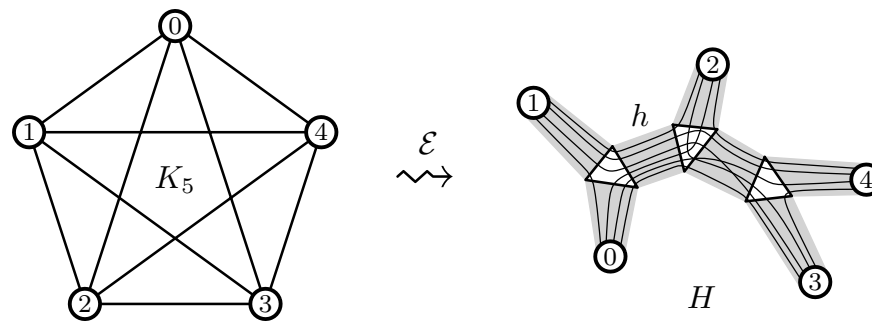[17]An *unrooted binary tree* is a tree in which each node is incident to either one or three arcs.

Figure 2: 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 12 (Bienstock, p. 108–111 of [4]). *Given a graph G with maximum degree $\Delta(G)$, we have*[18]
$$\max\left\{\tfrac{2}{3}(\mathrm{tw}(G) + 1), \Delta(G)\right\} \leq \mathrm{cng}(G) \leq \Delta(G)(\mathrm{tw}(G) + 1).$$

Example 13 (The Petersen graph). One of the most widely used examples in graph theory is the *Petersen graph*, denoted P, see Figure 3(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.

- **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.

- **pw$(P) = 5$**. A minimal-width path decomposition (which we computed using the module 'Vertex separation' of *SageMath* [21]) is the following.

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

- **tw$(P) = 4$**. An optimal tree decomposition (computed using *SageMath* [21]) is shown in Figure 3(ii).

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

---

[18]Only the right-hand side inequality of Theorem 12, $\mathrm{cng}(G) \leq (\mathrm{tw}(G) + 1)\Delta(G)$, is formulated explicitly in [4] as Theorem 1 on p. 111, whereas the left-hand side inequality is stated "inline" in the preceding paragraphs on the same page.
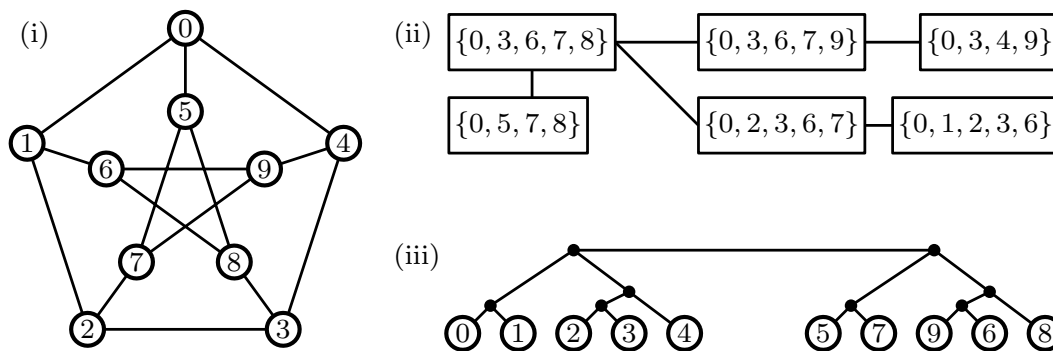
Figure 3: The Petersen graph (i), a tree decomposition realizing minimal treewidth (ii), and an unrooted binary tree realizing minimal congestion (iii)

4   Width-type parameters for 3-manifolds

The Heegaard genus (defined in Section 2.2) is a first example for a width-type 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 more general graph width, whose properties are essential for proving our results in Sections 5 and 6, respectively.

Linear width.   In [57] Scharlemann and Thompson extend the concept of thin position from knot theory [25] to 3-manifolds and define the linear width of a manifold.[19] For this they look at linear splittings, i.e., linear decompositions of a manifold into compression bodies. This setup is explained in Section 2.2.

Given a linear splitting of a 3-manifold $\mathcal{M}$ into $2s$ compression bodies with top boundary surfaces $\mathcal{R}_i$, $1 \leq i \leq s$, consider the multiset $\{c(\mathcal{R}_i) : 1 \leq i \leq s\}$, where $c(\mathcal{S}) = \max\{0, 2g(\mathcal{S}) - 1\}$ for a connected surface $\mathcal{S}$, and $c(\mathcal{S}) = \sum_j c(S_j)$ for a surface $\mathcal{S}$ with connected components S_j. This multiset $\{c(\mathcal{R}_i) : 1 \leq i \leq s\}$, when arranged in non-increasing order, is called the *width* of the (linear) splitting. 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 $\mathcal{M}$.[20] 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- and 3-handles as early as possible and 0- and 1-handles as late as possible in order to obtain a decomposition for which the "topological complexity" of the top bounding surfaces is minimized.

Theorem 14 (Scharlemann–Thompson, Rule 5 from [57])**.** *Let $\mathcal{M}$ be a 3-manifold together with a linear splitting into compression bodies $(\mathcal{N}_1, \mathcal{K}_1, \ldots, \mathcal{N}_s, \mathcal{K}_s)$ in thin position, and let $\{\mathcal{F}_i \subset \mathcal{M} : 1 \leq i \leq s - 1\}$, be the set of lower boundary surfaces $\mathcal{F}_i = \partial_-\mathcal{K}_i = \partial_-\mathcal{N}_{i+1}$.*

---

[19]Also see [33] and the textbook [56] for an introduction to generalized Heegaard splittings and to thin position, and for a survey of recent results.

[20]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 at this point from the definition of linear width in [57].

Then every connected component of every surface $\mathcal{F}_i$ is incompressible.

Theorem 14 has the following consequence.

Theorem 15 (Scharlemann–Thompson [57])**.** *Let $\mathcal{M}$ be irreducible, non-Haken. 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, the linear width of $\mathcal{M}$ is given by $\mathscr{L}(\mathcal{M}) = 2\mathfrak{g}(\mathcal{M}) - 1$.*

Proof (sketch). Let $(\mathcal{N}_1, \mathcal{K}_1, \mathcal{N}_2, \mathcal{K}_2, \ldots, \mathcal{N}_s, \mathcal{K}_s)$ be a thin decomposition of $\mathcal{M}$. By Theorem 14, all surface components of all bounding surfaces $\mathcal{F}_i$, $1 \leq i < s$, must be incompressible. Moreover, all $\mathcal{F}_i$ are separating and thus they are 2-sided. However, irreducible, non-Haken 3-manifolds do not have incompressible 2-sided surfaces. Hence $s = 1$ and therefore the decomposition $(\mathcal{N}_1, \mathcal{K}_1)$ must be a Heegaard splitting of $\mathcal{M}$. $\square$

Graph width. In [56] Scharlemann, Schultens and Saito further refine the concept of thin position to graph splittings of 3-manifolds, see Section 2.2. In particular, given a manifold $\mathcal{M}$ together with a graph splitting defined by a fork complex $\boldsymbol{F}$, let $\{\mathcal{R}_j : j \text{ grip of } \boldsymbol{F}\}$ be the set of top boundary surfaces of the decomposition. Then the width of the graph splitting coming from $\boldsymbol{F}$ is defined as the multiset $\{g(\mathcal{R}_j) : j \text{ grip of } \boldsymbol{F}\}$ 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 16 (Scharlemann–Schultens–Saito, [56] Corollary 5.2.5)**.** *Let $\mathcal{M}$ be a 3-manifold together with a thin graph splitting defined by fork a complex $\boldsymbol{F}$, and let $\{\mathcal{F}_i \subset \mathcal{M} : i \text{ tine of } \boldsymbol{F}\}$ be the set of lower boundary surfaces as defined in Section 2.2. Then every connected component of every lower boundary surface $\mathcal{F}_i$ is incompressible.*

Similarly to the linear width case, Theorem 16 implies that a thin graph splitting of an irreducible, non-Haken 3-manifold must be a Heegaard splitting. In particular, $\mathscr{G}(\mathcal{M}) = \mathfrak{g}(\mathcal{M})$ 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 17 (Agol, Theorem 3.2 in [1])**.** *There exist orientable, closed, irreducible, and non-Haken 3-manifolds of arbitrarily large Heegaard genus.*

Remark. The construction used to prove Theorem 17 starts with non-Haken n-component link complements, 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 [29] and Moriah–Rubinstein [48]. As can be deduced from the construction, the manifolds in question are closed and orientable.

5   An obstruction to bounded cutwidth and pathwidth

In this section we establish an upper bound for the Heegaard genus of a 3-manifold $\mathcal{M}$ in terms of the pathwidth of any triangulation of $\mathcal{M}$ (cf. Theorem 3). As an application of this bound we prove Theorem 1. That is, we show that there exists an infinite family of 3-manifolds not admitting triangulations of uniformly bounded pathwidth.

Theorem 18. *Let $\mathcal{M}$ be a closed, orientable 3-manifold of linear width $\mathscr{L}(\mathcal{M})$. Furthermore, let $\mathcal{T}$ be a triangulation of $\mathcal{M}$ with dual graph $\Gamma(\mathcal{T})$ of cutwidth $\mathrm{cw}(\Gamma(\mathcal{T}))$. Then we have $\mathscr{L}(\mathcal{M}) \leq 6\,\mathrm{cw}(\Gamma(\mathcal{T})) + 7$.*

Proof of Theorem 3 assuming Theorem 18. By Theorem 10, $\mathrm{cw}(\Gamma(\mathcal{T})) \leq 4\,\mathrm{pw}(\Gamma(\mathcal{T}))$ since dual graphs of 3-manifold triangulations are 4-regular. By Theorem 15, $\mathscr{L}(\mathcal{M}) = 2\mathfrak{g}(\mathcal{M}) - 1$ whenever $\mathcal{M}$ is irreducible and non-Haken. Combining these relations with the inequality provided by Theorem 18 yields the result. $\square$

Theorem 1 is now obtained from Theorem 3 and Agol's Theorem 17. It remains to prove Theorem 18. We begin with a basic, yet very useful definition.

Definition 19. Let $\mathcal{T}$ be a triangulation of a 3-manifold $\mathcal{M}$. The *canonical handle decomposition* $\mathrm{chd}(\mathcal{T})$ of $\mathcal{M}$ associated with $\mathcal{T}$ is given by

- 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 $\mathcal{T}$.

Remark. In the above definition of a canonical handle decomposition of a triangulation 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.

The following lemma gives a bound on the complexity of boundary surfaces occurring in the process of building up a manifold $\mathcal{M}$ from the handles of the canonical handle decomposition of a given triangulation of $\mathcal{M}$.

Lemma 20. *Let $\mathcal{T}$ be a triangulation of a 3-manifold $\mathcal{M}$ and let $\Delta_1 < \Delta_2 < \ldots < \Delta_n \in \mathcal{T}$ be a linear ordering of its tetrahedra. Moreover, let $\mathcal{H}_1 \subset \mathcal{H}_2 \subset \ldots \subset \mathcal{H}_n = \mathrm{chd}(\mathcal{T})$ be a filtration of $\mathrm{chd}(\mathcal{T})$ where $\mathcal{H}_j \subset \mathrm{chd}(\mathcal{T})$ is the codimension zero submanifold consisting of all handles of $\mathrm{chd}(\mathcal{T})$ disjoint from tetrahedra Δ_i, $i > j$. Then passing from $\mathcal{H}_j$ to $\mathcal{H}_{j+1}$ corresponds to adding at most 15 handles. Let $\mathcal{K}$ be the codimension zero submanifold constructed from $\mathcal{H}_j$ by adding an arbitrary subset of these handles, then the sum of the genera of $\partial \mathcal{K}$ is no larger than the sum of the genera of the components of $\partial \mathcal{H}_j$ plus four.*

Proof of Lemma 20. 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 $\mathcal{H}_j$ but not disjoint from $\mathcal{H}_{j+1}$. 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 triangles of the tetrahedron), which are the only handles increasing the sum of the genera of $\partial \mathcal{H}_j$. $\qquad\square$

With the help of Lemma 20 we can now prove Theorem 18.

Proof of Theorem 18. Let v_j, $1 \le j \le n$, be the nodes of $\Gamma(\mathcal{T})$ with corresponding tetrahedra $\Delta_j \in \mathcal{T}$, $1 \le j \le n$. 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 $\mathcal{H}_j \subset \mathrm{chd}(\mathcal{T})$, $1 \le j \le n$, be the filtration from Lemma 20. Moreover, let C_j, $1 \le j < n$, be the cutsets of the linear ordering above. 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 $\mathcal{H}(C_j) \subset \mathrm{chd}(\mathcal{T})$ 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 $\partial \mathcal{H}(C_j)$ of $\mathcal{H}(C_j)$ decomposes into two parts, one of which coincides with the boundary surface $\partial \mathcal{H}_j$. Since $\mathcal{H}(C_j)$ is of the form "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 $\mathcal{H}(C_j)$ form a union of handlebodies with sum of genera at most $3k$.

To complete the construction of $\mathcal{H}(C_j)$, the remaining 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 $\partial \mathcal{H}_j \subset \partial \mathcal{H}(C_j)$ is bounded above by $3k$.

Hence, following Lemma 20, the sum g of genera of the components of any bounding surface for any sequence of handle attachments of $\mathrm{chd}(\mathcal{T})$ compatible with the ordering $v_1 < v_2 < \ldots < v_n$ is bounded above by $3\,\mathrm{cw}(\Gamma(\mathcal{T})) + 4$. It follows that $2g - 1 \le 6\,\mathrm{cw}(\Gamma(\mathcal{T})) + 7$, and finally, by the definition, $\mathscr{L}(\mathcal{M}) \le 6\,\mathrm{cw}(\Gamma(\mathcal{T})) + 7$. $\qquad\square$

6 An obstruction to bounded congestion and treewidth

The goal of this section is to prove Theorems 2 and 4, the counterparts of Theorem 1 and 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 21. *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 $\mathcal{M}$ has graph width $\mathscr{G}(\mathcal{M}) \le \frac{9}{2} \mathrm{cng}(\Gamma(\mathcal{T}))$, or $\mathcal{T}$ only contains one tetrahedron.*

Proof of Theorem 4 assuming Theorem 21. First note, that the only closed orientable 3-manifolds which can be triangulated with a single tetrahedron are the 3-sphere of Heegaard
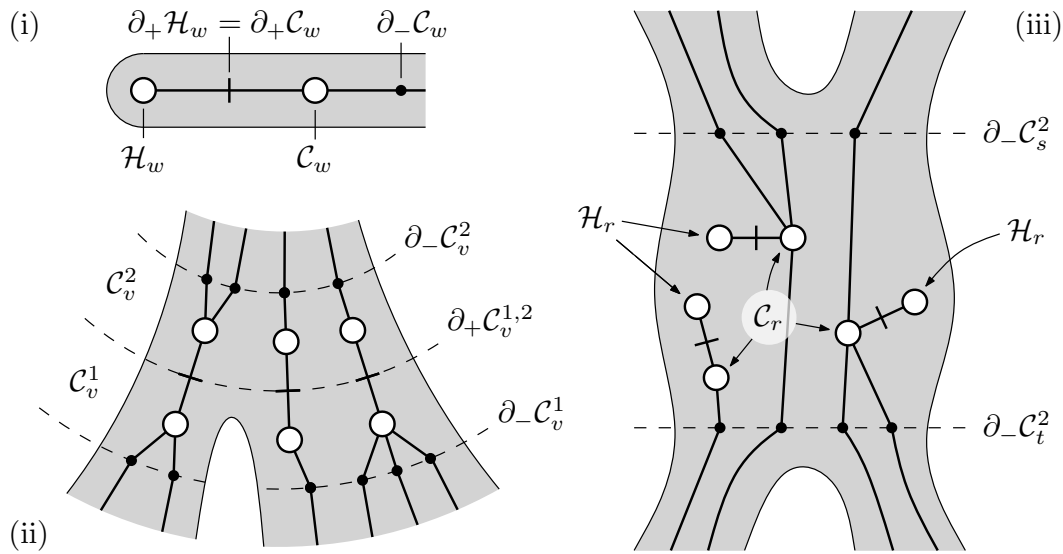
Figure 4: Local pictures of the fork complex $\boldsymbol{F}$ constructed in the proof of Theorem 21

genus zero, and the lens spaces of type $L(4,1)$ and $L(5,2)$ of Heegaard genus one (see Example 6) for which Theorem 4 holds. Otherwise, since dual graphs of 3-manifold triangulations are 4-regular, Theorem 12 yields $\mathrm{cng}(\Gamma(\mathcal{T})) \leq 4(\mathrm{tw}(\Gamma(\mathcal{T})) + 1)$. In addition, Theorem 16 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 21 we obtain Theorem 4. $\square$

Similarly as in Section 5, Theorem 2 immediately follows from Theorems 4 and 17. Hence, the remainder of the section is dedicated to the proof of Theorem 21.

Proof of Theorem 21. Let $\mathcal{M}$ be a closed, connected, orientable 3-manifold, $\mathcal{T}$ be a triangulation of $\mathcal{M}$ whose dual graph $\Gamma(\mathcal{T})$ has congestion $\mathrm{cng}(\Gamma(\mathcal{T})) \leq k$, and let H be an unrooted binary tree realizing $\mathrm{cng}(\Gamma(\mathcal{T})) \leq k$ (cf. Definition 11).

If $k = 0$, $\mathcal{T}$ must consist of a single tetrahedron, and the theorem holds. Thus we can 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 2.2), and then to analyze the genera of the top 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 splitting. Consider the canonical handle decomposition $\mathrm{chd}(\mathcal{T})$ of $\mathcal{M}$ associated with $\mathcal{T}$ as defined in Definition 19. 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 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:

Either, Δ_w forms a "snapped" 3-ball in $\mathcal{T}$, see Figure 5(i), in which case $\mathcal{H}_w$ is built from the 0-handle, the 1-handle, and the 2-handle of 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(ii), and then $\mathcal{H}_w$ consists of the 0-handle and of the 1-handle of chd($\mathcal{T}$), corresponding to Δ_w and to the face gluing, respectively.
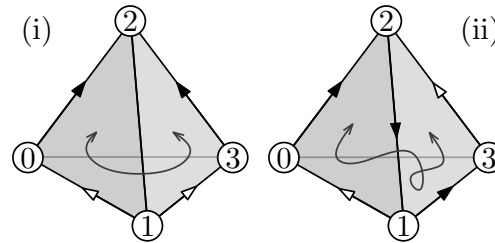
Figure 5: (i) A snapped 3-ball, and (ii) a one-tetrahedron solid torus

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 4(i).

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 chd($\mathcal{T}$) corresponding to a leaf of H below x,
- any 1-handle of chd($\mathcal{T}$) where both adjacent 0-handles are in $\mathcal{T}_x$
- any 2-handle of chd($\mathcal{T}$) where all adjacent 0-handles are in $\mathcal{T}_x$, and
- any 3-handle of 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 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 4(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 4(iii) shows a possible scenario around 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 4). 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 [56, Definition 5.1.4]), yielding a graph splitting of $\mathcal{M}$.

Bounding the width. Following [56, Section 5.1], cf. Section 4, 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

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

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 $g(\partial_+ \mathcal{C}_v^1)$. Note that, since $\text{cng}(\Gamma(\mathcal{T})) \leq k$, 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 $g(\partial_+ \mathcal{C}_v^1) \leq \frac{9k}{2}$. 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 $h \in \text{chd}(\mathcal{T}) \setminus \mathcal{T}_x$ that touches $\partial \mathcal{T}_x$ is adjacent or equal to a 1-handle of $\text{chd}(\mathcal{T})$ 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 $\text{chd}(\mathcal{T})$ that correspond to the leaves of H below x. Turning this around, every handle in $\text{chd}(\mathcal{T}) \setminus \mathcal{T}_x$ is either a 0-handle whose corresponding leaf is *not* below x, or is adjacent to at least one such 0-handle. Now let $h \in \text{chd}(\mathcal{T}) \setminus \mathcal{T}_x$ 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 $h' \in \text{chd}(\mathcal{T})$ 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 $h' \in \text{chd}(\mathcal{T}) \setminus \{h\}$ that are adjacent to at least one of these 0-handles. Thus there exists a 0-handle $h_1 \in \text{chd}(\mathcal{T})$ adjacent to h with corresponding leaf below x. Moreover, some other 0-handle adjacent to h, say h_2, must be in $\text{chd}(\mathcal{T}) \setminus \mathcal{T}_x$, 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 $h_2 \in \text{chd}(\mathcal{T}) \setminus \mathcal{T}_x$, there exists some even $q \in \{0, 2, \ldots, p-2\}$ for which $h^{(q)} \in \mathcal{T}_x$ and $h^{(q+2)} \in \text{chd}(\mathcal{T}) \setminus \mathcal{T}_x$. But then $h^{(q+1)}$ is a 1-handle of $\text{chd}(\mathcal{T})$ 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 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 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 chd($\mathcal{T}$) form a handlebody, the 2- and 3-handles of $\mathcal{X}$ form a union of handlebodies, denoted by $\mathcal{X}_{2,3}$, 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 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 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 $\mathscr{G}(\mathcal{M}) \leq \frac{9k}{2}$. $\square$

A   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 22. *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.[21] Conclude by the observation that graph manifolds modeled on non-isomorphic graphs are non-homeomorphic, cf. [43, Section 3].   □

Proposition 23. *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 [49] 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.   □

These propositions tell us that, as n grows larger, the number of 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 6? Theorem 2 answers this question in the negative in general.
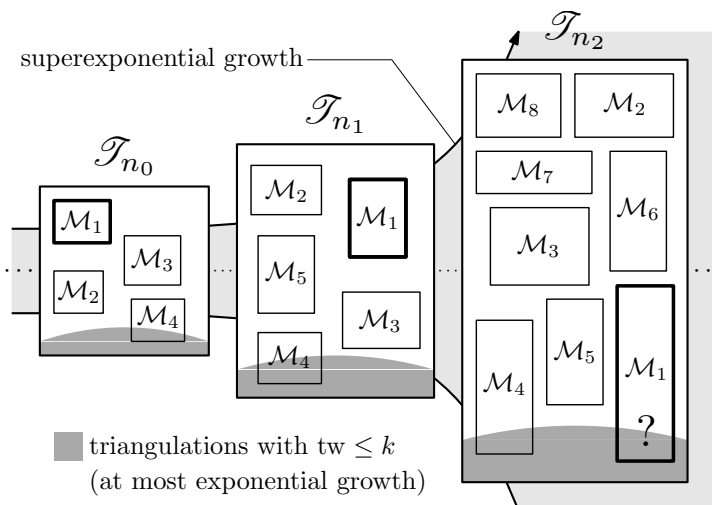
Figure 6

---

[21]See [50] for an overview on graph manifolds and Seifert fiber spaces.

B Complexity and fixed parameter tractability

In Table 1 we collect complexity and fixed parameter tractability properties of the problems 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 ($\checkmark$) 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 [2]	$\checkmark$ [9]	remains **NP**-complete [11]
pw	**NP**-complete [2]	$\checkmark$ [9]	remains **NP**-complete [47][22]
cw	**NP**-complete [26]	$\checkmark$ [64]	polynomial if tw bounded [63]
cng	**NP**-complete [61]	$\checkmark$ [62][23]	

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

We point out that there is a more detailed table in [6], showing the complexity of computing pathwidth and treewidth on several different classes of graphs.

C Computational aspects of different graph 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 8), 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.

---

[22]**NP**-completeness is shown for the vertex separation number which is equivalent to pathwidth [36].

[23]See the discussion in the introduction of [62].

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* [39]. Such a *nice tree decomposition*[24] 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 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 10 and 12), 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, [17]). 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.

---

[24]See the 3-manifold software *Regina* [14] for a visualization of a tree decomposition, and a nice tree decomposition of the dual graph of any given triangulation.

References

[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`.

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

[3] L. Bessières, G. Besson, S. Maillot, M. Boileau, and J. Porti. *Geometrisation of 3-manifolds*, volume 13 of *EMS Tracts Math.* European Mathematical Society (EMS), Zürich, 2010. `doi:10.4171/082`, `MR:2683385`.

[4] 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`.

[5] D. Bienstock. Graph searching, path-width, tree-width and related problems (A survey). In *Reliability Of Computer And Communication Networks*, volume 5 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 33–50. Amer. Math. Soc., Providence, RI, 1991. `doi:10.1090/dimacs/005/02`, `MR:1119138`.

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

[7] 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`.

[8] H. L. Bodlaender. Discovering treewidth. In *SOFSEM 2005: Theory and Practice of Computer Science, 31st Conference on Current Trends in Theory and Practice of Computer Science, Liptovský Ján, Slovakia, January 22–28, 2005, Proceedings*, pages 1–16, 2005. `doi:10.1007/978-3-540-30577-4_1`.

[9] H. L. Bodlaender. Fixed-parameter tractability of treewidth and pathwidth. In *The Multivariate Algorithmic Revolution and Beyond*, volume 7370 of *Lecture Notes in Comput. Sci.*, pages 196–227, 2012. `doi:10.1007/978-3-642-30891-8_12`.

[10] 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`.

[11] 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`.

[12] B. Burton, H. Edelsbrunner, J. Erickson, and S. Tillmann. Computational geometric and algebraic topology. *Oberwolfach Rep.*, 12(4):2637–2699, 2015. `doi:10.4171/OWR/2015/45`, `MR:3586035`.

[13] B. A. Burton. Detecting genus in vertex links for the fast enumeration of 3-manifold triangulations. In *Symbolic and Algebraic Computation, International Symposium, ISSAC 2011 (co-located with FCRC 2011), San Jose, CA, USA, June 7–11, 2011, Proceedings*, pages 59–66, 2011. `doi:10.1145/1993886.1993901`, `MR:2895195`.

[14] B. A. Burton, R. Budney, W. Pettersson, et al. Regina: Software for 3-manifold topology and normal surface theory, 1999–2017. URL: `http://regina-normal.github.io`.

[15] 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`.

[16] 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`.

[17] B. A. Burton, C. Maria, and J. Spreer. Algorithms and complexity for Turaev–Viro invariants. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6–10, 2015, Proceedings, Part I*, pages 281–293, 2015. `doi:10.1007/978-3-662-47672-7_23`, `MR:3382513`.

[18] B. A. Burton and J. Spreer. The complexity of detecting taut angle structures on triangulations. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6–8, 2013*, pages 168–183, 2013. `doi:10.1137/1.9781611973105.13`, `MR:3185388`.

[19] 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`.

[20] 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`.

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

[22] R. Diestel. *Graph theory*, volume 173 of *Grad. Texts in Math.* Springer, Berlin, fifth edition, 2017. `doi:10.1007/978-3-662-53622-3`, `MR:3644391`.

[23] R. G. Downey and M. R. Fellows. *Parameterized complexity*. Monogr. Comput. Sci. Springer, New York, 1999. `doi:10.1007/978-1-4612-0515-9`, `MR:1656112`.

[24] 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`.

[25] D. Gabai. Foliations and the topology of 3-manifolds. III. *J. Differential Geom.*, 26(3):479–536, 1987. `doi:10.4310/jdg/1214441488`, `MR:910018`.

[26] F. Gavril. Some NP-complete problems on graphs. In *Conference on Information Sciences and Systems, Proceedings*, pages 91–95. Johns Hopkins University, 1977.

[27] W. Haken. Theorie der Normalflächen. *Acta Math.*, 105:245–375, 1961. `doi:10.1007/BF02559591`, `MR:0141106`.

[28] 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`.

[29] 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`.

[30] 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`.

[31] 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`.

[32] D. Hoffoss and J. Maher. Morse functions to graphs and topological complexity for hyperbolic 3-manifolds, 2017. `arXiv:1708.04140`.

[33] 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`.

[34] K. Huszár, J. Spreer, and U. Wagner. On the treewidth of triangulated 3-manifolds. In *34th International Symposium on Computational Geometry (SoCG), June 11–14, 2018, Budapest, Hungary*, pages 46:1–46:15, 2018. `doi:10.4230/LIPIcs.SoCG.2018.46`.

[35] 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`.

[36] 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`.

[37] 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`.

[38] 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`.

[39] T. Kloks. *Treewidth. Computations and Approximations*, volume 842 of *Lecture Notes in Comput. Sci.* Springer, 1994. `doi:10.1007/BFb0045375`, `MR:1312164`.

[40] G. Kuperberg. Knottedness is in NP, modulo GRH. *Adv. Math.*, 256:493–506, 2014. `doi:10.1016/j.aim.2014.01.007`, `MR:3177300`.

[41] G. Kuperberg. Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization, 2015. `arXiv:1508.06720`.

[42] M. Lackenby. The efficient certification of knottedness and Thurston norm, 2016. `arXiv:1604.00290`.

[43] 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`.

[44] C. Maria and J. Spreer. A polynomial time algorithm to compute quantum invariants of 3-manifolds with bounded first Betti number. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16–19*, pages 2721–2732, 2017. `doi:10.1137/1.9781611974782.180`, `MR:3627909`.

[45] S. Matveev. *Algorithmic topology and classification of 3-manifolds*, volume 9 of *Algorithms Comput. Math.* Springer, Berlin, second edition, 2007. `doi:10.1007/978-3-540-45899-9`, `MR:2341532`.

[46] 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`.

[47] 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`.

[48] 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`.

[49] 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`.

[50] P. Orlik. *Seifert manifolds*, volume 291 of *Lecture Notes in Math.* Springer-Verlag, Berlin–New York, 1972. `doi:10.1007/BFb0060329`, `MR:0426001`.

[51] M. I. Ostrovskii. Minimal congestion trees. *Discrete Math.*, 285(1–3):219–226, 2004. `doi:10.1016/j.disc.2004.02.009`, `MR:2062845`.

[52] 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`.

[53] 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`.

[54] J. H. Rubinstein. An algorithm to recognize the 3-sphere. In *Proceedings of the International Congress of Mathematicians, Zürich, Switzerland, August 3–11, 1994*, volume 1, pages 601–611. Birkhäuser, 1995. `doi:10.1007/978-3-0348-9078-6_54`, `MR:1403961`.

[55] 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`.

[56] M. Scharlemann, J. Schultens, and T. Saito. *Lecture notes on generalized Heegaard splittings*. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2016. Three lectures on low-dimensional topology in Kyoto. `doi:10.1142/10019`, `MR:3585907`.

[57] 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`.

[58] 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`.

[59] 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`.

[60] 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`.

[61] P. D. Seymour and R. Thomas. Call routing and the ratcatcher. *Combinatorica*, 14(2):217–241, 1994. `doi:10.1007/BF01215352`, `MR:1289074`.

[62] D. M. Thilikos, M. J. Serna, and H. L. Bodlaender. Constructive linear time algorithms for small cutwidth and carving-width. In *Algorithms and Computation, 11th International Conference, ISAAC 2000, Taipei, Taiwan, December 18–20, 2000, Proceedings*, pages 192–203, 2000. `doi:10.1007/3-540-40996-3_17`, `MR:1858364`.

[63] 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 *Algorithms - ESA 2001, 9th Annual European Symposium, Aarhus, Denmark, August 28–31, 2001, Proceedings*, pages 380–390, 2001. `doi:10.1007/3-540-44676-1_32`, `MR:1914050`.

[64] 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`.

[65] 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`.

[66] W. P. Thurston. *Three-dimensional geometry and topology. Vol. 1*, volume 35 of *Princeton Math. Ser.* Princeton University Press, Princeton, NJ, 1997. Edited by S. Levy. doi:10.1515/9781400865321, MR:1435975.

[67] R. Zentner. Integer homology 3-spheres admit irreducible representations in $SL(2,\mathbb{C})$. *Duke Math. J.*, 167(9):1643–1712, 2018. doi:10.1215/00127094-2018-0004, MR:3813594.