On the Treewidth of Triangulated 3-Manifolds

Kristóf Huszár
Institute of Science and Technology Austria (IST Austria)
Am Campus 1, 3400 Klosterneuburg, Austria
kristof.huszar@ist.ac.at
https://orcid.org/0000-0002-5445-5057

Jonathan Spreer[1]
Institut für Mathematik, Freie Universität Berlin
Arnimallee 2, 14195 Berlin, Germany
jonathan.spreer@fu-berlin.de
https://orcid.org/0000-0001-6865-9483

Uli Wagner
Institute of Science and Technology Austria (IST Austria)
Am Campus 1, 3400 Klosterneuburg, Austria
uli@ist.ac.at
https://orcid.org/0000-0002-1494-0568

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 and of Scharlemann and Thompson, 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 $48(k+1)$ (resp. $4(3k+1)$).

2012 ACM Subject Classification Mathematics of computing → Geometric topology, Theory of computation → Fixed parameter tractability

Keywords and phrases computational topology, triangulations of 3-manifolds, thin position, fixed-parameter tractability, congestion, treewidth

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.46

Related Version A full version of the paper is available at [26], `http://arxiv.org/abs/1712.00434`.

[1] Research of the second author was supported by the Einstein Foundation (project "Einstein Visiting Fellow Santos") and by the Simons Foundation ("Simons Visiting Professors" program).

34th International Symposium on Computational Geometry (SoCG 2018).
Editors: Bettina Speckmann and Csaba D. Tóth; Article No. 46; pp. 46:1–46:15
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Acknowledgements We would like to thank Arnaud de Mesmay and Stephan Tillmann for helpful discussions and remarks, and the anonymous reviewers for useful comments and suggestions regarding the exposition. This work was initiated during a visit of the second author at IST Austria. The second author would like to thank the people at IST Austria for their hospitality.

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 [21], deciding whether a given 3-manifold is homeomorphic to the 3-sphere [40, 47], and, more generally (based on Perelman's proof of Thurston's geometrization conjecture [29]), deciding whether two given 3-manifolds are homeomorphic, see, e.g., [2, 31, 45]. The algorithm for solving the homeomorphism problem is still purely theoretical, and its complexity remains largely unknown [31, 33]. In contrast, the first two problems are known to lie in the intersection of the complexity classes **NP** and **co-NP** [22, 27, 30, 32, 43, 49].[2]

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* [9]) 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 [19]. This effort has proven to be highly effective and, today, there exist numerous FPT algorithms in the field [10, 12, 13, 14, 34]. More specifically, given a triangulation $\mathcal{T}$ of a 3-manifold $\mathcal{M}$ with n tetrahedra whose dual graph $\Gamma(\mathcal{T})$ has treewidth[3] at most k, there exist algorithms to compute

- taut angle structures[4] of what is called ideal triangulations with torus boundary components in running time $O(7^k \cdot n)$ [14];
- optimal Morse matchings[5] in the Hasse diagram of $\mathcal{T}$ in $O(4^{k^2+k} \cdot k^3 \cdot \log k \cdot n)$ [12];
- the Turaev–Viro invariants[6] for parameter $r \geq 3$ in $O((r-1)^{6(k+1)} \cdot k^2 \cdot \log r \cdot n)$ [13];
- 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 [10].[7]

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 [13, 34].

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

[2] The proof of **co-NP** membership for 3-sphere recognition assumes the Generalized Riemann Hypothesis.
[3] We often simply speak of the treewidth of a triangulation, meaning the treewidth of its dual graph.
[4] Taut angle structures are combinatorial versions of semi-simplicial metrics which have implications on the geometric properties of the underlying manifold.
[5] 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.
[6] 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.
[7] This result is analogous to Courcelle's celebrated theorem in graph theory [16].

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 [28], the Turaev–Viro invariant for parameter $r = 4$ is $\#\mathbf{P}$-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 [13], would constitute a polynomial time solution for a $\#\mathbf{P}$-hard problem. Furthermore, known facts imply that *most* triangulations of *most* 3-manifolds must have large treewidth.[8] However, while these arguments indicate that triangulations of small treewidth may be rare and computationally hard to find, it does 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.[9] 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 $\mathfrak{g}(\mathcal{M}) < 48(k+1)$.*

By a result of Agol [1] (Theorem 12 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.[10] Nonetheless, we provide separate proofs for each of the two statements. The motivation for this is that all the ingredients to prove Theorem 3 (and hence Theorem 1) already appear in the literature, while Theorem 4 needs extra work to be done.

---

[8] 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. See the full version of the article for more details.

[9] 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 [11, Problem 8] (formulated in the context of knot theory).

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

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. 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.

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 [17].

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 $k \in \mathbb{N}$.

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 [44] and to the seminal monograph [48].

A 3-*manifold with boundary* is a topological space[11] $\mathcal{M}$ 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 3-manifold is *bounded* (resp. *closed*) if it has a non-empty (resp. empty) boundary.

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, called *topological invariants*, such as "number of boundary components", or "connectedness".

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

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.[12] For instance, to build a closed 3-manifold from handles, we may start with

[11] More precisely, we only consider topological spaces which are second countable and Hausdorff.
[12] See Chapter 6 and Appendix B of [44] for a primer on handle decompositions and Morse functions.

a collection of disjoint 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 collection 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 from handles, the existing decomposition is a bounded 3-manifold and its boundary – called a *bounding surface* – separates the 3-manifold 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. Of course, in this sense, some embedded surfaces can be considered topologically more meaningful than others. In particular, we have the following notions.

Let $\mathcal{M}$ be a 3-manifold. A surface $\mathcal{S} \subset \mathcal{M}$ is said to be *properly embedded*, if it is embedded in $\mathcal{M}$, and we have for the boundary $\partial\mathcal{S} = \mathcal{S} \cap \partial\mathcal{M}$. 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 $\mathcal{M}$.[13]

A 3-manifold $\mathcal{M}$ is called *irreducible*, if every embedded 2-sphere bounds a 3-ball in $\mathcal{M}$.[14] Moreover, it is called P^2*-irreducible*, if it does not contain an embedded 2-sided[15] 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. 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*. In this article, we only apply the notion of a non-Haken manifold to manifolds that are (P^2-)irreducible. In the literature, such manifolds are also referred to as *small*.

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

A *handlebody* is a bounded 3-manifold which can be described as a single 0-handle with a number of 1-handles attached to it, or, equivalently, as a thickened version of a graph.

Let $\mathcal{S}$ be a closed (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 $\partial_+\mathcal{N} = \partial\mathcal{N} \setminus \partial_-\mathcal{N}$.[16]

Dual to this construction, we can also start with $\mathcal{S} \times [0,1]$, add 2-handles along $\mathcal{S} \times \{1\}$ and fill in some resulting 2-sphere boundary components with 3-balls. Again, $\mathcal{N}$ has two sets of boundary components: $\partial_-\mathcal{N} = \mathcal{S} \times \{0\}$ and $\partial_+\mathcal{N} = \partial\mathcal{N} \setminus \partial_-\mathcal{N}$.

[13] 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 $\mathbb{S}^2 \times \{x\} \subset \mathbb{S}^2 \times \mathbb{S}^1$.

[14] In particular, $\mathbb{S}^2 \times \mathbb{S}^1$ is not irreducible.

[15] 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.

[16] If $\mathcal{S}$ is a 2-sphere, and $\mathcal{S} \times \{0\}$ is filled in, i.e., $\partial_-\mathcal{N} = \emptyset$, then $\mathcal{N}$ is actually a handlebody.

Handlebodies very naturally occur when building up a (connected) 3-manifold $\mathcal{M}$ from an arbitrary but fixed set of handles. If (possibly after deforming the attaching maps) all 0- and 1-handles are attached before attaching any 2- or 3-handles, we obtain a decomposition of $\mathcal{M}$ into two handlebodies: the union of all 0- and 1-handles on one side, and its complement on the other side. Such a decomposition exists for every 3-manifold $\mathcal{M}$ [35] and is called a *Heegaard splitting of* $\mathcal{M}$. The genus of their (common) boundary surface is called the *genus of the decomposition*. It is equal to #1-handles minus #0-handles plus one. The smallest genus of a handle decomposition of $\mathcal{M}$, which is a topological invariant by definition, is called the *Heegaard genus of* $\mathcal{M}$ and is denoted by $\mathfrak{g}(\mathcal{M})$.

Compression bodies are central to Scharlemann and Thompson's definition of thin position [42], discussed in Section 4, as they can be used to describe 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.

Given a possibly bounded 3-manifold $\mathcal{M}$, a fixed set of handles, and an arbitrary sequence of handle attachments to build up $\mathcal{M}$, we look at the first terms of the sequence up to (but not including) the first 2- or 3-handle attachment. Let $\mathcal{S}$ be the surface given by the boundaries of all 0-handles in this subsequence plus potentially some of the boundary components of $\mathcal{M}$ (we want the boundary components of $\mathcal{M}$ to appear at the beginning or the end of this construction). We thicken $\mathcal{S}$ into a bounded 3-manifold $\mathcal{S} \times [0, 1]$, fill back in the 0-handles at the 2-sphere components of $\mathcal{S} \times \{0\}$ and attach the 1-handles from the subsequence to $\mathcal{S} \times \{1\}$. This is a compression body, say $\mathcal{N}_1$ (if no boundary components are added, $\mathcal{N}_1$ is merely the union of all 0- and 1-handles in the subsequence, and hence a union of handlebodies). In the second step we 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 $\mathcal{K}_1$. Iterating this procedure, $\mathcal{M}$ can be decomposed into a sequence of compression bodies $(\mathcal{N}_1, \mathcal{K}_1, \mathcal{N}_2, \mathcal{K}_2, \dots, \mathcal{N}_s, \mathcal{K}_s)$ satisfying $\mathcal{R}_i := \partial_+\mathcal{N}_i = \partial_-\mathcal{K}_i$, $1 \le i \le s$, and $\mathcal{F}_i := \partial_+\mathcal{K}_i = \partial_-\mathcal{N}_{i+1}$, $1 \le i < s$. By construction, the surfaces $\mathcal{R}_i$ have more handles and/or less connected components than the surfaces $\mathcal{F}_i$.

When talking about bounding surfaces in sequences of handle attachments, surfaces of type $\mathcal{R}_i$ and $\mathcal{F}_i$ are typically the most interesting ones as they form local extrema with respect to their topological complexity.[17] Thus, we refer to the $\mathcal{R}_i$ as the *large* and to the $\mathcal{F}_i$ as the *small* bounding surfaces.

Triangulations. In this article we typically describe 3-manifolds by *semi-simplicial triangulations* (also known as *singular triangulations*, here often just referred to as *triangulations*). That is, a collection of n 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[18] if no tetrahedral edge is identified with itself in reverse, and the boundary of a small neighborhood around each vertex is a 2-sphere.

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 $\mathcal{T}$.

[17] We specify precisely what we mean by the *topological complexity of a surface* whenever this is necessary.

[18] It is straightforward to extend the definition of a triangulation to include bounded 3-manifolds. However, for the purpose of this article it suffices to consider the closed case.

2.3 Parameterized complexity and fixed parameter tractability

There exist various notions and concepts of a refined complexity analysis for theoretically difficult problems. One of them, parameterized complexity, due to Downey and Fellows [18, 19], 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 $\mathcal{A}$, a parameter is a (computable) function $p\colon \mathcal{A} \to \mathbb{N}$. 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 $A \in \mathcal{A}$ with running time $O(f(p(A)) \cdot \mathrm{poly}(n))$, where n is the size of the input A, 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 [34] 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 possible 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 [24]. 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.

Treewidth and pathwidth. The concepts of treewidth and pathwidth were introduced by Robertson and Seymour in their early papers on graph minors [38, 39], also see the surveys [5, 7, 8]. Given a graph G, its *treewidth* $\mathrm{tw}(G)$ measures how tree-like the graph is.

▶ **Definition 5** (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* $\mathrm{tw}(G)$ is the minimum width among all possible tree decompositions of G.

▶ **Definition 6** (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* $\mathrm{pw}(G)$ 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 $\{v_1, \dots, v_n\} = V(G)$ of G on a line, the set of arcs running from a node v_i, $i \le \ell$, 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 7** (Bodlaender, Theorems 47 and 49 from [6][19]). *Given a graph G, we have*

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

Congestion. Bienstock introduced *congestion* [4], a generalization of cutwidth, which is a quantity related to treewidth in a similar way as cutwidth to pathwidth (compare Theorems 7 and 9).

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.

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., [37]). Here we work with the following.

▶ **Definition 8** (Congestion[20]). Let $T_{\{1,3\}}$ be the set of unrooted binary trees.[21] 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.

▶ **Theorem 9** (Bienstock, p. 108–111 of [3]). *Given a graph G, we have*[22]

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

[19] The inequality $\mathrm{cw}(G) \le \Delta(G)\,\mathrm{pw}(G)$ seems to be already present in the earlier work of Chung and Seymour [15] on the relation of cutwidth to another parameter called topological bandwith (see Theorem 2 in [15]). Pathwidth plays an intermediate, connecting role there. However, the inequality is phrased and proved explicitly by Bodlaender in [6].

[20] It is important to note that congestion in the sense of Definition 8 is also known as *carving width*, a term which was coined by Robertson and Seymour in [46]. 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.

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

[22] Only the right-hand side inequality of Theorem 9, $\mathrm{cng}(G) \le (\mathrm{tw}(G)+1)\Delta(G)$, is formulated explicitly in [3] as Theorem 1 on p. 111, whereas the left-hand side inequality is stated "inline" in the preceding paragraphs on the same page.

See the full version of the article for a comparison of cutwidth, pathwidth, treewidth and congestion of the Petersen graph. Also, see the full version for an explanation of how width-type parameters are used in FPT-algorithms in computational topology, and a comparison of different parameters for their potential computational advantages or disadvantages.

4 Thin position, and non-Haken 3-manifolds of large genus

In [42] Scharlemann and Thompson extend the concept of thin-position from knot theory [20] to 3-manifolds and define the linear width of a manifold.[23] For this they look at decompositions of a 3-manifold $\mathcal{M}$ into pairs of compression bodies, separated by so-called *large* boundary surfaces. This setup is explained in detail at the end of Section 2.2.

Given such a decomposition of a 3-manifold $\mathcal{M}$ into s pairs of compression bodies with large bounding 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 $\mathcal{S}$,[24] 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 a decreasing order, is called the *linear width* of the decomposition. The *linear width* of a manifold $\mathcal{M}$, denoted by $\mathcal{L}(\mathcal{M})$, is defined to be the lexicographically smallest width of all such decompositions of $\mathcal{M}$. A manifold $\mathcal{M}$ together with a decomposition into compression bodies realizing $\mathcal{L}(\mathcal{M})$ 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 large bounding surfaces is minimized.

▶ **Theorem 10** (Scharlemann–Thompson [42]). *Let $\mathcal{M}$ be a 3-manifold together with a decomposition into compression bodies $(\mathcal{N}_1, \mathcal{K}_1, \mathcal{N}_2, \mathcal{K}_2, \ldots, \mathcal{N}_s, \mathcal{K}_s)$ in thin position, and let $\mathcal{F}_i \subset \mathcal{M}$, $1 \leq i < s$, be the set of small bounding surfaces as defined in Section 2.2. Then every connected component of every surface $\mathcal{F}_i$ is incompressible.*

Theorem 10 has the following consequence (see the full version of the article for a sketch of the proof).

▶ **Theorem 11** (Scharlemann–Thompson [42]). *Let $\mathcal{M}$ be irreducible, non-Haken. Then the smallest width decomposition of $\mathcal{M}$ into compression bodies is a Heegaard decomposition of minimal genus $\mathfrak{g}(\mathcal{M})$. In particular, the linear width of $\mathcal{M}$ is given by a list containing only one element, namely $\mathcal{L}(\mathcal{M}) = (2\mathfrak{g}(\mathcal{M}) - 1)$.*

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

▶ **Theorem 12** (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 12 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 [23] and Moriah–Rubinstein [36]. As can be deduced from the construction, the manifolds in question are closed and orientable.

[23] Also see [25] and the textbook [41] for an introduction to generalized Heegaard splittings and to thin position, and for a survey of recent results.
[24] $g(\mathcal{S}) = 1 - \chi(\mathcal{S})/2$ denotes the (orientable) genus, and $\chi(\mathcal{S})$ is the Euler characteristic of $\mathcal{S}$.

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 13.** *Let $\mathcal{M}$ be a 3-manifold of linear width $\mathcal{L}(\mathcal{M})$ with dominant entry $\mathcal{L}(\mathcal{M})_1$. Furthermore, let $\mathcal{T}$ be a triangulation of $\mathcal{M}$ with dual graph $\Gamma(\mathcal{T})$ of cutwidth $\mathrm{cw}(\Gamma(\mathcal{T}))$. Then $\mathcal{L}(\mathcal{M})_1 \le 6\,\mathrm{cw}(\Gamma(\mathcal{T})) + 7$.*

Proof of Theorem 3 assuming Theorem 13. By Theorem 7, $\mathrm{cw}(\Gamma(\mathcal{T})) \le 4\,\mathrm{pw}(\Gamma(\mathcal{T}))$ since dual graphs of 3-manifold triangulations are 4-regular. By Theorem 11, $\mathcal{L}(\mathcal{M})_1 = 2\mathfrak{g}(\mathcal{M}) - 1$ whenever $\mathcal{M}$ is irreducible and non-Haken. Combining these relations with the inequality provided by Theorem 13 yields the result. ◀

Theorem 1 is now obtained from Theorem 3 and Agol's Theorem 12. It remains to prove Theorem 13. We begin with a basic, yet very useful definition.

▶ **Definition 14.** Let $\mathcal{T}$ be a triangulations of a 3-manifold $\mathcal{M}$. The *canonical handle decomposition* $\mathrm{chd}(\mathcal{T})$ of $\mathcal{T}$ 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 $\mathcal{T}$.

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 $\mathcal{M}$ (see the full version of the article for a proof).

▶ **Lemma 15.** *Let $\mathcal{T}$ be a (semi-simplicial) 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, with the maximum of the sum of the genera of the components of any of the bounding surfaces occurring in the process being no larger than the sum of the genera of the components of $\partial\mathcal{H}_j$ plus four.*

With the help of Lemma 15 we can now prove Theorem 13.

Proof of Theorem 13. 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 15. 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 collection 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 collection 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 15, 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 $2g-1 \leq 6\,\mathrm{cw}(\Gamma(\mathcal{T}))+7$, and finally, by the definition, $\mathcal{L}(\mathcal{M})_1 \leq 6\,\mathrm{cw}(\Gamma(\mathcal{T}))+7$. ◀

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 Heegaard genus.

▶ **Theorem 16.** *Let $\mathcal{M}$ be an orientable, closed, irreducible, non-Haken 3-manifold which has a triangulation $\mathcal{T}$ with dual graph $\Gamma(\mathcal{T})$ of congestion $\mathrm{cng}(\Gamma(\mathcal{T})) \leq k$. Then either $\mathcal{M}$ has Heegaard genus $\mathfrak{g}(\mathcal{M}) < 12k$, or $\mathcal{T}$ only contains one tetrahedron.*

Proof of Theorem 4 assuming Theorem 16. Since dual graphs of 3-manifold triangulations are 4-regular, Theorem 9 implies $\mathrm{cng}(\Gamma(\mathcal{T})) \leq 4(\mathrm{tw}(\Gamma(\mathcal{T}))+1)$. Moreover, the only closed orientable 3-manifolds which can be triangulated with a single tetrahedron are the 3-sphere of Heegaard genus zero, and the lens spaces of type $\mathrm{L}(4,1)$ and $\mathrm{L}(5,2)$ of Heegaard genus one. ◀

Similarly as before, Theorem 2 immediately follows by combining Theorems 4 and 12. Hence, the remainder of the section is dedicated to the proof of Theorem 16.

Proof of Theorem 16. Let $\mathcal{M}$ be an orientable, closed, irreducible, non-Haken 3-manifold which has a triangulation $\mathcal{T}$ whose dual graph $\Gamma(\mathcal{T})$ has congestion $\mathrm{cng}(\Gamma(\mathcal{T})) \leq k$, and let T be an unrooted binary tree realizing $\mathrm{cng}(\Gamma(\mathcal{T})) \leq k$. If $k = 0$, $\mathcal{T}$ must consist of a single tetrahedron, and the theorem holds. Thus we can assume that $k \geq 1$. Moreover, let $\mathrm{chd}(\mathcal{T})$ be the canonical handle decomposition of $\mathcal{T}$ as defined in Definition 14.

For every $e \in E(T)$, there exist arcs $\gamma_1, \gamma_2, \ldots, \gamma_\ell \in E(\Gamma(\mathcal{T}))$, $\ell \leq k$, running parallel to e, corresponding to triangles $t_1, t_2, \ldots, t_\ell$ in $\mathcal{T}$. Let $\mathcal{H}_e \subset \mathrm{chd}(\mathcal{T})$ be the submanifold built from the ℓ 1-handles of $\mathrm{chd}(\mathcal{T})$ corresponding to the triangles t_i, $1 \leq i \leq \ell$, and the $\leq 3\ell$ 2- and 3-handles each of $\mathrm{chd}(\mathcal{T})$ corresponding to the edges and vertices of these triangles.

▶ **Lemma 17.** *The codimension zero submanifold $\mathrm{chd}(\mathcal{T}) \setminus \mathcal{H}_e$ decomposes into a pair of disjoint components, denoted by $\mathcal{H}_e^+$ and $\mathcal{H}_e^-$.*

See Figure 1(i), and, for the proof, the full version of the article.

▶ **Lemma 18.** *At least one of $\mathcal{H}_e^+$ or $\mathrm{chd}(\mathcal{T}) \setminus \mathcal{H}_e^+$ can be built from at most $5k$ handles. The same statement holds for $\mathcal{H}_e^-$ and $\mathrm{chd}(\mathcal{T}) \setminus \mathcal{H}_e^-$.*

Idea of the proof of Lemma 18. Consider the decomposition $\mathcal{M} = \mathcal{H}_e^+ \cup_{\partial\mathcal{H}_e^+} (\mathrm{chd}(\mathcal{T}) \setminus \mathcal{H}_e^+) =: \mathcal{A}_0 \cup_{\mathcal{S}_0} \mathcal{B}_0$. Both the number of connected components and the genus of the surface $\mathcal{S}_0 = \partial\mathcal{H}_e^+$ is in $O(k)$, as $\mathcal{H}_e$ can be built from $O(k)$ handles (see above), and $\mathcal{S}_0 = \partial\mathcal{H}_e^+ \subset \partial\mathcal{H}_e$. Since $\mathcal{M}$ is irreducible, non-Haken, and all surface components of $\mathcal{S}_0$ are 2-sided by construction, none of them can be incompressible. Therefore, via a sequence of handle reductions along carefully chosen compressing disks we get a decomposition

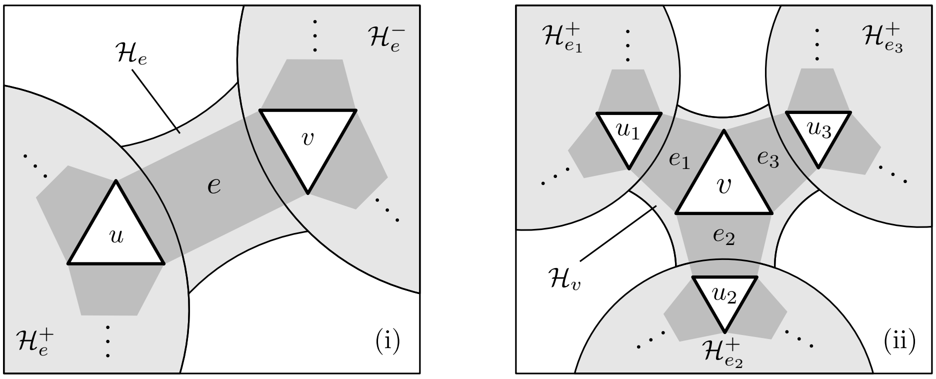

Figure 1 Schematic overview of the decompositions of $\mathcal{M}$ into (i) three and (ii) four submanifolds at arcs and degree-three nodes of the tree T, respectively

$\mathcal{M} = \mathcal{A}_* \cup_{\mathcal{S}_*} \mathcal{B}_*$ in $O(k)$ steps, where $\mathcal{S}_*$ consists of $m \in O(k)$ copies of the 2-sphere. It turns out that either $\mathcal{A}_*$ or $\mathcal{B}_*$, say $\mathcal{B}_*$, is a collection of punctured 3-balls with an overall number of punctures being equal to m, which implies that $\mathcal{B}_*$ can be built from m handles, from which $\mathcal{B}_0$ can be recovered via $O(k)$ handle attachments. We can make these bounds explicit by careful bookkeeping to obtain the result. See the full version of the article for a more detailed proof. ◀

Now at every degree three node $v \in V(T)$ we define a decomposition of $\mathcal{M}$ into four submanifolds (Figure 1(ii)). Three of them, $\mathcal{H}_{e_1}^+$, $\mathcal{H}_{e_2}^+$, and $\mathcal{H}_{e_3}^+$, are arising from the three connected components of T after removing the three arcs incident to v (cf. Lemma 17), and the fourth submanifold, $\mathcal{H}_v$, contains all remaining handles of $\mathrm{chd}(\mathcal{T})$.

By Lemma 18, at least two of the $\mathcal{H}_{e_i}^+$, $1 \leq i \leq 3$, can be built from at most $5k$ handles. For the fourth submanifold we have $\mathcal{H}_v = \mathcal{H}_{e_1} \cup \mathcal{H}_{e_2} \cup \mathcal{H}_{e_3}$. Observe that at most k arcs of $\Gamma(\mathcal{T})$ run parallel to each of the $e_i \in E(T)$. Counting those arcs along the e_i we encounter each of them twice, therefore at most $\frac{3}{2}k$ arcs of $\Gamma(\mathcal{T})$ meet v. It follows that $\mathcal{H}_v$ is a collection of at most $\frac{3}{2}k$ handlebodies with sum of genera at most $\frac{9}{2}k$ and an additional at most $\frac{3}{2}k$ 2-handles attached (see full version for details). Altogether, $\mathcal{H}_v$ can be built from at most $\frac{3}{2}k + \frac{9}{2}k + \frac{3}{2}k < 8k$ handles. Moreover, if v has a neighbor which is a leaf of T, then its corresponding submanifold is obtained from a 1-tetrahedron submanifold of $\mathcal{T}$ with at most three interior faces (the tetrahedron itself, at most one triangle, and at most one edge). Hence, this part corresponds to at most three handles of $\mathrm{chd}(\mathcal{T})$.

If all $\mathcal{H}_{e_i}^+$, $1 \leq i \leq 3$, are of size at most $5k$, $\mathcal{M}$ can be built using less than $(3 \cdot 5 + 8)k = 23k$ handles, thus $\mathfrak{g}(\mathcal{M}) < 12k$ and we are done. Hence, assume that exactly one of them, say $\mathcal{H}_{e_1}^+$, requires a larger number of handles to be built. Let u_1 be the other node of e_1.

Now we choose u_1 to be the new center (instead of v) and repeat the above process. Moving from v to u_1 merges $\mathcal{H}_{e_2}^+$, $\mathcal{H}_{e_3}^+$, and $\mathcal{H}_v$ (note that each of $\mathcal{H}_{e_2}^+$ and $\mathcal{H}_{e_3}^+$ can be built with $\leq 5k$ handles, and $\mathcal{H}_v$ can be built with $< 8k$ handles), and splits the remaining larger part into three submanifolds. Since these three submanifolds together form the larger part, they either cannot all be built from at most $5k$ / less than $8k$ handles – or $\mathcal{M}$ can be built from at most $23k$ handles. Similarly, the three merged parts must be one of the two new submanifolds which can be built from at most $5k$ handles – or $\mathcal{M}$ can be built from at most $23k$ handles. In both cases it follows that $\mathfrak{g}(\mathcal{M}) < 12k$ and we are done. Hence, assume one of the submanifolds coming from one of the two new subtrees must require a larger number of handles to be built. However, at the same time, the connected component of

$T \setminus \mathrm{st}_T(u_1)$ corresponding to this submanifold of unbounded size has fewer nodes than that corresponding to the previously largest part. Iterating the process must thus eventually lead to the conclusion that $\mathcal{M}$ is of Heegaard genus less than $12k$. ◀

References

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

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

3 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`.

4 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`.

5 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/starten.xml`.

6 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`.

7 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`.

8 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`.

9 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`.

10 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`.

11 B. A. Burton, H. Edelsbrunner, J. Erickson, and S. Tillmann, editors. *Computational Geometric and Algebraic Topology*, volume 12 of *Oberwolfach Reports*. EMS Publishing House, 2015. `doi:10.4171/OWR/2015/45`.

12 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`.

13 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`.

14 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`.

15 F. R. K. Chung and P. D. Seymour. Graphs with small bandwidth and cutwidth. *Discrete Mathematics*, 75(1–3):113–119, 1989. `doi:10.1016/0012-365X(89)90083-6`.

16 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`.

17 R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Berlin, fifth edition, 2017. `doi:10.1007/978-3-662-53622-3`.

18 R. G. Downey and M. R. Fellows. *Parameterized complexity.* Monographs in Computer Science. Springer-Verlag, New York, 1999. doi:10.1007/978-1-4612-0515-9.

19 R. G. Downey and M. R. Fellows. *Fundamentals of parameterized complexity.* Texts in Computer Science. Springer, London, 2013. doi:10.1007/978-1-4471-5559-1.

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

21 W. Haken. Theorie der Normalflächen. *Acta Math.*, 105:245–375, 1961. doi:10.1007/BF02559591.

22 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.

23 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.

24 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.

25 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.

26 K. Huszár, J. Spreer, and U. Wagner. On the treewidth of triangulated 3-manifolds, 2017. arXiv:1712.00434.

27 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.

28 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.

29 B. Kleiner and J. Lott. Notes on Perelman's papers. *Geom. Topol.*, 12(5):2587–2855, 2008. doi:10.2140/gt.2008.12.2587.

30 G. Kuperberg. Knottedness is in NP, modulo GRH. *Adv. Math.*, 256:493–506, 2014. doi:10.1016/j.aim.2014.01.007.

31 G. Kuperberg. Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization, 2015. arXiv:1508.06720.

32 M. Lackenby. The efficient certification of knottedness and Thurston norm, 2016. arXiv:1604.00290.

33 M. Lackenby. Some conditionally hard problems on links and 3-manifolds. *Discrete & Computational Geometry*, 58(3):580–595, 2017. doi:10.1007/s00454-017-9905-8.

34 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.

35 E. E. Moise. Affine structures in 3-manifolds. V. The triangulation theorem and Hauptvermutung. *Ann. of Math. (2)*, 56:96–114, 1952. doi:10.2307/1969769.

36 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.

37 M. I. Ostrovskii. Minimal congestion trees. *Discrete Mathematics*, 285(1–3):219–226, 2004. doi:10.1016/j.disc.2004.02.009.

38 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.

39 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.

40 J. H. Rubinstein. An algorithm to recognize the 3-sphere. In S. D. Chatterji, editor, *Proceedings of the International Congress of Mathematicians: August 3–11, 1994, Zürich, Switzerland*, volume 1, pages 601–611. Birkhäuser Verlag, 1995.

41 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.

42 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.

43 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.

44 J. Schultens. *Introduction to 3-manifolds*, volume 151 of *Graduate Studies in Mathematics*. Amer. Math. Soc., Providence, RI, 2014. doi:10.1090/gsm/151.

45 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.

46 P. D. Seymour and R. Thomas. Call routing and the ratcatcher. *Combinatorica*, 14(2):217–241, 1994. doi:10.1007/BF01215352.

47 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.

48 W. P. Thurston. *Three-dimensional geometry and topology. Vol. 1*, volume 35 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1997. Edited by Silvio Levy. doi:10.1515/9781400865321.

49 R. Zentner. Integer homology 3-spheres admit irreducible representations in $SL(2, \mathbb{C})$, 2016. arXiv:1605.08530.