Efficient Algorithms to Decide Tightness*

Bhaskar Bagchi[1], Benjamin A. Burton[3], Basudeb Datta[2], Nitin Singh[4], and Jonathan Spreer[3]

1 **Theoretical Statistics and Mathematics Unit, Indian Statistical Institute, Bangalore, India**
`bbagchi@isibang.ac.in`

2 **Department of Mathematics, Indian Institute of Science, Bangalore, India**
`dattab@math.iisc.ernet.in.`

3 **School of Mathematics and Physics, The University of Queensland, Brisbane, Australia**
`bab@maths.uq.edu.au; j.spreer@uq.edu.au.`

4 **IBM India Research Lab, Bangalore, India**
`nitisin1@in.ibm.com.`

Abstract

Tightness is a generalisation of the notion of convexity: a space is tight if and only if it is "as convex as possible", given its topological constraints. For a simplicial complex, deciding tightness has a straightforward exponential time algorithm, but more efficient methods to decide tightness are only known in the trivial setting of triangulated surfaces.

In this article, we present a new polynomial time procedure to decide tightness for triangulations of 3-manifolds – a problem which previously was thought to be hard. In addition, for the more difficult problem of deciding tightness of 4-dimensional combinatorial manifolds, we describe an algorithm that is fixed parameter tractable in the treewidth of the 1-skeletons of the vertex links. Finally, we show that simpler treewidth parameters are not viable: for all non-trivial inputs, we show that the treewidths of both the 1-skeleton and the dual graph must grow too quickly for a standard treewidth-based algorithm to remain tractable.

1998 ACM Subject Classification G.2.1 [Combinatorics] Combinatorial algorithms, F.2.2 [Non-numerical Algorithms and Problems] Computations on discrete structures, Geometrical problems and computations

Keywords and phrases discrete geometry and topology, polynomial time algorithms, fixed parameter tractability, tight triangulations, simplicial complexes

Digital Object Identifier 10.4230/LIPIcs.SoCG.2016.12

1 Introduction

The notion of *convexity* is very powerful in mathematics. Many theorems in many different mathematical fields only hold in the case of a convex base space. However, in geometry and topology, the concept of convexity has significant limitations: most topological spaces simply do not admit a convex representation. That is, most topological features of a space, such as handles or "holes" in the space, are an obstruction to convexity.

* This work is supported by the Department of Industry and Science, Australia under the Australia-India Strategic Research Fund (project AISRF06660). The third author is also supported by the UGC Centre for Advanced Studies.

32nd International Symposium on Computational Geometry (SoCG 2016).
Editors: Sándor Fekete and Anna Lubiw; Article No. 12; pp. 12:1–12:15
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Nonetheless, there is a distinct intuition that even for topologically non-trivial spaces, some representations look “more convex” than others. For example, for a solid torus, a doughnut shape is intuitively “more convex” than a coffee mug with one handle.

The idea of *tightness* captures this intuition in a mathematically precise way that applies to a much larger class of topological spaces than just balls and spheres. We give a precise definition in Section 3, but roughly speaking, a particular embedding of a topological space into some Euclidean space is said to be *tight* if it is “as convex as possible” given its topological constraints. In particular, a topological ball or sphere is tight if and only if it is convex.

Originally, tight embeddings were studied by Alexandrov in 1938 as objects that minimise total absolute curvature [1]. Later work by Milnor [34], Chern and Lashof [15] and Kuiper [30] linked the concept to topology, by relating tightness to the sum of the Betti numbers β_i of a d-dimensional topological space (these essentially count i-dimensional “holes”; see Section 2.2 for details). The framework was then applied to polyhedral surfaces by Banchoff [5], and later fully developed in the combinatorial setting by Kühnel [27]. Finally, work by Effenberger [20], and the first and third authors [3] made the concept accessible to computations.

There is a straightforward relaxation of tightness linking it to another powerful concept in geometry and topology: Morse theory (see Section 2.3 for a brief introduction). In a sense, *all* Morse functions on a tight embedding of a topological space must be perfect; that is, they must satisfy the Morse inequalities with equality. [1] Deciding whether a given embedding is tight is closely related to the much-studied problem of finding a perfect Morse function: the former asks if every Morse function is perfect, and the latter asks if there is at least one.

This article deals with tightness in the discrete setting, where our topological spaces are represented as simplicial complexes. Here there are only a finite number of “essentially distinct” ways of (i) embedding the complex into Euclidean space, as used for tightness; or (ii) defining a simplexwise linear “height” function on it, as used in Morse theory. For instance, in the latter setting, the critical points of a simplexwise linear height function are completely determined by the order of the heights of the vertices of the complex.

In this discrete setting, we can replace the notion of a tight embedding with a purely combinatorial notion of a *tight combinatorial manifold*. These are rare but very special objects. For instance, tight combinatorial manifolds are conjectured to be strongly minimal (i.e., to contain the minimum number of faces of every dimension) amongst all triangulations of the same manifold [29, Conjecture 1.3]. More generally, they make deep connections between the *combinatorial* condition of tightness and its *geometric and topological* properties of a manifold, which are still far from being fully understood. Few such connections are known, which emphasises the importance of tightness in discrete and computational topology.

In this discrete setting, both deciding tightness of an embedding and finding a perfect Morse function become (decidable) algorithmic problems. The question remains as to how hard these problems are. It is known that, for discrete Morse functions as defined by Forman [21], finding perfect Morse functions is NP-hard in general [25, 33, 36], but also fixed parameter tractable in the treewidth of the dual graph of the triangulation [10] (see Sections 2.1 and 2.4 for definitions of the dual graph and treewidth respectively). For the related piecewise-linear variant of Morse theory as formulated by Kuiper and Banchoff [5, 19, 27] the hardness remains unknown; however, both theories are closely related since they both discretise smooth Morse theory in similar (but technically different) ways.

---

[1] This statement only holds in general when applied within the right version of Morse-type theory; namely, the theory of *piecewise linear functions* and the *lower star filtration* [19]. This is sometimes also referred to as the theory of *regular simplex-wise linear functions* [27].

In this paper, we consider the hardness of deciding tightness in the discrete setting. In Section 5 of this article, we give a polynomial time solution in the case of 3-manifolds:

▶ **Theorem 1.** *Let M be a combinatorial 3-manifold with n vertices. There is an algorithm to decide whether or not M is tight with running time polynomial in n. Furthermore, the dominating term in the running time of the algorithm is the time required to compute the first Betti number $\beta_1(M, \mathbb{F}_2)$.*

This is a surprising result, given the close relation between deciding tightness and finding perfect Morse functions which, for Forman's discrete Morse theory, is known to be NP-hard (see above). Furthermore, this polynomial time solution links to a number of other problems in 3-dimensional computational topology where polynomial time procedures are unknown but conjectured to exist. For instance, finding a perfect Morse function on a 3-manifold solves specific instances of the 3-sphere and the unknot recognition problems, both of which are conjectured but not known to be polynomial time solvable [22, 23, 32, 35].

We next move to the more difficult setting of four dimensions. Here we do not obtain a polynomial time algorithm, but we do show that deciding tightness is fixed parameter tractable (see Section 2.4 for an overview of fixed parameter tractability):

▶ **Theorem 2.** *Let M be a combinatorial d-manifold, $d \leq 4$. Then deciding tightness for any field is fixed parameter tractable in the treewidth of the 1-skeletons of the vertex links of M.*

This essentially means that deciding tightness is polynomial time for inputs where the *parameter* – i.e., the treewidth of the 1-skeletons of the vertex links – is universally bounded.

This parameter is complex to describe, but our final result shows that this is necessarily so. In Section 7 we consider the simpler parameters of (i) the treewidth of the 1-skeleton of M; and (ii) the treewidth of the dual 1-skeleton of M, also known as the *dual graph* of M. We show that, for all non-trivial inputs, both of these parameters must grow with the input size, and therefore cannot be universally bounded. In other words, for these simpler parameters, the notion of fixed parameter tractability cannot help.

The results in this paper are not only relevant for computational geometry – they also hold significance for the study of tightness itself. One of the major difficulties in studying tightness is in obtaining explicit examples of tight combinatorial manifolds. Some infinite families are known for simple manifolds [27, 17, 8]; beyond these, only sporadic examples are known. However, these sporadic examples feature several of the most fascinating triangulations in the field, including minimal triangulations of the complex projective plane [28], the K3 surface [14], and a triangulation of an 8-manifold conjectured to be PL-homeomorphic to the quaternionic plane [6] (see [29] for a comprehensive overview of many more examples).

Most of these examples were proven to be tight using theoretical arguments specific to each case. Fast algorithms, such as those presented here in Theorems 1 and 2, will open the door to finding and exploring new tight combinatorial manifolds using powerful computer-assisted techniques.

The proofs in this paper are relatively simple – they use significant theoretical groundwork from the literature to do the heavy lifting, as outlined through Sections 2–4. The main contribution of this paper is to bring together significant results from combinatorial topology, computational complexity and Morse theory to build tractable and fixed parameter tractable algorithms for problems that have until now been beyond the realm of computation.

2 Preliminaries

2.1 Combinatorial manifolds

Let C be an abstract simplicial complex, i.e., a simplicial complex without a particular embedding. For any vertex $v \in C$, the *star of* v *in* C is the set of all faces of C containing v, as well as all of their subfaces, and is denoted by $\mathrm{st}_C(v)$. The boundary of $\mathrm{st}_C(v)$, that is, all faces of $\mathrm{st}_C(v)$ which do not contain v, is called the *link of* v *in* C, written $\mathrm{lk}_C(v)$. Given an abstract simplicial complex C, its underlying set $|C|$ is called the *geometric carrier of* C. An abstract simplicial complex is said to be *pure of dimension* d if all of its maximal faces (that is, faces which are not contained in any other face as a proper surface) are of dimension d.

A *combinatorial* d*-manifold* is an abstract pure simplicial complex M of dimension d in which all vertex links are combinatorial $(d-1)$-dimensional standard PL-spheres. A combinatorial $(d-1)$-dimensional standard PL sphere is a combinatorial $(d-1)$-manifold whose underlying space, i.e., its *geometric carrier*, is a PL standard sphere.

The f*-vector* of M is the $(d+1)$-tuple $f(M) = (f_0, f_1, \ldots f_d)$ where f_i denotes the number of i-dimensional faces of M. The set of vertices of M is denoted by $V(M)$, and the d-dimensional faces of M are referred to as *facets*.

We call M k*-neighbourly* if $f_{k-1} = \binom{f_0}{k}$, that is, if M contains all possible $(k-1)$-dimensional faces. Given a combinatorial manifold M with vertex set $V(M)$ and $W \subset V(M)$, the *sub-complex of* M *induced by* W is the simplicial complex $M[W] = \{\sigma \in M \mid V(\sigma) \subset W\}$, i.e., the simplicial complex of all faces of M whose vertices are contained in W.

A pure simplicial complex of dimension d is said to be a *weak pseudomanifold* if every $(d-1)$-dimensional face is contained in at most two facets. Naturally, any combinatorial manifold is a weak pseudomanifold. However, a weak pseudomanifold is not always a combinatorial manifold – it might allow singularities around faces of co-dimension ≥ 2 (e.g., the apex of a double cone). Given a weak pseudomanifold M, the *dual graph* of M is the graph whose vertices represent the facets of M, and whose edges represent gluings between the facets along common $(d-1)$-faces of M; this is denoted $\Gamma(M)$. Weak pseudomanifolds are the most general class of simplicial complexes for which a dual graph can be defined.

For brevity, we call a combinatorial manifold a *triangulation* of the underlying space.

2.2 Homology and $\mathbb{F}$-orientability

Given a d-dimensional topological space M and a field $\mathbb{F}$, the *homology groups* of M are a series of $\mathbb{F}$-vector spaces $H_\star(M,\mathbb{F}) = (H_0(M,\mathbb{F}),\ H_1(M,\mathbb{F}),\ \ldots,\ H_d(M,\mathbb{F}))$ associated with M. Roughly speaking, the ith homology group counts i-dimensional "holes" in M.

Homology can be defined on abstract simplicial complexes as follows. Let K be a d-dimensional abstract simplicial complex with an ordering on its vertices $V(K)$, and let $\mathbb{F}$ be a field. The *group of* i*-chains of* K $(0 \leq i \leq d)$, denoted $C_i(K,\mathbb{F})$, is the group of formal sums of ordered i-dimensional faces of K with coefficients in $\mathbb{F}$. The *boundary operator* is a linear operator $\partial_i : C_i(K,\mathbb{F}) \to C_{i-1}(K,\mathbb{F})$ that maps each i-face to its $(i-1)$-dimensional boundary: $\partial_i(v_0, \cdots, v_i) = \sum_{j=0}^{i}(-1)^j(v_0, \cdots, \widehat{v_j}, \cdots, v_i)$, where $(v_0, \ldots, v_i)$ denotes the ordered i-face with vertices $v_0, \ldots, v_i$, and $\widehat{v_j}$ means v_j is deleted from the tuple. Denote by $Z_i(K,\mathbb{F})$ and $B_{i-1}(K,\mathbb{F})$ the kernel and the image of ∂_i respectively. Observing $\partial_i \circ \partial_{i+1} = 0$, we define the ith (simplicial) homology group $H_i(K,\mathbb{F})$ of K to be the $\mathbb{F}$-vector space $H_i(K,\mathbb{F}) = Z_i(K,\mathbb{F})/B_i(K,\mathbb{F})$. The rank of $H_i(M,\mathbb{F})$ is called the *ith Betti number* of M, denoted $\beta_i(M,\mathbb{F})$. The field $\mathbb{F}$ is called the *field of coefficients*. This construction is a topological invariant, in that two distinct simplicial complexes triangulating the same topological space will have isomorphic homology groups.

As an example, the circle C and the solid torus T both have exactly one “1-dimensional hole”. Hence, they both have first homology group $H_1(C,\mathbb{F}) = H_1(T,\mathbb{F}) = \mathbb{F}$ (for any $\mathbb{F}$), and $\beta_1(T,\mathbb{F}) = 1$. The surface of a torus, however, has two “1-dimensional holes” and thus first homology group $H_1(T,\mathbb{F}) = \mathbb{F}^2$. The 2-dimensional sphere S has trivial first homology group $H_1(S,\mathbb{F}) = 0$. For more details about homology theory see [24].

A triangulation of a d-manifold is *orientable* if all of its d-faces Δ, and thus their boundaries $\partial_d\Delta$, can be endowed with a sign such that d-faces are glued together along their boundary $(d-1)$-faces with opposite signs in a globally consistent way. Otherwise the triangulation is called *non-orientable*. For instance, triangulations of the 2-sphere are orientable, but triangulations of the Möbius strip are not. Orientability is likewise a topological invariant.

Orientability can be generalised in terms of homology. A triangulation of a manifold is $\mathbb{F}$-*orientable* if it forms a d-dimensional void; that is, if $H_d(M,\mathbb{F}) = \mathbb{F}$ (instead of 0). All manifolds are orientable with respect to the field with two elements $\mathbb{F}_2$.

The homology groups of orientable manifolds satisfy a special relation known as *Poincaré duality*: if $\mathbb{F}$ is a field and M is an $\mathbb{F}$-orientable d-manifold, then $\beta_i(M,\mathbb{F}) = \beta_{d-i}(M,\mathbb{F})$. We make use of this fact in our algorithms to decide tightness for 3- and 4-manifolds.

2.3 Piecewise linear Morse theory

Smooth Morse theory studies functions $f : M \to \mathbb{F}$ from a manifold M to the real numbers. Here we introduce a piecewise linear analogue to this theory as defined in [19, 27].

Let M be a d-dimensional combinatorial manifold. A function $f : |M| \to \mathbb{R}$ is called *piecewise linear*, or *regular simplexwise linear (rsl)*, if $f(v) \neq f(w)$ for any two vertices $v \neq w$ of M and f is linear when restricted to any simplex of M.

Fix a field $\mathbb{F}$. A vertex $x \in V(M)$ is said to be *critical with respect to* $\mathbb{F}$ for an rsl-function $f : |M| \to \mathbb{R}$ if $H_\star(\mathrm{lk}_M(x)^-,\mathbb{F}) \neq (0,\ldots,0)$. Here $\mathrm{lk}_M(x)^-$ denotes the *lower link* of x with respect to f, defined as $\mathrm{lk}_M(x)^- = \mathrm{lk}_M(x)[\{y \in V(\mathrm{lk}_M(x)) \,|\, f(y) \le f(x)\}]$. Essentially, a point x is critical whenever an i-dimensional hole is contained in $M[\{y \in V(M) \,|\, f(y) \le f(x)\}]$, but not in $M[\{y \in V(M) \,|\, f(y) < f(x)\}]$. In the language of filtrations, the hole appears at value $f(x)$ in the filtration defined by f. Such a hole must intersect the lower link $\mathrm{lk}_M(x)^-$ in a hole of dimension $(i-1)$, yielding the homology relation described above.

We call a vertex x *critical of index i and multiplicity m* if $\beta_i(\mathrm{lk}_M(x)^-,\mathbb{F}) = m$. The vector $(\beta_0(\mathrm{lk}_M(x)^-,\mathbb{F}),\ldots,\beta_d(\mathrm{lk}_M(x)^-,\mathbb{F}))$ is called *multiplicity vector of x with respect to* $\mathbb{F}$. The sum $\mathbf{m}_i(f,\mathbb{F}) := \sum_{x\in M} \beta_i(\mathrm{lk}_M(x)^-,\mathbb{F})$ is called the *number of critical points of f of index i with respect to* $\mathbb{F}$. The sum of $\mathbf{m}_i(f,\mathbb{F})$ over all indices i is called the *number of critical points of f with respect to* $\mathbb{F}$. Note that in the theory of piecewise linear functions, higher order multiplicities cannot always be avoided.

▶ **Theorem 3** (Morse relations, [31]). *Let M be a combinatorial d-manifold, $f : |M| \to \mathbb{R}$ a piecewise linear function, and $\mathbb{F}$ a field. Then the following statements hold:*

i) $\beta_i(M,\mathbb{F}) \le \mathbf{m}_i(f,\mathbb{F})$,

ii) $\sum_{i=0}^{d} (-1)^i \mathbf{m}_i(f,\mathbb{F}) = \chi(M) = \sum_{i=0}^{d} (-1)^i \beta_i(M,\mathbb{F})$,

where $\chi(M)$ is an invariant of M called the Euler characteristic.

If there exists a field $\mathbb{F}$ such that $\mathbf{m}_i(f,\mathbb{F}) = \beta_i(M,\mathbb{F})$ for all $0 \le i \le d$, then f is called *perfect*. Geometrically speaking this means that M, seen from the direction described by f, is “as convex as possible”.

2.4 Parameterised complexity and treewidth

The framework of *parameterised complexity*, as introduced by Downey and Fellows [18], provides a refined complexity analysis for hard problems. Given a typically NP-hard problem p with A as its set of possible inputs, we choose some *parameter* $k : A \to \mathbb{N}$, which grades the inputs according to their corresponding parameter values. This *parameterised version* of p is then said to be *fixed parameter tractable (FPT)* with respect to parameter k if p can be solved in time $O(f(k) \cdot n^{O(1)})$, where f is an arbitrary (computable) function independent of the input size n and k is the parameter value of the input. Note that the exponent of n must be independent of k. In essence, if p is FPT in parameter k, then k (and not the problem size) encapsulates the hardness of p.

A priori, a parameter can be many things, such as the maximum vertex degree of a graph, the sum of the Betti numbers of a triangulation, or the output size. However, the significance of an FPT result strongly depends on specific properties of the parameter – the parameter should be small for interesting classes of problem instances and ideally efficient to compute.

In the setting of computational topology, the *treewidths* of various graphs associated with triangulations turn out to be such good choices of parameter [10, 12, 9]. Informally, the treewidth of a graph measures how “tree-like” a graph is. The precise definition is as follows.

▶ **Definition 4** (Treewidth). A *tree decomposition* of a graph G is a tree $T = (B, E)$ whose vertices $\{B_i \mid i \in I\}$ are called *bags*. Each bag B_i is a subset of vertices of G, and we require:

- every vertex of G is contained in at least one bag (*vertex coverage*);
- for each edge of G, at least one bag must contains both its endpoints (*edge coverage*);
- the induced subgraph of T spanned by all bags sharing a given vertex of G must form a (connected) subtree of T (*subtree property*).

The *width* of a tree decomposition is defined as $\max |B_i| - 1$, and the *treewidth* of G is the minimum width over all tree decompositions.

When describing FPT algorithms which operate on tree decompositions, the following construction has proven to be extremely convenient.

▶ **Definition 5** (Nice tree decomposition). A tree decomposition $T = (\{B_i \mid i \in I\}, E)$ is called a *nice tree decomposition* if T can be expressed as rooted tree for which:

1. Every bag of the tree T has at most two children.
2. The bag B_r at the root of the tree (called the *root bag*) has $|B_r| = 1$.
3. If bag B_j has no children, then $|B_j| = 1$ (in this case B_j is called a *leaf bag*).
4. If a bag B_i has two children B_j and B_k, then the sets B_i, B_j and B_k are identical (in this case B_i is called a *join bag*).
5. If a bag B_i has one child B_j, then either:
   a. $|B_i| = |B_j| + 1$ and $B_j \subset B_i$ (in this case B_i is called an *introduce bag*); or
   b. $|B_j| = |B_i| + 1$ and $B_i \subset B_j$ (in this case B_i is called a *forget bag*).

Nice tree decompositions are small and easy to construct by virtue of the following.

▶ **Lemma 6** ([26]). *Given a tree decomposition of a graph G that has width k and $O(n)$ bags, where n is the number of vertices of G, we can find a nice tree decomposition of G that also has width k and $O(n)$ bags in time $O(n)$.*

We make use of nice tree decompositions in Section 6.

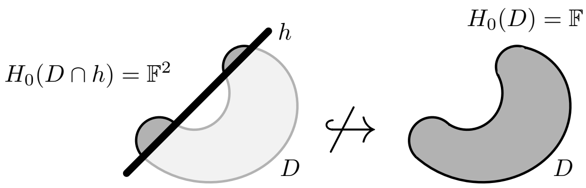

Figure 1 Non-convex embedding of the disk D in the plane. There exist a half space h such that $h \cap D$ has homological features (two connected components) which D has not.

3 Tightness

In its most general form, *tightness* is defined for compact connected subsets of Euclidean space. Here and in the following, $H_\star(M, \mathbb{F}) = (H_0(M, \mathbb{F}), H_1(M, \mathbb{F}), \ldots, H_d(M, \mathbb{F}))$ denote the simplicial homology groups of a d-manifold M with coefficients in the field $\mathbb{F}$.

Intuitively, a manifold M embedded into some Euclidean space E^d is tight if intersecting it with a half space $h \subset E^d$ does not introduce any topological (more precisely, homological) features in $h \cap M$ which are topologically (homologically) trivial in $M \subset E^d$. For instance, think of a non-convex embedding of the 2-dimensional disk D into E^2, as illustrated in Figure 1. We can always cut D with a half space $h \subset E^d$ such that $h \cap D$ is disconnected, or, in homology terms, $H_0(h \cap D, \mathbb{F}) = \mathbb{F}^2$ (note that multiple connected components are "0-dimensional holes"). But D itself is connected and thus $H_0(D, \mathbb{F}) = \mathbb{F}$. Hence a topological feature of $h \cap D$ disappears in D, and so D is not tight.

For a different type of non-tight embedding, imagine an embedding of the 3-dimensional ball B into E^3 with a dent. Because of the dent, we can find a half space $h \subset E^3$ such that $h \cap B$ is a solid torus. Thus, $h \cap B$ is still connected, but we have $H_1(h \cap B, \mathbb{F}) = \mathbb{F}$ whereas $H_1(B, \mathbb{F}) = 0$ and B is not tight. We formalise this intuition through the following definition.

▶ **Definition 7** (k-tightness and tightness [27])**.** A compact connected subset $M \subset E^d$ is called *k-tight with respect to a field* $\mathbb{F}$ if, for every open or closed half space $h \subset E^d$, the induced homomorphism $H_k(h \cap M, \mathbb{F}) \to H_k(M, \mathbb{F})$ is injective. If $M \subset E^d$ is k-tight with respect to $\mathbb{F}$ for all k, $0 \leq k \leq d$, then it is called $\mathbb{F}$-*tight*.

If a connected subset $M \subset E^d$ is referred to as *tight* without specifying a field, then it is usually understood that there exists a field $\mathbb{F}$ such that M is $\mathbb{F}$-tight. See Figure 2 on the left for a tight embedding of the torus into Euclidean 3-space.

If M is given as an abstract simplicial complex, tightness can be formulated as a combinatorial condition. In this setting, all half spaces from the geometric definition are replaced by subsets of vertices $W \subset V(M)$ of M, thought of as lying inside the half space. In particular, there are only $2^{|V(M)|}$ distinct subsets and thus only $2^{|V(M)|}$ "essentially distinct" half spaces to consider. The restriction $M \cap h$ from the geometric definition translates to combinatorics via the concept of induced subcomplexes $M[W]$.

▶ **Definition 8** (k-tightness and tightness [3, 27])**.** Let C be an abstract simplicial complex and let $\mathbb{F}$ be a field. We say that C is *tight with respect to* $\mathbb{F}$ if, for all subsets $W \subset V(C)$ and all $0 \leq k \leq d$, the induced homomorphism $H_k(C[W], \mathbb{F}) \to H_k(C, \mathbb{F})$ is injective.

See Figure 2 in the center and on the right hand side for the unique tight triangulations of the 2-sphere (the boundary of the tetrahedron) and the torus (the so-called Möbius torus, see [16] for an embedding of this complex into Euclidean 3-space with straight lines). Take a moment to verify that no induced subcomplex introduces holes which are filled in the full complex, or disconnected components which become connected in the full complex.

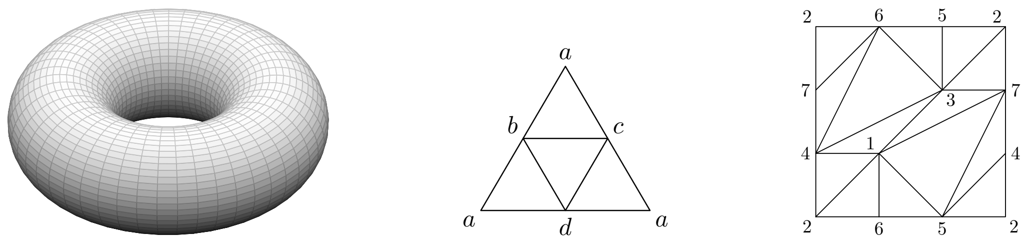

Figure 2 Left: tight embedding of the torus into E^3. Center: unique tight triangulation of the 2-sphere, the boundary of a tetrahedron. Right: unique tight triangulation of the torus with 7 vertices. Vertices with equal labels are identified.

Note that the above definition does not depend on a specific embedding. However, Definition 7 can be linked to Definition 8 by considering the standard embedding of a simplicial complex C into the $(|V(C)|-1)$-simplex (every vertex of C is sent to a unit vector in $E^{|V(C)|}$). Also, notice that in Definition 8, and as mentioned above, tightness is a definition depending on a finite number of conditions, namely $2^{|V(C)|}$, giving rise to an exponential time algorithm to decide tightness. Unless otherwise stated we work with Definition 8.

There are many criteria in the literature on when a simplicial complex is tight. See [29] for a more thorough survey of the field and [3, 4] for a summary of recent developments.

One of the most fundamental criteria says that an orientable triangulated surface is tight if and only if it is 2-*neighbourly* [27, Section 2A] (i.e., its set of edges forms a clique, see Section 2.1).

It is not difficult to see that the condition of 2-neighbourliness is necessary for tightness: Assume a connected simplicial complex C has two vertices u and v not connected by an edge, then the induced subcomplex $C[\{u,v\}]$ is not connected, and thus C cannot be 0-tight. Note that for dimensions greater than two, and for 2-dimensional complexes different from combinatorial surfaces, the converse of this statement no longer holds: there are 2-neighbourly complexes which are not tight.

In particular, for $d > 2$, no easy-to-check characterisation of tightness for general combinatorial d-manifolds is known (see [4] for a full characterisation of tightness of combinatorial 3-manifolds with respect to fields of odd characteristic).

Instead, the above condition for combinatorial surfaces generalises to combinatorial 3- and 4-manifolds in the following way.

▶ **Theorem 9** (Bagchi, Datta). *An $\mathbb{F}$-orientable combinatorial manifold of dimension ≤ 4 is $\mathbb{F}$-tight if and only if it is 0-tight and 1-tight with respect to $\mathbb{F}$.*

This directly follows from Theorem 2.6 (c) in [3]. The $\mathbb{F}$-orientability is necessary to apply Poincaré duality to the Betti numbers.

The two main results of this paper (presented in Sections 5 and 6) make use of this fact, yielding (i) a polynomial time procedure to decide tightness for combinatorial 3-manifolds, and (ii) a fixed parameter tractable algorithm for combinatorial 4-manifolds.

The σ- and μ-vectors

We briefly review the definition of the σ- and μ-vectors as introduced in [3] by the first and third authors. These *combinatorial invariants* build the foundation of a more combinatorial study of tightness of simplicial complexes. In essence they count the number of critical points of all piecewise linear functions on a complex simultaneously in order to detect a function with too many critical points (i.e., a non-perfect Morse function) which then can be used to conclude that the complex is not tight.

▶ **Definition 10** (Bagchi, Datta [3]). Let C be a simplicial complex of dimension d. The σ-vector $(\sigma_0, \sigma_1, \ldots, \sigma_d)$ of C with respect to a field $\mathbb{F}$ is defined as

$$\sigma_i = \sigma_i(C; \mathbb{F}) := \sum_{A \subseteq V(C)} \frac{\tilde{\beta}_i(C[A], \mathbb{F})}{\binom{f_0(C)}{|A|}}, \quad 0 \leq i \leq d,$$

where $\tilde{\beta}_i$ denotes the reduced ith Betti number ($\beta_i = \tilde{\beta}_i$ for $i > 0$, and $\beta_0 = \tilde{\beta}_0 + 1$). For $i > \dim(C)$, we formally set $\sigma_i(X; \mathbb{F}) = 0$.

In other words, $\sigma_i(C, \mathbb{F})$ adds the number of i-dimensional homological features, $0 \leq i \leq d$, in induced subcomplexes $C[A]$, weighted by the number of subsets $A \subset V(C)$ of size $|A|$. Definition 10 can then be used to define the μ-vector.

▶ **Definition 11** (Bagchi [2] and Bagchi, Datta [3]). Let C be a 2-neighbourly simplicial complex of dimension d, $n = |V(C)|$. We define

$$\mu_0 = \mu_0(C; \mathbb{F}) := 1 \qquad (1)$$
$$\mu_i = \mu_i(C; \mathbb{F}) := \delta_{i1} + \frac{1}{n} \sum_{v \in V(C)} \sigma_{i-1}(\mathrm{lk}_v(C)), \ 1 \leq i \leq d,$$

where δ_{ij} is the Kronecker delta.

Note that $\mu_1(C; \mathbb{F})$ only depends on the 0-homology of subcomplexes of C and 0-homology is independent of the field $\mathbb{F}$. Hence, we sometimes write $\mu_1(X)$ instead of $\mu_1(X; \mathbb{F})$.

Intuitively speaking, μ_i averages over the number of homological features in the vertex links of dimension $i - 1$. Thus, in some sense μ_i counts the average number of critical points of index i of a piecewise linear function on C (see the definition of critical points for piecewise linear functions in Section 2.3). The following result is thus essentially a consequence of the Morse relations, see Theorem 3.

▶ **Lemma 12** (Bagchi, Datta [3]). *Let M be an $\mathbb{F}$-orientable, 2-neighbourly, combinatorial closed d-manifold, $d \leq 4$. Then $\beta_1(M; \mathbb{F}) \leq \mu_1(M)$, and equality holds if and only if M is $\mathbb{F}$-tight.*

This is a special case of a much more general result, but suffices for the purpose of this work. To read more about many recent advances in studying tightness of simplicial complexes using the framework of combinatorial invariants, see [4].

4 Vertex links of tight 3-manifolds

Before we give a description of our polynomial time algorithm in Section 5, we first need to have a closer look at the vertex links of tight combinatorial 3-manifolds and a way to accelerate the computation of their σ-vectors.

▶ **Theorem 13** (Bagchi, Datta, Spreer [4]). *Let M be a tight triangulation of a closed combinatorial 3-manifold. Then each vertex link of M is a combinatorial 2-sphere obtained from a collection of copies of the boundary of the tetrahedron S^2_4 and the boundary of the icosahedron I^2_{12} glued together by iteratively cutting out triangles and identifying their boundaries.*

Furthermore, we have the following decomposition theorem for the σ-vector.

▶ **Theorem 14** (Bagchi, Datta, Spreer [4]). *Let C_1 and C_2 be induced subcomplexes of a simplicial complex C and $\mathbb{F}$ be a field. Suppose $C = C_1 \cup C_2$ and $K = C_1 \cap C_2$. If K is k-neighbourly, $k \geq 2$, then*

$$\sigma_i(C;\mathbb{F}) = (f_0(C)+1)\left(\frac{\sigma_i(C_1;\mathbb{F})}{f_0(C_1)+1} + \frac{\sigma_i(C_2;\mathbb{F})}{f_0(C_2)+1} - \frac{\sigma_i(K;\mathbb{F})}{f_0(K)+1}\right)$$

for $0 \leq i \leq k-2$.

Moreover, we call a simplicial complex C a *primitive simplicial complex* if it does *not* admit a splitting $C_1 \cup C_2 = C$ such that $K = C_1 \cap C_2$ is k-neighbourly, $k \geq 2$.

In particular, this theorem applies to the σ_0-value of a combinatorial 2-sphere S which can be split into two discs D_1 and D_2 along a common triangle K (which is 2-neighbourly), i.e. $S = D_1 \cup_K D_2$. In this case we write $S = S_1 \# S_2$, where S_i is the 2-sphere obtained from D_i, $1 \leq i \leq 2$, by pasting a triangle along the boundary. Note that whether or not we paste the last triangle into D_i does not change the σ_0-value of the construction. Now, Theorem 14, together with the explicit computation of the σ_0-value of the tetrahedron and the icosahedron, gives rise to the following statement.

▶ **Corollary 15** (Bagchi, Datta, Spreer [4]). *Let $k, \ell \geq 0$ and $(k,\ell) \neq (0,0)$. Then*

$$\sigma_0(\ kI^2_{12} \,\#\, \ell S^2_4\ ) = (9k+\ell+3)\left(\frac{617}{1716}k + \frac{1}{20}\ell - \frac{1}{4}\right).$$

5 Deciding tightness of 3-manifold triangulations in polynomial time

In this section we give a proof of Theorem 1, that is, we describe a polynomial time procedure to decide tightness for combinatorial 3-manifolds. The dominant part of the running time of the procedure accounts for computing the first Betti number of the combinatorial 3-manifold which runs in $O(n^6) = O(t^3)$ time, where n is the number of vertices and t is the number of tetrahedra of the triangulation.

The algorithm accepts any 3-dimensional simplicial complex M on n vertices, given as a list of abstract tetrahedra, i.e., subsets of size four of $\{1,\ldots,n\}$, together with a field $\mathbb{F}$ of characteristic $\chi(\mathbb{F})$ as input. It checks if M is an $\mathbb{F}$-orientable combinatorial manifold and, in the case it is, decides whether or not M is tight with respect to $\mathbb{F}$.

We use the fact that, by Lemma 12, a 2-neighbourly, $\mathbb{F}$-orientable combinatorial 3-manifold M is $\mathbb{F}$-tight if and only if $\beta_1(M;\mathbb{F}) = \mu_1(M)$.

First note that there is an $O(n^2\log(n))$ procedure to determine whether an n-vertex 3-dimensional simplicial complex is an $\mathbb{F}$-orientable 2-neighbourly combinatorial manifold:

- Check that M has $\binom{n}{2}$ edges and $(\binom{n}{2}-n)$ tetrahedra, which can be done in $O(n^2)$ time.
- Check that each of the triangles occurs exactly in two tetrahedra and store this gluing information (this can be done in almost quadratic time $O(n^2\log n)$).
- Compute all n vertex links of M. Since M is 2-neighbourly by the above, each vertex link must have $n-1$ vertices. Moreover, if M is a combinatorial manifold, each of the vertex links triangulates a 2-sphere. Since an $(n-1)$-vertex 2-sphere must have $2n-6$ triangles, we can either compute these vertex links in $O(n)$ time, or else conclude that M is not a combinatorial 3-manifold because some vertex link exceeds this size.

- Check that each vertex link is a 2-sphere. Since we know from the above that every triangle in M is contained in exactly two tetrahedra, and hence each edge in a link is contained in two triangles, this can be done by computing the Euler characteristic of each link, again a linear time procedure for each vertex link.
- If the characteristic of $\mathbb{F}$ is distinct from 2, use the gluing information from above to compute an orientation on M in quadratic time. Otherwise M is always $\mathbb{F}$-orientable.

Now, if M fails to be a combinatorial 3-manifold we stop. If M is a combinatorial 3-manifold, but either not 2-neighbourly or not $\mathbb{F}$-orientable we conclude that M is not tight (note that a non-orientable manifold can never be tight, see [3, Proposition 2.9 (b)]).

In case M is 2-neighbourly and $\mathbb{F}$-orientable we now have to test whether $\beta_1(M;\mathbb{F}) = \mu_1(M)$. More precisely, we must compute $\sigma_0(\mathrm{lk}_M(v))$ for all vertices $v \in V(M)$, a procedure which naïvely requires the analysis of $O(2^n)$ induced subcomplexes of $\mathrm{lk}_M(v)$.

However, using Corollary 15, computing the σ_0-value of the vertex link of a tight combinatorial 3-manifold simplifies to the following algorithm to be carried out for each vertex link S (note that because of the 2-neighbourliness, every vertex link needs to contain $n-1$ vertices, $3n-9$ edges, and $2n-6$ triangles):

- Enumerate all induced 3-cycles of S and split S along these cycles into separate connected components. This can be done in $O(n^3)$ time by first storing a list of adjacent edges for each vertex and a list of adjacent triangles for each edge, and then running over all vertex subsets of size three checking if the induced subcomplex is an empty 3-cycle.
- For each connected component, check if the 1-skeleton of the component is isomorphic to the graph of either S^2_4 or I^2_{12}. If not, return that M is not tight. Otherwise, sum up the number of connected components of each type. Note that each component can be processed in constant time, but that there is a linear number of components.
- Use Corollary 15 to compute the σ_0-value of the link.

Add up all σ_0-values and divide by n to obtain $\mu_1(M)$. Now, we know from Lemma 12 that M is tight if and only if $\mu_1(M) = \beta_1(M,\mathbb{F})$. Thus, if $\mu_1(M)$ is not an integer, M cannot be tight. This step overall requires a running time of $O(n^3)$.

Finally, suppose $\mu_1(M)$ is an integer. In this case, we must compute $\beta_1(M,\mathbb{F})$. For all $\mathbb{F}$, this information is encoded in $H_1(M,\mathbb{Z})$ which can be computed in $O(n^6)$ by determining the Smith normal form of the boundary matrix. Hence, this last step dominates the running time of the algorithm.

Altogether, checking for tightness can be done in the same time complexity as computing homology, a task which is considered to be easy in computational topology.

Furthermore, recent theoretical results in [4] show that, if the characteristic of $\mathbb{F}$ is odd, then M must be what is called a 2-neighbourly, *stacked* combinatorial manifold. Such an object can be identified in $O(n^2 \log(n))$ time (i.e., almost linear in the number of tetrahedra).

6 A fixed parameter tractable algorithm for dimension four

In the previous section, we prove that deciding tightness of 3-manifolds can be done efficiently. However, the algorithm for dimension three relies heavily on special properties of the vertex links. No such characterisation of the vertex links is known in higher dimensions.

However, both Lemma 12 and Theorem 14 can still be applied in the 4-dimensional setting. Hence, any computation of the σ_0-value of the vertex link of some combinatorial 4-manifold immediately reduces to computing the σ_0-value of the primitive components of that vertex link. For the remaining primitive pieces we have the following.

▶ **Theorem 16.** *Let C be a simplicial complex C whose 1-skeleton has treewidth $\leq k$. Then there exists an algorithm to compute $\sigma_0(C)$ in $O(f(k)\ n^5)$ time, where n is the number of vertices of C.*

Proof. We give an overview of the structure of this algorithm.

Note that for a simplicial complex C, $\sigma_0(C)$ only depends on the 1-skeleton C_1 of C. Thus, let $T = (B, E)$ be a nice tree decomposition of C_1. We write $B = \{B_1, B_2, \ldots, B_r\}$, where the bags are ordered in a bottom-up fashion which is compatible with a dynamic programming approach. In other words: the tree T is rooted; whenever B_i is a parent of B_j we have $i > j$; and our algorithm processes the bags in the order $B_1, \ldots, B_r$.

For each bag B_i, we consider the induced subgraph $C_1[B_i]$ spanned by all vertices in the bag B_i. Furthermore, we denote the induced subgraph of the 1-skeleton spanned by all vertices contained in bags in or underneath B_i by $C^-_{1,i}$.

Given a bag B_i: for each subset of vertices $S \subset B_i$, for each partition π of elements of S, and for all integers c, m with $c \leq m \leq n$, we count the number of induced subcomplexes of $C^-_{1,i}$ with m vertices and c connected components whose vertex set intersects bag B_i in precisely the set S, and whose connected components partition this set S according to π. Note that the count c includes connected components which are already "forgotten" (i.e., which do not meet bag B_i at all). Here, n denotes the number of vertices of C.

These lists can be trivially set up in constant time for each one-vertex leaf bag.

For each introduce bag, the list elements must be updated by either including the added vertex to the induced subcomplex or not. Note that in each step, the edges inside $C_1[B_i]$ place restrictions on which partitions of subsets $S \subset B_i$ can correspond to induced subcomplexes in $C^-_{1,i}$. The overall running time of such an introduce operation is dominated by the length of the list, which is at most quadratic in n multiplied by a function in k (for each subset and partition of vertices in B_i, up to $O(n^2)$ distinct list items can exist corresponding to different values of c, m).

For each forget bag, we remove the forgotten vertex from each list item, thereby possibly aggregating list items with equal values of S, π, c and m. This operation again has a running time dominated by the length of the lists.

Finally, whenever we join two bags, we pairwise combine list elements whenever the underlying induced subcomplexes are well-defined (i.e., whenever the subsets in the bag coincide and the partitions are compatible). This requires $O(n^4)$ time in total (the product of the two child list lengths).

After processing the root bag we are left with $O(n^2)$ list entries, labelled by the empty set, the empty partition, and the various possible values of c, m. Given the values of these list items it is now straightforward to compute $\sigma_0(C)$, as in Definition 10. Given that there are overall $O(n)$ bags to process, we have an overall running time of $O(f(k)\ n^5)$. ◀

The FPT algorithm to decide tightness for d-dimensional combinatorial manifolds M, $d \leq 4$, now consists of computing $\mu_1(M)$ by applying Theorem 16 to each vertex link, and comparing $\mu_1(M)$ to $\beta_1(M, \mathbb{F})$. Since the computation of $\mu_1(M)$ is independent of $\mathbb{F}$ and $\beta_1(M, \mathbb{F}) \leq \mu_1(M)$ for all $\mathbb{F}$, we can choose $\mathbb{F}$ to maximise β_1 and, by Theorem 9, M is tight if and only if

$$\max_{\mathbb{F}} \ \beta_1(M, \mathbb{F}) = \mu_1(M).$$

Following the proof of Theorem 16, the overall running time of this procedure is $O(\text{poly}(n) + n^6 f(k))$ for some function f.

Note that the 1-skeleton of M must always be the complete graph (or M is trivially not tight). It follows that each vertex link of M must contain $n-1$ vertices – but possibly only a linear number of edges. In this case the treewidth of the 1-skeleton of the links of M maybe very low and the above algorithm efficient. However, the 1-skeleton of the vertex link of a tight combinatorial 4-manifold M can be the complete graph, which has maximal treewidth. In this case, the algorithm presented in this section fails to give feasible running times. But in this case M must be a 3-neighbourly simply connected combinatorial 4-manifold which is tight by [27, Theorem 4.9]. Hence, worst case running times for the algorithm above are expected for combinatorial 4-manifolds of high topological complexity and with low, but strictly positive, first Betti numbers.

7 The treewidth of the dual graph of tight triangulations

It seems plausible that similar techniques as applied in Section 6 can be used to decide tightness in arbitrary dimensions. In fact, it can be shown that deciding j-tightness with respect to the field $\mathbb{F}_2$ is fixed parameter tractable in the treewidth of the dual graph of the combinatorial d-manifold. However, such an algorithm does not provide any real information about the hardness of the problem of tightness due to the following result (see [9] by the first author and Downey for a proof of a theorem implying this more special result).

▶ **Theorem 17.** *Let M be a connected combinatorial d-manifold, d fixed, with t facets and n vertices. Then either M is trivially not tight, or the treewidth k of the dual graph $\Gamma(M)$ of M satisfies*

$$k \in \Omega(t^{1/d} d^{-1}).$$

Proof. Let $\Gamma(M)$ be of treewidth k and let $T = (B, E)$ be a tree-decomposition of $\Gamma(M)$ of width k. For each bag $B_i \in B$, define B'_i to be the set of all vertices of M which are contained in some facet in B_i.

Claim: The tree $T' = (B', E)$ is a tree-decomposition of the primal graph of M.

Every vertex and every edge of M is contained in some facet of M, thus T' trivially satisfies the properties of *node coverage* and *arc coverage* of a tree-decomposition.

To verify the *subtree property*, let v be a vertex of M, let $B_v \subset B'$ be the set of bags containing v, and let $X, Y \in B_v$ be any two bags in this subset. By construction, both X and Y correspond to bags in B containing facets Δ_X and Δ_Y in the star of v in M. Since M is a combinatorial manifold, the star of v is a ball and hence there exists a sequence

$$\Delta_X = \Delta_0, \Delta_1, \ldots, \Delta_\ell = \Delta_Y$$

of facets in the star of v such that consecutive entries share a common $(d-1)$-face. Since $T = (B, E)$ is a tree-decomposition of $\Gamma(M)$ the set of bags containing a node corresponding to Δ_i, $0 \leq i \leq \ell$, is a subtree of T. Moreover, any two subtrees of consecutive entries Δ_i, Δ_{i+1} must intersect because of the edge coverage property of T. Hence, there is a path in T' from X to Y of bags containing v (i.e., bags in B_v). Since $X, Y \in B_v$ were arbitrary it follows that $T'[B_v]$ is connected and thus a subtree. This proves the claim.

Thus T' is a tree-decomposition of the 1-skeleton of M of width $(d+1)(k+1)-1$ and the treewidth of the 1-skeleton of M is at most this number.

To complete the proof, let us first assume that two vertices v and w of M are not connected by an edge. Then the induced subcomplex $M[\{v, w\}]$ is not connected but M is. It follows that M is not 0-tight and thus trivially not tight.

Now assume that any two vertices of M are connected by an edge. Hence the 1-skeleton of M is the complete graph on the set of vertices which has treewidth $n-1$, and we have $(d+1)(k+1) \geq n$. Since M is a combinatorial manifold it follows from the upper bound theorem that $t \in O(n^d)$, and thus $k \in \Omega(t^{1/d}d^{-1})$. ◀

The treewidth of the dual graph of a triangulation has recently been established as a standard tool to obtain fixed parameter tractable algorithms for problems in computational geometry and topology, see for example [10, 12, 9, 11]. Theorem 17 seems to significantly restrict the power of this approach, in particular in the simplicial setting where many typical input triangulations are 2-neighbourly, i.e., contain the complete graph in their 1-skeleton.

However, few problems in topology require us not to modify the input triangulation or even to stay in the simplicial category. In addition, there exist efficient algorithms to transform triangulations with a large number of vertices and edges into (PL-)homeomorphic triangulations with only a constant number of vertices. More precisely, there exist a polynomial time algorithm [7] due to the second author which turns any generalised triangulation of a (closed) 3-manifold M into a partial connected sum decomposition of M such that each summand is a triangulation with only one vertex. For higher dimensions there exist effective heuristics to produce crystallisations – which have a constant number of vertices – or even 1-vertex triangulations. For instance in [13] the second and fifth authors compute millions of distinct 1-vertex, 1-edge triangulations of the so-called $K3$-surface represented by a simplicial complex containing the complete graph on 16 vertices.

References

1. Aleksandr D. Alexandrov. On a class of closed surfaces. *Recueil Math. (Moscow)*, 4:69–72, 1938.
2. Bhaskar Bagchi. The mu vector, Morse inequalities and a generalized lower bound theorem for locally tame combinatorial manifolds. *European J. Combin.*, 51:69–83, 2016.
3. Bhaskar Bagchi and Basudeb Datta. On stellated spheres and a tightness criterion for combinatorial manifolds. *European J. Combin.*, 36:294–313, 2014.
4. Bhaskar Bagchi, Basudeb Datta, and Jonathan Spreer. Tight triangulations of closed 3-manifolds. *European J. Combin.*, 54:103–120, 2016.
5. Thomas F. Banchoff. Critical points and curvature for embedded polyhedral surfaces. *Amer. Math. Monthly*, 77:475–485, 1970.
6. Ulrich Brehm and Wolfgang Kühnel. 15-vertex triangulations of an 8-manifold. *Math. Ann.*, 294(1):167–193, 1992.
7. Benjamin A. Burton. A new approach to crushing 3-manifold triangulations. *Discrete Comput. Geom.*, 52(1):116–139, 2014.
8. Benjamin A. Burton, Basudeb Datta, Nitin Singh, and Jonathan Spreer. A construction principle for tight and minimal triangulations of manifolds, 2015. 27 pages, 2 figures. URL: `http://arxiv.org/abs/1511.04500`.
9. Benjamin A. Burton and Rodney G. Downey. Courcelle's theorem for triangulations. `arXiv:1403.2926 [math.GT]`, 24 pages , 7 figures, 2014.
10. Benjamin A. Burton, Thomas Lewiner, João Paixão, and Jonathan Spreer. Parameterized complexity of discrete morse theory. In *Proceedings of the Twenty-Ninth Annual Symposium on Computational Geometry*, pages 127–136. ACM SIAM, 2013.
11. Benjamin A. Burton, Clément Maria, and Jonathan 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 1*, pages 281–293. Springer, 2015. `doi:10.1007/978-3-662-47672-7_23`.

12 Benjamin A. Burton and Jonathan Spreer. The complexity of detecting taut angle structures on triangulations. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 168–183. ACM SIAM, 2013. URL: `http://arxiv.org/abs/1207.0904`.

13 Benjamin A. Burton and Jonathan Spreer. Computationally proving triangulated 4-manifolds to be diffeomorphic, 2013. 29th ACM Symposium on Computational Geometry, Young Researchers Forum, Collections of abstracts, 2013, pages 15–16.

14 Mario Casella and Wolfgang Kühnel. A triangulated $K3$ surface with the minimum number of vertices. *Topology*, 40(4):753–772, 2001.

15 Shiing-Shen Chern and Richard K. Lashof. On the total curvature of immersed manifolds. *Amer. J. Math.*, 79:306–318, 1957.

16 Ákos Császár. A polyhedron without diagonals. *Acta Univ. Szeged. Sect. Sci. Math.*, 13:140–142, 1949.

17 Basudeb Datta and Nitin Singh. An infinite family of tight triangulations of manifolds. *J. Combin. Theory Ser. A*, 120(8):2148–2163, 2013.

18 Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*, volume 3. Springer, 1999.

19 Herbert Edelsbrunner and John L. Harer. *Computational Topology: An Introduction.* Applied Mathematics. American Mathematical Society, 2010.

20 Felix Effenberger. Stacked polytopes and tight triangulations of manifolds. *J. Combin. Theory Ser. A*, 118(6):1843–1862, 2011.

21 Robin Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998.

22 Joel Hass and Greg Kuperberg. The complexity of recognizing the 3-sphere. *Oberwolfach Reports*, 24:1425–1426, 2012.

23 Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. The computational complexity of knot and link problems. *J. ACM*, 46(2):185–211, 1999.

24 Allen Hatcher. *Algebraic Topology.* Cambridge University Press, 2002.

25 Michael Joswig and Marc E. Pfetsch. Computing optimal Morse matchings. *SIAM J. Discrete Math.*, 20(1):11–25, 2006.

26 Ton Kloks. *Treewidth: Computations and Approximations*, volume 842. Springer, 1994.

27 Wolfgang Kühnel. *Tight polyhedral submanifolds and tight triangulations*, volume 1612 of *Lecture Notes in Math.* Springer-Verlag, Berlin, 1995.

28 Wolfgang Kühnel and Thomas F. Banchoff. The 9-vertex complex projective plane. *Math. Intelligencer*, 5(3):11–22, 1983.

29 Wolfgang Kühnel and Frank H. Lutz. A census of tight triangulations. *Period. Math. Hungar.*, 39(1-3):161–183, 1999. Discrete geometry and rigidity (Budapest, 1999).

30 Nicolaas H. Kuiper. Immersions with minimal total absolute curvature. In *Colloque Géom. Diff. Globale (Bruxelles, 1958)*, pages 75–88. Centre Belge Rech. Math., Louvain, 1959.

31 Nicolaas H. Kuiper. Morse relations for curvature and tightness. In *Proceedings of Liverpool Singularities Symposium, II*, volume 209 of *Lecture Notes in Math.*, pages 77–89, 1971.

32 Greg Kuperberg. Knottedness is in **np**, modulo GRH. *Adv. Math.*, 256:493–506, 2014.

33 Thomas Lewiner, Lopes Hélio, and Geovan Tavares. Toward optimality in discrete Morse theory. *Experiment. Math.*, 12(3):271–285, 2003.

34 John Milnor. On the relationship between the Betti numbers of a hypersurface and an integral of its Gaussian curvature (1950). In *Collected papers. Vol. 1, Geometry*, pages 15–26. Publish or Perish Inc., Houston, TX, 1994.

35 Saul 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., 2011.

36 Martin Tancer. Recognition of collapsible complexes is np-complete. `arXiv:1211.6254 [cs.CG]`, 21 pages, 13 figures, 2012.