

Efficiently Hex-Meshing Things with Topology

Jeff Erickson

Received: 10 September 2013 / Revised: 2 May 2014 / Accepted: 8 July 2014 /
Published online: 3 September 2014
© Springer Science+Business Media New York 2014

Abstract A topological quadrilateral mesh Q of a connected surface in \mathbb{R}^3 can be extended to a topological hexahedral mesh of the interior domain Ω if and only if Q has an even number of quadrilaterals and no odd cycle in Q bounds a surface inside Ω . Moreover, if such a mesh exists, the required number of hexahedra is within a constant factor of the minimum number of tetrahedra in a triangulation of Ω that respects Q . Finally, if Q is given as a polyhedron in \mathbb{R}^3 with quadrilateral facets, a topological hexahedral mesh of the polyhedron can be constructed in polynomial time if such a mesh exists. All our results extend to domains with disconnected boundaries. Our results naturally generalize results of Thurston, Mitchell, and Eppstein for genus-zero and bipartite meshes, for which the odd-cycle criterion is trivial.

Keywords Computational topology · Mesh generation · Cube complexes · Homology

Mathematics Subject Classification 57Q05 · 65N50 · 68U05 · 57M99

Portions of this work were done while the author was visiting IST Austria. Work on this paper was also partially supported by the National Science Foundation under Grant CCF 09-15519. A preliminary version of this paper was presented at the 29th Annual Symposium on Computational Geometry [28]. See <http://www.cs.uiuc.edu/~jeffe/pubs/hexmesh.html> for the most recent version of this paper.

J. Erickson
Department of Computer Science, University of Illinois at Urbana-Champaign,
Urbana, IL 61801, USA
e-mail: jeffe@illinois.edu

*Hexita, vexita,
When can we mesh wit' a
pure topological
complex of cubes?*

*Hexahedrizable
surface quad meshes are
null-cohomologous
inside the tubes!*

1 Introduction

Many applications in scientific computing call for three-dimensional geometric models to be decomposed into a mesh of geometrically simpler pieces. One of the most sought-after types of decomposition is a *hexahedral mesh*, or more colloquially, a *hex mesh*. In this context, a hexahedron is a convex polyhedron that is combinatorially equivalent to a cube: six quadrilateral facets, each in its own plane; twelve edges; and eight distinct vertices of degree 3. A hexahedral mesh is a set of hexahedra that cover a specified domain, such that any two hexahedra intersect in either the empty set, a vertex of both, an edge of both, or a facet of both.

Now suppose we are given a (not necessarily convex) polyhedron with convex quadrilateral faces. When can we construct a hexahedral mesh of the interior of the polyhedron, such that the boundary facets of the mesh are precisely the facets of the polyhedron? The hexahedral mesh can have an arbitrary finite number of vertices in the interior of the polyhedron, but no subdivision of the boundary is permitted. Despite decades of research, this problem is still wide open. No algorithm is known to construct hex meshes compatible with an arbitrary given quadrilateral mesh, or even to determine when a compatible hex mesh exists, even for the simple examples shown in Fig. 1. All known geometric meshes for these examples [16,67,74,75] include degenerate, inverted, and/or warped hexahedra.

The hex-meshing problem becomes much simpler if we ignore the precise geometry of the meshes and focus entirely on their topology. In this more relaxed setting, the input quadrilaterals and output hexahedra are not necessarily convex polygons or polyhedra, but rather topological disks and balls whose intersection patterns are consistent with the intersection patterns in geometric hex meshes. (We describe the precise intersection

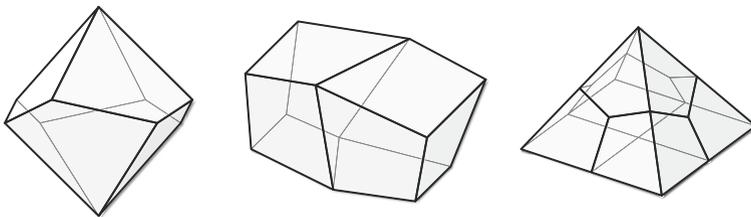


Fig. 1 An octagonal spindle [27], a bicuboid [8,9], and Schneiders' pyramid [61,62]

constraints in Sect. 2.) For example, any bicuboid has an obvious decomposition into two *topological* hexahedra sharing a single face.

This paper offers a complete solution to this topological hex-meshing problem. For example, let Q be a topological quadrilateral mesh (more colloquially, a *quad mesh*) of a connected surface in \mathbb{R}^3 . We prove that Q can be extended to a topological hex mesh of the interior domain Ω if and only if (1) Q has an even number of quadrilaterals and (2) no odd cycle in Q bounds a surface inside Ω . Moreover, if Q is given as a polyhedron in \mathbb{R}^3 with planar quadrilateral facets, we can either compute a compatible topological hex mesh or report correctly that no such mesh exists, in polynomial time. Our characterization and algorithm naturally generalize results of Thurston [68], Mitchell [44], and Eppstein [27] for genus-zero and bipartite meshes, for which the odd-cycle criterion is trivial.

Before describing our new results in detail, we first give an overview of these and other prior results.

1.1 Previous Results

1.1.1 Genus-Zero Surface Meshes

Thurston [68] and Mitchell [44] independently proved that any quadrilateral mesh of a topological sphere can be extended to a topological hex mesh of the interior ball if and only if the number of quadrilaterals is even. Thus, all the examples in Fig. 1 support topological hex meshes. Even parity is clearly a necessary condition for *any* hex mesh; each cube has an even number of facets, and facets are identified in pairs in the interior of the mesh. Thurston and Mitchell prove that evenness is also a sufficient condition in two stages. First, they show that the dual curves of the input quadrilateral mesh can be extended to a set of immersed surfaces in the interior of the ball, using a sequence of elementary moves related to *regular homotopy* [32, 72]; similar results were previously derived by Carter [17, 18] and by Csikós and Szűcs [22]. Second, they argue that the resulting surface immersion can be refined, by adding additional closed surfaces in the interior of the ball, into the dual complex of a hex mesh whose boundary is Q . We consider dual curves of quad meshes and dual surfaces of hex meshes in more detail in Sect. 2.

Mitchell's proof [44] can be translated into an algorithm that constructs a topological hex mesh with complexity $O(n^2)$, where n is the complexity of the input quad mesh. Eppstein [27] showed that carelessly choosing the elementary moves in the first stage of Mitchell's algorithm can lead to $\Omega(n^2)$ output complexity. More recent results of Nowik [55] imply the same $\Omega(n^2)$ lower bound even if all such choices are made optimally. With some effort, the output complexity can be reduced to $O(n)$ by using a larger vocabulary of elementary moves in the first stage of the algorithm; see related results of Carter [18] and Csikós and Szűcs [22].

Eppstein [27] described an algorithm that computes a topological hex mesh with complexity $O(n)$ in polynomial time, without first constructing a dual surface arrangement. Eppstein's algorithm extends the input quad mesh into a buffer layer of cubes, triangulates the interior of the buffer layer with $O(n)$ tetrahedra, subdivides each

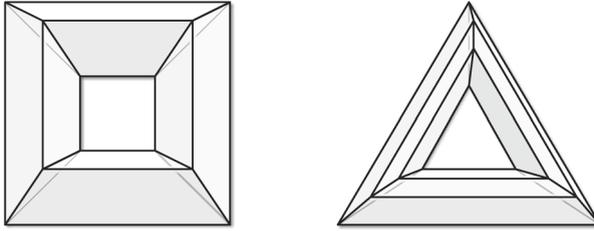


Fig. 2 Mitchell's 3×4 tori: only the mesh on the right supports an interior hex mesh [44]

interior tetrahedron into four cubes by central subdivision, and finally refines the buffer cubes into smaller cubes that meet the subdivided tetrahedra (and each other) consistently. Only the last stage is not explicitly constructive; the algorithm refines the boundary of each buffer cube into either 16 or 18 quadrilaterals and then invokes Thurston and Mitchell's proof. As Eppstein observes [27], it is straightforward but tedious to construct topological meshes for these remaining cases by hand, by following either Mitchell's algorithm or the later algorithms of Müller-Hannemann [48, 49] and Bern and Eppstein [8, 9].

1.1.2 Higher-Genus Surface Meshes

Both Mitchell and Eppstein consider the problem of hex-meshing domains with non-trivial topology, but neither offers a complete solution. Mitchell [44] gives both a necessary condition and a sufficient condition for a quad mesh with genus g to be the boundary of an interior hex mesh. First, a compatible hex mesh exists if one can find g disjoint topological disks in the interior body, each bounded by a cycle of *even* length in the quad mesh, that cut the interior body into a ball. Thus, at least in principle, Mitchell's algorithm can be applied to any *handlebody* in \mathbb{R}^3 . Second, a compatible hex mesh does *not* exist if there is a topological disk in the interior whose boundary is a cycle of *odd* length in the quad mesh.

Mitchell's conditions imply that the existence of a compatible hex mesh depends not only on the combinatorics of the surface quad mesh but also on its embedding in \mathbb{R}^3 . For example, a toroidal 3×4 grid of quadrilaterals has two natural embeddings in \mathbb{R}^3 , shown in Fig. 2. One embedding is the boundary of a hex mesh of the solid torus with three cubes; the other is not the boundary of a hex mesh, because there are interior disks bounded by triangles.

On the other hand, Eppstein [27] observed that his algorithm actually constructs an interior hex mesh for any *bipartite* quad mesh, regardless of its genus or its embedding in \mathbb{R}^3 . In particular, Eppstein's algorithm can be applied to bipartite quad meshes whose interior volumes are not handlebodies; consider, for example, a ball with a non-trivially knotted tunnel drilled through it.

In practice, hex meshes for complex domains are often constructed by partitioning into simply-connected subdomains and then meshing each subdomain independently [31, 48–50, 58]. In fact, this partitioning strategy is the main practical motivation for our requirement that the given boundary mesh cannot be refined. A more theoretical motivation is that topological hex meshing becomes trivial if we allow boundary

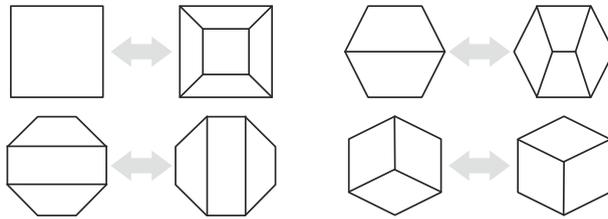


Fig. 3 The four 2-dimensional cube flips

refinement; Eppstein’s results [27] imply that *any* topological quad mesh can be extended to an interior hex mesh if we first split each quad into a 2×2 grid, because the refined quad mesh is always bipartite. In fact, in this setting, we can use Mitchell’s *Geode* template [45] to transition between the refined boundary and an interior triangulation [16], instead of Eppstein’s more complicated algorithm.

1.1.3 Cube Flips and Cubical Polytopes

Some time before 1995, Nathan Habegger asked [40, Problem 5.13] which pairs of PL-cubulations of the same manifold can be connected by a sequence of elementary operations variously called *bubble moves* [33,34] or *cube flips* [8,9,35,60]. A d -dimensional cube flip replaces a subcomplex B with another subcomplex B' , where B and B' are complementary simply connected subcomplexes of the boundary of a $(d + 1)$ -dimensional cube. The four two-dimensional cube flips are illustrated in Fig. 3. Cube flips were previously considered in the polytope literature, as a special case of a more general operation, by Billera and Sturmfels [10]. Reiner [60] asked whether any two cubical decompositions of the same zonotope are connected by cube flips.

Significant progress on Habegger’s question was reported by Funar [33–35]. In particular, Funar proved that two quadrangulations of the sphere S^2 are connected by cube flips if and only if they have the same number of facets modulo 2 [33]. It follows that any even quadrangulation of the sphere can be extended to a *shellable* hex mesh of the ball; the same theorem was proved independently by Müller-Hannemann [48,49] and Bern and Eppstein [8,9].

In a later paper [35], Funar proved that two quadrangulations Q_1 and Q_2 of the same orientable surface Σ are connected by cube flips if and only if they have the same number of facets modulo 2 and their dual curves Q_1^* and Q_2^* have the same homology class in $H_1(\Sigma; \mathbb{Z}_2)$; this theorem and its proof were the most direct inspiration for the new results in the current paper. Similar results were proved earlier by Nakamoto and Ota [52–54] for a different class of elementary moves, which they called *diagonal transformations*.

Building on earlier work on Banchoff [4], Babson and Chan [3] proved that any self-transverse piecewise-linear immersion of a $(d - 1)$ -manifold into the d -sphere can be refined into the dual immersion of a cube complex; we describe Babson and Chan’s refinement in our proof of Lemma 4. Babson and Chan call the natural cubical structure on the dual immersion the *derivative complex*. Using a generalization of the Hexhoop template developed by Yamakawa and Shimada for constructing hexahedral meshes [74], Schwartz [63] and Schwartz and Ziegler [64] proved that every

self-transverse piecewise-linear immersion of a $(d - 1)$ -manifold into the d -sphere can be refined into the dual immersion of a cubical d -polytope. In particular, they construct the first cubical 4-polytope with an odd number of facets, by refining Boy's classical immersion of the projective plane [13] (specifically, an orthogonal version of Boy's surface popularized by Petit [57]).

1.2 New Results

Our main theorem characterizes which surface quad meshes can be extended to interior hex meshes in terms of the *homology* (with \mathbb{Z}_2 coefficients) of certain subgraphs, either of the given quad mesh or its dual graph. Loosely, a subgraph is *null-homologous* in Ω if and only if it is the boundary of a possibly self-intersecting surface inside Ω . We define homology more carefully in Sect. 2.

Main Theorem *Let Ω be a compact connected subset of \mathbb{R}^3 whose boundary $\partial\Omega$ is a (possibly disconnected) 2-manifold, and let Q be a topological quad mesh of $\partial\Omega$ with an even number of facets. The following conditions are equivalent:*

- (1) Q is the boundary of a topological hex mesh of Ω .
- (2) Every subgraph of Q that is null-homologous in Ω has an even number of edges.
- (3) The dual curve arrangement Q^* is null-homologous in Ω .

Our characterization is a natural generalization of Thurston, Mitchell, and Eppstein's previous results. Bipartite quad meshes trivially satisfy condition (2), because every null-homologous subgraph is the union of disjoint closed walks. Similarly, genus-zero meshes trivially satisfy condition (3), because every closed walk on the sphere is null-homologous in the ball. We emphasize that our result does not require the domain Ω to have connected boundary or to be a handlebody.

Moreover, our condition (2) is closely related to Mitchell's necessary and sufficient conditions. On the one hand, it generalizes Mitchell's necessary condition (no disks in Ω are bounded by odd cycles in Q) by considering arbitrary immersed surfaces instead of only embedded disks. On the other hand, Mitchell's sufficient condition can be rephrased as follows: There are g disks in Ω , each bounded by an even number of edges in Q , such that the homotopy type of any cycle in Ω is determined by its crossing pattern with the disks. Our condition (2) has a similar reformulation: There are g immersed surfaces in Ω , each bounded by an even number of edges in Q , such that the homology type of any cycle in Ω is determined by its crossing pattern with the surfaces. See Sect. 3 for more details.

1.3 Outline

We prove our Main Theorem in several stages. The implication $(1) \Rightarrow (3)$ follows easily by considering the dual complex of the hex mesh, which we define carefully in Sect. 2. We prove the equivalence $(2) \Leftrightarrow (3)$ in Sect. 3 using Poincaré–Alexander–Lefschetz duality. We offer two proofs of the remaining implication $(3) \Rightarrow (1)$. In Sect. 4, we sketch a two-stage proof in the same spirit as Thurston and Mitchell's argument for

genus-zero meshes. First, essentially following arguments of Funar [35], we argue that if Q^* is null-homologous in Ω , then Q^* is the boundary of an *immersed* surface in Ω . Second, following arguments of Mitchell [44] and Babson and Chan [3], we show that any surface immersion can be refined to the dual of a hex mesh. Finally, in Sect. 5, we give a more direct proof, which uses a modification of Eppstein’s algorithm for bipartite meshes [27].

Both our proofs are constructive, although the first proof omits some implementation details because the resulting algorithm is less efficient. Given an arbitrary polyhedron in \mathbb{R}^3 with quadrilateral facets, our algorithms either construct a compatible topological hex mesh or correctly report that no such mesh exists, in polynomial time. Our second proof implies that the minimum number of hexahedra in within a constant factor of the minimum complexity of a topological triangulation of Ω that splits each quadrilateral facet in Q into exactly two triangles. For example, if Q is the boundary of a non-convex polyhedron with n convex quadrilateral facets that supports a hex mesh, then Q supports a hex mesh with complexity $O(n^2)$ [7, 21]. Since any hex mesh can be split into a triangulation with six simplices per hexahedron, this complexity bound is the best one can hope for.

2 Background

Before describing our results in detail, we first recall some standard definitions from combinatorial and algebraic topology. For more detailed background, we refer the reader to Edelsbrunner and Harer [26] and Hatcher [37].

2.1 Manifolds

A *d-manifold* is a Hausdorff topological space in which every point has a neighborhood homeomorphic to the Euclidean space \mathbb{R}^d . The integer d is the *dimension* of the manifold. For example, a 1-manifold is a disjoint union of circles. More generally, a *d-manifold with boundary* is a space Ω in which every point has a neighborhood homeomorphic to either \mathbb{R}^d or the closed halfspace $\{(x_1, \dots, x_d) \in \mathbb{R}^d \mid x_d \geq 0\}$. The subspace of points with halfspace neighborhoods is the *boundary* of the manifold, which we denote $\partial\Omega$. The boundary of any d -manifold is a $(d - 1)$ -manifold without boundary. We consider only compact 3-manifolds that are subsets of \mathbb{R}^3 in this paper, although most of our results apply to more general 3-manifolds without significant modification. We use the word *surface* to mean a compact 2-manifold, possibly with boundary and possibly with multiple components. The *genus* of a connected surface is the maximum number of disjoint simple cycles whose deletion leaves the surface connected.

2.2 Singularities and Immersions

We consider several types of well-behaved functions from surfaces to 3-manifolds with boundary. Fix a surface S , a 3-manifold $\Omega \subset \mathbb{R}^3$, and a continuous function

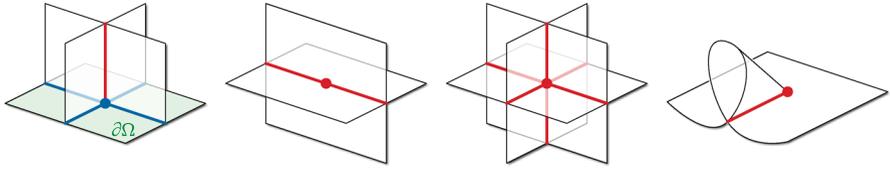


Fig. 4 Stable surface singularities: boundary double point, interior double point, triple point, branch point

$f: S \rightarrow \Omega$ that maps the boundary of S to the boundary of Ω . A point $f(x)$ in the image of f is *ordinary* if it lies in an open neighborhood U such that $\text{im } f \cap U$ is homeomorphic to either the plane (if x is in the interior of S) or a closed halfplane (if x is on the boundary of S), and *singular* otherwise.

For a *generic* surface map, every singular point has a neighborhood homeomorphic to one of four *stable singularities*, illustrated in Fig. 4. A *double point* has a neighborhood homeomorphic to two coordinate planes or two coordinate halfplanes; a *triple point* has a neighborhood homeomorphic to three coordinate planes; and a *branch point* has a neighborhood homeomorphic to a cone over a figure-8 (also known as a crosscap or a Whitney umbrella). A standard compactness argument implies that a generic surface map has a finite number of triple points, boundary double points, and branch points. Classical results of Whitney [73] imply that any surface map f can be continuously deformed to a generic surface map arbitrarily close to f .

A (topological) *immersion* is a generic surface map with no branch points. An *embedding* is a surface map with no singular points at all. Every immersion f is a local embedding; that is, every point $x \in S$ lies in an open neighborhood $U \subset S$ such that the restriction of f to U is an embedding.

2.3 Cube Complexes and Triangulations

As mentioned earlier, a *hexahedron* or *cube* is a convex polyhedron in \mathbb{R}^3 with eight vertices and six quadrilateral facets, combinatorially isomorphic to the standard cube $[0, 1]^3$. A *geometric cube complex* is a finite set of hexahedra, in which any the intersection of any two hexahedra is either a common facet, a common edge, a common vertex, or the empty set. The *underlying space* of a geometric cube complex X is the union of its hexahedra; we also call X a *geometric hex mesh* of its underlying space.

Formally, a *topological cube* is (the image of) a continuous injective function $q: [0, 1]^3 \rightarrow \mathbb{R}^3$. A *facet* of q is (the image of) a function $f: [0, 1]^2 \rightarrow \mathbb{R}^3$ obtained by restricting q to one of the facets of $[0, 1]^3$ and ignoring the fixed coordinate; edges and vertices of topological cubes are defined similarly. A *topological cube complex* is a finite set X of topological cubes in \mathbb{R}^3 such that the intersection of any two cubes is either a facet of both, an edge of both, a vertex of both, or the empty set. A topological cube complex X also called a *topological hex mesh* of its underlying space (the union of its constituent cubes).

Geometric and topological *quad meshes* are defined similarly, respectively using convex quadrilaterals or continuous injective maps from the square $[0, 1]^2$ instead of cubes; we colloquially refer to the constituent quadrilaterals as *quads*. A *boundary*

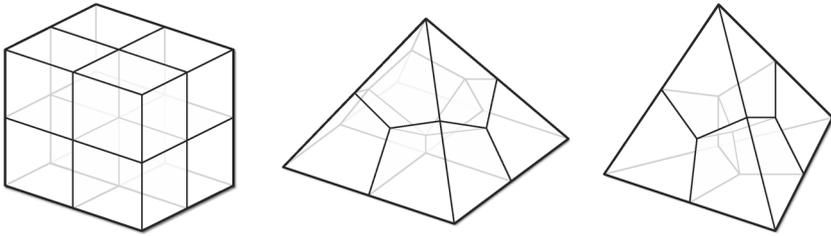


Fig. 5 Dual subdivision of a cube, a pyramid, and a tetrahedron

facet of a geometric cube complex X is a facet of exactly one cube in X ; the *boundary* ∂X of a (geometric or topological) hex mesh X is the (geometric or topological) quad mesh composed of all boundary facets of X . Geometric and topological *triangulations* are also defined similarly, with triangles or tetrahedra in place of quadrilaterals or cubes. Cube complexes and triangulations are examples of *polyhedral cell complexes*.

Some hexahedral meshing papers consider looser definitions of hexahedra and meshes, allowing, for example, inverted or twisted hexahedra, pairs of hexahedra that intersect in more than one common face, or hexahedra that are incident to the same face multiple times. However, we prefer the stricter definitions, in part to be consistent with geometric cube complexes, and in part because the looser definitions allow meshes that are useless in practice.

Our proofs implicitly rely on classical theorems of Moise [46,47] and Bing [11,12], which state that every 3-manifold with boundary is the underlying space of some topological triangulation, and therefore of some topological cube complex. Thus, although we rely on techniques from piecewise-linear topology, especially in Sect. 4, we need not assume that our input domain Ω is piecewise-linear or otherwise “tame”. Generalizations of our results to higher dimensions would require an explicit tameness assumption.

2.4 Dual Complexes

Any hex mesh X in \mathbb{R}^3 defines a natural *dual complex* X^* of the same underlying space, which can be constructed as follows. First subdivide each cube in X into eight smaller cubes by bisecting along each axis; the resulting subdivision X^{\boxplus} is a cube complex with the same underlying space as X . Then merge all subcubes in X^{\boxplus} incident to each vertex of X into a single (topological) polyhedron. The dual complex X^* has a 3-dimensional cell for every interior vertex of X , a 2-dimensional cell for every interior edge and boundary vertex of X , an edge for every interior facet and boundary edge of X , and a vertex for every cube and boundary facet of X ; see Fig. 5.

For any hex mesh X , the union of the interior 2-dimensional cells of the dual complex X^* is the image of a topological surface immersion. At the risk of confusing the reader, we use the same notation X^* to denote this dual immersion, which is variously called the *spatial twist continuum* [51], the *derivative*

complex of X [3,33,39,64], and the *canonical surface* of X [1,2]. The duality between cube complexes and immersed surfaces was already observed in the late 1800s, at least in preliminary form, in Fedorov's seminal study of zonotopes [29,30,65].

Similarly, the dual of any surface quad mesh Q is a cellular decomposition Q^* with a vertex for each facet of Q , an edge for each edge of Q , and a face for each vertex of Q . The vertices and edges of Q^* are the image of an immersion of one or more circles into the surface. In particular, if Q is the boundary of a hex mesh X , the dual curve immersion Q^* is the boundary of the dual surface immersion X^* .

We define the duals of more general cell complexes similarly. The dual subdivision of a polyhedron P has new vertices at the centroid of P , at the centroid of each facet of P , and at the midpoint of each edge of P , and these new vertices are connected with quadrilateral facets. For example, the dual subdivision of a tetrahedron consists of four cubes, and the dual subdivision of a pyramid consists of four cubes and an octagonal spindle; see Fig. 5. Subdividing every cell in a polyhedral complex X in this fashion gives us a subdivided complex X^{\boxplus} ; finally, merging all subcells in X^{\boxplus} that share a common vertex in X gives us the dual complex X^* .

2.5 Homology

Homology is an equivalence relation between cycles of various dimensions with a rich and useful algebraic structure. Our results specifically rely on cellular (or singular) homology with \mathbb{Z}_2 -coefficients. We give here only a brief overview of the most important definitions.

For the sake of concreteness, fix a cube complex X with underlying space Ω . For any integer k , a k -chain is a subset of the k -dimensional cubes in X . The *boundary* of a k -chain C , denoted ∂C , is the set of $(k - 1)$ -cubes in X that are facets of an odd number of k -cubes in C . A k -chain is called a k -cycle if its boundary is empty. In particular, a 1-cycle is a subgraph of the 1-skeleton of X in which every vertex has even degree. A k -cycle is *null-homologous* if it is the boundary of a $(k + 1)$ -chain, and two k -cycles are *homologous*, or in the same *homology class*, if their symmetric difference is null-homologous. The homology classes of k -cycles define an abelian group $H_k(X)$, called the k th *homology group* of X . Homology groups are finite-dimensional vector spaces over the finite field \mathbb{Z}_2 .

Now fix a subcomplex A of X . A k -chain in X is a *relative k -cycle* if its boundary is a $(k - 1)$ -chain in A , and two relative k -cycles are in the same *relative homology class* if their symmetric difference is homologous (in X) to a k -chain in A . The relative homology classes of relative k -cycles define the *relative homology group* $H_k(X, A)$. Relative homology groups are also finite-dimensional vector spaces over the finite field \mathbb{Z}_2 .

Homology and relative homology are topological invariants of the underlying space Ω , independent of the complex X . Thus, for example, we speak of graphs on the surface $\partial\Omega$ being null-homologous in Ω without specifying any particular cell decomposition of Ω or its boundary. This independence can be formalized using *singular* homology; we refer the interested reader to Hatcher [37] for further details.

3 No Odd Bounding Cycles

For the remainder of the paper, we fix a compact connected domain $\Omega \subset \mathbb{R}^3$ whose boundary is a 2-manifold, possibly with multiple components, and a topological quadrilateral mesh Q of $\partial\Omega$ with an even number of facets.

In this section, we prove that conditions (2) and (3) in our Main Theorem are equivalent. Our proof relies on standard tools from algebraic topology, most notably Poincaré–Alexander–Lefschetz duality, which are explained in detail by Edelsbrunner and Harer [26, Chapters IV and V] and Hatcher [37].

Lemma 1 *The dual graph Q^* is null-homologous in Ω if and only if every subgraph of Q that is null-homologous in Ω has an even number of edges.*

Proof Following Dey et al. [23–25], we call a subgraph of Q that is null-homologous in Ω a *handle cycle*. Let $H = \{\eta_1, \eta_2, \dots, \eta_g\}$ be a set of handle cycles in $\partial\Omega$ whose homology classes form a basis for the homology group $H_1(S^3 \setminus \Omega)$. We call H a *handle basis* for Q . Each handle cycle η_i is the boundary of a 2-chain σ_i in Ω , and the relative homology classes of $\sigma_1, \sigma_2, \dots, \sigma_g$ form a basis for the relative homology group $H_2(\Omega, \partial\Omega)$. Lefschetz duality implies that the homology class of any cycle γ in Ω is determined by the 2-chains σ_i that cross γ an odd number of times. In particular, a cycle γ is null-homologous in Ω if and only if each 2-chain σ_i crosses γ an even number of times. It follows that any cycle in $\partial\Omega$ is a handle cycle if and only if it crosses each handle cycle η_i an even number of times.

Suppose every handle cycle in Q has an even number of edges. Then in particular, every subgraph η_i in the handle basis H has an even number of edges. By definition, the dual graph Q^* crosses each edge of Q exactly once, so Q^* crosses every subgraph in H an even number of times. We conclude that Q^* is null-homologous in Ω .

On the other hand, suppose Q^* is null-homologous in Ω . Then every subgraph $\eta_i \in H$ crosses Q^* an even number of times and thus has an even number of edges. Every handle cycle in Q is the symmetric difference of a subset of subgraphs in H and a subset of facet boundaries (each with four edges). Thus, every handle cycle in Q has an even number of edges. \square

A similar—in fact simpler—argument implies that a quadrilateral mesh Q of a connected surface Σ is bipartite if and only if its dual graph Q^* is null-homologous on the surface Σ .

Recent work of Dey, Fan, and Wang [23] implies an efficient algorithm to test whether a *connected polyhedral* quad mesh Q satisfies the conditions of Lemma 1. Dey et al. describe an algorithm to compute a handle basis H for any polyhedron with connected boundary, genus g , and complexity n in $O(g^2n^2)$ time. The handle cycles in $\eta_i \in H$ may pass through the interiors of facets of Q , but we can easily deform them onto the 1-skeleton of Q in $O(gn)$ time if necessary. Q satisfies the conditions of Lemma 1 if and only if each basis cycle η_i has an even number of edges. Euler’s formula implies that the genus of Q is at most half the number of facets, so the running time of this algorithm is $O(n^4)$.

Theorem 2 *Given a connected polyhedron Q in \mathbb{R}^3 with genus g and n quadrilateral facets, we can determine whether Q^* is null-homologous in the interior of Q in $O(g^2n^2) = O(n^4)$ time.*

4 An Inefficient Proof

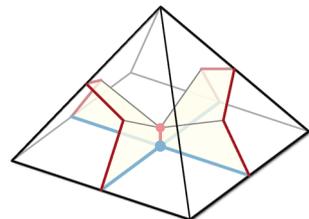
To complete the proof of our Main Theorem, it remains only to show that if Q^* is null-homologous in Ω , then Q is the boundary of a topological hex mesh of Ω . In this section we give a proof of this result in the spirit of Thurston's original argument for quad meshes of spheres [68]. Our proof is constructive, and it can be translated without much effort into an algorithm for constructing a topological hex mesh if one exists. However, in light of the more efficient algorithm in Sect. 5, we omit some implementation details and we do not attempt to analyze or reduce the complexity of the resulting mesh.

Lemma 3 *Suppose Q^* is null-homologous in Ω . Then Q^* is the boundary of a surface immersion into Ω .*

Proof Because Q^* is null-homologous in Ω , classical results of Whitney [73] and Papakyriakopoulos [56] imply that there is a generic surface map $f: S \rightarrow \Omega$ such that $f(\partial S) = Q^*$. In fact, we can construct a suitable generic surface map as follows. First, decompose Ω into a cell complex X that contains a pyramid over each facet of Q and otherwise contains only simplices, and let X^* be the dual complex of X . Let Σ be a 2-chain in X^* whose boundary is Q^* ; such a 2-chain must exist because Q^* is null-homologous in Ω . Without loss of generality, we can assume that $\Sigma \cap \partial\Omega = Q^*$. Each pyramid in X intersects exactly six facets of Σ , including all four facets touching the base of the pyramid, and each tetrahedron in X intersects either zero, three, or four facets of Σ ; see Figs. 6 and 7. It follows that Σ is the image of a generic surface map $f: S \rightarrow \Omega$ with no triple points and one branch point per vertex of Q^* .

We transform this generic surface map into an immersion following an argument of Funar [35], which is based in turn on results of Hass and Hughes [36] for immersions of surfaces without boundary. The singular points of f are the image of an immersion of a finite number of paths and cycles, where each path endpoint is either a vertex of Q^* or a branch point of f . (In particular, in our construction, the singular points consist of disjoint simple paths from vertices of Q^* to centroids of pyramids in X .) Because Q^* has an even number of vertices, f has an even number of branch points. We can transform f into a generic immersion (of a different surface) by canceling all

Fig. 6 Intersection pattern of Σ with a boundary pyramid of X



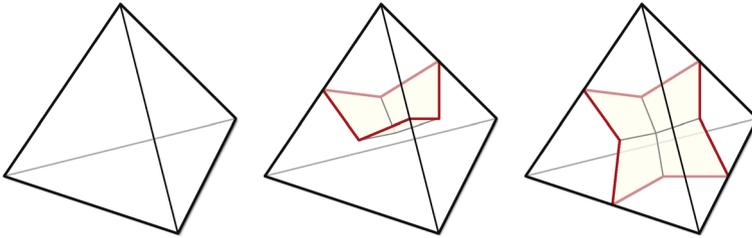
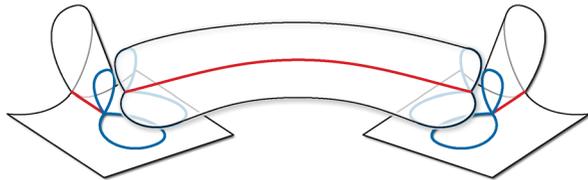


Fig. 7 Possible intersection patterns of Σ with interior tetrahedra of X

Fig. 8 Surgery to cancel two branch points



the branch points in pairs. To cancel any pair of branch points, we can delete small neighborhoods of both points and paste in a cylinder over a figure-8, which intersects the rest of the image of the surface transversely; see Fig. 8. □

Lemma 4 *Any generic surface immersion in Ω whose boundary is Q^* can be refined to the dual of a hex mesh whose boundary is Q .*

Proof We essentially follow the proof strategy suggested by Mitchell [44]. Fix a generic immersion $\phi: \Sigma \rightarrow \Omega$ such that $\phi(\partial \Sigma) = Q^*$. We extend ϕ to the dual of a hex mesh by adding a finite number of new surface components, each embedded transversely to each other and to the image of ϕ .

First we add a *buffer surface*, parallel to and just inside $\partial \Omega$, that separates every triple point of ϕ from the boundary of Ω . This buffer surface ensures that later modifications do not change the boundary curves Q^* . Let Ω° denote the portion of Ω inside the buffer surface.

Let T° be a topological triangulation of Ω° that contains (a triangulation of) $\text{im } \phi \cap \Omega^\circ$ in its 2-skeleton; such a triangulation always exists. To finish the construction, we “bubble-wrap” ϕ by inserting spherical bubbles around each simplex in T° , following an algorithm of Babson and Chan [3]. First add disjoint spheres around each vertex of the triangulation; then add disjoint spheres around the portion of each edge outside the vertex bubbles; then add disjoint spheres around the subset of each triangle outside the vertex and edge bubbles; and finally add disjoint spheres around the subset of each tetrahedron outside the vertex, edge, and triangle bubbles. See Fig. 9 for a two-dimensional example.

Straightforward case analysis implies that the bubble-wrapped surface immersion is dual to a topological hex mesh whose boundary is Q . See Babson and Chan [3], Bern and Eppstein [8,9], Schwartz [63], and Schwartz and Ziegler [64] for similar analysis; we omit further details here. □

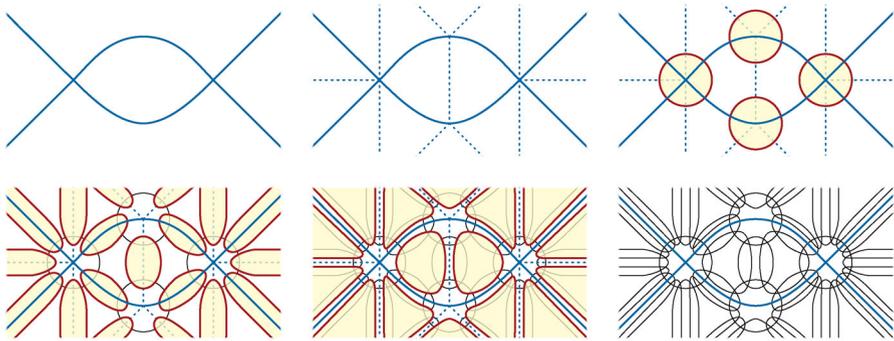


Fig. 9 Bubble-wrapping a curve immersion, after Babson and Chan [3]. *Top row* the original curves, a topological triangulation, bubble-wrapping vertices. *Bottom row* bubble-wrapping edges, bubble-wrapping triangles, the final immersion

As part of his proof for genus-zero meshes, Mitchell [44] describes an algorithm to transform a surface immersion into the dual of a hex mesh that introduces fewer new surfaces, and therefore produces smaller meshes, than Babson and Chan’s algorithm [3]. Unfortunately, his algorithm appears to have a subtle flaw. Specifically, Mitchell’s algorithm correctly computes a surface immersion X^* that is dual to a cube complex X , but the components of $\Omega \setminus X^*$ are not necessarily topological balls, which implies that X is not necessarily a mesh of *any* 3-manifold, including Ω . Several later papers make similar omissions [31, 42, 51, 66, 67].

5 An Efficient Proof

Instead of fleshing out the details of the previous argument, we give in this section a second constructive proof, which translates directly into an efficient algorithm to construct a topological hex mesh for any polyhedron in \mathbb{R}^3 with quadrilateral facets. This algorithm does not rely on the results in Sects. 3 and 4.

Our high-level strategy closely resembles Eppstein’s algorithm for bipartite quad meshes [27]. After some mild preprocessing, we first separate the boundary of Ω from its interior with a buffer layer B of cubes joining the input quad mesh Q to a parallel copy of Q just inside the boundary. We then compute a triangulation T of the inner domain $\Omega \setminus B$ that splits each inner boundary quad in B into two triangles. As in Eppstein’s algorithm, our final hex mesh is a refinement of the convex decomposition $B \cup T$, obtained by splitting each buffer cube in B and each tetrahedron in T into a bounded number of smaller cubes. However, our refinement strategy is different from Eppstein’s.

5.1 Connecting Odd Components

Recall that the input mesh Q has an even number of facets, and that the dual graph Q^* is null-homologous in Ω . For domains with disconnected boundaries, our meshing strategy requires a slightly stronger condition: Each *component* of Q must have an even

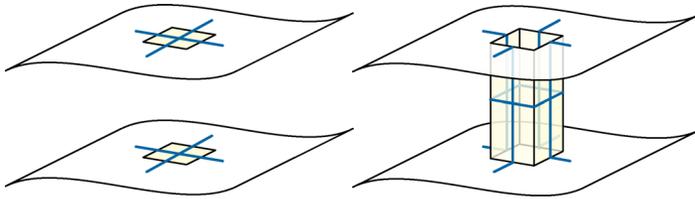


Fig. 10 Connecting two components of Q with a cubical tunnel

number of facets. Fortunately, this stronger condition is easy to enforce by connecting odd components of Q in pairs as follows.

Fix two odd components Q_1 and Q_2 of Q , and let q_1 and q_2 be arbitrary quads in those components. Let κ be a cube in Ω with opposite facets q_1 and q_2 ; such a cube must exist because Ω is connected. Let Ω' be the closure of $\Omega \setminus \kappa$, and let Q' be the quad mesh of $\partial\Omega'$ obtained from Q by deleting q_1 and q_2 and adding the other four facets of κ , as shown in Fig. 10. Given any hex mesh of Ω' whose boundary is Q' , inserting the cube κ yields a hex mesh of Ω whose boundary is Q .

Lemma 5 $(Q')^*$ is null-homologous in Ω' .

Proof Let $\partial\kappa^*$ denote the dual graph of the 1-skeleton of κ ; this graph consists of three cycles on the boundary of κ . The graph $(Q')^*$ is the symmetric difference of Q^* and $\partial\kappa^*$, which are both null-homologous in Ω . It follows that $(Q')^*$ is null-homologous in Ω .

As in the proof of Lemma 3, decompose Ω' into a cell complex X' that contains a pyramid over each facet of Q' and otherwise contains only simplices, and let $X = X' \cup \kappa$ be the corresponding cellular decomposition of Ω . Let Σ be a 2-chain in the dual complex X^* such that $\partial\Sigma = (Q')^*$; without loss of generality, we can assume that $\Sigma \cap \partial\Omega' = (Q')^*$. If Σ is a 2-chain in X' , the proof is complete.

Otherwise, $\Sigma \cap \kappa$ must be a disk δ that separates q_1 and q_2 . Let κ_1 be the component of $\kappa \setminus \delta$ with q_1 on its boundary, and let Σ' be the symmetric difference $\Sigma \oplus \partial\kappa_1 \oplus Q_1$. Then Σ' is a 2-chain in Ω' whose boundary is $(Q')^*$, so again the proof is complete. \square

More generally, let $\kappa_1, \kappa_2, \dots, \kappa_r$ be disjoint cubes in Ω , such that each cube κ_i has two opposite facets in Q and each odd component of Q contains exactly one facet of one cube κ_i . Let Ω' be the closure of $\Omega \setminus (\kappa_1 \cup \dots \cup \kappa_r)$ and let $Q' = Q \oplus \partial\kappa_1 \oplus \dots \oplus \partial\kappa_r$. Every component of Q' has an even number of facets, and Lemma 5 implies inductively that $(Q')^*$ is null-homologous in Ω' . Given any hex mesh of Ω' compatible with Q' , inserting the cubes $\kappa_1, \dots, \kappa_r$ yields a hex mesh of Ω compatible with Q .

Thus, for the remainder of this section, we assume without loss of generality that every component of the input mesh Q has an even number of facets.

5.2 Refining the Interior Triangulation

Let B be a buffer layers of cubes separating the boundary of Ω from its interior, obtained by joining the input quad mesh Q to a parallel copy of Q just inside $\partial\Omega$.

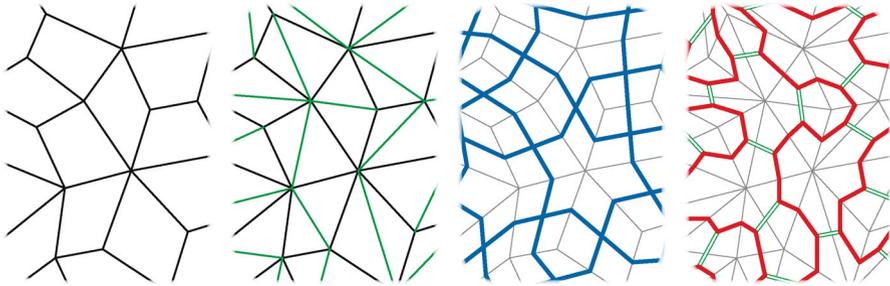


Fig. 11 A portion of the quad mesh Q , its triangulation ∂T , the dual curves Q^* , and the homologous curves $\Gamma \subset \partial T^*$

Then let T be a triangulation of the inner domain $\Omega \setminus B$ that splits each inner boundary quad in B into two triangles. Let ∂T denote the induced triangulation of the inner surface of B .

The dual complex T^* can be constructed by first splitting each tetrahedron in T into four cubes meeting at that tetrahedron's centroid, and then merging all subsets of cubes incident to each vertex of T , all quadrilaterals incident to each edge of T , and all pairs of segments meeting at a face of T . The restriction of T^* to the inner surface of B is the usual combinatorial dual of the surface triangulation ∂T ; that is, we have $(\partial T)^* = \partial(T^*)$.

Let Δ denote the diagonals used to refine Q into ∂T , and let Δ^* denote the corresponding edges of the dual graph ∂T^* . Finally, let Γ denote the subgraph $\partial T^* \setminus \Delta^*$; see Fig. 11. Every vertex in Γ has degree 2, which implies that Γ is a collection of disjoint simple cycles. Contracting all the edges in Δ^* transforms Γ into the dual complex Q^* of the original quad mesh Q . Thus, Γ is *homotopic* to a covering of Q^* by edge-disjoint circuits, which implies that Γ and Q^* are homologous. In particular, Γ is null-homologous in Ω .

Now let Σ be any 2-chain in T^* such that $\partial \Sigma = \Gamma$; if no such 2-chain exists, then Q^* is not null-homologous in Ω . We easily observe that Σ is the union of disjoint embedded quadrangulated surfaces. In particular, each interior vertex of Σ is incident to either three or four quadrangular facets of T^* , and the intersection of Σ with any tetrahedron in T is either empty or a disk. See Fig. 7.

Next, we refine T into a hex mesh Y by splitting each tetrahedron into either four, seven, or eight cubes, depending on whether the tetrahedron intersects zero, three, or four facets of Σ , as shown in Fig. 12. Equivalently, we partition each tetrahedron in T into four cubes by central subdivision, and then expand the surface Σ into a layer of cubes.

5.3 Refining the Buffer Cubes

It remains only to refine the buffer cubes in B to conform to the boundary of the refined triangulation Y . Each buffer cube has an *outer* facet in Q , an *inner* facet on the boundary of $\Omega \setminus B$, and four *transition* facets. The interior mesh Y subdivides

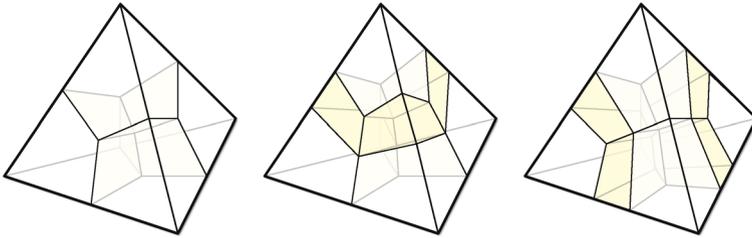
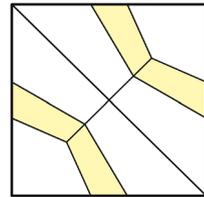


Fig. 12 Templates for refining tetrahedra into cubes; compare with Fig. 7

Fig. 13 Refinement of the inner facet of each buffer cube



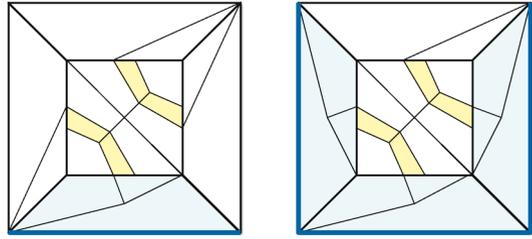
the inner facet of each buffer cube into ten quadrilaterals, as shown in Fig. 13, and each edge of that inner facet into three segments. Thus, the transition facets of B are combinatorially hexagons.

Following Eppstein [27, Lemma 1], we find a subgraph R of Q such that every facet of Q is incident to either one or three edges of R . We can compute such a subgraph in polynomial time by computing a perfect matching in the shortest-path metric on Q^* . (Here we require every component of Q to have an even number of facets.) Each edge in this matching is a path in Q^* . Let R^* be the subset of edges of Q^* that appear in an odd number of these paths. Because vertex is an endpoint of exactly one path, each vertex of Q^* lies on an odd number of edges in R^* . Finally, let R be the subgraph of Q dual to R^* .

Alternatively, we can apply a classical result of Berge [5,6, Chapter 18, Theorem 7], which states that every $(k - 1)$ -edge-connected k -regular graph with an even number of vertices contains a perfect matching. Each component of Q^* is 3-edge-connected, because every separating cycle (in fact, every cycle) in Q has length at least 3; it follows immediately that Q^* contains a perfect matching. (For a more recent independent proof, see Carbonera and Shepherd [15,16].) Let R be the subgraph of Q dual to this matching.

We then subdivide each transition facet of B into either two or three quadrilaterals, depending on whether that facet is bounded by an edge of R or not. Because each facet of Q is incident to either one or three edges of R , the boundary of each buffer cube is refined into either 20 or 22 quads, as shown in Fig. 14. To complete our construction, we refine each boundary cube into a hex mesh compatible with its boundary subdivision. The existence of such a hex mesh is guaranteed by Thurston and Mitchell’s original proof; alternatively, as in Eppstein’s proof [27], it is straightforward but tedious to construct topological hex meshes for these subdivided cubes by hand.

Fig. 14 Boundary refinement of the buffer cubes; *bold edges* are in the odd subgraph R



5.4 Analysis

We analyze both the running time of our algorithm and the complexity of the output hex mesh in terms of two parameters n and t , which respectively denote the number of quadrilateral facets in the input mesh Q and the number of tetrahedra in the interior triangulation T . We trivially have $t = \Omega(n)$, but without further assumptions, t cannot be upper-bounded by *any* function of n . For example, suppose Ω is homeomorphic to the complement of a knot K . The boundary of Ω is a torus, which can be decomposed into a 4×4 grid of quadrilaterals (to guarantee the existence of a compatible hex mesh). The hyperbolic volume of Ω is a lower bound on the complexity of any triangulation of Ω , and there are knots K whose complements have arbitrarily large hyperbolic volume [41, 59].

Our algorithm constructs a topological hex mesh with complexity $O(n+t) = O(t)$, which is optimal up to constant factors; any hex mesh can be refined to a triangulation by decomposing each cube into six tetrahedra.

Lemma 6 *Suppose Q is the boundary of a topological hex mesh of Ω . The minimum number of cubes in a hex mesh of Ω whose boundary is Q is within a constant factor of the minimum number of tetrahedra in a topological triangulation of Ω whose boundary splits each facet of Q into two triangles.*

Moreover, if we are given an interior triangulation T , we can compute a hex mesh with complexity $O(t)$, or determine correctly that no such mesh exists, in $O(t^3)$ time. Only two stages of the algorithm have nontrivial running time: computing the 2-chain Σ and computing the subgraph R .

We can compute the 2-chain Σ by solving a $O(t) \times O(t)$ system of linear equations in \mathbb{Z}_2 , with a variable for each 2-cell in T^* indicating whether that 2-cell is or is not in Σ , and an equation for each edge e in T^* indicating whether the number of 2-cells incident to e that lie in Σ is even (if $e \notin \Gamma$) or odd (if $e \in \Gamma$). We can solve this system of linear equations in $O(t^3)$ time using Gaussian elimination; this time bound can be improved using fast matrix multiplication [14, 38]. If this system of equations has no solution, then Γ is not null-homologous in Ω , which implies that there is no hex mesh compatible with Q .

To compute the odd subgraph R , it suffices to arbitrarily pair up the vertices of Q , compute the shortest path connecting each pair, and define R be the subset of edges in an odd number of these shortest paths. This construction requires $O(n^3) = O(t^3)$ time. Alternatively, we can find a perfect matching in Q^* in $O(n^{3/2}) = O(t^{3/2})$ time

using the classical algorithm of Micali and Vazirani [43, 70, 71], but this improvement is dominated by the time to compute Σ .

Theorem 7 *Let Ω be a compact connected subset of \mathbb{R}^3 whose boundary $\partial\Omega$ is a (possibly disconnected) 2-manifold. Suppose we are given a topological quad mesh Q of $\partial\Omega$ with n facets and a topological triangulation T of Ω with complexity t , such that T splits each facet of Q into two triangles. Then we can either compute a topological hex mesh of Ω whose boundary is Q , or report correctly that no such mesh exists, in $O(t^3)$ time. Moreover, if we return a hex mesh, its complexity is $O(t)$.*

We can modify our algorithm to return a null-homologous subgraph of Q with an odd number of edges when Q is not compatible with a hex mesh, in $O(t^3)$ additional time, as follows. If Q^* is not null-homologous, then Lemma 1 implies that every handle basis in Q contains at least one handle cycle of odd length. We can compute a handle basis in $O(t^3)$ time using standard homology algorithms [24, 25] or in $O(g^2n^2)$ time using the recent algorithm of Dey et al. [23] if Q is a connected polyhedron. (Again, the $O(t^3)$ time bound can be improved using fast matrix multiplication.)

When the facets of Q are actually planar convex quadrilaterals, results of Chazelle and Palios [20], Bern [7], and Chazelle and Shouraboura [21] imply that we can compute a triangulation with complexity $t = O(n^2)$ in $O(n^2 \log n)$ time.

Theorem 8 *Given a polyhedron Q in \mathbb{R}^3 with n quadrilateral facets, we can either compute a topological hex mesh of the interior of Q whose boundary is Q , or report correctly that no such mesh exists, in $O(n^6)$ time. Moreover, if we return a hex mesh, its complexity is $O(n^2)$.*

The upper bound $t = O(n^2)$ is tight for *geometric* triangulations of polyhedra, even with genus zero [19], but it appears to be an open question whether this bound can be improved for *topological* triangulations. A heuristic argument of Thurston and Thurston suggests that the lower bound $t = \Omega(n^{3/2})$ is plausible [69], but the only lower bound actually known is the trivial $t = \Omega(n)$.

If we only want to determine whether a compatible hex mesh exists, we can use the faster algorithm from Theorem 2.

Corollary 9 *Given a connected polyhedron Q in \mathbb{R}^3 with genus g and n quadrilateral facets, we can either report correctly that there is a topological hex mesh of the interior of Q whose boundary is Q , or compute a handle cycle in Q with odd length, in $O(g^2n^2) = O(n^4)$ time.*

6 Geometric Hex Meshing?

We conclude by describing some implications of our results for the construction of *geometric* hex meshes for domains with complex topology.

Bern and Eppstein [8, 9] reduce the geometric hex meshing problem from arbitrary genus-zero polyhedra to a specific family of polyhedra called *bicuboids*. A bicuboid is a convex polyhedron with ten quadrilateral facets, combinatorially isomorphic to the boundary of two cubes joined along a common facet; see the middle of Fig. 1.

Although Bern and Eppstein explicitly claim the following result only for connected genus-zero quad meshes, their argument applies verbatim for geometric quad meshes with arbitrary topology.

Lemma 10 (Bern and Eppstein [8, Theorem 2]) *If every bicuboid has a geometric hex mesh, then any polyhedron in \mathbb{R}^3 that is the boundary of a topological hex mesh is also the boundary of a geometric hex mesh.*

Corollary 11 *Let Q be a polyhedron in \mathbb{R}^3 with an even number of quadrilateral facets, whose dual graph Q^* is null-homologous in the interior. If every bicuboid is the boundary of a geometric hex mesh, then Q is the boundary of a geometric hex mesh.*

Bern and Eppstein's result does not imply any bounds on the *complexity* of geometric hex meshes. Eppstein [27, Sect. 5] sketches a different reduction of the genus-zero geometric hex-meshing problem to a larger finite set of polyhedra, each obtained by refining the boundary of a cube into at most 30 quads. If each of those boundary-refined cubes has a geometric hex mesh *with bounded complexity*, then Eppstein's reduction implies that any convex polytope with $2n$ quadrilateral facets has a geometric hex mesh with complexity $O(n)$.

Our algorithm supports a similar reduction to a more complex set of refined cubes, each with at most 40 facets. Following Eppstein, we compute a slightly more complex buffer layer B , containing one flat beveled cube incident to each facet of Q , two long skinny cubes incident to each edge of Q , and $2 \deg(v) - 2$ cubes incident to each vertex v of Q . We then compute a triangulation T of the inner volume $\Omega \setminus B$, which includes a triangulation ∂T of the inner buffer facets. The inner boundary of this more complex buffer has $O(n)$ facets, so this triangulation has complexity $O(n^2)$. Following our earlier algorithm, we find a set Γ of simple cycles in ∂T^* that is homologous with Q^* , compute a 2-chain Σ in T^* whose boundary is Γ , and then refine each tetrahedron in T into cubes using the templates in Fig. 12. This refinement splits each inner boundary facet of each buffer cube into at most 10 quads and introduces two new vertices to the boundary of each transition facet. Using Eppstein's matching technique, as in our earlier algorithm, we can subdivide each transition facet into a small number of quads, so that buffer cube has an even number of boundary quads. We omit further details, which are straightforward modifications of Eppstein's and our earlier arguments.

Theorem 12 *Let Q be a polyhedron in \mathbb{R}^3 with $2n$ quadrilateral facets, whose dual graph Q^* is null-homologous in the interior. If every subdivision of the boundary of a cube into k convex quadrilaterals, for any even integer $k \leq 40$, is the boundary of a geometric hex mesh with bounded complexity, then Q is the boundary of a geometric hex mesh with complexity $O(n^2)$.*

However, we conjecture that if these subdivided cubes (and the other polyhedra in Fig. 1) support geometric hex meshes at all, their complexity depends on the precise geometry of the polyhedra and can be arbitrarily large. We leave the proof or disproof of this conjecture as an interesting open problem.

References

1. Aitchison, I.R., Rubinstein, J.H.: An introduction to polyhedral metrics of non-positive curvature on 3-manifolds. In: Donaldson, S.K., Thomas, C.B. (eds.) *Geometry of Low-Dimensional Manifolds: 2: Symplectic Manifolds and Jones–Witten Theory*. London Mathematical Society Lecture Notes, vol. 151, pp. 127–161. Cambridge University Press, Cambridge (1990)
2. Aitchison, I.R., Matsumoto, S., Rubinstein, J.H.: Immersed surfaces in cubed manifolds. *Asian J. Math.* **1**(1), 85–95 (1997)
3. Babson, E.K., Chan, C.S.: Counting faces of cubical spheres modulo two. *Discret. Math.* **212**(3), 169–183 (2000)
4. Banchoff, T.F.: Triple points and surgery of immersed surfaces. *Proc. Am. Math. Soc.* **46**(3), 407–413 (1974)
5. Berge, C.: *Théorie des graphes et ses applications*. Dunod, Paris (1958). English translation in [6]
6. Berge, C.: *The Theory of Graphs*. Dover Publications, New York (2001). English translation of [5]
7. Bern, M.: Compatible triangulations. In: *Proceedings of the 9th Annual Symposium on Computational Geometry*, pp. 281–288 (1993)
8. Bern, M., Eppstein, D.: Flipping cubical meshes. In: *Proceedings of the 10th International Meshing Roundtable*, pp. 19–29 (2001). <http://www.andrew.cmu.edu/user/sowen/abstracts/Be812.html>. Preliminary version of [9]
9. Bern, M., Eppstein, D., Erickson, J.: Flipping cubical meshes. *Eng. Comput.* **18**, 173–187 (2002). Full version of [8]
10. Billera, L.J., Sturmfels, B.: Fiber polytopes. *Ann. Math.* **135**(3), 527–549 (1992)
11. Bing, R.H.: Locally tame sets are tame. *Ann. Math.* **59**, 145–158 (1954)
12. Bing, R.H.: An alternative proof that 3-manifolds can be triangulated. *Ann. Math.* **69**, 37–65 (1959)
13. Boy, W.: Über die Abbildung der projektiven Ebene auf eine im Endlichen geschlossene singularitätenfreie Fläche. *Nachr. Ges. Wiss. Göttingen*, pp. 20–33 (1901)
14. Bunch, J.R., Hopcroft, J.E.: Triangular factorization and inversion by fast matrix multiplication. *Math. Comput.* **28**(125), 231–236 (1974)
15. Carbonera, C.D., Shepherd, J.F.: On the existence of a perfect matching for 4-regular graphs derived from quadrilateral meshes. Technical Report UUSCI-2006-021. SCI Institute, University of Utah (2006)
16. Carbonera, C.D., Shepherd, J.F.: A constructive approach to constrained hexahedral mesh generation. *Eng. Comput.* **26**(4), 341–350 (2010)
17. Carter, J.S.: Extending immersions of curves to properly immersed surfaces. *Topol. Appl.* **40**(3), 287–306 (1991)
18. Carter, J.S.: Extending immersed circles in the sphere to immersed disks in the ball. *Commun. Math. Helv.* **67**, 337–348 (1992)
19. Chazelle, B.: Convex partitions of polyhedra: a lower bound and worst-case optimal algorithm. *SIAM J. Comput.* **13**(3), 488–507 (1984)
20. Chazelle, B., Palios, L.: Triangulating a nonconvex polytope. *Discrete Comput. Geom.* **5**, 505–526 (1990)
21. Chazelle, B., Shouraboura, N.: Bounds on the size of tetrahedralizations. *Discrete Comput. Geom.* **14**, 429–444 (1995)
22. Csikós, B., Szűcs, A.: On the number of triple points of an immersed surface with boundary. *Manuscr. Math.* **87**, 285–293 (1995)
23. Dey, T.K., Fan, F., Wang, Y.: An efficient computation of handle and tunnel loops via Reeb graphs. *ACM Trans. Graph.* **32**(4) (2013). Proceedings of SIGGRAPH 2013.
24. Dey, T.K., Li, K., Sun, J.: On computing handle and tunnel loops. In: *IEEE Proceedings of International Conference on Cyberworlds*, pp. 357–366 (2007)
25. Dey, T.K., Li, K., Sun, J., Cohen-Steiner, D.: Computing geometry-aware handle and tunnel loops in 3D models. *ACM Trans. Graph.* **27**(3), 1–9 (2008). Proceedings of SIGGRAPH 2008
26. Edelsbrunner, H., Harer, J.L.: *Computational Topology: An Introduction*. American Mathematical Society, Providence, RI (2010)
27. Eppstein, D.: Linear-complexity hexahedral mesh generation. *Comput. Geom. Theory Appl.* **12**, 3–16 (1999)
28. Erickson, J.: Efficiently hex-meshing things with topology. In: *Proceedings of the 29th Annual Symposium on Computational Geometry*, pp. 37–46 (2013)

29. Fedorov, E.S.: Načala učeniya o figurah (Elements of the study of figures). Not. Imp. Petersb. Mineralog. Soc. **24**, 1–279 (1885) (in Russian). Cited in [65]
30. Fedorov, E.S.: Elemente der gestaltenlehre. Z. Kryst. Mineral. **21**, 679–694 (1893)
31. Folwell, N.T., Mitchell, S.A.: Reliable whisker weaving via curve contraction. Eng. Comput. **15**(3), 292–302 (1999)
32. Francis, G.K.: Titus' homotopies of normal curves. Proc. Am. Math. Soc. **30**, 511–518 (1971)
33. Funar, L.: Cubulations, immersions, mappability and a problem of Habegger. Ann. Sci. École Norm. Supér. **32**(5), 681–700 (1999)
34. Funar, L.: Cubulations mod bubble moves. In: Nencka, H. (ed.) Low Dimensional Topology. Contemporary Mathematics, vol. 223, pp. 29–43. American Mathematical Society, Providence, RI (1999)
35. Funar, L.: Surface cubications mod flips. Manuscr. Math. **125**, 285–307 (2008)
36. Hass, J., Hughes, J.: Immersions of surfaces in 3-manifolds. Topology **24**(1), 97–112 (1985)
37. Hatcher, A.: Algebraic Topology. Cambridge University Press, Cambridge, UK (2002). <http://www.math.cornell.edu/~hatcher/AT/ATpage.html>
38. Ibarra, O.H., Moran, S., Hui, R.: A generalization of the fast LUP matrix decomposition algorithm and applications. J. Algorithms **3**(1), 45–56 (1982)
39. Jockusch, W.: The lower and upper bound problems for cubical polytopes. Discrete Comput. Geom. **9**, 159–163 (1993)
40. Kirby, R.: Problems in low-dimensional topology. In: Kazez, W.H. (ed.) Geometric Topology (Part 2), AMS/IP Studies in Advanced Mathematics, vol. 2, pp. 35–473. American Mathematical Society, Providence, RI (1997). <http://math.berkeley.edu/~kirby/problems.ps.gz>
41. Lackenby, M.: The volume of hyperbolic alternating link complements. Proc. Lond. Math. Soc. **88**, 204–224 (2004)
42. Ledoux, F., Shepherd, J.F.: Topological modifications of hexahedral meshes via sheet operations: a theoretical study. Eng. Comput. **26**(4), 433–447 (2010)
43. Micali, S., Vazirani, V.V.: An $O(\sqrt{v} \cdot |E|)$ algorithm for finding maximum matching in general graphs. In: Proceedings of the 21st Annual Symposium on Foundations of Computer Science, pp. 17–27 (1980)
44. Mitchell, S.A.: A characterization of the quadrilateral meshes of a surface which admit a compatible hexahedral mesh of the enclosed volume. In: Proceedings of the 13th Annual Symposium on Theoretical Aspects of Computer Science. Lecture Notes in Computer Science, vol. 1046, pp. 456–476. Springer, Berlin (1996)
45. Mitchell, S.A.: The all-hex Geode-template for conforming a diced tetrahedral mesh to any diced hexahedral mesh. In: Proceedings of the 7th International Meshing Roundtable, pp. 295–305 (1998)
46. Moise, E.E.: Affine structures in 3-manifolds, V. The triangulation theorem and Hauptvermutung. Ann. Math. **56**, 96–114 (1952)
47. Moise, E.E.: Affine structures in 3-manifolds, VIII. Invariance of the knot-type; local tame embedding. Ann. Math. **59**, 159–170 (1954)
48. Müller-Hannemann, M.: Hexahedral meshing by successive dual cycle elimination. Eng. Comput. **15**(3), 269–279 (1999)
49. Müller-Hannemann, M.: Shelling hexahedral complexes for mesh generation. J. Graph Algorithms Appl. **5**(5), 59–91 (2001)
50. Müller-Hannemann, M.: Quadrilateral surface meshes without self-intersecting dual cycles for hexahedral mesh generation. Comput. Geom. Theory Appl. **22**, 75–97 (2002)
51. Murdoch, P., Benzley, S., Blacker, T., Mitchell, S.A.: The spatial twist continuum: a connectivity based method for representing all-hexahedral finite element meshes. Finite Elem. Anal. Des. **28**, 137–149 (1997)
52. Nakamoto, A.: Diagonal transformations and cycle parities of quadrangulations on surfaces. J. Comb. Theory Ser. B **67**, 202–211 (1996)
53. Nakamoto, A.: Diagonal transformations in quadrangulations of surfaces. J. Graph Theory **21**, 289–299 (1996)
54. Nakamoto, A., Ota, K.: Diagonal transformations in quadrangulations and Dehn twists preserving cycle parities. J. Comb. Theory Ser. B **69**, 125–141 (1997)
55. Nowik, T.: Complexity of planar and spherical curves. Duke J. Math. **148**(1), 107–118 (2009)
56. Papakyriakopoulos, C.D.: On Dehn's lemma and the asphericity of knots. Ann. Math. **66**, 1–26 (1957)
57. Petit, J.P.: Le topologicon, Les aventures d'Anselme Lanturlu, vol. 12. Belin, Paris (1985). <http://www.savoir-sans-frontieres.com/JPP/telechargeables/Francais/topologicon.htm>

58. Price, M.A., Armstrong, C.G.: Hexahedral mesh generation by medial surface subdivision: Part II. Solids with flat and concave edges. *Int. J. Numer. Methods Eng.* **40**(1), 111–136 (1997)
59. Purcell, J.S.: Volumes of highly twisted knots and links. *Algebr. Geom. Topol.* **7**, 93–108 (2007)
60. Reiner, V.: The generalized Baues problem. In: Billera, L.J., Björner, A., Greene, C., Simion, R.E., Stanley, R.P. (eds.) *New Perspectives in Geometric Combinatorics*, vol. 38, pp. 293–336. MSRI Publications, Cambridge (1999)
61. Schneiders, R.: Open problem (1995). <http://www-users.informatik.rwth-aachen.de/~roberts/open.html>. Accessed 1 Aug 2013
62. Schneiders, R.: A grid-based algorithm for the generation of hexahedral element meshes. *Eng. Comput.* **12**, 168–177 (1996)
63. Schwartz, A.: Constructions of cubical polytopes. Ph.D. dissertation, Technical University, Berlin (2003). http://edocs.tu-berlin.de/diss/2004/schwartz_alexander.html
64. Schwartz, A., Ziegler, G.M.: Construction techniques for cubical complexes, odd cubical 4-polytopes, and prescribed dual manifolds. *Exp. Math.* **13**(4), 385–413 (2004)
65. Senechal, M., Galiulin, R.V.: An introduction to the theory of figures: the geometry of E. S. Fedorov. *Struct. Topol.* **10**, 5–22 (1984)
66. Shepherd, J.F., Johnson, C.R.: Hexahedral mesh generation constraints. *Eng. Comput.* **24**(3), 195–213 (2008)
67. Suzuki, T., Takahashi, S., Shepherd, J.F.: An interior surface generation method for all-hexahedral meshing. *Eng. Comput.* **26**(3), 303–316 (2010)
68. Thurston, W.P.: Hexahedral decomposition of polyhedra. Posting to Sci. Math. (1993). <http://www.ics.uci.edu/~eppstein/gina/Thurston-hexahedra.html>
69. Thurston, B., Thurston, D.: Complexity of random knot with vertices on sphere. *MathOverflow* (2011). <http://mathoverflow.net/questions/54417>
70. Vazirani, V.V.: A theory of alternating paths and blossoms for proving correctness of the $O(\sqrt{VE})$ general graph maximum matching algorithm. *Combinatorica* **14**(1), 71–109 (1994)
71. Vazirani, V.V.: A proof of the MV matching algorithm (2014). <http://www.cc.gatech.edu/~vazirani/new-proof.pdf>. Unpublished manuscript
72. Whitney, H.: On regular closed curves in the plane. *Compos. Math.* **4**, 276–284 (1937)
73. Whitney, H.: The singularities of a smooth n -manifold in $(2n - 1)$ -space. *Ann. Math.* **45**(2), 247–293 (1944)
74. Yamakawa, S., Shimada, K.: HEXHOOP: modular templates for converting a hex-dominant mesh to an ALL-hex mesh. *Eng. Comput.* **18**, 211–228 (2002)
75. Yamakawa, S., Shimada, K.: 88-element solution to Schneiders’ pyramid hex-meshing problem. *Int. J. Numer. Methods. Biomed. Eng.* **26**, 1700–1712 (2010)