Efficiently Hex-Meshing Things with Topology*

Jeff Erickson
Department of Computer Science
University of Illinois, Urbana-Champaign

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.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

Keywords: computational topology, hexahedral mesh generation, homology

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*. 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 hex 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 hexahedral 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 Figure 1. (All existing hexahedral meshes for these examples [10, 50, 55, 56] include degenerate, inverted, and/or warped hexahedra.)

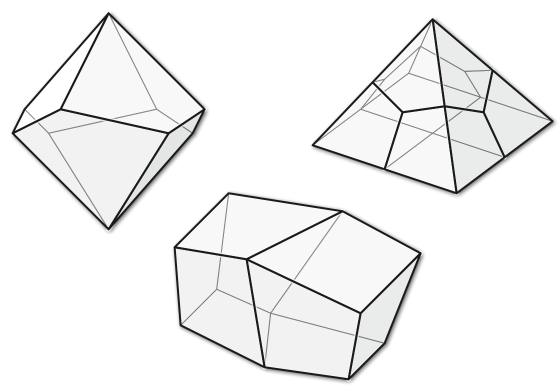

Figure 1. An octagonal spindle [18], a bicuboid [5, 6], and Schneiders' pyramid [44, 45].

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 are consistent with the intersection patterns in geometric hex meshes. (We describe the precise intersection constraints in Section 2.) For example, any bicuboid has an obvious decomposition into two *topological* hexahedra sharing a single face.

*Portions of this work were done while the author was visiting IST Austria. Work on this paper was also partially supported by NSF grant CCF 09-15519. See http://www.cs.uiuc.edu/~jeffe/pubs/hexmesh.html for the most recent version of this paper.

1.1 Previous Results

Genus-Zero Surface Meshes

Thurston [52] and Mitchell [32] independently proved that any quad mesh of a topological sphere can be extended to a topological hex mesh of the interior ball if and only if the number of quads is even. Thus, all the examples in Figure 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 quad 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* [22, 53]. Second, they argue that the resulting surface arrangement can be refined, by adding additional closed surfaces in the interior of the ball, to the dual complex of a hex mesh whose boundary is Q. (We describe dual curves of quad meshes and dual surfaces of hex meshes in more detail in Section 2.)

Mitchell's proof [32] 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 [18] showed that carelessly choosing the elementary moves in the first stage Mitchell's algorithm can lead to $\Omega(n^2)$ output complexity. More recent results of Nowik [40] imply the same $\Omega(n^2)$ lower bound even if all such choices are made optimally.

Eppstein [18] 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 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 quads and then invokes Mitchell's algorithm. As Eppstein observes [18], it is not difficult to construct meshes for these remaining cases by hand.

For other results related to topological hex meshes of balls, see Csikós and Szűcs [14], Funar [23, 24, 25], Bern *et al.* [5, 6], Babson and Chan [3], Schwartz [46], and Schwartz and Ziegler [47].

Higher-Genus Surface Meshes

Both Mitchell and Eppstein consider the problem of hex-meshing domains with nontrivial topology, but neither offers a complete solution. Mitchell [32] 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 an 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 $\mathbb{R}^3$; see Figure 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.

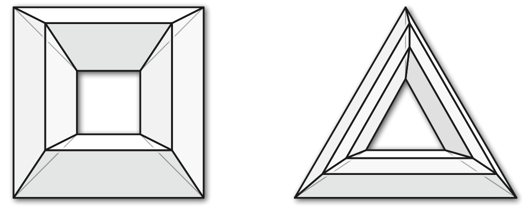

Figure 2. Mitchell's 3×4 tori: Only the mesh on the right supports an interior hex mesh [32].

On the other hand, Eppstein [18] observes 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 in which *no* cycle bounds an interior disk; consider, for example, a ball with a non-trivially knotted tunnel through it.

In practice, hex meshes for complex domains are often constructed by partitioning into simply-connected subdomains and then meshing each subdomain independently [21,36,37,38,42]. 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 *any* quad mesh can be extended to an interior hex mesh if we first partition 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 [33] to transition between the refined boundary and an interior triangulation [10], instead of Eppstein's more complicated algorithm.

1.2 New Results

This paper offers a near-complete solution to the topological hex-meshing problem, generalizing Thurston and Mitchell's results to any compact domain in $\mathbb{R}^3$ whose boundary is a 2-manifold. The domain boundary may be disconnected; each boundary surface may have arbitrary genus; and we do not require the domain to be a handlebody.

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 Section 2.

Main Theorem. *Let Ω be a compact 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 prove our Main Theorem in several stages. The implication (1)⇒(3) follows easily by considering the dual complex of the hex mesh, which we define carefully in Section 2. We prove the equivalence (2)⇔(3) in Section 3 using Poincaré-Alexander-Lefschetz duality. We offer two proofs of the remaining implication (3)⇒(1). In Section 4, we sketch a non-constructive proof in the same spirit as Thurston and Mitchell's argument for genus-zero meshes. Briefly, we argue that if Q^* is null-homologous in Ω, then Q^* is the boundary of an *immersed* surface in Ω, and any surface immersion can be refined to the dual of a hex mesh. Finally, in Section 5, we give a self-contained constructive proof, which uses a modification of Eppstein's algorithm for bipartite meshes. Given an arbitrary polyhedron in $\mathbb{R}^3$ with quadrilateral facets, our algorithm either constructs a compatible topological hex mesh or correctly reports that no such mesh exists, in polynomial time. Readers interested only in our algorithmic proof can safely skip Sections 3 and 4.

Our algorithmic 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 $O(n^2)$ [4, 13]. 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 [17] and Hatcher [27].

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 $\{(x_1, \ldots, x_d) \in \mathbb{R}^d \mid x_d \geq 0\}$. The subspace of points with halfspace neighborhoods is the ***boundary*** of the manifold, denoted $\boldsymbol{\partial\Omega}$. The boundary of any d-manifold is a $(d-1)$-manifold without boundary. We consider only compact manifolds that are subsets of $\mathbb{R}^3$ in this paper, although many of our results apply to more general 3-manifolds. A 2-manifold, possibly with boundary, is less formally called a ***surface***. The ***genus*** of a connected surface is the maximum number of disjoint simple cycles whose deletion leaves the surface connected.

Singularities and Immersions

We consider several types of well-behaved functions from surfaces to 3-manifolds with boundary. Fix a surface Σ, a 3-manifold $\Omega \subset \mathbb{R}^3$, and a continuous function $f : \Sigma \to \Omega$ that maps the boundary of Σ 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 $\operatorname{im} f \cap U$ is homeomorphic to either the plane (if x is in the interior of Σ) or a closed halfplane (if x is on the boundary of Σ), and ***singular*** otherwise.

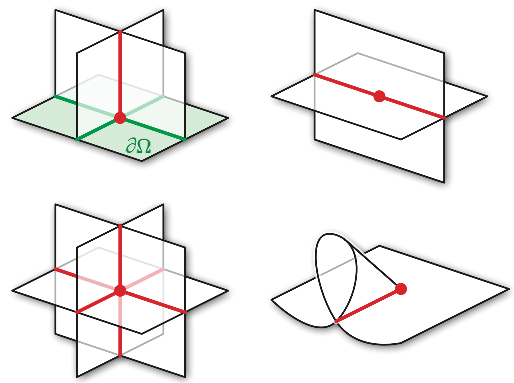

Figure 3. Stable surface singularities: boundary double point, interior double point, triple point, and branch point.

For a ***generic*** surface map, every singular point has a neighborhood homeomorphic to one of four ***stable singularities***, illustrated in Figure 3. 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 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 [54] imply that any surface map f can be continuously deformed to a generic surface map arbitrarily close to f.

A ***(topological, self-transverse) 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 $x \in \Sigma$ lies in an open neighborhood U such that the restriction of f to U is an embedding.

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 $q: [0,1]^3 \to \mathbb{R}^3$. A ***facet*** of q is (the image of) a function from $[0,1]^2 \to \mathbb{R}^3$ 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. A ***boundary facet*** of a geometric cube complex X is a facet of exactly one cube in X; the ***boundary*** $\boldsymbol{\partial X}$ 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 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 [34, 35] and Bing [7, 8], 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 Section 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.

Dual Complexes

Any hex mesh X in $\mathbb{R}^3$ defines a natural ***dual complex*** $\boldsymbol{X^*}$ 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 Figure 4.

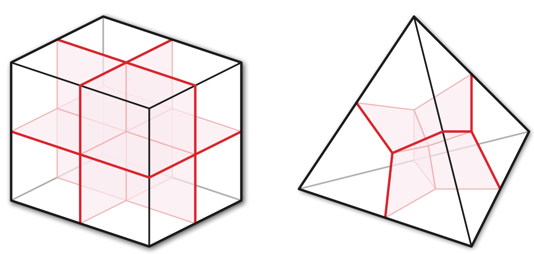

Figure 4. Dual subdivision of a cube and a tetrahedron.

For any hex mesh X, the union of 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* [39], the *derivative complex* of X [3, 23, 29, 47], 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 [19, 20, 48].

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^*.

The dual complex T^* of a triangulation T is defined similarly, by subdividing each tetrahedron in T into four subcubes to obtain a cube complex $T^{\boxplus}$, and then merging the cubes in $T^{\boxplus}$ incident to each vertex of T. See Figure 4.

Homology

Homology is an equivalence relation between cycles of various dimensions with a rich and useful algebraic structure. Our results specifically rely on either 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 $\boldsymbol{\partial C}$, 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*** if their symmetric difference is null-homologous. The homology classes of k-cycles define an abelian group $\boldsymbol{H_k(\Omega)}$, called the ***kth homology group*** of X. Homology groups are finite-dimensional vector spaces over the finite field $\mathbb{Z}_2$.

Homology is a topological invariant of the underlying space Ω, independent of the complex X; thus, we can speak of graphs in $\partial\Omega$ being null-homologous in Ω without specifying any particular cell complex. This independence can be formalized using *singular* homology; we refer the interested reader to Hatcher [27] for further details.

3. NO ODD BOUNDING CYCLES

For the remainder of the paper, we fix a compact domain $\Omega \subset \mathbb{R}^3$ whose boundary is a 2-manifold and a topological quad 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, which are explained in detail by Edelsbrunner and Harer [17, Chapters IV and V] and Hatcher [27].

Lemma 3.1. *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.* [15, 16], we call a subgraph of Q that is null-homologous in Ω a ***handle cycle***. Poincaré-Alexander-Lefschetz duality implies natural isomorphisms between the following homology groups:

$$H_1(S^3 \setminus \Omega) \;\cong\; H_2(\Omega, \partial\Omega) \;\cong\; H^1(\Omega).$$

Specifically, 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 of $H_2(\Omega, \partial\Omega)$. Lefschetz duality implies that the homology class of any cycle γ in Ω is determined by the 2-chains σ_i that γ crosses an odd number of times. In particular, a cycle γ is null-homologous in Ω if and only if γ crosses each 2-chain σ_i 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. □

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.

4. A NON-CONSTRUCTIVE PROOF

To complete the proof of our Main Theorem, it remains only to prove that if Q^* is null-homologous in Ω, then Q is the boundary of a topological hex mesh of Ω. In this section we briefly sketch a non-constructive proof of this result, in the spirit of Thurston's original argument for quad meshes of spheres [52]. Readers interested in a self-contained constructive proof can safely skip ahead to Section 5.

Lemma 4.1. *Suppose Q^* is null-homologous in Ω. Then Q^* is the boundary of a generic surface immersion in Ω.*

Proof (Funar [25, Lemma 4.1]): Because the dual graph Q^* is null-homologous in Ω, classical results of Whitney [54] and Papakyriopoulos [41] imply the existence of a generic surface map $f : \Sigma \to \Omega$ such that $f(\partial\Sigma) = Q^*$. (In fact, a suitable generic surface map can be constructed from a triangulation of Ω using the techniques in Section 5.)

The singular points of f can be viewed as an immersion of a finite number of paths and cycles. Each path endpoint is either a boundary double point of f (that is, a vertex of Q^*) or a branch point of f. 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 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 Figure 5. (Similar surgery operations are described by Hass and Hughes [26] for immersions of surfaces without boundary.) □

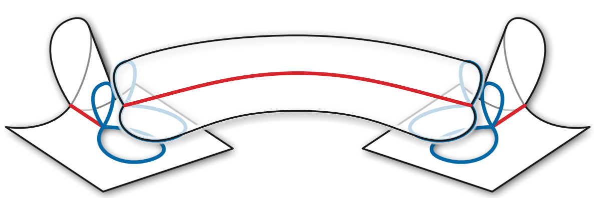

Figure 5. Surgery to cancel two branch points.

Lemma 4.2. *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 [32]. Let $\phi : \Sigma \to \Omega$ be a generic immersion 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) $\operatorname{im}\phi \cap \Omega^\circ$ 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]: For each integer i from 0 to 3, add disjoint bubbles around the subset of each i-dimensional simplex that is outside all previous bubbles. See Figure 6 for a two-dimensional example.[1]

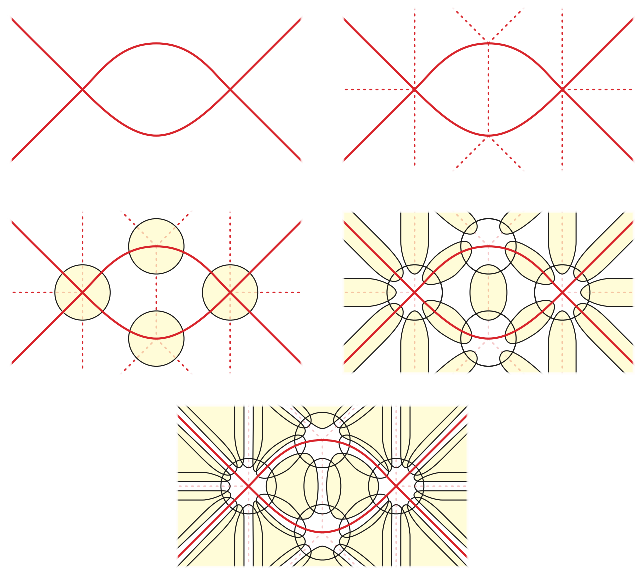

Figure 6. Bubble-wrapping a curve immersion, after Babson and Chan [3].

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 [5, 6], Schwartz [46], and Schwartz and Ziegler [47] for similar analysis; we omit further details here. □

5. AN ALGORITHMIC PROOF

Instead of fleshing out the details of the previous argument, we give in this section a self-contained 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. Our constructive proof does not depend on the results in Sections 3 and 4.

[1]Mitchell [32] describes a more efficient algorithm to transform a surface immersion into the dual of a hex mesh [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 $\Omega \backslash X^*$ 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 [21, 31, 39, 49, 50].

Our high-level strategy closely resembles Eppstein's algorithm for bipartite quad meshes [18]. 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 constant number of smaller cubes. However, our refinement strategy is different from Eppstein's.

5.1 Refining the Interior Triangulation

The dual complex T^* can be constructed by first splitting each tetrahedron 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. Let ∂T denote the induced triangulation of the inner surface of B. 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 $\partial T^* \setminus \Delta^*$. Every vertex in Γ has degree 2, so Γ is a collection of disjoint simple cycles. Contracting all the edges in Δ^* transforms ∂T^* 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.

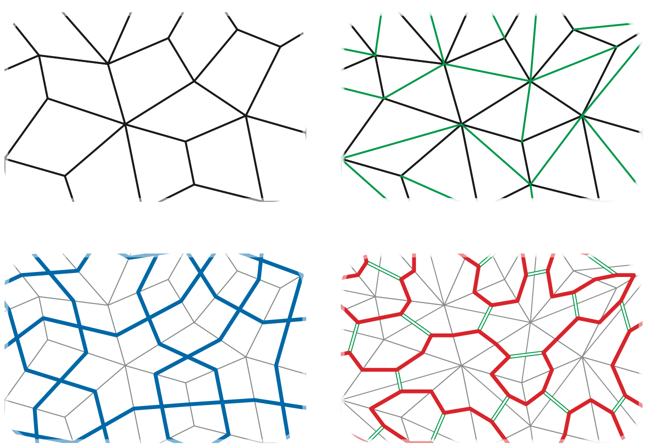

Figure 7. A portion of the surface quad mesh Q, its triangulation ∂T, the dual curves Q^*, and the homologous curves $\Gamma \subset \partial T^*$.

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 the top row of Figure 8.

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 the bottom row of Figure 8. Equivalently, we partition each tetrahedron in T into four cubes by central subdivision, and then expand the surface Σ into a layer of cubes.

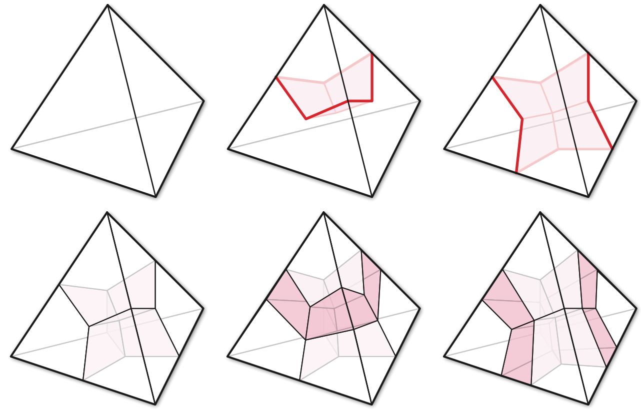

Figure 8. Intersection patterns of Σ with tetrahedra in T, and the corresponding templates for refining into cubes.

5.2 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 the inner facet of each buffer cube into ten quads, as shown in Figure 9, and each edge of that inner facet into three segments. Thus, the transition facets of B are combinatorially hexagons.

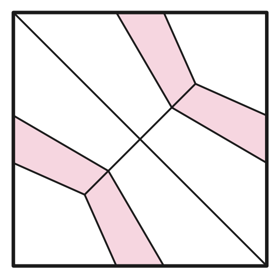

Figure 9. Refinement of the inner facet of each buffer cube.

Following Eppstein [18, 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^*. 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^*.

We then subdivide each transition facet of B into either two or three quads, 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. 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 [18], it is a straightforward exercise to construct explicit hex meshes for these subdivided cubes by hand.

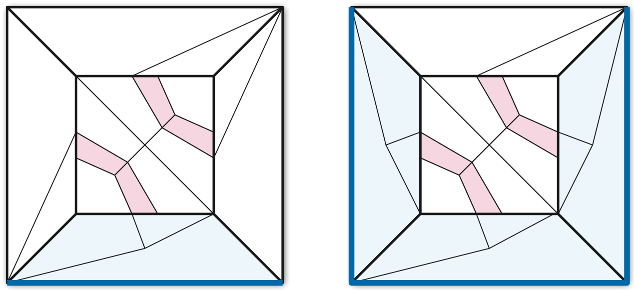

Figure 10. Boundary refinement of the buffer cubes; bold edges are in the odd subgraph R.

5.3 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 3×3 grid of quadrilaterals. On the other hand, 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 [30, 43].

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

Lemma 5.1. *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 [9, 28]. 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, we arbitrarily pair up the vertices of Q, compute the shortest path connecting each pair, and defining R be the subset of edges in an odd number of these shortest paths. This construction requires $O(n^3) = O(t^3)$ time.

Theorem 5.2. *Let Ω be a compact 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 triangulation T of Ω with complexity t, such that T that splits each facet of Q into two triangles. Then we can either compute a topological hex mesh of the interior of Q with complexity $O(t)$, or report correctly that no such hex mesh exists, in $O(t^3)$ time.*

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 3.1 implies that every handle basis in Q contains at least one cycle of odd length. We can compute a handle basis in $O(t^3)$ time using standard homology algorithms [17]; see also Dey *et al.* [15, 16]. (Again, this time bound can be improved using fast matrix multiplication.)

When the facets of Q are actually planar convex quadrilaterals, results of Chazelle and Palios [12], Bern [4], and Chazelle and Shouraboura [13] imply that we can compute a triangulation with complexity $t = O(n^2)$ in $O(n^2 \log n)$ time.

Theorem 5.3. *Given a polyhedron Q with n quadrilateral facets, we can either compute a topological hex mesh of the interior of Q with complexity $O(n^2)$, or report correctly that no such hex mesh exists, in $O(n^6)$ time.*

The upper bound $t = O(n^2)$ is tight for *geometric* triangulations of polyhedra, even with genus zero [11], 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 [51], but the only lower bound actually known is the trivial $t = \Omega(n)$.

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 [5, 6] 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 Figure 1. Although Bern and Eppstein explicitly claim the following result only for genus-zero quad meshes, in fact their argument applies verbatim for geometric quad meshes with arbitrary topology.

Lemma 6.1 (Bern and Eppstein [5, Theorem 2]). *If every bicuboid has a geometric hex mesh, then any geometric quad surface mesh in $\mathbb{R}^3$ that is the boundary of a topological hex mesh is also the boundary of a geometric hex mesh.*

Corollary 6.2. *Let Q be a geometric surface quad mesh in $\mathbb{R}^3$ with an even number of 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 [18] sketched a different reduction of the genus-zero hex meshing problem to a larger finite set of polyhedra, each obtained by refining the boundary of a cube into ether 16 or 18 quads. If each of those boundary-refined cubes has a geometric hex mesh *with constant complexity*, then Eppstein's reduction implies that any convex polytope with $2n$ quadrilateral facets has a geometric hex mesh with complexity $O(n)$. (It is possible that Eppstein's cubes support geometric hex meshes whose complexity depends on the precise geometry of the quadrilaterals.) Our algorithm supports a similar reduction to only a slightly more complex, but still finite, set of examples. We omit further details from this version of the paper.

Theorem 6.3. *Let Q be a geometric surface quad mesh in $\mathbb{R}^3$ with $2n$ facets, whose dual graph Q^* is null-homologous in the interior. If every subdivision of the boundary of a cube into at most 24 convex quadrilaterals is the boundary of a geometric hex mesh* ***with constant complexity****, then Q has a geometric hex mesh with complexity $O(n^2)$.*

References

[1] Iain R. Aitchison, Saburo Matsumoto, and J. Hyam Rubinstein. Immersed surfaces in cubed manifolds. *Asian J. Math* 1(1):85–95, 1997.

[2] Iain R. Aitchison and J. Hyam Rubinstein. An introduction to polyhedral metrics of non-positive curvature on 3-manifolds. *Geometry of Low-Dimensional Manifolds: 2: Symplectic Manifolds and Jones-Witten Theory*, 127–161, 1990. London Math. Soc. Lecture Notes 151, Cambridge Univ. Press.

[3] Eric K. Babson and Clara S. Chan. Counting faces of cubical spheres modulo two. *Discrete Math.* 212(3):169–183, 2000. ArXiv:9811085v1.

[4] Marshall Bern. Compatible triangulations. *Proc. 9th Ann. Symp. Comput. Geom.*, 281–288, 1993.

[5] Marshall Bern and David Eppstein. Flipping cubical meshes. *Proc. 10th Int. Meshing Roundtable*, 19–29, 2001. ⟨http://www.andrew.cmu.edu/user/sowen/abstracts/Be812.html⟩. Preliminary version of [6].

[6] Marshall Bern, David Eppstein, and Jeff Erickson. Flipping cubical meshes. *Eng. Comput.* 18:173–187, 2002. Full version of [5].

[7] R. H. Bing. Locally tame sets are tame. *Ann. Math.* 59:145–158, 1954.

[8] R. H. Bing. An alternative proof that 3-manifolds can be triangulated. *Ann. Math.* 69:37–65, 1959.

[9] James R. Bunch and John E. Hopcroft. Triangular factorization and inversion by fast matrix multiplication. *Math. Comput.* 28(125):231–236, 1974.

[10] Carlos D. Carbonera and Jason F. Shepherd. A constructive approach to constrained hexahedral mesh generation. *Eng. Comput.* 26(4):341–350, 2010.

[11] Bernard Chazelle. Convex partitions of polyhedra: A lower bound and worst-case optimal algorithm. *SIAM J. Comput.* 13(3):488–507, 1984.

[12] Bernard Chazelle and Leonidas Palios. Triangulating a nonconvex polytope. *Discrete Comput. Geom.* 5:505–526, 1990.

[13] Bernard Chazelle and Nadia Shouraboura. Bounds on the size of tetrahedralizations. *Discrete Comput. Geom.* 14:429–444, 1995.

[14] Balázs Csikós and András Szűcs. On the number of triple points of an immersed surface with boundary. *Manuscripta Math.* 87:285–293, 1995.

[15] Tamal K. Dey, Kuiyu Li, and Jian Sun. On computing handle and tunnel loops. *IEEE Proc. Int. Conf. Cyberworlds*, 357–366, 2007.

[16] Tamal K. Dey, Kuiyu Li, Jian Sun, and David Cohen-Steiner. Computing geometry-aware handle and tunnel loops in 3D models. *ACM Trans. Graphics* 27(3):1–9, 2008. Proc. SIGGRAPH 2008.

[17] Herbert Edelsbrunner and John L. Harer. *Computational Topology: An Introduction*. Amer. Math. Soc., 2010.

[18] David Eppstein. Linear-complexity hexahedral mesh generation. *Comput. Geom. Theory Appl.* 12:3–16, 1999.

[19] Evgraf Stepanovich Fedorov. Načala učeniya o figurah [Elements of the study of figures]. *Notices Imp. Petersburg Mineralog. Soc.* 24:1–279, 1885. In Russian. Cited in [48].

[20] Evgraf Stepanovich Fedorov. Elemente der gestaltenlehre. *Z. Krystalographie und Mineralogolgie* 21:679–694, 1893.

[21] Nathan T. Folwell and Scott A. Mitchell. Reliable whisker weaving via curve contraction. *Eng. Comput.* 15(3):292–302, 1999.

[22] George K. Francis. Titus' homotopies of normal curves. *Proc. Amer. Math. Soc.* 30:511–518, 1971.

[23] Louis Funar. Cubulations, immersions, mappability and a problem of Habegger. *Ann. Sci. École Norm. Supér.* 32(5):681–700, 1999. ArXiv:math/9804046.

[24] Louis Funar. Cubulations mod bubble moves. *Low Dimensional Topology*, 29–43, 1999. Contemporary Mathematics 233, Amer. Math. Soc.

[25] Louis Funar. Surface cubications mod flips. *Manuscripta Math.* 125:285–307, 2008. ArXiv:math/0501550.

[26] Joel Hass and John Hughes. Immersions of surfaces in 3-manifolds. *Topology* 24(1):97–112, 1985.

[27] Allen Hatcher. *Algebraic Topology*. Cambridge Univ. Press, 2002. ⟨http://www.math.cornell.edu/~hatcher/AT/ATpage.html⟩.

[28] Oscar H Ibarra, Shlomo Moran, and Roger Hui. A generalization of the fast LUP matrix decomposition algorithm and applications. *J. Algorithms* 3(1):45–56, 1982.

[29] William Jockusch. The lower and upper bound problems for cubical polytopes. *Discrete Comput. Geom.* 9:159–163, 1993.

[30] Marc Lackenby. The volume of hyperbolic alternating link complements. *Proc. London Math. Soc.* 88:204–224, 2004. With an appendix by Ian Agol and Dylan Thurston.

[31] Franck Ledoux and Jason F. Shepherd. Topological modifications of hexahedral meshes via sheet operations: a theoretical study. *Eng. Comput.* 26(4):433–447, 2010.

[32] Scott A. Mitchell. A characterization of the quadrilateral meshes of a surface which admit a compatible hexahedral mesh of the enclosed volume. *Proc. 13th Ann. Symp. Theoret. Aspects Comput. Sci.*, 456–476, 1996. Lecture Notes Comput. Sci. 1046, Springer-Verlag.

[33] Scott A. Mitchell. The all-hex Geode-template for conforming a diced tetrahedral mesh to any diced hexahedral mesh. *Proc. 7th International Meshing Roundtable*, 295–305, 1998.

[34] Edwin E. Moise. Affine structures in 3-manifolds, V. The triangulation theorem and Hauptvermutung. *Ann. Math.* 56:96–114, 1952.

[35] Edwin E. Moise. Affine structures in 3-manifolds, VIII. Invariance of the knot-type; local tame embedding. *Ann. Math.* 59:159–170, 1954.

[36] Matthias Müller-Hannemann. Hexahedral meshing by successive dual cycle elimination. *Eng. Comput.* 15(3):269–279, 1999.

[37] Matthias Müller-Hannemann. Shelling hexahedral complexes for mesh generation. *J. Graph Algorithms Appl.* 5(5):59–91, 2001.

[38] Matthias Müller-Hannemann. Quadrilateral surface meshes without self-intersecting dual cycles for hexahedral mesh generation. *Comput. Geom. Theory Appl.* 22:75–97, 2002.

[39] Peter Murdoch, Steven Benzley, Ted Blacker, and Scott A. Mitchell. The spatial twist continuum: A connectivity based method for representing all-hexahedral finite element meshes. *Finite Elements in Analysis and Design* 28:137–149, 1997.

[40] Tahl Nowik. Complexity of planar and spherical curves. *Duke J. Math.* 148(1):107–118, 2009.

[41] Christos D. Papakyriakopoulos. On Dehn's lemma and the asphericity of knots. *Ann. Math.* 66:1–26, 1957.

[42] Mark A. Price and Cecil G. Armstrong. Hexahedral mesh generation by medial surface subdivision: Part II. Solids with flat and concave edges. *Int. J. Num. Methods. Eng.* 40(1):111–136, 1997.

[43] Jessica S. Purcell. Volumes of highly twisted knots and links. *Alg. & Geom. Top.* 7:93–108, 2007. ArXiv:math.GT/0604476.

[44] Robert Schneiders. Open problem, 1995. ⟨http://www-users.informatik.rwth-aachen.de/~roberts/open.html⟩. Last accessed January 1, 2012.

[45] Robert Schneiders. A grid-based algorithm for the generation of hexahedral element meshes. *Eng. Comput.* 12:168–177, 1996.

[46] Alexander Schwartz. *Constructions of Cubical Polytopes*. Phd dissertation, Tech. Univ. Berlin, 2003. ⟨http://edocs.tu-berlin.de/diss/2004/schwartz_alexander.html⟩.

[47] Alexander Schwartz and Günter M. Ziegler. Construction techniques for cubical complexes, odd cubical 4-polytopes, and prescribed dual manifolds. *Experimental Math.* 13(4):385–413, 2004. ArXiv:math/0310269v3.

[48] Marjorie Senechal and Ravil Vagizovich Galiulin. An introduction to the theory of figures: The geometry of E. S. Fedorov. *Structural Topology* 10:5–22, 1984.

[49] Jason F. Shepherd and Chris R. Johnson. Hexahedral mesh generation constraints. *Eng. Comput.* 24(3):195–213, 2008.

[50] Tatsuhiko Suzuki, Shigeo Takahashi, and Jason F. Shepherd. An interior surface generation method for all-hexahedral meshing. *Eng. Comput.* 26(3):303–316, 2010.

[51] Bill Thurston and Dylan Thurston. Complexity of random knot with vertices on sphere. MathOverflow, February 10, 2011. ⟨http://mathoverflow.net/questions/54417⟩.

[52] William P. Thurston. Hexahedral decomposition of polyhedra. Posting to sci.math, 25 October 1993. ⟨http://www.ics.uci.edu/~eppstein/gina/Thurston-hexahedra.html⟩.

[53] Hassler Whitney. On regular closed curves in the plane. *Compositio Math.* 4:276–284, 1937.

[54] Hassler Whitney. The singularities of a smooth n-manifold in $(2n-1)$-space. *Ann. Math.* 45(2):247–293, 1944.

[55] Soji Yamakawa and Kenji Shimada. HEXHOOP: Modular templates for converting a hex-dominant mesh to an ALL-hex mesh. *Eng. Comput.* 18:211–228, 2002.

[56] Soji Yamakawa and Kenji Shimada. 88-element solution to Schneiders' pyramid hex-meshing problem. *Int. J. Num. Methods. Biomed. Eng.* 26:1700–1712, 2010.