CONF-960233--1
SAND95-1994C

A characterization of the quadrilateral meshes of a surface which admit a compatible hexahedral mesh of the enclosed volume

Scott A. Mitchell[1]

Comp. Mechanics and Visualization Dept., Sandia National Laboratories, Albuquerque, NM 87185

Abstract

A popular three-dimensional mesh generation scheme is to start with a quadrilateral mesh of the surface of a volume, and then attempt to fill the interior of volume with hexahedra, so that the hexahedra touch the surface in exactly the given quadrilaterals[3]. Folklore has maintained that there are many quadrilateral meshes for which no such *compatible* hexahedral mesh exists. In this paper we give an existence proof which contradicts this folklore: A quadrilateral mesh need only satisfy some very weak conditions for there to exist a compatible hexahedral mesh. For a volume that is topologically a ball, any quadrilateral mesh composed of an even number of quadrilaterals admits a compatible hexahedral mesh. We extend this to volumes of higher genus: There is a construction to reduce to the ball case if and only if certain cycles of edges are even.

Keywords: Computational Geometry, hexahedral mesh generation, existence.

1. Introduction

For some applications, a mesh composed of quadrilateral faces and hexahedral (i.e. cube-like) elements possess better numerical properties than a mesh composed of triangular faces and tetrahedral elements. Hence a sizable fraction of the mesh generation research conducted in recent years has been devoted to hexahedral meshes[19]. In some large-scale applications, the surface of an object is meshed before its interior. This is a requirement for meshing several adjoining objects independently. This requirement may also arise in parallel mesh generation, where the domain is first divided into many small regions, one for each processor. The problem is to produce a hexahedral mesh that fills the volume and touches the surface in exactly the given surface mesh. We say that such hexahedral and quadrilateral meshes are *compatible*, that the hexahedral mesh *respects* the quadrilateral mesh, and the quadrilateral mesh *admits* the hexahedral mesh.

For many years, meshing algorithm developers have tried to solve this problem, without even knowing if it could be done. The difficulty of this problem has lead many to conclude

---

1 *samitch@sandia.gov*. This work was supported by the U.S. Department of Energy under contract DE-AC04-76DO00789, by the Applied Mathematical Sciences program, U.S. Department of Energy Research.

that the surface mesh must be highly constrained in order for it to admit a compatible hexahedral mesh. Our result is that this is not the case. We show that for an input that is topologically a ball, all that is required is that the quadrilateral mesh has an even number of quadrilaterals! This condition is also necessary. We extend this result to a characterization of the surface meshes of a sphere with n-handles that admit compatible hexahedral meshes.

Bern[1] has previously given an algorithm that generates a compatible tetrahedral mesh of a polyhedron given any triangular mesh of its surface, settling the existence question for triangular/tetrahedral meshes. Schneiders has posed as an open problem the characterization of which quadrilateral surface meshes admit a compatible hexahedral mesh[17]. We do not give an algorithm, but we do settle the existence question. Our techniques are novel and entirely different from Bern's[1]. Our approach relies on the *spatial twist continuum*, or STC[2], a global interpretation of the connectivity of the dual of a mesh, and some theorems of topology concerning regular curves and manifolds[4].

Our result is entirely topological and combinatorial. We define the quadrilateral mesh, and the hexahedral mesh, purely in terms of their topology. We ignore the geometric embedding of the mesh. In particular, we have no guarantees on the shape of the elements of the mesh, except that certain combinatorial pathologies that a priori require bad elements are ruled out. For the most part, our techniques are elementary. The outline of the proof is as follows: First, we map the object to a ball, mapping the given surface mesh to a mesh of the sphere. Second, we form the two-dimensional STC of the surface mesh. This is an arrangement of regular curves on the sphere, whose induced cell complex is the combinatorial dual[8] of the quadrilateral mesh. Third, we use theorems of topology[4] to show that the curve arrangement can be extended into an arrangement of regular manifolds through the ball. Fourth, we add additional manifolds, entirely interior to the ball, so that the cell complex induced by the arrangement satisfies certain combinatorial constraints. Fifth, the induced cell complex of the manifold arrangement is dualized to create a hexahedral mesh, and mapped from the ball back to the original object. The combinatorial constraints on the manifold arrangement ensure that the hexahedral mesh is well defined. With the exception of showing the existence of the initial manifolds, the third part, the proof is constructive.

Most known hexahedral mesh generation codes don't address the problem of respecting a given quadrilateral surface mesh. Of those that do, all either change the surface mesh in some way, or add non-hexahedral elements. For example, Plastering[16], the current version of Whisker Weaving[3], and Algor's Hexagen coupled with Houdini[15], all allow the user to chose between changing the surface mesh, and having all hexahedral elements. This paper shows that these caveats are not usually necessary, so I hope developers will be inspired to remove them from their codes. In particular, Whisker Weaving[3] is a heuristic algorithm that attempts to create a valid STC, which is then dualized to a hexahedral mesh: The fix-up rules given in Section 6 should remove one of the few remaining difficulties for this algorithm.

The remainder of this paper is organized as follows. In Section 2 we describe our assumptions and requirements about quadrilateral and hexahedral meshes. In Section 3 we define the STC. In Section 4 we present the necessary conditions, and in Section 5 show that the two-dimensional STC of the surface mesh satisfying these conditions can be extended to a three-dimensional STC. In Section 6 we show how to add to the arrangement so that the induced cell complex dualizes to a valid hexahedral mesh. In Section 7 we extend our results to non-ball input. In Section 8 we present conclusions.

2. Mesh definitions and assumptions

Quadrilateral mesh definition

We suppose we are given a polyhedron P and a quadrilateral mesh of its surface. A quadrilateral mesh of the polyhedron will be a geometric face lattice[18] (or *cell complex*), composed of 0-dimensional nodes, 1-dimensional edges, and 2-dimensional faces. We require the following:

a. Each edge contains two distinct nodes.

b. Each facet is contained in at least one higher-dimensional facet. I.e. each node is in an edge, each edge is in a face.

c. Every edge is in exactly two distinct faces.

d. Each face is bounded by a cycle of four distinct edges.

e. Two nodes have at most one edge between them.

f. Two faces share at most one edge[5].

Mitchell et al. has recently given an algorithm called Pillowing Doublets, by which a quadrilateral mesh satisfying **a** through **e** can be locally refined to satisfy **f** as well[5].

Hexahedral mesh definition

We define a hexahedral mesh as follows. This definition of a hexahedral mesh may be too weak for certain numerical applications. Only the combinatorial aspects of the mesh are considered. For example, it may not be possible to embed the mesh in Euclidean space with straight edges. However, certain undesirable combinatorial structures are forbidden.

A hexahedral mesh is a three-dimensional geometric face lattice, or cell complex. The 0-dimensional entities are called *nodes*. The 1-dimensional entities are called *edges*. The 2-dimensional entities are called *faces*, and the 3-dimensional entities are called *hexahedra* or *hexes*. We require the following:

A. Each edge contains two distinct nodes.

B. Each facet is contained in at least one higher-dimensional facet. I.e. each node is in an edge, each edge is in a face, and each face is in a hexahedron.

C. Every face is contained in exactly two distinct hexes, except that those of the surface mesh are contained in exactly one hex. The edges of the face have the opposite ordering in the two hexes (i.e. the hexes are on "opposite sides" of the face).

D. Each face is bounded by a cycle of four distinct edges.

E. A hex is bounded by six distinct faces. Furthermore, these faces pairwise share edges in the following way. Face 0's ordered edge cycle is $\{abcd\}$, face 1 $\{aiel\}$, face 2 $\{bjfi\}$, face 3 $\{efgh\}$, face 4 $\{cjgk\}$, face 5 $\{dlhk\}$. Distinct letters represent distinct edges.

F. Two nodes have at most one edge between them.

These conditions may rule out additional pathologies that may not be immediately obvious. For example, they imply that the entire polyhedron is meshed, that there is not any "internal voids" left unmeshed.

By the technique of Pillowing Doublets[5], a mesh satisfying the above constraints can also be made to satisfy constraint **G** below. The surface mesh does not change as long as it satisfies condition **f**.

G. Any two faces share at most one edge. This also implies that any two hexes share at most one face.

We will use constructive techniques similar to Pillowing Doublets[5] in our existence proof below to force a cell complex to satisfy **A** through **F**.

3. STC definition

The *spatial twist continuum* (STC)[2] is a special structure superimposed on the combinatorial dual[8] of a quadrilateral or hexahedral mesh. Any quadrilateral and any hexahedral mesh induces an STC. By duality, it is also possible to derive a mesh from a given STC. The STC definition may be extended to meshes composed of d-cubes in any dimension d, but similar structures have not been found for meshes of other types of elements.

Two-dimensional STC

The two-dimensional STC is a non-degenerate arrangement of regular curves called *chords*. By non-degenerate, we mean that the chords are nowhere tangent, and that at most two chords meet at a point. For historical reasons[2], chords that are closed curves are called *loops* and the points of intersection are called *centroids*. The STC of a given quadrilateral mesh is any one of the of arrangements whose induced cell complex is the combinatorial dual of the mesh[8]. By combinatorial dual, we mean that only the combinatorial/topological structure is considered, the geometric embedding is ignored.

It appears to have been known for some time[9] that an arrangement of regular curves dualizes to a quadrilateral mesh, but the author is unaware of when this observation was first made, or if any similar observations about arrangements in three dimensions had been made prior to Murdoch et. al[2].

For simplicity, we first take the given input solid and map it to the ball, mapping the given quadrilateral mesh to a quadrilateral mesh of the sphere. We then form the two-dimensional STC of the quadrilateral mesh of the sphere. See Murdoch et. al[2] for a full description of the construction, but we outline it here for clarity: We first form the dual of the quadrilateral mesh; see Figure 1. We form chains of dual edges: the two edges that are dual to the two edges forming opposite sides of a quadrilateral are considered to lie in the same chain. Since the sphere is closed every quadrilateral edge is in exactly two faces, and each chain is actually a (self-intersecting) cycle of dual edges. For each chain we form a regular curve *(loop)* through its centroids.

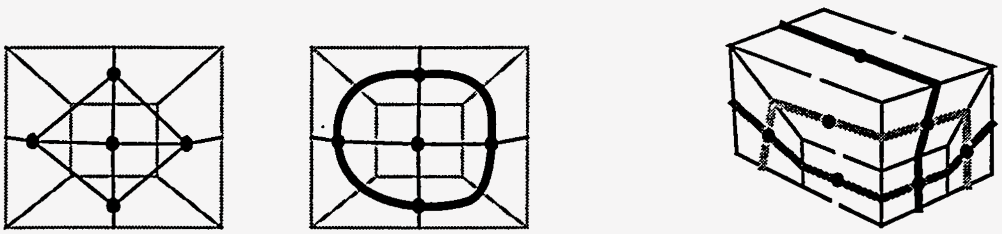

Figure 1. Left, the dual of a portion of a quadrilateral mesh. Center, a loop of the STC drawn with a wide line width. Left shows a three-dimensional view of the loops of a quadrilateral surface mesh.

Three-dimensional STC

The three-dimensional STC is a non-degenerate arrangement of regular manifolds called *sheets*[2]. By non-degenerate, we mean that the sheets are nowhere tangent, and that at most three sheets meet at a point. Figure 2 shows the STC for a mesh composed of a four hexahedra. For historical reasons[2], the curves of intersection are called *chords* and the points of intersection are called *centroids*. The STC of a hexahedral mesh is any one of the of arrangements whose induced cell complex is the combinatorial dual of the mesh. The intersection of the arrangement with the surface of the object is a two-dimensional STC. (The loops are closed curves but may be non-regular if e.g. the object surface is planar facetted. This is unimportant.) The two-dimensional STC's induced cell complex is the dual of the quadrilateral mesh of the surface. For brevity, we will call any non-degenerate arrangement of sheets an STC, regardless of whether its cell complex dualizes to a valid hexahedral mesh. The main thrust of this paper is to show the existence of an STC whose induced cell complex is the dual of some hexahedral mesh that meets the sphere in exactly the given quadrilateral mesh.

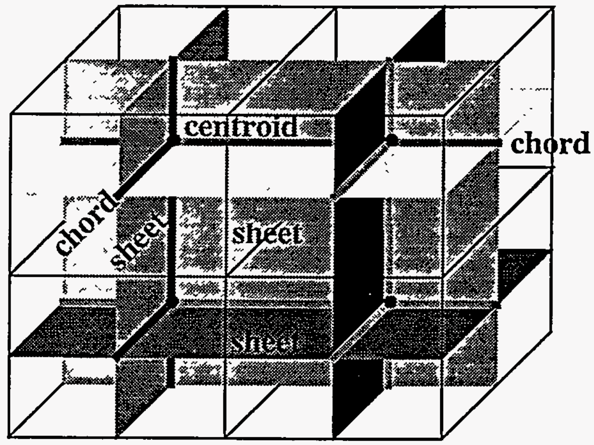

Figure 2. The spatial twist continuum (STC) for a mesh of four hexahedra.

4. Necessary condition

For any hexahedral mesh of any type of input, the surface faces are in only one hex and internal faces are in two hexes, so that there are an even number of faces on the boundary of the mesh. Consequently, for there to exist a hexahedral mesh that respects a given quadrilateral mesh, the quadrilateral mesh must have an even number of faces. For example, no hex mesh exists for the surface mesh in Figure 3.

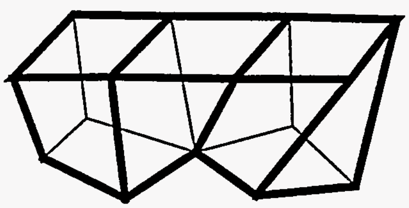

Figure 3. A surface mesh with an odd number of faces has no hex mesh respecting it. For example, this surface mesh has 13 faces.

The surprising result of this work is that this necessary condition is actually sufficient! The remaining sections of this paper are devoted to showing this.

***Theorem 1** Any hexahedral mesh has an even number of quadrilateral faces on its boundary.*

Proof. Let h be the number of hexes in the mesh, f the number of faces, and b the number of faces on the boundary of the mesh. Then b faces are in one hex, all other faces are contained in two hexes. Each hex contains six faces. Hence $6h = 2f - b$, so b is even. ∎

5. Constructing an initial STC

Given a surface mesh with an even number of quadrilaterals, we show that there exists an arrangement of regular manifolds that meets the sphere in exactly the two-dimensional STC of the given surface mesh. The even-ness condition is first translated into a condition on the parity of the number of loop self-intersections. An immediate consequence of an observation of Gauss[10], summarized in Rosenstiehl[9], is the following lemma.

Lemma 1 In the arrangement of loops of a two-dimensional STC of a quadrilateral mesh of the sphere, the parity of the number of loop self-intersections is equal to the parity of the number of surface faces.

Proof. Each quadrilateral of loop A is the intersection of loop A with some other loop, or A with itself. In Gauss[10] and Rosenstiehl[9], and by the Jordan curve theorem, we noted that two nowhere-tangent closed curves on the sphere (plane) intersect an even number of times. Hence the number of intersections of A with all other loops is even. So the parity of the number of quadrilaterals of A is the parity of the number of self-intersections of A.

For the entire surface mesh, there are an even number of intersections between loops A and B for all loops $A \neq B$. Hence the parity of the surface mesh is the parity of the sum over all loops of the number of self-intersections. ∎

A combinatorial description of arrangements of curves called Gauss codes has been studied for some time. It appears that the focus has been on algorithms to recognize which Gauss codes can be realized as arrangements[13], or classifying non-isomorphic arrangements[12], rather than on transforming one such arrangement to another as in the present work or in local mesh refinement[11]. A topological theorem of Smale[4] that is central to our work is the following:

Theorem 2 [Smale] Let x_0 be a point in the unit tangent bundle T of a Riemannian manifold M. Then there is a 1-1 correspondence between the set π_0 of classes (under regular homotopy) of regular curves on M which start and end at the point and direction determined by x_0 and $\pi_1(T, x_0)$.

What this theorem says for M equal to the sphere, is that there is a regular homotopy (i.e. smooth transformation) between any loop with an even number of self-intersections and a regular curve with no self-intersections. Similarly, there is a regular homotopy between any loop with an odd number of self-intersections and a regular curve with one self-intersection: Hence there is a regular homotopy between any two loops with an odd number of self intersections. The importance of the necessary parity condition, Theorem 1, is that, if there are an even number of surface quadrilaterals then the loops homotopic to a curve with one self-intersection can be taken in pairs.

We use the existence of these regular homotopys to show that there exist regular manifolds that respect the surface loops. By respect, we mean that every loop is contained in the boundary of one manifold, and the boundary of a manifold is exactly some number of surface loops.

Theorem 3 The two-dimensional STC of an even mesh of the sphere admits a compatible arrangement of regular manifolds through the ball.

Proof. We construct a manifold homeomorphic to a disk for each even loop L as follows. By Theorem 2 there exists a regular homotopy h on the sphere, where $h(0) = L$ and $h(1)$ is a circle. We then chose, for example, the manifold $(1 - t/2, h(t))$, where (r,s) denotes the curve s at radius r from the center of the ball. We can extend this to close the circle with a disk inside the ball.

We construct a manifold for each <u>pair</u> of odd loops L_1, L_2 as follows. There is a regular homotopy h from L_1 to L_2. We can then chose, for example, the manifold $\left(\frac{3}{4} + \left(t - \frac{1}{2}\right)^2, h(t)\right)$. This manifold has exactly loops L_1 and L_2 as its boundary.

If necessary, the manifolds are perturbed so that they are regular, nowhere tangent (and nowhere self-tangent), and at most three intersect at any point. Note that only one manifold contains each point of a loop, except that two contain each loop centroid (or one manifold contains the loop centroid twice in the case of a self-intersection). ▮

The dual of the cell complex induced by this arrangement (together with the sphere) is a cell complex that respects the surface mesh, but conditions **A** through **F** are not necessarily satisfied. We show below that by adding regular manifolds topologically equivalent to a sphere, conditions **A** through **F** may be satisfied. These additional sheets do not intersect the surface of the sphere, so the arrangement still respects the surface mesh.

6. Adding to the STC arrangement to satisfy the validity constraints

The dual conditions **A+** through **F+** for an arrangement's cell complex corresponding to **A** through **F** are the following. For simplicity, the sphere itself is considered to be part of the arrangement. We use the term *internal* to denote cells that are not on the sphere.

A+. Each internal 2-cell is contained in exactly two distinct 3-cells.

B+. Each facet contains at least one lower-dimensional facet (excepting centroids).

C+. Each STC edge has two distinct centroids. Every surface centroid has one internal STC edge, which connects it to an internal centroid.

D+. Each internal STC edge is contained in exactly four distinct 2-cells.

E+. Each internal STC centroid is contained in six STC edges. The edges are in twelve common 2-cells as follows. Edge 0 is in 2-cells $\{abcd\}$, edge 1 $\{aiel\}$, edge 2 $\{bjfi\}$, edge 3 $\{efgh\}$, edge 4 $\{cjgk\}$, edge 5 $\{dlhk\}$. Distinct letters represent distinct 2-cells.

F+. Two 3-cells have at most one 2-cell in common.

Theorem 4 For any non-degenerate arrangement of regular manifolds in a ball, whose boundary is a given two-dimensional STC on a sphere, there exists additional sheets interior to the ball, such that the combined arrangement satisfies conditions A+ through F+.

Proof. Many of these conditions are satisfied by the initial arrangement without modification. Most of the time when a condition fails to hold it is because of the requirement that facets are distinct. This arises because the arrangement is locally too coarse. Adding additional sheets alleviates this problem.These additional sheets are topologically equivalent to the sphere and do not intersect the surface mesh. Each added sheet surrounds some arrangement facet, at approximately distance ε from the facet. Each added sheet has its ε smaller than all previous ε's. The order in which the conditions are considered is important: When adding sheets to satisfy a condition, it is assumed that all previous conditions are already satisfied, and care is taken so that the previous conditions are still satisfied after the current modification.

Proof A+ Each internal 2-cell is contained in exactly two distinct 3-cells. A 2-cell is a subset of a sheet. The sheets constructed so far are orientable with boundary on the sphere. In the same way that the bounding curves divide the surface of the sphere into (at least) an inside and an outside, one cannot travel inside the ball from one "side" of the manifold to the other without crossing the manifold. Hence every point in a 2-cell of the manifold is in two distinct 3-cells.

Proof B+ Each facet contains at least one facet. An edge is a subset of a *chord*, a curve of intersection between two sheets or a sheet and itself. A chord either starts and ends at two surface centroids, or is a closed curve. Furthermore, a chord may be homotopic to the circle and contain no centroids; see the construction below. In all other cases, the edge contains a centroid. Every 2-cell contains an edge, since every sheet introduced so far has a boundary. Every 3-cell contains a 2-cell, since the ball is bounded.

Construction: If a closed chord of intersection C between sheets A and B exists without centroids, then we add two topological ball sheets Z_1 and Z_2 as in Figure 4, each one containing slightly more than half of C. The sheets may be made small enough that they do not intersect the arrangement otherwise. Chord C now has 4 centroids. Each of the two circles of intersection D between Z_1 and Z_2 are intersected orthogonally by A and B, as in Figure 4 right. These circular chords also have 4 centroids. **A+** is still satisfied because Z_1 and Z_2 are closed and orientable.

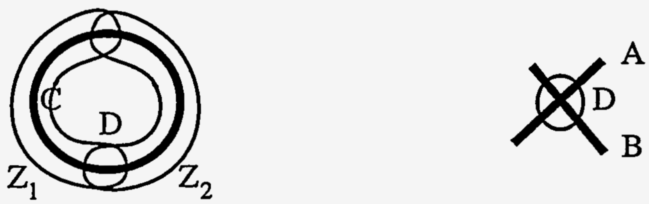

Figure 4. Fixing a chord C with no centroids: Right shows a cross section in the "plane" of C. Left shows a cross section in the "plane" of one of the curves of intersection between Z_1 and Z_2.

Proof C+ Each STC edge has two distinct centroids. An edge may contain the same non-surface centroid twice in the case that a sheet self-intersects at the centroid. Otherwise the centroids are distinct.

Construction: If an edge contains the same centroid twice, then we introduce a small ball sheet Z around that centroid. The intersection of Z with the arrangement is topologically equivalent to the intersection of the sphere with the three planes through the coordinate axis as in Figure 5 left.

Proof C+ Every surface centroid has one internal STC edge, which connects it to an internal centroid. Every loop centroid is in exactly one internal edge, but this edge may contain two (distinct) loop centroids.

Construction: If any edge contains two loop centroids, then we introduce a spherical sheet Z with radius $1-\varepsilon$ for some small ε. Thus, every 3-cell containing a loop is a prism whose base is a surface 2-cell A, and whose top is an interior 2-cell B. Every internal edge containing a centroid of A also contains a centroid of B. The previous conditions hold near Z because the surface mesh satisfies **a** through **f.**

Proof D+ Each internal STC edge is contained in exactly four distinct 2-cells. An edge is the non-tangent intersection of two sheets, or the same sheet with itself. In a small neighborhood around the edge, there are four distinct portions of 2-cells containing the edge. In the one sheet case, it may be that two non-opposite 2-cells are not distinct as in Figure 5.

Construction: If an edge E is contained in the same 2-cell A twice, then we surround the edge by a small ball sheet Z as in Figure 5. Z may be made small enough so that it doesn't otherwise intersect the arrangement. By **C+**, the two centroids of the edge are distinct, so each of the 2-cells containing E are distinct. Also, the 2-cell A that originally contained E twice, now instead contains two distinct edges that are the intersection of Z and the sheet containing A. None of the other edges of Z can be contained in a 2-cell twice: They are locally topologically equivalent to the edges of intersection between a unit ball and the three coordinate planes of $\Re^3$. The two 2-cells in Z containing each edge are distinct because they are separated by a "coordinate plane". Similarly the two other 2-cells containing an edge of Z are separated by Z.

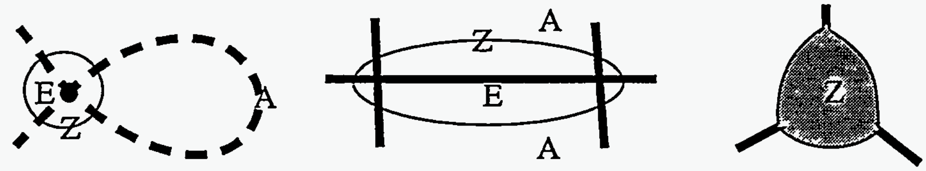

Figure 5. An edge E may appear in a 2-cell A twice. Left shows a cross section perpendicular to E. Center shows a cross section through E in the "plane" of A. Left, view in an octant near a centroid.

Proof E+ Each STC centroid is contained in six STC edges... E+ holds provided both **D+** and **C+** do. A centroid is the pairwise non-tangent intersection of three sheets. Hence every centroid has six edges. Provided **C+** holds, these edges are distinct. The 2-cells obey the required labelling, and are distinct if **D+** holds.

Proof F+ Two 3-cells have at most one 2-cell in common. This need not hold initially.

Construction: If two 3-cells contain 2-cells A and B in common, then we introduce a topological ball sheet Z containing A as in Figure 6. By the previous fix-up steps, each edge of A is contained in A only once. The sheet Z is made small enough that it intersects the arrangement only in a small neighborhood of each edge of A. Now, A is contained in

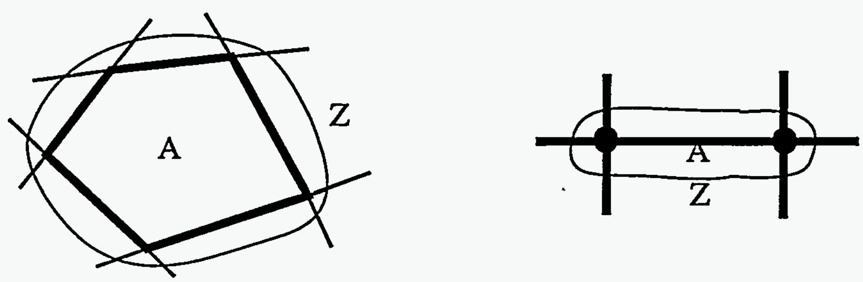

Figure 6. Removing a 2-cell A from containment in a 3-cell. Left, view in the "plane" of A. Right, view in the plane perpendicular to an edge of A.

the two 3-cells between A and Z. In each of the original 3-cells A has been replaced by a 2-cell on Z. All the previous conditions still hold.

This concludes the proof of theorem 4. ∎

Dualizing the cell complex of the constructed arrangement and the conditions **A+** through **F+** we have the following.

*Theorem 5 Any even quadrilateral mesh (satisfying **a** through **f**) of a surface topologically equivalent to a sphere admits a compatible hexahedral mesh (satisfying **A** through **G**) of the enclosed volume.*

7. Extensions to non-ball input

We now present necessary and sufficient conditions in order to extend our results to quadrilateralizations of a surface that is topologically equivalent to a sphere with n-handles. The basic idea is to reduce to the spherical case.

Lemma 2 A quadrilateral mesh of any surface is bounded by an even number of edges.

Proof. Let f be the number of faces, e be the number of edges, and b the number of boundary edges of the mesh. Then $4f = 2e - b$, so b is even. ∎

Consider the cycles of edges in a quadrilateral mesh of a sphere with n-handles.

Lemma 3 All simple cycles of edges in the same homotopy class have the same parity.

Proof. A cycle homotopic to 0 is the boundary of a pseudo-manifold (facetted surface) that is quadrilateralized. By Lemma 2 the cycle is even. Otherwise, take a reference cycle R, and take any other cycle K in the same homotopy class. By definition these bound a (collection of) quadrilateralized pseudo-manifolds. The sum of the number of edges of this boundary, the number of edges that they don't have in common, is even. Hence the sum of the total number of edges they have is even. ∎

Theorem 6 A hexahedral mesh respecting the surface mesh exists if and only if the surface mesh has an even number of quadrilaterals and every homotopy class of simple cycles of edges that can be contracted to 0 inside the volume has even parity.

Proof. In the previous proof we saw that cycles homotopic to 0 on the surface have even parity. Hence in the sphere with no handles case this degenerates to Theorem 5. Otherwise, for each handle in turn, we take some cycle R in its homotopy class. Since R has even parity, we can extend the quadrilateralization of the surface to include a disk-like pseudo-manifold M inside the volume P whose boundary is R. If we conceptually separate M and R into two, one for each side of M, i.e. "slicing" along R, we can recursively treat P as a manifold with n-1 handles, until the spherical case is reached. Thus the condition is sufficient.

This condition is necessary as well. In any hexahedral mesh respecting the surface mesh, it can be shown that the simple edge cycles of interest bound a pseudo-manifold, which is a quadrilateral sub-complex of the hexahedral mesh[6][7]. A quadrilateral mesh is always bounded by an even number of edges, hence the simple cycles of interest must be even. ∎

The following figure illustrates the difference between the simple cycles of edges that can be contracted to 0 in the volume and those that cannot.

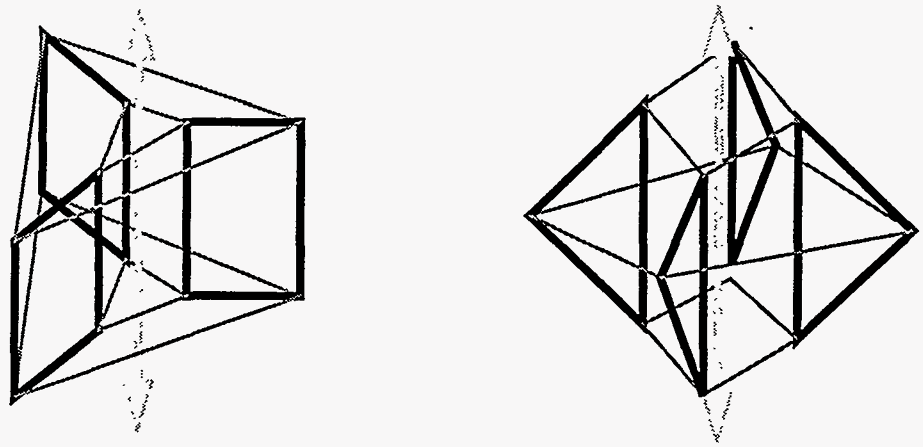

Figure 7. Two quadrilateral meshes of the surface of a torus. The axis of symmetry is vertical. Left, cycles in the non-trivial homotopy class have even parity and a hexahedral mesh of three elements is obvious. Right, cycles in the non-trivial homotopy class have odd parity, and no hexahedral mesh exists. Note that both surface meshes otherwise have the same combinatorial structure: the only difference is what is considered the inside and outside of the torus. The left figure has a homotopy class of simple cycles of odd length, but these surround the axis and are not contractible to 0 inside the volume, hence they do not enter into Theorem 6. Similarly, the right figure has a homotopy class of simple cycles of even length similar to the left figure, but these surround the axis and are not contractible to 0 inside the volume.

8. Conclusions

We have shown that given mild conditions on a surface mesh, there exists a hexahedral mesh filling the interior of the volume. The fact that the sufficient conditions are minor runs counter to the experience of most mesh generation algorithm developers. This is probably due to the fact that previous algorithms had no way to quantify the global connectivity constraints inherent in hexahedral meshes, while today we know that the STC captures these constraints beautifully and succinctly. Some steps of the proof are constructive, and may lead to practical algorithms.

CUBIT is a suite of mesh generation tools under development by Sandia National Laboratories and others under contract. Currently, Gasilov et al.[6] is developing a practical algorithm for CUBIT, along the lines of the proof of Theorem 6, to reduce the problem of constructing a compatible hexahedral mesh for a topological sphere with n-handles to the problem of constructing a compatible hexahedral mesh for a topological sphere. The sphere will then be meshed using the Whisker Weaving algorithm[3], or perhaps one of the other techniques available in CUBIT, such mapping identifiable subregions[14]. Whisker Weaving is based on the STC, and meshes a topological sphere by creating the arrangement of pseudo-manifolds in an advancing-front manner (in contrast to Section 5).

Acknowledgments

I would like to thank my CUBIT co-developers and sponsors for their valuable comments. I would like to thank V. A. Gasilov, E. L. Kartashova, and E. V. Sandrakova for their contribution to the CUBIT project, for putting Section 7 on a solid mathematical foundation, and for having a proof that the conditions in Theorem 6 are necessary. I would like to thank Stephen Smale for pointing me to his wonderful theorem (Theorem 2). I would like to thank Michael Hohmeyer, Jonathan Shewchuk and others at the 3rd International Meshing Roundtable for their interesting discussions concerning Figure 7. Finally, I would like to thank Robert Schneiders for posing this as an important problem on the Web.

References

1. M. Bern. Compatible tetrahedralizations, *Proc. 9th Annual Symp. on Computational Geometry,* 1993, 281-288.

2. P. Murdoch, S. Benzley, T. D. Blacker, S. A. Mitchell. The spatial twist continuum: a connectivity based method for representing all–hexahedral finite element meshes, submitted to Int. J. Numer. Methods Engrg. (1995).

3. T. J. Tautges, T. D. Blacker, S. A. Mitchell. Whisker Weaving: a connectivity based method for constructing all–hexahedral finite element meshes, submitted to Int. J. Numer. Methods Engrg. (1995).

4. S. Smale. Regular curves on Riemannian manifolds, Transactions of the American Mathematical Society 87:492-510, 1958.

5. S. A. Mitchell, T. J. Tautges. Pillowing doublets: refining a mesh to ensure that faces share at most one edge, *Proc. 4th International Meshing Roundtable*, to appear, 1995.

6. V. A. Gasilov, E. L. Kartashova, E. V. Sandrakova. Slicing doughnuts, report 1: problem outline, prepared for Sandia National Labs under contract, July 1995. Contact gasilov@imamod.msk.su or samitch@sandia.gov.

7. V. A. Gasilov. Personal communication, July 1995.

8. F. P. Preparata, M. I. Shamos. *Computational Geometry an Introduction,* Springer-Verlag, New York, 1985, pp. 24-26.

9. P. Rosenstiehl. Notes des membres et correspondants et notes présentées ou transmises par leurs soins, *Comptes Rendus* Acad. Sc. Paris, t. 283, Série A 551-553, October 1976.

10. C. F. Gauss. *Werke,* VIII, 1823, pp. 272, and 1844, pp. 282-286. Teubner Leipzig, 1900.

11. T. D. Blacker, S. A. Mitchell, T. J. Tautges, P. Murdoch, S. Benzley. Forming and resolving wedges in the spatial twist continuum, *Engineering with Computers,* to appear, 1995.

12. B. Grünbaum. Arrangements and Spreads, Reg. Conf. Ser. in Math. n° 10, Amer. Math Soc., 1972.

13. P. Rosenstiehl, R. E. Tarjan. Gauss codes, planar Hamiltonian graphs, and stack-sortable permutations, Journal of Algorithms 5:375-390, 1984.

14. D. R. White, L. Mingwu, S. E. Benzley, G. D. Sjaardema. Three dimensional meshing facilitators, *Proc. 4th International Meshing Roundtable*, to appear, 1995.

15. Algor Design World, December 1994, pp. 3. Promotional newsletter available from Algor.

16. T. D. Blacker, R. J. Meyers. Seams and wedges in Plastering: a 3-D hexahedral mesh generation algorithm, Engineering with Computers, 9:83-93, 1993.

17. R. Schneiders, http://www-users.informatik.rwth-aachen.de/~roberts/open.html.

18. R. P. Stanley. Enumerative Combinatorics, Volume 1, Wadsworth & Brooks/Cole Advanced Books and Software, Monterey, California, 1986.

19. *Proc. 3rd International Meshing Roundtable*, Sandia National Laboratories, Albuquerque, NM, October 24-25, 1994.