Conforming Hexahedral Mesh Generation via Geometric Capture Methods*

Jason F. Shepherd

Sandia National Laboratories
jfsheph@sandia.gov

Abstract. An algorithm is introduced for converting a non-conforming hexahedral mesh that is topologically equivalent and geometrically similar to a given geometry into a conforming mesh for the geometry. The procedure involves embedding geometric topology information into the given non-conforming base mesh and then converting the mesh to a fundamental hexahedral mesh. The procedure is extensible to multi-volume meshes with minor modification, and can also be utilized in a geometry-tolerant form (i.e., unwanted features within a solid geometry can be ignored with minor penalty). Utilizing an octree-type algorithm for producing the base mesh, it may be possible to show asymptotic convergence to a guaranteed closure state for meshes within the geometry, and because of the prevalence of these types of algorithms in parallel systems, the algorithm should be extensible to a parallel version with minor modification.

Keywords: Hexahedral, Mesh, Generation, Dual, Topology Modification.

1 Introduction

In this paper we explore more fully how to capture geometric boundary features using elements in the dual. We build on work completed in [1, 2, 3] and demonstrate conversion of non-geometry conforming hexahedral meshes to conforming hexahedral meshes followed by a mesh conversion to the fundamental mesh. The final quality of these meshes is largely dependent on the quality of the original non-conforming mesh; however, once a reasonable mesh is developed within a geometry it should be feasible to perform mesh conversions to optimize the structures within the mesh to improve geometric quality and mesh topology. The method described in this paper is extensible

* Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

to multi-volume meshes and geometry tolerant meshing paradigms. Additionally, because the concepts relating to fundamental meshes effectively localize mesh modification to geometric features, parallelizing the algorithm should be relatively straight-forward.

2 Background

The basic concepts that will be utilized throughout this paper utilize a 'dual' representation of a hexahedral mesh. The concept of the hexahedral mesh dual is foundational to many hexahedral mesh modification techniques that have been developed in recent years [4, 1, 5]. Defined by sheets and columns, the dual provides an alternate representation of a conforming hexahedral mesh. This alternate representation has supplied greater understanding about hexahedral mesh topology and has led to the creation of some basic mesh operations. Although the capture techniques introduced in this paper are novel, there is some similarity in the results produced by this method with grid-based methods. The reader is encouraged to review work in [6, 7, 8].

2.1 Dual Sheets and Columns

A hexahedral element contains three sets of four topologically parallel edges, as shown in Figure 1. Topologically parallel edges provide the basis for hexahedral sheets. The formation of a sheet begins with a single edge. Once an edge has been chosen, all elements which share that edge are identified. For each of these elements, the three edges which are topologically parallel to the original edge are also identified. These new edges are searched iteratively to find all connected hexahedra and the topologically parallel edges for each of these elements. This iterative procedure continues until no new adjacent elements are found. The set of all elements which are traversed during this process results in a layer of hexahedra, also known as a hexahedral sheet. Figure 2 shows a hexahedral mesh with a single hexahedral sheet highlighted.

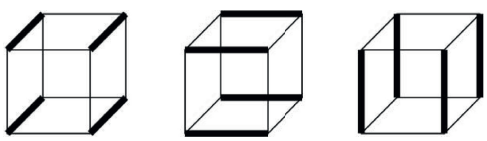

Fig. 1. A hexahedral element has three sets of four topologically parallel edges.

A hexahedral element also contains three pairs of topologically opposite quadrilateral faces, as shown in Figure 3. Topologically opposite faces provide the basis for hexahedral columns. The formation of a column begins with a single face. Once a face has been chosen, the two elements which share that face are identified. For each of these elements, the face which is topologically

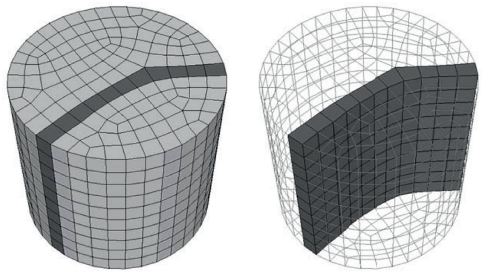

Fig. 2. A hexahedral mesh with one sheet highlighted.

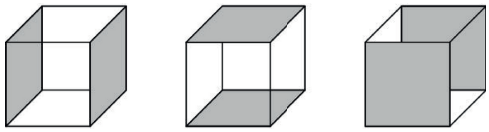

Fig. 3. A hexahedral element has three pairs of topologically opposite faces.

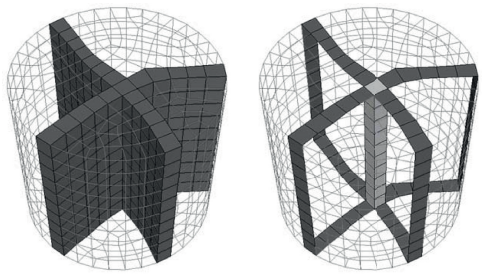

Fig. 4. The intersection of the two sheets shown on the left is defined by the column shown on the right.

opposite of the original face is also identified. These new faces are then used to find the incident hexahedra and topologically opposite faces on these adjacent elements. This process is repeated iteratively until no new adjacent hexahedra can be found. The set of all hexahedra which are traversed during this process makes up a hexahedral column. An important relationship between sheets and columns is that a column defines the intersection of two sheets. This relationship is illustrated in Figure 4.

2.2 Sheet Operations

The dual description of a hexahedral mesh is essentially an arrangement of surfaces satisfying specific criterion. It is, therefore, possible to modify an existing mesh simply by modifying the underlying arrangement of surfaces

describing the original mesh. The simplest form of modifying the mesh would involve adding or removing a surface from the arrangement of surfaces. This concept is often referred to as sheet insertion or sheet extraction [9]. Sheet extraction removes a sheet from a mesh by simply collapsing the edges that define the sheet, as shown in Figure 5.

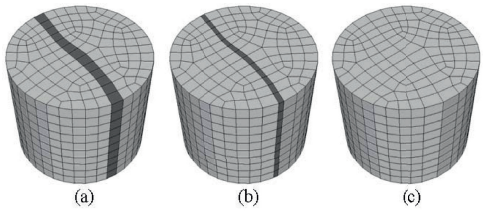

Fig. 5. Sheet extraction: (a) A sheet is selected. (b) The edges that define the sheet are collapsed. (c) The sheet is entirely removed from the mesh.

Using an inverse approach to sheet extraction, it is also possible to insert new sheets into an existing hexahedral mesh. The most common method for inserting a generalized sheet into a hexahedral mesh is pillowing [10]. Unlike sheet extraction, which removes an existing sheet from a mesh, pillowing inserts a new sheet into a mesh. As demonstrated in Figure 6, pillowing is performed on a set of hexahedral elements which make up a 'shrink' set. These elements are pulled away from the rest of the mesh and a new sheet is inserted by reconnecting each of the separated nodes with a new edge and creating new hexahedra utilizing all of the new created edges to fill in the gap. The new sheet surrounds the shrink set and maintains a conforming mesh.

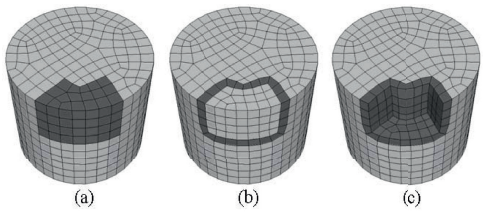

Fig. 6. Pillowing: (a) A shrink set is defined. (b) The shrink set is separated from the rest of the mesh and a sheet is inserted to fill in the gap. (c) The newly inserted pillow sheet.

2.3 Fundamental Hexahedral Meshes

Another concept of importance to the methods outlined in this paper is the notion of a fundamental hexahedral mesh. The definition of a fundamental

sheet is relative to the geometric object associated with the mesh [1]. The principle of a fundamental mesh is to have a single sheet for every surface, and a single chord for every curve on a surface (see examples in Figure 7). Such a mesh captures the geometry in such a way that every n-dimensional geometric cell is captured by one or many n-dimensional dual cells with particular restrictions. There is always at least one fundamental mesh for a given geometry, but the fundamental mesh definition allows for many different arrangements which satisfy the definition. If you translate this description in the primal mesh, you get the following definition given in [2].

Definition 2.1 *Let G and M be respectively a 3-dimensional geometric object and a hexahedral mesh. A hexahedral mesh M is a* **fundamental mesh** *with respect to G if and only if:*

1. *M is a strictly geometry-valid hexahedral mesh with respect to G;*
2. *For every geometric surface G_k^2, the number of hexahedral elements incident to G_k^2 is equal to the number of quadrilaterals classified on G_k^2;*
3. *For every curve G_k^1 and every surface $G_{k'}^2$, the number of quadrilaterals on $G_{k'}^2$ incident to edges on G_k^1 is equal to the number of edges classified on G_k^1.*

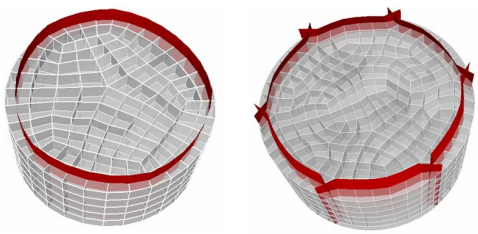

Fig. 7. Fundamental and non-fundamental meshes of a cylindrical geometric object. In the mesh on the right, multiple sheets are utilized to capture a single surface producing a non-fundamental mesh. The image on the left has a single sheet associated with the cylindrical surface and is fundamental.

Complementary explanations about this definition are given [2]. Note, for a geometric object, there is not one unique fundamental mesh, and, in fact, many geometry-valid hexahedral meshes will exist which are fundamental. These meshes are different for two reasons: there are many permutations of boundary sheets which satisfy the fundamental mesh requirements, and the number and configurations of non-boundary sheets within the hexahedral mesh is not restricted.

3 Theory and Assertions

In [1], the assertion is made that any non-fundamental hexahedral mesh can be converted to a fundamental hexahedral mesh. A proof of this assertion was

later given in [3]. In this paper, the goal is to show that a non-conforming hexahedral mesh whose boundary is topologically equivalent and geometrically similar to the composite boundary of a given solid geometry can be converted to a fundamental hexahedral mesh of the geometry and maintain reasonable quality metrics for the resulting hexahedra. We will discuss these two conjectures (i.e., conversion and quality) in more detail in this section.

3.1 Non-conforming Meshes to Fundamental Meshes

In [3], it was shown that a non-fundamental hexahedral mesh of a geometric object can always be converted to a fundamental hexahedral mesh of the geometric object. Therefore, it remains to be shown how to convert a non-conforming hexahedral mesh of a given object to a non-fundamental mesh of the same object.

The process of converting a non-conforming mesh to a conforming hexahedral mesh begins with two assumptions regarding the non-conforming mesh and the geometric object to be captured. The first assumption is that the boundary (composite) of the geometry and the boundary of the initial hexahedral mesh must be equivalent. That is, if the geometry is topologically spherical, then the boundary of the initial hexahedral mesh should also be spherical (i.e., it is not possible to convert a hexahedral mesh whose boundary is an n-toroid into a mesh of a spherical geometry). This assumption almost goes without saying, but depending on the method utilized to form the initial base mesh, may be commonly encountered.

The second assumption is more subjective. That is, the boundary of the initial base mesh should be geometrically similar to the boundary (composite) of the geometry to be captured. Satisfaction of this assumption is not binary (i.e., yes or no) like the topology equivalence assumption, and can involve a continuous range of satisfactory values. However, high geometric similarity between the base mesh and the geometry results in decreased modification to the base mesh and an associated higher probability that the quality of the resulting modified mesh will match the quality of the initial base mesh.

Given a base mesh that is non-conforming to the geometry, but topologically equivalent and geometrically similar, the process for converting the mesh to a conforming mesh involves the following steps:

1. Find an embedding of the geometric topology into base mesh.
2. Map the embedded geometric topology in the base mesh to the geometry

For the first item - finding an embedding of the geometric topology into the base mesh - we treat the boundary of the base mesh (e.g., the quadrilaterals, edges and nodes on the boundary of the mesh) as a graph (the mesh graph) and the geometric topology of boundary of the geometric object as a second graph (the geometric graph), and work to find an embedding of the second graph into the first graph. In some cases, this may require an enrichment of the the mesh graph when embedding the curves at high valent vertices

from the geometric graph. This enrichment can be accomplished utilizing a pillowing operation in the hexahedral mesh that will be described in the next section.

The second item - Map the embedded geometric topology in the base mesh to the geometry - is accomplished by finding an appropriate location for each of the nodes on the boundary of the base mesh to the appropriate geometric location determined from the embedding in the first step. At this point the mesh is now conforming to the geometry, although it likely has very poor quality (especially, near the new mesh boundary where most of the nodal movement took place in the last step). This quality will be improved during the conversion from the conforming, non-fundamental mesh to a fundamental mesh, and will be described in the next section.

3.2 Assertion on Hexahedral Quality

The second conjecture is in regards to the hexahedral element quality resulting from the conversion process outlined above. While it is very difficult to provide guarantees on potential hexahedral element quality, we refer back to an observation made in [11]. The ideal isotropic mesh contains perfectly planar sheets. Sheet curvature induces 'keystoning' of the hexahedral element (where the edges on one side of the hex are shorter than the opposite edges which are lengthened by the curvature). Sheets that do not intersect each other orthogonally induce element 'skewing' (see Figure 8). However, if only one of the three sheets defining the hex is subject to increased curvature or if there is only a single sheet with non-orthogonal intersections to the other sheets, then the feasible region for non-positive Jacobians remains large and it is likely that a smoothing or mesh optimization algorithm will be able to find a satisfactory nodal placement resulting in suitable hexahedral quality.

Fig. 8. Non-orthogonal intersections between sheets results in element 'skewing'.

We attempt to capitalize on this conjecture. That is, if the base mesh utilized consists of sheets with very low curvature and the sheets intersect each other with near orthogonality, then the sheets which are inserted during the conversion from non-fundamental to fundamental mesh will be the only sheets with any curvature or potential for non-orthogonality. If these inserted sheets are not interacting with each other, then the probability for creating a mesh which cannot be optimized to have reasonable quality is very low.

4 Algorithm

In this section, we describe an algorithm for performing the conversion of the non-conforming mesh to a conforming, non-fundamental mesh, and then finally to a fundamental mesh. We will also provide discussion on optimizing the mesh following the fundamental conversion. The algorithm described is presented for a generalized approach and various alterations can be made which limit generality but improve quality or desired mesh topology. These differences will be discussed in a later section.

The algorithm has four basic steps:

1. Establish the 2-manifolds for volumetric capture.
2. Convert the non-conforming mesh to a conforming mesh through geometric topology capture.
3. Convert the non-fundamental mesh to a fundamental mesh.
4. Mesh optimization to improve mesh quality.

4.1 Establishing 2-Manifolds

The first step in the algorithm is to capture a geometric solid from a pre-existing base mesh. We start by considering the boundary of the geometric solid as a closed, 2-manifold surface. We desire a base mesh whose boundary is topologically equivalent to the boundary of the geometric solid, as well as minimizing the geometric dissimilarities between these two boundaries. We only restrict this base mesh to topological equivalence, except to say, that the resulting quality of the final mesh will be heavily dependent on the geometric similarity between the base mesh and the geometric solid. Therefore, tailoring a base mesh (topologically and geometrically) to improve similarities between the base mesh and the geometric solid can provide dramatic quality differences in the final mesh.

There are several options for establishing the initial base mesh. Perhaps the easiest method is to create a structured grid in the bounding box representation of the solid geometry and eliminating hexahedra outside the boundary of the solid geometry (see Figure 9. This type of approach also allows base meshes to be created using standard octree-meshing techniques.

A similar approach can be utilized by creating simplified solid models and using standard hexahedral meshing algorithms on the simplified geometries and then eliminating hexahedra in the simplified geometry that do not match the original geometry (see example in Figure 10).

This approach can also be used to establish bounding 2-manifolds for multi-volume meshing. The only difference in multi-volume is to ensure that the topology of the all the volumetric base-meshes conform equivalent to the volumes in the solid model. Some initial work by Zhang, et al., demonstractes a method for establishing the base meshes for multi-volume biological models. A similar methodology can be utilized for solid models, although in some cases where the edges in the mesh do not have sufficient topology to match

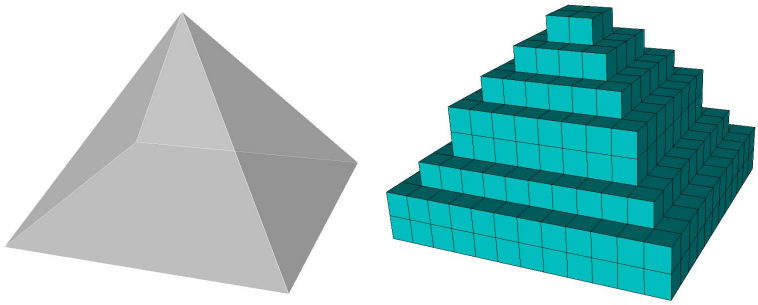

Fig. 9. An original geometry (on left) can be embedded in a regular mesh and the elements contained in the geometry is one alternative for defining a initial base mesh.

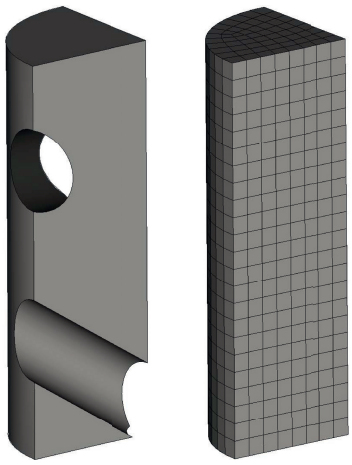

Fig. 10. A base mesh for an original geometry (on left) can also be created using standard hex primitive algorithms, including sweeping. The elements not contained in the mesh on the right can be removed and the resulting mesh would be suitable for a base mesh.

the curves in the solid model, a mesh enrichment step may be needed to provide the additional topology such that the base meshes are equivalent to the solid model.

4.2 Capturing Geometric Topology

Once the base mesh for the solid model has been established, the process of converting the non-conforming base mesh into a conforming base mesh can begin. Essentially, what we want to do in this step is to embed the curves and vertices from the boundary of the solid model into the boundary of the

base mesh. This also can be done in one of several ways. We will discuss one approach in this section.

Because the desire is to embed the geometric topology of the boundary into the base mesh, the logical first step is to first find an embedding of each of the geometric vertices into the boundary of the base mesh. In our case, we have done this by a geometric search, and find the closest node in the boundary of the base mesh to the geometric vertex. The number of edges emanating from each of the nodes is also taken into account and if the number of curves emanating from the geometric vertex is greater than the number of edges at the node and a nearby node captures the geometric topology more adequately, then a re-assignment may be made.

If the number of edges emanating from the closest node to a geometric vertex is fewer than the number of curves emanating from the vertex, a mesh enrichment procedure may be utilized to increase the nodal valence. Increasing the nodal valence consists of pillowing a collection of hexes in the neighborhood of the node resulting in additional edges emanating from the node (see Figure 11). This process can be repeated multiple times as needed to increase the nodal valence; however, care should be taken as possible, since the insertion of a small pillow will have reasonably high local curvature (the pillow in this case is essentially hemispherical). The curvature of the pillowed sheet may result in element quality reductions in the final mesh based on the conjecture discussed in Section 3.2.

Once an embedding of the geometric vertices is found and the nodal valence is equal to or greater than the vertex valence, we can begin to work on embedding each of the geometric curves into the mesh. If the quadrilateral boundary of the base mesh is treated as a graph, and the vertices have been embedded in this graph, then this problem can be viewed as similar to a collision-free network search. That is, we want to find a path between the embedded nodes which minimizes the geometric distance from the geometric curve and contains no collisions with the paths for each of the other curves. We demonstrate an example of this process in Figure 12.

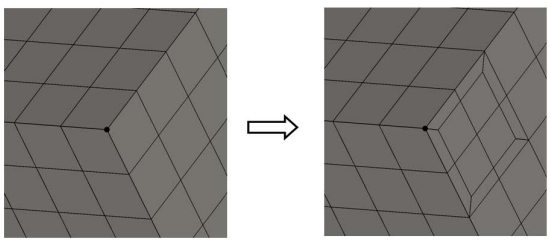

Fig. 11. Nodal valence can be increased using a simple pillowing operation. The original mesh is shown on the left and the resulting nodal valence increase after pillowing is displayed on the right.

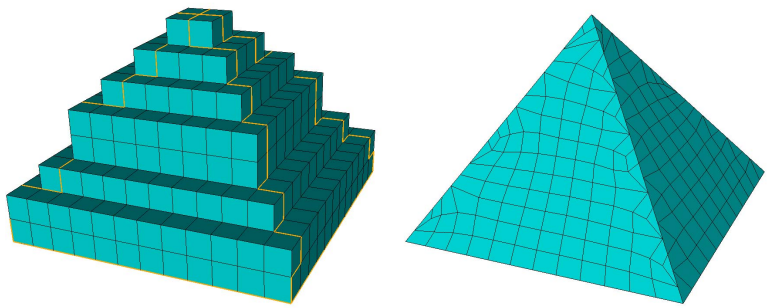

Fig. 12. The geometric graph is embedded into the base mesh using a conflict-free search of the mesh edges on the boundary of the base mesh to ensure topologic equivalence of the new embedding. Once the embedding has been established, the bounding mesh elements are 'snapped' to the appropriate geometry (image on right).

Once an embedding of the vertices and curves has been accomplished, the surface embedding is calculated by finding all of the quadrilaterals contained within the boundary of the embedded curves. This should complete the embedding of the geometric boundary into the boundary of the base mesh. At this point, all of the nodes on the boundary of the base mesh are moved to the correct geometric entity based on the previous embedding. This results in a conforming, but non-fundamental mesh of the geometry, from the previously non-conforming base mesh (see example in Figure 12).

4.3 Fundamental Conversions

Given a the embedding of the vertices and curves into the base mesh, we can begin the process of converting the base mesh to a fundamental mesh of the given geometry. A proof that this conversion can always occur is given in [2], and a construction of this proof can be utilized as follows in two steps:

1. Add a single sheet for each closed shell of boundary surfaces in the geometry. Assuming that the shell is manifold, this operation can be accomplished with a simple pillowing operation.
2. Add a single sheet for the collection of hexahedra contained by quadrilaterals associated with each surface of the geometry. Again, assuming that the collection of quadrilaterals is topologically equivalent to the associated geometric surface and the boundary of this collection of hexahedra if manifold, a simple pillowing operation will suffice for this sheet insertion. It should be noted that the first sheet insertion guarantees that there is a single hexahedron associated with each boundary quadrilateral, which is critical for this construction. Additionally, the embedding step described earlier can be used to guarantee topological equivalence of the quadrilateral collection with the geometric surface.

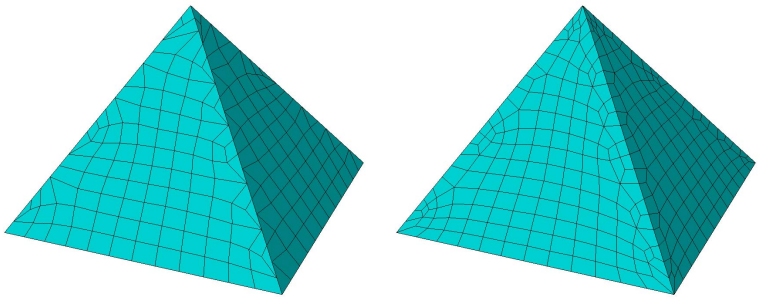

Fig. 13. The pyramid mesh following before (left) and after (right) conversion to a fundamental mesh.

The mesh following the fundamental conversion should be free of flattened (triangle-shaped) quadrilateral elements, and similar improvements will be noted in the interior hexahedra. The mesh for the pyramid example is shown in Figure 13 with the improved quality elements around the boundary.

This particular method for converting the mesh to fundamental is generalized, or in other words, it will work in all cases. However, it should be noted that fundamental meshes are not always required for satisfactory quality, and alternative manipulations of the mesh may be possible and desirable for improved quality. We will discuss this further in Section 6. Additionally, if the mesh in the second step is already fundamental, then no sheet insertion is required and the pillowing operation can be skipped.

4.4 Mesh Optimization

The conversion of the mesh to fundamental has the advantage of giving improved flexibility for mesh optimization near the boundary of the mesh. This is seen quite markedly in earlier papers [12, 1, 8], where dramatic improvements in the scaled Jacobian values are realized following the sheet insertion process and mesh smoothing/optimization. The pillowing process involves non-uniformly scaling elements in order to allow space for the newly inserted elements to occupy. This scaling often causes element inversions; however, the new topology introduced by the sheet insertion allows for greater flexibiliy by increasing the positive quality feasible regions for each of the elements near the boundary. In order to take advantage of this flexibility, mesh optimization algorithms with L2 and L-inf guarantees are recommended. In particular, we have heavily utilized the results of Knupp in mesh untangling [13], condition number optimization [14], and mean-ratio optimization [15, 16]. Additionally, it has been shown [17] that we can dramatically improve the speed of these algorithms by first utilizing a relaxation-based smoother (e.g. Laplacian smoothing), or using a focused-smoothing operation to reduce the number of elements being optimized. A recipe typically utilized following the conversion to fundamental is as follows:

1. Centroidal smoothing [18] each of the surfaces on the mesh boundary.
2. Untangling/optimization for each surface on the boundary (as needed).
3. Laplacian smoothing of the volumetric mesh.
4. Focused untangling of any pockets of mesh with inverted elements.
5. Focused L-inf optimization for any pockets of mesh elements with scaled Jacobian less than 0.2.
6. Additional mesh optimization as desired.

5 Examples

In this section, we demonstrate the process for one more example with slightly increased complexity (Figures 14 and 15). Following this example, we

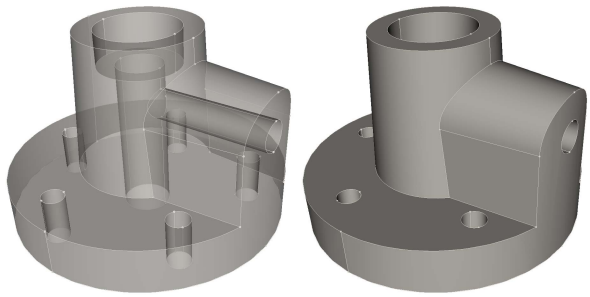

Fig. 14. The original 'sbase' geometric model shown in shaded and transparent modes. (Model provided courtesy of Ansys [19].)

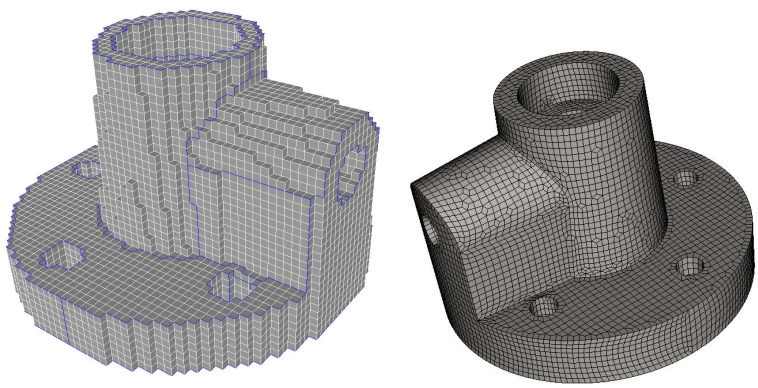

Fig. 15. A base mesh created for the 'sbase' geometry. The base mesh was created by creating a regular mesh in the bounding box of the geometry and eliminating all hexahedra located exterior of the geometry. This image also shows an embedding of the geometric graph in the boundary edges of the base mesh (the blue edges). The final mesh after conversion of the base mesh to a fundamental mesh is shown on the right.

Fig. 16. A mesh of a 6 generation airway of a human lung. The mesh on the right is a close-up view showing the sharp curves generated at the end caps of the model. (Model provided courtesy of Kwai Lam Wong, Oakridge National Laboratories.)

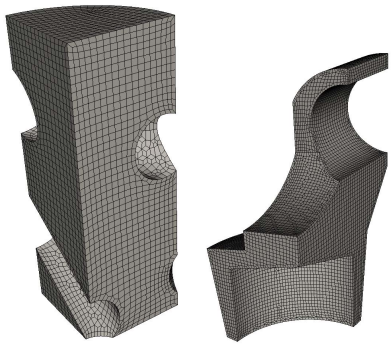

Fig. 17. Mesh of the valve model (Model courtesy of Kyle Merkley, Elemental Technologies, Inc.), and hook model (on right).

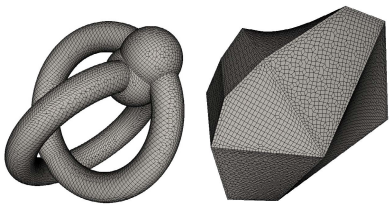

Fig. 18. Mesh of the 'a027' and 'ucp' models ('A027 Model courtesy of Ansys [19]).

demonstrate several other meshes that have been generated using this methodology (Figures 16, 17, and 18).

6 Alternate Methods

The method outlined in the previous sections of this paper are meant to be generalized. That is, these methods should work for all geometries with manifold boundaries. However, alternate conversions may produce better mesh quality, but may not be suitable for general cases. We will discuss some of these cases here, although it should be noted that the list provided below is not exhaustive, and additional methods may be developed and utilized.

6.1 Tri-valent Vertices

As indicated before, the method described in this paper is a generalized method. However, the price for generalization is decreased quality in some cases that could be improved with different sheet configurations that may not work generally. A tri-valent geometric vertex is a good example of this trade-off. In Figure 19(left) a mesh is shown that might be developed using the method outlined in this paper. At the trivalent vertex in the forefront of the image, the fundamental sheets are drawn (red, green, and yellow) for capturing the curves associated with this vertex. This configuration of sheets produces 2*v hexahedra, where v is the vertex valence. So, in this case, there are six hexahedra produced at this vertex location (i.e. the node at the vertex is contained in six hexahedra). An alternate sheet configuration (shown in Figure 19 (right) where the sheets are allowed to intersect one another is also shown. This configuration of sheets produces a single hexahedron at the vertex, and still fundamentally captures all of the curves associated with the

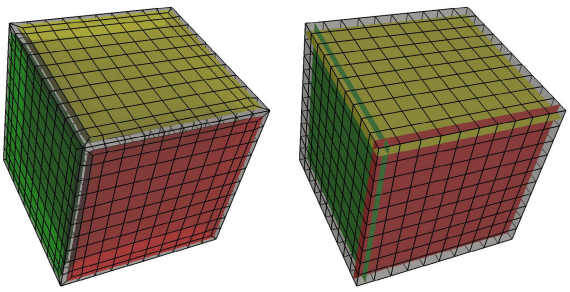

Fig. 19. Allowing changes in the sheet structure can improve the mesh quality. In the image on the left, the three sheets do not intersect resulting in six hexahedra at the vertex in the forefront of the image. On the right, the three sheets intersect and a single hexahedron is produced offering better opportunity for improved mesh quality at the vertex.

vertex. The resulting mesh has a higher potential for quality because of the change in the sheet topology.

Allowing the sheets to intersect one another is, in many cases, an improved strategy for inserting the sheets; however, the disadvantage to this alternate strategy is that it will only work if every vertex is trivalent. Therefore, applying the generalized algorithm is a better option for the original mesh generation algorithm, but it may be possible to 'clean-up' the mesh by re-arranging the sheet topology to improve the mesh quality locally at some of these tri-valent vertices. Additionally, it may be possible to develop similar recipes to reduce the element count and improve the quality potential while still maintaining a valid hexahedral mesh at each of the vertices by allowing some local sheet reconfigurations.

6.2 Regularizing Mesh Near Boundaries

Multiplying the number of sheets can also be utilized to increase the regularity of the mesh near the boundaries (see Figure 20). Balancing the arrangement of the sheets, can have a dramatic effect on the quality of the mesh. It will be advantageous to re-arrange the sheet topology following the initial mesh generation in order to improve regularity of the mesh and quality of the elements while still maintaining conformity with the geometry.

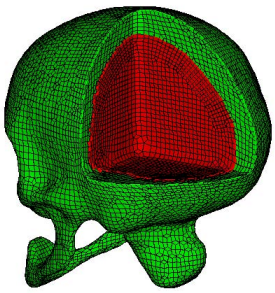

Fig. 20. Allowing for changes to sheet arrangements can improve the quality of the final mesh. In this image, several additional sheets were added near the interior cranial boundary to improve the regular structure of the mesh in this region. Other such additions, rearrangements, and removals of sheets may provide improved quality and topology of the final mesh.

7 Conclusion

We have outlined a method for building all-hexahedral meshes in arbitrary geometries. The procedure involves embedding geometric topology information into the given non-conforming base mesh and then converting the mesh to a fundamental hexahedral mesh. The procedure is extensible to multi-volume

meshes with minor modification, and can also be utilized in a geometry-tolerant form (i.e., unwanted features within a solid geometry can be ignored with minor penalty by the meshing procedure). Utilizing an octree-type algorithm for producing the base mesh, it may be possible to show asymptotic convergence to a guaranteed closure state for meshes within the geometry. Due to the prevalence of octree algorithms in parallel systems, the algorithm should also be extensible to a parallel version with minor modification.

References

1. Shepherd, J.F.: Topologic and Geometric Constraint-Based Hexahedral Mesh Generation. Published Doctoral Dissertation, University of Utah (May 2007)
2. Ledoux, F., Shepherd, J.: Topological and geometrical properties of hexahedral meshes. Submitted to Engineering with Computers (2008)
3. Ledoux, F., Shepherd, J.: Topological modifications of hexahedral meshes via sheet operations: A theoretical study. Submitted to Engineering with Computers (2008)
4. Murdoch, P.J., Benzley, S.E.: The spatial twist continuum. In: Proceedings, 4th International Meshing Roundtable, pp. 243–251. Sandia National Laboratories (October 1995)
5. Merkley, K., Ernst, C.D., Shepherd, J.F., Borden, M.J.: Methods and applications of generalized sheet insertion for hexahedral meshing. In: Proceedings, 16th International Meshing Roundtable, pp. 233–250. Sandia National Laboratories (September 2007)
6. Schneiders, R.: An algorithm for the generation of hexahedral element meshes based on an octree technique. In: Proceedings, 6th International Meshing Roundtable, pp. 183–194. Sandia National Laboratories (October 1997)
7. Shephard, M.S., Georges, M.K.: Three-dimensional mesh generation by finite octree technique. International Journal for Numerical Methods in Engineering 32, 709–749 (1991)
8. Zhang, Y., Hughes, T.J.R., Bajaj, C.: Automatic 3D meshing for a domain with multiple materials. In: Proceedings, 16th International Meshing Roundtable. Sandia National Laboratories (October 2007)
9. Borden, M.J., Benzley, S.E., Shepherd, J.F.: Coarsening and sheet extraction for all-hexahedral meshes. In: Proceedings, 11th International Meshing Roundtable, pp. 147–152. Sandia National Laboratories (September 2002)
10. Mitchell, S.A., Tautges, T.J.: Pillowing doublets: Refining a mesh to ensure that faces share at most one edge. In: Proceedings, 4th International Meshing Roundtable, pp. 231–240. Sandia National Laboratories (October 1995)
11. Shepherd, J., Johnson, C.: Hexahedral mesh generation constraints. Journal of Engineering with Computers (2007)
12. Shepherd, J.F., Zhang, Y., Tuttle, C., Silva, C.T.: Quality improvement and boolean-like cutting operations in hexahedral meshes. In: Proceedings, 10th Conference of the International Society of Grid Generation (September 2007)
13. Knupp, P.M.: Hexahedral mesh untangling and algebraic mesh quality metrics. In: Proceedings, 9th International Meshing Roundtable, pp. 173–183. Sandia National Laboratories (October 2000)

14. Knupp, P.M.: Achieving finite element mesh quality via optimization of the Jacobian matrix norm and associated quantities: Part II - a framework for volume mesh optimization and the condition number of the Jacobian matrix. International Journal for Numerical Methods in Engineering 48, 1165–1185 (2000)
15. Brewer, M., Freitag-Diachin, L., Knupp, P., Leurent, T., Melander, D.J.: The MESQUITE mesh quality improvement toolkit. In: Proceedings, 12th International Meshing Roundtable, pp. 239–250. Sandia National Laboratories (September 2003)
16. Diachin, L.F., Knupp, P., Munson, T., Shontz, S.: A comparison of inexact Newton and coordinate descent mesh optimization techniques. In: Proceedings, 13th International Meshing Roundtable, pp. 243–254. Sandia National Laboratories (September 2004)
17. Freitag, L.: On combining Laplacian and optimization-based mesh smoothing techniques. AMD Trends in Unstructured Mesh Generation, ASME 220, 37–43 (1997)
18. Jones, T.R., Durand, F., Desbrun, M.: Non-iterative, feature-preserving mesh smoothing. ACM Transactions on Graphics 22(3), 943–949 (2003)
19. ANSYS. ANSYS (January 2007), `http://www.ansys.com`