2008-07-28

Localized Coarsening of Conforming All-Hexahedral Meshes

Adam C. Woodbury
Brigham Young University - Provo

LOCALIZED COARSENING OF CONFORMING

ALL-HEXAHEDRAL MESHES

by

Adam C. Woodbury

A thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Master of Science

Department of Civil and Environmental Engineering

Brigham Young University

December 2008

BRIGHAM YOUNG UNIVERSITY

GRADUATE COMMITTEE APPROVAL

of a thesis submitted by

Adam C. Woodbury

This thesis has been read by each member of the following graduate committee and by majority vote has been found to be satisfactory.

Date	Steven E. Benzley, Chair

Date	Paul W. Richards

Date	Jason F. Shepherd

BRIGHAM YOUNG UNIVERSITY

As chair of the candidate's graduate committee, I have read the thesis of Adam C. Woodbury in its final form and have found that (1) its format, citations, and bibliographical style are consistent and acceptable and fulfill university and department style requirements; (2) its illustrative materials including figures, tables, and charts are in place; and (3) the final manuscript is satisfactory to the graduate committee and is ready for submission to the university library.

_____   _____
Date   Steven E. Benzley
                                          Chair, Graduate Committee

Accepted for the Department

                                          _____
                                          E. James Nelson
                                          Graduate Coordinator

Accepted for the College

                                          _____
                                          Alan R. Parkinson
                                          Dean, Ira A. Fulton College of Engineering
                                          and Technology

ABSTRACT

LOCALIZED COARSENING OF CONFORMING

ALL-HEXAHEDRAL MESHES

Adam C. Woodbury

Department of Civil and Environmental Engineering

Master of Science

Finite element mesh adaptation methods can be used to improve the efficiency and accuracy of solutions to computational modeling problems. For many finite element modeling applications, a conforming all-hexahedral mesh is preferred. When adapting a hexahedral mesh, localized modifications that preserve topologic conformity are often desired. Effective hexahedral refinement methods that satisfy these criteria have recently become available. However, due to hexahedral mesh topology constraints, little progress has been made in the area of hexahedral coarsening. This thesis presents a new method to locally coarsen conforming all-hexahedral meshes. The method works on both structured and unstructured meshes and is not based on undoing previous refinement. Building upon recent developments in quadrilateral coarsening, the method utilizes hexahedral sheet and column operations, including pillowing, column collapsing, and sheet extraction. A general algorithm for automated coarsening is presented and examples of models that

have been coarsened are shown. While results are promising, further work is needed to improve the automated process.

ACKNOWLEDGMENTS

TABLE OF CONTENTS

LIST OF TABLES

x

LIST OF FIGURES

1   Introduction

The efficiency and accuracy of a finite element computational modeling solution are greatly influenced by the distribution of elements in the finite element mesh. In a given model, there are usually regions that require greater mesh density than others. A higher concentration of elements in these regions may be necessary to reduce error in the finite element approximation, increase resolution where there are high gradients, or more accurately represent the model geometry. Regions where high accuracy is not critical or where gradients are low can generally be modeled with lower mesh density. Since the computational time required in a finite element analysis is directly related to the number of elements in the model being analyzed, it is advantageous to produce a mesh that has as few elements as possible. Therefore, in an ideal analysis, each region in the model should have enough elements to produce a good solution, but no more.

For most real-world finite element models, current mesh generation algorithms are unable to create an initial mesh that optimizes both accuracy and efficiency in the finite element solution. Although some control over mesh density is possible, an initial mesh will almost always contain regions that have too few elements, regions that have too many elements, or both. In addition, some finite element applications require mesh density to evolve throughout an analysis (see Figure 1-1) as areas of high and low activity change with time [1, 2, 3, 4]. For these reasons, much research has been devoted to the

development of mesh adaptation tools that make it possible to adjust element density in specific regions either before or during analysis.

Figure 1-1. Evolving mesh density in a triangle mesh.

Mesh adaptation consists of both refinement and coarsening. Refinement is the process of increasing mesh density by adding elements to a mesh, while coarsening is the process of decreasing mesh density by removing elements from a mesh, as illustrated in Figure 1-2. By refining areas that have too few elements and coarsening areas that have too many elements, a more accurate and efficient analysis can be performed.

Both tetrahedral and hexahedral elements are commonly used in three-dimensional finite element analyses. For complex model geometries, tetrahedral meshes are much easier to create than hexahedral meshes. In addition, localized modifications are generally much more straightforward in a tetrahedral mesh. However, in many modeling
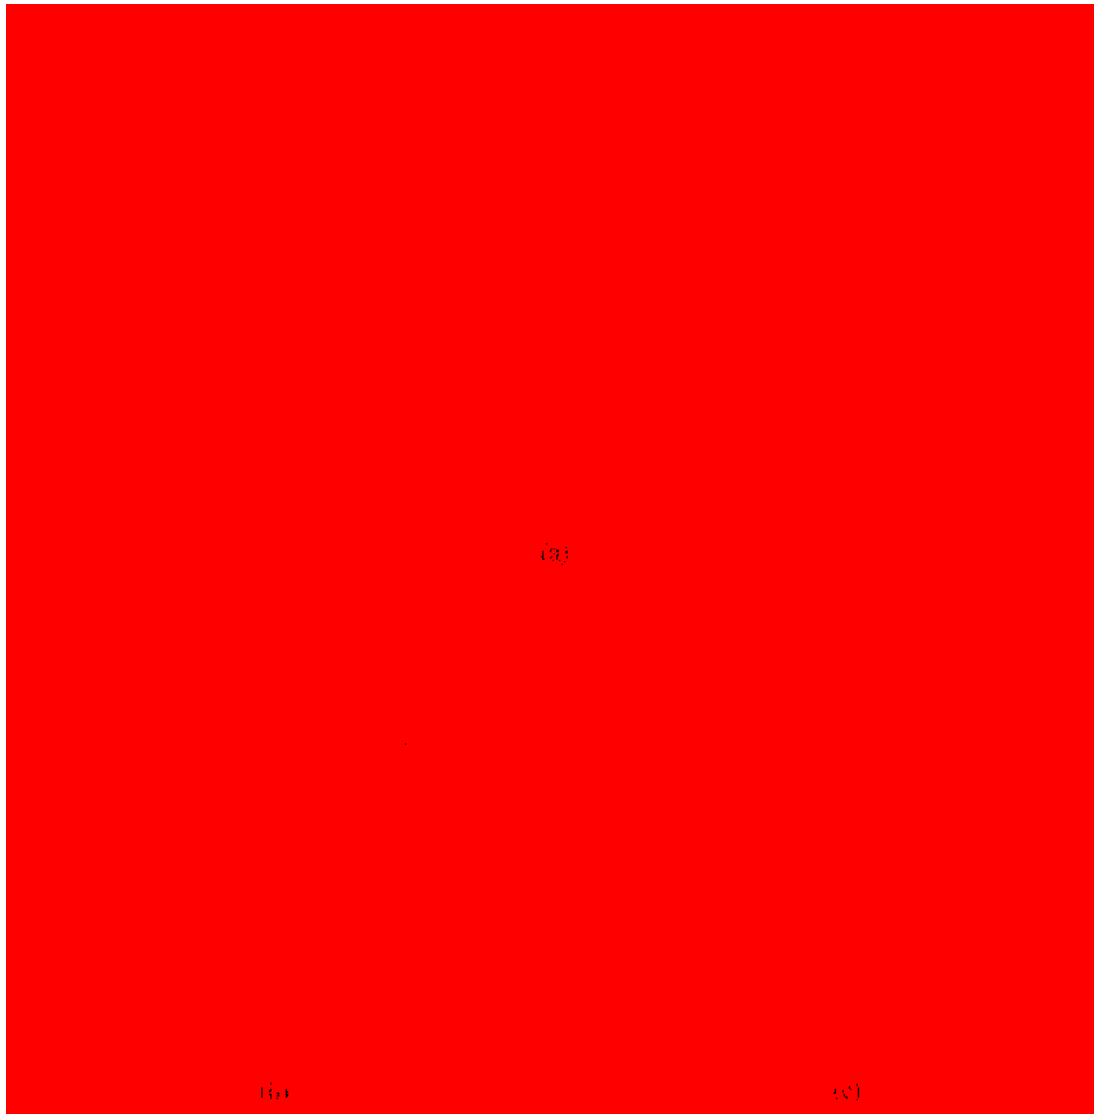
Figure 1-2. Mesh adaptation: (a) Original hexahedral mesh. (b) Refined mesh. (c) Coarsened mesh.

applications, hexahedral elements are preferred over tetrahedral elements because they provide greater efficiency and accuracy in the computational solution [1, 5]. For this reason, work has been done to improve hexahedral mesh adaptation methods. As a result, robust hexahedral refinement algorithms are becoming available [6, 7, 8, 9]. However, few developments have been seen in the area of hexahedral coarsening. The lack of

effective coarsening methods creates a major gap in the field of hexahedral mesh adaptation.

To satisfy the requirements of some finite element solvers and to effectively achieve the objectives of mesh adaptation, a truly general hexahedral coarsening algorithm should:

1. Preserve a conforming all-hexahedral mesh

2. Restrict mesh topology and density changes to defined regions

3. Work on both a structured and unstructured mesh (see Figure 1-3)

4. Not be limited to undoing previous refinement
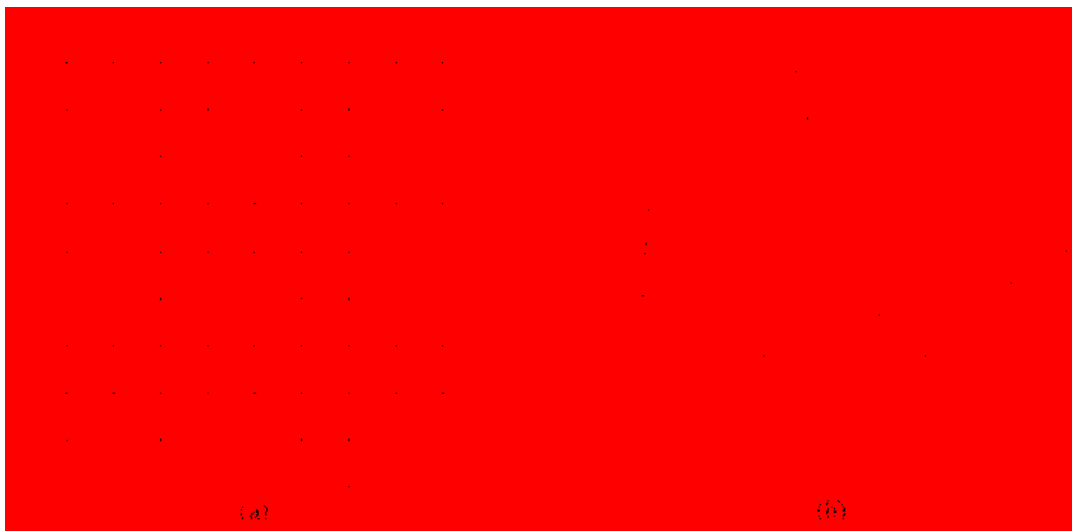
Figure 1-3. Two-dimensional structured and unstructured meshes: (a) A structured quadrilateral mesh where each interior node is shared by exactly four elements. (b) An unstructured quadrilateral mesh where interior nodes are shared by varying numbers of elements.

A mesh is conforming if all of the nodes and edges in the mesh are shared by each of their adjacent elements. Unlike a hybrid mesh, which contains more than one element

4

type, an all-hexahedral mesh contains only hexahedral elements. Figure 1-4 shows some two-dimensional examples of a non-conforming and hybrid mesh.
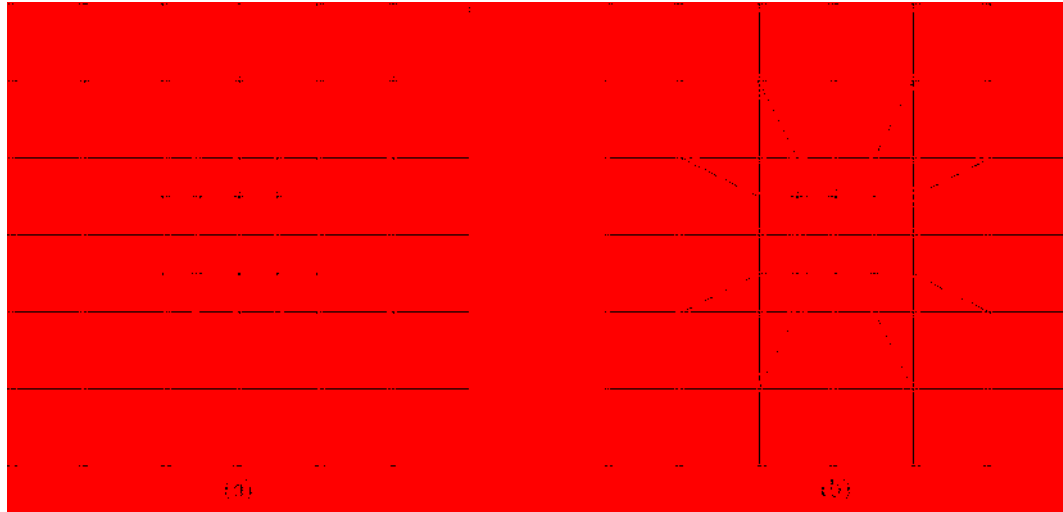
Figure 1-4. Two-dimensional non-conforming and hybrid meshes: (a) A non-conforming quadrilateral mesh. (b) A hybrid mesh with both quadrilateral and triangle elements.

Although hexahedral coarsening has been utilized in some modeling applications, no single algorithm has been developed that satisfies all the criteria listed above. This is, in large part, due to the topology constraints that exist in a conforming all-hexahedral mesh. These constraints make it difficult to modify mesh density without causing topology changes to propagate beyond the boundaries of a defined region [10, 11].

Since current hexahedral coarsening methods are unable to satisfy all the requirements listed above, they have limited applications. For example, to prevent global topology changes, some algorithms introduce non-conforming or non-hexahedral elements into the mesh [1, 2, 12, 13, 14]. While this is a valid solution for some types of analysis, not all finite element solvers can accommodate non-conforming or hybrid meshes. Other algorithms maintain a conforming all-hexahedral mesh, but they generally

require either global topology changes beyond the defined coarsening region [11, 15, 16], structured mesh topology where predetermined transition templates can be used [8, 15], or prior refinement that can be undone [2, 12, 13]. These weaknesses severely limit the effectiveness of these algorithms on most real-world models.

This thesis presents a new method to locally coarsen conforming all-hexahedral meshes. The method works on both structured and unstructured meshes and is not based on undoing previous refinement. The remainder of this thesis is organized as follows. Chapter 2 provides an overview of some basic hexahedral mesh operations. Chapter 3 shows how these operations have been combined to produce localized hexahedral coarsening and how the coarsening process has been automated. In Chapter 4, some examples of models which have been coarsened are shown. Finally, in Chapter 5, some areas of future work are discussed.

2   Hexahedral Mesh Operations

In recent years, a greater understanding of hexahedral mesh topology has led to the development of many new hexahedral mesh operations [17, 18, 19]. In this chapter, three operations which are useful for hexahedral coarsening are presented. These operations are based on hexahedral sheets and columns, which are topology-based groups of hexahedra that always exist in a conforming hexahedral mesh.

2.1   Hexahedral Sheets and Columns

A perfectly shaped hexahedral element contains three sets of four parallel edges, as shown in Figure 2-1. Regardless of its shape, a hexahedral element will always have the same topology. For this reason, it is convenient to describe the four edges in each set as being topologically parallel, even if they are not geometrically parallel. 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 then used to find another layer of elements and topologically parallel edges. This process is repeated until no new adjacent elements can be found. The set of elements which are traversed during this

process makes up a hexahedral sheet. Figure 2-2 shows a hexahedral mesh with one of the sheets in the mesh defined.
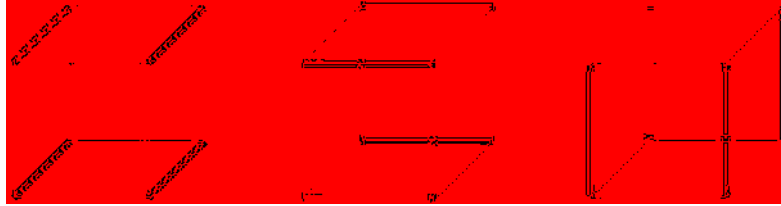
Figure 2-1. A hexahedral element's three sets of topologically parallel edges.
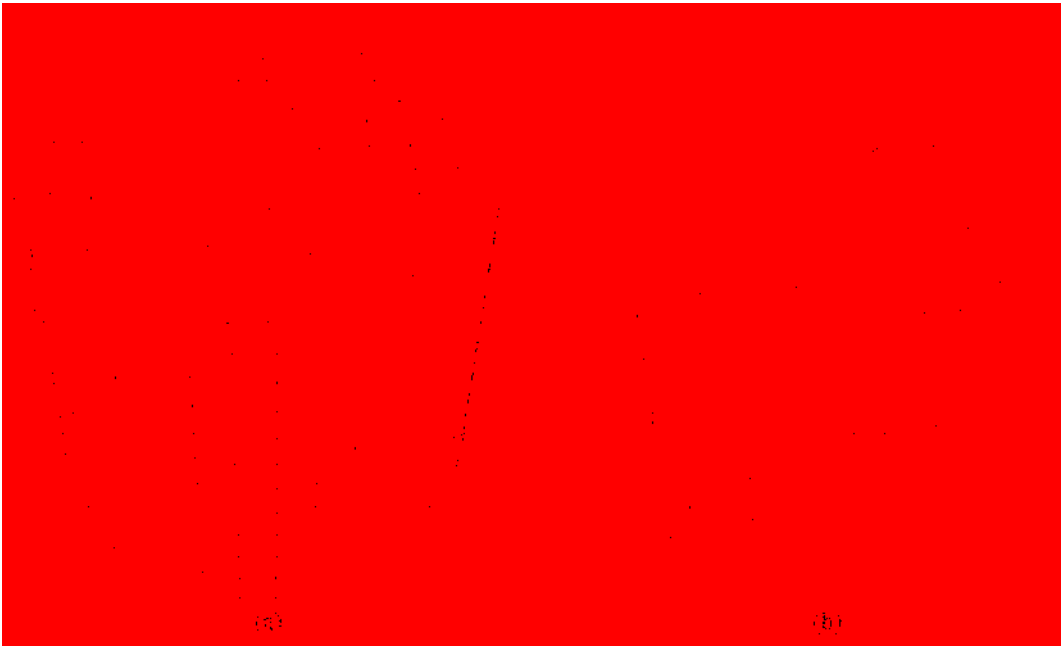
Figure 2-2. A hexahedral sheet: (a) A hexahedral mesh with one sheet defined. (b) A view of the entire sheet.

A hexahedral element also contains three pairs of topologically opposite quadrilateral faces, as shown in Figure 2-3. 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 elements which share that face are identified. For each of

8

these elements, the face which is topologically opposite of the original face is also identified. These new faces are then used to find another layer of elements and topologically opposite faces. This process is repeated until no new adjacent elements can be found. The set of elements 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 2-4.
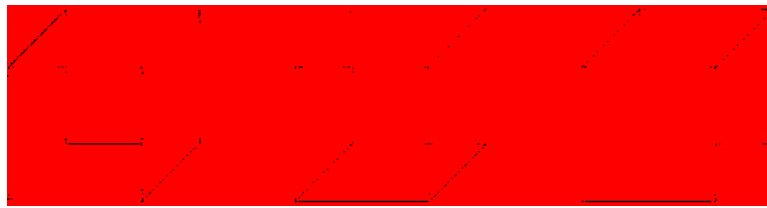
Figure 2-3. A hexahedral element's three pairs of topologically opposite faces.
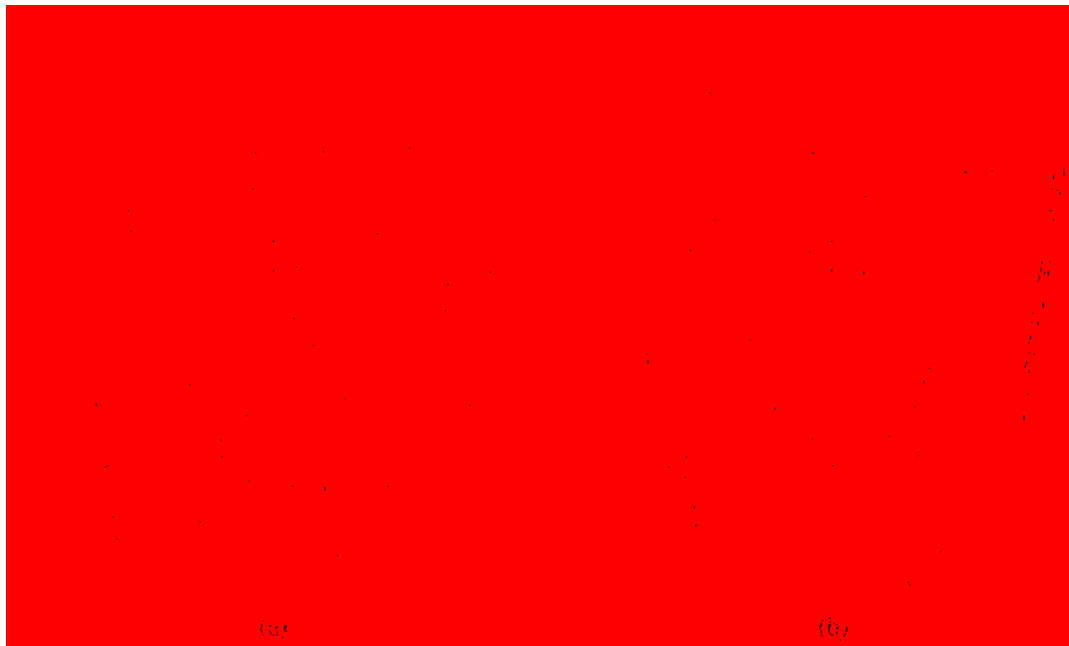
Figure 2-4. A hexahedral column: (a) Two intersecting sheets. (b) The column that defines the intersection of the two sheets in (a).

2.2   Sheet and Column Operations

Hexahedral sheet and column operations can be used to modify a hexahedral mesh without introducing non-conforming elements. One such operation is known as sheet extraction [16]. Sheet extraction removes a sheet from a mesh by simply collapsing the edges that define the sheet and merging the two nodes on each edge, as shown in Figure 2-5. Merging nodes in this manner decreases element density in the vicinity of the extracted sheet and guarantees that the resulting mesh will be conforming.

Figure 2-5. Sheet extraction: (a) A sheet is selected for extraction. (b) The edges that define the sheet are collapsed. (c) The two nodes on each edge are merged, which eliminates the sheet and preserves a conforming hexahedral mesh.

Another hexahedral mesh operation that involves sheets is pillowing [19, 20]. Unlike sheet extraction, which removes an existing sheet from a mesh, pillowing inserts a new sheet into a mesh. As demonstrated in Figure 2-6, pillowing is performed on a contiguous group of hexahedral elements which make up a "shrink" set. These elements are reduced in size and pulled away from the rest of the mesh, leaving a gap. A new sheet is then inserted into the gap by reconnecting each of the separated node pairs with a new

10

edge. The new sheet increases element density in the vicinity of the shrink set and ensures the preservation of a conforming hexahedral mesh.
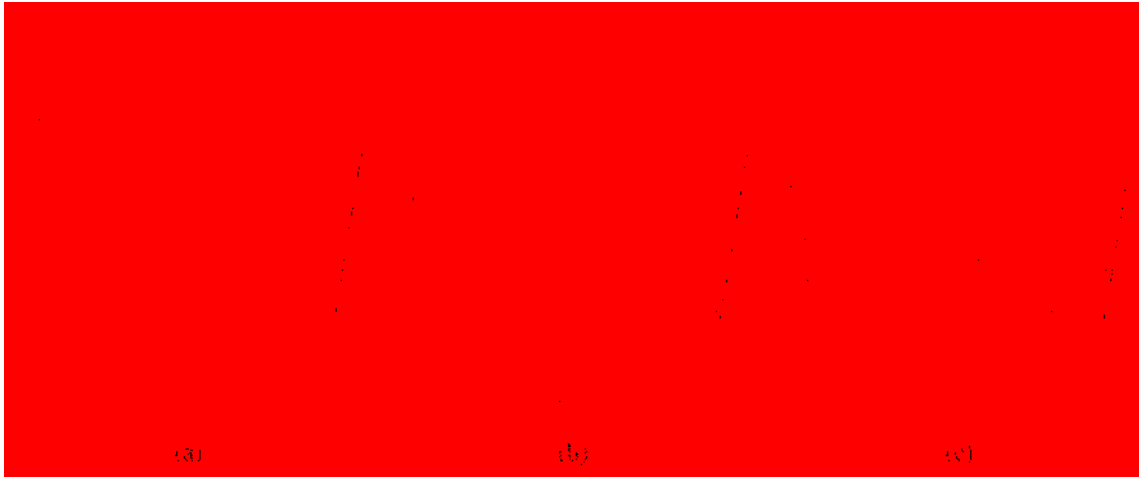
Figure 2-6. Pillowing: (a) A shrink set is defined. (b) The elements in the shrink set are reduced in size and separated from the rest of the mesh. A sheet is inserted to fill in the gap and preserve a conforming hexahedral mesh. (c) The newly inserted sheet.

A third hexahedral mesh operation is known as column collapsing [21]. A column is collapsed by merging diagonally opposite nodes in each quadrilateral face that defines the column, as shown in Figure 2-7. Since a quadrilateral face has two pairs of diagonally opposite nodes, a column can be collapsed in one of two different directions.

Figure 2-7. Column collapse operation.

11

As previously mentioned, a column defines the intersection of two sheets. When a column is collapsed, two intersecting sheets are altered such that they no longer intersect, as illustrated in Figure 2-8. The paths of the new sheets are determined by the direction of the collapse. Just like sheet extraction and pillowing, the column collapse operation always preserves a conforming hexahedral mesh. In addition, similar to sheet extraction, the column collapse operation decreases element density in the vicinity of the collapsed column.

Figure 2-8. Redirection of intersecting sheets through column collapsing: (a) Two intersecting sheets. (b) The column defining the intersection is collapsed. (c) The two sheets no longer intersect.

3 Hexahedral Mesh Coarsening

The hexahedral coarsening method presented in this chapter was discovered after many failed attempts to extend various quadrilateral coarsening methods to a hexahedral mesh. Not surprisingly, three-dimensional mesh modifications are much more difficult to localize than two-dimensional mesh modifications. While it is true that some quadrilateral coarsening operations can be directly extended to hexahedral coarsening, by themselves, these operations are not always able to prevent changes in element density from propagating beyond the boundaries of a defined hexahedral coarsening region.

Utilizing the sheet and column operations described in Chapter 2, the hexahedral coarsening method presented here builds upon recent developments in quadrilateral coarsening [21, 22]. Since two-dimensional coarsening operations are generally easier to visualize than their three-dimensional counterparts, an explanation of the related quadrilateral coarsening operations is provided in Appendix A. However, as will be seen in this chapter, entirely localized hexahedral coarsening often requires an additional step that is not necessary in quadrilateral coarsening.

3.1 Previously Developed Coarsening Techniques

As illustrated in Chapter 2, sheet extraction decreases mesh density by removing elements from a mesh. Therefore, sheet extraction is a very useful tool for hexahedral

coarsening. However, sheet extraction by itself is generally not sufficient when localized coarsening is desired. This is due to the fact that sheets are rarely contained entirely within a region that has been selected for coarsening. As shown in Figure 3-1, extracting a sheet that extends beyond the boundaries of a defined region decreases mesh density in areas where coarsening is not desired. Therefore, before sheet extraction can occur, it is often necessary to modify a mesh in such a way that produces sheets which are contained entirely within the boundaries of a defined coarsening region.
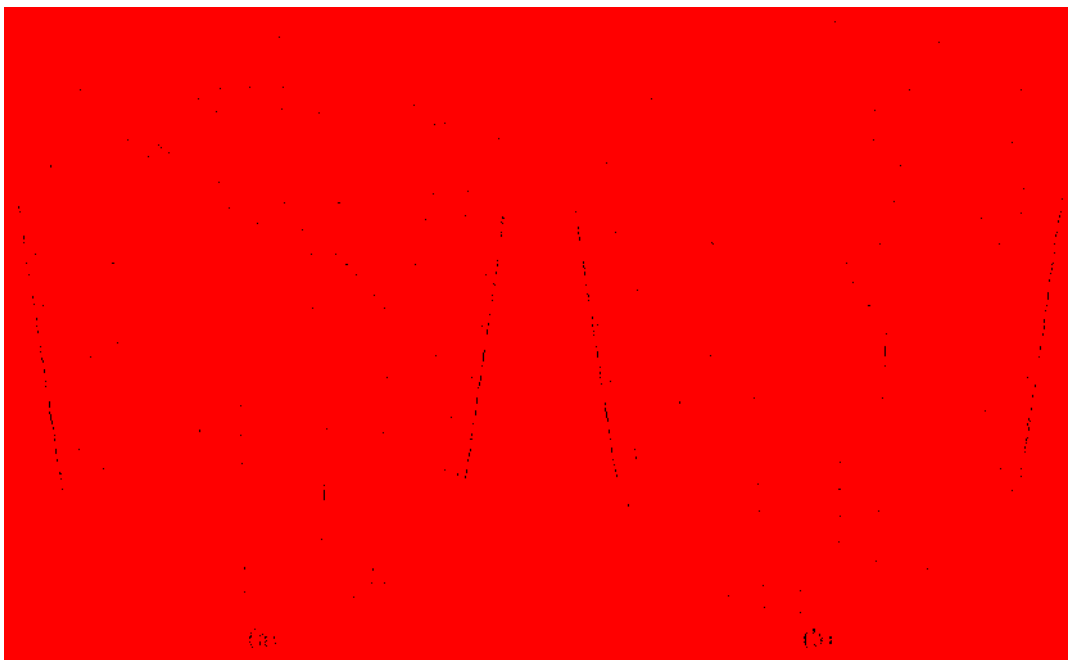
Figure 3-1. Global coarsening: (a) A sheet passes through a region selected for coarsening. (b) When the sheet is extracted, mesh density is decreased both inside and outside the defined coarsening region.

As described in Chapter 2, the paths of intersecting sheets can be altered using the column collapse operation. Figure 3-2 shows how this operation can be used to create a

sheet that is contained entirely within a defined region. Such a sheet can then be extracted

to coarsen the region without affecting any other part of the mesh.
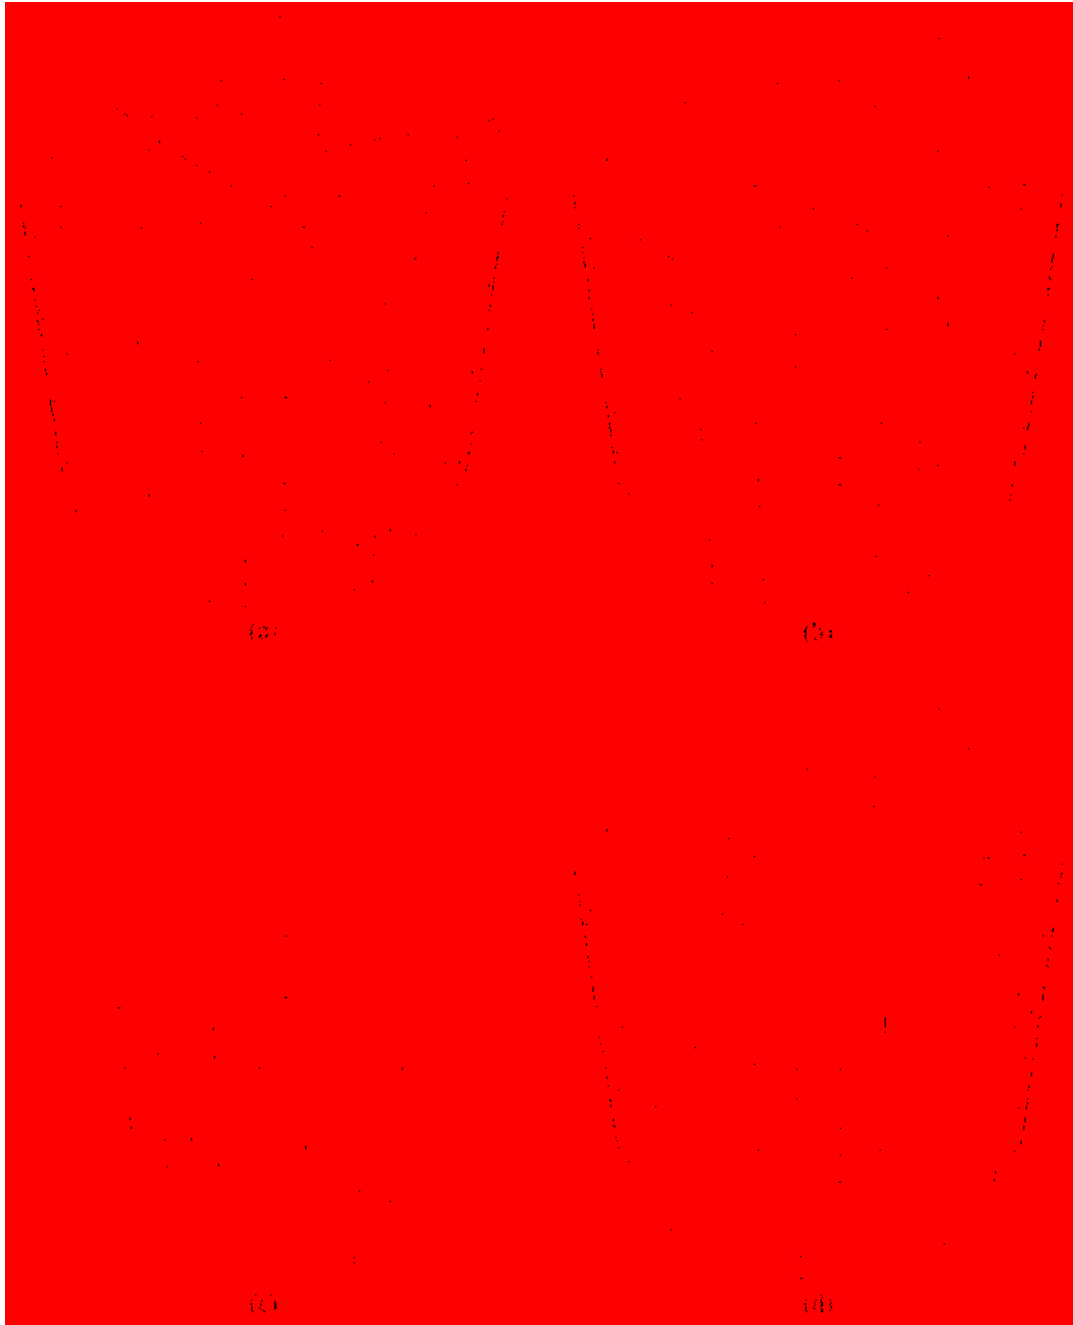
Figure 3-2. Localized coarsening: (a) Two intersecting sheets pass through a region selected for coarsening. (b) The column defining the intersection of the two sheets in (a) is collapsed to produce a sheet contained entirely within the coarsening region. (c) The sheet that will be extracted. (d) When the sheet in (c) is extracted, mesh density is only decreased within the defined coarsening region.

15

The coarsening region shown in Figure 3-2 extends from the top to the bottom of the mesh. Suppose the coarsening region is modified so that it only extends a few layers from the top of the mesh, as shown in Figure 3-3. In this case, the column collapse operation can be used twice to produce a sheet that is contained entirely within the coarsening region. However, as seen in the figure, the first collapse operation is performed on a column which extends beyond the boundaries of the region. Collapsing this column modifies mesh topology and density in areas where coarsening is not desired. This shows that entirely localized coarsening cannot always be accomplished with the column collapse and sheet extraction operations alone.

3.2   Entirely Localized Coarsening

The previous examples demonstrate that entirely localized coarsening requires all operations to take place within the boundaries of the selected coarsening region. Referring to Figure 3-3, it can be seen that the second collapse operation was performed on a column contained within the coarsening region. Collapsing this column produced a sheet contained within the region without affecting any other part of the mesh. Of course, the formation of this column was accomplished through a previous collapse operation that did affect areas outside the coarsening region. Therefore, a critical aspect of entirely localized coarsening is the creation of local columns. Such columns must be formed in the coarsening region without affecting areas outside the region.

One way to create columns that do not extend beyond the boundaries of the coarsening region without affecting areas outside the region is through pillowing. As illustrated in Chapter 2, pillowing is a form of refinement because it increases mesh
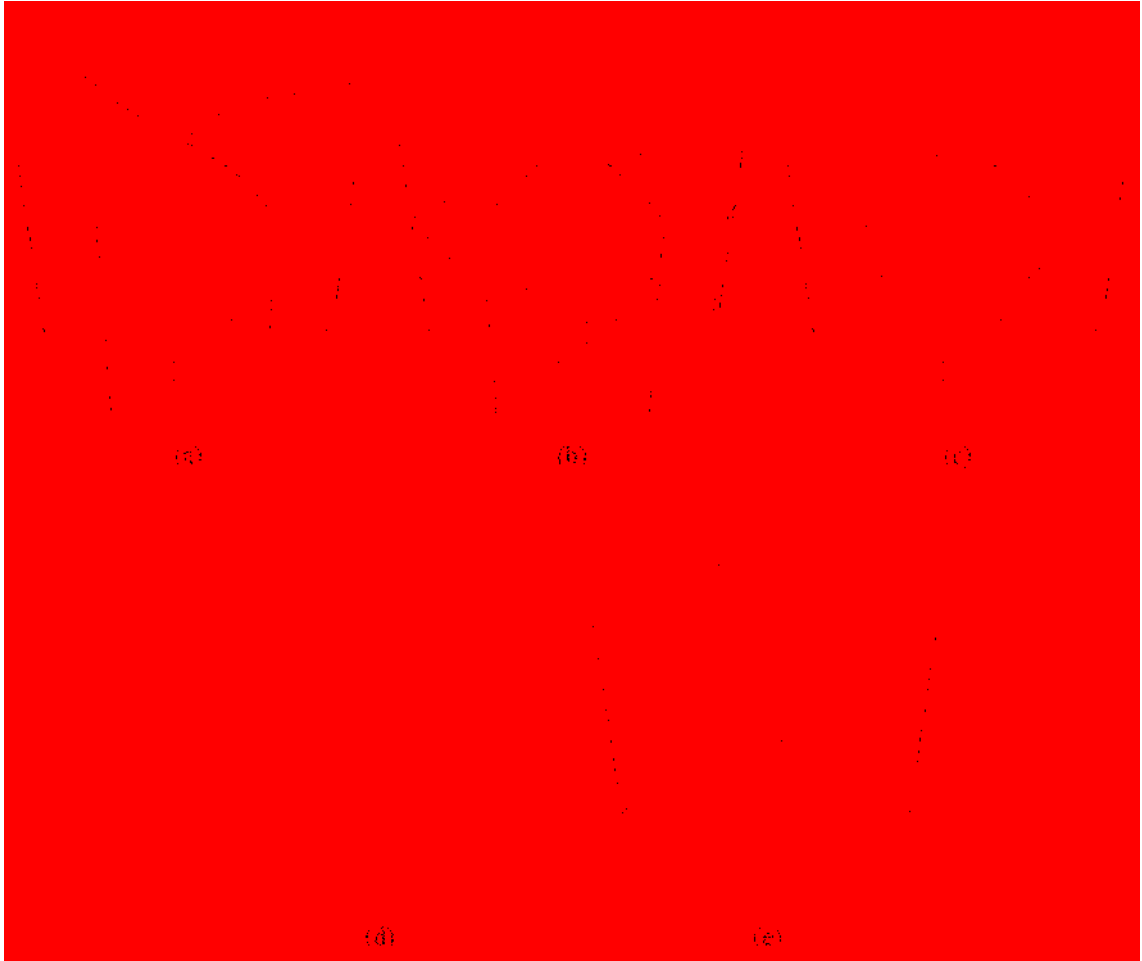
Figure 3-3. Semi-localized coarsening: (a) Two intersecting sheets pass through a region selected for coarsening. (b) The column defining the intersection of the two sheets in (a) is collapsed. A sheet formed by the collapse and another intersecting sheet are shown. (c) The column defining the intersection of the two sheets in (b) is collapsed to produce a sheet contained entirely within the coarsening region. (d) The sheet that will be extracted. (e) When the sheet in (d) is extracted, mesh density is only decreased within the defined coarsening region.

density in the vicinity of the shrink set. For this reason, pillowing is not an obvious solution for coarsening. However, due to the topology constraints that exist in a conforming all-hexahedral mesh, adding elements appears to be a necessary step when coarsening some regions.

Figure 3-4 shows how pillowing can be used to produce entirely localized coarsening. By pillowing the coarsening region, a sheet is inserted around the region.

17

This sheet intersects other sheets that pass through the coarsening region and provides columns which follow the boundary of the region. Such columns can be collapsed to form sheets contained within the coarsening region without modifying mesh topology or density in areas where coarsening is not desired. These sheets can then be extracted to locally coarsen the region. It should be noted that many of the elements added through pillowing are removed through sheet extraction. Only those elements which are necessary to transition from higher to lower mesh density are left in the mesh. As long as the number of elements removed through sheet extraction is greater than the number of elements added through pillowing, the final mesh density in the coarsening region will be lower than the initial mesh density.
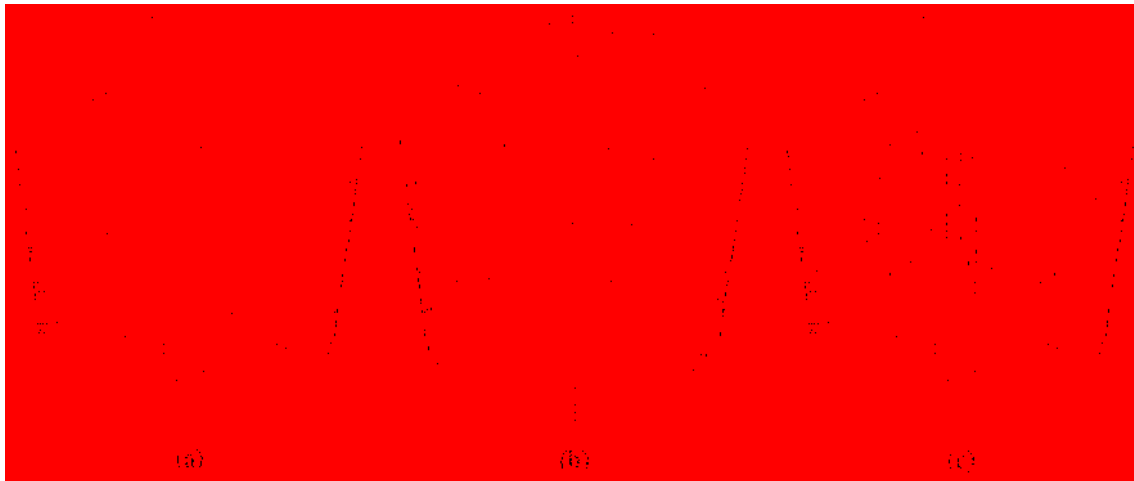
Figure 3-4. Entirely localized coarsening: (a) A coarsening region is defined. (b) The sheet that forms when the coarsening region is pillowed. This sheet provides columns which follow the boundary of the region. (c) Collapsing the columns in (b) produces sheets contained entirely within the coarsening region.

For a given region, the process of pillowing, column collapsing, and sheet extraction can be repeated multiple times to achieve various levels of coarsening. Each

iteration of this process requires a careful selection of column collapse and sheet extraction operations that ultimately lead to the desired mesh density. Depending on the size and structure of the region, it may not be possible to achieve an exact target mesh density, but a rough approximation is usually attainable.

3.3 Automated Coarsening Algorithm

To have any real value in finite element modeling applications, the coarsening process which has been presented would have to be performed automatically, without the need for a person to decide which columns to collapse, what direction to collapse them, etc. Fortunately, automation of this process is possible. A large part of the research which has been performed in relation to this new coarsening method has been devoted to the development of an algorithm which fully automates the coarsening process for a user defined region and level of coarsening. Implementing this algorithm has not only shown that automation is possible, but has led to more efficient discoveries of areas that need improvement.

Automation of the coarsening process has proven to be a fairly complex task. Much of the difficulty arises from three competing objectives that must each receive attention. The first objective is to achieve a specified mesh density within the region selected for coarsening. The second objective is to maintain high element shape quality in the region. The third objective is to complete the coarsening process in as little time as possible. Obviously, not all of these objectives can be achieved simultaneously. While consistent improvement has been made, finding an appropriate balance between these objectives continues to present significant challenges.

As already mentioned, achieving an exact element density may not be possible for some regions. Even when it is possible, it may not be feasible when trying to achieve the other two objectives. For this reason, a mesh density tolerance has been implemented into the automated coarsening algorithm. By allowing a final mesh density that is within a certain tolerance of the desired mesh density, some flexibility in meeting the other objectives is provided.

As with other mesh adaptation procedures, an unavoidable consequence of coarsening is a reduction in element shape quality. In general, the more coarsening a region undergoes, the lower the element quality in that region becomes. Usually, a region which has been selected for coarsening is in a less critical part of the model and can tolerate lower quality elements. However, even the least critical parts of a model must maintain an element quality above an acceptable threshold.

One way to improve element quality in a given region is through smoothing [23]. Smoothing is a process of optimizing the shape of each element by moving nodes to new locations without altering element connectivity. Although the mesh operations involved in hexahedral coarsening preserve a conforming mesh, they often leave behind very poorly shaped elements. As shown in Figure 3-5, smoothing a region which has been coarsened can dramatically improve the overall element quality in the region. While smoothing is an essential component of the automated coarsening algorithm, it is not guaranteed to produce acceptable element quality for any given mesh topology. Therefore, additional measures to improve element quality prior to smoothing have been implemented into the algorithm.
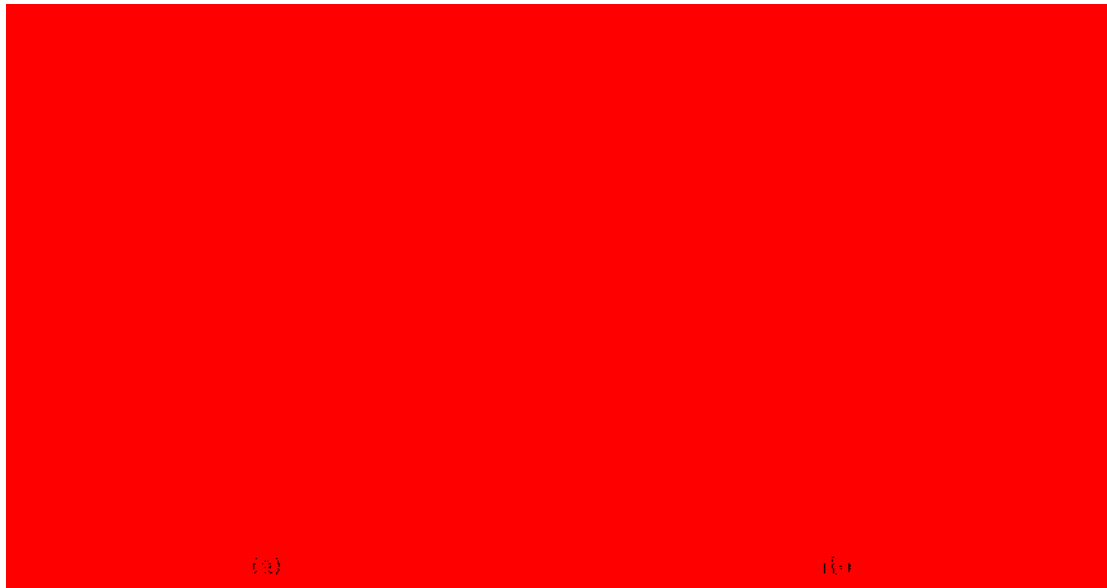
Figure 3-5. Smoothing: (a) A coarsened mesh that has not been smoothed. (b) The mesh in (a) after smoothing has taken place.

Before smoothing takes place, node locations are mostly affected by sheet extractions. As described in Chapter 2, when a sheet is extracted, pairs of nodes are merged into a single node. The location of the new node is the average of the two unmerged nodes' locations. Therefore, before a given sheet extraction occurs, it is possible to know where the nodes of adjacent elements will be located after the extraction. With this information, it can be known in advance how a sheet will affect the quality of the mesh if it is extracted.

In the automated coarsening algorithm, element quality is calculated using a shape quality metric, f_{shape}, proposed by Knupp [24]. This metric has a value of 1.0 if the element is a perfect cube and a value of 0.0 if the element is degenerate. The metric is mathematically defined in Equation 3-1.

21

$$f_{shape} = \frac{24}{\sum_{k=0}^{7}(\lambda_{11}^k + \lambda_{22}^k + \lambda_{33}^k)/(\alpha_k^{2/3})}$$

<div align="right">(3-1)</div>

where $\lambda_{i_j}^k$ = the ij^{th} component of the k^{th} metric tensor

α_k = the determinant of the k^{th} Jacobian matrix

Using this shape quality metric, sheets are ranked according to how they will affect the mesh if they are extracted. Sheets that will produce a higher quality mesh are given higher priority in the coarsening process.

Sometimes mesh operations can produce elements that are forced to have unacceptable shape quality. Such elements are usually confined by the geometry of the model or by the topologic connectivity they share with adjacent elements. One example of this is found in the formation of doublets [20]. In a quadrilateral mesh, doublets are formed when two quadrilateral elements share two edges, as shown in Figure 3-6. It can be seen from the figure that no matter how the nodes are moved, at least one of the elements will always be poorly shaped.
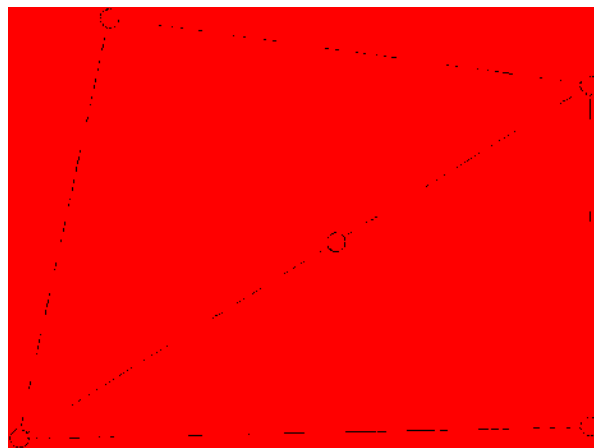
Figure 3-6. A Doublet.

Doublets and other constrained elements can easily form in a hexahedral mesh during the coarsening process if certain precautions are not taken. For this reason, several checks have been implemented into the automated coarsening algorithm in an attempt to prevent such elements from forming. These checks are generally effective, but in some cases they severely limit the amount of coarsening that can take place.

Each time the coarsening region is pillowed, columns are collapsed to produce sheets which are then extracted. In order to determine the very best combination of column collapse operations for a given iteration of the coarsening process, every possible combination would have to be analyzed. However, as shown in Equation 3-2, the number of possible column combinations, *N*, increases exponentially with each additional column that exists in the pillow sheet. Therefore, for most coarsening regions, an exhaustive search to determine the best combination is not feasible.

$$N = \sum_{i=0}^{n-1} 2^i \qquad (3\text{-}2)$$

where n = the number of columns that exist in the pillow sheet

In the automated coarsening algorithm, the number of column combinations analyzed never exceeds the number of sheets passing through the coarsening region. This relatively small sample of combinations gives priority to sheets that will produce higher element quality when extracted. While this method is somewhat limiting, it generally proves to be effective and provides a much more feasible approach than analyzing every possible combination.

The automated coarsening algorithm was initially developed for regions that require pillowing to produce entirely localized coarsening. However, to make the algorithm more general, the pillowing step is left out in cases where existing sheets are already contained entirely within the coarsening region. The overall structure of the algorithm is shown in Figure 3-7 and described by the following steps.

1. A coarsening region is defined and a target mesh density for that region is determined based on input given by a user.

2. Every sheet that passes through the coarsening region is found. Sheets contained entirely within the coarsening region are distinguished from those that extend beyond the region.

3. Due to a variety of geometry and mesh topology constraints, each sheet is examined to see if it will facilitate valid collapses and extractions during the coarsening process. Sheets that are unable to facilitate valid collapses and extractions are ignored from this point on.

4. For each acceptable sheet, the shape quality metric defined in Equation 3-1 is used to estimate how the quality of the mesh will be affected if that sheet, or the portion of that sheet contained in the coarsening region, is extracted. Sheets that will potentially produce a higher mesh quality are given higher priority.

5. If there are any sheets contained entirely within the coarsening region, then valid combinations of those sheets are analyzed. The combination that, when extracted, will produce a mesh density that is closest to the target mesh density without over-coarsening is saved. If no acceptable combination is found, the algorithm

24

moves to step 6. Otherwise, steps 6 through 8 are skipped because no other operations are needed before sheet extraction.

6. If there are any sheets that extend beyond the coarsening region, then valid combinations of those sheets are analyzed. For each combination, two coarsening options are possible, as shown in Figure 3-8. These two coarsening options are distinguished by which direction the columns are collapsed. The combination that will produce a mesh density that is closest to the target mesh density without over-coarsening is saved. If no acceptable combination is found, steps 7 through 9 are skipped.

7. A sheet is inserted around the boundary of the coarsening region through pillowing.

8. Columns in the pillow sheet are collapsed in directions which were previously determined when the best sheet combination was saved. These collapses form sheets which are contained entirely within the coarsening region.

9. Sheets contained entirely within the coarsening region are extracted.

10. Steps 2 through 9 are repeated until the target mesh density is achieved (within a certain tolerance) or no more valid sheet combinations are found.

11. If coarsening took place, the remaining elements in the region are smoothed to improve mesh quality.

Figure 3-7. Automated coarsening algorithm flowchart.
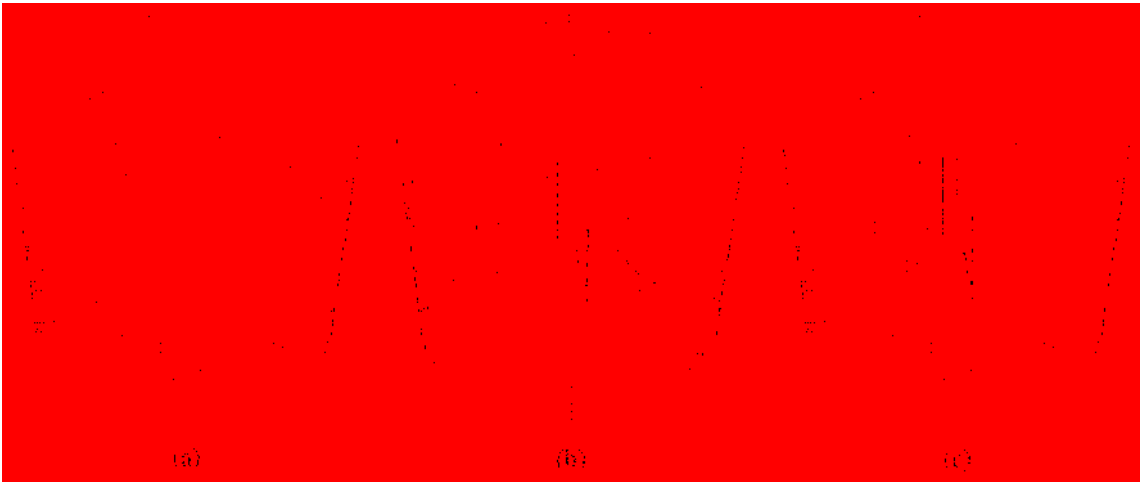
Figure 3-8. Two coarsening options: (a) Columns selected for collapsing. (b) The sheets that will form if the columns are collapsed one way. (c) The sheets that will form if the columns are collapsed the other way.

4 Examples

The following three examples show some results of the automated coarsening algorithm described in Chapter 3. In each example, the goal was to remove 25, 50, and 75 percent of the elements in the region selected for coarsening, while maintaining acceptable element quality. Quality was measured using the scaled Jacobian [25], which ranges from -1.0 to 1.0. An element is generally considered acceptable if it has a scaled Jacobian greater than 0.2.

The first example was performed on a structured mesh of a cube, as shown in Figure 4-1. The second example was performed on an unstructured multiple-source to single-target swept mesh of a mechanical part, as shown in Figure 4-2. The final example was performed on an unstructured mesh of a human head generated with an octree based, sheet insertion algorithm [26], as shown in Figure 4-3 and Figure 4-4. For both the mechanical part and human head models, refinement was performed prior to coarsening to create a higher starting mesh density.

Table 4-1, Table 4-2, and Table 4-3 provide element removal, element quality, and coarsening time results for each model. In almost every case, acceptable element quality was maintained and a density that very nearly reflects the target mesh density was achieved.

Table 4-1. Coarsening Results for Cube Model

Target % Removal	Elements in Region	Actual % Removal	Min. Scaled Jacobian	Coarsening Time (sec)
0	1331	--	1.00	--
25	1056	20.7	0.47	0.7
50	684	48.6	0.41	0.9
75	355	73.3	0.34	1.1

Figure 4-1. Structured cube example: (a) Original mesh with coarsening region defined. (b) 25 percent coarsening. (c) 50 percent coarsening. (d) 75 percent coarsening.

Figure 4-2. Unstructured mechanical part example: (a) Original mesh with coarsening region defined. (b) 25 percent coarsening. (c) 50 percent coarsening. (d) 75 percent coarsening.

Table 4-2. Coarsening Results for Mechanical Part Model

Target % Removal	Elements in Region	Actual % Removal	Min. Scaled Jacobian	Coarsening Time (sec)
0	7641	--	0.77	--
25	5807	24.0	0.59	5.3
50	4057	46.9	0.32	9.6
75	2205	71.1	0.22	12.5

Table 4-3. Coarsening Results for Human Head Model

Target % Removal	Elements in Region	Actual % Removal	Min. Scaled Jacobian	Coarsening Time (sec)
0	10080	--	0.48	--
25	7953	21.1	0.29	13.0
50	5129	49.1	0.17	17.9
75	2615	74.1	0.22	22.5

Figure 4-3. Unstructured human head example (side view): (a) Original mesh with coarsening region defined. (b) 25 percent coarsening. (c) 50 percent coarsening. (d) 75 percent coarsening.
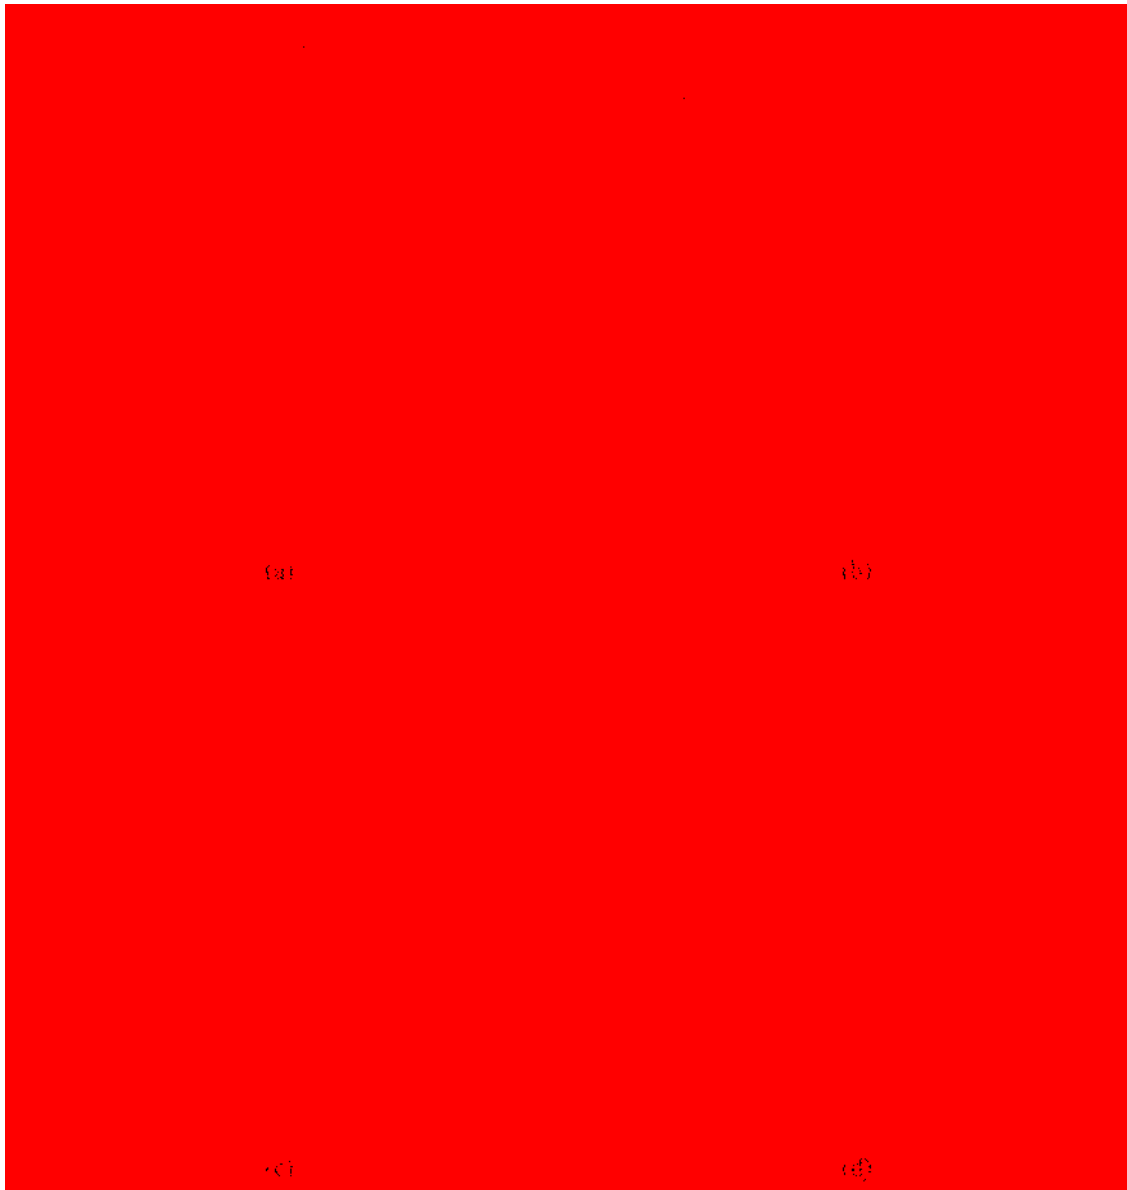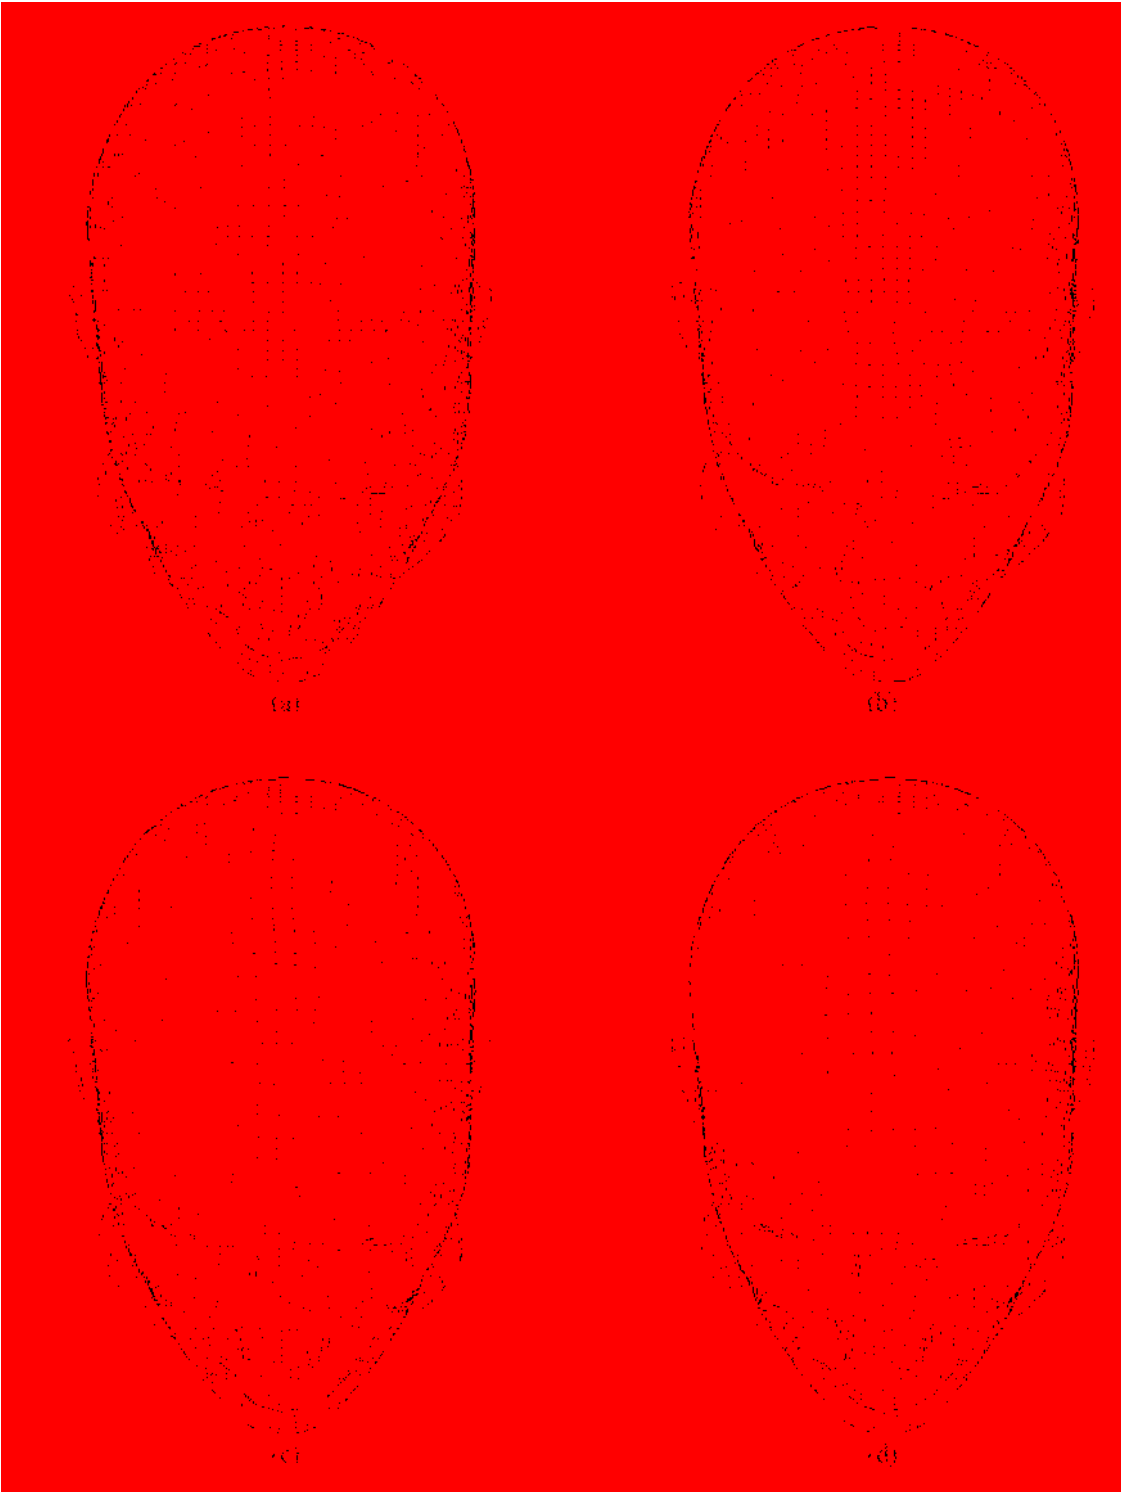
Figure 4-4. Unstructured human head example (top view): (a) Original mesh with coarsening region defined. (b) 25 percent coarsening. (c) 50 percent coarsening. (d) 75 percent coarsening.

5 Future Work

The automated coarsening algorithm described in Chapter 3 takes advantage of sheets already existing entirely within the coarsening region which can be extracted without any previous operations. However, it does not take advantage of columns already existing entirely within the coarsening region which can be collapsed without any previous operations. Modifying the algorithm to take advantage of such columns would improve the efficiency and effectiveness of the coarsening process in certain situations.

While the automated coarsening algorithm guarantees a topologically conforming mesh, it does not guarantee that the final quality of the mesh will be acceptable. Further research is needed to ensure that hexahedral coarsening does not degrade mesh quality below an acceptable threshold. This might be accomplished through more sophisticated methods which prevent poor quality elements from forming, or through cleanup operations which fix bad elements without significantly affecting mesh density. Many effective methods to cleanup a quadrilateral mesh have recently been developed [22]. It is hoped that further research will lead to similar methods for a hexahedral mesh.

The coarsening method presented in this thesis has been shown to work on unstructured meshes. However, even though these meshes are considered to be unstructured, they are usually structured in one dimension. Little work has been done to test this method on completely unstructured meshes. In theory, the method should work

for any hexahedral mesh. However, it is likely that some meshes cannot be coarsened without degrading element quality below an acceptable level.

While the coarsening method presented in this thesis works, it involves steps which may not be necessary. For example, most of the elements that are added to the mesh through pillowing are later removed. It is hoped that further research will lead to a more efficient method which only inserts elements that are necessary to transition from higher to lower mesh density.

Finally, the ultimate goal of this research is to develop an effective hexahedral coarsening method that can be combined with existing refinement methods in a fully automated mesh adaptation process. Currently, the automated coarsening algorithm relies on a user to define a coarsening region and level of coarsening. Further work is needed to link this algorithm with finite element software that can dictate which areas of a mesh need to be coarsened and by how much.

6 Conclusion

By utilizing sheet and column operations such as pillowing, column collapsing, and sheet extraction, entirely localized coarsening can be achieved in conforming all-hexahedral meshes. This method of coarsening works on both structured and unstructured meshes and is not based on undoing previous refinement. Although not fully developed, automation of this hexahedral coarsening method has already shown promising results. However, further work is needed to more effectively achieve the main objectives of an automated coarsening algorithm. These objectives include achieving a target mesh density, maintaining high element quality, and completing the coarsening process in as little time as possible. It is hoped that future research will lead to better methods of balancing these objectives. It is also hoped that the development of this new hexahedral coarsening method will lead to an effective coupling of refinement and coarsening in a fully automated mesh adaptation process.

References

1. R. Biswas and R. C. Strawn (1998). Tetrahedral and Hexahedral Mesh Adaptation for CFD Problems. *Applied Numerical Mathematics*, 26:135-151.

2. D. Morton and J. M. Tyler (2000). A New 3D Adaptive Finite Element Scheme with 1-Irregular Hexahedral Element Meshes. *Proceedings of the 2000 ACM Symposium on Applied Computing*, 99-104.

3. A. Tam, D. Ait-Ali-Yahia, M. P. Robichaud, M. Moore, V. Kozel, and W. G. Habashi (2000). Anisotropic Mesh Adaptation for 3D Flows on Structured and Unstructured Grids. *Computer Methods in Applied Mechanics and Engineering*, 189:1205-1230.

4. P. A. Cavallo, N. Sinha, and G. M. Feldman (2004). Parallel Unstructured Mesh Adaptation for Transient Moving Body and Aeropropulsive Applications. *42nd AIAA Aerospace Sciences Meeting and Exhibit*, 6555-6565.

5. S. E. Benzley, E. Perry, K. Merkley, B. Clark, and G. Sjaardema (1995). A Comparison of All-Hexahedral and All-Tetrahedral Finite Element Meshes for Elastic and Elasto-Plastic Analysis. *Proceedings of the 4th International Meshing Roundtable*, 179-191.

6. K. Tchon, J. Dompierre, and R. Camarero (2002). Conformal Refinement of All-Quadrilateral and All-Hexahedral Meshes According to an Anisotropic Metric. *Proceedings of the 11th International Meshing Roundtable*, 231-242.

7. N. J. Harris, S. E. Benzley, and S. J. Owen (2004). Conformal Refinement of All-Hexahedral Element Meshes Based on Multiple Twist Plane Insertion. *Proceedings of the 13th International Meshing Roundtable*, 157-167.

8. S. E. Benzley, N. J. Harris, M. A. Scott, M. J. Borden, and S. J. Owen (2005). Conformal Refinement and Coarsening of Unstructured Hexahedral Meshes. *Journal of Computing and Information Science in Engineering*, 5:330-337.

9. M. Parrish, M. J. Borden, M. L. Staten, and S. E. Benzley (2007). A Selective Approach to Conformal Refinement of Unstructured Hexahedral Finite Element Meshes. *Proceedings of the 16th International Meshing Roundtable*, 251-268.

10. T. J. Tautges and S. E. Knoop (2003). Topology Modification of Hexahedral Meshes Using Atomic Dual-Based Operations. *Proceedings of the 12th International Meshing Roundtable*, 415-423.

11. A. Sheffer and A. Ungor (2001). Efficient Adaptive Meshing of Parametric Models. *Journal of Computing and Information Science in Engineering*, 123:366-375.

12. C. K. Choi, E. J. Lee, and W. J. Yu (2004). Adaptive Mesh Refinement/Recovery Strategy for FEA. *Structural Engineering and Mechanics*, 17:379-391.

13. Y. Kallinderis and C. Kavouklis (2005). A Dynamic Adaptation Scheme for General 3-D Hybrid Meshes. *Computer Methods in Applied Mechanics and Engineering*, 194:5019-5050.

14. B. S. Kirk, J. W. Peterson, R. H. Stogner, and G. F. Carey (2006). libMesh: A C++ Library for Parallel Adaptive Mesh Refinement/Coarsening Simulations. *Engineering with Computers*, 22:237-254.

15. R. Taghavi (1996). Automatic, Parallel and Fault Tolerant Mesh Generation From CAD. *Engineering with Computers*, 12:178-185.

16. M. J. Borden, S. E. Benzley, and J. F. Shepherd (2002). Hexahedral Sheet Extraction. *Proceedings of the 11th International Meshing Roundtable*, 147-152.

17. P. Murdoch and S. E. Benzley (1995). The Spatial Twist Continuum. *Proceedings of the 4th International Meshing Roundtable*, 243-251.

18. J. F. Shepherd (2007). Topologic and Geometric Constraint-Based Hexahedral Mesh Generation. *Doctoral Dissertation – University of Utah*.

19. K. Merkley, C. Ernst, J. F. Shepherd, and M. J. Borden (2007). Methods and Applications of Generalized Sheet Insertion for Hexahedral Meshing. *Proceedings of the 16th International Meshing Roundtable*, 233-250

20. S. A. Mitchell and T. J. Tautges (1995). Pillowing Doublets: Refining a Mesh to Ensure that Faces Share at Most One Edge. *Proceedings of the 4th International Meshing Roundtable*, 231-240.

21. M. L. Staten, S. E. Benzley, and M. A. Scott (2008). A Methodology for Quadrilateral Finite Element Mesh Coarsening. *Engineering with Computers*, 24:241-251.

22. M. W. Dewey (2008). Automated Quadrilateral Coarsening by Ring Collapse. *Master's Thesis – Brigham Young University*.

23. P. M. Knupp (2003). A Method for Hexahedral Mesh Shape Optimization. *International Journal for Numerical Methods in Engineering*, 58:319-332.

24. P. M. Knupp (2003). Algebraic Mesh Quality Metrics for Unstructured Initial Meshes. *Finite Elements in Analysis and Design*, 39:217-241.

25. P. M. Knupp (2001). Algebraic Mesh Quality Measures. *SIAM Journal on Scientific Computing*, 23:193-218.

26. J. F. Shepherd, Y. Zhang, C. J. Tuttle, and C. T. Silva (2007). Quality Improvement and Boolean-like Cutting Operations in Hexahedral Meshes. *Proceedings of the 10th Conference of the International Society of Grid Generation*.

Appendix A.   Quadrilateral Coarsening Operations

The hexahedral coarsening method presented in this thesis grew out of recent developments in quadrilateral coarsening. Since two-dimensional coarsening operations are generally easier to visualize than their three-dimensional counterparts, an explanation of the related quadrilateral coarsening operations is given in this Appendix. While these operations can be directly extended to hexahedral coarsening, by themselves, they are not always able to prevent changes in element density from propagating beyond the boundaries of a defined hexahedral coarsening region. Entirely localized hexahedral coarsening often requires an additional step that is not necessary in quadrilateral coarsening.

A quadrilateral element can be represented by two line segments that connect the midpoints of opposite edges, as shown in Figure A-1. In a conforming quadrilateral mesh, these line segments combine to form chords, as seen in Figure A-2.

Using an operation known as chord extraction [21], the nodes associated with a given chord are merged together, as demonstrated in Figure A-3. Chord extraction is a valuable coarsening tool because it not only decreases mesh density, but it also guarantees the preservation of a conforming mesh. However, in many situations, chords extend beyond the boundaries of a defined region. The removal of such chords would decrease mesh density in areas where coarsening is not desired. Therefore, it is often
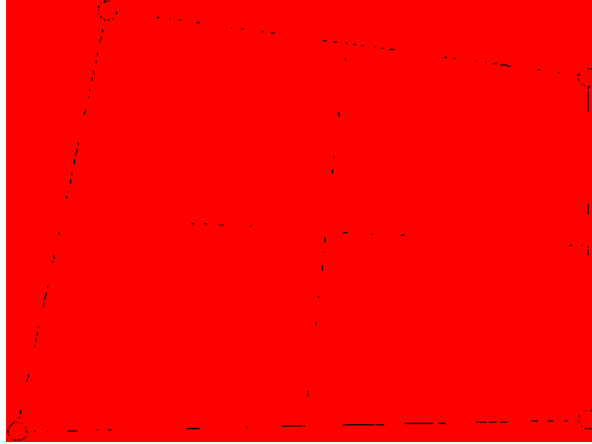
Figure A-1. A quadrilateral element with two line segments connecting the midpoints of opposite edges.
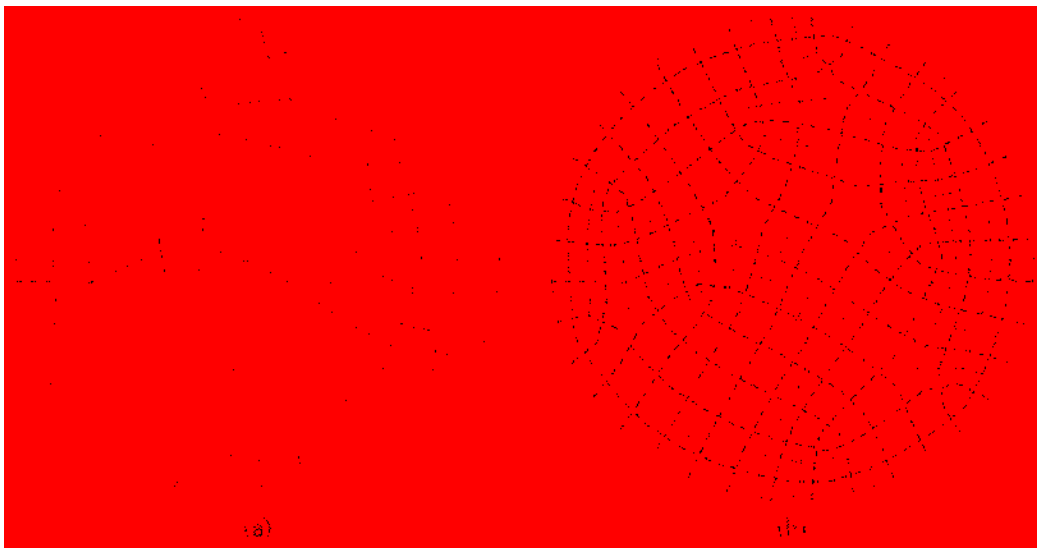
Figure A-2. Quadrilateral chords: (a) A quadrilateral mesh with one chord shown. (b) All the chords representing the mesh are shown.

necessary to modify the mesh in such a way that produces chords which are confined to the coarsening region.

Several operations have been developed to locally modify the topology of a quadrilateral mesh [21]. One such operation is an element collapse, which is illustrated in
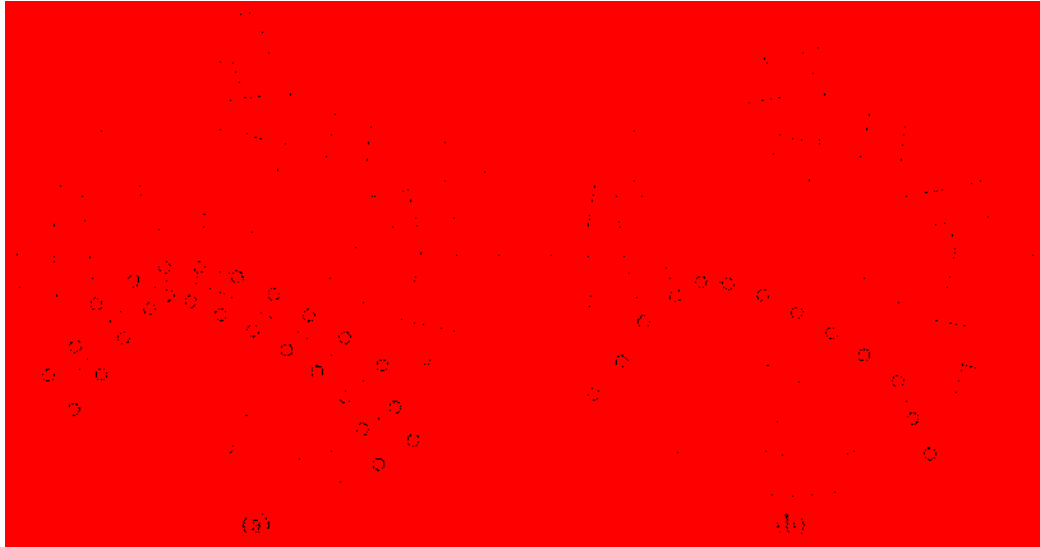
Figure A-3. Chord extraction: (a) A chord is selected for extraction. (b) The chord is extracted by collapsing the edges that define the chord and merging the two nodes on each edge.

Figure A-4. As seen in the figure, this operation alters the paths of two intersecting chords, preserves a conforming mesh, and only affects the elements immediately surrounding the collapsed element.

Figure A-4. Element collapse: (a) An element representing the intersection of two chords is selected for collapse. (b) The two nodes in (a) are merged and the two intersecting chords no longer intersect.

As shown in Figure A-5, the element collapse operation can be used at multiple chord intersections to form a chord which is contained entirely within the boundaries of a defined region. This chord can then be extracted to coarsen the region without affecting any other part of the mesh.

Figure A-5. Formation of a localized chord: (a) Three intersecting chords pass through a region selected for coarsening. (b) The elements representing the chord intersections in (a) are collapsed to produce a chord which is contained entirely within the boundaries of the coarsening region.