

RECEIVED

NOV 02 1999

OSTI

A METHOD FOR CONTROLLING SKEW ON LINKED SURFACES

Robert A. Kerr¹, Steven E. Benzley², Dave White³, Scott Mitchell⁴

¹*Brigham Young University, Provo, UT, U.S.A.* kerrr@et.byu.edu

²*Brigham Young University, Provo, UT, U.S.A.* seb@byu.edu

³*Sandia National Laboratories, Albuquerque, NM, U.S.A.* drwhite@sandia.gov

⁴*Sandia National Laboratories, Albuquerque, NM, U.S.A.* samitch@sandia.gov

Abstract

A new method for lessening skew in mapped meshes is presented. This new method involves progressive subdivision of a surface into loops consisting of four sides. Using these loops, constraints can then be set on the curves of the surface, which will propagate interval assignments across the surface, allowing a mesh with a better skew metric to be generated.

Keywords: mesh generation, skew measure, virtual vertex, mapped mesh

1 Introduction

One of the desired outcomes of automatic meshing is good element quality. It is well understood that a good quality mesh yields better results than does one with bad quality[1][2]. There are many different measures as to what constitutes "good quality". Among this group[3] are such metrics as jacobian, aspect ratio, taper, warpage, element area, stretch, maximum angle, minimum angle, oddy condition number, and skew. This paper will concentrate on the problem of controlling skew.

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

2 The Skew Problem

Skew is defined as the maximum absolute value of the cosine of the angle between edges at the center of the quadrilateral. In other words, it is the cosine of the angles formed by the two lines which pass through the midpoints of the sides of the quadrilateral[4]. In Figure 1, a quadrilateral element is shown, with one

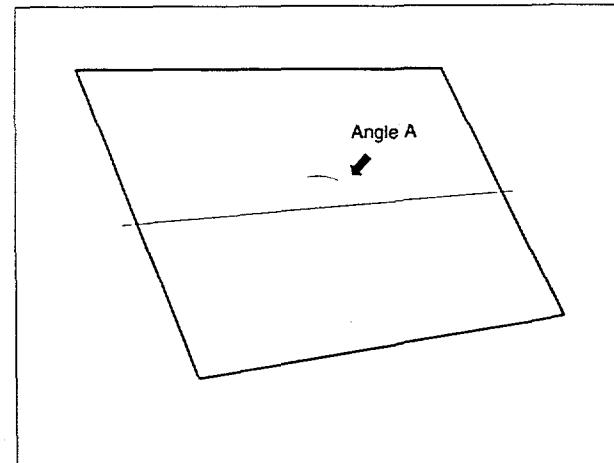


Figure 1: Quadrilateral mesh element showing representative angle 'A'.

of the angles labeled as angle A. The absolute value of the cosine of this angle would be one of the four pos-

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

sible measures of skew for this element. The cosines of the other three angles would be calculated, and the maximum of those four cosines would be the measure of skew for this face. As can be seen, the value of skew ranges between 0.0 (because of the absolute value operation) and 1.0, with the optimal value being 0. Mapped meshes, by their nature, depend on propagation of interval assignments[5][6]. Skew usually isn't a problem in small, simple models. However, when many volumes are multiply-connected, as shown in Figure 2, interval assignments propagate throughout

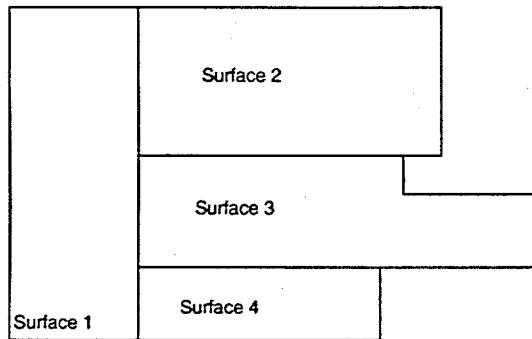


Figure 2: Multiply-connected set of surfaces which could lead to skew problems.

the model. This propagation of interval assignments can lead to skew problems. In Figure 2 interval settings on the ends of surfaces 2, 3, and 4 will propagate across those surfaces and affect the right-hand side of surface 1. Surface 1 will then have to set its left-hand edge to the same number of intervals as is on the right-hand side. If one of the right-hand side curves has a comparatively high interval count, the mesh on surface 1 could become greatly skewed. Figure 3 shows, for example, a possible set of intervals for the right-hand side of surfaces 2, 3, and 4. These intervals will propagate across the surfaces, as shown in Figure 4. After propagating across the surfaces, the intervals on the common edge between surface 1 and the other three surfaces are firmly set. A mesh then needs to be generated on surface 1. Most mesh

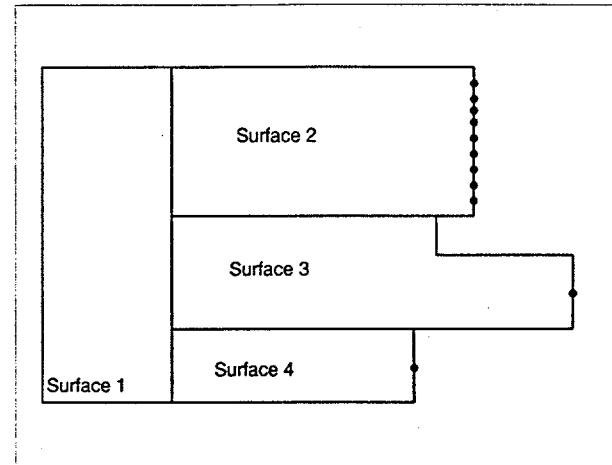


Figure 3: Set of surfaces with interval assignments on right-hand sides.

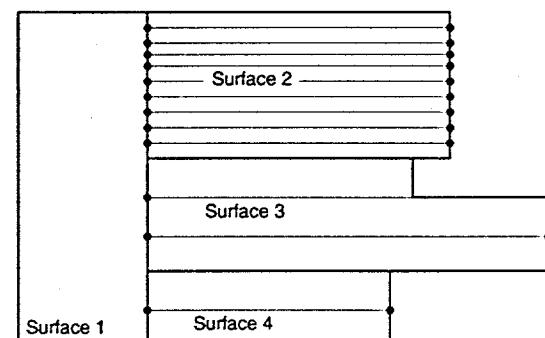


Figure 4: Interval propagation across surfaces.

generating software assigns equally-spaced intervals as the default. Therefore, the left-hand side interval assignment on surface 1 would look like that shown in Figure 5. The final meshing of surface 1 would then be done, yielding a final mesh as is shown in Figure 6. As can be seen, the mesh on this surface would result in a low measure of skew quality. What is required to reduce skew is to develop a method that will transfer the interval assignments from one side of

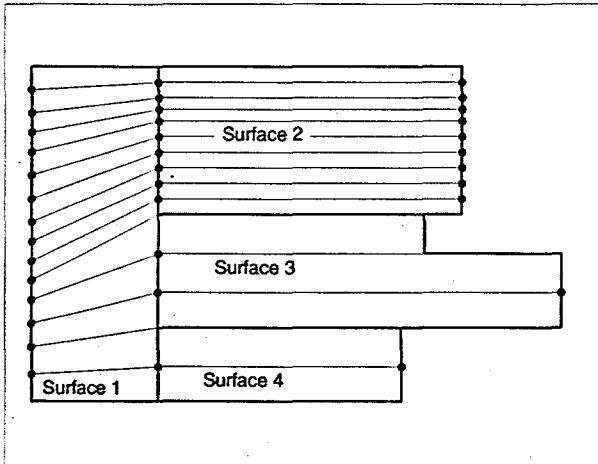


Figure 5: Default interval assignment on surface 1.

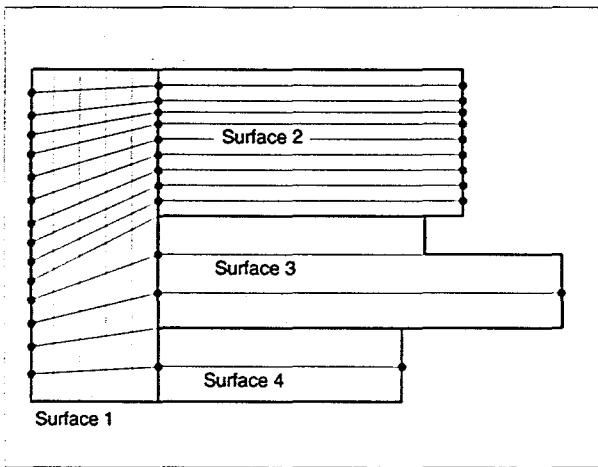


Figure 6: Final skewed mesh on surface 1.

the surface to the other in a manner that will preserve the relative interval spacings on different sections of the curves. Intervals on opposite sides cannot just be set to have the same intervals, because these opposite edges may be of different lengths. Manually setting intervals on surface 1 would still yield that same mesh as that shown in Figure 6. As can be seen, this is a severely limited solution to the problem of controlling skew—it isn't better than what would be done automatically. It is this circumstance which has prompted

the development of the skew control algorithm.

3 The Skew Control Algorithm

The purpose of this algorithm is to manage the interval settings on surfaces that exhibit unacceptable skew. Given a surface such as that discussed above, the skew control algorithm developed here will partition the edges of the surface and set up a system of matching edges across a surface or multiple surfaces. Once this is done, the interval count on corresponding edges is set to be equal, so that a created mesh has little or no skew. Instead of the skewed mesh in Figure 6, the final mesh would result as depicted in Figure 7.

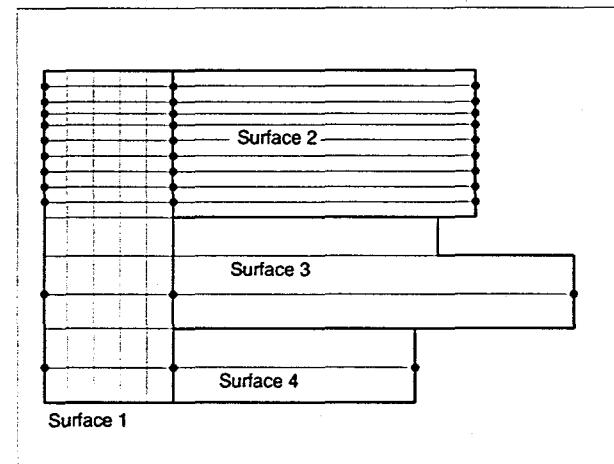


Figure 7: Surfaces showing a skew-controlled mesh.

The skew control algorithm has the following 7 steps:

1. Approximate the affected surfaces with pseudo geometry.
2. Create a loop of edges around the target surface.
3. Find the smallest projection on the surface.
4. Separate this small feature from the rest of the surface using a pseudo edge.

5. Continue separating small sections until all loops consist of only four edges.
6. Step through loops setting up interval assignments for opposite edges.
7. Clean up pseudo geometry.

These steps will be explained one-by-one with representative illustrations.

1. Approximate the affected surfaces.

The skew control algorithm depends on an approximation of the surfaces to be meshed. These surfaces are known as the target surfaces. The same set of surfaces shown in Figure 2 will be used in this example. The algorithm uses pseudo geometry known as skew control entities for this approximation. The curves and vertices which make up the real surface are used as templates to create skew control edges and skew control vertices. These skew control entities are the basis for almost all the work until the algorithm reaches step 7. Skew control entities only need to hold a little information. Each skew control edge knows what vertices define it, and each skew control vertex knows its position in three-space and its type, which will be defined later. Because of this sparsity of information, the memory overhead in using these pseudo-entities is negligible.

2. Create a loop of edges around the target surfaces.

The skew control edges and vertices that have been created from the target surfaces are now placed in lists which maintain their order. Each target surface is approximated by one list whose edges and vertices form a loop. These loops, which are the base for this algorithm, define the target surfaces throughout the rest of the algorithm. This example set of target surfaces would translate into four loops.

3. Find the smallest projection in the surface.

The skew control algorithm implements a type of "blocking" subdivision. "Blocking" refers to the process of dividing up the surface into blocks, or four-

sided figures. This blocking algorithm starts with the smallest areas, filling them with blocks, and expands to the larger ones. The algorithm could start with larger areas, and work to smaller ones, but by starting with the smallest, then the problem of intersection checking becomes much less pronounced. The starting step is to find the smallest end. For some geometries, this is not an easy location to find, and much depends on the definition of "end". An "end" is defined as a set of edges that are bounded by two vertices which are known as End_Type. As can be seen, the type of the skew control vertices has a great impact on the definition of an end. There are four types of vertices, based on the angle of the edges which share the vertex:

1. End_Type, with an angle close to 90° .
2. Side_Type, with an angle close to 180° .
3. Corner_Type, with an angle close to 270° .
4. Reversal_Type, with an angle close to 360° [5].

Figure 8 shows all four types. A skew control vertex

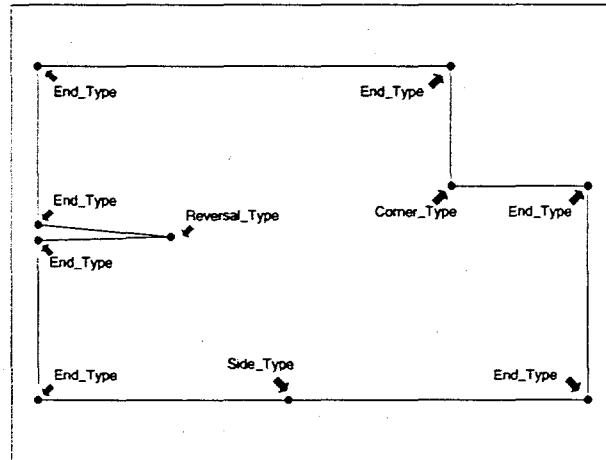


Figure 8: Surface showing the four vertex types.

is assigned a type based on the type of the underlying real vertex. If there is no underlying vertex, the type is computed based on the angle of the connected skew control edges. The target surface's loop is searched

for the “end” that is shortest, and that end is used for the next operation. As shown in Figure 9, the far right-hand end will be picked.

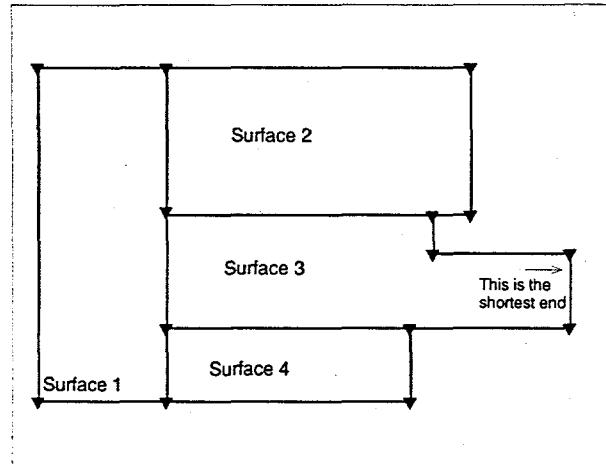


Figure 9: Surfaces showing which end will be picked first. Vertices are indicated by the dark triangles.

4. Separate the smallest projection from the rest of the loop.

When the shortest end has been found, the algorithm then looks at the two edges that form the sides of the block. At this point there are two possibilities: either one side will be longer than the other, or the two will be the same length (within a tolerance). If the first case applies, the algorithm splits the longer of the two sides so that the two sides are equal. Of course, this splitting can happen within a range of lengths, depending on local geometry and the desired behavior. If the pseudo edge to be split has an underlying geometry edge, that geometry edge is split too, and a virtual vertex is inserted. After the split of the pseudo edge, this first case becomes identical to the second case. When two vertices are at an equal distance from the end, a new pseudo edge is created between the two vertices. The edges and vertices comprising the block being replaced are then removed from the old loop, and the new edge is inserted in their place. A new loop is created using the newly independent edges and vertices and the new edge, as shown in Figure 10.

Figure 10. Where once there was one loop defining the

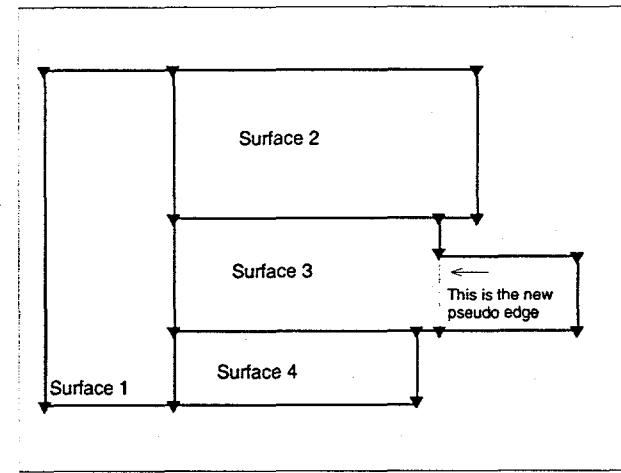


Figure 10: Surfaces after splitting curve into two curves and inserting a new pseudo edge. Note the new virtual vertex.

surface, now there are two. An important feature of this new loop is that it consists of only four edges. The goal is to separate all the loops into a series of small loops having only four edges apiece. When this goal is reached, then all loop subdivision should be finished, and the algorithm can proceed to the next step.

5. Continue separating small sections until all loops consist of only four edges.

Now there are five loops to consider, instead of the previous four. Each of these loops is checked to see if they have more than four edges. If one is found, the previous step is repeated and the next, smallest projection is separated from that loop and put into its own loop. Now, because of the possibility of an edge being in two loops at the same time, it is vital that the loops stay current regarding which edges belong to them. It would cause a severe problem if an edge was split but the loops containing the old edge still had pointers to that obsolete edge. Pointers to these owning loops are of necessity another piece of data contained in each skew control edge. Each edge keeps a list of loops it belongs to. With this infor-

mation, when an edge is subdivided, as in step 4, the algorithm can access this list of loops and update each one of them so that they contain the correct information. A side effect of this is that a loop that previously only had four edges might end up having more through the splitting of one of its included edges. This necessitates stepping through the set of loops multiple times until all the loops have only four edges. An example of this increase in edges is seen in Figure 11 where the center left-hand loop (which

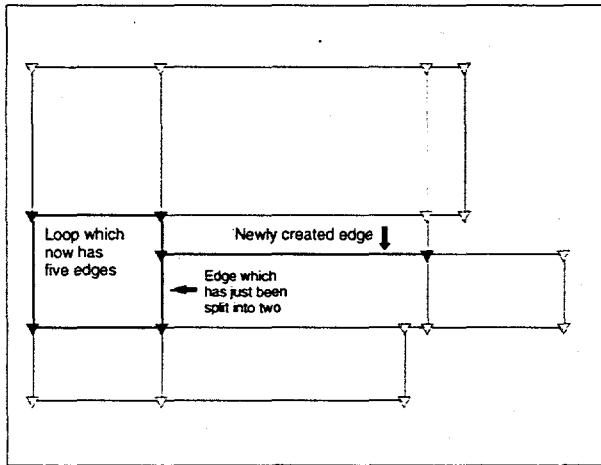


Figure 11: Surfaces showing loop about to be subdivided.

is shown in bold in the figure) is having its right-side curve subdivided. The next time this loop is examined, it will be subdivided again, because of the increase in number of edges. This is shown in Figure 12, where all the needed subdivisions have taken place.

6. Step through loops setting up interval assignments for opposite edges.

Once the loops have satisfied the requirement of having only four edges, the setting of constraints can be done. Each loop is examined to find an edge that has an underlying geometry edge (or owner). This edge is then marked and the loop is searched for the edge opposite the marked one. If the opposite edge has an underlying owner, then those two edges are set to

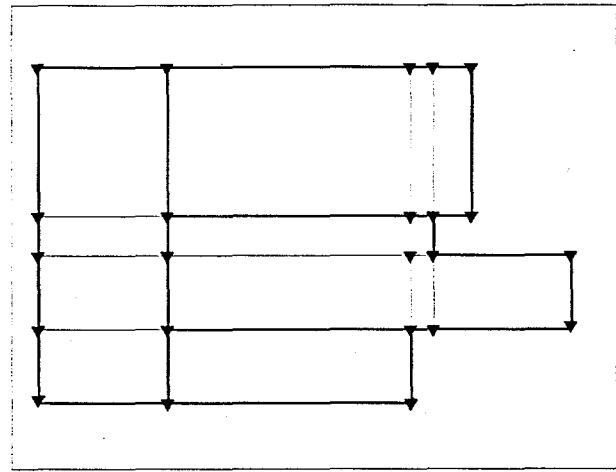


Figure 12: Surfaces showing final subdivision of target loops.

have equal intervals. If this opposite edge does not have an owner, the algorithm asks the edge for its list of owning loops and these loops are searched in the same manner. In this way, geometry edges that are at either end of a series of loops are found and then set to have equal intervals. Figure 13 shows the final

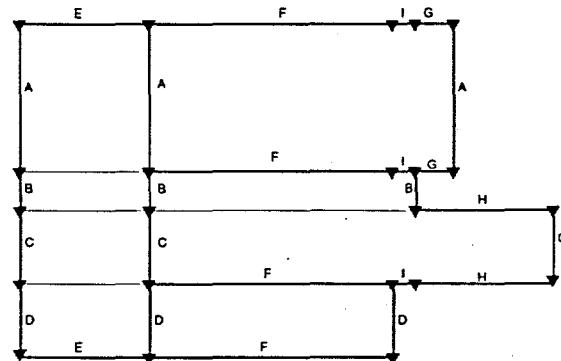


Figure 13: Surfaces showing final subdivisions with interval groups labeled.

state of the example surfaces, with each set of labeled edges having the same interval setting. Because each

edge can belong to any number of loops, the interval assignment on one edge may propagate to quite a few different edges on many different surfaces, but this is handled transparently. The algorithm doesn't know or care if the edges it is dealing with are on the same surface or different ones, all it sees are loops and owners.

7. Clean up pseudo geometry.

After the interval assignment has been done, the skew control entities can be safely deleted. The only changes done to the underlying geometry are that where edges had to be split, there are now two edges and a vertex (Figure 14), and interval assignments

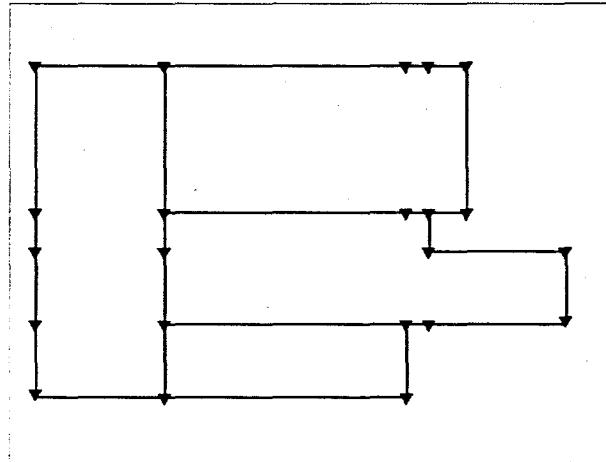


Figure 14: Surfaces after skew control cleanup.

have been made to all the curves in the target surfaces.

4 The Skew Control Algorithm: Results

The model shown in Figure 15 has some multiply-connected surfaces that can cause skew. The surface from this model shown in Figure 16 is rather complicated, with many other surfaces touching it. If meshed without skew control, the mesh would be created as shown in Figure 17. However, if the skew

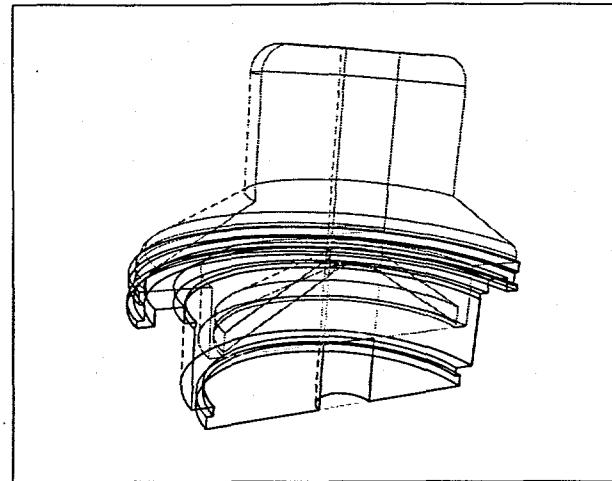


Figure 15: Model with potential skew problems.

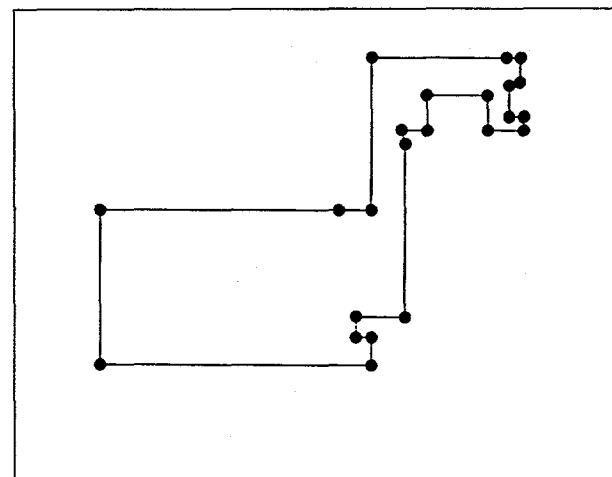


Figure 16: Hooked surface before meshing.

control algorithm is invoked on this surface, the resulting mesh demonstrates much less skew. The skew control algorithm processes this surface, inserting virtual vertices in places it deems appropriate. The surface in question now has new vertices, as shown in Figure 18. The mesh that results from the changed surface is shown in Figure 19. As this shows, the mesh is noticeably less skewed. Use of the skew control algorithm does result in an increase in the number of edges and vertices in the model, but this has

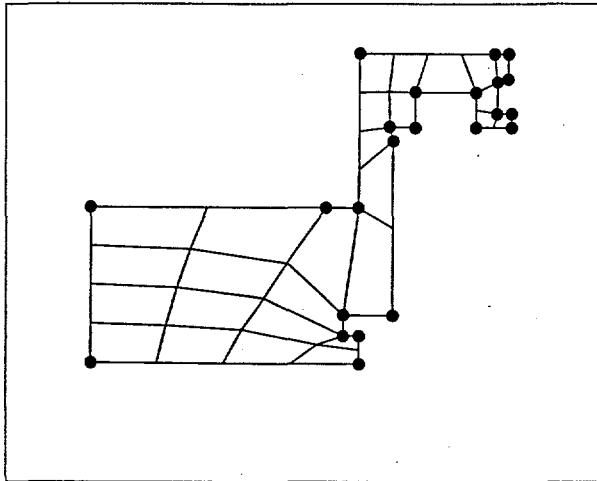


Figure 17: Hooked surface after meshing without skew control.

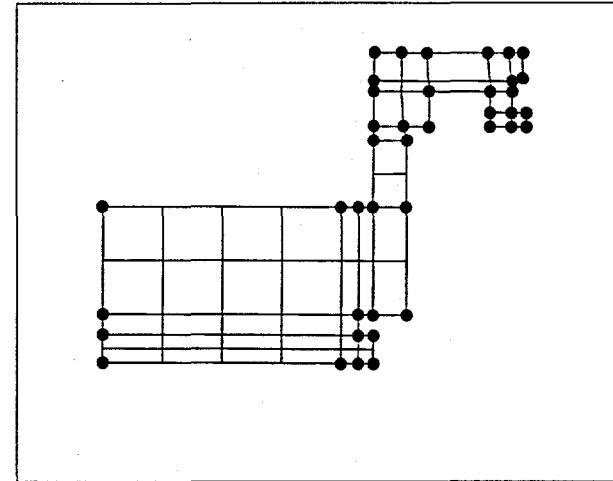


Figure 19: Hooked surface after meshing with skew control.

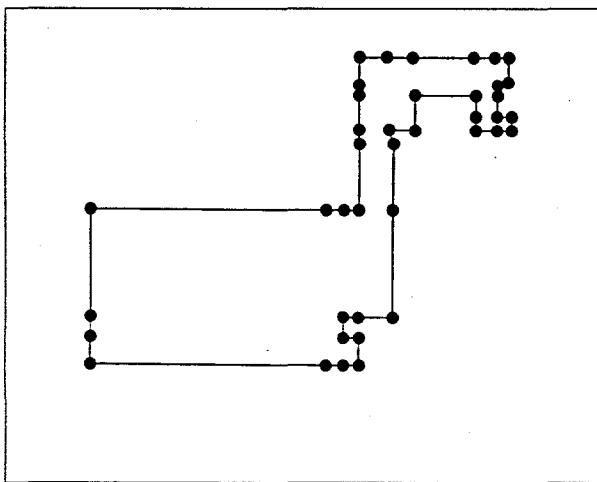


Figure 18: Hooked surface after processing by the skew control algorithm.

Hooked Surface	Non-controlled	Controlled
Maximum Skew	0.7081	0.1139
Mininum Skew	0.02564	0.00
Average Skew	0.2384	0.01282

Table 1: Comparison of skew for hooked surface with and without skew control.

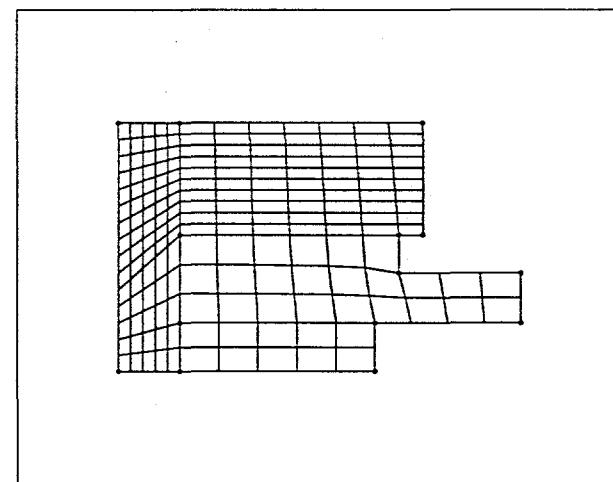


Figure 20: Skewed mesh generated on linked surfaces.

not proven to be a major concern. The mesh of the Hooked surface is of good quality, and the measure of skew has decreased dramatically, as is shown in Table 1.

Another example of skew control is shown in Figures 20 and 21. As can be seen, this is the example which was used to demonstrate the steps of the algorithm. Figure 20 demonstrates the mesh that is

generated without skew control. Although this paper

concentrated on the surface on the far left, called surface 1, it is also obvious that there are skew problems in other surfaces too. Figure 21 is the model with a

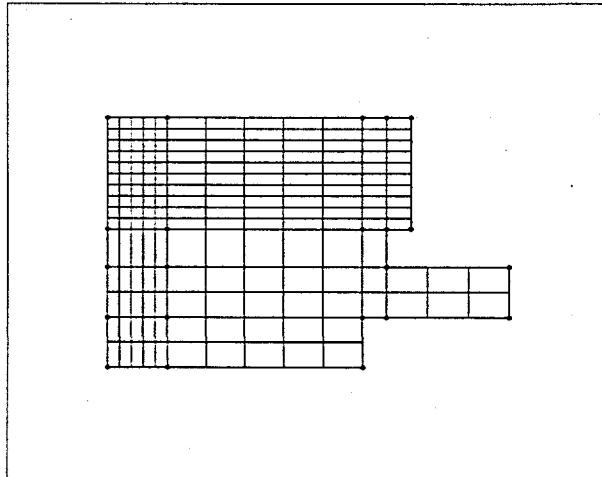


Figure 21: Non-skewed mesh generated with skew control algorithm on linked surfaces.

skew controlled mesh. As shown, there is very little skew on the surfaces. Table 2 compares the values of

Linked Surfaces	Non-controlled	Controlled
Maximum	0.6356	2.109×10^{-15}
Minimum	0.04329	0.00
Average	.3627	6.698×10^{-16}

Table 2: Comparison of skew for linked surfaces with and without skew control.

skew obtained with and without skew control. They are comparable to the values shown for the previous example.

5 The Skew Control Algorithm: Further work.

Although the skew control algorithm has been seen to provide good results in simple cases, there are some issues that need to be addressed to make this a universally applicable tool. Many of the remaining areas

of debate exist because of particular design decisions that need to be made.

- The skew control algorithm only works on surfaces that are submappable, ie. of blocky, roughly four-sided sections or subsections. This is a limitation that will probably stand, for the purpose of the algorithm is to enhance such algorithms as submapping—it is not meant to be used for surfaces that would need to be meshed with an unstructured mesh.
- Periodic surfaces are not handled correctly as of yet. Because the underlying geometry is not flat, the approximations done by the skew control Entities sometimes lead to errors in creating new vertices.
- User-set intervals are not propagated onto curves that have been split. This leads to some unanticipated results in the final mesh.
- Biased intervals settings are not propagated correctly from split curves.
- Surfaces with more than one loop, ie. surfaces with holes in them, are not handled correctly. There is some discussion as to how these surfaces should be dealt with, and the skew control algorithm's multi-loop handling is awaiting a design decision.

Another area of focus that deserves to be examined is the creation of virtual vertices. By creating vertices, the algorithm is able to easily set up the constraint equations for the final model. As can be seen by comparing Figures 16 and 18, the skew control algorithm does tend to insert a multitude of vertices. Although this creation of vertices hasn't proven to affect the model excessively, it is a goal of this research to develop a different way of setting up the constraint equations. Through use of the curve morphing algorithm[7], an algorithm that duplicates a curve mesh or set of curve meshes onto a target curve or set of curves, it is hoped to be able to correctly transfer the mesh from one side of a surface to another without the use of virtual vertices.

Perhaps the final area is that of testing. This tool needs to be tested extensively to decide such questions as tolerancing of curve intersections, and where to place target vertices on the target curve. In cases such as curved sides, the vertices need to have a better apparatus for deciding their final location--this will depend on the general shape of the loops.

In any case, the skew control algorithm seems to be a good area for further research in the hopes of lessening skew propagation caused by interval assignments, which will make large, complicated models much easier to mesh competently.

gorithms," Master's Thesis in Progress, Brigham Young University, December 1999.

References

- [1] Babuska, I., and Aziz, A., "On the Angle Condition in the Finite Element Method", SIAM Journal on Numerical Analysis, 13:214-226, 1976.
- [2] Fried, I., "Condition of Finite Element Matrices Generated from Nonuniform Meshes", AIAA Journal, 10:219-22, 1972.
- [3] Canann, S.A., Tristano, J.R., and Staten, M.L., "An Approach to Combined Laplacian and Optimization-Based Smoothing for Triangular, Quadrilateral, and Quad-Dominant Meshes", Proceedings of the 7th International Meshing Roundtable, Sandia National Laboratories, Oct., 1998.
- [4] Robinson, J., "CRE Method of Element Testing and the Jacobian Shape Parameters", Eng. Comput., Vol. 4, No. 2, p113-118, June 1987.
- [5] White, D.R., "Automatic, Quadrilateral and Hexahedral Meshing of Pseudo-Cartesian Geometries using Virtual Decomposition," Master's Thesis, Brigham Young University, August 1996.
- [6] Mitchell, S.A., "High Fidelity Interval Assignment", Proceedings, 6th International Meshing Roundtable, SNL, Albuquerque, N.M., Oct 1997, 33-44.
- [7] Kerr, R.A., "Improvement of Surface Meshes by Use of the Skew Control and Curve Morphing Al-