COMMUNICATIONS IN COMPUTATIONAL PHYSICS
Vol. 1, No. 4, pp. 716-735

Commun. Comput. Phys.
August 2006

A New Method for Efficient Generation of High Quality Triangular Surface Meshes

Desheng Wang*

School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 637616.

Received 24 August 2005; Accepted (in revised version) 30 November 2005

Abstract. A novel method for the generation of unstructured triangular surface meshes is presented. The method is based on remeshing techniques including edge splitting/contraction and edge swapping. Normalized edge lengths, based on a metric derived from curvature or from a user–specified spacing, are employed as the remeshing criterion. It is assumed that the geometry is input in the form of composite parametric surfaces, with Ferguson or Nurbs type multiple patch representation. Examples involving typical aircraft geometries and a ship model, are included to demonstrate how high quality meshes can be efficiently generated on surfaces with a high degree of geometric complexity.

Key words: Surface mesh generation; surface remeshing; edge split; edge contraction.

---

1 Introduction

Surface mesh generation is still a challenging task for numerical simulations involving large scale volumetric mesh generation. Existing approaches can essentially be categorized as being either direct or indirect [1–7]. In the direct approach, the mesh is generated directly on the 3D surface, using an octree based or an advancing front method [8, 9]. The indirect approach relies on generation in the parametric domain, using any 2D mesh generation procedure [4,5,7,10,11]. The mesh is then mapped onto the 3D surface. Direct methods have difficulty in checking the validity of the mesh, while indirect methods have difficulty in controlling the size and shape of the elements generated in the 2D domain. The development of a method for efficiently generating large-scale surface meshes with high quality motivates this work.

*Correspondence to: Desheng Wang, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 637616. Email: desheng@ntu.edu.sg

For engineering simulations, such as computational fluid dynamics, most boundary definitions come directly from CAD systems. The representation of a three dimensional geometry based on CAD patches, such as Ferguson patches or NURBS patches, will consist of both geometrical and topological data. The geometrical data is defined in the physical space and includes the basic parameters defining the shape of the support surfaces and the intersection curves. The topological data defines the generated regions and their support surfaces and bounding curves. In classical surface mesh generation, each region is meshed individually. This means that points must be generated on the bounding curves and these points will be part of the surface mesh for the region. In many cases, the bounding curves are of no physical significance and small thin patches often result in the creation of sliver or distorted triangles [3]. Therefore, the development of an approach that enables the generation of high quality patch independent surface meshes, which conform to the input surface representation, is the second objective of this work.

In this paper, we attempt to meet these objectives by describing a new method for the efficient generation of high quality surface triangular meshes using remeshing techniques [12–14]. The geometry definition formats that are supported by the current implementation of the method include composite parametric surfaces, with discrete composite parametric Ferguson curves and surfaces or NURBS. And the global surface consists of multiple regions. Initially, for each region, a triangulation of the points defining the intersection curves and those of the region's support surface which are lying within the region, is generated. These region-wise triangulations are then assembled into a global surface conforming mesh which can be regarded as an initial surface triangulation or a reference surface mesh to approximate the given surface with an adequate accuracy. Then, remeshing operations, such as edge splitting/contraction coupled with edge swapping, are performed for refinement or coarsening of this initial mesh. Metric controlled edge lengths are computed with respect to surface curvature and coupled with a user specified mesh spacing control function. Mesh points are added by edge splitting, and are located on the surface using the G1 interpolation [12, 14–16]. Finally, the surface mesh is optimized by utilizing nodal smoothing, node connectivity optimization and edge swapping techniques. The remeshing procedure requires only local information and is highly efficient. And the efficiency is demonstrated for various complicated aircraft configurations and a ship model.

The remainder of the article is organized as follows. The geometry modeling and the input data format are discussed in Section 2. An overview of the proposed surface mesh generation method is given in Section 3. Details of the construction of the initial surface triangulation are presented in Section 4, while Section 5 describes the mesh refinement and coarsening techniques that are employed. The optimization of the final grid is also briefly described in Section 5 and a few examples are given in Section 6.

2 Geometry modelling and input data format

In engineering design and analysis, CAD systems are normally used to generate the geometric definition of the system under consideration. Commercial and research solid modelers often represent boundaries as composite parametric surfaces, or employ a discrete representation in terms of triangular facets. For the purposes of this work, it will be assumed that the geometry is defined in terms of composite parametric surfaces.

Composite parametric surfaces are composed of sets of patches, together with intersection curves. Each patch is represented in terms of spline composite curves and tensor–product surfaces, such as Ferguson, Bezier or Nurbs [1,10,17,18]. This forms the so–called support surface [2,4,17], which may be shared by several regions. The patch technique utilizes a bi–cubic spline interpolation for the mapping. A set of intersection curves which bound the domain to be generated are also defined in the form of cubic splines. Standard data import formats, such as IGES and STEP, or the FLITE [2,4] fall into this category. Although other parametric curve and surface definitions may be used in practice, for the present implementation, the geometry definition must satisfy the following conditions:

1. The global surface consists of a number of conforming sub–surfaces, or patches, together with a set of intersection curves.
2. A mapping from physical space to a two–dimensional parametric space exists. This mapping is a tensor–product of splines.
3. The intersection curves are defined by parametric representations.
4. The mapping of each patch is bijective almost everywhere, with singular points appearing only on the boundary.
5. The boundary of a region is formed by one or more closed loops of intersection curves. Intersection curves forming a loop connect with each other only at their end points. If the boundary of a region is formed by more than one loop, then the loop enclosing the largest area in the parametric domain is taken to be the outer loop.

As an example, for an aircraft configuration, the support surfaces and the intersection curves bounding the generated surfaces are shown in Figs. 1(a) and 1(b) respectively.

3 Overview of the new method

The main stages in the proposed method are:

1. The generation of an initial surface triangulation. For each region, the points defining its intersection curves and its support surface are utilized to form a surface triangulation. Then, the triangulations of all regions are assembled into a global conforming surface triangulation by merging duplicate vertices. Fig. 2(a) shows such an initial surface mesh for a sphere.

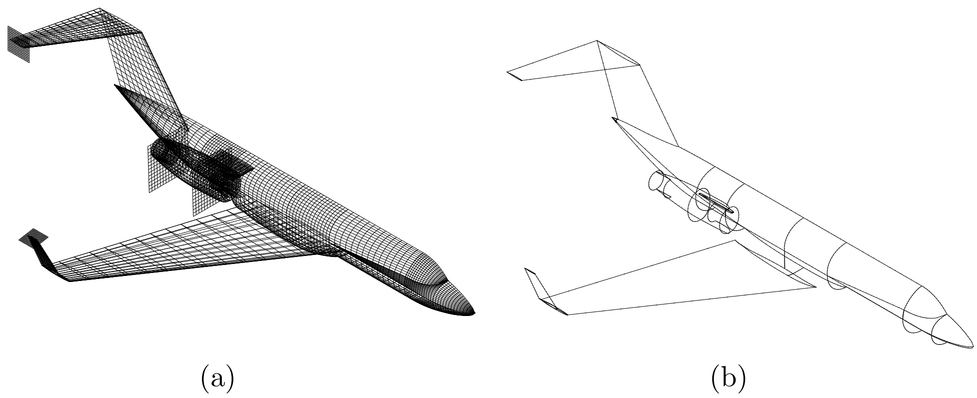

Figure 1: Modelling the surface of an aircraft showing: (a) the support surfaces; (b) the intersection curves.

2. The remeshing of the initial surface mesh to satisfy element spacing requirements. These requirements can either be user specified or determined automatically, based on the surface curvature. Remeshing is accomplished by utilizing local mesh modifications such as edge splitting/contraction and edge swapping. Figs. 2(b), (c) and (d) show the surface meshes achieved for a sphere after different numbers of remeshing steps.
3. The improvement of the element quality and point distribution by optimization. The triangulation is enhanced through point smoothing, node connectivity optimization and further edge swapping. Fig. 2(e) displays the final mesh for a sphere surface after optimization.

These stages are described in detail in the next three sections.

4 Initial surface triangulation

The generation of the initial surface triangulation is composed of three steps. These steps are the computation of the parametric coordinates of the support points of the intersection curves, the initial region–wise surface triangulation and the processes of assembly, orientation and ridge identification for the surface triangulation.

The Ferguson patch technique is utilized to obtain the continuous parametric mapping required for the support surface [4,10,17], with the intersection curves parameterized with spline interpolation in a conforming manner [17]. Using this representation, curvature data which can be used to compute isotropic or anisotropic curvature–controlled mesh size, may be readily evaluated.

The starting point of the procedure is the formation of an initial planar triangular mesh in the parametric domain. Before any two dimensional triangular mesh generation technique is applied, points defining the intersection curves of a region need to be mapped to the parametric domain.

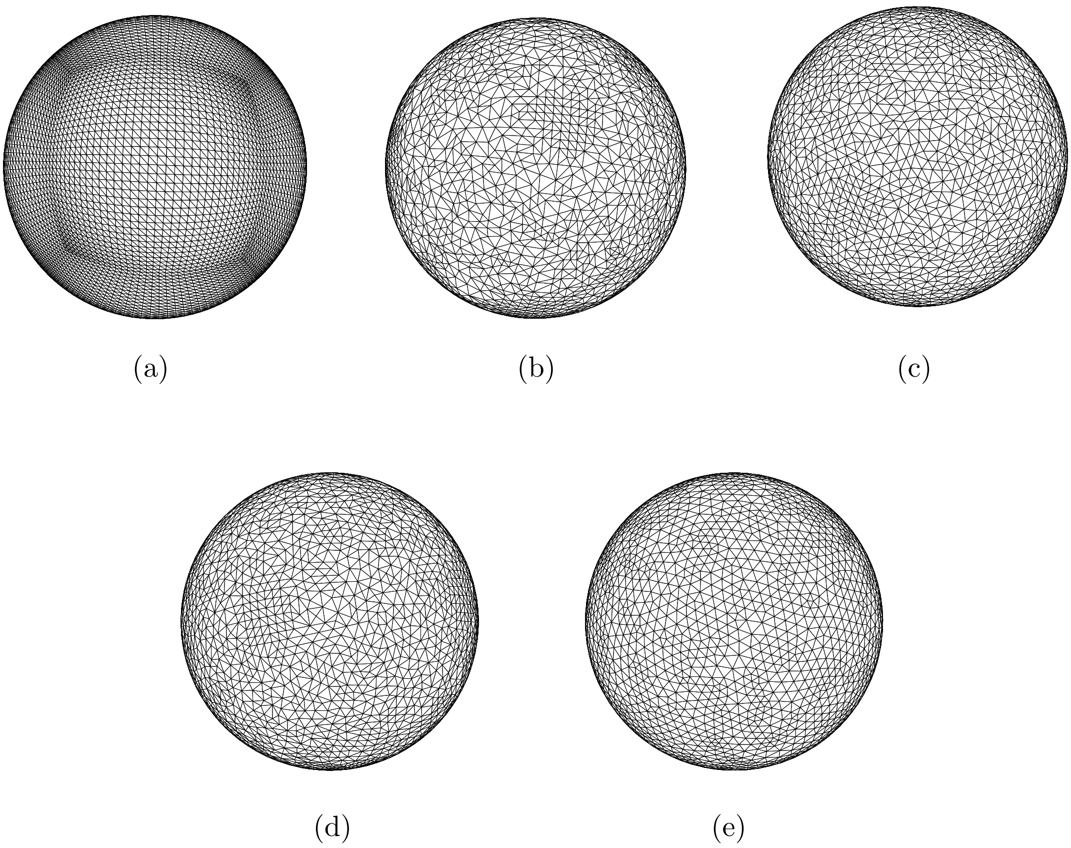

Figure 2: Illustration of the new method for a sphere: (a) initial surface triangulation; (b) surface triangulation after 1 remeshing step; (c) surface triangulation after 10 remeshing steps; (d) surface triangulation after 14 remeshing steps; (e) final mesh after optimization.

4.1 Computation of the parametric coordinates of the support points

The determination of the parametric coordinates of a vertex can be formulated as a point projection problem or, more precisely, a minimization problem. To determine the position of these intersection curve support points on the parametric plane, a simple steepest descent method is used [4]. The efficiency and success of this method depend upon the initial estimation and the search direction. The estimation process proceeds as follows:

1. A check is made to ensure that the point does not coincide with any singular point;
2. Starting from the four corners of the boundary of the support surface in turn, a search along the boundary is performed to find the minimum distance to the target point [19];
3. If the minimum distance is not within an acceptable tolerance, the point on the

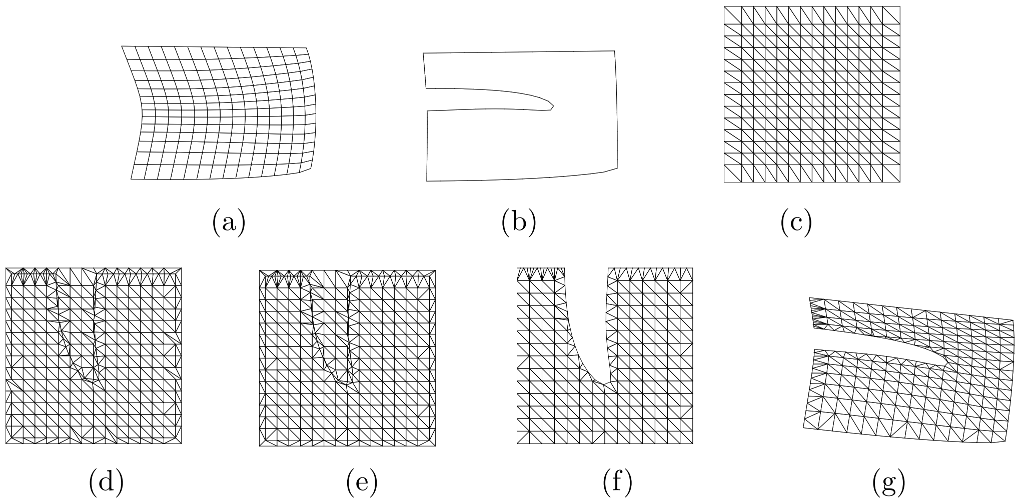

Figure 3: Illustration of the initial surface triangulation for a region: (a) the support surface of the region; (b) the intersection curves of the region; (c) the initial two dimensional triangulation of the images of all the points of the support surface(the boundary points are perturbed outwards to enclose all the image of the support points of the intersection curves); (d) the images of the support points of the intersection curves are inserted into the triangulation in a simple manner; (e) triangulation following edge swapping for boundary recovery; (f) triangulation for the image of the region; (g) final triangulation after mapping back to the surface.

boundary with the minimum distance is taken as the initial guess with a new search direction;

4. If the above steps fail to locate the target point, a slow projection procedure is performed to obtain the closest point to the target point.

4.2 Initial region — wise surface triangulation

The initial surface triangulation of each region contains both the support points of the intersection curves of the region and the points of the region's support surface which lie in the region. It may be achieved by employing, in the parametric domain of the support surface, any two dimensional triangulation method, such as a Delaunay or an advancing front method. Initially, as the input data is in the form of Ferguson patches (or Nurbs), the images of the points of the support surface in the parametric domain, i.e. the corresponding points in the parametric domain, provide a regular triangulation. Next, the boundary points of this triangulation are perturbed in a manner such that the triangulation covers all the images of the support points of the intersection curves of the region. Then, the images of all the support points of the intersection curves in the parametric domain are inserted into the triangulation in a simple manner. Finally, edge swapping is applied for boundary recovery and the triangulation of the region concerned is obtained by excluding all the triangles outside the image of the region.

Fig. 3 illustrates the procedure for generating an initial surface triangulation for a

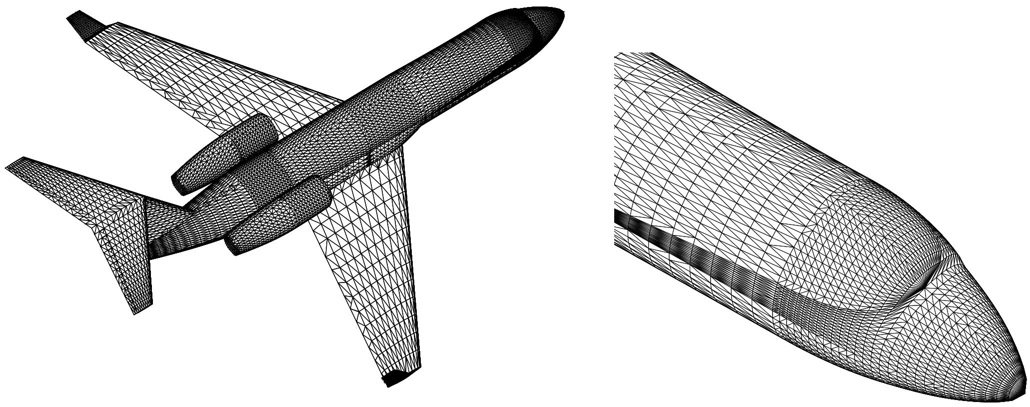

(a) (b)

Figure 4: Illustration of the initial surface triangulation of an aircraft configuration: (a) points of the support surfaces and the intersection curves are included; (b) local view of the surface triangulation of the area of the head.

typical region. After the planar triangulation, the surface triangulation of each region can be obtained by lifting the two dimensional vertices into three dimensions using the mapping determined by the patch.

4.3 Assembly, orientation and ridge identification for the surface triangulation

The triangulations of all the surface regions are assembled and common vertices are merged. A consistent surface orientation is constructed using surface integrals [1]. The surface which encloses the maximum volume is taken to be the outer boundary and all other closed loops are taken to be internal. The orientation is determined from the requirement that the normals to each triangle are pointing towards the domain.

All the edges in the triangulation are initially marked as non-ridges. Then, for each intersection edge, the dihedral angle of its two adjacent triangles is computed. If it is larger than a given criterion, the edge is considered to be a ridge. In this paper, 20 degrees is used for all examples. Of course, any user-specification of ridges can also be included in the above procedure. When all ridges are correctly identified, any point connecting more than two ridges is taken to be a corner point. In the subsequent remeshing, no ridges are allowed to be swapped and no corner points are allowed to be smoothed. But a ridge point can be allowed to move along the ridge.

For any internal point, i.e. non-ridge point or non-corner point, its normal is computed using the normals of its connecting triangles. Area weighted average is employed for this calculation. For the computation of the curvatures the method described in [12] is applied. This method utilizes quadratic surface interpolation through the connecting points of a vertex. With the curvatures, an isotropic or anisotropic sizing can be derived easily. But for a ridge point or a corner point, multi-normals and curvatures must be computed and stored. The triangles connecting the point are divided into several patches bounded by the ridges. Using the triangles forming each patch, a normal and curvature is computed. For isotropic meshing, if curvature is to be used to determine the spacing required at a given point, the minimum of the two spacings computed using the two principal directions of curvature is used. For ridge points, the spacing is taken to be the minimum of those obtained on the surrounding patches. These multiple treatments of normals and curvatures ensure that the refinement process will produce meshes which conform accurately to the given geometry.

An example of an assembled initial surface triangulation for a complete aircraft configuration, based upon the triangulation in the above method, is shown in Figs. 4(a) and (b), both for the global and local view (of the head).

This initial surface triangulation is obviously not suitable for use in a numerical simulation, as most triangles will not meet the mesh size or quality requirement. To produce an appropriate mesh, mesh enhancement is required. As the input Ferguson (or Nurbs) patch data approximates the surface with a very high accuracy, it is obvious that the initial triangulation is a very good approximation of the surface, hence it can be used as a reference mesh for further remeshing governed by a given sizing or surface curvatures. Now, the initial mapping defined by the Ferguson (or Nurbs) patches can be given up, and the required surface mesh can be generated by using the well developed remeshing techniques [1, 12, 13, 20] which will be discussed in the next section.

5 Remeshing based on local operations

The remeshing process employs edge splitting for mesh refinement and edge contraction for mesh coarsening, in addition to local mesh operations such as edge swapping. Node relocation or point smoothing is also applied to improve element quality [1, 12–14]. Here, only some unaddressed issues which are concerned about the robustness and efficiency of the procedure are discussed in detail.

The edge splitting and contraction are related to normalized edge lengths with respect to the computed spacing. The spacing can be defined by the use of sources [16], linear interpolation from boundary, octree-based interpolation, or some other kind of controls [1, 3]. Also, a curvature-controlled metric can be used for computation of the spacing to ensure certain geometric accuracy is achieved. Sizing derivation from curvature has been well developed in several papers [1, 5, 12, 21, 22]. Using the specified or computed spacing, the normalized edge length is computed [1,21]. A normalized edge length between $C_s = \sqrt{2}$ and $C_c = 1/\sqrt{2}$ is accepted while edges having normalized lengths larger that C_s are split

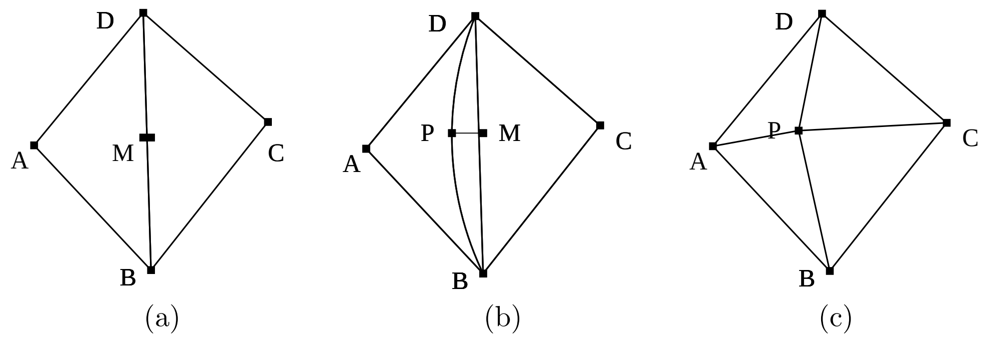

Figure 5: Edge splitting using G1 interpolation.

and edges with normalized lengths smaller that C_c are collapsed.

Before the remeshing, the sizing given or computed from curvatures can be smoothed with gradation, such as the H-correction proposed in [23]. This will reduce the ratio of neighboring element edge lengths. The smaller the gradation parameter, the smoother the final mesh.

In the following, four elementary mesh modification operators will be presented first, followed by a complete procedure of the mesh refinement and coarsening.

5.1 Edge splitting

Suppose an edge BD, shown in Fig. 5, has a normalized edge length larger than the given splitting criterion C_s, then a new point P on the surface will be introduced to split the edge into two smaller edges BP and PD. The triangulation will then be updated and the two adjacent triangles to the edge BD will be replaced by four new triangles which include the two new edges. The midpoint M of the edge BD is chosen and is projected onto the surface using the G1 interpolation [15, 16]. Fig. 5 shows the complete process of edge splitting. After the point insertion, neighboring edge swapping is locally performed to improve the geometry approximation as well as the element shape quality.

If the edge to be split is a ridge, the multi-normals are used to project the point onto the two adjacent surfaces using the G1 interpolation.

5.2 Edge contraction

When the normalized edge length of an edge is less than the collapse criterion C_c, the edge will be removed by merging the two points forming the edge under consideration. The merged point can be either one of the two end points or a new point located along the edge to be removed. This is illustrated in Fig. 6. In this paper, the merging technique shown in Fig. 6(c) is applied.

Special care is needed when collapsing edges which have one or more of their nodes on a ridge. In this case, the internal point will be moved to coincide with the ridge point.

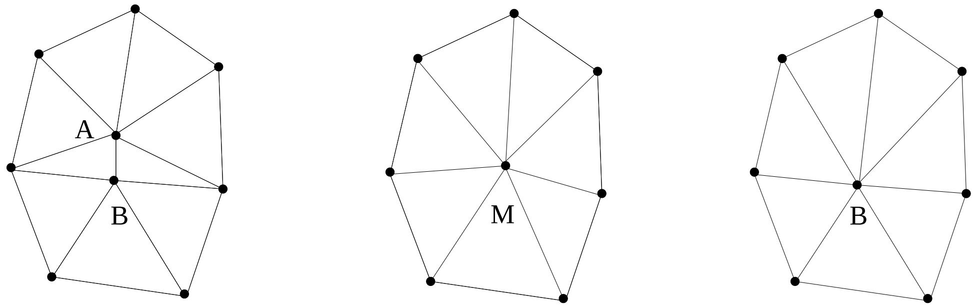

(a) (b) (c)

Figure 6: Two edge contraction schemes can be used: (a) AB is to be contracted; (b) the G1 interpolation of the midpoint is chosen; (c) B is chosen to be kept unchanged.

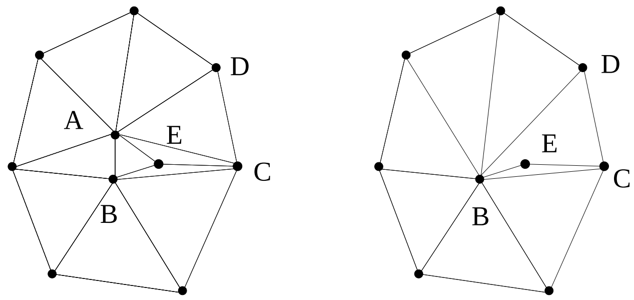

(a) (b)

Figure 7: An overlap is generated if the edge is contracted.

Furthermore, edges which have their two points on a ridge are considered only if the two points are two consecutive points on the ridge.

The difficulty in the edge collapse procedure for surface remeshing is highlighted in many publications [24, 25]. Continuous edge collapsing may result in self intersecting surfaces. Therefore, a rigorous validity check of the new local configuration has to be carried out before the collapse of an edge is accepted. Here, a robust procedure is proposed to ensure the consistency of the resultant triangulation:

1. Suppose AB is the edge to be collapsed and that vertex A is to be moved to B. The local configuration is initially checked to determine if there is a 3–division configuration neighboring the edge, as illustrated in Fig. 7. In this case, contraction is not allowed.
2. The signed areas of the triangles in the new configuration are calculated in the

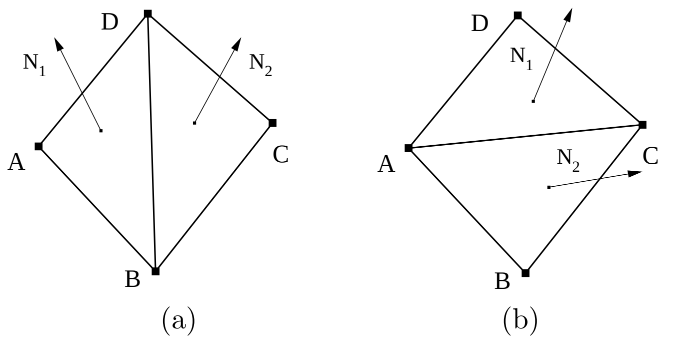

Figure 8: Edge Swapping: (a) the initial configuration; (b) the alternative configuration.

tangent plane at the point B. Any negative area indicates that overlapping has occurred and the contraction is not allowed.

3. A further check is performed in the physical space to ensure that the new configuration will not contain triangles with minimum angle tending to zero. In addition, the dihedral angles between all the neighboring triangles in the new configuration are determined and any configuration resulting in an angle larger than 45 degrees is rejected.

After contraction, neighboring edge swapping is applied to improve the shape of the new elements and the geometry approximation.

5.3 Edge swapping

The edge swapping process is illustrated in Fig. 8. The edge BD in Fig. 8(a) is identified as a non–ridge edge that connects the surface triangles ABD and BCD. The objective of edge swapping is to replace the connection BD by the diagonal edge AC. The configuration after the swapping is shown in Fig. 8(b). If the new configuration has better quality than the original one, the swapping operation is performed. Different measures may be adopted to analyze the quality of a configuration. Mesh quality measures can be classified as being either measures of element shape quality [1,26] or measures of geometric approximation [1,14,26]. Normally, in the initial refinement and coarsening, more attention is paid to the geometry approximation, while element quality becomes the main concern in the final optimization.

In the swapping procedure, the validity of the surface triangulation must be maintained. This is accomplished by using procedures similar to those employed when edge contraction is performed.

5.4 Point smoothing

Point smoothing, or relocation, modifies the position of a vertex without changing the topology of the triangles meeting at the vertex. The objective is to find an optimal position on the surface for the vertex, in the sense that this alternative configuration should have better quality, according to criteria such as minimal angle, shape quality, size conformity or others [1]. Here, only the element shape quality is considered and various smoothing methods can be used [1, 5, 27].

The smoothing takes place in the tangent plane constructed at the point under consideration [14]. Using the spring analogy method, a new point is located at the centroid of all the projected neighboring triangles. The triangle which contains the new node is identified and the node is projected onto the surface using G1 interpolation. Finally, the triangulation which will result from replacing the initial point by the new point are evaluated using the same geometric and shape measures used in the edge swapping. In order to ensure robust implementation of the smoothing procedure, a relaxation technique is implemented. A point located on a ridge is smoothed using the two adjacent points on the same ridge.

5.5 Refinement/coarsening procedure and final optimization

The refinement/coarsening procedure modifies an existing surface triangulation through edge splitting or contraction in an iterative manner. The objective is to ensure that elements are in better conformity with the size specification, with all the normalized edge lengths falling between the given splitting parameter C_s and contraction parameter C_c.

If the initial mesh is almost isotropic, sequential processing of the edges of the current mesh is feasible and efficient. However, the initial surface triangulation may contain very large aspect ratio triangles due to thin and narrow input patches. In this case, distorted local configurations with clusterings of many stretched triangles will be frequently produced. And performing further operations, such as edge swapping and edge collapsing, while maintaining the surface integrity, will not always be possible. To overcome this problem, a grouping technique is applied, in which the edges are sorted into groups according to their lengths. Within each group, there is no strict order. The group which contains the longest edges is processed first and each group is updated dynamically. This guarantees that the longest edges are refined first in global terms. The grouping technique results in a more stable and robust procedure for refinement/coarsening of the initial mesh.

As both splitting and coarsening are local operations, the computational time required increases linearly with the number of nodes added and edges merged. With the grouping technique, it is found that the total number of contraction operations is far less than the total number of splitting operations. Hence, the complete procedure is very efficient.

Finally, the surface triangulation is enhanced via a combination of edge swapping, point smoothing and node connectivity optimization. The combined application of these post-processing techniques is found to be very effective in improving the regularity and

Table 1: Small scale surface meshing: comparison of the CPU time requirements for an advancing front method and for the new method.

	Advancing Front		New Method	
Configuration	Elements	CPU	Elements	CPU
F6	125130	75s	134426	84s
Gulf2	63376	25s	79604	22s
F16	309672	120s	359964	125s
Ship	228608	97s	254212	107s

Table 2: Larger scale surface meshing: comparison of the CPU time requirements for an advancing front method and for the new method.

	Advancing Front		New Method	
Configuration	Elements	CPU	Elements	CPU
F6	703302	620s	714442	186s
Gulf2	521230	812s	512254	140s
F16	1024324	684s	1142216	297s
Ship	943326	647s	1224776	306s

smoothness of the surface mesh, with almost all points connected to six triangles. Element shape quality statistics for practical configurations are presented in the next section.

6 Application examples

The proposed new method has been applied to three aircraft and one ship configurations, which are all defined in terms of composite parametric surfaces in the FLITE format, using Ferguson patches [2]. Details of the surface meshes produced for generic Gulf 2, F6, F16 and Ship configurations are shown in Figs. 9, 10, 11 and 12 respectively. The element size distribution in each example is governed by prescribed sources, coupled with curvature control. And the sizing is smoothed by gradations.

To test the efficiency of the proposed method, the CPU time required for each example is compared with the time required by an advancing front based method [2](using the system FLITE3D developed by Prof.Oubay Hassan et al in the University of Wales Swansea) in Table 1 and in Table 2. For the relatively small scale meshes considered in Table 1, it is apparent that the method is almost as efficient as the advancing front method. For the F6 model, the advancing front method requires 75 seconds to generate 125 130 elements, while the proposed method needs 82 seconds to produce a mesh of 140 812 elements. The results for the Gulf 2, F16 and Ship models show similar trends. However, a different behaviour is produced when relatively larger surface meshes are generated, as demonstrated by the figures presented in Table 2. For example, for the F6 model, when the number of

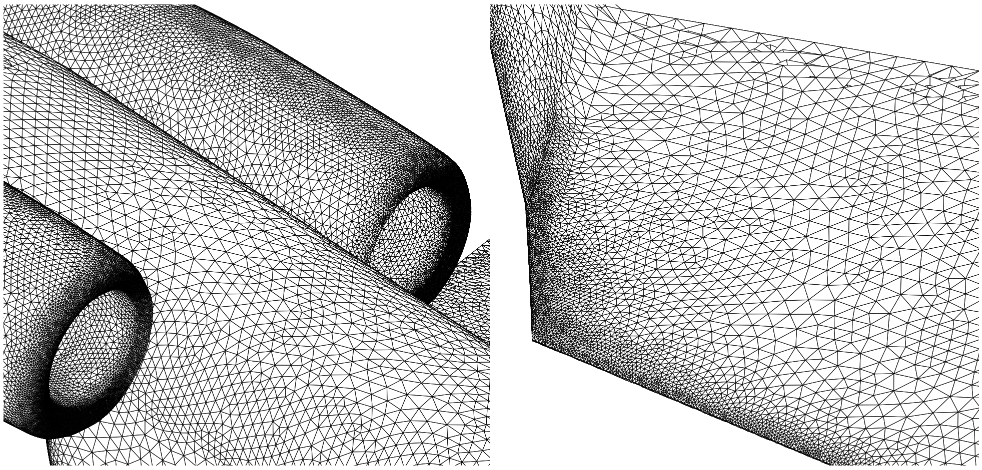

(a) (b)

Figure 9: Details showing a high quality surface triangulation for the Gulf 2 model.

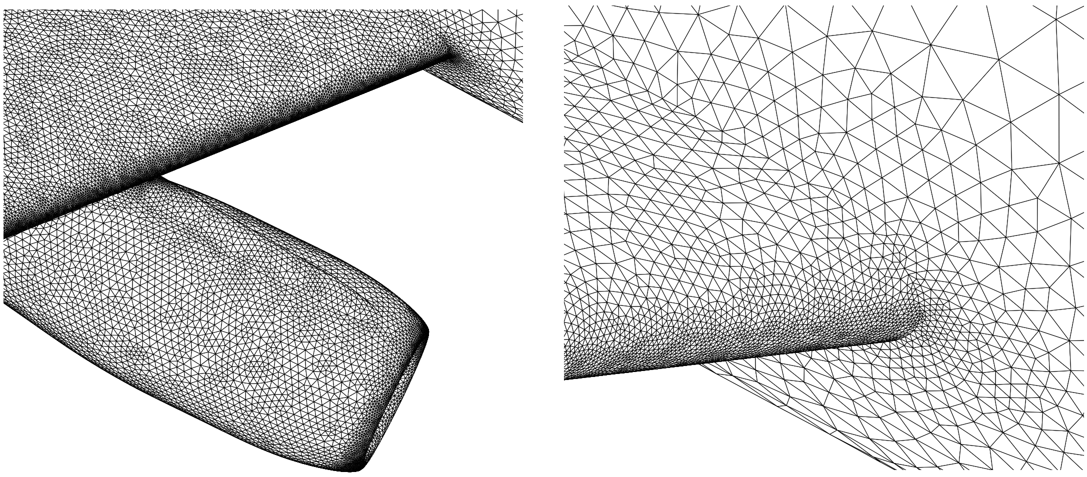

(a) (b)

Figure 10: Details showing a high quality surface triangulation for the F6 model.

elements generated is increased to 703 302, the CPU time needed by the advancing front method is 620 seconds. This is more than 8 times the figure shown in Table 1 for the smaller mesh for the same configuration. However, the time required to generate 714 442 elements for this model by the new method is only twice the corresponding figure presented in Table 1. Similar behavior is apparent for the other models. This difference in meshing efficiency behavior can be explained by recognizing that the number of operations required by the advancing front method in generating n elements can be approximated to be $\mathcal{O}(n \log n)$, where $\log n$ is related to the length of the current advancing front [1].

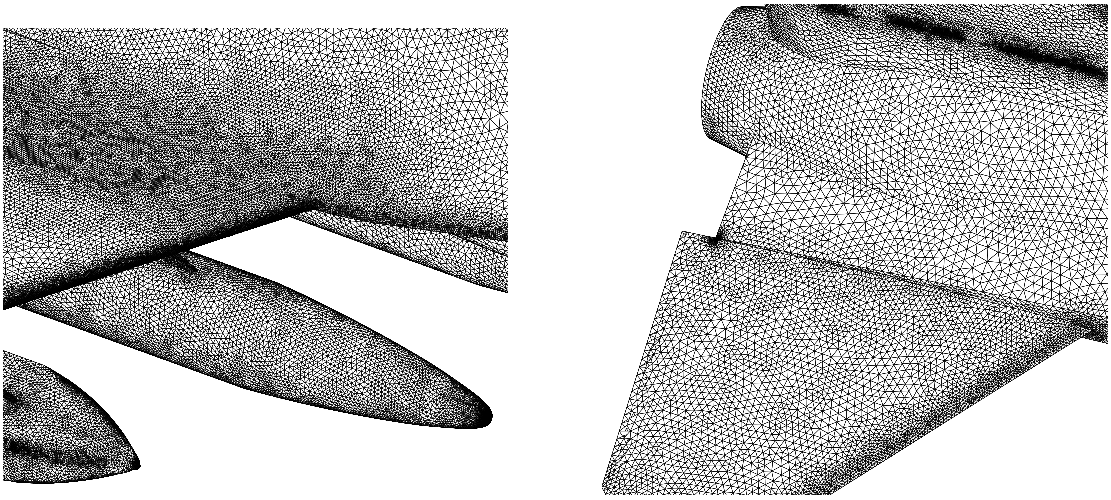

(a) (b)

Figure 11: Details showing a high quality surface triangulation for the F16 model.

And the operations required by the proposed new method is almost linear to the splitting operations(i.e. $\mathcal{O}(n)$), as the coarsening times is much less than that of edges splitting. So for small values of n, the number of operations required by the advancing front method is very similar to that required by the proposed method, once the initial surface triangulation has been accomplished. For larger values of n, the discrepancy in the operation count between the two methods becomes significant, especially for geometries involving regions which are meshed with a large number of elements, such as the Gulf 2 and the F6. For the F16, where there are 524 input patches which are all comparatively small, the difference in the efficiency of the two methods is not so great. For the Ship, where there are 1166 input patches among which there are more than 100 very small sliver patches, the meshing efficiency statistics is of the same characteristics as that of the F16. For general large scale surface grid generation involving an arbitrary number of patches, we can expect the proposed method to be significantly more efficient than the advancing front approach, independent of the size of the geometry.

The quality of the surface mesh elements generated by the two methods for the Gulf 2, F16 and the Ship configurations has also been analyzed. The meshes generated for the F6 configuration have been excluded from this comparison, as its surface definition contains fewer of the sliver–like patches which challenge the process of high quality meshing. The quality of a surface triangle e is measured in terms of the element quality Q_e defined as

$$Q_e = 4\sqrt{3}A_e / \sum_{i=1}^{3} \left(L_e^i\right)^2$$

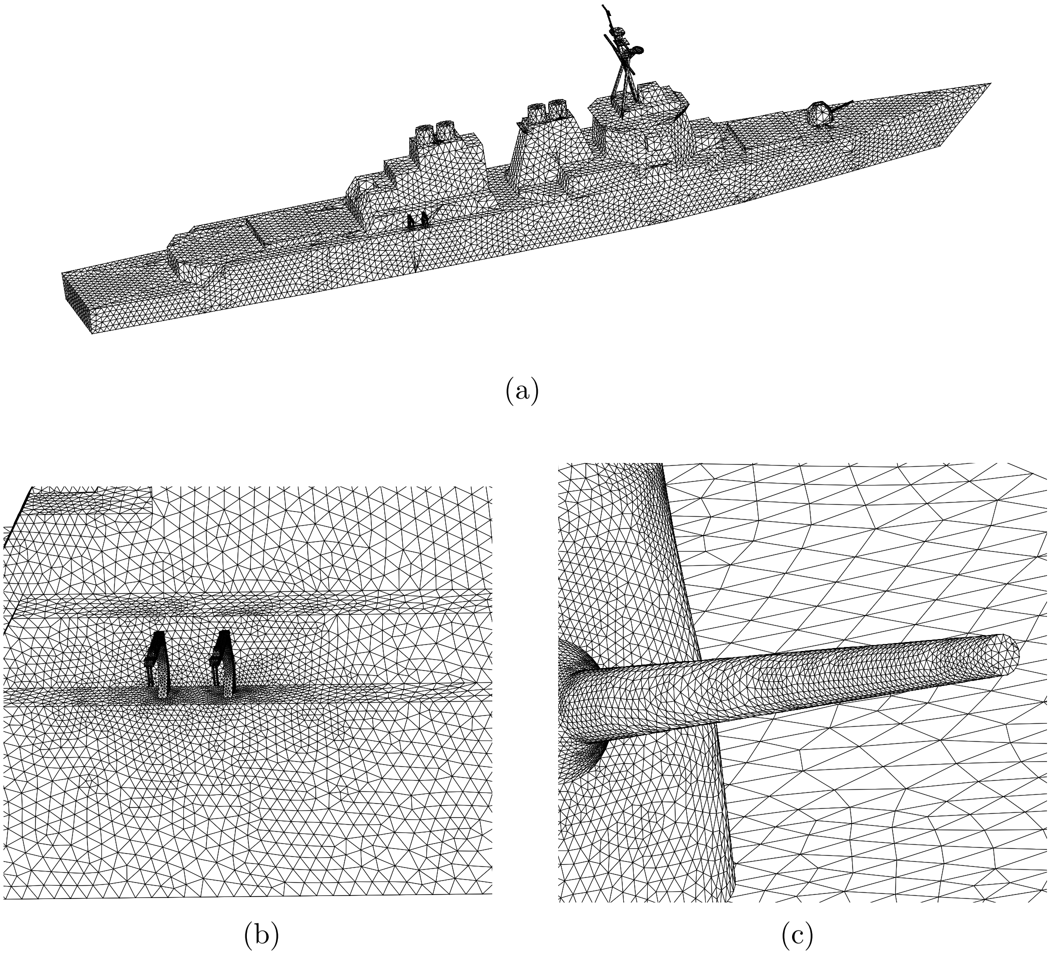

Figure 12: Surface mesh for a ship model: (a) Surface triangulation (with a lot of sliver triangles) generated in AFT; (b-c) Details showing a high quality surface triangulation generated in the new method.

where A_e is the area of triangle e and L_e^i, $i = 1, 2, 3$, denote the element edge lengths. It is well known that the advancing front method can produce a high quality distribution of points and, consequently, high quality surface meshes. These meshes are often superior to the meshes generated by alternative approaches, especially in terms of average element quality. For the four examples considered here, it is found that the average quality of the meshes generated by the new method is very similar to that for the meshes produced by the advancing front method.

For most numerical simulations, it is the element quality or the element angle that is significant when assessing the mesh quality. For this reason, we have computed the distribution of minimum element angles, in the range from 0 to 25 degrees, and the element quality. These distributions are shown in Fig. 13 for the Gulf 2 configuration, in Fig. 14

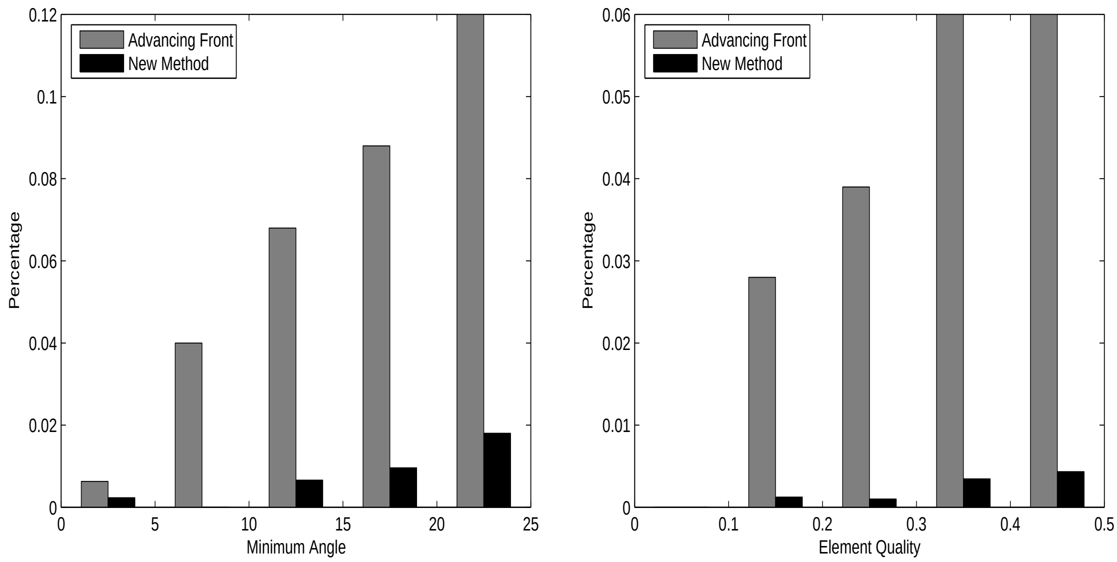

(a) (b)

Figure 13: Surface meshing of the Gulf 2 configuration by an advancing front method and the new method showing: (a) the minimum angle distribution; (b) the element quality distribution.

for the F16 configuration and in Fig. 15 for the Ship configuration. Due to the geometric constraint of input sliver patches, the surface mesh generated by the advancing front method has a number of small angled elements. This number is significantly reduced in the mesh created by the new method, which can be clearly seen in Figs. 13, 14 and 15.

7 Conclusions

A new method for high-quality triangular surface grid generation is proposed for composite parametric surfaces. The method is based on local mesh modifications involving edge splitting/contraction and edge swapping. It consists of three main steps: the generation of an initial surface triangulation which approximates the given surface with adequate accuracy and works as a reference mesh, the refinement/coarsening of the mesh applying local mesh modification, and mesh optimization. The proposed method is highly efficient as all operations are performed locally. The statistics presented for the application examples demonstrated that for complicated geometries the new approach provides an efficient robust procedure for the generation of high quality surface meshes. Future directions of research include the extension to the anisotropic case, the estimation of the approximation accuracy of the initial surface triangulation and the G1-interpolation, and the application to adaptive meshing including the treatment of moving boundaries.

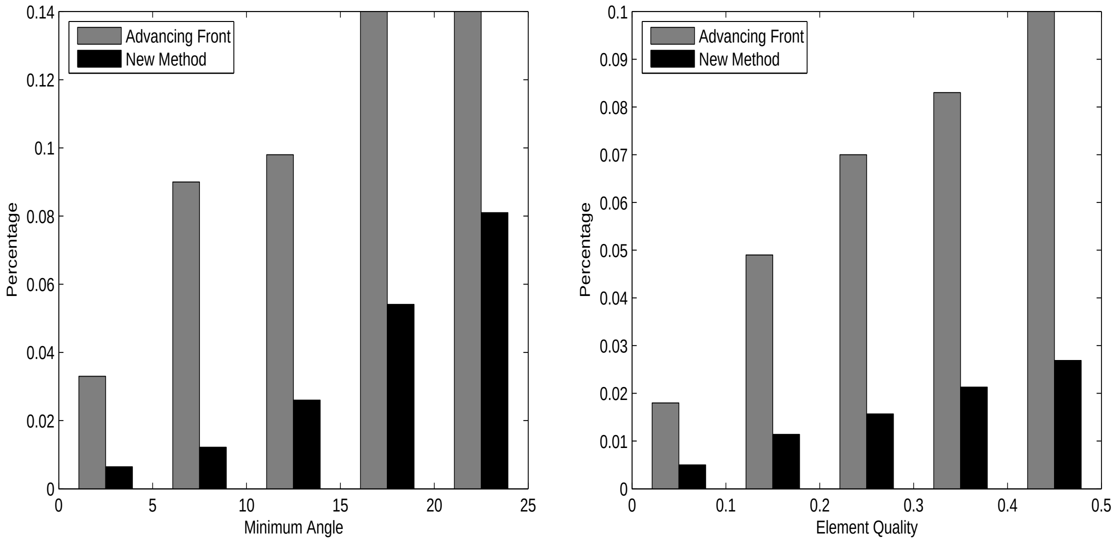

(a) (b)

Figure 14: Surface meshing of the F16 configuration by an advancing front method and the new method showing: (a) the minimum angle distribution; (b) the element quality distribution.

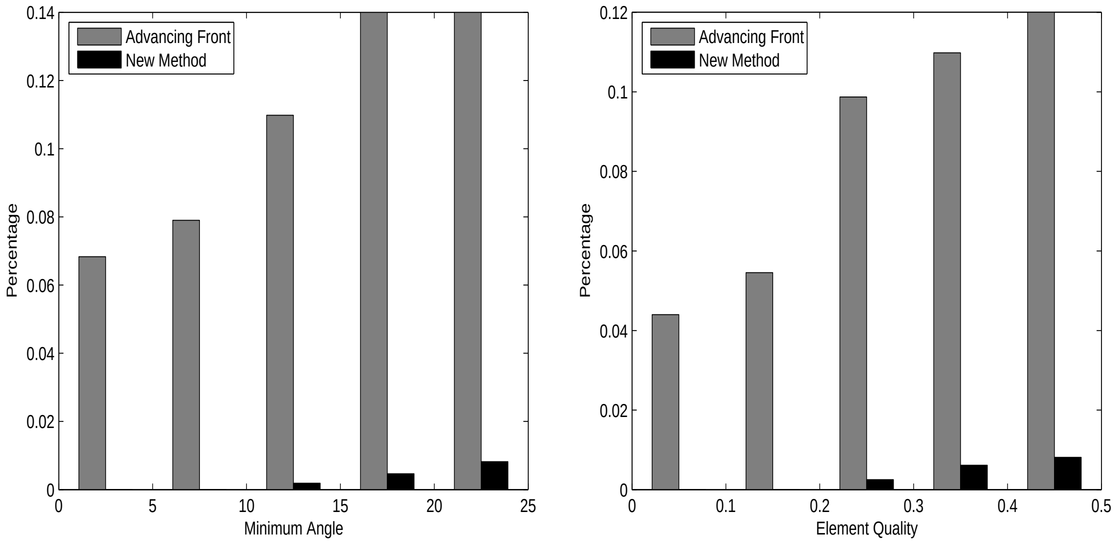

(a) (b)

Figure 15: Surface meshing of the Ship configuration by an advancing front method and the new method showing: (a) the minimum angle distribution; (b) the element quality distribution.

8 Acknowledgments

The author is grateful to the referees for the valuable suggestions and comments that helped to improve the presentation of this manuscript and would like to thank Prof.Oubay Hassan for providing the surface meshing code FLITE3D. This work was partially supported by the start-up grant M58110011 of NTU.

References

[1] J. Thompson, B. K. Soni and N. P. Weatherill, Handbook of Grid Generation, CRC Press LLC, 1999.

[2] J. Peiró, J. Peraire and K. Morgan, FELISA system reference manual. Part 1: Basic theory, University of Wales Swansea Report C/R/821/94, 1994.

[3] D. L. Marcum and J. A. Gaither, Unstructured surface grid generation using global mapping and physical space approximation, in: Proceedings of the 8th International Meshing Roundtable, Sandia National Laboratory, 1999, 397–406.

[4] Y. Zheng, N. P. Weatherill and O. Hassan, Topology abstraction of surface models for three-dimensional grid generation, Engrg. Comput., 17 (2001), 28–38.

[5] H. Borouchaki, P. Laug and P. L. George, Parametric surface meshing using a combined advancing front generalised Delaunay approach, Internat. J. Numer. Methods Engrg., 49 (2000), 233–259.

[6] S. H. Lo, Automatic mesh generation over intersecting surfaces, Internat. J. Numer. Methods Engrg., 38 (1995), 943–954.

[7] J. R. Tristano, S. J. Owen and S. A. Canann, Advancing front surface mesh generation in parametric space using Riemannian surface definition, in: Proceedings of 7th International Meshing Roundtable, Sandia National Laboratory, 1998, 429–445.

[8] K. Nakahashi and D. Sharov, Direct surface triangulation using the advancing front method, AIAA Paper 95–1686–CP, 1995.

[9] A. A. Shostko, R. Löhner and W. C. Sandberg, Surface triangulation over intersecting geometries, Internat. J. Numer. Methods Engrg., 44 (1999), 1359–1376.

[10] P. Laug and H. Boruchaki, Interpolating and meshing 3D surface grids, Internat. J. Numer. Methods Engrg., 58 (2003), 209–225.

[11] C. K. Lee, Automatic metric advancing front triangulation over curved surfaces, Engrg. Comput., 17 (2000), 48–74.

[12] P. J. Frey, About surface remeshing, in: Proceedings of the 9th International Meshing Roundtable, Sandia National Laboratory, 2000, 123–156.

[13] R. Löhner, Regridding surface triangulation, J. Comput. Phys., 126 (1996), 1–10.

[14] P. Frey and H. Borouchaki, Geometric surface mesh optimization, Comput. Vis. Sci, 1 (1998), 113–121.

[15] D. J. Walton and D. S. Meek, a trianular G^1 patch from boundary curves, Comput. Aided Design, 28 (1996), 113-123.

[16] N. P. Weatherill, M. J. Marchant, O. Hassan and D. L. Marcum, Grid adaptation using a distribution of sources applied to inviscid compressible flow simulations, Internat. J. Numer. Methods Engrg., 19 (1994), 739–764.

[17] J. Peters, C^1–surface splines, SIAM J. Numer. Anal., 32 (1995), 645–666.

[18] G. Farin, Curves and Surfaces for Computer Aided Geometric Design, Academic Press, New York, 1988.

[19] U. Tremel, Parallel Unstructured Adaptive Remeshing for Moving Boundary Problems, Ph.D. Thesis, University of Wales Swansea, 2005.

[20] H. L. de Cougny, Refinement and coarsening of surface meshes, Engrg. Comput., 14 (1998), 214–222.

[21] P. L. George and H. Borouchaki, Delaunay Triangulation and Meshing. Application to Finite Elements Methods, Hermès, Paris, 1998.

[22] C. K. Lee, On curvature element–size control in metric surface mesh generation, Internat. J. Numer. Methods Engrg., 50 (2001), 787–807.

[23] H. Borouchaki, F. Hecht and P. J. Frey, Mesh gradation control, Internat. J. Numer. Methods Engrg., 43 (1998), 1143–1157.

[24] Ito Y, Nakahashi K. Direct surface triangulation using Stereolithography(STL) Data, AIAA Paper, 2000-0924, 2000.

[25] Béchet C, Cuilliere JC, Trochu F. Generation of a finite element MESH from sterolithography(STL) files, Comput. Aided Design, 34 (2002), 1-17.

[26] P. J. Frey and H. Borouchaki, Surface mesh quality evaluation, Internat. J. Numer. Methods Engrg., 45 (1998), 101–118.

[27] P. M. Knupp, Achieving finite element mesh quality via optimization of the Jacobian matrix norm and associated quantitites. Part I–a framework for surface mesh optimization, Internat. J. Numer. Methods Engrg., 48 (2000), 401–420.