

SEMA SIMAI Springer series 30

Rubén Sevilla
Simona Perotto
Kenneth Morgan *Eds.*

Mesh Generation and Adaptation

Cutting-Edge Techniques

SēMA

SIMAI
SOCIETÀ ITALIANA DI MATEMATICA
APPLICATA E INDUSTRIALE




Springer

SEMA SIMAI Springer Series

Volume 30

Editors-in-Chief

José M. Arrieta, Departamento de Análisis Matemático y Matemática Aplicada, Facultad de Matemáticas, Universidad Complutense de Madrid, Madrid, Spain

Luca Formaggia , MOX–Department of Mathematics, Politecnico di Milano, Milano, Italy

Series Editors

Mats G. Larson, Department of Mathematics, Umeå University, Umeå, Sweden

Tere Martínez-Seara Alonso, Departament de Matemàtiques, Universitat Politècnica de Catalunya, Barcelona, Spain

Carlos Parés, Facultad de Ciencias, Universidad de Málaga, Málaga, Spain

Lorenzo Pareschi, Dipartimento di Matematica e Informatica, Università degli Studi di Ferrara, Ferrara, Italy

Andrea Tosin, Dipartimento di Scienze Matematiche “G. L. Lagrange”, Politecnico di Torino, Torino, Italy

Elena Vázquez-Cendón, Departamento de Matemática Aplicada, Universidade de Santiago de Compostela, A Coruña, Spain

Paolo Zunino, Dipartimento di Matematica, Politecnico di Milano, Milano, Italy

As of 2013, the SIMAI Springer Series opens to SEMA in order to publish a joint series aiming to publish advanced textbooks, research-level monographs and collected works that focus on applications of mathematics to social and industrial problems, including biology, medicine, engineering, environment and finance. Mathematical and numerical modeling is playing a crucial role in the solution of the complex and interrelated problems faced nowadays not only by researchers operating in the field of basic sciences, but also in more directly applied and industrial sectors. This series is meant to host selected contributions focusing on the relevance of mathematics in real life applications and to provide useful reference material to students, academic and industrial researchers at an international level. Interdisciplinary contributions, showing a fruitful collaboration of mathematicians with researchers of other fields to address complex applications, are welcomed in this series.

THE SERIES IS INDEXED IN SCOPUS

Rubén Sevilla • Simona Perotto • Kenneth Morgan
Editors

Mesh Generation and Adaptation

Cutting-Edge Techniques

 Springer

Editors

Rubén Sevilla
College of Engineering
Swansea University
Swansea, UK

Simona Perotto
Dipartimento di Matematica
Politecnico di Milano
Milan, Italy

Kenneth Morgan
College of Engineering
Swansea University
Swansea, UK

ISSN 2199-3041

ISSN 2199-305X (electronic)

SEMA SIMAI Springer Series

ISBN 978-3-030-92539-0

ISBN 978-3-030-92540-6 (eBook)

<https://doi.org/10.1007/978-3-030-92540-6>

© The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer Nature Switzerland AG 2022

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Foreword

This volume in the SEMA SIMAI Springer Series, entitled *Mesh Generation and Adaptation: Cutting-Edge Techniques*, is dedicated to Professor Oubay Hassan, on the occasion of his 60th birthday.

Oubay Hassan was born in Damascus, Syria, on May 3, 1960. He obtained his first degree in civil engineering from the University of Damascus in 1983. In 1985, he enrolled in the MSc course on the finite element method in the Department of Civil Engineering at Swansea University. He has remained at Swansea ever since.

As a practicing civil engineer, he was naturally attracted initially to the area of structural mechanics, and for his MSc thesis, he worked on the solution of non-linear problems involving reinforced concrete plates and shells [1]. However, he moved to a different area for his PhD studies, as he became interested in the unstructured mesh CFD research which was being carried out in the Department. Oubay addressed the problem of compressible viscous high-speed flow simulations, and this, with the accompanying difficulties of creating suitable meshes, initiated his interest in mesh generation and adaptivity. During his PhD studies, he combined the development of original ideas with skillful computer implementations. He became adept at manipulating unstructured meshes and devised a novel algorithm for creating continuous lines, made up of element sides, which pass once through each node of a general unstructured mesh. He was then able to use these lines as the basis for an implicit solution procedure in which the solution was achieved by line relaxation [2].

Although his initial work used the advancing front method, he made a major contribution to mesh generation with his research on the application of Delaunay



triangulation procedures to general three-dimensional configurations. This work resulted in a robust technique for creating a valid boundary conforming mesh of unstructured tetrahedra, with automatic point creation, for domains of arbitrary geometric complexity [3].

Following his appointment to the academic staff at Swansea, Oubay has continued to provide leadership, within what is now the Zienkiewicz Centre for Computational Engineering, on the development of CFD schemes able to operate effectively on unstructured grids. During this time, he has developed and enhanced his mesh-generation procedures [4] to such an extent that his advice and assistance is now regularly sought by companies and organizations in Europe, the USA, and the Far East.

The flexibility and generality of the mesh-generation tools that he has developed provided him with the possibility of making additional important contributions in the field of computational electromagnetics [5], as well as addressing various complex problems in CFD.



At his PhD award ceremony at Swansea with some distinguished colleagues.



Enjoying a coffee during a break at a conference in South Africa.

Oubay's research has been recognized in a number of different ways. In the 1990s, he used his CFD techniques to assist in the aerodynamic design process for the Thrust supersonic car. This car eventually took the World Land Speed Record beyond the speed of sound in October 1997. For his contribution to this project, Oubay was appointed a Member of the Most Excellent Order of the British Empire (MBE) by Queen Elizabeth. In 2012, he was appointed to the Fellowship of the UK Royal Academy of Engineering (FREng) and also to the Fellowship of the Learned Society of Wales (FLSW).

This volume represents a compilation of invited papers in the general area of mesh generation and adaptation, the research field in which Oubay has made the most profound and enduring contributions. The quality of the work and the range of material presented in these papers make this volume a fitting tribute to Oubay on the occasion of his 60th birthday.

MOX, Dipartimento di Matematica, Politecnico di Milano
Faculty of Science and Engineering, Swansea University

Luca Formaggia
Kenneth Morgan

References

1. Cervera, M., Hinton, E., Hassan, O.: Nonlinear analysis of reinforced concrete plate and shell structures using 20-noded isoparametric elements. *Comput. Struct.* **25**, 845–869 (1987)
2. Hassan, O., Morgan, K., Peraire, J.: An implicit finite–element method for high speed flows. *Int. J. Numer. Methods Eng.* **32**, 183–205 (1991)
3. Weatherill, N.P., Hassan, O.: Efficient three–dimensional Delaunay triangulation with automatic point creation and imposed boundary constraints. *Int. J. Numer. Methods Eng.* **37**, 2005–2039 (1994).
4. Xie, Z.Q., Sevilla, R., Hassan, O., Morgan, K.: The generation of arbitrary order curved meshes for 3D finite element analysis. *Comput. Mech.* **51**, 361–374 (2013)
5. Xie, Z.Q., Hassan, O., Morgan, K.: Tailoring unstructured meshes for use with a 3D time domain co–volume algorithm for computational electromagnetics. *Int. J. Numer. Methods Eng.* **87**, 48–65 (2011)

Contents

Mixed Order Mesh Curving	1
Steve L. Karman, Kristen Karman-Shoemake, and Carolyn D. Woeber	
A R&D Software Platform for Shape and Topology Optimization Using Body-Fitted Meshes	23
C. Nardoni, D. Danan, C. Mang, F. Bordeu, and J. Cortial	
Investigating Singularities in Hex Meshing	41
Dimitrios Papadimitrakis, Cecil G. Armstrong, Trevor T. Robinson, Alan Le Moigne, and Shahrokh Shahpar	
Intercode Hexahedral Meshing from Eulerian to Lagrangian Simulations	69
Nicolas Le Goff, Franck Ledoux, and Jean-Christophe Janodet	
Gmsh’s Approach to Robust Mesh Generation of Surfaces with Irregular Parametrizations	95
Jean-François Remacle and Christophe Geuzaine	
Adaptive Single- and Multilevel Stochastic Collocation Methods for Uncertain Gas Transport in Large-Scale Networks	113
Jens Lang, Pia Domschke, and Elisa Strauch	
HexDom: Polycube-Based Hexahedral-Dominant Mesh Generation	137
Yuxuan Yu, Jialei Ginny Liu, and Yongjie Jessica Zhang	
Mesh Adaptivity in the Framework of the Cartesian Grid Finite Element Method, cgFEM	157
Juan José Ródenas, Enrique Nadal, José Albelda, and Manuel Tur	
<i>h</i>- and <i>r</i>-Adaptation on Simplicial Meshes Using MMG Tools	183
Luca Arpaia, Héloïse Beaugendre, Luca Cirrottola, Algiane Froehly, Marco Lorini, Léo Nouveau, and Mario Ricchiuto	

Geometry and Adaptive Mesh Update Procedures for Ballistics Simulations	209
Saurabh Tendulkar, Fan Yang, Rocco Nastasia, Mark W. Beall, Assad A. Oberai, Mark S. Shephard, and Onkar Sahni	
High-Order Implicit Shock Tracking (HOIST)	233
Andrew Shi, Per-Olof Persson, and Matthew J. Zahr	
Breakthrough ‘Workarounds’ in Unstructured Mesh Generation	261
Rainald Löhner	
An Adaptive Conservative Moving Mesh Method	277
Simone Appella, Chris Budd, and Tristan Pryer	
A Global Optimization and Adaptivity-Based Algorithm for Automated Edge Grid Generation	301
Suzanne M. Shontz and David McLaurin	

Mixed Order Mesh Curving



Steve L. Karman, Kristen Karman-Shoemake, and Carolyn D. Woeber

Abstract Linear hybrid unstructured meshes are elevated to mixed-order meshes in response to geometry curvature. The linear meshes are elevated to the required degree on an element-by-element basis in regions of high geometry curvature. Weighted condition number mesh smoothing is used to untangle and improve the quality of the current mixed-order mesh. Periodically the mesh is tested for additional element elevation using a deviation criterion. Once the mesh smoothing is complete the mesh can be exported as a mixed order mesh or uniformly elevated to the desired degree. Details of the mesh elevation and smoothing process are described. Two three-dimensional examples are included that demonstrate the effectiveness of the method to produce high quality mixed-order meshes.

1 Introduction

High order mesh curving is an emerging technology that will greatly benefit those that utilize Finite-Element Methods (FEM) within the Computational Fluid Dynamics (CFD) solver community. Finite-element techniques offer increased accuracy with lower element counts over traditional CFD methods such as finite-volume and finite-difference methods. The increased accuracy is achieved by introducing additional vertices (new degrees of freedom) to edges, faces and interiors of linear elements. For elements adjacent to curved geometry these new degrees of freedom must lie on the geometry, thereby altering the shape of the original linear element. This process is more difficult when the mesh contains clustering of elements toward viscous boundaries. The edges and faces of interior elements must also be curved in response to the boundary element curvature to prevent element inversion.

S. L. Karman
Oak Ridge National Laboratory, Oak Ridge, TN, USA
e-mail: karmansl@ornl.gov

K. Karman-Shoemake · C. D. Woeber (✉)
Cadence Design Systems, Fort Worth, TX, USA
e-mail: kshoema@cadence.com; cwoeber@cadence.com

Research into mesh curving is taking place at a number of institutions. Radial basis function interpolation was investigated at Imperial College [1] and the University of Kansas [2, 3]. The more prominent mesh curving approaches tend to use solid mechanics analogies where the mesh is treated as an elastic solid that deforms due to forces acting on the boundaries [4, 5]. Other efforts focus on the solution to the Winslow equations to perform the interior mesh curving [6]. This approach is a natural application of Winslow smoothing techniques in the sense that a copy of the unperturbed, elevated mesh serves as the computational mesh. The solution to the Winslow equations then forces the interior of the physical mesh to take on the same character of the computational mesh. A novel mesh optimization approach with edge and face flips for moving P2 meshes and two-dimensional quadratic mesh adaptation to a Riemannian metric were developed by INRIA [7] and Gmsh [8]. Researchers at the Barcelona Supercomputing Center developed a mesh optimization method that attempts to minimize distortion [9]. Pointwise collaborated with researchers from the University of Tennessee Knoxville on an alternate approach for viscous mesh curving using weighted condition number (WCN) smoothing [10]. Subsequent updates to the technique were developed in 2018 and form the basis of the work presented in this chapter [11, 12].

Mesh generation applications with curving capabilities are now available on a limited basis to the CFD community. MeshCurve, developed as part of a Master's degree research project, is available for download [3]. Gmsh is a full featured mesh generation and visualization tool with curving capability to very high order [13]. Nektar++ has a meshing component, NekMesh, that has curving capabilities [14]. Pointwise recently released a version of their mesh generation software with an elevate-on-export capability [15].

Research on the WCN approach used by Pointwise has continued and permits mixed order meshes to resolve geometry curvature. The elements can be elevated to a maximum polynomial degree 4 (quartic) near highly curved geometry while far away from curved geometry, the elements remain linear. The mesh smoothing method uses a cost function to enforce desired element shapes and positive Jacobians across each element. Viscous mesh spacing is maintained as the elements are curved near the geometry. At completion, the mixed order mesh can be exported or a uniformly elevated mesh of the desired degree can be created. Results are shown for complex 3D configurations.

2 Mixed-Order Curving Framework

Elevating linear meshes and curving them in response to surface curvature requires easy access to the geometry and a robust initialization and smoothing process. Surface queries of the geometry are necessary to ensure the high-order nodes are accurately placed on the geometry during initialization and remain on the surface during mesh smoothing. The mesh smoothing process must be robust to ensure a valid mesh is produced that maintains the character of the input linear mesh with

respect to the distribution of nodes normal to the surface in the boundary layer region. Geometry access will be briefly described followed by a detailed description of the mixed order mesh curving process.

2.1 Geometry

Geometry access for elevating and smoothing is provided through the MeshLink API [16]. MeshLink is a library for managing geometry¹ and mesh data and provides a simple interface to query functions pertinent to mesh generation and mesh adaptation applications. Associations between geometric entities within a geometry data file and the mesh elements within a related mesh data file are stored in a MeshLink file. The complete geometry and mesh configuration can then be represented with the MeshLink file (XML), the geometry data file (CAD), and the mesh data file (NMB).

A key benefit of the MeshLink library that is integral to the elevation and smoothing process is the ability to define and use geometry groups. Geometry groups enable a mesh entity to be associated with multiple geometry entities that should be considered for a projection process. For instance, a surface mesh edge may be associated with a geometry face or a specific geometry curve. Alternatively, a surface mesh edge may be associated with several geometry curves. The mesh curving application program's efficiency depends on rapidly being able to query the correct geometry entity for a given mesh node projection without having to keep track of the details of the multiple geometry associations.

As an example, when the mesh curving program starts, the MeshLink API imports the CAD and XML file to create a database in memory that associates the surface elements of the mesh with the CAD entities. During element elevation and mesh smoothing the curving program makes node projection requests from the MeshLink API for nodes at surface mesh edges and faces. The queries include the end nodes of the edges and the corner nodes of the faces. The appropriate geometric entity is used by the library to project the requested query node. All of this is hidden from the mesh curving application program. All that was provided to the MeshLink function was the forming nodes for the mesh entity, edge or face, and the input query physical location. The process is more efficient than projecting to all geometry surfaces and more robust. If the linear mesh topology does not change, which is the case for this implementation of mesh curving, then the nodes are projected to the proper geometry entity. There is no ambiguity about projecting to the wrong surface, such as the lower wing surface from a node on the upper surface near a thin trailing edge.

¹ The intended application of this technique is mesh curving on NURB geometric surfaces but does not preclude use of discrete surfaces.

For the results presented within this chapter, the Pointwise meshing software was used to create the initial linear meshes. At completion the three files required by MeshLink were exported from Pointwise: the linear mesh in CGNS format, the geometry CAD file in NMB format, and the XML MeshLink file.

2.2 Mixed Order Curving Process

Mixed order mesh curving uses a process that begins with a valid linear mesh. The major components of the process are provided in the flowchart in Fig. 1. Within the flowchart and this chapter, note that the order or polynomial degree of an element is indicated using Q1 through Q4 nomenclature. Linear, quadratic, cubic, and quartic elements are Q1, Q2, Q3, and Q4 respectively. The high-order elements use Lagrangian basis functions to evenly distribute high-order nodes across the element's edges, faces and interior. These physical nodes are an integral part of the WCN method to enforce sub-element and element shapes.

The Initialization process, seen in the shaded box on the left in Fig. 1, uses the input linear mesh to begin walking through the element elevation process based on the user requested final degree, Q_{final} . The initialization process elevates elements in the mesh to the next higher degree depending on the deviation metric evaluation performed in Deviation Metric Testing process (see shaded box in the middle in Fig. 1). The process first elevates surface elements (2.2.1) and volume elements

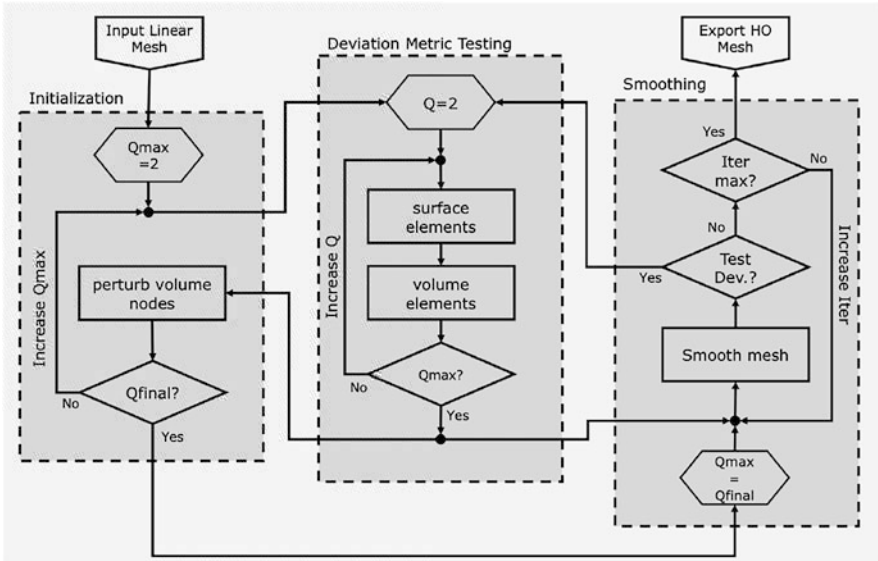


Fig. 1 Flowchart of the mixed order mesh curving process




Fig. 2 Leading edge of the Onera M6 wing at the symmetry plane

(2.2.3) to Q2. The new boundary nodes are placed on the geometry surface and the perturbations of these nodes from the initial linear surface are propagated into the interior using a simple transfer process. The initialization continues to Q3 and possibly Q4,² if requested, using the same deviation metric testing process and interior node perturbation process. Once the bootstrapping process is complete a mixed order mesh is produced that may include invalid elements near highly curved geometry.

The Smoothing process, seen in the shaded box on the right in Fig. 1, uses the WCN mesh smoothing method to correct any element inversions and improve the quality of the elements produced by the initialization process. Periodically each volume element is measured using the deviation metric testing in the middle box to determine whether additional elevation is warranted, not to exceed the specified maximum polynomial degree. The mesh smoothing phase is completed when all elements meet the deviation criterion and the mesh smoothing process converges.

The final output from the elevation and smoothing process is a mesh that contains high-order nodes that are shared between elements of the same order. Faces and edges shared between elements of different order will not share the same interface nodes. Shape conformity at these interfaces is imposed before export. At this point the mesh is exported in the appropriate high order mesh file format. An example of a mixed-order mesh created with this process is shown in Fig. 2 for a Q1–Q4

² The highest surface and volume element degree will be Q4 even if the surface polynomial degree is higher.

mesh on the Onera M6 wing leading edge: the light gray elements are Q1 and the dark gray elements are Q4. The elements at the leading edge, where the curvature is highest, are quartic and the element order decreases away from the leading edge as the curvature decreases. A quality constraint ensures the degree jump between elements is limited to one.

2.2.1 Surface Element Deviation Metric

A deviation metric is used to control the p-refinement (element elevation) process during initialization and as part of the mesh smoothing process. The deviation metric measures the displacement of test nodes on the edges and faces of an element adjacent to either a curved boundary or an adjacent volume element. If the element is on a curved boundary test nodes are computed at quadrature integration points of the surface element and projected to the geometry. The deviation amount is demonstrated in Fig. 3 where a test node at the centroid of a linear triangle (dark gray) is projected to a curved geometry surface (light gray). If the displacement of this test node exceeds a threshold distance for the adjacent volume element, then elevation is indicated. The threshold amount triggering elevation is the minimum linear edge length within the element multiplied by an input deviation threshold parameter, typically 1–5%.




Fig. 3 Test node at centroid of surface element projected to the geometry




Fig. 4 Sixth order Gauss points for reference (a) triangles and (b) quadrilaterals

Each surface element is examined at 6th order quadrature Gauss point locations, shown in Fig. 4a for a reference triangle and Fig. 4b for a reference quadrilateral. The physical coordinates at these Gauss points are computed using the 2D basis functions for the current element order and physical nodes in the surface element. If the deviation of any Gauss point from the surface geometry is farther than the threshold, then the surface element (and adjacent volume element) is marked for elevation to the next higher order.

2.2.2 Shape Conformity Metric

The surface element deviation metric described in Sect. 2.2.1 is also used to define the shape conformity metric, which measures how well the discrete curved surface matches the underlying geometry. It is defined as the integration of the difference between the mesh surface and the geometry surface over the surface triangular or quadrilateral element. Equation (1) integrates the distance from a mesh node to the geometry over the surface of an element using numerical integration. The numerator results in the volume of the space between the mesh surface and the geometry. The denominator is the surface area. Combined the quantity is the average distance between the mesh and the surface.

$$SC = \frac{\iint |\vec{d}\vec{r}| ds}{\iint ds} \quad (1)$$

The distance given by $|\vec{d}\vec{r}|$ is the surface element deviation, as described in Sect. 2.2.1 and illustrated in Fig. 3. Numerical integration is performed using

Gaussian quadrature over each surface element using the 6th order quadrature points shown in Fig. 4a,b for surface triangles and quadrilaterals.

The shape conformity metric produces a dimensional quantity in the units of the mesh length scale. For flat planar surfaces all mesh orders should produce machine zero values, indicating the mesh is on the planar surface. For curved boundaries the linear mesh should exhibit the largest error and increasing the mesh order should produce smaller error values.

2.2.3 Volume Element Deviation Metric

Volume elements with neighbors of a different order need deviation testing to ensure the shapes at the interface are similar. This process will also propagate high curvature regions of the geometry into the volume. At these interfaces, the deviation test is performed in one of two ways: either testing the lower order nodes against the higher order shape or vice versa. When performing mesh smoothing the nodes on the faces and edges of the lower order element are projected onto the adjacent, higher order shape. Figure 5a is used to illustrate. The high-order nodes on the face common to elements of different order are not shared by each element. The small dots along edges and in faces are the high-order nodes. At the edge between a quartic and cubic element there are two high-order nodes from the cubic element and three high-order nodes from the quartic element. Nodes on the cubic element must be forced to adhere to the quartic shape. Otherwise, the mesh smoothing process will drive the cubic nodes towards the original linear shape, as shown in Fig. 5b. For highly clustered meshes this would cross over the geometry, resulting in an invalid mesh. The same is true for the nodes at the interface of the Q3 and Q2 elements. Linear (Q1) elements have no face interior or edge interior nodes, so no shape enforcement is required.

Periodically, during mesh smoothing, the deviation test is performed where the nodes on the higher order side are tested against the lower order shape. The parametric coordinates of the higher order nodes are used to compute the




Fig. 5 The deviation test is used to ensure both that the presence of the curved surface is felt on the interior and that interfaces between lower order elements and higher order elements match. (a) Transition gaps. (b) Linear shape. (c) Lower order shape enforced

physical coordinates on the lower order shape, using the lower order basis function and physical nodes. If the distance from the current location to the lower order shape location exceeds the deviation threshold then the lower order element at that interface is marked for elevation. When the two shapes at element interfaces converge to within the tolerance then p-refinement stops. At the completion of mesh smoothing, before export, the higher order nodes at these mixed order interfaces are projected (not just measured) to the lower order shapes. The final example mesh, after additional p-refinement and mesh smoothing, is shown in Fig. 5c. All nodes on the interfaces between elements of different order appear to lie on the same shape.

2.2.4 Geometry Driven Mesh Perturbations

During the initialization process the elements near curved geometry are tested using the deviation criterion and the current maximum allowed degree. When an element is elevated the perturbations produced by the geometry are spread to other high-order nodes in the element using a simple transfer process with a linearly decaying rate. An iterative process of spreading these perturbations takes place where the deviation test is performed on adjacent elements. Neighboring elements use the volume deviation metric to sense the change in the shape of these newly elevated elements. This may trigger additional elements to elevate and the spreading continues. The requirement that the difference in degree between adjacent elements is limited to one continues to be enforced. These deviation and order difference tests quickly transfer the geometry perturbation into the volume. This initialization process may still result in mesh crossing near highly curved geometry, so mesh smoothing is required to ensure a valid mesh is produced.

2.2.5 Iterative Perturbation-Based Smoothing

Recent modifications to the smoothing method have resulted in a more robust technique for ensuring a valid computational mesh [12]. The basic smoothing method attempts to enforce shapes, derived from the original linear mesh, on the elevated high-order mesh. This is the weighted condition number (WCN) component of the cost function. The smoothing also imposes element size control through a normalized Jacobian-based component of the cost function. The overall cost function divides the normalized Jacobian by the weighted condition number, shown in (2) [17]. This function is computed on sub-triangles of each surface element and sub-tetrahedra of each volume element.

$$C = \frac{\min \left[1, \frac{J_p}{J_c} \right]}{WCN} \quad (2)$$

The numerator is the normalized Jacobian, a ratio of the determinants of the Jacobian matrices. The Jacobian matrix, given in (3), is computed at survey locations across the element using the appropriate Lagrangian basis functions. Subscript c refers to the computational mesh. This is the original mesh elevated but with no perturbations applied to any node, i.e. a copy of the physical mesh with straight-sides and no displacements. This mesh is assumed to be always valid with positive volumes and Jacobians. Subscript p refers to the physical mesh. When the physical Jacobian at a given survey point is less than the computational Jacobian the ratio will be less than one and the smoothing scheme will strive to increase it. This strongly influences node movement in tight viscous boundary layer regions of the mesh. The ratio is capped at 1, so larger ratios are permitted.

$$J = \begin{bmatrix} \sum \frac{\partial N_i}{\partial \xi} x_i & \sum \frac{\partial N_i}{\partial \xi} y_i & \sum \frac{\partial N_i}{\partial \xi} z_i \\ \sum \frac{\partial N_i}{\partial \eta} x_i & \sum \frac{\partial N_i}{\partial \eta} y_i & \sum \frac{\partial N_i}{\partial \eta} z_i \\ \sum \frac{\partial N_i}{\partial \zeta} x_i & \sum \frac{\partial N_i}{\partial \zeta} y_i & \sum \frac{\partial N_i}{\partial \zeta} z_i \end{bmatrix} \quad (3)$$

Mesh smoothing is applied to all nodes in the mesh. Surface nodal values of the cost function are computed using a biased average of sub-triangle cost value surrounding each surface node. Interior volume nodal values of the cost function are computed using a biased average of the sub-tetrahedra cost value surrounding each node. The surface (2D) cost function is defined first, followed by the volume (3D) cost function. Modifications to this smoothing method involved splitting the surface and volume mesh smoothing by leveraging separate weighted condition number and normalized Jacobian calculations for surface and volume elements.

Surface Weighted Condition Number

The surface weighted condition number enforces the shape of the surface element. The WCN component seen in the denominator of (2) serves as the weighted condition number for triangles given in (4). This quantity is computed for a sub-triangle of the high-order elements. All high-order surface elements, triangles and quadrilaterals, can be decomposed into sub-triangles.

$$WCN = \frac{\|AW^{-1}\| \|WA^{-1}\|}{2} \quad (4)$$

The bracketed quantities in (4) are the Frobenius norms of the matrix products. The W matrix is the weight matrix derived from computational coordinates for the sub-triangle. This defines the desired shape of the sub-triangle. Figure 6 shows the weight matrix where the entries of W are computed using the lengths of the three edges of a general triangle. The A matrix is the same as the W matrix using the physical coordinates. Even though these are two-dimensional elements, the matrices are defined using three-dimensional lengths of surface triangle edges.




Fig. 6 The weight matrix can be formed from the computational edge lengths of the triangle




Fig. 7 Quartic triangle. Computational element on the left. Physical element on the right

Surface Normalized Jacobian

The surface normalized Jacobian imposes size control for the surface element. For surface elements, it represents a ratio of surface normal vectors from computational and physical space at the same parametric coordinates of the element. To illustrate, a quartic triangle is shown in Fig. 7. The computational mesh is shown on the left. The physical, curved mesh is shown on the right. All 15 nodes of the elements are shown at the corners, edges and face interior. The rows in the Jacobian matrix from (3) represent the directions of the parametric coordinates in computational or physical space. For surface elements r_ξ (first row of the matrix) and r_η (second row of the matrix) are computed from the Lagrangian basis functions. This can be computed at any location in the element. The r_ζ vector for surface elements is computed as the cross product of the r_ξ and r_η vectors and represents the surface normal direction. The normalized Jacobian value used in (2) for surface elements is the ratio of magnitudes of these two vectors. The sign is taken from the dot product of the vectors to indicate a surface element inversion when the value is negative.

The normalized Jacobian part of the cost function can be computed in a number of ways. The simplest and least expensive computes the area ratio of the sub-triangles, shown in Fig. 8a. Only the nodes in the sub-triangle are involved in the calculation. This is consistent with the weighted condition number calculation. The assumption is that the sub-triangle area will provide enough influence to ensure the actual Jacobians remain positive throughout the entire element.




Fig. 8 The normalized Jacobian is computed using the area of sub-triangle or a biased weighting of the Jacobian evaluated at the quadrature points shown for 1 and 2 subdivisions. (a) Linear area. (b) One subdivision. (c) Two subdivisions

An alternate approach uses subdivision levels in the sub-triangle to create a grid of quadrature points. The normalized Jacobian is computed using all nodes of the element. Figure 8b,c show the quadrature point locations for 1 and 2 subdivision levels, respectively. Higher subdivision levels are possible. The normalized Jacobian value reported for the subdivision level is a biased average of the values at the quadrature points within the sub-triangle. The biased weighting is given by (5), where C and F are the normalized Jacobian values assuming the range is from 0 to 1.

$$F = F_{min} (1 - C_{min}) + F_{avg} (C_{min}) \quad (5)$$

The weighting biases the minimum value over the simple averaged value. If a negative value is detected the minimum value is returned as the cost function without computing the WCN component. In these cases, the mesh smoothing scheme is forced to correct element inversions first before enforcing element shape.

Volume Weighted Condition Number

The volume weighted condition number enforces the shape of the volume element. The WCN component seen in the denominator of (2) serves as the weighted condition number given in (6). This quantity is computed for a sub-tetrahedron of the high-order element. All high-order elements can be decomposed into smaller hexes, prisms, pyramids and tetrahedra. For sub-elements other than tetrahedra, the corners of the sub-element are used to form the tetrahedra in the cost function calculation.

$$WCN = \frac{\|AW^{-1}\| \|WA^{-1}\|}{3} \quad (6)$$

The bracketed quantities in (6) are the Frobenius norms of the matrix products. The W matrix is the weight matrix derived from computational coordinates for the same sub-tetrahedron. This defines the desired shape of the tetrahedron. Figure 9 shows the weight matrix where the entries of W are computed using the lengths of




Fig. 9 The weight matrix can be formed from the computational edge lengths of the tetrahedron

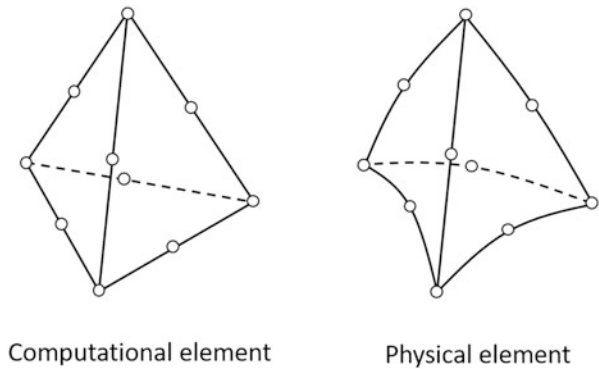


Fig. 10 Quadratic tetrahedron. Computational element on the left. Physical element on the right

the six edges of a general tetrahedron. The A matrix is the same as the W matrix using the physical coordinates.

Volume Normalized Jacobian

The volume normalized Jacobian imposes size control for the volume element. For volume elements, it represents the ratio of the determinant of the Jacobian matrices from the computational and physical space. To illustrate, a quadratic tetrahedron is shown in Fig. 10. The computational mesh is shown on the left. The physical, curved mesh is shown on the right. All 10 nodes of the elements are shown at the corners and mid-edges.

The normalized Jacobian part of the cost function can be computed in a number of ways. The simplest and least expensive computes the volume ratio of the sub-tetrahedrons, shown in Fig. 11a. Only the nodes in the sub-tetrahedron are involved




Fig. 11 The normalized Jacobian is computed using the volume of sub-tetrahedra or a biased weighting of the Jacobian evaluated at the quadrature points shown for 1 and 2 subdivisions. (a) Linear volume. (b) One subdivision. (c) Two subdivisions

in the calculation. Again, this is consistent with the weighted condition number calculation. The assumption is that the sub-tetrahedron volume will provide enough influence to ensure the actual Jacobians remain positive throughout the entire element.

An alternate approach uses subdivision levels in the sub-tetrahedron to create a grid of quadrature points. The normalized Jacobian is, again, computed using all nodes of the element. Figure 11b,c show the quadrature point locations for 1 and 2 subdivision levels, respectively. The value reported for the subdivision level approach is a biased average of the values at quadrature points within the sub-tetrahedron. The biased weighting is given by the same formula shown earlier (5), where C and F are the normalized Jacobian values assuming the range is from 0 to 1.

The first smoothing pass uses the simplest computational method. If negative Jacobians are detected after that pass is complete, a second smoothing pass is initiated with one subdivision level. Additional smoothing passes are possible but are rarely required.

Marching Direction

The smoothing is a perturbation method and requires a marching direction that will improve the value of the cost function locally. The marching direction for the mesh nodes is computed using the sensitivity of the cost function with respect to the X, Y and Z directions. The sensitivity of the corner nodes of the sub-tetrahedron is determined using C++ operator overloading of the math functions in a dual number framework [18]. This is essentially the numerical chain-rule differentiation of the cost function. Then the mesh nodal values of these derivatives are computed using the biased averaging formula in (5) where C is the cost value and F is the derivative

vector. The cost function routine for a specific sub-tetrahedron will return the cost value and 4 vectors at the corners comprised of 3 doubles for a total of 13 doubles. The biased averaging will produce directions that focus on improving the worst cost value of the surrounding sub-tetrahedra but will blend smoothly with the average cost as the minimum cost improves.

Marching Step Size

For the perturbation method, a meaningful distance is needed in addition to the marching direction already described. To determine this distance, the computed displacement of a given node is computed first using the minimum inscribing radius of the surrounding sub-tetrahedra. The inscribing radius of a sub-tetrahedron is shown in Fig. 12 for one corner of a Q2 tetrahedron. The minimum radius is computed for each moving node using the computational mesh. In collapsed elements the physical inscribing radius approaches zero which would be a poor selection for the step size. Conversely, the computational mesh is always valid and unchanging. Using the computational radius ensures a non-zero step size is used. The minimum radius is then multiplied by a user relaxation parameter, typically 0.05. This is further reduced by the minimum of one and the current difference between the nodal cost function value, which will approach a zero displacement as the node approaches the ideal position.

During mesh smoothing only nodes whose cost value is below a user specified convergence threshold, such as 0.95, are moved. This greatly reduces the overall computational expense, especially as the mesh smoothing converges. Most of the nodes in the mesh become inactive. Only nodes associated with lower cost values remain active towards the end of mesh smoothing.

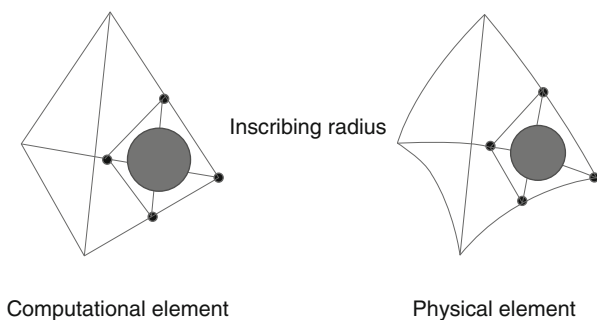


Fig. 12 Step size determined by inscribing radius of sub-tetrahedral elements

3 Results

Two realistic, complex cases are included that demonstrate the ability of the WCN-based mesh curving approach. These cases start with linear meshes generated by Pointwise. The linear mesh, geometry file and MeshLink XML file are exported and used by the curving code to produce mixed order meshes containing linear, quadratic, cubic and quartic elements.

3.1 *Juncture Flow Model*

The Juncture Flow Model (JFM) is a popular case for validating CFD methods on wing root viscous separation. It has been the focus of numerous studies and workshops, including the 3rd American Institute of Aeronautics and Astronautics (AIAA) Geometry and Mesh Generation Workshop (GMGW-3) [19]. Several mesh families were generated for this model for the purposes of the workshop. These include linear, mixed order and uniform order meshes up to quartic. Shown below in Fig. 13 is a representative mesh for the coarsest mesh in the sequence. Cubic elements are shown on the highest curvature regions of the body while some elements on the flat portion of the fuselage remain linear. Quadratic elements transition linear to cubic, enforcing the one order difference constraint between




Fig. 13 Mixed order mesh of the Juncture Flow Model using Linear, Quadratic, and Cubic elements

neighboring elements. The highest curvature in the geometry occurs at the wing tip trailing edge, shown in Fig. 14. The wing tip is rounded and the coarse grid shown here has only 4 elements wrapped around 180° turn near the trailing edge. The wall normal spacing is large enough to allow viewing of the highly curved tetrahedral elements at the surface. Finer meshes in the series have more elements spanning the wing tip, but also have finer wall normal spacing which challenges the curving process.

At the completion of each mesh smoothing phase the shape conformity metric is evaluated for all boundaries except planar boundaries, such as symmetry planes. The error values achieved for the shape conformity metric for the JFM fuselage is listed in Table 1. The left, middle, and right columns are the element order, average error, and maximum error respectively. As expected, as the element order is increased the errors reduce significantly. Notice that the maximum error for the mixed order Q1–Q3 and uniform order Q3 meshes are equal. The same is true for the Q1–Q4




Fig. 14 Cut at the wing tip trailing edge of the Juncture Flow Model showing mixed order Q1–Q4 elements

Table 1 Shape conformity for juncture flow model fuselage

Elevation order	Average error	Maximum error
Q1	0.471557	3.45525
Q1–Q2	0.00454984	0.180504
Q2	0.00454117	0.180398
Q1–Q3	0.000556685	0.0413742
Q3	0.000553584	0.0413742
Q1–Q4	0.000207317	0.0190914
Q4	0.000183944	0.0190914

Table 2 Shape conformity for juncture flow model wing

Elevation order	Average error	Maximum error
Q1	0.0293635	0.785052
Q1-Q2	0.000796288	0.0753876
Q2	0.000794145	0.0757585
Q1-Q3	0.00013674	0.0145939
Q3	0.0000920631	0.0145939
Q1-Q4	0.0000744936	0.0056608
Q4	0.0000252808	0.0056608

and Q4 meshes. This indicates that the maximum error is occurring on an element of maximum order, Q3 and Q4 respectively. The average error difference between the mixed order and the fully elevated order for the same maximum order mesh only vary slightly. This is expected with the user specified deviation metric of 0.01. Smaller values of the deviation metric will reduce this difference further at the cost of elevating more elements in the mesh to higher order.

The shape conformity metric for the wing surface is displayed in Table 2. Similar trends can be seen for the wing surface. The error levels are smaller than those reported for the fuselage due to the finer resolution mesh on the wing.

3.2 NASA High Lift Common Research Model

Another case studied at GMGW-3 was the NASA High Lift Common Research Model (CRM-HL) configuration. This case was the focus of the 4th AIAA High Lift Prediction Workshop (HLPW-4) [20] co-located with GMGW-3. Several mesh families were generated for this configuration also, including linear, mixed order and uniform order up to quartic. The meshes shown below represent the coarsest level meshes from the sequence. The symmetry plane mesh is mostly linear, as shown in Fig. 15a. The center section of the fuselage is quadratic. Cubic and quartic elements are shown for the forward and aft fuselage regions. Most of the underside of the wing, shown in Fig. 15b, is quadratic. Cubic and quartic elements exist at the leading edges of the wing and nacelle. This is also true for the topside view of the slats and nacelle pylon shown in Fig. 15c.

An axial cut at the wing tip trailing edge is shown in Fig. 16a. This very coarse mesh has only two triangle elements spanning the 180° turn of the rounded wing tip. The wall normal spacing, equivalent to an approximate Y^+ value of 100, is extremely coarse. Much finer wall normal spacing was used for other meshes in the series, but those are more difficult to visualize. Also shown in the figure are the nodes of the mesh. Notice the quartic elements contain 5 points (four segments) along each edge. When the adjacent element is cubic these mid-edge nodes are not shared. The adjacent elements have a different set of edge internal points. The




Fig. 15 Mixed order mesh of the Common Research Model using Linear through Quartic elements. (a) Symmetry plane. (b) Underside of wing. (c) Nacelle

enforcement of shape conformity at this interface ensures the curves represented on each side are the same, eliminating gaps in the mesh.

A cut through the volume mesh at the engine nacelle/pylon is shown in Fig. 16b. The high curvature of the nacelle leading edge is resolved with quartic elements. Quadratic elements cover most of the nacelle. The majority of the elements away from the curved geometry remain linear.

4 Conclusions

A method for generating curved, mixed order meshes has been presented. Geometry access is provided through the MeshLink API. A deviation metric is used to indicate when surface and volume elements need elevation. Elements up to 4th order are possible. Iterative perturbation-based smoothing is used to ensure a valid, high quality mesh is produced. The cost function for the smoothing is comprised of a normalized Jacobian component that ensures positive Jacobians and a Weighted




Fig. 16 Cutting planes showing Q1–Q4 elements on the CRM. (a) Axial cut at wing tip trailing edge. (b) Cut at engine nacelle and pylon

Condition Number component that enforces element shape. The combination allows for elevation and smoothing of meshes that include clustering to viscous boundaries. Shape conformity is imposed between elements of different order and used to evaluate the error between the elevated surface mesh and the underlying geometry. Two examples of realistic geometries were presented for configurations studied in the 4th AIAA High Lift Prediction Workshop and 3rd AIAA Geometry and Mesh Generation Workshop.

References

1. Chen, C.H.: A Radial Basis Functions Approach to Generating High-Order Curved Element Meshes for Computational Fluid Dynamics. Master's thesis, Imperial College, London (2013)
2. Stees, M., Shontz, S.M.: Spectral and High Order Methods for Partial Differential Equations, p. 229 (2018)
3. Ims, J., Duan, Z., Wang, Z.J.: In: 22nd AIAA Computational Fluid Dynamics Conference (AIAA 2015-2293)
4. Persson, P.O., Peraire, J.: In: 47th AIAA Aerospace Sciences Meeting Including The New Horizons Forum and Aerospace Exposition (AIAA 2009-0949). <https://doi.org/10.2514/6.2009-949>
5. Moxey, D., Ekelschot, D., Keskin, Ü., Sherwin, S.J., Peiró, J.: Comput. Aided Des. **72**, 130 (2016). <https://doi.org/10.1016/j.cad.2015.09.007>
6. Fortunato, M., Persson, P.O.: J. Comput. Phys. **307**, 1 (2016)
7. Feuillet, R., Loseille, A., Alauzet, F.: In: International Meshing Roundtable, pp. 3–21 Springer, Berlin (2018)
8. Zhang, R., Johnen, A., Remacle, J.F.: In: International Meshing Roundtable, pp. 57–69. Springer, Berlin (2018)
9. Ruiz-Gironés, E., Sarrate, J., Roca, X.: Procedia Eng. **163**, 315 (2016)
10. Karman, S.L., Erwin, J.T., Glasby, R.S., Stefanski, D.: In: 46th AIAA fluid dynamics conference (AIAA 2016-3178), p. 3178. <https://doi.org/10.2514/6.2016-3178>
11. Karman, S.L.: In: International Conference on Spectral and High Order Methods (2018)