Available online at www.sciencedirect.com
SCIENCE DIRECT®
Applied Numerical Mathematics 52 (2005) 251–271
APPLIED NUMERICAL MATHEMATICS
www.elsevier.com/locate/apnum

Adaptive mesh generation for curved domains

Mark S. Shephard [a,*], Joseph E. Flaherty [a], Kenneth E. Jansen [a], Xiangrong Li [a], Xiaojuan Luo [a], Nicolas Chevaugeon [a], Jean-François Remacle [b], Mark W. Beall [c], Robert M. O'Bara [c]

[a] *Scientific Computation Research Center, Rensselaer Polytechnic Institute, Troy, NY 12180, USA*
[b] *U. catholique de Louvain, Louvain-la-Neuve, Belgium*
[c] *Simmetrix, Inc., 10 Halfmoon Executive Park Drive, Clifton Park, NY 12065, USA*

Abstract

This paper considers the technologies needed to support the creation of adaptively constructed meshes for general curved three-dimensional domains and outlines one set of solutions for providing them. A brief review of an effective way to integrate mesh generation/adaptation with CAD geometries is given. A set of procedures that support general *h*-adaptive refinement based on a mesh metric field is given. This is followed by examples that demonstrate the ability of the procedures to adaptively construct anisotropic meshes for flow problems. A procedure for the generation of strongly graded, curved meshes as needed for effective *hp*-adaptive simulations is also given.

Keywords: Adaptive meshes; Anisotropic meshes; Curved meshes

1. Introduction

Adaptive methods to control the discretization errors associated with the application of finite element methods in the numerical solution of partial differential equations have been under continuous development since the pioneering work of Babuska and others began in the 1970s. Over the years these methods have been refined and for several classes of equations have matured to the point that effective adaptive analysis procedures can be developed and delivered. It is obvious to those involved in the development

* Corresponding author.
E-mail address: shephard@scorec.rpi.edu (M.S. Shephard).

0168-9274/$30.00
doi:10.1016/j.apnum.2004.08.040

of these technologies that the applications of adaptive methods can greatly increase the reliability of the simulations performed in the process of engineering design, thus allowing a greatly increased use of simulation to design superior products and lower costs. However, with the exception of a small number of specific examples, these methods are not being supported by the computer aided engineering industry that provide industry with the simulation tools they use.

There are several reasons why adaptive methods are not yet commonly applied in engineering practice. One reason is that most commercial finite element and finite volume codes employ software structures that are not easily extended for the efficient application of adaptive methods. In addition, these codes do not preclude users from "executing variational crimes" such as using point loads or constraints, which are in fact common practice, but will yield meaningless results when an adaptive method is applied on such problem specifications [3]. Another reason is that the procedures for the a posteriori estimation of the discretization errors are based on norms that are not directly related to the quantities of engineering interest. Recent efforts on the development of goal oriented error estimators are beginning to address this need (see, e.g., Ref. [21]). A final reason is a lack of procedures to use the a posteriori error estimates to specify what the adaptive discretization should be and to execute the processes to construct those discretizations. This paper focuses on the last of these issues which is the construction of the adapted discretizations in terms of a properly defined mesh. The procedures reviewed here represent an integrated set of tools for general three-dimensional domains that address general anisotropic mesh adaptation that properly account for the approximation issues associated with the representation of curved geometries.

The ability to support the automatic generation and adaptation of general curved geometries requires an appropriate definition of the geometric domain. Section 2 reviews the geometry functions needed to support these processes when the domains are defined within CAD systems. Section 3 overviews a set of procedures for the construction of adaptively refined anisotropic meshes based on the application of mesh modifications operations while Section 4 demonstrates the application of these procedures to two applications, one employing discontinuous Galerkin discretizations and one employing stabilized finite element discretizations. Section 5 then considers procedures for the generation of meshes appropriate for hp-adaptive discretizations of elliptic equations over general domains including re-entrant geometry where singularities must be resolved.

2. Geometry information from CAD systems and relationship to meshing

A major bottleneck to the effective application of adaptive discretization control is the lack of reliable means to automatically generate and adapt meshes directly from the domain definition information. Increasingly the complete definition of the domain of interest is defined by the solid modeling procedures within commercial computer-aided design systems. A majority of efforts to integrate mesh generation and adaptation procedures with solid models have employed standard file exchange methods. This approach has been found to have a low reliability for the simple reason that these file structures do not maintain information on the geometric tolerances and tolerancing methods used by the solid modeler. This information is required to support the consistent determination of how the geometric entities defining the solid model interact. Access to the tolerance information and methods can be obtained by directly integrating the meshing procedures with the modeler. This approach is supported by the majority of the solid modeling systems which provide a library of functions supporting a broad range of geometric in-

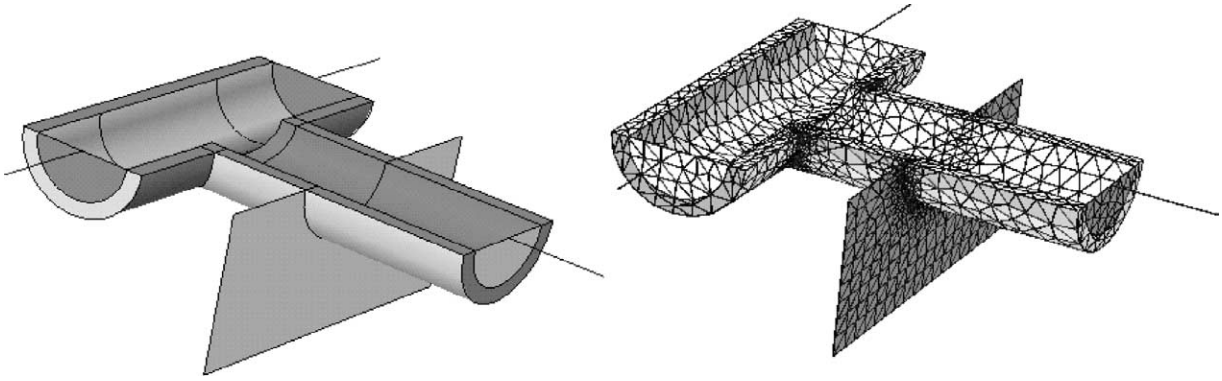

Fig. 1. Example of a non-manifold model (left) and a mesh of that model (right).

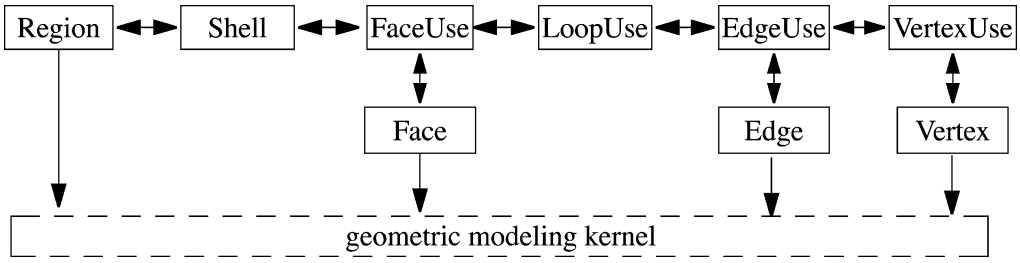

Fig. 2. Model topological adjacency information and relation to model geometry.

terrogations. These interrogations are keyed via the topological entities of a boundary representation of the solid model.

An effective approach to take advantage of the capabilities afforded by solid modelers is to employ the topological entities of the model and their adjacencies to provide the mesh generation and modification operations with the geometric information they need. Although each of the solid modelers maintain a topological representation, there are variations in the implementations that could complicate the integration with the meshing processes. Since they do provide sufficient information and functions to construct any selected topological entity and adjacency structure, a viable approach is to load and maintain an appropriate copy of the topological model within the mesh generation procedures. Since the solid models needed for simulation processes are often general combinations of solids, surfaces, curves and points (Fig. 1), a complete non-manifold representation in the form of the radial-edge data structure [28] (Fig. 2) is used. With this structure it is straightforward to link to the operator libraries provided by the solid modeler. This approach (see Refs. [7,25,26]) has been found to be highly reliable to mesh very complex domains such as automobiles with anisotropic meshes as needed for thermal-flow simulations (Fig. 3).

The topological model entities of the geometric model boundary representation also provide a convenient means to describe the additional information, commonly referred to as attribute information, needed to define a simulation [20]. The attribute information includes material parameters, boundary conditions, loadings and initial conditions. When coupled with the domain definition and knowledge of the PDEs to be solved this provides a complete problem specification that can be used to drive an automated adaptive analysis procedure.

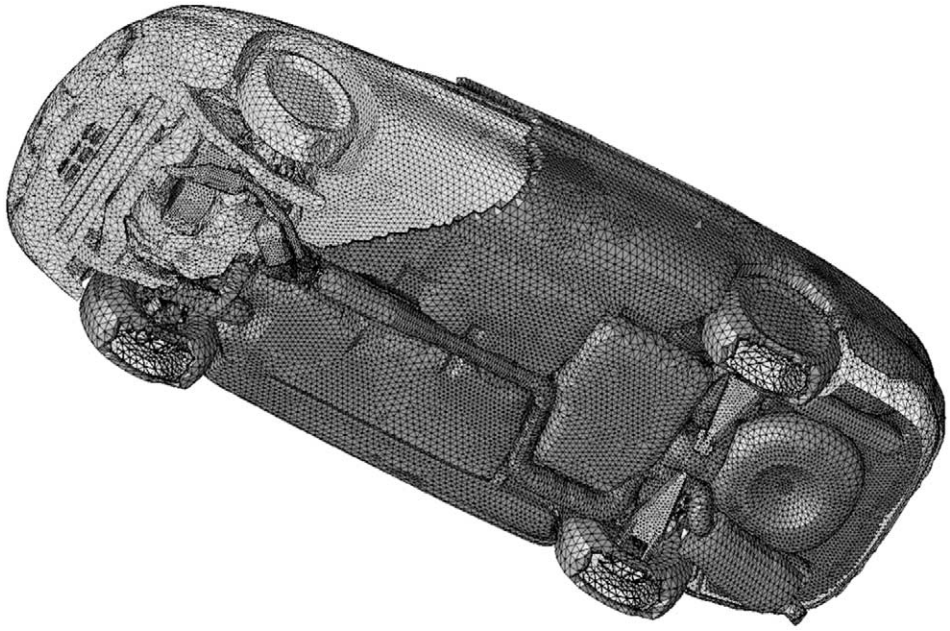

Fig. 3. Mesh of a complete automobile. (The light gray shows the surface mesh on portion of the automobile while the darker gray shows the mesh on the outer limits of the boundary layer mesh.)

Maintaining the relationship between the mesh and geometric model is critical to supporting mesh adaptation operations. For example, when the edge of a linear element which is on a curved surface of the model is split, the new vertex needs to be placed at an appropriate point on that surface. Maintaining these relationships is easily supported if the mesh is also defined in terms of a set of topological entities and their adjacencies. Under the assumption that each topological mesh entity of dimension d, M_i^d, is bounded by a set of topological mesh entities of dimension $d-1$, $\{M_i^d\{M^{d-1}\}\}$, the full set of mesh topological entities are:

$$T_M = \{\{M\{M^0\}\}, \{M\{M^1\}\}, \{M\{M^2\}\}, \{M\{M^3\}\}\}, \tag{1}$$

where $\{M\{M^d\}\}$, $d = 0, 1, 2, 3$ are respectively the set of vertices, edges, faces and regions which define the primary topological elements of the mesh domain. With both the mesh and model defined in terms of topological entities it is straightforward to maintain the association of the mesh entities to the model entities [8,23,25,26]. This association is referred to as classification in which the mesh topological entities are classified with respect to the geometric model topological entities upon which they lie.

Definition (*Classification*). The unique association of mesh topological entities of dimension d_i, $M_i^{d_i}$ to the topological entity of the geometric model of dimension d_j, $G_j^{d_j}$ where $d_i \leqslant d_j$, on which it lies is termed classification and is denoted $M_i^{d_i} \sqsubset G_j^{d_j}$ where the classification symbol, $\sqsubset$, indicates that the left hand entity, or set, is classified on the right hand entity.

Definition (*Reverse classification*). For each model entity, G_j^d, the set of equal order mesh entities classified on that model entity defines the reverse classification information for that model entity. Reverse classification is denoted as

$$RC\left(G_j^d\right) = \{M_i^d \mid M_i^d \sqsubset G_j^d\}. \tag{2}$$

Functions to provide the reverse classification are useful in various steps within an adaptive simulation process.

Mesh shape information can be effectively associated with the topological entities defining the mesh. In many cases this is limited to the coordinates of the mesh vertices and, if they exist, higher order nodes associated with mesh edges, faces or regions. In addition, it is possible to associate other forms of geometric information with the mesh entities. For example, associating Bezier curves and surface definitions with mesh edges and faces in p-version finite elements can effectively address the appropriate geometric approximation of those elements [11]. The mesh classification can be used to obtain other needed geometric information such as the coordinates of a new mesh vertex caused by splitting a mesh edge classified on a model face or to support the calculation of the geometric Jacobian information when doing an element stiffness integration.

3. Anisotropic mesh adaptation

The anisotropic mesh adaptation procedure employs a set of mesh modification procedures [17] that alter the given mesh to satisfy the anisotropic mesh sizes given by the adaptively defined mesh metric field.

3.1. Mesh metric field and use in adaptive mesh modification

The goal of the mesh metric field is to provide a spatially-based description of the anisotropic mesh sizes and orientations to be applied at that point in time. A convenient means to specify this information is in terms of a 3×3 mesh metric tensor [9,13]. A convenient definition of the anisotropic mesh metric field is one that defines the mapping of an ellipsoid into a unit sphere in terms of a diagonal distortion matrix, where the diagonal terms correspond to the lengths of the principal axes of the ellipsoid, times a rotation matrix that accounts for the orientation of the ellipsoid. When used for constructing the anisotropic mesh size field, lengths of the principal axes are interpreted as the desired mesh edge lengths in the principal directions at that location defined as

$$Q(x, y, z) = \begin{bmatrix} \vec{e}_1 \\ \vec{e}_2 \\ \vec{e}_3 \end{bmatrix} \begin{bmatrix} 1/h_1 & 0 & 0 \\ 0 & 1/h_2 & 0 \\ 0 & 0 & 1/h_3 \end{bmatrix}, \tag{3}$$

where h_1, h_2 and h_3 are lengths of the three axes of an ellipsoid, and $\vec{e}_1$, $\vec{e}_2$ and $\vec{e}_3$ are the orthogonal unit row vectors associated with the principal axes.

In an adaptive analysis process directionally sensitive error indicators are used to construct the mesh metric over the domain, typically in some piecewise manner such as the specification of nodal metrics on the current mesh. A commonly applied method to construct the mesh metric field is based on the Hessian matrix of constructed second derivatives of the solution field [15]. Such an approach is appropriate when

piecewise linear finite element approximations are used since from a simple approximation theory perspective, the discretization error is related to the second derivatives of the solution. A key area for the effective use of anisotropic adaptive refinement needing further development is the correction indication procedures that can construct appropriate mesh metric tensors that will optimally control the errors in the desired norms for the given finite element discretization method being used.

Given a mesh metric field defined over the domain of interest, the mesh modification procedures examines the mesh entities to determine if they adequately represent the mesh metric field. If an element does not adequately satisfy the given mesh metric, the mesh metric information is used to drive the application of mesh modification operations aimed at producing mesh entities that do acceptably satisfy the mesh metric field. The execution of this process focuses most of its attention on the mesh edges with a specific step included to consider the volume of mesh regions so as to avoid the creation of elements that meet the edge criteria, but have a volume far below that appropriate for the given mesh metric. This is necessitated by the well-known fact that it is possible to construct a zero volume mesh region with mesh edge lengths that do not dramatically vary from the ideal edge length.

The degree of the satisfaction of a mesh edge to the given mesh size field is measured in the transformed space. Considering a mesh edge that runs from vertex A to vertex B, the length of this edge in the transformed field is [13]

$$L_{AB} = \int_A^B \sqrt{\vec{e} Q Q^{\mathrm{T}} \vec{e}^{\,\mathrm{T}}}\, \mathrm{d}x, \tag{4}$$

where $\vec{e}$ is a unit row vector along the edge in physical space.

Since it is not possible to ensure all mesh edges are the correct length and still have the mesh remain compatible, it is necessary to accept edge lengths within an acceptable range. It is also necessary that the range of acceptable edge lengths are selected to be large enough that the mesh modification operations do not enter into an infinite loop of refining and coarsening edges. Defining the interval of acceptable edge length to be $[L_{\mathrm{low}}, L_{\mathrm{up}}]$ the values must be selected such that $L_{\mathrm{low}} \leqslant 0.5 L_{\mathrm{up}}$ and $L_{\mathrm{low}} \leqslant 1.0 \leqslant L_{\mathrm{up}}$. A mesh edge is considered "short" if its transformed length is less than the lower bound, L_{low}, of the interval, and a mesh edge is considered "long" if its transformed length is greater than the upper bound, L_{up}.

Sliver tetrahedra (poorly-shaped tetrahedra not bounded by any short mesh edge in transformed space) may exist even if the edge length criteria is met, so consideration is needed to determine and eliminate sliver tetrahedra. One of the standard non-dimensional shape measures, the cubic of mean ratio [18], is used in the transformed space for this purpose. Let Q be the associated transformation matrix of the tetrahedron,[1] the cubic of mean ratio in transformed space, η is

$$\eta = \frac{15552(|Q|V)^2}{\left(\sum_{i=1}^{6} (\vec{l}_i Q Q^{\mathrm{T}} \vec{l}_i^{\,\mathrm{T}})\right)^3}, \tag{5}$$

where $|Q|V$ is the volume of the tetrahedron in the transformed space ($|Q|$ represents the determinant of the transformation and V is the volume of a tetrahedron in physical space), and $\vec{l}_i$ ($i = 1, 2, \ldots, 6$) are the row direction vectors associated with the six edges of the tetrahedron in physical space. η has

[1] When the transformation is not constant over the tetrahedra, the one with the maximum length major axis is used.

been normalized to interval [0, 1] with 0 for flat tetrahedron and 1 for equilateral tetrahedron in transformed space. When the shape measure of a tetrahedron is below a specified limit mesh modifications are performed to eliminate it.

3.2. Mesh modification procedure

Given a geometric domain, a current mesh and a mesh metric field defined in a piecewise manner over that mesh, a series of controlled mesh modification steps are executed to convert the given mesh into one that satisfies the given mesh metric field [17]. These mesh modifications are executed such that the resulting mesh properly approximates the geometric model by ensuring the mesh entities classified on the model boundary are properly positioned on the boundary [16].

To ensure the adapted mesh will properly satisfy the required metric field, the mesh modification algorithms are carefully constructed and executed in three stages of (i) mesh coarsening to eliminate short edges, (ii) shape correction to provide better structured mesh for the third stage, and (iii) intelligent refinement that includes operations for ensuring the proper geometric approximation of the mesh and proper shape to match the mesh metric field.

The first stage is focused on the elimination of edges in the mesh that are shorter than L_{low} through coarsening operations. The local mesh operations used for this process include edge collapse, compound operators that combine swaps and collapses, and vertex relocation. The algorithm operates by identifying all the short edges to be eliminated and then eliminates them, one at a time, trying the possible modifications in an efficient order. In some cases coarsening operations are not allowed because this would yield an invalid mesh (e.g., collapsing an edge with vertices classified on two different model faces). An important consideration in this process where a substantial number of edges in a region are to be coarsened is being sure those collapsed are well distributed. This is accomplished by being sure to not follow the collapse of one edge by collapsing an immediate neighbor. In a small number of cases, none of the available operators are successful in eliminating a given short edge, in which case the edge is maintained. This occurs only a small percentage of the time and can be accepted since the consequence is the mesh is locally finer than it needs to be.

Since the mesh coarsening step is focused on the creation of elements with edge lengths that satisfy the mesh metric it is possible for poorly shaped elements in the transformed shape (as measured by Eq. (5)) to be created. Since these elements have satisfactory edge lengths, they must be sliver elements of one of the two classic configurations. The determination of the mesh modification operations most likely to be successful is a strong function of the specific configuration of the sliver element. By projecting one vertex of a sliver tetrahedra on the plane defined by the other three (the choice of which vertex is projected is arbitrary) it is possible to identify the mesh entity most appropriate for elimination as well as the most appropriate operations to eliminate that entity, thus leading to the elimination of the sliver element. Since this shape correction process is applied again at the very end of the mesh adaptation process, it is not strictly necessary to apply it at this time. However, not applying it at this time will give the refinement process a small number of poorly shaped elements that will be refined into elements of equally bad or worse shape which will have to be dealt with later. Therefore, it is more efficient to apply shape correction at this time.

The third stage focuses on the refinement of mesh edges longer that L_{up}. Given the set of mesh edges that are too long, the first step of this stage is the application of refinement of the edges and mesh entities bounded by those edges using a template procedure that does include selection of optimal diagonals (in

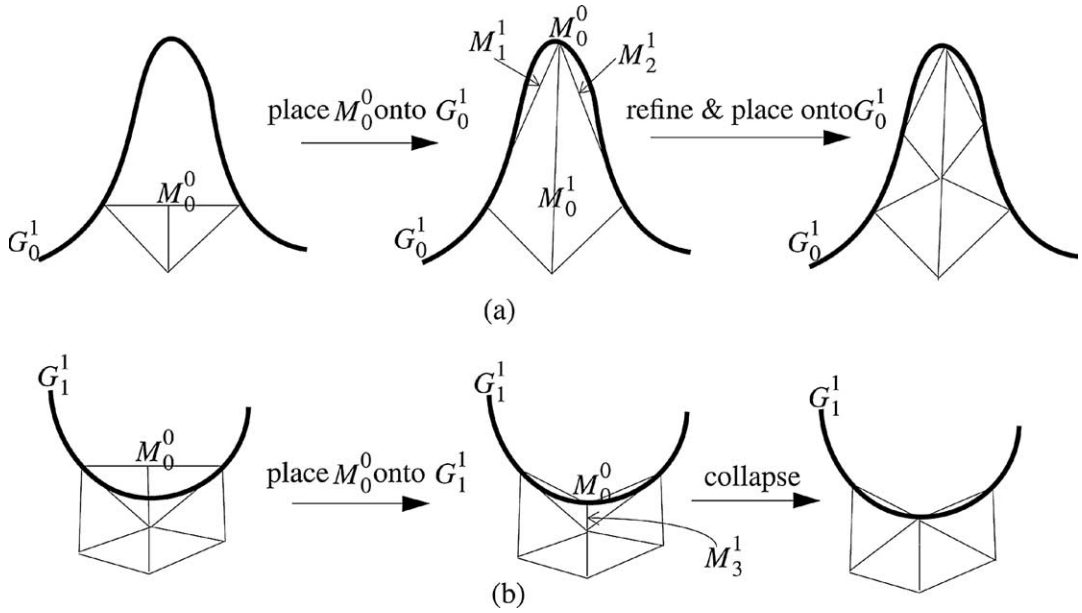

Fig. 4. Dealing with long and short edges introduced when refinement vertices are moved to curved boundaries.

the transformed space) when there are options with respect to diagonal selection [10]. Unlike coarsening, it is critical that each required mesh refinement be executed. Since the edge, face and region splitting operations associated with the application of these refinements is always possible, all refinements are executed.

Since the refinement process can introduce new mesh vertices that are classified on curved boundary entities, it is necessary to move those vertices onto an appropriate location on that boundary. (In the case of curved mesh entities, this process must also consider the split mesh edges and faces classified on curved model boundaries.) In general this process can cause connected elements to become invalid. In those cases a specific process that includes mesh modification and possibly local cavity triangulation is applied [16]. Since the process of placing the vertices on the appropriate curved boundaries will change the length of mesh edges, it is important to check the mesh entities connected to any moved vertices, or edges involved in a local mesh modification operation needed to place the vertex on the boundary, for satisfaction of the mesh metric. Fig. 4 shows examples where moving a refinement vertex to the boundary can make edges that are too long or too short. When mesh edges become too long, additional refinement is applied, and when they become too short appropriate collapses are applied (Fig. 4). Although one may consider that the level of geometry approximation of the meshes in Fig. 4 to be extreme and impractical, it should be noted that examples of this level of approximation in real 3D meshes are in fact common and must be dealt with. One alternative is to use curvature based mesh refinement during the initial meshing process to ensure adequate geometric approximation. However, if these geometric features happen to be located in unimportant portions of the domain, with respect to the current simulation goals, then ensuring that level of geometric approximation control is wasteful.

The process of applying the mesh refinement templates can introduce new edges that are locally shorter than required by the mesh metric field thus producing meshes that are finer than needed to the point in some cases of providing an inefficient representation of the requested mesh metric field. Therefore, as the

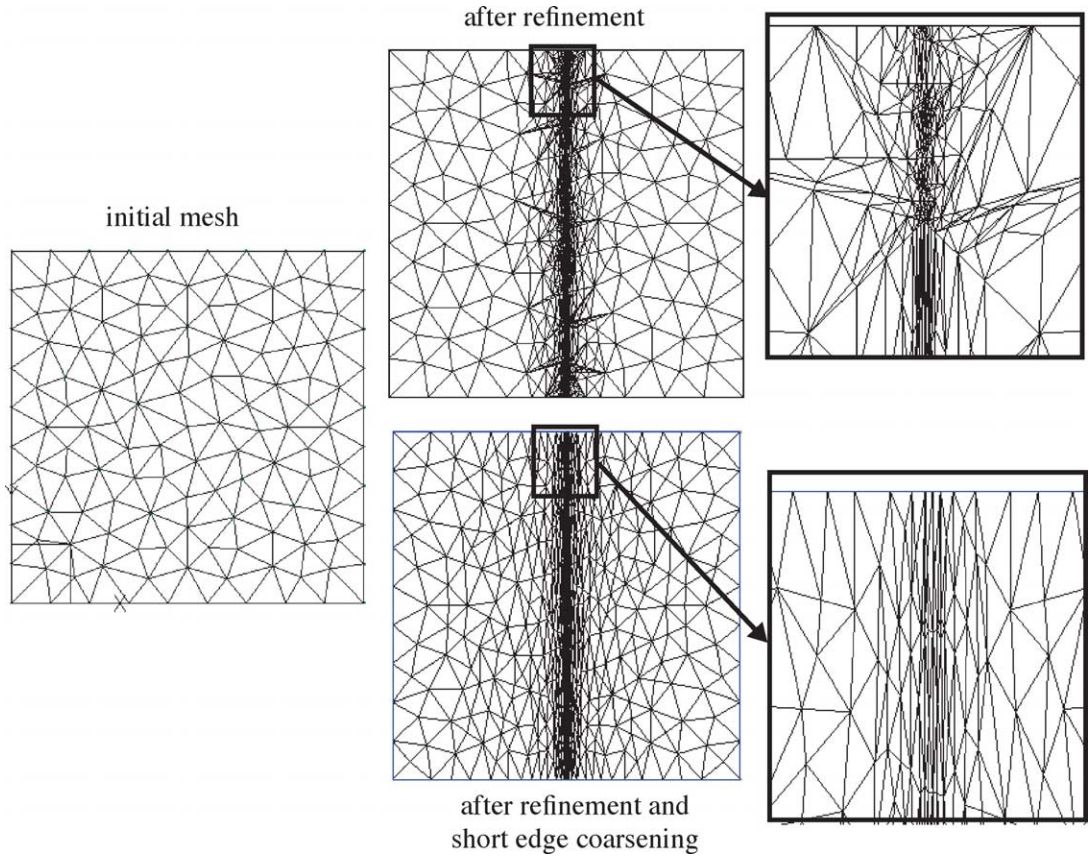

Fig. 5. Importance of collapsing short edges introduced during refinement.

refinement templates are being applied a list of these short edges is created. The edges in that list are then processed by the coarsening procedure to improve the mesh distribution. The importance of this process is demonstrated in Fig. 5 where an initially uniform mesh, shown on the left of Fig. 5, was refined to anisotropically capture a strong feature along a vertical line in the middle of the domain producing the mesh shown on the top right of Fig. 5. By applying the coarsening procedure to the disproportionately short edges created during refinement, the mesh shown on the bottom right of Fig. 5 was produced. This mesh provides a much more effective representation of the requested mesh metric field.

Since the above steps focus attention on mesh edges, they can again introduce a limited number of sliver elements. Therefore, the last step in this stage it to again apply the shape correction process to eliminate any sliver elements that have been created.

4. Applications of anisotropic mesh adaptation

4.1. Inviscid flow problems with shocks discretized by discontinuous Galerkin method

Consider the problem of a 2D approximation to the cannon blast in a tube with a 155 mm diameter that has a set of perforation holes of 28.6 mm near the exit of the tube. The initial conditions for the

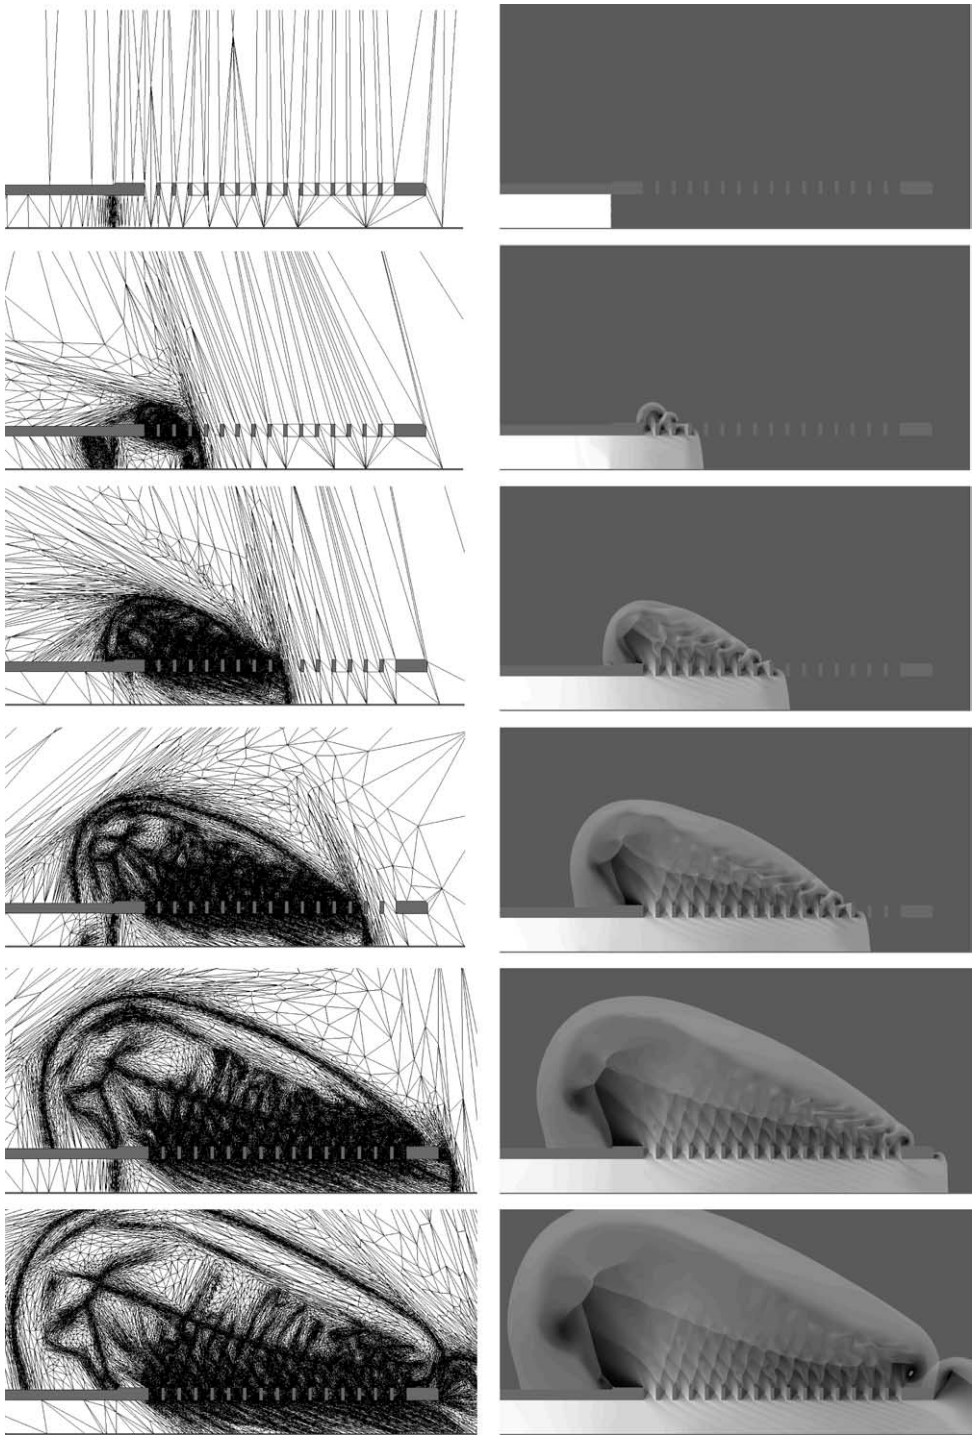

Fig. 6. Evolution of adaptive mesh (left) and density contours (right) for 2D cannon blast problem.The rows of images correspond to times of $t = 0.0$, 0.0001, 0.0002, 0.0003, 0.0004 and 0.0005 seconds.

problem are the one of a shock tube with a shock placed to the left of the perforations. The pressure on the upstream side of the shock is 565 times the downstream and freestream values. The initial temperature of the air inside the tube is 2111.5 degree K and the initial velocity is zero. The initial location of the virtual interface is indicated by the jump density contours for the initial condition shown in the top right image of Fig. 6.

A discontinuous Galerkin (DG) finite element approximation using linear basis functions was used to discretize the spatial component of the Euler equations being used to model this problem [24] and a second order Runge–Kutta time integration scheme was used in the simulation which is run for 5.0×10^{-4} seconds. The time steps applied during the simulation were adaptively set to maintain a CFL limit of 1.0. Initial time steps were 5.0×10^{-8} seconds and the final time steps were about 1.5×10^{-8} seconds. The mesh was adapted every 10^{-6} seconds giving a total of 501 mesh adaptation steps.

The mesh metric field for this simulation was constructed using a process that explicitly accounts for the discontinuities within the solution field. The procedure for adaptively computing the mesh metric field employs the following steps [24]:

- Apply a discontinuity isolation procedure to isolate the discontinuities in the solution field. The procedure takes advantage of the superconvergence properties of the DG method [1] and monitors estimates of the convergence rate over the domain. Element interfaces where the estimated rate of convergence is below the expected rate are marked as locations of a discontinuity.
- Construct the mesh metric field in the vicinity of the discontinuities using local solution gradient information on each side of the discontinuity to estimate the desired mesh sizes normal and tangential to the discontinuity.
- Calculate the mesh metric field in regions away from the discontinuities using the Hessian matrix of second derivatives of the solution. Since a DG discretization is being used, specific care is required to construct second derivatives. A variationally based procedure is used for this process.
- Merge and smooth the full mesh metric field over the computational domain.

Fig. 6 shows the mesh and density contours for every 10^{-4} seconds. The initial mesh, which was anisotropically refined based on a given mesh metric field to capture the initial interface discontinuity, has 5,556 degrees of freedom. After the 501 adaptive mesh modification steps there are 778,488 degrees of freedom in the mesh which captures the complex shock interactions caused by the perforations. Fig. 7 shows a zoom near the front shock which must be accurately predicted and tracked as it moves through the computational domain. Note the alignment of the anisotropic elements and the front shock.

4.2. Flow problem discretized by stabilized finite element method

The second example considers the flow within a pipe with a symmetric bifurcation (see Fig. 8). The velocity profile on inflow boundary on the left side is specified as

$$u_3 = \min\bigl(25(1-r), (1-r)^{1/7}\bigr) \tag{6}$$

where u_3 is the velocity in z direction, and $r = \sqrt{x^2 + y^2},\ r \leqslant 1$. A zero pressure boundary condition is set on the two outflow boundaries, and no slip wall boundary conditions are enforced.

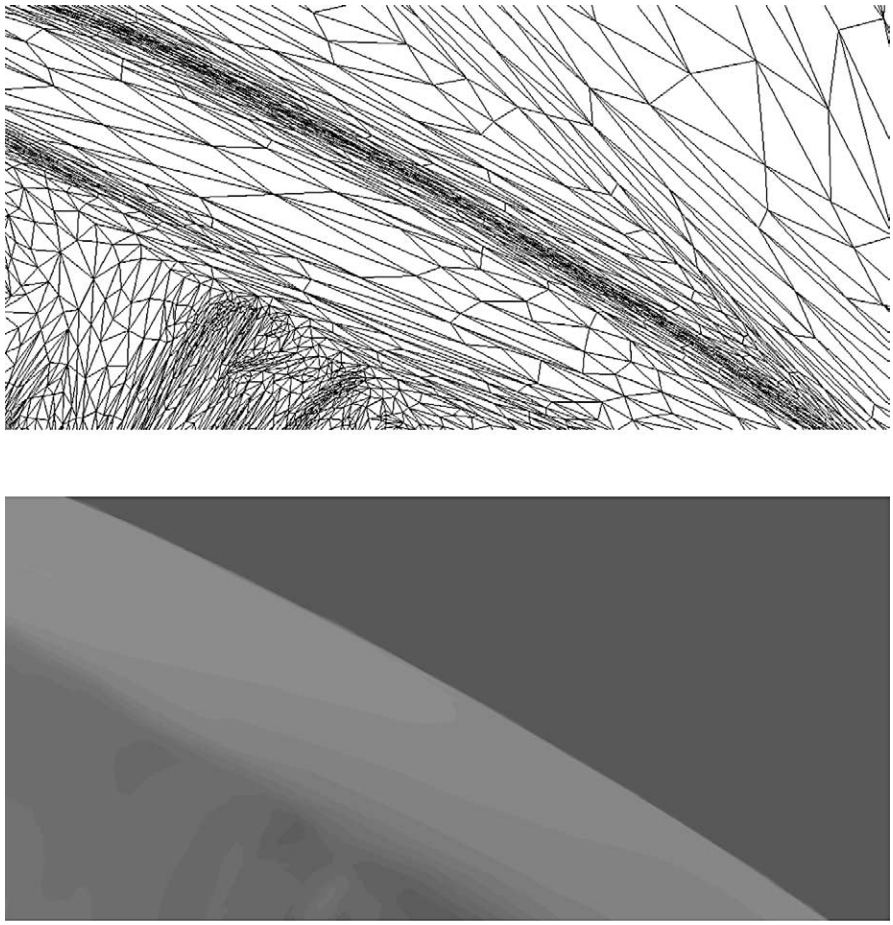

Fig. 7. Close-up of the mesh and density contours of the front shock at $t = 0.0005$ seconds.

For this problem the flow is modeled using time dependent, incompressible Navier–Stokes equations and discretized using a stabilized finite element [29]. The adapted mesh size field is constructed using a scaled Hessian matrix of the second derivatives of the flow field.

Fig. 9 gives the side view (looking down the y-axis as in Fig. 8) of initial mesh and the anisotropically refined mesh after the fourth application of mesh adaptations. The initial mesh is uniform and isotropic. It is used to solve the flow problem in solution steps from zero to 50 (0.5 seconds for each solution step). The refined mesh is achieved after the application of four anisotropic mesh adaptations at solution step 50, 80, 110, 140 respectively, and it is used to solve the flow problem from step 140 to step 170. The number of elements increased from 38,903 to 270,753. Fig. 10 shows the interior mesh faces interacting with the plane of $y = 0$ and flow speed contours on that plane, including close-up views to the area near the bifurcation point. It can be seen that mesh adaptation has produced small isotropic elements at the bifurcation point which then become anisotropic in the boundary layers downstream of the bifurcation point. Fig. 11 shows the refined mesh and its associated flow speed contours on two cross sections of the model (see Fig. 8 for definition of the two sections). The top figures show the surface mesh and the flow

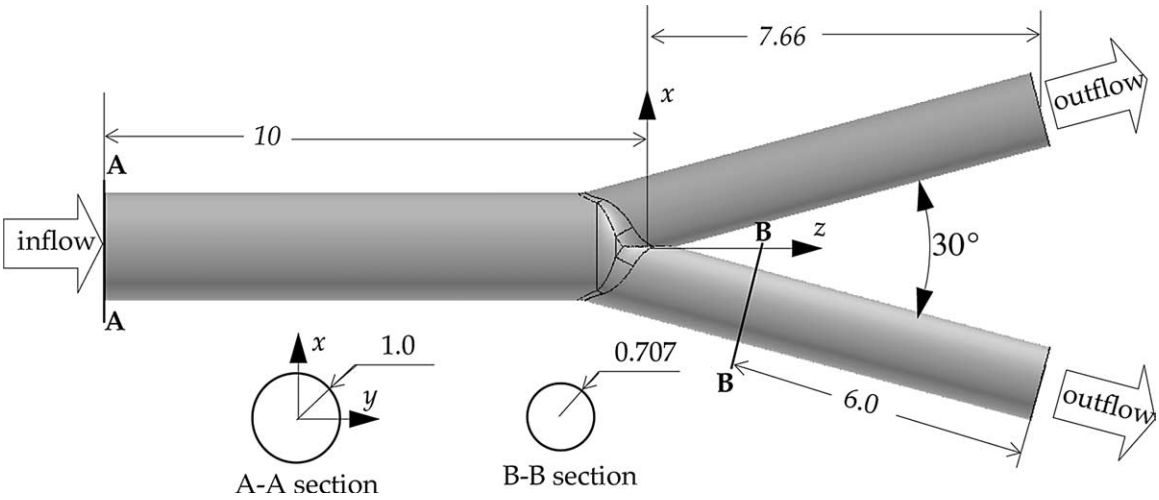

Fig. 8. Geometry of pipe bifurcation.

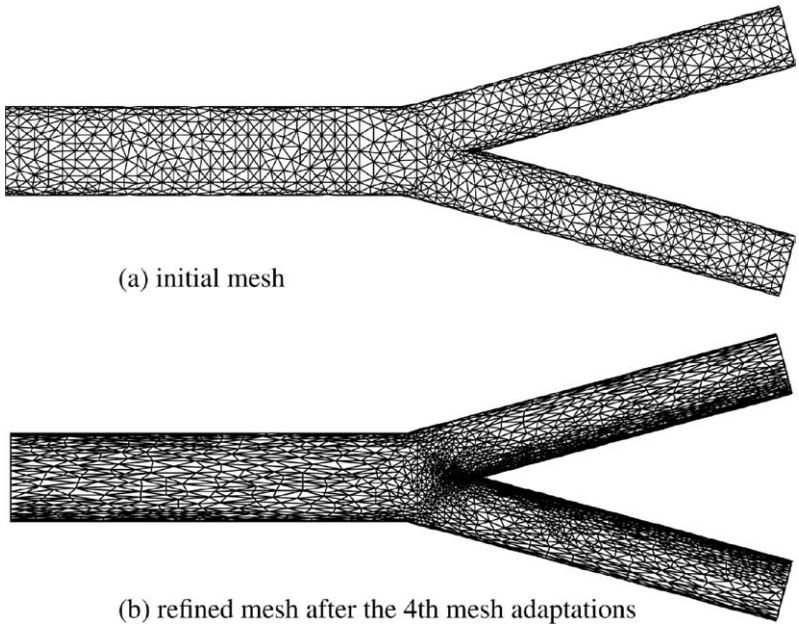

Fig. 9. Side view of the initial and adaptively refined mesh.

speed contour at inlet while the lower meshes show the interior mesh faces and speed contour related to the section B–B within one of the branches (see Fig. 8).

5. Mesh generation for *hp*-adaptive methods

Theoretical results that show the possibility of exponential rates of convergence of a finite element solution with the use of high polynomial order elements on properly refined meshes, the so-called *hp*-

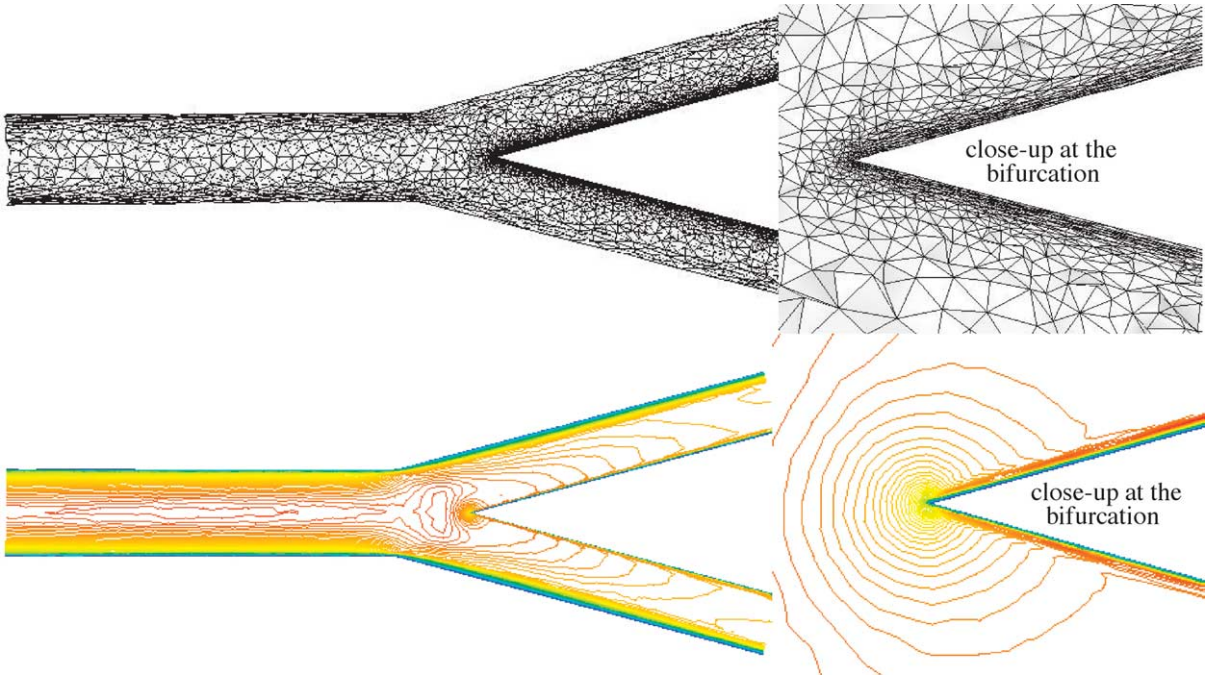

Fig. 10. The adapted mesh and flow speed contours on a plane through the center of the pipe. (The mesh images show the mesh faces the intersection plane pass through.)

finite element method, are well known and documented [4,27]. It has also been shown that through the proper construction of mesh configurations that it is possible to attain these high rates of convergence in the application to problems of interest [2,14,22]. This section overviews progress being made on the development of a procedure capable of generating meshes for general curved three-dimensional domains that will meet the requirements of hp-finite element methods.

5.1. Mesh requirements

The meshes required to attain the high rates of convergence and levels of accuracy possible with hp-version finite elements must meet severe requirements on both the gradation of the elements in the mesh and the shape of the individual elements. The discussion given here is focused on the requirements of solving elliptic equations on curved three-dimensional domains with piecewise smooth boundaries, loadings and boundary conditions.

A basic requirement of any mesh enrichment method is that as the finite element basis is improved, the approximation of the mesh to curved geometric domains is represented properly. In the case when a mesh of linear basis elements is refined, this requirement corresponds to placing all new refinements vertices classified on curved boundaries on the correct boundary. In the case when higher order finite element basis are used, the geometric approximation of all the mesh entities classified on the boundary that are modified must be properly improved to the correct order. In the cases of isoparametric elements defined in terms of standard interpolating Lagrange polynomials this requirement is met by being sure that all nodes at mesh vertices, on mesh edges and on mesh faces classified on curved boundaries are placed

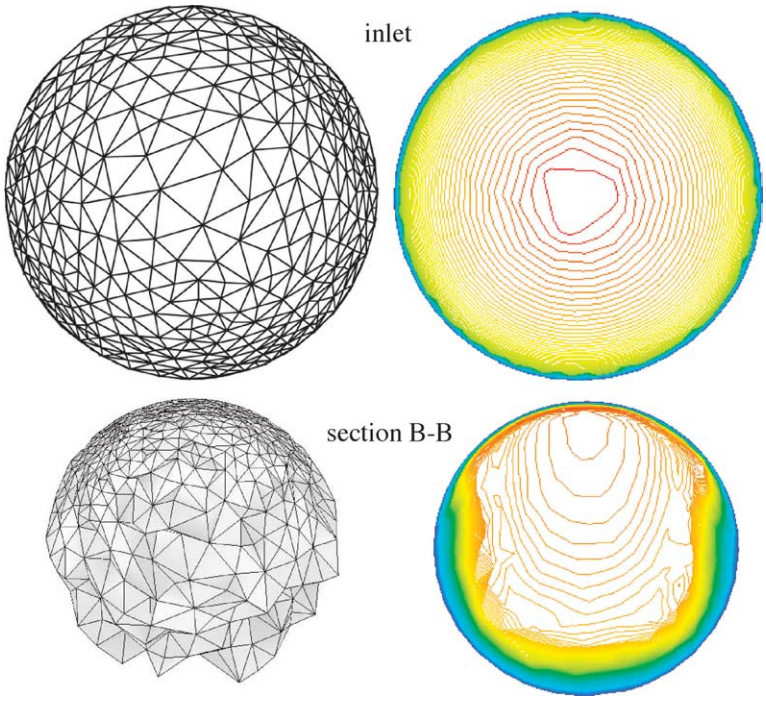

Fig. 11. Mesh and flow speed contours at the inlet and section B–B on one of the bifurcation branches. (The lower mesh image the mesh faces the intersection plane pass through.)

on the appropriate boundary. Since it is common to use different basis functions for the finite element approximation and element geometries in hp-finite element methods, one has to more carefully consider how to satisfy the geometric approximation requirements for these elements by the proper improvement of the mesh edge and face shapes.

There is limited theoretical information available on the level of geometric approximation required to maintain convergence to the correct solution for the given curved domain. A simple analysis based on the relation of approximation theory to the convergence of the error in the energy norm indicates that the energy norm will converge so long as the geometric approximation of the mesh is within one order of that used in the finite element basis. Information obtained from various numerical studies does not provide a clear picture of exactly how well the geometry must be approximated over a variety of mesh configurations to provide proper results for norms of interest. However, a simple study of coarse meshes on test problems does clearly demonstrate the loss of convergence in the energy norm and pointwise norms of engineering interest when the geometric approximation is not increased as the order of the finite element basis is increased [19]. Since, in areas away from singularities, it has been shown that the most effective means to improve the solution results is to increase the polynomial order of the elements [22, 27], one does want very large elements over those portions of the domain. Although additional theoretical results are desired, numerical results that have examined both energy and local pointwise norms indicate that properly maintaining the geometric approximation to the same order as that of the finite element basis leads to convergence of both.

In portions of the domain where the exact solution is smooth the most effective mesh is one that is as coarse as possible. However, the vicinity of singularities in the mesh must be refined. In particular, the optimal mesh requires a geometrical gradation in the direction normal to the singularity, with a grading factor that is a function of the strength of singularities [5,6,22]. Since in the case of three-dimensional domains the singularities can be along curves or at points, these geometrically graded meshes will be required in a "cylindrical shape" around the singular edges and "spherically" around the singular vertices.

5.2. Generation of meshes for hp-version finite elements

Since the requirements on the meshes for *hp*-version meshes are more complex than *h*-version mesh generation procedures, the creation of the appropriate meshes needs to influence the initial mesh generation process. In the case where the solution over the entire domain is smooth, the appropriate mesh is the coarsest one possible. Since most mesh generation procedures are oriented toward the generation of meshes of lower order elements (linear and, maybe, quadratic) they typically create meshes that are finer than desired. This is because they require a completed valid mesh of piecewise linear geometry elements and typically employ algorithms where there are usually multiple elements across any curved geometric model entity. In the case where the domain includes edge and vertex singularities, these procedures are not designed to create the appropriate mesh layouts near the singularity. The construction of the needed mesh layouts near the singularity would also be difficult to accomplish by the application of mesh modification operations.

The approach being developed for the generation of initial meshes appropriate for *hp*-analysis considers the requirements of coarse curved mesh entity creation during the meshing process and includes procedures that construct the geometrically graded meshes from singular edges and vertices. Conceptually, the ideal approach to accomplish this mesh generation process is to "carve-out" curved elements of the size and shape desired one at a time. Although such an approach would provide the most flexibility in the construction of the meshes, the lack of algorithms to support its operation and high level of computational effort on a per-element basis that would be required are both currently prohibitive. Therefore, a compromise approach that begins to account for the existence of curved mesh entities and the creation of appropriate mesh gradations as early in the process as practical is under development.

The steps in the automatic mesh generation procedure for *hp*-meshes currently under development are:

(1) Isolate all of the edges and vertices in the model that will have singularities.
(2) Generate a coarse linear mesh on the boundary of the model accounting for the isolated features.
(3) Generate coarse linear mesh with appropriate geometric gradation towards the isolated singular features.
(4) Generate a coarse linear mesh to fill the remainder of the domain.
(5) Curve the singular feature isolation mesh to ensure a proper curved mesh isolation of the singular feature.
(6) Curve the remaining mesh entities (edges and faces) classified on the curved boundaries to the currently required order of approximation. Apply mesh modifications as required to the surface mesh connected entities.
(7) Apply mesh modification, included mesh curving and curved mesh splits, collapses and swaps, as needed to interior mesh entities to ensure a valid mesh of acceptably shaped elements.

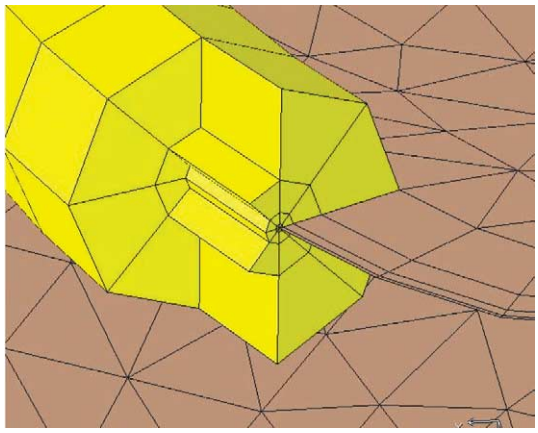

Fig. 12. Example of a geometrically graded mesh isolating a singular edge.

Given a geometry-based problem specification in terms of the geometric model and analysis attributes of loads, material properties and boundary conditions, it is possible to preprocess the geometric model to mark all the model edges and vertices at which the solution will be singular. For example, the geometric singularities can be detected by the execution of the appropriate geometric interrogations of surface normals which is supported by the solid modeling system the domain is defined within.

Given the isolated singular edges, a linear geometry surface mesh is generated using a general surface meshing procedure in conjunction with a procedure to construct a geometrically graded mesh on faces bounded by any singular edges and vertices.

Geometrically graded volume elements are then generated around the isolated edge and vertex singularities taking appropriate account of the local surface mesh. The procedures that create the geometrically graded surface and volume mesh around the isolated singularities (see Fig. 12) employs the functionalities of a generalized boundary layer mesh generation procedure [12].

The remainder of the interior is meshed with a coarse linear mesh.

Given this mesh the process of curving the appropriate mesh entities is executed. To maximize the quality of the mesh for hp-analyses, the curving process is carried out working from the most critical portions of the mesh to the less critical portions. This is done so that in those cases when mesh modifications other than curving entities are required, they are applied in the least critical areas where fine control over the local mesh configuration is not critical. The mesh edges and faces used to isolate the singular features are curved first. This process directly uses knowledge of the layers of geometrically graded elements and begins by curving the mesh edges classified on the curved model edges and then curving the mesh edges and faces in the layers isolating the edge in a manner to maintain the gradation of the mesh over those mesh entities. The result of this process is shown in Fig. 13 where a re-entrant curved edge has been isolated. The mesh on the left shows the local surface mesh before curving and the mesh on the right shows it after curving. The interior mesh entities isolating the edge are also curved in a similar manner.

The last two steps involve curving the surface mesh entities as needed to properly approximate the curved surfaces and properly interact with the curved surface mesh entities isolating singular edges and vertices that have also been properly curved to the boundary. Mesh modification processes of curved entity split, collapse and swap may be required during this process since curving can cause connected mesh

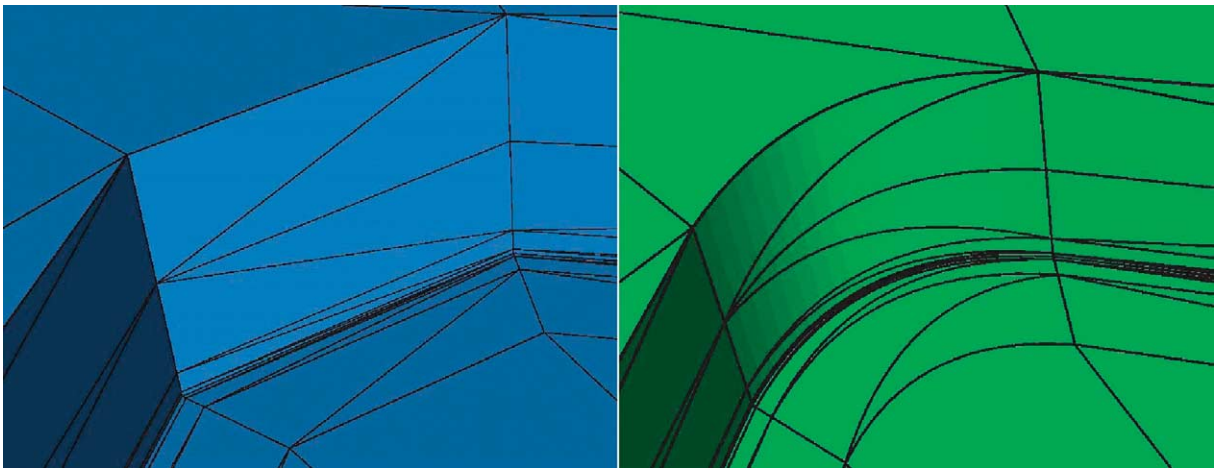

Fig. 13. Mesh around a portion of an isolated singular edges before (left) and after (right) curving.

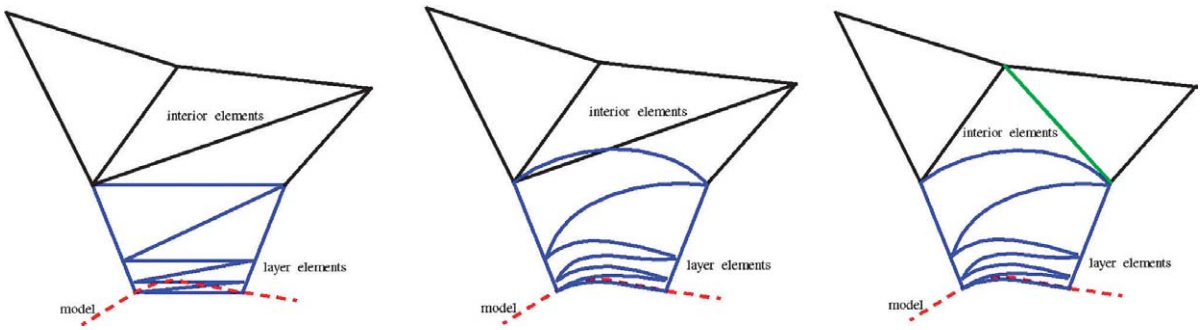

Fig. 14. The application of a swap to account for curving and edge in the edge isolation region.

entities to become invalid. Fig. 14 demonstrate this possibility for a simple 2D example. The linear mesh generated in both the singular edge isolation process and meshing the remainder of the domain is shown on the left side image in Fig. 14. When the mesh entities in the singular edge isolation region are curved, a connected element just outside that region becomes invalid (shown in the middle image of Fig. 14). By swapping one of the edges, the mesh becomes valid again. Although in this simple example, it would have also been straightforward to regain validity of the mesh by curving one mesh edge, it is common to have situations where this is not possible and combinations of curved mesh modification operations must be applied. After the surface mesh entities have been curved and a valid surface triangulation obtained, the final step is performing any required mesh modification to the interior mesh entities to regain validity of any elements that have become invalid. Fig. 15 shows an example of the procedure applied to a mechanical component to produce a mesh appropriate as the initial mesh for an adaptive hp-analysis. In this case the mesh is constructed to isolate the potential singular features with the remainder of the mesh being as coarse as possible.

The process of making the coarsest possible mesh for a model with geometric features (e.g., edge lengths) of substantially different sizes will create elongated elements, such as some of those on faces of the model in Fig. 15, are created. The mesh quality based on a priori geometric mesh entity shape

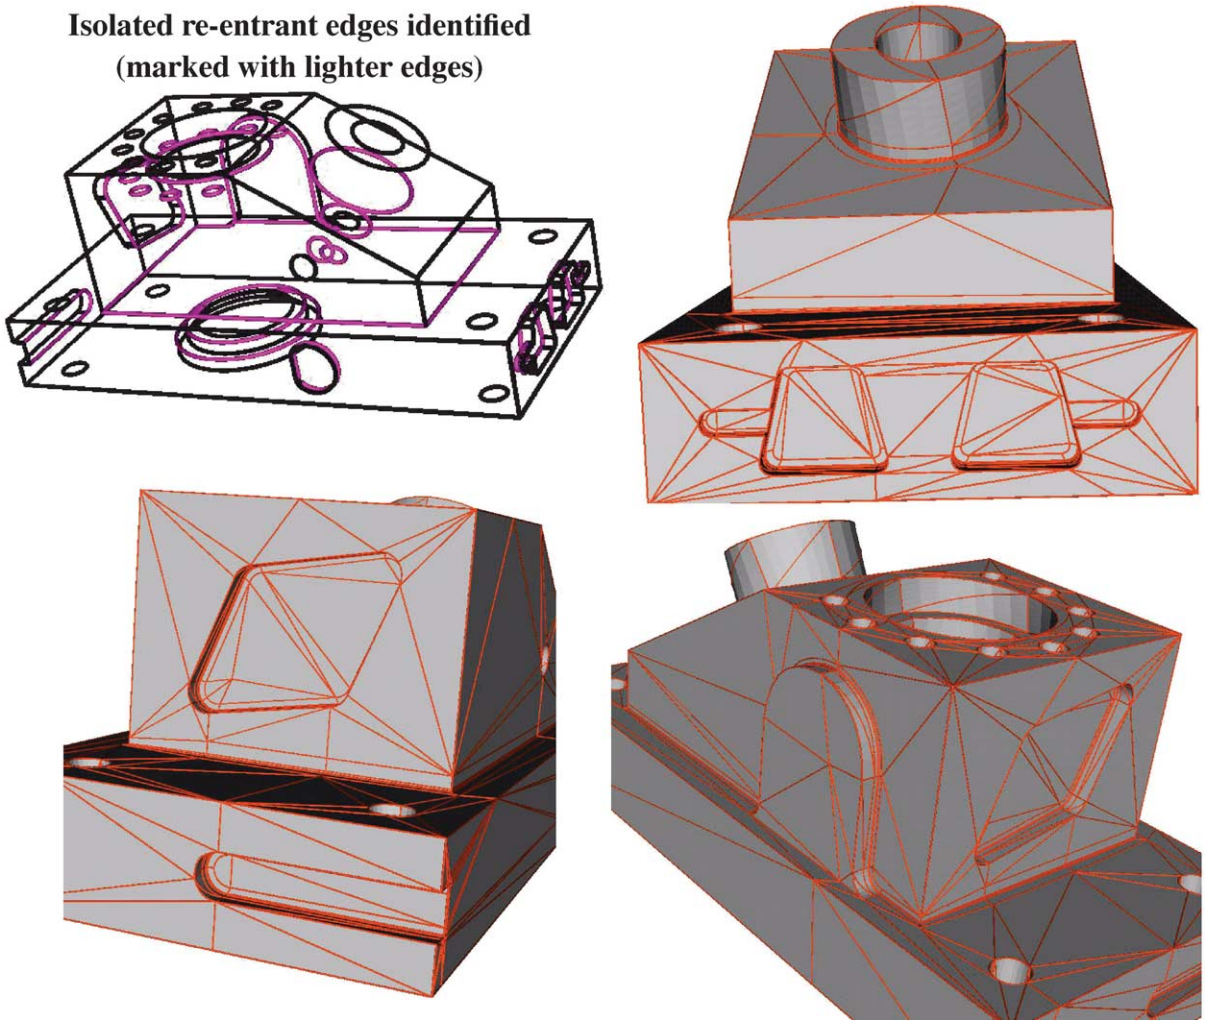

Fig. 15. Example hp-version mesh.

measures might not be considered good. However, if the mesh generator was required to create more well shaped, with respect to a regular element shape, the mesh would contain many more elements and would not be as likely to be able to attain the accelerated rates of convergence of an hp-adaptive method. Note that elements with shapes on the order shown here are a far way from causing any numerical conditioning problem, and that if the level of h-refinement in any direction is not sufficient the adaptive procedure should determine and correct it as part of the adaptive process.

6. Closing remarks

This paper has discussed the automatic generation of adaptively controlled meshes for general three-dimensional domains. The specific procedures presented include a mesh modification based method

capable of producing anisotropic mesh configurations, and a mesh generation procedure to produce curved meshes suited for hp-adaptive analysis.

As the adaptive results demonstrate, the anisotropic adaptive mesh modification procedures are capable of providing highly effective adapted meshes that can account for anisotropy of the solutions and the proper representation of curved domains given an adaptively defined anisotropic mesh metric field. To be most effectively applied these procedures need a new generation of error estimators and correction indicators that can account for the anisotropy of the solution field to construct the appropriate mesh metric fields.

Efforts on developing mesh generation tools for properly configured meshes for hp-adaptive meshes demonstrate the possibility of constructing meshes that can yield exponential rates of convergence for analyses over general curved 3D domains. As these procedures begin to mature it will also be important to develop hp-adaptive procedures that optimally indicate the appropriate mesh refinements and polynomial orders over the domain.

Acknowledgements

The procedures and results summarized in this paper have been derived from a number of projects supporting the authors of this paper. Sponsors of these projects include the National Science Foundation through a ITR grant No. ACI-0205741 and SBIR grant No. DMI-0132742, the Army Benet Laboratory, and the DOE ASCI Flash Center at the University of Chicago under contract No. B341495.

References

[1] S. Adjerid, K.D. Devine, J.E. Flaherty, L. Krivodonova, A posteriori error estimation for discontinuous Galerkin solutions of hyperbolic problems, Comput. Methods Appl. Mech. Engrg. 191 (2002) 1097–1112.
[2] B. Andersson, U. Falk, I. Babuska, T.V. Petersdorff, Reliable stress and fracture mechanics analysis of complex components using a $h-p$ version of fem, Internat. J. Numer. Methods Engrg. 38 (1995) 2135–2163.
[3] I. Babuska, Uncertainties in engineering design: Mathematical theory and numerical practice, in: J.A. Bennett, M.E. Botkin (Eds.), The Optimum Shape: Automated Structural Design, Plenum, 1986, pp. 171–197.
[4] I. Babuska, M. Suri, The p and $h-p$ versions of the finite element method, basic principles and properties, SIAM Rev. 36 (4) (1994) 578–632.
[5] I. Babuska, B. Anderson, B. Guo, J.M. Melenk, H.S. Oh, Finite element method for solving problems with singular solutions, J. Comput. Appl. Math. 74 (1996) 51–70.
[6] I. Babuska, T.V. Petersdorff, B. Anderson, Numerical treatment of vertex singularities and intensity factors for mixed boundary value problems for the Laplace equation, SIAM J. Numer. Anal. 31 (5) (1994) 1265–1288.
[7] M.W. Beall, J. Walsh, M.S. Shephard, Accessing CAD geometry for mesh generation, in: 12th International Meshing Roundtable, Sandia National Laboratories, SAND-2003-3030P, 2003, pp. 33–42.
[8] M.W. Beall, M.S. Shephard, A general topology-based mesh data structure, Internat. J. Numer. Methods Engrg. 40 (1997) 1573–1596.
[9] H. Borouchaki, P.L. George, F. Hecht, P. Laug, Saltel, Delaunay mesh generation governed by metric specifications, part I: algorithms, and part II: applications, Finite Elem. Anal. Des. 25 (1997) 61–109.
[10] H.L. de Cougny, M.S. Shephard, Parallel refinement and coarsening of tetrahedral meshes, Internat. J. Numer. Methods Engrg. 46 (1999) 1101–1125.
[11] S. Dey, R.M. O'Bara, M.S. Shephard, Curvilinear mesh generation in 3D, Computer Aided Geom. Design 33 (2001) 199–209.

[12] R. Garimella, M.S. Shephard, Boundary layer mesh generation for viscous flow simulations in complex geometric domains, Internat. J. Numer. Methods Engrg. 49 (1–2) (2000) 193–218.
[13] P.L. George, F. Hecht, Non isotropic grids, in: J. Thompson, B.K. Soni, N.P. Weatherill (Eds.), CRC Handbook of Grid Generation, CRC Press, Boca Raton, 1999, pp. 20.1–20.29.
[14] M.K. Georges, M.S. Shephard, Automated adaptive two-dimensional system for the hp-version of the finite element method, Internat. J. Numer. Methods Engrg. 32 (1991) 867–893.
[15] G. Kunert, Toward anisotropic mesh construction and error estimation in the finite element method, Numer. Methods Partial Differential Equations 18 (2002) 625–648.
[16] X. Li, M.S. Shephard, M.W. Beall, Accounting for curved domains in mesh adaptation, Internat. J. Numer. Methods Engrg. 58 (2003) 246–276.
[17] X. Li, M.S. Shephard, M.W. Beall, 3D anisotropic mesh adaptation by mesh modifications, Comput. Methods Appl. Mech. Engrg., 2005, in press.
[18] A. Liu, B. Joe, Relationship between tetrahedron shape measures, BIT 34 (1994) 268–287.
[19] X.J. Luo, M.S. Shephard, J.F. Remacle, R.M. O'Bara, M.W. Beall, B.A. Szabo, R. Actis, p-version mesh generation issues, in: 11th International Meshing Roundtable, Sandia National Laboratories, 2002, pp. 343–354.
[20] R.M. O'Bara, M.W. Beall, M.S. Shephard, Attribute management system for engineering analysis, Engrg. Comput. 18 (4) (2002) 339–351.
[21] S. Prudhomme, J.T. Oden, T. Westermann, J. Bass, M.E. Botkin, Practical methods for a posteriori error estimation in engineering applications, Internat. J. Numer. Methods Engrg. 56 (2003) 1193–1224.
[22] E. Rank, I. Babuska, An expert system for the optimal mesh design in the hp-version of finite element method, Internat. J. Numer. Methods Engrg. 24 (1987) 2087–2106.
[23] J.-F. Remacle, M.S. Shephard, An algorithm oriented mesh database, Internat. J. Numer. Methods Engrg. 58 (2003) 349–374.
[24] J.-F. Remacle, X. Li, M.S. Shephard, J.E. Flaherty, Anisotropic adaptive simulation of transient flows, Internat. J. Numer. Methods Engrg., submitted for publication.
[25] M.S. Shephard, Meshing environment for geometry-based analysis, Internat. J. Numer. Methods Engrg. 47 (1–3) (2000) 169–190.
[26] M.S. Shephard, M.K. Georges, Reliability of automatic 3D mesh generation, Comput. Methods Appl. Mech. Engrg. 101 (1992) 443–462.
[27] B.A. Szabo, I. Babuska, Finite Element Analysis, Wiley, New York, 1991.
[28] K. Weiler, The radial edge structure: a topological representation for non-manifold geometric modeling, in: M.J. Wozny, H. McLaughlin, J. Encarnacao (Eds.), Geometric Modeling for CAD Applications, IFIP WG5.2 Working Conference, Rensselaerville, NY, May 12–14, 1986, 1988, pp. 3–36.
[29] C.H. Whiting, K.E. Jansen, A stabilized finite element method for incompressible Navier–Stokes equations using a hierarchical basis, Internat. J. Numer. Methods Fluids 35 (2001) 93–116.