ORIGINAL ARTICLE

Automatic decomposition and efficient semi-structured meshing of complex solids

Jonathan E. Makem · Cecil G. Armstrong ·
Trevor T. Robinson

Abstract In this paper, a novel approach to automatically sub-divide a complex geometry and apply an efficient mesh is presented. Following the identification and removal of thin-sheet regions from an arbitrary solid using the thick/thin decomposition approach developed by Robinson et al. [1], the technique here employs shape metrics generated using local sizing measures to identify long-slender regions within the thick body. A series of algorithms automatically partition the thick region into a non-manifold assembly of long-slender and complex sub-regions. A structured anisotropic mesh is applied to the thin-sheet and long-slender bodies, and the remaining complex bodies are filled with unstructured isotropic tetrahedra. The resulting semi-structured mesh possesses significantly fewer degrees of freedom than the equivalent unstructured mesh, demonstrating the effectiveness of the approach. The accuracy of the efficient meshes generated for a complex geometry is verified via a study that compares the results of a modal analysis with the results of an equivalent analysis on a dense tetrahedral mesh.

Keywords Automatic decomposition ·
Geometric reasoning · Efficient meshing · Metric field ·
Anisotropic meshing · Structured meshing

1 Introduction

A major advantage of applying unstructured tetrahedral meshes to complex geometries for structural problems is the robustness of the current automatic tetrahedral mesh generators. One disadvantage is their limited capability for generating appropriate anisotropic stretched meshes. Element aspect ratio is a measure of the longest dimension of an element or region to the shortest dimension. For some CFD problems, meshes with very large aspect ratios can be generated and adaptively refined [2]. However, in structural problems a mesh density of one or two elements through the thickness of a thin sheet, or across the cross-section of a long-slender region, is often sufficient. This means that local, mesh-based approaches to anisotropic refinement are difficult. Alternatively, by generating a "mixed" or "hybrid" mesh consisting of structured meshes on specific regions of the model, combined with unstructured meshes on more complex areas, it is possible to achieve a dramatic reduction in degrees of freedom (dof) and thereby improve the efficiency of the analysis. This requires a process to partition the model into sub-regions where the different meshing strategies may be applied. This may be achieved by identifying easily mappable regions, and sweepable volumes such as thin sheets or long-slender sections, to which structured anisotropic pentahedral or hexahedral meshes may be applied. The remaining volumes in the model are referred to as being complex as structured meshes cannot be used to represent them. These may be filled with unstructured tetrahedral elements, which when combined with the structured meshes in the remainder of the model produce an efficient, semi-structured, mesh of the component.

Geometric properties of solids such as the medial axis transform have been used in the past to decompose complex geometries into thick and thin sheet sub-regions [1]. A structured shell mesh was applied to the thin-sheet regions and merged with isotropic unstructured elements in the adjoining thick regions. Even though a significant reduction in dof was achieved when using this approach for

J. E. Makem · C. G. Armstrong · T. T. Robinson (✉)
School of Mechanical and Aerospace Engineering,
Queen's University, Belfast BT9 5AH, UK
e-mail: t.robinson@qub.ac.uk

thin-walled component models, it was less than expected as the long-slender areas of the thick region consumed a lot of nodes when meshed with tetrahedra. However, if these long-slender regions could be identified, instead of being tetrahedral meshed they could be filled with a structured, swept mesh comprising anisotropic pentahedral or hexahedral elements, thereby reducing the dof even further.

The objective of this research is to develop an automatic approach that uses an a *priori* knowledge of shape properties to identify thin-sheet and long-slender regions in complex solids for the application of an efficient semi-structured mesh. Local sizing measures such as edge length and curvature, face width and curvature and local 3D thickness are employed to generate metric tensor fields which identify meshable sub-regions within complex volumes. Intelligent routines interrogate the geometry and automatically partition the body, isolating the long-slender regions from residual complex regions. Appropriate meshing strategies are applied to the respective sub-regions and the metric fields are used to grade the mesh in the structured mesh regions, ensuring a smooth transition with isotropic elements in residual complex areas. The effectiveness of the approach is demonstrated, with a substantial reduction in dof achieved.

The remainder of this paper is organized as follows: Sect. 2 reviews related work on geometric reasoning for meshing applications; Sect. 3 explains how the metric fields are generated using local sizing measures; Sect. 4 describes the approach for identifying the long-slender regions and partitioning the body into a non-manifold assembly of meshable sub-regions. It also details the mesh produced for two benchmark models of reasonable complexity and makes a comparison between the semi-structured and the equivalent unstructured tetrahedral meshes; Sect. 5 presents work on the automatic decomposition and efficient meshing of a gas turbine compressor inter-casing model; Sect. 6 reports on the results of a validation study that compares the results of modal analyses conducted on the efficient mesh, a coarse tetrahedral mesh with a similar number of dof and a dense tetrahedral mesh; Sect. 7 gives a discussion of the results and Sect. 8 presents conclusions.
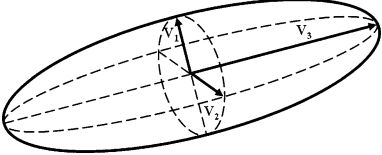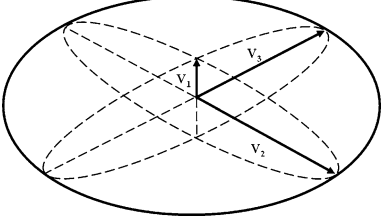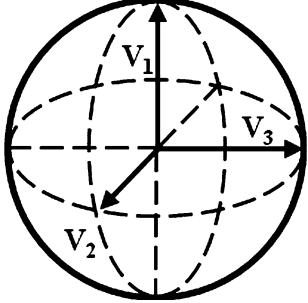
2 Literature review

Thakur et al. [3] gives a comprehensive summary of the state of the art in CAD model simplification techniques for meshing. Research of particular note to the work in this paper is the prototype thick/thin decomposition process developed by Robinson et al. [1], which provided a capability for locating thin-sheet regions in complex solid geometries. This functionality is now available in the commercial CAE tool CADfix, by TranscenData [4]. The approach uses the 3D medial object (MO) [5] of the CAD geometry to determine the local thickness followed by a 2D MO computation on the midsurfaces. By comparing the diameter of the inscribed disc which traced the 2D MO (which is an approximation of the lateral dimensions of the object) to the diameter if the inscribed sphere which traced the 3D MO (which is an approximation of object thickness), it is possible to determine if the region is a thin sheet or not. If the diameter of the inscribed circle on the 2D MO is large in comparison to the thickness, the region will be considered a thin sheet. Robinson et al. [1] also devised a method of predicting the reduction in dof which can be achieved using elements of large aspect ratio (m) by computing a "dof saving ratio". For thin sheets, a reduction in the order of m^2 is expected, whereas for long-slender volumes the reduction is proportional to m. Luo et al. [6] addressed the similar problem of finding thin sections for the generation of prismatic p-version finite element meshes, using an octree-based approach to identify medial surface points and local thin sections.

Although efficiently meshing the thin-sheet regions significantly reduces the dof in the model, it was noticed that there were still a lot of dof consumed when long-slender regions such as flanges were filled with an unstructured tetrahedral mesh. Price et al. [7, 8] proposed a logical approach to firstly partition a complex 3D object into meshable regions and secondly apply a structured mesh to these regions. The medial surface was used to sub-divide the solid into hex-meshable sub-regions known as primitives. After all the subregions have been formed, each primitive is meshed using a midpoint subdivision technique [9]. Mesh compatibility between the primitives can be controlled by an integer programming technique described by Tam et al. [10]. However, the approach has limitations as it relies on a robust 3D medial object computation which is currently not achievable for all geometries. Moreover, a comprehensive treatment of all possible shape features has not been developed.

Tchon et al. [11] used an a *priori* knowledge of model geometry to generate a Riemannian metric based on local curvature and thickness. An isotropically refined octree grid is used as a support medium to generate the shape metrics. The anisotropic sizing information provided by the metrics is then used to refine the mesh. Some impressive results have been achieved with this approach with a dramatic reduction in element count. However, the quality of the elements generated after the refinement process is questionable as their shape is not perfectly cubical. Other authors have reported similar geometry-

Table 1 Ellipsoid classification

Ellipsoid type	Diagram	Criteria
Long-slender		$V_3 \gg V_1$ $V_3 \gg V_2$ $V_1 \approx V_2$
Thin sheet		$V_3 \gg V_1$ $V_2 \gg V_1$ $V_2 \approx V_3$
Isotropic		$V_1 \approx V_2 \approx V_3$

based methods for mesh adaption and refinement [12, 13]. However, all of them relate to the relocation or refinement of an initial mesh. The focus here has been to use similar shape measures to identify regions of the model on which to apply structured meshing, rather than local mesh adaption.

Zhao et al. [14, 15] proposed a technique for the adaptive generation of an initial mesh based on the geometric features of a solid model. A special refinement field based on curvature and thickness was constructed to control mesh size and density distribution. The boundary of the adapted mesh produced was matched to the boundary of the solid model using a threading method. Although the resultant mesh on a complex model possessed fewer dof than a structured swept mesh applied to a manually sub-divided version, the quality of the mesh produced was not satisfactory as a significant number of mesh singularities were evident.

White et al. [16] developed an algorithm for automatically decomposing multi-sweep volumes into volumes that can be swept meshed. This is achieved by discretising the linking faces of volumes with quad elements which provide a layering system in the transverse direction. Target faces are pushed through the volume onto opposing source faces providing an imprint that governs the decomposition of the solid. The newly generated volumes are all sweep meshable with a single target face. However, the major limitation with the approach is that it can only be applied to sweepable geometries and cannot be used to decompose more complex models. Consequently, this identifies a requirement for a more robust technique for the automatic decomposition of complex models that will successfully partition any dumb solid geometry into an assortment of sub-regions, including sweepable volumes, where various structured meshing strategies may be applied.

3 Geometric reasoning using local sizing measures

The approaches and methodology reported in the forthcoming sections have been implemented in the CADfix software package using an API based on the Python programming language. Assuming that the thin-sheet regions have been identified and partitioned out using the thick/thin decomposition tool available in CADfix [4], decomposition of the remaining thick region into long-slender and
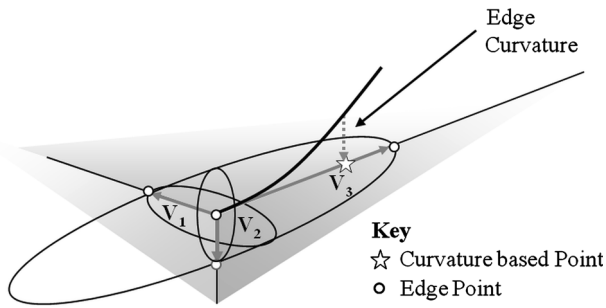
Fig. 1 1D metric generated using edge length and curvature

complex bodies may be achieved using a shape metric sized using the following steps.

3.1 Metric classification

The goal is to generate, within the CAD model, ellipsoids representing the shapes of different regions. The principal axes of an ellipsoid may be used to define the target element size in three directions [17].

If an ellipsoid has one principal direction that is much greater than the other two, this identifies a long-slender region where the mesh can be extruded in the direction of the largest principal axis. Areas of the model where each axis is similar in magnitude, and the ellipsoid is approximately spherical, identify regions where the target element size will be similar in all directions, and should be meshed using an unstructured isotropic tetrahedral mesh. The three main types of ellipsoids are displayed in Table 1. The criteria described assumes that

$V_3 > V_2 > V_1$. Note that regions represented by thin-sheet ellipsoids will have been removed from the model by the thick/thin tool prior to implementing the approach presented here.

Within the procedures an ellipsoid is generated on every edge with its centre at the midpoint of the edge. The edge length and curvature are used to determine the length of the ellipsoid axis along the edge direction. In terms of long-slender ellipsoids, this sizing measure is normally employed to gauge the size of the largest principal axis which provides the sweep direction for the extruded mesh. The other sizing measures of face width and curvature, and local 3D thickness are employed to determine the extent of the remaining two principal axes. If a face is planar, the length of an axis may be calculated using the 2D medial axis on the face or an approximation to it. The axis length on non-planar surfaces may be determined using a curvature-based sag value, d. For other scenarios ray casting is used to assess axis length by giving an approximation of the local thickness of the part.

3.2 Metric sizing using edge length and curvature

The use of local sizing measures based on edge length and curvature for mesh sizing applications is well documented [18]. In this work, if the edge is straight, then half the edge length is used for the extent of the ellipsoid axis by default, as shown in Fig. 1. However, if an edge is curved, it is necessary to employ a curvature sensitive approach to check the size of the ellipsoid axes along curved edges and surfaces. Along an edge, the change in tangent vector, $\mathbf{T_i}$, quantified by the angle α defined in Eq. (1), is used to

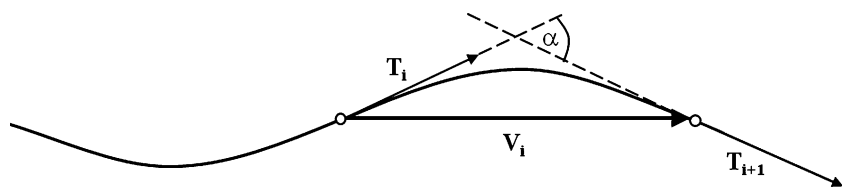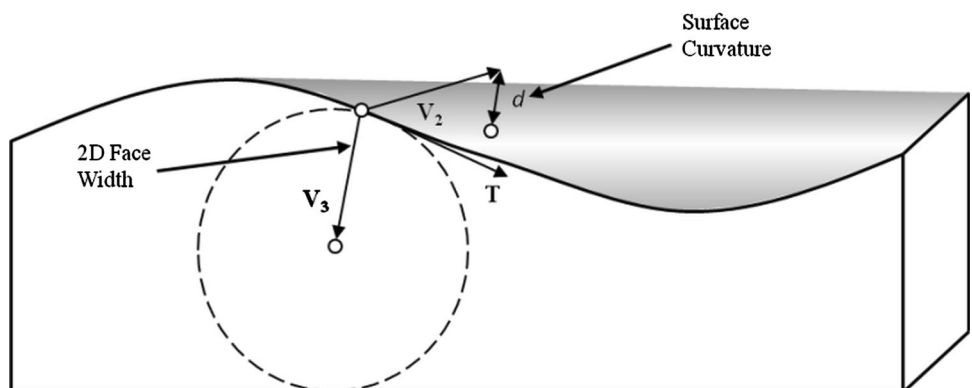**Fig. 2** Metric sizing using edge curvature

Fig. 3 Sizing of ellipsoid axis using surface curvature and face width

determine the size of the first ellipsoid axis, $\mathbf{V_i}$, shown in Fig. 2. When α reaches a user-defined tolerance value the ellipsoid axis vector $\mathbf{V_i}$ can be determined.

$$\alpha = \cos^{-1}\left(\frac{T_i \cdot T_{i+1}}{|T_{i1}||T_{i+1}|}\right). \tag{1}$$

3.3 Metric sizing using face width and curvature

If the faces adjoining an edge are curved, a similar strategy is applied to assess the size of the second and third axes. Methods for metric generation based on surface curvature and proximity are discussed by Quadros et al. [19]. The surface curvature is used to determine the size of the ellipsoid axis based on a sag value, d. For a pre-defined d value, the size of the vector orthogonal to the tangent vector at the centre point on the edge can be calculated, shown as $\mathbf{V_2}$ in Fig. 3. In cases where the surface is planar, or for surfaces with slight curvature with a maximum sag value less than d, the 2D medial object is used to determine the local face width and the length of the ellipsoid axis in this direction.
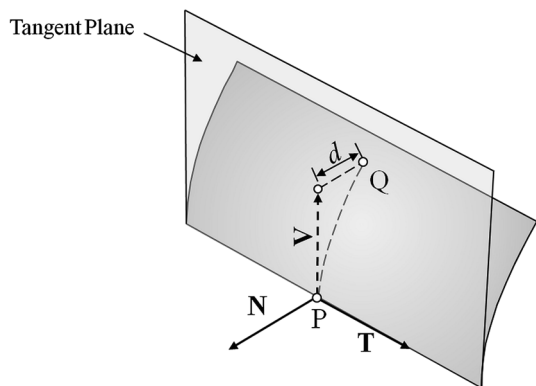

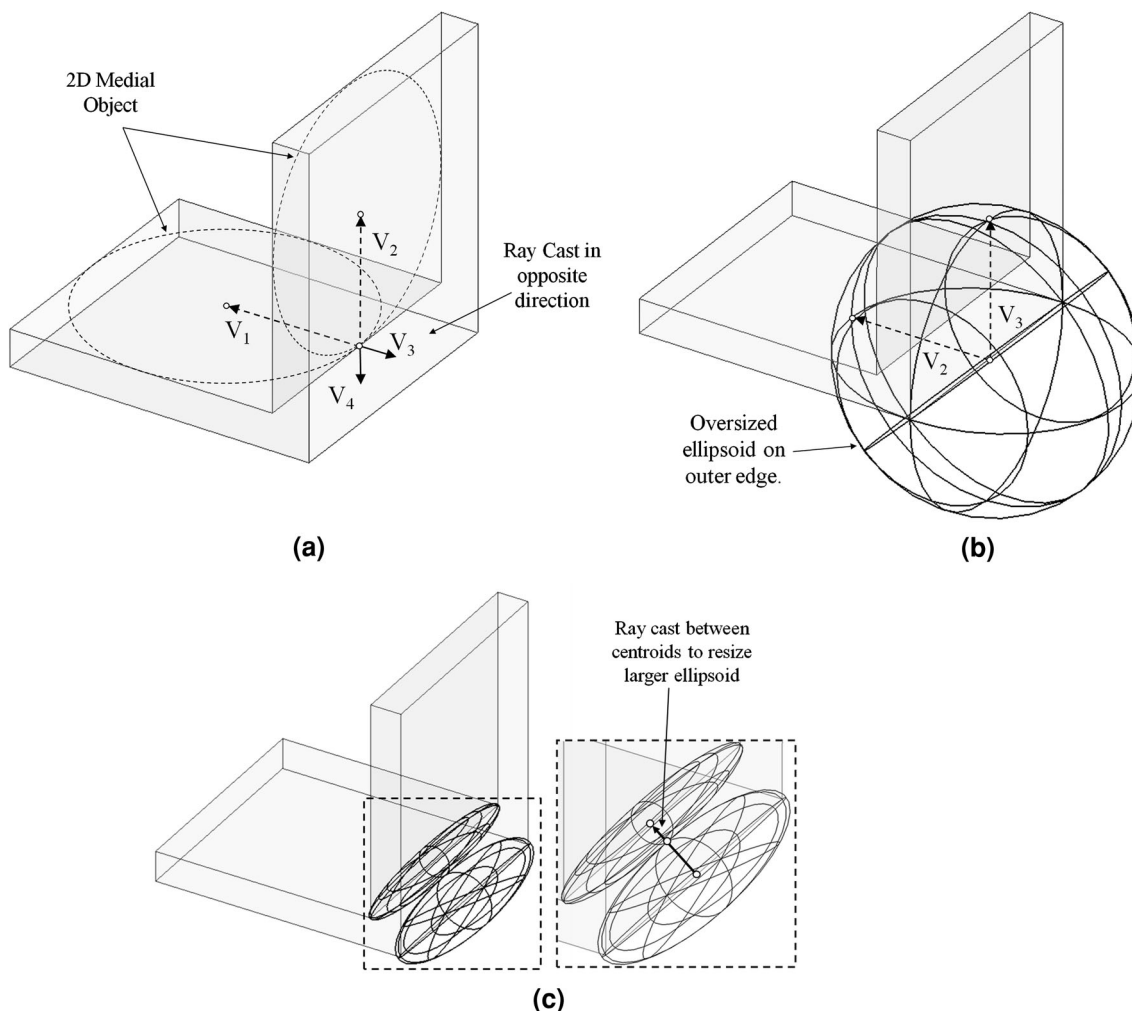Fig. 4 Metric sizing for a given sag value on a parametric surface

Fig. 5 Metric sizing by assessing local 3D thickness. **a** Ray casting through the thickness to establish the shortest ellipsoid axis. **b** Oversized ellipsoid causing interference with adjacent ellipsoid. **c** Resizing of large ellipsoid to remove interference

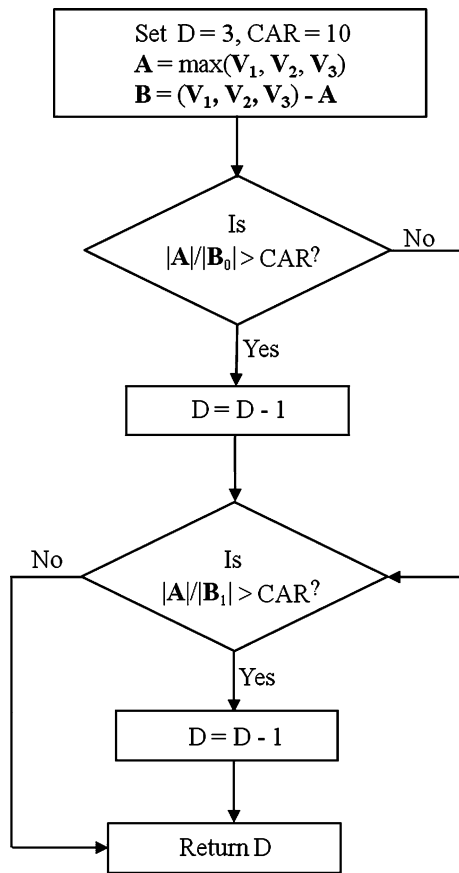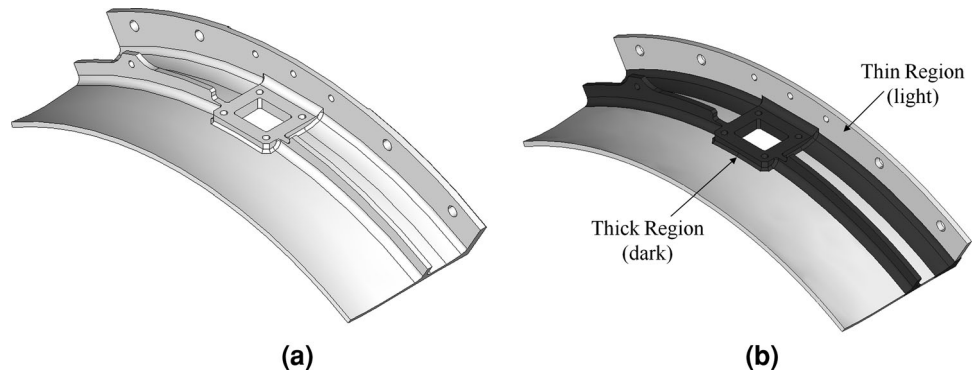Fig. 6 Thick/thin decomposition of complex model. **a** Dumb solid model, **b** sub-divided model

Fig. 7 Selection process for 1D ellipsoids using a critical aspect ratio

On a parametric surface, as shown in Fig. 4, for a given tangent vector, **T** and normal vector **N** at point P on a given edge, the sag value, d at point Q on the surface is equivalent to the projection of **PQ** onto **N**, as

$$d = PQ \cdot N. \tag{2}$$

d can also be described as the perpendicular distance from Q to the tangent plane at P. In parametric terms, P and Q

can be represented by the points $x(u, v)$ and $x(u + du, v + dv)$. Consequently, using Taylor's theorem [20], Eq. (2) can be modified to give

$$d = \frac{1}{2}d^2x \cdot N + O(du^2 + dv^2). \tag{3}$$

Thus, for a pre-defined sag value, d, Eqs. (2) and (3) can be used to determine the location of point Q and the extent of the ellipsoid axis, **V**.

3.4 Metric sizing using local thickness

Local 3D thickness is also considered as a sizing measure to gauge the extent of the ellipsoid axis. In the past, Quadros et al. [21] employed disconnected 3D skeletons to measure 3D proximity in metric generation. However, for this approach, ray casting is employed to assess the local thickness of the body.

As shown in Fig. 5a, ray casts are fired in the direction V_3 and V_4, which are in the opposite directions to axes V_1 and V_2 generated using the 2D medial object. An axis length based on the thickness of the body ($|V_3|$ or $|V_4|$) is compared to an axis length generated using the alternative sizing measures ($|V_1|$ or $|V_2|$), and the shortest axis is selected for the ellipsoid in that direction. However, the ellipsoids generated using the aforementioned sizing measures may be too large resulting in the interference between adjacent ellipsoids, as shown in Fig. 5b. In this case, the size of the ellipsoid is reduced to eliminate interference, as shown in Fig. 5c.

4 Automatic partitioning of complex geometry

After a thick/thin decomposition has been performed on a complex geometry, as shown in Fig. 6a, a search is performed on the resulting thick body (darker shading in Fig. 6b) for long-slender regions by generating the ellipsoids described in Sect. 3 at the midpoint of every edge.

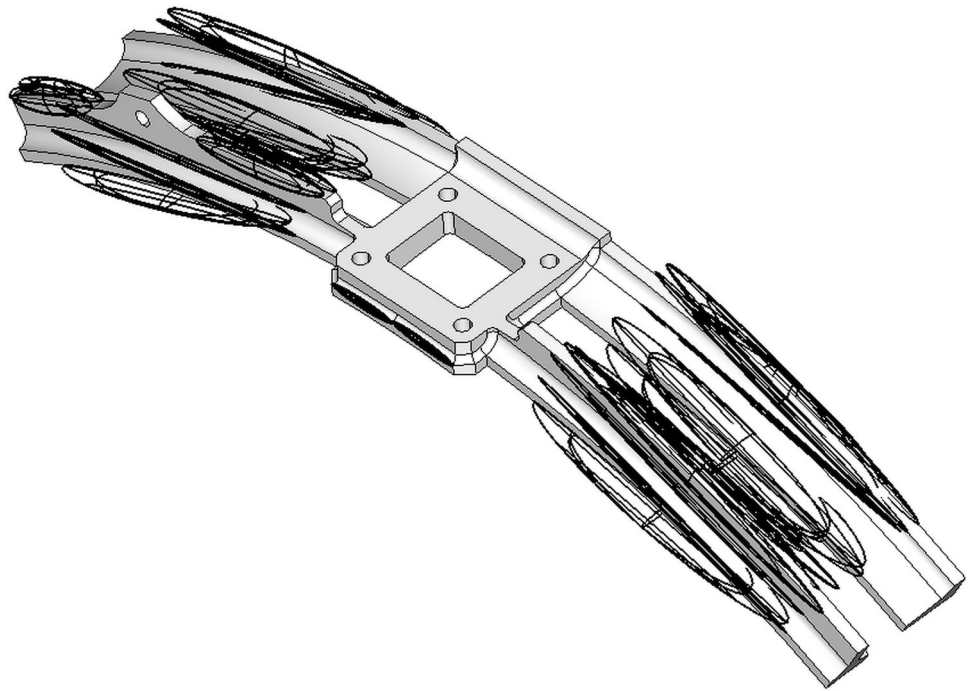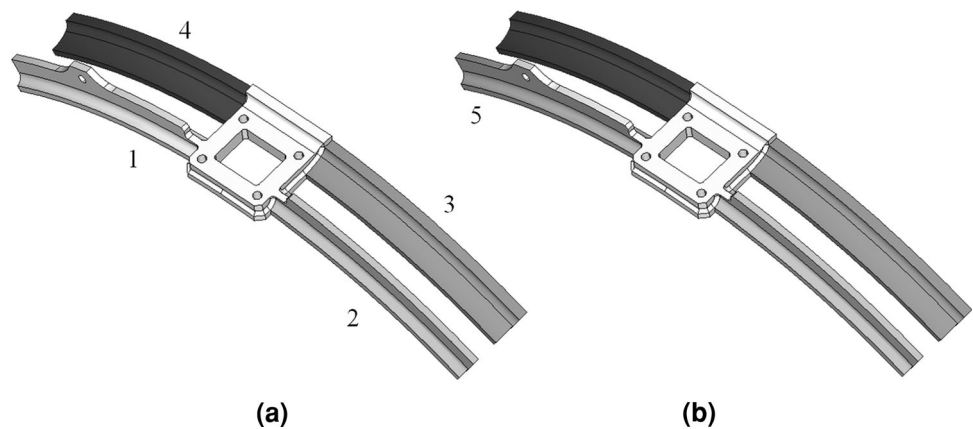**Fig. 8** Long-slender ellipsoids on the thick region of complex model

Fig. 9 Long-slender regions identified on complex model

(a)   **(b)**

4.1 Identification of long-slender regions

The aspect ratio of an ellipsoid compares the length of its longest axis to the length of its other axes. Critical aspect ratio (CAR) is a user-specified ratio value that is used to determine when an ellipsoid axis is considered to be much longer than the other two. For a pre-defined critical aspect ratio, all long-slender ellipsoids are identified using the algorithm detailed in Fig. 7. Here D represents the dimension of the ellipsoid (for long-slender ellipsoids $D = 1$), and V_1, V_2 and V_3 are the vectors defining the three principal axes of the ellipsoid. A is the vector with the greatest magnitude and B defines a set of vectors excluding the vector with the greatest magnitude. The thick region on

the model shown in Fig. 6 is shown in Fig. 8 along with all of the long-slender ellipsoids computed.

After all the long-slender ellipsoids are generated, a "closed-loop" searching algorithm is initiated that uses the ellipsoids to look for closed loops of surfaces which each bound a long-slender section. For any given edge on which a long-slender ellipsoid is located, the algorithm searches for an adjoining surface that has another long-slender ellipsoid on one of its bounding edges. If such an ellipsoid exists, this establishes a direction in which to continue the search, in either a clockwise or anti-clockwise fashion. The search is repeated for the other adjacent surface and if successful the process continues until a surface is located that contains an edge with the ellipsoid that was used for

Fig. 10 Closed-loop searching algorithm. **a** Select an ellipsoid, **b** locate an ellipsoid on the edge of an adjoining surface, **c** continue the search in the search direction, **d** stop when a closed loop is formed
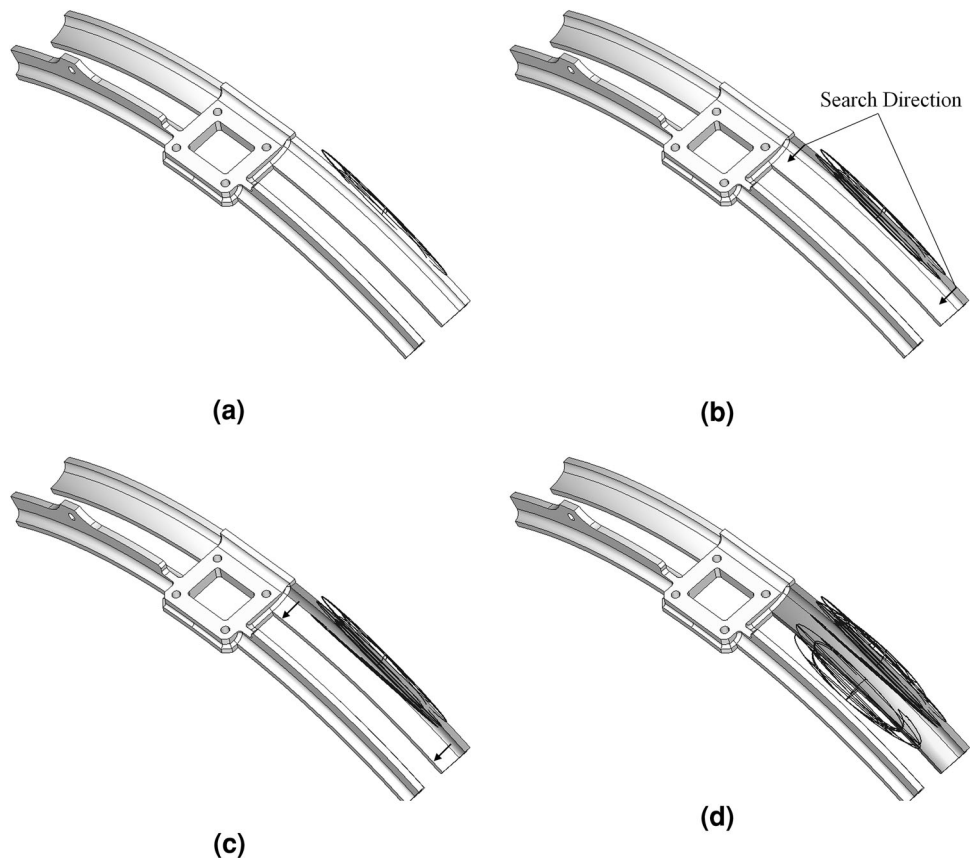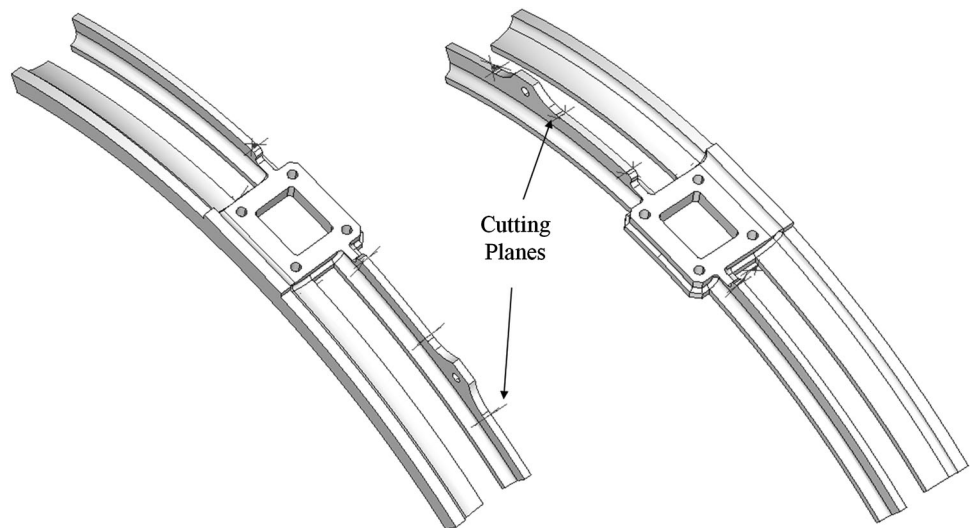

(a)

(b)

(c)

(d)

Fig. 11 Automatic generation of cutting planes

the initial iteration. At this point, a long-slender region comprising a closed loop of surfaces is identified and these surfaces are then removed from the search. A new edge is selected and the process is repeated until all closed-loop surface groups have been identified in the model. The results of the closed-loop search are shown in Fig. 9. Figure 10 describes the closed-loop searching algorithm for long-slender extruded regions. The algorithm successfully

identifies five long-slender regions shown in Fig. 9, where each region is highlighted in a different shade of grey. Note that the bottom-left arm of the model is composed of two separate long-slender regions denoted by 1 and 5. The different surfaces identifying those regions are in different colours on the left and right image. All possible long-slender regions have been automatically identified for this example. Since the closed-loop search is initiated at all

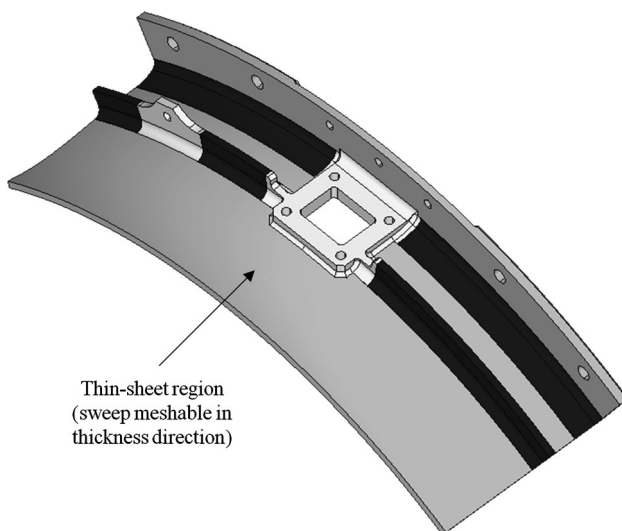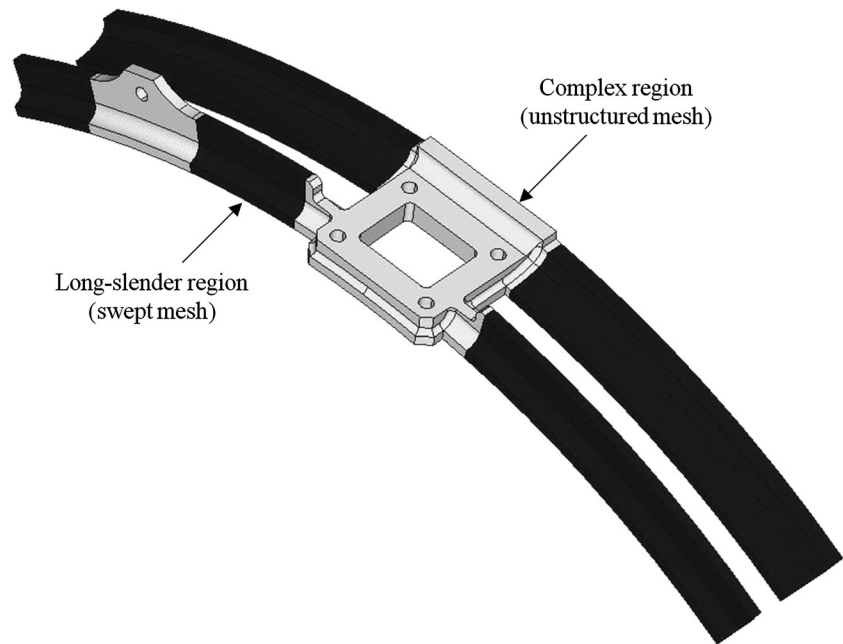**Fig. 12** Automatic partitioning into long-slender and complex sub-regions

Fig. 13 Full decomposition of complex model

edges having long-slender ellipsoids, the search for long-slender regions is comprehensive.

4.2 Automatic partitioning into meshable sub-regions

After the groups of surfaces that define each long-slender region have been identified, the next stage of the process involves generating meshable long-slender bodies by partitioning the surface groups using cutting planes. Open ended long-slender regions, which do not require any cutting planes, can be identified because the 'capping' surface shares an edge with all of the surfaces bounding the

sweepable region. Cutting planes are generated using the edge tangent and surface normal at an offset (calculated as a fraction of the shortest edge length) away from the long-slender-complex region interface, as shown in Fig. 11. The cutting plane(s) for each edge are grouped into a set and given an ID based on the surface group.

After associating each surface group with a cutting plane, the body is automatically partitioned into a non-manifold assembly of long-slender and residual complex bodies using a series of cutting commands. The cutting plane ID's are then used to identify if a body is complex or long-slender. Any body that lies between the cutting planes in a set is deemed to be a long-slender body. The result of the automatic partitioning process is shown in Fig. 12.

Finally, by combining the partitioned thick region generated using the procedures here with the thin region generated from the thick/thin decomposition it is possible to achieve the complete decomposition of the model into long-slender, thin-sheet and residual complex sub-regions, as shown in Fig. 13.

4.3 Detection of non-sweepable regions

On application of the aforementioned automatic partitioning technique to more complex geometries, such as that shown in Fig. 14, limitations were exposed in the process used to position the cutting planes. From the thick/thin decomposition of the above model, it became evident that certain long-slender regions of the thick body exhibit end features which are not sweepable due to their complex topology, as displayed in Fig. 15. To detect complex ends on long-slender regions the 2D medial object on the

Fig. 14 Thick/thin decomposition of casing section. **a** Thick regions between mounts and thin sheet. **b** Thick regions on leading and trailing edge
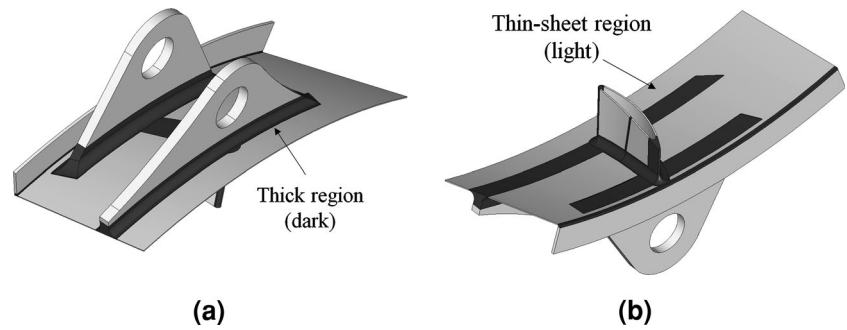
Thin-sheet region (light)

Thick region (dark)

(a)

(b)

Fig. 15 A non-sweepable, complex end on a long-slender region
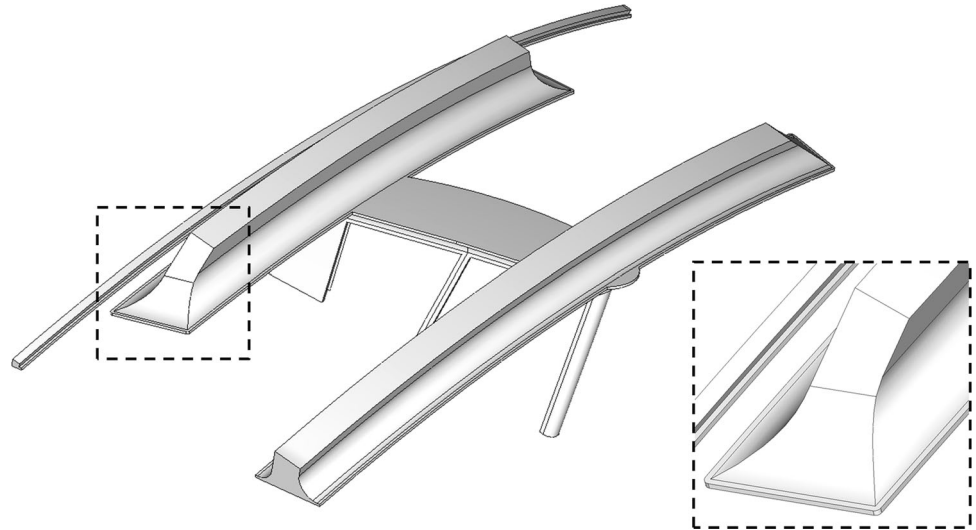
Fig. 16 Complex end detection via 2D medial bracelet search

Search Direction

Connecting face not part of closed loop

Shortest Edge

Start Position (mid point of shortest edge)
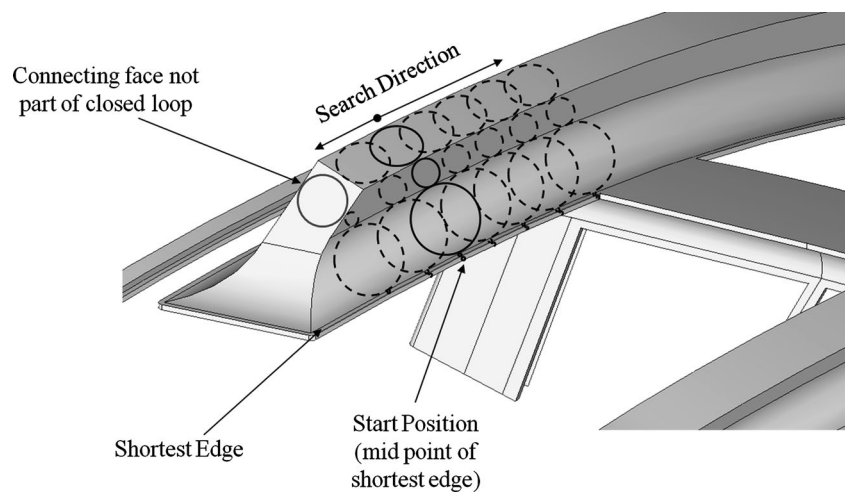
surfaces comprising each closed loop was used, as shown in Fig. 16.

Given a starting point at the midpoint of the shortest edge along the length of a closed loop of surfaces, a chain of inscribed circles on the 2D medial object of each surface was built up around the periphery of the 1D region. This chain of inscribed circles eventually forms a so-called "bracelet", when the touching point on the last surface meets the original starting point. The bracelet is then re-formed at intervals along the length of the 1D region in opposite directions away from the edge midpoint. As the bracelet traverses along the 1D region, if it encounters a complex piece of geometry (a surface that is not bounded by all closed-loop surfaces) an inscribed circle will be

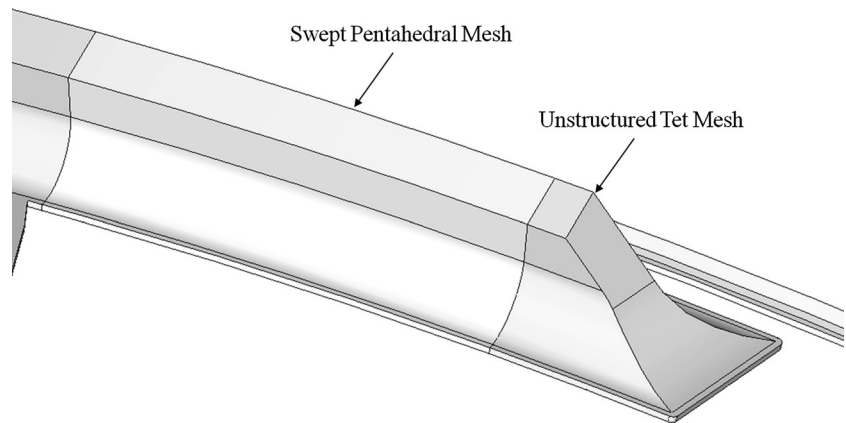Fig. 17 Partitioned complex end on long-slender region

Swept Pentahedral Mesh

Unstructured Tet Mesh

Fig. 18 Decomposition of complex model into thin-sheet, long-slender and residual complex volumes

(a)				**(b)**
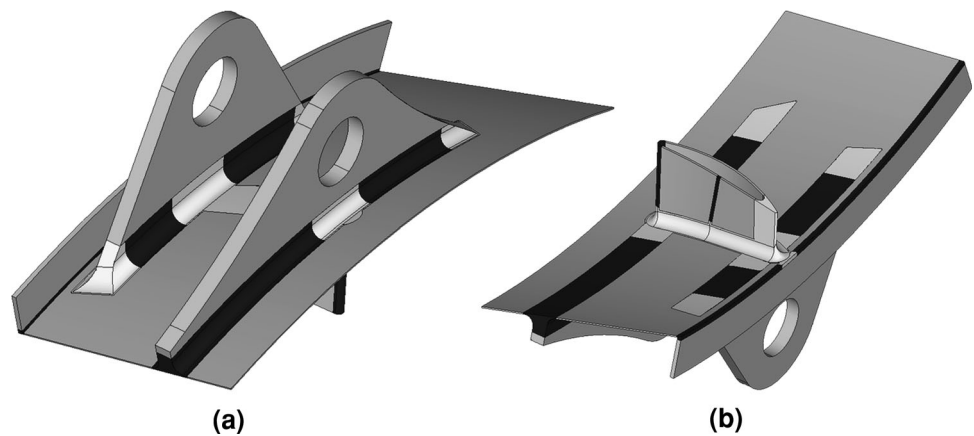
Fig. 19 Efficient meshing of decomposed complex model 1

Long, slender region 45% reduction in degrees of freedom

Thin-sheet region 86% reduction in degrees of freedom
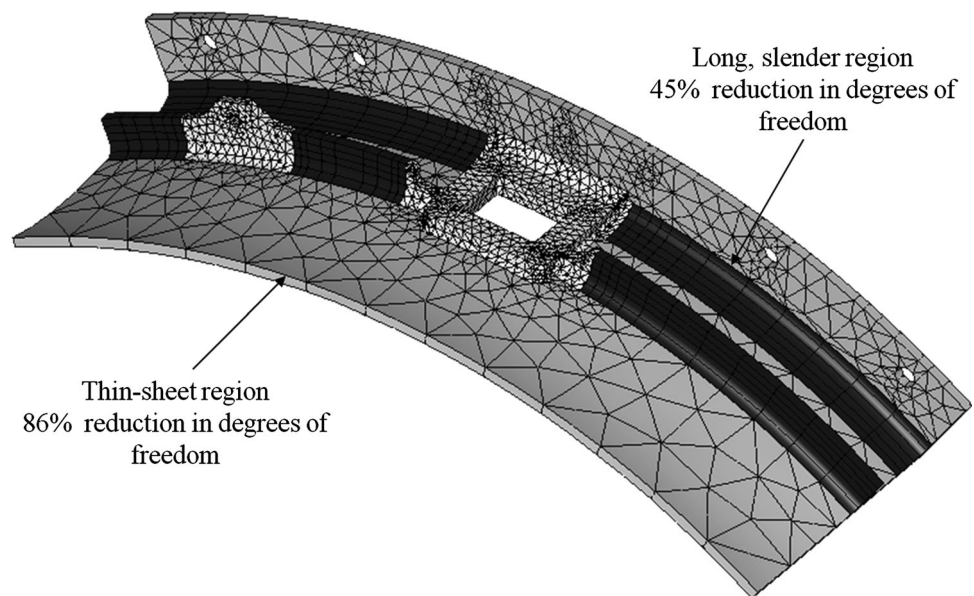
formed on a non-closed-loop surface and the bracelet will be broken. At this point it is considered that a complex end has been detected, which will be partitioned out at a given offset, as shown in Fig. 17. The offset distance is user defined, typically as a fraction of a characteristic size of the cross-section.

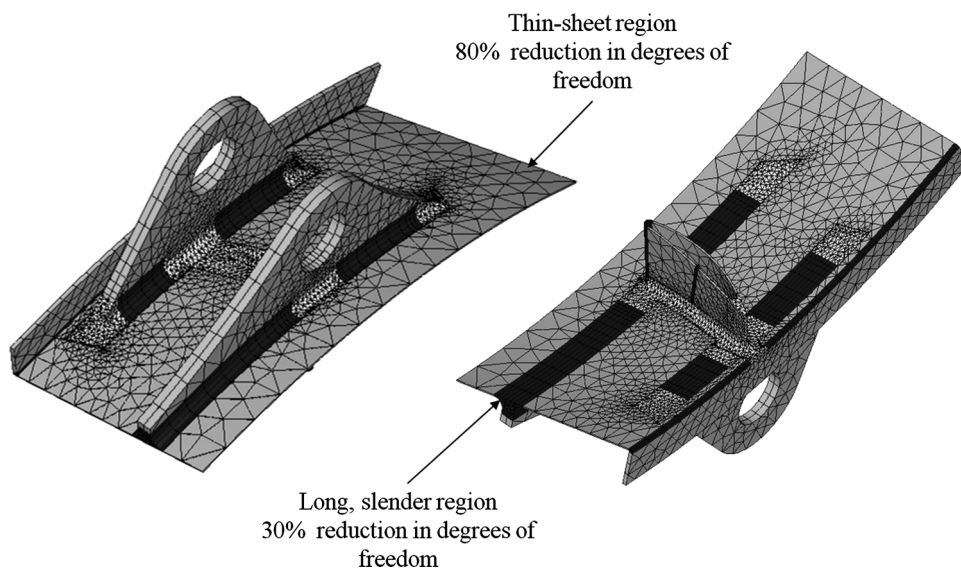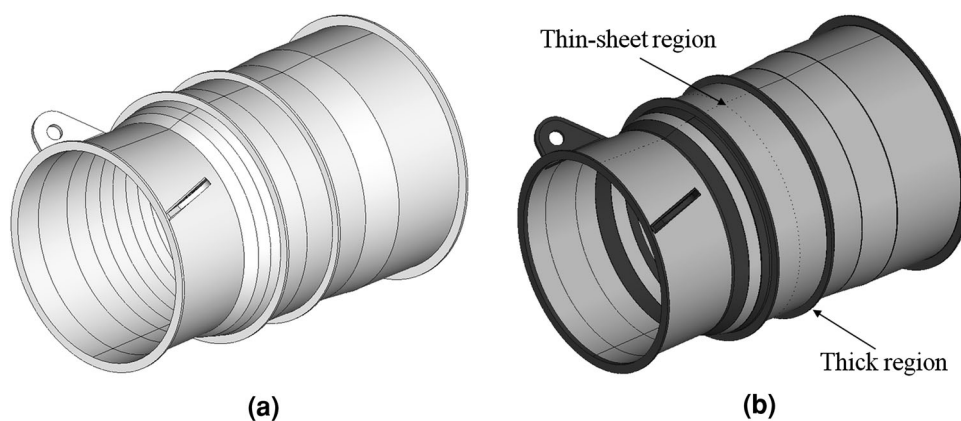Fig. 20 Efficient mesh of decomposed complex model 2

Thin-sheet region
80% reduction in degrees of freedom

Long, slender region
30% reduction in degrees of freedom

Fig. 21 CRESCENDO compressor inter-casing. **a** Original model, **b** thick/thin decomposition

Thin-sheet region

Thick region

(a)   **(b)**

By applying this approach to the rest of the model, it is possible to detect and partition out all non-sweepable complex features on the long-slender regions of the thick body. Also, combining the 1D/3D decomposition of the thick region with the adjacent thin-sheet bodies results in a full decomposition of the complex casing model into an assembly of long-slender, thin-sheet and complex bodies, each of which has a different meshing strategy applied, as shown in Fig. 18. This illustrates the approach as a valid means to automatically subdivide complex geometries into cells for the application of semi-structured mesh.

4.4 Meshing

After the model has been sub-divided a structured swept mesh can be applied to the long-slender and thin-sheet bodies. The complex volumes can be meshed with unstructured isotropic tetrahedral elements. In transition areas where non-conforming mesh types meet (e.g. hexahedral meet tetrahedral), pyramid elements are automatically inserted to ensure full conformity. Figures 19 and 20

display meshes which were automatically generated on the models described earlier in the section.

The ellipsoids are not only used to partition the model but can also grade the mesh in the long-slender regions. Each ellipsoid on the edge of a 1D sweepable region is used to define a "source field" on the model in CADfix. These source fields control the growth of the mesh and ensure a smooth transition between the 1D and 3D regions. The number of dof in the swept meshes was compared with an unstructured tetrahedral mesh with the element size set to the thin-sheet thickness. Overall, for model 1 shown in Fig. 19, a 69 % reduction in dof is achieved. For model 2 shown in Fig. 20, a 63 % reduction in dof is observed.

5 Automatic decomposition and efficient meshing of aero-engine casing geometry

The approach was applied to a representative industrial model. The CRESCENDO [22] compressor inter-casing casing, shown in Fig. 21a, is an example of an aero-engine

Fig. 22 Full decomposition of the CRESCENDO compressor inter-casing into long-slender, thin-sheet and residual complex regions

Complex region (light)

Long, slender region (dark)

Thin-sheet region (medium)
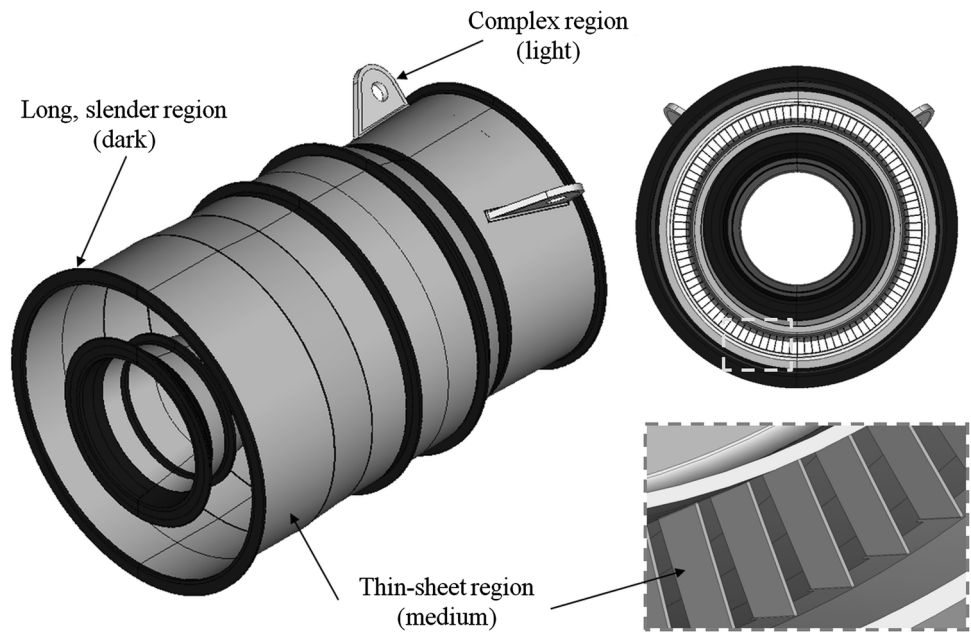
Fig. 23 An efficient, structured mixed solid mesh of the CRESCENDO compressor inter-casing

Thin-sheet region (medium) 93% reduction in degrees of freedom

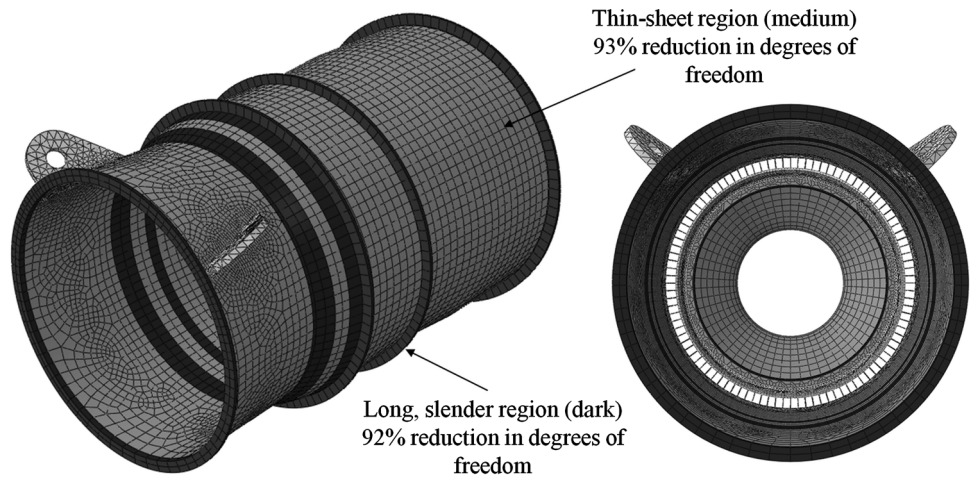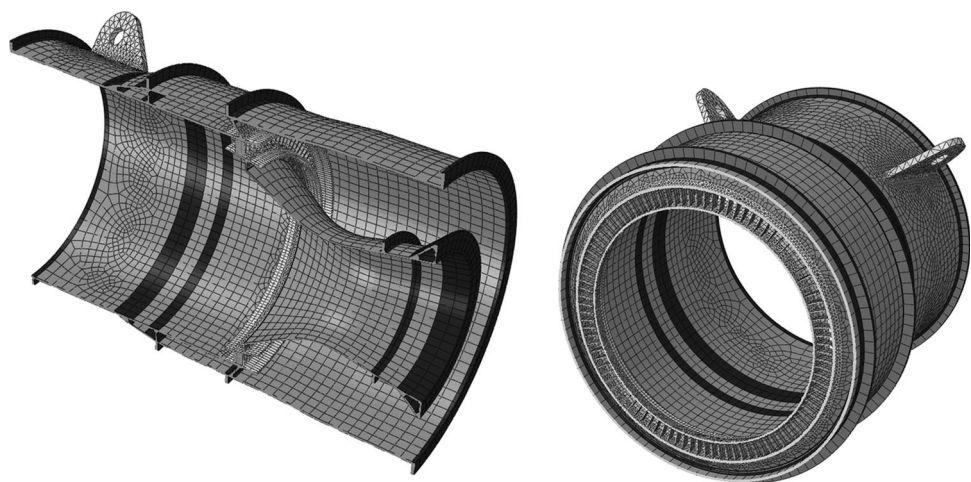Long, slender region (dark) 92% reduction in degrees of freedom

Fig. 24 Sectional views of the efficient mesh

component with large areas of thin-sheet, long-slender regions and complex details. As before, the first stage of the decomposition process was to partition the model using CADfix's thick/thin sub-division tool, the result of which is shown in Fig. 21b.

After performing a thick/thin decomposition on the model, the approach successfully identified and partitioned a total of 18 long-slender regions in the thick body as shown in Fig. 22. Here dark shaded areas represent long-slender regions, medium shaded areas represent thin-sheet regions and light shaded areas represent residual complex regions.

Again a structured swept mesh could be applied to the long-slender and thin-sheet regions of the model. The

Table 2 Material properties

Material property	Value
Density	7,850 kg/m^3
Poisson's ratio	0.30
Young's modulus	200 GPa

Table 3 Mass properties

Model	Mass (kg)	Centre of mass (mm)
Dense tetrahedral mesh	191.35	(4,542.2, 0.00090, 4.8088)
Efficient mesh	192.08	(4,542.4, 0.00078, 4.7897)
Coarse tetrahedral mesh	191.50	(4,542.1, 0.018931, 4.8048)

Table 4 Number of degrees of freedom (dof) for long-slender, thin-sheet and complex volumes

Mesh type	Long-slender region	Thin-sheet region
Efficient mesh	128,490	248,970
Dense tetrahedral mesh	1,473,150	3,939,294

resulting mesh of the assembly was generated in Abaqus/CAE [23], and consisted of several second-order element types. All thin-sheet volumes were meshed with hexahedral elements. The majority of long-slender regions were meshed with pentahedral elements, and all complex regions with tetrahedral elements, as shown in Figs. 23 and 24.

Overall an 89 % reduction in dof is observed, from approximately 2,000,000 nodes in an unstructured dense tetrahedral mesh (with an element size of the order of the thickness of the thin sheet) to approximately 215,000 nodes in the efficient mixed solid mesh. To verify the accuracy of the structured mesh, a modal analysis was conducted on the respective models.

6 Analysis and validation of efficient mesh

A free–free modal analysis was conducted on both the efficient structured mesh, a coarse tetrahedral mesh with approximately the same number of dof, and a dense tetrahedral mesh (with an element size of the order of the thickness of the thin sheet) using Abaqus/CAE v6.11. The efficient structured mesh was composed of C3D20R hexahedral elements, C3D15 pentahedral elements and C3D10 tetrahedral elements. The tetrahedral meshes used in the complex regions were entirely composed of C3D10 elements. The material properties used are shown in Table 2. The mass and centre of mass of the respective meshes are given in Table 3. Table 4 compares the number of dof which occupy the long-slender and thin-sheet regions of the inter-casing for both the efficient mesh and reference dense tetrahedral mesh.

Tie constraints were used at non-conforming 1D–3D and 2D–3D mesh interfaces. A comparison was made between the modal shape and frequency for the first 20 elastic modes. All mode shape plots, such as that shown in Fig. 25, show no noticeable difference in shape. In each
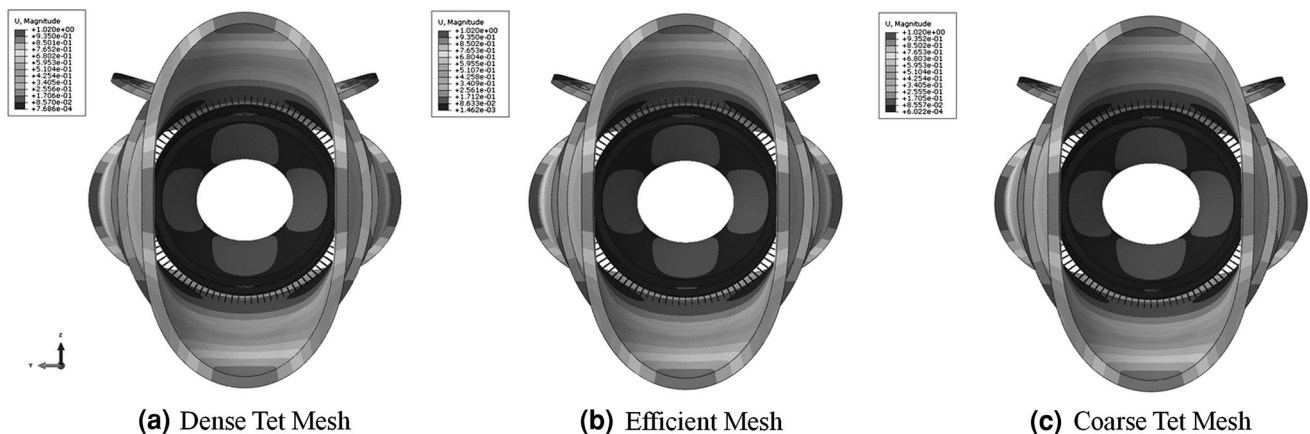

(a) Dense Tet Mesh   **(b)** Efficient Mesh   **(c)** Coarse Tet Mesh

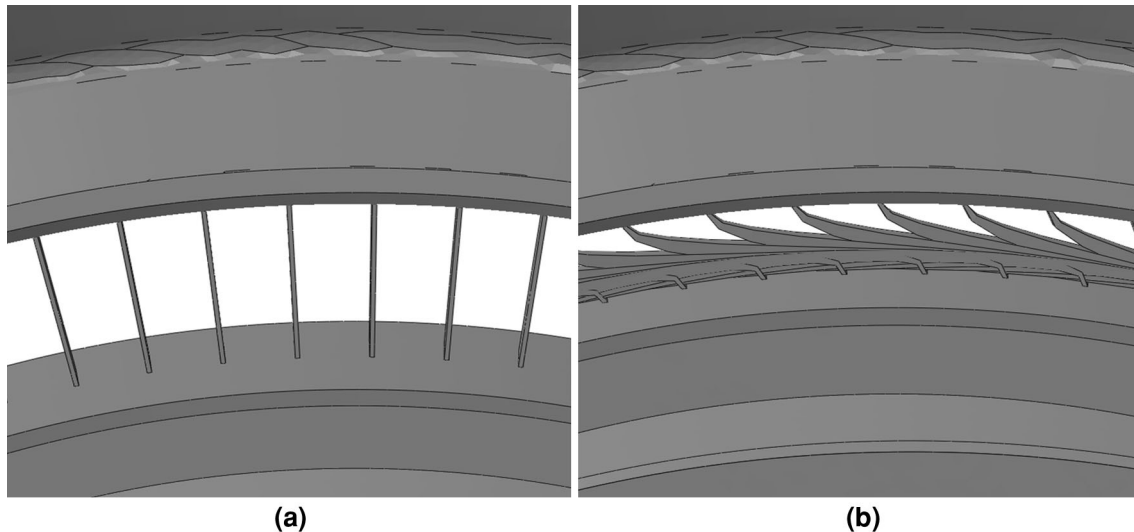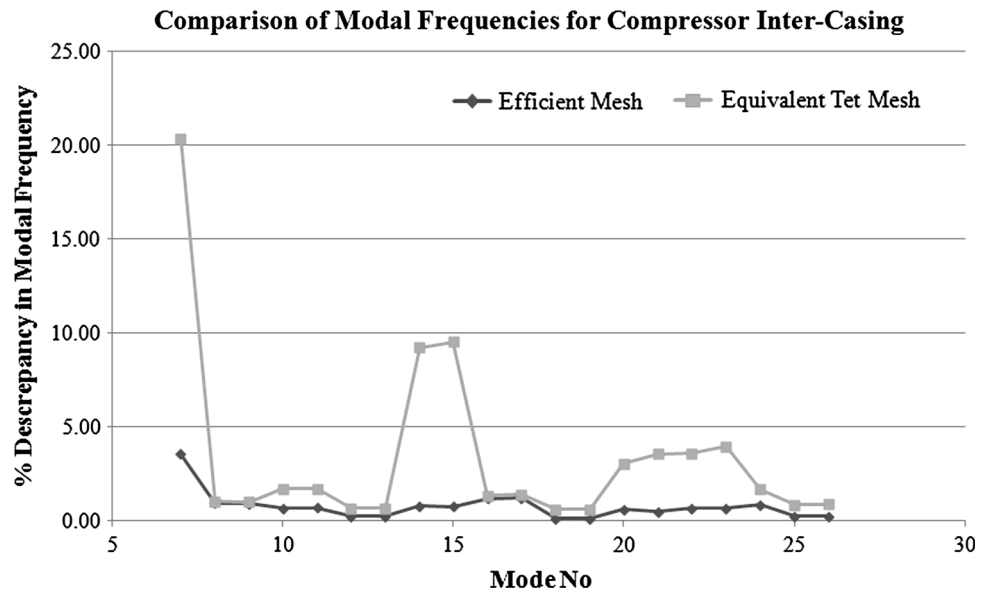**Fig. 25** Comparison of mode shapes for mode 7

Fig. 27 Deformation of the efficient mesh of the inter-casing vanes at mode 7. **a** Undeformed mesh, **b** deformed mesh

case, the deformation in the mode shape plot is normalised such that the maximum component of displacement is unity. A graph is displayed in Fig. 26 comparing the discrepancy in modal frequency of the efficient mesh and the coarse tetrahedral mesh relative to the dense tetrahedral mesh of the inter-casing model.

7 Discussion

The approach developed here uses a priori geometric reasoning about shape to subdivide a complex 3D CAD model into thin-sheet, long-slender and residual complex sub-volumes. For elliptic partial differential equations such as elasticity, the solution variation through thin sheets or along

the length of a long-slender shape is relatively simple compared to the variation in the other dimensions. Therefore, for these problems it can be assumed that elements with high aspect ratio can satisfactorily capture the solution. Where a long-slender region ends in a complex region, the elements need to become isotropic in shape in the locality of the complex region. Similarly, where a thin-sheet region is bounded by a long-slender one, the shape of the elements needs to transition from being large in the long-slender length direction only, to being large in both lateral dimensions towards the centre of the thin sheet, see Fig. 20.

The process was tested on realistic aero-engine geometry components. In general, a good correlation exists between the results of the modal analyses for the dense tetrahedral mesh and the efficient mesh of the inter-casing.

The modal frequencies correspond with the greatest discrepancy of 3.5 % recorded at mode 7. There are significant differences in modal frequencies between the reference dense tetrahedral mesh and coarse tetrahedral meshed models, particularly for modes 7, 14 and 15 with a maximum discrepancy of 20 % recorded for mode 7. On inspection of the corresponding mode shapes, it is evident that a high degree of deformation of the vanes is apparent for each mode, and in particular for mode 7, as shown in Fig. 27.

Consequently, the unusually large disparity at mode 7 and to a lesser extent at modes 14 and 15 may be attributed to poor element shape on the vanes which connect the outer and inner casings.

In terms of mode shapes, for the vast majority of modes the different casing meshes exhibit a similar shape, with the exception of mode 23 where a slightly different shape is observed for the dense tetrahedral model. There is an insignificant deviation of 0.4 and 0.1 % in mass of the efficient mesh and coarse tetrahedral meshes models, respectively, from the dense tetrahedral mesh. The centre of mass deviates by 0.2 and 0.1 mm for the efficient mesh and coarse tetrahedral mesh, respectively, which is negligible considering the diameter of the casing is approximately 800 mm. According to the statistics reported in Table 4, for an aspect ratio, m, of 4.1, there appears to be a reduction in the number of dof in the thin-sheet regions of the order of m^2 which is in agreement with Robinson et al. [1]. In the long-slender regions the "dof saving ratio" is significantly different to the m predicted by Robinson et al. [1]. In fact, there appears to be an even greater reduction in dof in the long-slender regions with over 60 % fewer dof than initially predicted using $1/m$. This may be attributed to a better mesh structure in the long-slender regions of the efficient mesh compared to the dense tetrahedral mesh. Overall, the modal analyses reveals that the efficient mesh is significantly more accurate than the coarse tetrahedral mesh with the same number of dof.

In terms of simulation analysis time, using an 8-core processor workstation with 72 GB RAM, the efficient mesh proved the most time efficient, as the modal analysis was completed in <8 min. The coarse tetrahedral mesh had a slightly longer analysis time of 11 min. As expected, the dense tetrahedral mesh model took the longest time to converge to a solution, with an analysis time in excess of 40 min.

The time taken to produce a fully decomposed model comprising long-slender, thin-sheet and complex sub-regions is largely dependent on the complexity of the source geometry. As described, the procedures detailed here make use of the CADfix thin/thick tool to identify and subdivide out the thin-sheet bodies. As the underlying technology for the thick/thin subdivision process is the 3D medial axis transform, which is computationally expensive, the vast majority of the process time for automatic partitioning is attributed to its calculation. It should be noted that any other tool which is capable of identifying and partitioning out the thin-sheet regions can be used instead. In contrast, the process time for the partitioning of the thick region is much quicker as the approach does not require a 3D medial object. For the inter-casing model, the thick/thin sub-division took approximately 3 h in comparison to the partitioning of the thick bodies into long-slender and complex regions which took approximately 10 min. Despite the relatively large processing time required to produce the full decomposition of a complex model, the approach is automatic and therefore reduces the manual effort and preparation time during the pre-processing stages of an analysis. This also reduces the skill set required to mesh the model efficiently.

The performance of the coarse tetrahedral mesh was a lot better than expected with 17 out of 20 modes exhibiting a discrepancy in modal frequency of <5 % in comparison to the reference dense tetrahedral mesh. There are other types of analysis which have not yet been explored where the discrepancy in the results could be much greater. The meshes generated using this method were efficient but it still relies on tetrahedral meshes for the complex regions. There are certain applications where hexahedral meshes are essential. For example in transient dynamic simulation applications such as "Fan Blade-Off" analysis, all areas of interest must be meshed with good-quality hexahedral. The efficient meshes generated from the automatic decomposition process presented in this paper provide a partial solution to this problem, as the majority of thin-sheet and long-slender regions have a structured hexahedral mesh applied.

Apart from the obvious advantages a fully decomposed model offers to mixed solid meshing applications, the decomposition can also be idealised into dimensionally reduced representations for automatic shell-beam meshing applications. In whole engine modelling, at the preliminary design stage, there is great demand for technologies which will enable rapid creation of inexpensive, idealised models that are cheap to run in processes requiring quick solutions for a large number of design variations. The relative lack of detail in the CAD model at this stage of the design process also means that the MO computation required for the thick/thin tool should be more robust and inexpensive, as the problematic areas which lead to difficulties in generating the MO are no longer present. For this particular application, all thin-sheet bodies can be easily replaced with a mid-surface shell mesh, and all long-slender bodies can be represented by a series of beam elements. Any complex volumes can be approximated by a point.

8 Conclusions

This paper has described an automatic approach for sub-dividing complex geometries into meshable sub-regions, to which efficient, structured meshes may be applied. It has been shown that:

- ellipsoids, sized using shape metrics based on local sizing measures, can be used to identify long-slender regions in a model,
- by applying a different meshing strategy to each region in the assembly of thin-sheet, long-slender and complex bodies, a semi-structured mesh can be generated,
- meshes produced in this way have significantly fewer dof, enhancing the efficiency of the numerical analysis,
- results of the modal analyses conducted on a realistic gas turbine engine compressor inter-casing indicate a good correlation between the efficient and dense tetrahedral meshes, thereby validating the effectiveness of the approach.

References

1. Robinson TT, Armstrong CG, Fairey R (2011) Automated mixed dimensional modelling from 2D and 3D CAD models. Finite Elem Anal Des 47:151–165. doi:10.1016/j.finel.2010.08.010
2. Loseille A, Dervieux A, Alauzet F (2010) Fully anisotropic goal-oriented mesh adaptation for 3D steady Euler equations. J Comput Phys 229:2866–2897
3. Thakur A, Banerjee AG, Gupta SK (2009) A survey of CAD model simplification techniques for physics-based simulation applications. Comput Aided Des 41:65–80
4. CADfix, TranscenData. http://www.transcendata.com/products/cadfix/. Accessed 4 July 2011
5. Blum H (1967) A transformation for extracting new descriptors of shape. In: Wathen-Dunn W (ed) Models for the perception of speech and visual form. MIT Press, Cambridge, pp 362–380
6. Luo X-J, Shephard MS, Yin L-Z, O'Bara RM, Nastasi R, Beall MW (2010) Construction of near optimal meshes for 3D curved domains with thin sections and singularities for p-version method. Eng Comput 26:215–229
7. Price MA, Armstrong CG (1995) Hexahedral mesh generation by medial surface subdivision: part I. Solids with convex edges. Int J Numer Methods Eng 38:3335–3359
8. Price MA, Armstrong CG (1997) Hexahedral mesh generation by medial surface subdivision: part II. Solids with flat and concave edges. Int J Numer Methods Eng 40:111–136
9. Li TS, McKeag RM, Armstrong CG (1995) Hexahedral meshing using midpoint subdivision and integer programming. Comput Methods Appl Mech Eng 124:171–193
10. Tam THK, Armstrong CG (1993) Finite element mesh control by integer programming. Int J Numer Methods Eng 36:2581–2605
11. Tchon K, Khachan M, Guibault F, Camarero R (2005) Three-dimensional anisotropic geometric metrics based on local domain curvature and thickness. Comput Aided Des 37:173–187
12. Shimanda K, Mori N, Kondo T (1999) Automated mesh generation for sheet metal forming simulation. Int J Vehicle Des 21:278–291
13. Frey PJ (2000) About surface meshing. In: Proceedings of 9th international meshing roundtable. New Orleans, LA, p 123–126
14. Zhao G, Zhang H (2007) Adaptive hexahedral mesh generation based on local domain curvature and thickness using a modified grid-based method. J Finite Elem Anal Des 43:691–704
15. Zhao G, Zhang H, Cheng L (2008) Geometry-adaptive generation algorithm and boundary match method for initial hexahedral element mesh. J Eng Comput 24:321–339
16. White DR, Saigal S, Owen SJ (2004) Sweep: automatic decomposition of multi-sweep volumes. J Eng Comput 20:222–236
17. Frey PJ, George PL (2008) Mesh generation. Wiley, London, p 337
18. Quadros WR, Shimada K, Owen SJ (2004) Skeleton-based computational method for generation of 3D finite element mesh sizing function. Eng Comput 20(3):249–264
19. Quadros WR, Owen SJ, Brewer M, Shimada K (2004) Finite element mesh sizing function for surfaces using skeleton. In: Proceedings of 13th international meshing roundtable. Williamsburg, VA, p 389–400
20. Lipschutz MM (1969) Schaums outline of theory and problems of differential geometry. McGraw-Hill, New York, p 176
21. Quadros WR, Vyas V, Brewer M, Owen SJ, Shimada K (2010) A computational framework for generating mesh sizing function in assembly meshing via disconnected skeletons. Eng Comput. Published online: Jan 12, 2010
22. CRESCENDO, Collaborative and robust engineering using simulation capability enabling next design optimisation. http://www.crescendo-fp7.eu. Accessed 13 Mar 2012
23. Abaqus/CAE, Simulia. http://www.3ds.com/products/simulia/portfolio/abaqus. Accessed 13 Mar 2012