Available online at www.sciencedirect.com
ScienceDirect
Journal of Computational Physics 227 (2008) 2781–2793
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Local remeshing for large amplitude grid deformations

Keri R. Moyle *, Yiannis Ventikos

Department of Engineering Science and Institute of Biomedical Engineering, University of Oxford, Parks Road, Oxford OX1 3PJ, UK

Received 25 April 2007; received in revised form 6 November 2007; accepted 9 November 2007
Available online 22 November 2007

Abstract

Fluid-structure interaction (FSI) modelling can involve large deformations in the fluid domain, which could lead to degenerating mesh quality and numerical inaccuracies or instabilities, if allowed to amplify unchecked. Complete remeshing of the entire domain during the solution process is computationally expensive, and can require interpolation of solution variables between meshes. As an alternative, we investigate a local remeshing algorithm, with two emphases: (a) the identification and remedy of flat, degenerate tetrahedra, and (b) the avoidance of node motion, and hence associated interpolation errors.

Initially, possible topological changes are examined using a dynamic programming algorithm to maximise the minimum local element quality through edge reconnection. In the 3D situation it was found that reconnection improvements tend to be limited to long edges, and those with few (three or four) element neighbours. The remaining degenerate elements are classified into one of four types using three proposed metrics – the minimum edge-to-edge distance (EE), the minimum node-to-edge distance (NE), and the shortest edge length (SE) – and removed according to the best manner for their type. Optimised thresholds for identifying and classifying elements for removal were found to be $EE < 0.18$, $NE < 0.21$, $SE < 0.2$.
© 2007 Elsevier Inc. All rights reserved.

MSC: 65D18; 74F10; 76M12; 92C35; 74L15

Keywords: Dynamic meshing; Mesh quality; Unstructured mesh; Tetrahedron; Aortic dissection

1. Introduction

The use of unstructured tetrahedral meshes in computational fluid dynamic (CFD) simulations allows easy representation of arbitrary domain shapes [1], making these meshes attractive when modelling biological flows. Tetrahedral meshes are also common in applications involving gross flow field deformations, especially where the motion is not predictable a priori. The ability of tetrahedral elements to provide a geometric representation does not automatically correlate to an arrangement acceptable for a computational simulation, and mesh quality issues involving the size [2] and shape [3] of the elements used are an important aspect of

* Corresponding author. Tel.: +44 1865 283 452; fax: +44 1865 273 010.
E-mail address: keri.moyle@eng.ox.ac.uk (K.R. Moyle).

0021-9991/$ - see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2007.11.015

CFD. Deterioration of quality is seen in locations where boundary motion occurs, reinforcing the importance of resolving numerical accuracy in such regions, and also highlighting one of the challenges that modelling these particular processes involves.

The motivation for this study stems from an attempt to model the fluid and solid dynamics of an aortic dissection: a situation involving pulsatile fluid flow and the free motion of a thin membrane flap, both within a tortuous and compliant vessel [4]. The large possible degree of motion of the flap effectively describes a grossly deforming fluid region, at the same time as the solver requires congruent faces and nodes between the fluid and solid domains.

The quality of an element is usually measured by defining an optimum shape, and measuring – in some direction – the departure of the element from that shape [3,5,6]. Hence variations in the idea of a 'good' shape and the rate at which quality deteriorates with shape give rise to different mesh quality measures. Most traditional mesh quality measures agree on what is good, and the equilateral tetrahedron is generally regarded as the ideal shape, as it maximises the minimum internal angle (which is used indirectly in calculating fluid fluxes).

Mesh quality measures are useful for locating degenerate elements within a mesh, but cannot generally identify the features within each element that make it so. It is these features within the degenerate element that must be fixed, and therefore must also be recognisable. In the meshing situation driving the current study, a technique was required that would not only identify 'bad' elements, but also propose the best route towards their repair or removal. For this to be possible, the measure must be able to differentiate between different types of poor element shapes, even when their mesh quality measures are equal.

Fig. 1 shows isosurfaces of mesh quality for four common shape measures: the shape factor (SF), the normalised shape ratio (NSR), QS [5], and the equi-angle skew (EAS). Also shown are the edge-to-edge distance (EE) and the node-to-edge distance (NE) which will be introduced and discussed later. Definitions of these measures given in Eqs. (1)–(6), and further shape measures for tetrahedra may be found in the works of Field [5], and Knupp [6],

$$\mathrm{SF} = \mathrm{sign}(V)\frac{12(9V^2)^{1/3}}{\sum_{i=1}^{6} l_i^2} \tag{1}$$

$$\mathrm{NSR} = \frac{3r}{R} \tag{2}$$

$$\mathrm{QS} = \frac{V}{\sum_{i=1}^{6} l_i^2} \tag{3}$$

where V is the tetrahedron volume, r is the inradius, R is the circumradius, and l is the edge length.

$$\mathrm{EAS} = 1.0 - \max\left\{\frac{\theta_{\max} - \theta_{\mathrm{ideal}}}{\pi - \theta_{\mathrm{ideal}}}, \frac{\theta_{\mathrm{ideal}} - \theta_{\min}}{\theta_{\mathrm{ideal}}}\right\} \tag{4}$$

where θ_{ideal} for a tetrahedron is 1.2312 or 70.54°, and $\theta_{\max}$ and $\theta_{\min}$ are the maximum and minimum angles between face planes.

The plots in the figure have been generated using an equilateral triangle as the base of a tetrahedron, perturbing the apex from the equilateral (in this case, 'perfect') location, then calculating and plotting the quality according to each measure using isosurfaces. From the figure it can be seen that the near-spherical isosurfaces of QS and SF are in fairly close agreement with one another, and likewise the EAS and NSR as regards increases in the height of the tetrahedron. However, as the element is flattened and the apex moves towards the basal plane, none of these measures is able to discern whether, for example, the apex node now lies near one of the other edges (see Fig. 2(b), far from any edge (see Fig. 2(a), within the basal triangle (see Fig. 2(c), or nearby to one of the other nodes (see Fig. 2(d). While each of these shapes is degenerate and requires removal, the manner of removal that is optimal – or even possible – differs between cases, and depends on tetrahedron shape. In order to identify these differences a further quantity is needed: this quantity must be able to be evaluated in the presence of degenerate elements, and must also provide information on the best manner of repair of the element. In the following section we define four types of flat tetrahedral elements, and describe the method of their repair.

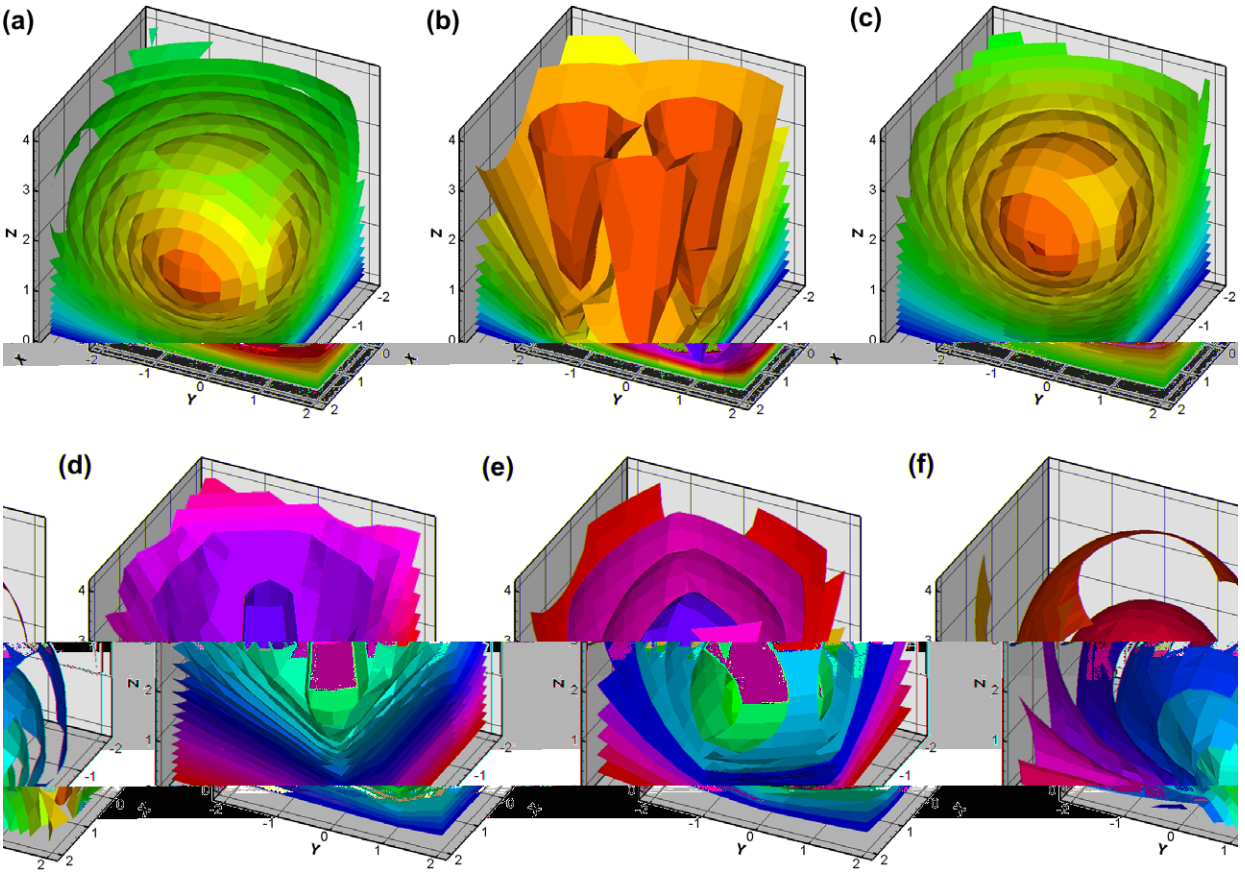

Fig. 1. Isosurfaces of quality measures for a tetrahedron with a fixed equilateral triangle base, with vertices at (200), $(1\sqrt{}20)$, (000) and a variable apex position. The measures plotted are (a) SF, (b) NSR, (c) QS, (d) EAS, (e) EE, and (f) NE.

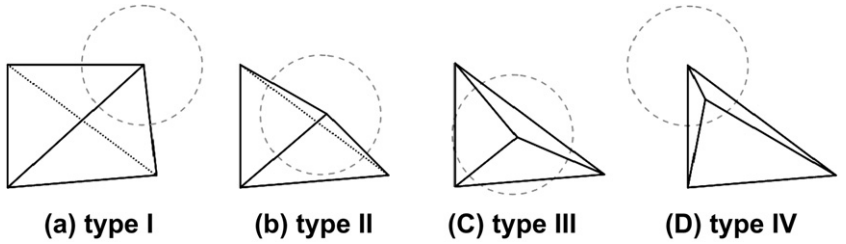

Fig. 2. Examples of degenerate tetrahedra and their type classification, shown with a threshold edge length line dotted around selected nodes. Note that Type IV is the only one in which a node falls within another's circle, so is the only type eligible for removal via edge splitting or edge collapse methods.

A survey of existing shape measures shows that the majority involve the circumradius (R), the inradius (r), and/or tetrahedron volume (V) in some combination. As the apex of any tetrahedron moves nearer to the plane of the opposite face, these terms become undefined and render the measure unable to distinguish between any of the degenerating tetrahedra shown in Fig. 2. Many existing techniques for remeshing [7–13], alter the mesh topology by comparing the local edge length to a predefined desired edge length, and splitting (or collapsing) edges that are too long (or too short). Mesh resolution (as measured by edge length) is then able to be locally maintained despite gross and ongoing deformations of the geometry. With reference to the elements in Fig. 2(a)–(c), we see that the degenerate elements do not meet the criterion for either edge splitting

or collapse, as their edge lengths are fairly uniform. Elements such as these could perhaps be remedied through reconnection. Other techniques, such as Delaunay meshing, can identify nodal configurations that do not conform to the Delaunay criteria and fix this aspect, but in 3D are prone to the formation of sliver elements [1], and thus the Delaunay criteria does not provide adequate control over mesh quality, or inherent correlation with numerically 'good' meshes.

Note that throughout this article the following connotations are used: topology refers to the connectivity relationship between groups of nodes, to form a group of elements; geometry refers to the 3D spatial location of nodes, faces and surfaces. Thus, stretching a mesh changes its geometry but not its topology, and for a given geometry there are many topological arrangements possible.

2. Methods

2.1. Mesh quality measures

In general, we seek a strategy to measure, assess and remedy meshes involving flat and degenerate tetrahedra, and more specifically we wish to avoid nodal displacement if possible. The movement of nodes outside of an Arbitrary Lagrangian–Eulerian (ALE) [14] solver introduces errors during interpolation of solution variables between the old and new meshes. Laplacian or spring-based smoothing may improve the local mesh quality, but these improvements must be weighed against the subsequent interpolation errors incurred.

We propose the following classification strategy, where typical examples of each class are shown in Fig. 2. Four types of degenerate tetrahedra are identified, and their classification can be explained in terms of a fixed basal triangle and a variable apex, positioned near to, or on, the basal plane. The element orientation and shape of the basal triangle are obviously arbitrary, but in practice we have found this classification sufficiently broad. We start by defining some simple geometric parameters for each element: the edge-to-edge distance (EE), being the minimum distance between any two points on an unconnected edge pair, and the node-to-edge distance (NE), being the minimum distance between any node and any edge unconnected to it. For two unconnected edges, we can calculate the minimum edge-to-edge length as

$$\mathrm{EE} = \min \left\{ \frac{|s \cdot (e_1 \times e_2)|}{|e_1 \times e_2|} \right\}_{\text{opposite edge pairs}} \tag{5}$$

where the opposite edges are e_1 and e_2, and s is a line between the ends of each edge. For any line and node that together define a face, the minimum node-to-edge distance is

$$\mathrm{NE} = \min \left[\frac{|e \times (e_0 - x)|}{|e|} \right]_{\text{edges, nodes}} \tag{6}$$

where the edge is given by e, the position of the opposite node is given by x, and e_0 is one end of the edge.

Type I elements are characterised by a small EE formed by the crossing diagonal edges, a quadrilateral footprint regardless of the choice of basal face, and a NE value near to the edge length. As the floating apex node is moved nearer to the basal triangle, the NE value decreases. When NE is less than some fraction of the average edge length, the element is labelled a Type II. Moving the apex still nearer produces an element whose footprint changes from a quadrilateral (as in the Type I and II elements) to the triangular footprint of Type III. The NE values of Type III elements are typically small, and because they have no crossed diagonal edges the EE value becomes meaningless. When the apex lies close to one of the other nodes, a Type IV element is formed. Although this element is degenerate, the short edge length meets the criteria for collapse, and would be routinely removed. Type IV elements do not represent the unresolved degenerate element shape of the other types, and are only included here for completeness.

2.2. Mesh reconnection routines

The most expensive component of mesh regeneration involves decisions about topology. Even for a small group of nodes, the number of possible connectivity outcomes is typically too large to exhaust iteratively. Fortunately, the topology changes can be broken down into subsets of simpler operations. For tetrahedral

reconnection, Shewchuk noted that for the elements surrounding a single edge, the viable topological changes can be reduced to a series of 2–3 flips, followed by a single 3–2 flip; each of these is defined below, and illustrated in Fig. 3. A 2–3 flip is so called because it changes the connective topology of two elements to form three; a 3–2 flip is the reverse process. Such decomposition provides a means for generalising an optimisation code, so that dynamic programming can be used to efficiently find the optimum configuration. [J.R. Shewchuk, Two discrete optimization algorithms for the topological improvement of tetrahedral meshes, unpublished manuscript, available from <http://www.cs.cmu.edu/~jrs/jrspapers.html>].

There are times when even optimal reconnection cannot sufficiently repair a mesh, and nodal density and/or nodal positions need to change. Two operations accomplish this: edge splitting (the addition of a node to the midpoint of an edge, effectively doubling the number of elements in the original edge's kernel), and edge collapse (the deletion of a node through merging with a partner along an edge). These are illustrated in Figs. 4 and 5, and form the building blocks of the second part of the mesh improvement algorithm.

Although the operations described are general enough that they could – with some caveats – be applied throughout a domain, we have chosen to prevent mesh alterations at boundaries and interfaces. Many software packages involving FSI modelling require a congruent boundary between the faces and nodes at the interface of the solid and fluid domains. Alterations of connectivity or nodal position in either domain would violate this criterion, and are therefore forbidden.

2.3. Specialised reconnection routines

In a similar manner to the formulation of reconnection strategies from simpler operations, combinations of edge collapse and splitting can be used to give operations that are optimised for the repair of specific types of

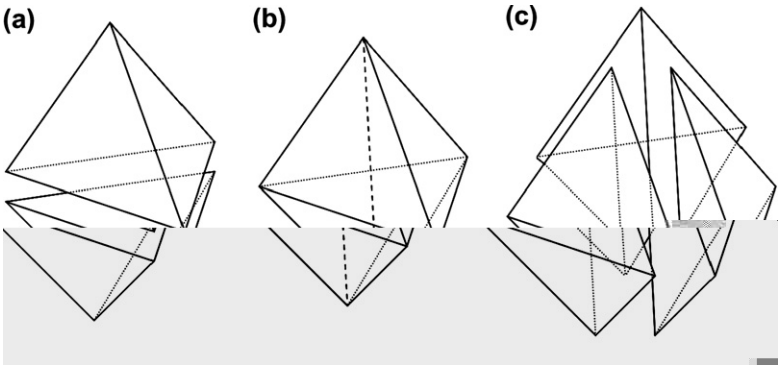

Fig. 3. The 2–3 reconnection operation, showing (a) the starting two elements, (b) the construction of a new edge between opposite vertices, and (c) the final three element connectivity. The 3–2 flip is the reverse process.

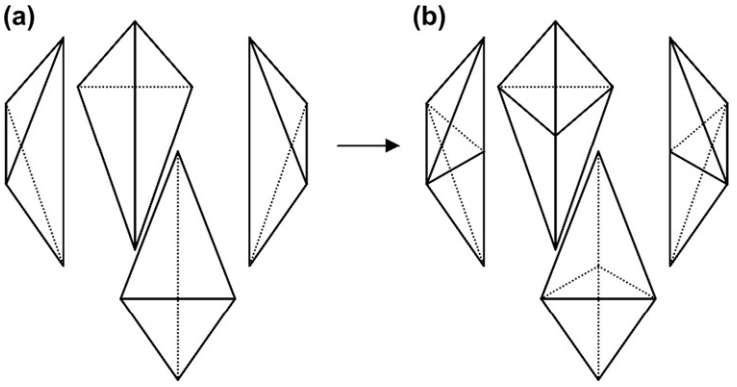

Fig. 4. Edge splitting operation, showing (a) the original elements with a long edge to be split, and (b) a new node added at the midpoint of the split edge, increasing the element count by the number of original element neighbours of the split edge.

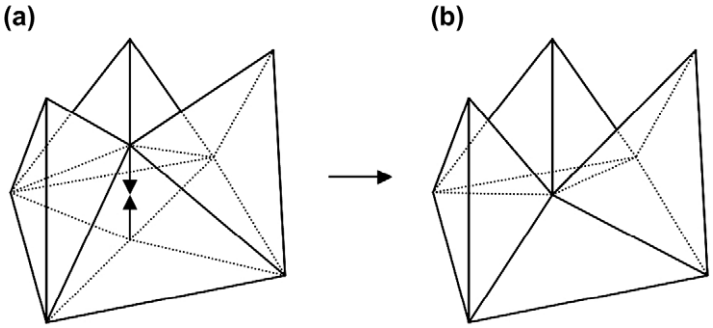

Fig. 5. Edge collapse operation, showing (a) the original short edge (arrowed) and its surrounding elements, and (b) the final connectivity when the short edge has been collapsed, removing one node and all elements sharing the collapsed edge.

Table 1
Typical properties of the types of degenerate tetrahedra

Classification	Type I	Type II	Type III	Type IV
Minimum edge-to-edge distance	Lim → 0	Lim → 0	(No crossed edges)	Small
Minimum node-to-edge distance	Large	Small	Small	Small
Minimum node-to-face distance	n/a	n/a	Lim → 0	n/a

element configurations. After using the element classification to identify the type of element, the appropriate combination can then be prescribed.

Table 1 summarises the typical parameters of each type of degenerate element. From the table and the illustrations in Fig. 2, it can be seen that some elements could fall into more than one category; having, for example, a small NE (indicating Type II) but a triangular footprint (indicating Type III). In practice, the priority order of classification is IV, II, III, I; the reasons becoming clear when the manner of removal of each is examined.

Type I tetrahedra are characterised by a quadrilateral footprint, and have crossed diagonals which approach intersection as the element approaches planarity. These elements are not usually removed by reconnection, as the tests for reconnection involve a single edge, not the two edges found in this type of element. Fig. 6 shows the repair of a Type I element. Both diagonal edges are broken – either at their midpoint, or

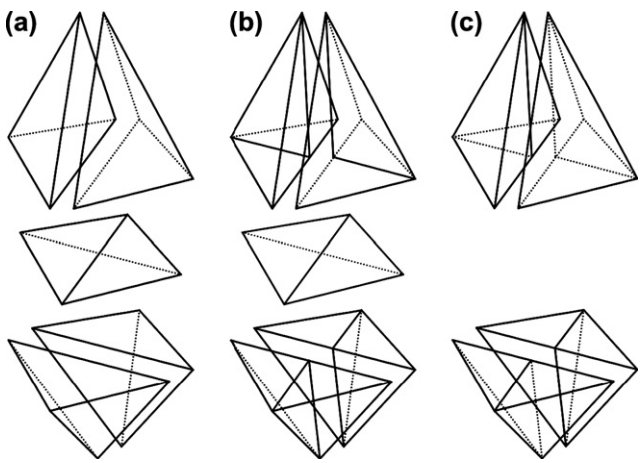

Fig. 6. Repair of a Type I element, showing (a) the original configuration, with the degenerate element shown in the middle, (b) the diagonal edges of the degenerate element are split, creating two transient points, and (c) the edge between the two transient points formed in (b) is collapsed, removing all children of the original degenerate element.

at the location nearest to the other edge – forming two new nodes. The edge connecting these two nodes is then collapsed, leaving only one new node in approximately the centre of the collection of affected elements. The accuracy of any interpolation onto the newly created node can be controlled by constraining the final node position to coincide with the midpoint of one of the original diagonal edges. Hence the interpolation is along an edge to its midpoint, rather than to an arbitrary point in space. The decision to improve interpolation must be weighed against the possible penalty incurred in the resultant mesh quality by breaking at the midpoints as opposed to the intersection point. In practice, we found that the better mesh quality occurred most often by breaking both edges at their midpoint, rather than the point nearest intersection.

The configuration (and hence the remedy) of Type II elements is similar to Type I. If an element with a small NE were broken (as for a Type I) at the intersection of its diagonals, the resultant elements would have faces with unacceptably small angles (shown in Fig. 7), producing child elements of low quality. Instead, the Type II element is repaired by breaking the edge involved in the smallest NE measurement to form a new transient point, and then collapsing the edge between the transient point and the node opposite. After collapse, the new location of the node becomes that of the transient point, and lies along a previously existing edge.

As mentioned earlier, there is a certain degree of overlap between the classifications. This is especially true in Type III elements, which tend to fall into shared categories of Type IV (having one short edge), or Type II (having a small NE distance). Once again, the strategy employed must account for the final configuration after repair. Because the shape of the faces other than the basal face is generally poor, the best general option for an uniquely classified Type III element is to locate the shortest edge connected to the apex, and collapse it towards the basal node position; effectively removing all elements sharing the poor faces, and leaving only the basal face remaining in the mesh (see Fig. 8). This is the method implemented during this study, and was found to be sufficient in our test cases. Finally, those elements classified as Type IV are removed by the simple collapse of their shortest edge (shown earlier in Fig. 5), with the remaining node placed at the midpoint of the old edge.

2.4. Mesh motion routines

Most FSI solvers require a face-to-face matching between the solid and fluid mesh domains. In the applications considered here, the fluid domain undergoes significantly more deformation than the solid domain, and so the solid domain topology is held constant and only the fluid domain is remeshed. Initially, the fluid mesh was deformed within the solution timestep by a pseudo-elastic solver, which treats the fluid domain as a

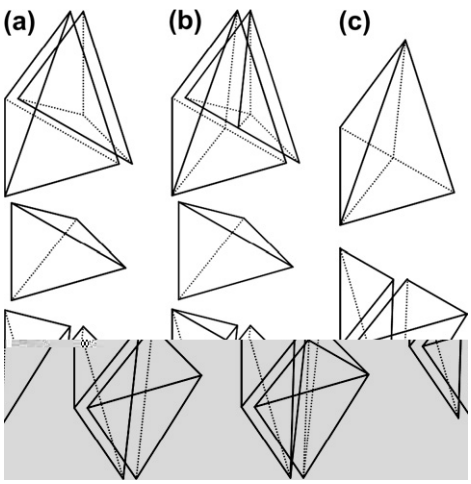

Fig. 7. Repair of a Type II element, showing (a) the original configuration, with the degenerate element in the middle, (b) splitting the edge nearest to an opposite node, and (c) the short edge formed by the split in (b) between the new point and the nearest node is collapsed, removing all children of the degenerate middle element, as well as children of others sharing the face formed between the original edge and nearest node.

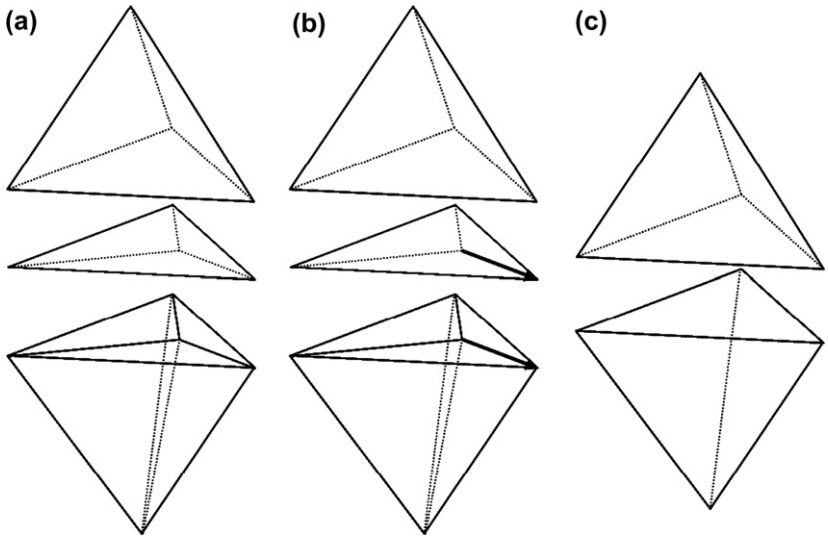

Fig. 8. Repair of a Type III element, showing (a) the original connectivity with the flat element in the middle, (b) the collapse of the shortest edge on the flat element connected to the apex (arrowed), and (c) the final elements formed.

hyper-elastic solid with Poisson's ratio of zero, and solves for the mesh deformation. As expected, the fluid regions nearest to the deforming solid undergo the most deformation, and usually yield the worst final mesh quality. To circumvent this problem, we have imposed restrictions on the system whereby a defined region of mesh near to any boundary is held rigid during mesh motion, and therefore undergoes minimal distortion. Using this method ensures that the elements near to the boundaries (where remeshing is forbidden, and quality is most critical) maintain a high quality, and those elements that undergo greatest deformation are moved into the volume mesh, where they may be operated upon by the reconnection and remeshing algorithms described earlier.

3. Results and discussion

3.1. Test cases

Initially, simple test cases were used to assess the viability of the dynamic remeshing algorithm. A FSI simulation of two rectangular channels separated by a flexible membrane was run using the pseudo-elasticity routine for grid deformation until it failed owing to poor element quality. Because that routine only alters nodal position, the 'failed' and 'original' meshes provided the extrema of quality for the given topology. Meshes of intermediate quality were then artificially created by linearly interpolating node positions between the 'failed' and 'original' meshes, thus providing a range of initial quality in these test cases.

The reconnection algorithm was set to maximise EAS, and the mesh was swept by inspecting only those elements below a starting threshold, then successively raising the threshold until no further improvements were possible. Elements with quality lower than the threshold were plotted before and after reconnection at each increment, and shown for the stretched membrane test case (Fig. 9). During compression, reconnection removed 100%, 93%, 73% and 37% of elements, and during stretching reconnection removed 100%, 95%, 90% and 60% of elements below the EAS thresholds of 0.1, 0.2, 0.3, 0.4. Obviously as the threshold is raised, the number of elements falling below it increases, and the possibility for improvement by reconnection alone diminishes.

The reconnection algorithm alters mesh topology only when the minimum element quality increased. Thus, although no degradation in minimum quality is seen, negligible improvements may have induced a totally new local connectivity; and this new configuration may have lowered an individual element quality in order to increase the local minimum. Fig. 10 shows the effect of reconnection on local mesh quality, from which the

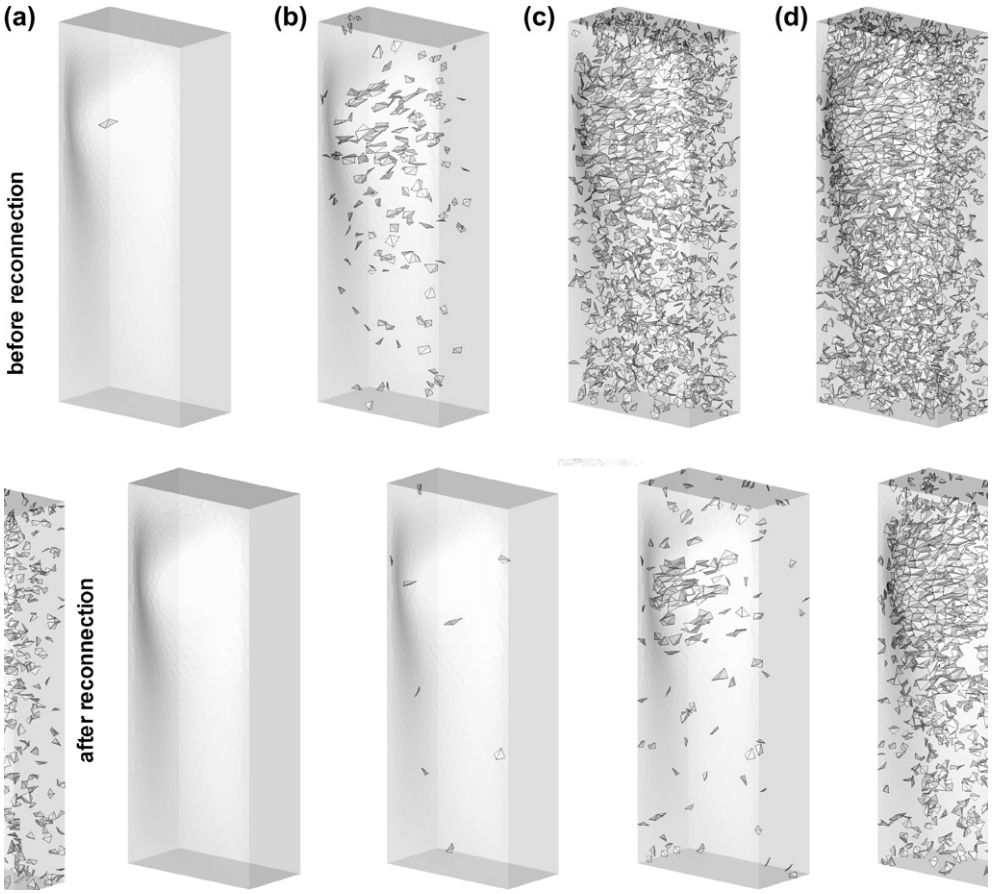

Fig. 9. Effect of reconnection on the element quality, applied to the stretched test case 1. The upper row is before reconnection, the lower row is afterwards. From left to right (a) EAS < 0.1, 100% (1/1) removed, (b) EAS < 0.2, 95% (170/179) removed (c) EAS < 0.3, 90% (1600/1781) removed, and (d) EAS < 0.4, 60% (2176/3635) removed.

effect of applying a minimum improvement condition on the algorithm can be deduced. It can be seen that local reconnection can result in significant improvement, but also that much of the improvement was below $\Delta EAS = 0.1$. Comparing the number of edges meeting the criteria for reconnection to the improvement given, we found that setting a minimum improvement threshold of $\Delta EAS = 0.2$ would halve the number of reconnections performed. Though the incorporation of such a threshold into the routine could reduce overall computational time, the most computationally intensive part of the routine is in the calculation of mesh quality itself, and hence the efficiency improvements are more likely to scale with the number of possible reconnections tested, rather than the total number performed.

It is possible to estimate the typical starting topology for which reconnection can give the greatest improvement from Fig. 10. The algorithm tested reconnections of edges with up to eight neighbouring elements, and found that improvements tended to occur more frequently on the edges with fewer (three or four) neighbours. However, this observation needs to include a caveat regarding element shape: when a group of elements is stretched along its common edge axis, the advantages of removing that axis via reconnection increase, and the opposite as an edge is compressed. The reconnection routine can be thought of as a method of removing the central edge, and replacing it with edges diagonal to its satellite nodes. In this case, the advantages of reconnection become apparent when the edge to be removed is too long, and the diagonals it forms are nearer to the required length. Obviously, when the central edge is too short, reconnection about that edge cannot help the problem, and a new central edge must be chosen. Heuristically restricting the reconnection search to edges more likely to provide improvement (for example, edges with three or four neighbours only) could improve

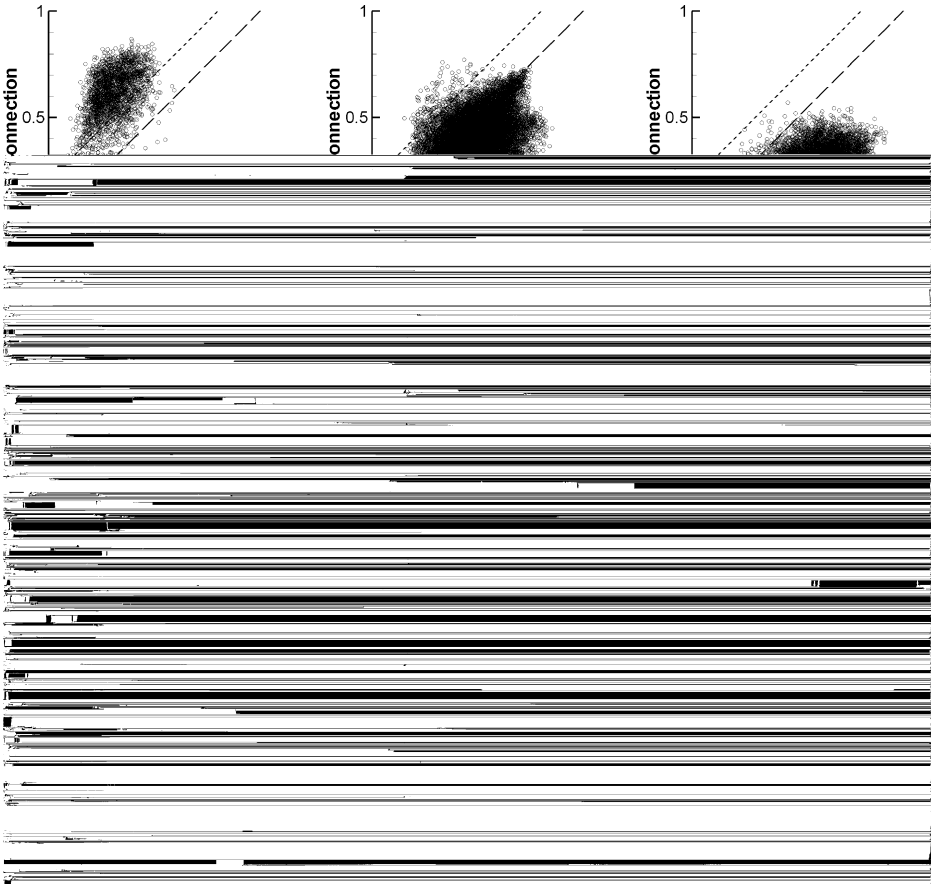

Fig. 10. The effect of reconnection on mesh quality, as influenced by the number of original element neighbours of the reconnected edge. The solid horizontal line represents the lower limit of acceptable mesh quality, points above the dashed line have improved in quality, and above the dotted line have a quality improvement greater than 0.2.

computational efficiency significantly, as it avoids the expensive tests for determination of reconnected element quality. This must be weighed against the need for the remaining mesh elements to be altered via subsequent operations instead.

The observation made above that three-neighboured edges give the best (that is, the highest and most consistent) improvement upon reconnection can be explained as follows. In order for three elements to completely enclose an edge, the angle that is formed between faces on that edge must be obtuse. Hence, the 3–2 flip on that edge replaces the three obtuse tetrahedra with two that must be (see Fig. 3) acute, can dramatically improve the EAS measure for that element kernel. In the 2–3 and 3–2 flips performed on a four-(or more)-neighboured edge, the same reasoning does not necessarily apply, and locally specific improvements (that depend on node placement, or edge length) are manifest instead.

To assess the applicability of the routine on a simple test case, a mesh of a translating sphere in a square duct was created, and the mesh deformed. The algorithm was able to remedy the mesh as the sphere moved, giving a mesh free from degenerate elements and of a quality acceptable for fluid solution. These results are shown in Fig. 11. If this mesh had been used for a FSI simulation, the degenerating mesh quality (from as early as the first timestep) would have caused unacceptably high levels of numerical error in the solution, though the solver may not necessarily stop. From the figure it can be seen that the areas of lowest quality are near to the fluid-structure interface, highlighting the particular importance of mesh quality checks in fluid-structure or moving mesh applications, as the area of greatest interest is also the most prone to mesh induced errors. In this example the timestep is greater than that which would be used for a FSI simulation, but the fact that degenerate elements appear almost immediately indicates that total domain remeshing (as opposed to local)

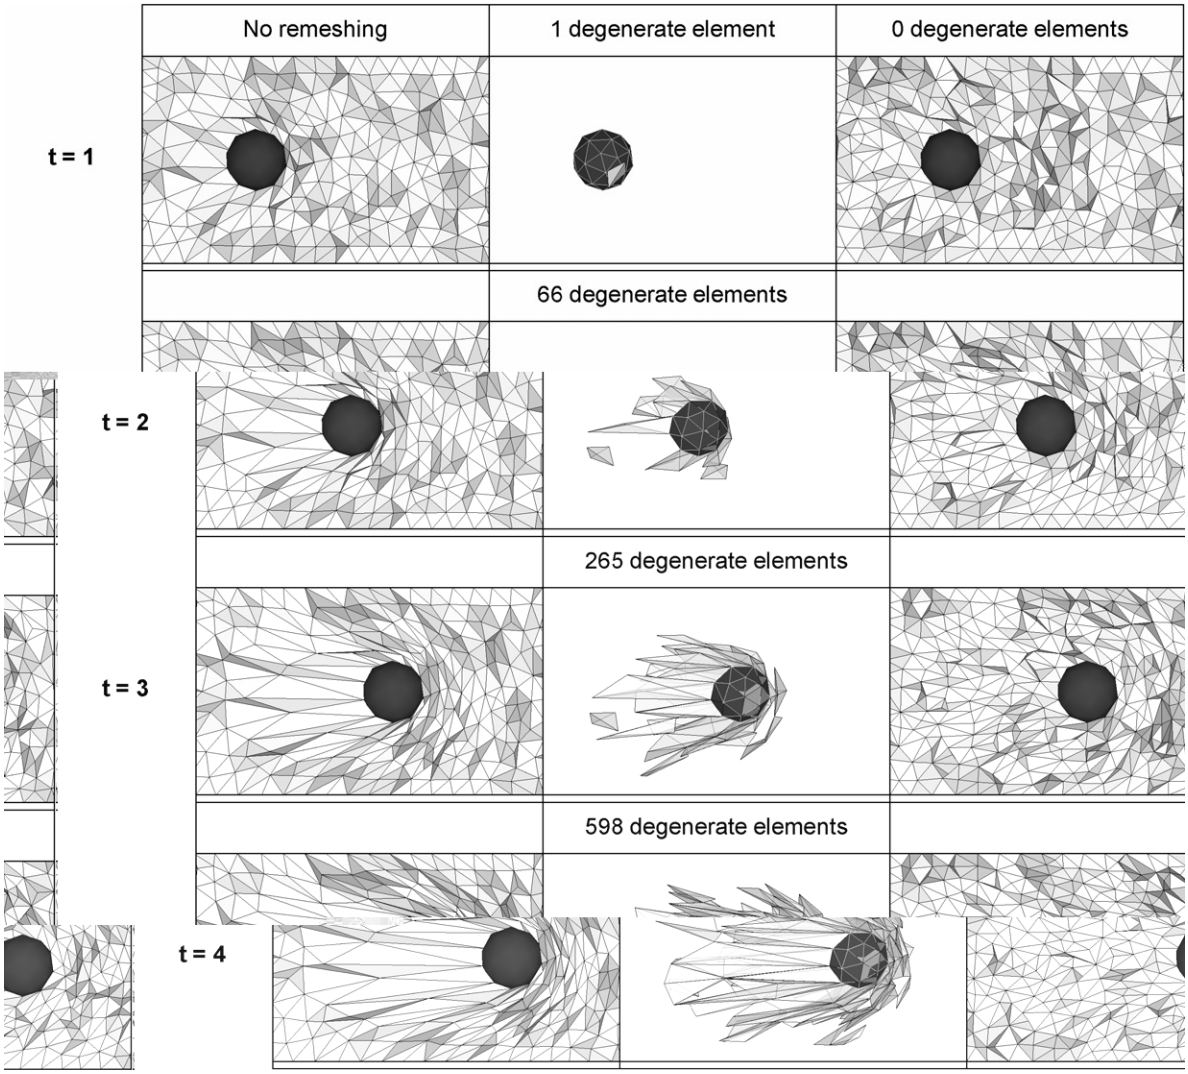

Fig. 11. Remeshing results for a translating sphere in a square duct, showing (a) no remeshing, (b) degenerate elements without remeshing, and (c) including remeshing, no degenerate elements.

would significantly increase the computational cost of the simulation, quite apart from the old-to-new solution interpolation issues that arise.

3.2. Degenerate element removal

To most effectively classify elements for removal, the interaction between their defining parameters and the removal method needs to be more fully understood. In contrast to traditional smoothing methods where a badly shaped element remains (albeit with higher quality) in the mesh, the complete removal of a degenerate element means that final mesh quality depends on the surrounding elements, rather than on the degenerate element itself. It makes sense that the limiting parameters set should be those that maximise the final mesh quality, and are (perhaps) independent of the element to be removed.

To investigate the relationship of the control parameters (NE, EE, SE and minimum shape quality), we performed a test using an unstructured tetrahedral mesh of a cube, which was distorted to create three lower quality meshes with distinctive element types. The distortions used were extension ($z' = 3z$), compression ($z' = z/3$),

Table 2
Element removal type giving optimal local mesh quality for distorted cube meshes

Element type	I	II	IV
All meshes	243	12	77
Compressed	75	1	0
Stretched	72	1	33
Sheared	96	10	44

The original mesh had 6797 tetrahedra and 1576 nodes.

and shearing ($z' = 1.5y$). Each element with EAS < 0.2 was removed by one of the three methods listed in Table 1 in turn, regardless of its true classification and the possible reduction in final mesh quality. The results were plotted against the control parameters to find a relationship between the optimal final mesh, the original element type, and the parameters that correctly define that type. The numbers of degenerate elements that were successfully (that is, in a manner that benefited the final mesh) removed by each method are shown in Table 2. It would seem that the removal method for Type I is the most robust, giving the greatest consistency in improvement between all the removal methods.

A weighting system was used to optimise the EE, NE and SE thresholds used. The ideal quality is given by the operation that was found to increase the local mesh quality the most. In some cases, removing the element resulted in a decrease in mesh quality; in these situations the ideal quality became the uncorrected mesh quality. Weighting factors equal to the difference between this measured ideal quality and the final quality following element removal (as classified by the thresholds) were applied, and the sum of these penalty weights was minimised. Elements that were correctly identified and removed incurred no penalty. This minimisation procedure gave optimised thresholds of

$$\mathrm{EE} < 0.18, \text{Type I}$$
$$\mathrm{NE} < 0.21, \text{Type II}$$
$$\mathrm{EAS} < 0.2, \text{NF real}, \text{Type III}$$
$$\mathrm{SE} < 0.2, \text{Type IV}$$

for a total of 525 elements removed from the three box meshes. These thresholds when applied to all test cases successfully removed all degenerate elements.

4. Conclusions

We have presented two metrics for mesh quality (the minimum edge to edge distance, EE, and the minimum node to edge distance, NE) and demonstrated their use in the classification and removal of degenerate tetrahedral elements. These metrics are needed because traditional quality measures involve properties (such as the circumradius, the inradius, and the volume) that become undefined as the element approaches planarity. Four types of degenerate tetrahedral configurations were defined, and their remedy described in terms of combinations of the simple edge splitting and edge collapse operations found in other existing algorithms. Elements are removed when their minimum edge-to-edge distance EE < 0.18 of the average edge length, when their node-to-edge distance NE < 0.21 of the average edge length, and when the shortest edge length SE < 0.2 of the average edge length.

We have also implemented a dynamic programming algorithm to investigate reconnection, showing that reconnection can be used to give dramatic increases in element quality without node repositioning. Edges with three element neighbours gave the most dramatic improvement, which decreased for edges with more element neighbours. No improvement through reconnection was seen for edges with more than five neighbours. In areas where reconnection improvement was not possible, edge splitting or edge collapsing routines were used to provide the correct node density, again without node motion. Finally, degenerate tetrahedra were successfully removed according to the criteria above, to provide meshes that have an acceptably high quality while minimising node motion.

Acknowledgment

KRM is supported by EPSRC Grant No. EP/C526309/1.

References

[1] D.J. Mavriplis, Unstructured grid techniques, Annu. Rev. Fluid Mech. 29 (1997) 473–514.
[2] P.J. Roache, Quantification of uncertainty in computational fluid mechanics, Annu. Rev. Fluid Mech. 29 (1997) 123–160.
[3] M. Berzins, Mesh quality: a function of geometry, error estimates or both? Eng. Comput. 15 (1999) 236–247.
[4] D. Mukherjee, K.A. Eagle, Aortic dissection – an update, Curr. Prob. Cardiol. 30 (2005) 287–325.
[5] D.A. Field, Qualitative measures for initial meshes, Int. J. Numer. Meth. Eng. 47 (2000) 887–906.
[6] P.M. Knupp, Algebraic mesh quality metrics for unstructured initial meshes, Finite Elem. Anal. Des. 39 (2003) 217–241.
[7] A. Anderson, X. Zheng, V. Cristini, Adaptive unstructured volume remeshing – I: the method, J. Comp. Phys. 208 (2005) 616–625.
[8] M. Dai, D.P. Schmidt, Adaptive tetrahedral meshing in free-surface flow, J. Comp. Phys. 208 (2005) 228–252.
[9] X. Zheng, J. Lowengrub, A. Anderson, V. Cristini, Adaptive unstructured volume remeshing – II: application to two- and three-dimensional level-set simulations of multiphase flow, J. Comp. Phys. 208 (2005) 626–650.
[10] Y. Zhao, A. Forhad, A general method for simulation of fluid flows with moving and compliant boundaries on unstructured grids, Comput. Meth. Appl. Mech. Eng. 192 (2003) 4439–4466.
[11] P. Le Tallec, J. Mouro, Fluid structure interaction with large structural displacements, Comput. Meth. Appl. Mech. Eng. 190 (2001) 3039–3067.
[12] A.K. Slone, K. Pericleous, C. Bailey, M. Cross, Dynamic fluid-structure interaction using finite volume unstructured mesh procedures, Comput. Struct. 80 (2002) 371–390.
[13] K. Nakahashi, Y. Ito, F. Togashi, Some challenges of realistic flow simulations by unstructured grid CFD, Int. J. Numer. Meth. Fluids 43 (2003) 769–783.
[14] C.W. Hirt, A.A. Amsden, J.L. Cook, An arbitrary Lagrangian–Eulerian computing method for all flow speeds, J. Comput. Phys. 14 (1974) 227–253.