NPL
National Physical Laboratory

Report

NPL Report
DEM-ES-021

Report to the National Measurement System Directorate, Department of Trade and Industry

From the Software Support for Metrology Programme

Testing Automatic Mesh Generation Software

L Wright

NOT RESTRICTED

March 2007

dti
National Measurement System

National Physical Laboratory | Hampton Road | Teddington | Middlesex | United Kingdom | TW11 0LW

Switchboard 020 8977 3222 | **NPL Helpline** 020 8943 6880 | **Fax** 020 8943 6458 | **www.npl.co.uk**

The DTI drives our ambition of 'prosperity for all' by working to create the best environment for business success in the UK. We help people and companies become more productive by promoting enterprise, innovation and creativity.

We champion UK business at home and abroad. We invest heavily in world-class science and technology. We protect the rights of working people and consumers. And we stand up for fair and open markets in the UK, Europe and the world.

The National Physical Laboratory is operated on behalf of the DTI by NPL Management Limited, a wholly owned subsidiary of Serco Group plc

Testing Automatic Mesh Generation Software

L Wright
Mathematics and Scientific Computing Group

March 2007

ABSTRACT

This report describes the results of a set of tests carried out on software used for pre-processing continuous models. The work reported here has focussed on automatic mesh generators for finite element (FE) analysis. The aim of the report is to explain some of the factors that affect mesh quality and how quality may be assessed. Results of tests on software packages in common use are described. The report includes a practical example from the field of corrosion modelling.

ISSN 1744–0475

National Physical Laboratory,
Hampton Road, Teddington, Middlesex, United Kingdom TW11 0LW

We gratefully acknowledge the financial support of the UK Department of Trade and Industry (National Measurement System Directorate)

Approved on behalf of the Managing Director, NPL
by Jonathan Williams, Knowledge Leader for the Electrical and Software team

Contents

1 Introduction

This report describes the results of a set of tests carried out on software used for pre-processing continuous models. It is aimed at users of finite element packages who wish to understand more about assessing mesh quality before solution calculation. The work reported here has focussed on automatic mesh generators for finite element analysis. The aim of the report is to explain some of the factors affecting mesh quality, to discuss some methods for assessing mesh quality, and to illustrate the points made using tests run on software packages in common usage. Previous Software Support for Metrology (SS*f*M) programmes have carried out testing of continuous model solution packages [10, 11], and the lessons learnt from this work have been incorporated into the SS*f*M guide to testing algorithms and software [7].

Automatic mesh generation tools are also used for contouring, interpolation, and visualisation of three-dimensional data. The work reported here concentrates solely on finite element mesh generators, but much of what is said applies equally well to mesh generation for other applications.

Finite element (FE) analysis is used throughout UK industry for testing and assessing new designs and processes and for simulation and understanding of unexpected problems or behaviour. Industrial sectors that use FE analysis include aerospace, automotive, civil engineering, oil, gas, power, manufacturing, medical, and communication. FE methods are also used throughout metrology for design of experimental equipment and procedures, interpretation of measured results, and improved understanding of experiments. Applications exist within many areas of metrology. A number of earlier SS*f*M reports have addressed issues connected with FE models [10, 11, 14, 12, 13].

As the use of FE analysis has become more widespread, FE software packages have become more user-friendly and are increasingly operated via a graphical user interface rather than via users creating typed input files. The increase in user-friendliness is a consequence of the automation of many of the steps involved in developing, solving, and post-processing an FE model, and in particular the use of automatic mesh generation techniques. FE methods require a mesh of interconnected points to calculate the solution to a problem, and the generation of the points and their connectivity can be a time-consuming task, particularly for the non-expert. The use of automatic mesh generation has decreased model development time and has significantly broadened the use of FE analysis.

The accuracy of the solution to an FE model is strongly dependent on the mesh used to generate the solution. Hence the quality of the mesh generation software is key to the quality of the model solution, and it is as important to test the mesh generation software as it is to test the model solution software.

1.1 Structure of the report

The report is divided into 6 sections. This section addresses the aims and motivation of the work. Section 2 explains which features of a mesh affect solution accuracy, discusses what makes a good mesh, and defines the scope of the testing work that has been done. Section 3 defines the quality measures that are used in the work and the quantities that are required to calculate them, as well as discussing some of the references that have led to the choice of measures used. Section 4 describes some of the most common algorithms used to generate meshes, focussing on the algorithms used by the packages tested here. Section 5 describes the model problem used for testing and reports the results of the tests. Finally, section 6 applies the testing procedure to a set of meshes used in a corrosion problem as an example with a more complicated geometry.

1.2 Acknowledgements

This guide was produced under the Software Support for Metrology (SS*f*M) programme, which is managed on behalf of the DTI by the National Physical Laboratory. For more information on the SS*f*M programme, contact the programme manager, Mr Bernard Chorley (phone +44 20 8943 7050, email ssfm@npl.co.uk, National Physical Laboratory, Hampton Road, Teddington, Middlesex, UK TW11 0LW), or visit the website at http://www.npl.co.uk/ssfm.

The author would also like to thank Dr Louise Crocker and Dr Alan Turnbull for provision of the mesh of a corrosion problem described in section 6.

2 Mesh quality and solution accuracy

The accuracy of the solution to a FE model strongly depends on the mesh used in the calculation. The mesh is a method of moving from a partial differential equation (pde) plus an associated domain (region of space and time), material properties, and boundary conditions, to a set of linear simultaneous equations that is solved using numerical algorithms. The solution to the pde is parameterised in terms of values of the unknown quantities at points defined by the mesh, and these parameters are the unknowns in the simultaneous equations.

Solution accuracy of a set of simultaneous equations depends in part on the **condition number** of the matrix defining the equations. For a non-singular matrix, the condition number is defined as the ratio of the modulae of the maximum to minimum eigenvalues of the matrix. If a matrix has condition number C, then $\log_{10} C$ is an upper bound on the number of decimal digits lost when solving a system of linear equations involving that matrix. A matrix with a very large condition number is called an **ill-conditioned** matrix, and a matrix with a reasonably-sized condition number is called **well-conditioned**. The quantitative definitions of these terms depends on the accuracy required by the application of interest.

The condition number of a matrix describes a property of the problem that is independent of the choice of algorithm used to solve the system of equations. If the problem is written $A\mathbf{x} = \mathbf{b}$, then the condition number of A times the relative error in A and $\mathbf{b}$ is an upper bound on the relative error in the solution $\mathbf{x}$. Hence the condition number of A quantifies the sensitivity of the solution $\mathbf{x}$ to the data A and $\mathbf{b}$.

For an accurate solution, it is necessary that

1. the mesh approximates the domain and the boundary conditions accurately,
2. the solution to the pde is well approximated by the parameterisation using mesh points (in some cases it is also important that the gradient of the parameterised solution is a good approximation to the gradient of the true solution), and
3. the matrix that defines the simultaneous equations (the **stiffness matrix**) is well-conditioned, so that rounding errors will not affect the accuracy of the algorithms used to solve the equations and the solution is not sensitive to perturbations in the problem data.

The first of these points is the most straightforward to determine, but does not always have a great impact on the accuracy of the solution. If the domain has a lot of fine detail in comparison to its overall size, then automatic meshing of the full, detailed, domain can generate very small elements. The use of such elements can lead to excessive solution times and ill-conditioning of the stiffness matrix, and so an overly precise description of the domain can lead to an unnecessarily time-consuming solution process and inaccurate results. Often the fine details of the domain do not affect the result of interest and so they do not need to be included in the mesh. It should be noted that some details, particularly rounding and filleting of corners, can have important effects [14] such as removing singularities. Careful consideration of the necessary level of detail needs to take place during mesh development.

The second point is difficult to quantify before the solution has been calculated. Various rules of thumb exist that identify areas of the domain where a large number of elements should be used, based on knowledge of the likely behaviour of the solution, e.g. where the solution is varying most. A general rule is that the addition of more elements will make the solution more accurate. Most automatic mesh generation software does not take likely solution behaviour into account, because much of the knowledge required to include this factor is difficult to formalise and hence automate. Ideally a user will define the mesh parameters and subdomains so that the mesh density is appropriate for the solution behaviour, but this process is not an intrinsic part of automated meshing. Various bounds exist for the errors incurred when approximating a continuous function and its gradients on a set of discrete points and triangles. These error bounds lead to usable quality measures defined in terms of geometric properties of the elements, which are discussed in section 3.

The third point is largely dependent on the sizes and shapes of the elements that make up the mesh. The stiffness matrix is assembled by calculating a smaller stiffness matrix for each element, and combining these smaller matrices by considering the formulation of the governing equations at the points where the various elements meet. It can be shown that the condition of the large stiffness matrix depends on the condition of the stiffness matrices of the individual elements [17]. Hence a good mesh will contain elements that have well-conditioned stiffness matrices. The condition number of the element stiffness matrices depends on the shape of the element, the material properties, and the governing pde. In addition to shape properties, the range of element sizes can affect the matrix conditioning, because refinement of small elements can lead to a reduction in the eigenvalue of minimal modulus and thus an increase in the condition number.

The work presented here has concentrated on the second and third points, since the element shape is determined by the automatic mesh generator (and so is not under the user's control) and has a strong effect on solution accuracy. The detailed effects of the element shape on the matrix condition will not be presented here because the stiffness matrix is dependent on the pde of interest, but the ideal shape for elements as outlined below will be valid for most equations (provided any anisotropies in the material properties are taken into account). In addition, the model problem has compared the volume enclosed by the mesh to the volume of the original domain, in order to consider the first point.

2.1 Element shape

For a given pde and material properties, each type of finite element has an ideal shape that will lead to a well-conditioned stiffness matrix. In general, this ideal shape is equilateral and equiangular, but for problems involving anisotropic materials the lengths of the sides may need rescaling to account for the anisotropy. In the remainder of this work, it will be

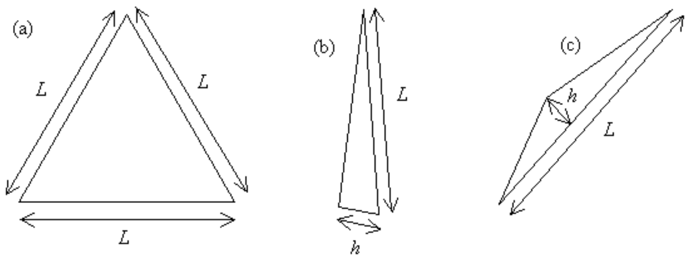

Figure 1: The perfect linear triangular element (a), and its two possible distortions: needles (b) and flattened triangles (c). Problems occur when $h/L \to 0$.

assumed that materials are isotropic. The majority of finite element models use isotropic properties (although applications with orthotropic materials such as piezoelectric elements and composite materials are becoming more common), and rescaling the axes and variables of an anisotropic problem will make the work presented here applicable to many cases.

In general, problems with solution accuracy occur when either one of the side lengths or one of the internal angles approaches zero. These cases are problematic because the calculation of the stiffness matrix of each element is carried out in a local coordinate system, and the transformation between local and global coordinates becomes singular if a side length or an internal angle is zero.

Figure 1 shows the ideal linear triangular element and the two main ways in which it can be distorted. Following the nomenclature of Pébay and Baker [16], triangles with one small angle and two angles close to $90°$, as shown in figure 1(b), will be called **needles**, and triangles with two small angles and one angle close to $180°$, as in figure 1(c), will be called **flattened** triangles.

The ideal linear quadrilateral element and some of its possible distortions are shown in figure 2. As stated above, aspect ratio distortion is not always problematic, but it has been considered in the results that follow. All these distortions will lead to degenerate elements if the quantities h or θ equal zero: either the element's area becomes zero or (for the tapering case) the element becomes a triangle. If a two-dimensional quadrilateral element is used in three dimensions (for example, for modelling a thin structure such as a membrane), the element can also be distorted by warpage. Warpage occurs when the nodes are distorted a long way in an out-of-plane direction. As an example, consider moving the nodes of a square with corners $(0,0,0), (1,0,0), (0,1,0), (1,1,0)$, i.e., a square lying in the plane $z = 0$, to $(0,0,-1), (1,0,1), (0,1,1), (1,1,-1)$. The element will still be equilateral and equiangular, but all the nodes lie a long way from their mean plane.

In three dimensions, the possible distortions are too complicated to draw clearly, but can be imagined as the three-dimensional equivalents of the two-dimensional distortions described above. It is reasonable to assume that if all the sides of a three-dimensional element are

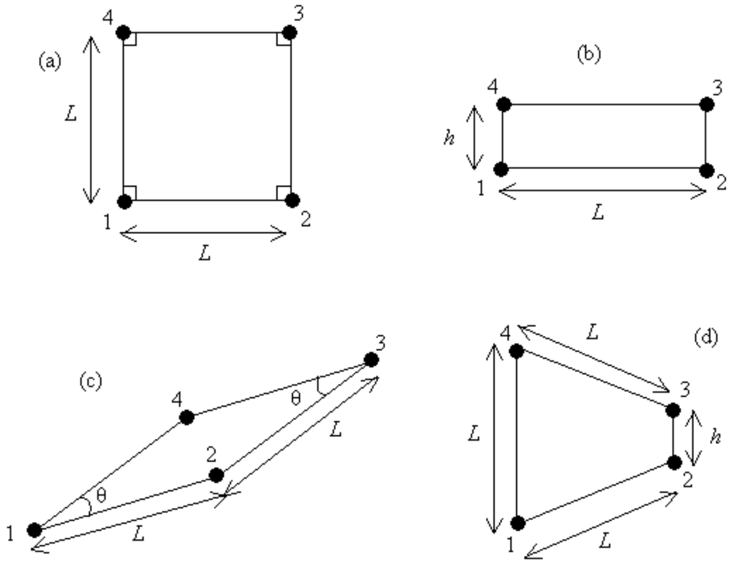

Figure 2: The perfect linear quadrilateral element (a), and its possible in-plane distortions: low aspect ratio (b), skewing (c), and tapering (d). Problems occur when $h/L \rightarrow 0$ or $\theta \rightarrow 0$. Other distorted quadrilaterals either can be created by combining several of these distortions, or are unlikely to be created during automatic mesh generation (e.g. concave shapes such as arrows).

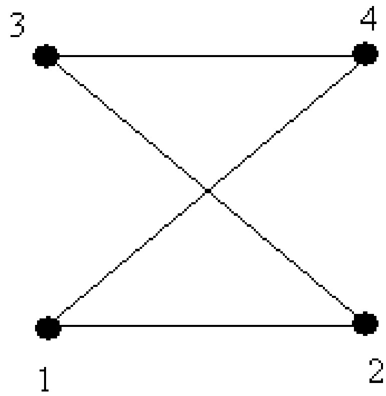

Figure 3: A degenerate element. This element has probably had its node numbers defined in the wrong order for the FE pre-processor and will be interpreted as having a negative area.

undistorted two-dimensional elements, then the three-dimensional element will lead to a well-conditioned stiffness matrix. The drawback of this assumption is that the calculations of the face properties can take a long time, and so a quality measure based on properties of the whole element is preferable.

The examples above illustrate the possible basic distortions of linear elements. More complicated distortions can be generated by a combination of several of these basic distortions. The addition of mid-side nodes to produce quadratic and higher order elements leads to a further range of possible distortions. An ideal element of any order should be close to equilateral and equiangular, but since no quality measures that take mid-side nodes into account were identified, only linear elements have been considered.

In addition to the distortions described above, it is possible to generate degenerate elements. Degenerate elements can take several forms, but in general such elements have either a side of zero length, an internal angle of zero size, or some other feature that makes the element unsuitable for use in an approximation. As an example, consider figure 3. The quadrilateral element shown has had its nodes defined in the wrong order and so has a degenerate shape. Similarly, some packages require triangular elements to have their nodes defined in anti-clockwise order, and calculate negative areas for triangles that do not obey this condition. This type of distortion is dependent on the requirements of the software package pre-processor since there is no universal standard order for listing the nodes of an element. Since this type of distortion is specific to FE model solution packages, it has not been considered in this work.

3 Quality measures

Section 2 described the ways in which element shape can affect the accuracy of FE solutions. This section defines some quantities associated with element shape, and introduces the quality measures that have been used to assess element quality.

3.1 Geometric definitions

Consider the triangle shown in figure 4. The edge connecting nodes i and j has length l_{ij}, the internal angle at node i is θ_i, and the triangle has area A. The **half-perimeter**, p, is defined as $p = \frac{1}{2}(l_{12} + l_{23} + l_{31})$. The root mean square (RMS) edge length is given by $l = \sqrt{(l_{12}^2 + l_{23}^2 + l_{31}^2)/3}$. Call the maximum and minimum edge lengths as l_{max} and l_{min} respectively, and call the other length l_{med}. For a tetrahedron, the edge lengths are identified in the same way as those of the triangle, the area of the face opposite node i is A_i, and the volume of the tetrahedron is V.

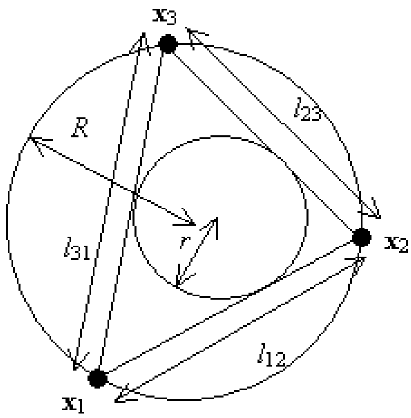

Figure 4: Useful properties of a triangle. Note that the two circles are not concentric, and that the larger circle would not pass through all three vertices if the largest angle was greater than 90°.

The **minimum containment circle**, radius R, is the smallest circle entirely enclosing the triangle. The **maximum inscribed circle**, radius r, is the largest circle that lies entirely within the triangle. The **minimum containment sphere** and **maximum inscribed sphere** are the corresponding three-dimensional equivalents for a tetrahedron. It can be shown that for a triangle, $A = \sqrt{p(p - l_{12})(p - l_{23})(p - l_{31})}$ and $r = A/p$. Similarly, for a tetrahedron $r = 3V/(A_1 + A_2 + A_3 + A_4)$.

For a triangle, there are two possible formulae for R. If the largest angle of the triangle is greater than 90°, $R = l_{max}/2$, i.e., the longest side is a diameter of the circle. If all angles of the triangle are acute, it can be shown that $R = (l_{12}l_{23}l_{31})/(4A)$. For a tetrahedron, the possibilities are more complex and so they will not be listed in full here (see Shewchuck [17]).

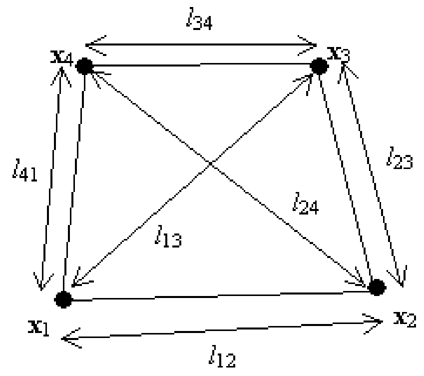

Figure 5: Useful properties of a quadrilateral. Out-of-plane properties have not been illustrated.

For quadrilaterals, the numbering convention for the corners used in this work is as shown in figure 5, and the lengths of the lines connecting the vertices, including diagonal lengths, are labelled in the same way as those of the triangle. Let l_{max} and l_{min} be the maximum and minimum side lengths (excluding diagonals) respectively, and let θ_{max} and θ_{min} be the maximum and minimum angles. W is the measure of out-of-plane warpage, given by

$$\mathbf{n} = \frac{(\mathbf{x}_1 - \mathbf{x}_3) \times (\mathbf{x}_2 - \mathbf{x}_4)}{|(\mathbf{x}_1 - \mathbf{x}_3) \times (\mathbf{x}_2 - \mathbf{x}_4)|}, \quad (1)$$

$$W = \frac{|4\mathbf{x}_1.\mathbf{n} - \sum_{i=1}^{4} \mathbf{x}_i.\mathbf{n}|}{\sum_{i=1}^{4} \mathbf{x}_i.\mathbf{n}}. \quad (2)$$

This value gives a measure of how non-planar the element is, with a perfectly planar element having a value of 0. The use of $\mathbf{x}_1$ in this formula is arbitrary: the same value will be obtained no matter which node is used.

3.2 Background to the quality measures

The quality measures described here came from several sources. Some software packages use measures like these to assess element quality, or to check the mesh prior to solution. The measures used for quadrilaterals are the most obvious measures for the possible types of distortion. The majority of the measures for triangles have been taken from papers on mesh quality.

Two useful papers on triangular element quality [6, 8] address the conditions under which a set of triangular elements gives a good approximation to a function. The first, by Babuška and Aziz [6], shows that the maximum difference between a twice differentiable function and its approximation on a linear triangle depends on the value of the largest angle of the triangle, and as long as the largest angle is not too close to 180° (i.e., the element is not a flattened triangle), the approximation will match the function well. It had been thought previously that the main angle criterion for triangles was that the smallest angle must not be too small, but the authors showed that the largest angle is of more importance. The paper generalises its results to higher-order elements, but does not provide any practically applicable quality measures for such elements.

The other paper, by Berzins [8], aims to assess the quality of a triangular or tetrahedral mesh after the solution to a pde problem has been calculated on that mesh. Initially, quality measures are derived in terms of the gradients of the solution along the edges of an element, but these measures are then generalised to geometric measures. The measure for triangles is illustrated by application to a two-dimensional anisotropic fluid flow problem designed to show the suitability of anisotropic elements for some problems. The geometric measures derived in this paper have not been used in this work because they are not bounds on the approximation error, unlike the measures derived in Shewchuck [17].

A paper by Pébay and Baker [16] analyses several of the commonly-used quality measures for linear triangular elements, assesses them to ensure that they have a single unique extreme value achieved only when the triangle is equilateral and investigates their behaviour for near-degenerate triangles such as those shown in figures 1(b) and 1(c). The study concludes that measures M_{11}, M_{12}, and M_{13} (see section 3.3 for definitions) would identify both classes of distorted triangle and would have a unique maximum value for an equilateral triangle.

Jonathan Shewchuck, a former winner of the Wilkinson prize for his automatic mesh generator *Triangle*, provides [17] definitions of M_{14}, M_{15}, M_{32}, and M_{33} (see section 3.3). Two versions of this paper exist. The full-length version includes derivations and proofs of the validity of these measures, an extensive list of other element quality measures, and a discussion of the effects of a single bad element. The shorter version explains the reasoning that went into the development of the measures without providing detailed proofs, and indicates which sources of numerical error the various measures identify. Shewchuck shows that M_{15} and M_{33} (see section 3.3) are designed to identify elements that have ill-conditioned stiffness matrices when applied to the heat equation (Poisson's equation). This may mean that these measures are unsuitable for other pdes.

Of the remaining tetrahedral measures (see section 3.3), M_{31} is used by Comsol Multiphysics to assess the quality of meshes generated by the Comsol Multiphysics mesh generator, and M_{34}, an adaptation of the equivalent triangular condition M_{11}, is cited in several places but was originally developed by Cavendish [9].

3.3 Quality measure definitions

The following quality measures have been taken from references or from the manuals of automatic meshing packages. Some automatic mesh generators check the element quality for the user, but different packages use different measures which makes intercomparison of their quality results difficult. The origins and strengths of the various measures are discussed in section 3.2.

A good quality measure should

- have a single unique maximum value, achieved by an equilateral equiangular "perfect" element, and no local maxima nor minima,
- be independent of element orientation, and
- be independent of linear scaling: if the lengths of all sides are doubled, the quality measure should not change.

All the measures in this section meet these criteria. The measures listed below have been scaled so that they will have a value of 1 for an equilateral equiangular perfect element shape and a value of 0 for degenerate elements. The quadrilateral warping measure is scaled so that it is 1 when the element is perfectly planar. Note also that each measure is a dimensionless quantity, so the above scaling requirement is fulfilled.

No attempt has been made during this work to define which values of the quality measures separate a "bad" element from a "good" one. This delineation will depend on the application, the equation to be solved, and the solver accuracy required. The quality measures are probably at their most useful when they are used to compare different meshes. In general it seems reasonable to assume that any value less than 0.1 is bad enough to warrant further investigation, and any value above 0.9 is good enough to be accepted. These are arbitrary limits and may be unsuitable for some practical applications. The Comsol Multiphysics manual suggests that any triangular element with a value of M_{13} greater than 0.3 will not affect solution accuracy, and any tetrahedron with a value of M_{31} greater than 0.1 will be acceptable.

The triangular quality measures used are:

$$\begin{aligned}
M_{11} &= \frac{2r}{R}, &(3)\\
M_{12} &= \frac{3\sqrt{3}r^2}{A}, &(4)\\
M_{13} &= \frac{4\sqrt{3}A}{3l^2}, &(5)\\
M_{14} &= \frac{(4(\sqrt{3}+2)^{2/3})A}{3^{5/6}(l_{max}l_{med}(l_{min}+4r))^{2/3}}, &(6)\\
M_{15} &= \frac{(4\sqrt{3})A}{3l^2+\sqrt{|9l^4-48A^2|}}. &(7)
\end{aligned}$$

The first three measures are fairly straightforward to interpret. If a triangle has a small value of r and a large value of R, then it must resemble one of the deformed triangles shown in figure 1. Similarly, a large perimeter and a small value of r mean that the triangle must be distorted. Finally, a triangle with an area that is shrinking more quickly than l^2 is likely to be a flattened triangle. All are quality measures for the interpolation error of the values of a function. M_{14} is a quality measure for the interpolation error of the gradients of a function, and M_{15} is a quality measure for the conditioning of the stiffness matrix generated when the element is applied to Poisson's equation.

As a guide to how these measures behave, figures 7 and 8 show the variation in quality measure value with angle, for the triangles shown in figure 6. The triangles tested are one set of right-angled triangles and one set of isoceles triangles. There are several points of interest about these graphs. The first is that the right-angled triangle results are symmetric about $\theta = 45°$, as would be expected (if this were not the case, the measure would depend on element orientation). It is also interesting to note that M_{11}, M_{12}, and M_{13}, the most straightforward criteria to understand, are quite similar in overall shape for the types of triangle shown.

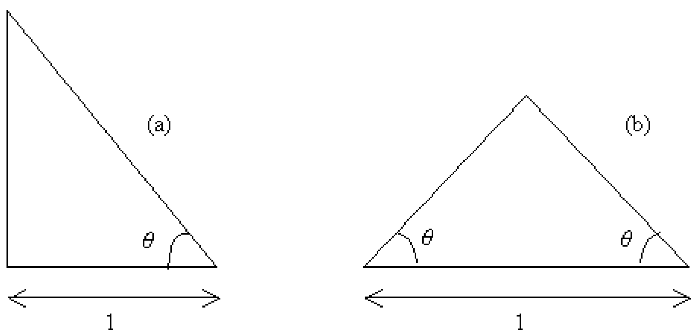

Figure 6: Triangle shapes and angles used to calculate values shown in figures 7 and 8.

Another point to note is that measure M_{15} is the least forgiving of the measures: for a given element, the value of M_{15} will be less than or equal to all other measures, and conversely for a given value of a measure, the range of triangles that can achieve that value is smaller for M_{15} than for the other measures. By the same criteria, M_{14} is the most forgiving. As further evidence of this, figures 9 and 10 show contour plots of how M_{14} and M_{15} (respectively) vary for a general triangle with one side of unit length (i.e., if the two angles marked in figure 6(b) were not necessarily equal).

The horizontal and vertical axes show the two angles of the triangle at either end of the unit length side, and each contour line represents $M_{14} = 0.1k, k = 0, 1, \ldots, 10$, with each line being marked with its value on the plot. The top right-hand region of the plot has zero value for both plots because the angles represented in that area add up to more than $180°$. The central area of the two plots represents $0.9 \leq M \leq 1$ (because the value of 1 is only

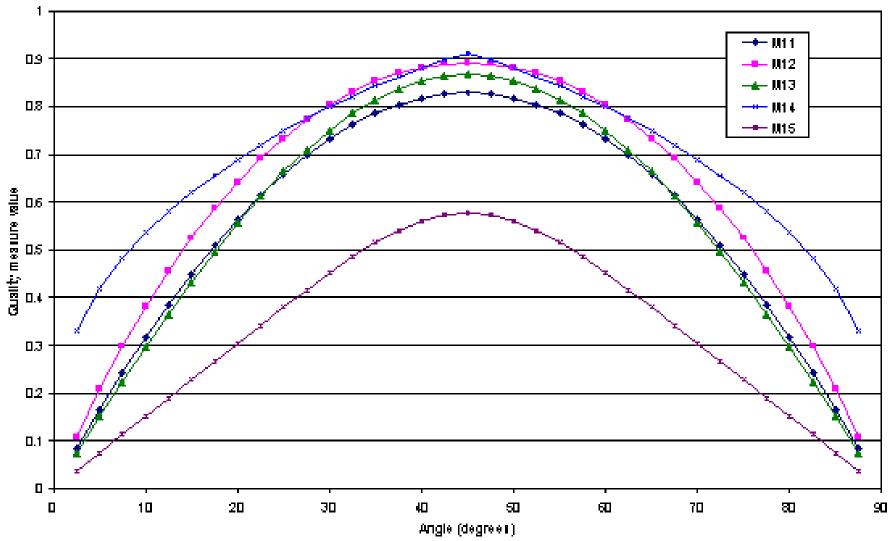

Figure 7: Variation of the quality measures with angle for a right-angled triangle.

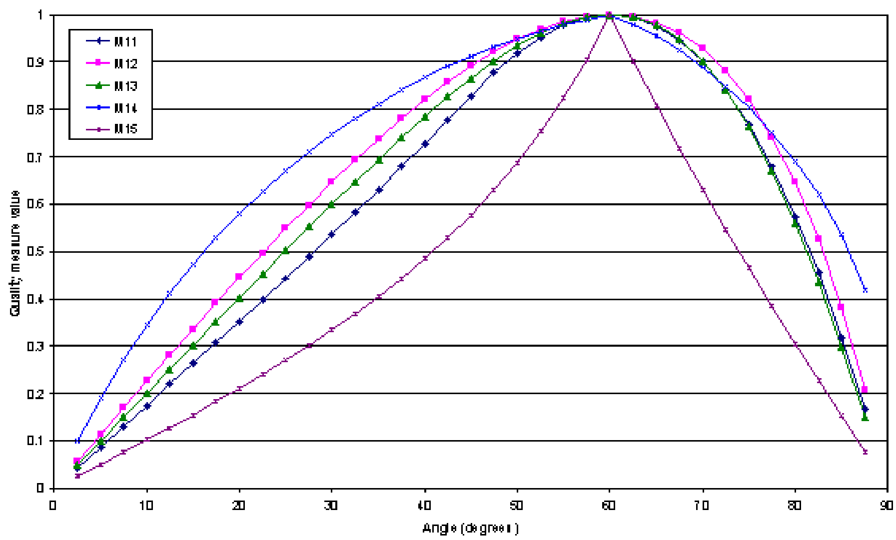

Figure 8: Variation of the quality measures with angle for an isoceles triangle.

achieved for an equilateral triangle), and is much smaller for the M_{15} plot: the area covered by $0.9 \leq M_{14} \leq 1$ is approximately equivalent to that covered by $0.6 \leq M_{15} \leq 1$.

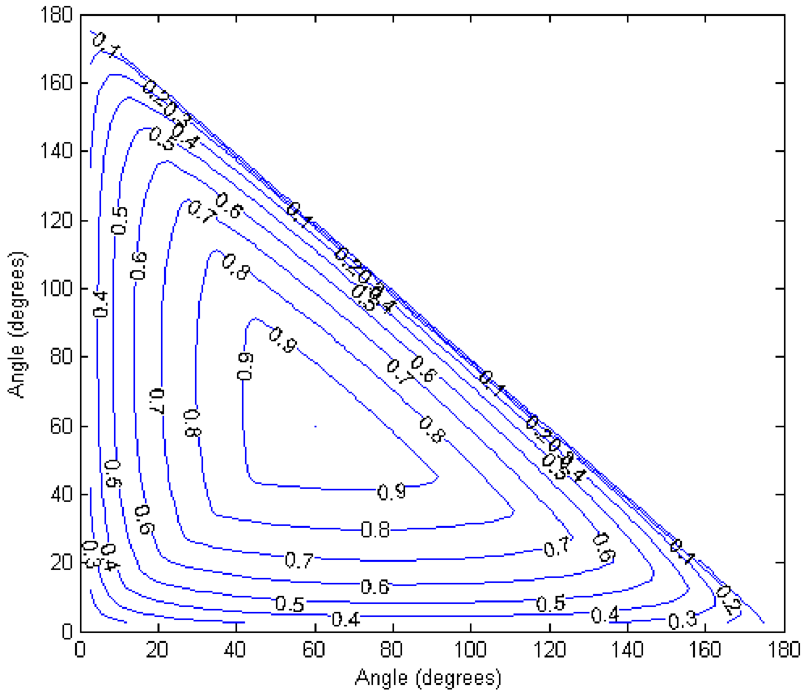

Figure 9: Variation of M_{14} with angle for a general triangle with one side of unit length. Each contour band represents a range of 0.1, and the measure is 0 at the edges and 1 for both angles equal to 60°.

The quadrilateral measures are

$$M_{21} = 1 - W, \tag{8}$$

$$M_{22} = \frac{l_{min}}{l_{max}}, \tag{9}$$

$$M_{23} = \frac{\theta_{min}}{\theta_{max}}. \tag{10}$$

These measures aim to identify warping, aspect ratio problems, and tapering/skewing, respectively. Their interpretation is much simpler than for the triangular quality measures. It is clear from these formulae that the quality measures for triangles are more complicated than those for quadrilaterals. The most likely cause for this is that triangles have been studied more widely. Triangles are more widely used for automatic mesh generation than quadrilaterals because they are more flexible and generally more straightforward to generate. This widespread usage may have led to more study of what constitutes a good quality triangle.

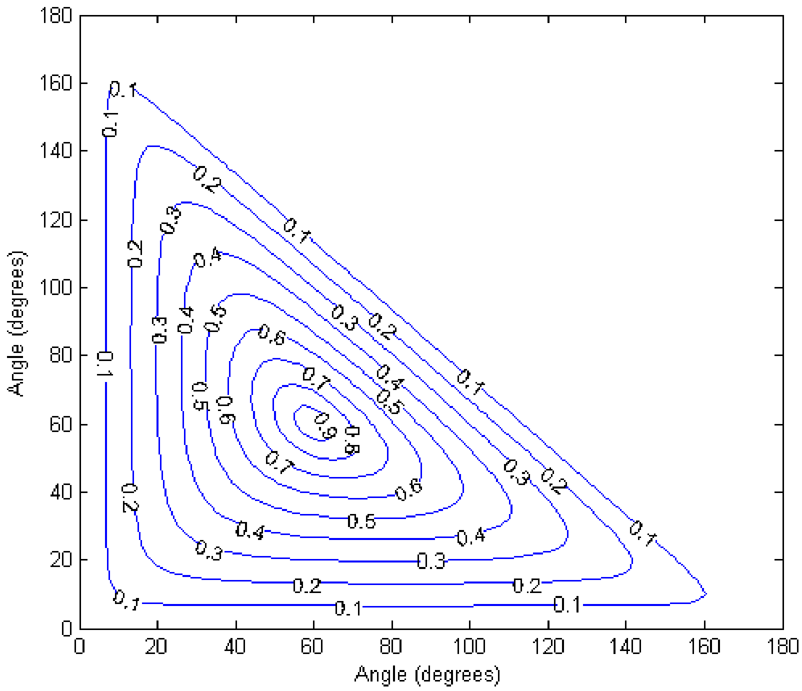

Figure 10: Variation of M_{15} with angle for a general triangle with one side of unit length. Each contour band represents a range of 0.1, and the measure is 0 at the edges and 1 for both angles equal to 60°.

Tetrahedra are widely used in three-dimensional mesh generation, and a set of suitable quality measures have been identified:

$$
\begin{aligned}
M_{31} &= \frac{72\sqrt{3}V}{(\sum_{1\le i<j\le 4} l_{ij}^2)^{3/2}}, && (11)\\
M_{32} &= V\left[\frac{\sum_{i=1}^{4} A_i}{\sum_{1\le i<j\le 4} A_iA_jl_{ij}^2 + 6V\max_{1\le i\le 4}\sum_{j\ne i} A_jl_{ij}}\right]^{\frac{3}{4}}, && (12)\\
M_{33} &= \frac{1}{\sqrt{3}}\left(\frac{V}{\lambda_{max}^3}\right)^{\frac{1}{4}}, && (13)
\end{aligned}
$$

where λ_{max} is the solution of

$$z^3 - \frac{\sum_{i=1}^{4} A_i^2}{9V}z^2 + \frac{\sum_{1\le i<j\le 4} l_{ij}^2}{36}z - \frac{V}{9} = 0 \qquad (14)$$

with the largest modulus, and

$$M_{34} = \frac{3r}{R}. \qquad (15)$$

The only one of these measures that can be interpreted easily is the final one. A tetrahedron with a small value of r and a large value of R will be almost flat in one direction, which will lead to a poorly-conditioned element stiffness matrix. The other measures are less easy to interpret. M_{31} is a measure of the aspect ratios of the element used by the Comsol Multiphysics FE package. M_{32} is a quality measure for the interpolation error of function gradients, and M_{33} is a quality measure for the conditioning of the stiffness matrix generated when the element is applied to Poisson's equation.

Hexahedral and wedge elements were tested by applying the triangular and quadrilateral quality measures to each face. This approach is taken by some finite element software packages when checking element quality. It seems reasonable to expect that the faces of a good regular three-dimensional element would all make good two-dimensional elements.

4 Automatic meshing

This section introduces some of the ideas and algorithms that are used in automatic mesh generation. An excellent introduction to meshing techniques [15], which has been the starting point for much of the material in this section, can be found here:

http://www.andrew.cmu.edu/user/sowen/survey/index.html.

The pages linked to this introduction include recent research on mesh generation, and useful conferences and publications in the area. The references section of the introduction to meshing provides an extensive reading list for anyone interested in further information.

Automatic meshing techniques are used for visualisation and interpolation as well as for mesh generation for FE techniques. In some cases the aim is to connect a set of points by defining a set of polyhedra whose vertices are members of the set of points, with the purpose of approximating some underlying unknown surface on which all the points lie. In the more general case, the aim is to fill a given area, surface, or volume with polyhedra and associated vertices so as to minimise the difference between the original shape and the polyhedral approximation. The descriptions of algorithms that make up the rest of this section concentrate on the second of these aims, because the packages that have been tested all aim to fill a given shape with elements rather than interpolate specified points.

The software may have a secondary aim of generating "good quality" polyhedra, i.e., those that are near-equilateral and/or near equiangular, or of generating a "smooth" surface, i.e., one in which the gradients of the surface are continuous at the vertices. In both cases it is common to generate an initial mesh and then attempt to improve it by moving the vertices. This work does not consider smoothing, quality improvement, nor other such post-generation algorithms.

Several stopping criteria can be used to control meshing algorithms, and the choice of criterion is usually determined by the implementation rather than the algorithm. Some implementations require the user to specify an approximate number of nodes, which is convenient for finite element applications as it allows the user to control the number of degrees of freedom and hence the computational effort. Others require the user to specify a maximum element dimension, and the algorithm stops when all element sides are smaller than this length.

Other criteria control behaviour on the boundaries. Sometimes the user is required to specify the number of elements that lie along each of the outer edges of the space to be meshed. Some software allows these criteria to be applied to subsets of the whole space, so that the user can specify a finer mesh in areas that require a more accurate solution, and some packages allow the user to specify regions of an initial mesh for refinement.

Meshes can be divided into two broad types: structured and unstructured. The interior points of a **structured** mesh are all vertices of the same number of elements, but the interior points of an **unstructured** mesh can be part of any number of elements. As an example, consider the two meshes of a unit square shown in figure 11. Mesh (a) is a structured mesh because all interior points are vertices of exactly four elements. Mesh (b) is an unstructured mesh, because different interior points are vertices of different numbers of elements.

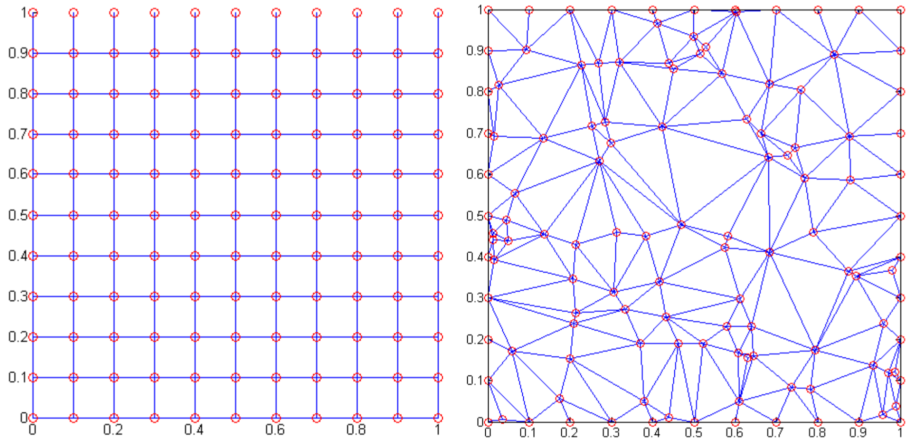

(a) (b)

Figure 11: Two meshes of a unit square. Mesh (a) is structured, mesh (b) is unstructured.

This section will include a short description of some of the algorithms that have are used by the packages that have been tested in section 5. Three different algorithms are known to have been used, and a list of which package uses which algorithm can be found in section 5.1. The classes of method that are described here are Delaunay methods, advancing front methods, and mapped meshing. The first two classes create unstructured meshes, and mapped meshing creates structured meshes.

4.1 Delaunay methods

A Delaunay triangulation of a set of points is a set of triangles connecting those points such that the triangles all obey the Delaunay, or “empty circle”, criterion. The empty circle criterion requires that the circle that passes through all three of the triangle’s vertices (the **circumcircle**) does not contain any of the other points in the set. An illustration is shown in figure 12. The two triangles in 12(a) obey the empty circle criterion, but the triangles in 12(b) do not because each triangle’s circumcircle surrounds a node that is not part of that triangle. A similar criterion can be defined in three dimensions by considering tetrahedra and their circumspheres. The rest of this section will refer to “triangulation” throughout, but all remarks apply equally to tetrahedral meshes (unless stated otherwise).

Delaunay triangulations maximize the minimum angle of all the angles of the triangles in the triangulation. Hence they are a reasonable basis for a finite element mesh, because they minimise the distortions shown in figure 1(b). It can be shown that the Delaunay triangulation of a set of points always exists and is unique, provided no three points lie on a

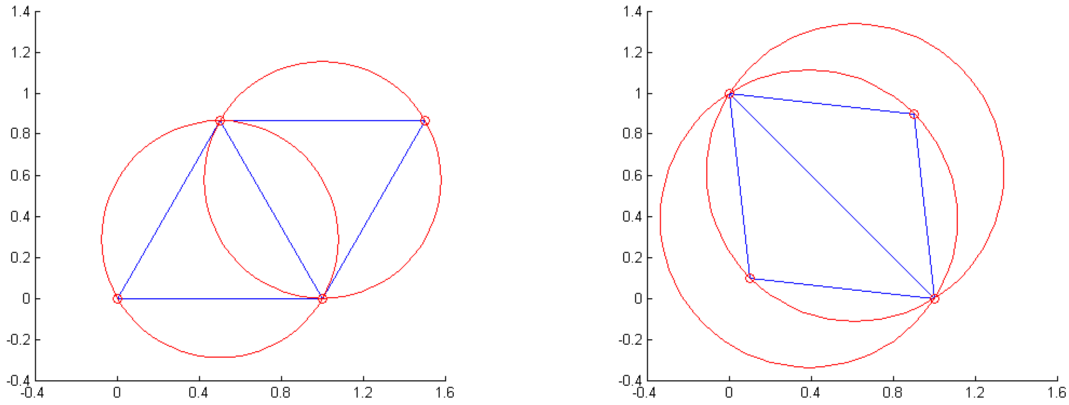

(a) (b)

Figure 12: Triangles that obey the empty circle criterion (a), and triangles that do not (b).

line and no four points lie on a circle.

Since Delaunay triangulation is a method for connecting a given set of points, any method based on creating such a triangulation requires a point generation algorithm. Most implementations create points along the boundary of the domain, often based on a user-specified spacing, triangulate the boundary points, and then add more points according to some (possibly iterative) algorithm and refine the triangulation based on the new points. The point definition method is the factor that most strongly influences the mesh characteristics such as uniformity and structure.

Popular methods for defining points on the interior of a triangulation include:

- Defining a grid of points based on a user-defined spacing. This option will only require a single triangulation to be calculated, but could lead to distorted elements.
- Placing new points at the centroid of existing triangles. This is an iterative method that can be repeated until some stopping criterion based on number of points or maximum element size is reached.
- Placing points at the centre of the circumcircle of the triangle. It can be shown that careful use of this method can enforce a lower bound on the minimum angle of all triangles in the mesh.
- Voronoi methods, where the new point is placed along the line linking the centres of the circumcircles of two adjacent elements. This method leads to regular, almost structured meshes with six triangles meeting at each node.
- Advancing front methods, which identify the best new location for a node close to each edge of the boundary, redefine the local mesh so that the empty circle criterion

is fulfilled, and repeat the process using the newly-defined nodes as a boundary. This method leads to structured meshes with an interior structure that follows the shape of the boundary. The advancing front method for element generation is discussed in more detail in the next section.

- Defining points at some user-defined increment along the non-boundary edges of the initial triangulation, taking care to remove any points that are too close together.

4.2 Advancing front methods

Advancing front methods (also called moving front methods) create unstructured triangular and tetrahedral meshes. The methods discretise the boundary into facets, then make each facet into an element by calculating the "best" position for an internal node to complete the shape. The "best" position could lead to the shape closest to an equilateral triangle, or to preservation of some minimum length criterion, or to some other desirable mesh quality such as one of the mesh quality measures discussed in section 3. Once all the boundary facets have been turned into elements, the new nodes are tested to see if any more elements can be formed, and the resulting internal faces of the newly-created elements (the new front) are regarded as boundary facets and the process is repeated.

Figure 13 shows the three stages described. The data shown were generated using the NAG library routines D06BAF (boundary mesh generation) and D06ACF (advancing front meshing of a region given a boundary mesh). Figure 13(a) shows the boundary points and associated linear facets, (b) shows the first set of interior points and elements, and (c) shows the next layer of elements including the new front.

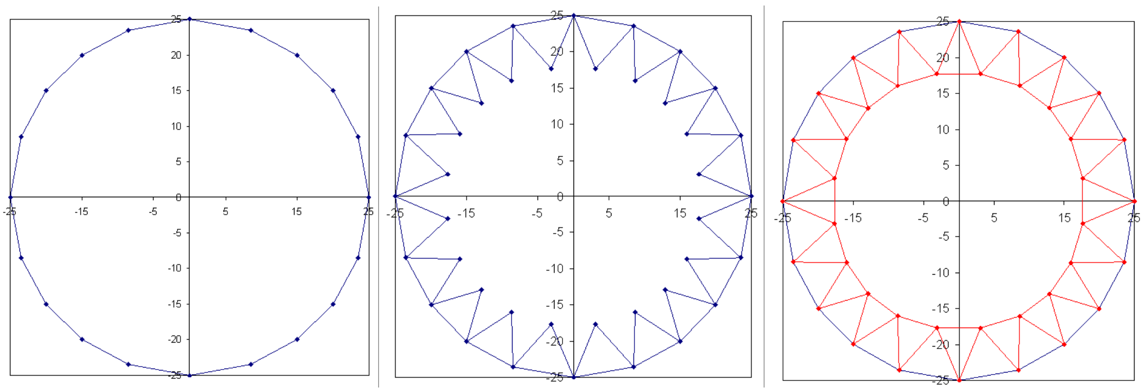

(a) (b) (c)

Figure 13: The three stages in advancing front mesh generation: (a) placement of nodes on the boundary, (b) generation of internal nodes and elements, and (c) formation of a new front by forming more elements (shown in red)

For comparison, figure 14 shows two meshes of the same geometry created using different techniques implemented within the NAG library. Figure 14(a) shows an advancing front mesh, and figure 14(b) shows a Delaunay triangulation created using an incremental point insertion method. The two meshes were derived from the same boundary mesh. The Delaunay triangulation has 62 elements and the advancing front method generated 86 elements. Note that the advancing front method has several inner circular boundaries that are similar to the initial outer boundary.

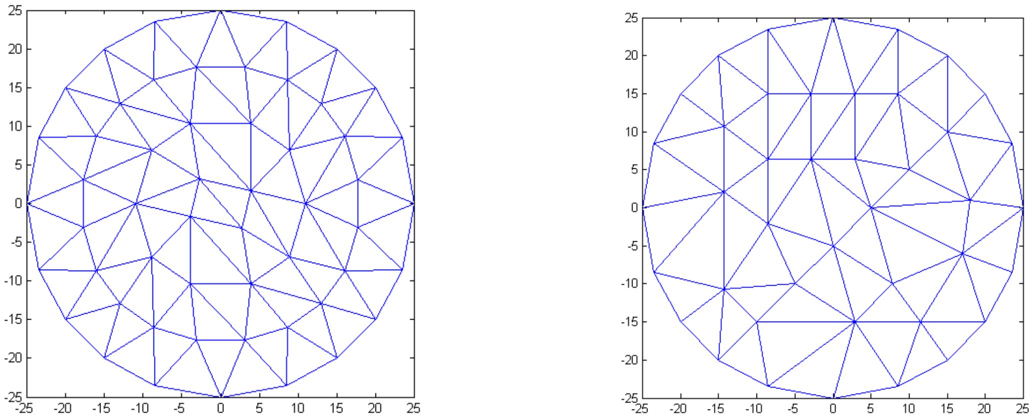

(a) (b)

Figure 14: Two meshes of the same geometry created using (a) the advancing front method, and (b) a Delaunay method with an incremental point insertion technique.

4.3 Mapped methods

Mapped meshing generates structured meshes and is often used to generate quadrilateral and hexahedral meshes. The methods require regions of the boundary that are "opposite" to one another in some sense to be subdivided into the same number of boundary facets. This requirement limits the range of geometries that are suitable for this type of meshing, but in many cases regions built up from simple polyhedra can be subdivided into regions that can use mapped methods.

At its simplest, mapped meshing leads to evenly-sized quadrilateral meshes as shown in figure 11(a), but biaising and subdivision can be used to generate a higher mesh density in regions of importance. In addition, "opposite" as used in the description above need not mean "parallel to", and regions being regarded as opposite one another need not consist of a single line segment. As an example, consider the decagon shown in figure 15. The mesh was created using a mapped method, regarding the face marked 1 as opposite the four faces marked 3, and the two faces marked 2 as opposite the three faces marked 4. This illustrates the potential flexibility of the method.

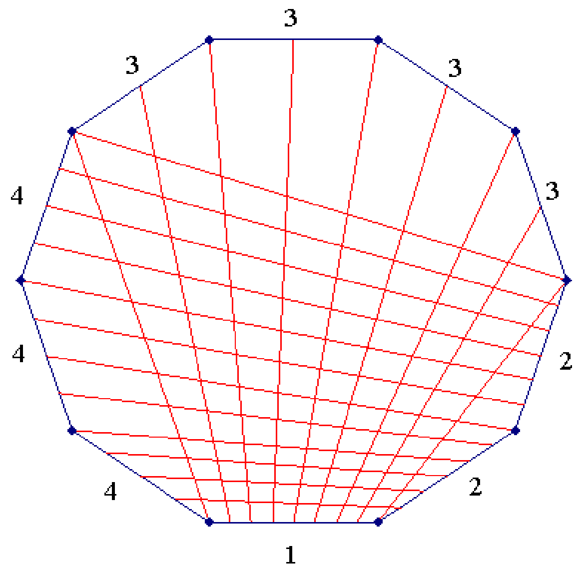

Figure 15: A mesh created using a mapped method, with faces marked 1 and 3 regarded as "opposite" one another and faces marked 2 and 4 regarded as "opposite" one another.

Hexahedral meshes can be created in a similar way by mapping pairs of quadrilateral boundary meshes onto one another. Alternatively, a quadrilateral mesh can be swept along some user-defined path and subdivided along the path length to create a hexahedral mesh. As with two-dimensional mapped methods, there is some flexibility regarding which pairs of meshes can be mapped onto one another.

4.4 Other methods

The methods described above are those implemented in the software described in this report, but other meshing techniques exist. It is possible to create unstructured quadrilateral meshes. A common method of achieving this is to create a triangular mesh and then split every triangle into three quadrilaterals, although this can lead to poor quality elements. An alternative is to merge adjacent pairs of triangles to form quadrilaterals. This method usually results in a mesh with a mixture of quadrilaterals and a few triangles, called a quad-dominant mesh. Algorithms exist to maximise the quality of the resulting elements by careful choice of triangle pairs for merging. These methods are called indirect methods because they rely on alteration of a triangular mesh, rather than spontaneous creation of quadrilateral elements. Direct methods for quadrilateral meshing include a form of the advancing front method called the paving method, and a decomposition method that splits the area into regions with shapes that are known to be meshable using quadrilaterals, then fits the regions together at the boundaries.

Unstructured hexahedral meshing is an ongoing area of research, because unstructured meshes are geometrically more flexible, and because, in some cases, hexahedral finite elements give more accurate solutions than tetrahedral ones. Two indirect methods are theoretically possible. A tetrahedron could be subdivided to create four hexahedra, but this method generally leads to such poor element quality that it is not in common usage. Similarly, five tetrahedra could be combined to form a hexahedron but no examples of practical implementations have been found in the literature. A variety of direct methods also exist, including equivalents of domain decomposition methods and the paving method, but many approaches seem to have reliability or mesh quality problems.

This lack of reliable hexahedral generation methods has led to implementations of hex-dominated methods where as much of the space as possible is filled with hexahedra using one of the direct methods whilst maintaining good element quality, and the more complex regions are filled with tetrahedra instead.

5 Model problems

Two model problems were used to test automatic mesh generators, one in two dimensions and the other in three dimensions. The two-dimensional test was to mesh a circle with a radius of 25 units, and the three-dimensional test was to mesh a cylinder of radius 25 units and height 5 units. These simple geometries were used so that the true area/volume was known, and the area/volume enclosed by the mesh could be calculated and compared to the true value.

In order to make comparisons between the different mesh generators fair, meshing refinement was set such that the two-dimensional mesh used about 2000-3000 nodes, and the three-dimensional mesh used about 20000-25000 nodes. Linear elements were used in all cases. Every shape of element available was used to generate a mesh, and some meshes included several element types.

No attempts have been made to link mesh quality to solution accuracy because it is very difficult to isolate mesh distortion from the other sources of numerical error (e.g. poor choice of solution algorithm for the linear equations, poor choice of numerical integration algorithm, etc.).

5.1 Software packages

Seven different packages were used to generate meshes for the model problems. They included:

- Two mesh generators that are part of finite element software pre-processors (Comsol Multiphysics [1] and Abaqus [2]).
- One software tool for mesh generation that is not tied to a specific FE solution package (FEMGV [3]).
- Three mesh generators that are included in multi-purpose software tools/libraries (two NAG [4] routines, one Matlab [5] function).
- One piece of in-house software called FD.

The algorithms used by each package were:

- Delaunay: Comsol, FEMGV, NAG, Matlab.
- Advancing front: NAG.
- Mapped: FEMGV, FD.

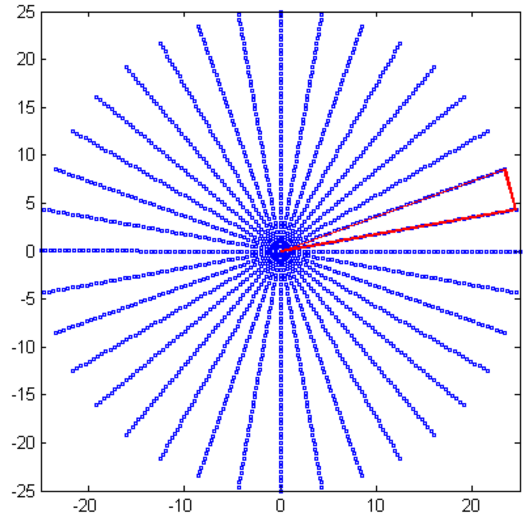

Figure 16: Typical cross-section through a mesh created by the FD software. A strip of elements is outlined.

At the time of writing this report, no details were available regarding the algorithm used by Abaqus. The in-house software was designed to generate a grid for use with cylindrical polar coordinates, so its elements are constructed using constant spacings in the radial, angular, and through-thickness directions.

The FD software is part of a routine developed to model laser flash thermal diffusivity measurement on cylindrical samples [12]. The model's transient heat flow calculation is carried out in cylindrical coordinates, and so the mesh is designed to create mesh points at equally spaced intervals of r, θ, and z. A cross-section through a typical FD mesh is shown in figure 16. The mesh is made up of a set of identical circular faces, each face is made of a number of identical strips (one strip is outlined in the figure), and each strip consists of a number of quadrilateral elements and a single triangular element at the centre. This structure means that all of the triangular wedge type elements are identical in size and shape. Hence the element identified as having the lowest value of a quality measure over all wedges in mesh FD1 is an arbitrary one and should be ignored.

Five different element types were used in the meshes. As well as triangles, quadrilaterals, tetrahedra and hexahedra, triangular prisms or "wedges" were used. These elements were used either on their own or in combination with hexahedra in a "hex-dominated" mesh. All hexahedral, hex-dominated, and wedge meshes used a swept method in the through-thickness direction.

The packages were used to produce a total of 16 meshes, labelled as follows:

- Abaqus: hexahedral (A1), tetrahedral (A2), wedge (A3), hex-dominated (A4), triangular (A5).
- Comsol: tetrahedral (C1), triangular (C2).
- FD: hex-dominated (FD1).
- FEMGV: triangular (FG1), two quadrilateral (FG2, FG3), hexahedral (FG4), wedge (FG5).
- Matlab: triangular (M1).
- NAG library: two triangular (N1, N2).

Each mesh was analysed using the quality measures defined in section 3. The hexahedral and wedge elements had each of their faces tested as two-dimensional elements. For each mesh, the maximum, minimum, mean and standard deviation of each of the quality measures was calculated for each element type. In addition, all element areas or volumes were calculated and summed to compare the meshed area or volume to the true area or volume.

5.2 Triangle quality measure results

Table 1 lists the test results for all triangular meshes. As well as listing the maximum, minimum, mean, and standard deviation values for each quality measure, the table lists the element numbers at which the minimum value was achieved. The inclusion of the element numbers makes it possible to check whether the same elements achieved the smallest values of the different measures. Table 2 shows the same results for the triangular faces of the wedge elements from wedge and hex-dominated meshes. Figure 17 plots the minimum values of each of the quality measures for all the meshes, and figure 18 plots the mean values.

In table 1 it can be seen that the average values of quality measures M_{11} to M_{14} are all between 0.87 and 0.97, and that the standard deviations of these measures are small in comparison with the mean values. These observations indicate that most of the triangles are not severely distorted and will not lead to ill-conditioned stiffness matrices. As was observed in section 3, M_{15} is more severe than the other measures, and so its average and minimum values are lower than those of the other measures, and its standard deviation is larger because the gradient away from a perfect triangle is steeper. However, even the most stringent measure has led to average values greater than 0.64 in all cases (and greater than 0.73 in four cases), so it is reasonable to say that all of the triangular meshes are of good quality.

Table 2 shows that the quality of the triangular faces of the wedge elements is generally as good as the quality of the two-dimensional triangular meshes. This similarity is likely to be caused by the mesh generation method. The wedge elements are created by forming a

Mesh label	Quality measure	Maximum value	Minimum value	Element number	Mean value	Standard deviation
A5	M_{11}	1.000	0.626	1937	0.956	0.057
A5	M_{12}	1.000	0.731	1937	0.971	0.038
A5	M_{13}	1.000	0.682	1937	0.962	0.047
A5	M_{14}	1.000	0.773	1937	0.957	0.037
A5	M_{15}	1.000	0.394	1937	0.807	0.130
C2	M_{11}	1.000	0.529	1858	0.878	0.093
C2	M_{12}	1.000	0.634	1858	0.917	0.065
C2	M_{13}	1.000	0.575	1858	0.896	0.078
C2	M_{14}	0.997	0.696	1858	0.906	0.051
C2	M_{15}	0.983	0.316	1858	0.650	0.127
FG1	M_{11}	1.000	0.555	3984	0.872	0.096
FG1	M_{12}	1.000	0.665	3984	0.913	0.067
FG1	M_{13}	1.000	0.620	3984	0.890	0.080
FG1	M_{14}	0.999	0.757	4144	0.902	0.053
FG1	M_{15}	0.991	0.347	3984	0.646	0.137
M1	M_{11}	1.000	0.790	227	0.954	0.034
M1	M_{12}	1.000	0.866	227	0.969	0.023
M1	M_{13}	1.000	0.837	227	0.960	0.030
M1	M_{14}	0.998	0.875	2067	0.950	0.027
M1	M_{15}	0.989	0.541	227	0.772	0.091
N1	M_{11}	1.000	0.639	75	0.934	0.052
N1	M_{12}	1.000	0.742	75	0.956	0.035
N1	M_{13}	1.000	0.677	299	0.943	0.045
N1	M_{14}	0.999	0.756	299	0.938	0.036
N1	M_{15}	0.994	0.390	299	0.735	0.110
N2	M_{11}	1.000	0.610	3545	0.952	0.066
N2	M_{12}	1.000	0.715	3545	0.967	0.046
N2	M_{13}	1.000	0.662	4090	0.957	0.056
N2	M_{14}	1.000	0.756	4090	0.950	0.044
N2	M_{15}	1.000	0.379	4090	0.796	0.138

Table 1: Results of the quality tests on the triangular meshes.

Mesh label	Quality measure	Maximum value	Minimum value	Element number	Mean value	Standard deviation
A1	M_{11}	1.000	0.624	5656	0.960	0.055
A1	M_{12}	1.000	0.730	5656	0.973	0.036
A1	M_{13}	1.000	0.685	5656	0.966	0.045
A1	M_{14}	1.000	0.779	2502	0.961	0.036
A1	M_{15}	1.000	0.396	5656	0.821	0.130
A4	M_{11}	0.999	0.620	13	0.960	0.049
A4	M_{12}	0.999	0.726	13	0.972	0.035
A4	M_{13}	0.999	0.678	13	0.964	0.043
A4	M_{14}	0.996	0.773	13	0.949	0.035
A4	M_{15}	0.962	0.391	13	0.791	0.103
FD1	M_{11}	0.318	0.318	17644	0.318	0.000
FD1	M_{12}	0.382	0.382	17644	0.382	0.000
FD1	M_{13}	0.296	0.296	17644	0.296	0.000
FD1	M_{14}	0.535	0.535	17644	0.535	0.000
FD1	M_{15}	0.152	0.152	17644	0.152	0.000
FG5	M_{11}	0.929	0.753	14922	0.855	0.038
FG5	M_{12}	0.950	0.822	14922	0.897	0.027
FG5	M_{13}	0.933	0.772	14922	0.866	0.034
FG5	M_{14}	0.918	0.813	14922	0.872	0.023
FG5	M_{15}	0.686	0.472	14922	0.581	0.045

Table 2: Results of the quality tests on the triangular faces of the wedge and hex-dominated meshes.

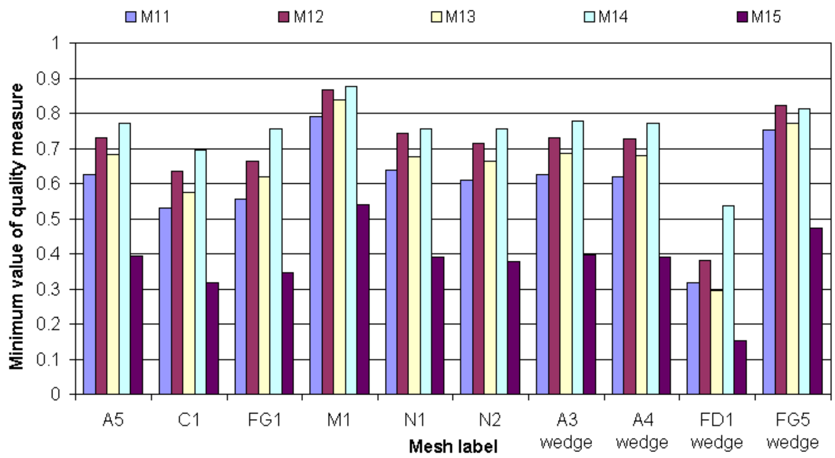

Figure 17: Minimum value of each of the triangular quality measures for all meshes involving triangles or wedges.

triangular mesh of the upper circular face of the cylinder, creating long prisms by projecting this mesh onto the lower circular face of the cylinder, and subdividing the long prisms into shorter elements. Hence the quality of the triangular faces of the mesh is dependent on the quality of the initial triangular mesh, which is two-dimensional. The maximum and minimum values for the FD software are identical because every triangular face is identical, and hence all have the same value for each quality measure and the standard deviation is zero for all measures.

Figure 17 shows that the triangular faces created by FD1 are of significantly lower quality than all the other meshes. This poor quality is caused by the mesh generation method being designed for cylindrical polar coordinate calculations. The stiffness matrices of such calculations would not be adversely affected by the low quality measures. This example shows that all the details of the calculation must be considered when constructing a mesh and when considering its quality. The calculation method for which this mesh has been designed takes advantage of the cylindrical shape and the symmetry of the mesh to improve computational efficiency.

5.3 Quadrilateral quality measure results

Table 3 lists the test results for all quadrilateral meshes. Table 4 shows the same results for the quadrilateral faces of the hexahedral elements from hexahedral and hex-dominated meshes, and table 5 shows values for wedge elements from wedge and hex-dominated meshes.

The first point to note from table 3 is that the minimum value of quality measure M_{23} is very small for mesh FG2. M_{23} is a measure of the degree of skewing or tapering and is

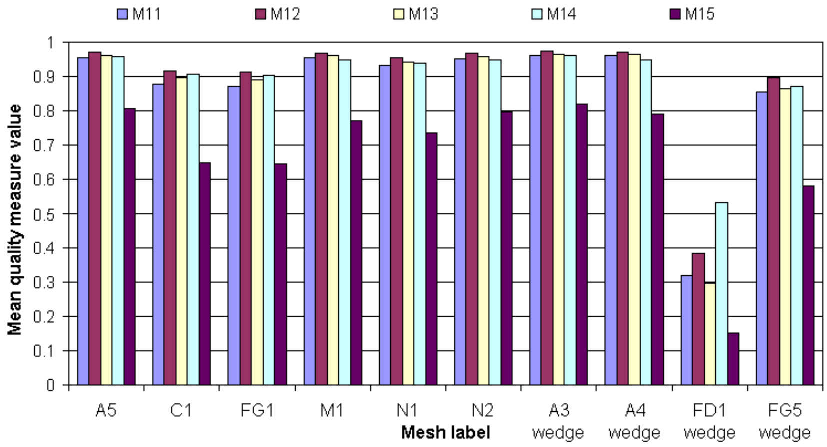

Figure 18: Mean value of each of the triangular quality measures for all meshes involving triangles or wedges.

Mesh label	Quality measure	Maximum value	Minimum value	Element number	Mean value	Standard deviation
FG2	M_{21}	1.000	1.000	1	1.000	0.000
FG2	M_{22}	0.999	0.915	1914	0.973	0.020
FG2	M_{23}	0.999	0.026	1936	0.673	0.249
FG3	M_{21}	1.000	1.000	1	1.000	0.000
FG3	M_{22}	0.988	0.204	1667	0.819	0.157
FG3	M_{23}	0.968	0.324	1564	0.789	0.186

Table 3: Results of the quality tests on the quadrilateral meshes.

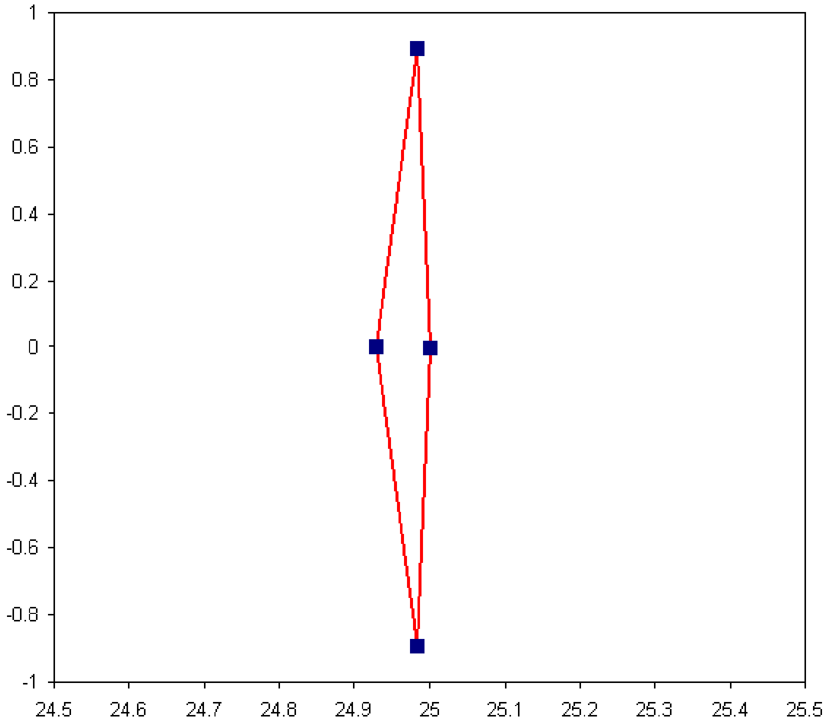

Figure 19: Quadrilateral with a low value of M_{23}, created as part of mesh FG2.

determined from the internal angles of the element. The element that has generated this value is shown in figure 19, and is clearly very different from the ideal square element. If this mesh were used in a calculation it is likely that the results would be inaccurate, particularly in the region immediately surrounding the element. This example illustrates the difficulty of meshing a curved shape with just quadrilaterals.

FG2 was generated using a mapped algorithm, which implies that the elements have been generated by drawing lines between parts of the boundary that are "opposite" one another in some sense. The boundary of the example here has no corners, which is likely to lead to generation of elements similar to that shown in figure 19, since the angle between regions of the boundary that are separate but not opposite is close to 180°.

Mesh FG3 has quite a low minimum value for M_{22}, the measure of aspect ratio. Some FE packages issue warnings for values of M_{22} between 0.2 and 0.07, and error messages for anything below that, so the worst element is slightly better than that warning limit. Another point of interest from table 3 is that each minimum value is generated by a different element, unlike the triangular quality measures. The reason for this is that each of the quadrilateral quality measures checks for a different type of distortion, and so it is comparatively unlikely that a single element should be the worst in all three categories.

The hexahedral and wedge element faces show similar behaviour to the quadrilateral meshes.

Mesh label	Quality measure	Maximum value	Minimum value	Element number	Mean value	Standard deviation
A1	M_{21}	1.000	1.000	9	1.000	0.000
A1	M_{22}	1.000	0.231	11433	0.854	0.121
A1	M_{23}	1.000	0.329	3159	0.938	0.124
A4	M_{21}	1.000	1.000	10719	1.000	0.000
A4	M_{22}	1.000	0.418	11987	0.829	0.110
A4	M_{23}	1.000	0.393	2174	0.946	0.103
FD1	M_{21}	1.000	1.000	291	1.000	0.000
FD1	M_{22}	1.000	0.115	1733	0.540	0.370
FD1	M_{23}	1.000	0.895	1337	0.965	0.050
FG4	M_{21}	1.000	1.000	4327	1.000	0.000
FG4	M_{22}	1.000	0.500	494	0.900	0.126
FG4	M_{23}	1.000	0.020	2705	0.899	0.208

Table 4: Results of the quality tests on the quadrilateral faces of hexahedral elements.

Mesh label	Quality measure	Maximum value	Minimum value	Element number	Mean value	Standard deviation
A3	M_{21}	1.000	1.000	12940	1.000	0.000
A3	M_{22}	1.000	0.349	4000	0.852	0.139
A3	M_{23}	1.000	1.000	1	1.000	0.000
A4	M_{21}	1.000	1.000	4	1.000	0.000
A4	M_{22}	0.999	0.475	229	0.738	0.128
A4	M_{23}	1.000	1.000	1	1.000	0.000
FD1	M_{21}	1.000	1.000	1	1.000	0.000
FD1	M_{22}	1.000	0.174	17644	0.725	0.389
FD1	M_{23}	1.000	1.000	1	1.000	0.000
FG5	M_{21}	1.000	1.000	10818	1.000	0.000
FG5	M_{22}	0.998	0.348	50	0.771	0.129
FG5	M_{23}	1.000	1.000	12198	1.000	0.000

Table 5: Results of the quality tests on the quadrilateral faces of wedge elements.

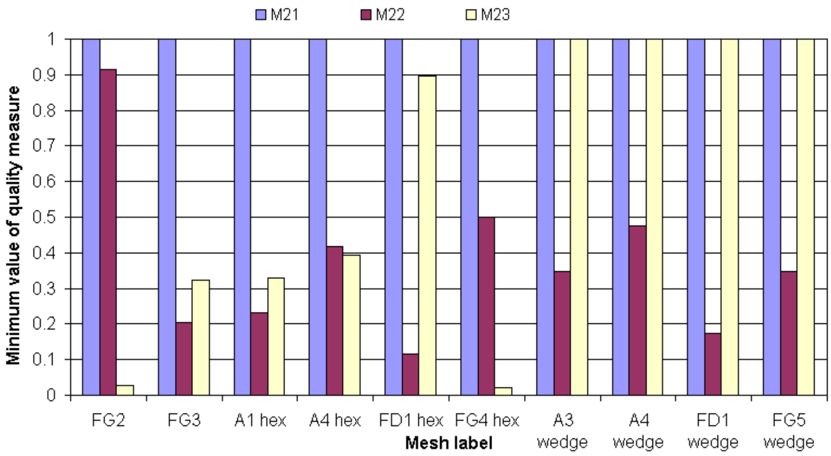

Figure 20: Minimum value of each of the quadrilateral quality measures for all meshes involving quadrilaterals, hexahedra or wedges.

In particular, the minimum value of M_{23} for mesh FG4 is very small because the mesh was created by using a mapped method on the upper surface of the cylinder and so an element face similar to that in figure 19 is produced. It is interesting to note that no out-of-plane warping has been produced. This is an effect of the meshing methods and the problem geometry. Since a right-angled cylinder was used, the meshing method generates prisms based on a planar front face with other faces perpendicular to the front face, and so no out-of-plane distortion will ensue.

Figure 20 plots the minimum values of each of the quality measures for all relevant meshes, and figure 21 plots the mean values. As with the triangular measures on the wedge elements, mesh FD1 generally has lower values of M_{22} than the other packages due to its basis in cylindrical polar coordinates.

5.4 Tetrahedral quality measures

Table 6 shows the values of the tetrahedral quality measures, and figures 22 and 23 show the minimum and mean values respectively. The values suggest that the two meshes are of similar quality since their mean, standard deviation, and minimum values are broadly the same. These results also offer an opportunity to consider the behaviour of the different tetrahedral quality measures, which are difficult to interpret by examining the formulae.

The average and minimum values of both meshes have $M_{32} \geq M_{34} \geq M_{31} \geq M_{33}$. Further analysis of the elements within the meshes suggests that around 90% of the elements have this ordering of the quality measures, and the exceptions tend to have approximately equal values for the measures that are not in this order. In general, the values of M_{31}, M_{32}, and M_{34} are quite close to one another.

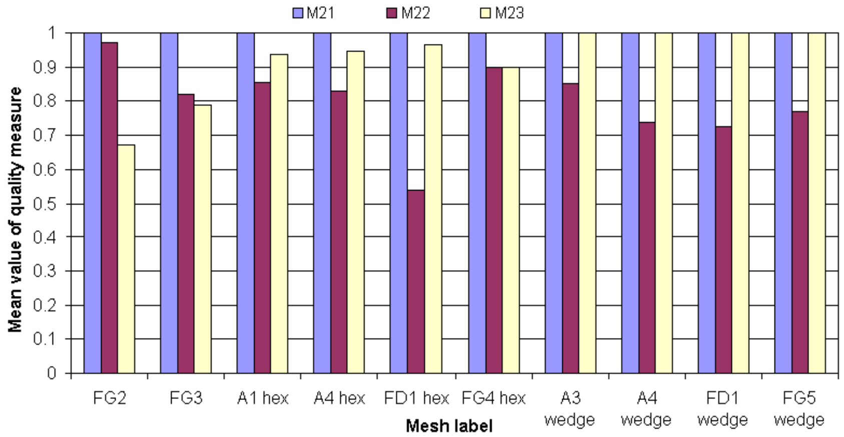

Figure 21: Mean value of each of the quadrilateral quality measures for all meshes involving quadrilaterals, hexahedra or wedges.

A2	M_{31}	0.999	0.327	30706	0.778	0.103
A2	M_{32}	0.997	0.507	95318	0.857	0.064
A2	M_{33}	0.984	0.203	95318	0.584	0.117
A2	M_{34}	0.999	0.397	30706	0.816	0.089
C1	M_{31}	0.998	0.294	76384	0.788	0.107
C1	M_{32}	0.995	0.465	6783	0.863	0.065
C1	M_{33}	0.974	0.184	10504	0.592	0.120
C1	M_{34}	0.999	0.341	6783	0.823	0.093

Table 6: Results of the quality tests on the tetrahedral meshes.

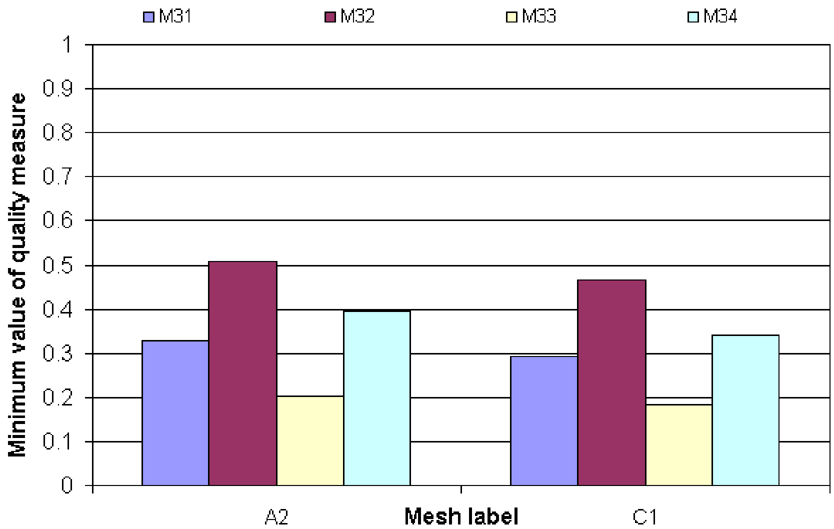

Figure 22: Minimum value of each of the quality measures for tetrahedral meshes.

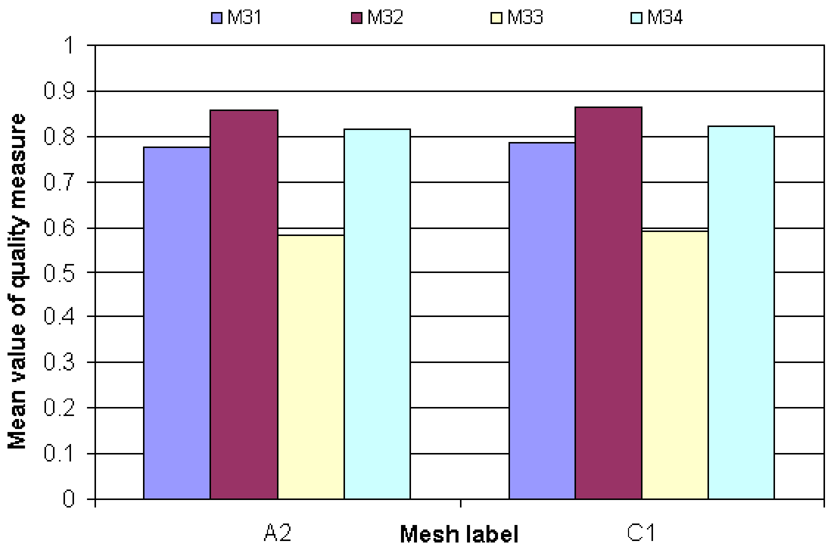

Figure 23: Mean value of each of the quality measures for tetrahedral meshes.

It is interesting to note that the measures generally found to be smallest for tetrahedra and triangles (M_{15} and M_{33}) are those designed to identify elements with ill-conditioned stiffness matrices. This may indicate a need to look beyond geometric parameters when assessing mesh quality.

5.5 Area and volume results

The area and volume results showed that all the packages generated meshes that were within 0.5% of the target area/volume, and that all except one package generated meshes that were within 0.05% of the target. The results are not listed in full because they are so small as to be negligible. The in-house package FD showed the largest discrepancy (0.5%). The discrepancy occured because the FD package is designed to generate meshes for use with polar coordinates rather than Cartesian grids. In cylindrical polar coordinates, the FD mesh has exactly the target volume, which is true of all of the meshes: the discrepancy between the target volume and the mesh volume is caused by linear approximations to a circular boundary.

5.6 Conclusions

The majority of the automatically generated meshes were of reasonably good quality. In particular, the triangular and tetrahedral meshes were generally of good quality. The testing of a mesh designed for cylindrical polar coordinates showed that the coordinate system needs to be considered during quality tests, since its elements performed poorly according to the measures used, despite being adequate for the intended purpose.

There was a high degree of correlation between the various measures used to assess triangular elements, as there was between the quality measures used to assess tetrahedra. This correlation may mean that t is reasonable to assess triangular and tetrahedral meshes using a single measure. There was no correlation between the measures used for quadrilateral (and by extension hexahedral and wedge) elements, since each of the measures looks for a different type of distortion.

The quadrilateral meshes were considerably less well-behaved than the triangular meshes. In particular, some of the elements around the perimeter of the circle were very distorted and had areas close to zero. Similarly, some of the hexahedral meshes generated by the same software package were distorted. It is likely that distortions like these would be identified by a good pre-processor, but often warning messages for distorted elements are overlooked or ignored.

All the packages that were tested generated meshes that were with 0.5 % of the target area or volume. It should be noted that a mesh generated to reproduce an area or volume to a high degree of accuracy without considering element quality can easily contain distorted elements.

6 Application to a mesh used to model a corrosion problem

6.1 Problem definition

The meshes described in this section were used to investigate the stress and strain distributions associated with a single corrosion pit in a cylindrical sample under tension. The basic sample geometry is shown in figure 24. The aim of the work was to identify possible causes of cracks associated with corrosion pits. The work is described in full in a recent NPL report [18], but the main points will be summarised here. Interested readers are referred to [18] for more details of the problem and the model results.

The work looked at two different sample diameters and three different pit geometries, and used a variety of different meshes. All samples were 25.4 mm long. The various properties of the sample geometries and the associated meshes are listed in table 7. The "U-shaped" pit had a diameter of 666 microns, and all hemispherical pits had a diameter of twice their depth. Note that the larger diameter sample was always modelled with a mixed mesh, but that all mixed meshes had more tetrahedra than hexahedra and wedges put together. All meshes were created in Abaqus version 6.5-7.

One of the meshes was created using second-order elements in order to check that the results of the model with linear elements were reliable, but all elements have been treated as linear in the mesh quality tests. This checking of results showed that the linear elements gave the same results as the higher-order elements.

Mesh Label	Element Type	Sample Diameter (mm)	Pit Depth (microns)	Pit Geometry	Number of Nodes	Number of Elements
1	Tetrahedra	6.4	100	Hemisphere	15989	82916
2	Tetrahedra	6.4	100	Hemisphere	33061	173898
3	Tetrahedra	6.4	500	Hemisphere	48732	259342
4	Mixed	32	100	Hemisphere	95371	165453
5	Mixed	32	500	Hemisphere	98639	185670
6	Tetrahedra	6.4	500	U-shaped	60502	325916
7	Tetrahedra	6.4	100	Hemisphere	139922	92161

Table 7: Properties of the various meshes.

Typical meshes of the two different sample sizes are shown in figures 25 and 26. Since the results of most importance were in the immediate vicinity of the corrosion pit, and since this region is also where the stresses and strains would be expected to be changing most rapidly, the mesh density was concentrated around the pit. Figures 25 and 26 are too large-scale to show the detailed mesh around the pit: instead the pit appears as a black region mid-way along the sample. In order to make the results of the different model comparable, every effort was made to ensure that the meshes in the region of the pit used elements of similar sizes independent of the other features of the models. The details of a typical mesh in the

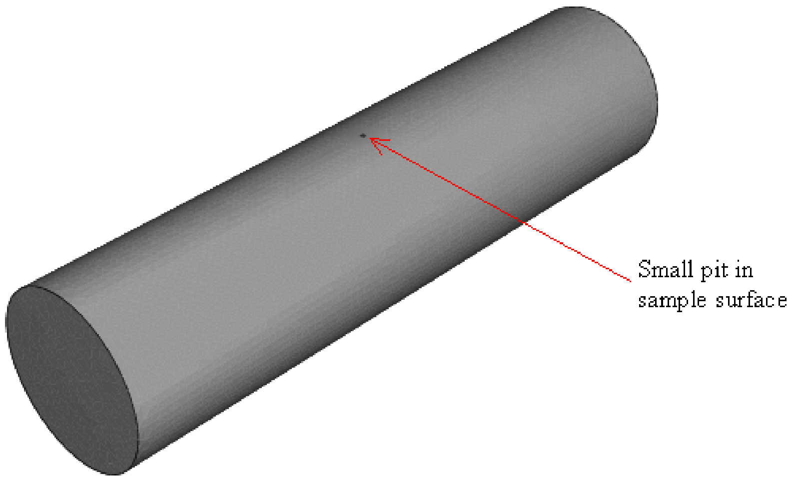

Figure 24: Cylindrical sample with a small pit.

vicinity of the pit are shown in figure 27. This requirement for detail in the vicinity of the pit has led to a range of different element sizes within each mesh. Provided that each of a given element's dimensions are approximately equal, this should not lead to problems: a good finite element mesh does not necessarily have to consist of identically-sized elements.

6.2 Results

The results of the mesh testing are shown in tables 8, 9, and 10. These tables show the results for tetrahedra and quadrilateral and triangular faces of other elements respectively. The mean and minimum values of the tetrahedral quality measures of each mesh are shown in figures 28 and 29.

The mean values of the quality measures indicate that on average, there is little to choose between the various meshes, and they are all of good quality: none of the measures has an average less than 0.5. However, the minimum values plotted in figure 29 show that meshes 1 and 6 include some poor-quality elements. The minimum values for M_{33} for both of these meshes is approximately 0.1, which could lead to conditioning problems.

The results obtained from the tetrahedral meshes were analysed further in order to get a clearer idea of the distribution of values, and in particular whether there were many elements with values close to the minimum. A frequency plot for each measure on each mesh was created by counting how many values lay in each interval of the form $0.01k \leq M \leq 0.01(k+1), k = 0, 1, \ldots, 99$, and then normalising by the total number of elements in each mesh. The results of this process are shown in figure 30.

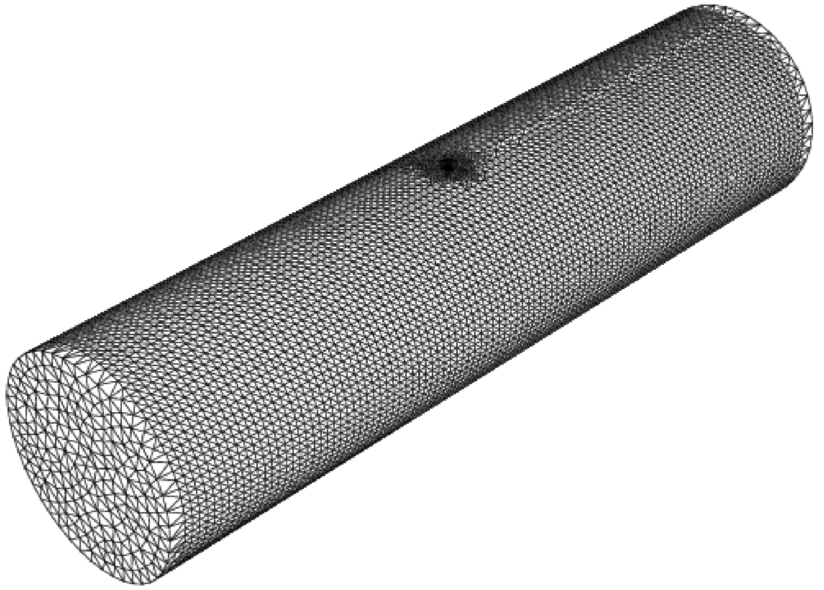

Figure 25: A typical tetrahedral mesh.

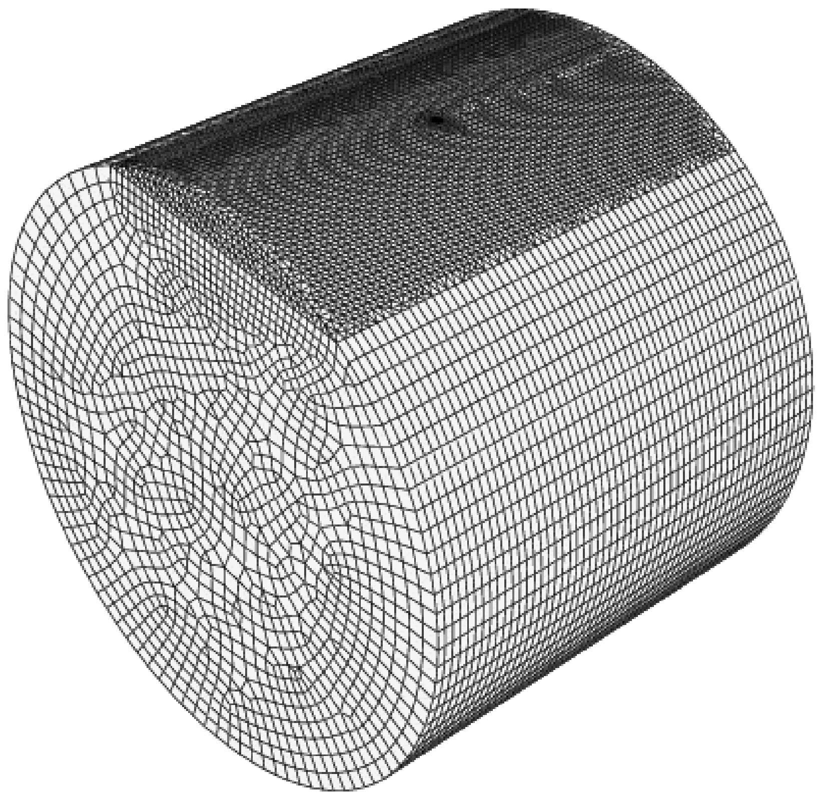

Figure 26: A typical mixed mesh.

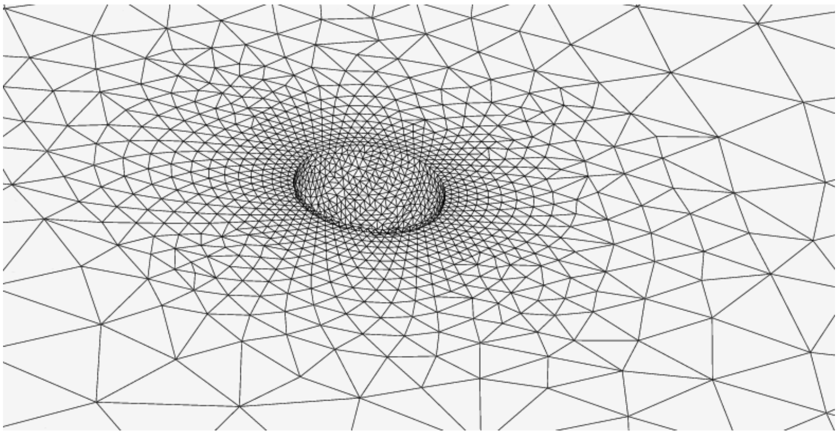

Figure 27: Typical detail of mesh around the pit.

Figure 30 shows that the frequency distributions for each quality measure vary little from mesh to mesh. The values for mesh 7 are slightly lower than those for the other meshes, which is reflected in the mean values for mesh 7 that are listed in table 8 being slightly lower for each quality measure than the mean values of the other meshes. This figure also indicates that the minimum values obtained appear to be isolated instances rather than frequent occurrences.

It is also interesting to note that the quality measures obey $M_{32} \geq M_{34} \geq M_{31} \geq M_{33}$ for most of the elements, as they did for the model problem results. This feature is shown in the mean values, the minimum values, and the frequency distribution. For all of the meshes, the most likely values of the measures are approximately: $M_{32} \approx 0.9$, $M_{34} \approx 0.85$, $M_{31} \approx 0.85$ (with a larger spread than M_{34}), and $M_{32} \approx 0.6$.

The results on the triangular faces of the wedge elements in table 10 show that the minimum values are above 0.3, suggesting that these element faces are unlikely to cause a problem. Mesh 5 seems to have particularly good quality triangular faces.

The quadrilateral faces of the mixed meshes show little or no signs of out-of-plane warping, and the minima of the other quality measures are high enough to mean that the hexahedra and wedge elements are of sufficiently high quality that the solution will not be adversely affected by them. The aspect ratio measure is within acceptable bounds for all elements, and the angle quality measure value is acceptable.

6.3 Conclusions

The meshes used for the finite element work are generally of good quality, with few (if any) badly-formed elements that could lead to inaccurate results. The problem of designing a mesh to give accurate results over a geometry which requires fine detail over a small region within a much larger volume has led to the element sizes varying widely within the mesh,

Mesh number	Quality measure	Maximum value	Minimum value	Element number	Mean value	Standard deviation
1	M_{31}	0.999	0.173	17260	0.771	0.108
1	M_{32}	0.998	0.273	17260	0.852	0.067
1	M_{33}	0.998	0.097	17260	0.580	0.119
1	M_{34}	0.999	0.193	17260	0.809	0.093
2	M_{31}	0.998	0.265	117468	0.772	0.106
2	M_{32}	0.995	0.410	101143	0.853	0.066
2	M_{33}	0.999	0.159	101143	0.583	0.122
2	M_{34}	0.999	0.310	101143	0.809	0.091
3	M_{31}	0.999	0.332	23374	0.774	0.105
3	M_{32}	0.998	0.486	23374	0.855	0.065
3	M_{33}	0.998	0.140	41737	0.584	0.122
3	M_{34}	1.000	0.368	23374	0.811	0.090
4	M_{31}	0.998	0.336	134661	0.763	0.108
4	M_{32}	0.994	0.476	134661	0.848	0.068
4	M_{33}	0.998	0.192	134661	0.578	0.129
4	M_{34}	0.998	0.386	134661	0.802	0.094
5	M_{31}	1.000	0.298	156794	0.762	0.109
5	M_{32}	0.997	0.443	156794	0.847	0.068
5	M_{33}	0.999	0.173	150410	0.571	0.125
5	M_{34}	1.000	0.334	156794	0.800	0.094
6	M_{31}	0.999	0.168	5217	0.773	0.105
6	M_{32}	0.998	0.282	100853	0.854	0.065
6	M_{33}	0.999	0.106	100853	0.583	0.122
6	M_{34}	0.999	0.196	100853	0.810	0.090
7	M_{31}	0.998	0.223	17176	0.746	0.112
7	M_{32}	0.995	0.514	17176	0.840	0.069
7	M_{33}	0.998	0.204	10380	0.560	0.122
7	M_{34}	0.999	0.351	17176	0.787	0.096

Table 8: Quality measure results for the tetrahedral elements in all meshes.

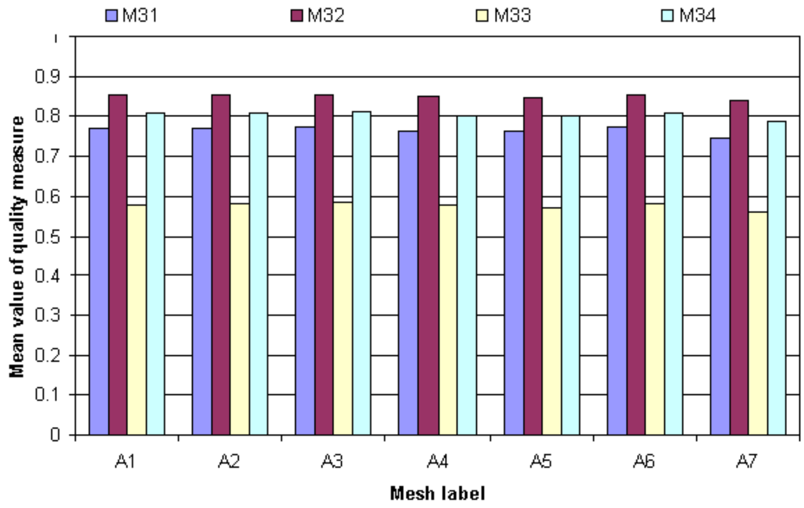

Figure 28: Mean values of the quality measures calculated for the tetrahedral elements within each mesh.

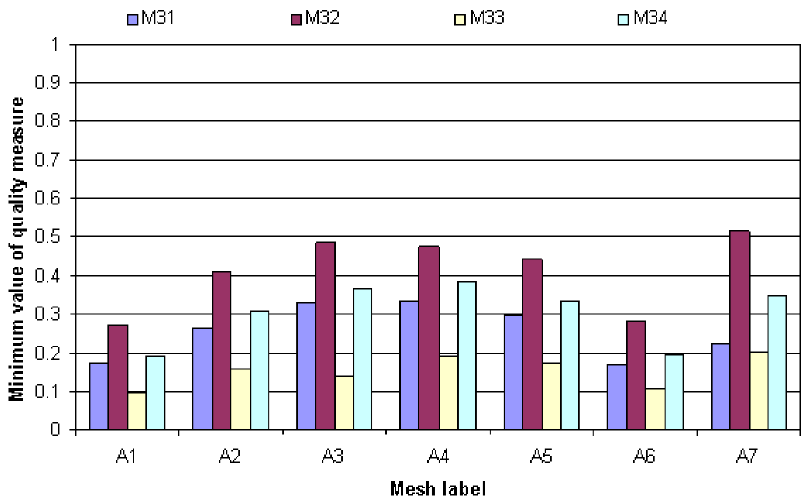

Figure 29: Minimum values of the quality measures calculated for the tetrahedral elements within each mesh.

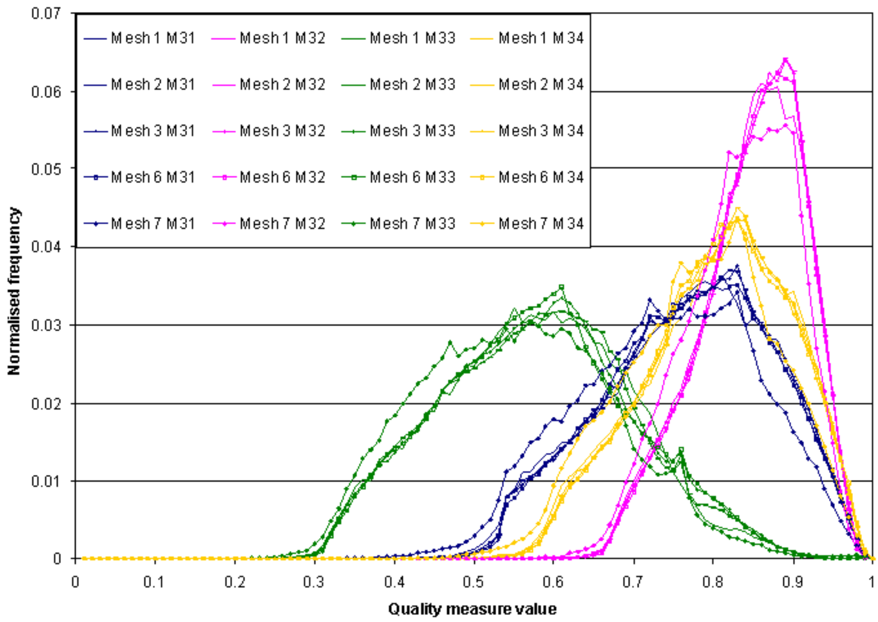

Figure 30: Frequency distribution of the tetrahedral quality measures on each of the tetrahedral meshes.

Mesh number	Quality measure	Maximum value	Minimum value	Element number	Mean value	Standard deviation
4W	M_{21}	1.000	1.000	3	1.000	0.000
4W	M_{22}	0.994	0.440	38	0.685	0.143
4W	M_{23}	1.000	0.994	34	0.999	0.001
4H	M_{21}	1.000	0.997	72940	1.000	0.000
4H	M_{22}	0.999	0.344	66167	0.699	0.180
4H	M_{23}	1.000	0.360	3180	0.936	0.114
5W	M_{21}	1.000	1.000	8	1.000	0.000
5W	M_{22}	0.978	0.494	1146	0.699	0.137
5W	M_{23}	1.000	0.960	1159	0.997	0.007
5H	M_{21}	1.000	0.960	4922	1.000	0.002
5H	M_{22}	1.000	0.323	64627	0.691	0.177
5H	M_{23}	1.000	0.351	8175	0.936	0.114

Table 9: Quality measure results for the quadrilateral faces of elements in the mixed meshes. W denotes wedge elements, H denotes hexahedral elements.

Mesh number	Quality measure	Maximum value	Minimum value	Element number	Mean value	Standard deviation
4	M_{11}	0.998	0.605	17	0.951	0.065
4	M_{12}	0.999	0.712	17	0.965	0.048
4	M_{13}	0.998	0.666	17	0.955	0.057
4	M_{14}	0.994	0.769	17	0.942	0.043
4	M_{15}	0.945	0.381	17	0.771	0.111
5	M_{11}	1.000	0.882	73	0.958	0.034
5	M_{12}	1.000	0.915	109	0.971	0.024
5	M_{13}	1.000	0.886	109	0.961	0.032
5	M_{14}	0.995	0.879	109	0.945	0.033
5	M_{15}	0.973	0.606	109	0.779	0.101

Table 10: Quality measure results for the triangular faces of wedge elements in the mixed meshes.

but none of the individual elements have been catastrophically distorted.

7 Summary

This report has investigated quality measures for assessment of FE meshes, and has tested meshes generated by some popular mesh generation packages. The work was motivated by the increased availability of automatic mesh generation tools, which has led to more widespread usage of FE by non-expert users.

The main requirements of a good mesh are that it must approximate the problem solution (and usually its gradient) well, and that the matrix of linear equations generated during the model solving process must be well-conditioned. Different quality measures are available to address these requirements. The measures make it possible to compare different meshes. Many FE packages supply a quality estimate for their meshes, but it is not always clear whether the suggested measure assesses functional approximation, gradient approximation, or matrix conditioning.

The software packages that were tested generally created good quality elements, particularly when triangular or tetrahedral elements were used. This result may be due to the simplicity of the test shape. The poorest quality elements were generated when the entire shape was mapped with quadrilaterals, leading to some elements with very large internal angles.

There was found to be a high degree of correlation between the various measures used to assess triangular elements, and between those used to assess tetrahedral elements. Within a given mesh, the element that had the lowest value of one of the measures would have the lowest value of all the measures. This correlation suggests that elements that approximate functions poorly also approximate gradients poorly, and can lead to ill-conditioned matrices.

The meshes used in modelling a corrosion problem were all found to be fit for purpose. The geometry used in this model was challenging, as it needed to accurately simulate a micrometre-sized defect in a millimetre-sized object.

It is expected that a software tool for automatic mesh quality assessment will be developed as part of the next SS*f*Mprogramme. This tool will allow users to assess FE meshes before or after solution of a problem, and will help users to improve their mesh quality and thus their solution accuracy.

References

[1] http://www.comsol.com/.

[2] http://www.abaqus.com/products/products_overview.html.

[3] http://www.femsys.co.uk/.

[4] http://www.nag.co.uk/.

[5] http://www.mathworks.co.uk/.

[6] I. Babuška and A. K. Aziz. On the angle condition in the finite element method. *SIAM J. Numer. Anal..*, 13(2):214–226, 1976.

[7] R. M. Barker, P. M. Harris, and L. Wright. Good Practice Guide No. 16. Testing algorithms and software. Technical report, National Physical Laboratory, Teddington, UK, 2005. Available for download from the SS*f*M website at www.npl.co.uk/ssfm/download/bpg.html.

[8] M. Berzins. A solution-based triangular and tetrahedral mesh quality indicator. *SIAM J. Sci. Comput.*, 19(6):2051–2060, 1998.

[9] J. C. Cavendish, D. A. Field, and W. H. Frey. An approach to automatic three-dimensional finite element mesh generation. *Int. J. Num. Meth. Eng.*, 21:329–347, 1985.

[10] T. J. Esward, E. G. Johnson, and L. Wright. Testing continuous modelling software. Technical report, National Physical Laboratory, Teddington, UK, 2004. Available for download from the SS*f*M website at www.npl.co.uk/ssfm/download/bpg.html.

[11] T. J. Esward, K. Lees, D. Sayers, and L. Wright. Testing continuous modelling software: Three case studies. Technical report, National Physical Laboratory, Teddington, UK, 2004. Available for download from the SS*f*M website at www.npl.co.uk/ssfm/download/bpg.html.

[12] T. J. Esward, G. Lord, and L. Wright. Model validation in continuous modelling. Technical Report CMSC 29/03, National Physical Laboratory, Teddington, UK, 2003.

[13] T. J. Esward and L. Wright. Guide to the use of finite element and finite difference software. Technical Report CMSC 30/03, National Physical Laboratory, Teddington, UK, 2003.

[14] T. J. Esward and L. Wright. Good Practice Guide No. 15. Continuous modelling. Technical report, National Physical Laboratory, Teddington, UK, 2006. Available for download from the SS*f*M website at www.npl.co.uk/ssfm/download/bpg.html.

[15] S. Owen. A survey of unstructured mesh generation technology. In *Seventh International Meshing Roundtable*, pages 239–267, 1998. Available for download from http://www.andrew.cmu.edu/user/sowen/survey/index.html.

[16] P. P. Pébay and T. J. Baker. Analysis of triangle quality measures. *Mathematics of Computation*, 72(244):1817–1839, 2003.

[17] J. R. Shewchuck. What is a good linear element? Interpolation, conditioning, and quality measures. In *Eleventh International Meshing Roundtable (Ithaca, New York)*, pages 115–126, 2002. Available for download in full and as a shorter version from http://www.cs.cmu.edu/j̃rs/jrspapers.html.

[18] A. Turnbull and L. Crocker. Finite element analysis of the stress and strain associated with a corrosion pit in a cylindrical specimen under applied tensile loading. Technical report, National Physical Laboratory, Teddington, UK, 2006. NPL Report DEPC-MPE 034.