

# Finite Element Analysis Methods

---

## 1.1. Introduction

The finite element method (FEM) rapidly grew as the most useful numerical analysis tool for engineers and applied mathematicians because of its natural benefits over prior approaches. The main advantages are that it can be applied to arbitrary shapes in any number of dimensions. The shape can be made of any number of materials. The material properties can be non-homogeneous (depend on location) and/or anisotropic (depend on direction). The way that the shape is supported (also called fixtures or restraints) can be quite general, as can the applied sources (forces, pressures, heat flux, etc.). The FEM provides a standard process for converting governing energy principles or governing differential equations into a system of matrix equations to be solved for an approximate solution. For linear problems, such solutions can be very accurate and quickly obtained. Having obtained an approximate solution, the FEM provides additional standard procedures for follow up calculations (post-processing), such as determining the integral of the solution, or its derivatives at various points in the shape. The post-processing also yields impressive color displays, or graphs, of the solution and its related information. Today, a second post-processing of the recovered derivatives can yield error estimates that show where the study needs improvement. Indeed, adaptive procedures allow automatic corrections and re-solutions to reach a user specified level of accuracy.

However, very accurate and pretty solutions of models that are based on errors or incorrect assumptions are still wrong.

When the FEM is applied to a specific field of analysis (like stress analysis, thermal analysis, or vibration analysis) it is often referred to as finite element analysis (FEA). An FEA is the most common tool for stress and structural analysis. Various fields of study are often related. For example, distributions of non-uniform temperatures induce non-obvious loading conditions on solid structural members. Thus, it is common to conduct a thermal FEA to obtain temperature results that in turn become input data for a stress FEA. FEA can also receive input data from other tools like motion (kinetics) analysis systems and computation fluid dynamics systems.

## 1.2. Basic Integral Formulations

The basic concept behind the FEM is to replace any complex shape with the union (or summation) of a large number of very simple shapes (like triangles) that are combined to correctly model the original part. The smaller simpler shapes are called finite elements because each one occupies a small but finite sub-domain of the original part. They contrast to the infinitesimally small or differential elements used for centuries to derive differential equations. To give a very simple example of this dividing and summing process, consider calculating the area of the arbitrary shape shown in Figure 1.1 (left).

If you knew the equations of the bounding curves you, in theory, could integrate them to obtain the enclosed area. Alternatively, you could split the area into an enclosed set of triangles (cover the shape




Fig. 1.1. An area crudely meshed with linear and quadratic triangles.

with a mesh) and sum the areas of the individual triangles:

$$A = \sum_{e=1}^n A^e = \sum_{e=1}^n \int_{A^e} dA.$$

Now, you have some choices for the type of triangles. You could pick straight sided (linear) triangles, or quadratic triangles (with edges that are parabolas), or cubic triangles, etc. The area of a straight-sided triangle is a simple algebraic expression. Number the three vertices in a counter-clockwise order, then the area is  $A^e = [x_1(y_2 - y_3) + x_2(y_3 - y_1) + x_3(y_1 - y_2)]/2$  and its centroid is located at

$$x_{cg}^e = \frac{x_1 + x_2 + x_3}{3}, \quad y_{cg}^e = \frac{y_1 + y_2 + y_3}{3}.$$

Similar expressions give the moment of inertia components. Thus, you just have to extract (gather) the element vertices coordinates from the mesh data in order to compute the area of a straight sided triangle. The area of a curved triangle is also relatively easy to calculate by numerical integration, but is computationally more expensive to obtain than that for the linear triangle. The first two triangle mesh choices are shown in Figure 1.1 for a large element size. Clearly, the simple straight-sided triangular mesh (on the left) approximates the area very closely, but at the same time introduces geometric errors along a curved boundary. The boundary geometric error in a linear triangle mesh results from replacing a boundary curve by a series of straight line segments. That geometric boundary error can be reduced to any desired level by increasing the number of linear triangles. But that decision increases the number of calculations and makes you trade off geometric accuracy versus the total number of required area calculations and summations.

Area is a scalar, so it makes sense to be able to simply sum its parts to determine the total value, as shown above. Other physical quantities, like kinetic energy, strain energy and mechanical work, can be summed in the same fashion. Indeed, the very first applications of FEA to structures was based on minimizing the energy stored in a linear elastic material. The FEM always involves

some type of governing integral statement which is converted to a matrix system by assuming how items vary within a typical element. That integration is also converted to the sum of the integrals over each element in the mesh. Even if you start with a governing differential equation, it gets converted to an equivalent integral formulation by one of the methods of weighted residuals (MWR). The two most common weighted residual methods, for FEA, are the Galerkin method and the Method of Least Squares.

The development of the necessary matrix relations will be covered in more detail later. For now the matrix representation of the kinetic energy of a part is presented, for the straight sided triangle element, as an example. Recall that the kinetic energy of a mass particle is  $KE = mv^2/2$ , where  $m$  is the mass and  $v$  is its velocity. The kinetic energy of the planar body, of thickness  $t$ , in Figure 1.1 is obtained by integrating over the differential masses

$$KE = \frac{1}{2} \int v^2 dm = \frac{1}{2} \int v^2 \rho dV = \frac{t}{2} \int v^2 \rho dA$$

and it can be obtained by summing the element integrals

$$KE = \sum_{e=1}^n KE^e = \frac{t}{2} \sum_{e=1}^n \int_{A^e} v^2 \rho dA.$$

Assume that the velocity of all mesh points is known. Then you can extract the velocity of the three vertices of each linear triangle element. In FEA you define the velocity of any point inside the element by interpolating between the element's vertex values. Denote the interpolation function of node  $j$  by  $N(x, y)_j$ . Then the velocity is

$$\begin{aligned} v(x, y) &= N(x, y)_1 \cdot v_1 + N(x, y)_2 \cdot v_2 + N(x, y)_3 \cdot v_3 \\ &= \sum_{j=1}^3 N(x, y)_j \cdot v_j. \end{aligned}$$

This linear spatial interpolation is usually written in a matrix notation

$$v(x, y) = [N(x, y)]\{v^e\} = \{v^e\}^T [N(x, y)]^T$$

so the element kinetic energy is a matrix integral

$$KE^e = \frac{1}{2} \int_{A^e} \{v^e\}^T [N(x, y)]^T \rho [N(x, y)] \{v^e\} t dA$$

but the nodal velocities are known constants (or functions of time alone) that move outside the integral

$$KE^e = \frac{1}{2} \{v^e\}^T \int_{A^e} [N(x, y)]^T \rho [N(x, y)] t dA \{v^e\}$$

so that the remaining square matrix integral is called the mass matrix,  $[m^e]$ :

$$KE^e = \frac{1}{2} \{v^e\}^T [m^e] \{v^e\}, \quad [m^e] = \int_{A^e} [N(x, y)]^T \rho [N(x, y)] t dA.$$

This result shows the similarity to a particle in that the body's kinetic energy is half the product of a mass matrix and the square (pre- and post-multiplication) of the nodal velocities. For the straight two-node bar element of length  $L^e$  and cross-sectional area  $A^e$ , the mass matrix is

$$[m^e] = \frac{\rho L^e A^e}{6} \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}.$$

The total element mass,  $m = \rho L^e A^e$ , is present but somehow distributed (coupled) between its two nodes, and their velocities,  $\{v^e\}^T = [v_1^e \ v_2^e]$ . If the two end velocities are the same (it is moving as a rigid body such that  $v_1^e = v_2^e = v$ ) then the matrix products yield the expected scalar answer,  $KE^e = mv^2/2$ . The point of this illustration is to show that any FEA converts scalar integrals to matrix expressions by assuming a spatial interpolation between the nodes of a typical element for items of interest, such as positions, displacements, velocities, or temperatures. Those spatial interpolations also define the spatial gradients that occur in most finite element integral forms.

### 1.3. Gather and Scatter Operators

An integral evaluation for an FEA requires a mesh. Typically it is a triangular mesh for surfaces and a tetrahedral mesh for solids.



The reverse of a gather operation is the scatter or assembly operation. It is a partial summation of element data to the matrices associated with the mesh data. A scatter takes something associated with the local nodes of an element and adds them to the corresponding matrix item at the full mesh level. Scatter, or assembly, operations fill the entries in the matrix equations that must be solved for the problem unknowns. These two common operations are sketched graphically in Figure 1.2. The required element connectivity data (the two nodes on each linear line element) are displayed in the third line of that figure. Gather and scatter will be illustrated in detail in the section on compound elastic bars.

## 1.4. Geometric Boundary Errors

You may think that the geometric boundary error cited for the linear triangles is eliminated by choosing to use the mesh of curved quadratic triangles (on the right of Figure 1.1). The parabola segments pass through three points lying exactly on the boundary curve, but can degenerate to straight lines in the interior. (To speed plotting of small elements, most systems draw all the parabolas as two straight line segments, as on the right in Figure 1.1.) Thus, the boundary shape error is indeed reduced, at the expense of more complicated area calculations, but it is not eliminated. Some geometric error remains because most engineering curves are circular arcs, splines, or nurbs (non-uniform rational B-splines) and thus are not matched by a parabola. The most common way to reduce mesh geometric error is to simply use many smaller elements. The default element choice in SolidWorks (SW) Simulation is the quadratic element. Other systems offer a wider range of edge polynomial degree (e.g., cubic), as well as other shapes like quadrilaterals or rectangles. In three-dimensional solid applications some systems offer dozens of choices for the edge degree polynomial order, and shapes including hexahedral, wedges, and tetrahedral elements. Hexahedral elements are generally more accurate, but can be more challenging to mesh automatically. Tetrahedral elements can match hexahedral element performance by using more (smaller) elements, and tetrahedral




Fig. 1.3. Linear or parabolic elements never match circular shapes.

elements are much easier to mesh automatically. SW Simulation uses only tetrahedral elements for solid studies. An example of the small two-dimensional geometric boundary error due to different curved shapes is seen in Figure 1.3 where a arc and a parabola pass through the same three points.

## 1.5. Stages of Analysis and Their Uncertainties

An FEA always involves a number of uncertainties that impact the accuracy or reliability of each stage of an FEA and its results. The book, *Building Better Products with Finite Element Analysis* by Adams and Askenazi [1] gives an outstanding detailed description of most of the real-world uncertainties associated with solid mechanics FEA. All engineers conducting stress studies should read it. That book also points out how poor solid modeling skills can adversely affect the ability to construct meshes for any type of FEA. Here, the most important FEA uncertainties are highlighted.

The typical stages of an FEA study are listed below:

- (1) Construct the part(s) in a solid modeler. It is surprisingly easy to accidentally build flawed models with tiny lines, tiny surfaces or tiny interior voids. The part will look fine, except with extreme zooms, but it may fail to mesh. Most systems have checking routines that can find and repair such problems before you move on to an FEA study. Sometimes you may have to export a part, and then import it back with a new name because imported parts are usually subjected to more time consuming checks than “native” parts. When multiple parts form an assembly, always mesh and study the individual parts before studying the assembly. Try to plan ahead and introduce

split lines into the part to aid in mating assemblies and to locate load regions and restraint (or fixture or support) regions. Today, construction of a part is probably the most reliable stage of any study.

- (2) Defeature the solid part model for meshing. The solid part may contain features, like a raised logo, that are not necessary to manufacture the part, or required for an accurate analysis study. They can be omitted from the solid used in the analysis study. That is a relative easy operation supported by most solid modelers (such as the “suppress” option in SW) to help make smaller and faster meshes. However, it has the potential for introducing serious, if not fatal, errors in a following engineering study. This is a reliable modeling process, but its application requires engineering judgment. For example, removing small radius interior fillets can greatly reduce the number of elements and simplifies the mesh generation. But, that creates sharp reentrant corners that can yield false infinite stresses. Those false high stress regions may cause you to overlook other areas of true high stress levels. Small holes lead to many small elements (and long run times). They also cause stress concentrations that raise the local stress levels by a factor of three or more. The decision to defeature them depends on where they are located in the part. If they lie in a high stress region you must keep them. But defeaturing them is allowed if you know they occur in a low stress region. Such decisions are complicated because most parts have multiple possible loading conditions and a low stress region for one load case may become a high stress region for another load case.
- (3) Combine multiple parts into an assembly. Again, this is well automated and reliable from the geometric point of view and assemblies “look” as expected. However, geometric mating of part interfaces is very different for defining their physical (displacement, or temperature) mating. The physical mating choices are often unclear and the engineer may have to make a range of assumptions, study each, and determine the worst case result. Having to use physical contacts makes the linear

problem require iterative solutions that take a long time to run and might fail to converge.

- (4) Select the element type. Some FEA systems have a huge number of available element types (with underlying theoretical restrictions). The SolidWorks system has only the fundamental types of elements. Namely, truss elements (bars), frame elements (beams), thin shells (or flat plates), thick shells, and solids. The SW simulation system selects the element type (beginning in 2009) based on the shape of the part. The user is allowed to convert a non-solid element region to a solid element region, and vice versa. Knowing which class of element will give a more accurate or faster solution requires training in finite element theory. At times a second element type study is used to help validate a study based on a different element type.
- (5) Mesh the part(s) or assembly, remembering that the mesh solid may not be the same as the part solid. A general rule in an FEA is that your computer never has enough speed or memory. Sooner or later you will find a study that you cannot execute. Often that means you must utilize a crude mesh (or at least crude in some region) and/or invoke the use of symmetry or anti-symmetry conditions. Local solution errors in a study are proportional to the product of the local element size and the gradient of the secondary variables (i.e., gradient of stress or heat flux). Therefore, you exercise mesh control to place small elements where your engineering judgment estimates high stress (or flux) regions, as well as large elements in low stress regions. The local solution error also depends on the relative sizes of adjacent elements. You do not want skinny elements adjacent to big ones. Thus, automatic mesh generators have options to gradually vary adjacent element sizes from smallest to biggest.

The solid model sent to the mesh generator frequently should have load or restraint (fixture) regions formed by split lines, even if such splits are not needed for manufacturing the parts. The mesh typically should have refinements at source or load regions and support regions.

A mesh must look like the part, but that is not sufficient for a correct study. A single layer of elements filling a part region is almost never enough. If the region is curved, or subjected to bending, you want at least three layers of quadratic elements, but five is a desirable lower limit. For linear elements you at least double those numbers.

Most engineers do not have access to the source code of their automatic mesh generator. When the mesher fails, you frequently do not know why it failed or what to do about it. Often you have to re-try the mesh generation with very large element sizes in hopes of getting some mesh results that can give hints as to why other attempts failed. The meshing of assemblies often fails. Usually the mesher runs out of memory because one or more parts had a very small, often unseen, feature that causes a huge number of tiny elements to be created (see the end of this chapter). You should always attempt to mesh each individual part to spot such problems before you attempt to mesh them as a member of an assembly.

Automatic meshing, with mesh controls, is usually simple and fast today. However, it is only as reliable as the modified part or assembly supplied to it. Distorted elements usually do not develop in automatic mesh generators, due to empirical rules for avoiding them. However, distorted elements locations can usually be plotted. If they are in regions of low gradients you can usually accept them.

You should also note that studies involving natural frequencies are influenced most by the distribution of the mass of the part. Thus, they can still give accurate results with meshes that are much cruder than those that would be acceptable for stress or thermal studies.

- (6) Assign a linear material to each part. Modern FEA systems have a material library containing the “linear” mechanical, thermal, and/or fluid properties of most standardized materials. They also allow the user to define custom properties. The property values in such tables are often misinterpreted to be more accurate and reliable than they actually are. The reported




Fig. 1.4. Typical distributions of properties of steel (left) and cast iron.

property values are accepted average values taken from many tests. Rarely are there any data about the distribution of test results, or what standard deviation was associated with the tests. Most tests yield results that follow a “bell shaped” curve distribution, or a similar skewed curve. The tests for stainless steel tend to have narrow distributions, like that on the left in Figure 1.4, while the results for cast iron have wider distributions.

When you accept a tabulated property value as a single number to be used in the FEA calculation remember it actually has a probability distribution associated with it. You need to assign a contribution to the total factor of safety to allow for variations from the tabulated property value.

The values of properties found in a material table can appear more or less accurate depending on the units selected. That is an illusion often caused by converting one set of units to another, but not truncating the result to the same number of significant figures available in the actual test units. For example, the elastic modulus of one steel is tabulated from the original test as 210 MPa, but when displayed in other units it shows as 30,457,924.92 psi. Which one do you believe to be the experimental accurate value; the 3 digit value or the 10 digit one? The answer affects how you should view and report stress results. The axial stress in a bar is equal to the elastic modulus times the strain,  $\sigma = E\varepsilon$ . Thus, if  $E$  is only known to three

or four significant figures then the reported stress result should have no more significant figures.

Material data are usually more reliable than the loading values (considered next), but less accurate than the model or mesh geometries.

- (7) Select regions of the part(s) to be loaded and assign load levels and load types to each region. In mathematical terminology, load or flux conditions on a boundary region are called Neumann boundary conditions, or non-essential conditions. The geometric regions can be points (in theory), lines, surfaces, or volumes. If they are not existing features of the part, then you should apply split lines to the part to create them before activating the mesh generator. Point forces, or heat sources, are common in undergraduate studies, but in an FEA they cause false infinite stresses, or heat flux. If you include them do not be misled by the high local values. Refining the mesh does not help since the smallest element still reports near infinite values.

In reality, point loads are better modeled as a total force, or pressure, acting over a small area formed by prior split lines. Saint Venant's Principle states that two different, but statically equivalent, force systems acting on a small portion of the surface of a body produce the same stress distributions at distance large in comparison with the linear dimensions of the portion where the forces act. That also implies that concentrated sources quickly become re-distributed, as seen in Figure 1.5. There a single axial force on the right end has been replaced by a small region of constant pressure. The other end stresses are essentially zero. Within the distance of about one depth the



Fig. 1.5. St. Venant's principle: local effects quickly die out.

axial stress has re-distributed to an essentially constant value in the remainder of the part.

In undergraduate statics and dynamics courses engineers are taught to think in terms of point forces and couples. Solid elements do not accept pure couples as loads, but statically equivalent pressures can be applied to solids and yield the correct stresses. Indeed, a couple at a point is almost impossible to create, so the distribution of pressures is probably more like the true situation.

The magnitudes of applied loads are often guesses, or specified by a governing design standard. For example, consider a wind load. A building standard may quote a pressure to be applied for a given wind speed. But, how well do you know the wind speed that might actually be exerted on the structure? Again, there probably is some type of “bell curve” around the expected average speed. You need to assign a contribution to the total factor of safety to allow for variations in the uncertainty of the load value or actual spatial distribution of applied loads.

Loading data are usually less accurate than the material data, but much more accurate than the restraint or supporting conditions considered next.

- (8) Determine (or more likely assume) how the model interacts with the surroundings not included in your model. These are the restraint (support, or fixture) regions. In mathematical terminology, these are called the essential boundary conditions, or Dirichlet boundary conditions. You cannot afford to model everything interacting with a part. For many decades engineers have developed simplified concepts to approximate surroundings adjacent to a model to simplify hand calculations. They include roller supports, smooth pins, cantilevered (encastre, or fixed) supports, straight cable attachments, etc. Those concepts are often carried over to FEA approaches and can over simplify the true support nature and lead to very large errors in the results.

The choice of restraints (fixations, supports) for a model is surprisingly difficult and is often the least reliable decision

made by the engineer. Small changes in the supports can cause large changes in the results. It is wise to try to investigate a number of likely or possible support conditions in different studies. When in doubt, try to include more of the surrounding support material and apply assumed support conditions to those portions at a greater distance from critical part features.

You need to assign a contribution to the total factor of safety to allow for variations in the uncertainty of how or where the actual support conditions occur. That is especially true for buckling studies.

- (9) Solve the linear system of equations, or the eigenvalue problem. With today's numerical algorithms the solution of the algebraic system or eigen-system is usually quite reliable. It is possible to cause ill-conditioned systems (large condition number) with meshes having bad aspect ratios, or large elements adjacent to small ones, but that is unlikely to happen with automatic mesh generators.
- (10) Check the results. Are the reactions at the supports equal and opposite to the sources you thought that you applied? Are the results consistent with the assumed linear behavior? The engineering definition of a problem with large displacements is one where the maximum displacement is more than half the smallest geometric thickness of the part. The internal definition is a displacement field that significantly changes the volume of an element. That implies the element geometric shape noticeably changed from the starting shape, and that the shape needs to be updated in a series of much smaller shape changes. Are the displacements big enough to require re-resolution with large displacement iterations turned on? Have you validated the results with an analytic approximation, or different type of finite element? Engineering judgments are required.
- (11) Post-process the solution for secondary variables. For structural studies you generally wish to document the deflections, reactions and stresses. For thermal studies you display the temperatures, heat flux vectors and reaction heat flows. With natural frequency models you show (or animate) a few mode

- shapes. In graphical displays, you can control the number of contours employed, as well as their maximum and minimum ranges. The latter is important if you want to compare two designs on a single page. Limit the number of digits shown on the contour scale to be consistent with the material modulus (or conductivity, etc.). Color contour plots often do not reproduce well, but graphs do, so learn to use them in your documentation.
- (12) Determine (or more likely assume) what failure criterion applies to your study. This stage involves assumptions about how a structural material might fail. There are a number of theories. Most are based on stress values or distortional energy levels, but a few depend on strain values. If you know that one has been accepted for your selected material then use that one (as a contribution to the overall factor of safety). Otherwise, you should evaluate more than one theory and see which is the worst case. Also keep in mind that loading or support uncertainties can lead to a range of stress levels, and variations in material properties affect the strength and unexpected failures can occur if those types of distributions happen to intersect, as sketched in Figure 1.6.
- (13) Optionally, post-process the secondary variables to measure the theoretical error in the study, and adaptively correct the solution. This converges to an accurate solution to the problem input, but perhaps not to the problem to be solved. Accurate garbage is still garbage.




Fig. 1.6. Overlap of stresses and material strengths can cause failure.




Fig. 1.7. Relative uncertainty of major modeling stages.

- (14) Document, report, and file the study. The part shape, mesh, and results should be reported in image form. Assumptions on which the study was based should be clearly stated, and hopefully confirmed. The documentation should contain an independent validation calculation, or two, from an analytical approximation or an FEA based on a totally different element type. Try to address the relative uncertainties of the main analysis stages, as summarized in Figure 1.7.


Technical communication and documentation is always important. In America, engineers are supposed to retain their calculations for at least seven years. Will your report be clear and helpful if you have to defend it years later? Paper hardcopies are the most reliable for long term storage. (Can you read the electronic media you used five years ago?)

You usually assume that the materials are linear. If not (creeping, hyperelastic, inelastic, plastic, viscoelastic, etc.), define the appropriate material data and the nonlinear equations to be solved. Then the matrix system becomes non-linear. Your original results check may lead you to conclude that the problem is actually an iterative one due to large displacements, or the need to insert physical contact interfaces.

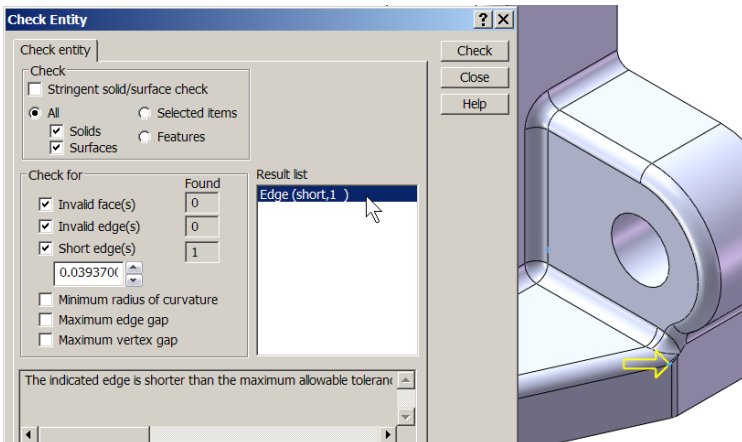
## 1.6. Part Geometric Analysis and Meshing Failures

Before attempting meshing your part, for a finite element analysis, you should check your solid model for potentially fatal geometric flaws that may not be noticed except at greatly magnified views. Within SolidWorks this is called a **Geometric Analysis**. To utilize that feature, a geometric analysis check the Angle\_Connector part will be outlined:

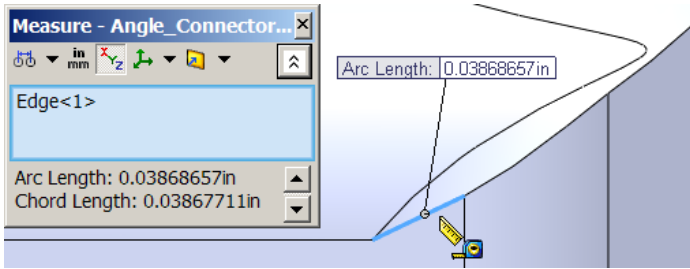
- (1) Select **Tools** → **Check** will open the **Check Entity** panel.




- (2) In that panel check the boxes for most entities, select **Check**.
- (3) Highlight each item in the **Result List**. As you scroll down the Result list the short edge location on the part is illustrated by a yellow arrow. Either the feature needs to be eliminated, or the mesh will need to be finer.



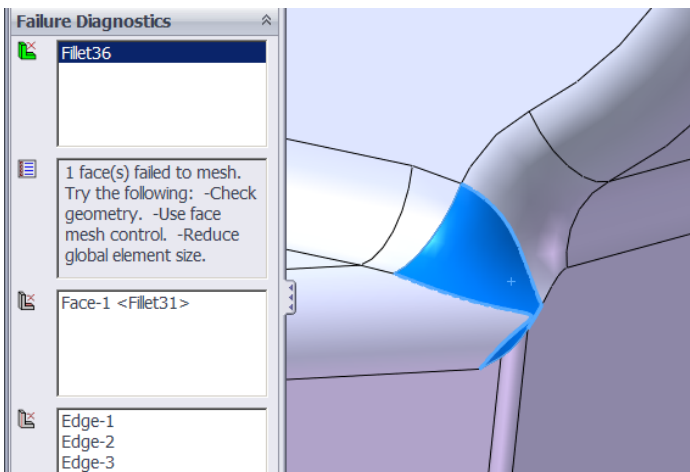
- (4) To consider a potential mesh refinement you should determine the size of the small feature. Use **Tools** → **Measure** to open up the **Measure** panel. Select the **XYZ** option, click on a edge of the feature to see its length.




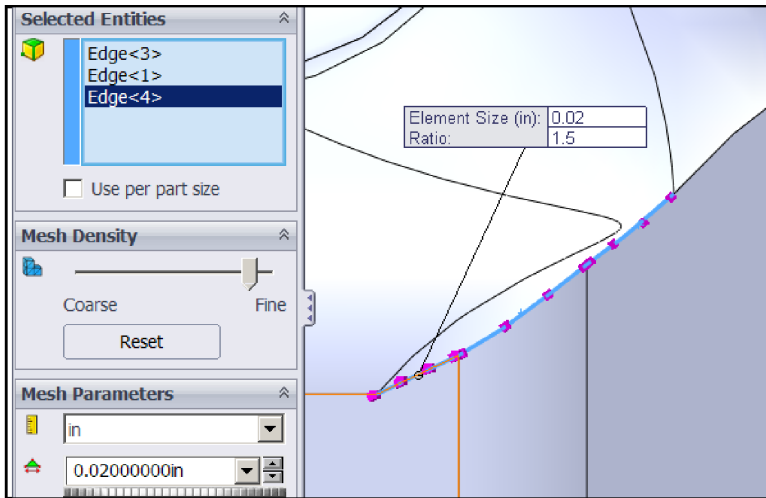
- (5) Attempt to create a mesh: **Mesh** → **Create Mesh**. As expected, that process fails and a failure diagnostic message appears:



- (6) Right click on **Mesh** to open the **Failure Diagnostics** panel. Scroll down the lists of faces or edges that caused the meshing failure. In this case, there is a highly distorted surface that formed with the fillets. Sometimes this type of surface can be removed by suppressing the fillets, or by simply building the fillets in a different order. Sometimes the surface can split by inserting split lines to make more manageable regions. Repairing the surface is better than having to *try* to control the mesh.




- (7) First, try to get some type of mesh output by specifying a small element size along the edges of the distorted region **Mesh** → **Apply Mesh Control**. Specify a local element size that will assure that one or two elements will fit along the smallest edge. Surprisingly, this worked. But it yielded a distorted mesh in the region of the small edge. Ideally, the surface triangles (one face of the tetrahedron) would be isosceles. That gives an element “aspect ratio” (say the ratio of the long side divided by the short one) of unity. Here the triangles are curved. A few are also badly distorted and not desirable for analysis if they are in an expected high stress region.




One measure of the quality of an element is its aspect ratio. Think of that as the ratio of the diameter of the enclosing sphere to the diameter of the enclosed sphere. Alternately, use the ratio of the longest element edge length to its shortest. An ideal aspect ratio should be near unity. Check the mesh quality by looking at a plot of the aspect ratio of the elements. Select **Mesh** → **Create Mesh Plot** → **Aspect Ratio**.


Try to improve this mesh by removing the bad surface, or subdividing it into two regions. At the narrow region, insert a **split line** that avoids very small intersection angles with both curves.




The small slender partition will need very small elements, but the larger partition can have larger ones. Especially if you use the transition control ratio to give five or more growth layers at an enlargement ratio of about 1.2 instead of the default value of 1.5. Use **Mesh** → **Apply Mesh Control** to specify element sides of *0.02* and *0.05 inches*, respectively in the **Mesh Control** panel. They give a much better mesh in this region.



Another part, the *Five\_Hole\_Link*, shows a tangency that gives very bad element aspect ratios. A common cause of failure in mesh generation is to have two solid regions or two joining surfaces meet at a near zero angle. That often happens in practice and requires



intervention to be able to create a mesh for analysis. If a tangency condition is really required in the part, then you must force smaller element sizes there via the **Mesh Control** option.



If the part can be modified to avoid the tangency, then meshing becomes much easier. That is illustrated below where it was feasible to avoid a tangency requirement in this application.