

Rock-Solid Innovation.

See page 8 for more information.



March 2011 / deskeng.com

DE

Desktop Engineering®

TECHNOLOGY FOR DESIGN ENGINEERING

PLM for SMBs P. 22

**Engineering the
3D-Printed Car** P. 34






Visualization's TRUE COLORS

**How to defend your finite
element stress data results.** P. 16

**INCORPORATE POWER
HARVESTING SENSORS** P. 37

**KNOWLEDGE-BASED
ENGINEERING** P. 46

**REVIEWS OF DELL AND
EUROCOM WORKSTATIONS** P. 26



See Analysis Data's True Colors

Why you don't always get what you want—and how to get what you need.

BY GEORGE LAIRD AND PAMELA J. WATERMAN

“Whenever you see a stress contour plot, just assume that it is wrong,” says Mark Sherman, head of the Femap Development Team for Siemens PLM Software Solutions. Although Sherman’s comment sounds a bit dramatic, it’s par for the course in computer modeling, where a common saying is “garbage in, gospel out (GIGO).” The questions that these comments raise are simple to pose, but are somewhat vexing to answer, even for the specialists.

How does a user quickly check finite element (FE) stress data validity or accuracy? How can one astutely defend these colorful, high-resolution results against the casual interloper who is throwing darts? Is the data gospel or garbage? This article will give you the background and ammunition to defend your results.

How Stresses are Calculated in FEA

Let’s look behind the scenes at the general approach to FE analysis. Once you have prepped and launched your FE model, the software analysis engine kicks into action, taking each element and breaking it down into a series of simple stiffness equations. Whether linear or non-linear, the equations are assembled into a massive matrix, solved by one of several intensely mathematical methods and evaluated for the given constraints and loads.

With a quick menu-click and perhaps some rotations, you have a colored contour graph showing ... *what exactly?*

FIGURE 1: Gaussian integration points are shown for common finite element shapes: plate, brick and tetrahedral. Images courtesy of Predictive Engineering

Time to back up. What is glossed over in most explanations is how the analysis engine generates stiffness equations from oddly shaped elements defined in two or three dimensions. It is an amazingly beautiful process of simplification done over several stages.

First, each element is broken down into quadrants. A weighting formula is used to calculate the approximate volume of each element using simple polynomials. Known as Gaussian integration, this step is the cornerstone of all FEA technology. Without this simplification, FEA would not exist.

Second, the software generates stiffness equations and the analysis engine applies the given constraints and loads. Then, the engine calculates all of the displacements at the element’s corner points or nodes, leading to the next question: How does it generate stresses from displacements away from the element’s corners?

The answer lies in another great trick of the FEA process. Theoreticians have determined that the best place to calculate stresses in finite elements lies at the Gaussian integration points (the center of the quadrants). The software then takes these displacements and uses Gaussian integration to calculate first the strain and then stress within each “Gaussian volume.” Most finite elements are analyzed using four Gaussian integration points, and thus the analysis engine generates stresses at four discrete

points within the element. (See Figure 1.)

But how do these Gaussian integration points relate to what we see on our real-world contour plots? Because most users have no use for raw element data, one final processing step is done. Values for the Gaussian point stresses are interpolated into the element's center, and also extrapolated out to the corner points or nodal points. At this stage, the analysis engine has done its work—and the visualization process starts.

De-bugging Jagged Stress Contours

You are more knowledgeable now than many stress analysts about how element stresses are computed, but what exactly are you seeing? These millions of extrapolated points have been loaded into the software's graphic display system and presented as a dazzling, yet smoothed spectrum of colors. Smoothed is a key word here.


In the default mode, FEA programs average the corner point stresses from each element and only present the averaged value to the user. This little smoothing technique has its good and bad points. Overall, it is a good thing because it smoothes out the stresses into a cleaner pattern. Numerically, this process adjusts for the non-physical variations in stresses that stem from minor variations in the element's shape.

Because physical stress is not a step-function (it wants to flow smoothly), numerically smoothed results display the more accurate solution. If they don't, you should definitely worry! This is the garbage part. Alarm bells should ring, which brings us to the task of developing your FEA eyes.

When contours look jagged, with lots of red spots (see Figure 2a) or have extremely irregular shapes, three causes are likely:


- poorly applied loads and constraints (likely);
- complicated or bad CAD geometry (tricky); or
- poor mesh quality. (See "Judging Good and Bad FE Shapes" to the right.)

Because we are using finite elements to approximate a continuum, sometimes it's best just to accept a few discrepancies when they can be easily explained as a reality of the modeling process. For example, one could quote Saint-Venant's principle and say that stress and displacement contours away from load application points, and do not depend on how the load was applied (concentrated or distributed) because forces and moments are always conserved. In other words, if the region of interest is far away from the load application point, you can ignore the less-than-smooth stresses around the load and constraint points.



← **FIGURE 2A:** Stress is never jumpy. It should always vary smoothly, so if your stress results don't appear right, then most likely they are not. Here, warning lights should come on when you see the "spots" of high stress repeatedly popping up along the curved edge of this part.

→ **FIGURE 2B:** These stress results are 20% higher than those in Figure 2a. The values reflect smooth transitions across the geometry. Visual inspection will tell you that the higher stresses reflect results much closer to reality.



From a geometric point of view, designers hand over CAD files that include every manufacturing detail down to 0.005-in. sharp edge-breaks or small diameter oil-feed passageways. Typically, we can (and should) ignore these small regions by invoking another useful relationship from Jean Claude Barré Saint-Venant. (See "Saint-Venant's Principle of Decreasing Load Effects," page 18.) He discovered that small features only create localized disturbances in the stress field. The extent of this disturbance is no more than three times the characteristic dimension of this small feature. For example, for holes of radius R , this size is $3 \cdot R$. Thus, if your objective is to determine the overall stress of the structure, localized excursions in the stress field will not affect your final answer. This can be

Judging Good and Bad FE Shapes

Gaussian integration is an exact method only for elements that have near-perfect shapes—side-to-side ratios of 1x1 for plates or 1x1x1 for solids (bricks and tetrahedrals, for example). Once the element shape starts to degrade, so does the quality of your results.

A standard rule of thumb is that the element shape should be "pleasing to the eye" and maintain this regularity across the model. Although it's a subjective approach, if the mesh appears to be close to having 1x1x1 proportions, you are good to go.




FIGURE 3A (LEFT): A poor meshing job. Judging good and bad elements is as easy as looking at the mesh. If areas of the mesh look angular, sharp and irregular, then most likely your stress results will be inaccurate. This part is a completely symmetric hemisphere (with the addition of one small hole at the lower left). The mesh should be virtually uniform across the part—but clearly is not, nor are the resulting stresses. The problem comes from an artifact of the meshing algorithm, and can be solved by dividing the part into subsections before meshing.

FIGURE 3B (RIGHT): Another trick to help you identify bogus meshing regions is that you should be able to turn off the mesh outline and not tell where the elements start or stop. See the pronounced variations in stress levels on the left side of this same hemispherical part, in spite of the part's physical symmetry.

easily proved by running two analyses on a part, both with and without a few holes and verifying that the results will be essentially identical.

If the small features are truly relevant and cannot be simplified away, then you will have to set the mesh density to be increased so that the resulting stress field looks smooth and realistic. (See Figure 2b.) Your existing software may handle this with no problems, or you may find you need to divide the geometry into subsections for the mesher to properly deal with the peculiarities of your part. (See Figures 3a and 3b.)

Saint-Venant's Principle of Decreasing Load Effects

In 1855, French elasticity theorist Jean Claude Barré de Saint-Venant concluded that mechanical load effects on a body rapidly diminish as you go away from the point of application—at distances that are large compared to the dimensions of the part. In 1945, Richard Edler von Mises put a mathematical description to this effect using partial differential equations.

The principle allows FE algorithms to replace complicated stress distributions with simpler ones as long as the boundaries are small. You can compare this situation to electrostatic field effects, where the field decays as $1/(r_i+2)$, where i relates to the number of poles.

Interpreting Stress Results

One more back-to-basics reminder: Stress is force divided by area. This basic tenet of stress analysis sounds simplistic, but it's extremely useful. For example, if you have two bars with the same cross-sectional area—one aluminum and the other steel—and the same force is applied to each bar, then each bar will undergo the same stress. In the age before computers, this is why stress guys could build scale models of complicated structures in plastic and get useful information.

In other words, if you are applying force or pressure to a structure, the resultant stress will only be a function of the geometry of the structure and not its material. One can then do material selection on the fly, since the model's maximum stress value is valid whether the material is plastic, aluminum or steel. The only thing that would change is the deflection of the structure, which scales with the elastic modulus.

A common occurrence in the interpretation of stress results is when the model returns a very high stress value—say, 500 megapascals (MPa)—when the user knows that the material of choice for this structure (low-grade steel, for example) has an actual yield stress of 300 MPa.

How does this happen? Because the FE calculation is linear, it has no idea that the material would plastically deform or break above the material's yield stress; it just makes the calculation based on the load and the geometry. This is where the user must interpret the result. Perhaps the load is unrealistically high and in reality, the load is half this value.

