

The Focus



A Publication for ANSYS Users

Contents

Feature Articles


- [New Multi-Layered Thermal Shell 132](#)
- [Coup d'Tet: Rethinking Thin Wall Meshing For Linear Static Stress Analysis](#)
- [U Convergence Criteria & Contact](#)

On the Web

- [PADT: ANSYS 7.0 Seminar](#)
- [ANSYS: ANSYS 7.0 New Features](#)

Resources

- [PADT Support: How can we help?](#)
- [Upcoming Training at PADT](#)
- [About *The Focus*](#)
 - *The Focus* Library
 - Contributor Information
 - Subscribe / Unsubscribe
 - Legal Disclaimer
- [Next in *The Focus*](#)



ANSYS
Support Distributor

New Multi-Layered Thermal Shell 132

by [Joe Woodward](#)

I know that many people wait as long as they can to upgrade to the newest version of ANSYS, but if you do thermal analysis, you may have a new reason for taking the plunge. ANSYS has introduced two new layered thermal shell elements. These new elements, the 4-node SHELL 131, which I will focus on, and the 8-node SHELL 132, were beta tested in ANSYS 6.1 and released in ANSYS 7.0. Of course you should always read the manuals before using any new elements, but this article should give you a good overview of these elements and describe some of the pitfalls and mistakes that are easily made when first starting to use them.




Figure 1. Shell 131 - 4 Node Layered Thermal Shell.

Shell 131 is a four-node layered thermal shell element with up to 32 temperature degrees of freedom at each node. This allows up to 15 layers with a quadratic gradient or 31 layers with a linear gradient. The node is placed at the mid-thickness of the shell and at each node there is a TTOP and TBOT degree of freedom. Then, depending on the number of layers, and whether the linear or quadratic gradient is used, the degrees of freedom TE2 to TE30 are added as shown in the lower right picture of Figure 1, which shows the placement for a quadratic gradient. Each layer of the shell can have a different material and fiber direction. Along with temperatures, heat generation can be specified for each layer, and the elements also support convection and heat flux on the surfaces and the edges. Radiation effects can be modeled on the top and bottom faces only.

The shell uses the SECTION set of commands. The piece of APDL code below creates a shell with 3 layers. Layers 1 & 3 will be .005 thick using material 1 & 2. Layer 2 will be 0.050 thick using material 2. The number of layers specified in KEYOPT(4) of the element-type command must match the number of SECDATA commands listed.

```
et,1,131,0,0,0,3 !Keyopt(4) is the number of layers
sectype,1,shell
```

The Focus



A Publication for ANSYS Users

```
secdata,.005,1,0!Material #1
secdata,.050,2,0!Material #2
secdata,.005,2,0!Material #2  May be necessary to match layers elsewhere.
type,1
sec,1
amesh,1
```

The help manual in version 7.0 says that you can leave KEYOPT(4) blank, and ANSYS will add element types as needed to match all of the section types with different number of layers. At the time of this article, however, that feature was not working correctly. So for now, you must specify an element type for each section type that has a unique number of layers. Sections can be plotted using the SECPLLOT command in the pre-processor, as shown in Figure 2.

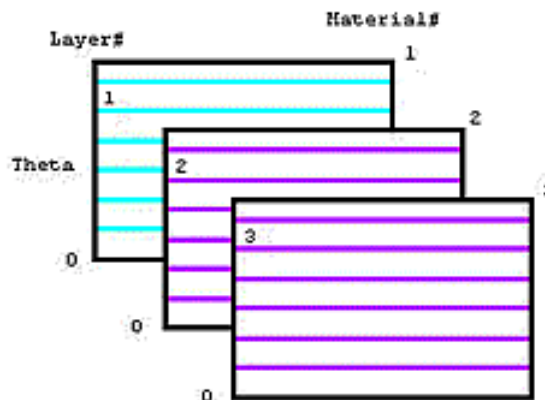


Figure 2. Section plot shows material and layer angle.

The shells are post-processed using the PLNSOL command. This allows the temperature distribution of each layer to be examined by simply plotting the corresponding degree of freedom. For example, PLNSOL, TE2 would plot the mid-thickness temperature of Layer 1 in Figure 1.

Benefits:

These elements exhibit the usual benefits of using shells instead of solid elements for thin walled applications. There is an increase in computational speed, you can have very thin aspect ratios, and with the layers it is simple to incorporate material or orientation changes though the thickness. These elements also have an option that allows them to be meshed directly to solid thermal elements. By setting the KEYOPT(6)=1, the TBOT degree of freedom is replaced with TEMP, as shown in Figure 2, which then matches the TEMP degree of freedom of the solid elements. The red dots represent the same node used on both elements. This paint option allows the ESURF command to be used to quickly simulate multi-layered insulation, or paint, added to an existing 3D solid thermal model.




Figure 3. 'Paint' option for Shell 131.

Another benefit is that the temperatures obtained with the SHELL 131 and SHELL 132 elements can be transferred directly to the corresponding structural shell elements using the LDREAD command. If your structural shell only uses two temperatures through the thickness, like SHELL 43, then only TTOP and TBOT from each node will be transferred. For

The Focus



A Publication for ANSYS Users

layered shells, like SHELL 181, all the temperatures will transfer, but the number of layers must match between the thermal and structural models.

Pitfalls:

Although there are many benefits, and using the shell is rather simple, there are some things of which a new user needs to be aware. One of them is that although the elements can have a quadratic gradient within each layer, this may not be sufficient for results containing a 3D temperature distribution. Not only must care be given to refine the in-plane mesh density, but more layers may be needed to capture the through-thickness gradient. In the case of the shell defined by the piece of code above, the center layer is 0.050 thick. For the model shown in Figures 4-6, this created a through-thickness aspect ratio greater than 20 for that layer. The resulting thermal distribution along the mid-plane of the 0.050 thick material layer is shown in Figures 4-6, using 1, 3, and 8 element layers respectively. You can see in the figures that for the same mesh density, the contours smooth out with the more layers used. There was very little change in the results when more than 8 layers were used. Not shown, however, is that for the same number of layers, say five, the distortion of the thermal distribution increases with an increasing mesh density, which is also an increase in through-thickness aspect ratio. This leads to the conclusion that the results of 3D gradients are sensitive to the aspect ratio of the thickness of each layer to the element size. This affect was not present when the load resulted in a 1D gradient, such as caused by a heat flux over the entire surface. After some trials, it seems that keeping the through-thickness aspect ratio below 4 seems to be a good rule of thumb for 3D temperature fields.

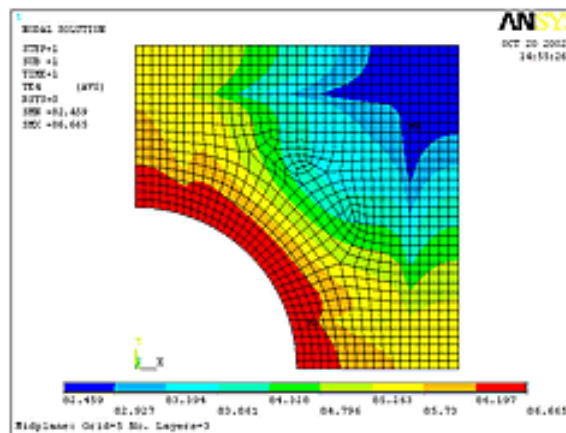


Figure 4. 1 Layer for 0.050".




Figure 5. 3 Layers for 0.050".




Figure 6. 8 Layers for 0.050".

The Focus



A Publication for ANSYS Users

Another problem, encountered during testing, is in the connecting of shells of different number of layers. If you have a continuous mesh of shells, some with two layers and some with three layers, you might expect to get the model shown in Figure 7. But because of the different degrees of freedom added for each layer, the actual result, if a linear gradient is used, would be that shown in Figure 8. The red dots signify the same node of a continuous mesh. At the connecting node, the TTOP degree of freedom is connected and the TE3 degree of freedom will be free from any connection. This would definitely cause incorrect results, which might be overlooked if you are not specifically checking for this type of situation. If you plot the results for TTOP, you would still get continuous contours. You might not even see the problem when plotting TE3, since there would be a discontinuity anyway, and the temperatures might be reasonable. That is why I am discussing the possible problems now, so that you will know what to watch out for.




Figure 7.

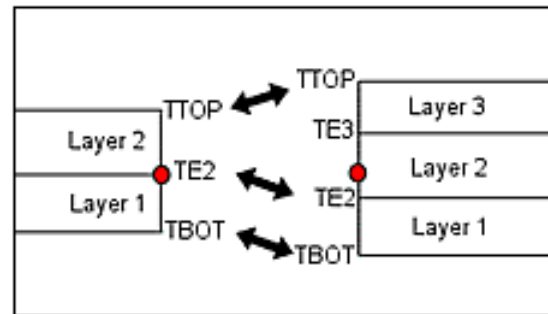


Figure 8.

There are two ways to fix this situation. The first would be to create coincident nodes at the interface of the two-layer and three-layer elements. The proper degrees of freedom can then be coupled together as shown in Figure 9. The second method would be to create a continuous mesh with the two-layer elements and then add the third layer where needed. By defining a one-layer element type, and creating a coincident mesh, the TBOT degree of freedom can be coupled to the TTOP degree of freedom as shown in Figure 10. In this same way, many elements can be stacked on each other if more layers are needed.




Figure 9.




Figure 10.

There are several considerations when choosing which method to use. If you need to minimize run time, you should choose the method that minimizes the number of constraint equations necessary. If you have many layers and only a small area has the extra layer, then you would want to use method two since it requires only one couple per node. If you have a large area to cover with the extra layer, method one may use less couples since there would be fewer nodes

The Focus



A Publication for ANSYS Users

at the connection interface. If you were concerned more with making it easy to apply boundary conditions, as in a parametric study, you would choose method one. Then the TTOP boundary conditions can be placed on all the elements. With method two, you have to take care not to put boundary conditions on the elements that are underneath the extra layer. If you have the need for nice continuous contour plots, however, you would want to use method two. You could then plot TTOP for all of the two-layer elements and the plot would be correct. With method one, a plot of TTOP would still show the top of the three-layer elements and the top of the two-layer elements, which are not on the same layer.

Conclusion:

I hope that this article has shed some light on these powerful new elements. When you know what to watch out for, it is fairly easy to avoid the pitfalls discussed, and the benefits are numerous. With a little practice, these new elements should expand the capability, and productivity, of your thermal analyses.

The Focus



A Publication for ANSYS Users

Coup d Tet Rethinking Thin Wall Meshing For Linear Static Stress Analysis

by [Jeff Stressin Strain](#), ANSYS Supporter and User

First off, a confession. I initially approached this article with the intent of demonstrating the importance of having at least two tetrahedral elements through the wall thickness in any given finite element model to capture bending behavior and achieve accuracy. After running a few test cases, I found myself rethinking the necessity of doing so, as I will demonstrate in this article. I've also included studies on brick meshes, both with and without midside nodes, for good measure.

The Problem

A common rule of thumb when meshing a model is to make sure you have a minimum of two elements across any dimension. This is particularly emphasized with regards to tet elements. This can become rather problematic when thin walls are involved. With a brick mesh, this can result in poor aspect ratios and/or a high number of elements taking up disk space and CPU time. With a tet mesh, this can be nearly impossible to achieve without setting the global element size to a value that results in an ultra-fine mesh throughout the entire model, unless an excessive number of local controls (e.g. AESIZE, LESIZE) are employed. Given the frustration in applying this rule of thumb, a valid question to ask is whether or not it is absolutely necessary. A simple study indicates otherwise for many, if not most, cases.

The Model

The test model in this case was a thin three-dimensional beam, 10 x 1 x 0.1 units, simply supported at each end and subject to a constant pressure of 1 unit. The Young's modulus was set to a value of 10×10^6 units. This scenario was chosen to avoid singularities at the location of maximum stress. The behavior was assumed to be linear to maintain consistency with closed form equations. See Figure 1 for an illustration.




Figure 1. Pressurized Slender Beam Model.

The Focus



A Publication for ANSYS Users

Theoretically, the maximum stress in the beam is $6M/bt^2 = 6pwL^2/8wt^2 = 6 \times 1 \times 1 \times 10^2 / 8(1 \times 0.12) = 7500$ units. The maximum deflection is $5pwL^4/384EI = -5 \times 1 \times 1 \times 10^4 / 384(10 \times 10^2 \times 1 \times 0.13/12) = -0.15625$ units. (Reference: *Roark's Formulas for Stress and Strain*, Table 3, Case 2e)

The Tet Mesh

The first two cases modeled used SOLID92 tetrahedral elements. Both meshes had 100 elements along the length and 10 elements across the width for consistency. The only parameter varied was the number of elements through the thickness. In the first test, a single element through the thickness was used. In the second test, the beam was meshed with two elements through the thickness. See Figure 2 and Figure 3 for an illustration of each mesh.




Figure 2. Tet Mesh One Element Through Thickness.

The Focus



A Publication for ANSYS Users





Figure 3. Tet Mesh Two Elements Through Thickness.

The resulting stresses and displacements for each model are plotted in Figure 4 and Figure 5. Note that the results are virtually identical, and that the discrepancy between the theoretical and finite element results is negligible (zero discrepancy in the stress results, in fact). This is not entirely surprising in that the only tetrahedral elements available have midside nodes, therefore exhibiting quadrilateral displacement functions.



A Publication for ANSYS Users




a) Stress Results

Figure 4a. Results for SOLID92 Tet Mesh One Element Through Thickness.



A Publication for ANSYS Users




b) Deflection Results

Figure 4b. Results for SOLID92 Tet Mesh One Element Through Thickness.



A Publication for ANSYS Users




a) Stress Results

Figure 5a. Results for SOLID92 Tet Mesh Two Elements Through Thickness.



A Publication for ANSYS Users



b) Deflection Results

Figure 5b. Results for SOLID92 Tet Mesh Two Elements Through Thickness.

The Hex Mesh

The behavior of a thin beam meshed with hexahedral elements, both with and without midside nodes, was also considered. Both sets of models had the same element sizes as in the tet models. The stress and deflection for each are listed in Table 1. Tet mesh results are also included for side-by-side comparison.

| Element Type | Elements | Thick | Max Stress | Max Deflection |
|--------------|----------|-------|------------|----------------|
| SOLID95 | 1 | 7500 | -0.156275 | |
| | 2 | 7500 | -0.156280 | |
| SOLID45 | 1 | 7499 | -0.156250 | |
| | 2 | 7499 | -0.156251 | |
| SOLID92 | 1 | 7500 | -0.156272 | |
| | 2 | 7500 | -0.156273 | |

Table 1. Stress and Displacement Results for Hex Mesh.

Note that even the SOLID45s, which have no midside nodes, exhibit very little error. This is largely due to the extra shape terms that have been part of the element shape functions since the early 1990s.

The Focus



A Publication for ANSYS Users

The Conclusions

Based on the test cases I've run for this study, it is evident that one tet element through a section is sufficient to capture bending behavior in models where there are no boundary conditions applied in-plane and where thin wall theory applies (a common rule of thumb as to whether a wall is thin is a major dimension to thickness ratio of 10:1). Regardless of what element you use, it is still important to watch out for sharp deflection and stress gradients due to more general mesh discretization errors. Discretization problems, in thin sections or otherwise, can result from complex loading and stress concentrations.

The Focus



A Publication for ANSYS Users

U Convergence Criteria & Contact

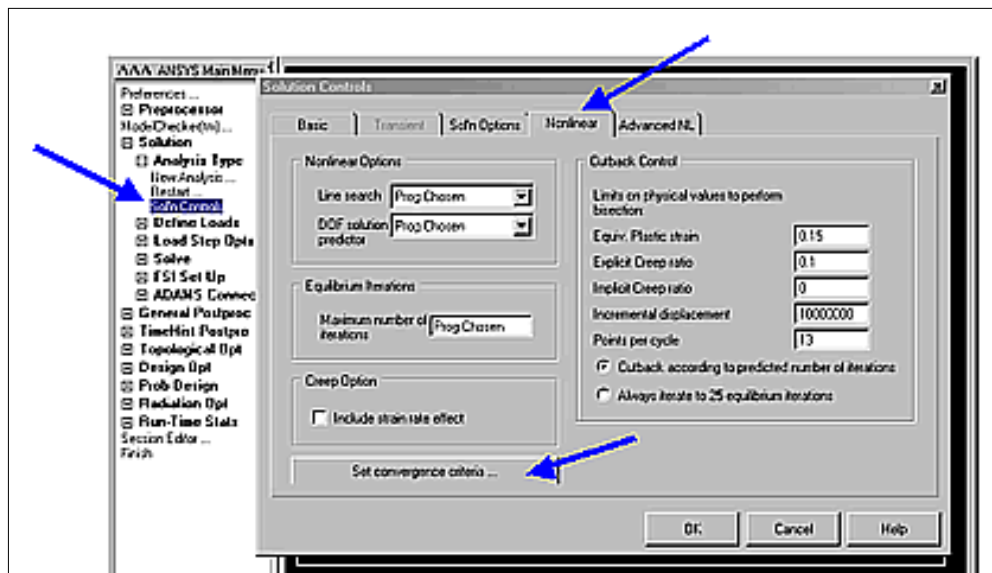
by [Rod Scholl](#)

When struggling for convergence in a contact analysis, you will want as many tricks in your bag as possible. When I have difficulty with contact convergence, I make adjustments to the following:

1. Time Step Size
2. Contact Stiffness (FKN)
3. Switch to U Convergence Criteria
4. I am told that LNSRCH,ON is also very helpful&

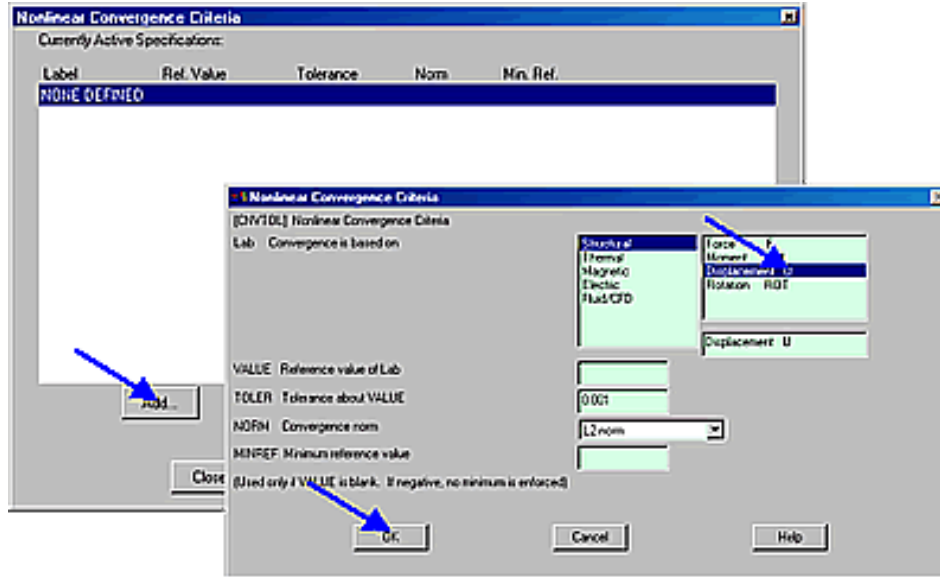
Set time steps sufficiently low. Set FKN as low as your accuracy requires. I find for most analyses, convergence is quickest with an FKN of about 0.01 to 0.001. This softness often is negligible in my analyses, especially when I don't care about stresses at the interface.

But switching to Displacement (U) convergence Criteria is my ace in the hole. Normally ANSYS tries to converge based on Force (F) and Moment (M). By using Displacement (U) some analyses converge more quickly (and some more slowly!).






A Publication for ANSYS Users



Defaults values are usually sufficient.

Test Case


To compare performance we will examine the following test case (download [input file](#)).



The Focus



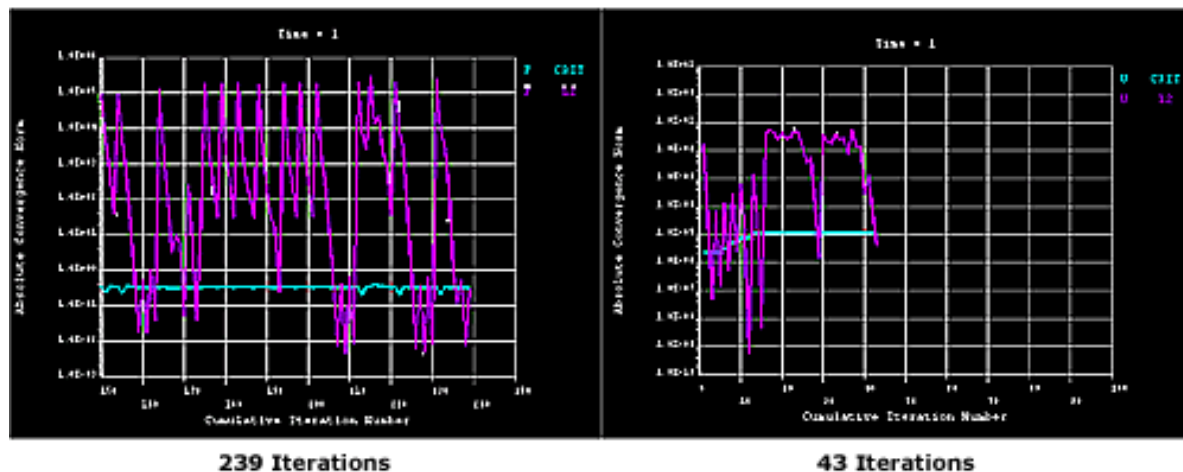
A Publication for ANSYS Users



In the above depiction of contact status, it is clear that only a node or two is making contact during the solution.

F Convergence – 39.5 seconds

U Convergence – 7.4 seconds



So the U convergence was 5X faster!

For rigorousness, one could add the F and M criteria after it converged, and resolve from the deflected shape. In this case, the stresses matched within 0.001% between the F and U convergence criteria-based runs.

REMEMBER: This is the exception, not the rule. Most the analyses (4 out of 5?) run faster with the default setting of F convergence norm. I had to spend some time crafting this analysis such that U was faster. Rigid parts with little deflection won't usually benefit from the U convergence criteria. I only switch to U if F/M is iterating too much, or I have some time to kill

Related Training

Learn more about contact and convergence in the [Basic Non-Linearities](#) course.



A Publication for ANSYS Users

About *The Focus*

The Focus is a periodic electronic publication published by PADT, aimed at the general ANSYS user. The goal of the feature articles is to inform users of the capabilities ANSYS offers and to provide useful tips and hints on using these products more effectively. *The Focus* may be freely redistributed in its entirety. For administrative questions, please contact [Rod Scholl](#) at PADT.

The Focus Library

All past issues of *The Focus* are maintained in an online [library](#), which can be searched in a variety of different ways.

Contributor Information

Please don't hesitate to send in a contribution! Articles and information helpful to ANSYS users are very much welcomed and appreciated. We encourage you to send your contributions via e-mail to [Rod Scholl](#).

Subscribe / Unsubscribe

To subscribe to or unsubscribe from *The Focus*, please visit the PADT e-Publication [subscriptions](#) management page.

Legal Disclaimer

Phoenix Analysis and Design Technologies (PADT) makes no representations about the suitability of the information contained in these documents and related graphics for any purpose. All such document and related graphics are provided as is without warranty of any kind and are subject to change without notice. The entire risk arising out of their use remains with the recipient. In no event, including inaccurate information, shall PADT be liable for any direct, consequential, incidental, special, punitive or other damages whatsoever (including without limitation, damages for loss of business information), even if PADT has been advised of the possibility of such damages.

The views expressed in *The Focus* are solely those of PADT and are not necessarily those of ANSYS, Inc.