

**AN INCOMPLETE GUIDE TO
ANSYS WORKBENCH**

Small Satellite Research lab
The University of Georgia

Version 1.0

Casper Versteeg, Christine Yang, John Courtney

May 14, 2018

Contents

- 1 Introduction** **3**

- 2 Creating Models** **3**

- 3 Model Setup** **4**
 - 3.1 Engineering Data 5
 - 3.2 Mechanical Model 6

- 4 Structural Simulations** **7**
 - 4.1 Modal 8
 - 4.2 Harmonic Response 11
 - 4.3 Response Spectrum 12
 - 4.4 Random Vibration 14
 - 4.5 Inertial Acceleration 14

- 5 Thermal Analysis** **17**
 - 5.1 Model Setup 17
 - 5.2 Transient Thermal 18
 - 5.3 Steady-State Thermal 22
 - 5.4 Results Interpretation 23

- 6 Conclusion** **24**

List of Figures

1	Geometry Comparison	3
2	Workbench Start-up	4
3	Modular Model Building	5
4	Engineering Data	6
5	Mechanical Model Start-up	7
6	Mesh Results	8
7	Structural Analysis Project Schematic	8
8	Boundary Condition Geometry Definition	9
9	Modal Total Deformation Results	10
10	Harmonic Response Results	11
11	Component Shock Spectrum	12
12	Response Spectrum Results	13
13	Random Vibration Directional Deformation Results	15
14	Mass Acceleration Curve	16
15	Inertial Acceleration Results	16
16	Geometry for Thermal	17
17	Transient Thermal Project Schematic	19
18	Transient Thermal Setup	19
19	Command Window	20
20	Transient Thermal Results	22
21	2- σ Corrected Temperatures	23

List of Tables


1	Component Shock Spectrum–Tabular Data	13
2	Random Vibration Qualification Profile–Tabular Data	14

1 Introduction

This guide is meant to help those completely new to ANSYS get acquainted with its major features and get started on building a project. By no means is this guide complete (hence the "incomplete guide") but although it skips over a lot of features, it should be comprehensive enough to give you the knowledge required to figure it out. The very first thing you will learn about ANSYS, it that it is incredibly picky, and can really make you regret life if you do not define constraints correctly. While this guide is meant to help put you on track, every problem is different, and different methodologies might apply.

2 Creating Models

The meshing tool in ANSYS is incredibly powerful, and its solver methods are well optimized. For this reason, creating models for an ANSYS simulation is reasonably straight forward. As long as there are not any terribly extravagant details, like threads or tiny holes for pin connections, your model will most likely be fine. Really, a rule of thumb should be, that if the feature is about a millimeter in side length or diameter, it is probably safe to remove it, unless it is structurally significant.



(a) Not Ok

(b) Probably Ok

Figure 1: Geometry Comparison

Figure 1 shows a comparison of acceptable geometry that will probably mesh well to a geometry that will probably not mesh well. The trick, here, is obviously all those holes that are normally for the PC/104+ connector pins. While they might be structurally significant, the ANSYS solver will have significant problems converging to a solution, and the mesh will likely contain thousands more nodes than are maybe necessary.

Structural significance is usually determined by either high elasticity in material, brittleness of material or low strength, but every instance is different and should be considered based on context rather than rules.

Any 3D model can be created in Inventor, or other CAD software for simulation purposes. For this guide, we will be using a Z-Magnetorquer board, as shown in Figure 1 (**note:** this component

is not actually used, but is good to show you how to operate ANSYS, simulations shown here are for demonstration). The best practice is to export any CAD models to .stp, STEP file, which are essentially text instructions for CAD software or interpreters—like ANSYS—to recreate geometry.

3 Model Setup

When you open ANSYS Workbench for the first time, it will look a lot like Figure 2. The project schematic area is where models are built and set up. The "Analysis Systems" on the left may look different for you depending on what license of ANSYS you are running, but the main components discussed here should all be present.




Figure 2: Workbench Start-up

It is generally good practice to keep ANSYS models modular. This means, instead of just dragging and dropping the analysis types into the schematic and working them all separately, instead building the model with common elements going to different analyses. This would make the model look something like Figure 3. The components A and B can be found by scrolling down to the "Component Systems". Components can be connected after dropping by simply clicking and holding the parts that need to be connected, and dragging the connection to the next component.

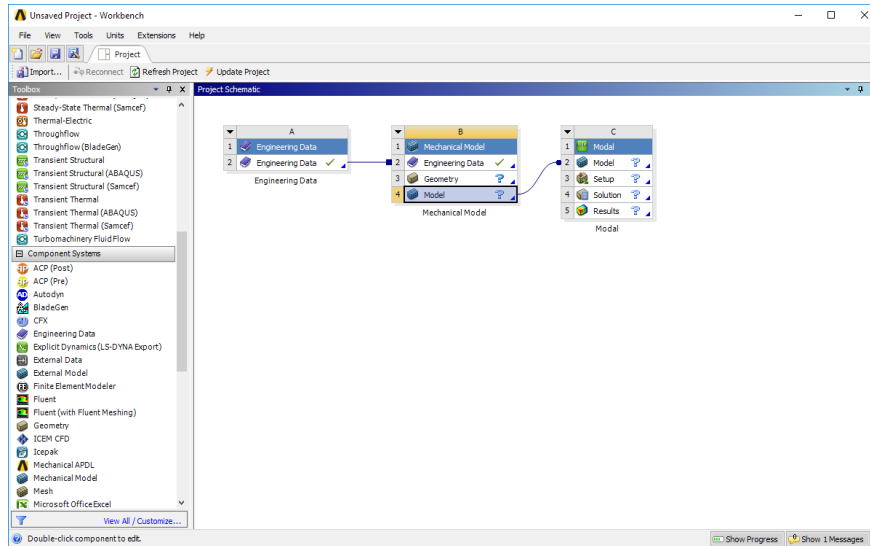


Figure 3: Modular Model Building

3.1 Engineering Data

The "Engineering Data" component in ANSYS is where material and fluid properties can be defined. As opposed to all the other modules we will be seeing in this guide, the engineering data simply opens a new tab in the workbench, named according to what parts are connected to—and thus dependent on—the engineering data (this name can get pretty obnoxiously long for large models).

By default, for a completely new Workbench model, the only material defined is "Structural Steel". Just remove that. By clicking on "Click here to add a new material" you can—you guessed it—add a new material and define its physical properties. We are trying to simulate the Z-Magnetorquer, which, for all intents and purposes, is just FR-4. Because we will initially only be looking at structural models, the only thing that needs to be defined here is density and linear, isotropic elasticity. These can be found under the "Toolbox" menu to the left by expanding the relevant menus. If you want to run a thermal model, you will need density, isotropic thermal conductivity and specific heat, as well.

Fields that require data, under the "properties" will be highlighted in yellow. For now, just define a density $\rho = 1850 \text{ kg} \cdot \text{m}^{-3}$. As for the isotropic elasticity, because the various properties in a linear material are linked, you can tell ANSYS to calculate any of them based on just inputting two, just depending on what data is readily available. This can be changed by changing the input in the drop-down menu next to "Derive from". For now, define the Young's Modulus (or elasticity) $E = 22000 \text{ MPa}$ and Poisson's Ratio $\nu = 0.118$. The units can be changed by clicking on the units next to the quantity for any of these, and selecting desired units. Once the quantity has been entered, ANSYS will automatically do the conversion if the units are changed. Your window should look like Figure 4

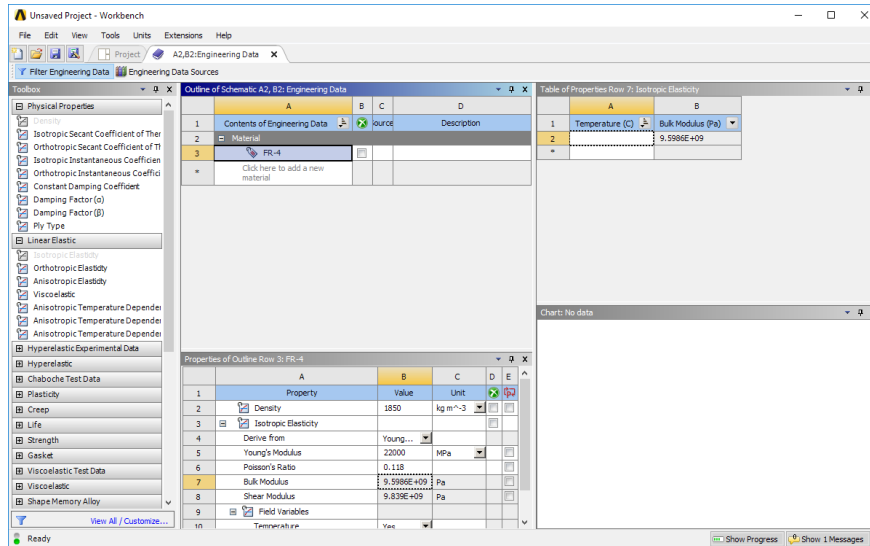


Figure 4: Engineering Data

Of course, if desired, more material data can be defined here, but the analysis in this guide will be fairly simple just to expose you to the various options available. Now, do **not** close the window, this will exit out of ANSYS. Instead, close the Engineering Data tab, this should bring you back to the project schematic.

3.2 Mechanical Model

While ANSYS has the ability for you to draw geometry, this generally does not work very well, and it is recommended to just import CAD made in a program that is designed better than ANSYS' DesignModeler. Trust me, don't hurt yourself, just import. Always import.

To import geometry, right click on the "Geometry" field under Mechanical Model. Select **Import Geometry**, and then "Browse...". Navigate to the **.stp** of the geometry you wish to analyze. The icon next to "Geometry" will change, to indicate you have added an external **.stp** file.

Double click on "Model" in the project schematic. This will open the Mechanical Model. Some dialogs will appear to tell you that ANSYS is importing the geometry you have previously attached, and then the window should look something like Figure 5.

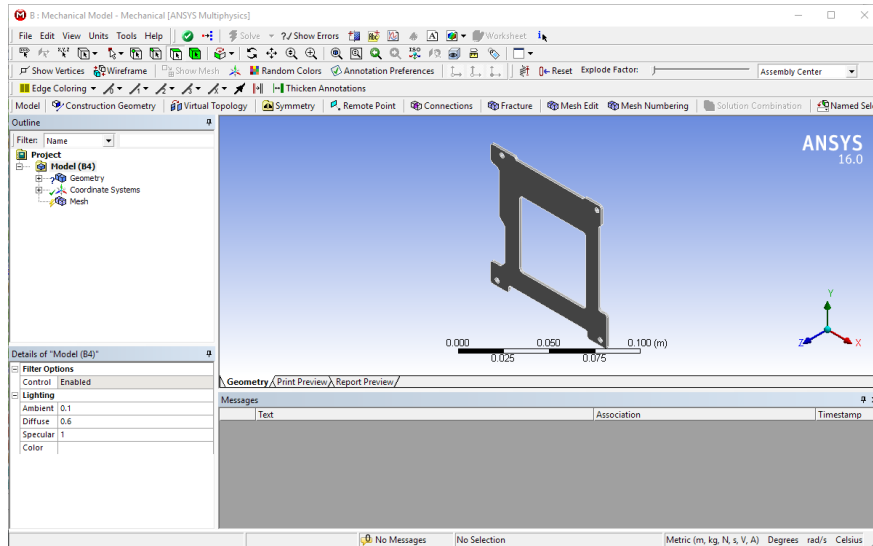


Figure 5: Mechanical Model Start-up

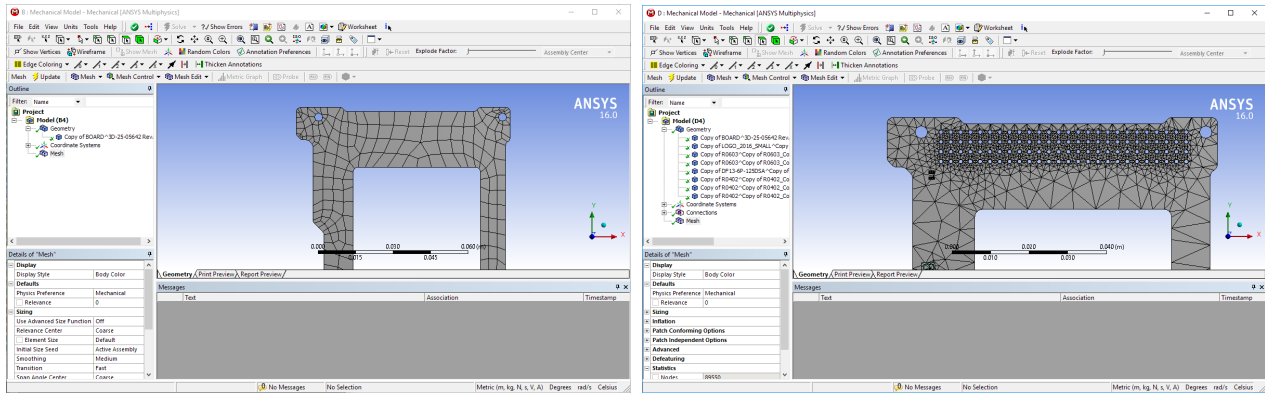
Notice the question mark (?) next to "Geometry". This means that some user input is required. If we expand the geometry menu, there will be a list of all the components that make up the assembly that is being modeled. Currently, this is only one part, but there could be, like, a lot. By clicking on the parts that have question marks next to them, this brings up the details of the part in the bottom left. Under "Material", note it is likely missing "Assignment". This just means the part needs a material to be assigned to it. If you click on the yellow field, this will bring up a drop-down menu of all the materials created in the Engineering Data section.

Once all the parts have materials assigned to them, the Geometry menu should have a nice, satisfying green tick next to it. We can now move on to creating the mesh of the part. This is really why simplifying the geometry is definitely recommended, because this can get seriously out of control if you don't. Click on "Mesh"; under "Details" we can see a bunch of meshing options that will either make your mesh more course or more fine. If the mesh is giving your trouble, your geometry is probably the culprit—ANSYS is smart enough to determine the optimal meshing parameters, so if errors pop up, it is probably your fault. Either right-click on "Mesh" and click "Update", or find the button in the top ribbon with the yellow lightning bolt that says "Update". This will generate the mesh.

Your Mechanical window should look something like Figure 6(a). The view here is to show you why simplifying the model is important. Figure 6(b) shows what the mesh would have looked like if the original model (Figure 1(a)) had been used. For unsimplified models, for one, the meshing will take significantly longer. For comparison, the unsimplified model has almost 90,000 nodes (41,000+ elements), whereas the simplified model has less than 3,000 nodes and about 350 elements. While more nodes and elements will give you a more exact result (which is why you can mess with the mesh parameters), computing the results of the analysis will take your entire life. Here, speed is preferred over accuracy, but exact settings should always depend on your model's requirements.

4 Structural Simulations

Structural simulations always come in four parts: Modal, Response Spectrum, Random Vibration, Harmonic Response and Inertial. The modal analysis should **always** precede the response spectrum, random vibrations and harmonic response analyses, as the dynamic results from the modal analysis



(a) Simplified

(b) Unsimplified

Figure 6: Mesh Results

are used in the proceeding analyses. The setup for each of these analyses will be different, so they will all be covered separately. To create the overall model, simply drag the additional components, and drop them on the "Solution" field of the modal analysis. The inertial (which is a static structural type of analysis) does not require the modal as a prerequisite, so it can be dropped on any "Model" field. You can always drag the analyses around to clean up the project schematic. For structural analyses, the model setup should look something like Figure 7.

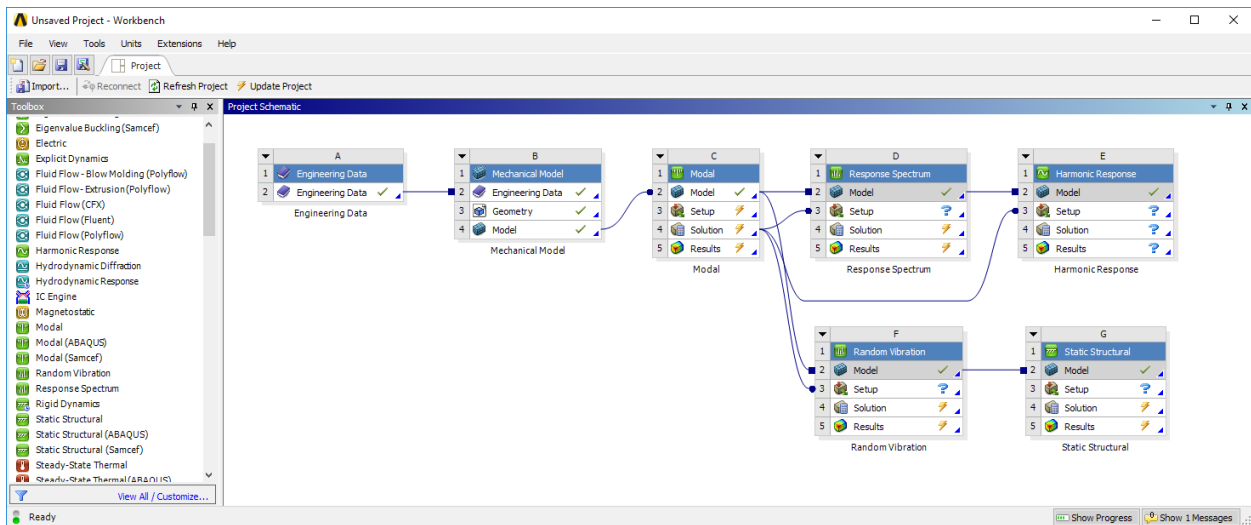


Figure 7: Structural Analysis Project Schematic

Note how all the analyses will either have the lightning bolts or question marks, denoting they require additional setup. These will all be covered separately here. This is a good time to save your project, and for the love of science, save frequently!

4.1 Modal

The modal analysis, as well as the response spectrum random vibration and harmonic response, are vibration-related analyses. A mode, in continuum mechanics, is analogous to the natural (or fundamental) frequency of a spring-mass system. In essence, all finite element analyses are,

spring-mass-like analyses. The process of meshing pretty much reduces a continuous system to a bunch of masses, connected by springs which the underlying solver will then use to compute specific kinds of results. A mode, much like a natural frequency, is the **forcing frequency** at which your structure experiences a local maximum in deflection, at some point in the structure. You will see what this means when we complete the analysis.

Open the modal analysis "setup". Another Mechanical window will appear, but now showing all the different types of analyses that are also connected to the modal analysis. We can stay on this window for a good while.

In general, the first thing you should always be thinking about at this point, is "what are the boundary conditions of this system", meaning, what specific kind of displacement, or temperature, or current or flow or *whatever* do you know about the system at this point. When talking about structural analyses, the boundary conditions are *generally* "Fixed Supports". A fixed support is essentially where the displacement of any node on the model gets set to 0 (zero).

The best way to add any kind of support, is to right-click on the analysis, and select "Insert" > "Fixed Support". Now, this fixed support is going to require a definition based on the geometry of the part, essentially *where* do we set the displacement to zero. For this, we select the inside surfaces of the four PC/104+ holes on the board, as shown in Figure 8.

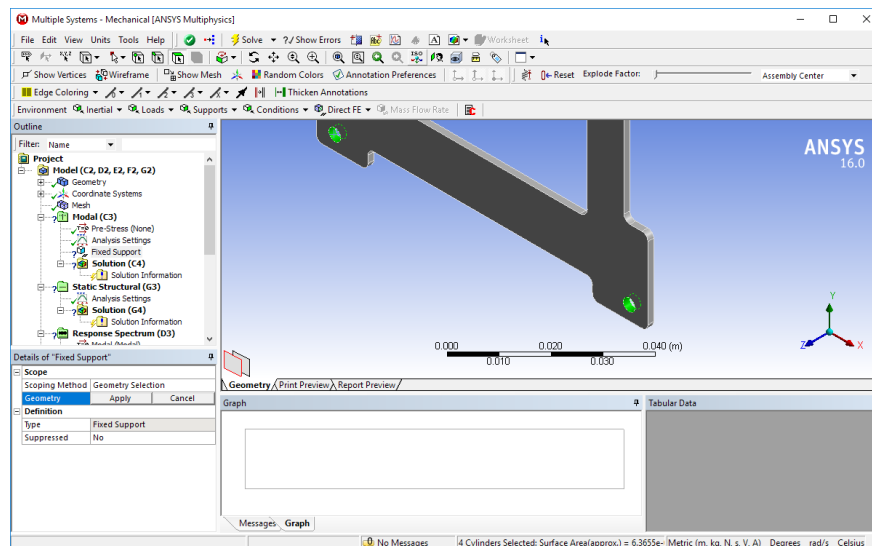


Figure 8: Boundary Condition Geometry Definition

Sometimes, you might need to define a "pre-stress", for something like fasteners, which are already under tension before they are subjected to launch stresses, but we will not bother with that now (you can generally get away with not modeling fasteners in ANSYS at all, there are equations you can use for fastener analysis to give you some conservative estimates). Under the analysis settings, you can select "Max Modes to Find". This defaults to six, which is a pretty good number. A higher number of modes will take more time to solve, but will also make the subsequent analyses give better results. Six is a happy middle-ground.

Realistically, we could call it done here, as the modal analyses will produce a bunch of dynamic results mostly used for further analyses, but let's set a solution type here, too. Under modal, right-click on solution, and insert a "Total Deformation" solution type. Now, **this is important:** the results you get from this, are completely meaningless. They are what is called *mass normalized*, which is a solution technique used in vibrations used to decouple degrees of freedom. If you want

to know more, take the class. Suffice to say, that this *will* tell you the shape of the structure's deformation under vibration. You can add multiple "Total deformation" solutions, and make them correspond to different modes in the details menu, up to as many modes as you have the analysis find (i.e. if you are solving up to the sixth mode, you can only have six total deformation solutions). **Beware: you will only need to extract the frequencies from this analysis!**

You can add whatever kind of solution types you want to this analysis to find—although the most important ones are covered—right-click on the solution under modal, and click "Solve". This will evaluate the modal results for this structure. Depending on complexity, number of modes, and number of elements, this can take a while. When it finishes, in the bottom right, there will be a table with the mode number, and what frequency it occurs at. For this particular model, the first mode is at 684.67Hz, which is fairly high—in modal analysis, high frequencies are good. It means that you will need to excite the structure with very short period, which takes a lot of energy, and is hard to do. Depending on how many total deformations you have the analysis find, you can see what the deflection would look like. The results for this model are shown in Figure 9.

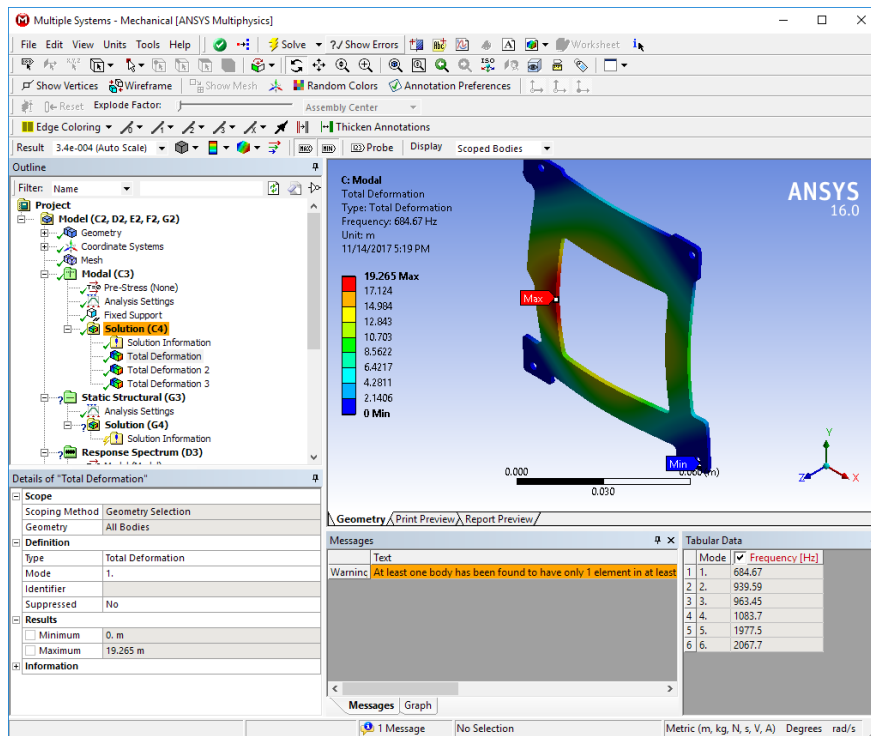


Figure 9: Modal Total Deformation Results

Note that, in the ribbon directly above the project tree, there are several drop-down menus next to "Result". This is where you can change the way solutions are displayed. You can change the scale of the deformation shown, change the type of contour, and add max/min markers to the model. Also note, that the "min" deflection marker is at one of those holes we set as fixed supports, which makes sense because the deflection will be zero there, which is the lowest magnitude of deflection in this model. With this done, we can move on to the next solution, which is related to the modal solution: the response spectrum.

4.2 Harmonic Response

The harmonic response will give you a much better idea of how your structure is going to behave under uniform vibration than the modal analysis, and the two should be complementary. The harmonic response does require the modal analysis as a prerequisite, but it returns deflections of your structure that are not mass normalized. Rather, it will use the dynamic behavior profile from the modal as input data, and returns the actual deflection of the structure under vibration. What the harmonic response does, is force excitation at the fixed supports defined in the modal analysis at a range of frequencies. Essentially, it uses the equation $f(t) = F \sin(ft)$ to vary the forcing sine wave over a range of frequencies. You will see, over different frequencies, you will get different responses.

The first thing you should see in your harmonic response, is that the modal data is being transferred in. There is then also a question mark by the analysis settings menu. The only input required here, is a range of frequencies you want to see evaluated in your harmonic response. For demonstration purposes, the range here will be set from 0 to 1000Hz, to be able to show you at least the first three modes. This analysis can be quite slow, so you should evaluate how many solution intervals (in Herz) you want. Here, it is set to 10000 (more points will make this take longer, but will give you a smoother curve).

The next thing is to define a forcing boundary condition. Basically, we will add an inertial acceleration here. Go up to the environment bar to the top left, and add Inertial > Acceleration. Under the definition for your loading, set "Define by" to "Components", and we're going to set the z component to $1\text{mm} \cdot \text{s}^{-2}$, however, you should perform this analysis on all axes separately (x , y and z) to get a complete picture of all possible responses. The act of setting this to $1\text{mm} \cdot \text{s}^{-2}$, essentially normalizes our response to the forcing amplitude, meaning that the actual response amplitude A will increase linearly with the amplitude of the forcing sine wave. We will see how this affects our results later.

Under solution, add a "Frequency Response" deformation solution. This will require a piece of geometry to parse for deformation, you can select all geometry in your part by using **Ctrl+B** to enable the body select tool, and then **Ctrl+A** to select all bodies in your analysis. Or, alternatively, you can just select the ones you are interested in, but since your entire body will be under excitation, a good first step will be to select all bodies, and trim it down from there if it's taking too long. Apply this geometry, and your simulation should be ready to go. Right-click the solution, and click Solve. This is going to take your structure, and evaluate it with a forcing amplitude $F = 1\text{mm} \cdot \text{s}^{-2}$, at the range of frequencies f you have specified.

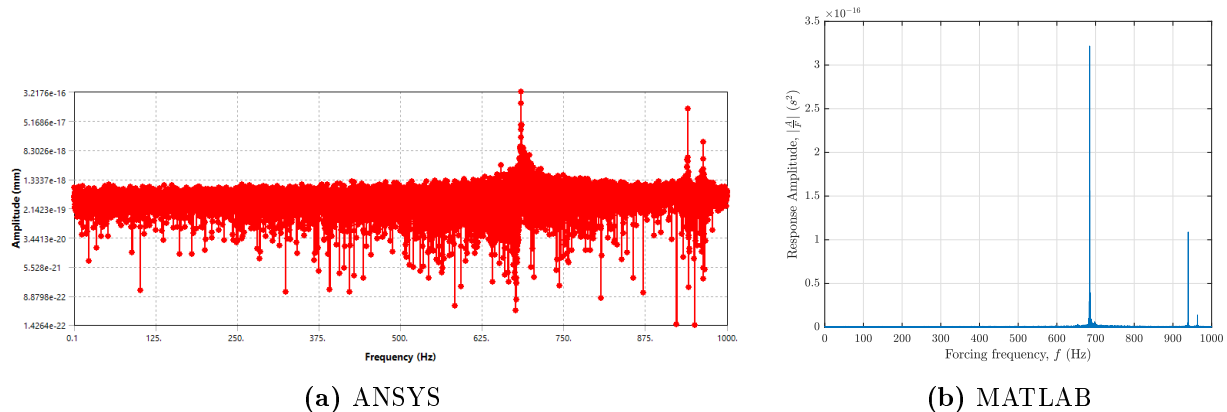


Figure 10: Harmonic Response Results

The results of this analysis are presented in Figure 10. The left is produced directly by ANSYS, which has the y -axis on a log-scale (you can turn this off, but presentation of graphs using MATLAB is always preferred, and less messy), and the right is produced with MATLAB on a linear scale. From the left graph, you can clearly see a spike in amplitude at just under 700Hz (684.67Hz, to be exact). This is the response amplitude from the first mode. The other spikes you can observe at slightly higher frequencies, are the second and third modes.

Note from the y -axis label on the MATLAB curve, that the response amplitude looks like $|\frac{A}{F}|$. This is what "normalized by the forcing amplitude" means. Essentially, because the frequency response "*transfer function*" is a linear operation, the response amplitude scales linearly with the forcing amplitude. So if this system were forced with $2\text{mm} \cdot \text{s}^{-2}$, the response amplitude would scale with the same factor 2. If you can quantify what the acceleration of the system is, you can quantify what the response amplitude will be.

4.3 Response Spectrum

The response spectrum analysis is somewhat hard to explain. Essentially, it is a type of shock analysis, and shock is a type of vibrational excitation. Essentially, given an excitation profile, this will give you the peak deflection your structure will see in the environment described by your profile. It is not like modal deflection, but its results depend on what the modal analysis looks like.

To begin the setup, there is a pre-determined shock spectrum supplied to us. Start by adding a "RS Acceleration" boundary condition to the analysis. Because the analysis should be linked to modal, it will take the fixed supports from the modal analysis. In the RS Acceleration details, next to "Boundary Condition" select "All BC Supports". This essentially forces the system at the bases, which is conventionally how a shock analysis is set up. Next, the "Load Data" field is also yellow, this means we will have to supply a table of values of frequency and acceleration. These values are given by the graph in Figure 11. For convenience, the data for the table is also provided in Table 1, where the acceleration in G's from the graph has been converted to $\text{m} \cdot \text{s}^{-1}$.

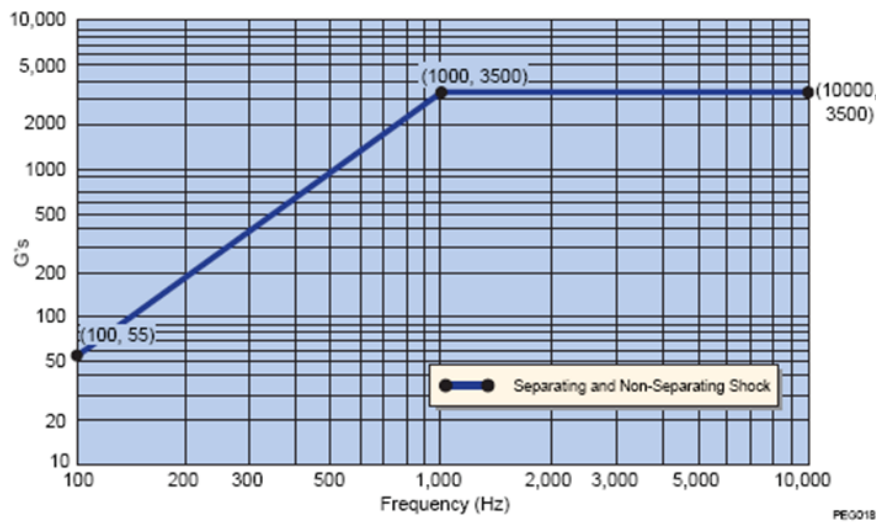


Figure 11: Component Shock Spectrum

Frequency (Hz)	Acceleration ($\text{m} \cdot \text{s}^{-2}$)	Acceleration ($\text{mm} \cdot \text{s}^{-2}$)
100	539.55	$593.55 \cdot 10^3$
1000	34335	$34.335 \cdot 10^6$
10000	34355	$34.335 \cdot 10^6$

Table 1: Component Shock Spectrum–Tabular Data

Lastly, the details for RS Acceleration require a direction in which the forcing load is applied. Because launch orientation is not generally known beforehand, the analysis (including all other analyses that are orientation dependent) will need to be run for forcing in all directions, x , y and z . For now, this example will be run with the forcing direction along the z -axis.

Under the solution for Response Spectrum, add a "Total Deformation" solution, and an "Equivalent Stress (Von Mises)". The stress will allow us to determine whether the part is likely to fail under shock loading. No other input is required at this point, so you can solve the model.

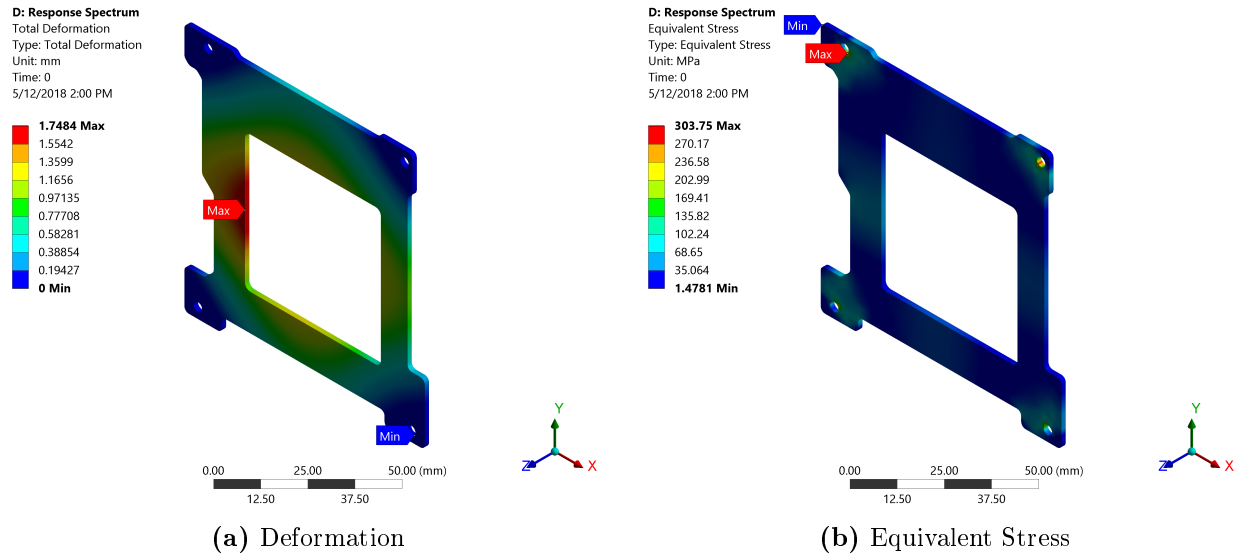


Figure 12: Response Spectrum Results

The results are shown here in Figure 12. Note how the deformation figure has the exact same shape as the modal analysis' first mode result. This should tell you, that the structure responds at a frequency lower than the first mode. Interpreting the data from this analysis can be quite tricky. It is not as straight-forward as saying that this is exactly what is going to happen to your structure. You have to remember, shock will only really occur a handful of times on your structure, and neither you, nor ANSYS, know at what frequency this will occur. If your modal analysis is predicting a high frequency for the first mode, you can rest assured that the response predicted by your response spectrum analysis will be an absolute worst-case scenario; your structure will likely not see these kinds of stresses, or deflections. The big take-away, is that response spectrum results will give you large numbers, but it should be on you as an engineer to interpret how much you ought to worry about that.

4.4 Random Vibration

The phenomenon of random vibration is a statistical one. Think about a series of random numbers in some interval, plotted against time. If you were to extract the frequency information from that (using something like a Fourier analysis) you can't. All frequencies are present in a randomly generated signal. This means, that if your structure is in a scenario where it is not being uniformly forced, but instead has some random forcing component (very common in rocket launches), your structure is going to respond to every single forcing frequency at once. Let that sink in for a second. In random vibration, your system responds to every forcing frequency simultaneously, including the modes.

The way to quantify this interval of forcing amplitudes (the range of random numbers), we use *Power Spectral Density*, or PSD. This can be given in terms of acceleration, displacement or velocity, but essentially quantifies the amount of average power present in the frequencies (don't worry about what that means). In your analysis, add a "PSD G Acceleration" boundary condition, and set it to all fixed supports. Then, add the data specified in Table 2 in the tabular data section, and set the direction to the axis you want to evaluate (again, you will need to do all of these separately).

Frequency (Hz)	G Acceleration ($G^2 \cdot \text{Hz}^{-1}$)
20	0.026
50	0.16
800	0.16
2000	0.026

Table 2: Random Vibration Qualification Profile–Tabular Data

The random vibration analysis, because of its statistical nature, produces statistical results, evaluated along a single axis. So, even though you have set the forcing axis to one, you will have to add *three* "Directional Deformation" solution types. Each of these, will be set to one of the three axes. So the first will be set to x , the second to y and the third to z , or whatever order you want, as long as all axes are represented. Because we are dealing with probabilities, you will have to define your confidence interval. To assure the structure is going to behave well under random vibration, set the "scale factor" to "3 Sigma". This means that you have $3\text{-}\sigma$ confidence, or 99.73%, that the actual deflection will be lower than predicted by this analysis.

Once you are set up, run the analysis. The results for this demonstration are shown in Figure 13. Note how there are three separate figures, each representing an axis. Generally, random vibration results are rather small if the first mode is very high. For a different system with a low first mode, this could be in the order of several millimeters, so it depends entirely on the modes. Also note, the z -axis deformation shown here, has the exact same shape as the modal analysis showed. These are good checks to make sure your results make sense, and you have not made a mistake somewhere.

4.5 Inertial Acceleration

The last type of structural simulation is probably the simplest one, plain old acceleration. Since we know that accelerating applying a force on a mass (like thrust from a rocket) causes it to accelerate, and vice versa, we can quantify the acceleration from the rocket on our structure, and determine the deflection.

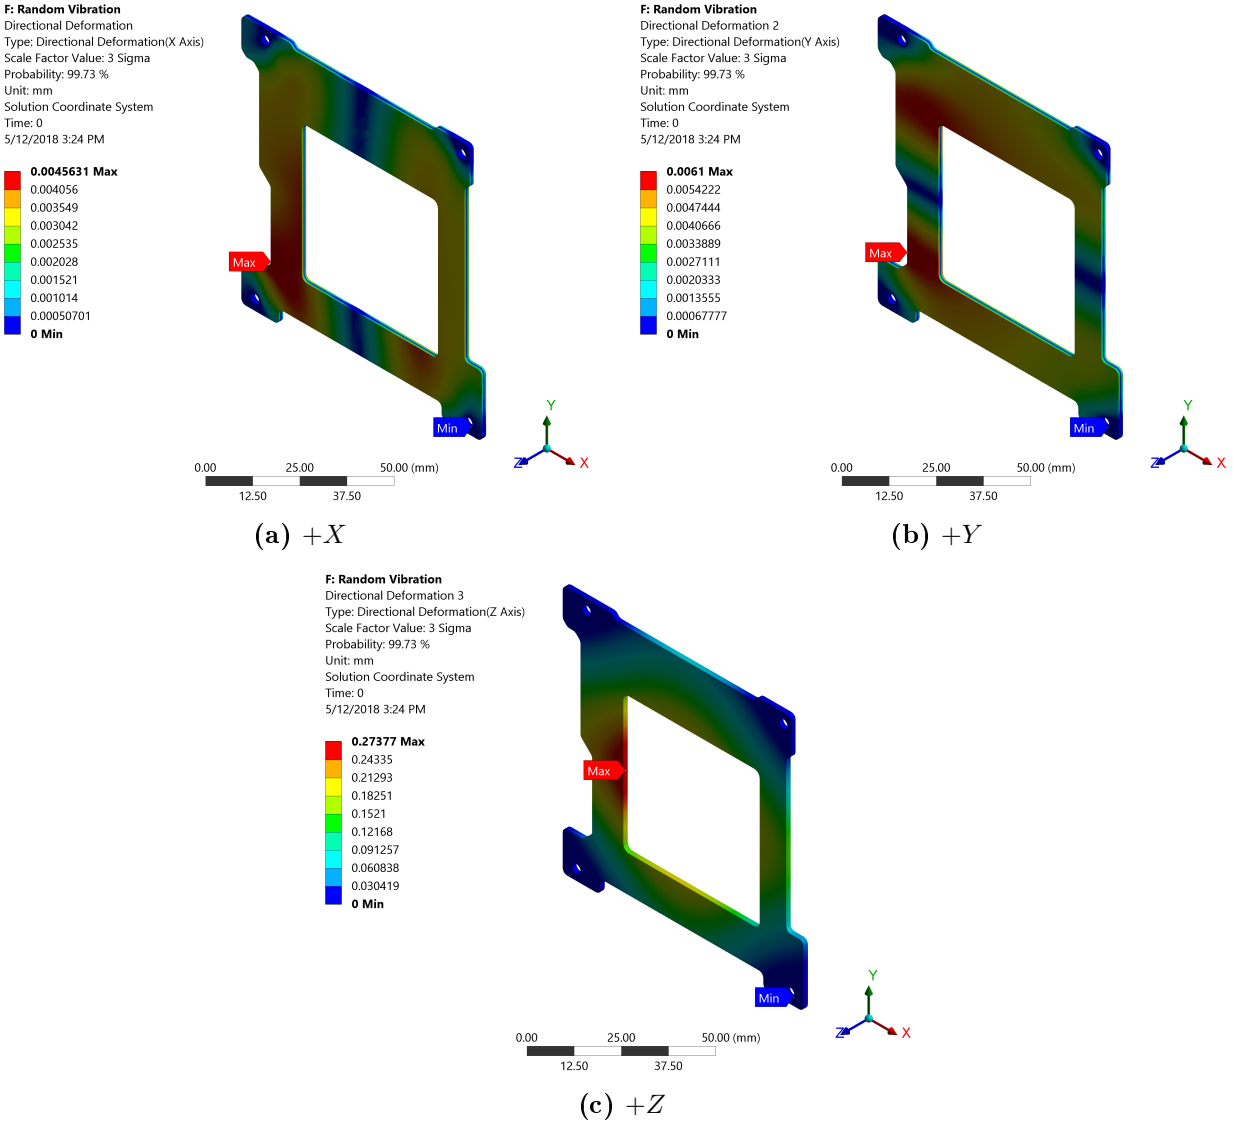


Figure 13: Random Vibration Directional Deformation Results

We will use the Static Structural analysis type for this, since inertial acceleration is also called static acceleration. Add the same boundary condition as for the harmonic response, and set it to define by components. Now, if you thought this was straight-forward, you would be wrong. As it turns out, structures of different mass will experience different accelerations in the same rocket. Meaning, your 4kg CubeSat will accelerate differently from the Hubble Space Telescope, which has a mass of about 11 tonnes. This acceleration is specified by the Mass Acceleration Curve (MAC), shown in Figure 14.

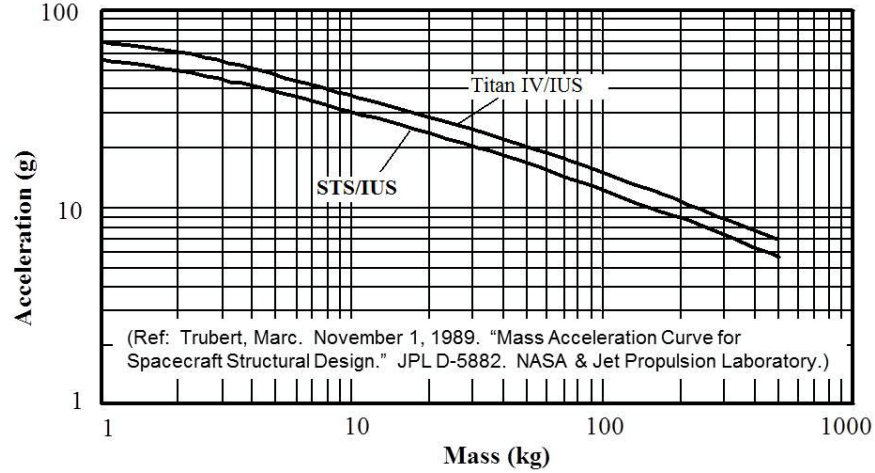


Figure 14: Mass Acceleration Curve

However, the MAC can be very inaccurate at low masses. It is recommended to use the acceleration corresponding to 8kg to anything that is lighter (or approximately $33G = 323.73m \cdot s^{-2} = 323730mm \cdot s^{-2}$), even though a 4kg satellite should accelerate closer to 42G's ($42G = 412.02m \cdot s^{-2} = 412020mm \cdot s^{-2}$). Whatever the satellite mass, as predicted by the Mass Budget, should be used first to evaluate the acceleration per the MAC. The larger acceleration should be used as an initial, more conservative benchmark, although both can be checked if necessary.

Because this is not a vibrations analysis, the fixed supports from the modal do not carry over, and will need to be redefined. You are free to copy the fixed supports from the modal, and paste them in the inertial. Then, all there is left to do, is add a total deformation and equivalent stress solution type, and run it (again, remember you have to run this analysis with the acceleration along each different axis, and present all results).

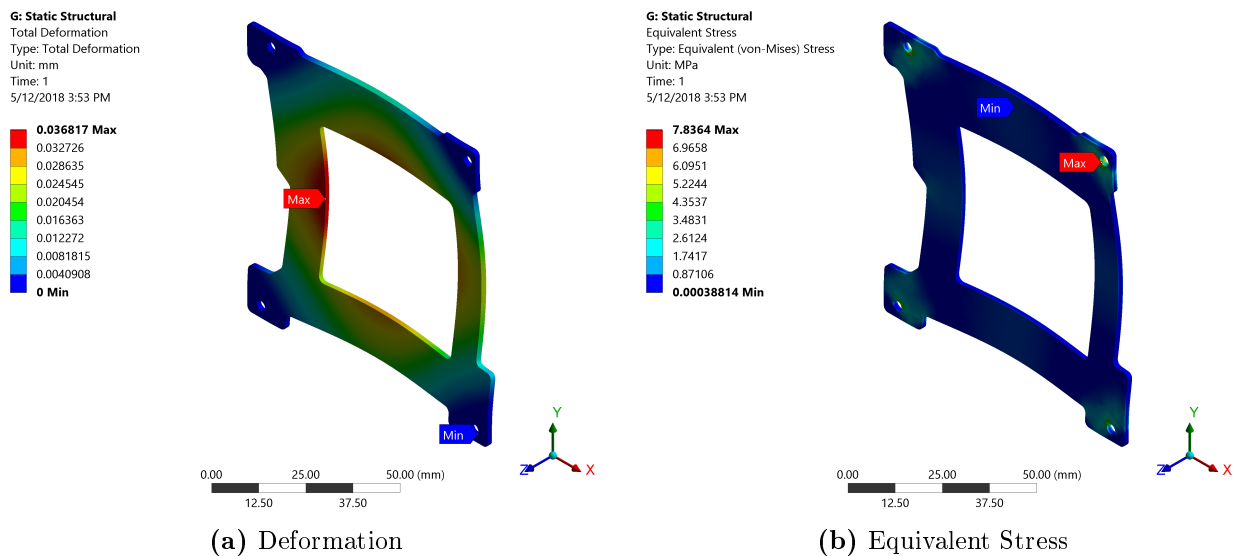


Figure 15: Inertial Acceleration Results

The results are presented in Figure 15. Note the shape of the deformation is completely different from the vibration analyses, which should make sense. The inertial deformation and stresses are a

component on top of everything that happens due to vibrations, but combining all these results is not as straight forward as just adding them. It should be up to you as the engineer to determine whether these results are acceptable. The usual way to do this, is by using a factor of safety. this is given by the equation $n_d = \frac{\text{predicted stress}}{\text{loss of function stress}}$. The loss-of-function stress, is usually the yield stress of the material, although the ultimate strength is also used. A minimum n_d of 2.0 is required for yield, and 2.6 for ultimate. Material strength specification can be readily found online.

5 Thermal Analysis

Thermal analyses are somewhat different from their structural counterpart. In structural models, we can see things deform and have a much better intuition for something being correct. Thermal models are quite different, since it deals with the more abstract quantity of temperature. Especially in satellite applications, where convection is not present, results can be somewhat counter-intuitive.

We must begin, by adding the thermal material property data to the engineering data. As discussed in Section 3.1, we will need density (which we have from the structural model), isotropic thermal conductivity (technically FR-4 is anisotropic because it is a composite material, but isotropic is a good enough approximation for now) and specific heat. Once these have been added to the material data, set the conductivity to $k = 0.294\text{W} \cdot \text{m}^{-1} \cdot ^\circ\text{C}^{-1}$, and the specific heat $c_p = 1150\text{J} \cdot \text{kg}^{-1} \cdot ^\circ\text{C}^{-1}$. With the material data in place, we can begin the thermal analysis.

5.1 Model Setup

Besides the steps needed in model setup defined for structural analyses, some more might have to be done for thermal analyses. For structures that have *enclosures*, meaning any surface of material that "sees" another surface that is not at the same temperature as it, these must be properly defined so that all the view factors are accounted for. Your working geometry might not require this, but in the case that it does, this tutorial will go over the setup of how to work with it.

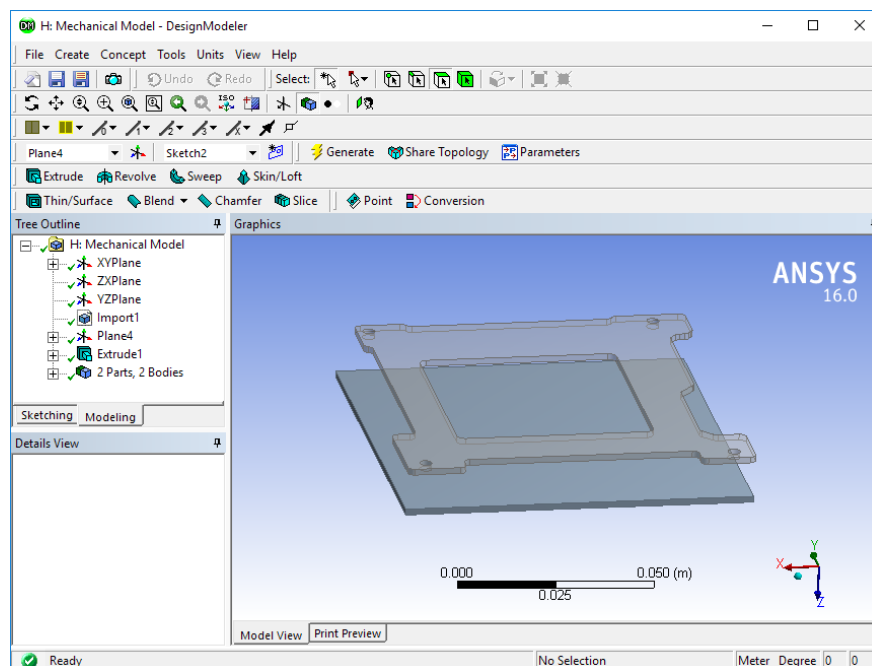


Figure 16: Geometry for Thermal

Figure 16 shows the new working geometry, which simply consists of a new piece of material (could be another board, who knows) offset from the Z-magnetorquer. They will have the exact same thermal properties, for simplicity.

Once the mesh has been created, we will need to create named selections. To do this, go into the model where you would normally be creating your mesh, right click on "Model" under Project in the Mechanical window, go to Insert > Named Selection. From here, select *all* surfaces that make up a single enclosure. If there are multiple enclosures, you will have to make multiple named selection.

If you have more than one enclosure, you will have to do some fancy stuff. Namely, you will have to add a few lines of APDL command code. This can be accessed by right-clicking on your analysis, and inserting a "Commands" environment. There is a supplementary document, provided by ANSYS, that tells you how to do this, but it is summarized here. Each enclosure gets a number. Since we only have two surfaces that see each other (therefore one enclosure), and we are assigning them the same emissivity (the fraction of thermal radiation a body emits relative to a perfect emitter, i.e. a black body), we can do this with two lines:

```
SF, <named selection>, rdsf, <emissivity>, <enclosure number>.
```

The `SF` command specifies surface loads on nodes (i.e. a surface-to-surface type radiation, and the `rdsf` label specifies the surface-to-surface condition. By default, the named selection will be called "selection" but can be named whatever you want. With a single surface-to-surface condition, you do not have to do this, because it is automatically designated enclosure 1.

Once all enclosures have been designated with the command line above (all of them have to live in the same command environment as separate lines), you can specify radiation boundary conditions as surface-to-surface using the appropriate enclosure number. Do note, that each surface may have **only one** radiation condition assigned to it. This means that a surface-to-surface condition cannot overlap with a "to ambient" condition, since the surface-to-surface condition should be specified as "open" if it also sees the environment.

5.2 Transient Thermal

A transient thermal analysis essentially tells you what the temperature profile of a structure looks like over time. Unlike structural models, which use structural mechanical boundary conditions and forces as loads, thermal models will have temperature-specific boundary conditions and heat flow as loads. The governing equations for structural and thermal models, are quite similar, but different in certain subtle ways that warrant separate explanation in this guide.

Ignoring the structural models, and making the edits to the geometry and engineering data outlined in section 5.1, the model should look something like Figure 17. Your geometry might not look exactly like Figure 16, but this guide should serve as a general setup procedure for similar project.

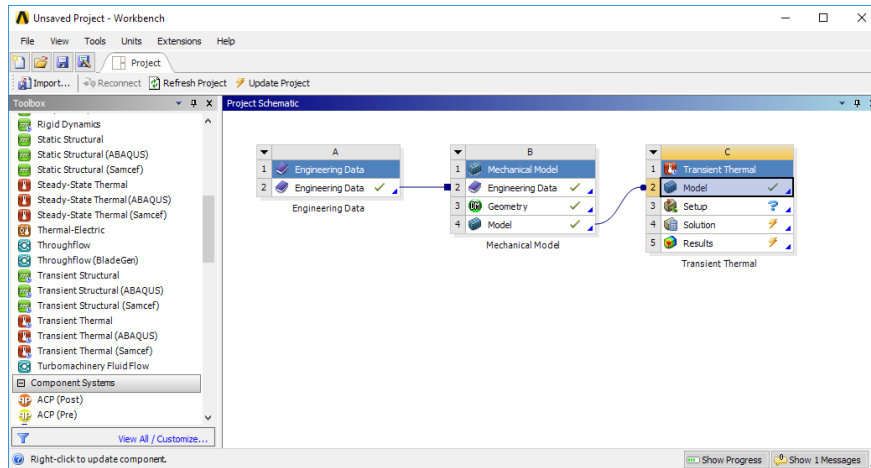


Figure 17: Transient Thermal Project Schematic

Having followed along with section 5.1, and done the named selection setup portion, your screen should look something like Figure 18. The first thing you should be thinking about, is what is the initial temperature of your simulation. The initial temperature is going to determine how much time it is going to take for your model to converge to a steady temperature. For this model, we will set the initial temperature at 20°C.

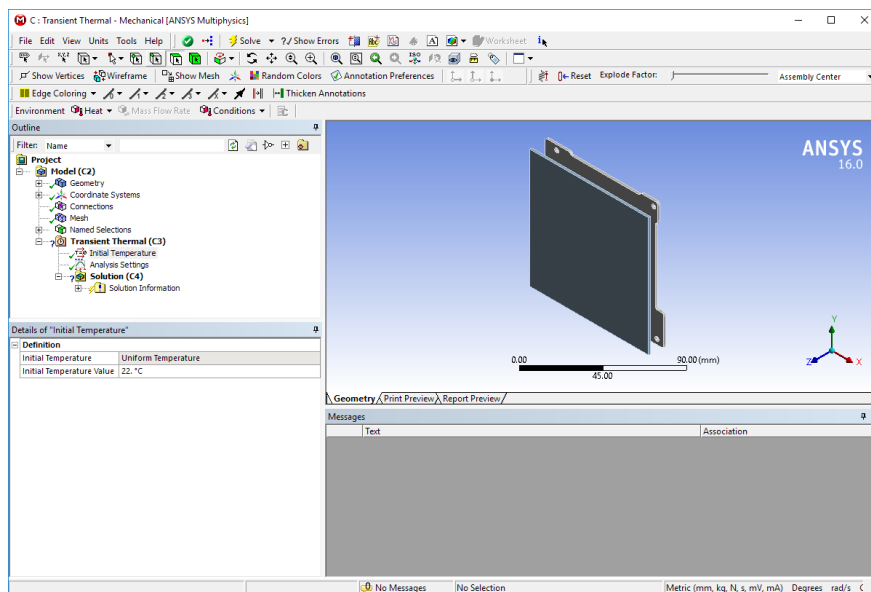


Figure 18: Transient Thermal Setup

Even though this model will only have one enclosure, this will be a good opportunity to show you how this might work. Right-click on the transient thermal project, and insert a "Commands (APDL)" setting. This should bring up a text editor window for the scripting. It should look something like Figure 19. Since the default name for a named selection is "Selection", and this was left as default for the purpose of this demonstration, enter the following line in the command window (again see Figure 19, if you have multiple enclosures, you *have* to do this, just keep adding lines):

SF, Selection, rdsf, 0.9, 1.

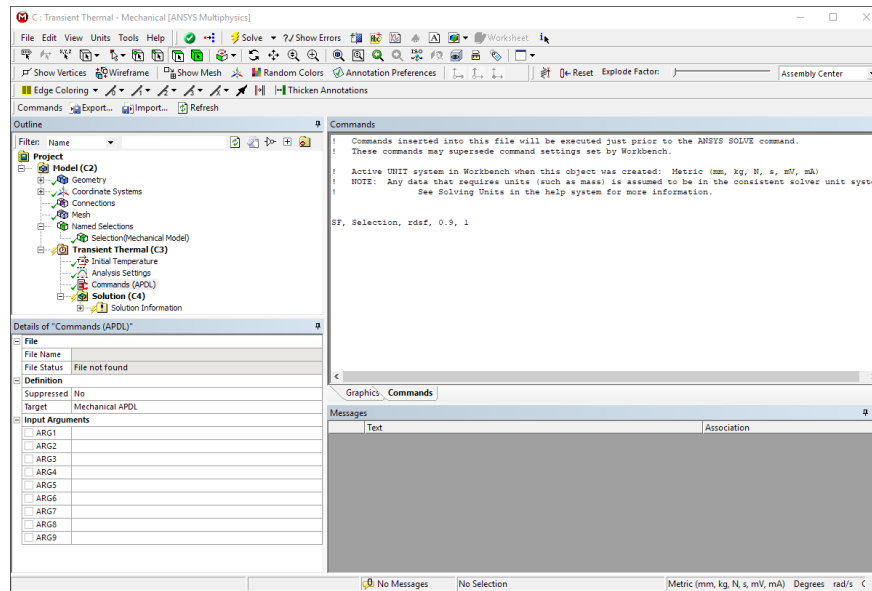


Figure 19: Command Window

Next, we will think about the boundary conditions before we work on the analysis settings. The reason for this, is that it is likely easier to quantify your loads first, and come up with a prediction of what your temperature profile might look like marching through time. This will then make it easier to set your step and solver controls to the proper settings. It is unlikely that you will get the analysis settings right the first time, but this will minimize the amount of time you might spend messing around with it.

The thing about ANSYS, much unlike Thermal Desktop, is that it will come up with conduction models for you. Thermal Desktop requires you to specify all the contacts, and give them some interface conductance, which ANSYS does not. ANSYS basically takes the mean thermal conductivity value (k), and assigns that as the interface conductance. This might be inaccurate, and so you might have to change this in the "Connections" menu to the left.

The other kind of boundary condition that exists, is radiation (since convection is not present in space, although adding convection conditions is very easy if you get the general setup of boundary conditions). Add in a radiation boundary condition, by either clicking the "radiation" button in the environment bar, or right clicking on the transient thermal project, and inserting it that way. Let's start with the surface-to-surface radiation condition. Since we already made the named selection earlier, change the "scoping method" under scope in details, to "Named Selection." Then, simply select the named selection you want this to apply to. Change "Correlation" from ambient to "surface to surface", and make sure the enclosure number and emissivity are the same as assigned in the script line. If the enclosure is "open" (meaning some of the surfaces can still "see" the ambient temperature, make sure the enclosure type is set to "open" as well. This will require you to set an ambient temperature, since the net radiation heat transfer depends on this. The ambient temperature is a difficult parameter to quantify correctly. A lot of back-and-forth between simulations might need to happen for you to manually converge this to a single number (or set of numbers). The recommended way to do this, is make a relatively simple, but representative model

in Thermal Desktop (or MATLAB, if you are not yet in this stage of modeling), and take a low bound, and high bound of your temperatures, and run this simulation twice, once for each bound. This will be similar to the way your hardware should be tested in thermal vacuum, but can be time consuming. It is what it is. For the purposes of this model, the ambient temperature will be set to 0°C, which is fairly common for a low-bound temperature extreme.

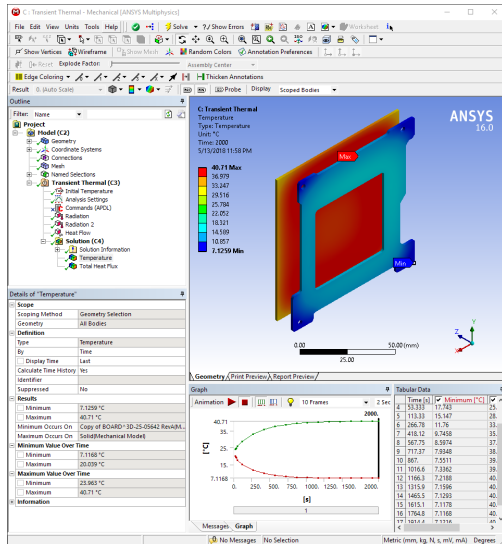
Since we do not have any more enclosures, the rest of the surfaces will be set up as radiating to ambient. It should be noted, that the entire satellite structure is not easily modeled in ANSYS, and should therefore be reserved for Thermal Desktop (i.e. only model components *inside* the frame in ANSYS thermal simulations). Similarly to the surface-to-surface radiation, the ambient temperature will be taken as 0°C with 0.9 emissivity (specialized black paint). When selecting the geometry this applies to, select every surface on the model, *except* the ones you have already assigned to another radiation boundary condition. You are not allowed two boundary conditions on the same surface, just like you cannot have a fixed support and a displacement on the same face of a structure. Your solver will not know which one to chose, and it will yell at you if you try to do this.

Lastly, we will add a load to our structure. Heat loads come in a few different flavors. There is "heat flow" which is generally the simplest to use, you basically apply some amount of Watts (or heat power) to a surface, and it will equally distribute it (this is very nice for electronics and components), there is "perfect insulation" which basically sets the heat transfer on some area to zero (you might think of radiation shielding for this), "heat flux" which is heat power per unit area (think of the solar heat flux incident on the satellite when it is not in eclipse, i.e. heating rates depend on the profile of the satellite shown to the sun), and "internal heat generation" which is heat power per unit volume (you might think this is easiest to use for electronics, but it can be obnoxious to keep track of all the volumes of your parts to normalize them, and they are usually so small that it does not matter if you use heat flow). Go ahead and add a "heat flow" do the bottom of the second board (i.e. not the Z-magnetorquer). Make sure that the magnitude of the heat load is set to "constant" (so that it says "step applied") and not tabular (unless you somehow wish to do this, but that will be slightly more advanced). Specify the value to 3W.


There is another type of boundary condition, which is a fixed temperature. This is not commonly used for real applications, but could be useful for troubleshooting your model. If you were to draw a box around your model, and gave that a fixed temperature, and included it as a surface-to-surface radiation type, instead of to ambient on your model, the results should be the same.

In transient thermal, there are a few solution types that might make sense for you to add. Temperature is the obvious one, but for design and possibly troubleshooting reasons, you might also want to add a "total heat flux" solution type.

Now that all the boundary conditions have been set, the model is almost ready to go. Remember, that we skipped over the analysis settings earlier, so we need to return to that. The main thing you should be doing here, is messing with the amount of time you are running the solution for. You can change the solver controls and convergence criteria, but the default should be adequate for most models, unless you are trying to speed it up while sacrificing some accuracy. For this model, we will specify the "Step End Time" to be 2000s (which was previously determined through some trial and error). You can decide how the solver marches your model forwards in time by setting auto time stepping on or off. If you set it to "on", ANSYS will determine the time step based on the rate of change of the temperature (if the temperature's "slope" starts to change, then it will start to use more points to evaluate your model, whereas it will use fewer points if the temperature slope is somewhat linear), but we will see what this means. If you want your temperature and heat flux returned, say, every second, you can turn the auto time stepping off, and specify the time step manually. We will leave this on for now. Your model is ready to go, so go ahead and solve it.



(a) Colormap Plot



(b) Time Response

Figure 20: Transient Thermal Results

The results of this model are shown in Figure 20. On the left is the colormap plot generated by ANSYS. You will notice, at the bottom of your solution screen in ANSYS, there is the exact same curve as it plotted in MATLAB here, along with some tabular data. The tabular data was used to create the MATLAB plot shown to the right. You can use the graph to control which time step is shown in your colormap plot, and you can play this out as a movie, and save it as a motion-picture file format if you so desire.

Some important things to think about here: what would have happened if we did not specify the surface-to-surface radiation condition? ANSYS would basically not know that the two pieces of material were able to "see" each other. So the Z-magnetorquer would basically cool down to the ambient temperature, which is clearly not the case in this model (it is kept slightly warmer by the radiation coming from the other board). What would have happened if we specified the initial temperature to be different? Say we set the initial temperature also to 0°C , then we would basically see both parts temperatures increase, since they will always reach their steady-state at the same temperature, regardless of initial conditions, it might just take them longer to get there. You are encouraged to play around with all the different parameters here (just as you are with Thermal Desktop), to get a feel for how a certain value changing one way or another, changes the temperature profile you observe.

5.3 Steady-State Thermal

Once you understand the gist behind transient thermal models, setting up and running a steady-state thermal model is quite straight forward. Actually, it is completely the same. You can use the exact same boundary conditions in a steady-state model, as you can a transient model.

However, there are some subtle differences in steady-state models you should be aware of. The main one being that you will not get to see how your model's temperature changes over time. As you might be able to tell from the name, a steady-state model basically just tells you the temperature of your model as time goes to infinity (or goes sufficiently large that your solver stops seeing variations in temperature and will call it steady-state to within some convergence criteria).

With your model going to steady-state, your initial temperature condition does not matter. Since time is a removed component (it will not even tell you how long it will take for this to happen, it just returns temperatures) all the transient effects, like initial temperature, will no longer be present, rather it will simply be whatever your very final temperature will look like.

The main reason for running a steady-state thermal model over a transient, is speed. A steady-state model will give you less information about your model than a transient simulation, but steady-state models can be orders of magnitude faster than transient. A transient simulation basically has to solve the same equations as the steady-state model, over a specified time-step, just tens to hundreds of times over.

A few notes about the steady-state results. If you let your transient model run for long enough, you should get the same profile, and temperatures as a steady-state model. However, the actual numbers might differ ever so slightly. You have to keep in mind, that ANSYS is a numerical solver, and is thus subject to error and conversion intervals. This might mean that your transient temperatures are actually *slightly higher*, than the steady-state predicted temperatures. As long as this is within a hundredth of a degree (higher or lower), this is likely just that numerical error rearing its ugly head, but if it is in the order of degrees, you might need to check you models.

5.4 Results Interpretation

It was mentioned earlier in this guide, that quantifying your ambient temperatures can be quite difficult. In general, thermal models are quite hard to run, and specify your boundary conditions correctly. If you are taking data from programs like Thermal Desktop, or even MATLAB, the actual temperatures are very sensitive to the time-step used in your model, as well as as a multitude of other factors.

To account for the things you might not be certain about, it is common to "pad" your temperature data with some empirically determined interval to make sure that your data will fall within a certain confidence interval. Yeah, you didn't expect statistics to pop up here too, did you...




Figure 21: 2- σ Corrected Temperatures

The extra step here, is rather simple. Essentially, per this paper from the Aerospace Corporation on testing for space vehicles, the interval you should use for your temperature data to bring it within acceptance-test levels, is 11°C. This means, take the highest temperature predicted in your model (easily acquired from steady-state thermal), and add 11°C to it. Then also take the lowest temperature predicted, and subtract 11°C from it. This will be the 2- σ confidence interval corrected temperature range you might expect your hardware to see, based on the models you have run. This interval is to account for errors in the model at every step of the way.

The corrected data will look something like Figure 21, as applied to the transient thermal results from Figure 20. There are other ways you might present your data, based on what you need. If you are running steady-state thermal models, you can get away with simply adding or subtracting 11°C. *Data interpretation is an important aspect of doing good engineering work*, so always ask yourself if your data makes sense, and if the way you have set up your simulation makes sense.

6 Conclusion

If you made it to the end of this guide, congratulations. I hope you are slightly more comfortable with ANSYS because of it. As with the Thermal Desktop guide, this guide is by no means designed to introduce you to every single option available in ANSYS, but it should expose you to enough stuff, and give you the tools to go and explore for yourself.

ANSYS will most likely yell at you for doing something wrong at some point in your time using it. This happens. ANSYS is finicky and particular about how it wants things. If it yells at you, ask someone how to fix it, or toss a quarter into the Google-o-matic, and hope to science someone on the internet will tell you how to fix it.

You should beware. ANSYS is an incredibly powerful tool. It *will* solve any system you give it. However, whether that model you give it is correct to begin with might is up to you to determine. If you give ANSYS an incorrect model, it will give you incorrect results. Use your intuition. If your results seem wrong, try to evaluate why they are wrong, and take careful, procedural steps to try to fix your results. Talk to experts. Do what you need to do to make sure your models are correct.