TR500017

Lightning Strikes Twice: Revisiting Generative Design for Mass Production

Peter Simpson
Autodesk

Richard Hatfield
Lightning Motorcycles

Nick Markovic
Autodesk

Learning Objectives

- Learn how to apply generative design manufacturing methods to create designs suitable for mass production
- Learn how to implement an integrated generative design and validation loop within Fusion 360
- Learn how to apply generative design to multiple bodies with a singular assembly
- Learn how to create a workflow to enhance the cost-effective manufacturing of generative design outcomes

Description

Lightning Motorcycles has been breaking records in the world of e-bikes since the company entered the market. The LS-218 is pushing electric and gasoline-powered motorcycles to be the fastest production bike on the market. But pure performance is not the only goal; Lightning is focused on providing consumers with the best quality and value in every product it creates. Key to meeting these goals and staying ahead of the competition are new technology and ways of thinking. Lightning was an early adopter of generative design. Collaboratively, we produced a part-consolidation and lightweighting workflow that saved 30% of target component mass. In this class, we'll show you how developments within generative design can be used to gain these benefits, while ensuring cost-effective manufacturing. This talk will cover workflows used within Fusion 360 software, maximizing the integrated simulation, generative design, and manufacturing tools to better understand the part and its performance.

Speaker(s)

Peter attended the University of Birmingham, graduating with a master's degree in mechanical engineering. He began his career with Autodesk during a summer internship and has since re-joined Autodesk as a Graduate Technical Consultant working in the Birmingham office, recently taking a full-time role in the Process Specialist Team. He has worked on a variety of projects, often focusing on the utilization of Generative Design within different industries, helping to drive the adoption of the platform and further develop the software. In his spare time, Peter is a keen sportsman, playing football, rugby and golf on a regular basis.

Richard is the founder and CEO of Lightning motorcycles. He founded the company after being invited to work on some leading projects within the field of electric vehicles. As a life-long motorcycle rider, he spotted a gap in the market that would allow the fields of electric vehicles and high-performance motorbikes to push the possible performance of these bikes. Lightning has since gone on to hold multiple speed records and race victories, often being the first electric powered bikes to even compete, let alone win these prestigious events.

Nick is a well-rounded individual with 11 years' experience as an Aerospace engineer coupled with a strong academic background. Nick has also gained key exposure in the oil and gas, wind, nuclear, and manufacturing engineering and automotive markets. He is a specialist in computer-aided engineering, with an emphasis on multiphysics optimization. Nick has joined Autodesk Research as a Senior Researcher in Manufacturing Industry Futures focusing research on manufacturing digital twin, creating novel Generative Design workflows, and smart design. Nick is currently managing Project MOnACO which is funded by the CleanSky program aiming to manufacture the world's biggest laser powder bed fusion jet engine part.

Fusion 360 Generative Design

Generative Design within Fusion 360 is a design exploration tool, aiming to cut down the necessary steps within the standard design exploration cycle. It is problem focussed, with real world loading and geometry inputs driving the design iterations and allows multiple viable designs to be created simultaneously. *Figure 1* shows the standard design to production workflow, in comparison with the Generative Design workflow. The grey workflow shows a few ideas being created, some key ones being chosen to explore further and then a complex loop of validation via design for manufacturing, simulation etc. The Generative Design workflow can be seen in colour. Large quantities of designs are created and iterated upon within the software itself. These all have a problem defined via loads and geometry, as well as manufacturing bias within the set-up of the problem. This can allow for vast productivity increases, as well as creating new designs that the human mind would not be able to produce in such a time, making it an incredibly powerful tool.

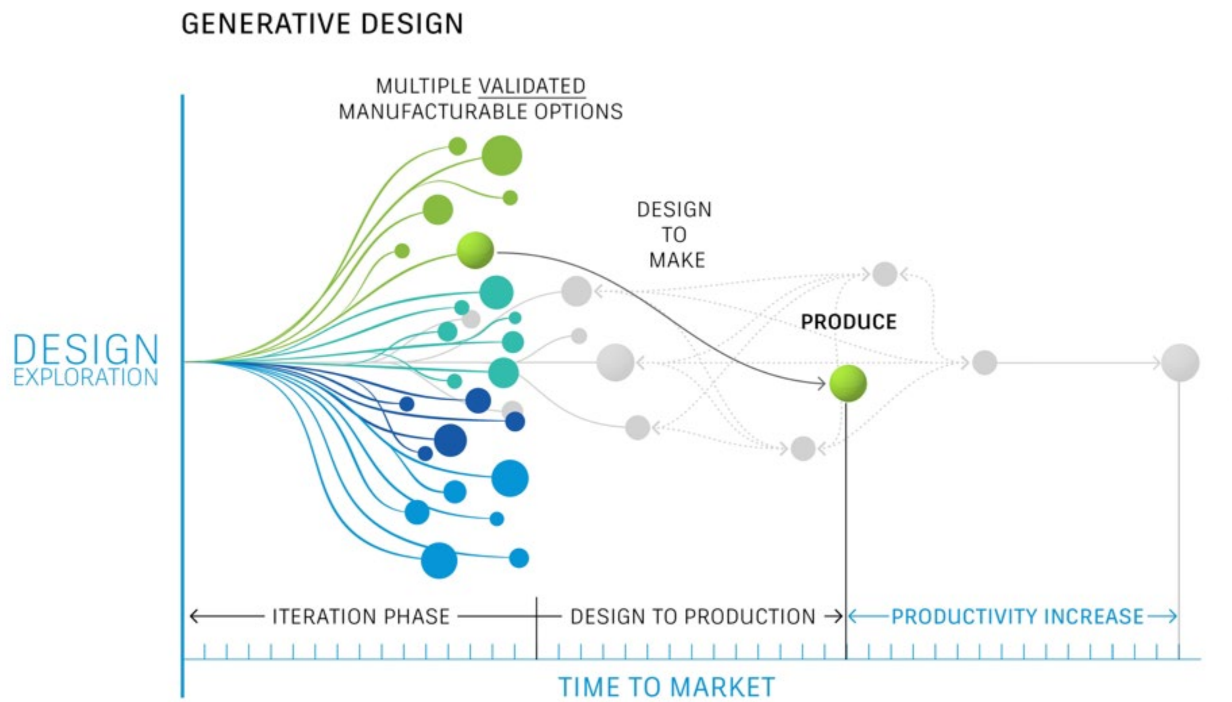

Figure 1 – Differences in Design Exploration Workflow.

Key learning objectives

- Understand Generative Design advancements allowing the revisiting of the project to be successful
- Understand the Generative Design set up
- Evaluate relevant trade-offs associated with Generative Design

Generative Design Advancements

One perceived shortcoming of Generative Design in years gone by is that it creates excessively complex geometry. Our initial work with Lightning motorcycles shows a good example of this. The software at the time was very new, yet still incredibly powerful. It allowed all these complex geometries to be created yet there was a lot of refinement needed for the workflow to succeed from the very early outcomes.

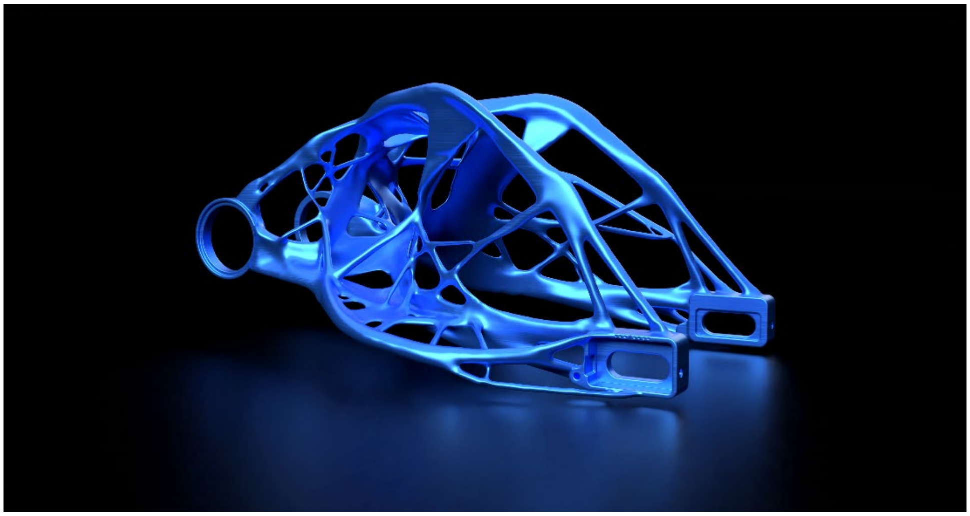

Figure 2 – Initial Lightning Motorcycles Generative Design Outcome.

Figure 2 shows this initial work with Lightning Motorcycles. You can see that the part has a high level of detail, and the original assembly was solved as one part. This was down to the limitations of the software at the time. There were far fewer solver options which meant that it was far harder to set up the real world loading of the part within the generative design software. As a result, the loading would be approximated. Nowadays, we have more complex solvers and a greater variety of load case attributes that we can fully utilise to better represent the real-world loading. Subsequently we were able to split our design into the same 3 parts as the original design, 2 end pieces and a central component acting almost as a torsional brace.

The other shortcoming of Generative Design at the time was the fact that it was strongly focussed on additive manufacturing. This was an area that was always going to be developed upon as it is key that the software allows for a wide variety of design problems to be applicable. Within Generative Design you can now include 2.5, 3 or 5 axis milling, 2 axis cutting and die casting, along with the traditional additive and unrestricted manufacturing methods. This allows the complexity of the part to be controlled. The inbuilt cost estimation, powered by aPriori, gives even more knowledge to the user about the cost-effectiveness of their designs, leading to a new wave of generative design outcomes to be possible.

Generative Design Set Up

For any problem-based solution to succeed, you need to be able to accurately represent the problem in order to make sure all outcomes are matching the real-world conditions that the part is going to be under. As you will see in the simulation section, we had a very well-defined problem due to the simulation criteria Lightning Motorcycles provided.

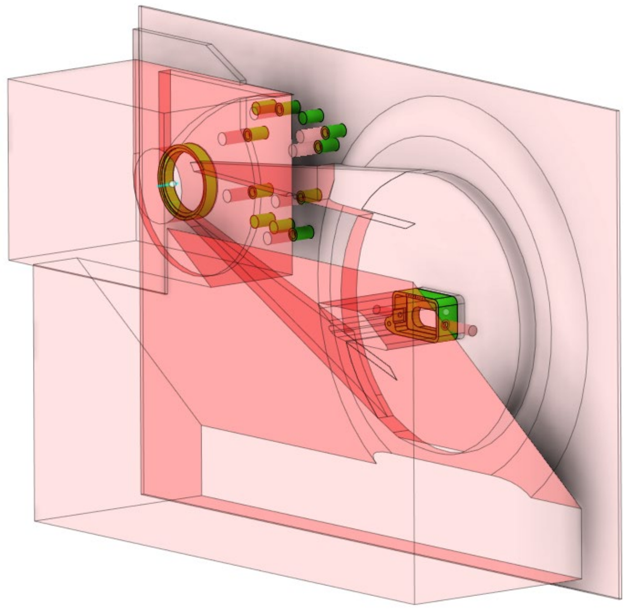

Figure 3 – Riders Left Generative Design Geometry.

For both of the side components, the set up was very similar. *Figure 4* shows the geometry used in the set up on the Riders Left hand side part (NB: riders left dictated as if rider was sat on bike facing forward). The geometry can be seen, with the red obstacle regions set to 40% opacity for ease of viewing. The obstacle regions for Generative Design represent areas that the solver will not add material in. These may be due to other parts of a wider assembly, or potential ranges of motion. It is worth noting that the left-hand side features the obstacle for the sprocket. This was key to represent well as Lightning vary the sprocket size on their bikes. As a result of this, we oversized our obstacle from our original model to fit all sprocket sizes. The other obstacles are dictated from parts of the bike assembly, such as the wheel and motor.

The preserves are shown in *Figure 3* in the characteristic green, as found within Generative Design. These are taken from the original design to maintain a similar assembly line to reduce the need to change the product line that Lightning already have in place. These are chosen by evaluating the critical fixture points, as well as the loading regions that allow the part to complete its purpose. It is possible that we could look at these in more depth, rearranging fixturing etc. but this may cause a large increase in the assembly time and may not provide much benefit.

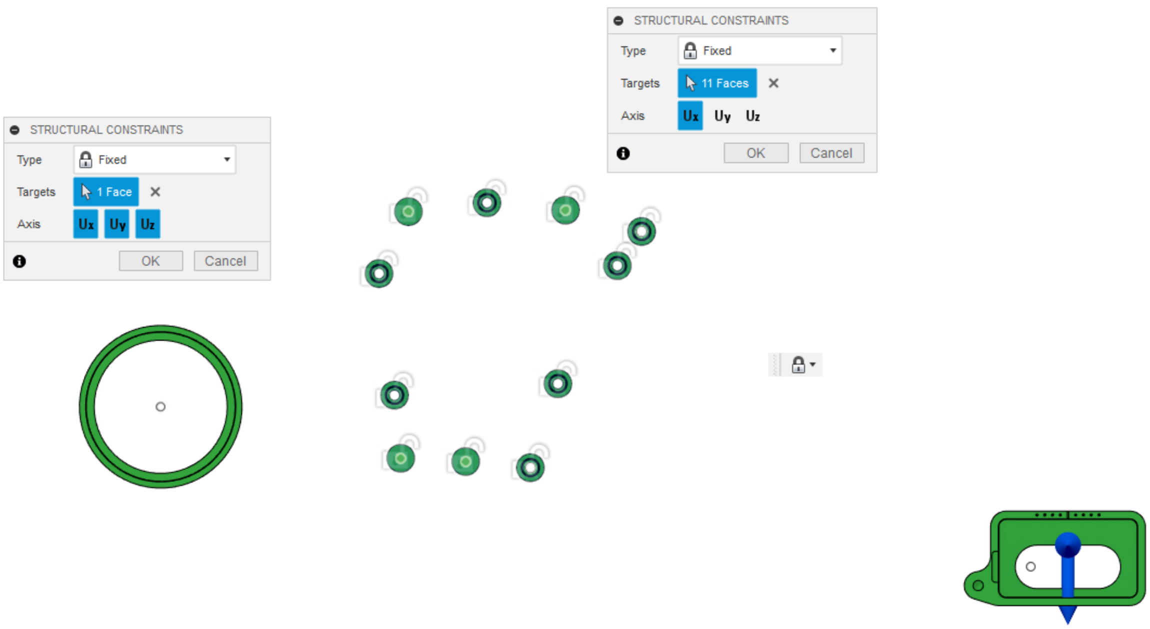

Figure 4 – Riders Left Generative Design Setup.

Figure 4 shows the actual set up of the Generative Design. To represent the part as a subset of the assembly, we fixed certain geometries that were not loaded to represent the fixturing to the other assembly parts. We then extracted the relevant loads from our simulation model and applied these to the loaded bodies. This allows each part of the assembly to be relevant to the overall assembly model whilst still maintaining the 3-part assembly overall. This was the same for both side parts of the assembly, with the only difference being slight changes in the obstacle geometry. As a result, I will not show the set-up of the rider's right-hand side.

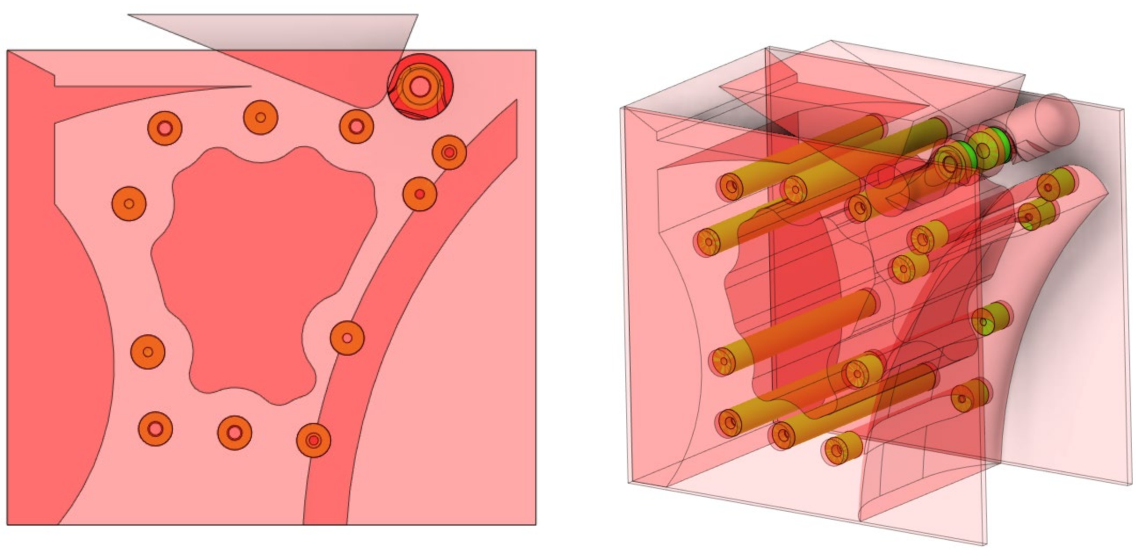

Figure 5 – Central Component Generative Design Geometry.

The geometry and set up for the central component were slightly more complex that the side parts. The geometry can be seen in *Figure 5*. This geometry was very complex due to the location of the component, and specifically its proximity to other components. As a result, there was a very well-defined build volume to avoid the tire, motor, shock absorber etc. that is represented by all of the obstacle geometry, again shown in 40% opacity red colouration.

The preserve regions were also slightly complex on this part. The part consisted of a series of through holes or capped holes that allowed the 2 side parts to be connected, via the central component. Due to our decision to keep the fixturing the same, these were all taken from the original geometry and utilised within our set up. The only non-fixture driven preserves are the lugs at which the shock absorber mounts to the swing arm assembly. This is very important to the loading of the swing arm as it provided the location of the main reaction force to the loading in the overall assembly. (NB: the approximation of the shock absorber will be discussed later)

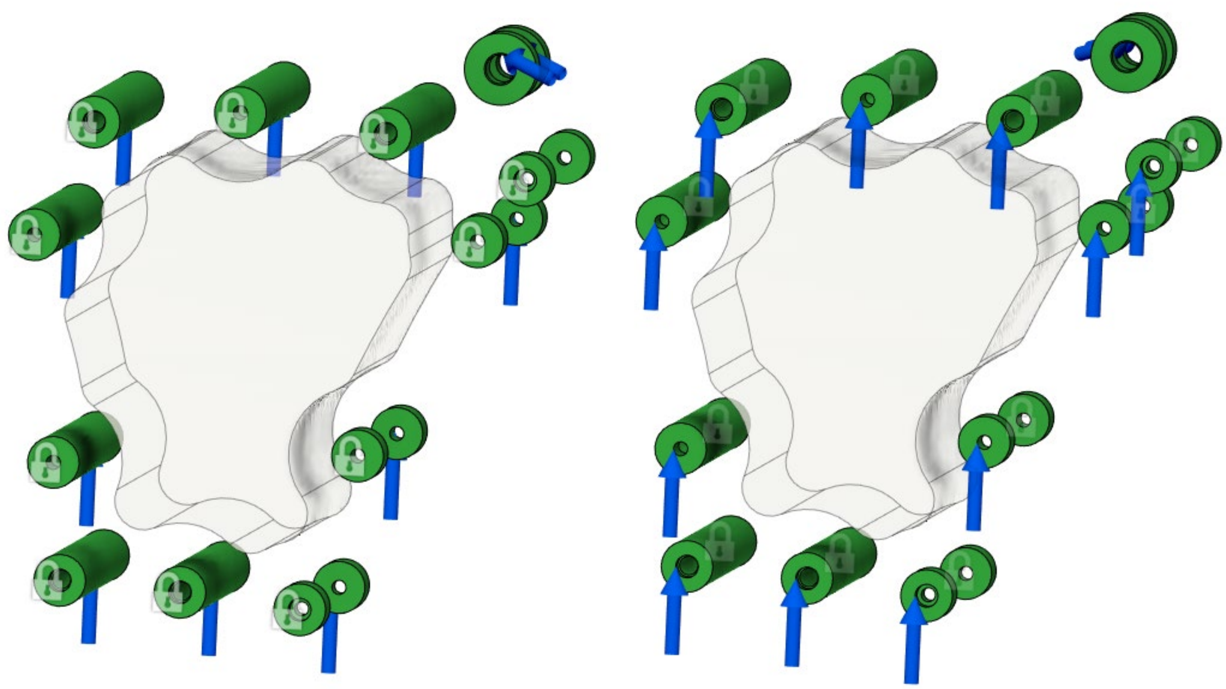

Figure 6 – Central Component Generative Design Set Up.

The loading of the central component was dictated by a combination of the forces on either side relating to the forces of the side components, and the reaction force from the shock absorber. *Figure 6* shows a crude representation of this. The setup of this central component had to be created due to a lack of loading within the assembly definition of the real-world loading. Nick worked hard to extract different forces and apply them to the remaining preserves.

Generative Design Trade Offs

As is the case with any design process, Generative Design is full of trade-offs. For the most part, these are trade-offs of performance (mass, factor of safety, displacements) against ease/cost of manufacture. Due to the nature of our previous work being mostly focussed on the mass savings possible, Lightning were very keen to see if we could harness that whilst focussing more on manufacturability and cost. *Figure 7* illustrates how this trade off was performed in a simple yet effective matrix. We compare the original design to our Generative Design outcomes from 2018 and 2021. Although there was an extra 30% mass saving in the 2018 project, we believe that the focus on manufacturability will allow this part to be built upon, and a production part is soon!

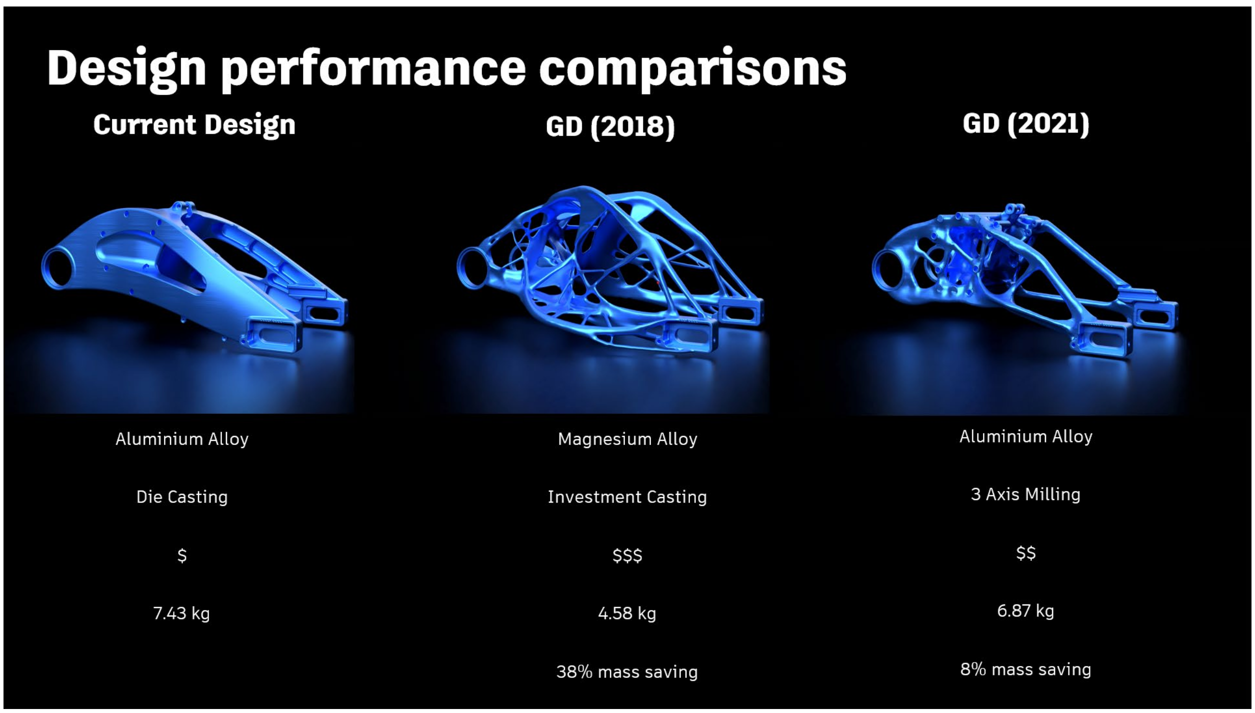

Figure 7 – Performance Comparison.

Fusion 360 simulation

Fusion 360 Simulation is a validation tool to help you understand how a design performs under certain conditions. A highly trained specialist could spend much time performing a detailed analysis to obtain the exact results of real-world conditions. However, you can often predict and improve a design based on the trending and behavioural information you obtain from fundamental analysis. If you perform this basic analysis early in the design phase, you can substantially improve the overall engineering process.

Use the analyses in the Simulation workspace to determine how loads lead to deformation and failure, so you can understand if and how a part will fail. Or you can determine natural vibration frequencies to avoid resonance. You can identify temperature distributions and thermally induced stresses.

Save time-to-manufacture, in the Simulation workspace, as you experiment with virtual design variations or adapt your model to changing design requirements. Use the tools in the Simulation workspace to minimize physical prototyping and destructive testing requirements. Fusion 360 Simulation studies that run in the cloud, rely on cloud computational services.

Key learning objectives

- Simulation selection
- Simplifying geometry
- Meshing
- Materials
- Load cases & boundary conditions
- Results exploration

Simulation selection

Static stress analyses are one of the most common types of finite element structural analyses. The component or assembly is subjected to a range of load conditions and the resultant stress, strain, and deformation results are analyzed to determine the likelihood of failure of the design.

Linear static stress analyses assume that:

- the structure returns to its original form
- there are no changes in loading direction or magnitude
- the material properties do not change
- deformation and strain are small

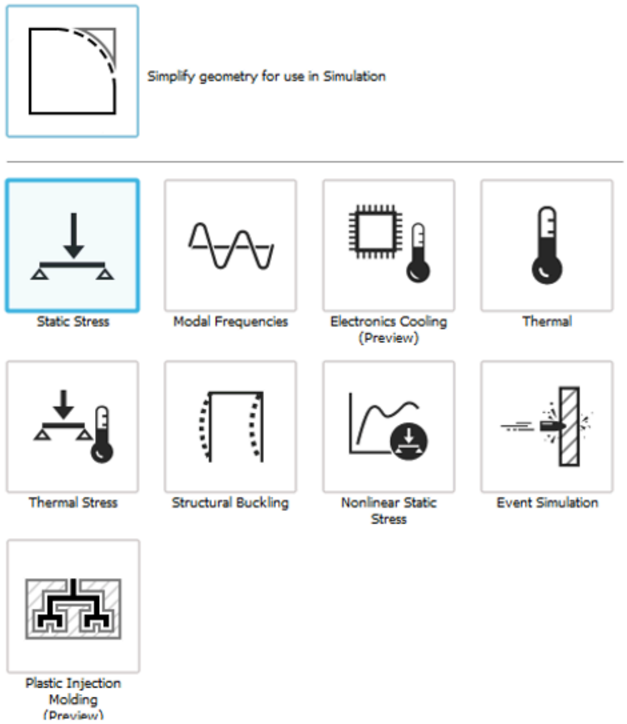

Figure 8 – Fusion 360 simulation capabilities.

Geometry

Remove unnecessary model features that complicate a simulation.

A simulation model can be less detailed than a manufacturing model. There are often features in a production design that is unimportant concerning stress, modal, or thermal stimulation. However, they can greatly increase the complexity of the mesh (producing a high element count). Therefore, the file size and solution time increase. Examples of potentially unnecessary features are embossed or raised part numbers and fillets on external corners.

Removal is supported for the following types of features:

- Fillet
- Hole
- Chamfer
- Extrude
- Revolve
- Other: Features that are small in comparison to the selected body that wasn't created with one of the above commands.

You can remove small holes, notches, or protrusions if all of the following statements are true:

- The feature is well removed from the critical stress region, so stress concentration effects are not a concern
- The feature does not have a significant effect on the overall stiffness of the structure
- The feature is not needed for a constraint or load application edge or face
- If running a modal frequencies analysis, the feature does not alter the mass enough to significantly affect the natural frequencies or mode shapes.

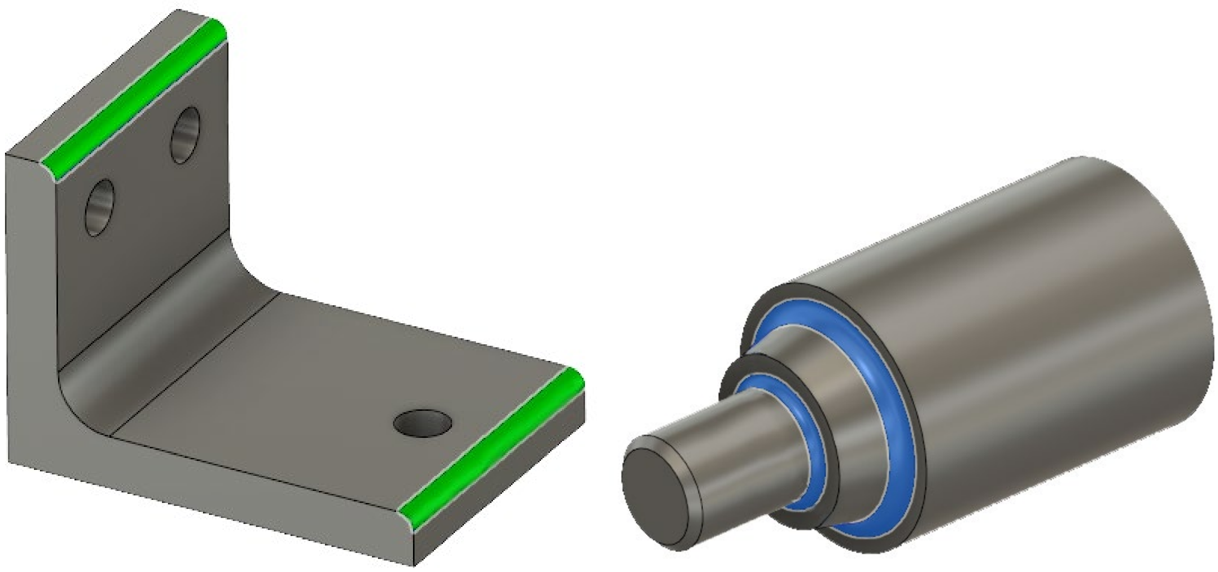

Figure 9 - Fillets and Chamfers you can suppress (green faces). Fillets you should not suppress (blue faces).

Meshing

The quality of a surface mesh and the shapes of the solid elements affect the accuracy of your simulation results. Also, the mesh density (that is, the number of elements per unit volume) affects the accuracy of the results. Pay attention to the quality of the mesh, and local mesh refinement, to maximize the accuracy of your results through mesh convergence techniques.

Autodesk Fusion 360 provides automated mesh generation along with both global and local settings for the mesh size, mesh quality, element order, and other options.

Autodesk Fusion 360 provides automated mesh generation along with both global and local settings for the mesh size, mesh quality, element order, and other options.

- Mesh settings apply to all components.
- Local mesh controls apply to selected entities.
- Each simulation study has its mesh settings.

A good mesh balances precision and computing time. Quality meshes converge quickly, produce accurate results, and do not produce errors. The majority of the meshing process involves specifying the appropriate mesh settings.
When critiquing the suitability of a mesh, consider the following characteristics:

- Two different element order settings are available in Fusion 360:
  - Linear: Linear elements have nodes at the corners only, and they have straight edges. A linear tetrahedral element has four triangular faces, six edges, and four nodes.
  - Parabolic: Parabolic elements have corner nodes and an extra node at the midpoint of each edge. A parabolic tetrahedral element has ten nodes (four corner nodes and six mid edge nodes).
- The number of elements through the thickness:
  - There is a rule of thumb that parts subjected to flexural (bending) stresses should have a minimum of four elements through the thickness of the part.
  - This rule applies to linear elements. The multiple elements are required to reasonably represent the varying and reversing tensile stress magnitude through the thickness. However, most analysts tend to use parabolic elements.
  - When using parabolic elements, two elements through the thickness produce similar results (due to the additional nodes between the corner vertices).

For thin parts, it is often difficult to achieve a mesh that is dense enough to satisfy this guideline. You can use local mesh refinement to reduce the element size only in critical regions. In this way, you can achieve the recommended number of elements without producing an extreme total element count.
Not all thin parts require multiple elements through the thickness for acceptable accuracy. If the stress is mostly tensile or compressive membrane stresses (not bending stresses) one element through the thickness may be adequate. Consider the following comparison in Figure 10.

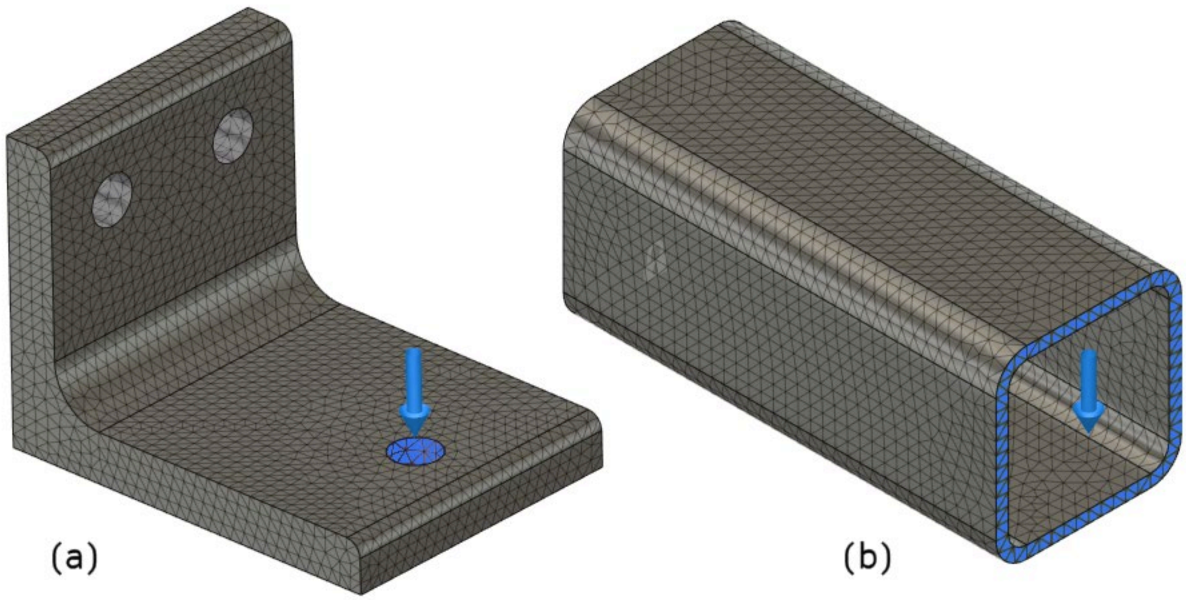

Figure 10 – (a) thick mesh body and (b) thin mesh body example

Both parts are fixed at the left end and have a downward force applied, which produces bending. In the horizontal portion of the angle, Part (a), the tensile stress reverses through the thickness of the part. The top surface is in tension, and the bottom surface is in compression. For such cases, it is recommended to have two to four parabolic elements is needed through the thickness to adequately capture the stress gradient.

In Part (b), the top and bottom walls of the square tube experience only a slight change in tensile stress magnitude through the thickness. The top wall is in tension, and the bottom wall is in compression. For the elements along the two side walls, the stress gradually changes from tensile to compressive as you move from top to bottom. In such cases, a single linear element through the thickness is sufficient. The stress gradient through the thickness is minimal, and there is no stress reversal.

Materials

The selection of materials is very important to the overall accuracy of your simulation. It is critical to correctly represent the physical properties in your simulation model. In many designs, however, the choice of materials is part of the design process. You can create multiple studies based on the same design model but assign different materials for each study.
Different studies have different material data requirements.

Mechanical properties

Displays material properties for the selected material. The following properties are displayed:

1. Density
2. Young's Modulus
3. Poisson's Ratio
4. Yield Strength
5. Ultimate Tensile Strength
6. Thermal Conductivity
7. Linear Expansion Coefficient
8. Specific Heat

Note: For static analysis, the first four bullet points are only required.

Contact

Automatic contacts are detected at the start of the analysis before meshing starts. You do not have to take direct action for automatic contacts to be generated.

Steps

To define the contact tolerance, click **Automatic Contacts** before meshing or running the simulation.

1. Click **(Simulation workspace > Setup tab > Contacts > Automatic Contacts)**.
2. Specify the desired **Contact Detection Tolerance**.
3. Click **Generate**. The model is evaluated.

Note: Automatic contacts are generated as a prerequisite to meshing whenever meshing is performed. The default contact pairs will be assigned as the bonded type. Bonding two parts together don't require the same mesh configuration.

Load cases & boundary conditions

Load cases

4. Click (Simulation workspace > Setup tab > Loads panel > Structural Loads).
5. In the Structural Loads dialogue, select the Type from the drop-down menu.
6. On your model, select the Target face(s), edge(s), or vertices. The Target entities you can select depend on the load Type you choose.

7. Note: As soon as you select your first model entity, the dialogue expands to provide more options. The options depend on the load type you select.
8. Where applicable, choose the load direction (Normal, Angle, or Vector).
9. Normal applies the load Normal to the selected face(s).
10. Angle applies the load according to the angle you select.
11. Either
12. In the Structural load's dialogue, click Select associated with Direction Reference, then on your model, select a face or edge for the reference direction.
    a. Or enter values for the X Angle, Y Angle, and Z Angle
    b. Or use the manipulators to alter the direction.
13. Vector) applies the specified load(s) along the vectors you choose.
14. Enter the required load in each vector direction.
15. Where applicable, choose the Reference, or Direction Reference.
16. Optionally, click Override Units to choose alternative units for input parameters.
17. Specify the magnitude and direction of the load or its global X, Y, and Z components.
18. Note: For Hydrostatic Pressures, specify the Fluid Type or Density and ensure that gravity is enabled and correctly oriented.
19. Optionally, click Flip Direction to reverse the load direction.
20. Click OK to apply the load and close the Structural Load dialogue.

Applying load example

The model is constrained in the traditional manner (statically stable). The small pin is fixed at its end faces, and the load is applied to the end faces of the large pin.

1. Apply a total force of 2,000 pounds in the -X direction to the end faces of the large pin (1,000 pounds per end).
2. Click (Simulation workspace > Setup tab > Loads panel > Structural Loads) to open the Structural Load dialogue.
3. In the Structural Load dialogue, confirm that Type is set to Force.
4. Click the top face of the large pin to select it as a Target, then rotate the model so you can see the underside of the pin and select the bottom face too.
5. In the Structural Lood dialogue, set the Direction Type to Vectors (x, y, z)
6. Type -2000 in the Fx input field.

Note: Since the Force per Entity option is not selected, the -2, pound-force once is divided among the selected faces. The areas of the two faces are equal, so each receives half of the total load (1,00lb force). Refer to Figure 11.

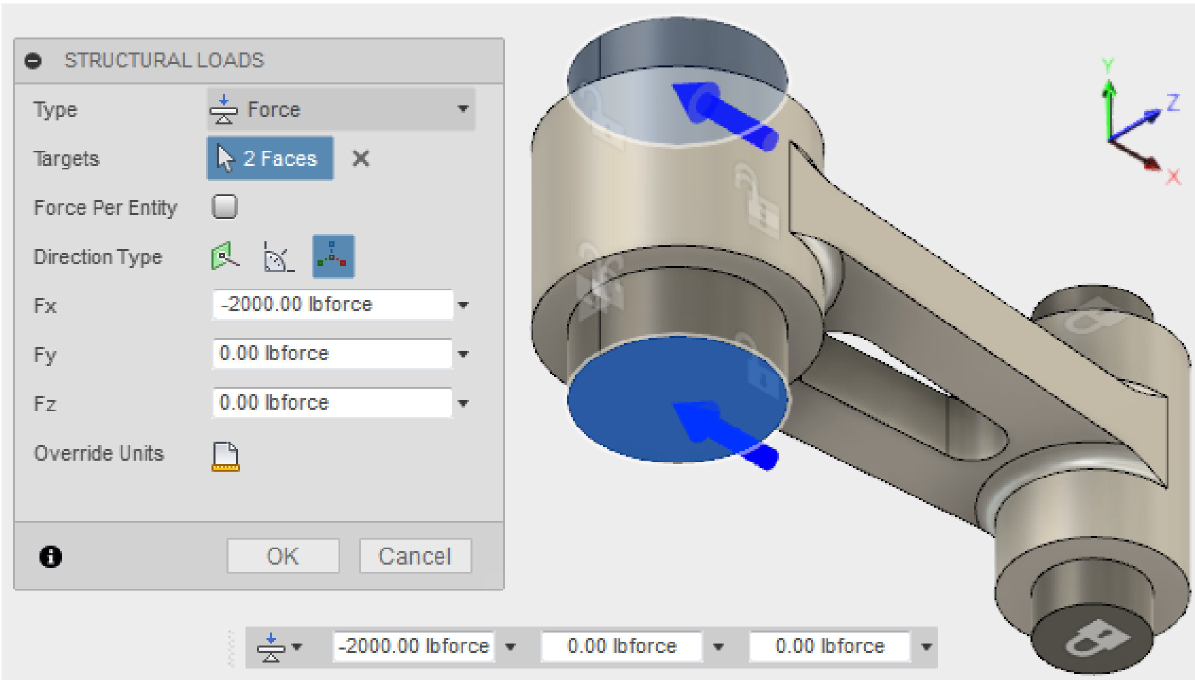

Figure 11 – applying force in the Fusion simulation environment.

Boundary condition

For static simulations, you need to prevent all rigid body motion, such as free translational and rotational movement. This process is known as making the model statically stable. To do that, you either apply a fixed constraint to a face, or combine partial constraints on faces, edges, or vertices.

Structural constraints include:

Constraint	Used to
Fixed	Prevent movement in selected directions. By default, all three global directions are constrained.
Frictionless	Prevent movement in the direction normal to the surface.
Pin	Prevent movement in radial, axial, or tangential directions. This constraint is only applicable to cylindrical surfaces (full cylinder or segment).
Prescribed Displacement	Prevent movement in selected directions, similar to a fixed constraint, except that the entities are held at a displaced position. You can specify displacement components separately in the three global directions.
Remote constraint	Prevent movement in selected directions, similar to a fixed constraint, except the constraint is located at a remote location. When all 6 degrees of freedom are fixed, a remote constraint is the same as a fixed constraint.

Figure 12 – constraint options in the Fusion static simulation

Applying fixture example

- Fully fix the end faces of the small pin.
    - Click  (Simulation workspace > Setup tab > Constraints panel > Structural Constraints) to open the Structural Constraints dialog.
    - In the Structural Constraints dialogue, confirm that Type is Fixed and Ux, Uy, and Uz Axis are all selected, in order to constrain all three directions.
    - Click the top face of the small pin to select it as a Target.
      Notice that the face turns blue to show it has been selected, and a white lock appears to show it has been constrained.
    - Rotate the model so you can see the underside of the small pin, then click on the bottom face to select it.

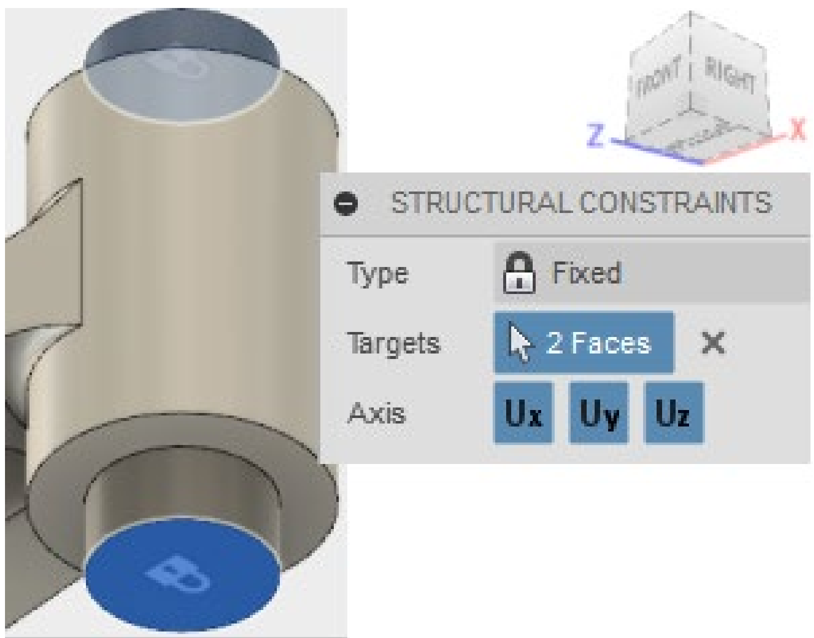

Figure 13 – applying fixture in the Fusion simulation environment

Results exploration

The output of a finite element analysis (FEA) solver is generally a substantial quantity of raw data. This data is difficult and tedious to interpret without graphical representations. Contour plots are **graphical** displays that represent the distribution of stresses, deformation, temperature, and other results, using various colours. A key aspect of a simulation analysis is the proper interpretation and evaluation of these results.

You perform a side-by-side comparison of the results of two or more studies. The Compare tab enables you to view the results of up to four simulation studies, at the same time, in tiled windows.

Displacement

The **Displacement** result is a colour contour plot that shows how much the model has moved relative to its original position. You can choose to view the total displacement of the components of the total displacement (in each of the three global directions):

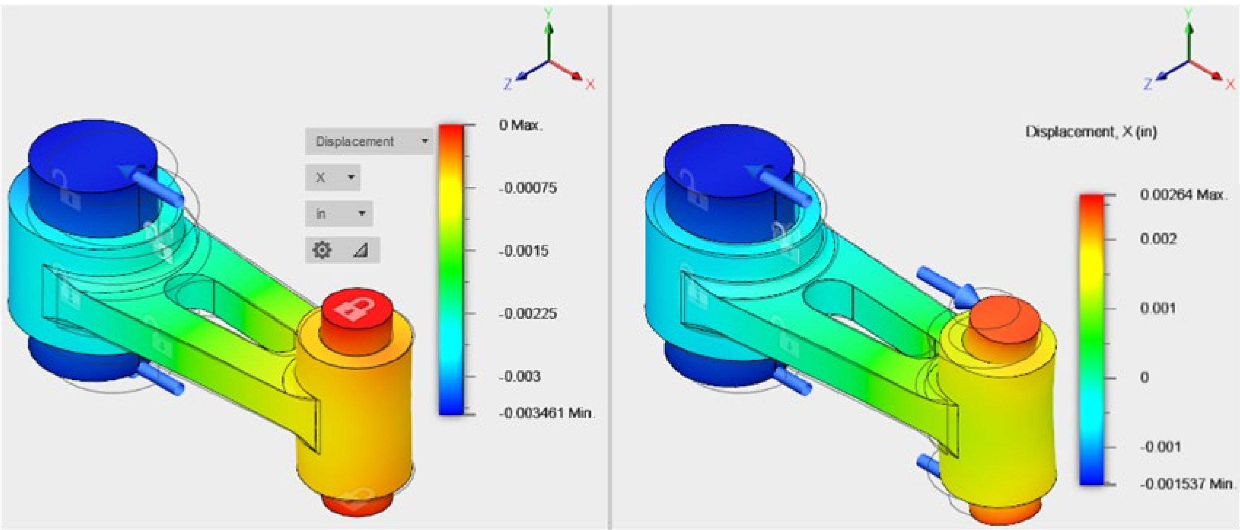

Figure 14 –Displacement X results between two studies

Stress

Stress is defined as the force acting per unit area. It is computed from strain (elongation or compression per unit length) and the material stiffness. The solver outputs six individual components (stress tensors) and three combined stress results.

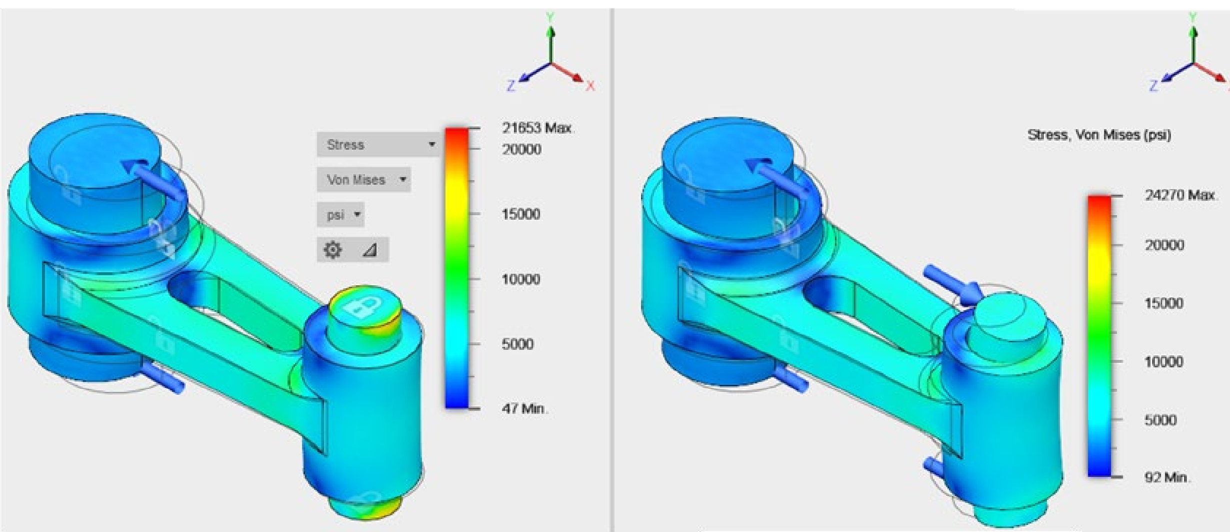

Figure 15 – von Mises stress results between two studies

Factor of safety

All objects have stress limits that are dependent upon their construction material. **Safety Factor** is an important result that you can use to evaluate how suitable a design is for its intended application. The safety factor indicates if a design is likely to survive unharmed, bend, or break when subjected to the applied loads. There are various criteria and considerations influencing what the safety factor should be for a given material, manufacturing process, and application.

Frequently, designers strive for a minimum safety factor between 1.5 and 6, depending on the application. Design safety factor compliance is based on the greatest expected loading scenario. However, lesser or greater safety factors might be targeted in certain cases.

Design safety factors typically exceed 1.0 by a significant margin. A safety factor less than 1 indicates that some sort of failure will occur (either permanent deformation or breakage). A safety factor of exactly 1.0 means that the actual stress equals the material strength limit, so the design is on the verge of failure.

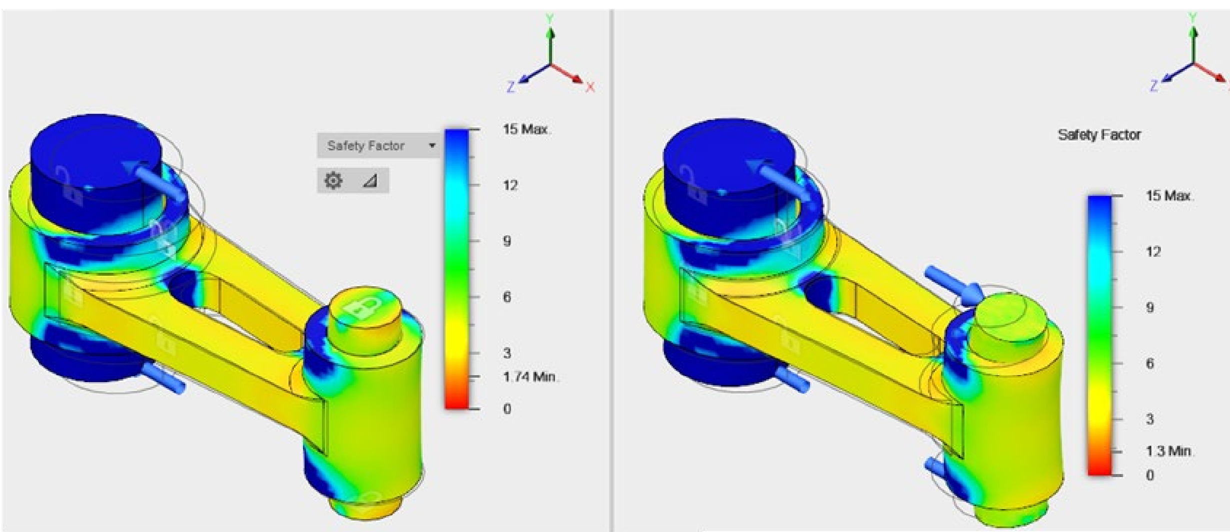

Figure 16 – Factor of Safety results between two studies

Observation: The safety factor results reflect the increased stress level in Study 2, where the safety factor is 25% lower than that of Study 1.