Oren Minsk, Josh Dalal, Eitan Seitchik

Professor Hu

MEAM 6900

April 21, 2026

Final Project Report

Introduction

Rocket engines function by ejecting mass through a nozzle which creates an equal and opposite reaction force to accelerate a vehicle. In order to maximize thrust and efficiency of an engine, its exhaust should be accelerated to as high of a velocity as possible as both parameters are directly proportional to exhaust velocity. With typical exhaust velocities of modern orbital class rocket engines nearing and exceeding 3 km/s, the incompressible flow assumption often made in many undergraduate fluids courses is no longer applicable. Supersonic flow requires an understanding of different types of shockwaves and compressibility effects which were discussed in this course. Rockets leverage chemical energy from propellant that is converted into kinetic energy in order to achieve high exhaust velocities. Without chemical energy, it is not known if humans would have ever sent satellites into space or people to the moon. While undergraduate fluids courses typically study homogenous flow of a single fluid that perhaps changes phase, this course covered combustion of multiple fluid propellants which are necessary for orbital class engines (though monopropellant engines do exist for less energy-intensive applications).

While some idealizations of compressible flow can be solved analytically (as demonstrated in the midterm project assignment), it requires many simplifications and assumptions. Analytical solutions can provide a meaningful estimate and provide validation for more complex models, but they cannot handle concepts like turbulence or wall friction very well,

which are important to model in rocket engines for accurate results before testing. Calculating how combustion and the release of chemical energy by hand or with analytical methods is also very difficult. Therefore, computational fluid dynamics (CFD) software is often used to model the combustion and compressible flow inside a rocket engine, and understanding how to model and simulate combustion and compressible flow in CFD is the motivation for this project.

This report is documentation of a project involving modeling of a simple rocket engine, meshing the geometry, simulating the flow and combustion in ANSYS Fluent, and discussing the results. It covers the modeling decisions that were made and how they were implemented in Fluent. This project independently varies back pressure, O/F mixture ratio, and the ratio of axial to tangential swirl inlet velocity of the propellants and then discusses the results.

Simulation and Fluent Implementation of Model of Rocket Motor

Geometry

The first step when modeling a rocket motor in ANSYS Workbench is to define the geometry. The geometry can be created in any CAD program and imported into Workbench, though ANSYS SpaceClaim was used for this project. The geometry consists of an inlet wall, a combustion chamber, and a converging-diverging nozzle. The inlet wall has designated areas for both fuel, oxidizer, and oxidizer film cooling towards the edge of the combustion chamber, acting as injectors. The sizes of these injector elements were taken from the Lab 3 provided geometry and were not a source of focus of this project (though entire projects and research papers could be done solely focussing on injector element sizing and spacing). Similarly, the combustion chamber length was also taken from the Lab 3 provided geometry. This could likely be improved upon because, as is clear later on from the simulation results, the fuel and oxidizer are not completely combustion by the time the nozzle begins to converge. While combustion

chamber length was not a parameter in this project, were it to be increased, more complete combustion and higher temperatures would be expected.

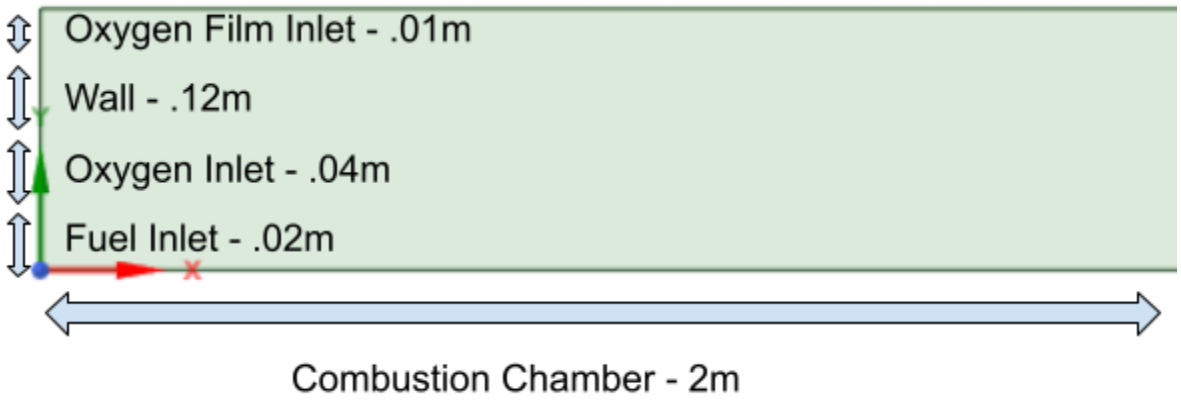

Image 1: Screenshot of the combustion chamber with dimensions overlaid. Arrows are not to scale. The small wall between the fuel inlet and oxygen inlet was not included in the visual.

The converging diverging nozzle is modeled as a parabolic nozzle with the equation $y = (200/1.64) * [(x/625)^2 + 1]$ and has an expansion ratio of 2.69. Converging-diverging nozzles can be modeled as conical, parabolic, or using the Method of Characteristics. A parabolic nozzle is a good choice because they represent a good balance of performance and complexity. Parabolic nozzles are more difficult to manufacture than conical nozzles, but they can easily be modeled in CAD using a simple equation. While the Method of Characteristics provides an ideal nozzle contour, a parabolic nozzle can achieve nearly the same efficiency and is much easier to compute. For a given expansion ratio, a parabolic nozzle is also much shorter than both a conical nozzle and a contour computed using the Method of Characteristics, making it lighter.

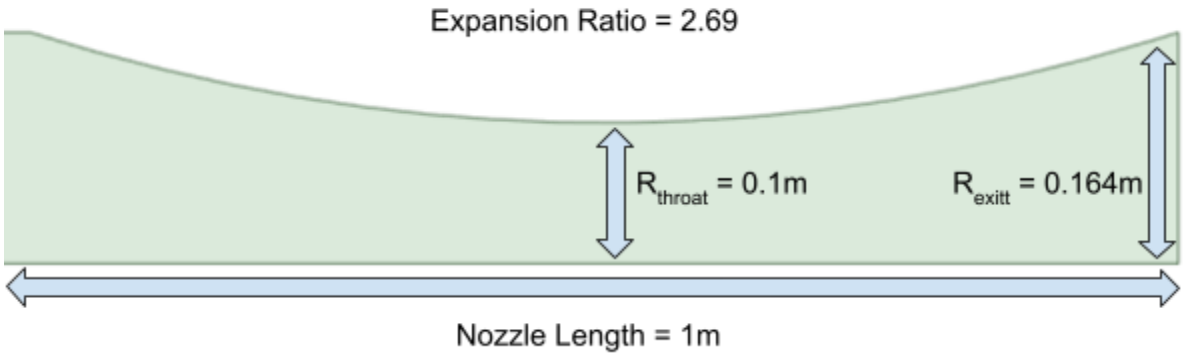

Image 2: Screenshot of the converging-diverging with dimensions overlaid.

The next step in the Workbench workflow is meshing the geometry, which will be discussed in a later section of the report. In Fluent, projects can be modeled as 2D, 2D Axisymmetric, 2D Axisymmetric Swirl, or 3D. 2D geometries assume an infinite depth, and 2D axisymmetric geometries replace (x,y) coordinate with (z,r) coordinates, assuming flow is identical at every angle around the central axis. This project uses 2D Axisymmetric Swirl because it requires less computation and is therefore faster than modeling a 3D geometry (and is also simpler to design in SpaceClaim), but it still allows propellant to swirl around the central axis which will be a parameter in this design project.

Model for Turbulence

The turbulence model implemented in this project is the standard kappa-epsilon viscosity with Fluent's default model constants. This model assumes isotropic turbulent viscosity. The kappa term represents the kinetic energy contained within turbulent eddies, and the epsilon term represents the rate at which viscous forces convert this kinetic energy into thermal energy. The flow through the combustion chamber has a Reynold's number on the order of 50,000 (obtained from the combustion chamber length, diameter, mass flow rates of propellant, average propellant density, and the methane-oxygen mixture viscosity), and the Reynold's number near the injector

is likely higher. This is a good model choice because it works best at high Reynold's numbers and is computationally efficient, requiring fewer iterations and less CPU time and memory to converge than other models.

Model for Combustion

The combustion model implemented in this project was the species transport model with volumetric reactions and eddy dissipation. The species transport model solves conservation equations for each chemical species in the flow, and the volumetric reactions setting enables chemical reactions to occur throughout the fluid domain rather than only at surfaces. The eddy dissipation model assumes that the reaction rate is controlled by the turbulent mixing rate rather than by the chemical kinetics, meaning that combustion occurs as fast as the turbulence can mix the fuel and oxidizer together. This is a reasonable assumption for this project because the combustion in a rocket engine occurs at very high temperatures where the chemical reaction rates are fast relative to the mixing timescale, making the overall reaction rate mixing-limited. The eddy dissipation model is also computationally efficient compared to finite-rate chemistry models, which would require solving stiff chemical kinetic equations at every cell and would significantly increase the simulation time.

Fuel Type

This project uses pure gaseous oxygen as an oxidizer and gaseous methane for fuel. Oxygen is the most common oxidizer for orbital class rockets, but the choice of methane for fuel is more interesting. Until 2023, no methalox engine had ever reached orbit, but it is becoming increasingly common in commercial aerospace for rocket engines to use methane as fuel (eg. SpaceX's Raptor, Blue Origin's BE-4, etc.) for a number of reasons. The first reason is that methane can be produced on Mars from its atmosphere using the Sabatier process and will be

necessary for any manned missions to the planet. Methane falls in between RP-1 and hydrogen, the two most common rocket fuels, when it comes to density of O/F ratio by mass. This results in methane tanks being nearly as small as RP-1 tanks and significantly smaller than hydrogen tanks for the same mass of oxygen, which is ideal for minimizing structural mass. It also falls between RP-1 and hydrogen when it comes to I_{sp}, as well as other metrics like boiling point, burning temperature, and others, making it a reasonable choice for considering all these tradeoffs.

Under the Species tab in Fluent, methane-air mixture was selected. Under materials, the properties of the methane-air mixture could be edited. For this project, the ideal gas method was selected for determining density because incompressible flow could not be assumed due to the high speeds and compressibility effects associated with a rocket engine. In order to model the oxidizer as pure oxygen and not air, the O_2 mole fraction of the air was manually input as 1 for both oxidizer inlets (the main and the film inlets). The choice of using pure oxygen instead of air made sense because air has lots of nitrogen that would not combust or release chemical energy and would lower the overall temperature of the combustion chamber.

Boundary Conditions

The next step in the Workbench workflow is to define boundary conditions. This project utilizes 5 different types of boundary conditions. The first is the axis boundary condition, used to define the central axis about which the solver will rotate the 2D geometry created in SpaceClaim. The next boundary condition imposed is the wall boundary condition. The wall boundary condition is imposed on the nozzle wall, the wall of the combustion chamber, and the part of the injector face from which the propellant doesn't flow in. This boundary condition makes the walls adiabatic, so no heat is transferred to or from the flow, and nonreactive with the fluid.

The nozzle outlet was set to a pressure outlet boundary condition. The value of this back pressure was nominally set to 0 kPa absolute as other parameters were varied, though there were a set of simulations where we varied the back pressure to explore different flow regimes. This 0 kPa absolute back pressure is representative of what the flow through this rocket engine would look like in the vacuum of space where there is no ambient pressure.

Mass flow rate boundary conditions were used for the propellant inlets. Mass flow rate boundary conditions were used instead of inlet velocity boundary conditions because it was easier to change the mixture ratios by just changing the mass flow rates of the oxygen and methane respectively instead of needing to think about their relative injector areas as well. Mass flow rate is also more commonly used as a metric in industry compared to inlet velocity. Nominal values of mass flow rate were 0.22 kg/s for oxygen and 0.055 kg/s of methane to achieve a mixture ratio of 4, which is very close to stoichiometric for this propellant combination. In a few simulations, these mass flow rates were varied to achieve a range of mixture ratios. A velocity inlet condition was used for the oxygen film inlet for simplicity.

Because a 2D Axisymmetric Swirl geometry was used, the mass flow rate boundary conditions allow the user to specify how much of the mass flow rate vector is directed in the radial and tangential directions. This project used a nominal swirl value of 0.1 or 10% when varying other parameters, though swirl % was also a parameter that was varied when holding other parameters constant.

Hybrid initialization was used to initialize the flow field before running the simulations. This method solves a simplified Laplace equation to interpolate between the multiple inlet boundary conditions, which provides a more physically realistic starting flow field than initializing from a single boundary. This is important for this project because the geometry has

multiple inlets with different species and flow rates, and a poor initial guess could cause the reactive flow solver to diverge.

Mesh Design and Convergence Test

In order to verify that the results of the simulations were not dependent on the mesh resolution, a mesh convergence study was conducted. Five meshes of increasing resolution were generated. The coarsest mesh contained approximately 5,000 elements with no inflation layers or local sizing applied. The second mesh added 2 inflation layers and local sizing on the inlet and outlet boundaries, bringing the element count to roughly 6,300. The third mesh further refined the boundary layer resolution with 5 inflation layers and finer edge sizing for a total of approximately 9,300 elements. The fourth mesh increased the number of inflation layers to 10 with additional edge refinement, reaching about 15,200 elements, and the fifth and finest mesh used the same 10 inflation layers with the finest edge sizing at approximately 48,200 elements. To assess convergence, the centerline Mach number profiles were compared across all five meshes, as shown in figure 1. The profiles for Meshes 3, 4, and 5 were nearly identical, while Meshes 1 and 2 showed noticeable deviations, particularly in the injector region and near the nozzle throat. Based on this, Mesh 3 with 9,300 elements was selected for all subsequent simulations because it produced results that matched those of the finer meshes without the additional computational cost. This is a reasonable approach because the goal of a mesh convergence study is to find the coarsest mesh that still produces grid-independent results, as using a finer mesh than necessary wastes computational resources without meaningfully improving accuracy.

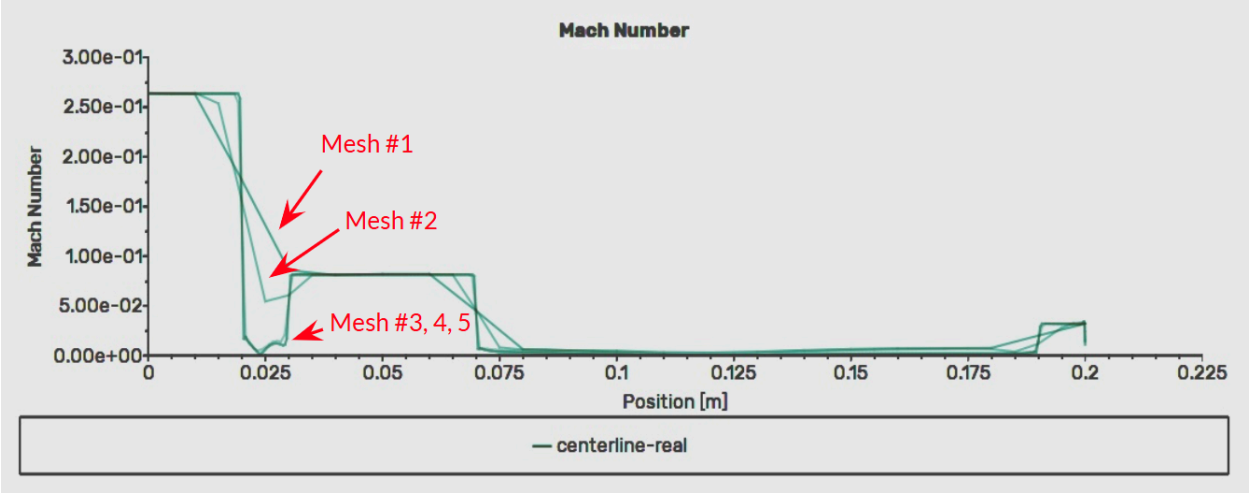

Figure 1: Plot showing mach number for 5 different meshes, showing convergence for meshes 3, 4, & 5

Experimental Design

This project explored the impact of varying 3 parameters independently while holding all else constant: back pressure, swirl %, and mixture ratio. As explored in the midterm project assignment, the back pressure in a converging-diverging nozzle is critical in determining the regime of the flow, and from varying the back pressure from 0 kpa to 10 kpa absolute, it was expected that subsonic, choked flow with a normal shock inside the nozzle, and underexpanded flow would be observed. In a second batch of simulations, swirl % was varied from 10% to 70%. It was expected that increasing the tangential velocity of the injected propellant would increase the amount of mixing and facilitate more combustion, resulting in higher temperatures and higher maximum mass fractions of combustion products in the combustion chamber. In a third batch of simulations, the O/F mixture ratio was varied from 2 to 6 by changing the magnitudes of the mass flow rate inlet conditions. Similar to the second batch of simulations where swirl % was varied, it was expected that the closer the mixture ratio was to stoichiometric, the hotter the maximum temperature and the higher the maximum mass fraction of combustion products. This

is because the closer the reaction is to stoichiometric, the more energy is released and the less unreacted propellant there is.

Calculation Activities

Each simulation was run for 2500 iterations, which was sufficient for convergence as verified by comparing the total mass flow rate at the inlets to the total mass flow rate at the outlet. Each simulation took approximately 4 minutes to solve. Because the parameters being varied in this project (back pressure, swirl percentage, and mixture ratio) are all specified as boundary conditions in Fluent, no changes to the geometry or mesh were required between simulations. This meant that the same Workbench file could be reused for all simulations, and only the relevant boundary condition values needed to be updated before re-running the solver.

Simulation Results and Discussion

Varying Back Pressure

The first set of simulations run in this project varied the nozzle back pressure while holding everything else constant. Back pressure was varied from 0 kPa to 10 kPa absolute. The results of this set of simulations was similar to what was observed in the midterm project assignment.

When the back pressure was 0 kPa absolute, the flow accelerated from subsonic in the combustion chamber to sonic at the throat to supersonic at the nozzle exit. Because the combustion chamber length is 2m and the nozzle length is 1m, the throat is located a distance 2.5m from the injector face, and as seen in Figure 2, this is where the Mach number equals 1, which is expected for choked flow. Figure 2 shows the Mach number contour for the entire geometry. The geometry beyond the nozzle exit is not modeled in this project, though if it was, it would show that the flow continues to expand beyond the nozzle exit because there is no ambient

pressure to compress it. This means that some of the exhaust travels not only axially but also radially outwards once it exits the nozzle, resulting in thrust losses. This is an inherent inefficiency of rocket engines in the vacuum of space and why expansion ratios for vacuum optimized nozzles are as large as possible.

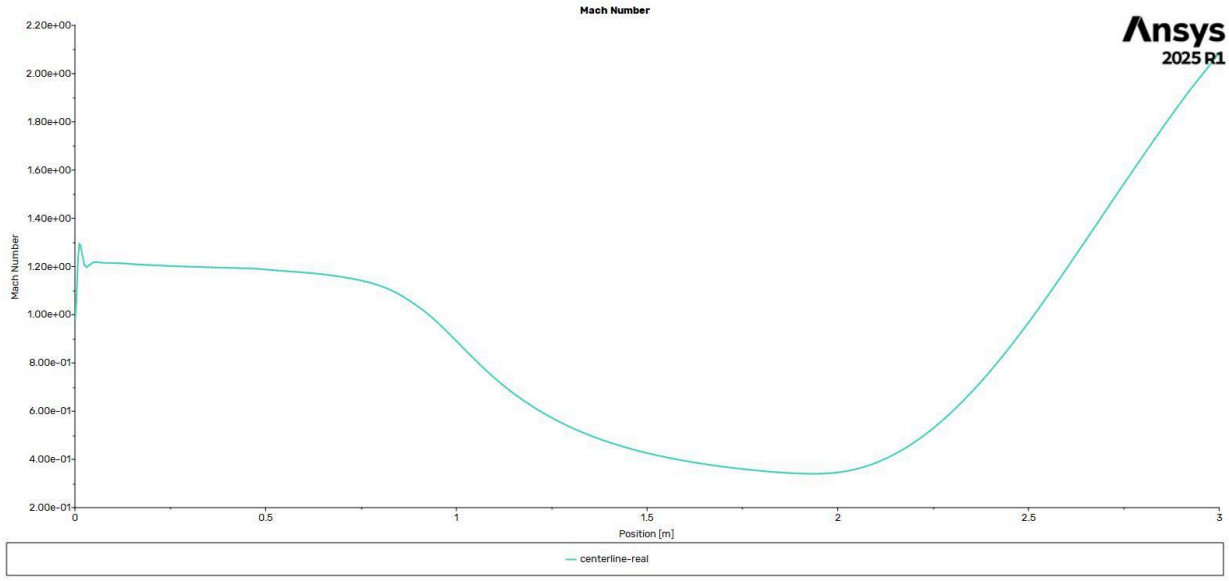

Figure 2: Mach number along the central axis of the combustion chamber and nozzle when back pressure is set to 0 kPa absolute.

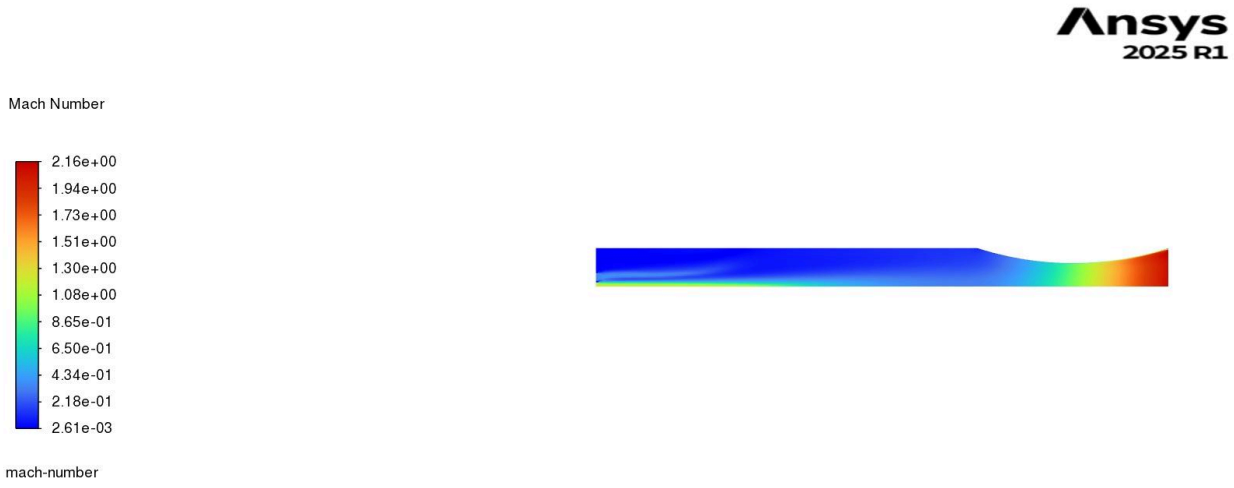

Figure 3: Mach number contour when back pressure is set to 0 kPa absolute.

When the back pressure was set to 7 kpa absolute, the flow was still choked, so it accelerated from subsonic in the combustion chamber to sonic at the nozzle and then to supersonic speeds as it entered the diverging section of the nozzle. However, because of the nozzle's expansion ratio, the flow would ideally exit the nozzle at a pressure meaningfully lower than the exit pressure. In order to equalize the flow's pressure with the ambient back pressure, a normal shock must appear. Across the shock, the static pressure of the flow spikes and the Mach number drops suddenly to subsonic velocities, which it maintains until the flow exits the nozzle. These can be seen by looking at the Mach number along the central axis in Figure 4 and the Mach number contour for the entire geometry in Figure 5.

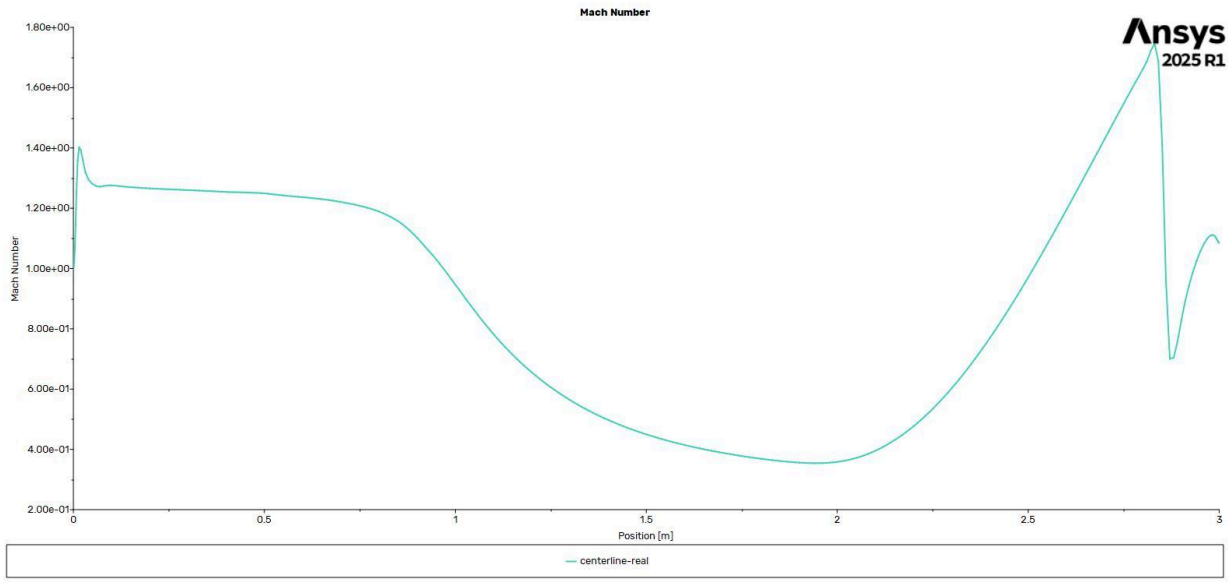

Figure 4: Mach number along the central axis of the combustion chamber and nozzle when back pressure is set to 7 kPa absolute.

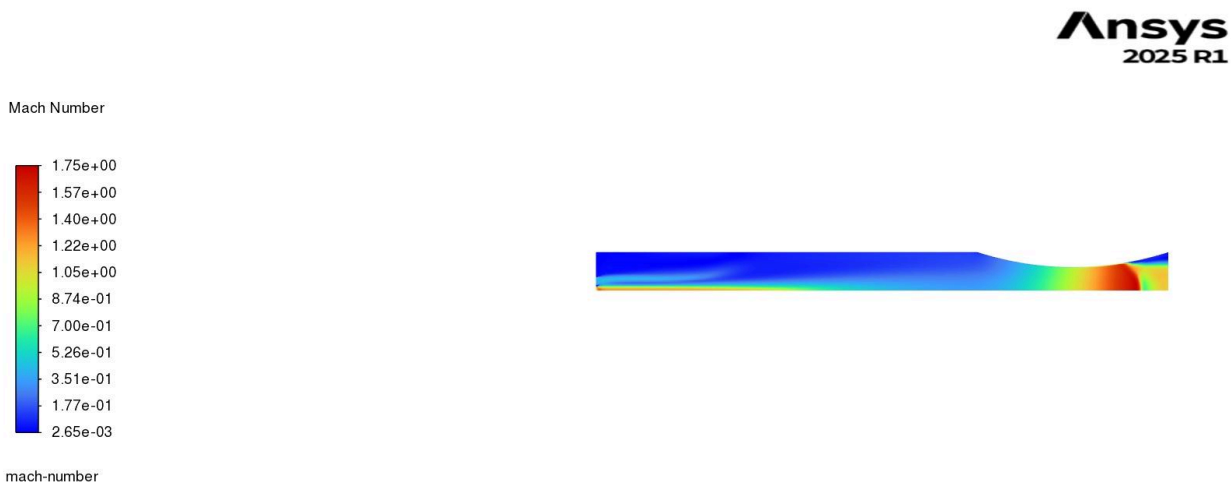

Figure 5: Mach number contour when back pressure is set to 7 kPa absolute.

When the back pressure was set to 10 kPa absolute, the pressure ratio was high enough such that the flow was not choked. The flow still accelerated from the combustion chamber as it approached the throat, but it did not achieve a sonic velocity at the throat, so it decelerated as it continued through the diverging section of the nozzle because the cross-sectional area increased. The subsonic flow along the central axis of the engine can be seen in Figure 6, and the Mach number contour for the entire geometry can be seen in Figure 7.

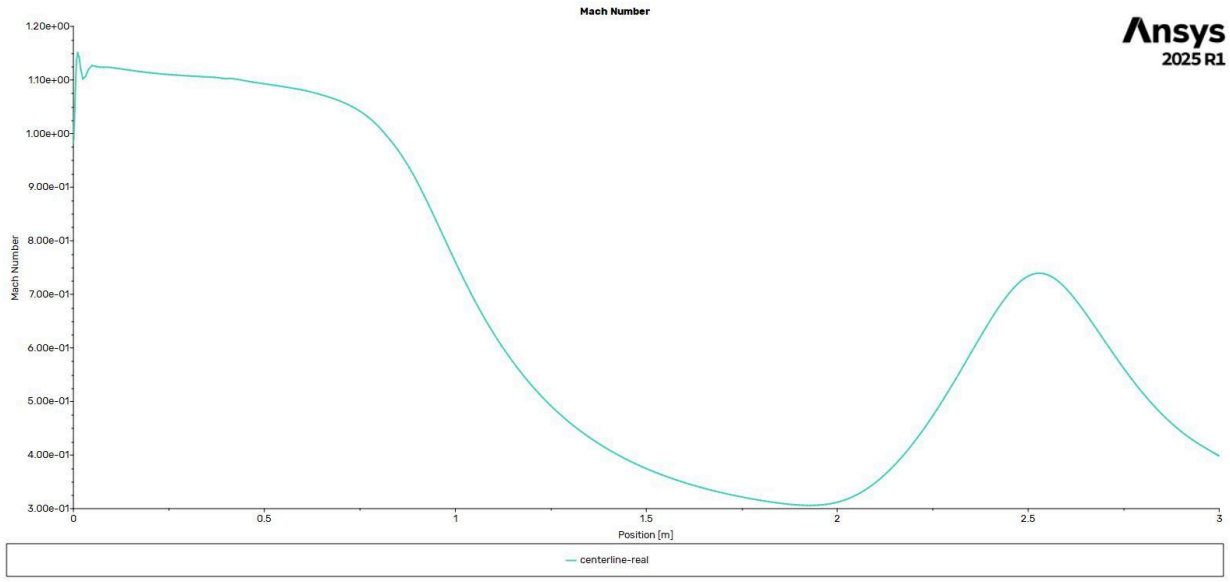

Figure 6: Mach number along the central axis of the combustion chamber and nozzle when back pressure is set to 10 kPa absolute.

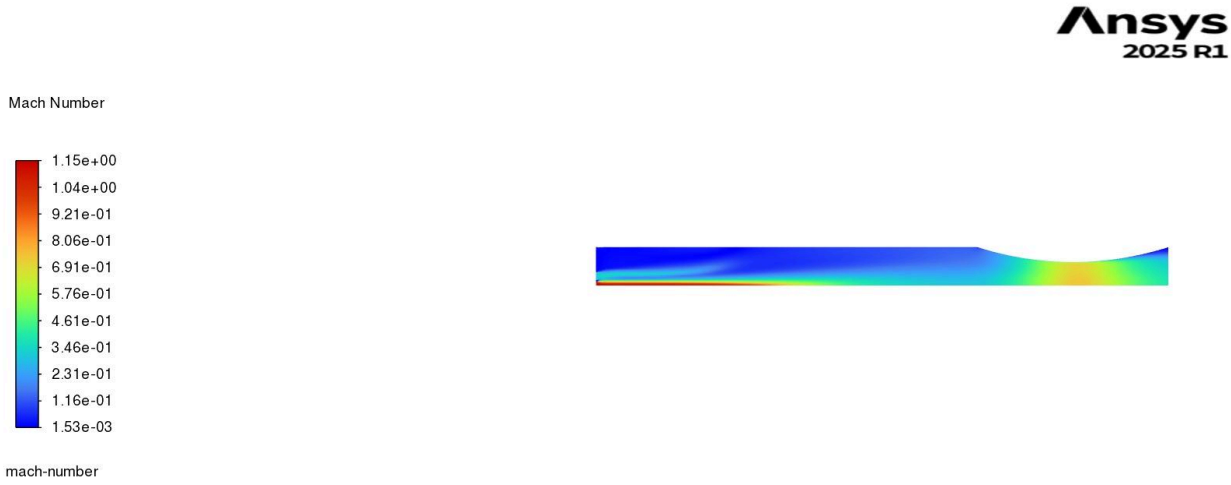

Figure 7: Mach number contour when back pressure is set to 10 kPa absolute.

It should be noted that the methane inlet velocity is seen as supersonic in these plots and contours, which is not realistic of actual rocket engines. This is only caused by the relative magnitude of the mass flow to the injector inlet area. Regardless of this, the bulk flow velocity in

the combustion chamber is subsonic and the flow is subsonic when it enters the nozzle at an axial distance of 2m, so it does not have a significant effect on the rest of results. While not shown in this report, the contours of various fluid properties (temperature, static pressure, density, etc.) look the same for these flow regimes as the back pressure was varied as the contours of these properties from the midterm project assignment.

Varying Swirl Percentage

The second set of simulations run in this project varied the swirl percentage of the propellant flow, or the size of tangential velocity component relative to the axial velocity component of the propellant flow through the inlet boundary condition. Swirl percentage values of 10%, 40%, and 70% were tested, and they were always equal for both propellants (10% swirl for both oxygen and methane, then 40% swirl for both oxygen and methane, etc.). The expectation is that as the amount of swirl in the flow increases, there would be more interaction, mixing, and therefore more combustion of the propellants, resulting in higher maximum temperatures and higher maximum mass fractions of combustion products. Figures 7, 8, and 9 show the temperature contours of the flow when the swirl is 10%, 40%, and 70% respectively. Results are discussed below.

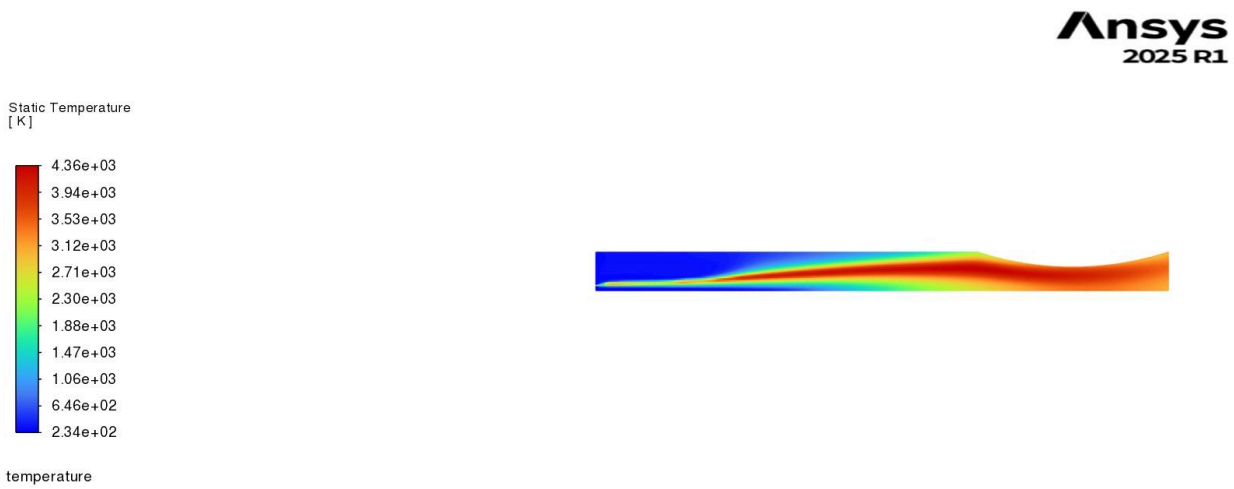

Figure 8: Temperature contour for 10% swirl.

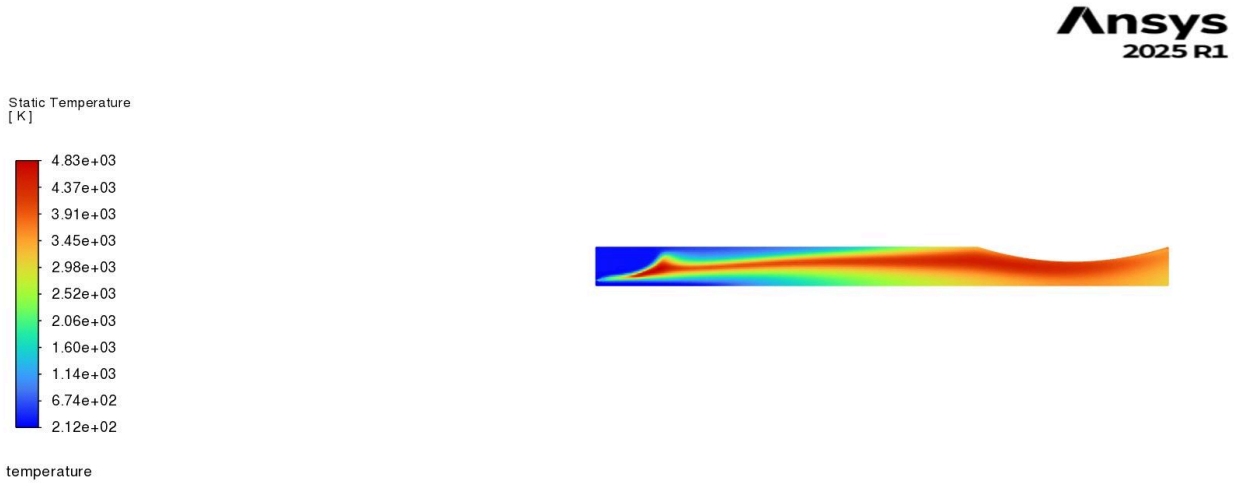

Figure 9: Temperature contour for 40% swirl.

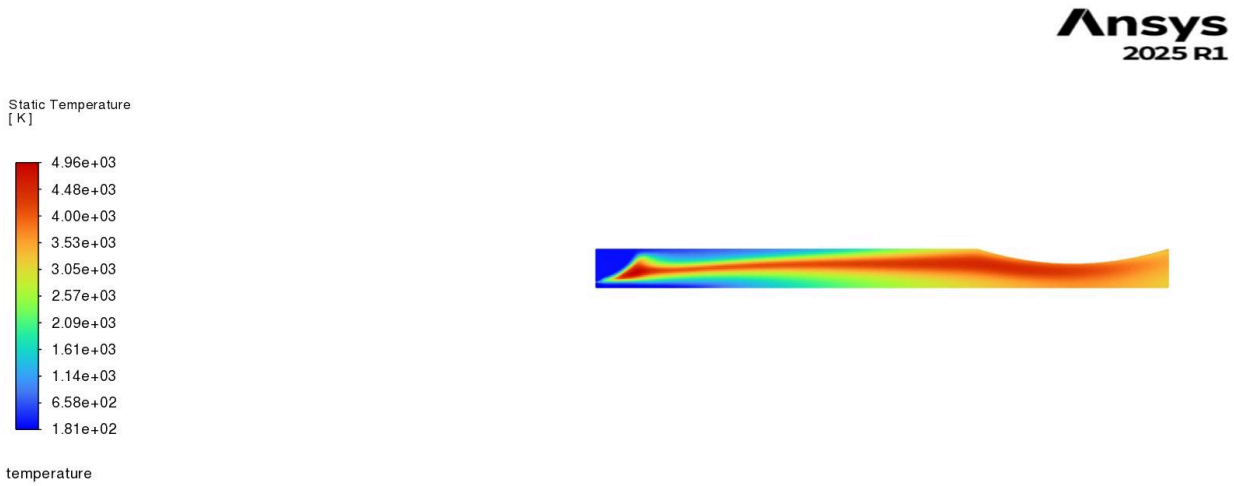

Figure 10: Temperature contour for 70% swirl.

Looking at the legend on the left hand side, it is observed that the maximum temperature of the flow does in fact increase as the swirl percentage increases. For 10% swirl, the maximum temperature is 4360 K, for 40% swirl, the maximum temperature is 4830 K, and for 70% swirl, the maximum temperature is 4960 K. As the swirl is increased, it promotes mixing between the fuel and oxidizer, resulting in more combustion and more chemical energy being released as thermal energy, increasing the temperature of the flow. What is perhaps more interesting is looking at the mass fractions of the combustion products as the swirl % changes. Figures 10, 11, and 12 show the CO_2 mass fraction contours of the flow when the swirl is 10%, 40%, and 70% respectively. Results are discussed below.

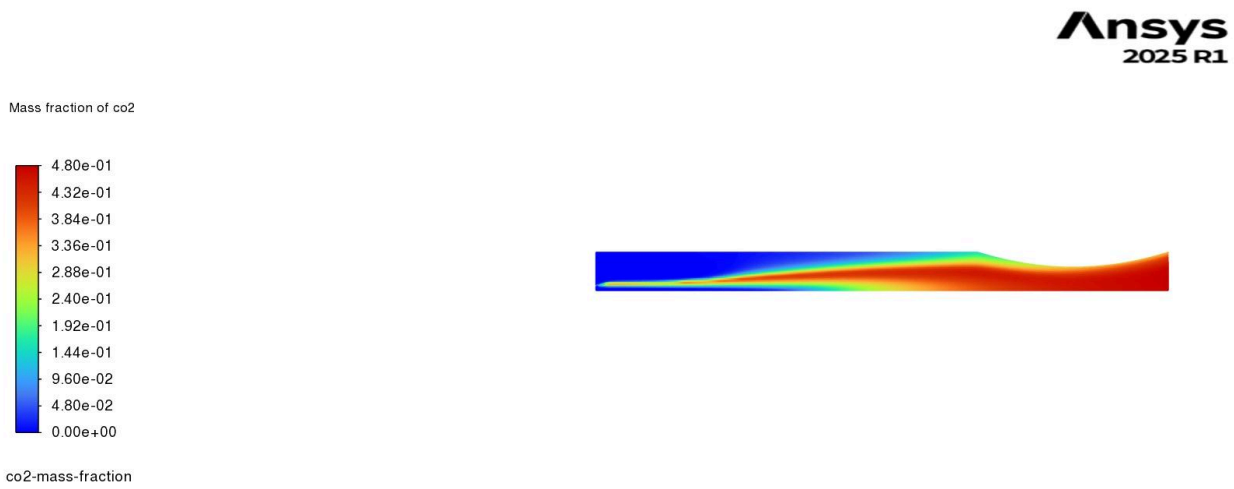

Figure 11: CO_2 mass fraction contour for 10% swirl.

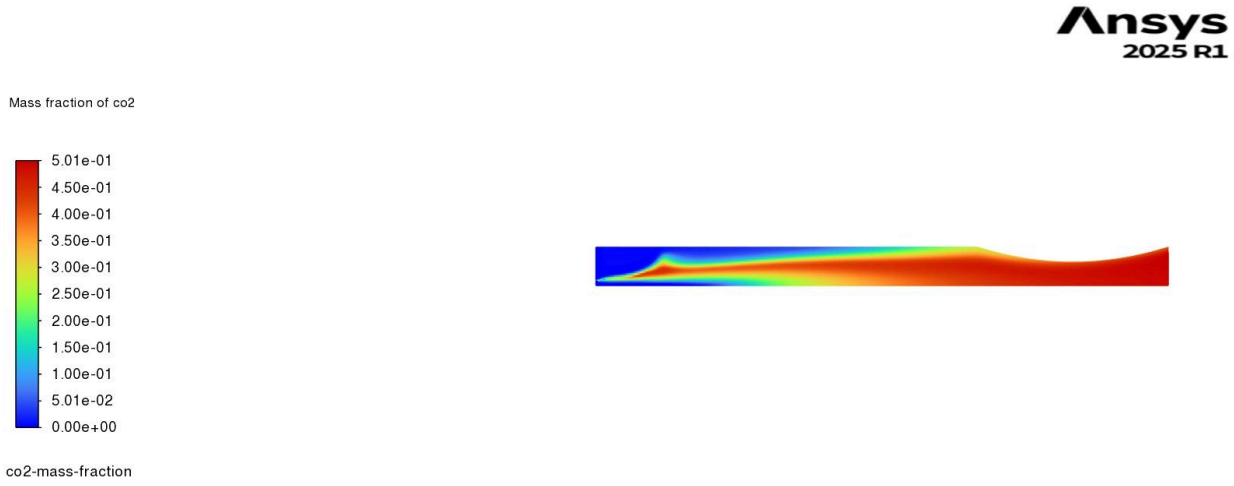

Figure 12: CO_2 mass fraction contour for 40% swirl.

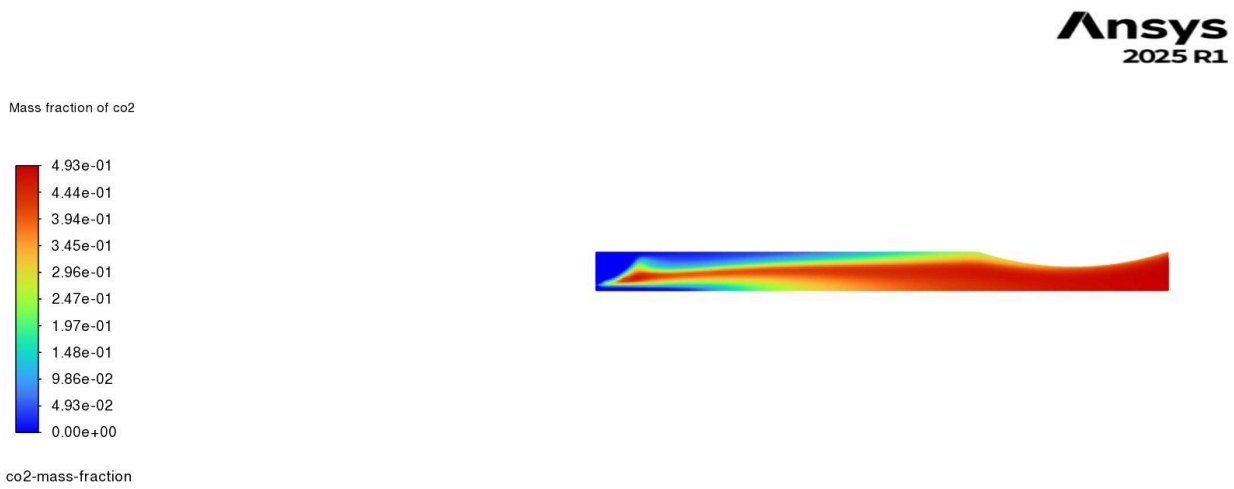

Figure 13: CO_2 mass fraction contour for 70% swirl.

Looking at the legend on the left hand side, it is observed that the maximum mass fraction of CO_2 increases from 0.480 to 0.501 when the swirl is increased from 10% to 40%, but it then slightly decreases to 0.493 when the swirl is further increased to 70%. This is interesting because it did not follow the expectation that the maximum mass fraction of combustion products would strictly increase with increase in swirl percentage. The reason that this expectation holds for temperature is that temperature is more of a cumulative metric that is reflective of the total chemical energy released during combustion that is converted into chemical energy. The mass fraction of combustion products at a specific point in the geometry, however, reflects a local concentration and depends on how spread out the combustion products are. It is possible that the further increase in swirl from 40% to 70% contributed to the combustion products spreading out more, resulting in a lower concentration of combustion products.

An important note, though, is that to maximize the exhaust velocity of a rocket engine and therefore the specific impulse, it is not necessarily the case that both temperature and mass fraction of combustion products should be maximized. Looking at the exhaust velocity equation

below, it is clear that the exhaust velocity is proportional to the ratio of T/M, where T is the total temperature of the combustion chamber and M is the average molar mass of the flow.

$$v_e = \sqrt{\frac{TR}{M} \cdot \frac{2\gamma}{\gamma - 1} \cdot \left[1 - \left(\frac{p_e}{p}\right)^{\frac{\gamma-1}{\gamma}}\right]}$$

Equation 1: This is the equation for the exhaust velocity of a rocket engine. Source: Wikipedia

It is possible, and is actually often the case, that the O/F ratio that maximizes exhaust velocity and therefore specific impulse is not the stoichiometric ratio. For example, in methane-oxygen rocket engines, a slightly fuel rich mixture ratio will decrease the combustion chamber temperature slightly (as seen in the next set of simulations), but the molar mass of the reaction products will be decreased by a larger amount because the excess uncombusted methane has a lighter molar mass than the combustion products. This results in an O/F ratio that maximizes specific impulse closer to ~3-4-3.6 for methane oxygen rocket engines, which is closer to the mixture ratio that the SpaceX Raptor 3 engine typically operates at. A mixture ratio that is not stoichiometric also reduces the combustion chamber temperature and therefore makes the combustion chamber easier to cool.

Varying O/F Mixture Ratio

The third set of simulations run in this project varied the O/F mixture ratio of the propellant by manually changing the magnitudes of the mass flow rate inlet boundary conditions. Mixture ratios of 3.5, 4, and 4.5 were tested with constant total mass flow rates. A mixture ratio of 4 is very close to stoichiometric for methane oxygen combustion. The expectation is that as the mixture ratio becomes further from stoichiometric, there will be more excess uncombusted

propellant and therefore less chemical energy released, resulting in lower maximum temperatures in the combustion chamber. Figures, 13, 14, and 15 show the temperature contours of the flow when the mixture ratio is 3.5, 4, and 4.5 respectively. Results are discussed below.

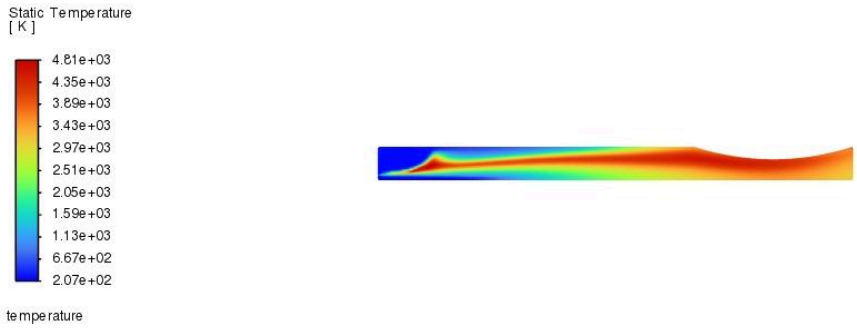

Figure 14: Temperature contour for mixture ratio of 3.5.

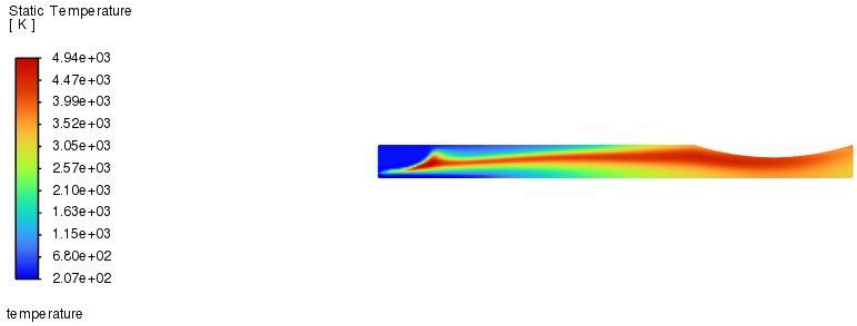

Figure 15: Temperature contour for mixture ratio of 4.0.

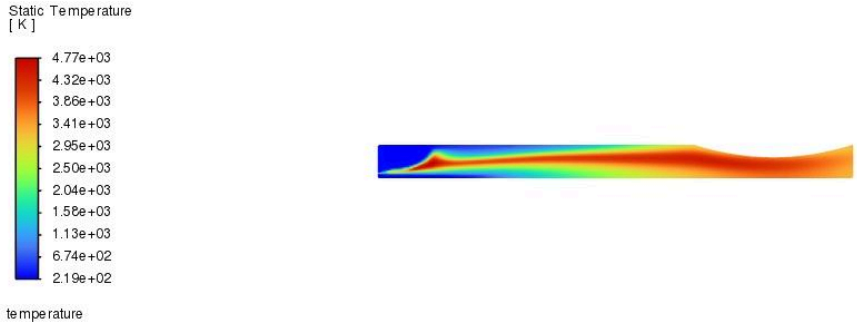

Figure 16: Temperature contour for mixture ratio of 4.5.

Looking at the legend on the left hand side, it is observed that the maximum temperature of the flow does in fact decrease as the mixture ratio gets further away from stoichiometric. For a mixture ratio of 4.0 (stoichiometric), the maximum temperature is 4940 K, for mixture ratio of 3.5, the maximum temperature is 4810 K, and for mixture ratio of 4.5, the maximum temperature is 4770 K. As the mixture ratio varied further from stoichiometric, there was less combustion and less chemical energy released, resulting in lower temperatures. The excess oxygen also could absorb some thermal energy and reduce the temperature as well.

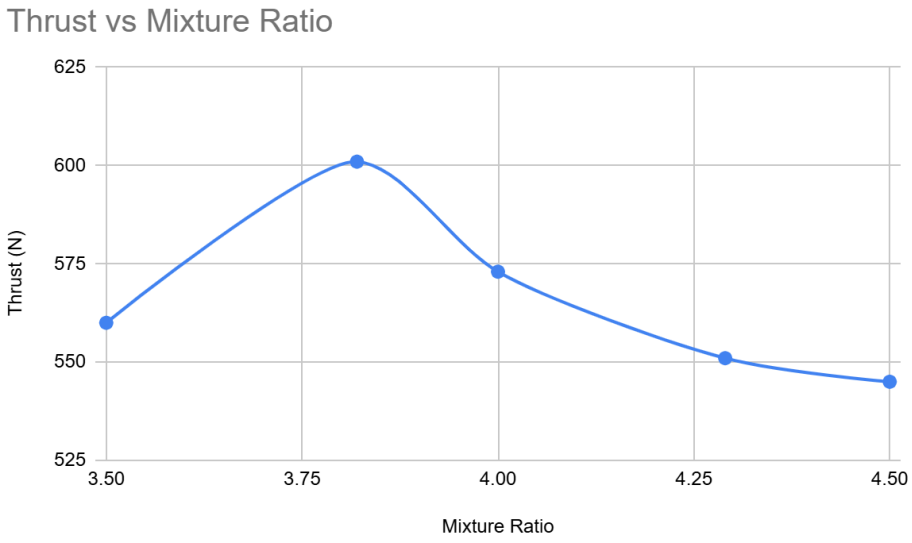

Figure 17: Thrust vs mixture ratio comparison for five different ratios

In addition to temperature, thrust was also computed for each mixture ratio simulation. As shown in Figure 17, the thrust peaks at approximately 601 N at a mixture ratio of 3.82 and decreases on both sides of this peak. This is consistent with the exhaust velocity equation discussed earlier, where exhaust velocity depends on the ratio of temperature to average molar mass. While the stoichiometric ratio of 4 produces the highest temperatures, a slightly fuel rich mixture reduces the average molar mass by a larger proportion because excess uncombusted methane is much lighter than the combustion products. This is why peak thrust occurs at a fuel rich condition rather than at stoichiometric, and is consistent with how methalox engines like the SpaceX Raptor typically operate.

Conclusion

This project demonstrated the ability to model and simulate a bipropellant rocket engine engine in CFD software, a critical skill to working with propulsion systems in industry. The results of the simulations demonstrated concepts that are commonly used in rocket engine design at aerospace companies. The first set of simulations demonstrated why high pressure ratios (chamber pressure / exit pressure) are used in rocket engines – because they allow for choked, supersonic flow and higher exit velocities. While not specifically demonstrated in this project, above the critical pressure ratio that chokes the flow, an increase in chamber pressure doesn't increase the exit velocity, but it will increase the mass flow rate, resulting in greater thrust. The second set of simulations demonstrated why injector elements often swirl the propellant flow or are designed to have the propellants impinge on one another – because promoting mixing of the propellants results in more combustion and chemical energy released which can then be converted into kinetic energy and expelled from the vehicle to generate thrust. The third set of

simulations demonstrated why rocket engines do not typically operate at stoichiometric mixture ratios – because operating at fuel rich or oxygen rich mixture ratios lowers the chamber temperatures and keeps the walls of the chamber from melting. Further experiments could be done to understand various parts of rocket engine design like varying the inlet temperatures of the propellants, changing the geometry of the combustion chamber, or using a 3D model geometry to analyze more complex injector element patterns. CFD software like Ansys Fluent provides a useful platform for conducting these experiments before moving forward with much more expensive and resource intensive hotfire tests.