Shape Memory Alloy-Augmented Thermal Protection Systems for Hypersonic Re-entry Vehicles

a project presented to
The Faculty of the Department of Aerospace
Engineering San José State University

in partial fulfillment of the requirements for the degree
Master of Science in Aerospace Engineering

by

Ralph R. Elias

May 2026

approved by

Dr. Yawo Ezunkpe
Faculty Advisor



© 2026
Ralph Elias
ALL RIGHTS RESERVED

ABSTRACT

Shape Memory Alloy-Augmented Thermal Protection Systems for Hypersonic Re-entry Vehicles

Ralph R. Elias

This study investigates the feasibility of a smart material based adaptive thermal protection system (ATPS) for representative hypersonic re-entry applications. The proposed system is not treated as a standalone smart material heat shield, but as a multilayer reusable ceramic thermal protection architecture augmented with an embedded adaptive interlayer. The baseline stack consists of a high-temperature emissivity coating, an outer reusable ceramic thermal protection layer, a shape memory alloy (SMA) interlayer, a ceramic or composite buffer layer, a strain isolation layer, adhesive bondlines, and a protected vehicle skin. Within this architecture, the outer ceramic materials provide the primary thermal protection function, while the SMA layer is positioned at an intermediate depth where it can undergo thermally induced phase transformation and generate recovery stress within the stack. The recovery stress generated by the SMA interlayer and transferred onto the stack is hypothesized to put the layers in a more compressed state and counteract the tensile forces that cause cracking and delamination in standard reusable ceramic based TPS.

A coupled computational workflow is developed using ANSYS SpaceClaim, ANSYS Meshing, ANSYS Fluent, and MATLAB. Two representative hypersonic geometries are evaluated: an axisymmetric sphere-cone representative of a blunt forebody or capsule heat shield, and a planar blunted wedge representative of a hypersonic leading edge or airfoil-like surface. Hypersonic CFD is performed in Fluent at a representative Mach 20, 65 km freestream condition using a pressure farfield boundary, density-based solver, energy equation, k-omega SST turbulence model, finite rate chemistry, species transport, and a five-species air model. The resulting surface heat fluxes are passed into a MATLAB thermal model that evaluates radiative equilibrium wall temperature, through thickness conduction, intermediate layer temperatures, SMA layer temperature, and vehicle skin temperature. A reduced order SMA recovery stress model is then used to estimate the effective transferred stress generated by the embedded SMA mesh.

The results show that the proposed ATPS architecture is feasible within a bounded thermal design space. The successful heating cases of 0.1 MW/m2, 0.62 MW/m2, and 1.2 MW/m2 remained within the approximate 150 mm practical stack thickness cap while reducing the vehicle skin temperature to protected ranges. When compared against literature heat fluxes for orbital and lunar return re-entry vehicles these cases reinforce the idea that the proposed architecture is viable for real world blunt body and acreage zone heating profiles. The highest successful case, 1.2 MW/m2, produced an outer ceramic temperature near 2201 K, reduced the TUFROC layer temperature to approximately 1822 K, placed the SMA interlayer near 686 K,

and maintained the vehicle skin near 364 K. In contrast, higher heat flux cases of 3.3 MW/m2, 8.0 MW/m2, and 13.8 MW/m2 exceeded the operating limits of the selected material stack and were treated as outside the practical design envelope. These results indicate that the proposed ATPS concept is most appropriate for blunt forebody acreage regions, relaxed heating zones, and downstream leading edge panel regions rather than severe leading edge stagnation points or localized geometry transition peaks.

The SMA recovery stress analysis indicates that the embedded adaptive layer can generate meaningful thermomechanical response when placed at the correct depth in the TPS stack. The successful thermal cases produced SMA layer temperatures suitable for partial or full phase transformation depending on material selection, with effective transferred stresses on the order of several MPa to tens of MPa depending on load transfer efficiency. Overall, this study supports the feasibility of an SMA-augmented reusable TPS architecture for selected moderate re-entry heating environments, while also identifying key limitations associated with outer ceramic temperature capability, isothermal wall assumptions in CFD, interface load transfer, and the need for future conjugate heat transfer and experimental validation.

Table of Contents

Table of Contents **5**
List of Figures **8**
List of Tables **10**
Nomenclature **11**
1.0 Introduction **1**
1.1 Motivation 1
1.2 Literature Review 2
1.2.1 High Speed Gas Dynamics and the Hypersonic Re-entry Environment 2
1.2.2 Thermal Loads and Material Constraints 6
1.2.3 Necessity of TPS and Historical Evolution 7
1.2.4 Categorization of TPS Requirements 9
1.2.5 Integrated/Multifunctional and Sandwich Structures for TPS 11
1.2.6 Advanced Computational Techniques in TPS Analysis 13
1.2.7 Definition and Classification of Smart Materials 15
1.2.8 Smart Materials and Multifunction Material Applications in Aerospace 17
1.2.9 Thermomechanical Mechanisms of Shape Memory Alloys 19
1.2.10 Other Smart Material Categories — An Overview 24
1.3 Project Proposal 25
1.4 Methodology 27
2.0 Materials & TPS Architecture **29**
2.1 Smart Material Candidates 29
2.1.2 Materials Discussion 35
2.2 ATPS Layering and Configuration 37
2.2.1 Layering Methodology 37
2.2.2 Functional Operation of Layered TPS Configurations 39
3.0 Geometry **43**
3.1 Overview 43
3.2 Dimensions and Profile 44
3.2.1 Leading Edge (Blunt Tipped Wedge): 46
3.2.2 Blunt Forebody (Axisymmetric Sphere-cone): 48
4.0 Meshing/Topology **51**
4.1 Mapped Mesh Approach 51
4.2 Leading Edge Mesh Topology 53
4.3 Blunt Forebody Mesh Topology 54
4.4 Edge Sizing and Biasing 54
4.5 Near Wall Resolution and y Plus Target 57
4.6 Wall Alignment and Heat Flux Accuracy 59
5.0 CFD **60**

5.1 Boundary Conditions..... 60
5.1.1 Inlet Conditions..... 60
5.1.2 Outlet Conditions..... 61
5.1.3 Wall Conditions..... 62
5.2 Model Physics Setup..... 65
5.2.1 General Physics Settings..... 69
5.2.2 Energy Equation and One Temperature Assumption..... 69
5.2.3 Turbulence Modeling..... 70
5.2.4 Species Transport and Chemical Nonequilibrium..... 70
5.2.5 Transport Properties..... 71
5.3 Solution Methods, Controls, and Initialization..... 71
5.3.1 Solver Controls..... 72
5.3.2 Spatial Discretization and Gradient Evaluation..... 72
5.3.3 Chemistry Solver Controls..... 73
5.3.4 Full Multigrid Initialization..... 73
5.3.5 Iteration Strategy and Calibration Cycles..... 74
5.3.6 Solution Monitoring..... 74
5.4 Results..... 75
5.4.1 Hypersonic Leading Edge (Blunt Tipped Wedge)..... 76
5.4.2 Blunt Re-entry Forebody (Axisymmetric Sphere Cone)..... 83
5.5 Discussion..... 90
5.5.1 Contour Plot Interpretation..... 90
5.5.2 Surface Heat Flux Behavior..... 91
5.5.3 Off Stagnation Shoulder Heating..... 92
5.5.4 Interpretation of Wall Shear Stress Alignment..... 93
5.5.5 Isothermal Wall Limitation..... 93
5.5.6 Segmented Wall Temperature Calibration..... 94
5.5.7 Theoretical Heatflux Ceiling..... 95
6.0 Thermal Analysis (Post Processing)..... 97
6.1 Scope and Modeling Assumptions..... 97
6.2 Thermal Boundary Conditions and CFD Coupling..... 99
6.3 Multilayer Conduction Model..... 100
6.3.1 Layering, coordinate, and notation..... 100
6.3.2 Governing Equations (Fourier conduction laws)..... 100
6.4 Back Face Temperature and Structural Thermal Constraint..... 102
6.5 Nitinol Layer Temperature Distribution and Transformation Window..... 104
6.6 MATLAB Solution Procedure and CFD–Thermal Calibration..... 105
6.7 Results..... 106
Case 1 - Leading Edge (Blunt Tip Wedge) Stagnation Point..... 106
Case 2 - Leading Edge (Blunt Tip Wedge) Acreage Relaxation Zone..... 107
Case 3 - Leading Edge (Blunt Tip Wedge) Near Truncation..... 108

Case 4 - Blunt Forebody (Axisymmetric Sphere-Cone) Stagnation Point..... 108
Case 5 - Blunt Forebody (Axisymmetric Sphere-Cone) Acreage Relaxation Zone..... 109
Case 6 - Blunt Forebody (Axisymmetric Sphere-Cone) Geometry Transition Off Stagnation Heating Spike.....110
6.8 Discussion.....110
7.0 SMA Recovery Stress Analysis.....116
7.1 Scope and Limitations.....116
7.2 Physical Basis of SMA Recovery Stress.....117
7.3 Transformation Temperatures and Viability Criteria.....117
7.4 Temperature-Path Representation.....118
7.5 Phase-Dependent Material Properties.....120
7.6 Installed Reference State and Constraint Treatment.....120
7.7 Mesh Geometry and Effective SMA Fraction.....122
7.8 Load Transfer Efficiency.....122
7.9 Stress-Temperature Envelope and Governing Stress.....123
7.10 MATLAB Implementation.....124
7.11 Results.....124
7.11.1 Case A Results.....125
7.11.2 Case B Results.....128
7.11.3 Case C Results.....130
8.0 Conclusion and Recommendations.....132
8.1 Summary of Work.....132
8.2 Major Findings.....133
8.3 Conclusions.....136
8.4 Recommendations.....138
9.0 Future Work.....140
References.....143

List of Figures

Figure 1.1 – Temperature behind a normal shock wave as a function of freestream velocity at a standard altitude of 52 km. Reproduced from ref. [1] …….. 4
Figure 1.2 – Maximum non-ablative heat transfer wrt. Velocity for past and future NASA missions. Reproduced from ref. [3] ..…… 5
Figure 1.3 – Re-entry maximum heating interval in altitude vs. time plot for Space Shuttle program. Reproduced from ref. [6] ..……...7
Figure 1.4 – Hierarchical sandwich TPS concept with embedded PCM in outer core and insulation in inner core. Reproduced from ref. [9] ..…… 11
Figure 1.5 – Sandwich TPS architectures: (a) Corrugated core with C/SiC, (b) Laminated composite with carbon foam and CFRP-aluminum honeycomb, (c) Lattice foam core made of CMC/SiC. Reproduced from ref. [6] ……. 12
Figure 1.6 – Various sandwich structure core designs. Reproduced from ref. [6] ..……. 12
Figure 1.7 – PIROM performance under material property perturbations. (a) Nominal and perturbed material properties; (b) Predicted trajectory comparisons for PIROM vs Fine FEM for both cases. Reproduced from ref. [13] ..……... 14
Figure 1.8 – Comparison of Method 1 G and CFD for modeling heating in hypersonic vehicles. Reproduced from ref. [3] ..…….. 15
Figure 1.9 – Stress-temperature phase diagram of shape memory alloys showing phase boundaries for martensitic and austenitic transformation. Reproduced from ref. [17] ..…….. 21
Figure 1.10 – Shape memory cycle of an SMA element under heating and cooling. Reproduced from ref. [17] ..…….. 22
Figure 3.1 – Leading edge (blunt tipped wedge) geometry ……. 47
Figure 3.2 – Leading edge (blunt tipped wedge) fluid domain ……... 48
Figure 3.3 – Blunt forebody (axisymmetric sphere-cone) geometry ……... 49
Figure 3.4 – Blunt forebody (Axisymmetric sphere-cone) fluid domain ..……. 50
Figure 4.1 – Leading edge (blunted wedge) full mesh ..……. 52
Figure 4.2 – Blunt forebody (axisymmetric sphere-cone) full mesh ….…….. 53
Figure 4.3 – Blunted wedge nearfield mesh ..…….. 55
Figure 4.4 – Blunted wedge near wall front ..……. 55
Figure 4.5 – Blunted wedge near wall aft ..……. 56
Figure 4.6 – Axisymmetric forebody nearfield mesh ..…… 57
Figure 4.7 – Blunted wedge first cell height …… 58
Figure 4.8 – Axisymmetric forebody first cell height …….. 59
Figure 5.1 – Wall catalycity illustration. Reproduced from ref. [49] ……...…..……... 63
Figure 5.2 – Partially catalytic models for TPS. Reproduced from ref. [49] ....……... 64
Figure 5.3 – Leading edge: contour of static temperature (full view) ….…… 76
Figure 5.4 – Leading edge: contour of static temperature (nearfield) ……. 77
Figure 5.5 – Leading edge: contour of static temperature (near wall) …… 77
Figure 5.6 – Leading edge: contour of static pressure (full view) ……... 78
Figure 5.7 – Leading edge: contour of static pressure (nearfield) ..…… 79
Figure 5.8 – Leading edge: contour of static pressure (near wall; altered for gradient) ….…… 79
Figure 5.9 – Leading edge: contour of mach number (full view) ….……... 80
Figure 5.10 – Leading edge: contour of mach number (nearfield) ……... 81
Figure 5.11 – Leading edge: contour of mach number (near wall) ….……. 81

Figure 5.12 – Leading edge: front surface heat flux vs. y Position …… 82
Figure 5.13 – Leading edge: front surface wall shear stress vs. y position …… 82
Figure 5.14 – Blunt forebody: contour of static temperature (full view) ……... 83
Figure 5.15 – Blunt forebody: contour of static temperature (nearfield) …… 84
Figure 5.16 – Blunt forebody: contour of static temperature (near wall) ……... 84
Figure 5.17 – Blunt forebody: contour of static pressure (full view) …….. 85
Figure 5.18 – Blunt forebody: contour of static pressure (nearfield) ...……... 86
Figure 5.19 – Blunt forebody: contour of static pressure (near wall) …….. 86
Figure 5.20 – Blunt forebody: contour of mach number (Full View) ...……... 87
Figure 5.21 – Blunt forebody: contour of mach number (Nearfield) ...……... 88
Figure 5.22 – Blunt forebody: contour of mach number (Near Wall) ...…….. 88
Figure 5.23 – Blunt forebody: surface heat flux vs. y position (front surface) ……... 89
Figure 5.24 – Blunt forebody: wall shear stress vs. y position (front surface) …… 89
Figure 7.1 – Case A: primary effective stack stress vs. temperature ...……..127
Figure 7.2 – Case A: eta sweep effective stack stress vs. temperature …….. 127
Figure 7.3 – Case B: primary effective stack stress vs. temperature ...…….. 129
Figure 7.4 – Case B: eta sweep effective stack stress vs. temperature …….. 129
Figure 7.5 – Case C: primary effective stack stress vs. temperature ...…….. 131
Figure 7.6 – Case C: eta sweep effective stack stress vs. temperature …….. 131

List of Tables

Table 1.1 – Smart material types with respective inputs and outputs. Reproduced from ref. [14]……17
Table 2.1 – Table of materials ……. 30
Table 2.2 – TPS layer visualization ..……... 39
Table 3.1 – Geometry dimensions ……... 45
Table 5.1 – General atmospheric conditions for LEO re-entry phase ……. 60
Table 5.2 – Physics set up ANSYS Fluent ..……. 65
Table 5.3 – Calibration results …… 75
Table 6.1 – Case 1 results ……. 106
Table 6.2 – Case 2 results ……. 107
Table 6.3 – Case 3 results ……. 108
Table 6.4 – Case 4 results ……. 108
Table 6.5 – Case 5 results ……. 109
Table 6.6 – Case 6 results ……. 110
Table 7.1 – SMA recovery stress response analysis: inputs …….. 124
Table 7.2 – SMA recovery stress response analysis: fixed inputs …… 125
Table 7.3 – Case A: primary effective stack stress …….. 125
Table 7.4 – Case A: secondary eta sweep comparison ..……... 126
Table 7.5 – Case B: primary effective stack stress …….. 128
Table 7.6 – Case B: secondary eta sweep comparison …… 128
Table 7.7 – Case C: primary effective stack stress …….. 130
Table 7.8 – Case C: secondary eta sweep comparison …… 130

Nomenclature

Roman Symbols

A	Area
A_f	Austenite finish temperature
A_s	Austenite start temperature
a_{ij}	Number of atoms of element j in species i
b_j	Total moles of element j
c_p	Specific heat at constant pressure
d	SMA mesh wire diameter
E	Young's modulus
E_A	Austenite Young's modulus
E_M	Martensite Young's modulus
$E(\xi)$	Young's modulus as a function of martensite content
G	Gibbs free energy
h	Heat transfer coefficient / enthalpy (context dependent)
h_0	Stagnation enthalpy
k	Thermal conductivity
L	Layer thickness or characteristic length
M	Freestream Mach number
M_f	Martensite finish temperature
M_s	Martensite start temperature
n_i	Number of moles of species i
p	Static pressure
p_∞	Freestream static pressure
Pr	Prandtl number
q	Heat flux
$q"$	Surface heat flux
q_{cond}	Conductive heat flux
q_{rad}	Radiative heat flux
Q	Total heat transfer rate
Q_{CFD}	Total heat transfer rate extracted from CFD
R	Universal gas constant / thermal resistance (context dependent)
R_{cond}	Conduction thermal resistance
R_n	Nose radius
Re	Reynolds number
Re_x	Reynolds number at position x
s	SMA mesh wire spacing
T	Temperature

T_0	Reference temperature
T_∞	Freestream static temperature
T_A	Austenite transformation temperature, general notation
T_{allow}	Maximum allowable material temperature
T_{back}	Backface temperature
T_f	Final temperature
T_M	Martensite transformation temperature, general notation
T_{peak}	Peak temperature reached during thermal cycle
T_{rec}	Recovery temperature estimate
T_s	Surface temperature
T_{wall}	Wall temperature
T_w	Wall temperature
V	Velocity
V_∞	Freestream velocity
$v_{f,mesh}$	Effective SMA mesh volume fraction
x	Streamwise coordinate
y	Wall normal or radial coordinate
y+	Dimensionless wall distance

Greek Symbols

α	Thermal diffusivity / coefficient of thermal expansion, (context dependent)
α_A	Austenite coefficient of thermal expansion
α_M	Martensite coefficient of thermal expansion
δ	Boundary layer thickness
ε	Total or mechanical strain
ε_{ref}	Reference or installed prestrain
ε_{tol}	Convergence tolerance
ε_{rad}	Surface emissivity
$\eta_{transfer}$	Load transfer efficiency factor
γ	Specific heat ratio / wall catalytic recombination coefficient (context dependent)
λ_c	Constraint or load transfer scaling coefficient
μ	Chemical potential / dynamic viscosity, (context dependent)
μ_i	Chemical potential of species i
ρ	Density
ρ_∞	Freestream density
σ	Stress / Stefan-Boltzmann Constant (context dependent)
σ_0	Reference stress
σ_{raw}	Raw SMA recovery stress
$\sigma_{transfer}$	Effective transferred stress into the TPS stack
τ_w	Wall shear stress

Θ	Thermal expansion modulus
ξ	Martensite volume fraction
ξ_0	Reference martensite volume fraction
ξ_{peak}	Peak martensite volume fraction during thermal path
ξ_T	Temperature-induced martensite volume fraction
ξ_S+, ξ_S-	Tension/compression-induced martensite volume fractions
Ω	Martensitic transformation coefficient
Ω_S+, Ω_S-	Martensitic transformation stress coefficients

Subscripts

0	Reference or initial state
∞	Freestream condition
A	Austenite phase
back	Backface or vehicle-side surface
CFD	Computational fluid dynamics result
cond	Conduction
eff	Effective value
f	Finish or final state
i	Species or layer index
j	Element index
M	Martensite phase
max	Maximum value
min	Minimum value
peak	Peak value during thermal cycle
rad	Radiation
ref	Reference or installed state
s	Surface / start condition (context dependent)
w	Wall condition

Acronyms and Abbreviations

AETB	Alumina Enhanced Thermal Barrier
ATPS	Adaptive Thermal Protection System
CFD	Computational Fluid Dynamics
CHT	Conjugate Heat Transfer
CMC	Ceramic Matrix Composite
CFL	Courant-Friedrichs-Lewy number
FEM	Finite Element Method
FMG	Full Multigrid
FRCI	Fibrous Refractory Composite Insulation

HIAD	Hypersonic Inflatable Aerodynamic Decelerator
ISA	International Standard Atmosphere
ITPS	Integrated Thermal Protection System
LEO	Low Earth Orbit
MLI	Multilayer Insulation
PCM	Phase Change Material
PICA	Phenolic Impregnated Carbon Ablator
PIROM	Physics-Infused Reduced-Order Modeling
RCC	Reinforced Carbon-Carbon
RCG	Reaction-Cured Glass
ROM	Reduced-Order Model
SIP	Strain Isolation Pad
SMA	Shape Memory Alloy
SME	Shape Memory Effect
SMP	Shape Memory Polymer
SST	Shear Stress Transport
TPS	Thermal Protection System
TUFI	Toughened Unipiece Fibrous Insulation
TUFROC	Toughened Uni-piece Fibrous Reinforced Oxidation-Resistant Composite
UDF	User-Defined Function
UHTC	Ultra-High Temperature Ceramic
YSZ	Yttria-Stabilized Zirconia

1.0 Introduction

1.1 Motivation

Hypersonic re-entry vehicles are required to operate in one of the most extreme environments encountered by any aerospace vehicle. During re-entry through Earth's atmosphere, spacecraft can reach speeds that exceed Mach 20 which in effect, subjects these vehicles to intense aerodynamic heating and pressure loads. These aerothermal effects cause rapid temperature increases around 1500 degrees celsius or more, often leading to structural degradation, ablation, or in some cases catastrophic failure. For this reason, hypersonic re-entry has been a major point of concern and critical failure for space missions over the past few decades.

To ensure structural integrity, Thermal Protection Systems (TPS) are a required design feature for all re-entry systems. Traditionally TPS designs have relied heavily on the use of ablative materials; these materials absorb heat and dissipate thermal energy primarily through material loss. While ablatives have legacy in NASA missions and have been proven to be largely effective, they are often single use, heavy, and lack adaptability to changing conditions. Passive and single use systems simply do not meet the demands of today's rapidly changing space technology market.

In the modern aerospace industry, rapid reusability is becoming an increasingly sought after and even required capability for proposed mission architectures. Current solutions involve using reusable ceramic tiles which are still limited by their passive nature and susceptibility to damage under repeated thermal stress. Thus the demand grows for more advanced solutions to thermal protection. Adaptive TPS technologies may be able to offer real-time thermal management, structural adaptability, and repeated use without major need for refurbishment.

In response to these demands Smart Materials particularly shape memory alloys have emerged as promising candidates for revolutionizing TPS design. These materials could be engineered into Adaptive Thermal Protection Systems (ATPS) offering significant advantages in aerothermal performance, structural resilience, and enabling rapid reusability. Although promising Smart Materials have yet to be employed in thermal shielding on flight tested re-entry vehicles. Thus there is still a major gap in the literature around ATPS and with a growing need for innovative solutions, it may be important for this generation of aerospace engineers to delve into this concept.

This project is motivated by the need to investigate and validate the feasibility of Smart Material-based ATPS for hypersonic re-entry applications. Specifically, this study will seek to address key technical challenges including the effective integration of smart materials into TPS architectures, modeling of adaptive thermo-structural responses, and computational validation through detailed numerical simulations under realistic hypersonic conditions. By addressing these challenges, the project aims to significantly advance the state-of-the-art in hypersonic

vehicle design, aligning closely with the aerospace industry's objectives of increased reliability, adaptability, and rapid reusability for next-generation hypersonic re-entry systems.

1.2 Literature Review

1.2.1 High Speed Gas Dynamics and the Hypersonic Re-entry Environment

Vehicles returning from low Earth orbit (LEO) or beyond must transition from the vacuum of space into Earth's dense, multi-layered atmosphere. This re-entry almost always occurs at hypersonic speeds. The term hypersonic is often loosely defined as flow velocities exceeding Mach 5. However, in reality, hypersonic flow is actually a much more complicated thermo-chemical change. Re-entry vehicles often reach velocities well into the range of Mach 20–25, which corresponds with freestream velocities of 6–8 km/s depending on entry trajectory and vehicle mass. At such velocities, the aerodynamic and thermodynamic characteristics of the flow are governed by complex and coupled phenomena that can heavily impact and threaten vehicle survivability.

As a hypersonic body enters the Earth's atmosphere it compresses the air in front of it which in turn generates an intense shock wave. The air's temperature and pressure after the shock wave rapidly increases, leading to extreme levels of aerodynamic heating through convective and radiative heat transfer. The post-shock gas can reach temperatures in excess of 10,000 K, which is sufficient to cause dissociation of diatomic oxygen (O_2) and nitrogen (N_2), as well as ionization of atomic species. The formation of this viscous shock layer traps a high enthalpy and chemically reactive gas between the shock wave and the vehicle's surface. [1]

At these extreme temperatures, air can no longer be treated as a calorically perfect gas. Instead, it behaves as a high temperature chemically reacting mixture composed of O, N, O_2, N_2, NO, and eventually electrons and ionized species such as O^+ and N^+. The dissociation of molecular oxygen and nitrogen occurs via endothermic reactions as shown below.

$$O2 + \Delta H \rightarrow 2O, \qquad N2 + \Delta H \rightarrow 2N \tag{1}$$

These reactions absorb energy from the flow and reduce the effective stagnation temperature at the surface, but they simultaneously create chemically reactive species capable of oxidizing or eroding TPS materials. At higher energies, ionization becomes significant. [1]

$$O + \Delta H \rightarrow O^+ + e^-, \qquad N + \Delta H \rightarrow N^+ + e^- \tag{2}$$

The progression through translational, rotational, vibrational, electronic, and finally chemical energy modes results in a flow that is thermodynamically and chemically nonequilibrium — particularly within the boundary layer near the stagnation point. This complexity often requires the use of finite-rate chemistry models and multi-temperature models in CFD simulations to resolve internal energy relaxation and species production rates accurately. [1]

To predict the composition of the post-shock gas at a given temperature and pressure, the species number densities (n_i) are typically calculated using equations 1 and 2 and by minimizing the Gibbs free energy of the system.

$$G = \sum_{i=1}^{N} n_i \cdot \mu_i \tag{3}$$

The number densities are also subject to species conservation in thermochemical equilibrium.

$$\sum_{i=1}^{N} a_{ij} n_i = b_j, \quad \text{for each conserved element } j \tag{4}$$

Figure 1.1 below adapted from Anderson's textbook on high temperature gas dynamics plots the shock layer temperature in K x 10^{-3} as a function of the Re-entry velocity in km/s for air as a calorically perfect gas (CPG) and an equilibrium chemically reacting gas or thermally perfect gas (TPG). This graphics aids in understanding the difference in behavior of air under these two assumptions. [1]

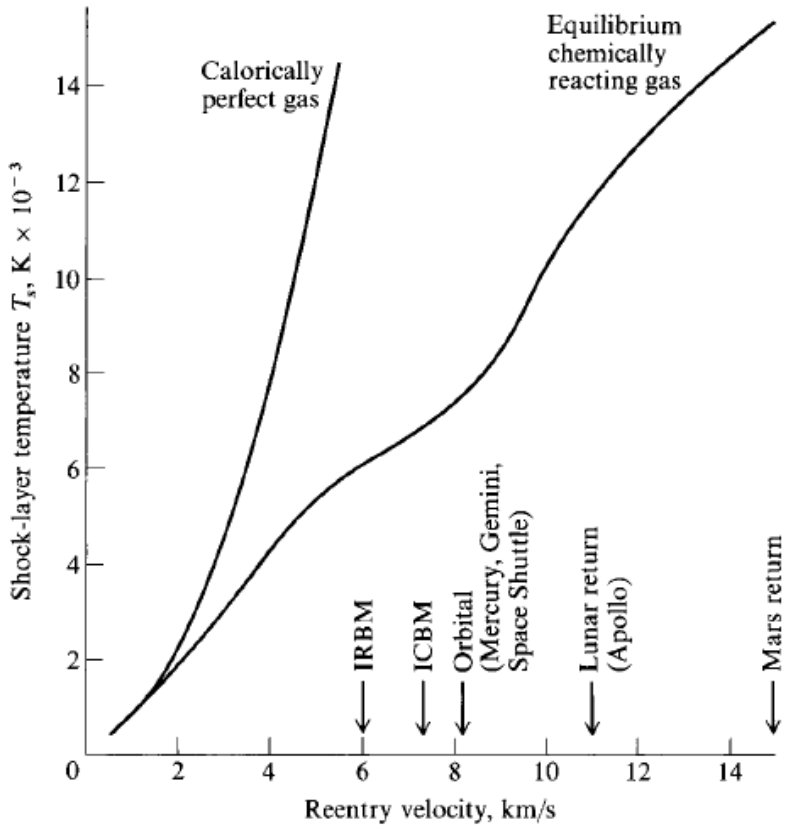


Figure 1.1 – Temperature behind a normal shock wave as a function of freestream velocity at a standard altitude of 52 km. Reproduced from ref. [1]

In addition to its chemical complexity, the hypersonic boundary layer undergoes significant viscous effects, especially in low density, high altitude flight regimes. The boundary layer thickness, δ, scales directly with freestream Mach number and inversely with local Reynolds number as seen below in equation 3. [1]

$$\delta \propto \frac{M_1^2}{\sqrt{\mathrm{Re}_x}} \tag{5}$$

This scaling effect leads to a larger shock layer which increases convective heating surface area, and alters shear stress distributions. This directly impacts surface heat flux, boundary layer stability, and ultimately TPS sizing and placement. Accurately predicting viscous interaction effects is a fundamental aspect of high temperature gas dynamics modeling for re-entry design. [1]

The stagnation point, typically the nose or leading edge of the vehicle, is where the airflow decelerates to zero relative velocity and all kinetic energy is converted into internal energy. The resulting stagnation enthalpy is the most critical driver of thermal load, as it represents the maximum energy per unit mass available for transfer to the surface. This stagnation enthalpy (h_0) can be approximated via equation 4 below. [2]

$$h_0 = \frac{V^2}{2} + c_p T \tag{6}$$

At re-entry velocities of 7500 m/s the kinetic energy alone contributes over 28 MJ/kg, dominating over the freestream thermal energy and resulting in stagnation temperatures of 2000–3000°C even without accounting for dissociation losses.To quantify the thermal loading, a widely accepted model is the Sutton-Graves stagnation point convective heat flux equation [2], which estimates the heat flux, q, to a blunt body during atmospheric entry as shown in equation 2. This relationship illustrates the extreme sensitivity of heat flux to velocity (V^3), as well as the inverse dependence on nose radius ($R_n^{-0.5}$), justifying the use of blunt geometries in re-entry vehicle design. For a typical capsule with R_n = 0.5m, re-entering at V_e = 7500m/s, and ρ_e = 0.02kg/m^3 (ISA for approx. 60 km), the resulting convective heat flux exceeds 800kW/m^2. This magnitude of heating is consistent with the 600–1000 kW/m² range cited in modern simulations of hypersonic re-entry, confirming the necessity of robust thermal protection systems [2]. Figure 1.2 illustrates heat flux in W/cm^2 as a function of velocity in m/s.

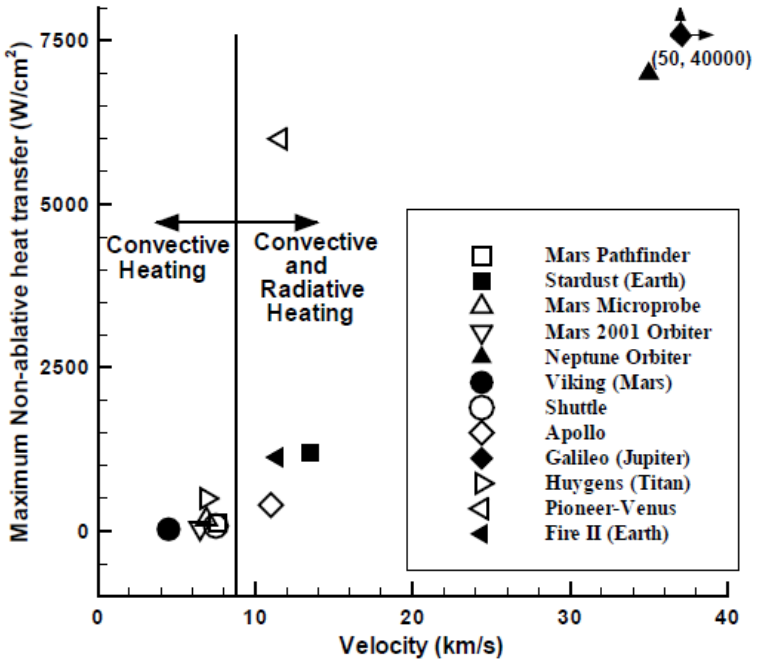


Figure 1.2 – Maximum non-ablative heat transfer wrt. Velocity for past and future NASA missions. Reproduced from ref. [3].

Furthermore, the heating environment is not static. The vehicle descends through varying layers of the atmosphere, where density increases exponentially, magnifying heat transfer despite decreasing velocity. The heating profile is often transient and peak-limited, with maximum heat flux occurring over a duration of approximately 200–400 seconds, depending on trajectory and vehicle orientation [4]. While early phases are dominated by high velocities and low density, mid-phase re-entry (30–60 km altitude) is where thermal loads peak, due to the specific combination of density and remaining kinetic energy. At these altitudes, the flow is generally in a thermochemical nonequilibrium, meaning that the air's molecular energy modes, translational, rotational, vibrational, are not in equilibrium. Dissociation and recombination reactions lag behind thermal and mechanical relaxation. This nonequilibrium condition reinforces the design requirement for TPS capable of withstanding not only high temperatures but chemically aggressive environments.

In addition to convective heating, radiative heating can become significant—especially for high enthalpy entries from lunar or interplanetary return missions where shock layer radiation can exceed 100 kW/m². This can be seen in figure 1.2 where the radiative healing begins to take effect above ~8 km/s. However, for most low Earth orbit (LEO) returns, convective mechanisms remain dominant (the simulation done in this study will assume a LEO return).

Lastly, aerodynamic heating is highly localized. Leading edges, control surfaces, and protrusions often experience much higher heat fluxes than nominal body surfaces. For instance, sharp fins or control flaps may concentrate heat fluxes 2–3 times higher than the stagnation point, requiring additional localized reinforcement.

1.2.2 Thermal Loads and Material Constraints

The thermal loads encountered by hypersonic re-entry vehicles are not only severe, they are complex in distribution as well as time evolution. It is vital to any study revolving around re-entry to examine how these extreme conditions can translate into internal vehicle constraints and material performance limitations. The role of TPS arises from how high heat fluxes interact with structural materials over time.

At high flux regions like those aforementioned in section A. surface temperatures during peak re-entry heating may reach between 1600°C and 3000°C. This will depend on several factors like vehicle geometry, angle of attack, and trajectory steepness. These temperatures have been directly observed and validated across legacy missions like the Apollo crew capsules as well as modern Earth return vehicles such as Artemis and Starship. For instance the reinforced carbon-carbon (RCC) panels on the Space Shuttle nose cap were subjected to surface temperatures exceeding 1650°C during re-entry, with no capacity for error in thermal response. Sharp leading-edge components such as fins may encounter even more localized heating and experience higher temperatures due to geometrical amplification. [4]

These surface temperatures are driven primarily by convective heating via the post shock boundary layer. While radiative heat transfer does become important in certain high enthalpy flows such as lunar or interplanetary return trajectories, LEO re-entry is dominated by convective effects. The way this heat propagates internally depends on the thermal properties of the materials beneath the surface. Thus either way TPS structures must resist thermal inputs across multiple modes — and at the very least, convection and conduction. To approximate the internal distribution of heat, 1D transient conduction into the TPS can be represented by equation 5 below which is derived from Fourier's law.

$$\frac{\partial T}{\partial t} = \alpha \frac{\partial^2 T}{\partial x^2} \tag{7}$$

Where α is the thermal diffusivity represented by equation 6 below.

$$\alpha = \frac{k}{\rho c_p} \tag{8}$$

The thermal loading in the re-entry environment is also highly non-uniform. As stated earlier stagnation points experience the highest heat flux, nut transition regions, control surfaces, and structural discontinuities often experience complex localized heating effects. The result of this is a naturally spatially heterogeneous thermal environment where certain areas may require far more thermal protection than others.

High thermal gradients encountered during peak heating introduce substantial amounts of thermal stress. As the outer surface begins to rapidly heat inner layers remain cooler and expansion mismatches occur. This generates internal stress fields that can exceed the mechanical

limits of the materials involved. These internal stresses lead to layer delamination and thermal fatigue which can then in severe cases result in TPS failure which can lead to catastrophic failure of the entire system. This concern becomes more severe in reusable vehicles where repeated thermal cycling exacerbates microstructural degradation.

The duration of peak heating is a constraint that must be considered and defined. It has been observed in both numerical analysis and empirical studies that the high flux period of re-entry typically lasts 200–400s with total heating durations around 600s (or about 10 minutes) [4], [5]. However heating time can well exceed these values as illustrated below in figure 1.3.

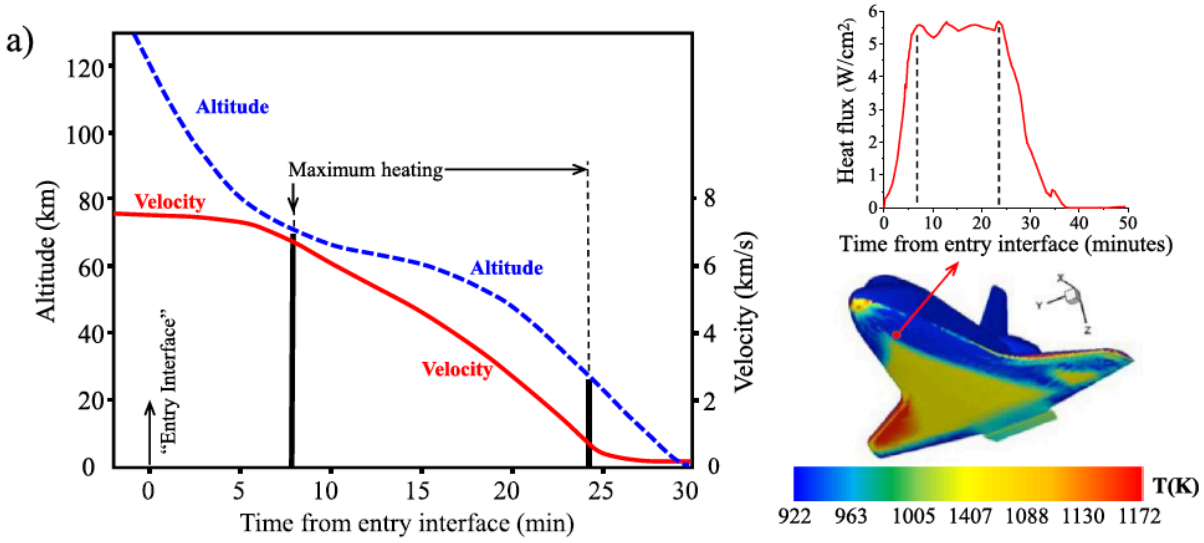


Figure 1.3 – Re-entry maximum heating interval in altitude vs. time plot for Space Shuttle program. Reproduced from ref. [6].

During this time the vehicle is receiving a prolonged thermal soak. Even materials that can momentarily survive temperatures of 2000°C may degrade or fail when exposed to sustained high temperatures over several minutes. Thus TPS must not only block heat but also delay and effectively distribute it to prevent a backface temperature rise.

1.2.3 Necessity of TPS and Historical Evolution

Most structural materials used in aerospace vehicles. Such as titanium and aluminum alloys, or carbon fiber composites, cannot withstand direct exposure to the re-entry environment. For example titanium alloys lose structural integrity above 600-700°C and aluminum melts around that temperature range. For prolonged heat exposure these alloys would degrade at much lower temperatures especially for multi-use and reusable vehicles. Even more advanced alloys would fail under sustained thermal gradients and oxidative stress encountered during re-entry. Without adequate and rigid TPS structures re-entry vehicles would encounter rapid structural

degradation and experience catastrophic failure within seconds. As such TPS is not a performance enhancer but an existential requirement.

The development and application of thermal protection systems has been widely recognized as a key enabler of space vehicle recovery. Early U.S. and Soviet capsule programs relied heavily on ablative heat shields composed of materials such as phenolic impregnated carbon ablator (PICA) and reinforced carbon-carbon composites (RCC). These materials protect the vehicle by deliberately degrading and carrying heat away by eroding and shedding material. These were used extensively on return vehicles in early pioneering missions such as the Apollo Command Module, Viking Mars Landers, and Mars Science Laboratory (MSL) [7],[8]. Ablatives like those aforementioned, according to former NASA USRP Intern Effe E. Greene are historically proven for TPS applications and classified at TRL (Technology Readiness Level) 9 with multiple successful flights and mission operations. However ablatives are in inherently single use which introduces major limitations for modern aerospace applications.

In a recent publication Piacquadio and colleagues note that ablative TPS carry significant mass penalties and require complete refurbishment after every mission due to irreversible material loss [9]. While these systems show clear reliability in the past, they do not meet the growing industry demand for reusability, rapid turnaround, and operational cost efficiency that define the next generation of hypersonic vehicles. These standards have been brought about largely by privatization or private company involvement in space (i.e. SpaceX, Blue Origin, Rocket Lab). NASA's own assessments, such as those published by Greene [7], emphasize that the logistics of repairing or replacing ablative shields after each flight impose significant burdens on mission scheduling and cost structure.

Furthermore, traditional ablative systems offer minimal structural support and must be backed by dedicated load-bearing substructures. This adds mechanical complexity and significantly increases system mass all the while unnecessarily introducing additional failure points. Moreover, predicting the behavior of ablatives remains a significant modeling challenge. Certain variables such as pyrolysis gas formation, material decomposition, and boundary-layer interactions all interact in nonlinear ways, making simulation difficult especially under transient or coupled flow conditions [9]. Thus making the behavior of ablative TPS highly unpredictable and subjective.

These vulnerabilities are not just theoretical, they have played out in real flights. The Space Shuttle Columbia disaster in 2003 is one of the most devastating TPS failures in history. During launch, a piece of foam insulation broke free from the external tank and hit the RCC leading edge of the left wing. This impact resulted in a breach that went undetected. During re-entry, plasma entered the wing structure, causing rapid unscheduled disassembly or catastrophic disintegration. This led to the loss of lives of all seven astronauts on board. The Columbia Accident Investigation Board (CAIB) identified TPS breach as the initiating failure, a clear example of how even a localized defect can compromise the entire vehicle [7].

Other missions have experienced critical TPS events during atmospheric entries. Although not an Earth entry the Galileo probe experienced complete ablation of its PICA heat

shield during its descent into Jupiter's atmosphere much earlier than expected [10]. Both the Viking and MSL entry systems experienced much higher than expected material erosion during their entry into the Martian atmosphere, causing NASA to require larger engineering safety margins. These examples reinforce the truth that even mature TPS technologies require constant redesign and adaptation to evolving flight conditions, mission architectures, and industry demands.

Today's programs and ambitions require vehicles that are reusable, low maintenance, and able to survive multiple re-entries with minimal intervention and refurbishment during turn around. In this context both ablative systems and legacy ceramic tiles like PICA and RCC respectively, fall short of meeting these standards. These challenges are clearly articulated in a 2020 publication by Uyanna and Najafi [4], who emphasize that mass efficiency, structural integration, and reusability are becoming a mandatory design criteria for TPS. Historically the reliance on ablatives and RCC tiles do not extend well to the future. Modern mission requirements highlight the need for more adaptive, multifunctional, and structurally integrated solutions which can

Solutions for reusability have been utilized and flight tested with good TRA level, specifically Toughened Uni-piece Fibrous Reinforced Oxidation-Resistant Composite (TUFROC) tiles which are an improved ceramic matrix composite (CMC) tile. These tiles originally developed by NASA Ames Research center were developed as a replacement for reinforced carbon-carbon (RCC) TPS [7] and have seen extensive use on launch vehicles such as Space X's Starship. Notably however, TUFROC and CMC tiles tend to detach during re-entry due to internal structural strain and can have significant gaps especially in spacecraft with movable joints and fins. Although they have shown proven reusability in more nominal flight conditions they also still limit turn-around time as they tend to need extensive maintenance before re-use. TUFROC may survive a few re-entries, but still shows erosion, spalling, bond degradation, and the same structural issues encountered by RCC.

1.2.4 Categorization of TPS Requirements

Developing a functional Thermal Protection System is not simply a matter of choosing a heat resistant material, rather it must satisfy a wide range of performance parameters including thermal, structural, and operational. As the aerospace industry grows its demand for vehicles that are increasingly complex, reusable, and dynamically operated, these requirements become more and more rigorous. Modern mission architectures require TPS to do far more than simply survive a single re-entry, they must integrate seamlessly with the overall system and ensure the vehicles long term operational reliability

The first fundamental requirement is of course thermal resistance. Any TPS must be capable of withstanding the peak surface temperatures which as previously stated can exceed well over 1600°C. The TPS must be able to endure sustained heat soak meaning it must withstand extreme temperatures for extended periods of time (5–10 minutes). Regardless of the mechanism of protection, the backface temperature of the vehicle must be respected to ensure

survival of the vehicle. If the backface is allowed to heat over its maximum temperature catastrophic failure likely in the form of rapid unscheduled disassembly or vehicle disintegration will follow in seconds.

Tied closely to thermal resistance is the need for structural support. TPS are subjected to intense structural loading under re-entry due to aerodynamic pressure, shockwave loading, and crucially internal stresses due to thermal expansion gradients. As noted by Piacquadio et al. [9], traditional systems offer minimal load–bearing capabilities and must be backed by heavy substructures. Ceramic tile based systems like RCC or TUFROC are prone to cracking, spalling, detachment and delamination under intense thermal stress. Ideally TPS must be able to maintain mechanical integrity and resist delamination, and interlayer failure in a highly complex and unpredictable thermal environment as well as in post-entry operations.

Modern mission frameworks also bring about a new requirement of rapid reusability. Legacy designs as aforementioned do not meet the cut. Both commercial and governmental spaceflight increasingly favor reusability, low refurbishment overhead, and quick turnaround times. Thus next generation TPS must exhibit minimum fatigue and performance degradation over several thermal cycles.

Because of these requirements there now arises a new requirement for bleeding edge TPS technology; this is adaptability or what can otherwise be called adaptive thermal protection systems (ATPS). As opposed to current technology which involves passive systems ATPS would in theory have the ability to dynamically adapt to the thermal environment of hypersonic re-entry. This would mean responding dynamically to spatial variability or localized heat flux and thermal gradients. Uyanna and Najafi [4] specifically highlight adaptability as a critical gap in existing TPS technologies and identify it as a principal target for innovation.

A lesser mentioned but equally important criteria for new age TPS systems is that they must support aerothermal prediction accuracy. In other words this is referring to how well the system can be modeled pre-flight as well as how reliably its behavior can be predicted under transient thermal-structural conditions. As shown in the computational work by Marley and Driscoll [5], high fidelity simulation of TPS performance requires accurate definitions of material behavior under multi-axis loading, complex geometries, and highly nonuniform boundary conditions.

Collectively these categorizations define the complex and multidimensional design space for which advanced TPS systems must excel in whereas traditional systems may have fallen short in several. There is a growing emphasis on developing systems that can converge these requirements into a single integrated solution. ATPS concepts based on smart and multifunctional materials aim to address this in one unified architecture. Section 1.2.6 will explore how these integrated and multifunctional systems are being pursued in hypersonic vehicle design.

1.2.5 Integrated/Multifunctional and Sandwich Structures for TPS

In response to the significant limitations of traditional ablative and passive TPS, the aerospace industry has increasingly emphasized the adoption of Integrated and Multifunctional Thermal Protection Systems (ITPS). ITPS offer the potential to combine structural load-bearing functions and dynamic thermal management capabilities, addressing both structural and thermal challenges in a unified manner.

Piacquadio et al. present a detailed comparison of various ITPS designs specifically targeted at reusable launch vehicles. Notably, their research extensively evaluates ceramic matrix composite (CMC) tiles and metallic sandwich structures with embedded phase-change materials (PCMs). The authors highlight PCMs embedded within metallic lattice cores produced through additive manufacturing, citing their high thermal energy storage capability, reduced thermal gradients, and resultant decreased thermal stresses. This approach significantly improves the structural integrity and extends the operational lifespan of TPS components under repeated thermal cycling, directly aligning with industry demands for rapid reusability and reduced maintenance [9].

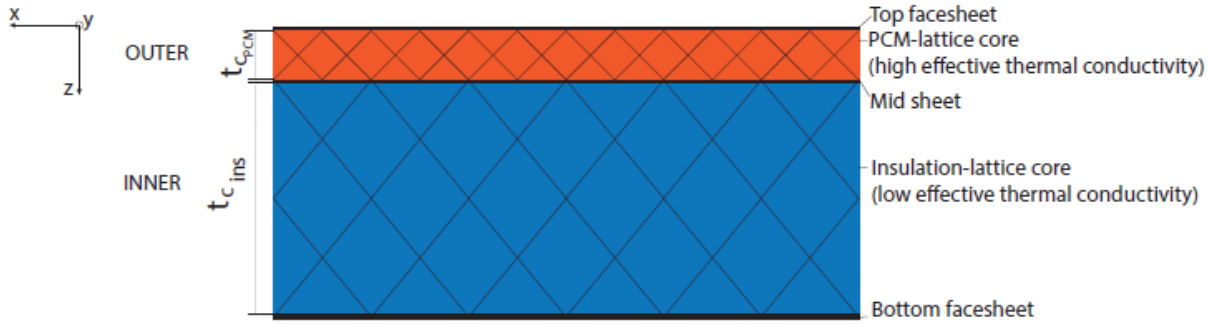


Figure 1.4 – Hierarchical sandwich TPS concept with embedded PCM in outer core and insulation in inner core. Reproduced from ref. [9].

Advanced sandwich structures reviewed by Le et al. have gained considerable attention due to their multifunctional capabilities. These structures typically consist of multilayered cores designed with ceramic or metallic materials, optimized for both mechanical strength and thermal insulation. Le et al. systematically outline how variations in core geometries, materials selection, and fabrication processes directly influence the structural and thermal performance. The authors underscore the promise of CMC sandwich structures for their excellent high temperature stability, lower thermal conductivity, and superior strength-to-weight ratios compared to traditional metallic counterparts. Despite these advantages, the authors also emphasize critical challenges such as precise manufacturing, potential delamination under extreme conditions, and the complexity inherent in achieving robust thermo-structural coupling [6].

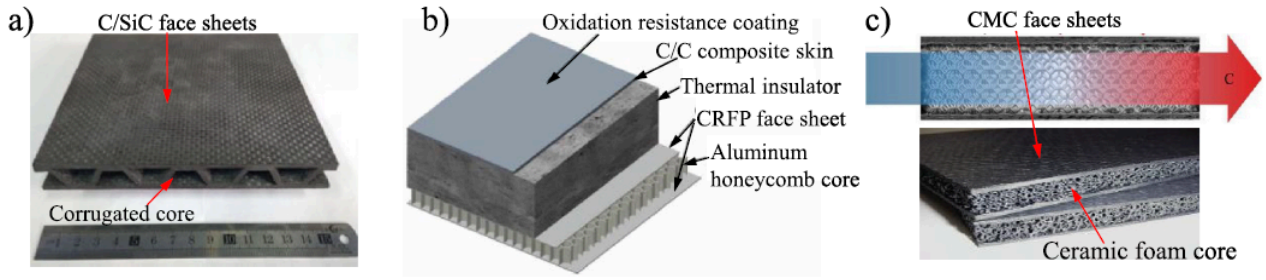


Figure 1.5 Sandwich TPS architectures: (a) Corrugated core with C/SiC, (b) Laminated composite with carbon foam and CFRP-aluminum honeycomb, (c) Lattice foam core made of CMC/SiC. Reproduced from ref. [6].

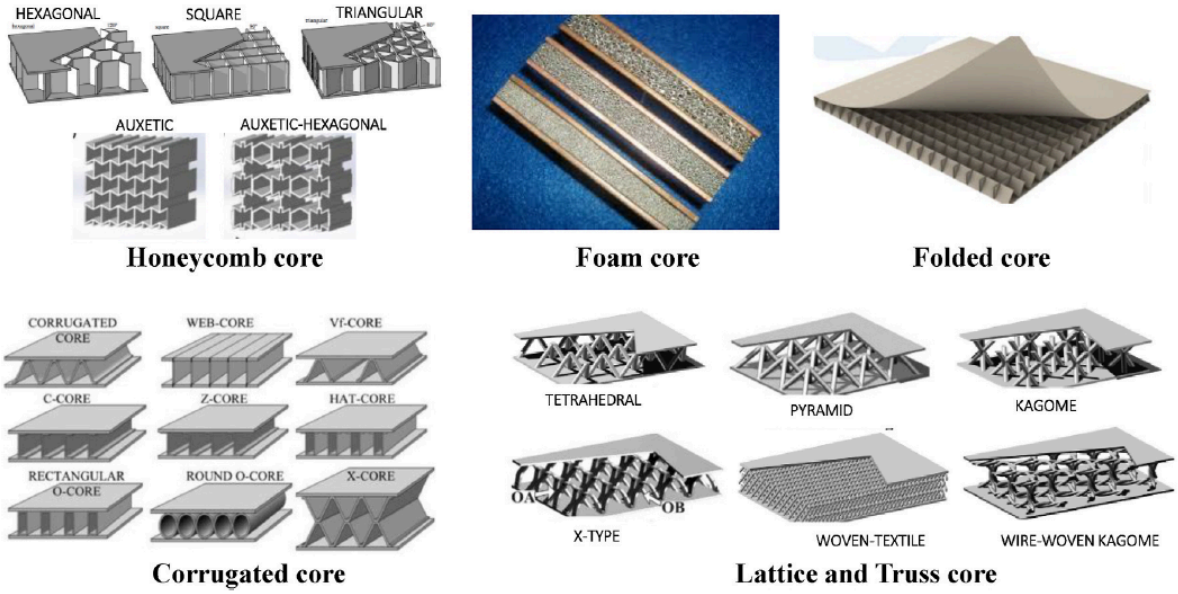


Figure 1.6 – Various sandwich structure core designs. Reproduced from ref. [6].

Recent advancements in the design of sandwich-structured thermal protection systems (TPS) further validate the integration of structurally adaptive and thermally responsive materials. Zhang J. et al. explored hybrid corrugated sandwich plate configurations designed explicitly for hypersonic applications, demonstrating significant improvements in mechanical stability and thermal insulation under coupled fluid-structure-thermal loads. Their work confirmed that corrugated core structures effectively reduced thermal deformation due to inherent geometric stiffness enhancements, while simultaneously managing thermal gradients through their multi-layered designs [11]. This finding strongly supports the rationale behind selecting structurally integrated smart-material layers, such as those proposed in this study, and underscores the importance of rigorous computational analysis of structural deformation and fluid-thermal coupling effects.

Both of the aforementioned studies collectively report the advancements ITPS have over legacy TPS. However, they also clearly identify ongoing technical hurdles with these proposed systems. Challenges in the encapsulation of PCMs, accurate predictive modeling of thermo-mechanical behaviors, and advanced manufacturing techniques required for complex sandwich structures remain an issue. Addressing these challenges through continued research and innovation is crucial to realizing the full potential of ITPS.

1.2.6 Advanced Computational Techniques in TPS Analysis

Developing Adaptive Thermal Protection Systems (ATPS) for hypersonic vehicles require the use of high fidelity computational modeling. Traditional analytical methods and simplified computational approaches fail to capture the highly complex interactions between aerodynamic heating, structural deformation, and material response under realistic flight conditions. In response to this need advanced computational methods, particularly coupled fluid-thermal-structural analyses and reduced-order modeling techniques, have emerged as important tools.

Computational accuracy and stability are critical for modeling adaptive TPS behaviors under extreme hypersonic conditions. Zhang J. et al. [12] introduced an implicit coupling framework between hypersonic nonequilibrium flows and material thermal response simulations. This method substantially improves numerical stability, mitigating issues common with explicit methods such as numerical oscillations and discretization errors at interfaces. Although their research primarily addresses ablative materials, their computational strategies directly inform the robust coupling and boundary condition methodologies applied in this thesis. Adopting such techniques is essential for accurately capturing dynamic smart material responses in computational simulations under realistic hypersonic thermal and structural loading conditions.

Marley and Driscoll present a framework for the computational analysis of both active and passive thermal protection systems using reduced-order modeling techniques [5]. Their approach divides the TPS structure into multiple 1D thermal zones and solves transient heat conduction through each region based on external heating condition inputs. Convective heat fluxes are estimated using relations derived from boundary layer theory, while radiative heat transfer is modeled using the Stefan–Boltzmann law. For active systems, internal heat removal is simulated using coolant flow modeling through internal channels. This type of modular framework allows for efficient simulation of multilayered systems, including the combined effect of outer insulation and internal active cooling. Marley and Driscoll's work proves how reduced-order models can enable faster but accurate prediction of wall temperature evolution with significantly less computational cost.

Addressing the significant computational resources typically required for full-scale simulations, Vargas Venegas et al. propose an innovative Physics-Infused Reduced-Order Modeling (PIROM) approach specifically tailored for multi-layered hypersonic TPS [13]. Their work combines Lumped-Capacitance Models (LCM), with data-driven correction methods based on the Mori-Zwanzig formalism. This type of hybrid modeling more efficiently simulates critical

physical and thermal phenomena, enabling significant computational savings without sacrificing accuracy. The PIROM approach, notably, demonstrated prediction accuracy within approximately 1% of traditional high fidelity simulations, while still providing a computational speedup of roughly two orders of magnitude. The PIROM approach allows for generalization beyond the specific material properties used during modeling. This is useful for hypersonic TPS design, where material property uncertainties can impact thermal response heavily. Figure 1.7 illustrates PIROM's prediction accuracy under nominal and perturbed material conditions compared with full ordered FEM modeling.

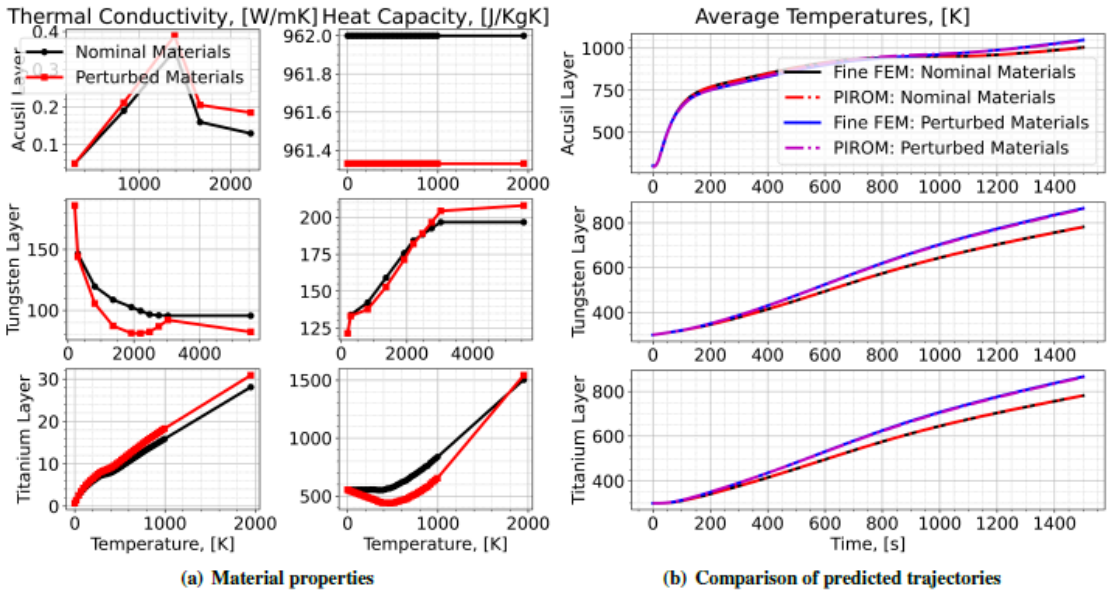
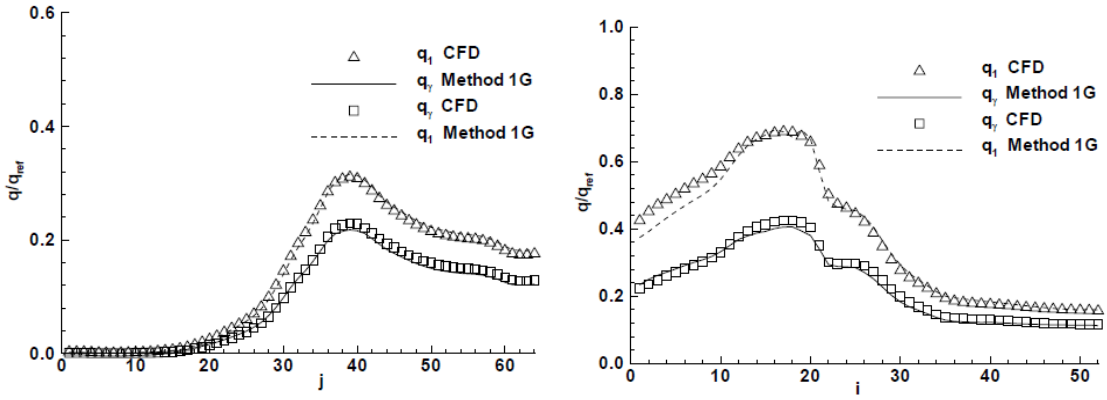


Figure 1.7 – PIROM performance under material property perturbations. (a) Nominal and perturbed material properties; (b) Predicted trajectory comparisons for PIROM vs Fine FEM for both cases. Reproduced from ref. [13].

This substantial improvement in efficiency would enables rapid exploration of this concept in larger design spaces and allows for performing extensive parametric and optimization studies.

The techniques discussed in this section represent essential capabilities in contemporary TPS research, offering a balance between computational accuracy and efficiency. While legacy systems relied on extrapolated data from simplified models such as the archaic Method 1G modern CFD methods offer far greater fidelity and accuracy. The difference between these methods is highlighted in the work of Gnoffo et al. and is illustrated in figure 1.8 below.



a) $i = 40$, circumferential distribution

b) $j = 64$, centerline distribution

Figure 1.8 – Comparison of Method 1 G and CFD for modeling heating in hypersonic vehicles. Reproduced from ref. [3].

The ability to precisely predict coupled thermal and structural behavior and rapidly evaluate diverse configurations is critical for the successful implementation of ATPS and will allow for accurate modeling of smart material based thermal protection systems systems.

1.2.7 Definition and Classification of Smart Materials

The term smart materials refers to a class of materials that exhibit adaptive or responsive behaviors when subjected to external stimuli. These are unlike conventional materials which maintain fixed physical properties under operating conditions. The term smart materials generally encompasses materials that can dynamically alter their internal structure, mechanical state, or surface characteristics in response to inputs such as temperature, stress, magnetic field, electric potential, or light intensity. These transitions are typically reversible and provide opportunities for active control, real-time adaptation, and multifunctional operation within structural systems. [14]

Smart materials can be broadly grouped into two functional categories: passive and active systems. Though they are capable of environmental response, passive smart materials do not enable actuation or direct energy transfer. These include materials such as thermochromics or optical fibers. Active smart materials generate force, displacement, or signal output through controlled energy conversion. These include shape memory alloys (SMAs), piezoelectric ceramics, magnetostrictive compounds, and electroactive polymers among several others. [14]

Classification schemes for smart materials are typically based on the stimulus response pair and the underlying physical mechanism. Some of the most commonly studied and applied categories include those in the list below [14]:

- **Shape Memory Materials (SMMs):** undergo phase transformation that enables large recoverable strain and shape recovery. Includes:

    - **Shape Memory Alloys (SMAs)**: metallic, e.g., Nitinol, Cu-Al-Ni.
    - **Shape Memory Polymers (SMPs)**: thermally triggered, lighter and cheaper than SMAs.
    - **Shape Memory Ceramics (SMCs)**: high temp capable but brittle.
- **Piezoelectric Materials:** convert mechanical strain into electrical charge and vice versa. Useful in aerospace sensing and actuation systems.
- **Magnetostrictive Materials:** deform in response to magnetic fields; capable of high force actuation, typically used in vibration damping and precision control.
- **Electrostrictive Materials:** deform in electric fields via dipole realignment; applicable in precision lenses and soft robotics.
- **Chromic Materials (Electrochromic, Photochromic):** change optical transparency or color in response to voltage or light intensity; useful in energy-efficient windows and avionics displays.
- **Rheological Fluids (ERFs/MRFs):** smart fluids whose viscosity changes under electric or magnetic fields. Widely used in adaptive damping and isolators.
- **Thermoelectric Materials:** convert thermal gradients into electrical potential. Applied in energy harvesting and thermal sensing.
- **Hydrogels:** swell or shrink in response to temperature, pH, or moisture. More relevant in biomedical and microfluidic systems.

A complete classification of these material categories along with their stimulus/response mechanisms is shown in table 1.1 below adapted from Chen et al. [14].

Table 1.1 – Smart material types with respective inputs and outputs. Reproduced from ref. [14].

Smart material types with respective inputs and outputs.

Material	Input	Output
Piezoelectric	Electric field Mechanical load	Mechanical strain Electric potential
Photostrictive	Incident light	Mechanical strain
Magnetostrictive	Magnetic field Mechanical load	Mechanical strain Magnetization
Electrostrictive	Electric field Mechanical load	Mechanical strain Electric potential
Thermoelectric	Thermal load Electric field	Electric potential Temperature change
Shape memory materials	Thermal load Magnetic load	Mechanical strain
Photovoltaic	Incident light	Electric potential
Magneto- Electro- } rheological fluids	Magnetic field Electric field	Mechanical strain
Photo- Thermo- Magneto- Electro- } chromic	Incident light Thermal load Magnetic field Electric field Mechanical load	Color change

Table 1.1 shows an overview of input-output relationships across the full range of smart material types and can be useful in determining suitability for thermal protection systems for hypersonic re-entry vehicles.

1.2.8 Smart Materials and Multifunction Material Applications in Aerospace

In recent decades smart materials have emerged as a dynamically adaptive solution to modern aerospace structural problems. Materials classified as "smart" are typically characterized by their ability to respond to external and environmental stimuli such as temperature, stress, and electric or magnetic fields in real time. They represent a promising approach in developing adaptive thermal protection systems (ATPS) for hypersonic re-entry vehicles. Smart materials differ from passive materials in that they exhibit controlled changes in their physical properties in real-time. This could potentially allow them to actively mitigate adverse environmental conditions experienced during atmospheric entry.

Advancements in smart materials have portrayed their significant potential for adaptive aerospace applications. Lyujia Gu provides an extensive review, highlighting various smart material classes such as shape memory alloys (SMAs), piezoelectric ceramics, and

magnetostrictive materials, and explores their unique capability to adapt structural and thermal responses actively [15]. For example, SMAs, characterized by their ability to recover original shapes after deformation through heating or stress changes, offer considerable promise for adaptive TPS applications by enabling real-time structural adjustments to changing aerodynamic heating and mechanical loads. Piezoelectric materials similarly exhibit potential through their inherent capacity to generate electrical charges in response to mechanical stress, enabling integrated sensing and active thermal management within structural elements. Gu emphasizes that the integration of these responsive capabilities significantly enhances the adaptability and reliability of aerospace structures, particularly under unpredictable and severe hypersonic conditions [15].

Shape memory alloys are particularly promising for their potential use as an adaptive structural reinforcement against cracking, delamination, and damage due to thermal expansion in heat shields. The viability of shape memory alloys, particularly NiTi, or nitinol, as adaptive structural components in thermal protection systems is strongly supported by Farber et al. (2020). Their comprehensive review on additive manufacturing of NiTi alloys demonstrates the ability to precisely tailor transformation temperatures and shape-memory behaviors through controlled fabrication techniques. This capability is essential for realizing predictable adaptive responses during hypersonic thermal cycling. The insights provided by Farber et al. reinforce the feasibility of computationally modeling these adaptive responses, validating the selection of NiTi alloys as primary candidates for this thesis's adaptive TPS configurations [16].

While the majority of research on shape memory alloys (SMAs) focuses on systems that permit significant geometric actuation, recent studies have demonstrated that constrained SMAs can still exert beneficial mechanical and thermal effects when embedded within multilayer composite structures. In particular, constrained SMA layers, when pre-strained and thermally activated, generate internal recovery stress without undergoing visible macroscopic deformation. This effect is of particular relevance to multi-layered TPS configurations, where SMA layers may be mechanically sandwiched between various other materials.

In his 2024 dissertation, author Zhang Q. of Purdue University provides a detailed analysis of this behavior, noting that restoring stress and modulus changes due to phase transformation in thermally activated SMAs contribute significantly to the stiffening and dynamic response of hybrid composite systems. The study confirms that even when physical shape change is limited, the internal stress generated by the SMA enhances interlayer coupling, improves structural integrity under thermal gradients, and supports adaptive damping. [17]

This internal stress-driven interaction is particularly valuable in the ATPS context, where precise mechanical movement is undesirable due to packaging constraints and system complexity. Instead, thermally responsive stiffness modulation via SMA preloading offers a passive and highly reliable means of adaptively enhancing the mechanical coherence of the TPS stack. Such findings validate the design strategy adopted in later sections, in which Nitinol layers are embedded beneath surface tiles to reinforce thermal interfaces and dynamically manage internal stress distributions during flight.

Further practical exploration into smart material applications has been presented by Mehaan Mehra, who discusses specific aerospace applications and underscores the potential of smart materials in real world aerospace contexts. Mehra identifies SMAs, phase-change materials (PCMs), and self-healing composite materials as promising candidates for adaptive aerospace structures [18]. PCMs exhibit profound capabilities in absorbing and releasing thermal energy during phase transitions, enabling effective thermal regulation within aerospace structures subjected to extreme temperature fluctuations. Self-healing composites, which autonomously repair microcracks induced by thermal and mechanical stresses, offer considerable improvement in reliability and structural integrity, significantly reducing maintenance efforts and extending operational lifetimes. This property of self-healing composites could support rapid-reusability needs. Mehra concludes that the effective integration of these smart materials into aerospace structures could provide considerable benefits in terms of structural efficiency, safety, and rapid reusability—aligning closely with contemporary aerospace mission demands [18].

In the context of PCM integration, Yang et al. investigated a composite PCM embedded within active cooling channel walls tailored for hypersonic conditions. Their simulations quantified the PCM's capability to actively mitigate transient thermal loads, effectively reducing peak temperatures and thermal gradients across structural layers. This research provides concrete computational evidence of PCM effectiveness, directly validating PCM as candidates for the multilayered TPS configurations proposed in this thesis. Furthermore, their detailed methodology for thermal analysis under realistic boundary conditions aligns well with the simulation approach adopted herein, strengthening the credibility of expected thermal management benefits [19].

However, significant challenges remain for the widespread implementation of smart materials in aerospace applications. Gu specifically addresses critical hurdles such as ensuring long term material reliability under repeated thermal cycling, reliable actuation mechanisms, and the complexity of integration of these materials into existing aerospace structural systems [15]. Additionally, Mehra highlights manufacturing difficulties, high initial costs, and strict performance validation requirements that currently hinder broader application [18].

Despite these challenges, the ongoing research and continued development of smart materials clearly demonstrate their transformative potential in aerospace engineering. The capability to dynamically adapt structural and thermal responses in real-time positions smart material-based TPS as a critical technological advancement. Smart materials overall could fundamentally reshape the landscape of thermal protection systems and potentially enable moving toward more adaptive, multifunctional, and rapidly reusable re-entry vehicles.

1.2.9 Thermomechanical Mechanisms of Shape Memory Alloys

Smart material responses can include changes in shape, stiffness, stress state, or energy dissipation. The fundamental mechanisms that govern these behaviors are dependent on material type. They are generally reversible internal transformations at the molecular level or in crystalline structure. In the context of thermal protection systems smart materials that exhibit thermomechanical responsiveness are of the most immediate reliance. Shape Memory Alloys

(SMAs) are well suited for integration into ATPS due to their ability to undergo solid state phase transformations under combined thermal and mechanical loading [14].

SMA behavior, at its core, is characterized by reversible transformation between what are known as the martensite and austensite phases. These phases are reversible solid state phases and are essentially referring to differing types of crystal structures or atomic arrangement. At low temperatures SMAs exist in the martensitic phase, which is a relatively soft, low symmetry crystal structure — typically monoclinic or orthorhombic depending on the alloy. Within the martensitic phase the internal structure forms twinned variants, which can be rearranged into detwinned martensite when mechanical stress is applied. This change occurs without the need for diffusion and allows for large recoverable inelastic strains. When the material is subjected to heating above a critical temperature, the internal structure shifts into the austenitic phase which has a more ordered and higher symmetry configuration like a body centered cubic. This transformation is also diffusionless, highly reversible and results in the recovery of the materials original shape and mechanical stiffness. This transformation is temperature dependent and stress driven and allows for SMAs to function as sort of thermomechanical actuators, dampers, and adaptive structural reinforcements. The total martensite volume fraction is often represented as seen below in equation 6.

$$\xi = \xi_T + \xi_S^+ + \xi_S^- \tag{9}$$

This relationship is fundamental to predicting the transformation path of SMA materials under constrained, nonuniform heating conditions. The transitions described are also known as the shape memory effect and are graphically represented in the stress-temperature phase diagram shown in figure 1.9 below adapted from Zhang Q. [17].

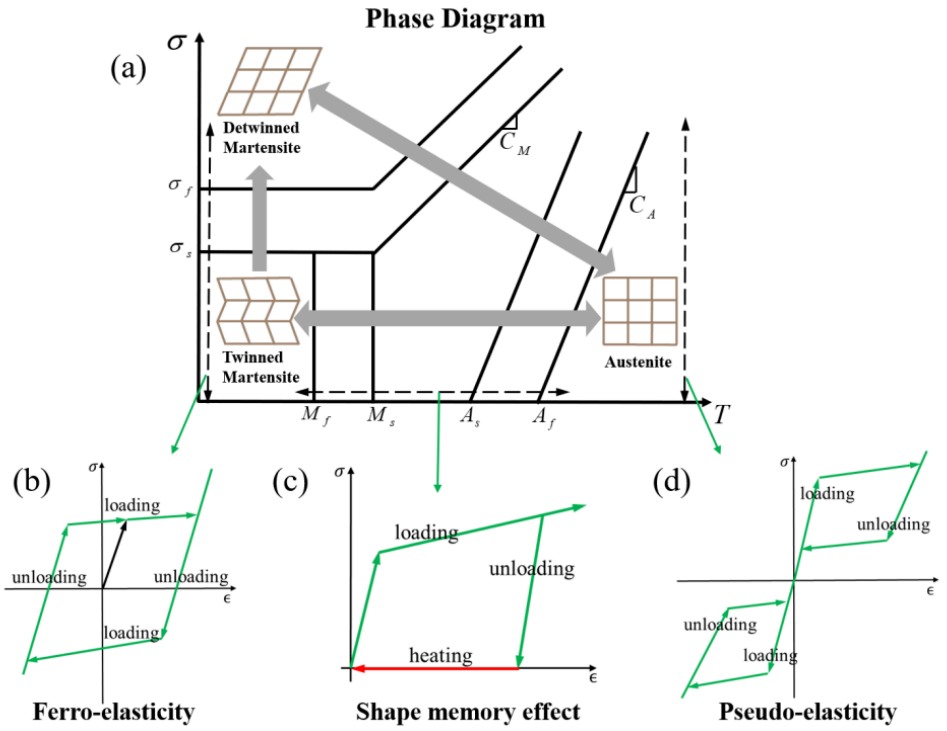


Figure 1.9 – Stress–temperature phase diagram of shape memory alloys showing phase boundaries for martensitic and austenitic transformation. Reproduced from ref. [17].

As shown in figure 1.9 diagram (d) SMAs also exhibit pseudoelasticity or superelastic behavior when loaded adobe the austenite temperature. In this regime they can undergo larger strains and recover them once unloaded due to the reversion of detwinned martensite. The energy absorbed during this cycle is dissipated through hysteresis which makes SMAS effective as dampers. The SME cycle is shown in Figure 1.10 and illustrates the transition from a deformed martensitic state to a recovered austenitic shape upon heating, as well as the return to martensite upon cooling [17].

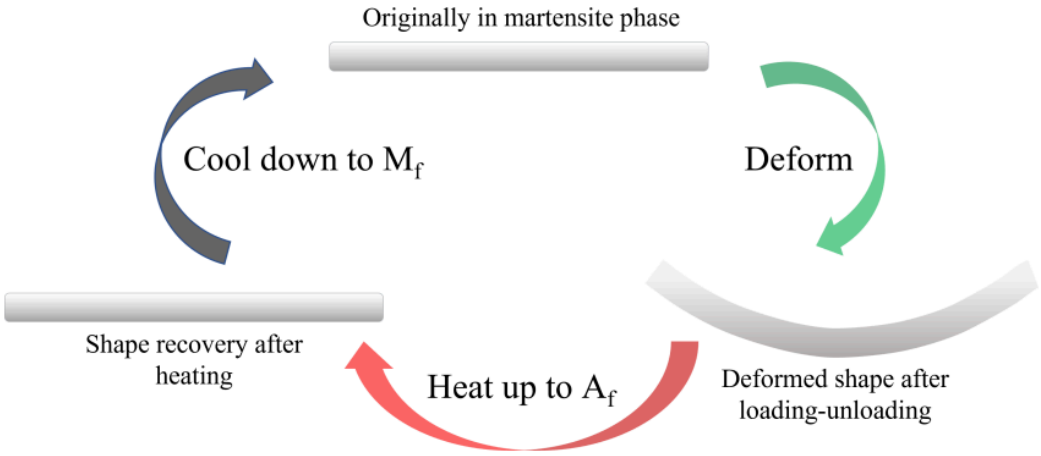


Figure 1.10 – Shape memory cycle of an SMA element under heating and cooling. Reproduced from ref. [17].

The phase transformation results in measurable macroscopic mechanical behavior, including modulus change, internal stress development, and strain recovery. When SMA layers are spatially confined, such as inside of a composite structure, the transformation to austenite generates significant restoring stress [17]. This transformation induced stress can be described by the mathematical model shown in equation 10 which accounts for both thermal and mechanical inputs.

$$\sigma = \sigma_0 + E(\xi)\varepsilon - E(\xi_0)\varepsilon_0 + \left[\Omega_S^+(\xi)\xi_S^+ + \Omega_S^-(\xi)\xi_S^-\right] - \left[\Omega_S^+(\xi_0)\xi_{S0}^+ + \Omega_S^-(\xi_0)\xi_{S0}^-\right] + \Theta(T - T_0) \tag{10}$$

Where:

- σ is the Total stress
- σ_0 is the Reference stress
- ε is the Total strain
- ε_0 is the Reference strain
- ξ is the Martensite volume fraction
- ξ_S^+, ξ_S^- are the Tension/compression induced martensite volume fractions
- Ω is the Martensitic transformation coefficient
- Ω_S^+, Ω_S^- are the Martensitic transformation stress coefficients
- Θ is the Thermal expansion modulus
- $E(\xi)$ is Young's Modulus as a function of martensite content

This effect can potentially improve stiffness, mechanical resistance, and very notably limit interfacial delamination during thermal loading cycles rendering them very useful in TPS architectures. As the TPS material heats during peak heating, localized heat flux can cause uneven material expansion leading to delamination and even detachment. An intermediate SMA layer could respond to this by applying internal forces as they transition into the austenitic phase in these higher temperature areas. This would allow the TPS to actively mitigate delamination in critical areas and help to close gaps in the structure avoiding detachment of TPS tiles or panels.

Ferroelastic behavior shown in figure 1.9 diagram (b) dominates SMA response in the low temperature martensitic phase. In this regime, SMAs can undergo large recoverable strain through twin boundary motion without full phase transformation. This ferroelastic reorientation results in high energy dissipation during cyclic loading and is particularly effective at suppressing vibration and structural resonance. On the other hand, in the pseudoelastic regime, the SMA returns to austenite upon unloading, with reduced hysteresis and higher stiffness [17]. The ability to tune between these regimes by controlling local temperature give promising opportunities in TPS design for adaptive stiffness modulation and thermal load redistribution.

The characteristic stress–strain loops for ferroelastic and pseudoelastic cycling show strong hysteresis, which directly translates to internal damping capacity. This behavior is advantageous in the hypersonic context, where transient thermal and structural loads can induce rapid oscillations, shock loading, or vibrational instability. Embedded SMA layers can absorb and dissipate this energy without requiring active control or external sensors.

SMAs can also be programmed to exhibit one-way or two-way shape memory effects. One-way SME, the most common, allows the material to return to a fixed shape upon heating but does not retain memory of the deformed shape upon cooling. Two-way SME, achievable through thermomechanical training, allows reversible shape change between two geometries during heating and cooling cycles. While two-way SME is less common in aerospace-grade alloys due to complexity, it holds potential for morphing structures or passive deployment systems in environments where temperature gradients are cyclic and predictable [14].

In aerospace thermal protection, the constrained SMA configuration is most applicable. As discussed in Zhang Q's detailed investigation of SMA hybrid composites, embedding SMA layers into sandwich structures enables phase-transformation-induced stiffness enhancement during heating and stress recovery during cooling [17]. This cyclical mechanical behavior could allow a TPS structure to adapt passively to each phase of flight, improving both structural integrity and thermal resistance.

Although other smart materials—such as piezoelectrics or magnetostrictives—are capable of dynamic response, their thermal limits, brittleness, and integration complexity limit their viability in high temperature TPS systems. SMAs, by contrast, have demonstrated functional reliability in high cycle, high temperature, and mechanically constrained environments. Gu [15] and Mehra [18] both identify SMAs as a leading candidate for adaptive TPS due to their reusability, passive control capability, and strong thermomechanical coupling behavior.

The physical mechanisms of SMAs therefore align directly with the demands of hypersonic thermal protection. These characteristics underscore the adaptive layering strategies used throughout this thesis and directly inform the multi-layer sandwich architecture developed in later sections.

1.2.10 Other Smart Material Categories — An Overview

While shape memory alloys represent the most directly applicable class of smart materials for hypersonic thermal protection, several other categories of smart materials have been explored in aerospace engineering. Each offers distinct advantages based on its underlying physical mechanism, but most have limitations that constrain their direct use in thermal protection systems. Nevertheless, some of these materials may serve complementary roles, either as auxiliary subsystems, embedded sensing layers, or adaptive components within an integrated TPS architecture.

Piezoelectric materials are among the most mature and widely applied smart materials in aerospace, particularly in sensing and actuation applications. These materials generate an electric charge when subjected to mechanical stress (direct effect) and conversely deform when an electric field is applied (inverse effect). This makes them highly effective for structural health monitoring, vibration control, and precision actuation. However, many high performance piezoelectrics, such as lead zirconate titanate (PZT), are ceramic in nature and thus exhibit brittleness, low fracture toughness, and limited thermal resilience [14], making them less suitable for high temperature, mechanically dynamic environments like re-entry TPS. Their role in thermal protection is therefore more aligned with integrated sensing, not load-bearing or adaptive heat management.

Magnetostrictive materials operate through the reorientation of magnetic domains in response to external magnetic fields, resulting in dimensional changes that can be harnessed for actuation. Terfenol-D is a well-known magnetostrictive material used in energy harvesters and high force actuators. These materials offer high output force and fast response, but tend to be dense, difficult to integrate at scale, and are highly sensitive to eddy current losses at high frequencies. Additionally, their operational bandwidth is narrow under the combined thermal and vibrational loads typical of hypersonic flight, limiting their TPS relevance [14],[12].

Shape memory polymers (SMPs) share some functional similarities with SMAs but rely on temperature-induced transitions in polymer chain structure to achieve shape memory behavior. They are often classified as thermoresponsive materials and can exhibit large deformations and low density. SMPs have found use in deployable space structures, self-folding mechanisms, and adaptive aerodynamic surfaces. However, they generally offer lower mechanical strength, reduced thermal stability, and limited actuation force compared to SMAs. Their application in TPS is mostly confined to non-load-bearing deployment systems or low heat flux regions, where their geometric adaptability and low mass may still offer advantages [14],[15].

Thermoelectric and chromic materials operate by converting thermal gradients into electrical potential (thermoelectrics) or by modulating optical properties in response to heat or voltage (chromic materials). Thermoelectric materials have potential utility in energy harvesting or active cooling subsystems by converting otherwise wasted heat into usable electrical energy. Chromic materials—such as electrochromic and thermochromic coatings—have seen use in spacecraft windows and radiation control surfaces. However, both are typically limited to passive or secondary functions, as they lack structural capacity and cannot mitigate the intense heating rates present during atmospheric re-entry. [14]

Phase-change materials (PCMs) are a type of multi-functional material that may serve auxiliary roles within a layered TPS. PCMs store latent heat during melting, buffering downstream structures from thermal spikes. While not "smart" in the strict sense, they are often integrated into advanced TPS systems for their passive heat absorption properties. Encapsulated PCMs are promising for TPS use, especially in the form of sandwiched structures and multi layered integrated structures described in detail in section 1.2.5 and by Le et al. [6].

Self-healing composites offer the ability to autonomously repair microcracks or delamination via thermally activated bonding agents or microencapsulated healing chemistries. As noted by Mehra [18], such materials can extend the lifetime and reliability of reusable TPS tiles or panels, especially in modular systems where inspection and replacement are time-consuming. However these materials are extremely difficult to model and highly unpredictable in behavior.

On their own, smart materials are not capable of fulfilling all of the core requirements of ATPS as most of them would not survive direct surface exposure to temperatures exceeding 1600°C. Rather these systems will be investigated as intermediate layers. TPS architectures may utilize hybrid systems where smart, multifunctional, and traditional materials when combined into a layered structure can potentially support areas where other layers would fail on their own, helping to meet the fundamental requirements for TPS.

1.3 Project Proposal

Existing Thermal Protection System (TPS) solutions such as carbon-based porous ablators [20], Toughened Uni-piece Fibrous Reinforced Oxidation-Resistant Composite (TUFROC) tiles [7], and various Integrated Thermal Protection Systems (ITPS) [9], [6] have demonstrated considerable progress compared to historical systems. However, despite these improvements, substantial limitations persist. These systems are passive in nature and do still exhibit reliability concerns especially under repeated thermal cycling which may potentially limit rapid reusability needs. These risks are associated with susceptibility to delamination, material and bondline degradation, and progressive cracking and fracturing. This sort of progressive damage over time can result in extensive maintenance demands after flight significantly inhibiting turn around time.

To address these persistent challenges, adaptive material concepts that can improve thermal and structural response within a reusable TPS architecture are required. In this study, the

proposed Adaptive Thermal Protection System (ATPS) is not treated as a single smart material heat shield, but as a multilayer stack augmented with a smart material sublayer. As developed in chapter 2 this will consist of an outer ceramic thermal protection layer, a shape memory alloy (SMA) adaptive interlayer, a ceramic or composite buffer layer, a strain isolation layer, adhesive bondlines, and a protected vehicle skin. Within this architecture, high temperature ceramic materials provide the primary thermal barrier, while the embedded SMA layer provides the adaptive thermomechanical function. Shape memory alloys, including Nitinol-based and high temperature NiTi alloy systems, are of particular interest because their transformation behavior can generate recovery stress, redistribute strain, bridge cracks, or reduce local delamination tendency when placed within an appropriate temperature range inside the TPS stack.

While the theoretical benefits of smart and adaptive materials for aerospace structures have been extensively discussed in the literature [15], [18], robust validation of their practical application within realistic hypersonic TPS environments remains limited. Previous studies have often focused on isolated smart material behavior or simplified thermal loading conditions, rather than evaluating the material response within a representative multilayer TPS exposed to realistic aerodynamic heating. As a result, there remains a gap between the theoretical promise of adaptive TPS concepts and their practical computational validation under hypersonic re-entry conditions.

The primary objective of this project is to address this gap by developing and evaluating a computational framework for an SMA-based Adaptive Thermal Protection System. The study combines hypersonic CFD, thermal analysis, and thermomechanical modeling to evaluate how an embedded adaptive interlayer may influence TPS performance under representative re-entry heating conditions. The CFD portion of the study resolves the external hypersonic flowfield around simplified representative geometries and provides surface heat flux data for the thermal model. The thermal model then evaluates through-thickness heat transfer, surface radiation, back-face temperature, and the temperature at the SMA interlayer. The structural model evaluates the resulting recovery stress and thermomechanical response of the multilayer stack.

To achieve this, representative TPS stack configurations will be modeled using different ceramic, buffer, SMA, strain isolation, adhesive, and vehicle skin material options with a final composition being chosen for each heating case. The primary comparison will focus on how the selected material stack influences surface temperature, back-face temperature, through-thickness gradients, wall heat flux response, and SMA activation potential. The study will also compare the aerothermal loading generated by two simplified hypersonic geometries: an axisymmetric blunt forebody representative of a capsule or sphere-cone heat shield, and a blunted wedge leading edge representative of a hypersonic airfoil or leading-edge surface. Through this computational approach, the study will evaluate whether an embedded SMA layer can provide a meaningful adaptive thermomechanical benefit within a reusable TPS architecture and identify the key limitations associated with material temperature capability, layer placement, and thermal protection requirements.

1.4 Methodology

To achieve the objective of evaluating an SMA-based Adaptive Thermal Protection System under representative hypersonic re-entry conditions, this study uses a coupled computational workflow that combines geometry development, CFD, thermal analysis, and thermomechanical SMA modeling. The methodology is organized so that the external aerodynamic heating environment is first resolved using ANSYS Fluent, then transferred into MATLAB based thermal and SMA response models. This allows the external flowfield and internal TPS stack behavior to be evaluated without explicitly modeling the full multilayer solid structure inside Fluent.

The geometry development begins in ANSYS SpaceClaim. Two simplified representative hypersonic geometries are created: an axisymmetric blunt forebody and a blunted wedge leading edge. The blunt forebody represents a capsule like or sphere cone re-entry heat shield, while the blunted wedge represents a simplified hypersonic leading edge or airfoil section. These geometries are not intended to represent complete flight vehicles. Instead, they are designed to capture the main aerothermal loading environments relevant to TPS design, including stagnation heating, shoulder or leading edge heating, acreage panel heating, and downstream surface heat flux behavior.

After geometry creation, the surrounding fluid domains are prepared and meshed using ANSYS Meshing. A mapped quadrilateral mesh strategy is used so that the cell structure can be controlled near the wall, around the expected shock region, and through the wake. The mesh is refined along the body surfaces to support accurate wall heat flux and wall shear stress prediction. The near wall spacing is selected based on the required wall resolution for hypersonic CFD, while the farfield and wake regions are sized to reduce boundary influence on the near body flowfield.

The CFD simulations are then performed in ANSYS Fluent. Fluent is used to resolve the external hypersonic flowfield around each geometry and compute the local surface heat flux distribution. The CFD setup includes compressible flow physics, energy transport, turbulence modeling, and finite rate air chemistry. The solid TPS stack is not modeled directly inside Fluent. Instead, the wall is assigned a fixed temperature during each CFD run. This fixed wall temperature acts as an initial estimate of the outer TPS surface temperature and is later updated through the thermal calibration process.

A MATLAB thermal analysis is used to connect the Fluent surface heat flux to the multilayer TPS stack response. The MATLAB model solves the through thickness thermal behavior of the proposed ATPS stack, including the outer ceramic layer, emissivity coating, buffer layer, SMA interlayer, strain isolation layer, adhesive layers, and structural backface. The model evaluates the surface energy balance, internal temperature distribution, SMA layer temperature, and backface temperature. This is necessary because the physical wall temperature is not prescribed independently in reality. Instead, the outer surface temperature is governed by a

balance between aerodynamic heating, radiation from the hot surface, and conduction into the TPS stack.

The CFD and thermal models are therefore coupled through an iterative calibration process. Fluent is first run using an estimated fixed wall temperature. The resulting surface heat flux is exported from Fluent and applied in the MATLAB thermal model. The MATLAB model then computes an updated wall temperature from the thermal energy balance. This updated wall temperature is returned to Fluent for the next CFD run. The process is repeated until the imposed CFD wall temperature and the MATLAB thermal response are consistent. This provides a practical way to approximate radiative equilibrium and through thickness TPS behavior without requiring a full conjugate heat transfer model inside Fluent.

Once the thermal response is obtained, the MATLAB model is also used to evaluate the SMA thermomechanical response. The peak temperature reached within the SMA layer is extracted from the thermal simulation and used as the thermal input to the SMA recovery stress model. This allows the study to estimate whether the embedded SMA reaches a useful transformation temperature range and whether it can generate a meaningful recovery stress or adaptive thermomechanical response. In this way, the SMA is evaluated as an internal adaptive interlayer rather than as an exposed TPS surface material.

The final analysis is organized into six primary evaluation cases. Each geometry is evaluated at three surface locations. For the axisymmetric blunt forebody, one case is taken near the stagnation region and two additional cases are taken at selected points of interest along the acreage or shoulder region. For the blunted wedge leading edge, one case is taken near the leading edge stagnation region and two additional cases are taken along the downstream acreage panel surfaces. These six cases allow the study to compare how the ATPS stack responds to different local heating environments across both representative geometries.

The results from these cases are used to evaluate surface temperature, internal layer temperatures, backface temperature, SMA activation potential, and estimated SMA recovery stress. This workflow provides a consistent process for moving from external hypersonic flow simulation to internal TPS thermal response and then to adaptive SMA behavior. The detailed geometry development, mesh strategy, Fluent physics setup, thermal calibration method, and SMA response model are developed in later chapters.

2.0 Materials & TPS Architecture

2.1 Smart Material Candidates

The material candidates considered in this study are organized according to their intended function within the proposed adaptive thermal protection system. The objective is not to identify a single material capable of withstanding the full re-entry heating environment by itself. Instead, the proposed ATPS is treated as a multi-layer stack in which each material performs a specific thermal, structural, or adaptive function. The general architecture consists of an outer ceramic thermal protection layer with an emissive coating, an adaptive shape memory alloy interlayer, a ceramic or composite buffer layer, a strain isolation layer, adhesive bondlines, and a structural vehicle skin. This approach follows the broader sandwich-structure concept discussed in Section 1.2.5 and is consistent with prior studies on multi-material adaptive and smart TPS architectures by Picadio et al. and Le et al. [9], [6].

This distinction is important because many adaptive or smart materials are not suitable for direct exposure to the maximum surface temperatures associated with Earth re-entry. For example, conventional shape memory alloys and polymer-based strain isolation materials cannot typically survive temperatures above approximately 1000 deg C as exposed outer surfaces. Their value is instead associated with their placement within the stack, where the outer ceramic layer and insulation reduce the temperature to a range where the adaptive material can remain functional. In this study, the shape memory alloy is therefore considered as an intermediate adaptive layer rather than as a primary heat shield material.

The table below lists candidate materials for each major layer of the ATPS stack. The material categories include reusable ceramic tiles, ultra-high-temperature ceramic face layers, high emissivity coatings, ceramic buffer materials, shape memory alloys, strain isolation materials, adhesives, and representative vehicle skin materials. For each material, approximate thermal conductivity and maximum operating temperature are included for preliminary screening. For the shape memory alloy candidates, transformation temperature ranges are also included because the martensite and austenite transformation behavior determines whether the material can provide a useful adaptive response at the expected layer temperature.

The values listed in the table are intended for initial material screening and model development. They should not be interpreted as final design allowables. Many of the properties are dependent on density, processing route, alloy chemistry, heat treatment, coating system, and vendor-specific grade. Once a final material stack is selected, the thermal and structural models should be updated using material data for the specific commercial grade or experimentally characterized material system.

Phase change materials, self-healing materials, and piezoelectric materials are not included as primary candidates in this table. Although these materials are often discussed in the broader field of smart thermal protection systems, they are not the focus of the present ATPS configuration. Phase change materials may provide thermal energy absorption, but they are not

inherently smart materials unless integrated into a broader adaptive system. Self-healing and piezoelectric materials were also excluded from the current material selection because they are less directly aligned with the thermal-mechanical objectives of this study. The rationale for excluding these material classes is discussed further in Section 2.1.2.

Table 2.1 - Table of materials

ATPS Layer / Function	**Candidate Material**	**Material Property**	**Maximum Operating Temperature**	**Transformation Temperatures**
Outer ceramic tile	LI-900 silica tile [21]	0.05 - 0.13 W/m-K	1260 - 1540°C 1533 - 1813 K	N/A
Outer ceramic tile	LI-2200 silica tile [21]	0.08 - 0.18 W/m-K	1260 - 1540°C 1533 - 1813 K	N/A
Outer ceramic tile	FRCI-12 tile [21]	0.06 - 0.15 W/m-K	1260 - 1480°C 1533 - 1753 K	N/A
Outer ceramic tile	AETB-8 / AETB-12 tile [21]	0.07 - 0.18 W/m-K	1400 - 1650°C 1673 - 1923 K	N/A
Outer ceramic tile	TUFI-coated AETB [21], [22]	0.08 - 0.20 W/m-K	1650 - 1730°C 1923 - 2003 K	N/A
Outer ceramic cap	TUFROC [23]	0.10 - 0.30 W/m-K	1700°C 1973 K	N/A

Outer ceramic face layer	ZrB2-SiC [24]	20 - 60 W/m-K	2000 - 2400°C 2273 - 2673 K	N/A
Outer ceramic face layer	HfB2-SiC [24]	25 - 70 W/m-K	2200 - 2500°C 2473 - 2773 K	N/A
Ceramic buffer layer	8YSZ thermal barrier ceramic [25]	1.5 - 2.5 W/m-K	1100 - 1200°C 1373 - 1473 K	N/A
Ceramic buffer layer	Alumina-silica fibrous board [26], [27]	0.08 - 0.20 W/m-K	1200 - 1430°C 1473 - 1703 K	N/A
Ceramic buffer layer	Mullite fibrous board [27]	0.10 - 0.25 W/m-K	1400 - 1600°C 1673 - 1873 K	N/A
Ceramic buffer layer	Zirconia fiber board [28]	0.12 - 0.30 W/m-K	1700 - 2000°C 1973 - 2273 K	N/A
SMA adaptive interlayer	Binary NiTi (Nitinol) [29]	10 - 18 W/m-K	100 - 150°C 373 - 423 K	Ms, Mf, As, Af commonly below 100°C for commercial grades
SMA adaptive interlayer	High temperature	10 - 18 W/m-K	150 - 200°C 423 - 473 K	Ms about 70 - 120°C; Af about 90 - 150°C

	binary NiTi (Nitinol) [29]			
SMA adaptive interlayer	Cu-Al-Ni SMA [30]	30 - 100 W/m-K	200 - 260°C 473 - 533 K	Mf about 118 - 134°C; Ms about 152 - 153°C; As about 149 - 157°C; Af about 163°C
SMA adaptive interlayer	NiTiHf very high temperature Nitinol [31]	10 - 20 W/m-K	300 - 500°C 573 - 773 K	Ms about 120 - 400°C; Af about 150 - 450°C
SMA adaptive interlayer	NiTiZr very high temperature Nitinol [32]	10 - 20 W/m-K	250 - 450°C 523 - 723 K	Ms about 150 - 350°C; Af about 180 - 400°C
SMA adaptive interlayer	NiTiPd very high temperature Nitinol [31]	10 - 25 W/m-K	300 - 500°C 573 - 773 K	Ms about 200 - 350°C; Af about 250 - 400°C
SMA adaptive interlayer	NiTiPt very high temperature Nitinol [33]	10 - 25 W/m-K	500 - 700°C 773 - 973 K	Ms about 400 - 600°C; Af about 450 - 700°C
Strain isolation pad (SIP)	Nomex felt [34]	0.04 - 0.08 W/m-K	200 - 250°C 473 - 523 K	N/A

Strain isolation pad (SIP)	Saffil alumina fiber felt [27]	0.08 - 0.20 W/m-K	1000 - 1200°C 1273 - 1473 K	N/A
Emissivity coating / outer surface	Reaction-cured glass coating, RCG [22]	Emissivity: 0.80 - 0.90	~1260°C 1533 K	N/A
Emissivity coating / outer surface	Toughened unipiece fibrous insulation coating, TUFI [22]	Emissivity: 0.85 - 0.92	1650 - 1730°C 1923 - 2003 K	N/A
Emissivity coating / outer surface	Silicon carbide emissivity coating [22]	Emissivity: 0.85 - 0.95	1500 - 1700°C 1773 - 1973 K	N/A
Emissivity coating / outer surface	Borosilicate glass coating [22]	Emissivity: 0.80 - 0.90	1000 - 1300°C 1273 - 1573 K	N/A
Emissivity coating / outer surface	Hafnium carbide / zirconium carbide UHTC coating [24]	Emissivity: 0.60 - 0.90	2000 - 2500°C 2273 - 2773 K	N/A
Adhesive / bondline	RTV silicone adhesive [35]	0.15 - 0.25 W/m-K	200 - 315°C 473 - 588 K	N/A

Adhesive / bondline	high temperature silicone adhesive [36]	0.20 - 0.35 W/m-K	260 - 370°C 533 - 643 K	N/A
Adhesive / bondline	Alumina-silicate ceramic cement [36]	0.5 - 1.5 W/m-K	1000 - 1600°C 1273 - 1873 K	N/A
Adhesive / bondline	Zirconia ceramic adhesive [37], [38]	1.0 - 2.0 W/m-K	1600 - 2000°C 1873 - 2273 K	N/A
Adhesive / bondline	Phosphate-bonded ceramic adhesive [38]	0.8 - 2.0 W/m-K	1200 - 1700°C1473 - 1973 K	N/A
Vehicle skin	Aluminum 2024 [39]	121 - 151 W/m-K	About 177°C 450 K NASA protected structural limit [40]	N/A
Vehicle skin	Aluminum 7075 [41]	130 - 160 W/m-K	120 - 150°C 450 K NASA protected structural limit [40]	N/A

Vehicle skin	Ti-6Al-4V [42]	6.7 - 7.5 W/m-K	350 - 427°C 623 - 700 K	N/A
Vehicle skin	Inconel 625 [43]	9 - 17 W/m-K	650 - 980°C 923 - 1253 K	N/A
Vehicle skin	C/SiC ceramic matrix composite [44]	5 - 20 W/m-K	1200 - 1600°C 1473 - 1873 K	N/A

2.1.2 Materials Discussion

The material discussion is organized according to the final adaptive thermal protection system configuration used in this study. The proposed ATPS is modeled as a multilayer stack consisting of an outer ceramic thermal protection layer, a shape memory alloy adaptive interlayer, a ceramic or composite buffer layer, a strain isolation layer, adhesive bondlines, and a protected vehicle skin. Each layer is selected based on its functional role in the system rather than its ability to serve as a standalone heat shield.

The outer ceramic layer is responsible for withstanding the applied aerothermal environment and reducing heat transfer into the underlying structure. Candidate materials for this layer include silica-based reusable surface insulation, alumina-borosilicate tile systems, toughened ceramic tile systems, TUFROC-type hybrid reusable TPS, and ultra-high-temperature ceramic face materials. These materials are included because the external surface may experience temperatures well above the operating limits of shape memory alloys, adhesives, and structural metals. The ceramic surface layer therefore acts as the primary thermal barrier and controls how much of the external heating environment is transmitted into the adaptive layers below.

The shape memory alloy layer is the primary adaptive material in the proposed TPS stack. The SMA is not intended to function as an exposed heat shield material. Instead, it is embedded beneath the outer ceramic and buffer layers, where the local temperature is reduced to a range where thermally driven phase transformation can occur. Candidate materials include binary NiTi, Cu-Al-Ni, NiTiHf, NiTiZr, NiTiPd, and NiTiPt. These materials are considered because their martensite and austenite transformation temperatures can be selected through alloy composition and heat treatment. This allows the SMA layer to be tuned to activate within the expected internal temperature range of the tile backface. Due to the extreme aerothermal loads

present in this study, consequently, the four high temperature rated Nitinol alloys, NiTiHf, NiTiZr, NiTiPd, and NiTiPt, are the primary materials of study.

The purpose of the SMA layer is to provide a thermomechanical response that may improve the durability of the TPS stack. Potential functions include recovery stress generation, strain redistribution, crack bridging, reduction of local delamination tendency, and compensation for thermally induced deformation. These effects are most relevant near material interfaces where stiffness mismatch, thermal expansion mismatch, and through-thickness temperature gradients can produce high local stresses. For this reason, the SMA is treated as an embedded adaptive interlayer rather than as a surface material. high temperature SMA systems such as NiTiHf, NiTiZr, NiTiPd, and NiTiPt are especially relevant because conventional binary NiTi generally transforms at temperatures that are too low for placement near the hot side of a reusable TPS.

A ceramic or composite buffer layer is included between the outer TPS material and the SMA layer to control the temperature transmitted to the adaptive interlayer. Candidate buffer materials include 8YSZ, alumina-silica fibrous board, mullite fibrous board, zirconia fiber insulation, and similar high temperature ceramic insulation materials. This layer provides additional thermal resistance while also helping to smooth the thermal gradient between the hot ceramic surface and the lower-temperature adaptive layer. The buffer layer is important because the SMA transformation response must occur within a useful temperature range. If the buffer layer is too insulating, the SMA may not activate. If the buffer layer provides insufficient thermal resistance, the SMA may exceed its functional temperature limit.

The strain isolation layer is included to reduce stress transfer between the brittle ceramic layers and the underlying structural skin. This layer is important because the TPS stack contains materials with significantly different stiffness, thermal conductivity, and coefficients of thermal expansion. Candidate strain isolation materials include Nomex felt, ceramic felt, flexible alumina-silica insulation, and other compliant high temperature insulation layers. In the model, this layer is used to represent a compliant transition region that reduces interfacial stress and helps protect the lower structure from thermal and mechanical mismatch.

Emissivity coatings and adhesive bondlines are also included in the material selection because they influence both the thermal boundary condition and the structural integrity of the TPS stack. high emissivity surface coatings such as reaction-cured glass, TUFI-type ceramic coatings, silicon carbide coatings, and UHTC-based coatings improve radiative heat rejection from the outer surface while also providing oxidation and surface durability benefits. These coatings are especially important because the external wall temperature and radiative heat loss are strongly dependent on surface emissivity. Adhesive and ceramic bondline materials are included because each interface in the multilayer stack must maintain attachment while also limiting excessive thermal stress. Silicone-based adhesives provide compliance and are useful in lower-temperature regions, while ceramic adhesives and cements provide higher temperature capability at the cost of increased stiffness and potential interfacial stress. For this reason, both

emissivity coatings and adhesive layers are treated as supporting materials that can significantly affect the predicted thermal response, even though they are not the primary adaptive layer.

The vehicle skin represents the protected structural boundary of the TPS system. Aluminum alloys are included as baseline structural materials because they are common in aerospace applications and are representative of protected airframe structures. However, their allowable temperature is low compared with titanium alloys, nickel superalloys, and ceramic matrix composites. Ti-6Al-4V, Inconel alloys, and C/SiC are therefore included as higher-temperature structural alternatives. These materials provide different tradeoffs between thermal conductivity, structural temperature capability, and mass. The vehicle skin is not expected to experience the external re-entry surface temperature, but its allowable temperature provides an important back-face constraint for evaluating the effectiveness of the TPS stack.

Phase change materials, self-healing materials, and piezoelectric materials are not included as primary candidates in the material selection table. Phase change materials can provide passive thermal energy absorption, but they do not directly provide the thermomechanical adaptive behavior targeted in the present design. Self-healing materials remain promising for future TPS damage-tolerance concepts, but their behavior is difficult to model without specialized fracture, healing kinetics, and material recovery data. Piezoelectric materials provide sensing or actuation capability, but their thermal protection function is indirect and would require a coupled-field model outside the scope of this work. For these reasons, the material discussion focuses on ceramic TPS materials, high temperature buffer materials, shape memory alloys, strain isolation materials, bondlines, and vehicle skin materials.

2.2 ATPS Layering and Configuration

2.2.1 Layering Methodology

The present study investigates the thermomechanical behavior of a layered Adaptive Thermal Protection System (ATPS) architecture developed for hypersonic re-entry applications. The objective of the layered design is not only to reduce through-thickness heat transfer into the vehicle structure, but also to improve resistance to thermally driven interfacial damage mechanisms such as crack initiation, debonding, and delamination. In this work, the TPS is treated as a multifunctional assembly in which the dominant thermal protection function and the adaptive mechanical-response function are assigned to different layers that act cooperatively through the thickness.

The architecture adopted in this study is based on the functional logic of reusable ceramic tile systems developed for the Space Shuttle program, modified to incorporate a thin smart-material interlayer. The resulting stack consists of a high emittance protective surface coat, a primary low conductivity ceramic tile, a Nitinol shape memory alloy (SMA) interlayer located near the backface of the tile, a thin buffer layer, a strain-isolation pad (SIP), a compliant adhesive bondline, and the underlying metallic vehicle skin or structural substrate. This configuration reflects a deliberate separation of roles within the stack. The outer tile system is responsible for

attenuating the imposed aerodynamic heat load, while the embedded SMA is intended to improve the local mechanical integrity of the stack under severe thermal gradients and repeated thermal cycling.

The present methodology does not treat the SMA as a primary load-bearing member in the manner of a structural facesheet or metallic backer. Instead, it is introduced as a thin adaptive mesh interlayer whose purpose is to generate thermally activated recovery stress, improve resistance to crack opening near the tile backface, and redistribute local strain across the transition from the brittle ceramic region to the more compliant support layers beneath it. This distinction is important, because the structural role of the Nitinol in the current architecture is fundamentally interfacial and local rather than global and load-carrying.

Two thermal protection system concepts are considered within this framework. The first is a passive multilayer configuration in which the heat-shielding function is provided entirely by the conventional thermal protection layers. The second is an adaptive configuration in which a Nitinol mesh is embedded within the stack at a location selected to maximize its mechanical influence while minimizing thermal penalty. In both cases, the broader geometric setup, external heating environment, and boundary conditions remain consistent, allowing the influence of the SMA layer to be interpreted directly in terms of thermomechanical response rather than gross geometric change.

The layer ordering is selected to reflect both physical TPS practice and the intended functional role of each constituent. The outer surface coat is included to represent the environmentally resistant, relatively high emittance protective skin required at the exposed hot face. The primary ceramic tile beneath it provides the dominant thermal resistance of the system. The Nitinol mesh is located near the tile backface, rather than near the external surface or deep near the structure, because this region is thermally cooler, mechanically more favorable for smart-material survivability, and closer to the interface where gradient-driven cracking and interfacial opening are most likely to develop. Below the SMA, a thin buffer layer is introduced to moderate local mismatch and provide a transition into the lower compliant region. The SIP and adhesive beneath this region preserve the compliant support logic of reusable tile systems by isolating the brittle upper layers from deformation of the metallic substructure. Finally, the vehicle skin represents the protected structural element whose temperature must remain below allowable limits during flight.

This layer arrangement is used consistently in the reduced-order thermal analysis, in the later numerical modeling, and in the conceptual structural interpretation of the stack. The methodology therefore emphasizes not only thermal attenuation, but also the coordination of layer placement, material function, and interfacial mechanics. In this sense, the layered architecture is not treated merely as a set of materials arranged in series, but as a controlled thermomechanical system in which layer order directly determines both temperature distribution and damage resistance.

For clarity, the layered TPS concept investigated in this work may be summarized as follows:

Table 2.2 – TPS layer visualization

Exposed Surface to Protected Structure	Material Type
Layer 1	Emissive Optical Coating
Layer 2	Primary Ceramic Insulation Tile
Layer 3	Nitinol/SMA Mesh Interlayer
Layer 4	Buffer Layer
Layer 5	Strain Isolation Pad (SIP)
Layer 6	Compliant adhesive bondline
Layer 7	Vehicle Substructure

2.2.2 Functional Operation of Layered TPS Configurations

The functional behavior of the proposed layered TPS architecture is governed by the cooperative interaction of thermal insulation, mechanical compliance, and adaptive interfacial reinforcement. Each layer in the stack is assigned a distinct role, but the overall performance of the assembly depends on how these roles interact under the coupled thermal and structural loading environment of hypersonic re-entry. The objective is not simply to maintain a large temperature drop through the thickness, but to do so while preserving the integrity of the brittle ceramic and bonded interfaces that are typically most susceptible to thermally induced damage.

The outermost surface coat functions as the environmental barrier at the exposed hot face of the TPS. During re-entry, this layer is subjected directly to the imposed aerothermal loading and must tolerate elevated surface temperature, thermal radiation exchange, oxidation, and localized surface damage mechanisms. Because it is thin relative to the total stack thickness, its principal role is not to provide the majority of the thermal resistance, but to protect the porous insulating tile beneath it and to provide an outer surface with appropriate durability and emissive characteristics. In this sense, the surface coat acts as the first protective skin of the system, while the deeper layers perform the majority of the thermal attenuation.

Immediately beneath the surface coat is the primary ceramic tile, which serves as the dominant thermal barrier in the architecture. This layer carries the largest through thickness temperature drop thereby reducing the incoming effective heat load to levels compatible with the lower layers and the protected vehicle structure. Because the tile is highly insulating but inherently brittle, it can be mechanically delicate under aerothermal loads. Under rapid heating, the tile develops a steep temperature gradient from its hot exterior to its cooler backface. This gradient induces internal thermal expansion mismatch (CTE mismatch), which can generate

bending , localized tensile stress, and microcrack initiation, especially in the colder rear portion of the tile. Since ceramic materials are far more tolerant of compressive stress than tensile opening, the mechanical vulnerability of the tile is concentrated not only at the exposed surface, but also at the lower portion of the tile and at its adjacent interfaces.

For this reason, the Nitinol layer is placed near the backface of the ceramic tile rather than near the hot surface or deep near the structural substrate. This placement is chosen as a compromise between thermal survivability and mechanical effectiveness. By positioning the Nitinol below the majority of the ceramic thickness, the layer is protected from the most extreme temperatures in the stack and is more likely to remain within a useful transformation regime rather than being thermally overdriven. At the same time, this location places the SMA close to the region where tile backface cracking, interfacial opening, and thermally driven separation are most likely to occur. The Nitinol is therefore situated where it can exert the greatest mechanical influence without significantly undermining the primary thermal function of the ceramic tile.

The structural role of the Nitinol in this architecture is local and adaptive. When prestrained during fabrication and later heated into its transformation range, the constrained Nitinol attempts to recover its trained shape as it transforms toward the austenitic phase. Because the surrounding layers restrict large macroscopic motion, this recovery does not primarily manifest as visible shape change. Instead, it generates internal recovery stress within the layered assembly. In the present TPS concept, thermally induced recovery stress can impose a localized clamping or compressive effect near the tile backface. This is beneficial because the brittle ceramic layer is much more resistant to compressive loading than to tensile opening. The adaptive Nitinol interlayer therefore offers a mechanism for suppressing or moderating the tensile stress components that promote debonding or crack initiation during severe heating.

Another key contribution of the Nitinol layer is its potential to aid in interfacial crack control. The ceramic tile remains a brittle material even when it performs well thermally particularly under repeated thermal cycling. Reuse and cycling makes it more susceptible to local cracking or delamination near the cold side or along bonded interfaces. A bonded Nitinol mesh in this region can act as a crack bridging network. It is important to note that this does not transform the tile into a conventional reinforced structural laminate or replace the need for a compliant support system beneath the tile. Rather, it provides a restraining mechanism that can resist local opening displacement, reduce the rate at which interfacial separation develops, and increase the resistance of the stack to progressive damage accumulation over repeated heating cycles. This crack-control function is especially relevant for reusable TPS concepts, where long-term durability is influenced by cyclic degradation rather than by a single thermal event alone.

A potential function in current TPS models is that the Nitinol layer may also improve strain transfer through the thickness by redistributing local mismatch strain near the transition from the stiff ceramic region to the more compliant lower layers. In the present stack, the tile is relatively stiff and brittle, while the lower support region, particularly the SIP and adhesive, is considerably more compliant. This abrupt change in stiffness can create concentrated interfacial

stress under thermal expansion mismatch and structural loading. By embedding a thin Nitinol mesh near the tile backface, an intermediate layer is introduced that can alter how strain is transferred across this transition. If the SMA undergoes a temperature dependent change in modulus during phase transformation, then its role becomes even more useful, since it is able not only to provide a local internal stress bias but also to modify the effective stiffness distribution in response to the evolving thermal environment.

The Nitinol is modeled conceptually as a mesh or grid rather than as a continuous solid sheet. This geometric choice is important for both thermal and mechanical reasons. A continuous metallic interlayer would introduce an undesirable increase in through-thickness thermal conductivity and would risk undermining the insulating function of the tile. A mesh geometry reduces the effective conductive cross-sectional area while still allowing distributed adaptive stress generation. Mechanically, a mesh also provides greater in-plane compliance and is better able to accommodate local mismatch strain among ceramic, compliant support materials, and metallic layers. In addition, a multidirectional or quasi-isotropic mesh pattern allows the adaptive effect to be distributed more uniformly across the panel footprint and avoids strongly directional behavior in crack suppression or strain redistribution. The mesh form therefore represents the most practical means of extracting useful adaptive behavior from the SMA while limiting thermal penalty and mass increase.

Beneath the Nitinol lies the buffer layer, which serves as a transition region between the adaptive interlayer and the lower compliant support system. This layer contributes modestly to thermal resistance, but its primary purpose is to soften local mismatch between the Nitinol and the lower support layers and to prevent the adaptive interlayer from acting directly against a highly compliant substrate with no intermediate transition. Functionally, the buffer can be viewed as a stabilizing layer that helps distribute the local effects of SMA recovery stress before they are transmitted into the deeper support region.

The strain-isolation pad beneath this region retains the essential function it serves in reusable tile systems: it decouples the brittle upper tile assembly from deformation of the underlying structure. The SIP is compliant by design and allows differential motion to occur between the TPS and the vehicle skin without transferring excessive structural strain directly into the ceramic tile. In the present architecture, this function remains indispensable. The Nitinol is intentionally placed above the SIP rather than below it, because the adaptive action of the SMA is more valuable near the tile side region than at the deeper structural interface. This placement allows the SIP to continue serving as the principal strain-isolation element while the Nitinol acts as a localized interfacial reinforcement near the ceramic backface.

Below the SIP is an adhesive bondline, which joins the TPS assembly to the metallic vehicle skin or primary structure. The adhesive must accommodate residual mismatch strain, maintain attachment through thermal cycling, and tolerate local shear and peel deformation. In conventional TPS systems, this bondline is often a location of progressive degradation because it is repeatedly subjected to mismatch-driven strain and localized stress concentration. In the present architecture, the adhesive remains a compliant attachment layer, but it may benefit

indirectly from the presence of the Nitinol above. To the extent that the adaptive interlayer reduces tensile opening and moderates strain transfer in the upper stack, the stress transmitted into the lower compliant system may become less severe. The Nitinol therefore does not directly replace the adhesive's function, but rather supports a more favorable stress environment throughout the overall multilayer system.

The final layer is the metallic vehicle skin or structural substrate, which remains the ultimate load-bearing element of the assembly. The role of the TPS stack is to ensure that by the time the incoming heat reaches this layer, the temperature has been sufficiently attenuated to preserve structural integrity. The thermal function of the upper tile system and the mechanical function of the lower compliant support region therefore converge at this point: the stack must both reduce the thermal load and avoid transmitting harmful interfacial stresses into the structure. In this sense, the vehicle skin is not simply the last layer of the stack, but the protected structural boundary whose survivability defines the minimum required performance of the TPS.

Under hypersonic service conditions, the functional sequence of the architecture is therefore as follows. The surface coat and ceramic tile receive and attenuate the external aerodynamic heat load, producing a strong but controlled through-thickness thermal gradient. Because this gradient induces thermally driven expansion mismatch and local tensile tendencies near the colder side of the ceramic tile, the embedded Nitinol layer is positioned to respond once its transformation temperature range is approached. As the Nitinol activates, it develops internal recovery stress that acts locally to oppose interfacial opening, resist crack propagation, and redistribute strain across the transition from the brittle tile region to the compliant lower support region. The buffer, SIP, and adhesive further isolate the protected structure from both thermal and mechanical mismatch.

The resulting assembly is therefore a mechanically adaptive TPS concept intended to improve durability against the coupled thermal and interfacial damage mechanisms associated with repeated hypersonic service.

3.0 Geometry

3.1 Overview

The full ATPS system proposed would likely be composed of multiple discrete panels or tiles. Each panel would follow the same general material configuration described in Chapter 2. As previously discussed, reusable TPS architectures introduce risks associated with cracking, delamination, bond degradation, and CTE mismatch. The proposed configurations are intended

to reduce these risks through the compressive force and strain redistribution provided by the recovery stress of the embedded transformed SMA layer. These effects are outlined in detail in Section 2.2.2.

All geometries developed in this study are built within ANSYS Spaceclaim as part of a Fluent workflow in ANSYS Workbench. To evaluate the performance of the proposed system, this study employs simplified two dimensional CFD geometries that balance aerodynamic relevance, thermal loading fidelity, and computational efficiency. In practical implementation, the ATPS would be composed of modular tiles or panels conformally mounted to vehicle surfaces. However, the present analysis does not attempt to resolve individual tile gaps, fasteners, or panel to panel discontinuities. Instead, the simulations use continuous cross sectional representations of representative vehicle surfaces. This approach allows the external hypersonic flowfield, shock structure, boundary layer development, and surface heat flux distribution to be captured in a controlled manner while still preserving the internal multilayer ATPS stack needed for thermal and thermomechanical evaluation.

Two representative geometries are analyzed. The first is an axisymmetric blunt forebody, representative of a stagnation region heat shield such as a capsule nose, sphere cone, or rounded re-entry body. This geometry is used to capture detached bow shock behavior, stagnation region heating, and the variation of convective heat flux along a curved forebody surface. The second is a two dimensional blunt tip wedge leading edge, representative of sharp or semisharp hypersonic vehicle features such as fins, chines, control surface leading edges, or edge panels. This case is included because leading edge heating is highly localized and strongly dependent on nose radius, wedge angle, shock stand off distance, and boundary layer behavior near the stagnation region.

Both geometries incorporate the multilayer ATPS sandwich structures described in Section 2.2. For the axisymmetric blunt forebody case, the outer mold line defines the aerodynamic surface exposed to the hypersonic flowfield, while the internal material layers are offset inward from the external surface. Each layer maintains approximately constant thickness normal to the preceding layer, allowing the ATPS stack to conform to the local surface curvature while preserving the intended material arrangement. For the blunt tip wedge leading edge case, the outer layer follows the rounded leading edge radius and downstream wedge surfaces, with the lower layers conformally offset from the external profile. This is important because the thermal response of the system depends not only on the magnitude of the applied heat flux, but also on how that heat flux is distributed along the curved or angled external surface.

This geometry selection provides two distinct hypersonic heating environments relevant to reusable TPS design. The blunt forebody case evaluates the ATPS under a stagnation dominated heating environment with a detached bow shock and distributed heating over a curved nose region. The blunt tip wedge case evaluates the system under a leading edge heating environment where the thermal load is concentrated near a small radius nose and then convected downstream along the wedge surfaces. Together, these cases provide a more representative assessment of the ATPS concept than a simplified planar model, while still avoiding the computational cost and added uncertainty of a full three dimensional vehicle simulation.

This approach allows realistic hypersonic surface heat flux distributions to be applied to the ATPS thermal model while preserving computational efficiency. It also enables detailed tracking of through thickness thermal gradients, local back face temperature response, temperature dependent material behavior, phase transformation behavior within the SMA layer, and thermomechanical actuation effects within the proposed smart material architecture.

3.2 Dimensions and Profile

The two selected geometries are scaled using a common vertical characteristic length of 1 m. This characteristic length is defined as the maximum vertical height or diameter of the aerodynamic body in the two dimensional CFD profile. This scale was selected to keep the geometry large enough to represent a realistic re-entry surface while maintaining a computational domain that is practical for hypersonic CFD. The intent of this sizing is not to directly match the full vehicle scale, but to define representative aerodynamic surfaces that capture the dominant flowfield behavior relevant to TPS loading, including shock formation, shock stand off distance, stagnation heating, boundary layer development, and the surface heat flux distribution along the exposed outer mold line.

This study within ANSYS Fluent is looking at only the external aerodynamics of these geometries. Thus the multilayer ATPS structure is not explicitly modeled in the CFD domain and is deferred to the thermal analysis in chapter 6. Instead, the CFD model is used to resolve the external hypersonic flowfield around the representative body, and the resulting surface heat flux is later applied to the thermal model of the ATPS stack.

The solid ATPS stack used for the thermal and thermomechanical analysis maintains a maximum plausible total thickness cap of 150 mm for each geometry and layering configuration. Literature on the shuttle tiles cites a thickness of around 10-12 cm or up to 125mm on peak heating zones [45]. A conservative estimate could plausibly place a thickness ceiling at 150 mm. Each material layer follows the same general arrangement described in Chapter 2, with the local geometry defining the outer mold line and the remaining layers offset inward from the external surface.

The first geometry is the axisymmetric blunt forebody. This geometry is representative of a stagnation region heat shield such as a capsule nose, sphere cone, or rounded re-entry forebody. The model has a maximum vertical diameter of 1 m, corresponding to a 0.5 m body radius. The total axial body length is 2 m. The forebody nose uses a 0.5 m blunt radius, producing a smooth rounded stagnation region before transitioning into the downstream body profile. The curved forebody surface has an approximate total arc length of 1.877 m. In the axisymmetric Fluent model, this two dimensional profile is revolved about the centerline, meaning the effective surface area is generated by the axisymmetric formulation rather than by assigning an arbitrary exposed area. This geometry is intended to represent curved and contoured re-entry vehicle surfaces such as a command module heat shield, capsule forebody, nose cap, or rounded body section.

The second geometry is the two dimensional blunt tip wedge leading edge. This geometry is representative of sharp or semisharp hypersonic vehicle features such as fins, chines, control surface leading edges, and edge panels. The wedge uses the same 1 m vertical characteristic length, defined as the total height of the wedge section. The leading edge is modeled with a finite nose radius rather than an ideal sharp tip, allowing the formation of a detached or partially detached shock near the leading edge. The baseline leading edge radius is 100 mm. Downstream of the rounded nose, the upper and lower wedge surfaces follow a straight profile with a wedge half angle of 15 degrees measured from the centerline. For a 1 m total wedge height and a 15 degree half angle, the corresponding axial wedge length is approximately 1.5 m, with the rounded leading edge adding a small additional nose length depending on the exact tangent construction. This provides a total leading edge body length within the same order as the 2 m blunt forebody case.

Together, these geometries define two distinct aerodynamic heating cases. The axisymmetric blunt forebody captures a stagnation dominated heating environment where the peak heat flux is associated with a detached bow shock and curved forebody flow. The blunt tip wedge captures a leading edge heating environment where the thermal load is concentrated near the finite radius nose and then distributed downstream along angled wedge surfaces. Using a common 1 m vertical characteristic length keeps the two cases comparable in scale, while the different nose radius, body length, and wedge angle define different hypersonic TPS loading environments.

Table 3.1 - Geometry dimensions

Parameter	**Axisymmetric Blunt body**	**Blunt Tipped Leading Edge**
Geometry type	Axisymmetric sphere-cone	Blunt tipped wedge
CFD formulation	2D axisymmetric	2D planar
Characteristic vertical length	1.0 m	1.0 m
Maximum diameter / total height	1.0 m	1.0 m
Body radius / half height	0.5 m	0.5 m
Nose radius, (R_n)	0.5 m	0.1 m
Total axial body length	2.0 m	~1.575m

Curved forebody arc length	~1.9 m	N/A
Wedge half angle, (θ)	N/A	15°
Included wedge angle	N/A	30°
TPS Stack Maximum Plausible Thickness	150 mm	150 mm

3.2.1 Leading Edge (Blunt Tipped Wedge):

The blunted wedge is representative simplification of a typical hypersonic leading edge or airfoil. The geometry is artificially truncated at the aft end and features a nose radius of 100mm, a half angle of 15 deg and 1m characteristic length or exposed surface vertical height. It is mainly designed to be compliant with the meshing and computational limitations of the ANSYS student license rather than represent a real hypersonic body. Figure 3.1 shows a detailed dimensioned image of the leading edge geometry.

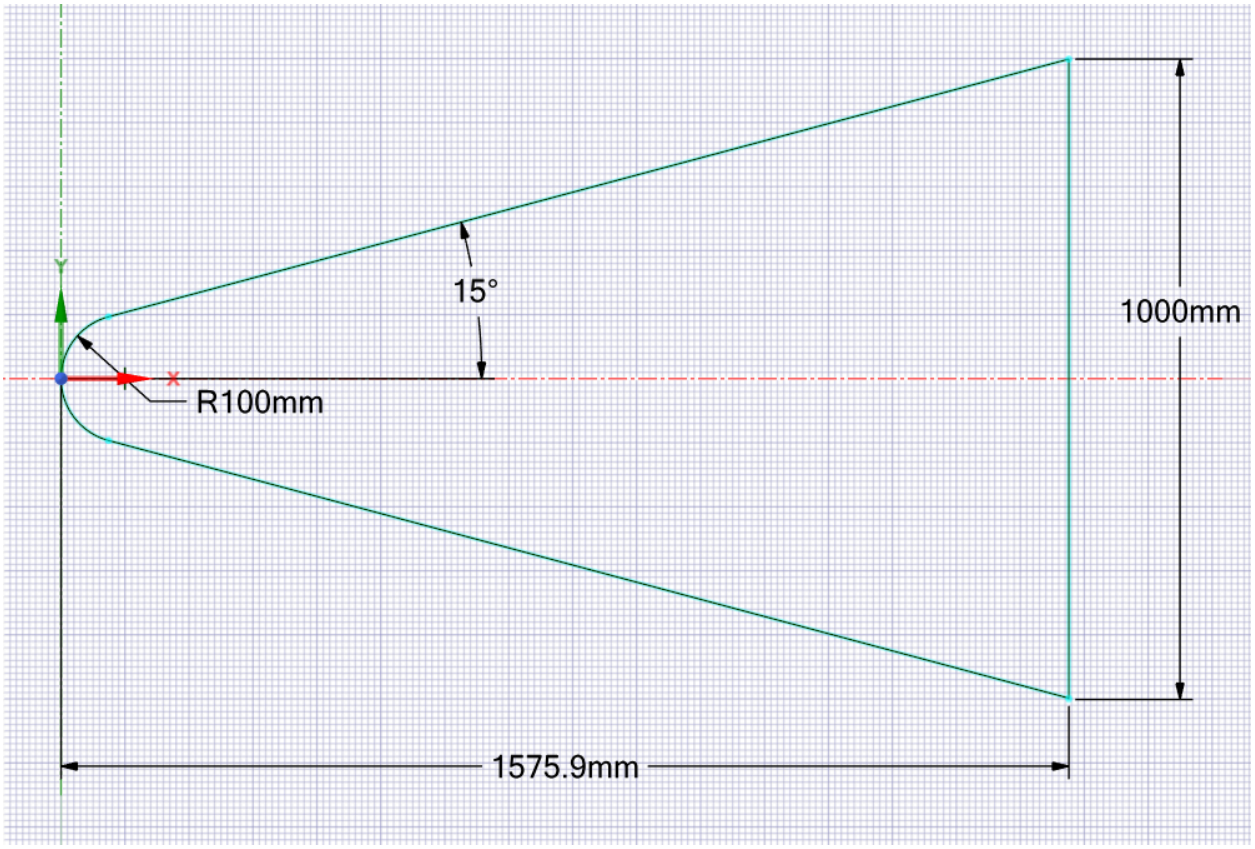


Figure 3.1 - Leading edge (blunt tipped wedge) geometry

The leading edge geometry features a C grid type fluid domain surrounding the body. This is done to ensure inlet flow does not intersect with the fairfield boundary. The topology is split into 4 quadrants. Both the forward and fairfield cushions are set at 10 body lengths, and the wake or streamwise cushion at 20 body lengths. This ensures that the disturbed flow does not meaningfully interact with the inlet, outlet, or fairfield. This sizing is semi consistent with the literature for 2D hypersonic CFD although most suggest cushions of at least 20 in all directions [46], [47]. However a key consideration here for this study was the mesh cell count cut off for the ANSYS student license at about 1,040,000 elements The selected sizings provided an adequate cushion while also allowing for an effective y+ within the limited element count (see chapter 4 for more detail on y+). Figure 3.2 shows the fluid domain for the leading edge/blunted wedge.

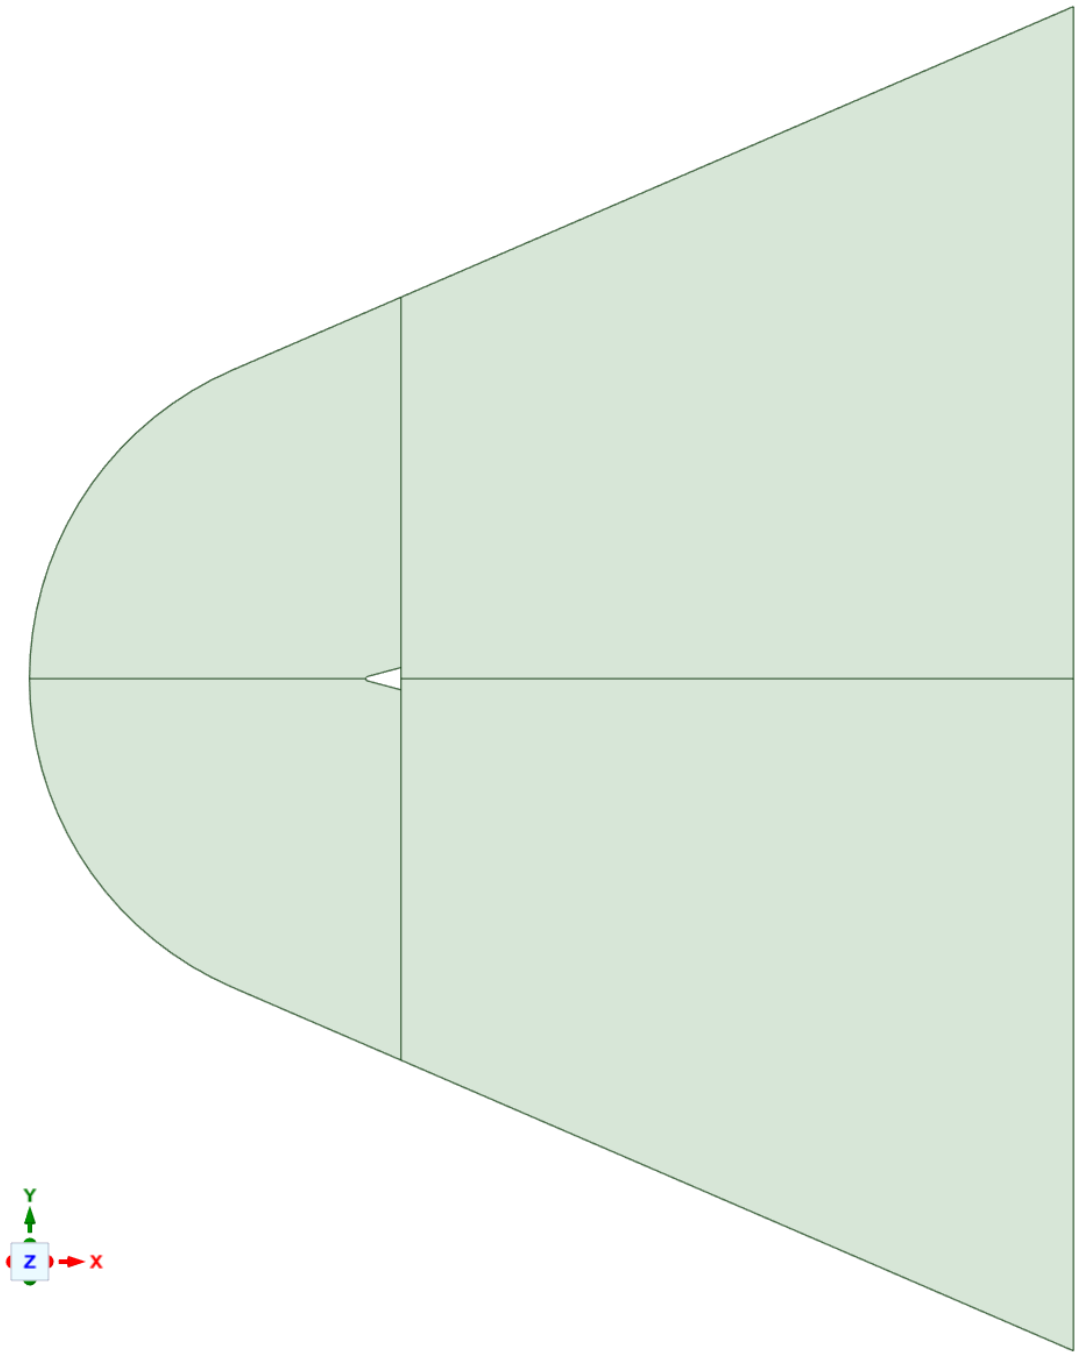


Figure 3.2 - Leading edge (blunt tipped wedge) fluid domain

3.2.2 Blunt Forebody (Axisymmetric Sphere-cone):

The axisymmetric sphere-cone geometry is a representative simplification of a blunt re-entry forebody or capsule type heat shield. The geometry features a nose radius of 0.5 m, a maximum vertical diameter of 1 m, and a total axial body length of 2 m. The curved forebody arc length is approximately 1.9 m. In the axisymmetric Fluent model, the two dimensional profile is

revolved about the centerline, allowing the body to represent a three dimensional blunt forebody while maintaining a simplified computational domain. This geometry is mainly intended to capture the dominant aerothermal behavior of a blunt re-entry surface, including detached bow shock formation, stagnation region heating, shoulder heating effects, and the surface heat flux distribution along the outer mold line. It is also designed to remain compliant with the meshing and computational limitations of the ANSYS student license rather than represent a complete flight vehicle as previously addressed in this section. Figure 3.3 shows a detailed dimensioned image of the axisymmetric blunt forebody geometry.

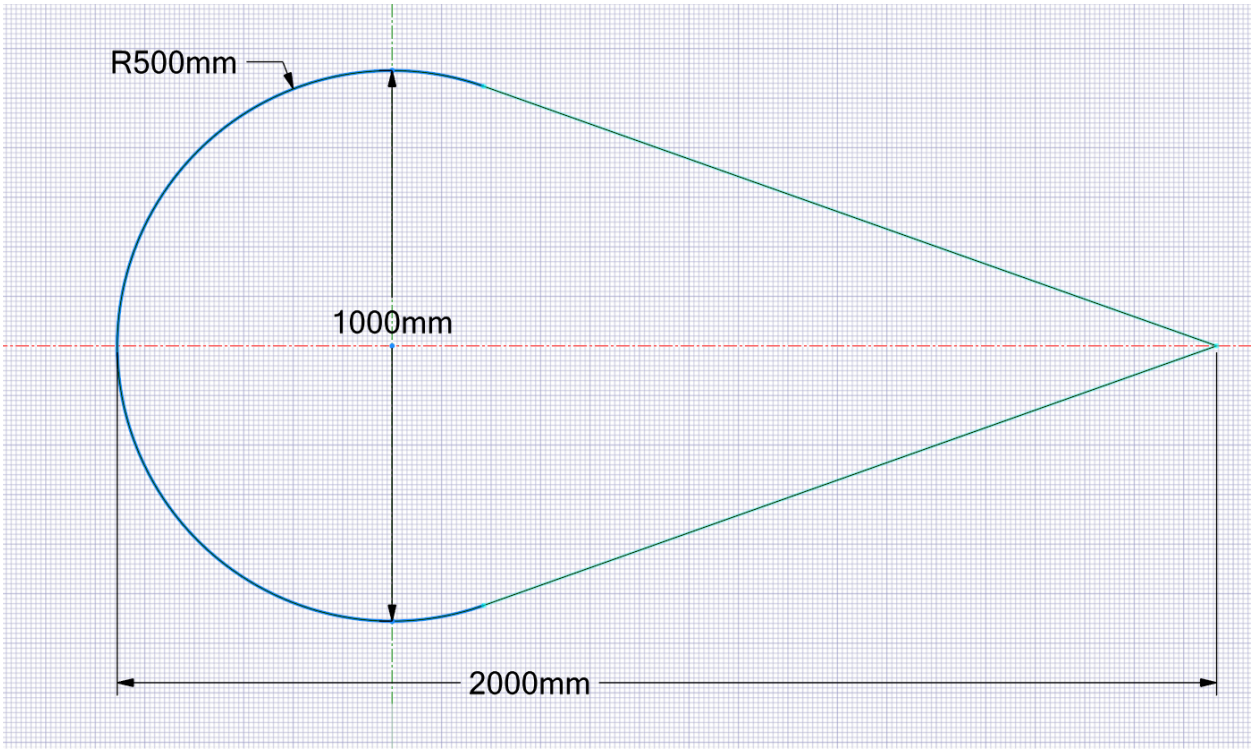


Figure 3.3 - Blunt forebody (axisymmetric sphere-cone) geometry

The axisymmetric blunt forebody geometry uses a C-grid type fluid domain surrounding the body. This domain topology was selected to provide a clear upstream farfield, upper and lower farfield boundaries, and a downstream wake region while avoiding unnecessary interaction between the disturbed flowfield and the outer domain boundaries. The topology is divided into 6 structured regions to support mapped meshing around the body and wake. The forward and radial farfield cushions are set at approximately 10 body lengths, while the downstream wake cushion is set at approximately 20 body lengths which is consistent with the domain sizing for the leading edge. This provides sufficient distance for the detached bow shock, shock layer, and wake structure to develop without strongly interacting with the inlet, outlet, or farfield boundaries. As previously stated in this section this sizing is generally aligned with common practices for lower fidelity two dimensional hypersonic external flow CFD while remaining

within the limitations of the Student License of ANSYS. Figure 3.4 shows the fluid domain for the blunt forebody.

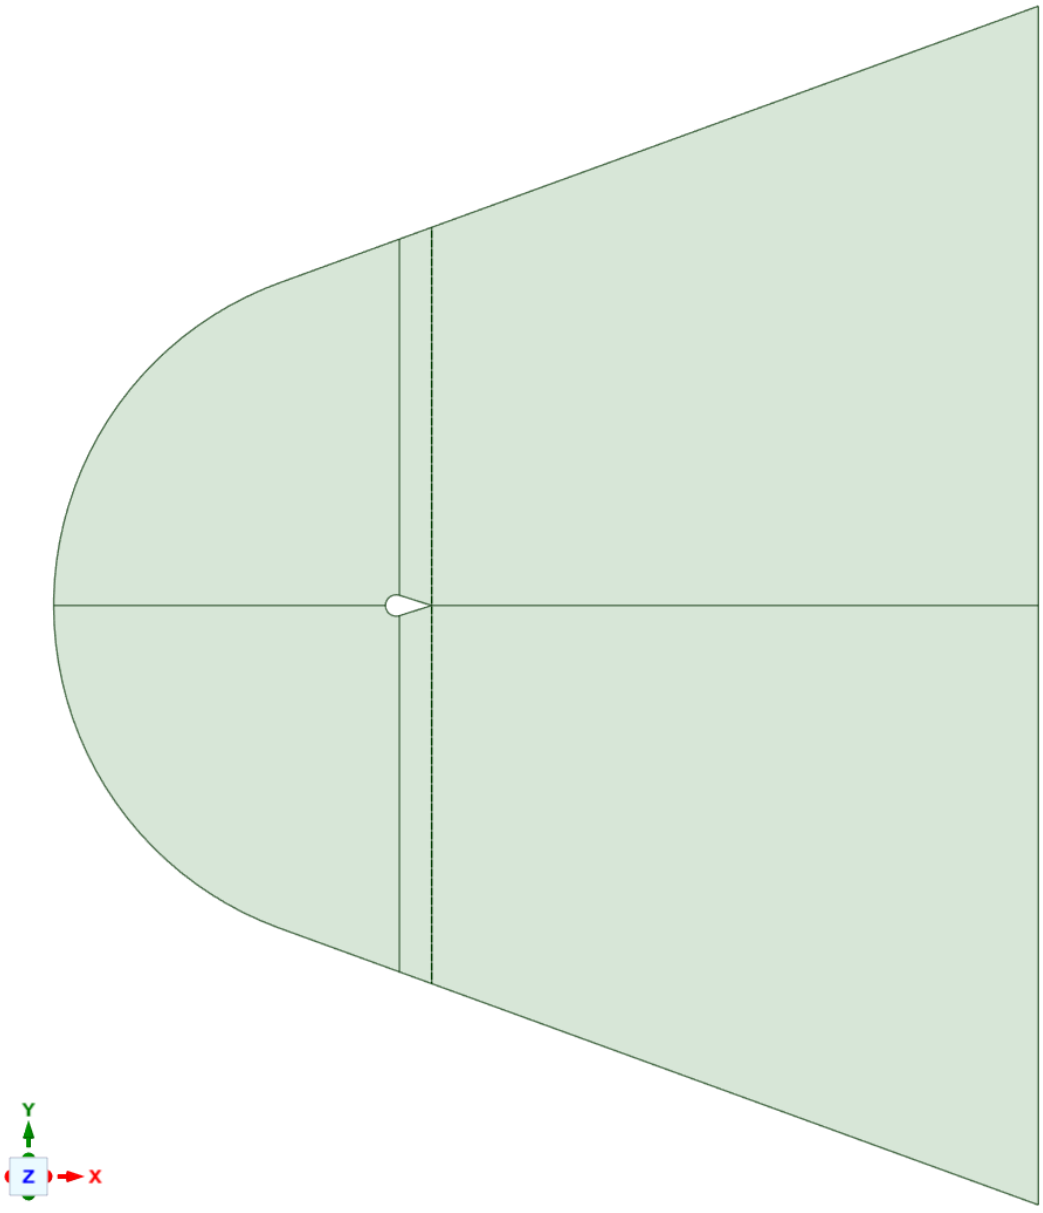


Figure 3.4 - Blunt forebody (axisymmetric sphere-cone) fluid domain

4.0 Meshing/Topology

This chapter describes the mesh development process used for the hypersonic CFD simulations. The purpose of the mesh is to discretize the external flow domain with sufficient resolution to capture the primary aerothermal features of interest, including detached bow shock formation, shock stand off distance, boundary layer development, wake structure, wall shear stress, and surface heat flux. Because the predicted wall heat flux is strongly dependent on the near wall temperature gradient, special attention is given to wall adjacent cell spacing and cell alignment along the body surface.

The two geometries considered in this study require different mesh topologies. The blunt tip wedge uses a four face mapped topology because the geometry consists of a rounded leading edge, two straight wedge surfaces, and an aft truncation surface. The axisymmetric sphere cone uses a six face mapped topology because the forebody shoulder introduces a stronger curvature transition that requires additional partitioning to maintain clean mapped regions. In both cases, the mesh is generated using structured quadrilateral face meshing to provide direct control over cell spacing, edge biasing, and growth away from the wall.

The mesh strategy is developed around three main requirements. First, the near wall region must be sufficiently refined to support a target y+ near unity for accurate wall heat flux and wall shear stress prediction. Second, the near field and wake regions must be resolved enough to capture the shock layer and downstream flow development. Third, the total cell count must remain compatible with the ANSYS Student license limitation of 1.04 million elements. The final mesh therefore represents a balance between physical resolution, numerical stability, and computational practicality.

The sections that follow describe the mapped mesh approach, the wedge and sphere cone mesh topologies, the edge sizing and biasing strategy, and the near wall resolution procedure used to estimate the first cell height. The chapter also discusses how the final meshes were checked for cell alignment and suitability for the high speed reacting flow simulations performed in Fluent.

4.1 Mapped Mesh Approach

The meshes for both geometries were generated using a mapped quadrilateral topology. This approach was selected to provide direct control over cell spacing, cell alignment, and near wall resolution while remaining within the ANSYS Student cell count limitation. Hypersonic CFD requires strong resolution in regions with steep pressure, temperature, species, and velocity gradients. For this reason, the mesh was designed to concentrate cells near the body surface, expected bow shock region, and downstream wake while allowing the outer farfield to remain coarser.

Mapped face meshing was applied to each partitioned region of the fluid domain. This ensured that the mesh remained structured and that the cells followed the main flow direction and the local body contour. The use of mapped quadrilateral regions also improved the control of

cell growth away from the wall, which is important for resolving wall heat flux and wall shear stress. The full mesh and farfield views for the wedge and axisymmetric forebody are shown in Figures 4.1 and 4.2.

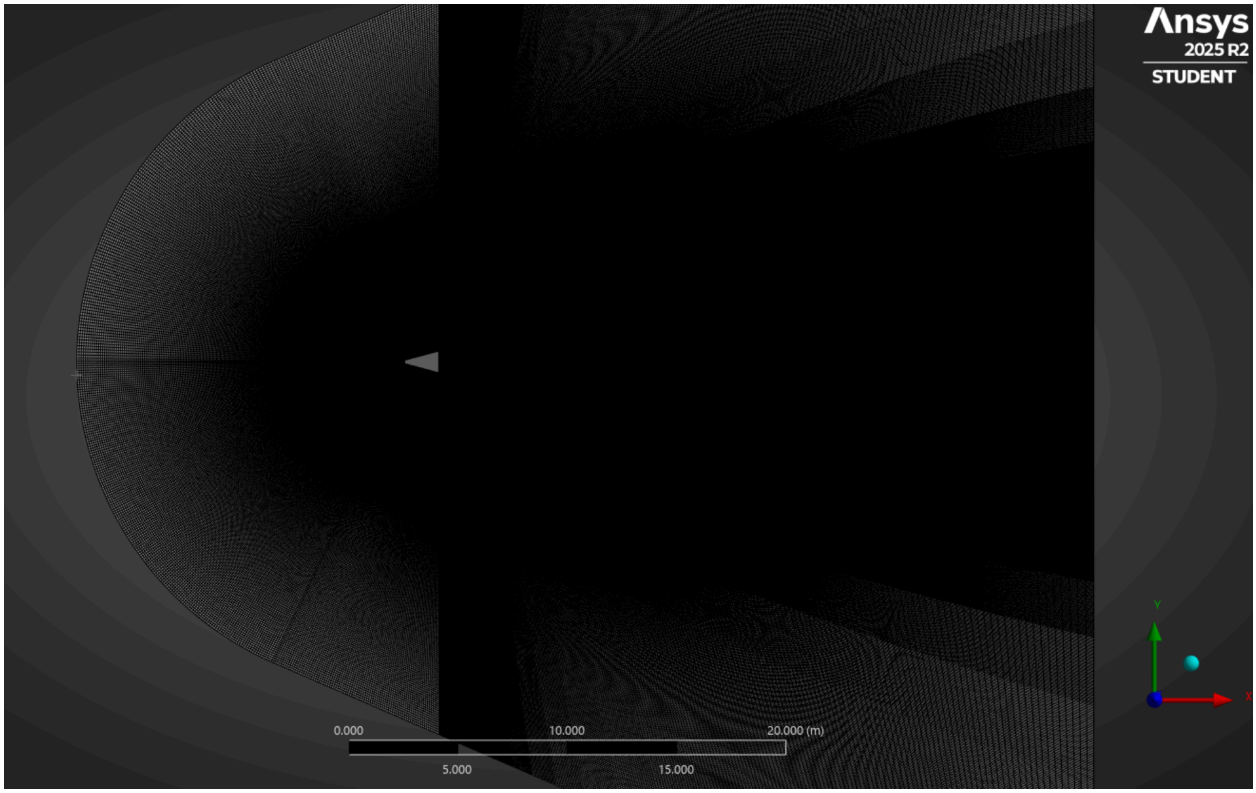


Figure 4.1 - Leading edge (blunted wedge) full mesh

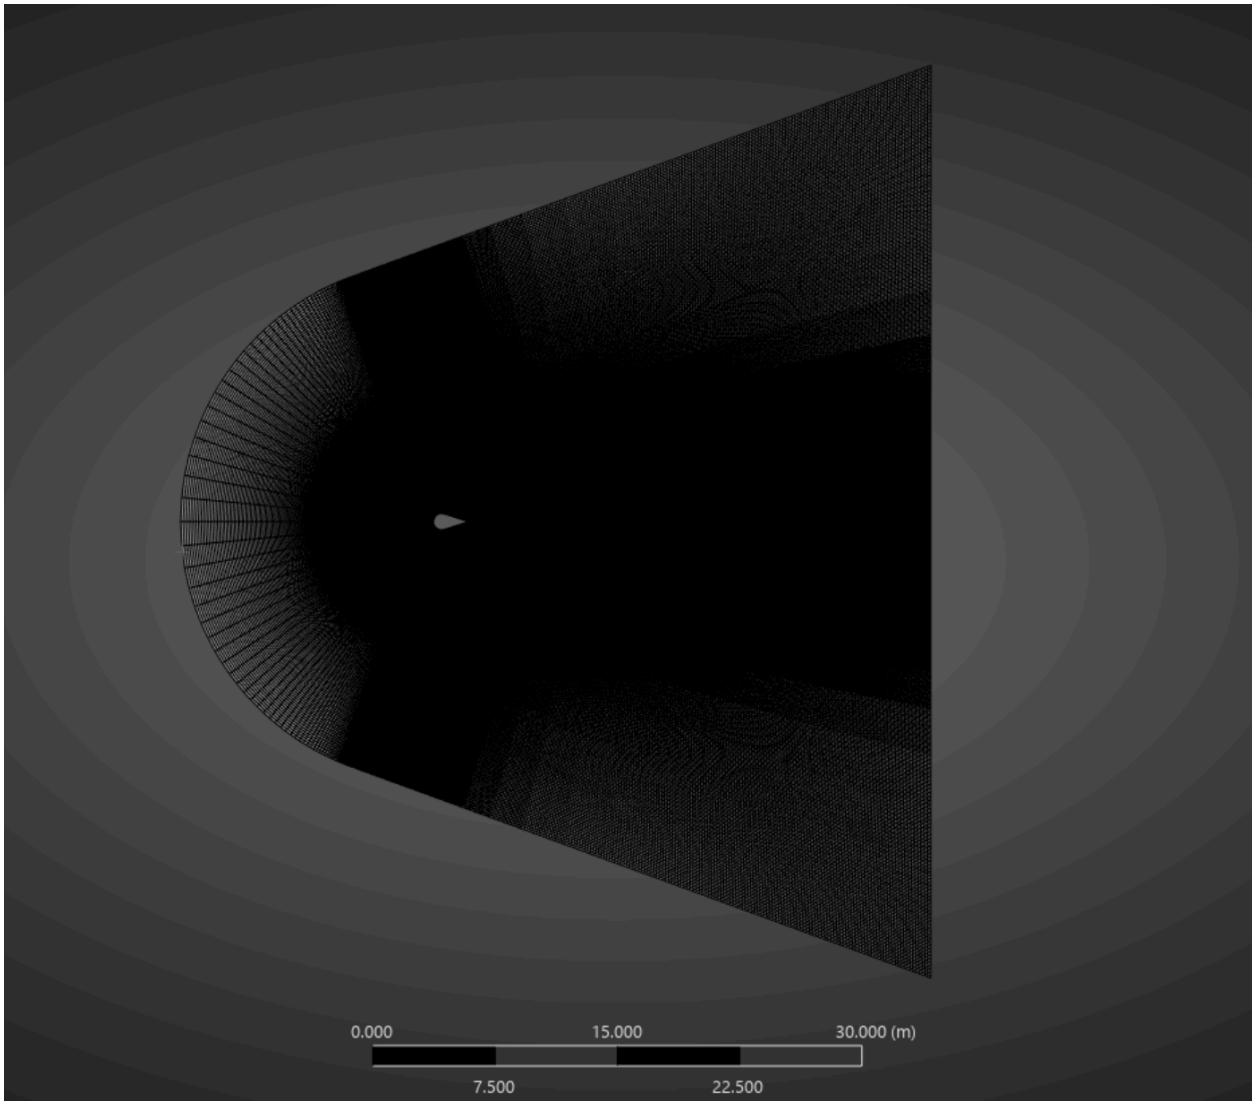


Figure 4.2 - Blunt forebody (axisymmetric sphere-cone) full mesh

4.2 Leading Edge Mesh Topology

The blunt tip wedge/leading edge geometry was divided into four mapped faces. This topology was sufficient because the geometry consists of a rounded leading edge, two straight wedge surfaces, and an artificially truncated aft wall. The surface transitions align relatively well with the outer boundary which allows the mapped mesh to remain organized without additional partitioning.

The mesh is most refined near the rounded leading edge and along the upper and lower wedge surfaces. These regions are expected to contain the strongest thermal and momentum gradients because the detached shock forms near the blunt nose and the boundary layer develops along the wedge surfaces. Downstream of the body, the wake region is resolved using

streamwise aligned cells so that the flow behind the truncated aft wall can develop without abrupt changes in cell size.

4.3 Blunt Forebody Mesh Topology

The axisymmetric blunt forebody required a slightly more detailed topology than the wedge. The sphere cone geometry contains a stronger shoulder transition between the rounded forebody and the downstream tapered body. This transition creates a more complex mapping requirement because the body wall introduces additional edge segments in the surrounding near body regions.

To maintain mapped quadrilateral faces, the axisymmetric forebody domain was divided into six mapped faces. The additional upper and lower partition regions were introduced near the shoulder to prevent distorted mapped faces and to keep each region four sided. This was necessary because the shoulder region contains a rapid change in surface curvature and is also one of the most important regions for heat flux prediction. The resulting topology produced improved cell alignment near the wall and reduced distortion around the forebody shoulder.

4.4 Edge Sizing and Biasing

Edge sizing was applied differently depending on the role of the edge within the mapped topology. For the main topology edges and face edges, division number sizing was used. The number of divisions generally ranged from approximately 500 to 1000 divisions. Higher division counts were assigned to edges near the body, bow shock region, and wake, while lower division counts were used on farfield edges where the solution gradients are weaker.

Bias factors were applied to selected topology edges so that the cell density increased toward the body and near field region. This allowed the mesh to provide higher resolution where the thermal and momentum gradients are largest without exceeding the available element count. The biasing was especially important for the curved farfield blocks and wake blocks because it provided smooth expansion from the near wall mesh into the outer domain.

For boundary edges that were short, split, or discontinuous, element size controls were used without bias. These local edges occur near topology partitions, body splits, shoulder transitions, and aft truncation regions. Applying unbiased element sizing to these edges helped avoid abrupt spacing changes across neighboring mapped faces. This was important for preserving mesh quality and preventing local distortion in regions where multiple mapped blocks meet. The near field mesh and close up images of the wall are shown in Figures 4.3 through 4.6.

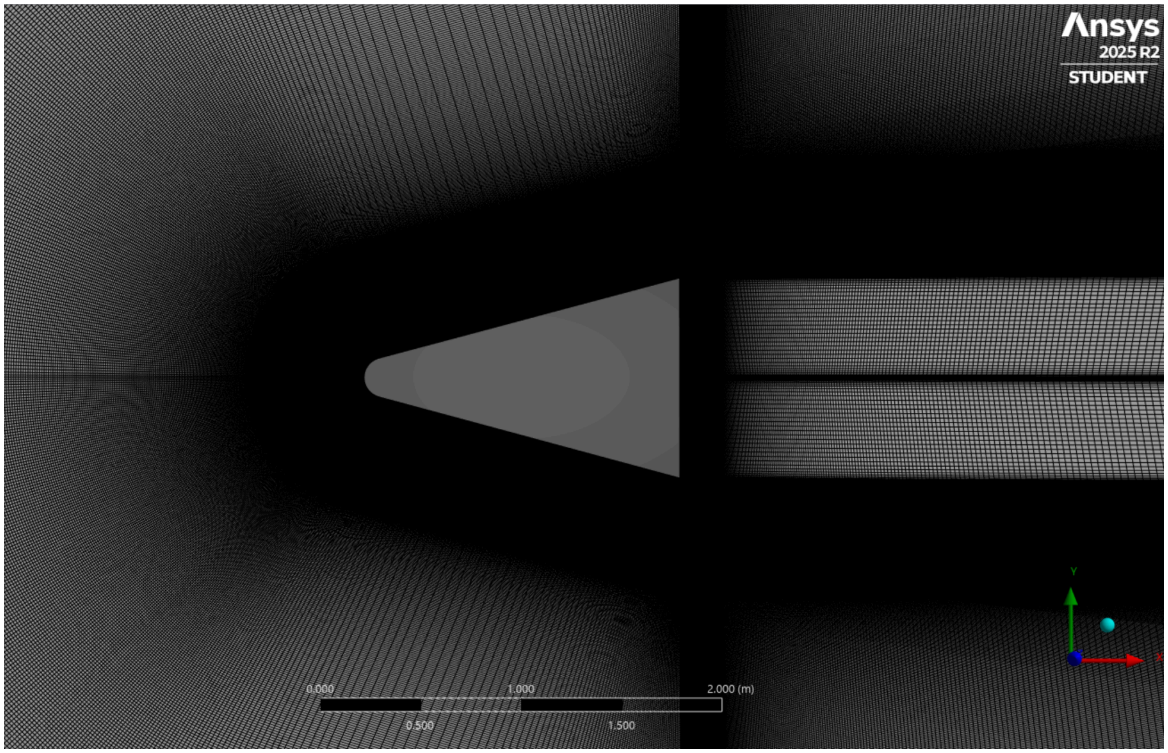


Figure 4.3 - Blunted wedge nearfield mesh

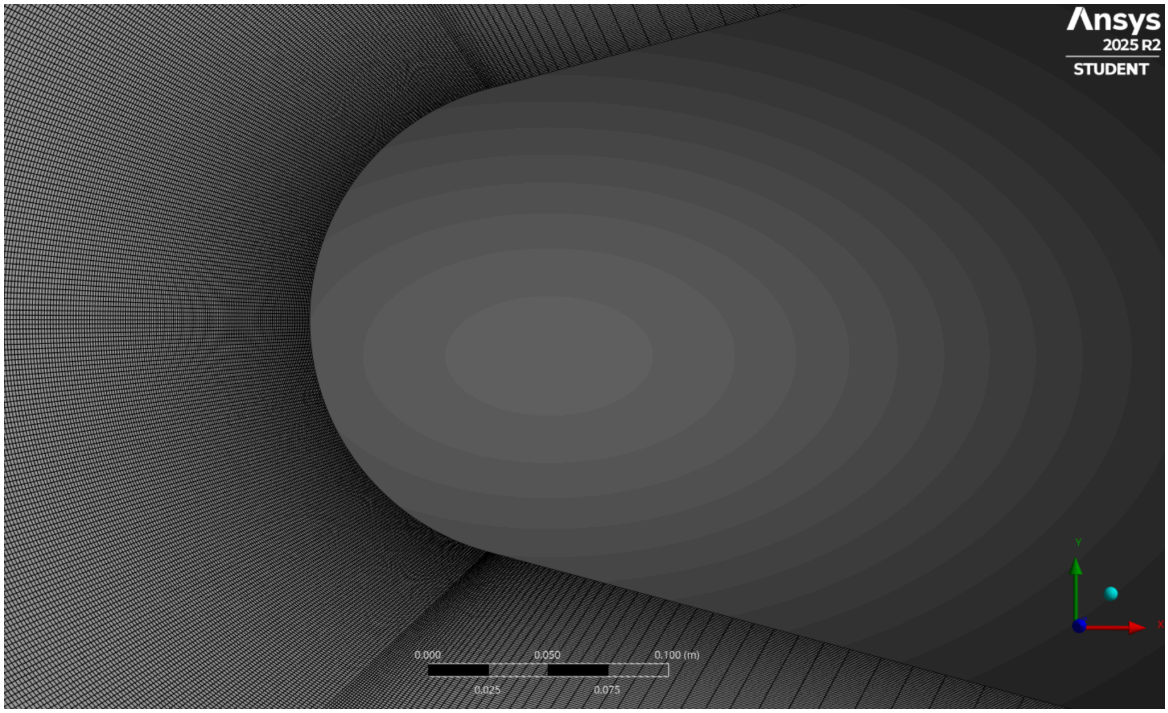


Figure 4.4 - Blunted wedge near wall front

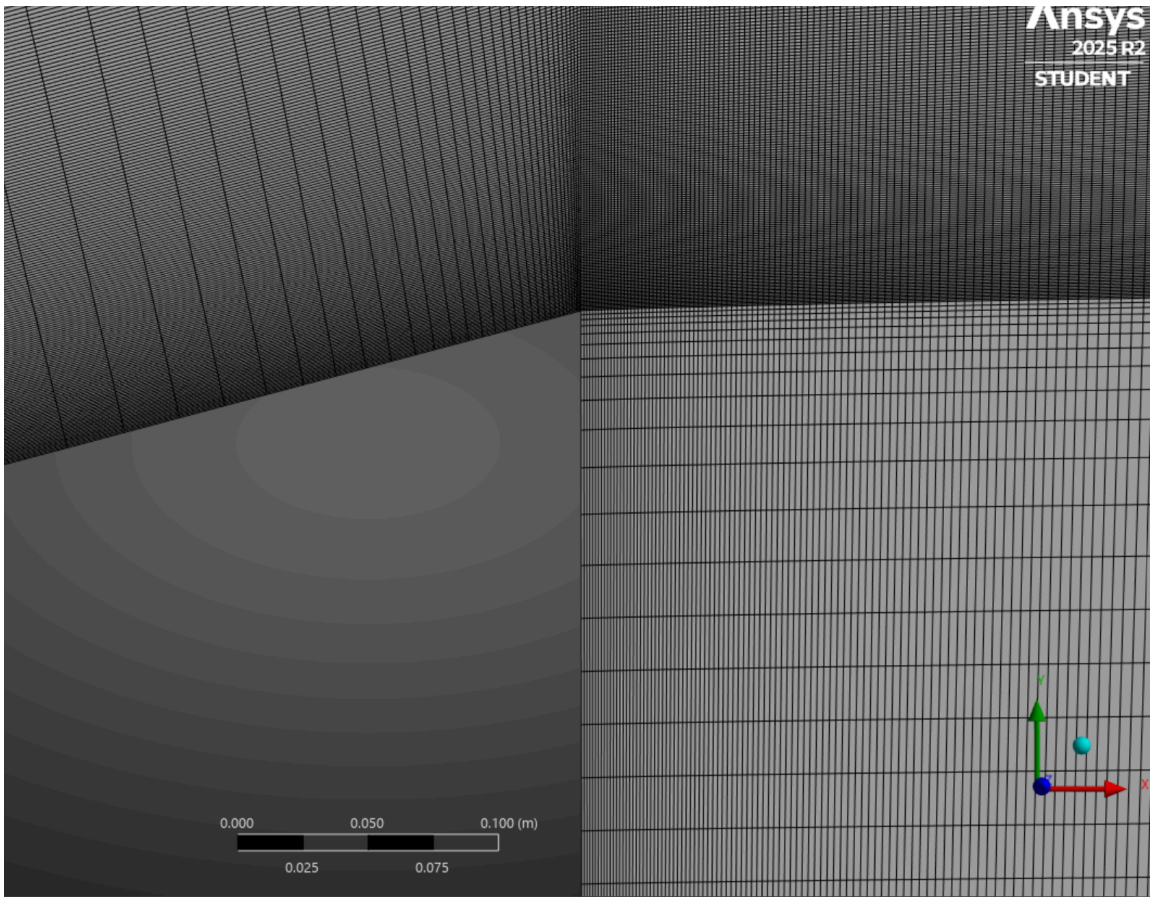


Figure 4.5 - Blunted wedge near wall aft

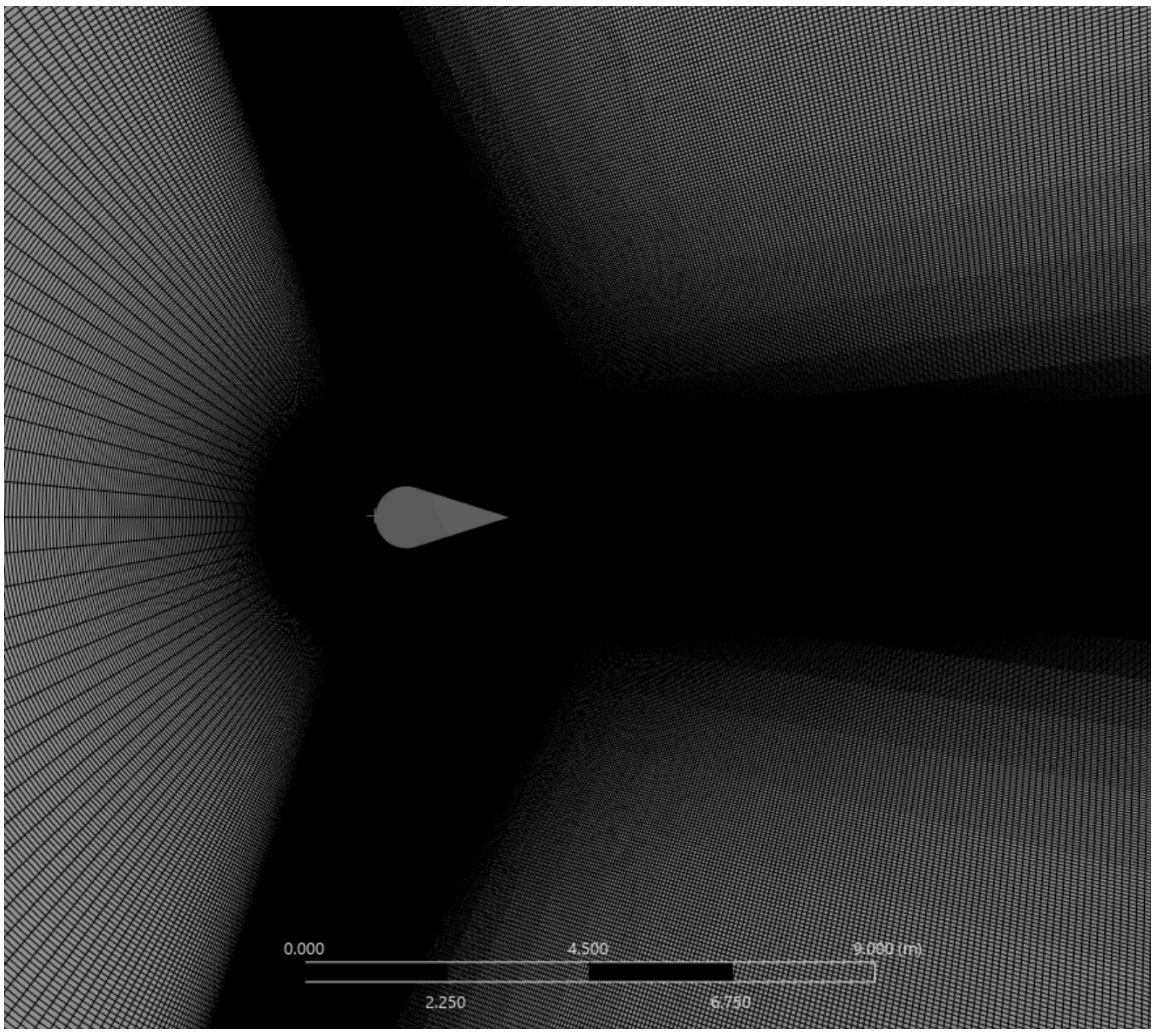


Figure 4.6 - Axisymmetric forebody nearfield mesh

4.5 Near Wall Resolution and y Plus Target

Near wall resolution was selected to support accurate prediction of wall heat flux and wall shear stress. Since the simulations include high speed compressible flow, energy transport, species transport, and wall heat transfer, the first cell height must be small enough to resolve the viscous and thermal gradients near the wall. The target for this study was $y+ \leq 1$. This target was selected so that the near wall region could be resolved directly rather than relying on wall function behavior. This is especially important because the wall heat flux is computed from the temperature gradient normal to the wall. The nondimensional wall distance is defined in equation 11 where y_1 is the first cell height normal to the wall, ρ is the near wall density, μ is the dynamic viscosity, and u_τ is the friction velocity.

$$y^+ = \frac{\rho u_\tau y_1}{\mu} \tag{11}$$

The friction velocity is defined in equation 12 where τ_w is the local wall shear stress. Rearranging the definition of y+ gives the required first cell height as shown in equation 13.

$$u_\tau = \sqrt{\frac{\tau_w}{\rho}} \tag{12}$$

$$y_1 = \frac{y^+ \mu}{\rho u_\tau} \tag{13}$$

The first cell height was estimated using the expected near wall properties from the hypersonic flow condition and then checked using the wall shear stress values from the CFD solution. Based on this estimate, the first cell height was selected to be on the order of 10^{-4}. This provided a practical wall spacing for maintaining y+ near or below unity while still satisfying the total element count limitation. The zoomed wall resolution images in figures 4.7 and 4.8 show the first cell spacing and the structured wall adjacent cells used in the final meshes.

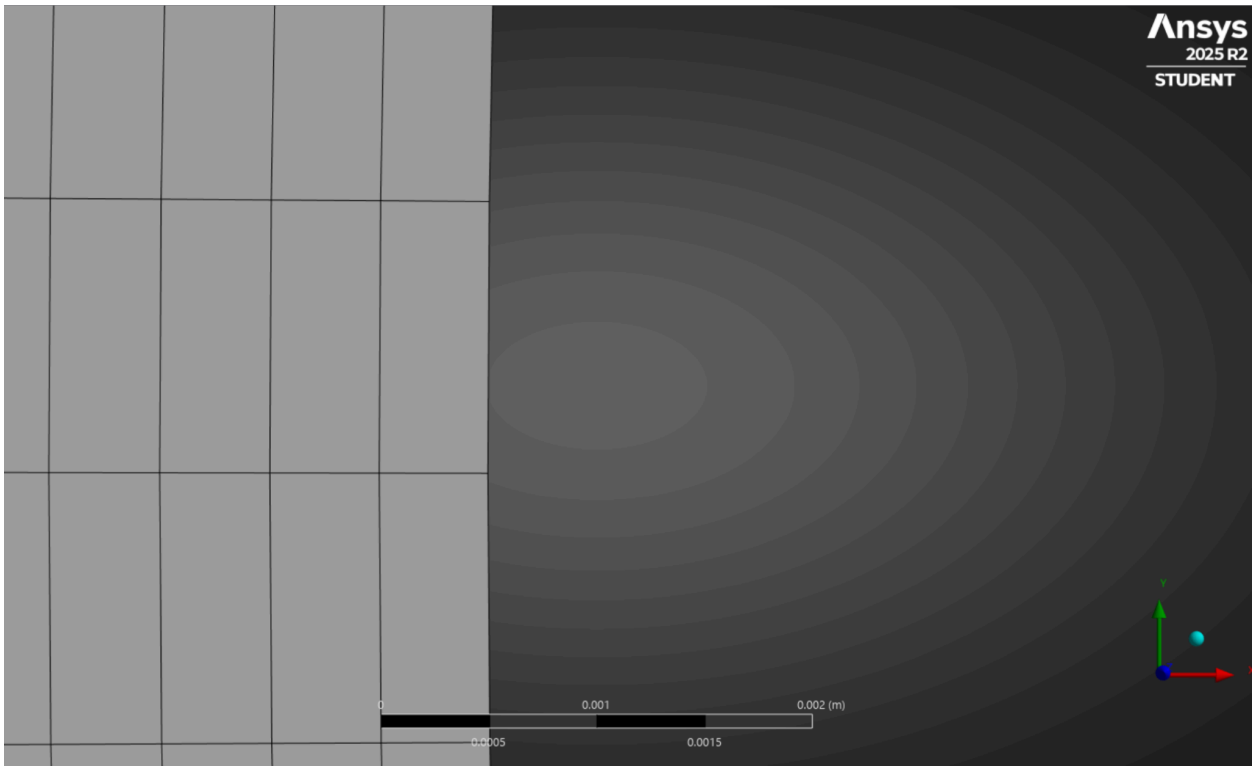


Figure 4.7 - Blunted wedge first cell height

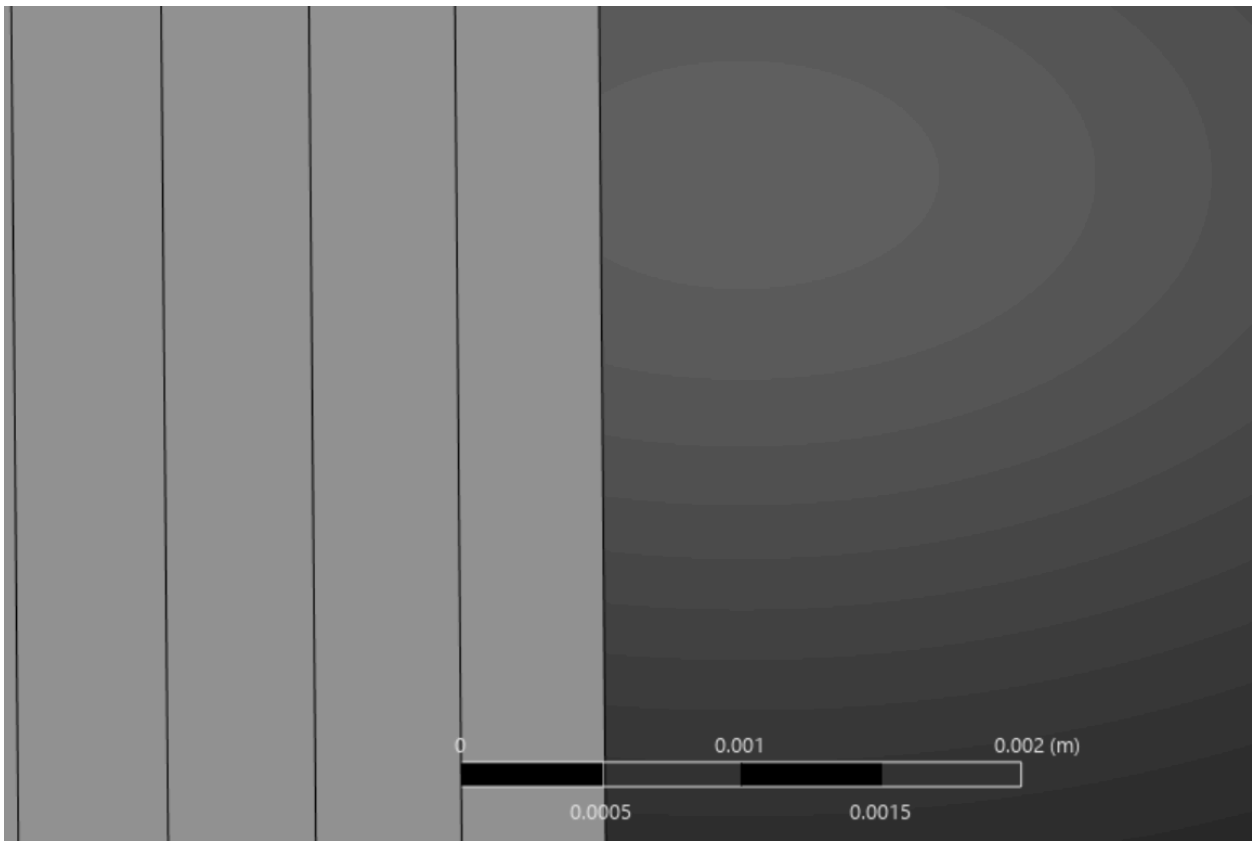


Figure 4.8 - Axisymmetric forebody first cell height

4.6 Wall Alignment and Heat Flux Accuracy

The mesh near the body was designed so that the wall adjacent cells remained aligned normal to the surface. This is important because Fluent calculates the wall gradients of velocity, temperature, and species mass fraction using the local near wall cell structure. Poor alignment near the wall can introduce errors in the computed wall shear stress and surface heat flux. For this reason, the body surface edges were given dedicated sizing controls and the surrounding mapped blocks were arranged to maintain smooth cell growth away from the wall.

The near wall mesh images show that the cells remain structured along the rounded wedge nose, wedge surfaces, blunt forebody, and shoulder region. This is especially important for the axisymmetric forebody because the shoulder region was later shown to produce localized heat flux and wall shear stress augmentation. Maintaining a clean and aligned mesh in that region helped confirm that the observed heating behavior was associated with the flowfield and geometry rather than local mesh distortion.

5.0 CFD

5.1 Boundary Conditions

Proper set up of boundary conditions comes from understanding the laws of fluid mechanics as well as the general operating conditions of re-entry vehicles. Table 5.1 presents flight phases with typical ranges for atmospheric conditions of a re-entry vehicle deorbiting from LEO (~7.8km/s). The values are derived and approximated from the International Standard Atmosphere (ISA) and Sutton & Graves Re-entry Modeling [2], [48].

Table 5.1 - General atmospheric conditions for LEO re-entry

Phase	Altitude (km)	Mach Range	T_∞ (K)	p_∞ (Pa)	ρ_∞ (kg/m³)	V_∞ Range (m/s)
Free-Molecular / Transitional	120–200	25–35	350–1000	$10^{-6} – 10^{-4}$	$10^{-11} – 10^{-9}$	7500–11000
Hypersonic Continuum	75–120	20–30	186–210	0.3 – 4	$5\times10^{-6} – 6\times10^{-5}$	6000–8000
Peak Heating*	55–70	12–22	219–260	4 – 40	$8\times10^{-5} – 8\times10^{-4}$	3500–6500
Supersonic → Subsonic	20–40	<5	216–250	300 – 5500	0.004 – 0.09	300–1500

**This study is particularly interested in the Peak Heating phase and all CFD simulations will be run at boundary conditions that represent the ranges listed in table 5.1.*

5.1.1 Inlet Conditions

The upstream and outer farfield boundary is selected as the front facing edge of the fluid domain and the edges wrapping around the upper and lower farfield. This boundary is defined as a pressure far field in ANSYS Fluent. A pressure far field boundary condition is more appropriate than a velocity inlet for this study because the flow represents a hypersonic external aerodynamic problem. A velocity inlet directly forces inflow behavior everywhere it is applied, which can over constrain a curved or wrapped farfield boundary. In contrast, a pressure far field boundary imposes the freestream Mach number, static pressure, static temperature, species composition, and flow direction while allowing the compressible boundary treatment to behave consistently with the local flow solution.

For all cases, the freestream conditions are defined at an altitude of 65 km with corresponding values from table 5.1 and the International Standard Atmosphere (ISA). The gauge pressure is set to 10 Pa, assuming an operating pressure of 0 Pa, the static temperature to 231 K and the Mach number is set to 20. The freestream direction is prescribed in the positive x direction by setting the x component of the flow direction equal to 1 and the y and z components equal to 0. Since the pressure far field uses Mach number rather than directly prescribing velocity magnitude, the freestream velocity is computed internally by Fluent from the local speed of sound and gas composition. For molecular air at 231 K, this corresponds to an approximate freestream velocity of 6.1 km/s.

For high altitude conditions, a low freestream turbulence level is assumed. Consequently, the turbulent intensity is set to 1 percent and the turbulent viscosity ratio is set to 1 for all cases. These values provide a low turbulence freestream condition without artificially increasing the incoming turbulent viscosity. More information on the turbulence modeling approach is provided in Section 5.X.

The thermal condition for the pressure far field is defined by the freestream static temperature. For the baseline case, this value is set to 231 K. This temperature represents the undisturbed atmospheric temperature at the selected 65 km freestream condition and is imposed at the farfield boundary before the flow is compressed and heated through the shock.

The species setup is required because the CFD model includes species transport and finite rate chemistry, as discussed in Section 5.X. Undisturbed dry air is commonly approximated as approximately 79 percent nitrogen and 21 percent oxygen by mole fraction. However, Fluent uses species mass fractions unless the mole fraction option is explicitly enabled. Therefore, the freestream species are specified using mass fractions. For molecular air, the oxygen mass fraction is set to approximately 0.233 and the nitrogen molecule mass fraction is approximately 0.767. In the five species air model used in this study, one species may be treated as the dependent species and calculated automatically as the remaining mass fraction. If nitrogen molecule, N2, is the dependent species, the visible species inputs are set such that atomic nitrogen, atomic oxygen, and nitric oxide are zero, while molecular oxygen is set to 0.233. Fluent then computes the remaining mass fraction as molecular nitrogen. This ensures that the freestream enters the domain as undissociated molecular air before high temperature reactions occur in the shock layer.

5.1.2 Outlet Conditions

The outlet is selected as the rear edge of the fluid domain downstream of the body and wake region. This boundary is defined as a pressure outlet. The outlet is placed sufficiently far downstream so that the wake and disturbed flowfield can develop before reaching the boundary. This reduces the likelihood that the outlet condition will influence the near body shock structure, wall heat flux, or wake behavior.

The conditions at the outlet must be set to match the conditions at the inlet to avoid boundary anomalies or excess backflow. The gauge pressure type is set to option static pressure

rather than total pressure because for this study the outlet is intended to represent the ambient downstream farfield pressure, rather than a stagnation reservoir or imposed total pressure condition. Maintaining the same conditions from section 5.1.1 the gauge pressure is 10 Pa and the backflow thermal condition is set to the freestream temperature of 231 K. The outlet species composition is also matched to the freestream molecular air composition, with molecular oxygen set to 0.233 by mass fraction, atomic species set to zero, and molecular nitrogen treated as the remaining mass fraction when it is the dependent species.

An outlet velocity is not directly prescribed. For a properly placed downstream outlet in a hypersonic external flow simulation, the flow should leave the domain without requiring strong backflow correction. However, backflow settings must still be physically consistent in case local recirculation or numerical backflow occurs near the outlet. The backflow pressure is specified as static pressure, and the backflow direction is treated using the normal to boundary option. This allows the outlet to remain consistent with the downstream boundary treatment without imposing an artificial streamwise velocity condition.

5.1.3 Wall Conditions

The representative body surfaces are assigned the wall boundary type. The wall surfaces are split into groups corresponding to the compression surface, aftbody surface, and any artificial truncation surfaces used to close the simplified geometry. The solid TPS stack is not explicitly modeled inside Fluent. This is done to reduce computational cost and maintain compatibility with the ANSYS Student license limitations. Instead, Fluent is used to resolve the external hypersonic flowfield and compute the surface heat flux applied to the wall.

The wall is modeled using a fixed temperature boundary condition during the CFD calculation. This fixed wall temperature represents an assumed TPS outer surface temperature for the external flow solution. Physically, the TPS surface is expected to approach a radiative equilibrium condition, where the aerodynamic heat input to the surface is balanced by heat rejected through surface radiation and heat conducted into the TPS stack. In simple terms, radiative equilibrium means that the wall temperature adjusts until the incoming convective and chemical heating is balanced by radiation emitted from the hot surface and by internal heat transfer through the material stack.

The full radiative equilibrium and through thickness TPS response are not solved directly in Fluent because the solid multilayer wall is not included in the CFD domain. Instead, the wall temperature is imposed as a fixed value for the CFD run, and the resulting surface heat flux is later passed to the thermal model. The thermal model and MATLAB solver are then used to update the wall temperature based on the energy balance between aerodynamic heating, surface radiation, and conduction into the ATPS stack. This calibration procedure is discussed in detail in Chapter 6. The fixed wall temperature in Fluent should therefore be interpreted as an iterative boundary condition used to couple the CFD and thermal analysis, rather than as a final prescribed material temperature.

The compression surface is treated as a partially catalytic wall. A partially catalytic wall allows some of the dissociated atomic species reaching the wall to recombine into stable molecules, but does not force complete recombination. This is important for high temperature hypersonic flow because the shock layer can contain chemically reacting species such as atomic nitrogen, atomic oxygen, and nitric oxide. Wall catalycity affects the chemical energy released or absorbed near the surface and can therefore influence the predicted wall heat flux. A fully non-catalytic wall may underpredict recombination effects, while a fully catalytic wall may overpredict the amount of recombination heating for many TPS materials. The partially catalytic assumption provides a more representative treatment for reusable thermal protection materials. The diagram in Figure 5.1 adapted from chapter 18 of the ANSYS Fluent User Guide for release 2025 R2 below provides an illustration for this concept.

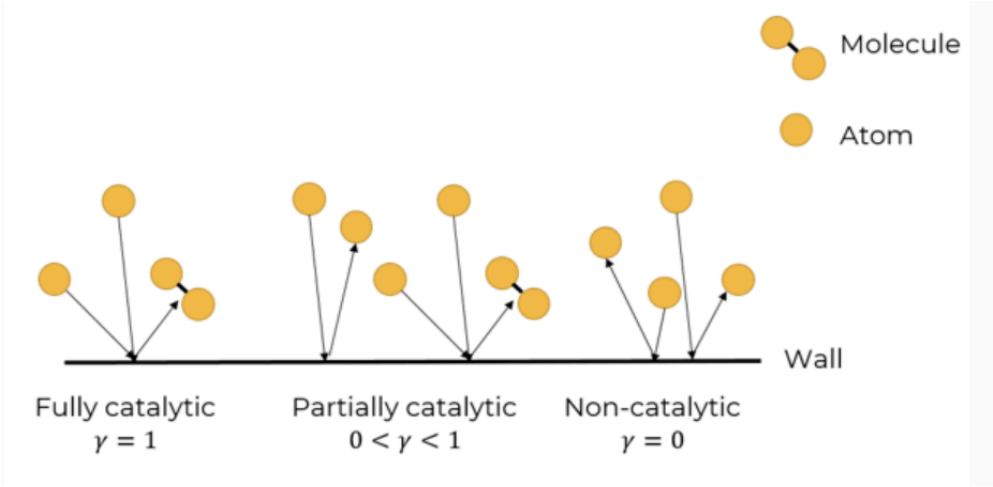


Figure 5.1 – Wall catalycity illustration. Reproduced from ref. [49].

Fluent provides several wall recombination models for thermal protection applications. In this study, the Recombination Model for Thermal Protection Systems is used. This model includes multiple submodels for catalytic recombination behavior and allows the wall recombination coefficients for atomic nitrogen and atomic oxygen to be specified using temperature dependent correlations. These submodels are shown in Figure 5.2.

$$Kolodziej: \begin{cases} \gamma_N = 6.1\times10^{-2} e^{\frac{-2480}{T}}, & 1410 < T < 1640K \\ \gamma_N = 6.1\times10^{-4} e^{\frac{-5090}{T}}, & 1640 < T < 1905K \\ \gamma_O = 40 e^{\frac{-11440}{T}}, & 1435 < T < 1580K \\ \gamma_O = 39\times10^{-9} e^{\frac{21410}{T}}, & 1580 < T < 1845K \end{cases}$$

$$Zoby: \begin{cases} \gamma_N = 0.0714 e^{\frac{-2219}{T}}, & 950 < T < 1670K \\ \gamma_O = 0.00941 e^{\frac{-658.9}{T}}, & 800 < T < 1400K \end{cases}$$

$$Scott: \begin{cases} \gamma_N = 0.0714 e^{\frac{-2219}{T}}, & 950 < T < 1670K \\ \gamma_O = 16 e^{\frac{-10271}{T}}, & 1400 < T < 1650K \end{cases}$$

Figure 5.2 – Partially catalytic models for TPS. Reproduced from ref. [49].

The selected catalytic model must be consistent with the imposed wall temperature range. If the fixed wall temperature is significantly outside the validity range of a given recombination correlation, the accuracy of the wall species boundary condition can decrease. For this reason, the catalycity model is checked against the wall temperature used for each CFD case. As the wall temperature is updated through the CFD and thermal calibration process, the applicable catalytic recombination correlation must remain valid for the final surface temperature range.

Recall that the objective of the CFD portion of this study is to obtain a total heat transfer rate from the fluid to the compression surface. This is given in units of J/s or Watts. Additionally the TPS material will not be modeled in fluent thus an approximate wall temperature must be computed before running the simulation. This effectively becomes an initial guess of the temperature on the surface of the first layer of the TPS stackup which later must be calibrated by the results of the CFD and Thermal Analysis. See chapter 6 for calculation steps and details on calibration methodology and MATLAB Solver.

The aftbody and artificial truncation surfaces are also assigned wall boundary conditions. For surfaces representing the same exposed TPS material, the same fixed wall temperature and species wall treatment are applied for consistency. Artificial truncation surfaces are not used to define the primary heat flux results, but they are still assigned physically consistent wall conditions so that the surrounding wake solution is not disrupted by an inconsistent thermal or species boundary condition. The primary reported heat flux values are taken from the relevant

exposed aerodynamic surfaces, such as the forebody compression surface or leading edge surface, depending on the geometry.

5.2 Model Physics Setup

This section describes the physics setup for both geometries in ANSYS Fluent. The set up is generally consistent with the guidelines in chapter 18 of the 2025 R2 release of the ANSYS Fluent User guide [49]. Both geometries share the same physics set up. Table 5.2 below lists all the settings developed in this section for summary.

Table 5.2 - Physics set up ANSYS Fluent

Category	Setting / Option	Value / Choice	Notes
Solver type	Solver	Density-based or pressure-based compressible setup depending on your Fluent case	Density based solver preferred for hypersonic and high speed flow regimes.
Analysis type	Flow regime	Hypersonic external flow	($M_\infty = 20$)
Turbulence Model	Model	k-omega-SST	High fidelity model very compatible for viscous effects and high speed flow
Turbulence Model	Compressibility effects	ON	Turned on for compressible flows M > 0.8
Gas model	Material / mixture	Air-5 species Park 93	N, N2, O, O2, NO
Energy	Energy Equation	ON	Required for hypersonic compressible reacting flow.
Energy	Two-temperatur e model	OFF / unavailable	Student license does not allow it. Finite rate chemistry with TPG assumption.
Species transport	Species Transport	ON	Required for finite-rate chemistry / dissociation.

Chemistry model	Finite-rate chemistry	ON	Non-equilibrium finite-rate chemistry through species transport,
Operating pressure	Operating pressure	0 Pa	Recommended so gauge pressure equals absolute static pressure for high altitude hypersonic setup.
Outer arc boundary/Inlet	Boundary type	Pressure far-field	Avoids forcing a free flow condition at the boundary and creating compression in the far field
Pressure far-field Boundary	Gauge pressure	10 Pa	Assuming operating pressure = 0 Pa.
Pressure far-field Boundary	Mach number	20	Defines freestream speed through Mach and mixture speed of sound.
Pressure far-field Boundary	Direction vector	(x=1, y=0, z=0)	Uniform freestream in +x direction.
Pressure far-field Boundary	Static temperature	231 K	Static freestream temperature.
Pressure far-field Boundary	Turbulent intensity	1%	low turbulence high altitude freestream assumption.
Pressure far-field Boundary	Turbulent viscosity ratio	1	Conservative low freestream turbulence level.
Pressure far-field species	Species input type	Mass fractions	Computes species based on mass fractions

Pressure far-field species	(Y_{O_2})	0.233	Molecular dry air by mass fraction.
Pressure far-field species	(Y_{N})	0	Atomic nitrogen should be zero in freestream.
Pressure far-field species	(Y_{O})	0	Atomic oxygen should be zero in freestream.
Pressure far-field species	(Y_{NO})	0	NO should be zero in freestream.
Pressure far-field species	(Y_{N_2})	0.767, dependent/remainder	Fluent computes value as the remaining mass fraction.
Downstream boundary	Boundary type	Pressure outlet	Right side outlet is separate from the outer arc.
Pressure outlet boundary	Gauge pressure	10 Pa	Same freestream static pressure if operating pressure = 0 Pa.
Pressure outlet boundary	Backflow pressure	Static pressure	Use static pressure, not total pressure.
Pressure outlet boundary	Direction method	Normal to boundary	Safer than manually forcing (+x).
Pressure outlet boundary	Temperature	231 K	Freestream static temperature.
Pressure outlet species	Backflow composition	Same as freestream molecular air	(Y_{O_2}=0.233), (Y_{N_2}=0.767), atoms/NO = 0.
Body boundary	Boundary type	Wall	Forebody and aftbody wall zones.
Wall thermal condition	Thermal BC	Fixed wall temperature (Isothermal)	Externally calibrated in MATLAB. Same for forebody and aftbody.

			Caused issues discussed later in chapter (see discussion)
Wall species / catalycity	Catalytic wall	Recombination Model for TPS	Important because wall recombination can affect heat flux in reacting flow.
Wall roughness	Roughness	Default / smooth	No roughness added to wall.
Methods/ Controls	Courant Number (CFL)	Startup: 1 Post render: 3	Switch from 1 to 3 after ~ 1000 iterations or first calibration.
Methods/ Controls	Gradient Method	Least Squares Cell-Based	Recommended over structured / mapped meshes
Methods/ Controls	Discretization	Second order upwind for flow, energy, species, and pressure	Can be reduced to first order for fast convergence during startup.
Methods/ Controls	High Speed Numerics	ON	Must be turned on for hypersonic flow
Initialization	Initialization source	Compute from Inlet (pressure far-field)	Computing initial values from the freestream conditions
Initialization	Initialization Type	Full Multi-Grid (FMG)	High fidelity initialization to aid with convergence for hypersonic flow
FMG Settings	Levels	5	Iterate on 5 levels
FMG Settings	CFL	0.5	Low courant number for convergence
FMG Settings	Cycles	Lvl1=200, lvl2=400,lvl3=800,lvl4=1600,lvl5=1600	Number of iterations per level
FMG Settings	Species Reactions	ON	Species reactions rendered during initialization

FMG Settings	Viscous Calculations	ON	Viscous calculations computed during initialization

5.2.1 General Physics Settings

The CFD simulations are performed in ANSYS Fluent using a density based steady state solver. A density based formulation is selected because the flowfield is hypersonic and highly compressible, with a freestream Mach number of 20. At this Mach number, density, pressure, temperature, and velocity are strongly coupled through the shock layer, and the solver must resolve large gradients across the shock and near wall boundary layer. The steady state assumption is used because the objective of the CFD analysis is to obtain representative peak heating surface conditions for comparison with the thermal protection system model rather than to resolve a time dependent entry trajectory.

The energy equation is enabled for all simulations. This is required because the flow contains strong compressibility effects, shock heating, wall heat transfer, and high temperature gas reactions. The energy equation allows Fluent to resolve the conversion of kinetic energy into internal energy through the shock and to compute the temperature field used for wall heat flux prediction. A one temperature formulation is used in this study. The two temperature model is not included because it is unavailable under the ANSYS Student license used for this work. Therefore, the translational, rotational, vibrational, and electronic energy modes are not solved as separate temperatures. This is a limitation for high enthalpy nonequilibrium hypersonic flow, but the one temperature model still allows finite rate chemical nonequilibrium to be included through the species transport formulation.

5.2.2 Energy Equation and One Temperature Assumption

The energy equation is required for the present simulations because temperature is one of the primary quantities of interest. The predicted wall heat flux depends on the local wall normal temperature gradient, while the chemical reaction rates and mixture properties also depend strongly on temperature. The energy equation is therefore enabled for every CFD case.

For hypersonic reacting flow, a more complete model would solve separate thermal energy modes using a two temperature formulation. This can be important when the vibrational and electronic energy modes are not in equilibrium with the translational and rotational modes. However, the present study uses the one temperature approach available in the ANSYS Student version. Under this assumption, all energy modes are represented by a single static temperature. The model therefore captures shock heating, wall heat transfer, and finite rate chemical composition changes, but does not separately resolve thermal nonequilibrium between vibrational and translational modes.

This limitation is acceptable for the scope of the present study because the CFD results are used primarily to compare surface heat flux trends between simplified geometries and to

provide external heating inputs to the thermal model. The absence of a two temperature model is acknowledged as a modeling limitation, especially because the freestream condition is hypersonic and the post shock gas temperature can be high enough for dissociation reactions to occur.

5.2.3 Turbulence Modeling

Turbulence is modeled using the k omega SST model with compressibility effects enabled. The SST model is selected because it combines the near wall behavior of the k omega formulation with improved outer layer behavior, making it suitable for wall bounded compressible flow with strong adverse pressure gradients. The model is also commonly used for aerodynamic heating problems because it provides stable near wall resolution when the mesh is refined to a low y plus value.

The freestream turbulence quantities are specified through the pressure far field boundary using a turbulence intensity of 1 percent and a turbulent viscosity ratio of 1. These values represent a low turbulence high altitude freestream and avoid introducing excessive turbulent viscosity at the inlet. The mesh is designed to resolve the near wall region directly with a target y plus of approximately unity, which is consistent with the use of the SST model for wall heat flux and wall shear stress prediction.

Compressibility effects are included because the flow is hypersonic and the turbulence model must account for high speed density variations and dilatation effects. The addition of compressibility corrections is intended to improve the treatment of turbulent kinetic energy behavior in regions with strong shocks and high Mach number gradients. The turbulence model is therefore used to provide an engineering level prediction of wall shear stress and heat transfer while maintaining a solver setup that is practical for the available computational resources.

5.2.4 Species Transport and Chemical Nonequilibrium

Species transport is enabled to model the finite rate dissociation and recombination of high temperature air. In hypersonic flow, the gas passing through the shock layer can reach temperatures high enough for oxygen and nitrogen molecules to dissociate. The ANSYS Fluent user guide describes chemical nonequilibrium in terms of the Damkohler number, which compares the characteristic flow time to the characteristic chemistry time [49]. When this value is near unity, the chemistry does not respond instantly to the flowfield and the gas is locally in chemical nonequilibrium. This condition is common in hypersonic shock layers, where reaction rates and convective time scales can be comparable.

The flow is therefore modeled using species transport with volumetric reactions and finite rate chemistry. Fluent identifies finite rate chemistry as a suitable approach for modeling chemically nonequilibrium hypersonic flow, provided that a species set and chemical mechanism are defined. For Earth atmospheric flight or entry below approximately 10,000 K, Fluent provides the predefined air 5 species Park 93 mixture. This mixture contains molecular nitrogen,

molecular oxygen, nitric oxide, atomic nitrogen, and atomic oxygen, along with the Park 93 chemical mechanism.

The air 5 species Park 93 mixture is used in this study because the expected flowfield temperature is high enough for dissociation to be relevant, but ionization is not the primary focus of the model. Fluent documentation indicates that when the flow temperature exceeds approximately 10,000 K, weak ionization may become important and an 11 species ionized air mixture may be more appropriate[49]. Since the present simulations focus on neutral air dissociation and surface heating trends, the 5 species model provides a practical balance between physical fidelity and computational cost.

The finite rate chemistry model is solved using volumetric reactions. This allows the local gas composition to evolve through the shock layer and along the body as a function of temperature, pressure, and reaction rate. The freestream composition is initialized as undissociated molecular air, with molecular oxygen and molecular nitrogen specified using mass fractions. Dissociated species are initialized as zero at the farfield boundary and are generated only where the local thermochemical state produces finite rate reactions in the flow.

5.2.5 Transport Properties

Temperature dependent transport properties are required because viscosity and thermal conductivity vary strongly through the shock layer. In the one temperature hypersonic modeling approach, Fluent uses Gupta curve fits as the default property method for air mixtures such as air 5 species Park 93. The Fluent user guide states that Gupta curve fits are applicable for hypersonic one temperature simulations up to 30,000 K and can be used for species viscosity and thermal conductivity [49].

The Gupta curve fit method provides temperature dependent species viscosity and thermal conductivity values using coefficient based curve fits. This is important because the local heat flux depends on thermal conductivity and the near wall temperature gradient, while wall shear stress depends on the local viscosity and velocity gradient. Using temperature dependent transport properties allows the gas mixture properties to respond to the high temperature shock layer rather than being treated as constants.

For the air 5 species Park 93 mixture, the material properties are therefore left consistent with the Fluent predefined hypersonic air model, using Gupta curve fit viscosity and Gupta curve fit thermal conductivity. This maintains consistency between the selected chemical mechanism, the one temperature energy model, and the transport property model.

5.3 Solution Methods, Controls, and Initialization

The solution controls were selected to improve stability for hypersonic reacting flow while maintaining sufficient accuracy for wall heat flux prediction. Hypersonic CFD contains strong coupling between density, pressure, temperature, species composition, and transport properties. The detached shock, high temperature shock layer, and near wall gradients can

produce convergence difficulties if the solution is advanced too aggressively. For this reason, the solver was initialized conservatively and then gradually advanced toward the final second order solution settings.

5.3.1 Solver Controls

The simulations use the density based implicit solver with high speed numerics enabled. The density based formulation is appropriate for high Mach number compressible flow because it solves the coupled conservation equations in a form suitable for shock dominated flowfields. High speed numerics were enabled to improve solver robustness in the presence of strong shocks and large pressure and temperature gradients. This setup is consistent with Fluent best practice guidance for hypersonic flow, which recommends the density based implicit solver and enabling high speed numerics for these types of simulations [49].

The Courant number was initially set to 1 for at least the first 1000 iterations. This value is lower than the Fluent default value of 5 and was used to prevent the solution from developing too aggressively during the early shock formation and species reaction stages. After the primary shock structure, wall heat flux, and residual behavior became more stable, the Courant number was increased to 3 to improve convergence rate while still remaining below the default value. Fluent guidance notes that lowering the CFL number from 5 to 1 can help stabilize difficult hypersonic reacting flow calculations [49].

The positivity rate limit was reduced to 0.1. This was done to improve numerical stability and prevent nonphysical updates in pressure, density, temperature, or species fields during the early development of the solution. Fluent best practices specifically identify reducing the positivity rate limit from the default value of 0.2 to 0.1 or 0.05 as a stabilization approach for difficult hypersonic calculations [49].

5.3.2 Spatial Discretization and Gradient Evaluation

The gradient evaluation method was set to Least Squares Cell Based. This method was selected because the final meshes are structured quadrilateral mapped meshes with strong alignment near the body and wake regions. The cell layout was designed to maintain high quality and smooth growth away from the wall, so the least squares cell based method provides a stable and efficient gradient calculation for the present mesh topology.

Second order spatial discretization was used for the final solution because the CFD results are used to evaluate wall heat flux, wall shear stress, shock structure, and near wall temperature gradients. First order discretization may be used during the initial part of the calculation to help establish the main shock and flowfield structure. Once the flow is developed and the residuals and monitored quantities become stable, the solution should be switched to second order to reduce numerical diffusion and improve the accuracy of the final flowfield. This approach is consistent with Fluent best practice guidance, which recommends starting with first order and switching to second order once the flow is developed if convergence difficulties are encountered [49].

5.3.3 Chemistry Solver Controls

Finite rate chemistry is included through species transport with volumetric reactions. For the reacting flow cases, the chemistry solver is set to None - Direct Source. This option was selected because the density based solver is used and the chemistry source terms are strongly coupled to the species transport equations. Fluent guidance recommends the None - Direct Source chemistry solver for reacting hypersonic cases using the density based solver [49].

The chemical source terms are solved with the air 5 species Park 93 mixture described in Section 5.2. The freestream is initialized as undissociated molecular air, while dissociation and recombination are allowed to develop inside the shock layer and near the wall according to the finite rate volumetric reactions. This approach allows the species field to evolve with the local temperature and pressure field rather than forcing a frozen or equilibrium composition throughout the domain.

5.3.4 Full Multigrid Initialization

Full Multigrid initialization was used after standard initialization to provide a more developed initial flowfield before beginning the main steady state calculation. This is especially useful for hypersonic simulations because the solution contains strong discontinuities and tightly coupled thermochemical behavior. Fluent guidance identifies Full Multigrid initialization as a best in class initialization approach for hypersonic flow because it produces an improved starting solution and can allow the solver to begin more robustly [49].

The FMG initialization was performed using a CFL number of 0.5. This is lower than the default FMG CFL value and was selected to increase robustness during initialization. Fluent guidance notes that the FMG CFL may be lowered to 0.5 or 0.25 for hypersonic cases to improve initialization stability [49].

Five FMG levels were used. The iteration schedule was increased with each level so that the coarse level solution could develop before being transferred to finer levels. The FMG schedule was defined as 200 iterations on level 1, 400 iterations on level 2, 800 iterations on level 3, and 1600 iterations on levels 4 and 5. The increased number of iterations was used to compensate for the reduced FMG CFL and to improve the development of the shock structure, temperature field, and reacting species field before the main calculation.

Both viscous FMG and species reaction FMG were enabled. Viscous FMG was included so that the initialization could begin developing boundary layer and wall heat transfer behavior along the isothermal walls. Species reaction FMG was included because the flow uses finite rate chemistry for air dissociation. Fluent guidance identifies these FMG options as useful for developing boundary layers, accounting for heat transfer effects along isothermal walls, and including finite rate reactions during initialization [49].

After FMG initialization, the initial solution was visually inspected before continuing the calculation. The pressure and temperature fields were checked to confirm that the main shock structure appeared in a realistic location and that the solution values were physically reasonable.

This follows Fluent guidance that the FMG initialized solution should be inspected to ensure that discontinuities are present in the expected locations and that pressure and temperature values are realistic before continuing the calculation [49].

5.3.5 Iteration Strategy and Calibration Cycles

The first CFD run in each wall temperature calibration sequence was advanced for 1000 iterations after initialization. This longer initial run was used to allow the shock layer, wall heat flux, wall shear stress, and species field to develop from the initialized state. After the first calibration run, subsequent CFD runs were advanced for 500 iterations for each wall temperature update. The reduced iteration count was used because later calibration cycles start from a more developed solution and primarily adjust the wall heat flux response to the updated fixed wall temperature.

The calibration procedure is described in detail in Chapter 6. In general, the CFD model provides the aerodynamic surface heat flux for a fixed wall temperature, while the thermal model and MATLAB solver update the wall temperature based on the balance between aerodynamic heating, surface radiation, and conduction into the ATPS stack. The CFD solution is then repeated with the updated wall temperature until the imposed wall temperature and the thermal model response are consistent.

5.3.6 Solution Monitoring

Convergence was evaluated using both residuals and physical monitors. For hypersonic reacting flow, residuals may not decrease smoothly because the solution includes shocks, strong gradients, species reactions, and tightly coupled energy behavior. Therefore, additional monitored quantities were used to judge whether the solution reached a stable steady state. Fluent best practices recommend monitoring aerodynamic loads, mass flux balance, total heat flux at isothermal walls, maximum pressure, maximum temperature, maximum Mach number, turbulent viscosity ratio, and reacting or produced species for hypersonic reacting cases.

For this study, the most important monitored quantities are total surface heat flux, wall shear stress, maximum static temperature, maximum Mach number, pressure field behavior, and species production in the shock layer. The surface heat flux monitor is especially important because the CFD solution is used as the aerodynamic heating input for the thermal model. A solution is considered suitable for post processing when the shock structure is stable, the heat flux distribution no longer changes significantly, and the monitored quantities approach steady values.

The final solution was also evaluated through contour inspection. Pressure and temperature contours were used to verify bow shock structure and shock layer development, while surface heat flux and wall shear stress contours were used to identify local heating behavior along the body. Fluent guidance recommends examining shock waves, post shock quantities, surface heat flux, species mixture, and wall resolution after obtaining a steady state

hypersonic solution. It also notes that shock fronts should appear smooth and regular and that wall heat flux contours should not show significant nonphysical oscillations [49].

5.4 Results

Section 5.4 presents the results of the CFD simulations. The desired outputs are as follows:

- Calibrated Total Heat Transfer Rate (W)
- Calibrated Average Wall Temperature (K)
- Surface Heat Flux (W/m^2) vs Y Position (m)
- Wall Shear Stress (Pa) vs Y Position (m)
- Contour of Static Temperature (K)
- Contour of Static Pressure (Pa)
- Contour of Velocity Magnitude (Mach)

Table 5.3 below shows the total heat transfer rates, that is heat lost from the fluid into the wall (values are negative), and average isothermal wall temperatures for the calibrated models.

Table 5.3 - Calibration results

Geometry	Total Heat Transfer Rate	Average Wall Temperature
Leading Edge (Blunt Tipped Wedge)	1.43×10^6 W 1.43 MW	1705.28 K
Blunt Re-entry Forebody (Axisymmetric Sphere Cone)	0.62×10^6 W 0.62 MW	1004.57 K

5.4.1 Hypersonic Leading Edge (Blunt Tipped Wedge)

Contour Plots

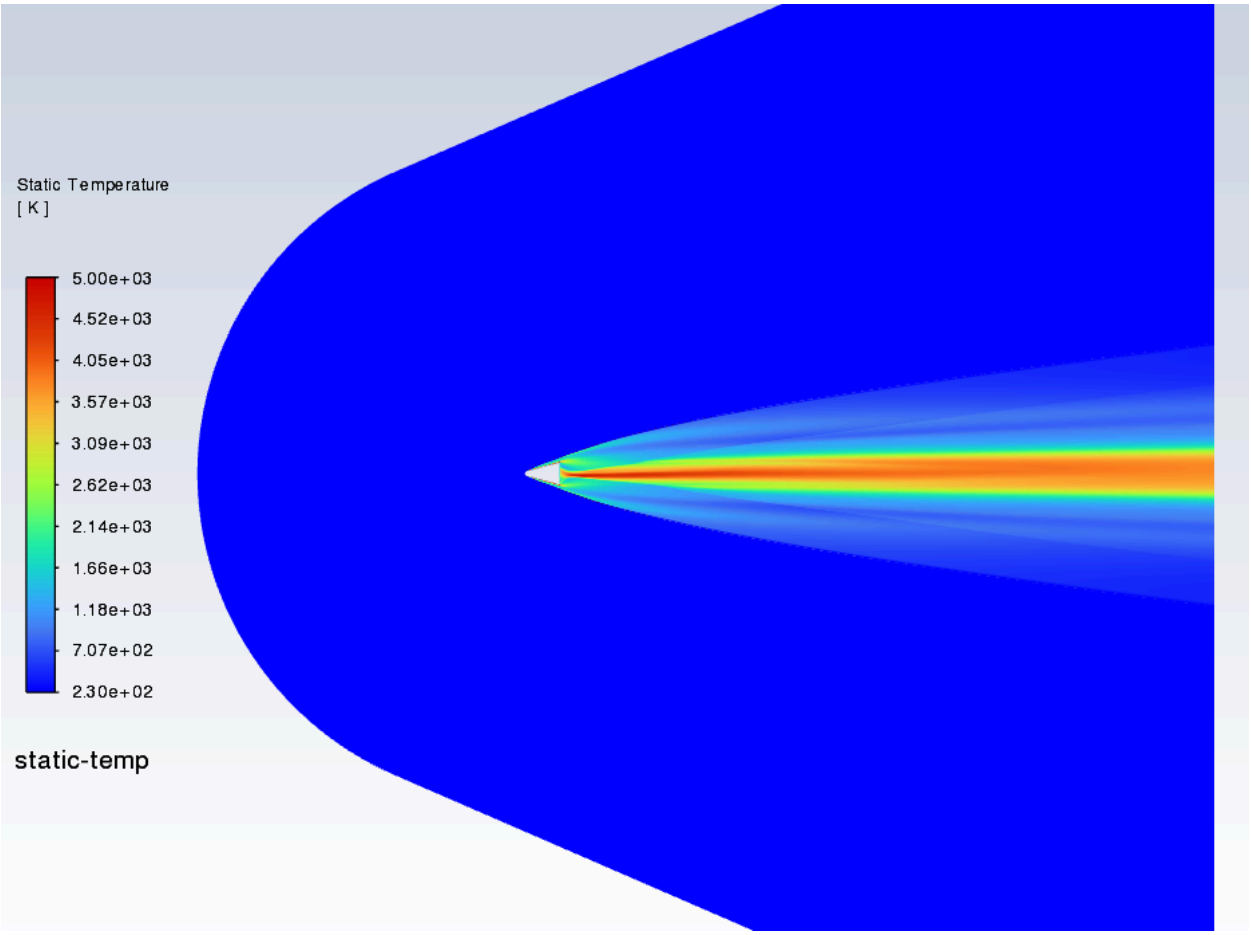


Figure 5.3 - Leading edge: contour of static temperature (full view)

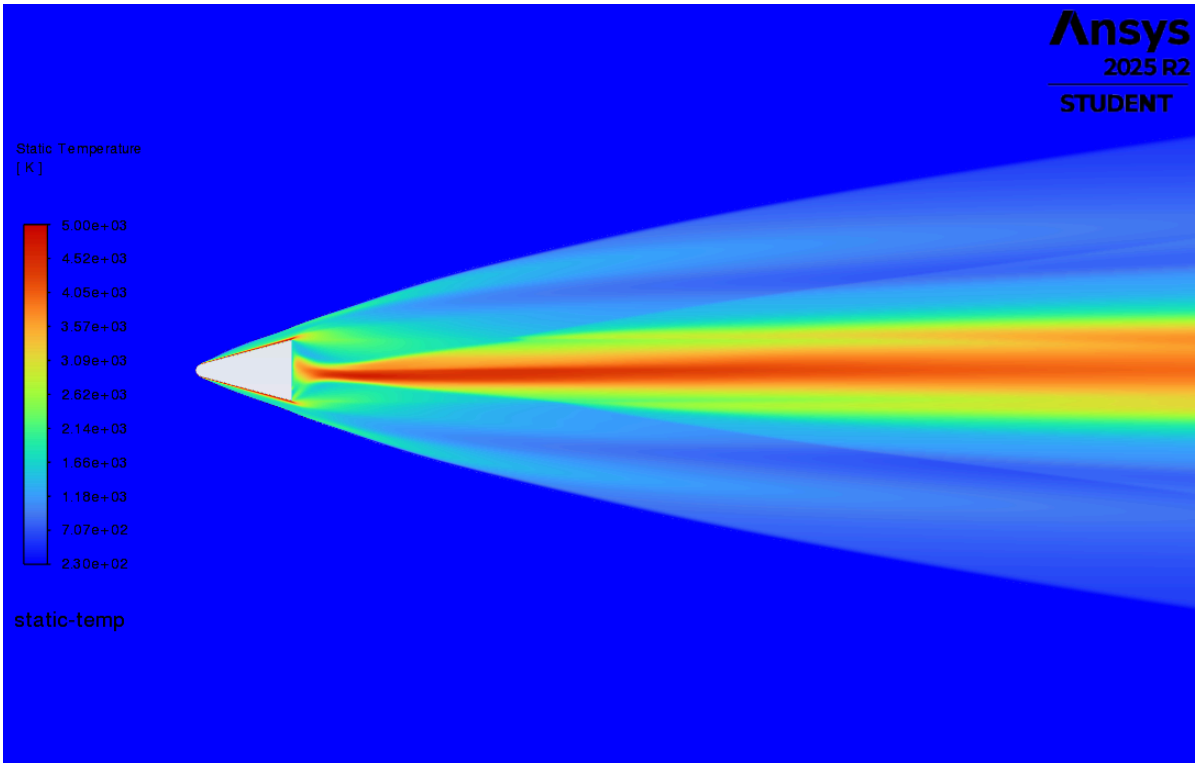


Figure 5.4 - Leading edge: contour of static temperature (nearfield)

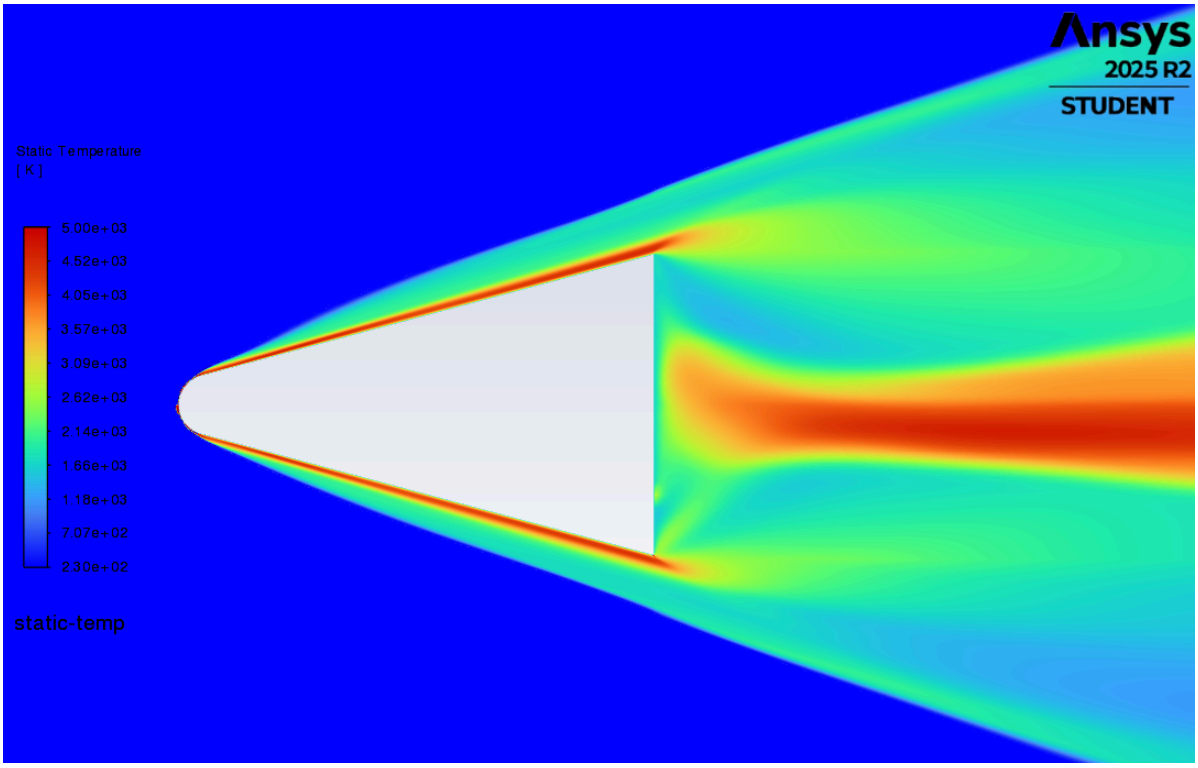


Figure 5.5 - Leading edge: contour of static temperature (near wall)

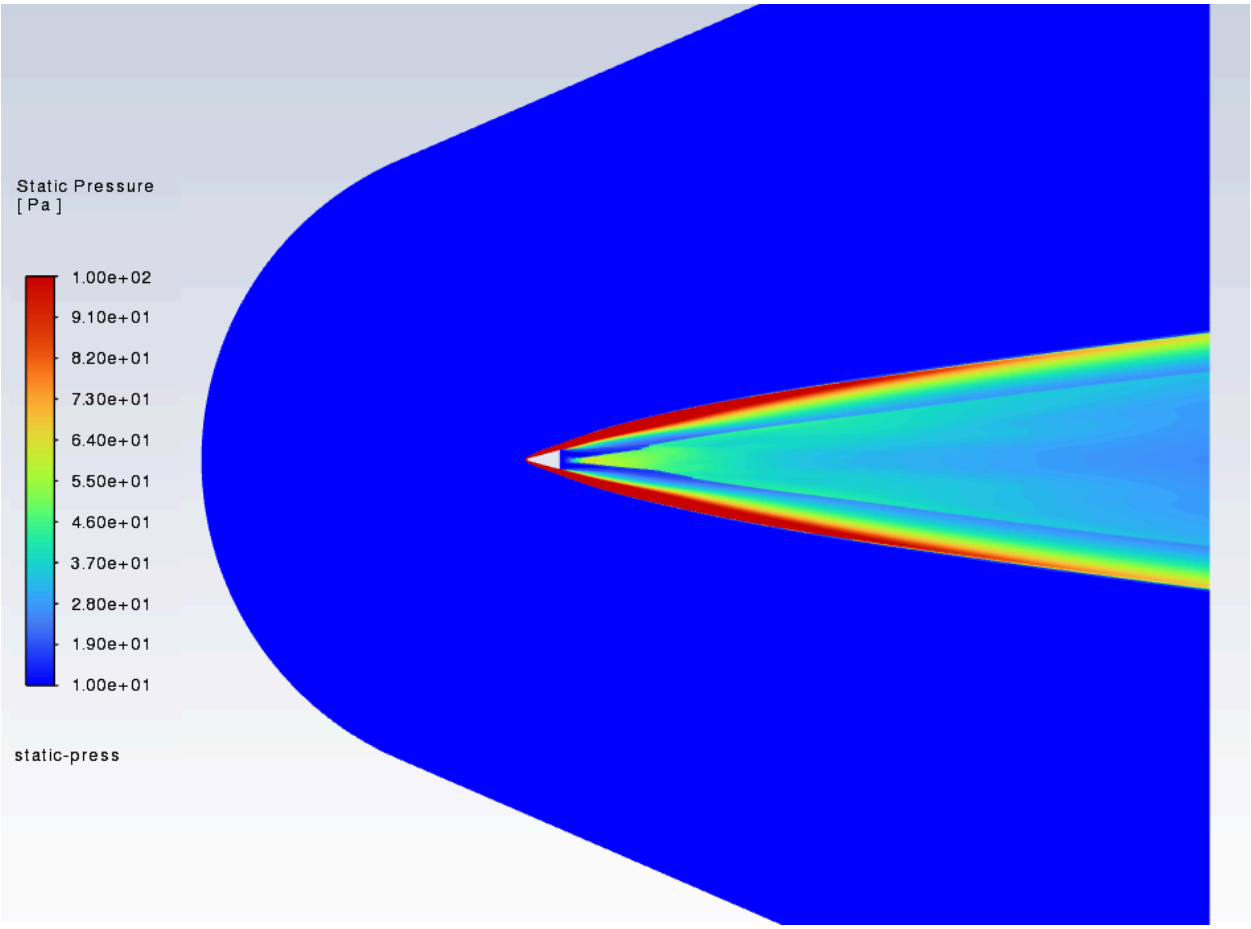


Figure 5.6 - Leading edge: contour of static pressure (full view)

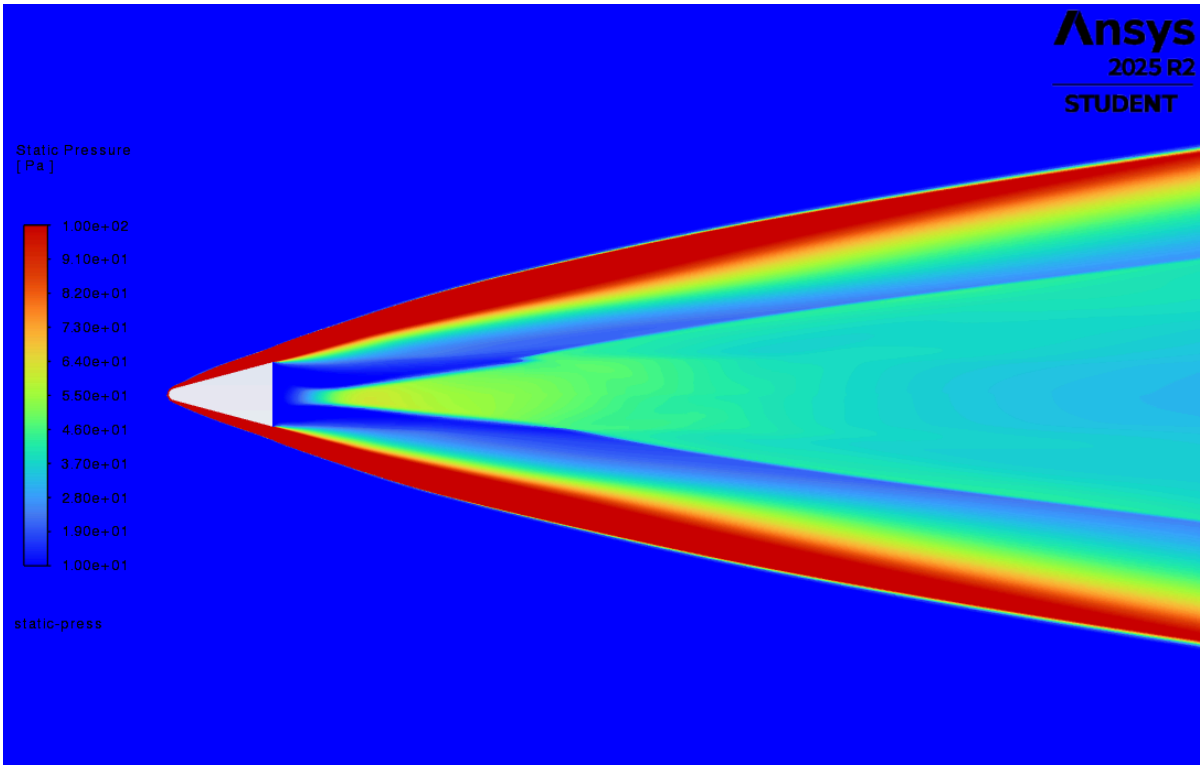


Figure 5.7 - Leading edge: contour of static pressure (nearfield)

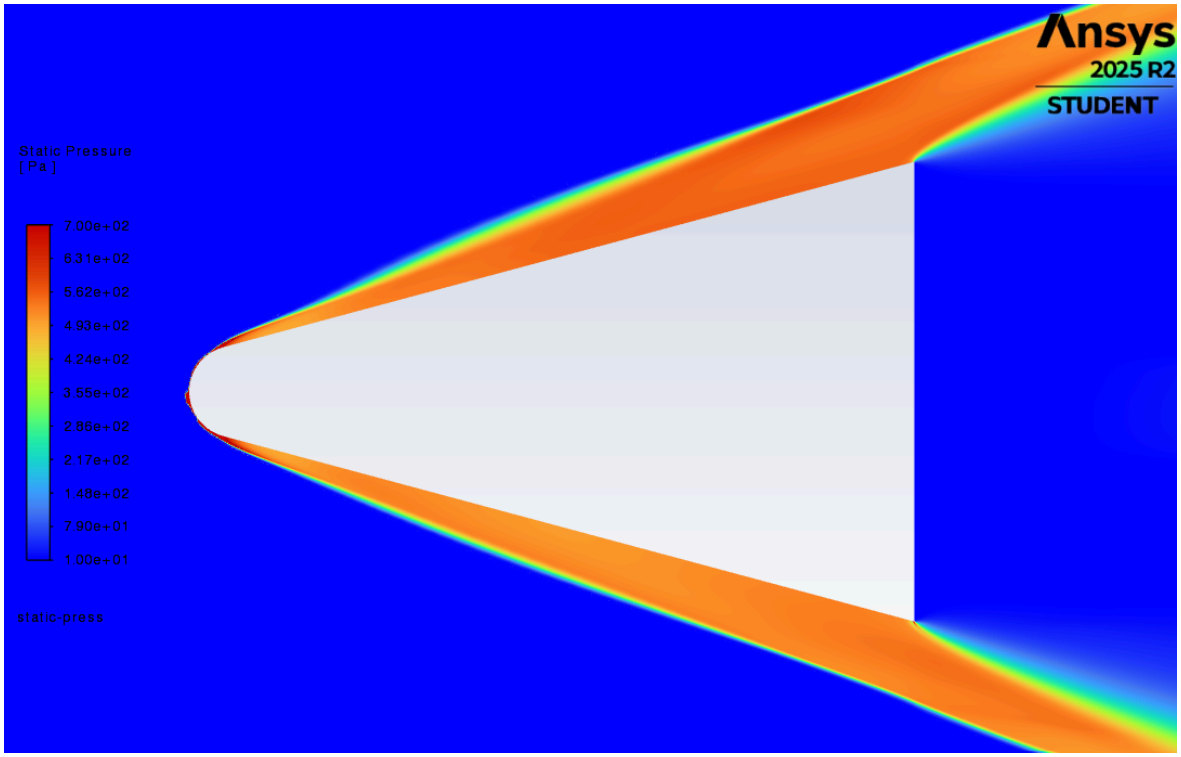


Figure 5.8 - Leading edge: contour of static pressure (near wall; altered for gradient)

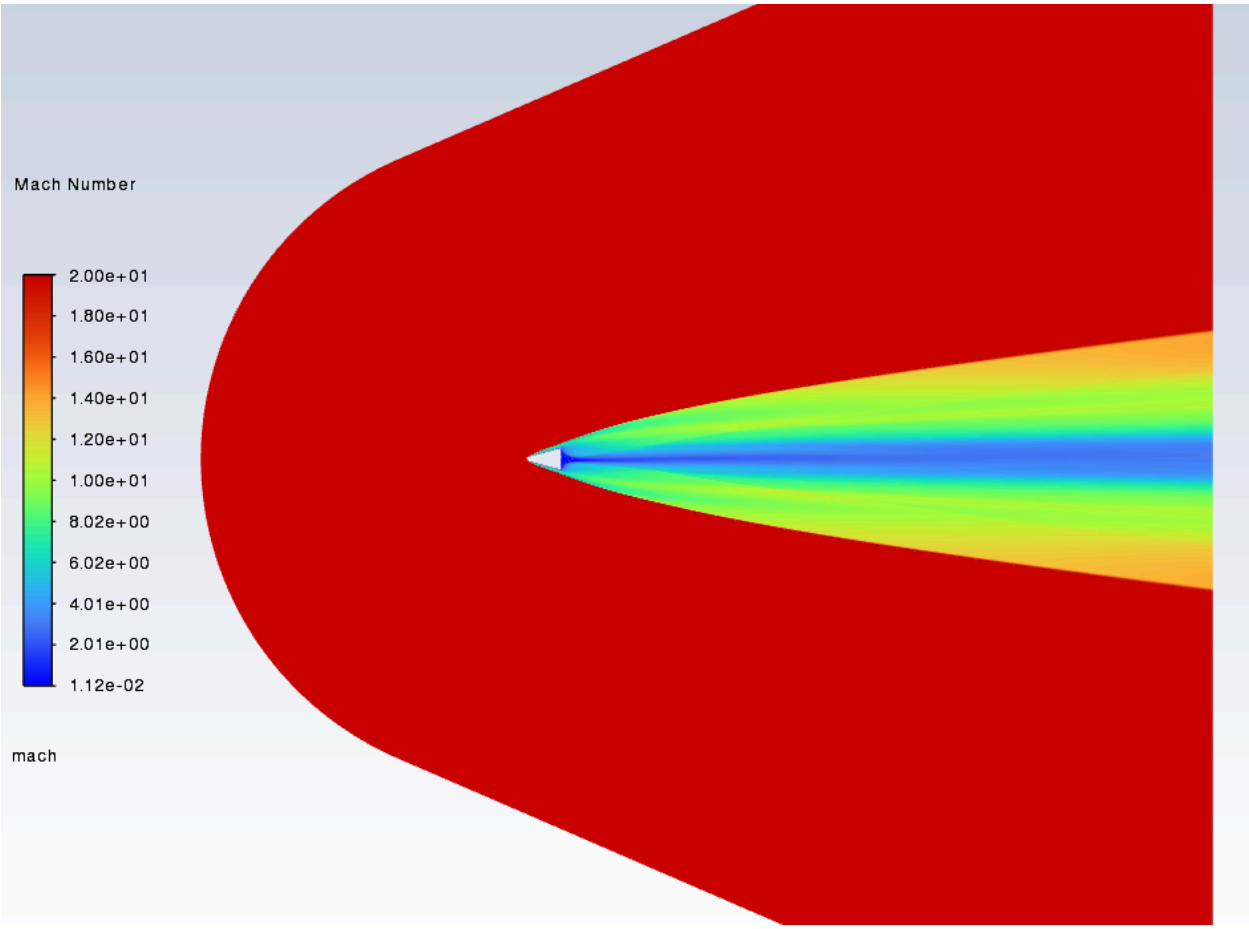


Figure 5.9 - Leading edge: contour of mach number (full view)

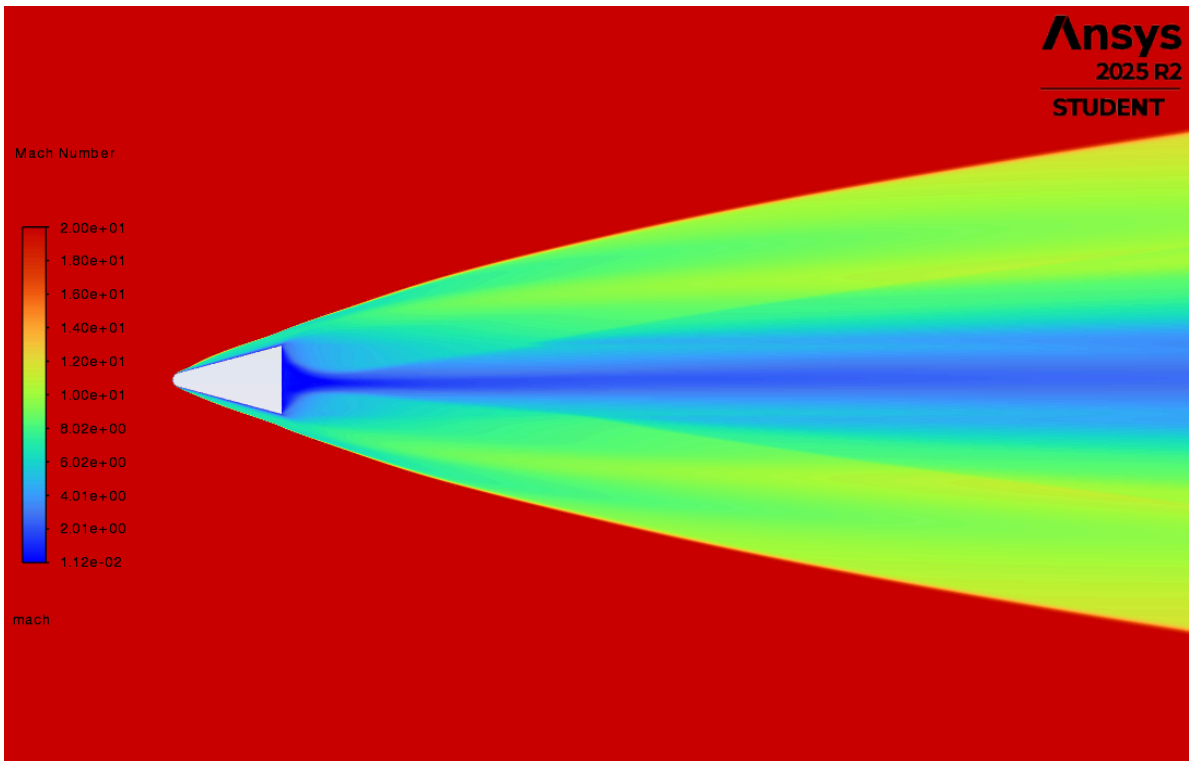


Figure 5.9 - Leading edge: contour of mach number (nearfield)

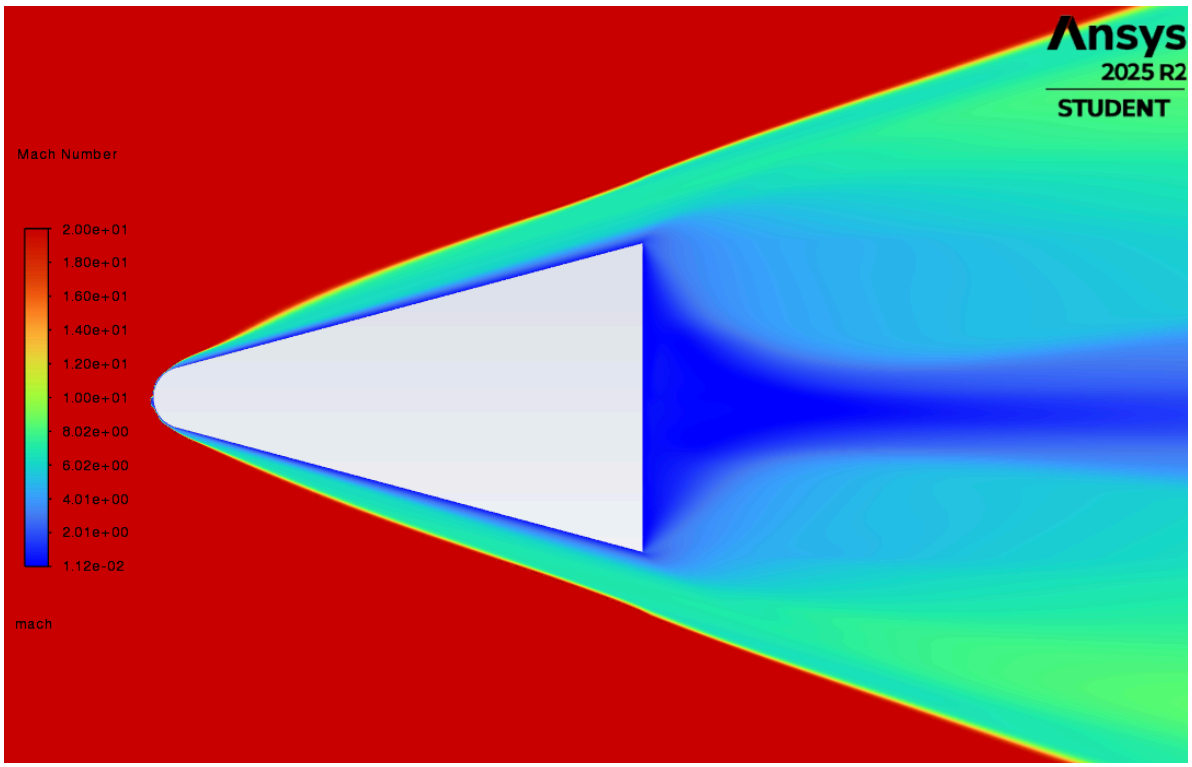


Figure 5.10 - Leading edge: contour of mach number (near wall)

XY Plots

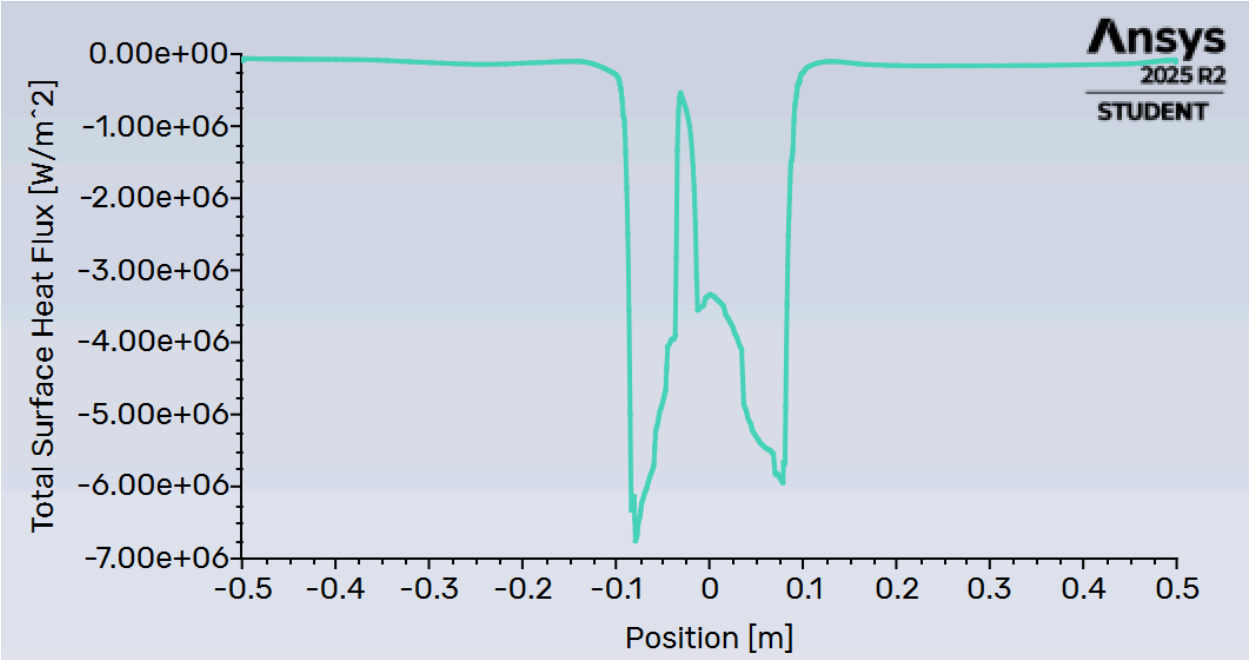


Figure 5.11 - Leading edge: front surface heat flux vs. y position

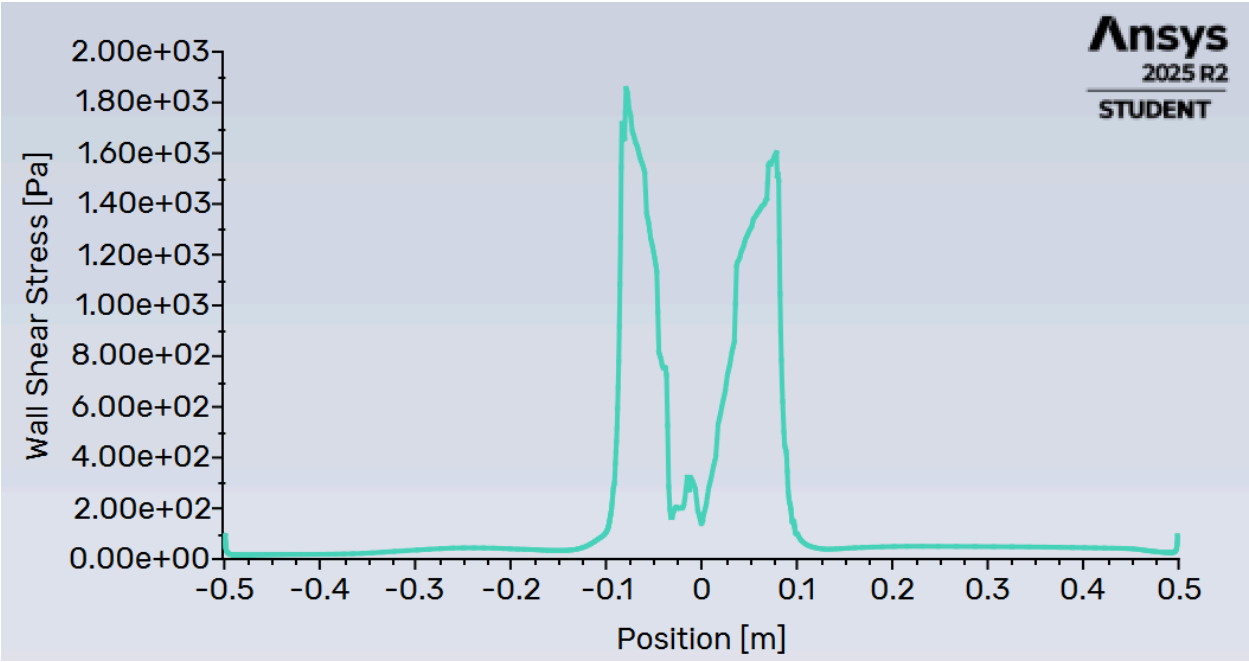


Figure 5.12 - Leading edge: front surface wall shear stress vs. y Position

5.4.2 Blunt Re-entry Forebody (Axisymmetric Sphere Cone)

Contour Plots

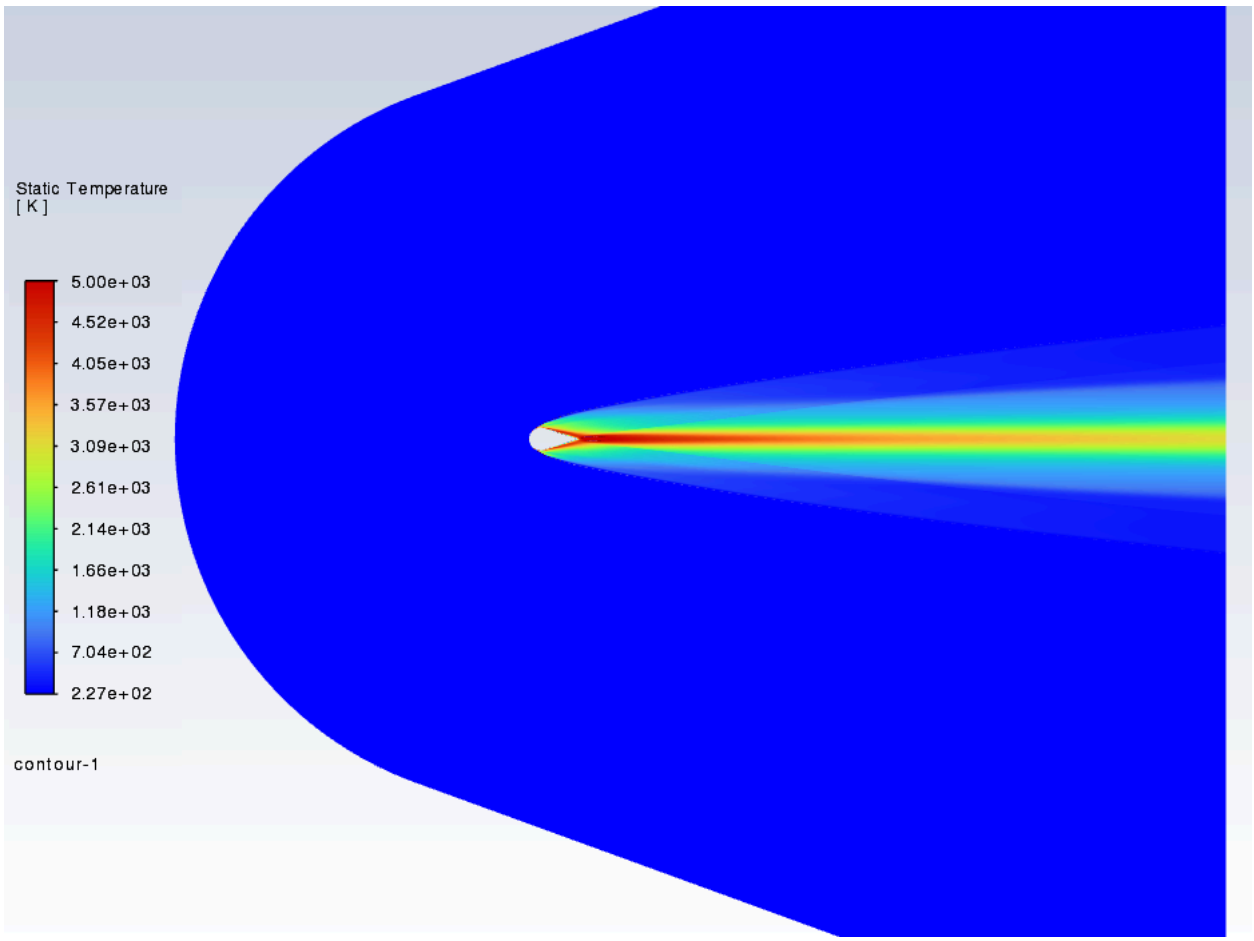


Figure 5.13 - Blunt forebody: contour of static temperature (full view)

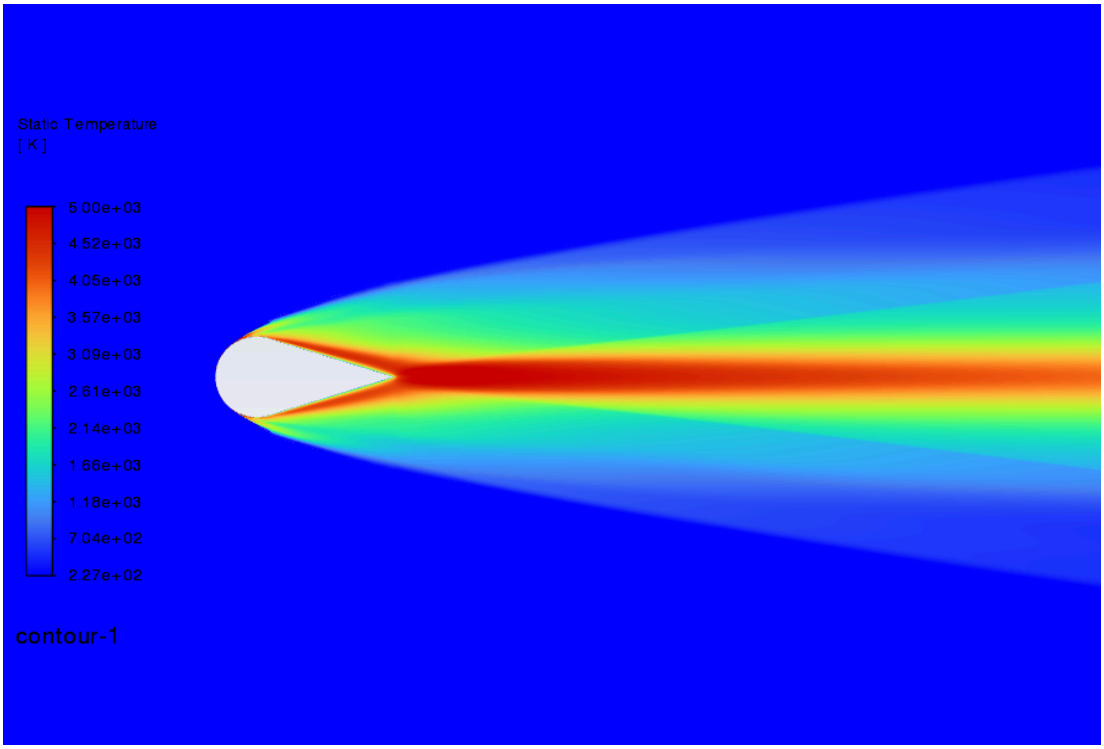


Figure 5.14 - Blunt forebody: contour of static temperature (nearfield)

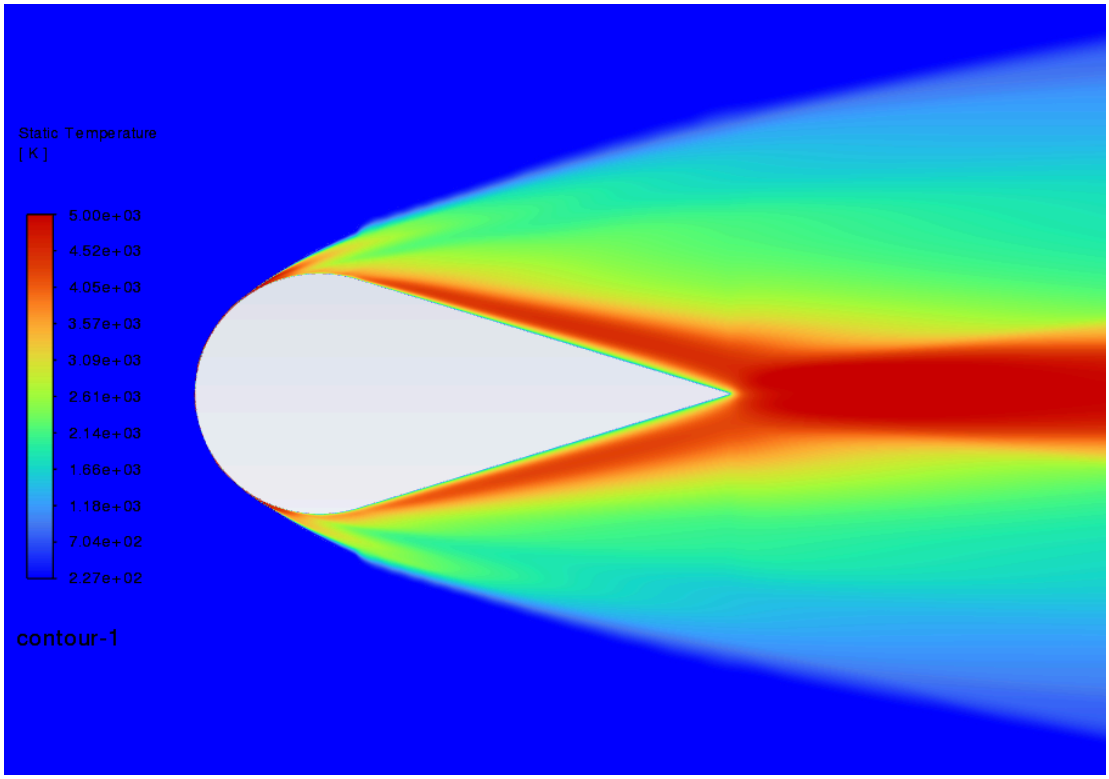


Figure 5.15 - Blunt forebody: contour of static temperature (near wall)

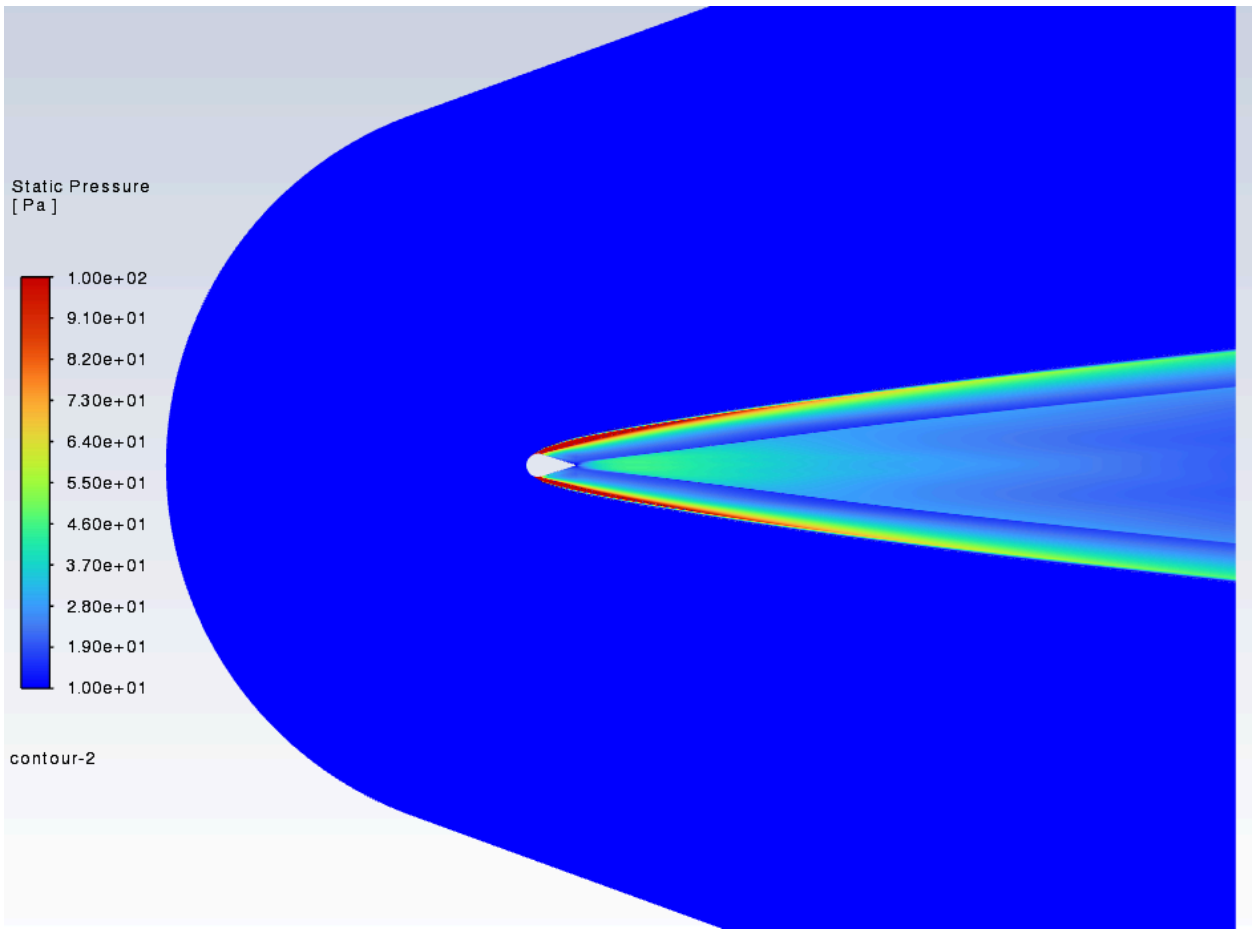


Figure 5.16 - Blunt forebody: contour of static pressure (full view)

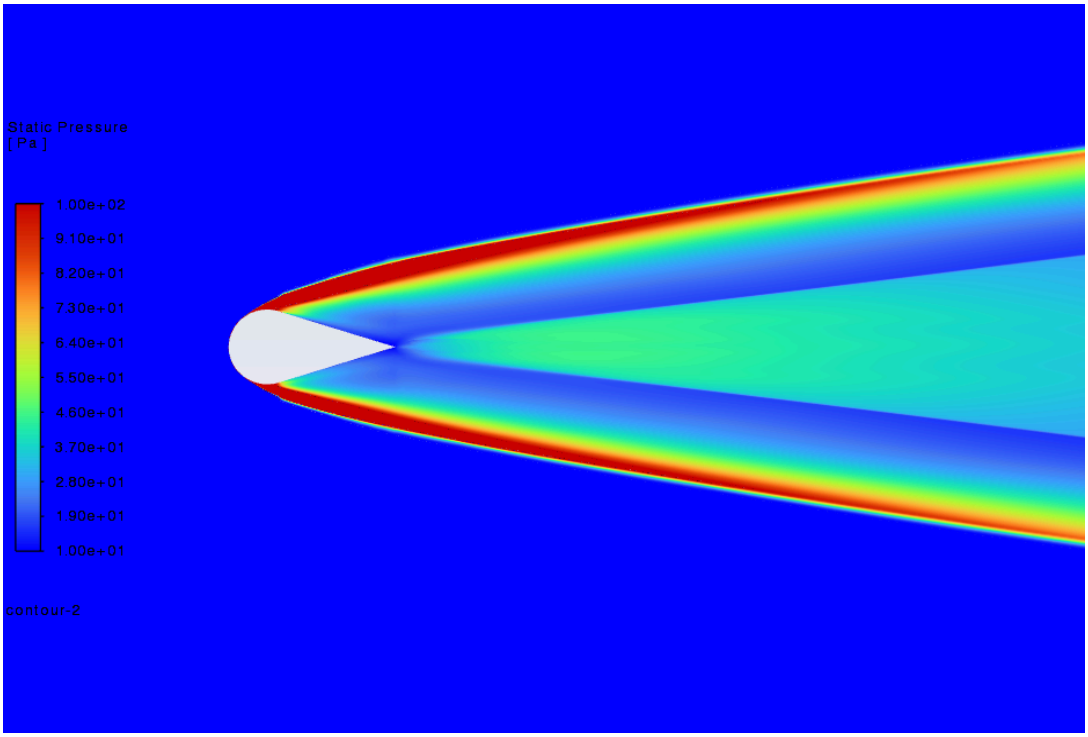


Figure 5.17 - Blunt forebody: contour of static pressure (nearfield)

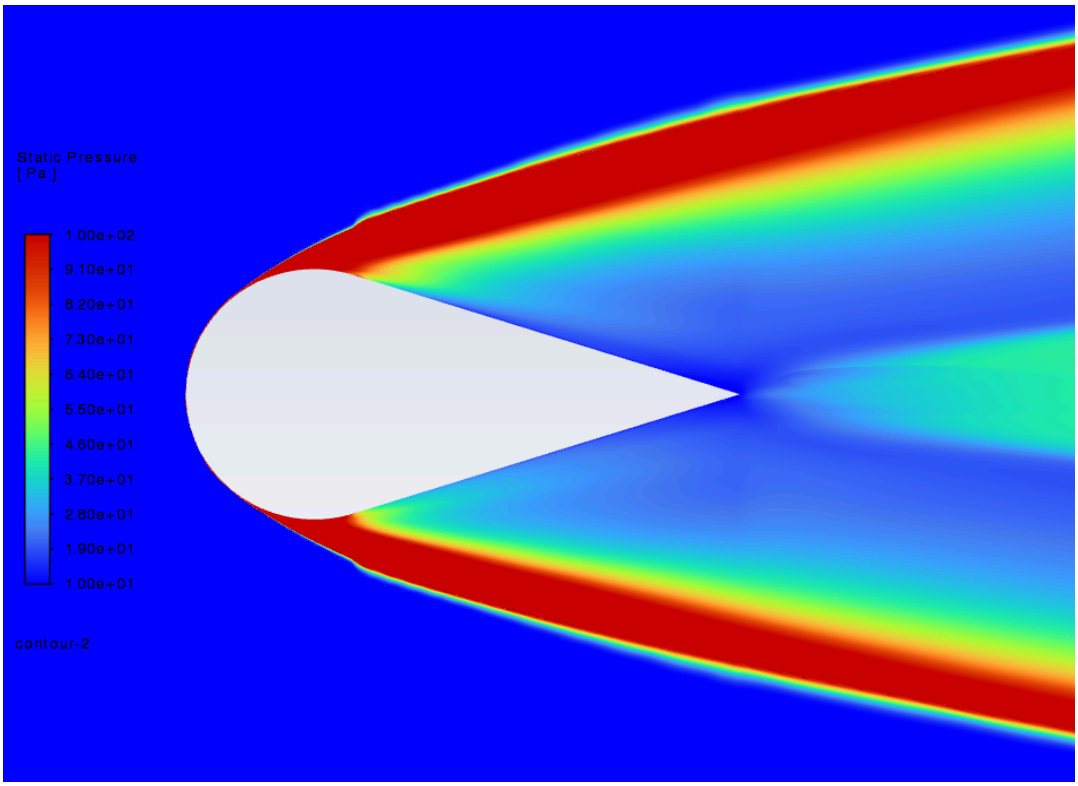


Figure 5.18 - Blunt forebody: contour of static pressure (near wall)

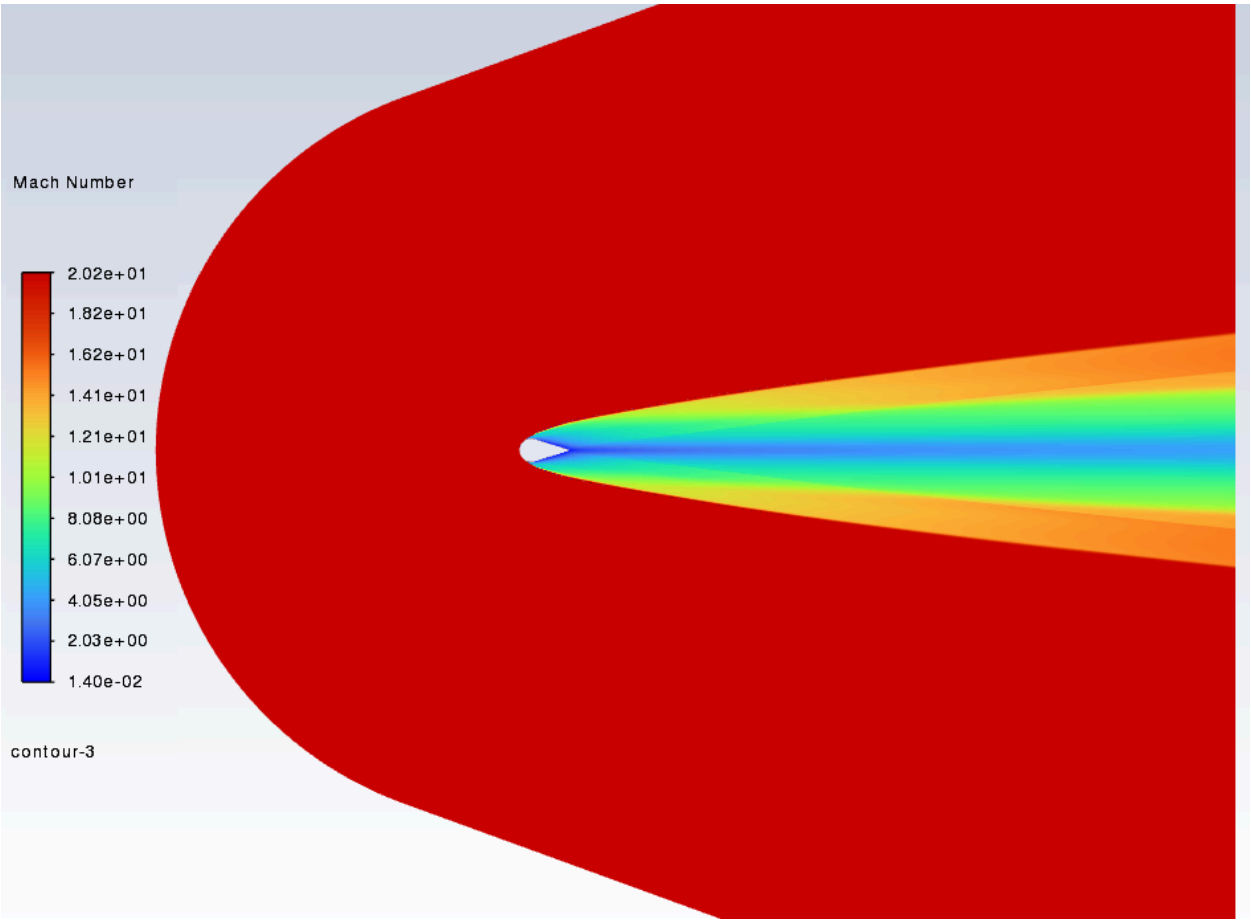


Figure 5.19 - Blunt forebody: contour of mach number (full view)

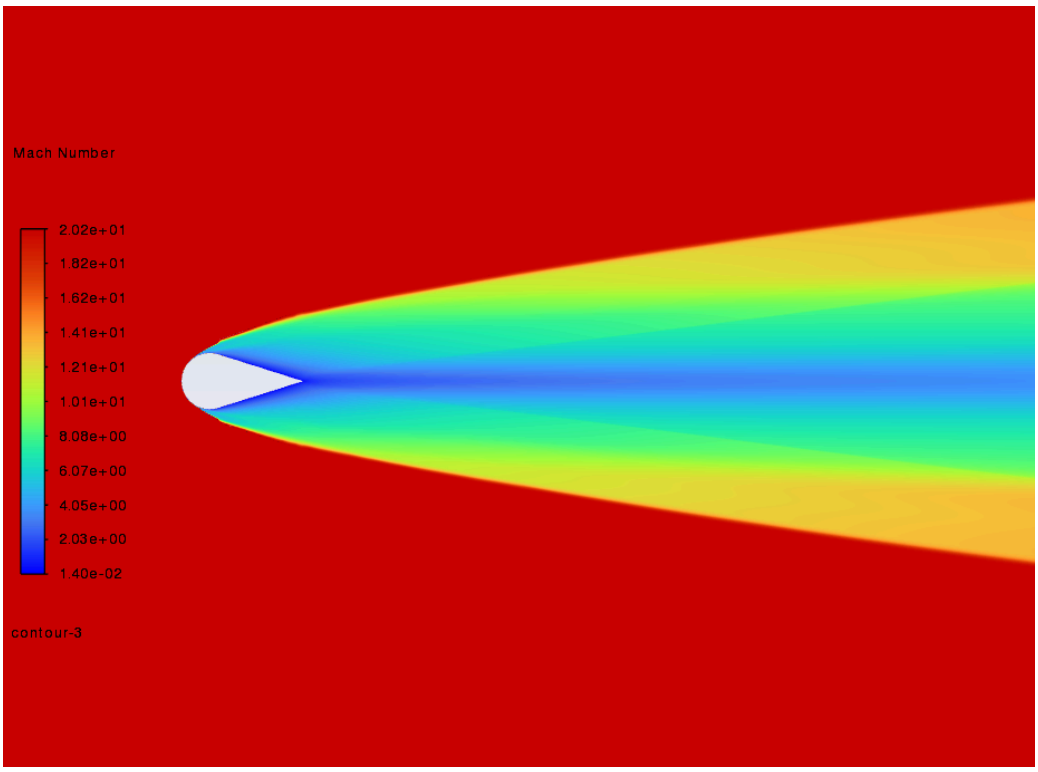


Figure 5.20 - Blunt forebody: contour of mach number (nearfield)

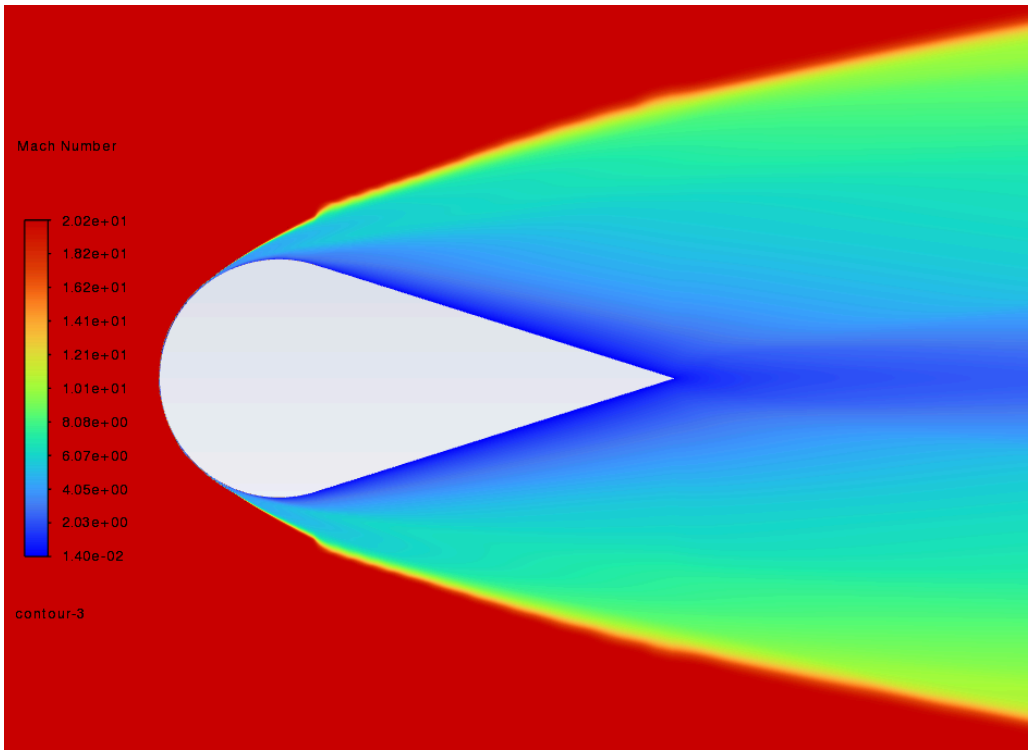


Figure 5.21 - Blunt forebody: contour of mach number (nearfield)

XY Plots

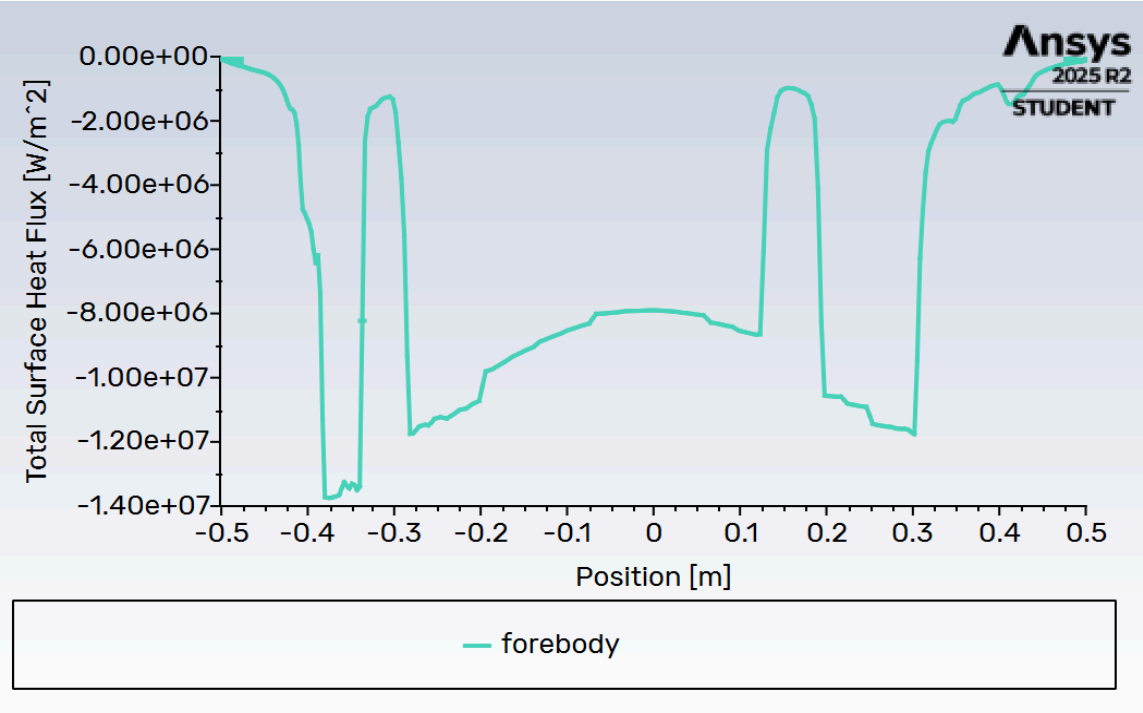


Figure 5.22 - Blunt forebody: surface heat flux vs. y position (front surface)

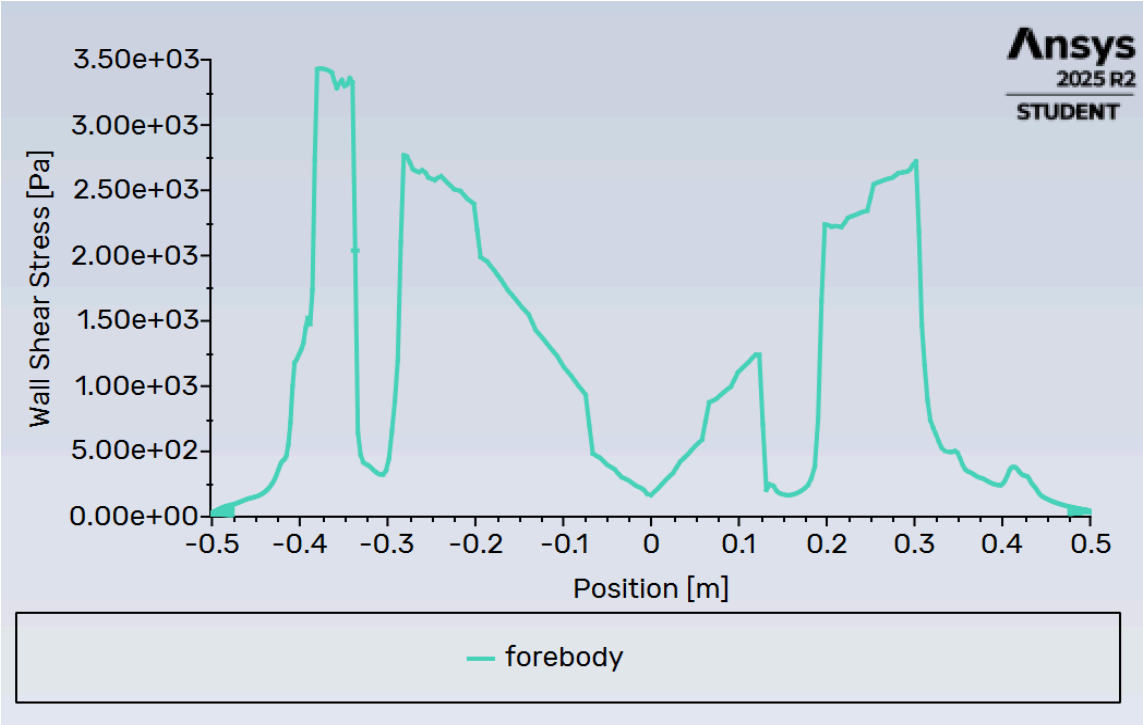


Figure 5.23 - Blunt forebody: wall shear stress vs. y yosition (front surface)

5.5 Discussion

5.5.1 Contour Plot Interpretation

The static temperature, Mach number, and pressure contours show that both geometries produce physically reasonable hypersonic flowfield behavior, but the shock structures differ due to the geometry of each body. For the axisymmetric blunt forebody, a broad detached bow shock is not clearly visible in the contour plots. This does not necessarily indicate that a detached shock is absent. Rather, the shock stand off distance is very small relative to the body scale, and the shock layer remains thin and close to the surface. The rounded forebody produces compression near the nose, but the downstream tapered body allows the flow to turn along the surface and develop into a strong oblique shock structure. As a result, the flowfield appears more like a thin compressed shock layer around the body rather than a widely detached capsule type bow shock.

This behavior is influenced by both geometry and freestream conditions. A blunt capsule with a broad forebody would be expected to produce a more visibly detached bow shock. In contrast, the present sphere cone has a rounded nose followed by a tapered downstream profile, which reduces the effective bluntness of the overall body and allows the shocked flow to remain closer to the surface. The 65 km freestream condition also contributes to the thin appearance of the shock layer because the low pressure and low density environment concentrates the high temperature and Mach number gradients into a narrow region in the continuum CFD solution. Therefore, the lack of a visually broad detached bow shock is not necessarily a numerical error, but is consistent with the selected geometry and high altitude flow condition.

For the blunted wedge, the contours show a more distinct oblique shock system. This is expected because the wedge imposes a finite flow turning angle through the inclined upper and lower surfaces. The rounded leading edge generates the initial shock structure, while the sloped wedge faces continue to compress the flow downstream. As a result, the Mach number reduction and static temperature increase are more pronounced along the wedge shock system. The shock structure is therefore stronger and more clearly defined than the thin shock layer observed around the axisymmetric blunt forebody.

The temperature contours support this interpretation. The wedge produces a sustained high temperature region behind the oblique shock and along the sloped body, while the sphere cone produces a thinner high temperature layer close to the forebody surface. Similarly, the Mach contours show a stronger and more extended low Mach region behind the wedge shock, while the sphere cone retains a thinner near body shock layer. The pressure contours follow the same trend, with the wedge showing strong compression bands along the inclined surfaces and the sphere cone showing concentrated compression near the forebody and shoulder transition.

Both geometries show similar downstream wake behavior because both bodies are artificially truncated at the aft end. The truncated aft wall produces a separated base region and a low Mach wake that extends downstream of the body. This wake structure is mainly a consequence of the simplified geometry and should not be interpreted as the primary TPS

loading environment. The forebody, shoulder, leading edge, and acreage surface regions remain the primary surfaces of interest for the thermal analysis.

Overall, the contour plots indicate that the axisymmetric sphere cone produces a thin near body shock layer with strong downstream oblique shock behavior, while the blunted wedge produces a clearer and stronger oblique shock system due to its finite wedge angle. The observed differences are expected based on the geometry and high altitude freestream condition. These contour results also support the heat flux interpretation because the local heating peaks occur near regions where the flowfield shows strong turning, compression, and near wall gradient development.

5.5.2 Surface Heat Flux Behavior

The primary output from the CFD simulations is the local surface heat flux along the exposed aerodynamic surfaces. This heat flux is used as the aerodynamic heating input for the thermal analysis described in Chapter 6. For each geometry, Fluent computes the surface heat flux based on the imposed wall temperature, local flowfield temperature, transport properties, and near wall thermal gradient. The heat flux is then exported from the relevant wall surfaces and used to evaluate the thermal response of the proposed ATPS stack.

For the axisymmetric blunt forebody, the heat flux distribution was found to be nonuniform along the surface. The stagnation region produced a high heat flux, as expected for a blunt body with strong shock behavior. However, the maximum local heat flux did not occur exactly at the stagnation point. Instead, localized off stagnation heat flux peaks developed near the forebody shoulder region, where the rounded forebody transitions into the downstream body profile. This indicates that the peak thermal loading for the selected sphere cone geometry is controlled not only by stagnation heating, but also by local curvature and flow acceleration near the shoulder.

The heat flux profile also showed local reductions in heating before and after the off stagnation peaks. Starting from the stagnation point and moving outward along the forebody, the heat flux first changes through a lower heating region before rising sharply near the shoulder. Downstream of the shoulder peak, the heat flux decreases again. This produces a nonmonotonic heat flux distribution rather than a simple profile where heating is maximum at stagnation and decreases smoothly along the surface. This behavior is important because it identifies the shoulder as a local thermal design driver for the blunt forebody case.

For the blunted wedge leading edge, the heat flux is concentrated near the finite radius nose and then develops along the upper and lower wedge surfaces. This behavior is consistent with a leading edge heating environment, where the highest thermal gradients are expected near the blunt nose and along the early developing boundary layer on the wedge faces. A localized heat flux increase was also observed near the region where the blunt tip transitions into the sloped wedge body. This feature is different from the shoulder heating observed on the axisymmetric forebody, but it is still associated with a geometric transition where the rounded nose blends into the straight inclined surface. The aft truncation surface is included as part of the

closed CFD body, but it is not used as a primary heat flux reporting surface because it represents an artificial geometric closure rather than the main leading edge TPS surface.

It is also important to note that the stagnation heat flux values between the two geometries are influenced by the imposed fixed wall temperature used in each CFD case. The blunt forebody case had a lower calibrated wall temperature of approximately 1007 K, while the blunted wedge case had a higher calibrated wall temperature of approximately 1700 K. This difference is likely due to the blunt forebody seeing a much more pronounced drop in temperature from the stagnation point to the aft end resulting in a lower wall averaged temperature, this is an expected behavior. Since the computed wall heat flux depends on the wall normal temperature gradient, a hotter imposed wall reduces the gas to wall temperature difference and can lower the reported convective heat flux. Therefore, the lower stagnation heat flux observed for the wedge should not be interpreted as a geometric effect. It is mainly influenced by the calibrated isothermal wall condition used for that case. Generally the magnitude of the heat flux at the stagnation point in either case is overpredicted and rendered invalid due to this problematic assumption that was forced by nature of the limitations of the ANSYS student license. However the behavior is still consistent with the geometry and still allows for a meaningful interpretation of the aerothermal performance of both bodies. The issues with the isentropic wall temperature assumption is discussed in more detail in section 5.4.4.

5.5.3 Off Stagnation Shoulder Heating

The off stagnation heating behavior observed on the blunt forebody is interpreted as a geometry induced shoulder heating augmentation. As the hypersonic flow passes through the shock and moves away from the stagnation region, it turns around the rounded forebody and accelerates tangentially along the surface. Near the shoulder region, the local geometry produces a region of increased wall shear stress. This indicates that the near wall velocity gradient is larger and that the local boundary layer is thinner in this region.

A similar local heating behavior was also observed on the blunted wedge, although the geometric cause is not completely identical. In the wedge case, the local increase occurs near the transition between the blunt nose radius and the sloped wedge body. As the flow turns around the finite radius leading edge and attaches to the inclined wedge surface, the local curvature and pressure gradient change rapidly. This can increase near wall shear and strengthen the wall normal temperature gradient in the transition region. Therefore, the wedge heating feature is better described as a nose to wedge transition heating peak rather than a forebody shoulder heating peak.

The alignment between the heat flux peaks and the wall shear stress peaks supports the conclusion that the heating augmentation is a physical flowfield feature rather than a mesh or boundary condition artifact. The wall heat flux is governed by the wall normal temperature gradient, while the wall shear stress is governed by the near wall velocity gradient. Thus when both quantities peak in the same surface region, it strongly indicates that the local flow is producing stronger near wall gradients in both momentum and thermal fields. In the blunt

forebody case, the shoulder curvature increases the local acceleration and concentrates the thermal loading away from the stagnation point. In the wedge case, the transition from the blunt tip to the sloped body produces a similar local increase in near wall gradients.

This behavior is not classified as a classical shock boundary layer interaction. A classical SBLI would require a clearly identifiable incident shock, reflected shock, separation region, or shock induced boundary layer disturbance. In the present case, the observed heating increase is more directly associated with the local forebody curvature and the transition into the downstream body profile. For the wedge, the observed increase is associated with the local transition from the rounded leading edge to the sloped wedge surface. The more appropriate description is therefore a geometry induced local heating peak, with the blunt forebody producing an off stagnation shoulder heating peak and the wedge producing a nose to wedge transition heating peak.

5.5.4 Interpretation of Wall Shear Stress Alignment

The wall shear stress distribution provides an important diagnostic for interpreting the local heat flux behavior. Wall shear stress is defined by the near wall velocity gradient and the local dynamic viscosity. Therefore, it is sensitive to both the flow acceleration near the wall and the local gas temperature. However, the fact that the wall shear stress peaks occur in the same locations as the heat flux peaks is significant because it shows that the off stagnation heating is linked to the local flow development rather than only to the thermal boundary condition.

The wall shear stress result suggests that the shoulder region produces a thinner or more energetic near wall boundary layer. This increases the velocity gradient at the surface and also increases the wall normal thermal gradient. Since the surface heat flux is computed from this thermal gradient, the heat flux peak is expected to occur where the boundary layer is most strongly compressed or accelerated along the surface. For the wedge, the same interpretation applies at the transition from the blunt tip to the sloped body, where the local geometry changes the near wall flow development and produces an increase in both wall shear stress and heat flux.

Notably, the wall shear stress is not completely independent of the imposed wall temperature. The local viscosity and boundary layer structure also depend on temperature, so a different wall temperature field can change the magnitude of both wall shear stress and heat flux. However, the locational alignment of the wall shear stress and heat flux peaks supports the interpretation that the location of the heating feature is driven primarily by the geometry and local flowfield. The imposed wall temperature may affect the magnitude of the peak, but it does not by itself explain why the peak occurs specifically near the shoulder in the blunt forebody case or near the blunt tip to sloped body transition in the wedge case.

5.5.5 Isothermal Wall Limitation

The CFD simulations impose a fixed wall temperature on the body surface during each run. This is necessary because the solid TPS stack is not modeled directly inside Fluent. Instead, the wall temperature is calibrated through the MATLAB thermal model described in Chapter 6.

Although this coupling approach allows the CFD heat flux and thermal response to be iterated, the Fluent wall condition remains isothermal over each selected wall zone.

This creates an important limitation. If the entire forebody is treated as a single fixed temperature wall, the imposed temperature represents an average or representative surface temperature rather than the true local wall temperature. In reality, the wall temperature would vary along the surface. Regions with higher heat flux, such as the forebody shoulder, would likely reach higher local surface temperatures than lower heating acreage regions. Similarly, in the wedge case, the blunt tip to sloped body transition may reach a different local temperature than the leading edge stagnation region or the downstream acreage panels. If a single uniform wall temperature is imposed, these local heating regions may be assigned a wall temperature that is lower than their true local radiative equilibrium temperature. This increases the local gas to wall temperature difference and can inflate the predicted local heat flux magnitude.

Therefore, the off stagnation shoulder peak should be interpreted in two parts. The location of the peak is likely physical because it aligns with the wall shear stress distribution and occurs near a geometric transition. However, the absolute magnitude of the peak may be conservative if the local shoulder temperature would be higher than the imposed average wall temperature. The same limitation applies to the wedge transition heating feature. The isothermal wall assumption is therefore useful for identifying the relative heating behavior and the location of thermal hot spots, but it may overpredict local heat flux in regions where the actual surface temperature would rise above the imposed value.

5.5.6 Segmented Wall Temperature Calibration

To reduce the limitations associated with a single isothermal wall boundary condition, the body surface could be divided into multiple wall zones in a future refinement of the CFD model. Each wall zone could then be assigned its own fixed temperature and calibrated separately using the local heat transfer rate from Fluent and the MATLAB thermal model. This would provide a regional approximation to a spatially varying wall temperature without requiring full conjugate heat transfer inside Fluent.

For the blunt forebody, possible wall zones could include the stagnation region, shoulder or transition regions, acreage regions, and aftbody or truncation regions. For the blunted wedge, possible wall zones could include the leading edge region, blunt tip to sloped body transition regions, upper and lower acreage panel regions, and the aft truncation surface. Fluent could then report the heat transfer rate or area averaged heat flux on each wall zone independently, allowing each region to be passed separately into the MATLAB thermal analysis.

This segmented calibration method was not implemented in the present study in order to reduce model complexity, preserve mesh robustness, and limit computational cost. Splitting the wall into multiple boundary zones would require additional edge partitioning in the geometry and could introduce new mapped meshing constraints near the body surface. Since the final meshes already rely on carefully controlled mapped faces and edge sizing, additional wall splits could increase the risk of local mesh distortion, mismatched edge divisions, or poor cell quality

near regions where wall heat flux is most sensitive. For this reason, the present CFD model uses a simplified isothermal wall treatment for each major body surface.

The single isothermal wall assumption is therefore treated as a modeling limitation. It allows the CFD and MATLAB thermal model to be coupled with a practical calibration procedure while remaining compatible with the ANSYS Student license and available computational resources. However, it does not fully capture the spatial variation in wall temperature that would occur over a real TPS surface. Regions with higher local heating, such as the blunt forebody shoulder or the blunt tip to sloped body transition on the wedge, may reach higher local radiative equilibrium temperatures than the imposed average wall temperature. As a result, the local heat flux magnitude in those regions may be conservative because the imposed wall temperature may be lower than the true local surface temperature.

The ultimate solution to this limitation would be to model the solid TPS stack directly in Fluent using a conjugate heat transfer approach. In that case, the wall temperature would not need to be prescribed as a uniform boundary condition. Instead, Fluent would solve the fluid and solid thermal response together, allowing the wall temperature to vary naturally along the surface based on local aerodynamic heating, conduction through the TPS, and radiative heat rejection. This approach would provide a more physically complete treatment of the coupled flowfield and wall response, but it would also greatly increase mesh size, material modeling complexity, and computational cost. Due to the ANSYS Student license limitations and the need to preserve adequate near wall CFD resolution, this full conjugate approach was not used in the present study.

5.5.7 Theoretical Heatflux Ceiling

We can apply a ceiling to our result by looking at typical peak heat flux on similar bodies. This information is widely available in the literature. A paper from NASA Langley on high-energy re-entry for a HIAD (Hypersonic Aerodynamic Decellerator) system cites heat fluxes on the order of 15-40W/cm^2 or 0.15-0.4 MW/m^2 with peak heat flux exceeding that and reaching around 0.6 MW/m^2[50]. HIAD is a well studied blunt conical heat shield so the geometry is semiconsistent with the axi-symmetric sphere cone model. Another paper: the "Assessment of Fencing on the Orion Heat Shield" discusses heat fluxes reaching 150W/cm^2 or 1.5MW/m^2 [51]. Crucially Orion data is largely coming from Lunar return missions while the present study is focused on high energy orbital re-entry. A conservative ceiling for the maximum plausible heat flux for a blunt forebody under orbital re-entry could likely be approximated to be somewhere around 1.0MW/m^2. Data is limited on leading edges for re-entry vehicles but a study from the University of California Santa Barbara and NASA Langley discusses Heat Pipe Cooled Leading Edges for Hypersonic Vehicles. The presentation cites heat fluxes on the order of 160-200W/cm^2 or 1.6-2.0MW/m^2. The body in study sees a Mach 8 freestream condition and has a relatively blunt tip radius [52]. We can infer that this could scale dramatically with sharper objects and at re-entry speeds. A conservative ceiling for a blunted leading edge under orbital re-entry conditions could be placed at 2.0-3.0MW/m^2 for a surface peak heat flux. These heating

environments will be evaluated in chapters 6 and 7 as additional literature backed cases to the resultant geometries.

6.0 Thermal Analysis (Post Processing)

This chapter develops the steady state one dimensional thermal model used to propagate the aerodynamic heating obtained from the CFD solution into the multilayer Adaptive Thermal Protection System (ATPS) stack. The purpose of this chapter is twofold. First, it establishes a transparent analytical framework relating the imposed aerothermal loading to the temperature distribution through the thickness of the TPS. Second, it provides a preliminary screening method for evaluating whether a given stackup can simultaneously maintain acceptable back-face temperatures and place the embedded Nitinol layer within the desired transformation temperature window. In this sense, the present model serves as the first thermal sizing step for the proposed smart material based TPS architecture.

The analysis is based on one dimensional steady conduction through the thickness of a localized flat panel representation of the TPS. The external thermal loading is taken directly from the CFD solution described in Chapter 5, where the computed heat transfer rate or corresponding wall heat flux is applied at the exposed surface of the multilayer stack. From this input, the model determines the outer wall temperature, the interface temperatures between layers, the temperature distribution through the stack, and the resulting back face or substructure temperature. Because of its simplicity and transparency, the model is well suited for rapid evaluation of material choices, layer thicknesses, and general thermal design trends during the early stages of TPS development.

The need to evaluate back-face temperature explicitly is fundamental in thermal protection system design. For reusable cool structure systems, the objective is not only to withstand high surface heating, but also to ensure that the underlying structural material remains below its allowable temperature. This requirement was central to the Shuttle thermal protection system, where the TPS was designed to limit the aluminum airframe to approximately 350°F or approximately 450K, in order to avoid structural softening and associated system level risks [40]. That precedent motivates the same thermal design philosophy in the present work: the ATPS must not only reduce heat flow from the hot outer surface, but must also maintain the protected back face or structural temperature within an acceptable operating range. In parallel, the thermal state of the embedded Nitinol layer must be assessed to determine whether the SMA remains martensitic, enters the transformation regime, or becomes fully austenitic under the applied heating. The present chapter therefore establishes the thermal framework needed to evaluate both structural protection and SMA activation within the proposed ATPS concept.

6.1 Scope and Modeling Assumptions

The analytical model developed in this chapter is based on a localized flat panel idealization, in which a small surface patch of the overall heat shield or vehicle wall is treated as a representative section. Each TPS panel is modeled as a one dimensional stack of homogeneous layers, and heat flow is assumed to occur only in the direction normal to the panel surface. This approximation is considered appropriate when the in-plane thermal gradients over a local surface

patch are small compared with the through thickness gradient imposed by external aerodynamic heating and internal heat rejection. Accordingly, the CFD derived heat transfer rate is spatially averaged over localized regions, allowing each one-dimensional calculation to represent the mean thermal loading of a patch for which lateral conduction effects are treated as second order. The analysis is repeated at representative thermal extremes, including stagnation point, acreage, and trailing edge conditions, to capture bounding hot and cold cases. Initial averaged heat flux case along the entire heat shield surface is used for calibration to estimate the corresponding average wall temperature.

Each material layer is assumed to be homogeneous and isotropic for the purposes of the present analysis. Thermal conductivity is treated as constant within each layer, and no internal volumetric heat generation is included. All interfaces are assumed to be perfectly bonded, such that interfacial thermal contact resistance is negligible in the present formulation not including the adhesive bondline at the vehicle skin. This is an idealization, but it allows the temperature field to be expressed in closed form and provides a useful lower bound estimate of overall thermal resistance. Any degradation in bond quality or imperfect interfacial contact would increase the actual thermal resistance and alter the local interfacial temperature distribution. Such effects are deferred to future numerical analyses.

The thermal model in this chapter is steady state as derived from steady state CFD results. Accordingly, the through thickness heat flux is uniform through all solid layers, and each layer exhibits a linear temperature profile. The present chapter therefore does not account for transient thermal soak, finite heating duration, or thermal lag within the stack. This treatment is consistent with the intended role of the present analysis as a preliminary screening model and feasibility check rather than a complete thermal history solution.

The Nitinol layer is similarly treated here from a thermal standpoint only. The present chapter determines the temperatures at the hot and cold faces of the Nitinol layer, along with the one dimensional temperature distribution across its thickness. These temperatures are then compared against the relevant transformation temperatures in order to determine whether the layer is expected to remain martensitic, enter the transformation regime, or become fully austenitic under the applied thermal loading. The detailed constitutive recovery stress model introduced earlier in the report is not solved in this chapter, but the thermal state obtained here provides the direct input needed for that later thermomechanical calculation.

Finally, this chapter treats the external CFD heat transfer result as the thermal forcing function for the solid stack. In its simplest form, the CFD result is converted into a surface heat flux and applied directly as the outer boundary condition. In a more complete closure, the outer surface temperature can be solved from an energy balance that includes both conduction and radiative heat rejection from the hot surface. Both viewpoints are useful. The imposed flux model is the simplest formulation for hand calculations, while the surface energy balance formulation is better suited for calibration against the CFD wall temperature assumption and for coupling to the MATLAB solver developed for this work.

6.2 Thermal Boundary Conditions and CFD Coupling

The thermal input to the solid TPS stack is obtained from the aerodynamic heating solution computed in Chapter 5. The CFD solution is reported in terms of total heat transfer rate Q_{CFD} over a surface patch of area A. It is important to additionally clarify that the CFD total heat transfer rate includes convective losses back into the fluid. In a sense the entire convective part of the thermal analysis is handled by the CFD model. The corresponding average applied heat flux is written as shown in equation 14.

$$q''_{\mathrm{CFD}} = \frac{Q_{\mathrm{CFD}}}{A} \tag{14}$$

This quantity represents the net rate at which thermal energy is delivered from the shock layer flow to the outer TPS surface per unit area. For the present one dimensional steady state model, q_{CFD} is taken as the driving thermal input.

In practice not all of the incoming aerothermal load inputs directly into the TPS by conduction. At elevated surface temperatures, part of the imposed heating can be rejected by thermal radiation from the hot face. Reusable TPS tiles are typically given an optical coating as radiation serves as one of the primary methods of heat rejection for heat shields. The outer surface energy balance may therefore be written as shown in equation 15.

$$q''_{\mathrm{CFD}} = q''_{\mathrm{rad}} + q''_{\mathrm{cond}} \tag{15}$$

with the radiative contribution represented by the Stefan-Boltzmann relation

$$q''_{\mathrm{rad}} = \varepsilon_w \sigma \left(T_w^4 - T_{\mathrm{rad}}^4 \right) \tag{16}$$

and the conductive contribution into the stack represented by

$$q''_{\mathrm{cond}} = \frac{T_w - T_b}{R''_{\mathrm{tot}}} \tag{17}$$

Where T_w is the wall temperature, T_b is the back face temperature at the end of the TPS stack, ε_w is the surface emissivity, σ is the Stefan-Boltzmann constant, and R''_{tot} is the total through thickness thermal resistance per unit area. Combining the preceding relations gives the nonlinear wall temperature closure as specified in equation 18.

$$q''_{\mathrm{CFD}} = \varepsilon_w \sigma \left(T_w^4 - T_{\mathrm{rad}}^4 \right) + \frac{T_w - T_b}{R''_{\mathrm{tot}}} \tag{18}$$

This relation is especially useful in the MATLAB implementation, because it allows the wall temperature to be solved iteratively for a specified CFD heat input and specified back face condition.

A wall temperature must often be assumed before running the CFD itself, since the TPS material is not fully resolved in the flowfield model. That assumed wall temperature is not the final answer, but rather an initial boundary-value guess used to obtain a CFD heating estimate. After the CFD heat transfer result has been extracted, the energy balance model is solved to obtain an updated wall temperature. The guessed and computed wall temperatures are then compared, and the process can be repeated until the wall temperature assumed in CFD is consistent with the thermal solution. This calibration loop gives the analytical model an important role beyond simple post-processing: it serves as the reduced order thermal closure linking the fluid solution to the TPS stack.

6.3 Multilayer Conduction Model

6.3.1 Layering, coordinate, and notation

Let x measure distance inward from the hot face, with x = 0 at the surface. The eight solid layers are numbered i = 1…n. Each layer, i, has thickness L_i and thermal conductivity k_i. The total solid thickness, L_{tot}, is the sum of each layer thickness, L_i. Because convection is excluded in this chapter, absolute temperature levels are referenced either to T_0 (if provided by a separate hot surface closure) or to T_n (if a back face reference is prescribed elsewhere). Otherwise, results are reported as temperature drops and interface temperatures relative to T_0.

6.3.2 Governing Equations (Fourier conduction laws)

Within each homogeneous layer, one dimensional steady conduction with no internal heat generation is governed by Fourier's law,

$$q'' = -k_i \frac{dT}{dx} \tag{19}$$

The model assumes steady state, no internal heat generation, perfect thermal contact between layers. Additionally the layers are arranged in series. These assumptions imply that the same q″ passes through every layer. The hot-side boundary condition for this chapter is the CFD-supplied wall heat flux; consequently, q″ is treated as a known input and is spatially uniform through the stack.

The per-unit-area thermal resistance of layer R_i'' is introduced as the quotient of the layer length L_i and layer conductivity k_i with an overall conduction resistance equal to the sum of each layer's thermal resistance as shown in equations 20 and 21. Because q″ is uniform, the temperature drop across layer i can be represented by equation 22.

$$R_i'' = \frac{L_i}{k_i} \tag{20}$$

$$R_{\text{tot}}'' = \sum_{i=1}^{N} R_i'' = \sum_{i=1}^{N} \frac{L_i}{k_i} \tag{21}$$

The total per-unit-area thermal resistance of the entire stack is then

$$\Delta T_i = q'' R_i'' = q'' \frac{L_i}{k_i} \tag{22}$$

If the hot-surface temperature is denoted by T_0, and the interfacial temperature after layer j is denoted byT_j, then the temperature sequence through the stack can be written in telescoping form as

$$T_j = T_0 - q'' \sum_{m=1}^{j} R_m'', \quad (j = 1, \ldots, 8) \tag{23}$$

In particular, the back face temperature is

$$T_N = T_0 - q'' R_{\text{tot}}'' \tag{24}$$

If, instead, the back face temperature is known or constrained, then the hot-surface temperature can be written as

$$T_0 = T_N + q'' R_{\text{tot}}'' \tag{25}$$

These equations model the central closed-form result of the steady one dimensional conduction model. For a given heat flux and known material stackup, every interfacial temperature follows directly once either the hot-face or back face temperature is specified.

The resistance-network representation is especially useful because it provides a direct way to interpret the thermal role of each material layer. Layers with large L_i/k_i ratios dominate the total resistance and therefore control the overall temperature drop across the stack. In practical TPS design, low conductivity porous ceramic and aerogel-like layers are expected to contribute the largest share of the temperature reduction, while metallic substructures and compliant layers contribute much smaller thermal resistance but may be required for structural or adaptive functionality. This is consistent with reusable TPS practice more broadly, where low conductivity porous silica layers are used specifically to isolate the substructure from the external heating environment. [40]

6.4 Back Face Temperature and Structural Thermal Constraint

From the standpoint of vehicle survivability, the most critical scalar output of the present analytical model is often not the hot-surface temperature, but the back face temperature. The back face represents the thermal state of the deepest TPS boundary and, depending on the stack definition, may correspond either to the backplate of the sandwich panel or to the temperature of the protected structural substrate. Even if the outer surface is allowed to become very hot, the TPS remains acceptable only if the underlying structure is kept below its allowable temperature.

This requirement can be written simply as

$$T_b \leq T_{\text{allow}} \tag{26}$$

where T_b is the back face or substructure temperature and T_{allow} is the allowable temperature of the protected structural material. For aluminum-backed reusable structures, the widely cited benchmark is approximately 450K, which was the Shuttle design limit used to prevent annealing and loss of structural integrity in the aluminum airframe. That limit provides a useful engineering reference for the present study whenever the proposed ATPS is interpreted as protecting a conventional cool structure vehicle. [40]

In the simplest form of the present model, the back face temperature may be prescribed directly as a known boundary condition. Under that assumption, the wall energy balance is closed using the imposed interface or structural temperature, and the outer wall temperature is solved such that the incoming CFD heating is balanced by front face reradiation and in-depth conduction. In that case, the conductive heat rate through the stack is written as

$$Q_{\text{cond}} = \frac{T_w - T_b}{R_{\text{cond}}} \tag{27}$$

where T_w is the outer wall temperature and R_{cond} is the total conductive thermal resistance of the multilayer stack. The corresponding front face energy balance becomes

$$Q_{\text{CFD}} = Q_{\text{rad}} + Q_{\text{cond}} \tag{28}$$

with the radiative loss from the outer wall written as

$$Q_{\text{rad}} = \varepsilon_w \sigma A \left(T_w^4 - T_{\text{rad,front}}^4 \right) \tag{29}$$

so that the nonlinear wall-temperature closure is as shown below.

$$Q_{\text{CFD}} = \varepsilon_w \sigma A \left(T_w^4 - T_{\text{rad,front}}^4 \right) + \frac{T_w - T_b}{R_{\text{cond}}} \tag{30}$$

This formulation is useful when the analyst wishes to enforce a structural temperature requirement directly and determine what outer wall temperature and through-thickness gradient are required to maintain that condition. In that sense, the model acts as a constrained design calculation rather than a free prediction of back face temperature.

A more physical treatment, however, is obtained when the back face temperature is no longer prescribed directly, but instead emerges from the thermal interaction between the TPS, the vehicle skin, and the internal thermal environment. In the present work, this more realistic representation is obtained by extending the stack to include the vehicle skin as a final conductive layer and by applying a backside boundary condition behind that structural skin. Under this formulation, the tile to vehicle interface temperature and the backside temperature of the vehicle skin are both outputs of the solution rather than imposed inputs.

If the total conductive resistance of the full stack, including the structural skin, is denoted R_{cond}, then the outer wall temperature and backside skin temperature are related through

$$T_w = T_{b,\mathrm{skin}} + Q_{\mathrm{cond}} R_{\mathrm{cond}} \tag{31}$$

where $T_{b,skin}$ is the temperature at the back side of the structural skin. The conducted heat reaching this location must then be rejected into the interior environment through an appropriate back side boundary condition. If the internal environment is modeled using convection alone, then

$$Q_{\mathrm{cond}} = h_b A \left(T_{b,\mathrm{skin}} - T_{\infty,b}\right) \tag{32}$$

where h_b is the effective backside convection coefficient and $T_{\infty,b}$ is the ambient temperature seen by the back side of the skin, such as an internal cavity or vehicle interior environment.

If backside radiation is also included, then the back side boundary condition becomes

$$Q_{\mathrm{cond}} = h_b A \left(T_{b,\mathrm{skin}} - T_{\infty,b}\right) + \varepsilon_b \sigma A \left(T_{b,\mathrm{skin}}^4 - T_{\infty,b}^4\right) \tag{33}$$

where ε_b is the emissivity of the back side of the structural skin. In this case, the conducted heat reaching the structure is partitioned into internal convective rejection and backside radiative rejection. The front face energy balance remains

$$Q_{\mathrm{CFD}} = Q_{\mathrm{rad,front}} + Q_{\mathrm{cond}} \tag{34}$$

with

$$Q_{\mathrm{rad,front}} = \varepsilon_w \sigma A \left(T_w^4 - T_{\mathrm{rad,front}}^4\right) \tag{35}$$

so that the full steady state system may be interpreted as a coupled outer wall radiation problem, multilayer conduction problem, and back side heat rejection problem.

Under this more physical formulation, the tile to vehicle interface temperature is no longer imposed and instead follows directly from the computed temperature distribution through the stack. This is important because the interface temperature is itself a quantity of design interest. If it is too high, structural or adhesive degradation may occur even if the final external skin temperature criterion is formally satisfied. Conversely, if the stack is made excessively insulating in order to keep the vehicle skin cold, the Nitinol interlayer may never reach its desired transformation temperature range. The thermal design of the stack must therefore satisfy two coupled requirements simultaneously: it must preserve the structural substrate below its allowable temperature while also allowing the embedded SMA layer to enter a thermally useful actuation regime.

The present one dimensional model therefore supports two complementary design checks. The first is a smart-material functionality check, namely whether the Nitinol layer reaches the desired transformation temperature range. The second is a structural-protection check, namely whether the temperature of the protected structure or vehicle skin remains below its allowable limit. These two conditions are not necessarily aligned. Increasing thermal resistance above the Nitinol may better protect the structural substrate, but it may also prevent the SMA from heating into its desired actuation band. Conversely, placing the Nitinol too close to the hot face may favor transformation and recovery-stress generation while allowing excessive heat penetration into the deeper support layers and structure. The analytical model developed in this chapter is therefore especially useful because it exposes this tradeoff in a transparent manner.

6.5 Nitinol Layer Temperature Distribution and Transformation Window

Because the adaptive function of the proposed ATPS depends on the embedded Nitinol layer, it is necessary to determine not only the overall back face temperature but also the local temperature field within the SMA itself. Let the Nitinol layer be indexed as layer n, with thickness L_n and conductivity k_n. The temperature at the hot-side face of the Nitinol layer is then obtained from the cumulative resistance of all layers above it,

$$T_{n,\mathrm{in}} = T_0 - q'' \sum_{m=1}^{n-1} R''_m \tag{36}$$

and the temperature at the cold-side face of the Nitinol layer is

$$T_{n,\mathrm{out}} = T_{n,\mathrm{in}} - q'' R''_n \tag{37}$$

Because the present model is one dimensional and steady, the temperature distribution within the Nitinol layer is linear. If the local coordinate inside the Nitinol is denoted by ξ, measured from the hot-side face of the Nitinol so that $0 \leq \xi \leq$ Ln, then

$$T_n(\xi) = T_{n,\text{in}} - \frac{q''}{k_n}\,\xi \tag{38}$$

A useful representative temperature for phase-state screening is the average Nitinol temperature,

$$\overline{T}_n = \frac{T_{n,\text{in}} + T_{n,\text{out}}}{2} \tag{39}$$

The thermal state of the Nitinol layer can then be compared against the transformation temperatures M_s, M_f, A_s, and A_f. A simple first-pass criterion is to interpret the layer as fully martensitic when its entire thickness remains below M_f, fully austenitic when its entire thickness exceeds A_f, and in a mixed or transforming state when its temperature range overlaps the interval bounded by these transformation temperatures. We can represent these as such:

- Fully austenitic: $T_{n,out} \geq A_f$
- Fully martensitic: $T_{n,in} \leq M_f$
- Austenitic transformation regime: $T_{n,in} \geq A_s$ and $T_{n,out} < A_f$
- Martensitic transformation Regime: $T_{n,out} \leq M_s$ and $T_{n,in} > M_f$

The above representation provides the thermal screening logic needed to determine whether the SMA is likely to activate under a specified external heating condition. In the present chapter this is used as a temperature-window check only. The actual recovery stress and stiffness modulation generated by the constrained SMA are functions not only of temperature but also of pre-strain, martensite fraction, constraint condition, and constitutive parameters, and will therefore be treated in the later thermo-mechanical analysis. Even so, the temperature distribution derived here is the essential thermal input to that later constitutive calculation.

6.6 MATLAB Solution Procedure and CFD–Thermal Calibration

The analytical relations developed above are implemented in MATLAB in order to provide a rapid and traceable solution procedure for the multilayer TPS stack. The numerical workflow begins with an assumed outer wall temperature used as the TPS boundary condition in the CFD model. The CFD solution then yields a total heat transfer rate or an equivalent area averaged wall heat flux. That heating result is imported into MATLAB and applied to the one dimensional thermal resistance model.

If the hot surface temperature is prescribed directly, the interface temperatures follow immediately from the closed-form resistance relations. If instead the surface temperature is unknown, then the MATLAB routine solves the nonlinear outer surface energy balance,

$$f(T_w) = q''_{\mathrm{CFD}} - \varepsilon_w \sigma_{\mathrm{SB}} \left(T_w^4 - T_{\mathrm{rad}}^4\right) - \frac{T_w - T_b}{R''_{\mathrm{tot}}} = 0 \tag{40}$$

for the physically consistent wall temperature T_w. In this case, the initial wall temperature supplied from the recovery temperature estimate is used only as the numerical starting value for the root-finding algorithm; it does not determine the final thermal solution. Once T_w is obtained, the conduction heat flux into the stack and the interfacial temperatures are computed, including the hot side and cold side temperatures of the Nitinol layer.

The thermal solution is then compared with the wall temperature that was originally assumed in the CFD model. If the two values differ significantly, the wall boundary condition in CFD is updated and the process is repeated. The iteration is considered converged when the relative change in wall temperature between successive coupled iterations satisfies

$$\frac{\left|T_w^{(m+1)} - T_w^{(m)}\right|}{|T_w^{(m)}|} < \varepsilon_{\mathrm{tol}} \tag{41}$$

where ε_{tol} is the prescribed convergence tolerance. This coupled CFD–thermal loop allows the one dimensional model developed in this chapter to serve as a reduced-order closure model between the flow solver and the multilayer solid stack. It also provides a computationally inexpensive way to evaluate the influence of layer thickness and material choice before the full three-dimensional ANSYS simulations are performed.

6.7 Results

Case 1 - Leading Edge (Blunt Tip Wedge) Stagnation Point

For case 1 the stagnation point on the blunt tip wedge is analyzed. From figure 5.11 this is the local surface heatflux reading at the position 0 m on the x axis. Below are the results of this study.

Table 6.1 - Case 1 results

Metric Type	Selected Material	Result Value
Heat Flux	N/A	3.3 MW/m^2
Optical Coating	Toughened unipiece fibrous insulation coating, TUFI	2819.85 K *
Outer Ceramic Tile	TUFI-coated AETB	2819.15 K*

SMA Interlayer	NiTiZr very high temperature Nitinol	625.49 K
Buffer	Alumina-silica fibrous board	625.31 K
SIP	Saffil alumina fiber felt	559.50 K
Adhesive	RTV silicone adhesive	419.11K
Vehicle	Aluminum	414.72 K

*The results show that a heat flux of 3.3 MW/m^2 is outside the practical design envelope for this study. This heating case results in a wall temperature beyond the operating limits of the highest valid temperature rated ceramic tile from table 2.1.

Case 2 - Leading Edge (Blunt Tip Wedge) Acreage Relaxation Zone

For Case 2, the acreage relaxation zone on the blunt tip wedge is analyzed. From Figure 5.11, this is the local surface heat flux reading taken downstream of the leading edge after the initial stagnation region and nose-to-wedge transition heating have relaxed along the sloped body surface. Below are the results of this study.

Table 6.2 - Case 2 results

Metric Type	Selected Material	Result Value
Heat Flux	N/A	0.62 MW/m^2
Optical Coating	Toughened unipiece fibrous insulation coating, TUFI	1855.86 K
Outer Ceramic Tile	TUFROC	1855.33 K
SMA Interlayer	NiTiHf very high temperature Nitinol	587.82
Buffer	Alumina-silica fibrous board	587.15 K
SIP	Nomex felt	418.03 K
Adhesive	RTV silicone adhesive	350.41 K
Vehicle	Aluminum	348.72 K

Case 3 - Leading Edge (Blunt Tip Wedge) Near Truncation

For Case 3, the near truncation region on the blunt tip wedge is analyzed. From Figure 5.11, this is the local surface heat flux reading near the downstream end of the sloped wedge body before the artificial aft truncation surface. This case is included to evaluate the TPS response in a lower downstream heating region away from the leading edge peak. Below are the results of this study.

Table 6.3 - Case 3 results

Metric Type	Selected Material	Result Value
Heat Flux	N/A	0.1 MW/m^2
Optical Coating	Borosilicate glass coating or RCG	1198.70 K
Outer Ceramic Tile	LI-900 silica tile	1198.45 K
SMA Interlayer	Cu-Al-Ni	448.42 K
Buffer	Alumina-silica fibrous board	448.36 K
SIP	Nomex felt	409.30 K
Adhesive	RTV silicone adhesive	345.23 K
Vehicle	Aluminum	345.21 K

Case 4 - Blunt Forebody (Axisymmetric Sphere-Cone) Stagnation Point

For Case 4, the stagnation point on the axisymmetric sphere-cone blunt forebody is analyzed. From Figure 5.22, this is the local surface heat flux reading at the stagnation region on the centerline of the forebody. Below are the results of this study.

Table 6.4 - Case 4 results

Metric Type	Selected Material	Result Value
Heat Flux	N/A	8.0 MW/m^2
Optical Coating	NULL*	~3600K
Outer Ceramic Tile	NULL*	NULL*

SMA Interlayer	NULL*	NULL*
Buffer	NULL*	NULL*
SIP	NULL*	NULL*
Adhesive	NULL*	NULL*
Vehicle	NULL*	NULL*

*Similarly to case 1 the wall temperature exceeds allowable temperatures and material operating limits for the highest temperature rated materials in this study. Therefore this heating case is outside the practical design envelope for the current proposed ATPS structure.

Case 5 - Blunt Forebody (Axisymmetric Sphere-Cone) Acreage Relaxation Zone

For Case 5, the acreage relaxation zone on the axisymmetric sphere-cone blunt forebody is analyzed. From Figure 5.22, this is the local surface heat flux reading taken away from the stagnation point and away from the localized geometry transition peak. This case represents a reduced heating acreage region on the forebody surface. Below are the results of this study.

Table 6.5 - Case 5 results

Metric Type	Selected Material	Result Value
Heat Flux	N/A	1.2 MW/m^2
Optical Coating	Hafnium carbide / zirconium carbide UHTC coating	2201.49 K
Outer Ceramic Tile	Zirconia Fiber Board + TUFROC (Composite)	2201.36 K (Zirconia) 1822.42 K (TUFROC)
SMA Interlayer	NiTiHf very high temperature Nitinol	685.62 K
Buffer	Alumina-silica fibrous board	683.71K
SIP	Nomex Felt	457.34 K
Adhesive	RTV silicone adhesive	366.40 K
Vehicle	Aluminum	364.13 K

Case 6 - Blunt Forebody (Axisymmetric Sphere-Cone) Geometry Transition Off Stagnation Heating Spike

For Case 6, the geometry transition region on the axisymmetric sphere-cone blunt forebody is analyzed. From Figure 5.22, this is the local surface heat flux reading near the shoulder region where the rounded forebody transitions into the downstream body profile. This region produced the localized off-stagnation heating peak observed in the CFD results. Note the localized off stagnation heating peak for the leading edge was not analyzed since the results are redundant. Below are the results of this study.

Table 6.6 - Case 6 results

Metric Type	Selected Material	Result Value
Heat Flux	N/A	13.8 MW/m^2
Optical Coating	NULL*	4054.89 K
Outer Ceramic Tile	NULL*	NULL*
SMA Interlayer	NULL*	NULL*
Buffer	NULL*	NULL*
SIP	NULL*	NULL*
Adhesive	NULL*	NULL*
Vehicle	NULL*	NULL*

*Similarly to both case 1 and case 3 the wall temperature far exceeds allowable temperatures and material operating limits for the highest temperature rated materials in this study. Therefore this heating case is outside the practical design envelope for the current proposed ATPS structure.

6.8 Discussion

The results show that the proposed ATPS stack is feasible only for a subset of the heating environments evaluated in this study, however they do show potential feasibility for real world typical peak heating environments that are reported in the literature. The lower and moderate heat flux cases 2, 3 and 5 were able to maintain reasonable material temperatures through the multilayer stack, while the highest heat flux cases exceeded the operating limits of the candidate materials. This distinction is important because the objective of this study is not only to determine whether the stack can reduce backface temperature, but also whether each material layer remains within a practical operating range. The results therefore need to be

interpreted using both the thermal response of the stack and the material limits listed in Table 2.1. A successful case is characterized as a case where both the vehicle skin temperature and SMA transformation temperature requirements are met and no maximum operating temperatures are exceeded for any layer.

The NULL cases identify heating environments that are outside the practical design envelope of the current ATPS configuration. Case 1, the blunt tip wedge stagnation point, used a heat flux of 3.3 MW/m2 and produced an outer surface temperature of approximately 2819 K. Although the internal layers were reduced to more manageable temperatures, the wall temperature exceeded the operating range of the selected TUFI coating and TUFI coated AETB tile. The highest temperature rated materials were the HfB2-SiC and the Hafnium Carbide Ceramic UHTC coating both at a maximum operating temperature of 2773 K which is still 46 K below the reported wall temperature. HfB2-SiC was not chosen as a selected candidate for any of the cases due to its high thermal conductivity. For this reason, Case 1 was marked as outside the practical design envelope even though the lower stack temperatures were not immediately unreasonable. The SMA interlayer reached approximately 625 K, the SIP reached approximately 560 K, the adhesive reached approximately 419 K, and the aluminum vehicle skin reached approximately 415 K, but the outer ceramic surface temperature was too high for the chosen reusable tile system.

Case 4 and Case 6 were more severe and were therefore marked as fully NULL cases. Case 4 corresponds to the axisymmetric sphere cone stagnation point and uses a heat flux of 8.0 MW/m2. This case produced an estimated wall temperature of approximately 3600 K, which is above the operating limits of the highest temperature materials considered in this study. Because the surface temperature exceeded the practical capability of the optical coating and ceramic layers, the remaining stack was not treated as a viable material selection. Case 6 was even more severe, corresponding to the off stagnation geometry transition peak on the blunt forebody. This case used a heat flux of 13.8 MW/m2 and produced a wall temperature of approximately 4055 K. This is far beyond the allowable range of the coating, ceramic tile, SMA, SIP, adhesive, and vehicle skin materials considered in the current material table. These cases show that the local CFD heating peaks cannot be treated as feasible reusable TPS design points within the current 150 mm stack thickness cap.

The high temperature NULL cases also connect directly to the CFD limitations discussed earlier in Chapter 5. The CFD model imposed a fixed isothermal wall temperature rather than solving a spatially varying wall temperature field through conjugate heat transfer. This was necessary because the solid TPS stack was not modeled directly in Fluent, which reduced mesh size and computational cost and helped keep the model compatible with the ANSYS Student license. However, this also means that the heat flux peaks can be conservative in localized high heating regions. If a real TPS surface were allowed to reach its own local radiative equilibrium temperature, the hottest regions would likely rise to a higher local wall temperature than the imposed representative value. This would reduce the local gas to wall temperature difference and lower the computed convective heat flux. Therefore, the NULL cases should be interpreted as

conservative or bounding results for the present methodology rather than direct proof that every similar physical configuration would fail.

This limitation does not remove the significance of the NULL cases. Even if the local heat flux values are inflated by the isothermal wall assumption, the predicted wall temperatures are so far above the material limits that these cases are still outside the intended design range of the current ATPS stack. A heat flux of 8.0 MW/m2 or 13.8 MW/m2 would require a different class of TPS architecture, likely involving a dedicated leading edge or stagnation region system, actively cooled structure, ablative material, or a much thicker and more specialized ultra high temperature ceramic hot structure. These cases are therefore useful because they define the boundary of the current concept. The proposed ATPS stack is not intended to solve every possible peak heating environment, but to evaluate whether a reusable adaptive stack can be feasible in lower to moderate local heating regions.

The successful cases show that the proposed material strategy is more appropriate for acreage and relaxed heating zones than for the most severe stagnation or geometry transition peaks. Case 2, the blunt tip wedge acreage relaxation zone, used a heat flux of 0.62 MW/m2 and produced an outer surface temperature of approximately 1856 K. This required a high temperature reusable ceramic material, so TUFROC was selected as the outer ceramic tile with a TUFI type optical coating. The SMA interlayer reached approximately 587 K, which falls into a range where a high temperature NiTiHf alloy is more appropriate than conventional binary NiTi. The buffer layer remained near 587 K, the Nomex felt SIP reached approximately 418 K, the RTV adhesive reached approximately 350 K, and the aluminum vehicle skin reached approximately 348 K. These values show that the stack can reduce the temperature from a high surface value to a backface temperature that is much closer to the protected structure range.

Case 3 represents the lowest heating case, corresponding to the near truncation region of the blunt tip wedge. The applied heat flux was 0.1 MW/m2, and the predicted wall temperature was approximately 1199 K. Under this condition, a lower temperature reusable tile option such as LI-900 silica tile was sufficient, with a borosilicate glass or RCG coating. The SMA interlayer reached approximately 448 K, which is more compatible with a high temperature binary NiTi option than with the very high temperature SMA systems. The SIP reached approximately 409 K, the RTV adhesive reached approximately 345 K, and the vehicle skin remained near 345 K. This case is the strongest indication that the proposed stack can be thermally feasible in lower heating acreage regions using more mature and lower cost materials.

Case 5 is the most important feasible case because it represents a higher heat flux acreage relaxation region on the blunt forebody. The applied heat flux was 1.2 MW/m2, which is significantly higher than Case 2 and Case 3, but still below the severe stagnation and off stagnation peak cases. This case required a more capable outer layer than a single reusable tile. A hafnium carbide or zirconium carbide UHTC coating was selected, and the outer ceramic layer was modeled as a composite arrangement using zirconia fiber board with a TUFROC layer. The surface side of this composite region reached approximately 2201 K in the zirconia layer, while the TUFROC layer was reduced to approximately 1822 K. This indicates that the additional

zirconia based layer was required to keep the TUFROC portion within a more reasonable operating range while still allowing the stack to fit within the maximum practical thickness cap. The SMA interlayer reached approximately 686 K, the buffer was approximately 684 K, the Nomex felt SIP reached approximately 457 K, the RTV adhesive reached approximately 366 K, and the aluminum vehicle skin reached approximately 364 K.

The Case 5 result supports the idea that the ATPS concept can remain feasible near the upper end of the moderate heating range, but only with careful material selection. A heat flux of 1.2 MW/m2 is not low, and it forced the use of an additional high temperature ceramic layer rather than a conventional acreage tile alone. This result shows why the material stack was selected based on temperature margin, material maturity, cost, and targeted temperature ranges. The outer surface required a high temperature coating and ceramic system to survive the applied heating. The SMA layer required a high temperature NiTiHf material because the predicted interlayer temperature was much higher than the useful range of conventional binary NiTi. The SIP and adhesive layers had to remain far enough below their allowable ranges to avoid thermal degradation, and the aluminum skin had to remain near a protected backface temperature. Within the 150 mm practical thickness cap, this case appears feasible, but it also shows that the design margin is controlled primarily by the outer ceramic system.

The material choices in the feasible cases were not selected only to minimize temperature. They were also selected to maintain a practical balance between performance, cost, reliability, and maturity. For lower temperature regions, more mature and lower cost options such as LI-900, RCG or borosilicate coating, binary or high temperature binary NiTi, Nomex felt, RTV adhesive, and aluminum are preferred because they are more established and easier to justify. For higher temperature regions, materials such as TUFROC, TUFI, zirconia fiber board, Saffil felt, and NiTiHf are introduced only when the temperature environment requires them. The very high end material options, such as hafnium carbide or zirconium carbide coatings and composite ceramic tile arrangements, are reserved for cases where standard reusable tile materials would exceed their operating limits. This approach avoids selecting exotic materials where a lower temperature and more mature material would be sufficient.

The SMA material selection is controlled by the peak temperature reached inside the interlayer, not by the external surface temperature. This is important because the SMA is not intended to function as the exposed heat shield. In Case 3, the SMA temperature was approximately 448 K, which supports the use of a high temperature binary NiTi option. In Case 2, the SMA temperature was approximately 546 K, which shifts the preferred material toward NiTiHf. In Case 5, the SMA temperature reached approximately 686 K, which also supports the selection of NiTiHf rather than lower temperature binary NiTi. This shows that the SMA layer can be thermally activated in feasible cases if the stack is designed so that the interlayer temperature falls within the transformation range of the selected alloy. However, the SMA layer must also remain below its functional temperature limit, which prevents it from being placed too close to the external surface in severe heating cases.

The backface temperatures also show why the multilayer concept is useful. In the feasible cases, the backface temperatures were much lower than the surface temperatures. For Case 2, the surface temperature was approximately 1856 K while the vehicle skin temperature was approximately 389 K. For Case 3, the surface temperature was approximately 1199 K while the vehicle skin temperature was approximately 345 K. For Case 5, the surface side of the stack reached approximately 2201 K while the vehicle skin remained near 364 K. These results show that the stack can provide substantial through thickness thermal attenuation when the outer material remains within its allowable temperature range. The key failure mechanism in the NULL cases was not that the lower layers immediately overheated, but that the external surface temperature exceeded the allowable range of the highest temperature outer TPS materials considered in this study.

When compared against the literature values discussed in Chapter 5, the successful heat flux cases of 0.1 MW/m², 0.62 MW/m², and 1.2 MW/m² suggest that the proposed ATPS architecture is potentially feasible with real world blunt body re-entry vehicles. The 0.1 MW/m² case is at the low end of the HIAD level heating range discussed earlier and represents a low severity acreage condition. The 0.62 MW/m² case is slightly above the approximately 0.4 MW/m² peak value cited for higher heating HIAD conditions, while the 1.2 MW/m² case approaches the higher blunt body heat flux range and remains just below the approximately 1.5 MW/m² Orion heat shield value discussed in Chapter 5. This is important because HIAD and Orion are both blunt heat shield systems, making the comparison relevant to the axisymmetric sphere cone and capsule-like TPS application. Notably however the Orion capsule is subjected to lunar return re-entry conditions which is a more severe aerothermal environment than a typical orbital re-entry trajectory. Based on this comparison, the successful cases support practical feasibility for a real world blunt forebody, capsule, or heat shield region where the local heating is on the order of 0.1 to 1.2 MW/m². The strongest successful result is the 1.2 MW/m² case, where the stack remained within the approximate 150 mm thickness cap while reducing the vehicle skin temperature to approximately 364 K and keeping the SMA interlayer near the phase transformation temperature envelope. This suggests that the architecture can remain viable near the upper end of realistic blunt body acreage heating, provided that the outer ceramic system is selected with sufficient temperature margin.

The leading edge comparison is more restrictive. The literature values discussed in Chapter 5 place a conservative blunted leading edge heating range near 2.0 to 3.0 MW/m², based partly on reported heat pipe cooled leading edge environments of about 1.6 to 2.0 MW/m² at Mach 8 and the expectation that orbital re-entry conditions could scale higher. Case 1 produced a heat flux of 3.3 MW/m², which is at or slightly above this conservative leading edge range, and the resulting wall temperature was approximately 2819 K. This exceeded the operating capability of the reusable tile and coating systems considered in the current material table. Therefore, the analysis indicates that the proposed SMA based ATPS stack is not well suited for the most severe real world leading edge stagnation point environments without a different hot structure, active cooling approach, ablative layer, or more specialized UHTC leading edge system. In contrast, the

successful 0.62 MW/m² and 0.1 MW/m² wedge cases suggest that the architecture may still be feasible downstream of the leading edge, where the heating has relaxed into a more acreage-like panel environment.

The results therefore support a bounded conclusion. The SMA based ATPS concept is not a complete replacement for all TPS regions on a hypersonic vehicle. It is not suitable, in its current form, for the most severe stagnation point or geometry transition heating cases predicted by the CFD model. However, it does show feasibility in both the CFD generated moderate heating regions and nominal real world blunt re-entry body peak heating conditions where the surface temperature can be controlled using a realistic high temperature ceramic layer and where the SMA interlayer reaches a useful activation temperature without exceeding its functional range. This is consistent with the expected role of adaptive TPS materials: they are most realistic as part of a layered reusable TPS architecture, not as exposed heat shield materials or as universal replacements for ablators and ultra high temperature leading edge systems.

Overall, the results show that the proposed ATPS stack can be made thermally viable for selected heating regions, especially for heat fluxes between approximately 0.1 and 1.2 MW/m2. Higher fidelity modeling and real world testing would still be needed to confirm this concept for industry wide application.

7.0 SMA Recovery Stress Analysis

The thermal model developed in Chapter 6 determines the through-thickness temperature distribution of the Adaptive Thermal Protection System (ATPS) stack under aerodynamic heating. While that analysis establishes whether the underlying vehicle structure remains thermally protected, it does not by itself determine whether the embedded smart-material layer provides a useful mechanical function. The purpose of the present chapter is therefore to evaluate the first-order structural response associated with the embedded Nitinol shape-memory alloy (SMA) mesh interlayer.

The central objective of this chapter is to estimate the recovery-stress response of the Nitinol interlayer when it is heated from its installed martensitic reference condition toward the austenitic transformation regime. In the proposed ATPS architecture, the Nitinol mesh is assumed to be prestrained and embedded near the ceramic tile backface. When the mesh is heated into its transformation range, it attempts to recover its trained shape. Because it is constrained by the surrounding tile, buffer, and bonded support layers, this recovery strain is partially blocked. The blocked recovery produces internal stress within the SMA, which is then transferred into the surrounding TPS stack as a reduced effective stress. The aim is that the recovery stress of the SMA puts the structure of the TPS in a more compressed state, preventing cracking and bond delamination by working against the tensile forces caused by thermal expansion.

7.1 Scope and Limitations

The present analysis is not intended to represent a full three-dimensional fracture or delamination simulation. Instead, it provides a reduced-order constitutive screening model that links the thermal state of the Nitinol layer to an estimated recovery-stress envelope. The model is implemented in MATLAB and evaluates the SMA response over the expected temperature excursion from the installed ambient state to the peak thermal state and back through cooling. This allows the governing stress state to be identified even when it does not occur exactly at the peak temperature.

The structural analysis therefore serves three functions. First, it determines whether the Nitinol layer reaches a temperature at which austenitic transformation and recovery response are possible. Second, it estimates the raw constitutive stress generated within the SMA material itself. Third, it reduces that raw material-level stress into an effective transferred stress acting on the surrounding TPS stack using mesh geometry and load-transfer efficiency factors. This structure preserves the distinction between the stress that exists inside the SMA and the stress that is mechanically transmitted into the surrounding TPS architecture.

7.2 Physical Basis of SMA Recovery Stress

Nitinol and related temperature NiTi-family alloys exhibit a reversible solid-state phase transformation between martensite and austenite. Martensite is the lower-temperature phase and is comparatively more deformable. In this phase, the material can be prestrained by reorienting its martensitic variants. Austenite is the higher-temperature phase and corresponds to the material's trained or recovered configuration. When the SMA is heated from martensite into austenite, it attempts to recover the deformation introduced during prestraining.

For a free SMA element, this transformation produces a recoverable shape change. For a constrained SMA element, however, the same transformation produces recovery stress instead of free deformation. In the present ATPS concept, the Nitinol mesh is embedded within the TPS stack and is therefore not allowed to recover freely. Its attempted recovery is resisted by the surrounding material system, resulting in internal stress generation within the SMA and an equal-and-opposite mechanical reaction on the adjacent TPS layers.

The recovery-stress mechanism can be described conceptually by distinguishing between stored recoverable strain and realized internal stress. Prestrain introduced in the martensitic state stores recoverable deformation, but the installed reference state is not assumed to impose a large initial stress on the tile. Instead, the installed prestrained martensitic condition is treated as the zero-stress reference state. Stress develops only as the temperature increases away from this reference condition and the material stiffness, thermal expansion state, and phase fraction evolve.

This distinction is important for the reduced-order model. The SMA should not be modeled as imposing a large initial load on the tile at ambient temperature simply because it was prestrained. The useful adaptive stress is generated during heating, when the SMA attempts to recover and the surrounding architecture partially constrains that recovery. The governing physical idea is therefore not prestrain alone, but constrained transformation of a prestrained material.

7.3 Transformation Temperatures and Viability Criteria

The first structural screening step is thermal rather than mechanical. The Nitinol layer must reach a temperature at which austenitic transformation begins in order for the recovery-stress mechanism to become active. The peak Nitinol temperature obtained from the Chapter 6 thermal analysis is denoted by T_{peak}. For a selected SMA alloy with austenite start temperature A_s, a minimum activation requirement is

$$T_{\text{peak}} \geq A_s \tag{42}$$

If this condition is not satisfied, the SMA does not enter its intended transformation regime and the selected alloy or stack architecture is not considered viable for the adaptive function. Reaching A_s indicates the onset of transformation, but full austenitic transformation requires the

temperature to reach the austenite finish temperature A_f. Thus, the stronger condition for full activation is

$$T_{\text{peak}} \geq A_f \tag{43}$$

The model also includes an upper allowable temperature for the selected SMA alloy, denoted T_{allow}. This limit is used to prevent the material from being thermally overdriven beyond its assumed useful operating range. The corresponding upper bound criterion is

$$T_{\text{peak}} \leq T_{\text{allow}} \tag{44}$$

The overall thermal viability condition for the selected SMA is therefore

$$A_s \leq T_{\text{peak}} \leq T_{\text{allow}} \tag{45}$$

If $T_{peak}<A_s$ the MATLAB model terminates the SMA stress calculation because the layer does not reach the desired transformation regime. If $T_{peak}>T_{allow}$ the model also terminates because the selected alloy is being exposed beyond its intended operating range. Only when the peak temperature lies within the allowable activation window does the model proceed to compute the stress-temperature response.

The heating branch of the transformation is governed by the austenite start and finish temperatures, A_s and A_f. The cooling branch is governed by the martensite start and finish temperatures, M_s and M_f. These two pairs of temperatures describe different directions of the hysteresis loop. During heating, the material transforms from martensite toward austenite. During cooling, it transforms from austenite or partially transformed material back toward martensite. Accordingly, A_s and A_f are used for the heating path, while M_s and M_f are used for the cooling path.

7.4 Temperature-Path Representation

The present analysis uses a steady-state peak Nitinol temperature from the thermal model rather than a full transient temperature-time history. Therefore, the MATLAB model does not predict the exact time evolution of SMA stress during flight. Instead, it constructs a constitutive stress-temperature path over the expected thermal excursion. The path begins at the installed ambient reference temperature T_0, heats to the peak temperature T_{peak}, and then cools back toward ambient temperature.

The heating temperature path is represented as

$$T_{\text{heat}} \in [T_0,\ T_{\text{peak}}] \tag{46}$$

and the cooling path is represented as

$$T_{\text{cool}} \in [T_{\text{peak}},\ T_0] \tag{47}$$

This approach does not resolve dwell time, heating rate, or transient thermal lag. However, it is useful for identifying the stress envelope associated with the full range of temperatures that the embedded SMA is expected to experience. This is important because the governing effective stress may occur at a transformation boundary or during cooling rather than exactly at T_{peak}.

The martensite fraction is denoted by ξ, where ξ=1 corresponds to fully martensitic material and ξ=0 corresponds to fully austenitic material. For the heating branch, the model assumes that the material remains martensitic below A_s, transforms linearly between A_s and A_f, and becomes fully austenitic above A_f. The heating-branch martensite fraction is therefore approximated by the piecewise function (43).

$$\xi_{\text{heat}}(T) = \begin{cases} 1, & T \le A_s \\ \frac{A_f - T}{A_f - A_s}, & A_s < T < A_f \\ 0, & T \ge A_f \end{cases} \tag{48}$$

During cooling, the material must begin from the phase state actually reached at the end of heating. This is necessary to avoid an artificial discontinuity between the heating and cooling branches. Let the martensite fraction at peak heating be

$$\xi_{\text{peak}} = \xi_{\text{heat}}\left(T_{\text{peak}}\right) \tag{49}$$

The cooling branch then retains this peak-heating phase state until the material cools below M_s, after which it transforms back toward martensite until M_f. The cooling-branch martensite fraction is approximated as the piecewise function (45).

$$\xi_{\text{cool}}(T) = \begin{cases} \xi_{\text{peak}}, & T \ge M_s \\ \xi_{\text{peak}} + (1 - \xi_{\text{peak}})\,\dfrac{M_s - T}{M_s - M_f}, & M_f < T < M_s \\ 1, & T \le M_f \end{cases} \tag{50}$$

This piecewise representation is a reduced-order approximation of SMA hysteresis. It captures the basic directional behavior of the heating and cooling transformations while avoiding the complexity of a full thermomechanical hysteresis model. The resulting stress temperature

curve should therefore be interpreted as a constitutive screening envelope rather than a time-resolved transient prediction.

7.5 Phase-Dependent Material Properties

The SMA constitutive response depends on the phase state of the material. The elastic modulus is modeled as a phase-fraction-weighted mixture of the martensite modulus E_M and austenite modulus E_A. The effective modulus is written as

$$E(\xi) = \xi E_M + (1 - \xi)\, E_A \tag{51}$$

At the installed reference condition, the material is assumed to be fully martensitic, so ξ_0=1 and therefore $E(\xi_0) = E_M$.

The thermal-expansion coefficient is also modeled as phase dependent. The martensitic and austenitic coefficients of thermal expansion are denoted by α_M and α_A, respectively. The effective coefficient of thermal expansion is approximated as

$$\alpha(\xi) = \xi \alpha_M + (1 - \xi)\, \alpha_A \tag{52}$$

The corresponding thermoelastic coefficient used in the constrained thermal-stress term is

$$\Theta(\xi) = \lambda_c E(\xi)\alpha(\xi) \tag{53}$$

where λ_c is the constraint factor. This factor represents the degree to which the surrounding TPS architecture prevents the SMA from freely accommodating thermal and transformation-induced strain. A value of λ_c=1 corresponds to fully constrained recovery, while λ_c=0 corresponds to no effective stress generation due to complete release. Intermediate values represent partial constraint.

The use of λ_c in the thermal-stress term is important because a partially constrained interlayer should not generate the same thermal stress as a perfectly blocked member. The same physical idea is applied to the elastic state-change and transformation terms: the model estimates the stress that would be generated by phase and temperature changes, then scales that stress by the assumed degree of mechanical constraint.

7.6 Installed Reference State and Constraint Treatment

The installed reference state is defined as the prestrained martensitic state at ambient temperature. This reference condition is assigned zero stress in the reduced-order model. The prestrain is therefore treated as stored recoverable deformation rather than as an initial load applied to the TPS stack. This avoids the nonphysical result of imposing a large initial stress at the time of installation.

The installed prestrain is denoted by ε_{ref}. At the reference state,

$$T = T_0, \qquad \xi = \xi_0, \qquad \sigma = \sigma_0 \tag{54}$$

with σ_0=0. The model is formulated so that all stress terms vanish at the reference state. The elastic contribution is written as an increment from the installed reference state,

$$\sigma_{\text{elastic}} = \lambda_c \left[E(\xi) - E(\xi_0) \right] \varepsilon_{\text{ref}} \tag{55}$$

This form ensures the below relation.

$$\sigma_{\text{elastic}} = 0 \quad \text{when} \quad \xi = \xi_0 \tag{56}$$

The thermal contribution is similarly written as an incremental term from the reference temperature,

$$\sigma_{\text{thermal}} = \lambda_c E(\xi) \alpha(\xi) \left(T - T_0 \right) \tag{57}$$

Which satisfies

$$\sigma_{\text{thermal}} = 0 \quad \text{when} \quad T = T_0 \tag{58}$$

If stress-induced martensite terms are included in a later calibrated model, the transformation contribution can be written as

$$\sigma_{\text{transform}} = \lambda_c \left[\Omega_p \xi_p + \Omega_m \xi_m - \Omega_p \xi_{p0} - \Omega_m \xi_{m0} \right] \tag{59}$$

where Ω_p and Ω_m are transformation coefficients associated with tension- and compression-induced martensite, and ξ_p and ξ_m are the corresponding martensite fractions. In the present first-pass implementation, these terms are left inactive unless later calibrated, such that $\Omega_p = \Omega_m = 0$. The total raw SMA constitutive stress is therefore

$$\sigma_{\text{SMA}} = \sigma_0 + \sigma_{\text{elastic}} + \sigma_{\text{transform}} + \sigma_{\text{thermal}} \tag{60}$$

This equation is the core of the reduced-order stress model. It estimates the material-level stress generated inside the SMA due to constrained phase-dependent stiffness change and constrained thermal expansion. The result is not directly interpreted as the stress acting on the ceramic tile or bondline. Additional reduction is required to account for the fact that the SMA is present as a mesh and that not all of its internal stress is transferred into the surrounding structure. This is done in the sections proceeding.

7.7 Mesh Geometry and Effective SMA Fraction

The embedded Nitinol layer is represented as a mesh rather than a continuous sheet. The mesh is modeled as a homogenized effective interlayer whose active material fraction is determined by the wire diameter and wire spacing. This representation preserves the first-order effect of SMA areal density without resolving the individual wire geometry in the structural model.

For a simple orthogonal square mesh with wire diameter d and center to center spacing s, the effective mesh fraction is approximated by adding the area covered by the two orthogonal wire families and subtracting their overlap at intersections. The resulting expression is

$$v_{f,\mathrm{mesh}} = \frac{2d}{s} - \left(\frac{d}{s}\right)^2 \tag{61}$$

where $v_{f,mesh}$ is the effective SMA fraction of the interlayer region. This parameter accounts for the fact that only a portion of the interlayer area is occupied by active Nitinol. A denser mesh, larger wire diameter, or smaller spacing increases $v_{f,mesh}$, thereby increasing the effective stress that can be transferred to the surrounding TPS stack. A more open mesh decreases $v_{f,mesh}$, reducing both the mechanical contribution of the SMA and the interfacial load introduced by the mesh.

The mesh representation is therefore an important design variable. It allows the SMA response to be tuned without changing the intrinsic material behavior. The raw SMA stress may be large at the material level, but the effective mechanical action on the stack is reduced by the fraction of active material present in the mesh.

7.8 Load Transfer Efficiency

The second reduction applied to the raw SMA stress is the load-transfer efficiency factor, denoted by $\eta_{transfer}$. This parameter represents how effectively the stress generated inside the Nitinol mesh is transmitted into the surrounding TPS stack. It is not a material property of Nitinol. Instead, it is a reduced-order representation of unresolved local mechanics, including bond quality, local compliance, embedding condition, slip, mesh anchoring, and stress-transfer nonuniformity. The effective transferred stress on the surrounding TPS stack is modeled as

$$\sigma_{\mathrm{stack,eff}} = -\eta_{\mathrm{transfer}} v_{f,\mathrm{mesh}} \sigma_{\mathrm{SMA}} \tag{62}$$

The negative sign indicates that tensile stress in the constrained SMA produces an equal-and-opposite compressive reaction on the surrounding TPS stack. If the SMA attempts to shorten during heating and is prevented from doing so, the SMA itself develops tensile recovery stress, while the adjacent material system experiences a compressive bias. The limiting values of

$\eta_{transfer}$ have clear physical interpretations. If $\eta_{transfer} = 0$ then the mech is mechanically ineffective and no useful stress is transferred into the surrounding structure. If $\eta_{transfer} = 1$ then the model assumes ideal load transfer from the SMA mesh into the adjacent TPS layers. Practical values are expected to lie between these limits. In the MATLAB implementation, the user specifies a primary value of $\eta_{transfer}$ for the main stress calculation, while additional values may be plotted as a sensitivity sweep.

The load-transfer efficiency is retained as a tunable parameter because the present thesis does not resolve the detailed local bond mechanics of the SMA mesh. A future higher fidelity analysis could calibrate this parameter from a local finite-element submodel or from an experimental coupon test. In the present reduced order model, however, $\eta_{transfer}$ provides a practical way to assess how strongly the assumed bonding and embedding architecture influences the effective structural response.

7.9 Stress-Temperature Envelope and Governing Stress

Because the present model evaluates a full temperature path rather than only a single peak condition, it can identify the governing stress over the entire heating and cooling excursion. The stress at peak temperature is denoted by $\sigma_{stack,eff}(T_{peak})$. However, the maximum stress magnitude may occur elsewhere along the temperature path. Therefore, the governing effective transferred stress is defined as shown below.

$$\sigma_{\text{gov}} = \max_{T \in \mathcal{T}} |\sigma_{\text{stack,eff}}(T)| \tag{63}$$

Where T is the full evaluated temperature path,

$$\mathcal{T} = [T_0 \rightarrow T_{\text{peak}} \rightarrow T_0] \tag{64}$$

The corresponding governing temperature is

$$T_{\text{gov}} = \arg\max_{T \in \mathcal{T}} |\sigma_{\text{stack,eff}}(T)| \tag{65}$$

This definition is useful because the largest transferred stress can occur near transformation boundaries or during cooling, depending on the phase state reached at peak heating and the assumed constraint condition. Therefore, the peak temperature and the peak mechanical stress are not necessarily identical.

The MATLAB model reports both the stress at T_{peak} and the governing stress over the full temperature excursion. It also reports the effective transferred stress at key transformation temperatures, including A_s, A_f when reached, M_s, and M_f. These values allow the user to determine whether the critical mechanical demand is associated with heating into austenite, partial transformation at peak temperature, or cooling back toward martensite.

7.10 MATLAB Implementation

The MATLAB implementation follows the thermal screening workflow established in Chapter 6, but applies it to the structural response of the Nitinol mesh. The peak Nitinol temperature is imported from the thermal analysis and assigned as T_{peak}. The code first evaluates whether Tpeak lies within the selected SMA activation and operating window. If the peak temperature is below A_s, the design is rejected for insufficient SMA activation. If the peak temperature exceeds T_{allow}, the design is rejected for over-temperature exposure.

After this thermal viability check, the code constructs the heating and cooling temperature vectors. At each temperature step, it evaluates the martensite fraction, effective modulus, phase-dependent thermal expansion coefficient, raw SMA constitutive stress, and effective transferred stress. The heating and cooling paths are made continuous by initializing the cooling branch from the martensite fraction reached at the end of the heating branch. This prevents the model from artificially assuming that the material is fully austenitic during cooling if T_{peak} did not reach A_f.

The computational workflow may be summarized as follows. The thermal input Tpeak is first checked against the SMA transformation window. The heating branch is then evaluated from T_0 to T_{peak}, using A_s and A_f to define the martensite fraction. The cooling branch is then evaluated from T_{peak} back to T_0, using M_s and M_f while preserving the phase state reached at peak heating. Finally, the raw SMA stress is converted into effective stack stress using the mesh fraction and load-transfer efficiency.

The reduced order MATLAB model therefore provides four primary outputs: the raw SMA stress at peak temperature, the effective transferred stack stress at peak temperature, the governing effective transferred stress over the full temperature excursion, and the temperature at which this governing stress occurs. These outputs provide the structural-response counterpart to the thermal interface temperatures computed in Chapter 6.

7.11 Results

3 Cases will be evaluated for the SMA recovery stress response analysis. Each case is coupled with one of the 3 successful thermal analysis cases. Those being case 2, 3, and 5. The load transfer efficiency and mesh geometry are kept at fixed values but the peak temperature and transformation temperatures are case dependent. These values are reported in tables 7.1 and 7.2 with the peak SMA temperature coming from the upstream thermal analysis and the listed transformation temperatures coming from literature data. The outputs are reported for each case.

Table 7.1 - SMA recovery stress response analysis: inputs

SMA case	Thermal case	Heat flux	Peak SMA temperature	Selected SMA	Transformation temperatures

SMA Case A	Case 3	0.1 MW/m²	448.42 K	Cu-Al-Ni polycrystal	Mf = 391 K; Ms = 425 K; As = 422 K; Af = 436 K
SMA Case B	Case 2	0.62 MW/m²	587.82 K	NiTiHf	Mf = 466 K; Ms = 530 K; As = 546 K; Af = 604.65 K
SMA Case C	Case 5	1.2 MW/m²	685.62 K	NiTiHf	Mf = 466 K; Ms = 530 K; As = 546 K; Af = 604.65 K

Table 7.2 - SMA recovery stress response analysis: fixed inputs

Input	Value
Wire Diameter (d_{wire})	0.2 mm
Wire Spacing (s_{wire})	2.5mm
Load Transfer Efficiency ($\eta_{transfer}$)	0.2

7.11.1 Case A Results

Table 7.3 - Case A: primary effective stack stress

Metric	**Result**
Transferred stress at T_peak	-11.285 MPa
Governing stress occurs at	448.20 K
Transferred stress at As	-0.659 MPa
Transferred stress at Af	-11.023 MPa

Transferred stress at Ms	-10.942 MPa
Transferred stress at Mf	-0.453 MPa

Table 7.4 - Case A: secondary eta sweep comparison

Transfer Case	**eta_transfer**	**Transferred Stress at T_peak**	**Governing Stress occurs at**
Low transfer	0.10	-5.643 MPa	448.20 K
Moderate transfer	0.30	-16.928 MPa	448.20 K
High transfer	0.50	-28.213 MPa	448.20 K

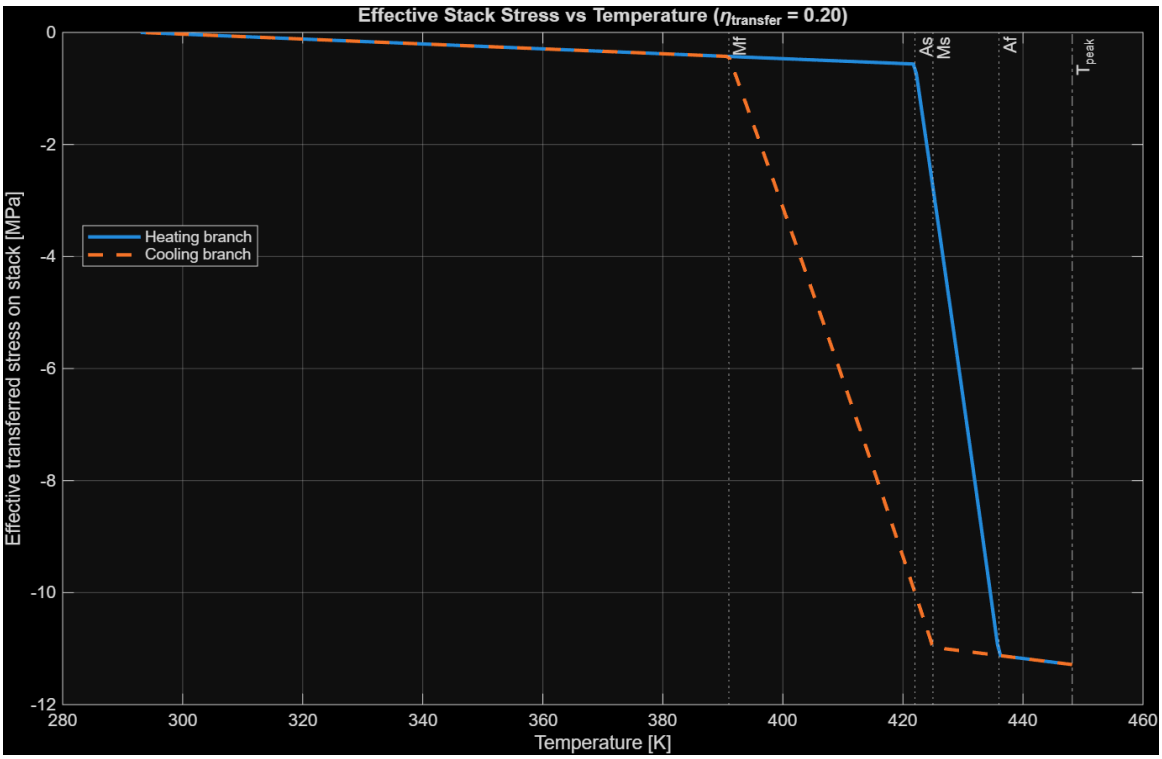


Figure 7.1 - Case A: primary effective stack stress vs. temperature

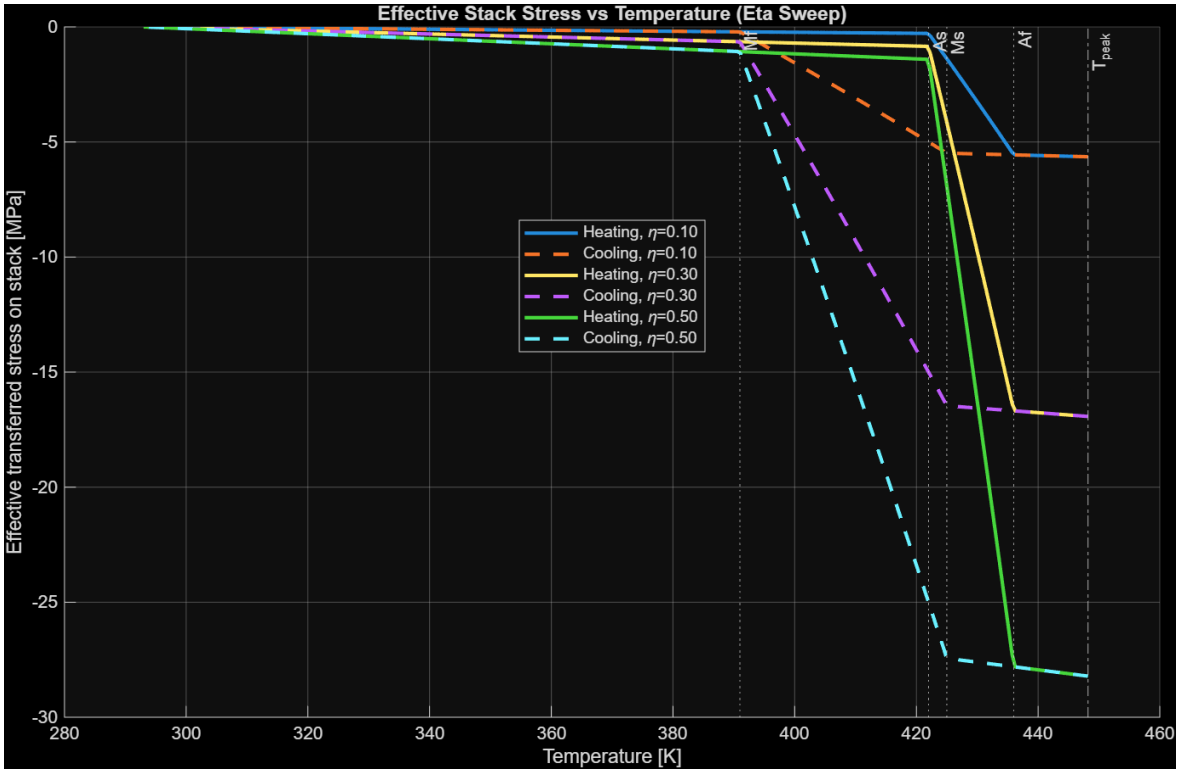


Figure 7.2 - Case A: eta sweep effective stack stress vs. temperature

7.11.2 Case B Results

Table 7.5 - Case B: primary effective stack stress

Metric	Result
Transferred stress at T_peak	-9.586 MPa
Governing stress occurs at	587.82 K
Transferred stress at As	-1.152 MPa
Transferred stress at Af	Not reached, T_peak < Af
Transferred stress at Ms	-8.958 MPa
Transferred stress at Mf	-0.786 MPa

Table 7.6 - Case B: secondary eta sweep comparison

Transfer Case	eta_transfer	Transferred Stress at T_peak	Governing Stress Occurs At
Low transfer	0.10	-4.793 MPa	587.82 K
Moderate transfer	0.30	-14.379 MPa	587.82 K
High transfer	0.50	-23.965 MPa	587.82 K

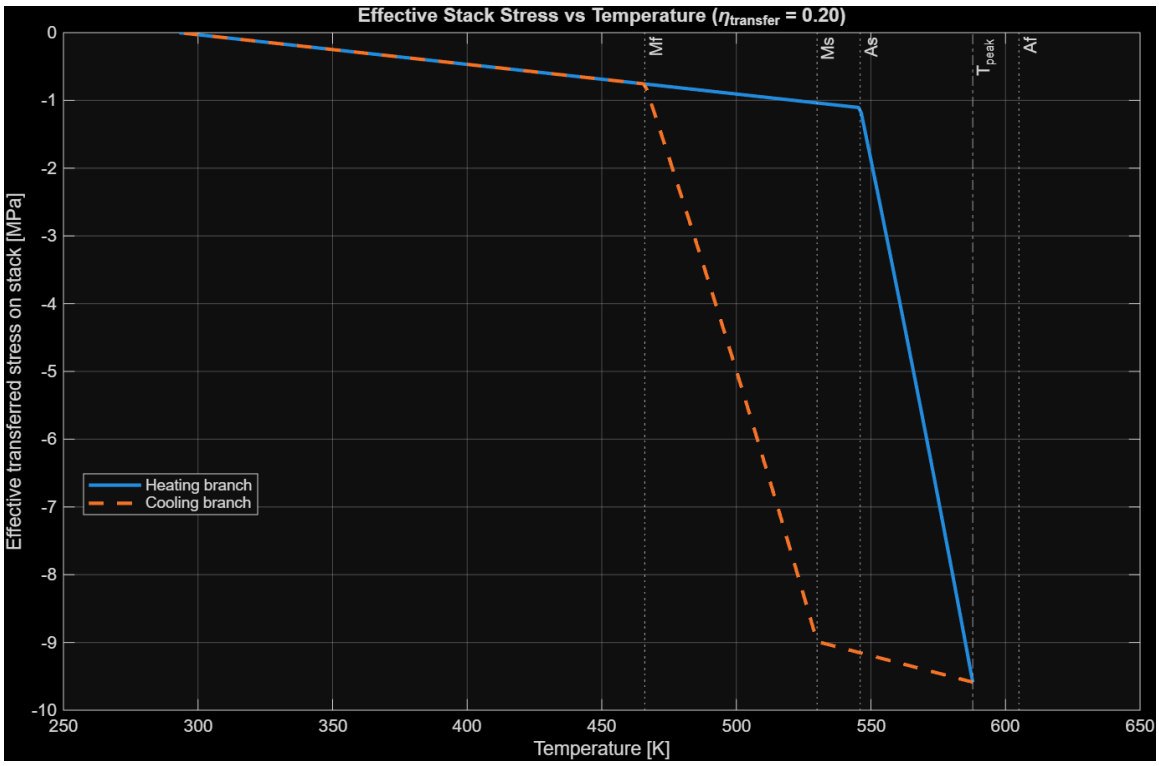


Figure 7.3 - Case B: primary effective stack stress vs. temperature

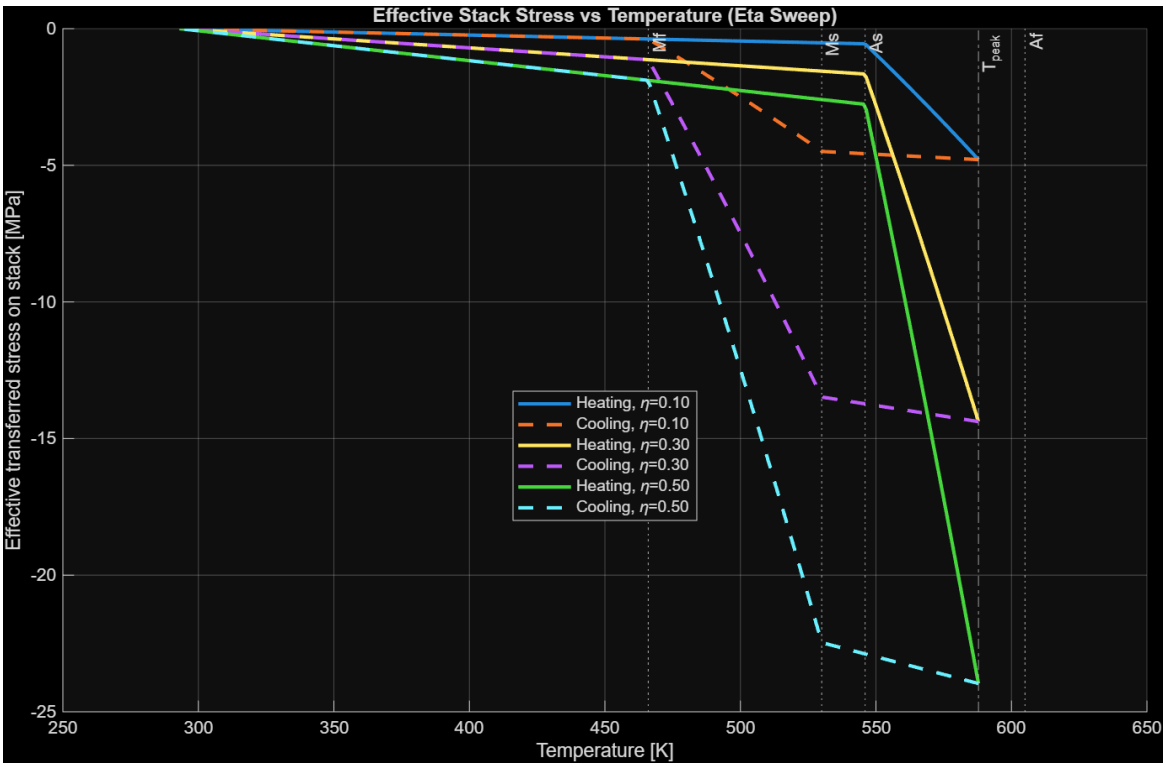


Figure 7.4 - Case B: eta sweep effective stack stress vs. temperature

7.11.3 Case C Results

Table 7.7 - Case C: primary effective stack stress

Metric	Result
Transferred stress at T_peak	-14.451 MPa
Governing stress occurs at	685.62 K
Transferred stress at As	-1.162 MPa
Transferred stress at Af	-13.309 MPa
Transferred stress at Ms	-12.316 MPa
Transferred stress at Mf	-0.798 MPa

Table 7.8 - Case C: secondary eta sweep comparison

Transfer Case	eta_transfer	Transferred Stress at T_peak	Governing Stress Occurs At
Low transfer	0.10	-7.225 MPa	685.62 K
Moderate transfer	0.30	-21.676 MPa	685.62 K
High transfer	0.50	-36.126 MPa	685.62 K

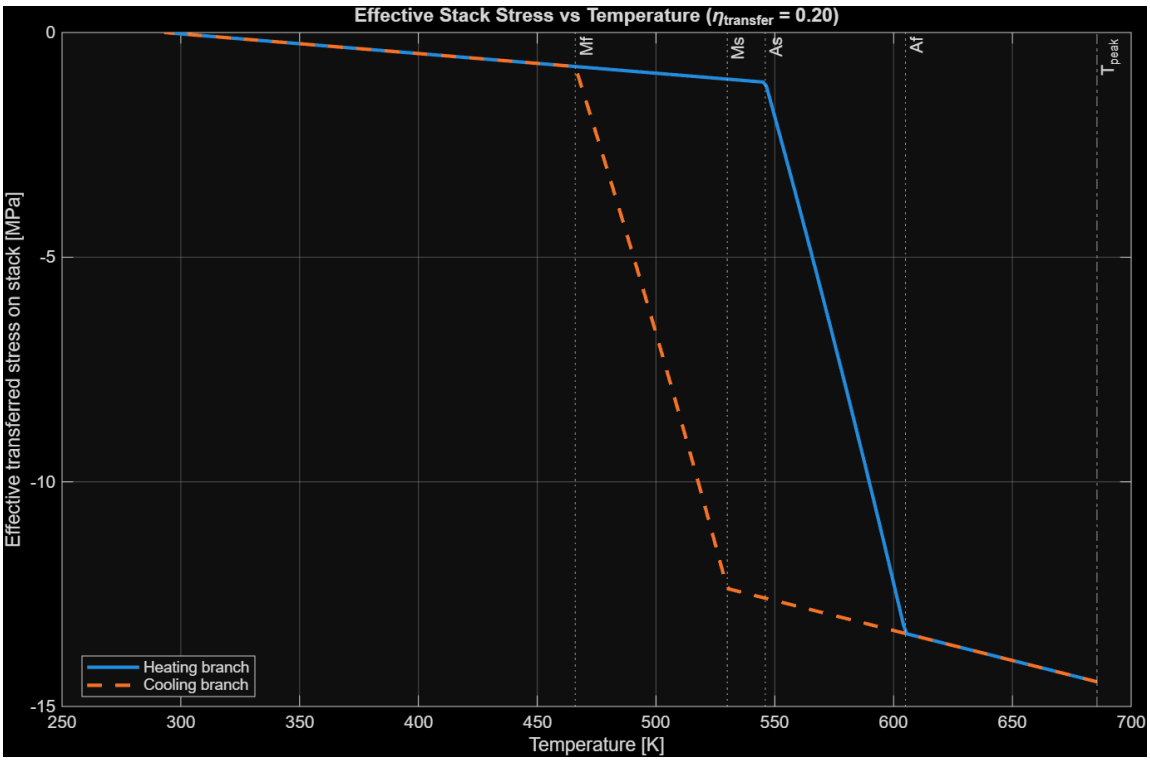


Figure 7.5 - Case C: primary effective stack stress vs. temperature

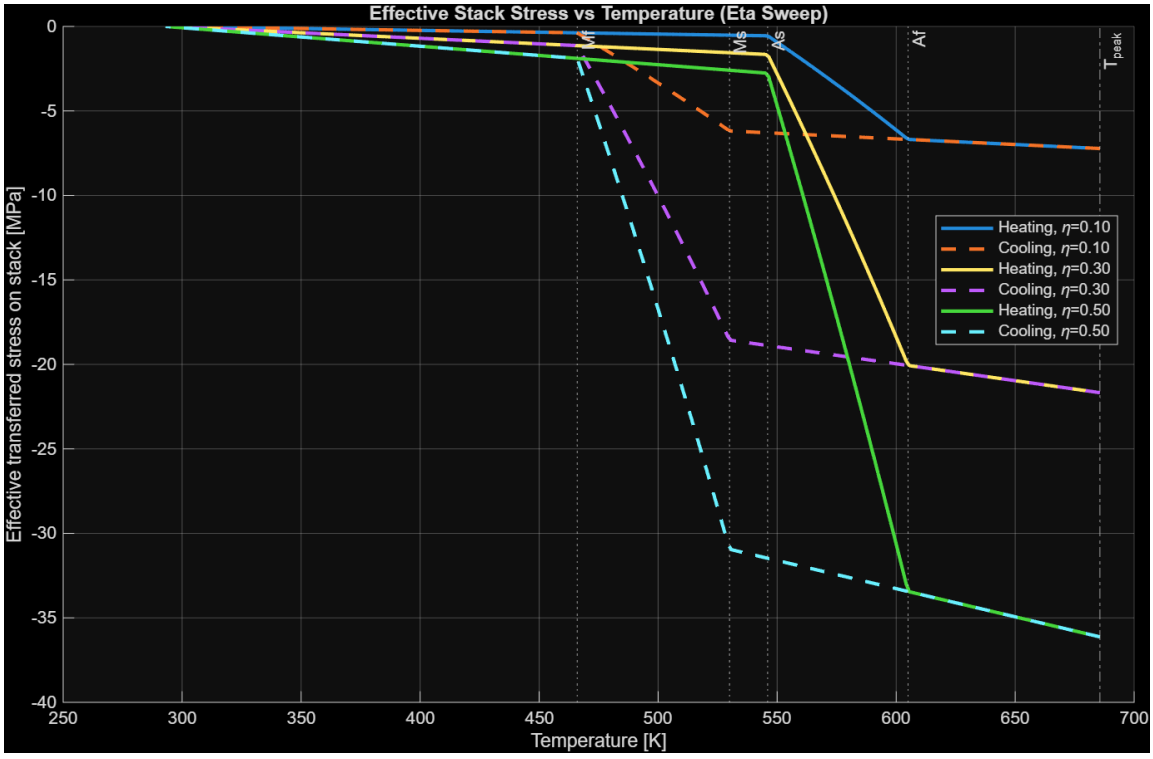


Figure 7.6 - Case C: eta sweep effective stack stress vs. temperature

8.0 Conclusion and Recommendations

8.1 Summary of Work

This study investigated the feasibility of an SMA-based Adaptive Thermal Protection System (ATPS) for representative hypersonic re-entry applications. The proposed system was not treated as a standalone smart material heat shield. Instead, it was developed as a multilayer thermal protection architecture in which the primary thermal protection function is carried by a ceramic outer system, while the adaptive thermomechanical response is provided by an embedded shape memory alloy interlayer. This distinction is important because the SMA is not intended to survive direct re-entry exposure or act as the primary thermal barrier. Its intended role is to operate at an intermediate temperature inside the stack, where it can undergo phase transformation, generate recovery stress, and potentially improve resistance to crack opening, debonding, or local delamination.

The overall workflow combined geometry development, structured CFD meshing, hypersonic flow simulation, reduced-order thermal analysis, and SMA recovery stress modeling. Two representative external aerodynamic geometries were evaluated. The first geometry was an axisymmetric blunt forebody representative of a capsule, sphere-cone, or rounded re-entry heat shield. The second geometry was a planar blunted wedge representative of a hypersonic leading edge or airfoil-like surface. These simplified geometries were selected to capture two different thermal loading environments: a blunt forebody stagnation and acreage heating environment, and a leading edge heating environment where the flow is forced around a finite radius nose and then over sloped wedge panels.

The CFD portion of the study was performed in ANSYS Fluent using a density-based hypersonic setup with energy enabled, finite rate chemistry, species transport, and the Park 93 five-species air model. The freestream condition was selected to represent a high altitude Mach 20 re-entry environment at approximately 65 km, with freestream static temperature near 231 K and freestream static pressure near 10 Pa. The CFD model used pressure farfield boundaries rather than a velocity inlet because the pressure farfield formulation is more appropriate for compressible external hypersonic flow. The pressure farfield condition allows Fluent to impose the freestream Mach number, pressure, temperature, and flow direction in a manner consistent with external aerodynamics, while avoiding some of the nonphysical behavior that can occur when a velocity inlet is forced onto a large hypersonic external domain.

A mapped mesh strategy was used for both geometries to preserve element quality and near-wall resolution under the ANSYS Student license cell-count limitation. The wedge geometry was divided into four mapped faces, while the sphere-cone required six mapped faces because of its stronger shoulder and geometry transition region. Edge sizing and biased division controls were used to concentrate cells toward the body, shock layer, and wake-aligned regions. Near-wall sizing was selected using a y+ based estimate to provide adequate boundary layer

resolution for the k-omega SST turbulence model. This meshing approach provided a practical balance between wall resolution, shock capture, wake resolution, and computational cost.

The CFD results were then passed into a MATLAB thermal analysis framework. The thermal model solved a simplified through-thickness heat transfer problem for a multilayer ATPS stack. The model included surface radiation, conduction through each material layer, and backface temperature prediction. The wall temperature was calibrated by balancing the incoming CFD heat flux with radiation from the external surface and conduction into the TPS stack. This allowed the CFD generated aerodynamic heating to be converted into material temperatures through the stack, including the outer coating temperature, ceramic tile temperature, SMA interlayer temperature, buffer layer temperature, SIP temperature, adhesive temperature, and vehicle skin temperature.

The final part of the work evaluated the SMA recovery stress response using the peak SMA interlayer temperatures predicted by the thermal model. The SMA model was not intended to prove full structural integrity or directly simulate delamination suppression. Instead, it was developed as a reduced-order screening model to estimate the effective transferred recovery stress that could be generated by a constrained SMA mesh embedded within the TPS stack. The model accounted for phase transformation temperature ranges, prestrain, effective mesh volume fraction, and a load-transfer efficiency factor. This allowed the study to estimate whether the SMA layer reached a useful activation temperature and whether the resulting recovery stress was large enough to provide a meaningful adaptive mechanical response.

8.2 Major Findings

The CFD results demonstrated that the local heat flux distribution is strongly controlled by geometry and prescribed wall conditions. The axisymmetric sphere-cone did not produce a simple heat flux profile where the maximum heating occurred only at the stagnation point and then smoothly decreased along the body. Instead, the CFD results showed localized off-stagnation heat flux peaks near the shoulder region where the rounded forebody transitions into the downstream body profile. The wall shear stress distribution aligned with these heat flux peaks, supporting the conclusion that the local heating feature was driven by the flowfield and geometry rather than being purely numerical. The shoulder region increased the local near-wall velocity gradient and thermal gradient, which resulted in elevated heat flux away from the geometric stagnation point.

The blunted wedge showed a related but slightly different behavior via the same mechanism. The main leading edge heating was concentrated near the finite radius nose, but localized heating behavior was also observed near the region where the blunt tip transitioned into the sloped wedge body. This was interpreted as a nose-to-wedge transition heating feature rather than the same shoulder heating seen on the sphere-cone. In both geometries, the localized heat flux increases were associated with geometric transition regions where the flow turns, accelerates, or adjusts rapidly along the wall. These results support the conclusion that local TPS sizing cannot rely only on the nominal stagnation point or a smooth acreage heating assumption.

Even simple geometries can generate local heating features that become important thermal design drivers.

The contour plots further supported these observations. The axisymmetric sphere-cone produced a thin near body shock layer and strong oblique shock behavior downstream rather than a large visually detached bow shock. This does not mean that the blunt body physics were absent. Rather, the shock standoff distance was small relative to the body scale and the selected high altitude Mach 20 condition. The geometry also allowed the flow to turn around the rounded forebody and transition into a downstream oblique shock structure. The blunted wedge produced a clearer and stronger oblique shock system because the inclined wedge faces imposed a sustained flow turning angle. The Mach number drop and temperature increase were more pronounced along the wedge shock structure. Both geometries also produced similar separated wake structures downstream of the artificial aft truncation, which was expected because both CFD bodies were closed using simplified truncated aft surfaces.

A major limitation of the CFD methodology was the use of an isothermal wall boundary condition. The wall temperature was calibrated through the MATLAB thermal model, but Fluent still treated each major wall surface as having a fixed imposed temperature during each run. This simplified the coupling procedure and reduced computational cost, but it did not allow the wall temperature to vary locally along the body. In reality, higher heat flux regions would likely reach higher local radiative equilibrium temperatures than lower heating regions. A hotter local wall would reduce the gas to wall temperature difference and could reduce the predicted local convective heat flux. Therefore, some of the localized heat flux peaks, especially the most extreme shoulder and transition values, should be interpreted as potentially inflated by the isothermal wall assumption. Furthermore the stagnation point heat flux on the blunt forebody was higher as compared with the leading edge. This observed behavior is best explained by the fact that calibrated wall temperature prescribed on the sphere-cone was lower than on the blunt wedge, around 1004 K compared to 1705 K. This limitation is directly tied to the computational constraints of the study and the decision not to model the solid TPS stack directly in Fluent using full conjugate heat transfer.

The thermal results showed that the proposed ATPS architecture was feasible only within a bounded heating range. Three heating cases derived from local heat fluxes from the CFD results were concluded to be thermally successful: 0.1 MW/m2, 0.62 MW/m2, and 1.2 MW/m2. In these cases, the material stack was able to reduce the outer surface temperature to acceptable internal layer temperatures, while simultaneously maintaining SMA Interlayer temperatures within or slightly above their phase transformation envelope, and keeping the vehicle skin within a protected range below the cited 450 K [40]. The lowest heating case, 0.1 MW/m2, produced an outer surface temperature near 1199 K and reduced the vehicle skin temperature to approximately 345 K. The moderate wedge acreage case, 0.62 MW/m2, produced an outer surface temperature near 1856 K and reduced the vehicle skin temperature to approximately 348 K. The higher blunt forebody acreage case, 1.2 MW/m2, required a more capable material stack but remained feasible, producing an outer ceramic temperature near 2201 K, a TUFROC layer

temperature near 1822 K, an SMA interlayer temperature near 686 K, and a vehicle skin temperature near 364 K.

The results also identified heating cases that were outside the practical design envelope of the proposed ATPS stack. The 3.3 MW/m2 leading edge stagnation case produced a wall temperature near 2819 K, which exceeded the operating capability of the reusable ceramic tile and coating systems considered in the material table. The 8.0 MW/m2 blunt forebody stagnation case produced a wall temperature near 3600 K, and the 13.8 MW/m2 geometry transition case produced a wall temperature near 4055 K. These cases were marked as NULL because the outer thermal protection materials exceeded their allowable operating ranges. The key failure mechanism was not that every lower layer immediately overheated, but that the external surface temperature exceeded the capability of the coating and ceramic material system. This means that the proposed architecture is not appropriate for all re-entry heating environments. It is more appropriate for acreage, relaxation-zone, and moderate peak heating regions than for severe leading edge stagnation points, extreme geometry transition peaks, or highly concentrated local hot spots.

When the successful heat flux cases are compared with literature level heating values, the results support practical feasibility for selected blunt-body TPS regions. The successful values of 0.1 MW/m2, 0.62 MW/m2, and 1.2 MW/m2 cover a range that is relevant to blunt heat shield and capsule-like heating environments. The 0.62 MW/m2 case is comparable to higher HIAD-level and typical orbital re-entry peak heating values discussed earlier, while the 1.2 MW/m2 case approaches the higher end of blunt body heating and remains below the approximately 1.5 MW/m2 heat flux value discussed for Orion lunar return heat shield conditions. This is significant because the 1.2 MW/m2 case remained viable within the approximate 150 mm stack thickness cap. It suggests that the proposed ATPS architecture may be practically relevant for real world blunt forebody, capsule, or heat shield acreage regions where the local heat flux is on the order of 0.1 to 1.2 MW/m2.

The leading edge results were more restrictive. Literature values discussed in Chapter 5 place a conservative blunted leading edge heating range near 2.0 to 3.0 MW/m2, while the leading edge stagnation case in this study reached 3.3 MW/m2. Since this case produced a wall temperature beyond the selected reusable coating and ceramic tile limits, the current architecture should not be considered suitable for the most severe leading edge stagnation environments. However, the successful wedge cases at 0.62 MW/m2 and 0.1 MW/m2 suggest that the architecture may still be feasible downstream of a leading edge, where the heat flux has relaxed into a more acreage-like panel environment.

The SMA recovery stress analysis showed that the embedded SMA layer can generate a meaningful thermomechanical response when the thermal stack is viable and the SMA layer reaches the correct transformation temperature range. The recovery stress cases were based on the thermally successful heating cases. Case A, corresponding to the low heating case, produced a transferred stress at peak temperature of approximately -11.285 MPa. Case B, corresponding to the moderate heating case, produced a transferred stress at peak temperature of approximately

-9.586 MPa. Case C, corresponding to the highest successful heating case, produced a transferred stress at peak temperature of approximately -14.451 MPa. The $\eta_{transfer}$ sensitivity studies showed that the effective stack stress is strongly dependent on load transfer efficiency. For Case C, the transferred stress ranged from approximately -7.225 MPa at $\eta_{transfer}$=0.10 to approximately -36.126 MPa at $\eta_{transfer}$=0.50. This confirms that the interface mechanics and load transfer assumptions are critical to the actual effectiveness of the SMA layer.

The SMA results support the central concept of the study. The outer ceramic stack can reduce the external re-entry heating environment to an intermediate temperature range where an embedded SMA layer can activate. The SMA layer is therefore not exposed to the full external wall temperature, but it can still reach a temperature high enough to generate recovery stress. This is the main physical logic behind the proposed ATPS concept. The thermal stack protects the vehicle, while the SMA layer is positioned to respond mechanically inside the structure. However, the recovery stress results should be interpreted as reduced-order screening values. They do not prove that crack growth, debonding, or delamination will be eliminated. A complete demonstration of that effect would require detailed interface mechanics, cohesive-zone modeling, and experimental validation.

8.3 Conclusions

Overall, the results indicate that the proposed SMA-based ATPS architecture is feasible within a bounded thermal design space. It is not a universal replacement for ablative TPS, reinforced carbon-carbon, hot structures, actively cooled leading edges, or ultra-high-temperature ceramic leading edge systems. The architecture is most realistic for acreage regions, relaxed heating zones, and blunt body TPS surfaces where the local heat flux remains within a range that can be managed by a ceramic outer layer while still allowing the embedded SMA to activate. More work will need to be done to determine whether or not the proposed concept is viable for leading edge stagnation heating. Based on the results of this study, the practical heating range for the current architecture is approximately 0.1 to 1.2 MW/m2. Although composite ceramic structures such as the one proposed for the heating case 5 (1.2 MW/m2) could potentially support higher heating environments with a conservative estimate of up to 2.0 MW/m2. However the proposed system involved bonding a layer of Zirconia fibrous board and a UHTC emissivity coating to a reusable ceramic insulator both of which have limited heritage and bonding these layers together may introduce additional risk into the system.

The thermal screening results also show that the highest CFD heat flux cases are not feasible with the current material set and stack thickness. The 3.3 MW/m2 leading edge stagnation case, the 8.0 MW/m2 blunt forebody stagnation case, and the 13.8 MW/m2 geometry transition case exceeded the operating capability of the selected materials. These cases are still useful because they define the boundary of the architecture. They show where a reusable ceramic and SMA-based stack is no longer sufficient and where a different TPS solution would be required. These regions would likely require a dedicated UHTC leading edge, a hot structure, an ablative surface, active cooling, or a substantially different geometry and material architecture.

From an aerothermal standpoint, geometry has a strong influence on local heating behavior. The blunt forebody produced localized off-stagnation shoulder heating, while the blunted wedge produced a local heating feature near the transition from the blunt tip to the sloped body. These local peaks aligned with wall shear stress behavior, indicating that the heating features were connected to local flow acceleration, boundary layer thinning, and near-wall gradient development. This means that simplified stagnation-only heat flux estimates may miss important localized thermal design drivers. Local geometry transitions must be evaluated carefully in any reusable TPS design.

The external CFD-to-thermal coupling approach was effective for the scope of this study, but the isothermal wall assumption remains a major limitation. The method used in this study was practical and computationally efficient, but it did not fully capture the spatial variation of wall temperature along the body. This may have inflated some of the localized heat flux peaks, especially in the shoulder and transition regions. The results should therefore be interpreted as conservative for localized hot spots, particularly where the imposed wall temperature was lower than the true local radiative equilibrium temperature would likely be. This does not invalidate the results, but it does mean that the most severe local heat flux values should be treated as bounding values rather than final flight ready predictions. The same external thermal coupling approach but with added spatial discretion along the wall was proposed as a potential solution to this problem. However, this method should still not be interpreted as a full replacement for higher fidelity CFD-to-thermal coupling within the ANSYS Suite with wall effects and energy balance modeled directly in the Fluent solver.

Within the feasible thermal cases, the SMA interlayer can produce useful recovery stress if it is placed at the correct depth in the TPS stack. The SMA must be buried far enough below the surface to avoid exceeding its functional temperature range, but not so deep that it remains below its transformation window. The successful cases show that this is possible. In the lower heating case, the SMA layer reached approximately 448 K. In the moderate case, it reached approximately 587 K. In the higher successful case, it reached approximately 686 K. These temperatures allowed the SMA model to evaluate partial and full transformation behavior depending on material selection. This supports the feasibility of the adaptive layer concept.

The mechanical contribution of the SMA layer depends strongly on how effectively the recovery stress is transferred into the surrounding stack. The raw SMA response alone is not enough to determine the usefulness of the interlayer. The effective stack stress depends on the amount of SMA material present in the mesh, the degree of mechanical coupling to the surrounding layers, the assumed prestrain, and the temperature-dependent transformation behavior. The eta-transfer sweep showed that a more effective load path could significantly increase transferred stress, while weak bonding or poor mechanical coupling could greatly reduce the benefit. This reinforces that future ATPS work must treat the SMA layer as an interface design problem, not just a material selection problem.

Taken together, the proposed ATPS concept is best understood as a multifunctional supplement to reusable TPS, not as a complete replacement for existing systems. The ceramic

outer layers still perform the dominant heat shielding function. The SMA layer provides an added adaptive mechanical function by generating recovery stress within a useful temperature range. This layered role division is consistent with the literature on integrated and multifunctional TPS architectures. The study therefore supports the concept of an SMA-augmented reusable TPS stack, but it does not support the idea that SMAs alone can serve as re-entry thermal protection materials.

8.4 Recommendations

Based on the thermal limits identified in this study, future designs should treat the SMA-based ATPS as an acreage or moderate heating TPS technology. The architecture should be targeted toward blunt forebody acreage regions, capsule heat shield regions away from severe local peaks, and downstream surfaces where the heating has relaxed from a leading edge or stagnation point. It should not be used as the primary solution for the most severe leading edge stagnation point, extreme off-stagnation transition peak, or highly concentrated local heating region without additional protection.

Material selection should begin with the outer ceramic system before the SMA system. The results show that the primary feasibility limit is the surface temperature capability of the coating and ceramic layer. If the outer surface exceeds its allowable temperature, the case becomes invalid regardless of whether the SMA layer reaches a useful transformation temperature. Therefore, material selection should begin by verifying that the coating and ceramic tile can survive the wall temperature with margin. Only after the outer system is thermally viable should the SMA transformation range be selected.

For the adaptive layer, material selection should be based on the embedded interlayer temperature rather than the external heat flux or outer surface temperature. The SMA is not directly exposed to the re-entry flow, so its design temperature is the temperature predicted at the SMA layer inside the stack. For lower temperature cases, Cu-Al-Ni or high-temperature binary NiTi may align with the required transformation range, although Cu-Al-Ni should be treated carefully because of brittleness and cyclic durability concerns. For the higher successful cases, NiTiHf is the most appropriate option because it provides a higher transformation temperature range and is more relevant to elevated-temperature aerospace applications.

For future high temperature ATPS development, NiTiHf should be prioritized as the primary SMA candidate. The 0.62 MW/m2 and 1.2 MW/m2 cases place the SMA layer in a temperature range where conventional binary NiTi is not ideal. NiTiHf provides a more realistic high-temperature SMA option and is better aligned with the thermal results. It also provides a stronger basis for future aerospace-relevant development than Cu-Al-Ni because of its greater relevance to high-temperature NiTi-based actuator systems.

To improve the CFD-to-thermal coupling when using the external calibration method developed in this study, segmented or spatially discretized wall temperature calibration should be implemented in future CFD runs. The body wall could be divided into multiple wall zones, such

as stagnation, shoulder, acreage, transition, and aftbody regions for the blunt forebody, and leading edge, nose-to-wedge transition, acreage panel, and truncation regions for the wedge. Each zone could then be calibrated separately using its own heat transfer rate and MATLAB-predicted wall temperature. This would reduce the limitations of a single isothermal wall and provide a better approximation of the spatially varying radiative equilibrium temperature.

9.0 Future Work

The present study provides a first-order computational evaluation of an SMA-based Adaptive Thermal Protection System, but several areas require additional development before the concept could be considered mature. The most important future work is the implementation of full conjugate heat transfer. In the current study, the solid TPS stack was not modeled directly inside Fluent. Instead, the wall temperature was imposed in Fluent and calibrated through the MATLAB thermal model. This approach was necessary for computational efficiency and compatibility with the ANSYS Student license, but it limited the ability of the wall temperature to vary spatially. A full conjugate heat transfer model would allow Fluent to solve the external hypersonic flow and the internal TPS conduction problem together. This would provide a more accurate surface temperature distribution, reduce uncertainty in local heat flux peaks, and improve the physical coupling between the flowfield and the wall response.

A related future improvement is the implementation of segmented wall temperature calibration. If full conjugate heat transfer is not computationally practical, the wall can still be divided into multiple regions and calibrated independently. For the blunt forebody, these regions could include the stagnation zone, shoulder transition zone, acreage region, and aftbody region. For the wedge, these regions could include the leading edge, blunt tip to sloped body transition region, upper and lower acreage panels, and aft truncation region. This would provide a more realistic approximation to a spatially varying wall temperature while remaining less expensive than full conjugate modeling.

Future work should also include transient trajectory-based heating. The present study used steady CFD-derived heating values and applied them to representative thermal cases. A real re-entry vehicle experiences a time-varying heat flux as altitude, density, velocity, Mach number, and angle of attack change throughout the trajectory. A transient thermal model using a full heat flux time history would allow the study to evaluate thermal soak, time-to-peak backface temperature, post-peak cooling behavior, and repeated thermal cycling. This would be especially important for reusable TPS because material damage is often driven by thermal cycling and accumulated exposure rather than a single steady-state peak.

The CFD model should also be extended to higher fidelity gas physics. The current model used a one-temperature finite rate chemistry approach with the air-5species-park93 mixture. This was appropriate for the scope of the present study and the software limitations, but future work should evaluate a two-temperature thermochemical nonequilibrium model when available. A two-temperature model would allow vibrational-electronic energy modes to relax separately from translational-rotational temperature, improving the prediction of high-enthalpy shock layer behavior. If the post-shock temperatures exceed the threshold where ionization becomes relevant, an 11-species air model should also be evaluated.

Additional geometric fidelity is also recommended. The two geometries used in this study were intentionally simplified to isolate representative heating environments and maintain computational feasibility. Future models should evaluate more realistic capsule forebodies,

shoulder radii, leading edge radii, flap geometries, panel gaps, tile steps, and local protrusions. Since the present study showed that geometric transition regions can produce strong local heating peaks, future work should place special focus on the local geometry features that often drive TPS failure. Three-dimensional simulations would also be valuable because real vehicles include crossflow, local curvature, asymmetric heating, and finite-span effects that cannot be captured in a two-dimensional or axisymmetric model.

The thermal model should be expanded from a steady one-dimensional reduced-order approach into a transient multilayer model with temperature-dependent properties. Many TPS materials have thermal conductivity, heat capacity, emissivity, and mechanical properties that vary strongly with temperature. Including these effects would improve prediction accuracy, especially in cases where the surface temperature approaches the upper capability of the ceramic material. Future work should also account for contact resistance, bondline degradation, coating oxidation, and radiation exchange between internal gaps if present.

The SMA recovery stress model should be expanded into a more complete thermomechanical simulation. The present model used a reduced-order formulation to estimate transferred recovery stress as a function of temperature, prestrain, mesh geometry, and load transfer efficiency. This provided useful screening information, but it did not directly model the detailed mechanics of each SMA wire or its interaction with the surrounding ceramic and adhesive layers. Future work should model the SMA mesh explicitly or develop a homogenized finite element material model that captures temperature-dependent modulus, transformation strain, hysteresis, and constrained recovery stress.

A major future task is direct delamination and crack-growth modeling. The purpose of the SMA layer is to improve interfacial integrity, but the current study does not directly prove crack suppression. A future model should include cohesive-zone elements or fracture mechanics based interfaces between the ceramic tile, SMA layer, buffer, SIP, adhesive, and vehicle skin. This would allow the study to calculate whether SMA recovery stress actually reduces crack opening displacement, interfacial shear, or delamination growth under thermal gradients. This would be the most direct way to validate the structural benefit of the adaptive interlayer.

Experimental validation should also be performed. A coupon-level experiment could be built using a representative ceramic tile, high-emissivity coating, embedded SMA mesh, buffer layer, SIP, adhesive, and metallic backplate. The coupon could be exposed to controlled radiative or plasma heating while thermocouples measure the outer surface, SMA layer, and backface temperatures. Strain gauges or embedded fiber optic sensors could be used to measure the mechanical response during heating and cooling. Post-test inspection could then be used to evaluate cracking, debonding, adhesive degradation, and permanent deformation. Such testing would provide the data needed to validate both the thermal model and the SMA recovery stress model.

Future work should also include repeated-use cycling. Since one of the primary motivations for ATPS is reusability, it is not enough to show that a stack can survive one heating event. The stack must be evaluated over multiple heating and cooling cycles. Repeated cycling

can change ceramic microstructure, degrade bondlines, alter SMA transformation behavior, reduce recovery stress, and introduce accumulated damage at interfaces. Testing and modeling should therefore include cyclic thermal exposure representative of multiple re-entry events.

Finally, future work should explore optimization of the full ATPS stack. The current study selected material stacks based on temperature margin, material maturity, cost, and practical operating ranges. A formal optimization could vary coating emissivity, ceramic tile material, ceramic thickness, SMA placement, SMA alloy, mesh density, buffer material, SIP thickness, adhesive type, and vehicle skin material. The objective should be to minimize mass and thickness while maintaining surface survivability, SMA activation, backface temperature margin, and interfacial structural integrity. This type of optimization would help move the concept from a feasibility study toward a design tool for future reusable hypersonic vehicles.

References

[1] Anderson, J. D., *Hypersonic and High-Temperature Gas Dynamics*, 2nd ed., AIAA Education Series, AIAA, Reston, VA, 2006.

[2] Sutton, K., and Graves, R. A., "A General Stagnation-Point Convective-Heating Equation for Arbitrary Gas Mixtures," NASA TR R-376, Nov. 1971.

[3] Gnoffo, P. A., Weilmuenster, K. J., Hamilton, H. H., II, Olynick, D. R., and Venkatapathy, E., "Computational Aerothermodynamic Design Issues for Hypersonic Vehicles," *Journal of Spacecraft and Rockets*, Vol. 36, No. 1, 1999, pp. 21–43.
doi:10.2514/2.3430

[4] Uyanna, O., and Najafi, H., "Thermal Protection Systems for Space Vehicles: A Review on Technology Development, Current Challenges, and Future Prospects," *Acta Astronautica*, Vol. 176, Nov. 2020, pp. 341–356.
doi:10.1016/j.actaastro.2020.06.047

[5] Marley, C. D., and Driscoll, J. F., "Modeling an Active and Passive Thermal Protection System for a Hypersonic Vehicle," AIAA Paper 2017-0118, Jan. 2017.
doi:10.2514/6.2017-0118

[6] Le, V. T., Ha, N. S., and Goo, N. S., "Advanced Sandwich Structures for Thermal Protection Systems in Hypersonic Vehicles: A Review," *Composite Structures*, Vol. 262, April 2021, Paper 113660.
doi:10.1016/j.compstruct.2021.113660

[7] Greene, E. E., "Thermal Protection and Control," NASA KSC-2013-201, Kennedy Space Center, FL, July 2013.

[8] Wright, M. J., Hollis, B. R., and Venkatapathy, E., "Aerothermal Environments for Mars Entry Vehicles," NASA Ames Research Center, Moffett Field, CA, 2002.

[9] Piacquadio, S., Pridöhl, D., Henkel, N., Bergström, R., Zamprotta, A., Dafnis, A., and Schröder, K.-U., "Comprehensive Comparison of Different Integrated Thermal Protection Systems with Ablative Materials for Load-Bearing Components of Reusable Launch Vehicles," *Aerospace*, Vol. 10, No. 3, March 2023, Paper 319.
doi:10.3390/aerospace10030319

[10] Mitcheltree, R. A., and Braun, R. D., "Aerothermal Heating Predictions and Flight Results for the Galileo Probe," NASA TM-110440, Jan. 1996.

[11] Zhang, J., Luo, Z., Wang, X., and Xiong, J., “Mechanical Stability of Hybrid Corrugated Sandwich Plates Under Fluid-Structure-Thermal Coupling for Novel Thermal Protection Systems,” *International Journal of Mechanical Sciences*, Vol. 217, March 2022, Paper 107029. doi:10.1016/j.ijmecsci.2021.107029

[12] Zhang, J., Nie, C., Cai, J., and Pan, S., “An Implicit Coupling Framework for Numerical Simulations Between Hypersonic Nonequilibrium Flows and Thermal Responses of Charring Materials in the Presence of Ablation,” *Aerospace Science and Technology*, Vol. 158, March 2025, Paper 109915. doi:10.1016/j.ast.2025.109915

[13] Vargas Venegas, C. A., Huang, D., Blonigan, P., and Tencer, J., “Physics-Infused Reduced-Order Modeling for Analysis of Multi-Layered Hypersonic Thermal Protection Systems,” Sandia National Laboratories, SAND2025-06421O, Albuquerque, NM, 2025.

[14] Chen, Q., Kalpoe, T., and Jovanova, J., “Design of Mechanically Intelligent Structures: Review of Modelling Stimuli-Responsive Materials for Adaptive Structures,” *Heliyon*, Vol. 10, 2024, Paper e34026. doi:10.1016/j.heliyon.2024.e34026

[15] Gu, L., “Advancements and Challenges of Smart Materials in Aerospace: Applications, Mechanisms, and Future Prospects,” *Proceedings of the 3rd International Conference on Mechatronics and Smart Systems*, 2025. doi:10.54254/2755-2721/147/2025.22163

[16] Farber, E., Zhu, J.-N., Popovich, A., and Popovich, V., “A Review of NiTi Shape Memory Alloy as a Smart Material Produced by Additive Manufacturing,” *Materials Today: Proceedings*, Vol. 30, 2020, pp. 761–767. doi:10.1016/j.matpr.2020.01.563

[17] Zhang, Q., “Shape-Memory-Alloy Hybrid Composites: Modeling, Dynamic Analysis, and Optimal Design,” Ph.D. Dissertation, School of Mechanical Engineering, Purdue University, West Lafayette, IN, 2024.

[18] Mehra, M., “Smart Materials and Their Applications in Aerospace Engineering,” *IOSR Journal of Mechanical and Civil Engineering*, Vol. 21, No. 6, Nov.–Dec. 2024, pp. 24–26. doi:10.9790/1684-2106012426

[19] Yang, F., Xu, X., Liu, H., and Feng, L., “Structure Design and Heat Transfer Performance Analysis of a Novel Composite Phase Change Active Cooling Channel Wall for Hypersonic Aircraft,” *Applied Thermal Engineering*, Vol. 213, Aug. 2022, Paper 118880. doi:10.1016/j.applthermaleng.2022.118880

[20] Poloni, E., Bouville, F., Schmid, A. L., Pelissari, P. I. B. G. B., Pandolfelli, V. C., Sousa, M. L. C., Tervoort, E., Christidis, G., Shklover, V., Leuthold, J., and Studart, A. R., "Carbon Ablators with Porosity Tailored for Aerospace Thermal Protection During Atmospheric Re-entry," *Carbon*, Vol. 195, Aug. 2022, pp. 80–91.
doi:10.1016/j.carbon.2022.03.062

[21] Williams, S. D., and Curry, D. M., "Thermal Protection Materials: Thermophysical Property Data," NASA RP-1289, Lyndon B. Johnson Space Center, Houston, TX, Dec. 1992.

[22] Johnson, S. M., "Coatings and Surface Treatments for Reusable Entry Systems," NASA ARC-E-DAA-TN30202, Ames Research Center, Moffett Field, CA, March 2016.

[23] Stewart, D. A., Leiser, D. B., Smith, M., and Goldstein, H. E., "Lightweight Thermal Protection System for Atmospheric Entry," NASA Ames Research Center, Moffett Field, CA, 2007.

[24] Fahrenholtz, W. G., Hilmas, G. E., Talmy, I. G., and Zaykoski, J. A., "Refractory Diborides of Zirconium and Hafnium," *Journal of the American Ceramic Society*, Vol. 90, No. 5, May 2007, pp. 1347–1364.
doi:10.1111/j.1551-2916.2007.01583.x

[25] Clarke, D. R., Oechsner, M., and Padture, N. P., "Thermal-Barrier Coatings for More Efficient Gas-Turbine Engines," *MRS Bulletin*, Vol. 37, No. 10, Oct. 2012, pp. 891–898.
doi:10.1557/mrs.2012.232

[26] ZIRCAR Ceramics, Inc., "Alumina-Silica Type ECO-27," ZIRCAR Ceramics, Inc., Florida, NY, Product Data Sheet.

[27] Morgan Advanced Materials, "Inorganic Boards," Morgan Thermal Ceramics, Windsor, United Kingdom, Technical Data Sheet, 2024.

[28] Final Advanced Materials, "Zirconia Fibre Felts," Final Advanced Materials, Dardilly, France, Product Data Sheet.

[29] Otsuka, K., and Ren, X., "Physical Metallurgy of Ti-Ni-Based Shape Memory Alloys," *Progress in Materials Science*, Vol. 50, No. 5, July 2005, pp. 511–678.
doi:10.1016/j.pmatsci.2004.10.001

[30] Hautcoeur, A., Fouché, F., and Sicre, J., "Cu-Al-Ni Shape Memory Single Crystal Wires with High Transformation Temperature," *Proceedings of the 43rd Aerospace Mechanisms Symposium*, NASA Ames Research Center, Moffett Field, CA, May 2016, pp. 185–192.

[31] Elahinia, M., Shayesteh Moghaddam, N., Amerinatanzi, A., Saedi, S., Toker, G. P., Karaca, H., Bigelow, G. S., and Benafan, O., "Additive Manufacturing of NiTiHf High Temperature Shape Memory Alloy," *Scripta Materialia*, Vol. 145, March 2018, pp. 90–94. doi:10.1016/j.scriptamat.2017.10.016

[32] Carl, M., Smith, J., Van Doren, B., and Young, M., "Effect of Ni-Content on the Transformation Temperatures in NiTi-20 at.% Zr High Temperature Shape Memory Alloys," *Metals*, Vol. 7, No. 11, Nov. 2017, Paper 511. doi:10.3390/met7110511

[33] Noebe, R. D., Padula, S. A., and Garg, A., "Properties and Potential of Two (Ni,Pt)Ti Alloys for Use as High-Temperature Shape Memory Alloys," NASA Glenn Research Center, Cleveland, OH, NASA Technical Memorandum, 2005.

[34] Smithsonian National Air and Space Museum, "Shuttle Strain Isolation Pad, STS-1," Smithsonian Institution, Washington, DC, Collection Record.

[35] Momentive Performance Materials, "RTV100 Series Silicone Rubber Adhesive Sealants," Momentive Performance Materials, Waterford, NY, Technical Data Sheet.

[36] Aremco Products, Inc., "High Temperature Solutions," Aremco Products, Inc., Valley Cottage, NY, Product Catalog, 2022.

[37] Final Advanced Materials, "Zirconia Ceramic Adhesives," Final Advanced Materials, Dardilly, France, Product Data Sheet.

[38] Applied Ceramics, Inc., "High Temperature Sealants and Adhesives," Applied Ceramics, Inc., Doraville, GA, Catalog Section 11.

[39] ASM International, "Aluminum 2024-T3," *MatWeb Material Property Data*, ASM International, Materials Park, OH.

[40] Moser, T. L., and Schneider, W. C., "Strength Integrity of the Space Shuttle Orbiter Tiles," AIAA Paper 81-2469, Nov. 1981.

[41] ASM International, "Aluminum 7075-T6," *MatWeb Material Property Data*, ASM International, Materials Park, OH.

[42] ASM International, "Titanium Ti-6Al-4V, Grade 5, Annealed," *MatWeb Material Property Data*, ASM International, Materials Park, OH.

[43] Special Metals Corporation, "INCONEL Alloy 625," Special Metals Corporation, Huntington, WV, Technical Bulletin.

[44] Kiser, J. D., Bhatt, R. T., and DiCarlo, J. A., "SiC/SiC Ceramic Matrix Composites Developed for High-Temperature Applications," NASA Glenn Research Center, Cleveland, OH, NASA Technical Report, 2005.

[45] National Aeronautics and Space Administration, "Structures and Materials: Space Shuttle Tiles," *Museum in a Box Series*, Aeronautics Research Mission Directorate, NASA, Washington, DC, 2023.

[46] Bartels, R. E., Rumsey, C. L., and Biedron, R. T., "CFL3D Version 6.4—General Usage and Aeroelastic Analysis," NASA TM-2006-214301, March 2006.

[47] Pezzella, G., "Aerodynamic and Aerothermodynamic Analysis of Space Mission Vehicles: A Suitable Investigation Approach at Phase-A Design Level," *Hypersonic Vehicles*, IntechOpen, London, 2017.
doi:10.5772/intechopen.69598

[48] International Organization for Standardization, "Standard Atmosphere," ISO 2533:1975, International Organization for Standardization, Geneva, Switzerland, 1975.

[49] ANSYS, Inc., "Modeling Hypersonic Flows," *ANSYS Fluent User's Guide*, Release 2025 R2, ANSYS, Inc., Canonsburg, PA, 2025, Chap. 18.

[50] Mazaheri, A., "High-Energy Atmospheric Re-entry Test Aerothermodynamic Analysis," *Journal of Spacecraft and Rockets*, Vol. 50, No. 2, March–April 2013, pp. 270–281.
doi:10.2514/1.A32241

[51] Alunni, A. I., and Gökçen, T., "Assessment of Fencing on the Orion Heatshield," AIAA Paper 2016-3128, June 2016.
doi:10.2514/6.2016-3128

[52] Glass, D. E., "Heat-Pipe-Cooled Leading Edges for Hypersonic Vehicles," Workshop on Materials and Structures for Hypersonic Flight, Santa Barbara, CA, July 2006.