Aero-thermo-elastic Analysis and Optimization of Scramjet Inlets

by

J. A. Horing

B.S., Harvard University, 2018

M.S., University of Colorado Boulder, 2023

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Aerospace Engineering

2025

Committee Members:

Iain D. Boyd, Chair

Kurt K. Maute

Brian Argrow

Peter Hamlington

Kevin Bowcutt

Horing, J. A. (Ph.D., Aerospace Engineering)

Aero-thermo-elastic Analysis and Optimization of Scramjet Inlets

Thesis directed by Prof. Iain D. Boyd

ii. Abstract

Although supersonic combustion ramjets — scramjets — provide a fuel-efficient method for propulsion at hypersonic speeds, current challenges with the engine prohibit the robustness necessary for space accessibility and trans-atmospheric flight. One such challenge the engine faces is that the inlet is a highly elongated compression surface that is prone to deformations due to the combined high thermal and structural loads. This dissertation begins by developing a high-fidelity aero-thermo-elastic model to quantify the impacts of the deformation on the inlet behavior; the model subsequently couples to a low-fidelity engine model to understand the ramifications on the downstream components and the engine as a whole. The system is found to be extremely sensitive to the changes in deformation, leading to increased flow separation and heating and to deviations of the engine performance and efficiency from the original design point. Additionally, the deformations impact the vehicle's aerodynamic performance due to the integrated airframe/inlet design. Therefore, the design of the inlet must include the implications of these deformations and to do so, an aero-thermo-elastic sensitivity analysis and optimization process are developed for this work. By leveraging high-fidelity solvers, a more accurate assessment of the fluid-thermal-structural interactions is possible; however, there is a high computational cost associated with these analyses, especially when optimizations are considered. The current implementation reduces the cost by leveraging a semi-analytical aero-thermo-elastic sensitivity analysis, which combines a direct method approach for the fluid dynamics and a finite-difference approach for the thermo-structural response. With the computationally efficient aerothermodynamic and aero-thermo-elastic sensitivity analyses in place, hypersonic optimizations are performed to redesign a re-entry vehicle and a scramjet inlet for better performance. Together, these chapters demonstrate the development and application of advanced analysis and optimization techniques to improve the performance of hypersonic vehicles, offering insights into the coupling of aerothermodynamics, thermal response, and structural deformation.

"It's the questions we can't answer that teach us the most. They teach us how to think. If you give a man an answer, all he gains is a little fact. But give him a question and he'll look for his own answers."

Patrick Ruthfuss, *The Wise Man's Fear*

Acknowledgements

I would like to begin by thanking my advisor, Prof. Iain Boyd, for his guidance and mentorship throughout my time in graduate school. The support and insights he provided were invaluable to the ideation and execution of my dissertation. I'd also like to thank my entire committee for their encouragements and discussions that helped drive the direction of my work. I would especially like to thank my unofficial co-advisor, Prof. Kurt Maute, who's help (along with many of his students) over the years has been crucial to pursuing my research. I was also supported by Smead Program and would like to thank Ann Smead, Michael Byram, Chris Muldrow and Carrie Ann Ngai for making the program feel like a family. Additionally, I'm so thankful to my research group for answering unending questions, the Smead Scholars for hanging out over limitless amounts of food and the FSM cohort making graduate school such a fun time! I'd also like to thank my friends outside of the office, who provided endless excitement and support of my work. Finally, I am incredibly grateful for my supportive family and partner, who had no choice but to listen to every step of the process. I couldn't have done it without all of you.

This work has been supported by a National Defense Science and Engineering Graduate fellowship funded by Air Force Office of Scientific Research. Additional support was provided by the Naval Air Warfare Center Aircraft Division under Contract N004212310001 and most recently by NASA under Grant 80NSSC24K1739.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction

1.1 Background

In 1947, Charles "Chuck" Yeager piloted a bright orange, rocket-powered Bell X-1 aircraft to break the sound barrier and achieve supersonic flight. In the years since, technology has continued to break through that barrier and beyond, all the way to hypersonic flight, or velocities over five times the speed of sound. Less than 15 years later, the X-15 aircraft reached altitudes (354,200 feet) and speeds (Mach 6.7) never before experienced by a manned aircraft [115]. As this high-speed vehicle rammed through the air, shock waves appeared, which converted the kinetic energy to thermal energy and caused intense temperatures of the air surrounding the vehicle. At these high temperatures, energy is exchanged through several mechanisms in the fluid, including excitation of internal energies, dissociation of molecules and potentially ionization of the gas. These phenomena that occur in the hypersonic regime pose challenges for all aspects of the vehicle, including the aerodynamics, high temperature materials, propulsion systems, structures, and guidance and control, which is why manned, powered hypersonic flight still remains a challenge almost 80 years after Chuck Yeager's first supersonic flight [117].

Despite these challenges, there are other limited classes of vehicles that are able to achieve hypersonic flight today, such as re-entry vehicles, missiles and research aircraft. However, some classes of vehicles are just beginning to emerge, such as cruise missiles, hypersonic aircraft and space planes, all of which require a high-speed airbreathing propulsion system [117]. A large portion of the progress for these propulsion systems, and much of hypersonics in general, can

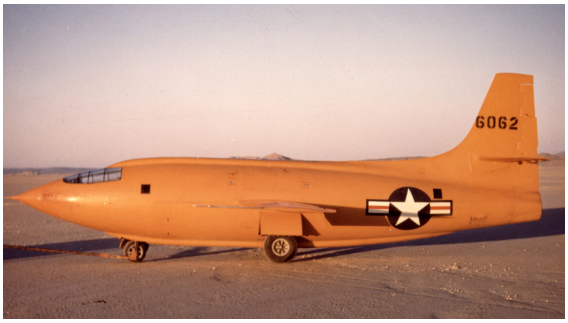

(a) Bell X-1.

(b) North-American X-15.

Figure 1.1: First supersonic and hypersonic manned aircraft.

be traced back to the National AeroSpace Plane (NASP) program, which had the aspirational goal of reaching space using a single-stage-to-orbit (SSTO) vehicle. Although the lofty goal was never achieved, the program advanced research in computational fluid dynamics, material science, airbreathing propulsion and many more areas, paving the way for the next generation of high-speed airbreathing vehicles [122].

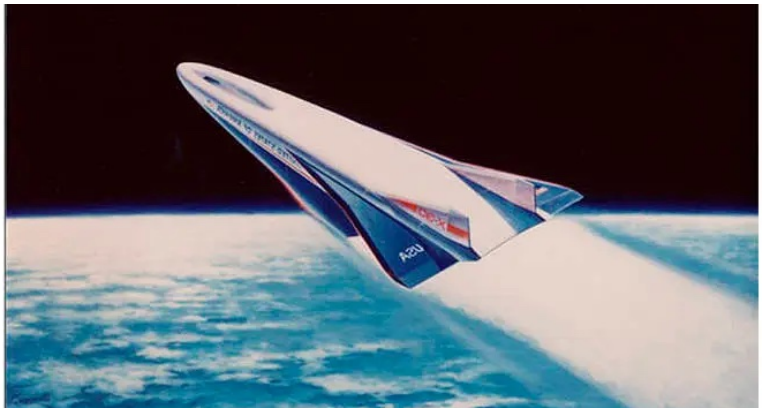

Figure 1.2: Rendering of the Rockwell X-30 technology demonstrator created for the NASP program.

The purpose of the demonstrator was a proof of concept that airbreathing propulsion can surpass the efficiency of rocket propulsion systems and allow for a single vehicle to reach space, rather than a multistaged system, as seen in rockets today. By using the air around the vehicle as an oxidizer, hence the name "airbreathing" this propulsion system has an advantage over rockets,

which must carry the oxidizer onboard. The drawback for airbreathing systems, though, is that there is less density higher in the atmosphere, so in order to capture enough oxygen for combustion, the vehicle needs to fly at hypersonic speeds. At these altitudes and speeds, using a typical turbojet is no longer possible, since they rely on rotary compressors. Instead, a different type of engine was developed for the NASP program: a supersonic combustion ramjet, or scramjet for short.

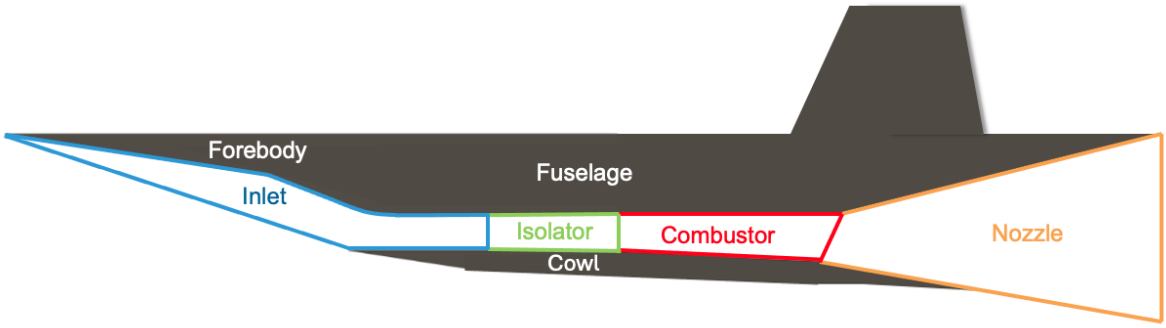

Figure 1.3: Scramjet components.

A scramjet achieves compression of the air by use of forebody and inlet shock waves instead of a mechanical compressor, as shown in Figure 1.3. The inlet is highly integrated into the forebody of the vehicle, so the vehicle's outer-mold-line (OML) design must take into account the engine design. The combustion process then occurs within supersonic flow, hence the origin of the name. The goal of the scramjet inlet is to compress the gas to a high enough pressure that it is able to combust fuel. A shock wave appears when there is a sudden disturbance in supersonic flow, such as the leading edge of a vehicle. A normal shock is a shock wave that is perpendicular to the direction of the flow, which will increase the pressure of the air, but also increase the temperature. That increase of the temperature and entropy of the gas causes inefficiencies in the combustion chamber and nozzle of the engine. Therefore, scramjets are designed instead to use a series of oblique shock waves, which are at an acute angle from the flow direction and do not increase the temperatures and pressures as greatly as a normal shock. This allows for the same increase in pressure across a train of shock waves, but fewer losses in entropy and efficiency. Also, although these engines must ram through

the air to capture enough oxidizer for the engine to run properly, they cannot experience too high of forces that surpass the structural limits of the vehicle. Therefore, these scramjet-powered vehicles are limited to a narrow band of flight envelopes, which is shown in Figure 1.4.

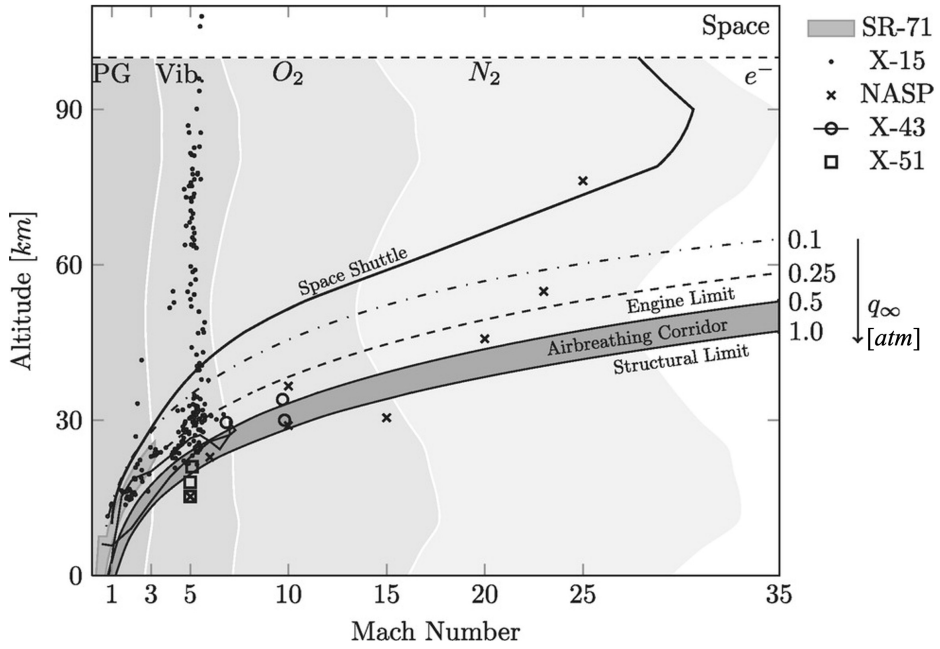

Figure 1.4: Flight envelopes of existing and potential hypersonic vehicles [74].

During the NASP program, both computational tools and ground tests were used to build confidence that these types of engines could potentially facilitate an efficient and robust method for accessing space. Years later, the first-generation development and testing of these engines has shown great success in previous and current flight campaigns such as X-43A[34], X-51 [45], HIFiRE [60], and HAWC. However, additional improvements are necessary for scramjet engines to be widespread and incorporated into vehicles with a larger variety of flight conditions.

1.2 Motivation

There are many obstacles currently in the way of advancing first-generation developments of scramjet engines into more versatile and robust systems. Focusing on the engine inlet, there are three main challenges: (1) understanding and controlling the complicated flow structures (e.g.

shock-boundary-layer interactions, separation, turbulence and transition), (2) integrating the inlet into the vehicle forebody and downstream engine components (e.g., burner and nozzle), and (3) addressing the highly coupled nature of the flowfield and the vehicle/engine. Each of these challenges are difficult and must be addressed in part through analysis since there are limited ground and flight test opportunities.

The first challenge creates a need for high-fidelity modeling of the flowfield using computational fluid dynamics (CFD), which is often a computationally expensive endeavor. Although there are low-cost, low-fidelity models, their accuracy can be limiting; therefore, more efficient methods of exploring the design space with high-fidelity CFD are necessary. The second challenge creates a need for assessing the inlet's impact on system-level (e.g., vehicle and engine) quantities of interest as well as downstream component-level quantities of interest. The analysis of the inlet flowfield alone is already a computationally expensive task, so modeling the downstream components and vehicle together requires a multi-fidelity approach.

The final challenge mentioned above is the highly coupled physics, which occurs due to the flowfield surrounding the vehicle imparting thermal and mechanical loads onto the surface that will heat and deform the vehicle. As the vehicle distorts, the flowfield within the scramjet inlet changes; this could have several implications on the inlet behavior, potentially causing stronger shocks, greater separation and varied shock impingement locations [11]. The deformation can also lead to modified total pressure recovery (TPR), which is a measure of efficiency of the inlet based on the change in total pressure from freestream to isolator, and can potentially cause unstart, which is an unstable choking of the engine [51]. Such changes to the body and flowfield introduce a need for a coupled approach: aero-thermo-elastic analysis. Due to the integrated nature of the vehicle and engine – challenge (2) – the fluid-thermal-structural interactions (FTSI) of the inlet impacts the subsequent engine components, the total engine performance and the aerodynamic behavior of the vehicle. Therefore, accounting for these deformations when designing and analyzing is critical to finding an optimal configuration.

With those challenges in mind, the motivation behind this dissertation is to develop a high-

fidelity aero-thermo-elastic computational model that can be used to explore the complicated design space of an inlet to enhance the performance and efficiency of the scramjet (addressing challenges 1 and 3). This high-fidelity analysis is then coupled to a low-fidelity engine model to understand the impacts to the engine and its components (addressing challenge 2). For hypersonic vehicles, there is a strong interdependency between the engine and airframe, which often produces conflicting requirements that must be addressed [15]. Especially when deformation effects are considered, the design of the vehicle – including the shape, materials, structure, thermal management, etc. – requires a multidisciplinary analysis and optimization (MDAO) technique to account for the realistic conditions occurring in flight. Therefore, this dissertation not only develops tools to assess the impacts of FTSI on an inlet/engine system but also develops an efficient sensitivity analysis to enable design exploration through MDAO.

1.3 Review of Related Work

1.3.1 Aero-thermo-elastic Analysis

Aero-thermo-elasticity has been a challenge for hypersonic flight in the past and has, therefore, been previously studied in depth since the late 1950s and 1960s [86]. During the X-15 program, examples of these fluid-thermal-structural interactions were shown to cause failures to the wings and control surfaces of the vehicle [38]. Because of the difficulties capturing this behavior in ground tests, there has been a significant amount of effort to model these interactions. Work by McNamara et al. [86] in aero-thermo-elasticity concluded that interactions between the three physics – aerothermodynamics, thermal response, and structural response – have varying degrees of influence. Some of the "weaker" couplings were chosen to be neglected to simplify the analysis in these initial works; the coupling often singled out as "weak" was the interaction between the elastic response and the thermal response and was chosen to be modeled using one-way coupling. The approximation of the coupling as one-way relies on the assumption that "static aeroelastic coupling (static elastic deflections due to steady-state pressure and thermal loading) is insufficient to alter the tem-

perature distribution from the reference condition" [25]. Later, Culler et al. worked to "identify key characteristics that determine the level of coupling needed for different situations" specific to skin panels, including Mach number, materials, boundary conditions, and other characteristics [25, 27, 26]. Often times, a *two-way* coupling approach was necessary between the aerothermal and aeroelastic response when the deformation sufficiently influenced the internal temperature distribution. These works highlighted the necessity of tight coupling for these high-fidelity models in order to accurately predict the aero-thermo-elastic behavior of these geometries.

The focus for aero-thermo-elastic analytical research in hypersonics has often been on simpler geometries such as panels [116] and leading edges [31]. The majority of the experimental work has also focused on simpler geometries such as panels [110, 94]. Although these results are informative, they are also limited in their applicability to more complicated geometries such as compression surfaces and inlets. The combination of panels and leading edges that make up an inlet leads to complicated flowfield physics such as shock-shock interactions, shock-boundary layer interactions and separation, all of which can vary with deformation and impact the heating to the vehicle. There is a current push to explore more complicated geometries both experimentally and analytically [10, 118], including compression surfaces similar to inlets.

Applying this coupled FTSI analysis approach to inlets has only been prevalent for the last five years or so; this area is still in a nascent period of research [10, 67, 74, 127, 77, 92, 112]. Prior work has established that there is a need for coupled analysis of inlet and engine performance (deformations cannot be ignored), but the implications and mitigations for these deformations still need to be explored further. For example, Bhattrai et al. [11] detailed the changes in the shock structures internal to the inlet when elastic deformation is included, resulting in stronger shock-shock and shock-boundary layer interactions. Similarly, Guangyue et al. [42] and Ye et al. [127] also explored the influence of FTSI on the internal flowfield, concluding that deformation can cause changes to the shock impingement location and larger separation bubbles, respectively. Ye et al. then extended this understanding of the flowfield behavior to quantities of interest for the inlet, such as the inlet compression ratio and total pressure recovery coefficient, which deviated by -8.8%

and 10.5%, respectively, when deformation was considered.

Since the inlet compression is the first step in the engine cycle, these deviations have an impact on the subsequent steps and components. In addition, the forebody of the vehicle has a dual-purpose: it serves as an inlet to direct and compress the airflow and a lifting-surface to provide the necessary aerodynamic performance of the vehicle. The integrated nature of the engine into the vehicle is an additional challenge for scramjets, unlike their subsonic counterparts (turbojets), the majority of which can be designed separately from the aircraft [16]. Previous research by Kline et al. [67] and Riley et al. [92] has begun to incorporate the engine-level and vehicle-level effects including outputs such as lift, pitching moment, stream-thrust, and mass capture. Both studies included simplifications to the aero-thermo-elastic analysis, only using loose coupling, in order to explore more realistic vehicles/engines and flight conditions. Kline et al. concluded that the temperature of the structure had a substantial impact on the deformation and the thrust, where a 20 K increase in average temperatures resulted in a loss of 1% thrust for the 3D inlet studied. Meanwhile, Riley et al. found that certain metallic material choices could result in a 60% decrease in thrust coefficient while high-temperature composites could reduce that change to under 5%. These studies indicate that the inlet and engine system are sensitive to the material and thermal choices, which will be further studied in this work as potential avenues to mitigate deformation.

To date, there has only been one other study on mitigating the impact of the deformations using the OML by Li et al. [79]. This work made use of pre-deforming the geometry; in an ideal situation, the geometry would then deform to the original design under the thermal and structural loads. Limited success was shown using this pre-deforming strategy, since the geometry was not always restored to the ideal shape after deformation. Additionally, this study did not assess the impact of pre-deforming the vehicle on the off-nominal conditions, which would cause non-ideal deformations. However, the design of the outer-mold-line of the vehicle can still be used as a tool to alleviate the repercussions of deformation, specifically using a multidisciplinary shape optimization approach.

1.3.2 Hypersonic Optimization

Because of interdisciplinary challenges like scramjets, experts previously recommended that multidisciplinary analysis and optimization be a 2030 Grand Challenge for CFD in order to "advance aerospace engineering" [105]. Hypersonics is one of many fields that would benefit from the advancement of MDAO capability. As a complement to costly experiments and flight tests, an MDAO approach enables engineers to further explore the design space. According to Bowcutt [15], due to the strong interdependence between the vehicle and engine, "the conceptual design of this class of vehicle requires a multidisciplinary design optimization."

As mentioned previously, the hypersonic regime loosely begins around Mach 5, but the more significant indicator differentiating supersonic and hypersonic flow is the beginning of complex physics occurring within the air. A strong shock will generate high pressures and high temperatures that can cause the internal energies of the air molecules (such as translational, rotational, and vibrational) to stray from perfect gas behavior and potentially equilibrium. At even higher temperatures, the air molecules can chemically react, dissociating and potentially ionizing around a hypersonic vehicle. The chemical and thermal energy mode excitations and potential nonequilibrium occurring within the flow critically impacts the aerodynamic performance and heating of the vehicle. Due to the cost and complexity of flight and ground tests, accurate and efficient computational fluid dynamics codes have been developed to resolve the physical phenomena, or aerothermodynamics, occurring within the flow at hypersonic speeds [83, 22, 126, 40].

MDAO has been used for hypersonic vehicles in the past, but early design techniques for hypersonic vehicles leveraged fast analytical methods to produce preliminary optimal configurations due to the computational costs [114, 120]. Analytical methods such as the Taylor-Maccoll equation, modified Newtonian theory, and oblique shock relations have been previously used to optimize waveriders [17], the Mars Science Laboratory (MSL) [114], and scramjet inlet designs [109], respectively. These analytical techniques allow for rapid aerodynamic computations, which are integrated into optimization algorithms for design space exploration. Similarly, advances in aero-

thermo-elastic optimization have been studied using analytical techniques for the flowfield. For example, Munk et al. [87] focused on topology optimization of an aircraft's wing using a supersonic panel method. Since these methods are computationally inexpensive, gradient-based and gradient-free optimization algorithms have been employed. The trade-off for this efficient calculation is that these techniques make several assumptions that can cause inaccuracies at high temperature, nonequilibrium conditions.

As computational resources have increased, more expensive calculations such as CFD are now utilized earlier in the design process, which provides much more accuracy for the flowfield [105]. The integration of such analyses into optimization algorithms has been shown to have a large benefit for the design. Neville and Candler optimized the MSL aeroshell [88], similar to Theisinger and Braun [114], but using CFD instead of Newtonian theory. Neville and Candler optimized an aeroshell at one specific trajectory point to minimize drag and heat flux. In order to reduce the number of CFD simulations necessary for optimization, a gradient-based method was chosen. The results determined that a spherical Apollo-like capsule geometry is optimal for the particular flight conditions considered in this study, potentially increasing the lift-to-drag ratio by 12% from the original MSL configuration. This geometry differs from the traditional design of a 70° sphere-cone, indicating the benefit of incorporating these types of analyses into the optimization process.

With the integration of CFD into the design optimization process, the type of optimization approach becomes even more important. One such decision is whether to use gradient-free (i.e. Nelder-Mead, genetic, particle swarm, etc.) or gradient-based (conjugate gradient, method of moving asymptotes, sequential quadratic programming, etc.) algorithms [84]. Often, a gradient-free method is associated with a global search and the potential for finding the global optimum, whereas a gradient-based method is associated with a local search converging to a local optimum. Due to the inherent benefits of a gradient-free method, many studies have incorporated CFD into that type of optimization approach. For example, Rodi studied the optimization of a cowl leading edge under an Edney Type-IV shock using a particle swarm technique [95], and Seager and Agarwal optimized the re-entry of a blunt body using a genetic algorithm [101]. But, even for

these simpler geometries, the number of iterations is often much larger for gradient-free methods than gradient-based. A gradient-free method often requires a large number of evaluations of the CFD solver that increases with the number of design variables and the computational expense can be prohibitive. Therefore, other studies have focused on gradient-based techniques to reduce the number of optimization iterations. For example, Neville and Candler [88] used a quasi-Newton gradient-based method to optimize MSL, and Eyi et al. [35] used a modified method of feasible directions to reduce the drag produced by the nose of the vehicle. In both studies, the gradients of the objective and constraints with respect to the design variables are computed using a finite-difference (FD) approach. Still, even with gradient-based methods, more complicated geometries parameterized with a large number of design variables are infeasible to optimize when using an FD approach as the number of CFD simulations scales linearly with the number of design variables. Furthermore, the FD approximation of the gradients may be sensitive to the perturbation size and the accuracy imposed on the convergence of the CFD solver.

More recent efforts have focused on leveraging analytical or semi-analytical sensitivity analysis techniques into the hypersonic CFD solvers in order to reduce the cost of computing the gradient as compared to a finite-difference technique. There are few codes currently capable of producing these sensitivities for the hypersonic flow solvers. There are two approaches for creating these types of codes: (1) beginning from a subsonic flow solver designed specifically for ease of computing these gradients and subsequently incorporating hypersonic features into the code [80, 24, 85] or (2) beginning from an established hypersonic CFD code that is extended to include the computation of these gradients [29]. The limitation of the former approach is that some flow phenomena, such as chemistry and thermal nonequilibrium, have not been fully incorporated into the fluid solver or the sensitivity analysis. The limitation of the latter approach, specifically for the Eilmer CFD solver used in Ref. [29], is that the computation of the Jacobian can be expensive when only used for the sensitivity analysis.

By leveraging these semi-analytical sensitivities, the ability to include hypersonic CFD in MDAO, specifically, aero-thermo-elastic optimization, becomes more tractable. Initial work on this

field includes Kline et al., who leveraged SU2 to study shape deformation of an inlet [67] and, separately, the optimization of an inlet geometry [68, 66]. Additionally, Kamali et al. developed an aero-thermo-elastic analysis framework [62] and a thermo-elastic sensitivity solver for hypersonic problems [63] using the in-house CFD solver, NSU3D. Kamali et al. then built on these two methods to develop an aero-thermo-elastic sensitivity analysis [61]. However, several assumptions were included in the analysis such as omitting the heat flux sensitivity and modeling the flowfield without real gas effects. Finally, the application was limited to flow over a panel. Therefore, although there have been steps towards a hypersonic aero-thermo-elastic optimization approach, more work is still needed to develop a comprehensive procedure that can then be applied to vehicle design.

Focusing the scope to scramjets, a similar history to the hypersonic optimization is observed for scramjets. The shape of the inlet was initially designed using oblique shock relations that formed isentropic compression inlets, such as multi-ramp inlets. Later, methods such as streamtracing became prevalent when designing inlets [12], for example, the Busemann inlet [20]; the streamtracing technique led to 3D inlet shapes such as the Rectangular-to-Elliptical Shape Transition inlet [108]. Subsequently, optimization techniques were introduced, initially using lower fidelity inviscid oblique shock methods [109, 5]. Finally, more in-depth optimization techniques have been used to optimize a scramjet inlet including high-fidelity CFD analysis [18, 32, 68]. However, the optimization work-to-date for scramjet inlet design has yet to include the aero-thermo-elastic deformations occurring with the vehicle.

1.4 Dissertation Scope and Outline

The goal of this research is to advance the current understanding of fluid-thermal-structural interactions for a scramjet inlet and assess the repercussions of the resulting deformations on the inlet, engine and vehicle performance. Another aim of this work is to propose mitigation techniques when designing the inlet, either through material choice, OML decisions, or thermal management. To assess the OML influence, a coupled aero-thermo-elastic sensitivity analysis and

shape optimization technique is developed and performed; an additional goal for this optimization is to show the benefit of incorporating FTSI earlier in the design process. The subsequent chapters of this dissertation delve into the specifics for these goals and a brief description of the remaining chapters are provided below.

The following chapter, Numerical Methods, introduces the pre-existing and newly developed tools needed for the goals of this work. The initial development focused on an aero-thermo-elastic analysis, which couples a hypersonic CFD tool and a thermo-structural finite element solver. Additionally, a low-fidelity analytical engine model is incorporated into the FTSI tool to assess component- and engine-level implications. The solution to these coupled solvers is considered a forward analysis. The second development effort introduced in Chapter 2 outlines the optimization process and sensitivity analysis, which the forward analysis is incorporated into. Because a gradient-based method is chosen for the optimization technique, a multi-physics semi-analytical method is developed to compute the necessary gradients efficiently. Limited updates to the pre-existing tools are necessary for the forward analysis, but major updates are needed for the sensitivity analysis, all of which are detailed in Chapter 2.

Although there is some previous work for the aero-thermo-elastic response of scramjets, there are still several unanswered questions, as mentioned in Section 1.2 and 1.3.1. Therefore, Chapter 3 details several aero-thermo-elastic studies. This includes (1) understanding how heating and thermal management impact deformation and performance, (2) how and when the FTSI coupling can be more efficiently modeled, (3) how the deformations impact the sensitivity and performance of the engine/vehicle, and (4) how the thermal transient impacts the performance of the engine. Prior to those studies, a validation effort is performed against existing experimental data for an aero-elastic coupling to provide some confidence in the implementation of the numerical methods.

Chapter 4 transitions to focusing on benchmarking the sensitivity analysis. The semi-analytical method for computing the sensitivities that is employed for this work is an intrusive method that was not previously implemented in the hypersonic CFD solver. Therefore, the first part of this chapter details the benchmarking efforts for just the aerothermodynamic sensitivity

analysis on a simplified nose-cone geometry of a re-entry vehicle. Next, an integrated FTSI sensitivity analysis is performed on the same geometry, again to be used as a benchmark to confirm proper implementation. After that initial benchmarking effort, the sensitivity analysis is applied to the scramjet, which includes many more complex flow features such as shock-boundary layer interactions and turbulence. Once again, both an aerothermodynamic and aero-thermo-elastic sensitivity analysis are performed to confirm the flow features are properly resolved. Once the sensitivity analysis is benchmarked, the gradients are integrated into the optimization procedure.

Two major optimization studies are included in Chapter 5: a re-entry vehicle and a scramjet inlet. The chapter begins with a benchmarking case for the simplified re-entry vehicle nose, this time confirming that the sensitivities are properly integrated into the optimization. Next, a more detailed study is performed for the re-entry vehicle in order to minimize drag along its trajectory. This initial optimization is only performed using the aerothermodynamic forward and sensitivity analyses (i.e. no deformation is considered). Subsequently, a scramjet inlet optimization is performed with both aerothermodynamic-only cases and an aero-thermo-elastic case. The two types are compared to understand the influence of accounting for the deformations in the design process.

Finally, Chapter 6 summarizes the work performed for this dissertation and the primary contributions to the current state-of-the-art research. Additionally, recommendations/suggestions for additional future work on this topic are provided. Although many more studies are possible for scramjet deformation, the tools built in this dissertation are not limited in scope to scramjet inlets, and can have many hypersonic applications for the future. Finally, a summary of the publications for various components of this work are detailed for further reference.

Chapter 2

Numerical Methods

A scramjet inlet is highly integrated with both the engine and the vehicle. Any design changes to the inlet (i.e. shape, material, structure, etc.) will have impacts to the aerodynamic behavior of the vehicle since the forebody is used as both a compression surface for the inlet and a lifting body for the vehicle. Also, flow conditions produced by the inlet will strongly affect the conditions within the isolator, combustion chamber and nozzle of the engine. By perturbing the inlet flowfield, off-design conditions in the subsequent components can occur; this could lead to efficiency losses or in extreme cases, component failures. Changing the inlet, combustor and nozzle conditions will also impact the total engine performance, the accuracy of which guidance and control engineers rely upon. Therefore, a tightly-coupled, high-fidelity aero-thermo-elastic analysis is crucial. This chapter begins by describing the physics solvers necessary to compute the forward aero-thermo-elastic analysis; this includes a computational fluid dynamics code, LeMANS, a finite-element tool, MORIS, and an analytical system model of the engine. Including a model of the completed engine, even with a lower fidelity approach, highlights important implications of the inlet's FTSI. All of these tools are coupled in an iterative analysis described in the following section.

Because FTSI is so influential for scramjet inlets, *designing* the inlet with the deformations in mind can allow for better performance. One method for incorporating FTSI into the design process is using a multiphysics optimization tool. The numerical methods and implementation of an optimization approach and the sensitivity analysis are included in this chapter. Much of the details provided focus on the new implementation of the aerothermodynamic sensitivity analysis.

In addition, that sensitivity analysis is integrated with a thermo-structural sensitivity analysis to produce accurate coupled gradients. The thermo-structural sensitivity analysis is detailed in this chapter, as well as the FTSI sensitivity analysis coupling. Finally, a general description of the optimization procedure and its implementation are introduced and will be further elaborated upon for specific problems in Chapter 5.

2.1 Physics Solvers

2.1.1 Computational Fluid Dynamics: LeMANS

LeMANS is a CFD code that solves the compressible Navier-Stokes equations. Specifically, LeMANS is used to solve the partial differential equations that govern the thermal and chemical nonequilibrium occurring in hypersonic flow applications. LeMANS has been benchmarked, verified, and validated against other hypersonics solvers and experiments [83, 49]. A brief overview of LeMANS is included in this section, but no significant modifications are made to the fluid solver portion of the code, so more detail on the capabilities of LeMANS can be found in prior work [82, 99]. The only minor additions to the code have been updating the boundary condition at the wall to allow for a varying temperature to accommodate the coupling scheme and creating an adaptive CFL ramping scheme, both of which will be described at the end of this section.

The conservation equations are solved with LeMANS as follows:

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot (\mathbf{F} - \mathbf{F}_d) = \mathbf{C} \tag{2.1}$$

where

$$\mathbf{U} = \begin{pmatrix} \rho \mathbf{Y}^T \\ \rho \mathbf{v}^T \\ E \\ E_{ve} \end{pmatrix} \qquad \text{and} \qquad \mathbf{C} = \begin{pmatrix} \dot{\mathbf{w}}^T \\ \mathbf{0}^T \\ 0 \\ \dot{w}_v \end{pmatrix} \tag{2.2}$$

are the conserved variables and source terms, respectively, and are originally developed from Refs. [76, 39]. Within the conserved variables, ρ is the total density; $\mathbf{Y} = (Y_1, ..., Y_5)$ is the density vector

for each species; $\mathbf{v} = (v_x, v_y, v_z)$ is the velocity vector; E and E_{ve} are the total and vibrational-electron-electronic energies per unit volume of mixture, respectively; $\dot{w}$ is a source term for each species; and $\dot{w}_v$ is the vibrational energy source term. Additionally, the inviscid ($\mathbf{F}$) and diffusive ($\mathbf{F}_d$) flux matrices in the x-direction are given by

$$F_x = \begin{pmatrix} \rho Y_1 v_x \\ \vdots \\ \rho Y_5 v_x \\ \rho v_x^2 + p \\ \rho v_x v_y \\ \rho v_x v_z \\ (E+p)v_x \\ E_{ve} v \end{pmatrix} \quad \text{and} \quad F_{d,x} = \begin{pmatrix} -J_{x,1} \\ \vdots \\ -J_{x,5} \\ \tau_{xx} \\ \tau_{xy} \\ \tau_{xz} \\ \tau_{xx} v_x + \tau_{xy} v_y + \tau_{xz} v_z - (q_{tr,x} + q_{ve,x}) - \Sigma(J_{x,s} h_s) \\ -q_{ve,x} - \Sigma(J_{x,s} e_{ve,s}) \end{pmatrix}, \tag{2.3}$$

where p is the pressure, $\mathbf{J}$ is the directional species diffusion flux tensor, $\boldsymbol{\tau}$ is the viscous tensor, $\mathbf{h}$ is the species enthalpy vector, $\mathbf{e}_{ve}$ is the vibrational-electron-electronic energy vector, and $\mathbf{q}_{tr}$ and $\mathbf{q}_{ve}$ are the directional translation-rotational and vibrational-electron-electronic heat flux vectors, respectively. The mixture pressure, p, is derived from Dalton's law of partial pressures and the perfect gas law for each species [39]. More information on the conservation equations and equations of state can be found in Ref. [39].

The code has a multi-temperature modeling capability for thermal nonequilibrium; this research utilizes Park's two-temperature model [91]. In the two-temperature model, accounting for the energy modes assumes the translational and rotational temperatures are in equilibrium and the vibrational, electronic and electron translational temperatures are in equilibrium. Chemical nonequilibrium is evaluated using a finite-rate model with modified Arrhenius rate coefficients. A 5-species model is used (N_2, O_2, NO, O, N), which neglects ionization. LeMANS includes various Reynolds Average Navier Stokes (RANS) models for turbulent flows, including the Menter Base-

line, Menter Shear Stress Transport (SST) and the Spalart-Allmaras models. The viscous tensor is calculated assuming a Newtonian fluid and Stokes hypothesis:

$$\boldsymbol{\tau} = \mu(\nabla \mathbf{v} + (\nabla \mathbf{v})^T) + \lambda(\nabla \cdot \mathbf{v})\mathbf{I} \qquad , \tag{2.4}$$

where $\mathbf{I}$ is the identity matrix, λ is the second viscosity and μ is the mixture coefficient of viscosity. The species mass diffusion fluxes are modeled using Fick's law [113]. The transport properties of a general gas mixture are obtained using the semi-empirical Wilke mixing rule [124]; each individual species viscosity is obtained by Blottner curve fits [14]. The conductivity is calculated using Eucken's relation [119] and the conductive heat fluxes for each individual temperature are calculated using Fourier's law. All species use an equal diffusion coefficient that is calculated assuming a constant Lewis number of 1.4.

LeMANS uses a second-order accurate finite-volume method to solve the partial differential equations in the spatial domain. A parallel line-implicit method is used to solve the first-order accurate time integration. Due to the large gradients normal to the wall, LeMANS solves the system of equations in that direction using a tridiagonal solver and relaxes the system in the other directions for efficient convergence [126]. The parallelization allows for using domain decomposition, which accelerates convergence. The inviscid fluxes are computed using a modified form of the Steger-Warming flux vector splitting across the cell faces and the viscous fluxes are computed using a second-order accurate central scheme; this scheme is chosen for its accuracy within the boundary layer and low dissipation.

The two major updates for accommodating a scramjet inlet and coupling to a FEM code are adaptive CFL ramping and variable temperature boundary conditions. When analyzing scramjet inlets, a trend appeared where the initial ramping of the CFL often caused a divergence or plateauing issue of the residuals. The cause was due to the interaction of the forebody shock with the cowl, which required a reduced CFL to capture. Therefore, an adaptive CFL ramping scheme is implemented to detect any divergences or plateauing and then to ramp down the CFL until the issue is resolved. After resolving the shock interaction with the cowl, the CFL is ramped up once

again. The second addition is that a variable wall temperature must be applied as a boundary condition based on the FEM temperature results. To do so, LeMANS reads in a new file based on the temperatures interpolated from the FEM solution and applies it to each wall node boundary. The application of the temperature is similar to the isothermal boundary condition, except a vector assigns the temperatures to each node rather than a single constant.

2.1.2 Finite Element Analysis: MORIS

The structural and thermal analyses are performed using the multi-physics finite element code MORIS, developed by Noël et al. [89]. The research code uses standard and immersed finite element methods to solve partial differential equation constrained optimization – focusing on topology and shape optimization – using isogeometric analysis. MORIS employs the eXtended IsoGeometric Analysis (XIGA) methodology (leveraging B-spline basis functions), which is an augmentation of the traditional eXtended Finite Element Method (XFEM) methodology (leveraging Lagrangian basis functions). The forward analysis employs only certain features of the code, beginning with the thermo-structural FEM analysis with a body-fitted mesh approach.

A staggered approach is taken for this analysis between the thermal and structural response. The thermal response is solved using a bulk diffusivity model, an aerothermodynamic heat flux and radiation boundary condition, and an internal fixed temperature to represent the effects of active cooling. The conservation energy equation for this thermal analysis is given by

$$\rho C_P \frac{\partial T}{\partial t} = \nabla \cdot (k \nabla T) \qquad , \tag{2.5}$$

where T is temperature, ρ is the density of the material, C_P is the specific heat, and k is the thermal conductivity; since this analysis assumes steady state, the left-hand side of the equation is zero. The weak form of this equation can then be obtained using the divergence theorem and the domain can be discretized for the finite element method:

$$\mathbf{KT} = \mathbf{q} \qquad . \tag{2.6}$$

In this expression, $\mathbf{K}$ is the thermal conductance matrix, $\mathbf{T}$ is a vector of temperatures and $\mathbf{q}$ is a vector of applied heat fluxes. Also note that the thermal diffusivity for the materials are temperature dependent, which is shown in Chapter 3.

The structural analysis includes a plane-strain linear elastic model with a fixed boundary condition at the base of the scramjet inlet. Since the deformations are relatively minor, a linear-elastic model is used, but future work could include research into non-linear modeling. The static elastic equation for the structure solved by MORIS is written as

$$\mathbf{K}_s\mathbf{u} = \mathbf{f}_s + \mathbf{f}_T \quad , \tag{2.7}$$

where $\mathbf{K}_s$ is the traditional stiffness matrix and $\mathbf{u}$ is the displacement vector. The final terms, $\mathbf{f}_s$ and $\mathbf{f}_T$, are the aerodynamic force and thermal loading vectors, respectively, applied along the outer surface derived from the aerothermodynamic analysis results. Finally, the material properties (e.g. Young's Modulus, coefficient of thermal expansion, conduction, etc.) used to compute the stiffness matrices are temperature dependent and will be discussed in further detail in Chapter 3; all materials are assumed to be isotropic.

2.1.3 Engine System Analysis

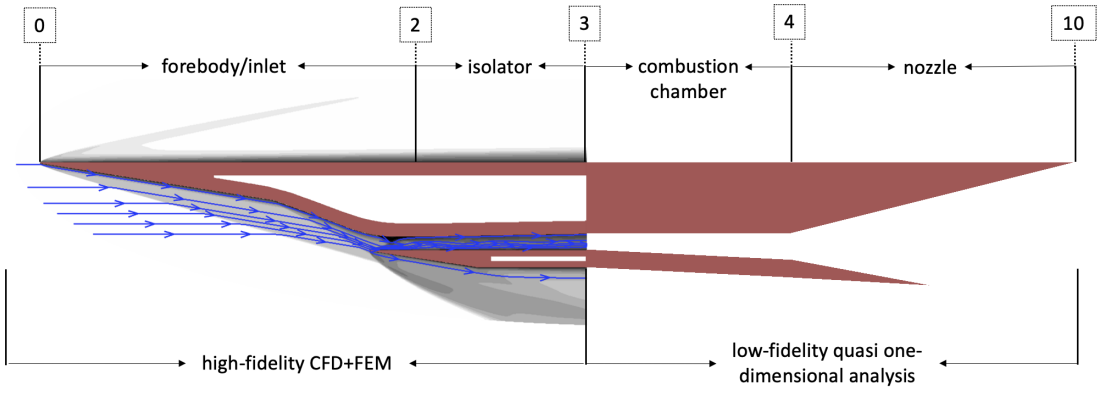

Figure 2.1: Flow stations for a scramjet inlet.

For the remaining engine components, a lower-fidelity approach is utilized to reduce computational cost while still accounting for the effects of the inlet on the engine performance as a whole.

Figure 2.1 (based on Heiser [46] and Smart [107]) depicts the various components along the flow path of the scramjet. The remaining sections will use the station labels corresponding to Figure 2.1. The high-fidelity CFD modeling includes both the inlet and the isolator; therefore, the engine system analysis begins at the interface between the isolator and the combustion chamber at station 3.

The combustion process and the nozzle expansion both are modeled using a quasi-one-dimensional approach. Therefore, it is necessary to extract flow properties from the two-dimensional (2D) CFD results at the end of the isolator and convert them to one-dimensional (1D) properties for the system model. The first task of the system model extracts the one-dimensional flow properties using a mass-flux weighted average technique. For this method, each one-dimensional flow property (represented by a bold-face $\boldsymbol{\psi}$) is calculated using Equation 2.8:

$$\boldsymbol{\psi} = \frac{\int \psi \rho (\boldsymbol{v} \cdot \boldsymbol{n}) dA}{\int \rho (\boldsymbol{v} \cdot \boldsymbol{n}) dA} \quad , \tag{2.8}$$

where ψ is an arbitrary property, ρ is the density, $\boldsymbol{v}$ is the velocity vector, $\boldsymbol{n}$ is the normal vector over the surface extracting from and dA is the differential area of that surface.

The system model then uses these one-dimensional results to approximate Mach number, pressure, total pressure, temperature, density, stagnation enthalpy and species mass fractions at station 3 to initialize the combustion chamber. A quasi-1D method that incorporates area change, heating due to the exothermic reactions along the chamber, friction and loss of heat to the wall is employed, following a similar analysis developed by Smart [106]. The combustion chamber area linearly expands with an area ratio of 2, beginning at the original height of the inlet. Skin friction is accounted for by a constant coefficient of friction, C_f and the heat lost to the wall, dQ, is accounted for by using the Reynolds analogy. Finally, the combustion heat release model uses an equilibrium chemistry assumption at each step, recalculating the gas constants using the tool CANTERA [41]. The model assumes a mixing efficiency curve that has been empirically found to calculate the total change in enthalpy shown in the equations below:

$$dH_t = h_{pr} f_{st} d\phi - dQ \tag{2.9}$$

$$\eta_m = \eta_{m,e} \frac{\Theta X}{1 + (\Theta - 1)X} \tag{2.10}$$

$$X = \frac{(x - x_3)}{(x_4 - x_3)} \quad , \tag{2.11}$$

where dH_t is the change in total enthalpy, h_{pr} is the heat of combustion, f_{st} is the stoichiometric fraction of fuel to air, and $d\phi$ is the equivalence ratio of fuel that combusts along a differential length [46]. In Equation 2.10, η_m is the mixing efficiency curve (found empirically), η_{m-e} is the mixing efficiency at the end of the combustor, X is a linear function of position and Θ is an empirical constant between 1 and 10. Hydrogen or kerosene fuel is assumed for the subsequent analyses.

Knowing the area change, total temperature change and friction distribution, the following ordinary differential equation is solved for the properties along the combustion chamber,

$$\frac{d(M^2)}{M^2} = \frac{-2\left(1 + \frac{\gamma-1}{2}M^2\right)}{1 - M^2}\frac{dA}{A} + \frac{(1 + \gamma M^2)\left(1 + \frac{\gamma-1}{2}M^2\right)}{1 - M^2}\frac{dT_0}{T_0} + \frac{\gamma M^2\left(1 + \frac{\gamma-1}{2}M^2\right)}{1 - M^2} 4C_f \frac{dx}{\tilde{D}} \quad , \tag{2.12}$$

where M is the Mach number, γ is the ratio of specific heats computed at every step along the reaction using CANTERA, A is the cross-sectional area, T_0 is the stagnation temperature, and $\tilde{D}$ is the hydraulic diameter of the section (which starts as the height of the isolator exit and increases with the area ratio) [4, 103]. By solving this ordinary differential equation, we can extract the flow properties along the length of the combustion chamber to compare the entire process for different inlet configurations. To incorporate equilibrium chemistry and real gas effects along the combustion process, the gas constants (e.g., γ, C_p and R) are calculated along the length of the chamber at each step as well, in a similar manner to Smart [107].

Leveraging the flow properties output from the combustion model, the nozzle module is then used to calculate the expansion and final gas properties. Isentropic expansion is employed

with a given area ratio. Typically, the nozzle efficiency is affected by under expansion, non-equilibrium/recombination effects, cross-flow and viscous/boundary layer effects; since these effects are difficult to model in a 1D analysis, an efficiency ratio is used to account for all losses.

Once each engine component is analyzed, the system-level quantities of interest are evaluated in terms of efficiency and capability. Initially, the component stream thrusts are calculated as metrics for each component's capability:

$$Sa_{inlet} = Sa_3 - Sa_0 \tag{2.13}$$

$$Sa_{burner} = (1 + f_{st})Sa_4 - Sa_3 \tag{2.14}$$

$$Sa_{nozzle} = (1 + f_{st})(Sa_{10} - Sa_4) \tag{2.15}$$

$$Sa_{area} = \frac{R_0 T_0}{v_0}\left(\frac{A_{10}}{A_0}\right) \quad , \tag{2.16}$$

where

$$Sa_{\#} = v\left(1 + \frac{RT}{v^2}\right) \quad . \tag{2.17}$$

Equation 2.16 refers to the impact of the area change across the engine, which is influential for deformation analyses. Then, the specific thrust ($F/\dot{m}_0$) is calculated as a metric for the total engine's capability using Equation 2.18.

$$\frac{F}{\dot{m}_0} = Sa_{inlet} + Sa_{burner} + Sa_{nozzle} + Sa_{area} \tag{2.18}$$

Various quantities such as thermal (η_{th}), propulsive (η_p) and overall (η_o) efficiencies are computed for the system using the following equations from Heiser [46]:

$$\begin{aligned} \eta_{th} &= \frac{v_{10}^2 - v_0^2}{2fh_{PR}} \\ \eta_p &= \frac{2}{v_{10}/v_0 + 1} \\ \eta_o &= \eta_{th} \cdot \eta_p \end{aligned} \tag{2.19}$$

Three inlet specific metrics are evaluated to assess capability and efficiency of the component: the compression ratio (PR) between the pressures at stations 3 and 0 represents the performance of the inlet; the total pressure recovery (π_c) between the stagnation pressures of same two stations showcases the efficiency of the compression; and the inlet kinetic energy efficiency (η_{KE}) defines the ratio of kinetic energy that would be achieved if expanded isentropically compared to the freestream kinetic energy. Assessing both the inlet quantities of interest and the system-level quantities of interest with a system-level analysis is necessary to estimate the impacts of the deformation on the vehicle.

2.1.4 Coupled Forward Analysis

Each of the different physics being modeled – aerothermodynamics, material response and structural response – exhibit different characteristic time scales. For general hypersonic problems, McNamara and Friedman assert that "the multidisciplinary nature of hypersonic vehicles requires inclusion of unsteady aerothermodynamics, temperature-dependent structural dynamics, and heat transfer analysis, as well as appropriate coupling mechanisms between each discipline" [86]. Note that McNamara and Friedman discussed the need for unsteady dynamics models, but since this vehicle is assumed to be at cruise, a steady-state model is used instead. Further, these authors state that the hypersonic aero-thermo-elastic problem can be split into two different disciplines: aerothermal and aeroelastic coupling. The coupling schemes between the two different disciplines are case dependent. Although some previous work in aero-thermo-elastic analysis points to the conclusion that the aerothermal coupling can be simplified to a one-way scheme without causing too large of overestimations for certain cases such as a flat plate, the analysis presented here includes a two-way coupling scheme for the aerothermal analysis [25]. A two-way coupling mechanism is used for two main reasons: (1) the heat flux from the aerothermodynamic analysis is essentially "free" and already computed and (2) with larger deformations, the changing aerothermal environment should have a pronounced effect on the thermo-structural deformation (this assumption is further studied in Section 3.4).

Therefore, because a large disparity in the flowfields appears between highly deformed cases, a two-way coupled analysis is employed between both the aerothermodynamics and the thermo-structural response. For the forward analysis, each response is assumed to reach a quasi-steady state, which is indicative of a vehicle at cruise conditions. The overall workflow for the forward analysis is depicted in Figure 2.2. The heat flux and tractions along the wall output from Le-MANS are fed into the MORIS thermo-structural analysis. The thermo-structural response uses a staggered approach, beginning with the thermal analysis. Once that portion is converged, the temperature field determined from MORIS is used as the wall temperature boundary condition for the CFD. An under-relaxation factor, ψ, is applied to the wall temperature passed from the thermal response solver to the flow solver:

$$T_{FS,n} = (1 - \psi)T_{FS,n-1} + \psi T_{TS} \qquad . \tag{2.20}$$

Here the subscripts FS and TS represent the flow solver and thermal response solver, respectively, and n represents the current iteration. A value of 0.5 is used for ψ, but additional work is needed to assess the convergence using other factors. Radial basis functions (RBF) are used to interpolate between the finite volume mesh of the CFD and the finite element mesh of the thermo-structural analysis. To inform the CFD solution of the thermo-elastic deformations, the fluid mesh is adapted so that the wall boundaries match the deformed geometry from MORIS. RBF interpolation methods are leveraged for the fluid mesh deformation since the method preserves orthogonality near the deformed wall and maintains the high quality of the mesh [102, 73]. Both uses leverage a thin plate spline as the RBF; more details on RBFs are outlined in the next section.

The coupled process outlined above is repeated until the deformations and vehicle temperature field reach convergence. An example of the convergence trend can be seen in Figure 5.14, where the maximum nodal difference between iterations is plotted for displacements and temperatures within the body. By the final iteration, both the propulsive and aerodynamic coefficients are changing by less than 0.1% and are deemed converged. The exception to this is when an unsteady flow

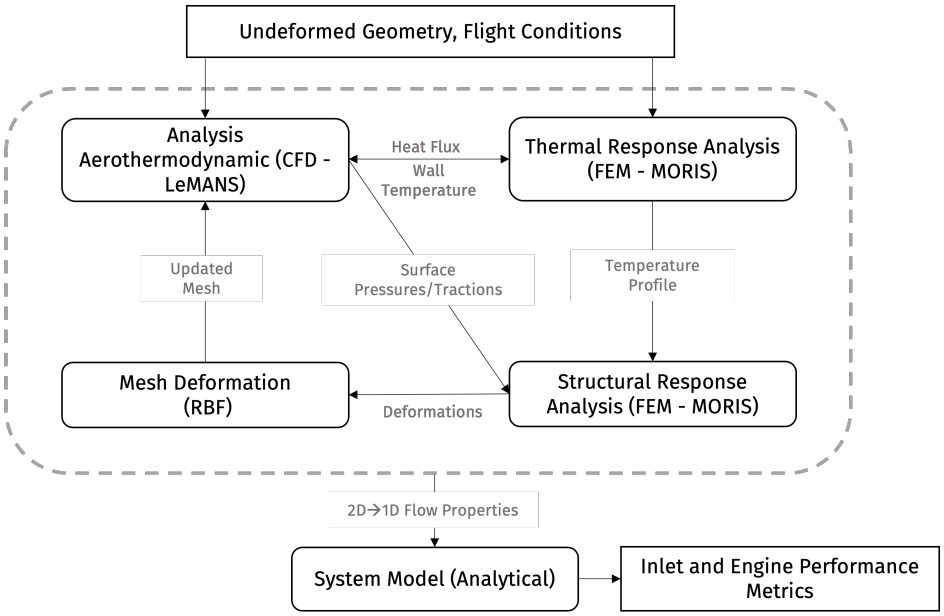

Figure 2.2: Aero-thermo-elastic two-way coupling scheme.

solution is found, in which case the presented results are just a snapshot of the aero-thermo-elastic response during that unsteadiness. Since the flowfield is dynamically changing in a repetitive manner, the flow solution does not fully converge to one steady-state solution; instead, the steady-state assumption is imposed by only analyzing a snapshot of the flowfield. Once the aero-thermo-elastic iterative analysis is converged, the necessary quantities are extracted from the isolator flowfield for the low-fidelity engine analysis.

2.2 Optimization Methods

Building on the forward analysis described above, a multidisciplinary optimization can allow for optimized designs that already account for the deformation. The focus of this work is on gradient-based optimization techniques that require some form of a sensitivity analysis. An analytical or semi-analytical sensitivity analysis reduces the computational cost of obtaining the gradient, allowing for more complex optimization problems with additional design variables. This makes solving high-fidelity multiphysics hypersonic problems much more feasible. The current section includes the formulation and setup of a gradient-based optimization technique, and the subsequent

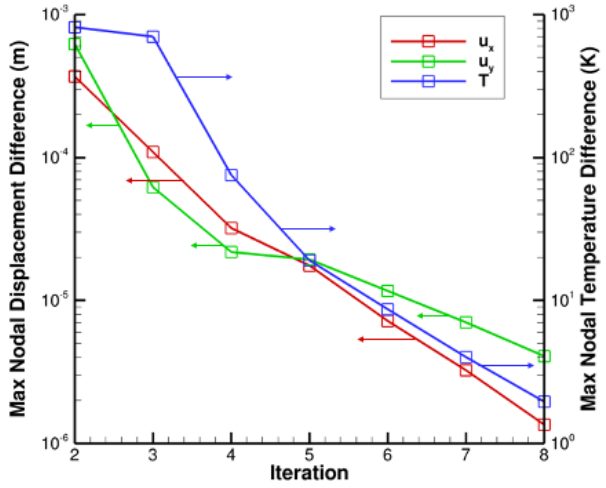

Figure 2.3: Maximum nodal displacement and temperature change for each coupled iteration.

section outlines the coupled sensitivity analysis processes to compute the gradient of the objective.

2.2.1 Optimization Approach

To reduce computational cost, a gradient-based optimization technique is implemented for this work. The algorithm chosen is a sequential quadratic program from Python's optimization toolbox, which allows for box constraints for the design variables. A generic aero-thermo-elastic optimization problem is formulated as follows:

$$\min_{s} \quad f(\mathbf{x}(\mathbf{s}), \mathbf{q}_F(\mathbf{s}), \mathbf{q}_T(\mathbf{s}), \mathbf{q}_S(\mathbf{s})) \tag{2.21}$$

$$\text{s.t.} \quad \mathbf{g}(\mathbf{x}(\mathbf{s}), \mathbf{q}_F(\mathbf{s}), \mathbf{q}_T(\mathbf{s}), \mathbf{q}_S(\mathbf{s})) \leq 0 \tag{2.22}$$

$$\underline{s}_i \leq s_i \leq \bar{s}_i, \qquad i = 1, ..., n_s \qquad . \tag{2.23}$$

The function, f, is the objective function, which can depend on the design variables, $\mathbf{s}$, through the mesh node locations, $\mathbf{x}$, and implicitly through the state variables of the fluid, $\mathbf{q}_F$, solid thermal, $\mathbf{q}_T$, and structural analyses, $\mathbf{q}_S$. The shape of the geometry is defined as a function of the design variables, $\mathbf{s}$. The optimization problem is also subjected to inequality constraints, $\mathbf{g}$, to maintain vehicle requirements such as maximum heating, temperature, unstart conditions, aerodynamic

performance, etc. Analytical methods leverage the knowledge of the governing equations, which maintains the physical feasibility of the solution:

$$\mathcal{R}_F(\mathbf{x}(\mathbf{s}), \mathbf{q}_F(\mathbf{s}), \mathbf{q}_T(\mathbf{s}), \mathbf{q}_S(\mathbf{s})) = 0 \tag{2.24}$$

$$\mathcal{R}_T(\mathbf{x}(\mathbf{s}), \mathbf{q}_F(\mathbf{s}), \mathbf{q}_T(\mathbf{s}), \mathbf{q}_S(\mathbf{s})) = 0 \tag{2.25}$$

$$\mathcal{R}_S(\mathbf{x}(\mathbf{s}), \mathbf{q}_F(\mathbf{s}), \mathbf{q}_T(\mathbf{s}), \mathbf{q}_S(\mathbf{s})) = 0 \qquad , \tag{2.26}$$

$$\tag{2.27}$$

where $\mathcal{R}_F$, $\mathcal{R}_T$, and $\mathcal{R}_S$ are defined as the residual vectors for the governing equations of the CFD, thermal solver and structural solver, respectively, indicating that the optimization requires the solution to adhere to a converged solution from each discipline.

2.2.2 Design Variables

2.2.2.1 Parameterization

Bézier curves are used to parameterize the surface of the vehicle for the geometry change process for the analysis and optimization. Therefore, the design variables for the shape optimization are defined as the control points for the Bézier curves. Specifically, the Bézier curves are used to define the *perturbation* of the surface from its original location shown as follows:

$$\mathbf{x}_w(\mathbf{t}) = \mathbf{x}_{w,org}(\mathbf{t}) + \sum_{i=0}^{n} B_{i,n}(\mathbf{t})\mathbf{P}_i \qquad , \tag{2.28}$$

where x_w are the coordinates of the altered body surface and $x_{w,org}$ are the coordinates of the original body surface, both in the global coordinate system. Additionally, $\mathbf{t}$ are parametric coordinates, $\mathbf{P}_i$ are the coordinates of the control points, n is the number of control points and $B_{i,n}(\mathbf{t})$ are the Bézier-Bernstein polynomials. The parametric coordinates, $\mathbf{t}$, are computed from the distance along the surface for each of the original node positions from the start of the curve.

2.2.3 Mesh Deformation

To account for the shape change of the surface, due to both the design variable changes and the structural deformation, the fluid mesh must deform to represent the correct wall boundary condition while still maintaining its high quality for the CFD analysis. The same technique to deform the mesh from the forward analysis is used here: RBF interpolation methods deform the mesh based on the wall's Bézier curve [73]. For this interpolation, the far-field nodes and any wall nodes that are not being altered by the design variables are fixed. The displacements of the internal nodes, $\mathbf{u}$, are computed as a sum of basis functions:

$$\mathbf{u}(\mathbf{x}) = \sum_{j=1}^{n_b} \alpha_j \phi(||\mathbf{x} - \mathbf{x}_{b_j}||) + \mathbf{p}(\mathbf{x}) \qquad , \tag{2.29}$$

where $\mathbf{u}$ are the displacements, n_b is the number of boundary nodes with known displacements, $\boldsymbol{\alpha}$ are coefficients, ϕ is the basis function and $\mathbf{p}$ is a linear polynomial. The wall positions, $\mathbf{x}_w$, are concatenated with the unchanged freestream, $\mathbf{x}_f$, and outlet positions, $\mathbf{x}_{out}$, to form $\mathbf{x}_b$, as shown in Equation 2.30. Similarly, $\mathbf{u}_b$ are the displacements of the nodes along the wall boundary, freestream and outlet, which are known from the Bézier curve evaluations or are fixed and zero:

$$\mathbf{x}_b = \begin{bmatrix} \mathbf{x}_w \\ \mathbf{x}_f \\ \mathbf{x}_{out} \end{bmatrix} \qquad \text{and} \qquad \mathbf{u}_b = \begin{bmatrix} \mathbf{u}_w \\ \mathbf{0}_f \\ \mathbf{0}_{out} \end{bmatrix} \tag{2.30}$$

To solve for the displacements for the entire mesh, $\mathbf{u}$, the coefficients vectors, $\boldsymbol{\beta}$, defining the polynomial $\mathbf{p}$, and $\boldsymbol{\alpha}$ must be determined by solving the following linear system:

$$\begin{bmatrix} \mathbf{u}_b \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} M_{b,b} & P_b \\ P_b^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha} \\ \boldsymbol{\beta} \end{bmatrix} \tag{2.31}$$

The matrix $M_{b,b}$ contains the evaluation of the basis function $\phi_{b_i,b_j} = \phi(||x_{b_i} - x_{b_j}||)$; for this analysis the thin plate spline is chosen as the basis function, where $\phi(x) = x^2 \log(x)$. Finally, $\mathbf{P}_b$ is a matrix for the linear polynomial with each row j given by $[1 \; x_{b_j} \; y_{b_j} \; z_{b_j}]$. With the coefficients computed, Equation 2.31 is evaluated to determine the displacements of all the nodes in the mesh.

For the current optimization procedure, the structural mesh is deformed in a similar manner. Radial basis functions perturb the mesh based on the new wall boundary dictated by the design variables. The mesh deformation tool is also used during the finite-difference process, when perturbations to the geometry are necessary to compute the thermo-structural state variable sensitivities (which will be elaborated on in Section 2.3.2).

2.3 Sensitivity Analysis

Because a gradient-based optimization procedure is chosen, an efficient coupled sensitivity analysis is necessary to drive the algorithm. An aerothermodynamic sensitivity analysis is built and integrated into the established thermo-structural optimization solver, MORIS. The sensitivity analysis for each individual discipline is first explained, followed by the coupling procedure for the multi-physics gradient computations. For this work, a semi-analytical direct method is implemented for the aerothermodynamics portion, since it is the more computationally expensive portion. Meanwhile, the thermo-structural portion is relatively inexpensive for the number of design variables used in this analysis. Therefore, a finite-difference approach is used to compute the necessary sensitivity for those physics. The current benchmarking strategy introduced here focuses on a hypersonic vehicle nose tip, since this has been previously studied in the literature using the same CFD solver [35, 36], which is detailed in Chapter 4. The following section details the definition and procedure of the sensitivity necessary to build the objective function's gradient, which is expanded as follows:

$$\frac{df}{d\mathbf{s}} = \frac{df}{d\mathbf{s}}^{expl} + \frac{df}{d\mathbf{q}_F}\frac{d\mathbf{q}_F}{d\mathbf{s}} + \frac{df}{d\mathbf{q}_T}\frac{d\mathbf{q}_T}{d\mathbf{s}} + \frac{df}{d\mathbf{q}_S}\frac{d\mathbf{q}_S}{d\mathbf{s}} \quad , \tag{2.32}$$

where $d\mathbf{q}/d\mathbf{s}$ is the sensitivity of the state variables for each discipline to the design variables. The terms, $df^{expl}/d\mathbf{s}$ (the explicit derivative of the objective to the design variables) and $df/d\mathbf{q}$ (the explicit derivative of the objective on the state variables) can be computed analytically or through inexpensive finite-difference techniques depending on the objective. A similar execution of

the chain rule is used to determine the sensitivities of the inequality constraints, $\mathbf{g}$. To determine the gradient of the state variable for the objective and constraints computations, the derivatives of the residual equations are leveraged:

$$\begin{bmatrix} \frac{\partial \mathcal{R}_F}{\partial q_F} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \frac{\partial \mathcal{R}_T}{\partial q_T} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \frac{\partial \mathcal{R}_S}{\partial q_S} \end{bmatrix} \begin{bmatrix} \frac{dq_F}{ds} \\ \frac{dq_T}{ds} \\ \frac{dq_S}{ds} \end{bmatrix} = \begin{bmatrix} \frac{d\mathcal{R}_F}{ds} \\ \frac{d\mathcal{R}_T}{ds} \\ \frac{d\mathcal{R}_S}{ds} \end{bmatrix} - \begin{bmatrix} \mathbf{0} & \frac{\partial \mathcal{R}_F}{\partial q_T} & \frac{\partial \mathcal{R}_F}{\partial q_S} \\ \frac{\partial \mathcal{R}_T}{\partial q_F} & \mathbf{0} & \mathbf{0} \\ \frac{\partial \mathcal{R}_S}{\partial q_F} & \frac{\partial \mathcal{R}_S}{\partial q_T} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \frac{dq_F}{ds} \\ \frac{dq_T}{ds} \\ \frac{dq_S}{ds} \end{bmatrix} , \tag{2.33}$$

where $\partial \mathcal{R}/\partial \mathbf{s}$ is the explicit sensitivity of the residuals on the design variables, holding the state variables constant. Also, $\partial \mathcal{R}/\partial \mathbf{q}$ is the Jacobian of the discrete conservation equations solved for each discipline. Adopting the so-called direct method [84], Equation 2.33 is solved for each optimization variable, s_i, and the resulting derivatives, $d\mathbf{q}/ds_i$ are substituted into Equation 2.32 to obtain df/ds_i. The first row of Equation 2.33 is solved by the fluid module, while the second and third rows are solved by the thermo-structural module. Equation 2.33 is solved via a staggered approach, specifically a linear block Gauss-Seidel algorithm[84].

2.3.1 Fluid Sensitivity Analysis

Prior to this work, LeMANS did not have a semi-analytical method for computing gradients. LeMANS had only ever been used with a blackbox finite-difference approach for optimization [35, 36]. Therefore, this section outlines the details of implementing the direct method into the current version of LeMANS. Two examples are shown in later chapters, a re-entry vehicle and a scramjet inlet, but many other applications for hypersonics can leverage this sensitivity analysis and optimization process.

Fluid Mesh Sensitivity

Prior to solving either Equation 2.32 or 2.33, the sensitivity of the fluid mesh node displacements to the design variable must be computed. For the shape optimizations of this one discipline, the objective, constraints and state variables exclusively depend on the design variables via the position of the mesh nodes; thus:

$$\frac{df^{expl}}{ds_i} = \frac{df^{expl}}{d\mathbf{u}}\frac{d\mathbf{u}}{ds_i} \tag{2.34}$$

$$\frac{\partial \mathcal{R}_F}{\partial s_i} = \frac{\partial \mathcal{R}_F}{\partial \mathbf{u}}\frac{d\mathbf{u}}{ds_i} \tag{2.35}$$

Therefore, the sensitivity of the mesh node displacements, $\mathbf{u}$, on the design variable (Bézier control points) must be computed. This begins with the sensitivity of the wall boundary nodes on the control points, $\frac{dx_w}{ds}$. The derivative of the Bézier curves is a simple calculation, resulting in the coefficient from the initial definition of the perturbation, B_{i,n_B}.

For the computation of the mesh deformation, the derivative of Equations 2.29 and 2.31 is calculated analytically as follows:

$$\frac{d\mathbf{u}}{ds_i} = \sum_{j=1}^{n_b} \frac{d\alpha_j}{ds_i}\boldsymbol{\phi} + \alpha_j\frac{d\boldsymbol{\phi}}{ds_i} + \frac{d\mathbf{p}}{ds_i} \tag{2.36}$$

$$\frac{d\boldsymbol{\alpha}}{ds} = \boldsymbol{M}^{-1}\left(\frac{d\boldsymbol{u}_b}{ds_s} - \boldsymbol{\alpha}\frac{d\boldsymbol{M}}{ds_i} - \boldsymbol{\beta}\frac{d\boldsymbol{P}}{ds_i} - \boldsymbol{P}\frac{d\boldsymbol{\beta}}{di}\right) \tag{2.37}$$

$$\frac{d\boldsymbol{\beta}}{ds_i} = \boldsymbol{M}_p\left(\boldsymbol{\alpha}\frac{d\boldsymbol{P}^T}{ds_i} + \boldsymbol{P}^T\boldsymbol{M}^{-1}\left[\frac{d\boldsymbol{u}_b}{ds_i} - \boldsymbol{\alpha}\frac{d\boldsymbol{M}}{ds_i} - \boldsymbol{\beta}\frac{d\boldsymbol{P}}{ds_i}\right]\right) \tag{2.38}$$

$$\text{where} \qquad \boldsymbol{M}_p^{-1} = \boldsymbol{P}^T\boldsymbol{M}^{-1}\boldsymbol{P} \qquad . \tag{2.39}$$

These analytical derivatives are used to evaluate the sensitivity of the deformation, $\frac{d\boldsymbol{u}}{ds_i}$, which is equivalent to the coordinate sensitivities, $\frac{d\boldsymbol{x}}{ds_i}$. The gradient $\frac{d\boldsymbol{M}}{ds_i}$ is dependent on the derivative of the basis function, $\boldsymbol{\phi}$, and the gradient $\frac{d\boldsymbol{P}}{ds_i}$ is computed for each row as $[1 \ \frac{dx_{b_j}}{ds_i} \ \frac{dy_{b_j}}{ds_i} \ \frac{dz_{b_j}}{ds_i}]$. The resulting mesh deformation sensitivity then informs the computation of the residual sensitivity within the flowfield analysis.

Fluid State Variable Sensitivity

Prior to performing a flowfield sensitivity analysis, a fully converged steady-state flow solution is first computed. Once the residuals have converged to the required precision, the first row of Equation 2.33 is solved. The directional, explicit sensitivity of the residual on the mesh change

seen in the Equation 2.33 can be computed by finite-differencing:

$$\frac{\partial \mathcal{R}}{\partial \mathbf{s}}^{expl} = \frac{\mathcal{R}(\mathbf{x}_0 + \epsilon \frac{d\mathbf{x}}{d\mathbf{s}}, \mathbf{q}_0) - \mathcal{R}(\mathbf{x}_0, \mathbf{q}_0)}{\epsilon} \quad . \tag{2.40}$$

By perturbing the mesh proportional to the gradient without updating the state variables, a new residual is computed. Note that this finite-difference step must be performed for each design variable, but this requires only one evaluation of the residual, not an additional converged flowfield solution. When coupled to the structural response sensitivity analysis, the displacement sensitivity, $d\mathbf{u}/d\mathbf{s}$, is summed with the mesh position sensitivity, $d\mathbf{x}/d\mathbf{s}$. A similar technique is used to compute the sensitivity of the residual to the wall temperature, the output state variable from the thermal sensitivity model. Instead of perturbing the mesh node locations, the wall temperature is perturbed in the direction of the gradient computed in the thermal sensitivity analysis, $d\mathbf{T}_{wall}/d\mathbf{s}$:

$$\frac{\partial \mathcal{R}}{\partial \mathbf{s}}^{expl} = \frac{\mathcal{R}(\mathbf{x}_0, \mathbf{T}_{0,wall} + d\mathbf{T}_{wall}/d\mathbf{s}) - \mathcal{R}(\mathbf{x}_0, \mathbf{T}_{0,wall})}{\epsilon} \quad . \tag{2.41}$$

The next step in computing the state variable sensitivity, $d\mathbf{q}_F/d\mathbf{s}$, is solving the top row of Equation 2.33. The simplest method to solve the linear equation would be to use the Jacobian computed within LeMANS with the tridiagonal solver for the time implicit solution already in place. However, the Jacobian in LeMANS is an approximation, which has proven effective in solving the nonlinear problems but is not suited for sensitivity analysis which requires a consistent linearization of the residual with respect to the state variables. This is discussed in more detail in Chapter 4. Instead, an iterative solver technique is implemented, which allows for a Jacobian-free approach: generalized minimum residual algorithm (GMRES) [97, 37]. Rather than computing a consistent Jacobian, the method leverages the estimation of the following matrix-vector product of the Jacobian:

$$J(\mathbf{q}_0)w = \frac{\mathcal{R}_F(\mathbf{q}_0 + \epsilon w) - \mathcal{R}_F(\mathbf{q}_0)}{\epsilon} \quad , \tag{2.42}$$

where the Jacobian, J, is equivalent to $\partial \mathbf{R} / \partial \mathbf{q}$. When using a finite-difference method, both secant errors and round-off errors are present [84]. To reduce the errors, the step size, ϵ, is determined by

a method from Knoll and McHugh [69].

In order to increase the efficiency and robustness of the GMRES algorithm, a preconditioner is applied to the linear system. By using a right-preconditioned GMRES, the solution process already built into LeMANS can be leveraged to evaluate the inverse of the preconditioner and vector product, applying the estimated Jacobian from LeMANS as the preconditioner. In addition, restart and parallelization capabilities are included in the solver to reduce the memory necessary, while maintaining accuracy and efficiency.

2.3.2 Thermo-structural Sensitivity Analysis

The thermo-structural sensitivity analysis is performed using a forward finite-difference method. The computational cost for the thermo-structural analysis is much less than the CFD, making this option computationally tractable. The input thermo-structural mesh is deformed according to the design variables, and the heat flux and tractions are perturbed according to the resulting aerothermodynamic sensitivity analysis to compute the deformation and wall temperature gradients. Note that the thermo-structural mesh sensitivity is included in the finite-difference process itself and does not need a separate analytical computation. The structural mesh is adapted to the new shape due to the design variables. The same RBF method is again applied, with the altered surface updated based on the Bézier curve while the others remain fixed.

2.3.3 Coupled Sensitivity Analysis

The strategy for integrating the three disciplines' sensitivity analyses (SA) is an iterative process and is shown in Figure 2.4. This process for the sensitivity analysis is encapsulated by the optimizer. A brief description of the process for obtaining the gradient of the objective with respect to the design variables is outlined below:

(1) The sensitivity analysis begins with a converged aero-thermo-elastic iteration.

(2) The optimizer determines the gradient of the wall shape, dx_w/ds, using the Bézier curve

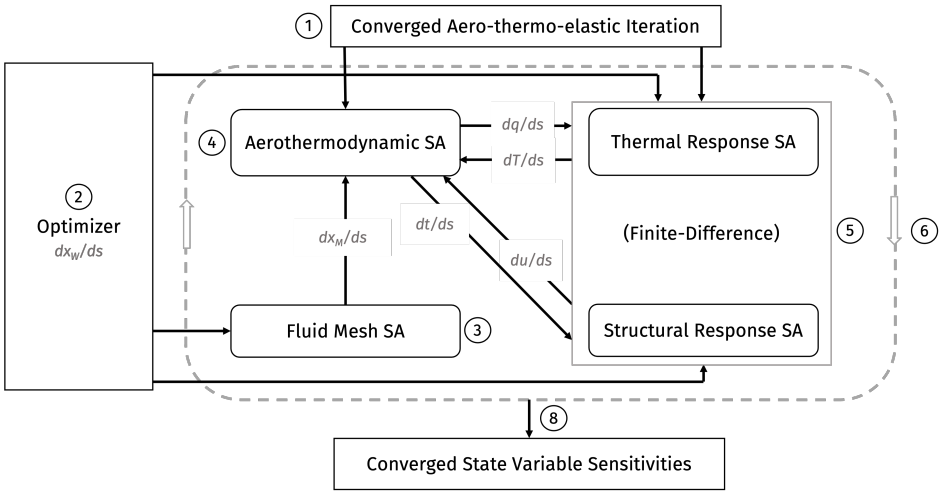

Figure 2.4: Coupled sensitivity analysis diagram.

definition and passes that to the fluid mesh, structural response and thermal response SA's.

(3) The fluid mesh SA computes dx_m/ds and passes that to the aerothermodynamic SA.

(4) The aerothermodynamic fluid SA computes the state variable sensitivities for the heat flux at the wall, dq/ds, and the traction along the wall, dt/ds.

- If this is the first iteration, the state variable sensitivities from the thermal/structural response are zero and the fluid mesh SA is the only driver,
- Else, the state variable sensitivities from the previous thermal/structural response are used (dT/ds and du/ds) along with the fluid mesh SA.

(5) The resulting sensitivities, dq/ds and dt/ds, are used to perturb the inputs to the thermo-structural finite-difference process. In addition, dx_w/ds is used to perturb the thermo-structural mesh using RBFs.

(6) The thermo-structural finite-difference is used to compute the state variable sensitivities for temperature, dT/ds, which are then passed back to the aerothermodynamic SA for the subsequent iteration with relaxation.

(7) The thermo-structural finite-difference is also used to compute the objective variable sensitivities for displacement, du/ds, and adds them to the fluid mesh sensitivities for the subsequent iteration with relaxation.

(8) Steps 4-7 are repeated until the state variable sensitivities have converged.

(9) The state variable sensitivities are then used to compute the objective function sensitivities using Equation 2.32.

2.4 Chapter Summary

This chapter lays the groundwork for the analysis in the upcoming three chapters. The following chapter, Aero-thermo-elastic Forward Analysis, utilizes the coupled forward analysis procedure that combines the two physics solvers: LeMANS and MORIS. These two high-fidelity techniques are able to compute the steady state of the fluid and structure of the inlet, which can then provide many important quantities for designing an inlet. This final state is also fed into a low-fidelity engine model to account for the performance and efficiency changes to the engine and study the effects of deformation on the conditions experienced in the combustion chamber and nozzle.

Chapters 4 and 5 then utilize the aerothermodynamic and coupled sensitivity analysis procedures outlined in this chapter. The direct method semi-analytical sensitivity analysis built into LeMANS for this work allows for robust and efficient gradient computations. In order to interface with this fluid sensitivity analysis, new features of MORIS were developed to allow for communication and proper computations of the surface mesh implementation. The multiphysics sensitivity analysis and optimization can be applied to several hypersonic applications, shown in Chapters 4 and 5.

Chapter 3

Aero-thermo-elastic Forward Analysis

3.1 Chapter Introduction

Because studying the deformation of a scramjet inlet is still in its nascent period, there are still many questions that need to be addressed. This chapter leverages the forward aero-thermo-elastic analysis and the integrated low-fidelity engine modeling to better understand the impacts, and possible mitigations of these deformations. This chapter begins to answer questions such as how thermal management may be leveraged to control these deformations (Section 3.3), how best to model the coupling when efficiency and accuracy are considered (Section 3.4), how these deformations may affect the vehicle's sensitivity to off-nominal conditions (Section 3.5), and how the transient behavior impacts the performance of the vehicle (Section 3.6).

Within this chapter, two inlet configurations are considered to evaluate these perceived gaps. There is limited available aero-thermo-elastic experimental data, especially on relevant geometries. Instead, previous studies, such as Ye et al. and Guangyue et al. [127, 42], have used the GK-01 inlet tested at the DLR Windtunnel in Cologne [43], which does not include any deformation. Therefore, the initial geometry is validated against the experiment, but the aero-thermo-elastic analyses are not validated against an experiment. More recently, Bhattrai et al. [10], from the University of New South Wales (UNSW), performed an experimental test campaign with an inlet to capture aero-elastic behavior at the University of Southern Queensland. Although the test capabilities were limited in their ability to capture thermal effects, the experiment provided important data for this work to validate the aero-elastic analysis against, which is explored in the following section.

3.2 UNSW Aero-elastic Validation

The campaign by Bhattrai et al. [10] provides one of the most comprehensive scramjet intake deformation datasets. The focus of the experimental campaign was to understand the influence of the aero-elastic response for a cantilevered compression surface on inlet quantities of interest such as TPR and flow spillage. The experiment was able to capture the dynamic behavior of the structure using the displacements of the ramp and the flowfield behavior using Schlieren imaging. Leveraging these previous ground-test experiments, the current work aims to validate the aero-elastic numerical modeling. Although thermal loads are shown to be an important factor on the distortion of the vehicle, this work focuses on just the aero-elastic influence for validation due to ground-based test facility capabilities (short duration and cold flow).

The design of the experiment performed by Bhattrai et al. [10] – experimentalists at UNSW – included a compression surface and inlet. The facility used was a free-piston compression-heated Ludwieg tube at the University of Southern Queensland, known as the TUSQ hypersonic test facility [21, 13]. The configuration has ample data to validate against, both for a fixed geometry and a deforming geometry. Note that the experiments were performed under cold flow conditions and a flow duration of 200 ms, so thermal strains are not able to be evaluated. Therefore, only an aero-elastic analysis is captured here. For more information on the experimental campaign, see Bhattrai et al. [10].

3.2.1 Test Setup

The methods used in this analysis are detailed in Section 2 for the aero-thermo-elastic forward analysis. An initial analysis, which included the thermal module, showed that there was only an increase in temperature of 2 K and no appreciable thermal expansion. Therefore, the thermal module is not included for this analysis and the CFD module uses an isothermal boundary condition. This updated aero-elastic forward analysis is then coupled to the analytical system model to understand the impacts to the engine.

The experimental configuration included a smooth, cantilevered compression surface, which acts as a compliant forebody surface of the inlet, and a fixed cowl. The compression surface is composed of a 3 mm thick, cantilevered, aluminum ramp extending into the flow, as shown in Figure 3.15. The nominal test conditions are shown in Table 3.1 and are used for the following aero-elastic analysis. Details on the geometry and design process for the experimental intake can be found in previous literature[10, 28]. To capture the necessary measurements, multiple instruments were used for the experiment. The cantilevered ramp was outfitted with two types of static pressure transducers; the isolator total pressure was measured with an additional pressure transducer; the flow visualization was captured using a high-speed Z-type Schlieren camera; and pressure distribution measurements were time-resolved using pressure-sensitive paint (PSP). Additionally, digital image correlations (DIC) measured the deflection of the ramp. Freestream measurements are not modeled in-situ in the core flow to avoid disturbances affecting the test. Instead, the pressure in the barrel upstream of the test section is measured and then the test section freestream conditions are extrapolated using isentropic relations.

Table 3.1: TUSQ freestream conditions.

$\mathbf{M}$	$\mathbf{p}_\infty$	$\mathbf{T}_\infty$	$\mathbf{Re}_\infty$	$\mathbf{T}_w$
5.85	755 Pa	75 K	7.1 $\times 10^6$ m^{-1}	300 K

3.2.2 Results

A two-dimensional (2D) representation of the geometry is depicted in Figure 3.15, which includes the Mach number contours with a freestream Mach number of 5.85. Impingement on the cowl nose from the leading edge shock of the forebody is shown, followed by an additional shock train within the inlet. A separation bubble is observed within the inlet, which causes increased pressures and temperatures in that region. Additionally, a numerical Schlieren image is shown in Figure 3.2, which closely matches the numerical results from both [10, 28]. However, the results do not accurately match the Schlieren images of the experiment from Bhattrai et al. [10] for the internal

compression region. The previous numerical results indicated that the experimental geometry had several three-dimensional (3D) effects and was not perfectly captured by a 2D analysis. For the experiment, the separation bubble was shown to be smaller, altering the angles, locations and strengths of the separation and reattachment shocks. The current CFD code, LeMANS, has been previously validated with other experiments [58], including a 2D inlet surface tested at the DLR Windtunnel in Cologne [43] (shown in Section 3.3); therefore, the aerothermodynamic difference is not likely a concern of the numerics, but instead an issue with ignoring the 3D effects of this experiment.

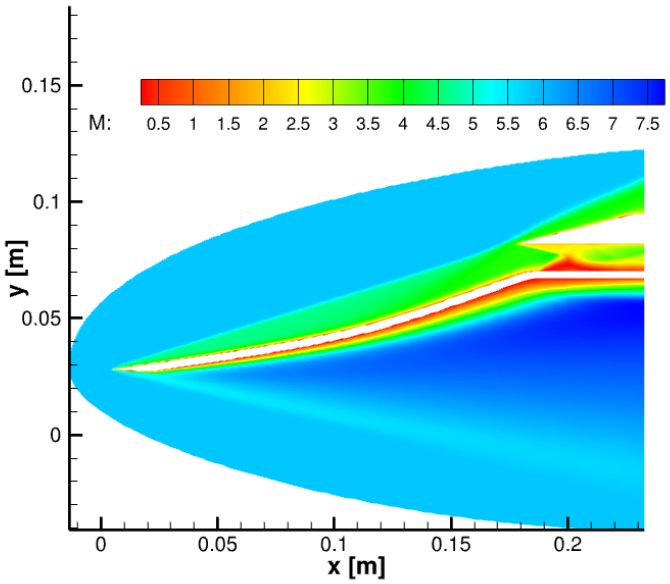

Figure 3.1: Mach number contours for compression surface.

The previous work by Bhattrai et al. [10] performed both an experimental campaign and CFD computations. The resulting pressure along the wall of the forebody for a fixed geometry is shown in Figure 3.3. Both a 2D and 3D analysis were performed; a large difference can be seen in the results, indicating that there are significant 3D effects for the tested configuration. This significance is further confirmed by the experimental data, which matches closely with the 3D analysis. Since only a 2D analysis is performed for this work, the 3D effects are neglected. However,

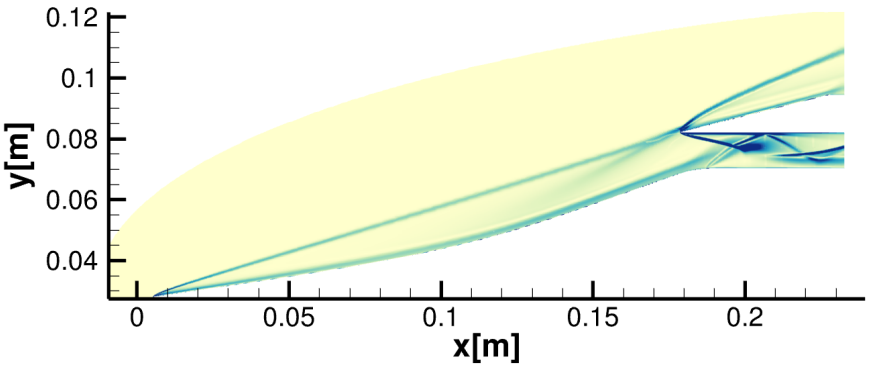

Figure 3.2: Synthetic Schlieren image generated using the magnitude of the density gradient obtained from the CFD solution for the experimental setup.

the 2D LeMANS results are shown to be in good agreement with the 2D results from the previous work.

Even with a 2D assumption, the forward aero-elastic analysis performs well compared to the experimental results. The deflections of the compression surface (normalized by the ramp length, S) are shown in Figure 3.4 for the test time. These deflections are captured by using both the Schlieren image results and digital image correlations. The experiment captures the transient dynamics of the compression surface, which is unable to reach a steady state during the test time, whereas this analysis only evaluates the aero-elastic behavior at steady state. Therefore, a comparison between the transient experiment and steady state solution is shown in Figure 3.4, indicating a close match in magnitude of the results despite the assumptions made. The accuracy is largely due to the higher contribution of the external compression region on the moment bending the geometry; the external region shows similar wall pressure results for both 2D and 3D representations. This entire validation effort is captured in Ref. [52]. With confidence that the aero-elastic coupling is implemented correctly, the rest of the chapter extrapolates on this and explores the aero-thermo-elastic effects on a scramjet inlet.

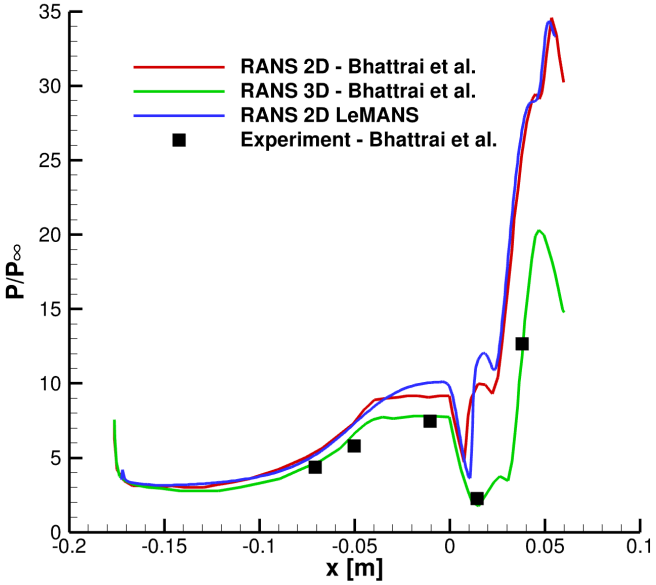

Figure 3.3: Comparison of wall pressure with previous CFD analysis and experiment[10].

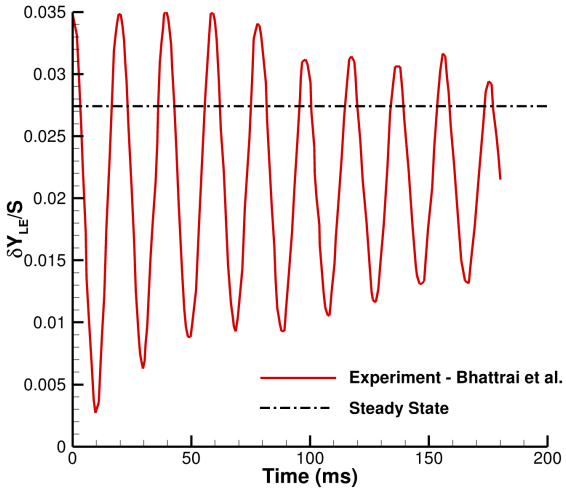

Figure 3.4: Comparison of leading edge deformation for a steady state solution against transient experiment[10].

3.3 Active Thermal Management

Previous studies on aero-thermo-elastic scramjet work concluded that the deformations caused unwanted changes to the flowfield behavior [11, 127, 77, 67]. Therefore, mitigation strategies should be explored to reduce the impact of the aero-thermo-elastic deformations on the engine performance. One such mitigation strategy works in tandem with thermal management: active cooling. Kline et al. showed that the thrust produced by the engine is highly sensitive to the temperature and deformation of the inlet [67]. Although active cooling comes at a cost of complexity and weight of the coolant (which can sometimes be repurposed fuel) and system, the benefits to the vehicle and engine can outweigh those drawbacks. Understanding the impacts of active cooling on the inlet and engine performance may tip the scales in favor of active cooling. Note that active cooling is used as an *example* of a potential thermal management strategy, but other techniques such as multi-material insulation, heat sinks or more complex structural layouts can also be used and will be discussed in the Recommendations for Future Work Section.

As mentioned in Section 1, the inlet performance impacts the subsequent engine components and the vehicle shape. Therefore, extending the effect of the deformation beyond the scope of the inlet is necessary to attain an understanding of how the vehicle reacts. Kline et al. and Riley et al. began studying this, both with a loosely coupled model [67, 92]. Both works concluded that there were impacts to the overall engine and vehicle performance. This work builds on the previous by including these engine and vehicle impacts as well as including more details about how the combustor and nozzle are affected. In addition, the flowfield in the inlet has a direct impact on the startability of the engine [106], which is also evaluated in this work.

3.3.1 Test Setup

The geometry chosen is representative of the inlet model GK-01, which was tested by Häberle and Gülham at the DLR Windtunnel in Cologne [43]. The campaign for this geometry has ample data to validate against for the undeformed geometry, and many previous aero-thermo-elastic anal-

yses have also used the GK-01 inlet [127, 42]. The two-dimensional geometry is shown in Figure 3.5. A fluid mesh of over 6×10^5 structured cells and a thermo-structural mesh of almost 7.5×10^4 unstructured cells are used in this analysis.

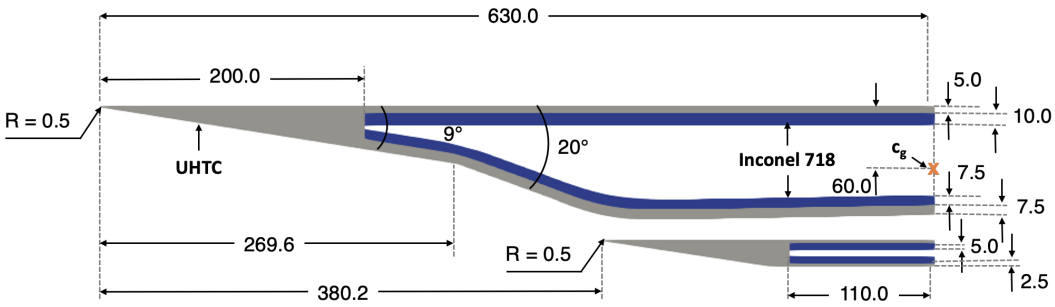

Figure 3.5: Schematic of GK-01 model for computation (units: mm).

A grid convergence study is performed for both the fluid mesh and the thermo-structural mesh. The convergence study for the fluid domain compares a coarse mesh (4×10^5 cells), medium mesh (6×10^5 cells) and a fine mesh (8×10^5 cells). Note that this convergence study is only done for the undeformed configuration. A comparison of the coefficient of pressure along the bottom surface of the forebody/inlet is shown in Figure 3.6. In addition, the convergence rate for the error in drag coefficient is plotted in Figure 3.7a, which is observed to match a 2nd order convergence rate. The equivalent mesh size is defined as $h = N^{-1/d}$, where N is number of cells and d is the spatial dimension. The difference in drag coefficient between the fine and medium mesh is 0.4%; since the error is minimal and in order to reduce computational expense, the medium grid is chosen. Similarly, a thermo-structural mesh discretization convergence study is performed; the aerodynamic boundary conditions for the thermal and structural problem use the undeformed solution from the CFD results. Three different meshes are compared again: coarse (2.5×10^4 cells), medium (4.0×10^4 cells), and fine (7.5×10^4 cells). The convergence rate for the error in displacement at the nose is plotted in Figure 3.7b; a 5th order convergence rate is observed. The error in displacement between the fine and medium is less than 4%, but since the thermo-structural analysis is a faster computation, the finest mesh is chosen.

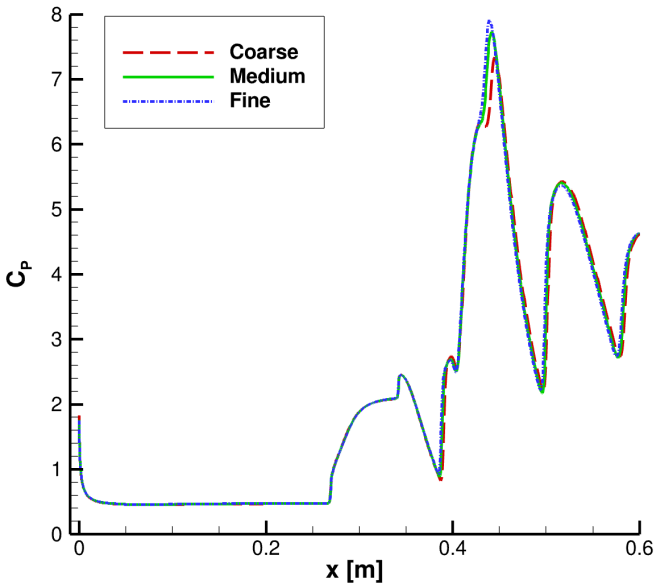

Figure 3.6: Coefficient of pressure along the interior surface of the forebody/inlet.

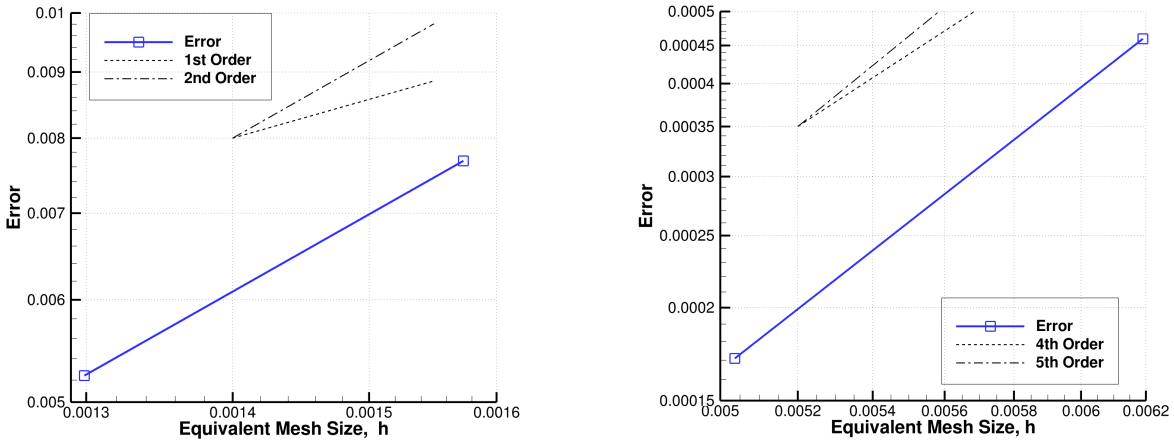

(a) Fluid.

(b) Thermo-structural.

Figure 3.7: Mesh discretization convergence for error in drag coefficient (left) and displacement at the nose (right).

A comparison of the computational results with the experiment for the original undeformed geometry is shown in Figure 3.8, which uses the cold-flow conditions from the experiment: a freestream temperature of 46 K, pressure of 170 Pa and Mach number of 7 [43]. The coefficient of

pressure (C_p) from the model is compared against the results from the original experiment. The two are shown to be in good agreement with each other, indicating that the wall properties of the model match up closely with the experimental results. The root-mean-square error is within 4% for the ramp and 13% for the cowl; the final data points are highly affected by the back-pressure, which is present in the experiment but is unknown and not enforced as a boundary condition for the CFD. Without the final three data point comparisons for the cowl, the root-mean-square error is reduced to less than 1%. The validation of wall properties prediction is especially important because the heat flux and traction from the CFD are inputs and driving factors for the thermal and structural analysis. The extracted quantities for the system analysis are taken slightly upstream within the isolator, where the CFD remains in good agreement with the experiment.

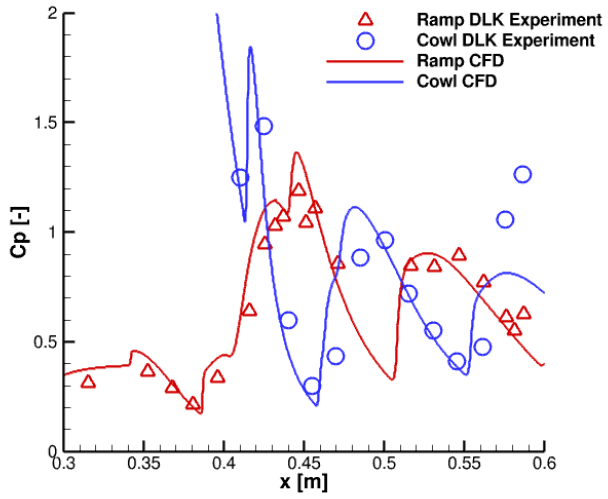

Figure 3.8: Comparison of the coefficient of pressure between experiment and CFD.

The flight conditions for the aero-thermo-structural analysis case are chosen to be at an altitude of 23.5 km and a Mach number of 7 to best compare against a previous aero-thermo-elastic study of the GK-01 engine [42]. As seen in Figure 3.5, the internal structure for the current analysis is kept very simple. The leading edge is built out, but then a panel-like profile is applied to the rest of the inlet. The material chosen for the leading edge is an ultra-high-temperature composite (UHTC), chosen for its ability to withstand high heat loads seen on sharp leading edges:

HfB_2/SiC composite. The material chosen for the paneling is Inconel 718 with a layer of the UHTC material serving as the thermal protection system. The properties for both materials vary with temperature as found in Refs. [47], and [111]; the properties necessary for the current analysis include the density, heat capacity, thermal conductivity, thermal expansion, Poisson's ratio and Young's modulus and are shown in Figure 3.49.

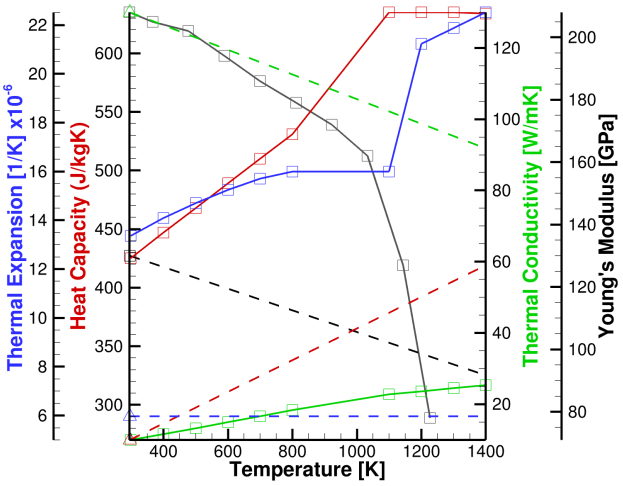

Figure 3.9: Material properties for Inconel 718 (solid lines) and UHTC (dashed lines).

To account for active cooling, a temperature is prescribed on the inside surface of the panels and nose of both the forebody and cowl. For the following active cooling analysis, the maximum cooling temperature is chosen based on keeping the Inconel safely below its melting temperature. But, this original calculation to determine the maximum cooling temperature only uses the *undeformed* heat flux results to show the change in temperature the vehicle can experience due to deformation. The minimum cooling temperature is chosen to be 300 K, since it would be beneficial for the Inconel to remain at room temperature and not thermally expand. The prescriptions along the surface can also be divided into the upper and lower paneling surfaces. The specific temperatures representing the active cooling will be discussed in the next section.

3.3.2 Deformation Results

Although not a fully converged coupled solution, the first iteration between the CFD and the FEM analyses gleans important information about how the structure performs under mechanical and thermal loads. The first iteration helps inform design decisions, such as choosing the cooling amount. In Figure 3.10, the first iteration deformation results, due to the initial aerothermodynamic analysis on the undeformed configuration, are compared for various Dirichlet boundary conditions (applied to all the inner surfaces). As expected, the case with the hottest internal boundary condition, 1200 K, shows the largest amount of deformation in the x-direction. The majority of the deformation in the negative x-direction is attributed to thermal expansion; since the 1200 K case provides the least amount of cooling to the vehicle, the temperatures and thermal expansion are the greatest.

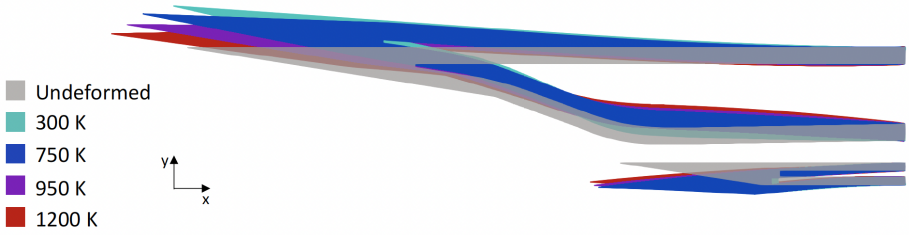

Figure 3.10: Surface contours of the deformed inlet with varying levels of active cooling after the initial coupled iteration (deformation scale factor: 5.0)

The deformation seen in the y-direction is more complicated since it is a combination of the structural displacements due to the imposed tractions and the thermal expansion due to the imposed heat flux. Figure 3.11, once again, is only from the initial iteration for the aero-thermo-elastic coupling. The analysis for these results applies the thermal-loading and the structural-loading separately to understand the resulting deformation of each. As mentioned earlier, the majority of the deformation in the negative x-direction is due to thermal expansion, but the majority of the deformation in the positive y-direction is also caused by thermal expansion. The structural-loading only contributes a small amount to the y-displacement. Instead, the thermal expansion plays an

important role, which is dependent on structure and material choices.

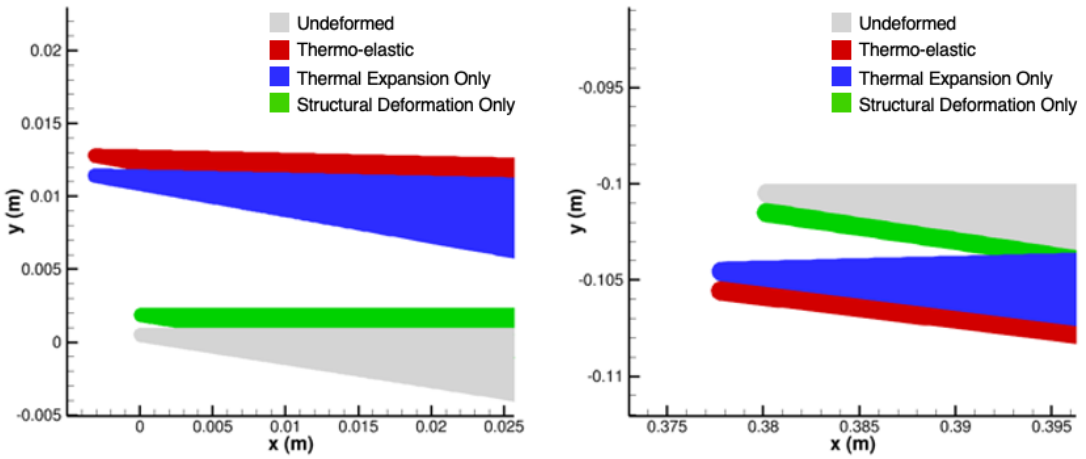

(a) Forebody leading edge

(b) Cowl leading edge

Figure 3.11: Thermal and structural deformation comparisons for 300 K active cooling case.

Focusing on the forebody for clarity (the same explanation applies to the cowl, as well), the top and bottom surface panels are heated unevenly from the flowfield. The air interacting with the top surface only experiences one weaker shock, while the air interacting with the bottom surface faces multiple stronger shocks and sometimes separation, as shown in Figure 3.12. The multi-shock system heats the airflow and the adjacent bottom surface. As the bottom surface heats up more than the top surface, the panels serve as lever arms. The bottom surface expands more due to the increased heat flux; this mismatched expansion of the top and bottom surfaces causes the forebody to pitch upwards in the y-direction, perpendicular to the freestream.

This opening behavior is more pronounced for the cases with a lower temperature boundary condition, as seen in Figure 3.13. Referring back to Figure 3.9, the thermal conductivity and heat capacity of the Inconel increase at higher temperature, promoting a lower temperature gradient across the panel and reducing the rise in temperature, respectively. The changes to the temperature-dependent thermal properties lead to less mismatched temperatures and expansion between the top and bottom surfaces, resulting in less of an opening effect for higher temperature active cooling

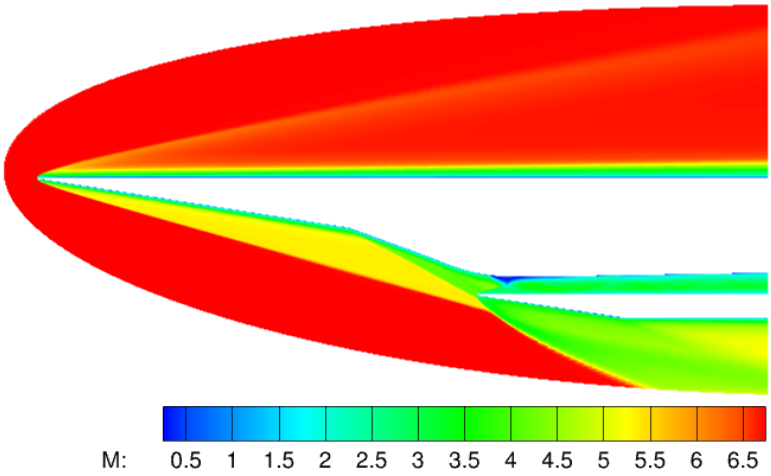

Figure 3.12: Flow solution contours of the Mach number for the undeformed configuration.

boundary conditions. Note that Inconel has a significantly larger thermal expansion coefficient, which is why the material drives the thermal deformation of the inlet more so than the UHTC.

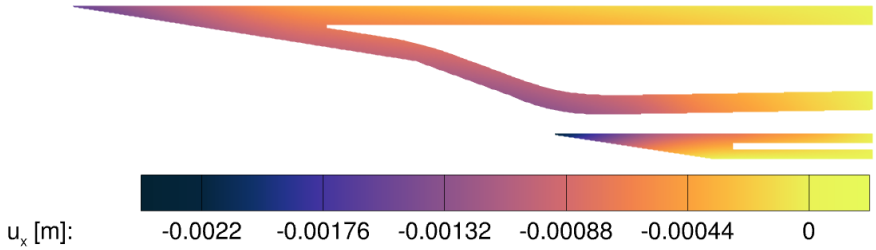

(a) 300 K active cooling case.

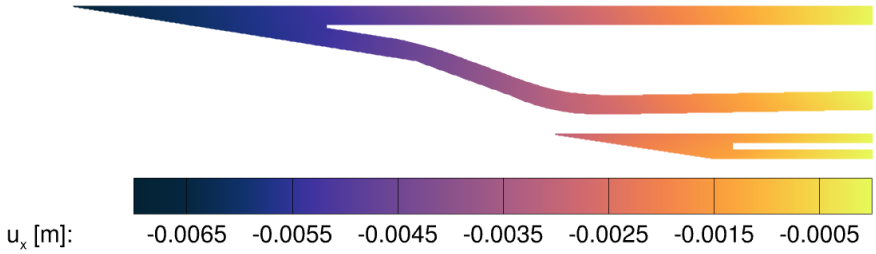

(b) 1200 K active cooling case.

Figure 3.13: Surface contours of displacement in the x-direction for the initial coupled iteration.

Overall, the repercussions of the mismatched thermal expansion lead to the behavior seen in Figure 3.10: the cooler interior boundary condition shows larger deformation in the y-direction. Similarly, the cowl will be deflected downwards in the negative y-direction, opening up the inlet. The opening of the inlet has a much more pronounced effect on the flowfield than the expansion upstream. Therefore, in an attempt to thwart this opening, different Dirichlet boundary conditions can be set to represent active cooling on the top and bottom surfaces of both the forebody and the cowl. By doing so, the opening of the inlet due to mismatched thermal expansion can be reduced.

Three uniform cooling cases are chosen to be studied for this work. The maximum and minimum temperatures are chosen based on the reasoning explained earlier in Section 3.3.1; the corresponding internal wall temperatures are 1200 K and 300 K, respectively. The third case is chosen to be in the middle of the two at 750 K. Finally, a fourth case is explored with distinct upper and lower wall temperatures to mitigate the opening effect. The top surface of the forebody, set warmer to increase the expansion on the top arm, is at 930 K and the bottom surface, set cooler to decrease expansion on the bottom arm, is at 750 K. The reverse is applied on the cowl for the same reasoning. The final case is meant to be compared against the constant 750 K case in order to show a mitigation strategy to the opening effect. The temperature of the top surface is chosen to eliminate the y-deformation using the heat flux from the *initial* coupled iteration.

Finally, with the temperatures for the four cases established, the deformation results from the fully-coupled analysis with multiple iterations are shown in Figure 3.14. The displacements for the forebody and cowl leading edges are listed in Table 3.2. The final deformations, although consistent with some of the trends shown in the initial deformation of Figure 3.10, are not identical to the initial deformation. The differences show the need for an iterative and tightly coupled approach to establish accurate results. The major differences are due to the changes that appear in the flowfield, which result in larger heating for the 750 K and 300 K cases. These differences are discussed in more detail in the next section.

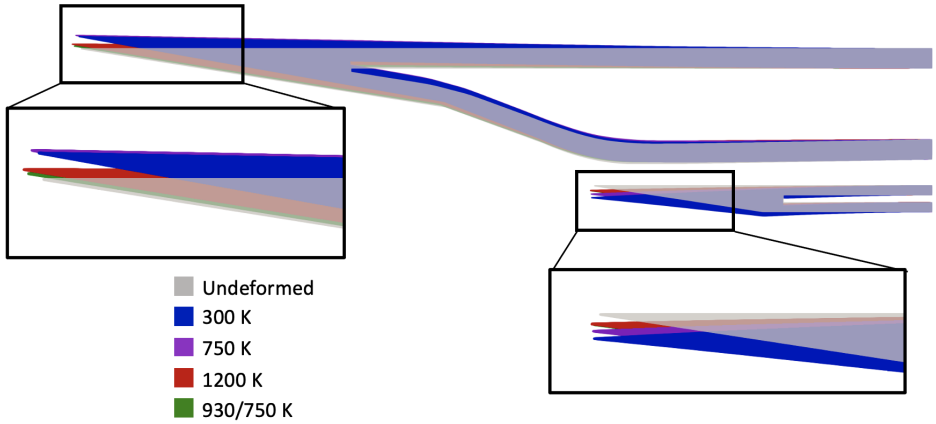

Figure 3.14: Surface contours of the deformed inlet with varying active cooling after the final coupled iteration.

Table 3.2: Maximum leading edge deformation.

		300 K	750 K	1200 K	930/750 K
Forebody	u_x	-1.5	-3.9	-6.7	-5.1
(mm)	u_y	9.0	9.9	3.4	2.0
Cowl	u_x	-2.2	-2.5	-3.1	-2.3
(mm)	u_y	-8.8	-6.1	-3.5	-3.7

3.3.3 Flowfield Results

Inlets are susceptible to many complicated flow structures such as shock-shock interactions, shockwave-boundary layer interactions (SWBLI), and boundary layer separation, all of which contribute to unsteady behavior, inlet losses, and increased heating, and can result in engine unstart; the original geometry is designed to minimize the negative impacts of such behaviors. The baseline GK-01 inlet is designed for the first shock from the forebody to remain ahead of the cowl at Mach 7, avoiding shock impingement on the cowl leading edge, as shown in Figures 3.12 and 3.15a. From the undeformed case, we can observe the internal shock from the cowl intersecting the shock from the separation along the top surface to form reflected and reattachment shocks within the channel. For the deformed geometries, the cowl deflects downwards and expands into the flow, altering the impingement on the cowl leading edge. High-fidelity modeling is needed for the deformed structure, not just the original geometry, to capture the changes to this flowfield seen in Figure 3.15.

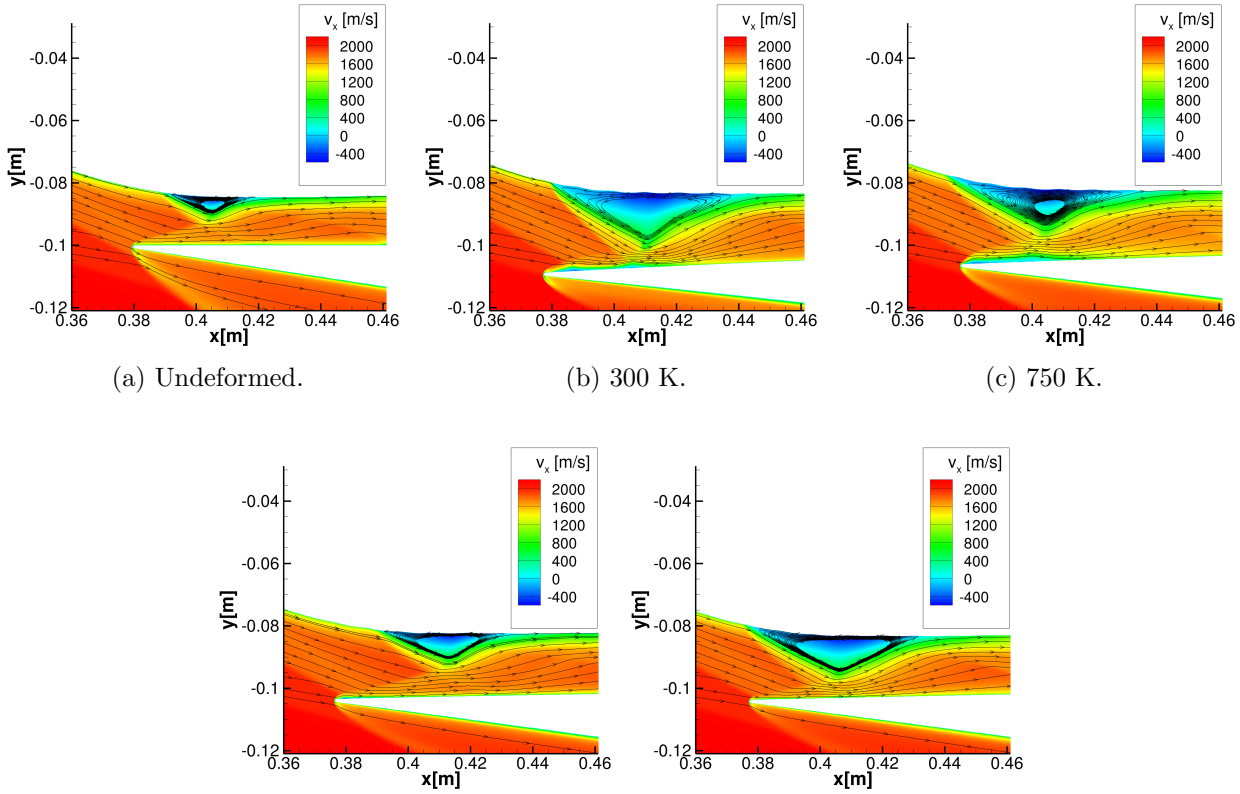

(a) Undeformed. (b) 300 K. (c) 750 K.

(d) 1200 K. (e) 930/750 K.

Figure 3.15: Contour plots of v_x around the cowl leading edge with streamlines.

Generally, for each of the cases, the deflection of the forebody upwards impacts the initial compression shocks. The flow deflection angle of the first ramp increases with the deformation, resulting in stronger oblique shocks. The 1200 K case and the 930/750 K case only show slight increases to the shock strength, while the 300 K and 750 K cases show a more drastic increase in shock strength, corresponding with the increased deformation in the y-direction. Additionally, the top surface, originally at zero angle of attack, deforms into an expansion surface as the nose is pitched upwards.

Starting with the 300 K case, the deformation to the body introduces new phenomena into the flowfield seen in Figure 3.15b. Not only is the secondary ramp's oblique shock now intersecting with the cowl's oblique shock, but also the first ramp's oblique shock is now impinging on the cowl's bow shock. The large deformation occurring for this case produces an Edney Type IV shock-shock

interaction between the impinging oblique shock and the bow shock [33]. Type IV interactions cause severe heating and boundary layer separation. The deformation also causes a more favorable pressure gradient along the expansion shoulder, lengthening the separation bubble upstream. The separation bubble (visualized as blue in the contour plots in Figure 3.15) on the top surface of the channel not only grows, but also enables a secondary separation along the bottom surface as that original separation shock increases in strength; this has implications such as reducing the overall velocity in the inlet, increasing the pressure and temperature, and potentially growing to the point of unstart.

Similar to the 300 K case, the 750 K case also shows a shock-shock interaction with the secondary ramp and a larger separation bubble along the shoulder that is strong enough to separate the flow along the bottom surface. These shock-shock interactions and boundary layer separations cause increased heating into the walls, which then influences the deformation through thermal expansion. Capturing these phenomena with the aerothermodynamics on the deformed body is why Figure 3.14 does not perfectly match Figure 3.10. Instead, the increased heating from the new flowfield phenomena further deflects the forebody upwards.

For both the 300 K and 750 K cases, the increased separation zones cause the flow to become unsteady, a very undesirable repercussion when trying to have a stable engine and thrust output. Therefore, the deformation and flowfield results for these two cases are just snapshots of the behavior within this unsteady regime. The flowfield and deformation oscillate as the separation zones interact and affect each other.

On the other hand, the 1200 K case remains steady, with only a small amount of separation occurring on the interior surface of the cowl. The shock from the secondary ramp still interacts with the cowl shock, but location and decreased intensity reduce the impact to the boundary layer. Of all the cases, the separation zone on the upper surface of the channel remains the closest to the undeformed case, indicating the favorable pressure gradient around the shoulder is most similar is strength, but the region still shows significant growth.

The two-temperature boundary condition case, 930/750 K, shows a similar result to the

1200 K case, since the deformation is mainly in the x-direction. The separation bubble on the top surface has expanded some, but the flow remains steady. Juxtaposing the 750 K case and the 930/750 K case, the two-temperature active cooling case is able to mitigate some of the effects from deformation. Not only has the separation region on the bottom surface disappeared, but the flow also remains steady. These benefits are in part due to the shock emanating from the first ramp remaining weaker than for the 750 K case as well as the shock from the secondary ramp interacting further upstream for the 930/750K case.

3.3.4 Thermal Results

The changes to the flowfield have many implications for the vehicle and engine design; one of which is the material and structural choices for the inlet. The maximum temperature active cooling case, 1200 K, was originally chosen to maintain the temperature of the Inconel paneling well below the melting point. But, this temperature was derived from the heat flux from the *undeformed* configuration. For obvious reasons, the temperature of the Inconel should be kept below the melting point, which could start as low as 1483 K. The original maximum Inconel temperature calculated from the undeformed geometry heat flux is 1275 K. The maximum Inconel temperature calculated from the converged deformed geometry heat flux is 1408 K, getting dangerously close to the melting point and surpassing the strength-to-weight drop-off seen in Figure 3.9 for Inconel.

All the cases show similar trends to the 1200 K case, with a stark increase in the maximum temperature of the Inconel. In Figure 3.16, the internal temperature contours for the 300 K case are shown. Not only has the entire vehicle heated substantially, but also there are now visible hot spot areas corresponding to the separation zones on the cowl and forebody. The 300 K case also shows an even more drastic change in maximum Inconel temperature from the initial iteration to the final, rising from 533 K to 1100 K.

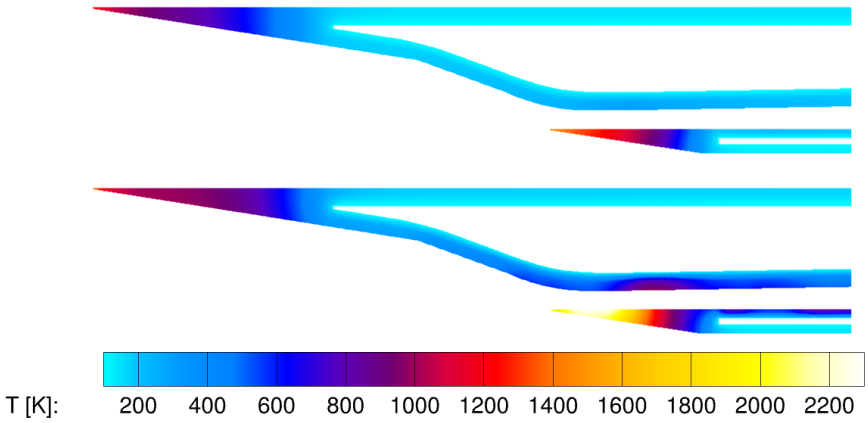

Figure 3.16: Surface contours of the vehicle temperature for the initial and final coupled iteration.

3.3.5 Inlet Performance

Other implications of the changing flowfield are seen in the performance of the inlet using capability and efficiency metrics. Using the system analysis tool, the flowfield results from the cases describe above – baseline, 300 K, 750 K, 1200 K, 930/750 K – shed some light onto the influence of the deformation on these metrics. The extracted flow properties at the exit of the isolator produce quantities of interest for the inlet, as shown in Table 3.3.

As a representation of capability, the compression of the inlet for the five cases is juxtaposed. Originally, the compression ratio hovers just over 27, which is lower than the recommended value by Smart for an efficient inlet [107]. The deformed cases show a significant increase in compression due to stronger external shocks and separation within the inlet. The increase in compression ratio is actually shown to be an improvement to engine performance in the current PR range. The PR's for each case are directly proportional to the degree of separation seen in Figure 3.15. Of the undeformed cases, the 1200 K case shows the smallest separation zone, and has the smallest PR. Meanwhile, the 750 K case shows the largest amount of separation, between the two zones on the forebody and cowl, and has the largest PR. These separation zones serve to stagnate the flow, driving up the pressure within the inlet.

In contrast, both the total pressure recoveries and the kinetic energy efficiencies – representa-

tions of the inlet efficiency – decrease for the majority of the deformation cases, with the exception of the total pressure recovery for the 1200 K case. For all but the 1200 K case, the Mach number at station three (M_3) decreases a significant amount, which greatly reduces the total pressure within the inlet. The 1200 K case is able to maintain a similar Mach number to the undeformed case since the capture area increases without too much change to the separation bubble size, enabling a more efficient total pressure recovery. The kinetic energy, only a function of the velocity, shows a decrease for all the deformed cases.

Finally, due to the widening of the capture area from the thermal and structural loading, the mass flow rate into the engine increases with deformation. The undeformed geometry shows the oblique shock from the forebody leading the cowl lip, which results in some overflow and a reduced mass flow efficiency. With the deformed geometries, the compression surface's secondary shock impinges past the cowl, allowing for a better mass flow efficiency. Depending on how the downstream components are designed, the increase in mass flow rate may cause heating/structural issues or incomplete reactions in the combustion chamber. Overall, although the deformations are on a much smaller scale than the engine and vehicle itself, they induce substantial effects on the inlet performance.

Table 3.3: Inlet quantities of interest.

	Undeformed	300 K	750 K	1200 K	930/750 K
PR	27.1	47.4	52.3	34.3	36.7
π_c	0.191	0.172	0.142	0.223	0.176
η_{KE}	0.610	0.506	0.469	0.591	0.540
$\dot{m}$ (kg/s)	0.497	0.733	0.749	0.688	0.635
M_3	2.65	2.27	2.13	2.59	2.43

Another integral performance metric for the inlet focuses on starting the engine. Establishing supersonic flow through the inlet is known as inlet starting, which is very sensitive to inlet contraction ratios. Many trends have been found through experiment that indicate when proper starting behavior should occur for an engine; one such trend for 2D inlets is the Kantrowitz limit [64]. This self-starting limit relates the ratio of the minimum engine area at the inlet (A_2) over the cross-

sectional area at the cowl closure point (A_t), or throat, to the Mach number at the cowl closure point. Figure 3.17 shows the Kantrowitz limit along with the isentropic compression limit. The two lines divide the plot into three regions: the region below the isentropic limit indicates where the inlet cannot work physically; the region above the Kantrowitz limit indicates where self-starting occurs; and the region between the two limits (in grey) indicates where the inlet will continue to run once already started, but will not self-start as is. The undeformed geometry is directly on the Kantrowitz limit, likely balancing the compression ratio with the ability to self-start. Meanwhile, the deformed configurations fall into the third region, indicating that the engine can continue to run, but cannot self-start any longer, which would limit the engine's restart capability for a given trajectory.

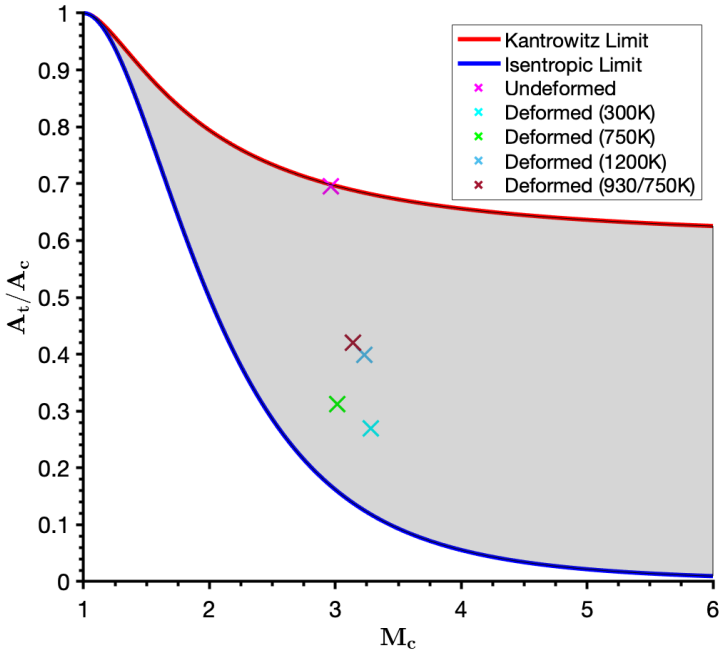

Figure 3.17: Self-starting criteria comparison for deformed geometries.

3.3.6 Engine Performance

Typically, the undeformed geometry is the design point for the engine, so any deviation from the original geometry will have impacts not only on the inlet, but the engine and vehicle as a whole.

The system analysis tool allows us to gather insight into the effects on the downstream components, starting with the combustion chamber. A comparison between the undeformed geometry results and the deformed geometries along the chamber is shown in Figure 3.18. Pressure, temperature and area are all normalized by the 1D quantities at station 3 for the undeformed geometry results.

In the most drastic case, 750 K, the pressure increases in the combustion chamber by over two-fold. The design of the combustion chamber would therefore need to accommodate the rise in pressure due to the inlet deformation. Additionally, the temperature increases within the chamber for all four deformed geometries. Therefore, the inlet deformation could lead to downstream repercussions such as design changes to the combustion chamber. As mentioned in the Inlet Performance subsection, the Mach number at station 3 for all the deformed cases is below the undeformed case. Similarly, the Mach number along the combustion chamber for those four cases continually remains below the undeformed case.

Although the inlet quantities of interest are very informative, the engine-level quantities are able to give a much deeper insight into the influence of inlet deformation on the engine performance. The breakdown of the specific thrusts for each component is shown in Table 3.4 to understand how each contribute to the total engine performance. The inlet stream thrust decreases with increasing deformation in the y-direction since the velocity at the exit is reduced due to stronger shocks and separation. Fortunately, the stream thrusts for the burner and the nozzle increase, which counteracts the impact of the inlet on the total specific thrust. Previous research by Smart [107] has shown that there is a positive correlation between thrust and compression ratios between 20 and 80 for this Mach number, so the higher PR seen in the deformed cases benefits the engine performance. As seen in Figure 3.18, the benefit comes at a cost of higher pressures and temperatures that the vehicle is not designed for.

Along with the engine capability, the engine performance can be compared for each case with the overall efficiency, taking into account the propulsive and thermal efficiency. Due to the increased separation occurring within the inlet, all deformed cases show a decreased efficiency compared to the original geometry. Interestingly, the 930/750 K case shows the greatest decrease since the

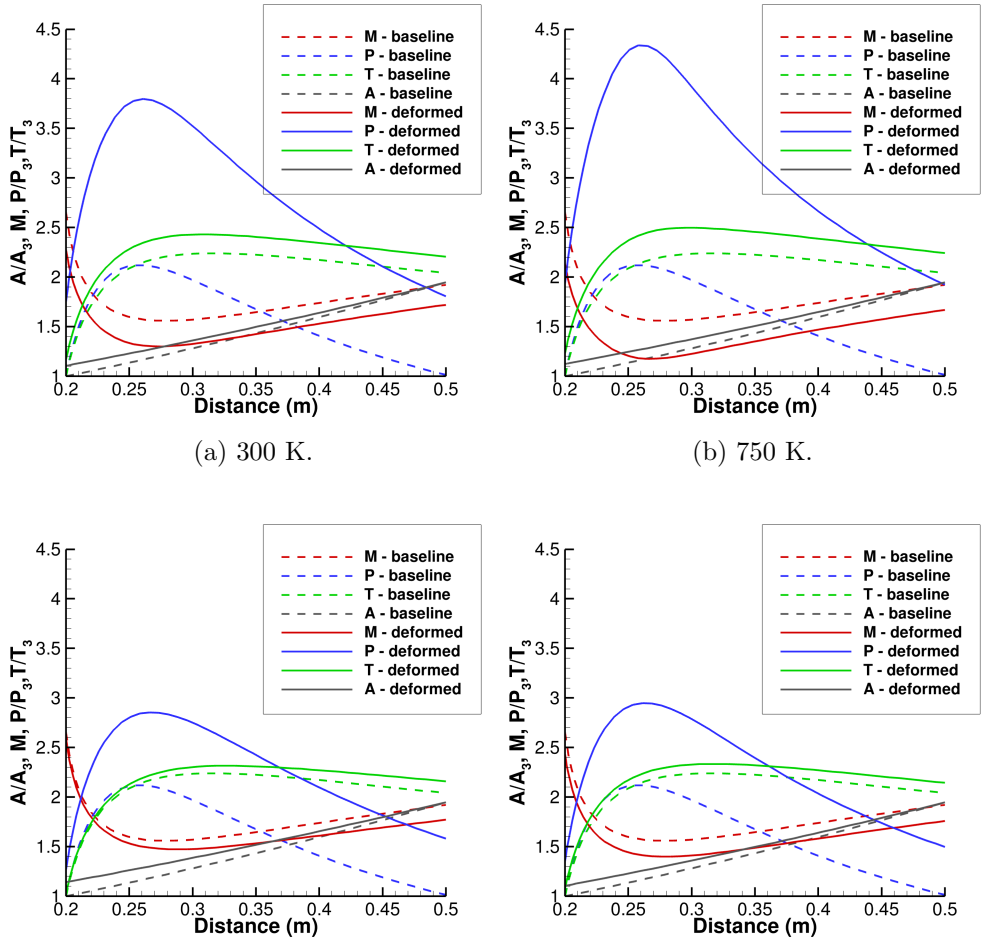

(a) 300 K.

(b) 750 K.

(c) 1200 K.

(d) 930/750 K.

Figure 3.18: Combustion chamber properties for undeformed (solid lines) and deformed (dashed lines) inlets.

Table 3.4: Engine-level quantities of interest.

	Undeformed	300 K	750 K	1200 K	930/750 K
Sa_{inlet} (N-s/kg)	-315	-407	-442	-330	-382
Sa_{burner} (N-s/kg)	358	391	411	338	366
Sa_{nozzle} (N-s/kg)	550	606	619	589	590
Sa_{area} (N-s/kg)	-81.1	-44.8	-43.2	-49.7	-56.4
Specific Thrust (N-s/kg)	512	545	545	547	517
η_0	0.527	0.512	0.507	0.523	0.496

flow decreases in Mach number within the inlet and does not increase substantially in temperature due to weaker shocks and fewer separation/reattachment shocks than the 300 K and 750 K cases.

Overall, the total specific thrust increases with increasing deformation; additionally, the efficiency for this engine is shown to decrease for the deformed geometries. As such, the engine has traded thrust for efficiency.

3.3.7 Vehicle Performance

Along with ramifications to the engine, the deformation of the inlet also affects the aerodynamic performance of the vehicle. The integrated airframe/engine structure assures that changes to the shape of the forebody/inlet affect the outer mold line of the vehicle. As this change occurs, the contribution of the forebody/inlet to the lift, drag and pitching moment is altered, which impacts the controllability and stability of the vehicle. Table 3.5 compares the original, undeformed configuration quantities of interest to those obtained with each active cooling case. Each of the aerodynamic coefficients – coefficient of drag (C_D), coefficient of lift (C_L), and the pitching moment coefficient (C_M) – are only in reference to the inlet and do not include the entire vehicle. Therefore, the change in those quantities is the relative change for the *inlet*. Also note that the pitching moment coefficient is evaluated using a center of gravity (cg) about the end of the modeled inlet at 0.63 m, shown in Figure 3.5.

Table 3.5: Vehicle performance quantities of interest.

	Undeformed	300 K	750 K	1200 K	930/750 K
C_D	0.0547	0.0798	0.0832	0.0602	0.0627
C_L	0.0537	0.00350	0.0205	0.0125	0.0337
C_M	0.0349	0.0217	0.0247	0.0253	0.0260

When the aero-thermo-elastic deformations are accounted for, a large increase in drag is observed, shown in Table 3.5. There are two contributions to drag: the pressure force and the skin friction. As the inlet opens up, the pressure force increases since the projected area of the inlet from the perspective of the flow has increased. Additionally, the skin friction increases as the body expands, exposing a larger surface. As an example, the 300 K case has a larger opening effect while the 1200 K case greater amount due to the increased temperature. Comparing the drag coefficients

of each case highlights that the pressure drag (or opening effect) of the 300 K case has a larger influence on the drag than the skin friction increase of the 1200 K case.

The majority of the changes to the lift are caused by the changing flowfield during the internal compression. Changes to the shock-shock interactions and separation impacts the inlet's contribution to lift. The combination of increased drag and decreased lift leads to a decrease in the inlet's contribution to the lift-to-drag ratio. Additionally, the pitching moment decreases when aero-thermo-elastic deformations are considered; this is in large part due to the shock from the forebody impinging on the cowl. To provide a safe and stable flight for the vehicle, the changes to the aerodynamic coefficients must be accounted for in some manner.

Overall, this study indicated the important effects of deformation for a scramjet on the inlet flowfield, inlet temperature, engine performance and vehicle performance and can be found in Refs. [58] and [51]. The current work shows an analysis for an experimental configuration already designed and tested. However, because of the important implications of FTSI, incorporating aero-thermo-elastic analysis into the design of the vehicle could help eliminate some of these issues. For example, this study included one potential avenue for mitigating the unwanted effects of deformation: thermal management. Further work such as material/structural choices and the inlet shape can also enable more robust designs. Unfortunately, a fully coupled simulation is often infeasible for an entire design space, so finding cost-effective methods for incorporating FTSI is necessary.

3.4 Coupling Approach

Powered hypersonic flight has the potential to revolutionize transport through the atmosphere. With any new vehicle, though, there are risks associated with the wide range of flight conditions the vehicle will experience, not all of which can be thoroughly vetted beforehand using ground tests. Not only are ground tests costly, but they are also limited in their ability to simulate the conditions of a particular flight regime, often trading high-enthalpy, high-speed and model size. Accurate, but computationally efficient, numerical analyses for these high-speed vehicles are necessary to better understand the flight conditions that have not been previously tested, reducing the

risk associated with a new vehicle. Often, as the physics becomes more complicated and coupled, accuracy and computational expense for numerical analyses work in opposition.

Fluid-thermal-structural interactions for a scramjet inlet is one instance of this opposition. A tightly coupled aero-thermo-elastic analysis is needed to accurately model the interactions between the flowfield and the vehicle, but often the computational expense is too high of a barrier, especially for a comprehensive analysis of the entire flight envelope. A high-fidelity aerothermodynamic analysis is necessary to capture flowfield changes, including shock-shock interactions between the vehicle leading edge and cowl, separation zones along expansion surfaces, and shock-boundary layer interactions within the inlet. CFD is used to capture these effects and FEM is used to predict the deformations, both of which can be computationally expensive analyses, even individually. Finding an efficient and accurate coupling strategy is one way to reduce the computational expense of these types of analyses.

As mentioned in Chapter 1, there is uncertainty about which of the couplings between the physics can be considered weak. Initial studies indicated that the connection between the deformation and heat flux was minimal enough that it could be ignored, decoupling an aero-thermo-elastic problem into an aero-thermal and an aero-elastic problem to be solved in series [86]. More recent work showed that the strength of this connection was highly dependent on the amount of deformation and in many cases could not be ignored [25]. The exact demarcation of when the deformation influences the heat flux "enough" is unknown and will therefore be explored in this work. The current work aims to address key factors for a hypersonic inlet, discovering where that "insufficient" deflection limit is and when that assumption breaks down.

In regards to scramjets, Ye et al. [127] have focused on the aero-thermo-elastic deformation of a scramjet inlet and have begun initial work on the accuracy and efficiency of various coupling strategies at a single trajectory point. Ye et al. included three coupling strategies: neglect heating, one-way coupling and two-way coupling. However, the deformed flowfield in Ye et al. [127] displayed only minor changes in the flow and temperature distribution due to the minimal deformation, unlike other research such as Guangyue et al. [42] and Horing et al.[51], where larger deformations caused

stronger shock-shock interactions and separation. The results from Ye et al. [127] suggest that for the one specific flight condition and geometry analyzed, the two-way coupling method induced the greatest deformation, but the one-way coupling showed similar results, indicating an insufficient amount of deformation to change the temperature distribution substantially for their case. Other works, such as Kline et al. [67], used a *sandwich* approach to reduce computational expense, which applies two CFD computations around a single FEM structural computation. Although a much faster overall computation, the accuracy of using such a coupling method needs to be addressed. The present work focuses on exploring key factors that influence the coupling strategy necessary for the aero-thermo-elastic response of a scramjet inlet along a certain flight path. Three different strategies – sandwich coupling, one-way coupling and two-way coupling – are employed to assess accuracy and efficiency of each analysis method.

3.4.1 Test Setup

Some previous work in aero-thermo-elastic analysis points to the conclusion that the aerothermal coupling can be simplified to a one-way scheme (where the heat-flux is *not* affected by the structural deformation) without causing too large of overestimations for certain cases such as a flat plate [25] or certain flight conditions for a scramjet inlet [127]. Therefore, a one-way aerothermal-aeroelastic analysis is also performed for this work. The diagram representing the flow of analyses is shown in Figure 3.19. Unlike the previous two-way coupling case, two separate iterative loops are shown with the dashed boxes. First, the aerothermal coupling between just the aerothermodynamics and the thermal response is converged for the undeformed body. Then the aeroelastic coupling between the aerodynamic and elastic response is converged. The elastic response uses the vehicle temperature profile found from the aerothermal coupling to calculate the thermal strains. Therefore, the inaccuracy of this analysis depends on the significance that the deformations have on the heating and temperature profile of the inlet surfaces.

The final coupling strategy employed is the sandwich method. The purpose of comparing this strategy to the previous two is to quantify the inaccuracy of the method for future use, since the

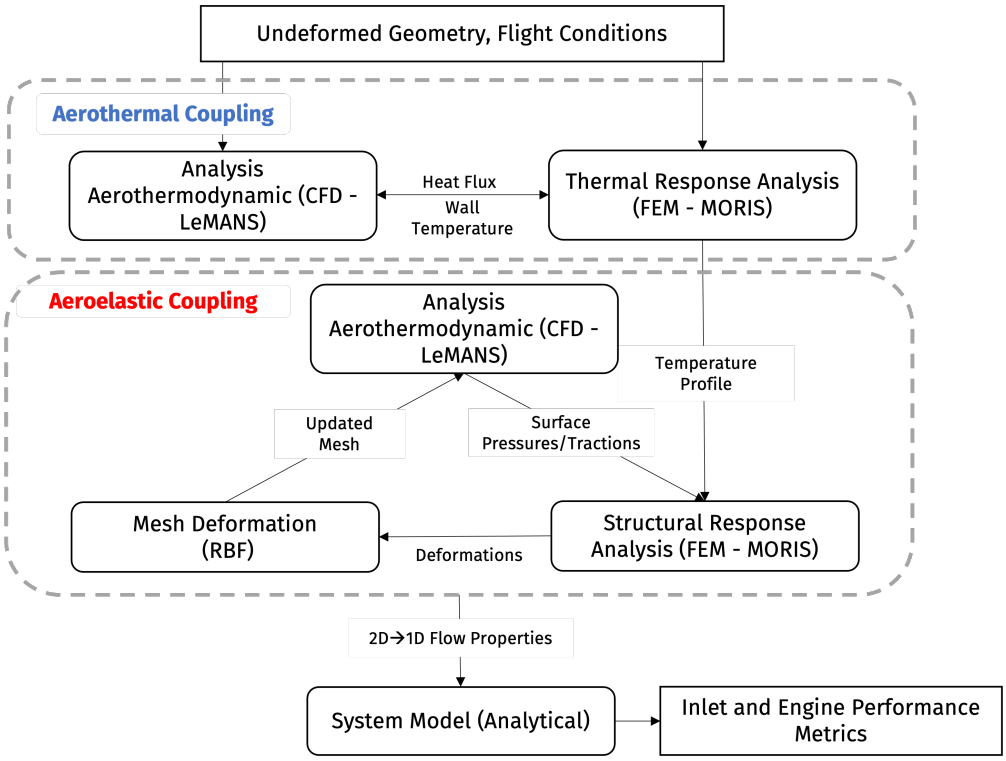

Figure 3.19: Aero-thermo-elastic one-way coupling scheme.

computational efficiency of a sandwich coupling is appealing. A diagram of the sandwich coupling process is shown in Figure 3.20, where the solid grey boxes now represent a single evaluation of a discipline, rather than an iterative process. The analysis begins with an aerothermodynamic evaluation of the undeformed geometry using a radiative-equilibrium boundary condition along the wall. The heat flux and tractions along the wall are then fed into a staggered thermo-elastic analysis, which computes the temperature profile and deformation of the structure. The deformations then inform the fluid mesh adaptation. Finally, using that deformed fluid mesh and the wall temperatures computed from the thermal response, a final CFD analysis is evaluated. This method is typically employed by researchers knowing that accuracy has been traded for much reduced computational expense [67, 92].

For all cases, once the final converged aerothermodynamic results are found, the fluid states from the exit of the isolator are fed into the system model to calculate the inlet and engine performance quantities of interest. Each coupling strategy is then compared for accuracy using these

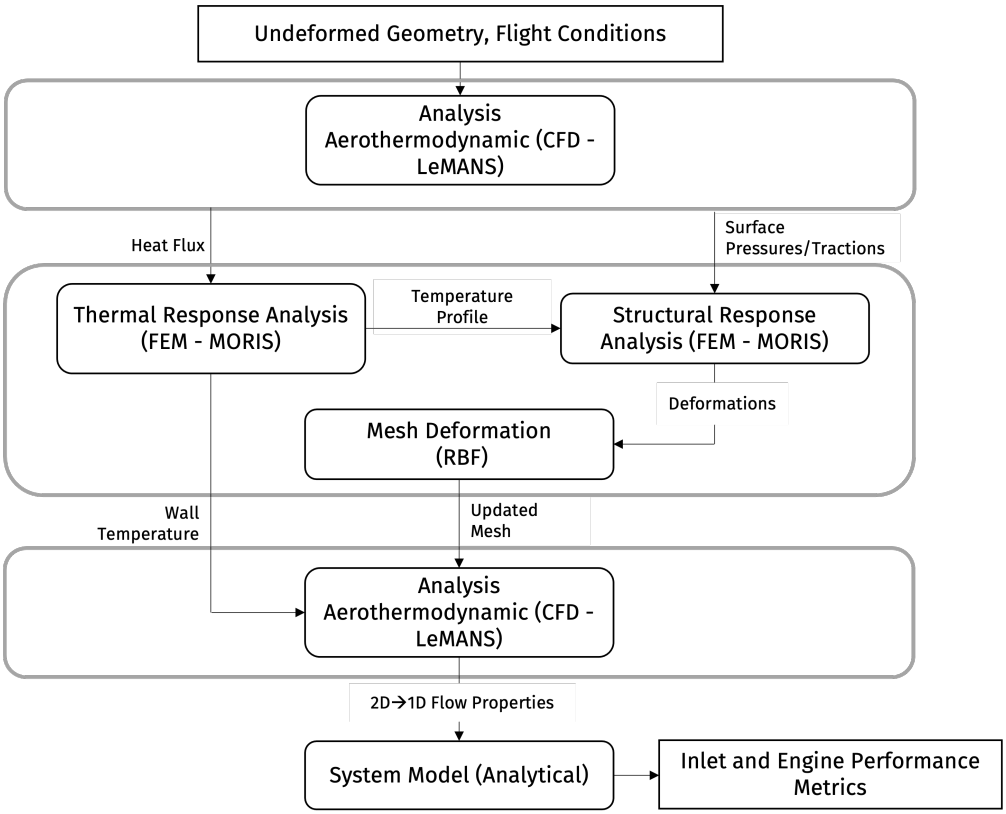

Figure 3.20: Aero-thermo-elastic sandwich coupling scheme.

metrics for the various conditions tested.

The geometry chosen is representative of the inlet model GK-01, which was tested at the DLR Windtunnel in Cologne [43, 48]. The campaign for this geometry has ample data to validate against, and many previous aero-thermo-elastic analyses have also used the GK-01 inlet [127, 42]. The two-dimensional planar geometry is shown in Figure 3.21. The convergence studies for the fluid and structural meshes are described in Section 3.3.1. Additionally, a validation of the computational results compared to experimental results for the undeformed configuration is also found in 3.3.1 and showed good agreement between the two. The root-mean-square errors for the coefficient of pressure for the ramp and cowl were 4% and 1%, respectively. Therefore, the same two meshes are used in this work.

As seen in Figure 3.21, the internal structure for the current analysis is kept simple. The leading edge is built out, but then a panel-like profile is applied to the rest of the inlet. Two materials

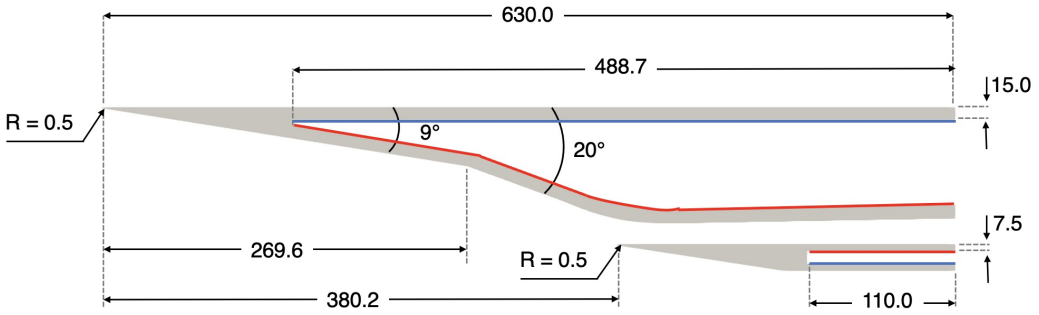

Figure 3.21: Schematic of GK-01 model for computation (units: mm).

are chosen for the coupling comparisons: Inconel 718 and TZM-C03 (a molybdenum alloy). The properties for both materials vary with temperature as found in Refs. [47], [3] and [65]; the properties necessary for the current analysis include the thermal conductivity, thermal expansion, Poisson's ratio and Young's modulus. The temperature-dependent material properties of the two materials are shown in Figure 3.22. There are several key differences between the properties of the two materials: the magnitudes of the thermal expansion, the temperature dependence of the thermal conductivity and the steep drop off of Young's Modulus for Inconel. These differences influence the resultant deformation of the inlet and the decision for the coupling strategy that will be shown in Section 3.4.2.

A Dirichlet temperature boundary condition is set on the internal surfaces to represent the thermal management of the vehicle. The TZM-C03 inlet serves as the baseline, so the internal temperatures are chosen to minimize the deformation perpendicular to the flow for the TZM-C03, as suggested by Horing et al. [58] or Section 3.3, in order to reduce unwanted flowfield behavior. Typically, the structural loads cause the inlet leading edges to bend outward. Meanwhile, the unequal aerodynamic heating on the top and bottom surfaces of the forebody and cowl results in an unequal temperature profile and amount of expansion between the surfaces. By properly choosing internal boundary conditions, the thermal expansion can be used to counteract deformation due to

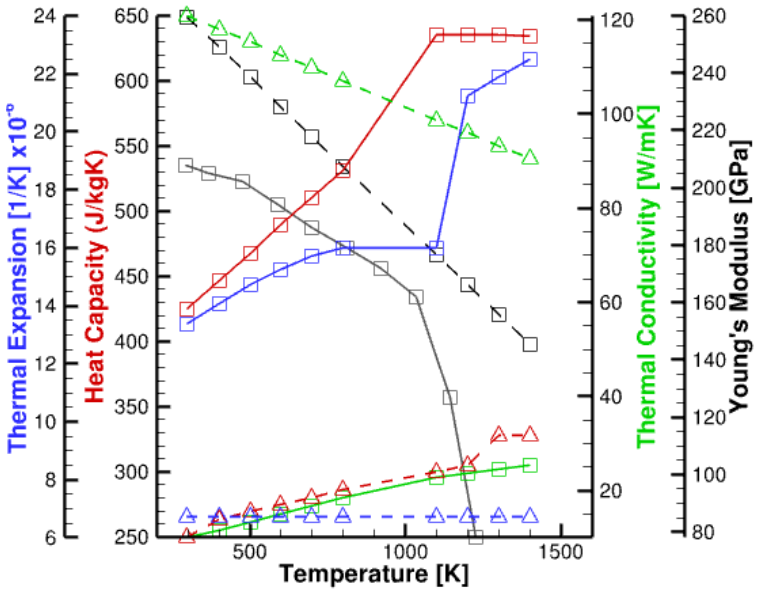

Figure 3.22: Temperature dependent material properties of Inconel 718 (solid lines) and TZM-C03 (dashed lines).

structural loading. To do so, temperatures of 1050 K and 1000 K are applied to the upper and lower internal forebody walls, respectively, indicated in Figure 3.21 by the red (1050 K) and blue (1000 K) lines. Similarly, the 1050 K and 1000 K boundary conditions are applied to the lower and upper internal surfaces of the cowl, respectively. The same boundary condition scheme is then applied to the Inconel case as well. Previous work has shown that the inlet deformations and resultant flowfield are sensitive to the internal temperature [58, 67]. Therefore, it is important to determine the effect of this flowfield change on the coupling strategy. To do so, the two-temperature case is compared against additional one-temperature cases: 1000 K and 1050 K. The one-temperature cases refer to the entire interior surface set to a single temperature, either 1000 K or 1050 K. These three internal temperature conditions are used to evaluate their influence on the coupling strategy.

Finally, multiple flight conditions are chosen to represent different stages of a typical powered hypersonic flight envelope. Scramjet flight envelopes – typically denoted as a flight corridor – are bounded by two isodynamic pressures: 0.5 atm and 1.0 atm [46, 74]. Below a dynamic pressure

of 0.5 atm, the engine is unable to function efficiently; above a dynamic pressure of 1.0 atm, the vehicle exceeds its structural limit that it can withstand due to harsh thermal and mechanical environments. In order to understand the effect of dynamic pressure on the strength of the coupling phenomena, both limits are evaluated. The flight corridor also spans potential Mach numbers. The GK-01 engine is originally designed for Mach 7 conditions; therefore, Mach 7 is chosen to be the cruise condition and the baseline for the comparisons. Additionally, a slower speed of Mach 6.5 and a higher speed of Mach 7.5 are evaluated to understand the influence of Mach number on the heating and deformation. Finally, both positive and negative angles of attack are modeled to represent ascent and descent, respectively. The three parameters – dynamic pressure, Mach number and angle of attack – are used to evaluate the key influential factors for choosing a coupling strategy. All flight conditions are shown in Table 3.6.

Table 3.6: Freestream flight conditions.

Trajectory Point		Mach	Altitude (km)	Dynamic Pressure (atm)	Velocity (m/s)	Angle of Attack
1	baseline	7.00	25.8	0.75	2090	0.00°
2	M	6.50	24.9	0.75	1940	0.00°
3		7.50	26.7	0.75	2250	0.00°
4	α	7.00	25.8	0.75	2090	1.00°
5		7.00	25.8	0.75	2090	-1.00°
6	q_∞	7.00	28.5	0.50	2110	0.00°
7		7.00	23.9	1.00	2080	0.00°

3.4.2 Results

The fluid-thermal-structural modeling procedure is evaluated for a 2D scramjet inlet under various conditions: material choice, thermal management, and flight conditions. The FTSI coupling strategy is evaluated both for accuracy and for computational efficiency in order to quantify benefits and drawbacks of the three approaches. The accuracy is evaluated using physical quantities of interest, such as temperature and deformation, and vehicle-level quantities of interest, such as the coefficient of drag and the specific impulse of the engine. The two-way coupling approach is

considered the most accurate, and the differences produced by the other approaches are considered to be errors. The efficiency is evaluated using the number of iterations and CPU wall time for running the three strategies to convergence.

3.4.2.1 Accuracy

As mentioned previously, TZM-C03 is used as the material for the baseline case. The influence of material choice is evaluated by comparing to another metallic option: Inconel 718. Under the same flowfield conditions, the two materials deform in different manners, as seen in Figure 3.23. From the current perspective, the undeformed case is indistinguishable from the TZM case. The Inconel shows additional expansion upstream and deflects outwards, perpendicular to the flow. The cause of this exacerbated distortion is the difference in material properties shown in Figure 3.22. Two major distinguishing factors are present: the magnitude of the thermal expansion and the steep drop-off in the Young's Modulus for Inconel at high temperatures. Both factors cause additional thermal and mechanical deformations for the Inconel inlet.

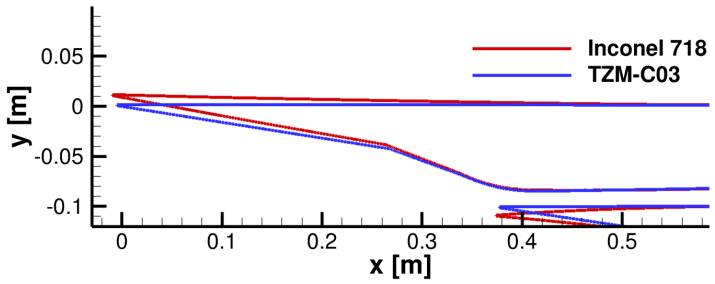

Figure 3.23: Two-way coupling deformation for Inconel 718 and TZM-C03.

Focusing on the forebody leading edge, the influence of the three different coupling strategies is seen for both materials. Figure 3.24a shows the resulting deformations for the baseline case using TZM-C03, applying each different coupling scheme. The limited deformation for the sandwich coupling case in the y-direction is evident; the two other coupling cases show additional deformation (approximately 1 mm). Both the two-way and one-way coupling methods produce very similar

resulting deformations, indicating that the aerothermal coupling is not significantly influenced by the deformation in this case. Meanwhile, the Inconel 718 case shows significant differences in the deformation for all three coupling strategies in Figure 3.24b. Again, the sandwich method produces the least amount of deformation for these cases, underestimating the values seen with the iterative FTSI processes. Also, there is a significant difference between the one-way and two-way coupling methods for the deformation in the y-direction. This variance indicates that the deformation causes a significant change in the temperature profile of the flowfield.

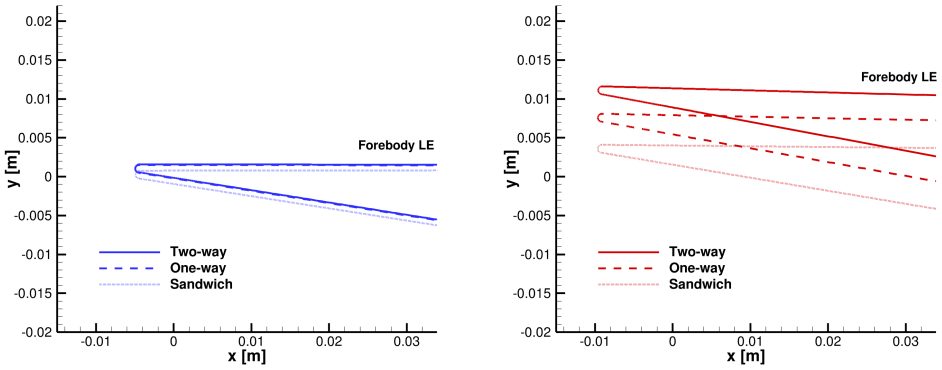

(a) TZM-C03 (b) Inconel 718

Figure 3.24: Deformation comparisons for the forebody leading edge for TZM-C03 and Inconel 718.

The difference in the temperature and heat flux experienced by the vehicle is due to the changes to the flow structures when deformation is introduced. Figures 3.25a-3.25c show minor changes to the flowfield around the cowl leading edge for the three cases involving TZM-C03, where the deformation is small. By comparison, Figures 3.25d-3.25f show drastic changes to the flowfield once larger deformations are introduced with Inconel 718. For the smallest deformation case, sandwich coupling, the flowfield already appears significantly altered from those seen for TZM-C03. The separation bubble along the top surface of the inlet has grown and extended upstream onto the expansion surface. Additionally, two new separation bubbles appear along the bottom surface of the inlet, caused by stronger shock-boundary layer interactions. The separation bubbles cause

additional spikes in heating to the surface of the vehicle. When more deformation is introduced with the two-way and one-way coupling schemes, seen in Figures 3.25d and 3.25e, respectively, the shock from the second forebody ramp impinges on the cowl. The shock impingement also creates an increase to the heating on the surface of the cowl. These changes to the flow structure lead to a significant enough dependence of the temperature profile on the deformation, indicating that the assumptions for the one-way coupling method no longer hold. The localized heating from the resulting flow structures contributes to an overall increase in the temperature of the inlet, specifically on the panels internal to the inlet. The increase in temperature then results in additional expansion of those surfaces, which contributes to displacement in both the x- and y-directions [58].

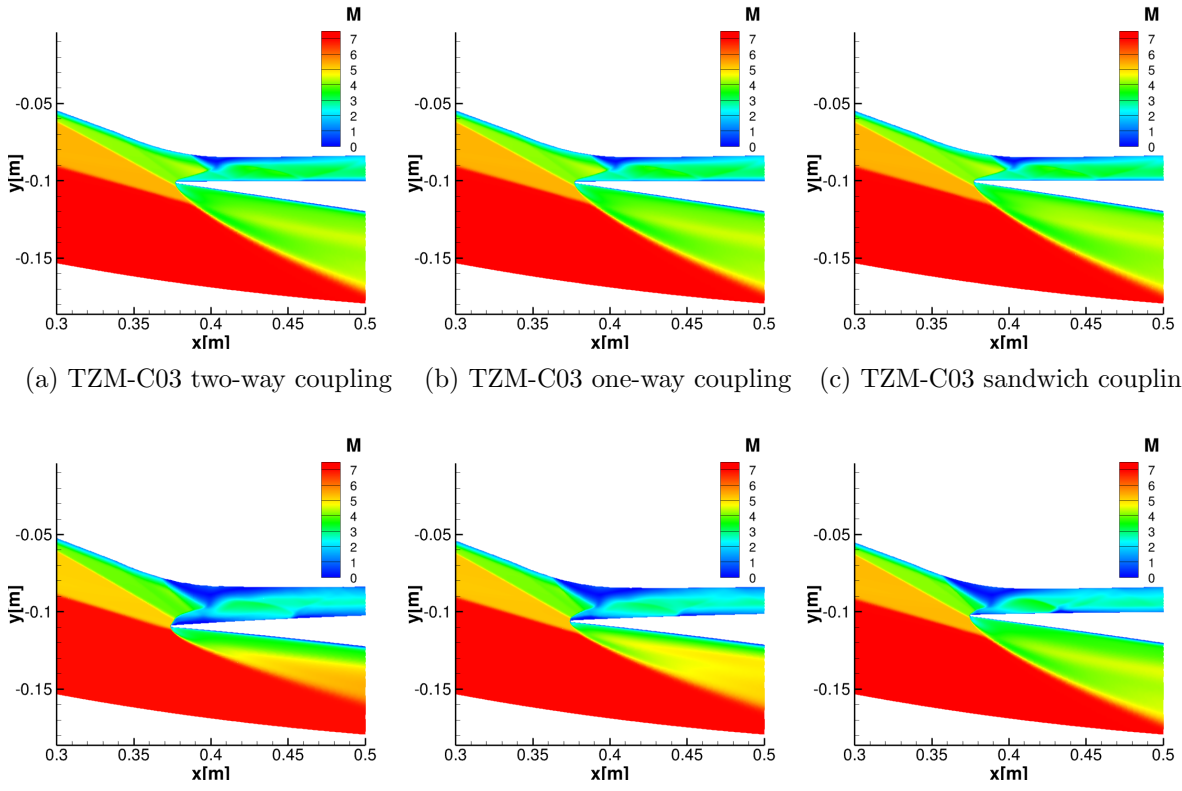

(a) TZM-C03 two-way coupling (b) TZM-C03 one-way coupling (c) TZM-C03 sandwich coupling

(d) Inconel 718 two-way coupling (e) Inconel 718 one-way coupling (f) Inconel 718 sandwich coupling

Figure 3.25: Mach number contours for different materials and coupling schemes.

The differences in the deformation and resultant flowfield not only impact the inlet, but also the vehicle and engine due to the integrated nature of these high-speed powered vehicles. An

example of a vehicle-level quantity of interest is shown in Figure 3.26a, which compares the coupling strategy analyses for the two materials with the undeformed vehicle analysis for the coefficient of drag. Across all conditions, the coefficient of drag has increased. Note that the TZM-C03 geometry shows a much smaller increase in the drag coefficient for all three coupling cases. The two-way and one-way coupling schemes produce very similar drag coefficients, due to the similar deformations seen in Figure 3.24a. Meanwhile, the sandwich method produces a slightly smaller coefficient of drag, more similar to the original geometry, since the deformation is more minor. For the Inconel cases, a much larger increase in drag coefficient is observed. The increase in drag scales with the increase in deformation in the y-direction across these three cases. The breakdown of the assumptions for the one-way coupling scheme is observed again in underestimation of the drag coefficient once the vehicle deforms.

The limits of one-way coupling are also apparent for the specific impulse of the vehicle, shown in Figure 3.26b. There is a larger difference between the two-way and one-way coupling cases for the Inconel (2.9%) than for the TZM-C03 (0.3%). Similarly, the sandwich method shows a larger error from the two-way coupling case for Inconel (7.5%) rather than for TZM-C03 (4.5%). Once again, there is a consistent trend that the two-way coupling method results in the highest specific impulse, followed by the one-way coupling and finally the sandwich method. This trend reflects the deformation results, observed in Figure 3.24. Despite the apparent increase in specific impulse with increased deformation, the efficiency of the inlet decreases. The specific impulse for the Inconel two-way coupling case is 3.9% higher than the TZM-C03, but the total engine efficiency has decreased by 6.0%.

The next factor influencing the decision for the FTSI coupling strategy is the thermal management of the vehicle and its imposed boundary condition on the analysis. Since these vehicles are prone to thermal expansion at extreme temperatures, the thermal management impacts the temperature field and expansion of the inlet. Figure 3.27 compares the three internal temperature boundary conditions for the two-way coupling case. The baseline case, 1050/1000 K, shows the smallest displacement in the y-direction for the forebody and cowl leading edges, indicating that

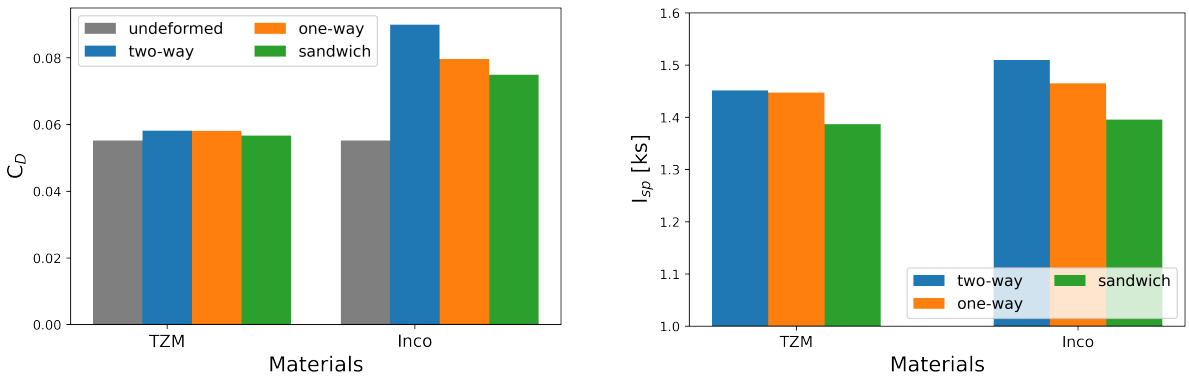

(a) C_D (b) I_{SP}

Figure 3.26: Vehicle-level property comparisons for the coupling methods for the two different materials.

the imposed boundary condition is successfully reducing the deformation, as intended. Meanwhile, both one-temperature boundary condition choices show additional displacements in the y-direction for both the cowl and forebody with only a 50 K difference on half the surfaces. By comparing the 1050 K case to the 1000 K case, it is observed that the 1050 K case expands upstream into the flow further due to the higher prescribed wall temperatures promoting thermal expansion. As mentioned previously, the structural loading forces the inlet to open. The two-temperature boundary condition exploits the unequal heating on the top and bottom surfaces of both the forebody and cowl to expand the inlet back towards the original position. The one-temperature boundary conditions allow a difference in thermal expansion between the top and bottom surfaces, which adds to the opening behavior [58].

Despite the difference in deformation, each case shows similar trends when comparing the three coupling schemes. In Figure 3.28, the temperature profile along the forebody compression surface for each case is shown above the deformed forebody leading edge geometry. The temperature profiles each show the two-way and one-way coupling methods converging to a similar result, with root-mean-squared errors of 0.37%, 0.59% and 0.52% for the 1050/1000 K, 1000 K and 1050 K cases, respectively. Overall, these temperature differences are relatively minor, indicating that the temperature profile within the flowfield has not been substantially altered by structural defor-

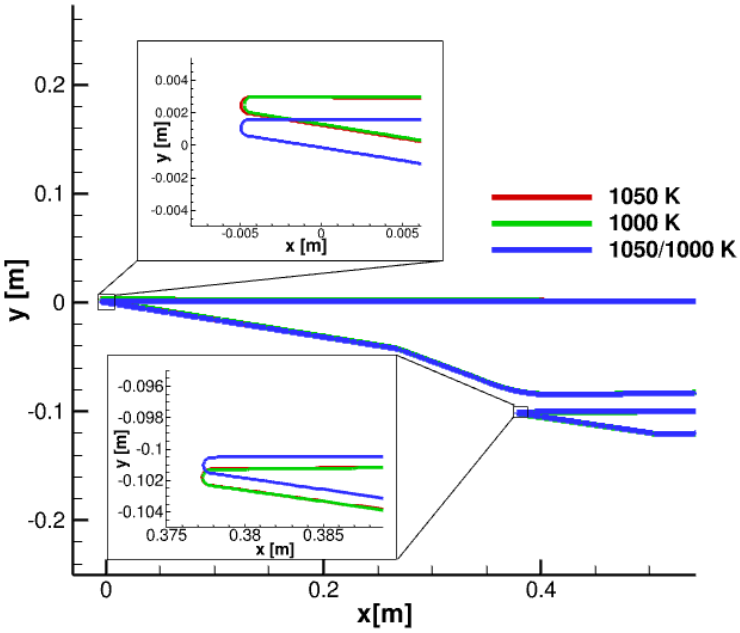

Figure 3.27: Two-way coupling deformation comparisons for different internal prescribed wall temperatures.

mations, and the one-way coupling method assumption is valid. Meanwhile, the sandwich method underpredicts the wall temperature by root-mean-squared errors of 4.12%, 4.43% and 3.95% for the 1050/1000 K, 1000 K and 1050 K cases, respectively. This underprediction is caused by the first CFD iteration on the undeformed geometry using a radiative-equilibrium boundary condition along the surfaces of the walls. That assumption overestimates the temperature and underestimates the heat flux into the vehicle, so when the thermo-structural computation is performed for the sandwich method, it uses this lower heat flux and results in a lower overall temperature. Less expansion occurs for this lower structural temperature, and therefore the sandwich method produces the least amount of deformation and resulting change to the vehicle-level quantities for these cases.

The difference in the three coupling methods can also be evaluated by comparing the resulting vehicle-level properties for the different boundary conditions. Figure 3.29a shows that the drag has increased for all conditions from the undeformed case. The drag coefficient for the 1000 K case and 1050 K case have increased by 10.8% and 10.3%, respectively, compared to the undeformed

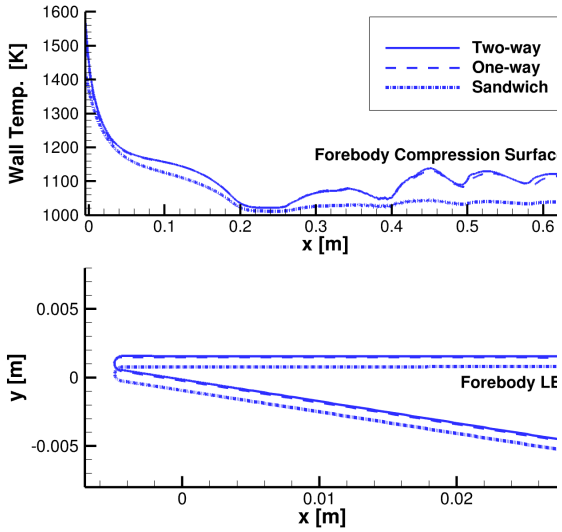

(a) 1050/1000 K

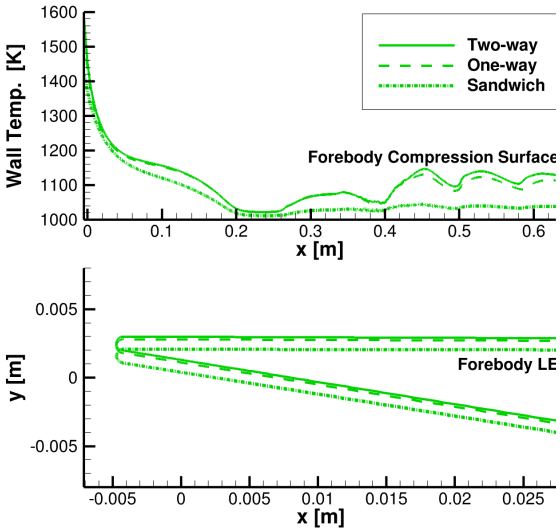

(b) 1000 K

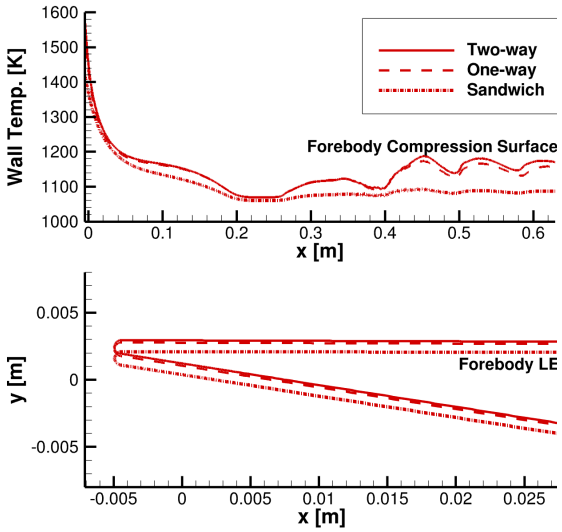

(c) 1050 K

Figure 3.28: Forebody wall temperature and position for the internal boundary condition cases.

case. Meanwhile, the two-temperature boundary condition increased only by 5% from the original undeformed geometry. Additionally, a trend appears such that the two-way coupling method produces the highest drag, followed closely by one-way and then sandwich coupling. The reverse trend is observed in Figure 3.29b for the lift-to-drag ratio, which shows the two-way coupling as the smallest ratio and the sandwich coupling as the largest. Once again the sandwich method shows the closest agreement with the undeformed method, indicating that using the sandwich method will underestimate deformation and its impact if a radiative-equilibrium boundary condition is used for the first CFD analysis. Also, similar to the drag coefficient case, the one-temperature boundary condition cases cause approximately double the change in lift-to-drag ratio than the two-temperature boundary condition case.

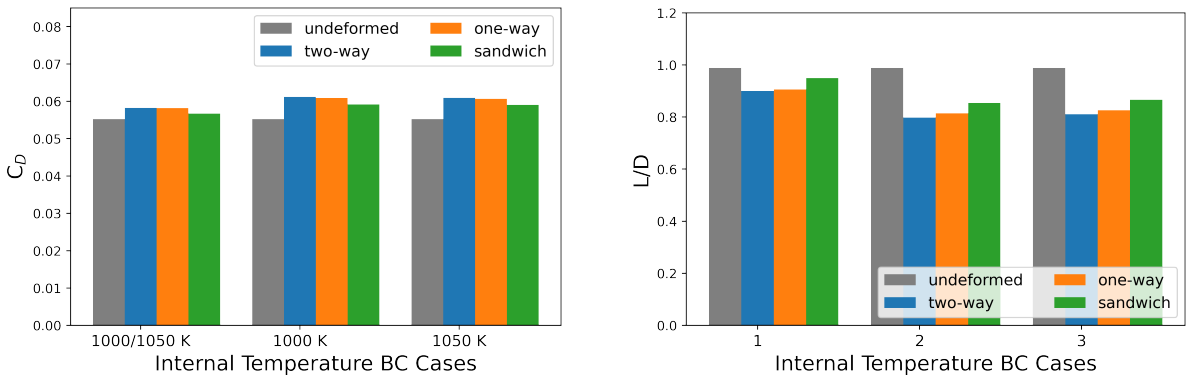

(a) C_D (b) L/D

Figure 3.29: Vehicle-level property comparisons for the coupling methods for each internal boundary condition case.

Finally, the coupling methods are compared across different freestream conditions, varying the Mach number, dynamic pressure and angle of attack. For reference, the trajectory points evaluated are defined in Table 3.6. In order to condense the comparison of physical changes for each of these cases, the angle of the first ramp on the forebody, the deformation of the cowl in the x- and y-directions and the maximum temperature along the body are collected in Figure 3.30. To start, the angle of the first ramp along the forebody's compression surface is originally designed to be 9°, but Figure 3.30a shows that FTSI influences the resultant ramp angle. Narrowing in on trajectory

point (TP) 1, the sandwich method is shown to produce an angle below the original geometry. This occurs because we have limited the deflection upwards in the direction perpendicular to the flow, but the vehicle has still expanded in the negative x-direction, which results in a reduced angle. Conversely, both the two-way and one-way methods produce an angle more similar to the original geometry because they deflect upwards as they expand upstream, as seen in Figure 3.24a. Similar trends are seen across the trajectory points, where the two-way coupling produces the largest ramp angle, followed by the one-way and then the sandwich method. Similarly, the deformation of the cowl leading edge in the x-direction follows the same trend with the two-way coupling resulting in the largest deformation. In this case, the deformation of the cowl in the x-direction is predominantly dependent on the thermal loads and thermal expansion of the vehicle.

Meanwhile, the deformation of the cowl in the y-direction is dependent on a combination of both the thermal and structural loading [58]. Therefore, the y-deformation does not follow the same trend of the two-way coupling producing the most amount of deformation for all trajectory points. As observed for TP6 in Figure 3.30c, the sandwich coupling produces the largest deformation. However, the two-way coupling does produce the most deformation in the negative y-direction, which opens up the inlet. The baseline case, TP1, is chosen to reduce the y-deformation for the sandwich case, which is seen in Figure 3.30c. Since both thermal and mechanical loads are applied to the walls, the internal boundary conditions temperatures are chosen to balance those two sets of loads. Meanwhile, for TP6, the mechanical loads (i.e. pressure and shear) are lower at a lower dynamic pressure. This adjustment to the loads causes the balance between the thermal expansion and the mechanical deflections to now favor the cowl bending into the inlet, indicated by a positive deformation, for the sandwich case. Even the two-way and one-way methods show a major reduction in deformation in the negative y-direction. Regarding the differences between the methods, the sandwich method is shown to have much larger error at higher Mach numbers (TP3) and higher dynamic pressures (TP7). Additionally, a negative angle of attack produces a substantial difference between the sandwich method and the two others.

Figure 3.30d juxtaposes the maximum temperature along the wall for the various cases. Over-

all, the one-way and sandwich methods produce similar results, indicating that the one evaluation of the thermal analysis produces a maximum temperature very similar to the iterative aerothermal evaluations on the undeformed geometry. On the other hand, the two-way coupling scheme consistently produces the highest maximum temperatures for all trajectory cases. Also, when comparing the two-way coupling method results to each other, all the cases at Mach 7 (TP1 and TP4-7) produce a similar maximum temperature. The major difference in temperature occurs when the Mach number is adjusted: the lower Mach number produces a lower temperature and a higher Mach number produces a higher temperature. This is to be expected but indicates that the thermal loads are highly dependent on the Mach number.

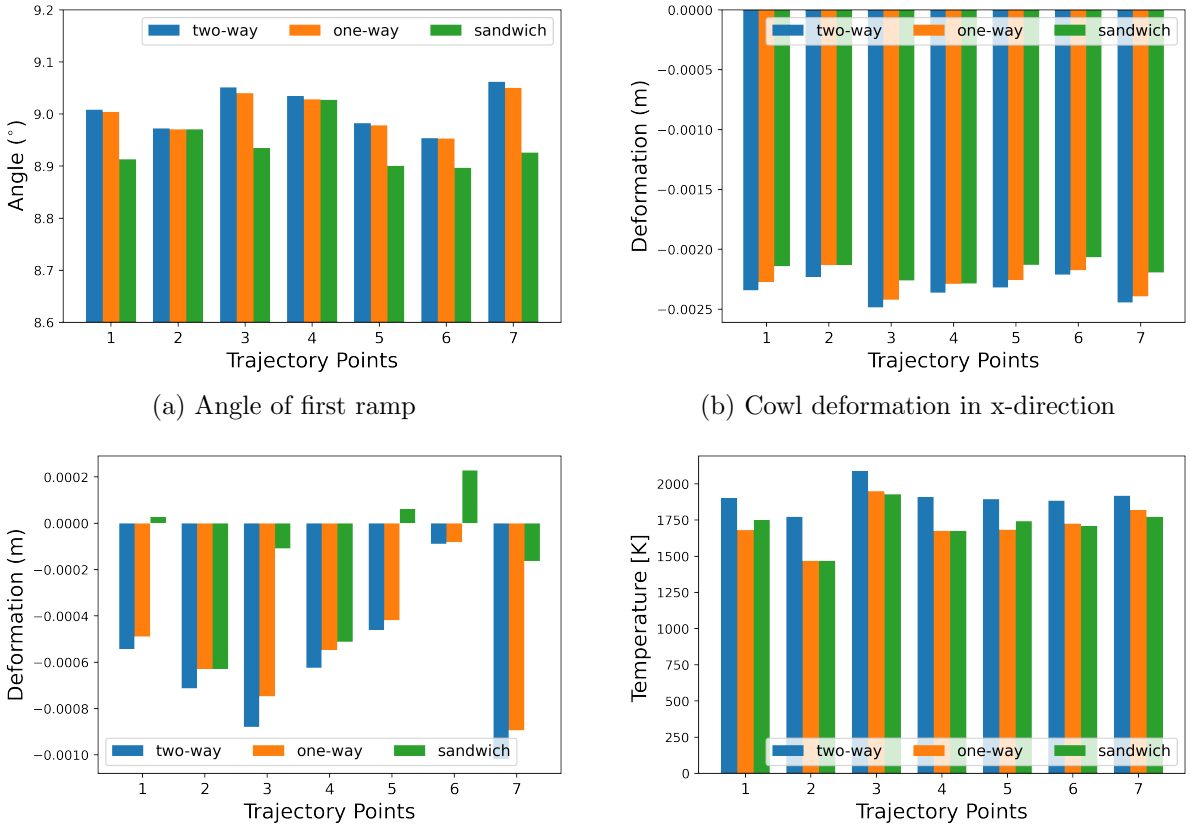

(a) Angle of first ramp

(b) Cowl deformation in x-direction

(c) Cowl deformation in y-direction

(d) Maximum temperature along the body

Figure 3.30: Physical property comparisons for the coupling methods at each trajectory point.

Since these high-speed vehicles are tightly integrated, the vehicle-level properties are depen-

dent on the shape change of the inlet, represented by the physical properties in Figure 3.30. To start, the aerodynamic coefficients are compared in Figures 3.31a-3.31c for each coupling method and each trajectory point. When considering the drag coefficient, the error is below 1% for all the trajectory points using the one-way coupling method. The error in drag coefficient for the sandwich coupling method is below 5% for all the trajectory points. The largest error for the sandwich coupling occurs at TP7, the highest dynamic pressure case, while the largest errors for the one-way coupling occur for TP2 and TP3, the varying Mach number cases. When considering the lift-to-drag ratio, the error is below 6% for the one-way coupling and 17% for sandwich coupling. The largest errors for both cases occur at the highest Mach number (TP3), followed by the highest dynamic pressure (TP7). Finally, when considering the pitching moment coefficient, the error is below 3% for the one-way coupling and 9% for the sandwich coupling. Once again, the largest errors for both cases occur at the highest Mach number (TP3), followed by the highest dynamic pressure (TP7).

In addition, the engine-level quantities are impacted by the deformation and temperature profile of the vehicle. The specific impulse is used as a representative quantity for the effects on the engine and is shown in Figure 3.31d. The error for the specific impulse across all the trajectory points is below 0.5% for the one-way coupling method and 3% for the sandwich method. The largest errors for both the one-way coupling and sandwich methods occur at the highest Mach number case (TP3), followed by the highest dynamic pressure case (TP7). Overall, for both aerodynamic and propulsive coefficients, the one-way method shows more agreement with the two-way coupling method than the sandwich method. Additionally, the largest disagreement for most of the quantities of interest occurs at higher Mach numbers and dynamic pressures.

3.4.2.2 Computational Efficiency

Although accuracy of the model is crucial for simulating the effects of FTSI, the efficiency and computational expense also influence the choice in coupling methods. To begin to gauge the computational cost of each method, the number of iterations for each analysis is compared in Table

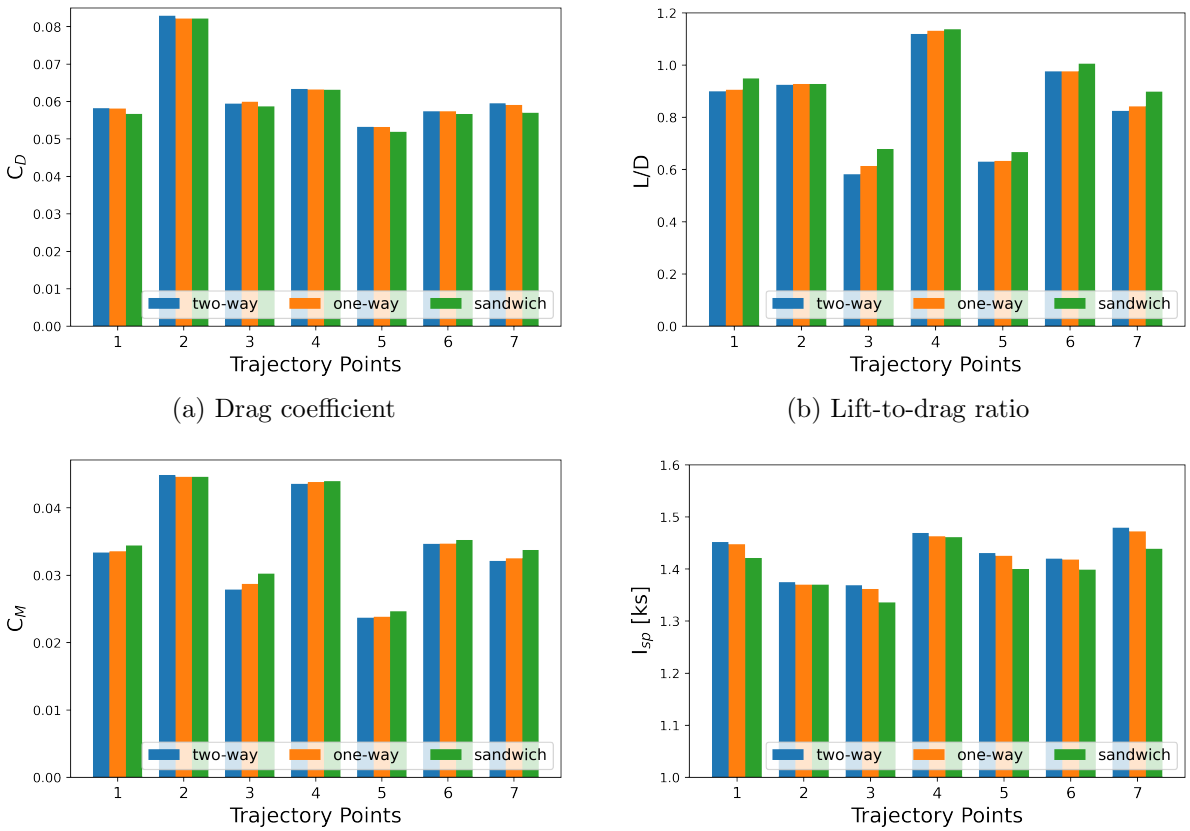

(a) Drag coefficient

(b) Lift-to-drag ratio

(c) Pitching moment coefficient

(d) Specific impulse

Figure 3.31: Vehicle-level property comparisons for the coupling methods at each trajectory point.

3.7. The two-way column refers to the aero-thermo-elastic iterative loop shown in the dashed box in Figure 2.2. The first one-way column refers to the aerothermal iterative loop shown as the top dashed box in Figure 3.19. Finally, the second one-way column refers to the aero-elastic iterative loop shown as the bottom dashed box in Figure 3.19. The sandwich method is not shown in Table 3.7 since there are consistently only two evaluations of the CFD routine and one evaluation of the FEM routine.

To assess convergence, both the absolute L2 norm residual of the temperature field of the structure and the displacements are reduced below 1×10^{-3} and 1×10^{-8}, respectively, which reduces the variation in propulsive and aerodynamic coefficients by less than 0.1%. The two-way coupling method checks for convergence of both, the aerothermal process of the one-way coupling

checks for convergence of the temperature and the aero-elastic process checks for convergence of the deformation. Table 3.7 shows the number of iterations necessary to achieve that convergence. Both the aero-thermo-elastic analysis and the aerothermal analysis need approximately 10 iterations. The exception for this is the Inconel case, where there is more deformation occurring. Meanwhile, the one-way case needs an additional two iterations for the aeroelastic analysis to converge. The number of iterations for an aeroelastic analysis is relatively small, indicating that the aerothermal analysis is the limiting factor for convergence.

Table 3.7: Number of iterations for each process.

Cases:	**Number of Iterations**		
	Two-way (Aero-thermo-elastic)	One-way (Aerothermal)	One-way (Aero-elastic)
Baseline/TP1	9	9	2
Inconel	12	9	3
1000 K BC	9	10	2
1050 K BC	9	9	2
TP2	10	9	2
TP3	10	9	2
TP4	9	10	2
TP5	9	9	2
TP6	9	9	2
TP7	10	9	2

The results of Table 3.7 may indicate that the two-way coupling needs fewer overall iterations than the one-way coupling, but the computation cost for each type of iteration differs. Figure 3.32 compares the CPU wall clock times for an example case. Generally, the aero-thermo-elastic cases use more computational time than the aerothermal cases. The aeroelastic cases are on a similar timescale to the aero-thermo-elastic cases, which indicates that considering structural deformations is more costly than the changing wall temperature for the CFD analysis. Note that the CFD evaluations are much more computationally expensive than the FEM evaluations and adapting the fluid mesh, so the difference in CPU time is predominantly determined by the CFD. Overall, for this example, the total computational cost is approximately 3200 CPU-hours for the two-way method and approximately 2600 CPU-hours for the one-way method. For the example shown in Figure

3.32, the sandwich method results in a total computational time of roughly 1000 CPU-hours, which is approximately shown by the first two iterations of the aero-thermo-elastic results. Therefore, the sandwich method is significantly more computationally efficient than either the two-way or one-way coupling approaches. Note that additional work is needed to determine the optimal value for the under-relaxation factor, so there is a possibility of reducing the computational cost of the two-way and one-way coupling methods.

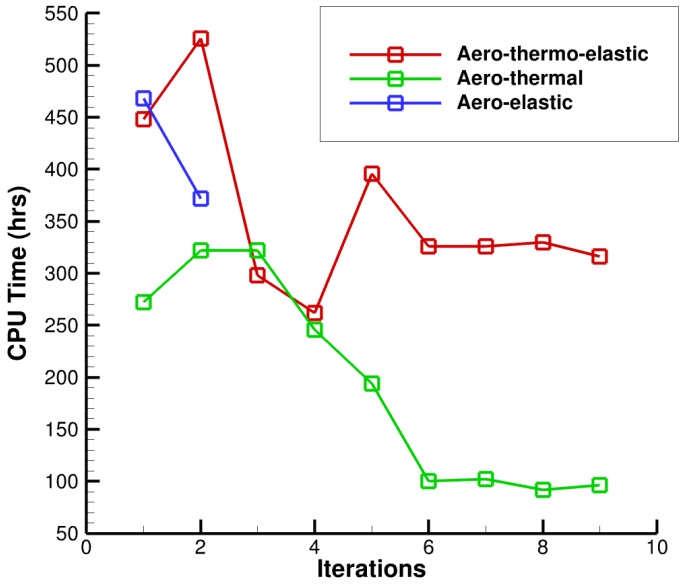

Figure 3.32: Example of the CPU time per iteration for each coupling process.

The work in this section, also found in Ref. [57], quantifies the error and computational savings of three different coupling methods under different situations: material choice, thermal management and freestream conditions. The goal is to provide information on the uncertainty that the less accurate methods accrue, which will be useful in Chapter 5 when aero-thermo-elastic optimizations are performed. Additionally, the results from this study indicate that the material choice can have a grave impact on the deformation in the case of Inconel, while high-temperature refractory metals such as TZM-C03 show much more promising results. Finally, the study also indicates that the deformation is highly sensitive to the freestream explored in this analysis.

3.5 Freestream Sensitivity

Deviations from the nominal freestream conditions of the trajectory are known, and expected, to occur for both flight and ground tests. Even for an undeformed geometry, the vehicle's aerodynamic and engine performances have been shown to be sensitive to these freestream variations, modifying the expected flowfield within the inlet and isolator of the scramjet [30]. This section aims to extend this understanding to capture the aero-thermo-elastic response of the vehicle. The sensitivity of the deformations to freestream variations could cause unexpected and appreciable changes to the aerodynamic and propulsive coefficients. A robust scramjet design will aim to reduce the sensitivity of the performance to varying flight conditions, even when considering aero-thermo-elastic distortions.

Kline et al. studied a three-dimensional inlet's sensitivity to aero-thermo-elastic deformation using polynomial response surface modeling [67], but additional work is needed on the subject. The previous work focused more on the deformation's sensitivity to the original shape and temperature of the inlet, rather than on the off-nominal flight conditions, and showed the influence both have on the engine performance. Although not the focus of the work, the results presented by Kline et al. showed that the various freestream conditions (two different Mach numbers and dynamic pressures) changed the response surface for several quantities of interest such as thrust. The results did not indicate whether the variations in the response surfaces were due to expected undeformed scramjet behavior as indicated by de Siqueira et al. [30] or due to changes to the deformation of the vehicle.

Therefore, the type of work by Kline et al. is expanded upon for this research. A similar analysis is performed to produce a quadratic surface response. When designing the computational experiment for the response surface, central composites or Box-Behnken methods are typically used [1]. The number of cases needed for a study depends on the design factors, or independent variables, and the chosen method. For three design factors, the Box-Behnken requires 15 cases while central composite methods require 20. Despite the advantage of requiring fewer cases, the Box-Behnken

method will often "miss corners" when combining extremes for the factors. Meanwhile, a central composite face-centered method (CCF) provides a high quality prediction over the entire design space.

Additionally, Lamorte et al. constructed polynomial response surface models for the aero-thermo-elastic response of a hypersonic vehicle to then use for an uncertainty quantification [75, 32]. Therefore, these types of analyses can lend themselves well to future work such as uncertainty quantification, extending the analysis beyond just the use of understanding the sensitivity of the engine to its flight conditions. Note that uncertainty quantification will not be presented in the current research plan, but is an idea for future work.

3.5.1 Test Setup

For this analysis, a two-way coupling strategy is used for the FTSI cases and an aerothermodynamic-only analysis with a radiative-equilibrium boundary condition at the wall is used for the "aero-only" cases. A geometry consisting of TZM-C03 paneling is considered and thermally-cooled boundary conditions of 1050 K and 1000 K are applied to the upper and lower internal forebody walls, respectively, as shown in Figure 3.21. Similarly, the 1050 K and 1000 K boundary conditions are applied to the lower and upper internal surfaces of the cowl, respectively. The two-temperature boundary condition is implemented to limit the deformation perpendicular to the flow, which has been shown to cause unwanted flowfield behavior in previous sections. The outer mold line of the inlet is consistent with the GK-01 inlet [43] and a kerosene fuel is used for the combustion process. The aerothermodynamic results for this geometry have been previously validated for the undeformed configuration against an experimental campaign performed at DLR Windtunnel [43] and can be found in Ref. [58] or Section 3.3.

3.5.1.1 Response Surface

To perform the sensitivity analysis, the design of the experiment is performed using a CCF response surface method. Three design factors are considered with the following mean and extremes:

Mach number (7.0 ± 0.5), dynamic pressure (0.75 ± 0.25 atm) and angle of attack (0° ± 1°). Initially, a linear response is fit to the data only to obtain a comparison of the significance for each factor. Additionally, an approximation of the resultant design space is then developed using a second-order polynomial model, fit using a least-squared regression:

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_{12} x_1 x_2 + \beta_{13} x_1 x_3 + \beta_{23} x_2 x_3 + \beta_1 1 x_1^2 + \beta_{22} x_2^2 + \beta_{33} x_3^2 \quad , \tag{3.1}$$

where y is the predicted output, $\mathbf{x}$ are the design factors and $\boldsymbol{\beta}$ are the coefficients. When appropriate, a third-order polynomial model is used to fit the data, as well. This applies mainly when the quadratic polynomial is shown to have a mediocre fit to the data and the cubic is able to improve that fit substantially.

3.5.2 Linear Response Model

To begin, the data is fit to a linear model. Although these vehicle responses are not necessarily fit best by a linear model, the general influence of each parameter can be established. Additionally, by comparing the undeformed aero-only cases with the deformed FTSI cases, the most influential design factor(s) on the deformation can also be established through the change in the design factor coefficients. To do so, the coefficients for the linear response model are shown in Table 3.8. Note that the coefficients for each quantity of interest (e.g. C_L) should only be compared to each other, not to the coefficients for other quantities of interest. Also, these aerodynamic coefficients are only representative of the inlet's contribution, not the entire vehicle.

Starting with the deformation of the cowl nose in the y-direction (perpendicular to the flow) there is an order of magnitude difference between each of the coefficients for the three design factors. The dynamic pressure is the largest, representing the most influential factor, followed by the Mach number. This trend will be seen in more detail when the quadratic response surface is shown in the following section. Conversely, the Mach number is shown to have the greatest influence for the maximum temperature of the inlet structure, while the dynamic pressure and angle of attack

are of a similar order of magnitude. These two quantities, deformation and temperature, represent physical attributes the deformed inlet will have when FTSI is included.

Next, vehicle-level quantities of interest for the deformed cases are compared against the undeformed cases. For the lift coefficient of the undeformed aero-only case, the Mach number has the largest coefficient or influence, followed by the angle of attack and finally the dynamic pressure. Once deformations are introduced using FTSI, the dynamic pressure influence/coefficient significantly increases to the same order of magnitude as the two other design factors. This result indicates that the deformation, and therefore its influence on the lift coefficient, is significantly affected by the dynamic pressure. A similar situation is observed for the drag and pitching moment coefficients, but the dynamic pressure coefficient has risen well over that for the angle of attack. This result indicates that the deformation caused by changing dynamic pressure is more influential than that of the angle of attack. Note that for lift, drag and pitching moment coefficients, the coefficients of influence for the Mach and angle of attack are also impacted by the deformation, but less significantly than the dynamic pressure coefficient. In addition, the specific impulse follows a similar trend, but the inclusion of FTSI raises the coefficient for the dynamic pressure to the same order of magnitude as that for the Mach number.

Finally, as a representative quantity for the deformation's impact on subsequent components, the maximum gas temperature within the burner is compared for the deformed and undeformed cases. Once again, a large increase in the coefficient of dynamic pressure is observed. An unexpected increase in the temperature within the combustion chamber could lead to conditions outside of the chamber wall material's capabilities and could even lead to a catastrophic failure, so understanding the dependence of these conditions on freestream variations is crucial.

3.5.3 Quadratic Response Model

The next step for the sensitivity analysis is to assess the validity of the response surface quadratic models. To do so, the accuracy of the regression is evaluated using the coefficient of determination of the prediction, (R^2), which reflects the proportion of variance in the output

Table 3.8: Coefficients for linear response model.

		M (-)	**Q (Pa)**	**α(°)**
$d_{Y,cowl-nose}$	*FTSI*	-5.11×10^{-4}	-1.91×10^{-3}	-8.62×10^{-5}
$T_{max,inlet}$	*FTSI*	354	53.8	10.9
C_L	*FTSI*	−0.0290	−0.0195	0.0190
	Aero-only	−0.0254	−0.0034	0.0187
C_D	*FTSI*	−0.0026	−0.0013	0.0057
	Aero-only	−0.0026	−0.0024	0.0047
C_M	*FTSI*	0.0185	0.0134	−0.0095
	Aero-only	0.0165	0.0032	−0.0009
I_{SP}	*FTSI*	−115	118	21.0
	Aero-only	−115	52.0	13.7
$T_{max,burner}$	*FTSI*	106	25.4	21.2
	Aero-only	98.7	−38.8	11.9

variable that is explained by the model. A value of 1.0 reflects a perfect fit between the model and the data, while a value of 0.0 is equivalent to a constant model predicting the expected value of the output. Therefore, a value closer to 1.0 indicates better predictive capabilities of the model. The R^2 value for each of the different model outputs are shown in Table 3.9.

Table 3.9: Coefficient of determination of the prediction.

	$d_{Y,cowl-nose}$	$T_{max,inlet}$	C_L	C_D	C_M	I_{SP}	$T_{max,burner}$
Aero-only	N/A	N/A	0.9786	0.6379	0.9567	0.7977	0.7861
FTSI	0.9093	0.9891	0.9824	0.6668	0.9700	0.8729	0.7605

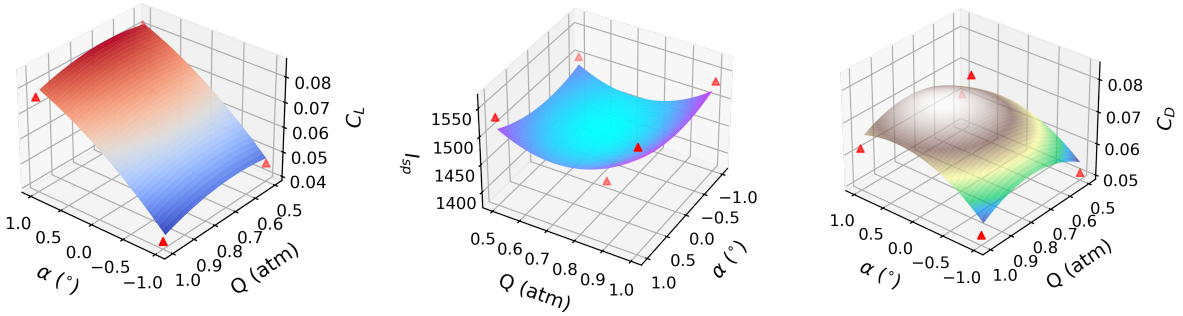

(a) Lift coefficient. (b) Specific impulse. (c) Drag coefficient.

Figure 3.33: 3D response surface contours with sample data points (triangles) at Mach 7.

The results from Table 3.9 indicate that the coefficient of lift and pitching moment coefficient for the inlet are well approximated by the second-order polynomial, while the coefficient of drag, maximum temperature of the gas in the burner, and specific impulse are only moderately well approximated. Still, general trends can be understood from these models, especially when comparing undeformed aero-only results with deformed FTSI results. To better understand these fits, examples of three surfaces and their corresponding sample data points are shown in Figure 3.33. The figures decrease in R^2 from left to right, but even for the drag coefficient, the general shape of the polynomial surface does match the data.

With the reliability of the surfaces understood, the contours can then be used to identify general trends appearing in the sensitivity of the inlet. To start, the displacement of the cowl nose in the y-direction is used as the representative value for the amount of distortion occurring for the inlet. The response surface of the cowl displacement, shown in Figure 3.34, indicates a strong dependency on the dynamic pressure. Higher pressure and shear loads will distort the vehicle further with larger dynamic pressures. Additionally, a slight dependence on angle of attack can be observed, which is more pronounced at higher dynamic pressures. The dependence is due to the change in the wall normal, which will affect both surface forces and heating. Also, the response is not symmetric about 0° since this is an asymmetric vehicle. Meanwhile, the amount of displacement is shown to increase with increasing Mach number. This trend is due to the increase in heating and temperature of the inlet with increased Mach number, as observed in Figure 3.35, which causes additional deformation due to expansion.

The deformations described above will then influence the trends observed in vehicle-level quantities of interest. Beginning with the lift coefficient, Figure 3.36 and Figure 3.37 juxtapose the response surface of the aero-only cases with the FTSI solutions, respectively, at various angles of attack. The aero-only cases show a decrease in lift coefficient with increased Mach number, but very little dependence on the dynamic pressure. In comparison, the FTSI cases show a much stronger dependency on dynamic pressure, especially at higher Mach numbers. This dependency on dynamic pressure appears due to the structural deformation's dependence on dynamic pressure

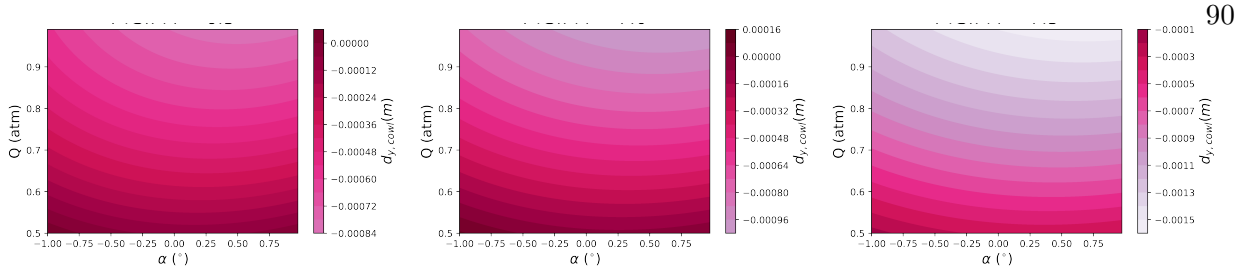

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.34: Response surface contours of cowl nose displacement in the y-direction at various Mach numbers.

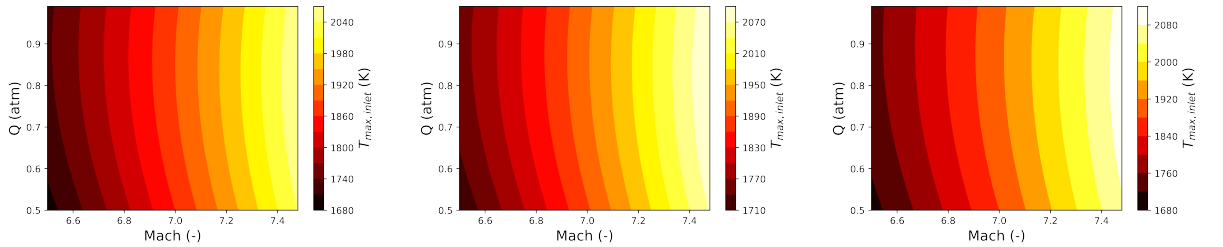

(a) $\alpha = -1°$. (b) $\alpha = 0°$. (c) $\alpha = 1°$.

Figure 3.35: Response surface contours of maximum inlet temperature at various angles of attack for the FTSI case.

shown previously.

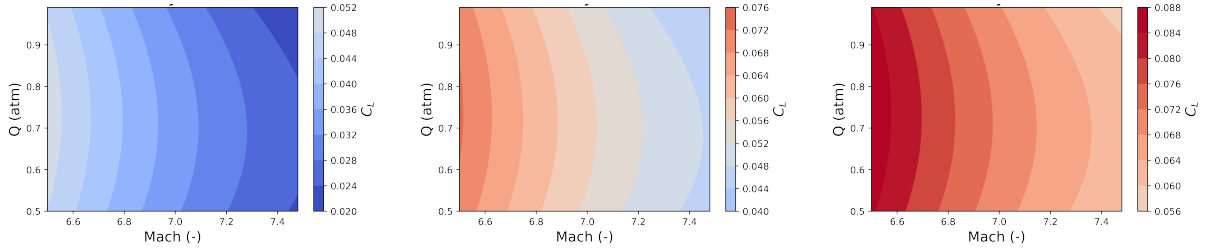

(a) $\alpha = -1°$. (b) $\alpha = 0°$. (c) $\alpha = 1°$.

Figure 3.36: Undeformed aero-only response surface contours of lift coefficient at various angles of attack.

The pitching moment coefficient, C_M, follows a similar trend, as observed in Figure 3.38 and Figure 3.39. The undeformed cases show a limited and roughly symmetric (about 0.7 atm) dependence on dynamic pressure. Note that this symmetry appears because of the aerodynamic

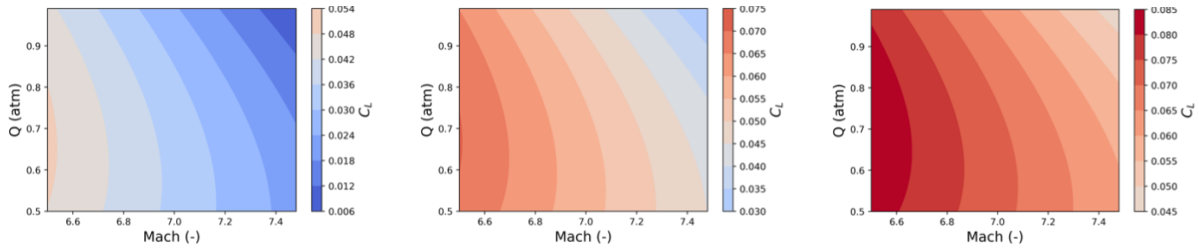

(a) $\alpha = -1°$. (b) $\alpha = 0°$. (c) $\alpha = 1°$.

Figure 3.37: Deformed FTSI response surface contours of lift coefficient at various angles of attack.

coefficients' normalization with the dynamic pressure. The deformed cases see much more dependence on the dynamic pressure and less symmetry at higher Mach numbers. According to Figure 3.34, deformation increases with increasing Mach number, which causes the Mach number dependence seen across Figure 3.39. The pitching moment coefficient minimum is shifted to lower dynamic pressures because the deformations deflect the forebody upwards (as observed in Figure 3.14). At higher dynamic pressures, the deflection will increase, pitching the nose more upwards and increasing the moment.

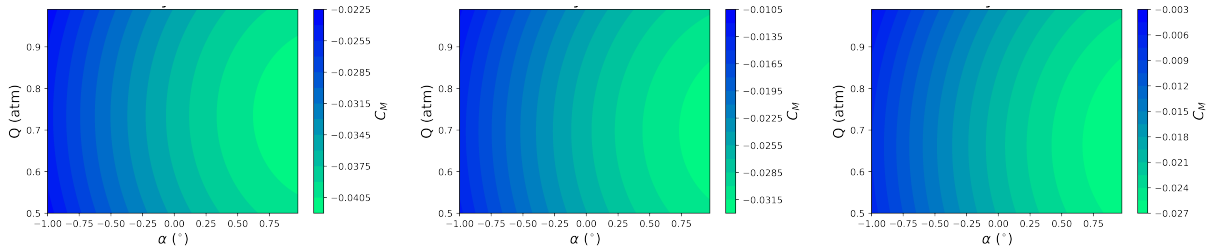

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.38: Undeformed aero-only response surface contours of pitching moment coefficient at various angles of attack.

Additional differences arise when comparing the specific impulse response surface. Figure 3.40 and Figure 3.41 similarly juxtapose the aero-only cases with the FTSI cases, respectively, but with varying Mach numbers. Two main differences arise in the shape of the convex surface: the location of the minimum and the eccentricity of the center. Beginning with the location of the minimum, when FTSIs are considered, the minimum I_{SP} shifts from a consistent dynamic pressure of 0.7 atm

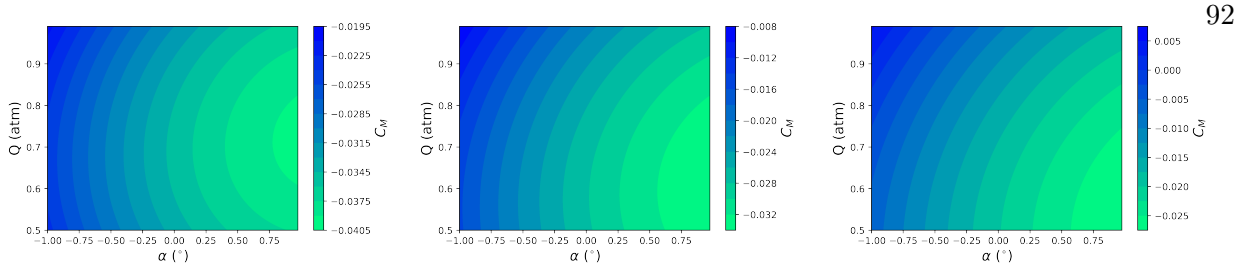

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.39: Deformed FTSI response surface contours of pitching moment coefficient at various angles of attack.

to a range between 0.6 ± 0.05 atm. Upon further inspection of the stream thrust components, this lowered dynamic pressure relates to the capture area changing and influencing the Sa_{area}. Figure 3.42 portrays the shift seen for Sa_{area}. The capture area is not only a function of the physical throat area of an inlet, but also the flowfield itself since it reflects the amount of flow being captured into the inlet [46], which are both dependent on the shape of the inlet. Therefore, this change to the I_{SP} and Sa_{area} reflect that a lower dynamic pressure causes a smaller increase in the capture area than higher dynamic pressures. Next, the eccentricity of the specific impulse contours is reduced towards a circle as FTSI is considered. This reduction in eccentricity mainly occurs along the angle of attack axis, indicating that the deformed geometry is more sensitive to angle of attack than the aero-only case. The angle of attack will impact the wall normal direction for the traction along the surface, dictating the deformations due to mechanical loading. Overall, the reaction of the specific impulse to freestream conditions has altered its behavior when considering deformations.

Meanwhile, the surface response of the maximum gas temperature in the burner does not change significantly when FTSIs are included. The location of the maxima across the various Mach numbers does change slightly, as seen in Figure 3.43 and Figure 3.44. For the undeformed case, the maxima are consistently located at 0.7 atm and 0.2° angle of attack across all Mach numbers. Once the deformations are introduced, the maxima begin shifting towards 0.78 atm and 0.3° angle of attack at higher Mach numbers. Additionally, an overall increase in temperature is observed when deformation is included, which should inform the design of the combustion chamber.

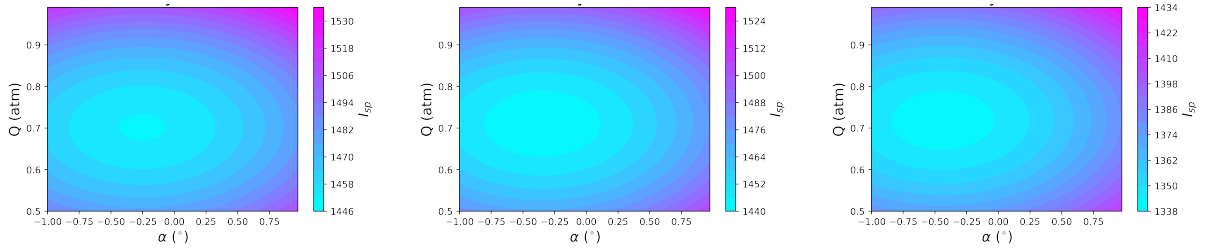

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.40: Undeformed aero-only response surface contours of specific impulse at various Mach number.

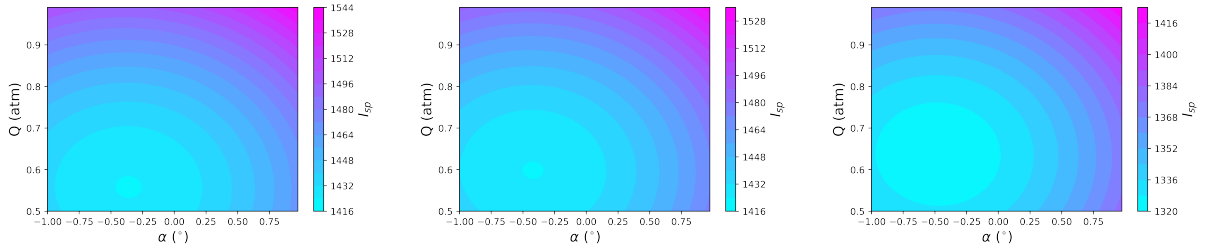

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.41: Deformed FTSI response surface contours of specific impulse at various Mach number.

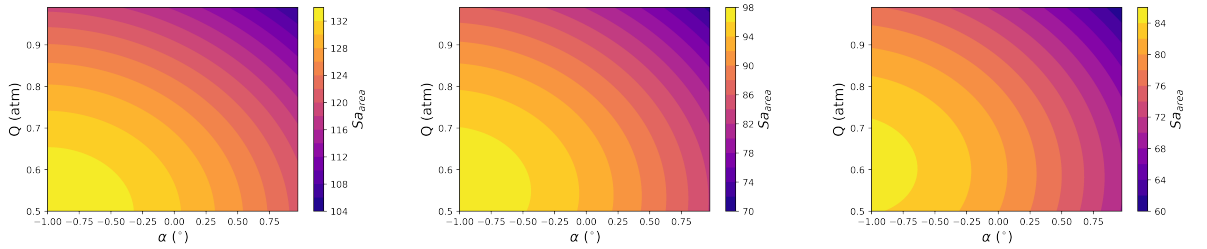

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.42: Deformed FTSI response surface contours of area stream thrust at various Mach number.

3.5.4 Cubic Response Model

As mentioned previously, the drag coefficient only has an R^2 value of approximately 0.65. Despite the poor fit, the surface results are portrayed in Figure 3.45 and Figure 3.46 for the undeformed and deformed configurations, respectively. Overall, the resulting surface response shapes

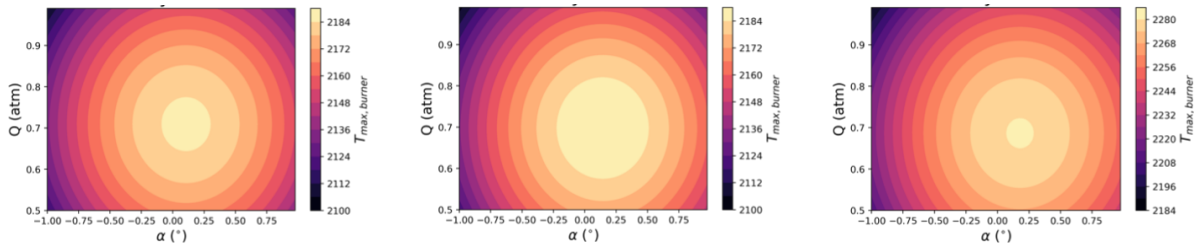

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.43: Undeformed aero-only response surface contours of maximum temperature in the burner at various Mach number.

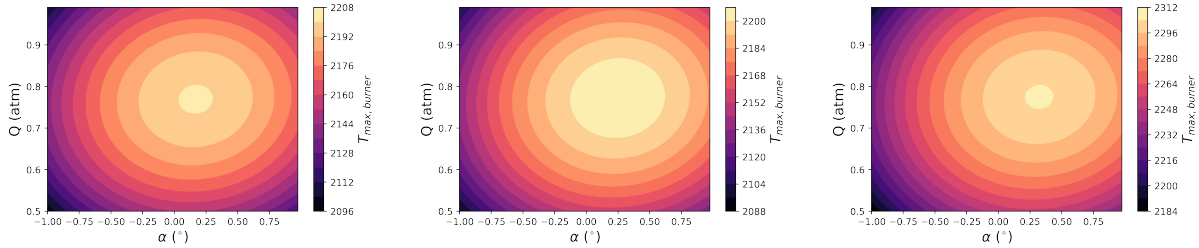

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.44: Deformed FTSI response surface contours of maximum temperature in the burner at various Mach number.

look very similar, with consistent minima locations and eccentricity. The magnitudes have increased when deformations are included, but there is a similar dependence on each of the design factors for both cases. This result is not consistent with the linear response model, indicating that the response surface is also not approximated well by the linear solution, which had an R^2 value of 0.35 for both cases.

In order to better capture the surface response of drag, a cubic polynomial is fit to the data. In doing so, the R^2 value reaches 0.9606 and 0.9630 for the aero-only and FTSI cases, respectively, indicating a much better fit for the data. The results using the cubic fit are shown in Figure 3.47 and Figure 3.48. Already, key differences appear that could not be observed with the quadratic fit case between the deformed and undeformed cases. The most clear difference is the dependence of the drag on angle of attack. The asymmetry of the inlet geometry causes a nonsymmetric behavior around the varied angles of attack. Also, by including deformation, the maxima across the various

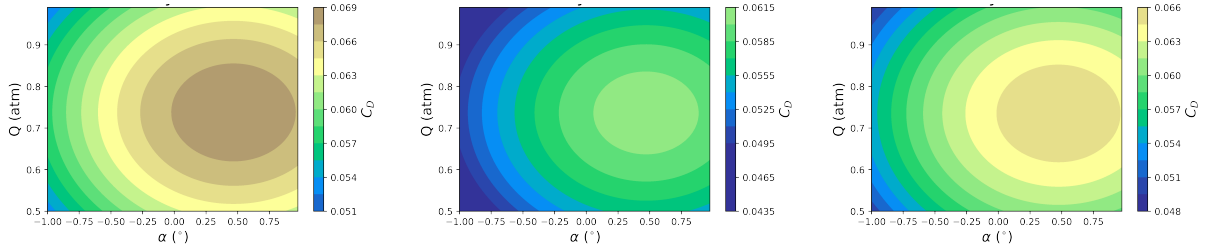

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.45: Undeformed aero-only response surface contours of drag coefficient at various Mach number.

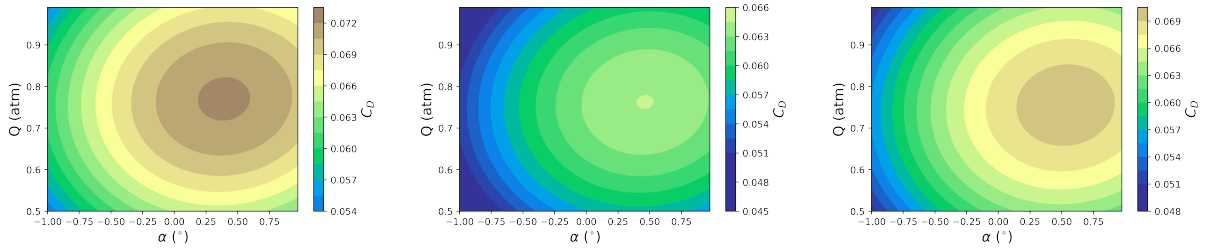

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.46: Deformed FTSI response surface contours of drag coefficient at various Mach number.

Mach numbers shift towards lower angles of attack. As seen in Figure 3.14, the forebody will deflect upwards, changing the angle of the forebody ramp and therefore introducing a new angle of attack, which will influence the drag. One additional conclusion seen when comparing the quadratic fit to the cubic is that the magnitudes, specifically the maxima, are not captured in the quadratic fit, further emphasizing the need to model the drag as a cubic function.

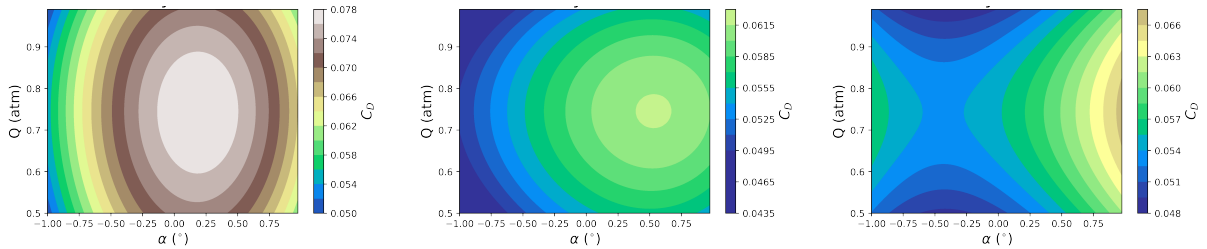

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.47: Undeformed aero-only cubic response surface contours of drag coefficient at various Mach number.

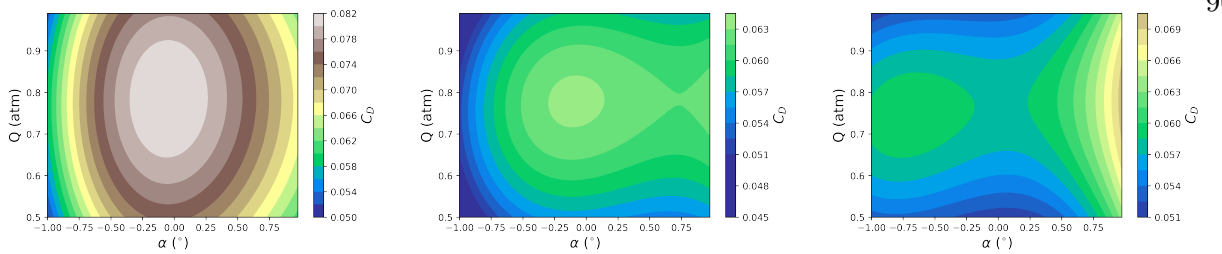

(a) Mach 6.5. (b) Mach 7.0. (c) Mach 7.5.

Figure 3.48: Deformed FTSI cubic response surface contours of drag coefficient at various Mach number.

Overall, these analyses and comparisons require the correct polynomial fit for the data in order to ascertain the impact of the deformations on the vehicle sensitivities. While some performance metrics are captured properly by quadratic fits, others benefit greatly with an additional degree to the polynomial. Once that response has been properly captured, the influence of deformations can be clearly seen for most quantities of interest. The results from this study can also be found in Ref. [56]. Freestream conditions are not the only unsteady behavior experienced by the vehicle; thermal transients can have a large influence on the deformation and resultant performance of the inlet.

3.6 Thermal Transient Analysis

As FTSI modeling improves, parallel efforts are needed for experimental work to be used as validation. One concern for experimental work is the time it takes to reach a thermal steady state, since that is indicative of the test time necessary. Therefore, FTSI models can help inform what the timing looks like as well as how the transient affects the problem. This work uses the outer-mold-line from Bhattrai et al. [10] and builds upon previous transient results from Currao et al. [28], adding some additional fidelity. Overall, building upon Section 3.5, this work aims to understand the uncertainty associated with a range of freestream conditions, which often occur in both flight and ground experiments.

3.6.1 Test Setup

A transient aero-thermo-elastic analysis is performed to extend the capabilities of the experimental inlet to realistic flight conditions. There is a large difference in the magnitudes of the time scales for the fluid and thermal response, and therefore the fluid response can be assumed to be quasi-steady at each evaluation [86]. Similarly, the structural dynamics and thermal response have disparate time scales and the structural oscillations have been shown previously to be damped-out within the first few seconds, again allowing for a quasi-steady assumption for the structural dynamics [28, 10, 86]. Initializing the geometry to the original design and the temperature to 300 K, the analysis steps through time until a steady thermal solution is reached. An explicit time integration scheme is performed, with quasi-steady fluid and structural dynamics. The heat flux and tractions along the wall output from LeMANS are fed into the thermo-structural analysis, which are performed as a staggered analysis beginning with the thermal component. The temperature field determined from MORIS is then used as the wall temperature boundary condition for the next step of the fluid solver and the displacements inform the fluid mesh. Constant freestream conditions are employed throughout the transient analysis for the fluid solver.

A molybdenum alloy, TZM-C03, is used for the entire intake structure due to its high-temperature strength. The temperature-dependent material properties for TZM-C03 are shown in Figure 3.49 and more information can be found in previous work by Kasen [65]. Following Currao et al.[28], the simulation represents a flight vehicle deploying from the protection of a fairing and entering a cruise condition. Therefore, the simulation begins with the entire structure equilibrated to 300 K.

The flight conditions for this analysis are for an altitude of 33.3 km, specifically over Boulder, Colorado in the United States. The nominal conditions for this location are shown in Table 3.10, found using the Earth Global Reference Atmospheric Model (GRAM) [123]. For this study, the freestream conditions are considered uncertain, since there are fluctuations in the atmosphere at this altitude. The previous study in Section 3.5 indicated that the dynamic pressure has a significant

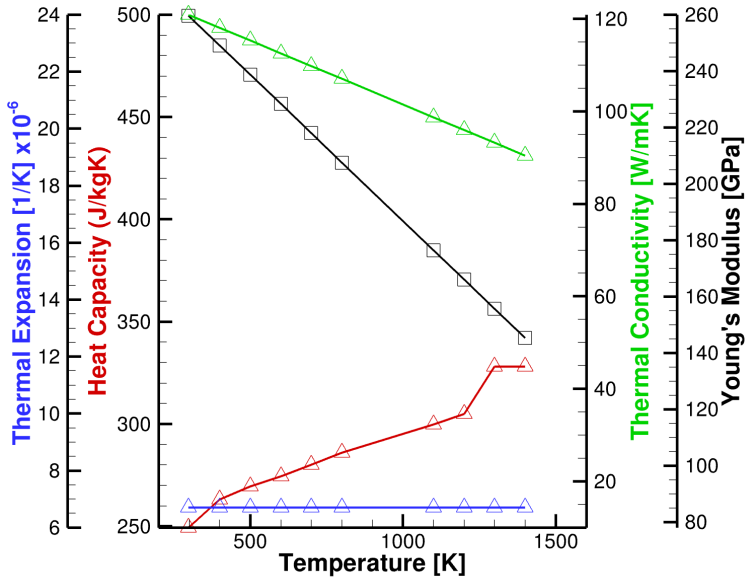

Figure 3.49: Temperature dependent material properties of TZM-C03.

influence on the aerodynamic and propulsive performance when deformation is considered. Based on the uncertainty in the density and the wind velocities, the freestream condition chosen to be uncertain in this analysis is the density, since the variation in density produces the largest difference in dynamic pressure. According to GRAM, the density has a standard deviation (σ) of 9.6×10^{-4} kg/m^3. In order to obtain the influence of the perturbed density, three conditions are applied across the entire transient analysis: nominal density, $+3\sigma$, and -3σ. In addition, an aero-thermal-only case is considered to compute the direct influence of the deformation on the temperature of the structure and performance of the inlet.

Table 3.10: In-flight freestream conditions at 33.3 km over Colorado, USA.

$\mathbf{M}$	$\mathbf{p}_\infty$	$\mathbf{T}_\infty$	ρ_∞	$\mathbf{T}_{w,0}$
5.85	747 Pa	231 K	1.11 $\times10^{-2}$ kg/m^3	300 K

3.6.2 Results

The nominal case considers aero-thermo-elastic behavior at the expected density. The transient analysis uses a step size of 5 seconds and is shown to reach a steady state ($\Delta T_{max} < 5$ K) at approximately 75 seconds. The increase in temperature along that period is shown in Figure 3.50a; the nose of the forebody and the enclosed portion of the intake (beginning at approximately 0.18 m) are the hottest regions along the ramp. Shown in Figure 3.50b, the structure is initialized as the original design, but very quickly adapts to the environment within the first few seconds. Additional deformation occurs as the vehicle heats up and expands, until the thermal response of the vehicle has reached a steady state. The steady state geometry bends over 2 mm from the original geometry in the y-direction and extends upstream over 1 mm in the x-direction. At steady state, very minor displacements (less than 5×10^{-5} m) are observed at 75 seconds, further indicating the structure has reached equilibrium, shown in Figure 3.50b.

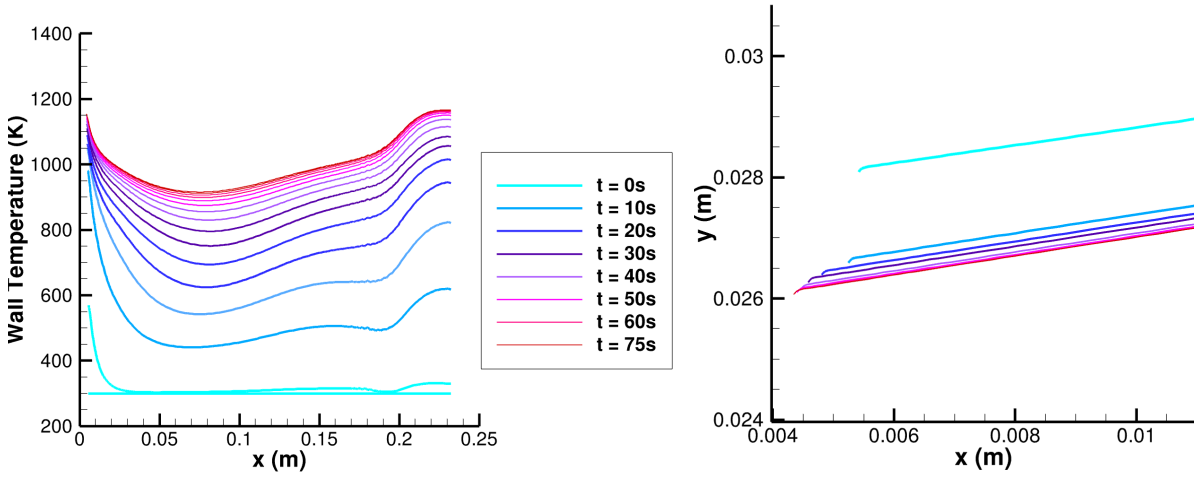

(a) Wall temperature.

(b) Ramp leading edge.

Figure 3.50: Compression surface transients for the nominal FTSI case.

The aero-thermal-only case considers the same ideal flight conditions, but ignores any de-

formations occurring. According to Figure 3.51a, very similar temperatures to the nominal FTSI case are produced along the ramp until the start of the enclosed region, again at 0.18 m. In the enclosed region, a lower temperature is computed at steady state when deformations are considered. This indicates that there is a significant enough change in the internal shock structure of the inlet. Similarly, the pressure is shown to be substantially higher for the nominal FTSI case in Figure 3.51b. Although the final profiles of the pressure are similar, the magnitudes are quite different, indicating that the total shock strength of the inlet has increased.

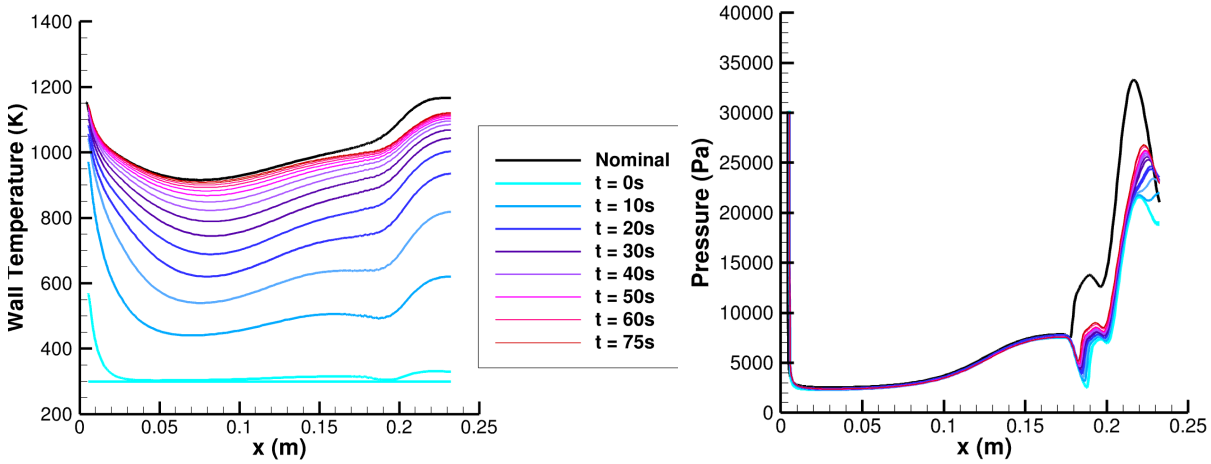

(a) Wall temperature.

(b) Wall pressure.

Figure 3.51: Compression surface transients for the nominal aero-thermal case.

One cause for this increase to the shock strength is the difference in the separation region occurring along the expansion corner of the ramp. The separation bubble, outlined in black in Figure 3.52, extends further upstream into the flow around the corner and outwards towards the cowl for the aero-thermo-elastic case. The deformation causes a stronger favorable pressure gradient at the expansion corner, which drives the separation. By doing so, the shock from the separation bubble is shifted forward, and the strength is increased, both of which can help contribute to the increase in temperature and pressure observed within the inlet.

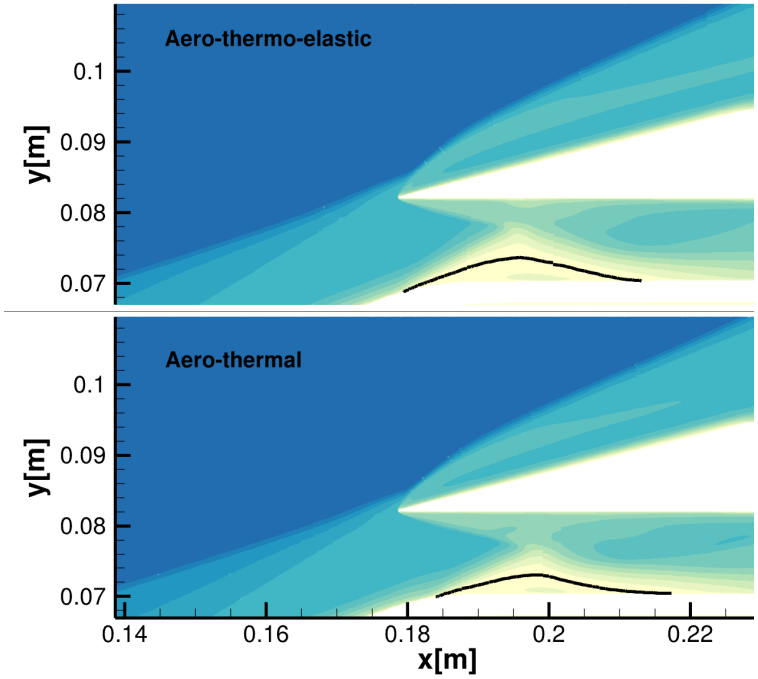

Figure 3.52: Mach contours of the cowl region for the aero-thermo-elastic and aero-thermal cases. The separation region is outlined in black.

To evaluate the uncertainty due to the density variations seen in this region of the atmosphere, the upper and lower 3σ bounds are updated for the densities, maintaining a constant freestream temperature and velocity. The transient wall temperature results for the upper and lower dispersed cases are shown in Figures 3.53a and 3.54a, respectively. Each can be compared to the final steady state temperature of the nominal case, depicted in black. Overall, both show very similar profiles, with slight changes to the temperature magnitudes in the direction of the change in density. Similarly, the pressure variations are captured in Figures 3.53b and 3.54b, respectively, which show similar trends. Narrowing the focus to the end of the inlet, the profile of the pressure evolves over time, indicating that the structure of the flowfield, shock-train and boundary layers within the inlet is time-dependent. Referring back to Figure 3.51b, the same evolution is observed in the pressure profile near the exit of the inlet for the aero-thermal case; this implies that the behavior is caused by a thermal disparity between the original and final states.

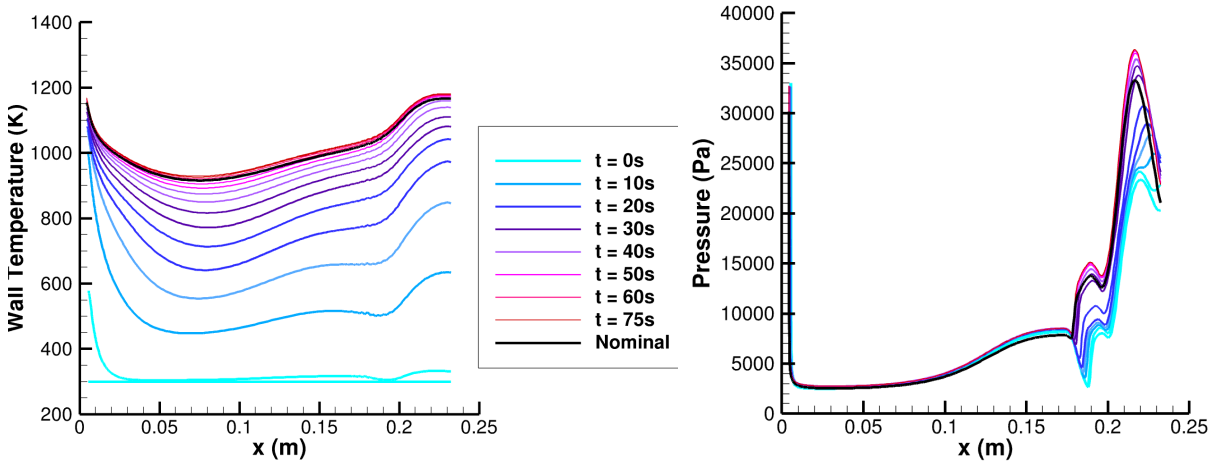

(a) Wall temperature.

(b) Wall pressure.

Figure 3.53: Compression surface transients for the $+3\sigma$ density FTSI case.

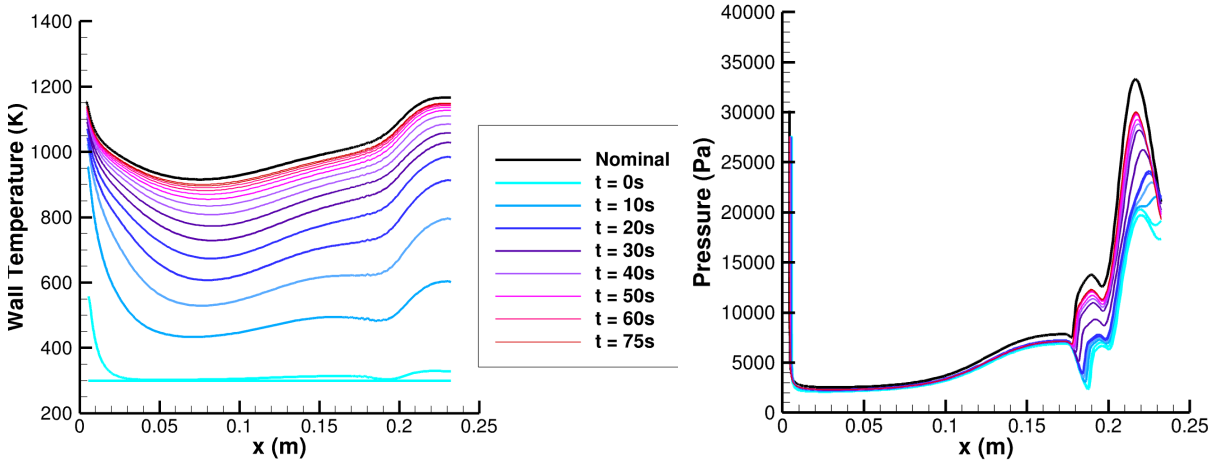

(a) Wall temperature.

(b) Wall pressure.

Figure 3.54: Compression surface transients for the -3σ density FTSI case.

The shock-train can be changed by a number of phenomena, including changes to the separation region along the shoulder of the ramp. Figure 3.55 shows the evolution of the separation bubble, outlined in black, at three states along the transient for the nominal FTSI case. The

separation region at earlier times begins just after the shoulder along the ramp, whereas the later period shows a growth of the separation region upstream along the shoulder. Additionally, the separation bubble protrudes further into the inlet towards the cowl, obstructing a larger area of the flow. The change to the separation region has a number of impacts, as mentioned previously, including increasing the shock strength and increasing the pressure within the inlet due to the larger obstruction. The bubble also impacts the shock-train due to the change in angle of the separation shock and moving the shock location forward.

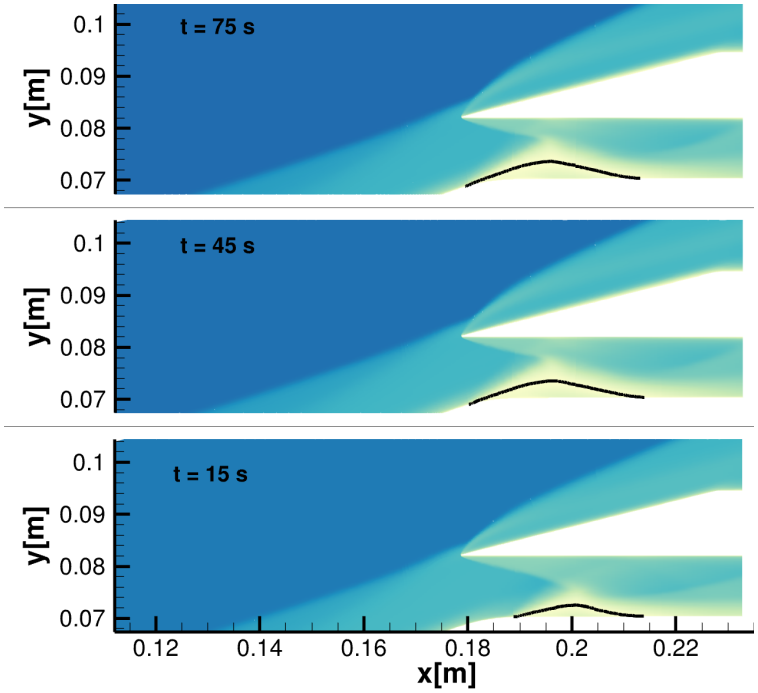

Figure 3.55: Mach contours of the cowl region for the nominal FTSI at three times. The separation region is outlined in black.

A summary of the final wall temperatures at steady state for each of the three cases is displayed in Figure 3.56a. The three FTSI cases show that the temperature directly depends on the density, as expected. The aero-thermal case, despite the same freestream conditions, produces lower temperatures than the nominal FTSI case, indicating that FTSI should be included to accurately capture the structure's temperature. Figure 3.56b includes the deformation of the three FTSI cases; once again, the amount of deformation is directly proportional to the freestream density.

Although only millimeters of deformation occur, a pronounced impact can be observed in the engine performance.

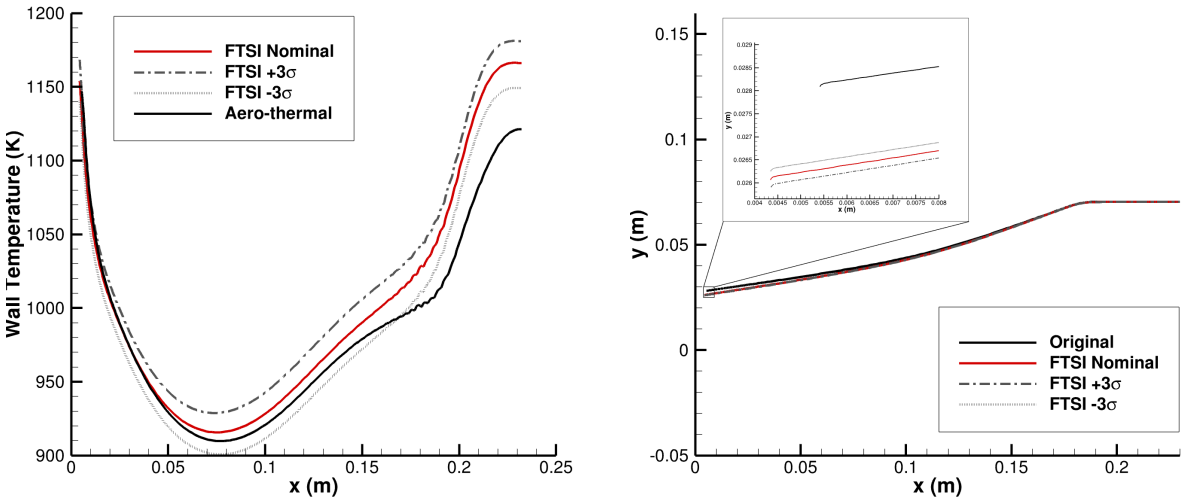

(a) Wall temperature.

(b) Geometry.

Figure 3.56: Steady state properties along the wall.

The specific impulse, I_{SP}, is often used as a measure of the engine's performance. Throughout the transient, both the elastic and thermal behaviors impact the I_{SP}, as shown in Figure 3.57. Beginning with the aero-thermal case, the thermal effects cause a rise in I_{SP} as the structure heats up. By decreasing the heat loss to the vehicle, more energy is maintained in the flow, which can be converted into usable energy or thrust, increasing the I_{SP} by over 6%. Next, by contrasting the I_{SP} observed in the aero-thermal analysis to the nominal FTSI analysis it is seen that the elastic impacts cause an additional rise in I_{SP}. This behavior is linked to the increase in pressure due to the strengthening of the shock emanating from the forebody and the widening of the separation bubble. The bounds around the I_{SP} are computed using the dispersed density cases. The bounds increase during the first part of the transient, indicating a higher variation in I_{SP} during that period, with a spread up to approximately 0.75%. But, as the structure approaches steady state, the bounds decrease slightly, producing less variation in the specific impulse, with a spread of about

0.33%. During the period of higher variability (between ∼20-40 seconds), the separation region is expanding, until it reaches a steady size and location at around 45 seconds. The freestream density impacts the exact timing of the separation bubble movement, which increases the uncertainty in I_{SP} during that transitional period.

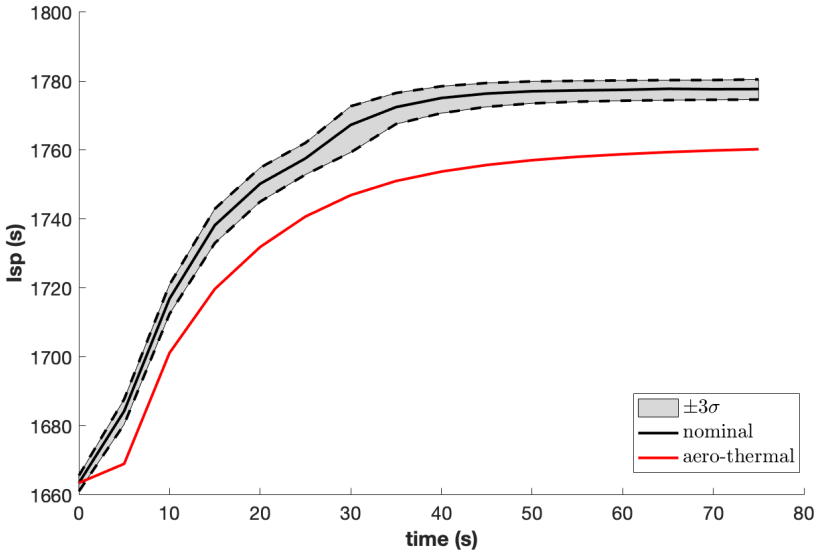

Figure 3.57: Transient behavior of I_{SP}.

The freestream density is just one of many uncertain parameters when modeling scramjets. In this case, the freestream density is an aleatoric uncertainty (cannot be reduced with more data because it is inherent to the system) of the Earth's atmosphere, along with other examples such as temperature and wind speed. Especially for hypersonic CFD and FEM modeling, there are several epistemic uncertainties (can be reduced with more data), including turbulence, material properties, and many more. These uncertainties are not included in this work but instead are discussed in the Recommendations for Future Work Section. The work for this section can also be found in Ref. [52].

3.7 Chapter Summary

This chapter looked at several studies using the aero-thermo-elastic method to understand the impacts of the deformations. The chapter began by validating the aero-elastic component of the tool against previous experiments. Although the experiment determined that 3D effects occurred, the 2D coupled analysis performed in this work showed good agreement with the overall deformation observed. After this result, the aero-thermo-elastic coupling, which has not been properly validated due to limited experimental data, was performed for the subsequent analyses.

The first study explored the impacts the deformation had on the inlet flowfield and structure, subsequent engine components, and vehicle performance. The study found that the deformation perpendicular to the freestream – opening of the inlet – can cause enough separation on the bottom and top surfaces that unsteadiness can arise in the flow. Larger separation regions and additional separation regions occurred when more aero-thermo-elastic deformations appeared in the y-direction. The increased separation, as well as increased shock strength, led to inefficiencies in the inlet and the engine and could cause unsteadiness in the flowfield. Meanwhile, the capability of both the inlet and the engine was seen to increase with deformation, although the undeformed configuration was shown to start below the recommended pressure ratio range. Both the inlet and the combustion chamber experienced increases in temperatures and pressures. Therefore, it was critical to consider the deformations when choosing materials for the engine to avoid catastrophic failures. Additionally, since the airframe and inlet are integrated, the deformations impacted the aerodynamic behavior of the vehicle, substantially decreasing the lift-to-drag ratio. The control authority of the vehicle was dependent on knowing the aerodynamic quantities, so as those quantities change with deformation, the controllability of the vehicle will also change.

The thermal expansion was shown to cause significant changes to the geometry, especially for an Inconel-based structure. As such, controlling the temperature and expansion through thermal management was studied. A two-temperature approach was developed to mitigate deformation in the y-direction, i.e. perpendicular to the freestream, that causes the inlet area opening, which

was shown to cause increases in separation and lead to unsteadiness of the flow. This method was successful in reducing many of the unwanted impacts, such as high temperatures in the combustion change and unsteadiness in the inlet flow; however, thermal management is not able to completely eliminate all deformation.

The second study explored the coupling approach necessary to capture the FTSI while balancing accuracy and computational cost. Three different FTSI coupling strategies were applied to a GK-01 scramjet inlet and assessed for accuracy and computational cost. These three coupling strategies – two-way, one-way and sandwich methods – were tested under various conditions including different materials for the inlet section of the scramjet, internal temperatures, and freestream conditions.

The first comparison was composed of the baseline material, TZM-C03, and Inconel 718. The material properties of each metal influence the resulting deformation. The case involving Inconel 718 saw much more deformation due to the higher thermal expansion coefficient and lower Young's Modulus at high temperatures. The second comparison was composed of three different internal temperature boundary conditions: 1050/1000 K, 1000 K and 1050 K. All three internal temperature boundary conditions showed similar errors for the one-way coupling and the sandwich coupling methods. Also, the temperature profile along the forebody showed very similar solutions for both the one-way and two-way coupling, indicating that the deformations did not substantially affect the temperature field of the fluid. The final comparison was composed of seven different freestream conditions, where Mach number, dynamic pressure and angle of attack were varied. Overall, a higher Mach number and a higher dynamic pressure were shown to cause the largest error between the two-way cases and the one-way/sandwich cases. Using a representative case, the CPU time necessary for the sandwich method was approximately 70% of the CPU time necessary for the two-way coupled case. In comparison, the one-way coupled case only reduced the computational cost by 20% compared with the two-way coupling cost. That substantial difference in computational efficiency could result in making more complicated FTSI analyses, such as 3D geometries or multidisciplinary optimizations, more realistic when accuracy is less necessary.

The third study illustrated that the consideration of FTSI was necessary across a variety of flow conditions, not just at the single design point, since the vehicle's response may vary significantly. The vehicle's sensitivity to changes in dynamic pressure, Mach number and angle of attack were altered due to the deformations occurring. Those deformations were also sensitive to changes in freestream conditions, introducing an additional complexity to understanding the sensitivity of the vehicle performance. As dynamic pressure increased, the surface forces along the inlet wall increased and caused larger deformations. In addition, the angle of attack dictates the normal direction of the surface forces, which impacted the deformation amount as well. Finally, the Mach number had a significant influence on the temperature of the inlet, leading to varied amounts of thermal expansion. The performance of the vehicle, which is dependent on the inlet shape, then inherited the dependencies on the freestream conditions that the deformations exhibit.

The response surfaces for various performance metrics of the vehicle such as C_L, C_D, C_M, I_{SP}, and maximum combustion gas temperature were compared for both the undeformed cases and deformed cases to pinpoint how impactful these deformations were. Not only did the deformations affect the magnitude of each of these quantities, but they also influenced the response surface's shape. The inclusion of FTSI showed the coefficients of lift and pitching moment and the specific impulse to be much more dependent on the dynamic pressure. Meanwhile, the drag coefficient no longer had the same dependency on angle of attack as it originally had for the undeformed configuration. The sensitivity of these vehicle-level quantities must be thoroughly understood to design and control a robust scramjet system capable of handling the changes due to inlet deformation.

The fourth and last study in this chapter focused on a transient analysis evaluated under atmospheric density uncertainty. The structure took approximately 75 seconds to reach a steady thermo-structural state when initializing the geometry from an isothermal condition of 300 K for all of the cases. The higher density case produced higher heat loads and dynamic pressures, which resulted in higher structural temperatures, deformations, and specific impulses; the opposite was true for the low density condition. A larger spread of uncertainty was observed midway through the transient, as the separation bubble expanded upstream, indicating that the uncertainty was

time-dependent. The aero-thermal-only case indicated similar temperatures on the external portion of the intake, but lower temperatures on the internal portion; in total, this caused a lower I_{SP} along the entire transient, with a difference of 17 seconds at steady state, emphasizing the need to account for the deformations at least in the uncertainty of the engine performance.

Chapter 4

Coupled Sensitivity Analysis

4.1 Chapter Introduction

The current chapter focuses on the semi-analytical method that is implemented to solve complex hypersonic multiphysics problems. Semi-analytical methods require intrusive measures, extending the physics solvers to include the use of direct or adjoint methods to compute the analytical gradients (also known as sensitivities within the context of optimization). Unlike blackbox finite-difference methods, these methods require knowledge of the governing equations and algorithms used to solve said equations, which is why the methods must be implemented into the physics solver. In short, adjoint and direct methods leverage the residual equations and their derivatives to compute the sensitivities of state variables on design variables; the state variable sensitivities are then used to compute the quantity of interest's gradient. The fundamental equations for solving the semi-analytical sensitivities were shown in Chapter 2.

Both adjoint and direct methods are based on Equation 2.33, but the direct method requires solving an equation for every design variable whereas the adjoint requires solving an equation for every objective [84]. Therefore, when many design variables are of interest, the adjoint can be more computationally efficient. Copeland et al. showed a CFD solver, SU2, with adjoint analysis that was built upon to handle hypersonic flow physics [24]. On the other hand, Damm et al. began with a fully developed hypersonic CFD solver and added features to implement an adjoint analysis [29]. One downside of an adjoint method is that the equations have a different structure than the physics equations, which can often make the implementation more complicated. For optimizations

where the design variables and optimization variables are on a similar order of magnitude, the direct approach still performs well, and will therefore be leveraged for these analyses.

For this research, the hypersonic solver utilized will be LeMANS (details of the solver were discussed in Section 2.1.1), which has been previously used for optimization and even multidisciplinary optimization that included shape change due to ablation [35, 36, 90]. But, this work was performed on a simplified geometry – a nose-cone – that allowed the blackbox finite-difference method to be computationally feasible. For more complicated geometries, such as a scramjet inlet, the blackbox method is too computationally expensive. Therefore, a direct method is implemented within LeMANS.

Because the sensitivity analysis is a new feature of LeMANS, this chapter begins with a benchmarking of the direct method to ensure the proper implementation using the IRV-2 vehicle (which can also be found in Ref. [55]). Similarly, the coupled sensitivity analysis integrates multiple codes together for the first time; therefore this chapter also includes a benchmarking effort to confirm the accuracy of the new coupled method. The later sections of this chapter include the benchmarking for the aerothermodynamic-only and aero-thermo-elastic sensitivities for the scramjet inlet case, using the geometry designed by UNSW [10], and can also be found in Ref. [53]. These cases, both the IRV-2 vehicle and scramjet inlet, provide confidence for the optimizations following this chapter.

4.2 Aerothermodynamic Benchmarking

When just considering the aerothermodynamic analysis, Equations 2.32 and 2.33 can be simplified as follows:

$$\frac{df}{ds_i} = \frac{df}{ds_i}^{expl} + \frac{df}{d\mathbf{q}}\frac{d\mathbf{q}}{ds_i} \tag{4.1}$$

$$0 = \frac{\partial \mathcal{R}}{\partial s_i} + \frac{\partial \mathcal{R}}{\partial \mathbf{q}}\frac{d\mathbf{q}}{ds_i} \quad . \tag{4.2}$$

The direct method implemented into LeMANS solves Equation 4.2. As a new feature of LeMANS, the sensitivity analysis must be benchmarked before use. This is performed on the nose region of the IRV-2 geometry [72]. To isolate the direct method implementation in LeMANS without the mesh sensitivity step, a design variable of the inflow velocity is used instead of the Bézier shape. This choice has two major advantages: (1) any potential errors due to the Bézier curve or RBF derivatives are not included and (2) the derivatives along the wall are all implicit-only through the state variables. A forward finite-difference method – two flow solutions with different inflow conditions – is compared to the direct method described above. Figure 4.1b juxtaposes the contours of the density multiplied by the y-component of velocity (ρv) sensitivity throughout the flowfield. The forward analysis is required to converge to an absolute residual of 1×10^{-14} and then a $+1\%$ perturbation on the velocity is applied. In order to compare the flowfields and wall properties, a root-mean-square-percentage (RMSPE) is used:

$$RMSPE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(\frac{\hat{y}_i - y_i}{y_i}\right)^2} \quad , \tag{4.3}$$

where n is the number of data points, $\hat{y}_i$ is the direct method value and y_i is the finite-difference value. The two sensitivity methods show identical flow features for the sensitivity, with an RMSPE of approximately $4 \times 10^{-3}\%$; while the finite-difference solution requires two converged flowfield solutions, the semi-analytical sensitivity method only requires one. To further highlight the agreement between the two methods, Figure 4.2 compares the shear stress along the wall for both the finite-difference (FD) and GMRES method, which produces an RMSPE of approximately $6 \times 10^{-6}\%$. Additionally, the solution using only the LeMANS-computed Jacobian and the tridiagonal solver is plotted and shows much worse agreement, indicating that the LeMANS-computed approximation of the Jacobian is not accurate enough to be used as the Jacobian for the direct method.

As mentioned in Chapter 2, a right-preconditioned linear system is implemented for the GMRES solver. This method is chosen to leverage the Jacobian and tridiagonal solver already included in LeMANS to speed up the convergence of the GMRES algorithm. Although not accurate enough to be used as the true Jacobian in the sensitivity equations, the LeMANS-computed Jacobian is

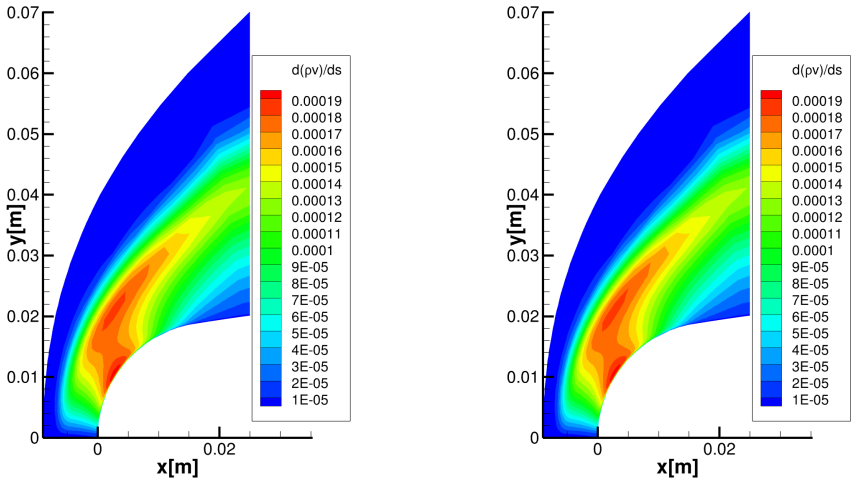

(a) Finite-difference solution (b) Semi-analytical method solution

Figure 4.1: Contours of the sensitivity of the state variable, ρv, to the design variable.

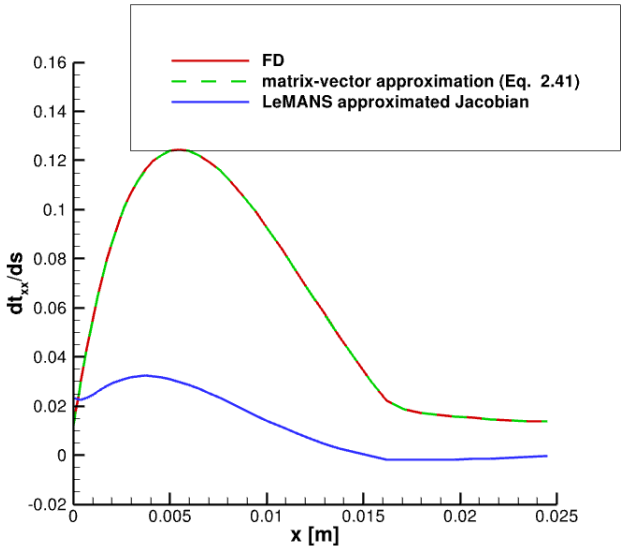

Figure 4.2: Sensitivity of the shear stress along the wall using three different solution methods.

a useful preconditioner. By applying it as such, the algorithm converges more reliably and more efficiently. A comparison of the convergence for the right-preconditioned system with two generic left-preconditioned systems is shown in Figure 4.3. The two other choices considered for the ease of calculating the inverse are Jacobi and Gauss-Seidel matrices, which only store a portion of the

Jacobian such as the diagonal or lower-triangular component. Both show much slower convergences than the algorithm leveraging the LeMANS Jacobian.

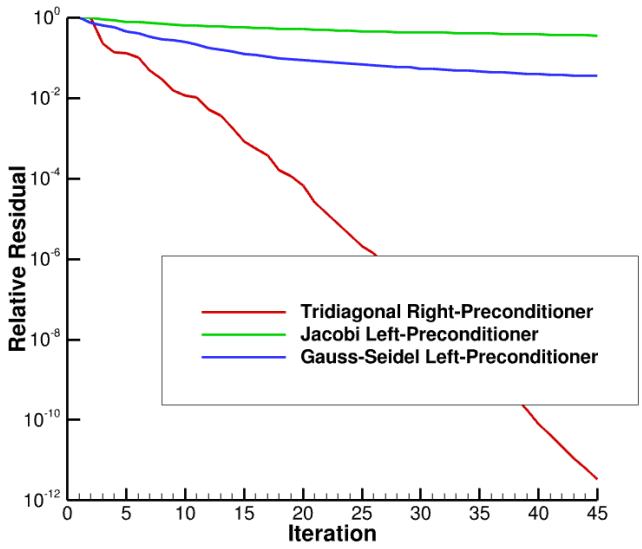

Figure 4.3: Convergence rates for three preconditioners implemented in the GMRES algorithm.

With the GMRES solver and direct method for the flowfield verified, the addition of the shape sensitivity can be benchmarked. Once again, the finite-difference method is used as the reference solution for the sensitivities. For this analysis, four control points/design variables are considered. Restricting the design changes to maintain a smooth leading edge, the first control point and the second control point's x-coordinate position are fixed. The y-component of the second control point is then the first design variable. The third control point can be adjusted in both the x- and y-directions. Finally, the last design variable is the x-component of the third control point, whereas the y-component is fixed to allow for a smooth transition to the rest of the vehicle. The resultant shape change from the initial shape for each design variable (dv), individually, are shown in Figure 4.4.

The current results focus only on the first design variable and the resulting perturbed wall geometry, since the other design variables show very similar comparisons. The sensitivity of the wall node locations is propagated through the entire sensitivity analysis described previously. Note that

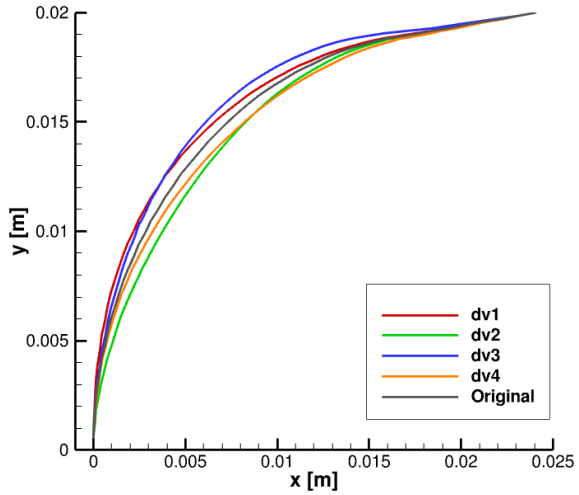

Figure 4.4: Deformation due to individual perturbations of each design variable for the IRV-2 nose.

this sensitivity analysis includes the computation of the mesh sensitivities through Equation 2.36. Focusing on the results along the wall, Figure 4.5 includes both pressure and shear stress sensitivities to the first design variable (a perturbation in the y-direction) to account for the verification of implicit and explicit derivatives. Figure 4.5 shows excellent agreement for both, with RMSPE values of less than 5% and 7% for the pressure and shear stress sensitivity, respectively, indicating the mesh sensitivity is properly integrated into the analysis.

Additionally, the gradient of the drag coefficient for the initial shape is compared for the new direct method against the finite-difference approach. The results are shown in Figure 4.6, where each design variable is plotted. The two gradients produce similar results (<5% RMSPE), further benchmarking this implementation for the gradient computation. With confidence that the direct method is producing accurate gradients, the aerothermodynamic sensitivity analysis can be introduced to an optimization process or an aero-thermo-elastic sensitivity analysis.

The benefit of the semi-analytical method introduced in the previous section is primarily the efficiency of computing these gradients. For a finite-difference analysis with four design variables, a baseline case is converged, followed by four perturbed cases, each requiring a few thousand iterations. In comparison, the direct method only needs one converged baseline solution per design

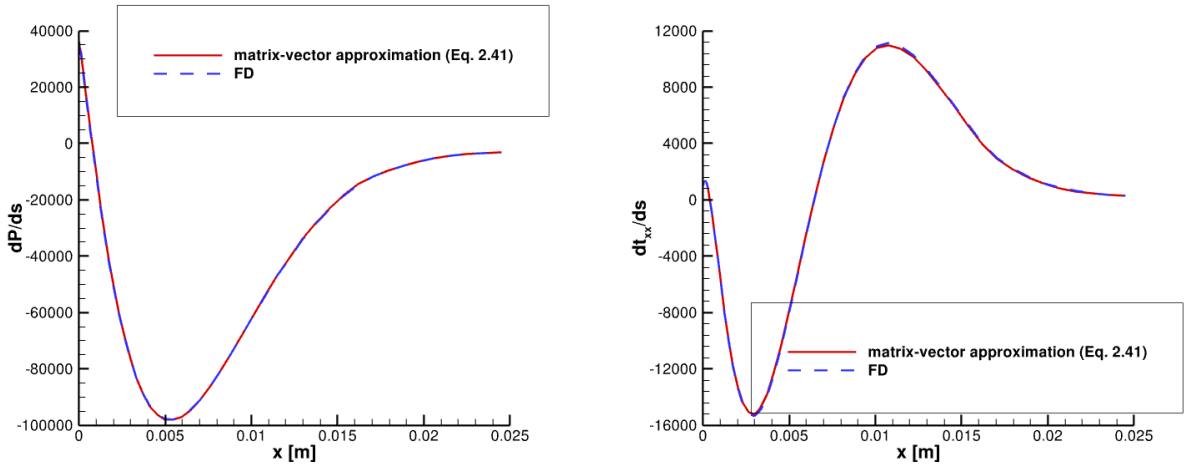

(a) Pressure.

(b) Shear stress.

Figure 4.5: Sensitivity of wall variables to the design variable.

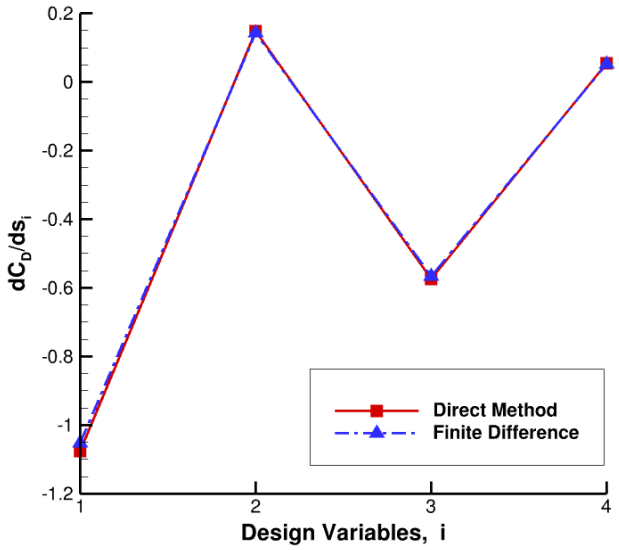

Figure 4.6: Comparisons of the drag coefficient gradient for the IRV-2 configuration.

iteration. From there, the drag sensitivity for each of the four design variables can be computed in around 40 GMRES iterations, which is approximately 40 residual calculations. Therefore, using the direct method allows for approximately 100 times faster computations of the gradients. The benefit of the direct method can already be observed with only four design variables. As this

analysis is scaled to more design variables, the benefit of the analysis will scale similarly. However, the sensitivity analysis may speed up the computation of the gradient, but the search algorithm will remain the same, so the total calls to the fluid solver may not scale linearly. Overall, the semi-analytical sensitivity analysis method reduces the cost of a gradient optimization method and improves the robustness and accuracy of the computation.

4.3 Aero-thermo-elastic Benchmarking

The current section characterizes the aero-thermo-elastic sensitivity analysis that was explained in Section 2.3. The integration and benchmarking of both methods with the aerothermodynamic sensitivity analysis is detailed below for a simple nose-cone geometry from the IRV-2 vehicle. The coupled sensitivity analysis is benchmarked to ensure the proper implementation. The IRV-2 nose outer mold line is used again, but TZM-C03 is used for the material [65]. An adiabatic boundary condition is assumed on the axis for symmetry and an isothermal temperature of 1000 K is assumed on the back face. A fixed Dirichlet condition is assumed on the back face, as well. Note that this is not a representative case of the true in-flight behavior of the IRV-2 vehicle, which was composed of reinforced carbon-carbon [72], but this case is instead used as a means of testing the new sensitivity analysis implementation.

The resultant deformation computed through the forward analysis is shown in Figure 4.7. After converging the temperature and deformation in the forward analysis, the sensitivity analysis begins by computing the aerothermodynamic sensitivities without any influence from the deformation and wall temperature sensitivities. This heat flux sensitivity solution along the wall for the first step in the sensitivity analysis is shown in Figure 4.8, along with the final FTSI solution. Although the aero-only first step produces a very similar curve to the finite-difference solution, the coupled SA agrees more closely, highlighted in the zoomed in region and the tail of the curve.

Extending the study to all four design variables, Figure 4.9 compares the wall pressure sensitivities for each design variable. Again, a close agreement for each design variable is seen between the coupled sensitivity analysis and the finite-difference method, all of which are within an RMSPE

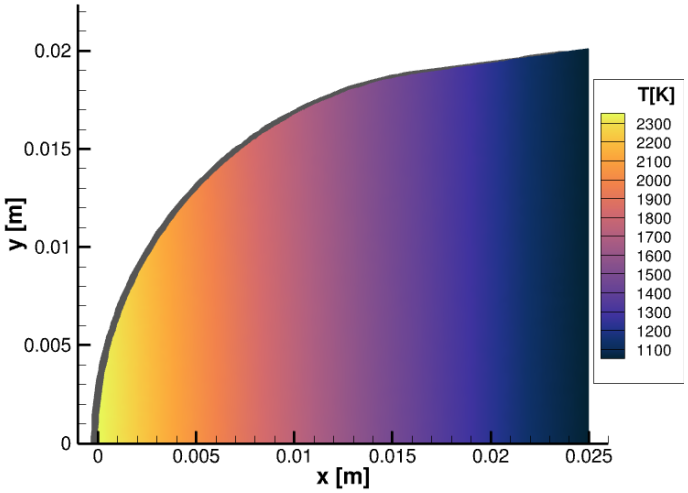

Figure 4.7: Temperature contour along with the thermo-elastic deformations shown in grey for the IRV-2 example.

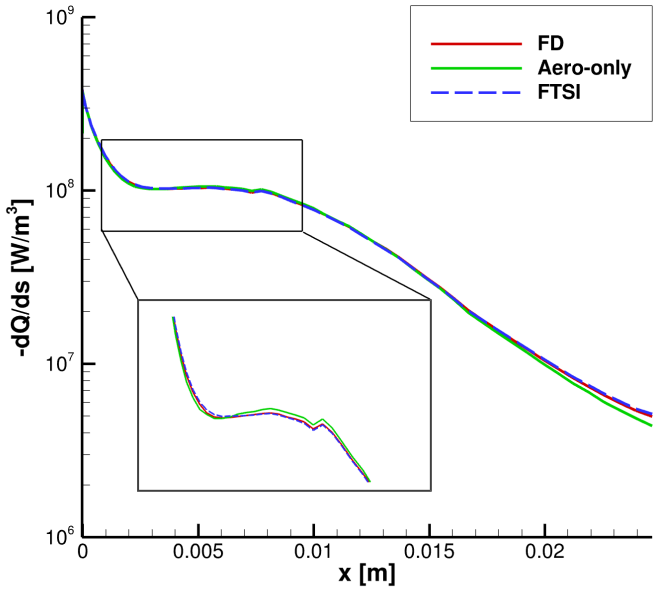

Figure 4.8: Heat flux sensitivities for the coupled aero-thermo-elastic sensitivity analysis.

of 7%. In addition, the sensitivity of the drag coefficient for each of the design variables is shown in Figure 4.10. The figure indicates that the 3rd and 4th design variables have a larger influence

on the drag coefficient. Overall, a good agreement between the coupled sensitivity analysis and the finite-difference solutions is shown, providing confidence in the method and implementation of this analysis.

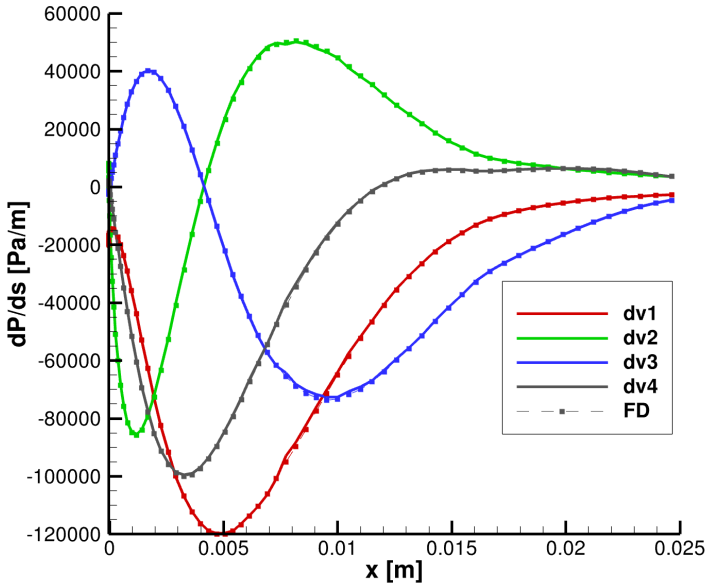

Figure 4.9: Comparisons of the optimized geometry using the finite-difference (dashed lines) and the direct method (solid lines) sensitivities.

4.4 Scramjet Sensitivity Analysis

The geometry chosen for the analysis is the compression surface and inlet built by Bhattrai et al. for aero-elastic experiments [10]. The validation for the cold-flow forward analysis is shown in Section 4.2 and uses the Spalart-Allmaras turbulence model. The current sensitivity analysis extrapolates to flight conditions, as shown in Table 3.10. Similar to the nose-cone case above, an effort to benchmark the sensitivity analysis for a scramjet inlet is provided here.

4.4.1 Aerothermodynamic

A preliminary aerothermodynamic sensitivity analysis is performed at the more realistic flight conditions with an isothermal wall temperature assumption of 300 K. The compression surface

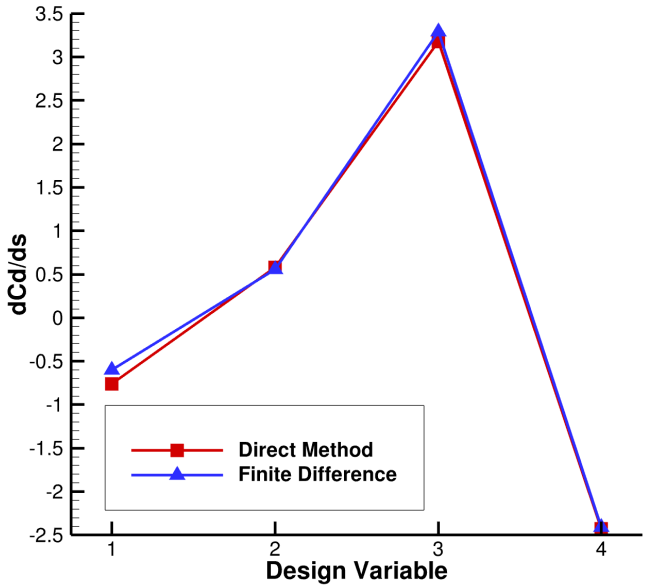

Figure 4.10: Comparisons of the drag coefficient gradient for the IRV-2 configuration.

forward of the isolator is chosen as the region to deform for the optimization procedure. Only four design variables are used to adjust this surface. Five control points are used for the Bézier curves. The first and last coordinates are fixed; the y-value of the second and last are also fixed to produce smoothness at the ends; the remaining four coordinates of the control points are defined by the design variables. As an example of the deformation, Figure 4.11 shows the original geometry and the perturbed geometry when each design variable is perturbed individually.

This initial analysis includes only the aerothermodynamic portion, with a converged, undeformed CFD computation. An aero-only sensitivity analysis is performed to verify the accuracy of the sensitivities for the scramjet inlet. The sensitivity of the wall pressure along the interior of the forebody and cowl to each design variable is shown in Figure 4.12. Both the finite-difference approach and the direct method are included; the direct method shows good agreement generally but with slight differences in the magnitudes of the peaks within the inlet. In this region, separation and a shock train are present, creating a complicated flowfield to resolve. The overall shape and general magnitudes fit very closely, despite the challenging flowfield occurring within

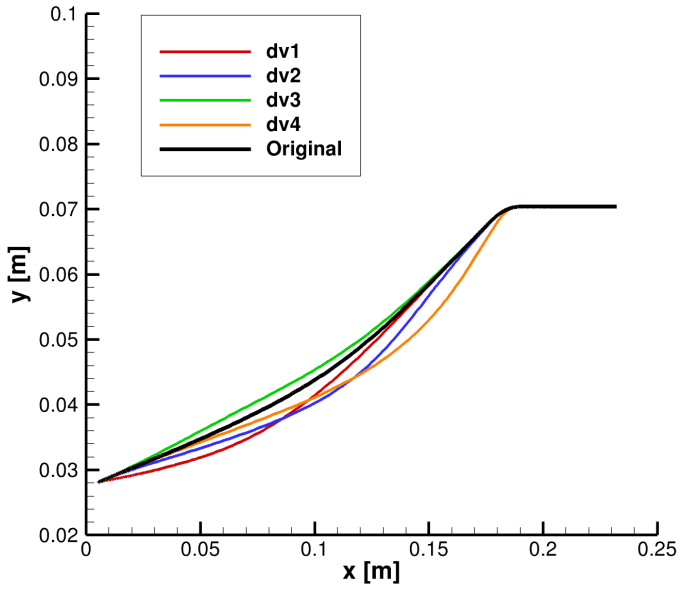

Figure 4.11: Deformed forebody for each design variable (deformation is scaled for visualization).

that region, resulting in RMSPE of 6%, 12%, 16% and 10% for the forebody wall pressure for each design variable, respectively. In particular, the compression surface of the forebody prior to the shoulder shows great agreement between the two methods. The cowl performs similarly, with slight differences in the magnitudes of the peaks, but the locations of such features are well resolved.

For the case of the scramjet inlet, the outflow into the isolator is extremely important for inlet and engine performance. The sensitivities along the wall are still necessary for the feedback between the coupled sensitivity analyses, but the outflow parameters, such as the compression ratio and total pressure recovery, are indicative of the performance of the inlet. Therefore, the sensitivities at the exit plane of the inlet are compared, as well. Figure 4.13 compares the direct method to a finite-difference sensitivity analysis for the pressure along the outflow plane within the inlet for the first two design variables. Focusing on the first design variable, the two show good agreement, producing an RMSPE 10%. However, when the 2D flowfield pressure sensitivity is converted into a 1D sensitivity to integrate into the engine model, there is less than a 2% difference in the results. The outlet sensitivity will inform objectives and constraints for the optimization

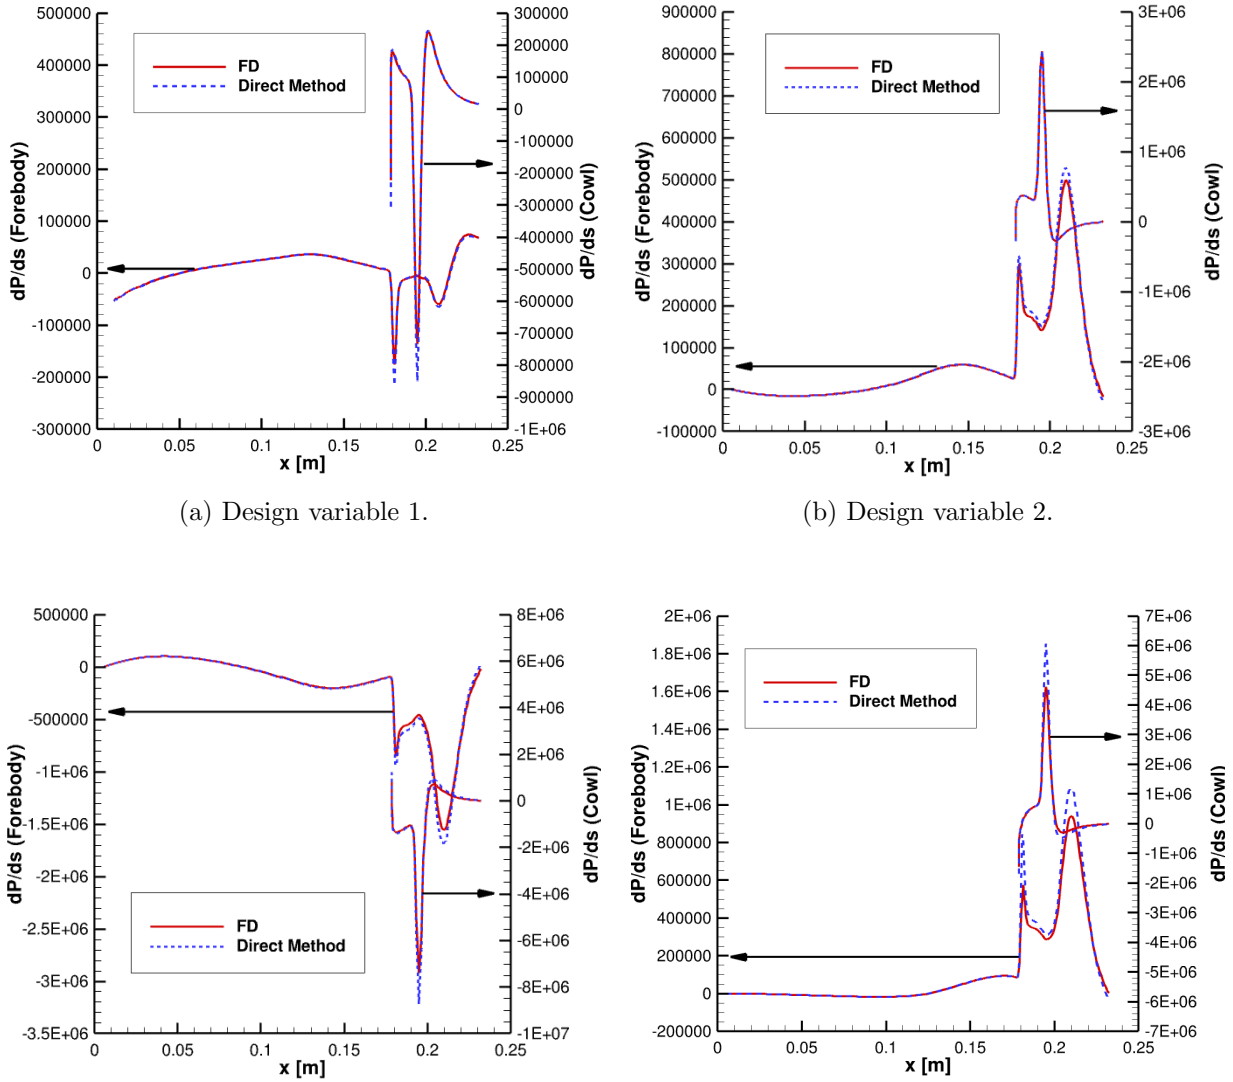

(a) Design variable 1.

(b) Design variable 2.

(c) Design variable 3.

(d) Design variable 4.

Figure 4.12: Sensitivity of wall pressure variables to the design variable.

process of the inlet.

If either the compression ratio or the total pressure recovery is considered the objective, the goal of the optimization would be to increase the objective. Therefore, the quantity must be reformatted as a minimization problem for the SLSQP minimization algorithm. For example, if the goal is to maximize the total pressure recovery (π_c), one way to define the objective as a minimization is to reformat the total pressure recovery as an expression with the reciprocal:

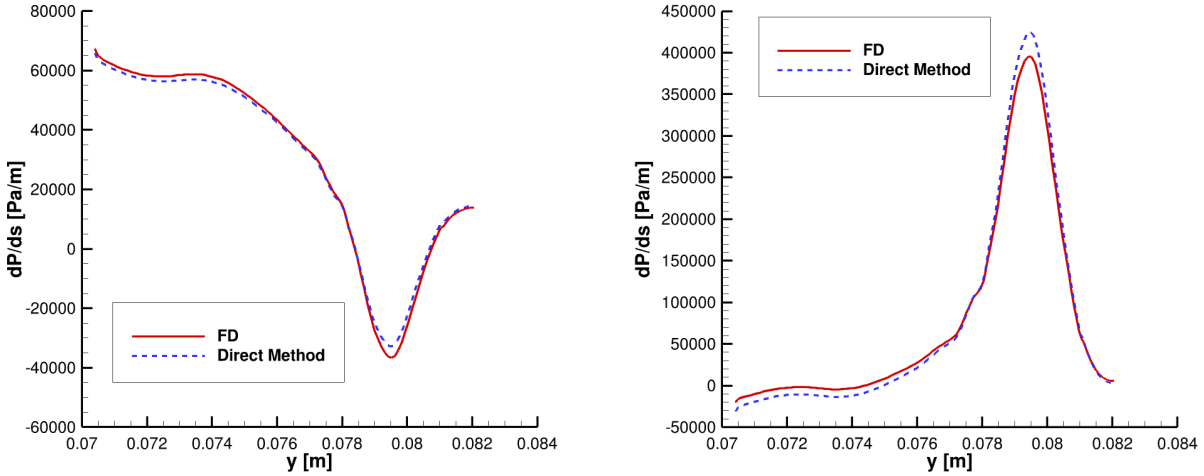

(a) Design variable 1.

(b) Design variable 2.

Figure 4.13: Sensitivity of pressure at the isolator outflow.

$$f = \pi_{c,0}/\pi_c \qquad , \tag{4.4}$$

where $\pi_{c,0}$ is the total pressure recovery of the original design. The resultant sensitivity is shown in Figure 4.14, comparing the direct method against the finite-difference method for each design variable. An RMSPE of less than 8% is shown between the two methods, indicating that the proper gradient is computed for the new implementation of the direct method. For this example, the third design variable is shown to have the largest influence on reducing the objective or increasing the total pressure recovery.

In terms of computational expense, the aerothermodynamic direct method sensitivity analysis is considerable faster than a finite-difference approach. A fully converged flowfield restarted from a similar case takes approximately 480 CPU-hours to complete. Therefore, the computational cost for approximating the objective gradient for all four design variable using a forward finite-difference solution is 1920 CPU-hours. For the direct method, each design variable takes approximately 10 CPU-hours to converge for a total of 40 CPU-hours for the entire objective gradient. That speed-up enables the optimization of this complex design, especially when the aerothermodynamic sensitiv-

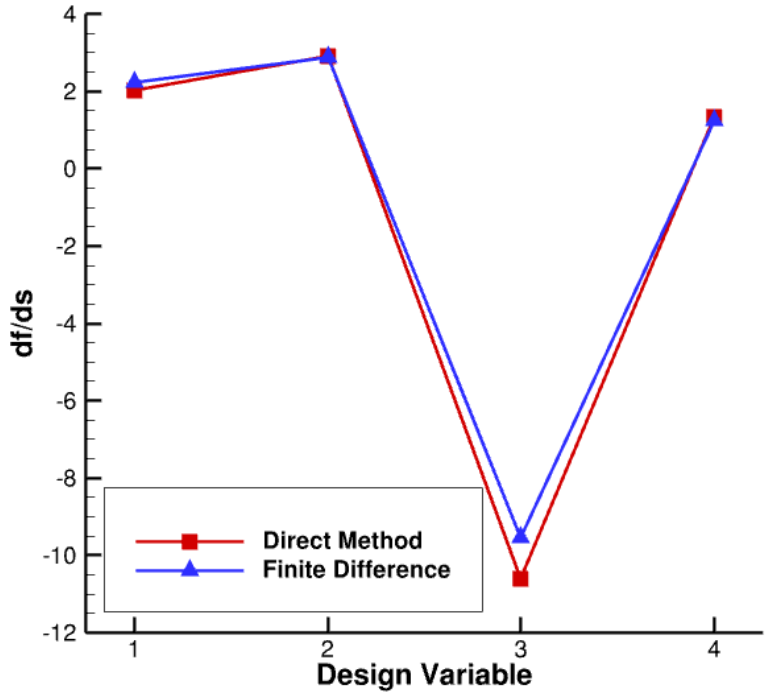

Figure 4.14: Comparisons of the objective gradient for the aero-only scramjet case.

ity analysis is executed multiple times during the iterative aero-thermo-elastic coupled sensitivity analysis.

4.4.2 Aero-thermo-elastic

Next, after assessing the standalone aerothermodynamic sensitivity analysis for a scramjet, the coupled forward analysis is used to determine the deformed configuration and resulting flow-field. Note that a finite-difference is used to compute the thermo-structural sensitivities. The molybdenum alloy, TZM-C03, is applied to the entire, filled interior geometry with temperature dependent material properties found in Ref. [65]. The resultant thermo-structural deformation computed through the forward analysis is shown in Figure 4.15. The majority of the deformation for the cowl and forebody is due to the increased temperature and thermal expansion within the body, resulting in a deflection of about 2 mm. Between the changes to the wall temperature and

the deformation, the separation bubble size and shock train have adjusted when FTSI is considered [52]. The wall pressure for the cold-wall aero-only case and this FTSI case are shown in Figure 4.16 as an indicator for the difference in shock-boundary layer interaction locations, often measured by the sharp increase in pressure in that area.

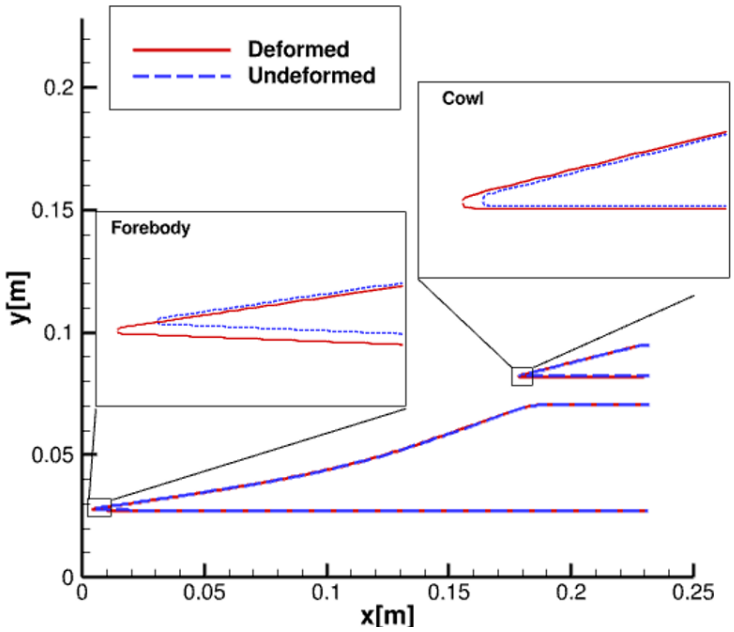

Figure 4.15: Deformation of the scramjet inlet for the coupled forward analysis.

Similar to the aero-only case described above, the wall variables are assessed to help benchmark this approach. Unlike the aero-only case, though, these quantities of interest are used in the coupled SA as inputs to the thermo-structural SA. This reinforces the need to have accurate sensitivities along the wall, since more than just an objective or constraint is dependent on it. Figure 4.17 shows the wall pressure sensitivities along the inlet for each design variable. The coupled direct method implementation is compared against a coupled finite-difference method for the original geometry. Overall, good agreement is found between the two methods – RMSPE of 5%, 14%, 20% and 10% for the forebody wall pressure sensitivity are found for each design variable, respectively – indicating the proper implementation of the direct method technique. There are still minor discrepancies in the magnitudes at certain peaks, but this is to be expected since this was

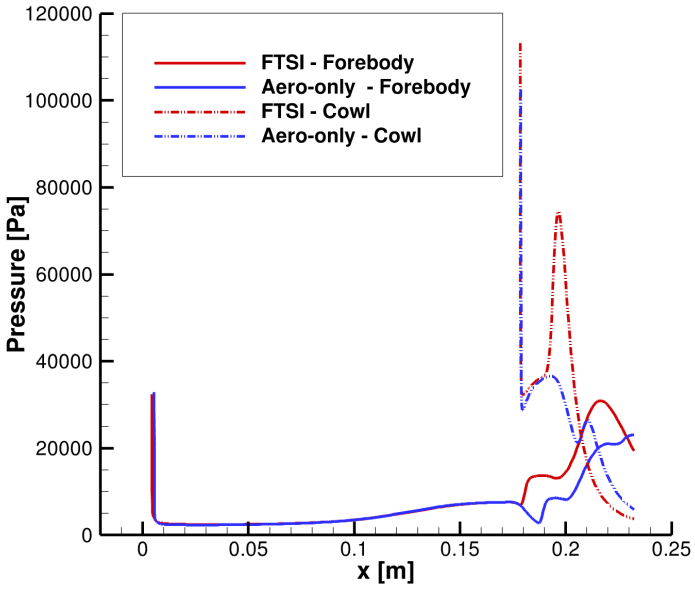

Figure 4.16: Comparison of the pressure along the wall for the aero-only and FTSI cases.

seen in the aero-only cases, which this method is built upon.

In addition to the wall quantity sensitivities being of high importance, the accuracy of the outflow for the inlet is imperative for an optimization. The outflow quantities are directly used to compute the performance of the inlet, the conditions of the combustion chamber and nozzle, and the engine performance. Therefore, the sensitivity of the pressure along the outflow is evaluated in Figure 4.18. Again, there is good agreement between the finite-difference solutions and the coupled direct method, indicating that the crucial flow features have been properly accounted for. The RMSPE for the four design variables are 13%, 9%, 12% and 8%, respectively. These pressures, along with other outflow values such as density and velocity, feed into the computation for the compression ratio and total pressure recovery, defining the performance and efficiency of the inlet.

Continuing the example objective from the previous section, the sensitivity of the reciprocal of the total pressure recovery is compared for the finite-difference and direct method solutions in Figure 4.19. The two methods are shown to again be in good agreement with an RMSPE of 9.5%, further endorsing the sensitivity analysis for the optimization algorithm.

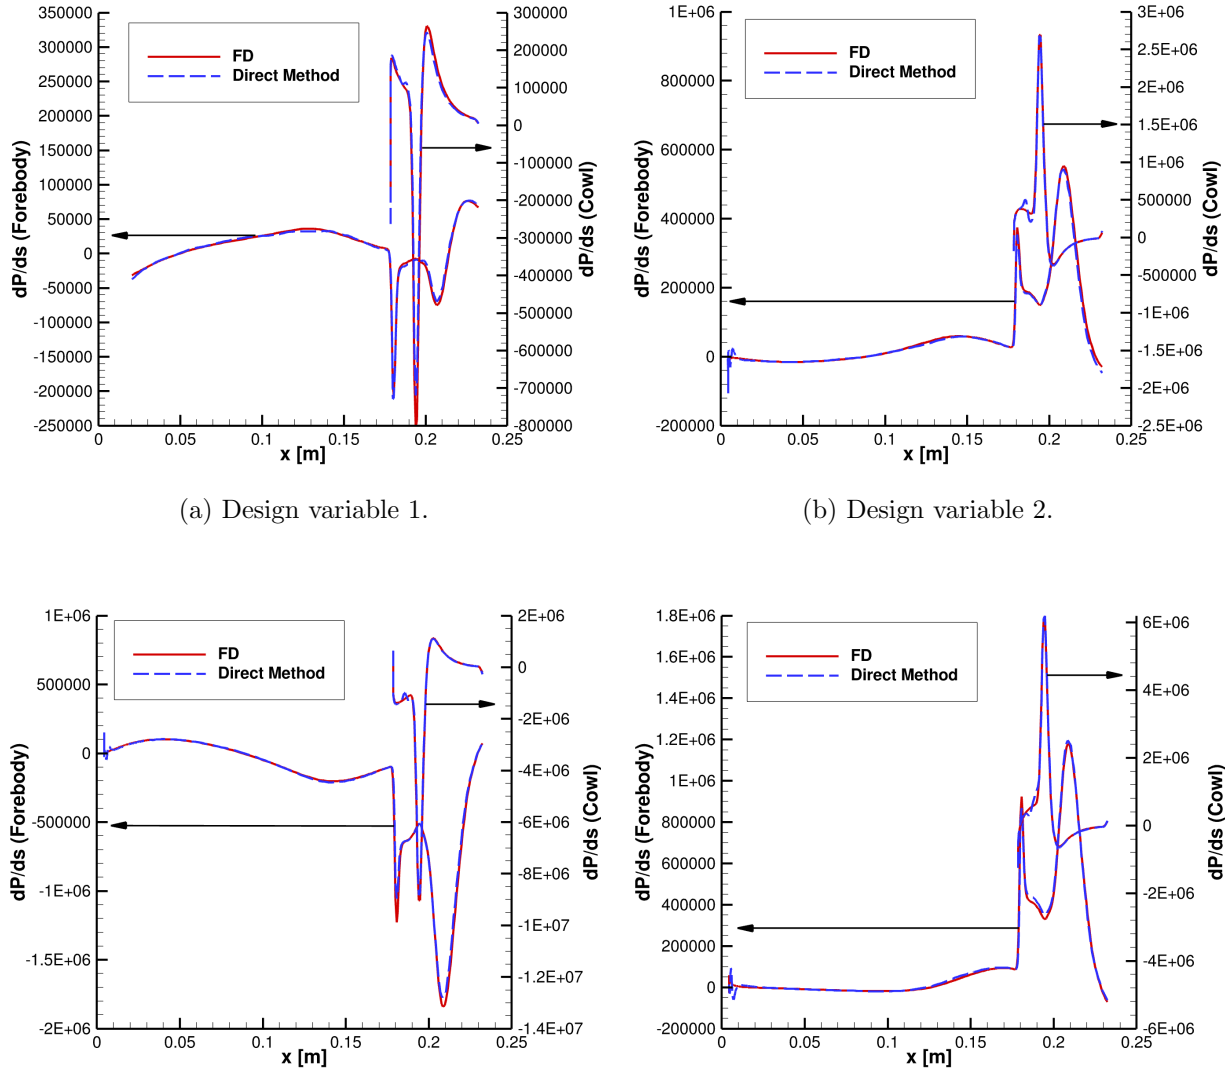

(a) Design variable 1.

(b) Design variable 2.

(c) Design variable 3.

(d) Design variable 4.

Figure 4.17: Sensitivity of wall pressure to the design variable.

4.5 Chapter Summary

The purpose of this chapter was to benchmark the new sensitivity analyses, both the aero-only and the aero-thermo-elastic, so they can be incorporated into an optimization process. This effort began by verifying that the direct method, when applied to a simple IRV-2 nose-cone geometry, produced the same results as a finite-difference method. The aerothermodynamic-only analysis

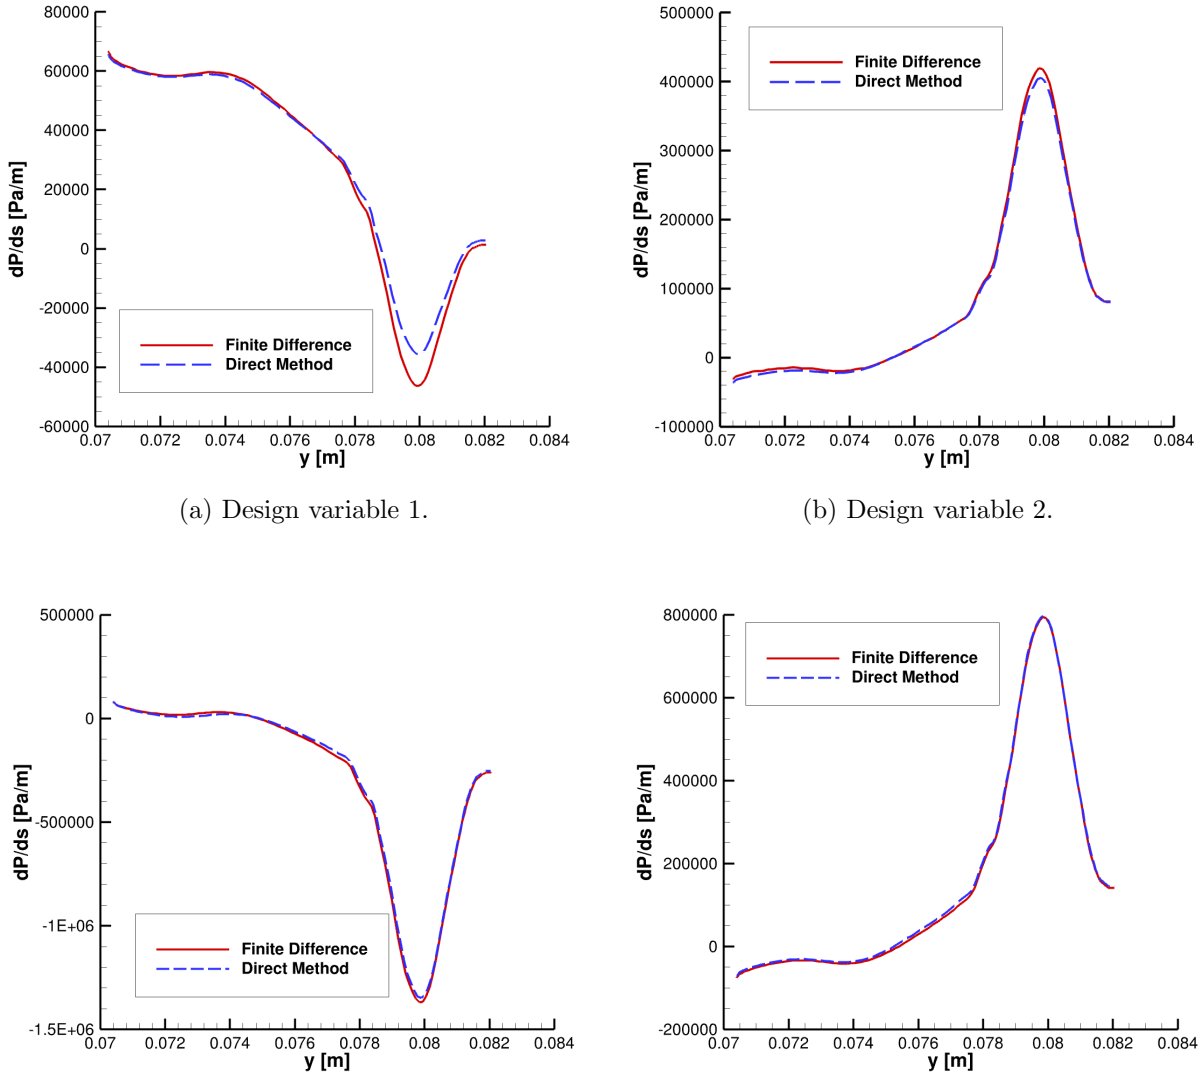

(a) Design variable 1.

(b) Design variable 2.

(c) Design variable 3.

(d) Design variable 4.

Figure 4.18: Sensitivity of outlet pressure to the design variable.

showed that the flowfield, wall properties, and objective sensitivities all agreed well with the finite-difference results. In addition, the study showed the computational benefits of leveraging the tools already available in the fluid solver, including the approximated Jacobian and tri-diagonal solver technique. By incorporating the two into the GMRES algorithm to assist in solving the preconditioned set of linear equations, the algorithm is able to converge 12 orders of magnitude in just 45 iterations for this case. Overall, the direct method computed the sensitivities 100 times faster than

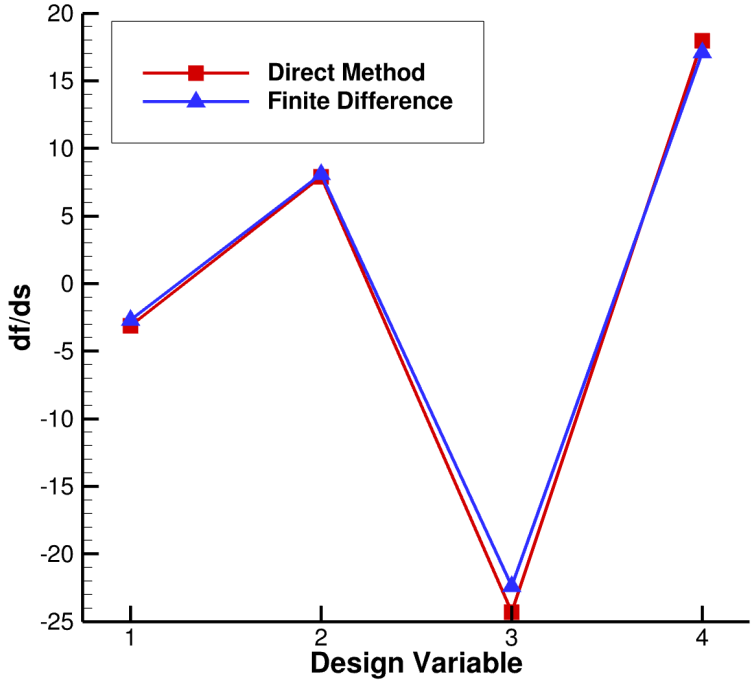

Figure 4.19: Comparisons of the objective gradient for the aero-thermo-elastic scramjet case.

the finite-difference approach. Next, the aero-thermo-elastic sensitivity analysis was benchmarked for both thermo-structural sensitivity methods for this same geometry. Minor deformations were observed, resulting in a need to incorporate the thermo-structural sensitivities into the analysis. Once again, the wall quantities and objective sensitivities were compared against a finite-difference solution and showed good agreement between the two for both methods.

One goal of this work is to perform a multidisciplinary optimization on a scramjet inlet, therefore the second geometry analyzed for the sensitivity analysis was an inlet. Both an aero-only and an aero-thermo-elastic analysis were performed to confirm that both properly match a finite-difference method. Pressures along the wall and at the outlet were compared for both sensitivity analyses. Additionally, an example of an objective gradient (the reciprocal of the total pressure recovery) was evaluated for both the aero-only and FTSI sensitivity analyses. All quantities showed very good agreement with the finite-difference approach, while reducing the computational cost by

almost 50×. Therefore, these sensitivities can be incorporated into an optimization procedure for a scramjet inlet, which is shown in Chapter 5.

Chapter 5

Hypersonic Shape Optimization

5.1 Chapter Introduction

One major complication with hypersonic vehicles is that many systems need to be designed in unison, e.g. the OML of the vehicle has a direct impact on the engine inlet. This often leads to contending or conflicting requirements that do not have a simple, intuitive solution. Instead, numerical optimization is increasingly popular to assist in designing these vehicles, allowing for the requirements to directly influence the design choice. Another issue with hypersonics in particular is that assumptions such as laminar flow or perfect gas break down, requiring CFD solutions to accurately model the flow instead of analytical solutions. Therefore, there is a need for integrating the high-fidelity CFD into the design and optimization processes, since the flowfield can substantially impact requirements, such as temperature and structural limits of a material, unstart of an inlet, and aerodynamic performance.

This chapter begins by applying a hypersonic aerothermodynamic optimization approach to a re-entry vehicle, IRV-2, in order to reduce the drag along a trajectory with imposed constraints acting as requirements. The vehicle is chosen because a previous optimization was performed using LeMANS for the nose-cone portion; however, a finite-difference approach was implemented and a one-temperature model was assumed for that previous study [35]. The optimizations performed in this chapter build on the previous work, but extend the analysis beyond the nose of the vehicle. The results provide several optimized geometries for different points along the trajectory; then each optimized solution is evaluated along the trajectory to help inform design decisions.

The second half of this chapter focuses on optimizing a scramjet inlet, both ignoring and including aero-thermo-elastic effects. The two are done in juxtaposition to highlight the importance of including the deformation in the design process. The UNSW experimental inlet is once again used and both the aerothermodynamic and aero-thermo-elastic sensitivities have been benchmarked for this configuration and flight condition in Chapter 4. Because the inclusion of aero-thermo-elastic effects is more expensive, two aerothermodynamic-only cases are evaluated to assess if there is a method capable of producing more similar results to the aero-thermo-elastic case. The cases differ by the thermal boundary condition: an isothermal wall boundary condition or a prescribed temperature profile based on the aero-thermo-elastic solution of the original geometry.

5.2 IRV-2 Optimization

Re-entry vehicles experience tremendous amounts of heating upon plunging into the thicker parts of the atmosphere. For certain types of re-entry vehicles, the range may be the primary goal, which would require reducing the drag and/or increasing the lift of the vehicle to maintain speed. At those speeds, though, the vehicle can experience increased amounts of heating and loads, which can lead to structural failures. Often, there are material and structural requirements counteracting the objective to decrease the drag. Therefore, these types of re-entry vehicles are a prime opportunity to employ optimization. Further, the high Mach number flows cause non-equilibrium behavior in the flowfield that must be captured with CFD.

A specific example is the IRV-2 vehicle, shown in Figure 5.1. This 1.4 meter-long, biconic configuration was flown and achieved Mach numbers of around 22. Many previous works have studied this vehicle to understand the ablation surrounding the nose [71, 72]. More recently, work by Eyi et al. utilized LeMANS to optimize the nose-region of the vehicle to minimize drag using the modified method of feasible directions algorithm and a finite-difference sensitivity analysis [35]. The study was able to successfully reduce the drag contribution of the nose by over 25% for several trajectory points, while including constraints on the total heat rate, maximum wall temperature, maximum wall pressure and maximum heat flux. The authors then built on this work to include

Figure 5.1: IRV-2 vehicle [72].

multiple objectives: reduce the drag and the heat flux of the nose [36].

The nose portion included in the study performed by Eyi et al. only included the first 2.5 cm of the vehicle. Although significant reductions in drag were seen by that portion of the vehicle, it is unclear how that translates to the drag contribution overall. Therefore, the current work expands on the research performed by Eyi et al., using a similar setup to the first study by optimizing for several trajectory points [35]. Eyi et al. used a one-temperature model, but Figure 5.2 shows that two-temperature modeling (which captures the vibrational nonequilibrium) has an influence on the heating. Therefore, a two-temperature model is used for this study and the optimization can not be exactly compared to the previous study.

5.2.1 Test Setup

A gradient-based optimization technique is implemented in this work due to the high computational cost of CFD. The algorithm utilized is a sequential quadratic program from Python's

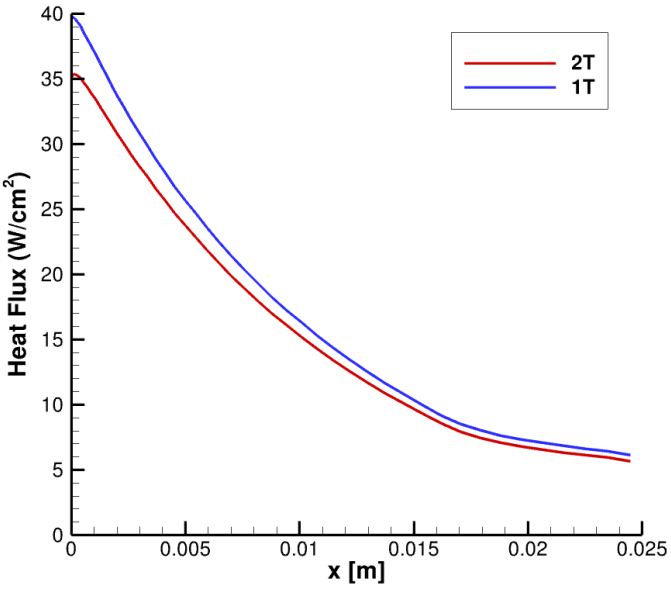

Figure 5.2: Heat flux along the IRV-2 nose-cone for one-temperature (1T) and two-temperature (2T) models.

optimization toolbox, which allows for box constraints for the design variables. The optimization problem solved in this work is summarized as follows:

$$\min_{s} \quad f(\mathbf{x}(\mathbf{s}), \mathbf{q}(\mathbf{s})) \tag{5.1}$$

$$\text{s.t.} \quad \mathbf{g}(\mathbf{x}(\mathbf{s}), \mathbf{q}(\mathbf{s})) \leq 0 \tag{5.2}$$

$$\underline{s}_i \leq s_i \leq \bar{s}_i \qquad i = 1, ..., n_s \qquad . \tag{5.3}$$

The function, f, is the objective, which can depend on the design variables, $\mathbf{s}$, through the mesh node locations, $\mathbf{x}$, and implicitly through the state variables, $\mathbf{q}$. For the current work, the objective is the drag, the design variables are control points for the Bézier curves used to alter the shape of the geometry, as described in Section 4.4, and the state variables are the flowfield variables. The optimization problem is also subjected to inequality constraints, $\mathbf{g}$, to maintain vehicle requirements such as maximum heating, pressure and temperature along the surface of the vehicle. Finally, the problem is subjected to bounded constraints, where $\underline{s}$ and $\bar{s}$ are the lower and upper bounds,

respectively, for all n_s design variables. In addition, the state variables, $\mathbf{q}(s)$, in Equations 5.1 and 5.2 are determined such that the governing equations are satisfied:

$$\mathcal{R}(\mathbf{x}(\mathbf{s}), \mathbf{q}(\mathbf{s})) = 0 \qquad , \tag{5.4}$$

where $\mathcal{R}$ is the residual vector for the governing equations of the CFD simulation. This formulation requires a converged solution, $\mathbf{q}$, from the flow solver for each design, $\mathbf{s}$.

The objective drag, D, is both explicitly dependent on the design variables due to the dependency of drag on the physical shape and implicitly dependent due to the dependency of drag on the state of the flowfield at the interface, as shown in Equation 5.5:

$$D = 2\pi \sum_{j=0}^{n_w} A_j[(P_\infty - P_j)n_{x,j} + \tau_{xx,j}n_{x,j} + \tau_{xy,j}n_{y,j}] \tag{5.5}$$

The area, A_j, and the component of the normal directions, $\mathbf{n}_j$, for all n_w wall segments are both explicitly dependent on the shape and, therefore, the design variables. The pressure, P_j, and the component of the shear forces, τ_j are functions of state variables at the wall and are implicitly dependent on the design variables. Additional inequality constraints are also implemented in the form of quadratic penalty functions:

$$f = D + \sum_{k=1}^{n_g} \rho_k g_k \qquad , \tag{5.6}$$

where ρ_k is a penalty magnitude, greater than zero, for all n_g constraints. These constraints include limiting the temperature (T), pressure (P), heat flux (q) and total surface heat transfer rate (Q) to the original configuration's maximum values of each. Additionally, a minimum volume (V) constraint is applied, limiting the volume to be at least that of the initial geometry. The pressure constraint is computed as $g_k = (\min(0, 1 - \frac{\max(P)}{\max(P_0)}))^2$; the temperature and heat flux constraints are computed similarly. The total surface heat transfer rate is computed as $g_k = (\min(0, 1 - \frac{Q}{Q_0}))^2$; the total volume constraint is computed similarly. The box constraints around the design variables are set at ([-1.0,1.0]) for all design variables and are only implemented to restrict the optimization

from beginning the line search with design variables that are too large; these constraints are not active for the optimized solutions.

5.2.2 Benchmarking: Optimization Approach

The benchmarking of the sensitivity analysis provides confidence in the gradient calculation needed for the optimization algorithm. The final check is that the flowfield sensitivity analysis is properly implemented within the optimization process. To begin, the drag is optimized under the constraint that the maximum pressure does not exceed that of the original geometry (5450 Pa) using only four design variables. Restricting the design changes to maintain a smooth leading edge, the first control point and the second control point's x-coordinate position are fixed. The y-component of the second control point is then the first design variable. The third control point can be adjusted in both the x- and y-directions. Finally, the last design variable is the x-component of the third control point, whereas the y-component is fixed to allow for a smooth transition to the rest of the vehicle. The optimization process uses two approaches for the gradient calculation: finite-difference and the direct method. The optimal solution to both is shown in Figure 5.3, which shows great agreement between the two methods (the percent difference in drag is <1%) and a drag reduction of approximately 21% for both. Additionally, the surface pressure constraint is satisfied for both cases. The shape begins to narrow around the nose to reduce the drag, but the pressure constraint limits the contraction. This optimal shape causes extra heating compared to the original geometry, so the addition of heating and temperature constraints may be needed for a more realistic design.

The addition of the constraint on the heat flux satisfies both the max heat flux and temperature requirements, since the two are linked through the radiative equilibrium boundary condition. The heat flux constraint is implemented similarly to the pressure constraint using a quadratic penalty term. The maximum heat flux is limited to $4.5 \times 10^6\ W/m^2$, which is the maximum heat flux of the original geometry. The inclusion of this constraint limits the shape change available to reduce the drag, as seen in Figure 5.4. Note that for this optimal configuration, the heat flux

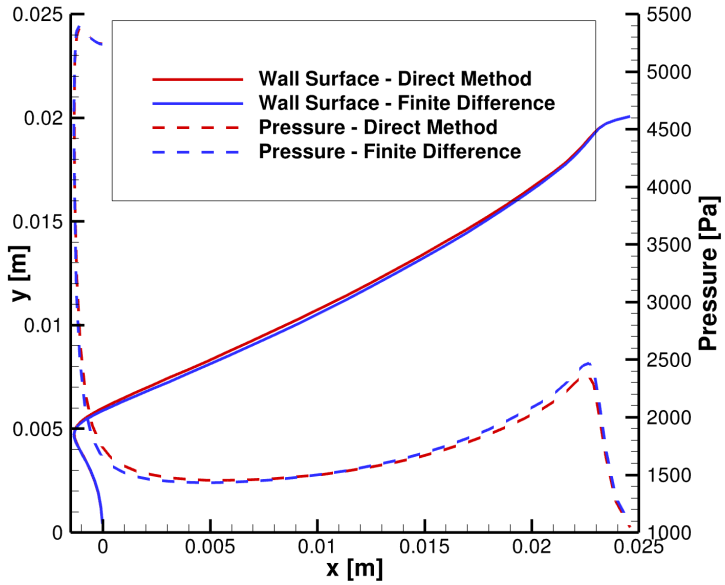

Figure 5.3: Comparisons of the optimized geometry using the finite-difference and the direct method sensitivities with a pressure constraint.

constraint is active while the pressure constraint is inactive. Additionally, the drag decreases by approximately 16% using this optimal geometry.

As an additional comparison between the two gradient methods – finite-difference and direct method – the drag reduction for each iteration is shown in Figure 5.5. Both follow a similar path in the course of the optimization process, with a slight divergence during the third iteration. This difference is likely due to numerical inaccuracies since both have features of finite-differencing in the analysis, such as the residual's sensitivity to the mesh in the semi-analytical approach. The two methods also converge to the same drag coefficient, with an error of 0.05% for the final iteration, giving further confidence in the semi-analytical direct method.

5.2.3 Re-entry Vehicle Optimization Results

The direct method sensitivity analysis approach facilitates a more computationally feasible approach, enabling optimization of entire vehicles rather than just portions of the vehicle, such as the nose region. The IRV-2 vehicle is explored here, which is a sphere-biconic geometry with a

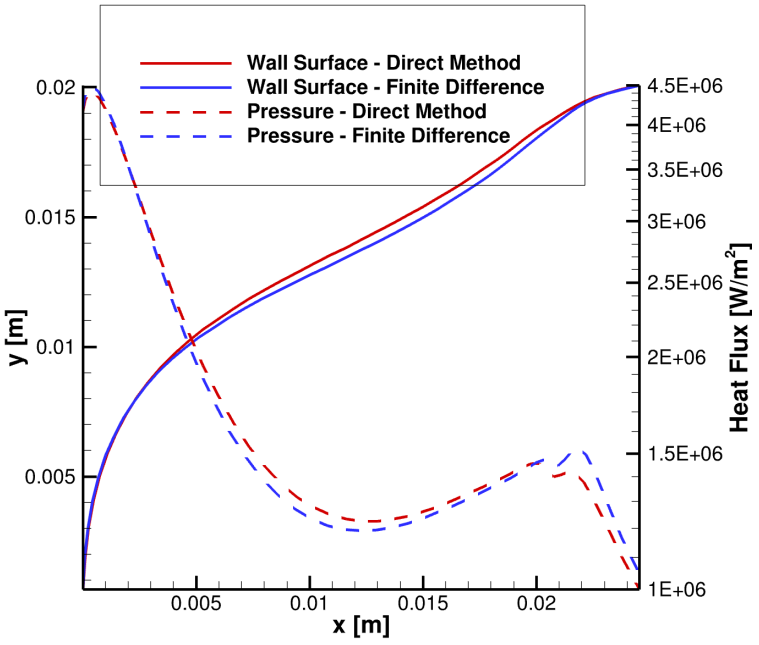

Figure 5.4: Comparisons of the optimized geometry using the finite-difference and the direct method sensitivities with a pressure and heat flux constraint.

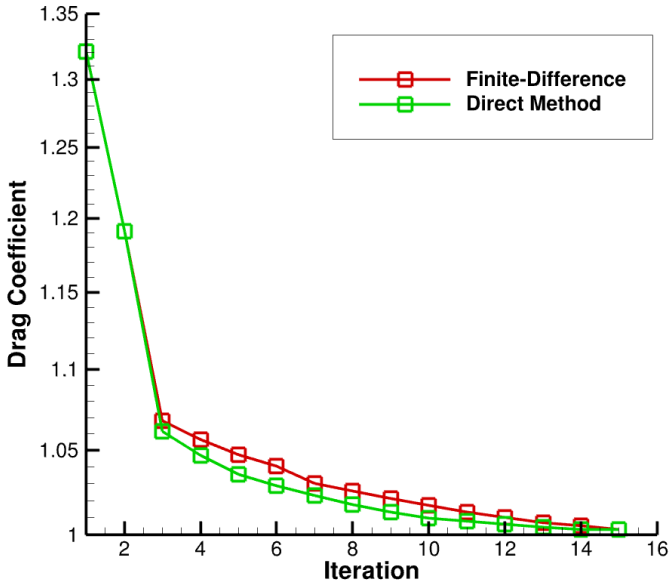

Figure 5.5: Convergence rates for the optimization procedure using finite-difference sensitivities and the semi-analytical direct method.

nose radius of 0.1905 m and a total length of 1.386 m. The first angle is 8.42 degrees, followed by 6.10 degrees at an axial location of 0.1488 m. Previously, Eyi et. al. [35] focused on reducing the drag contribution from the nose region of a re-entry vehicle [72]. Although the analysis showed promising results, the benefit of drag reduction may not be substantial for the entire vehicle. To first confirm Eyi et al.'s findings, only the nose region of the IRV-2 vehicle is optimized for drag. Four constraints are applied to the geometry as described in Section 5.2.1: the maximum wall pressure, temperature, heat flux and total surface heat transfer rate must be less than or equal to that of the original geometry. These initial results are obtained for trajectory point (TP) 1, with the first ten points shown in Table 5.1. According to Kuntz et al., these first ten trajectory points are all considered laminar and are treated as such in this analysis [71]. In order to determine the number of design variables necessary to capture the optimal shape, a study is performed with 4, 6, 8 and 10 design variables. The outcome for each is shown in Figure 5.6, each producing less than a 1% difference in drag coefficient compared with the 10 design variables case; 8 design variables are chosen for this initial study. The design variables are the x- and y-coordinates of the control points for the Bézier curves, except the x-coordinate of the first control point and the y-coordinate of the last control point, which are fixed.

Table 5.1: Freestream conditions for each trajectory point (TP) [72].

TP	**Time (s)**	**Altitude (m)**	**Velocity (m/s)**	**Temperature (K)**	**Density (kg/m^3)**	**Mach number**
1	0.00	66,935	6,780.6	227.81	1.2505×10^{-4}	22.41
2	4.25	55,842	6,788.3	258.02	5.0454×10^{-4}	21.08
3	6.75	49,290	6,785.2	270.65	1.1344×10^{-3}	20.57
4	8.75	44,042	6,773.0	261.40	2.2593×10^{-3}	20.90
5	10.25	40,108	6,752.4	250.35	3.9957×10^{-3}	21.29
6	11.50	36,836	6,722.0	241.50	6.4268×10^{-3}	21.58
7	12.50	34,229	6,684.3	234.30	9.5832×10^{-3}	21.78
8	13.25	32,283	6,644.9	228.76	1.3145×10^{-2}	21.91
9	13.95	30,480	6,596.7	226.91	1.7313×10^{-2}	21.84
10	14.75	28,236	6,527.1	224.73	2.4310×10^{-2}	21.71

The optimal shape of the nose region using various constraints is represented in Figure 5.7a. Additionally, the pressure, temperature and heat flux, which define three of the constraints, are plotted along the wall for the original geometry and each optimal solution in Figures 5.7b-5.7d.

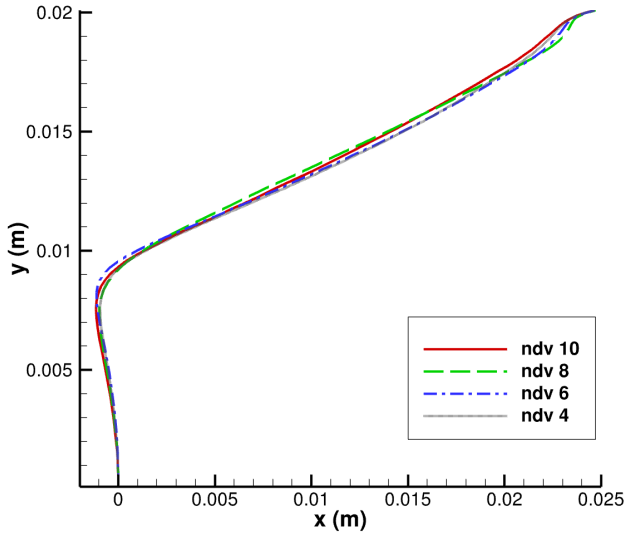

Figure 5.6: Optimized nose shapes varying the number of design variables.

Without constraints, the drag is reduced by 29.2%; with only a pressure constraint, the drag is reduced by 26.1%; with a pressure, heat flux and temperature constraint, the drag is reduced by 23.6%; and with a pressure, heat flux, temperature and total surface heat transfer rate constraint, the drag is reduced by 21.3%. Therefore, even with the constraints limiting the shape change, a substantial reduction in drag can be created in the nose region. The optimal shape with all the constraints is a flattened region at the front of the nose, followed by a more linear region connecting to the side wall. These nose shapes are consistent not only with the previous study using LeMANS by Eyi et al. [35], but also by Seagar and Agarwal [101]. Additionally, concave and flatter nose shapes were previously tested experimentally in the late 1950s at NASA Langley [78, 50, 8, 19].

Focusing on the optimal solution that includes all four constraints, a comparison of the pressure contours for the original and optimal geometry is shown in Figure 5.8. The blunted shape of the optimal geometry causes the shock standoff distance to increase, which can be seen by comparing the two black lines representing the sonic line. The increase to the shock standoff distance decreases the temperature of the flow surrounding the wall around the stagnation region. Meanwhile, the pressure along the blunt surface is higher than the original geometry, but quickly

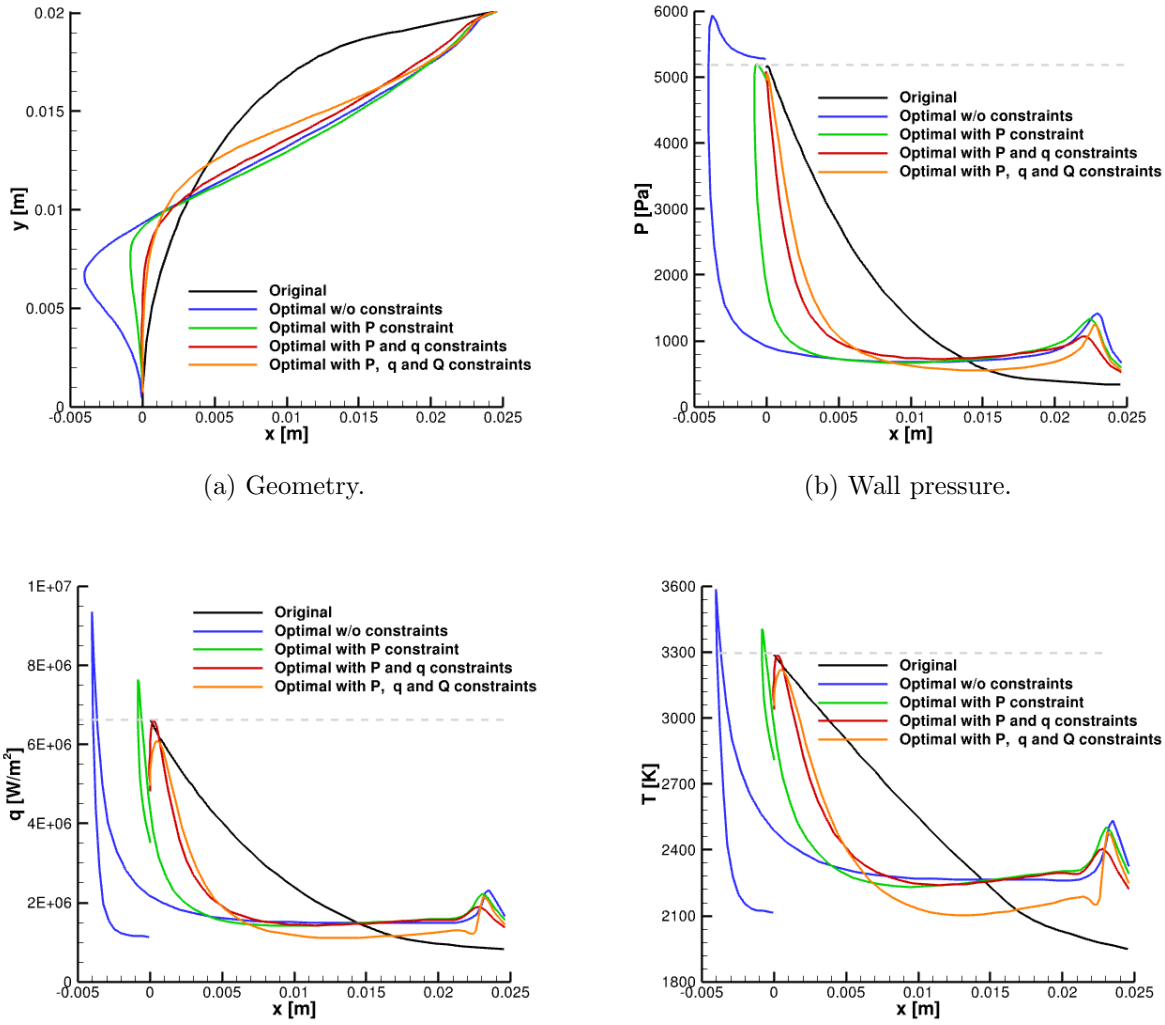

(a) Geometry.

(b) Wall pressure.

(c) Wall heat flux.

(d) Wall temperature.

Figure 5.7: Optimized nose geometry and wall properties.

declines as the flow expands around the sharper shoulder. Therefore, the roughly linear region between the shoulder and the end of the nose overall has a lower pressure than the original geometry. The exception to that is right at the end of the nose, where there is a slight non-linearity in the shape that causes a small increase in pressure, that will become increasingly important in the following results.

In order to assess whether the substantial decrease in drag for the nose translates to the

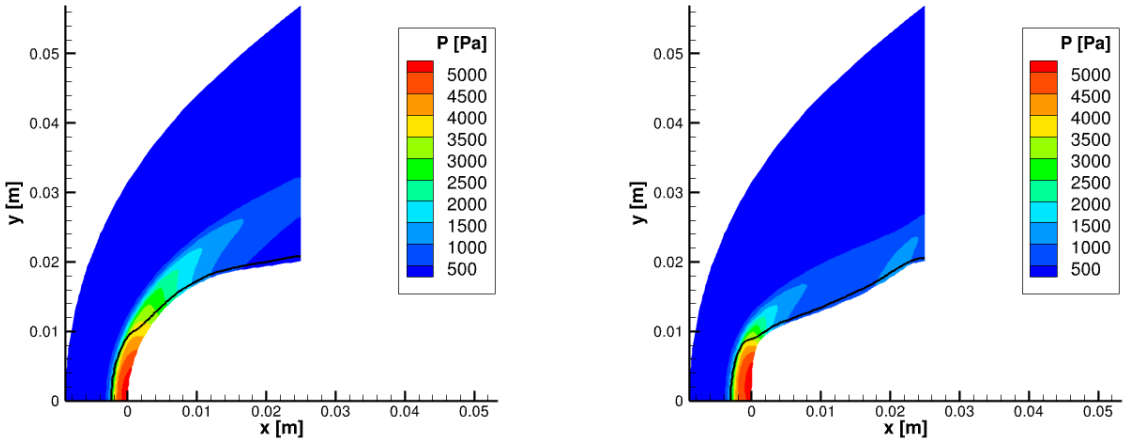

(a) Original geometry.

(b) Optimized geometry.

Figure 5.8: Pressure contours within the flowfield and the sonic line indicated as the black line.

entire vehicle, the optimal nose including the constraints is applied to the nose region of the IRV-2 vehicle. The entire vehicle is then analyzed and evaluated for the change in drag. The geometry with an optimized nose increases the drag by approximately 1.1% when the entire vehicle is taken into account. The leading contributor to this difference is the increase in pressure along the wall aft of the nose, as shown in Figure 5.9. The non-linearity at the end of the nose causes an increase to the pressure, which continues along the entire length of the vehicle. The consistent increase in pressure along the wall causes additional drag due to pressure, outweighing the initial loss of drag from the optimized nose. In order to ameliorate this concern, the entire length of the vehicle needs to be included in the optimization process to properly account for the contribution of drag along the sidewall.

To determine an appropriate mesh for the optimization study, a mesh convergence study is performed for the first trajectory point for the vehicle based on the Richardson extrapolation method [23]. Three sets of grids are used in the present study: 350 $\times$ 35, 350 $\times$ 100 and 600 $\times$ 100. The differences in the pressure along the wall can be seen in Figure 5.10, indicating that the coarse mesh is not able to accurately capture the pressure along the sidewall. In turn, the coarse

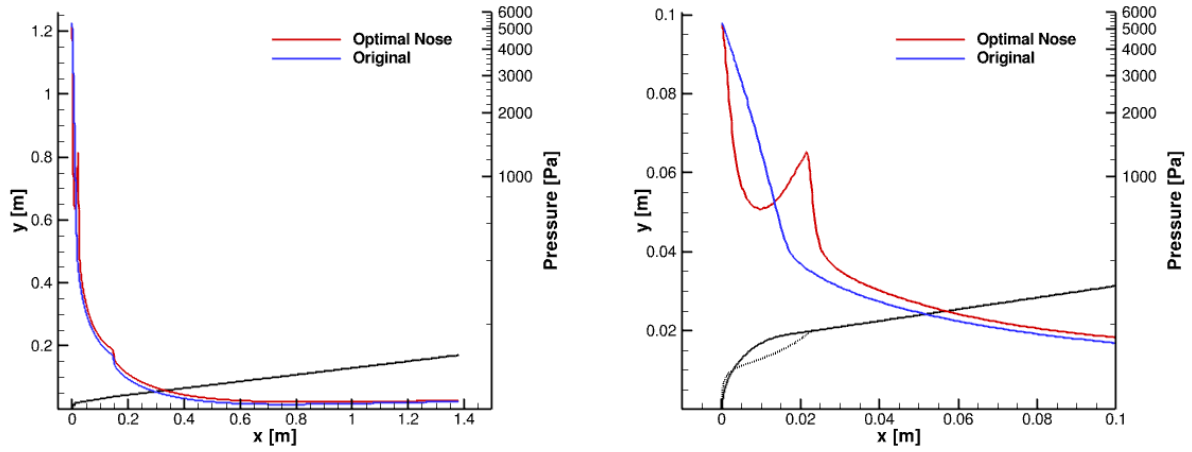

(a) Full body.

(b) Zoomed in.

Figure 5.9: Wall pressure comparisons between original geometry (black solid line) and optimized nose geometry (black dashed line).

mesh produces the largest error compared to the finest mesh when considering the drag coefficient as the quantity of interest. Overall, the approximate error for the 350 × 100 grid is 0.38% and the extrapolated error and grid convergence index (GCI) are less than 0.01%. Therefore, the 350 × 100 grid is chosen for the optimization study.

When including the entire vehicle, a larger portion of the vehicle can be modified to optimize the geometry. Three different optimal geometries are found using different regions of the vehicle that the design variables can control. Further restrictions are necessary to maintain a realistic vehicle, such as internal volume. Therefore, another constraint is included for these optimizations to maintain an adequate total volume for a potential payload, as described in Section 5.2.1. The previous constraints are also still included and the number of design variables is still set to 8. The first optimization only allows the nose region (indicated in blue in Figure 5.11 and defined as less than 0.025 m axially) to deform so that a comparison with the previous analysis is obtained. The results show that the optimal geometry is able to reduce the drag by 1.1% while still maintaining the constraints. The second optimization allows the front end of the IRV-2 vehicle to deform, which appends the green region addition indicated in Figure 5.11 to the nose region. This region is defined

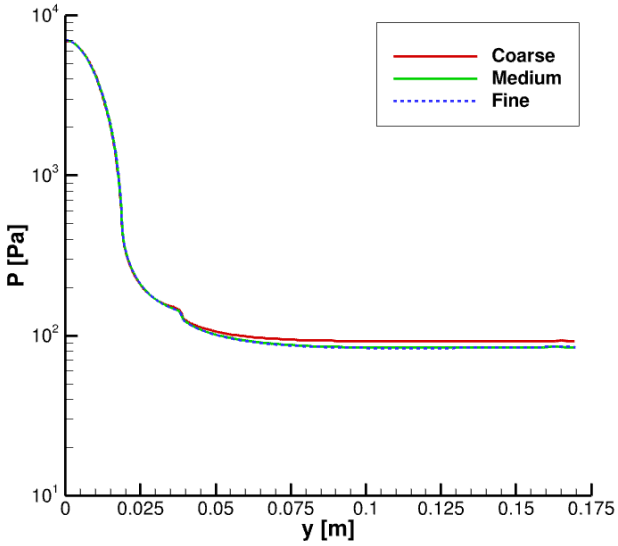

Figure 5.10: Mesh discretization results for wall pressure.

as less than 0.14 m axially, which is approximately where the change in angle occurs for the original geometry. By defining a larger region for the design variables and allowing more deformation along the vehicle, the optimal configuration is able to reduce the drag by 3.4%. Finally, the entire vehicle is allowed to deform, including the red region addition in Figure 5.11, resulting in a decrease of 5.7% in the drag.

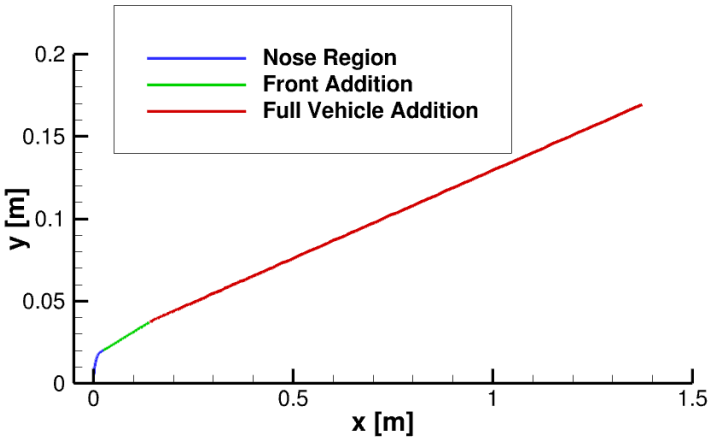

Figure 5.11: Regions of deformation for each optimization.

The three optimal configurations are shown in Figure 5.12, which illustrates a very different behavior for each optimal solution. When only the nose is allowed to deform, a blunter shape appears as compared to the previous results in Figure 5.7a. The major difference is that the new configuration results in a higher drag around the nose, but an overall lower drag once the sidewalls are accounted for. As an indicator of the drag, the wall pressure is shown in Figure 5.13, which shows a higher pressure at the nose but a lower pressure along the sidewalls. Similarly, when the front of the vehicle is allowed to deform, an even blunter nose is observed, and higher pressure/drag is produced at the nose. However, the lower pressures along the wall post-nose offset the increase in the drag due to the nose bluntness. Finally, when the entire vehicle is allowed to deform, the optimal configuration shows less bluntness at the nose compared to the other two optimal configurations. Instead, the decrease in drag is created downstream when the geometry curves to produce expansion along the aft half of the sidewall; this produces lower pressures and in turn, reduces the drag experienced by the vehicle. Because the entire vehicle is considered during the optimization process, a different (and more effective) mechanism for reducing drag is observed. Note that the non-smooth features in the wall pressure appearing in Figure 5.13 are signs of geometry transitions, such as the angle change at x=0.1488 m and the end of the deformation region.

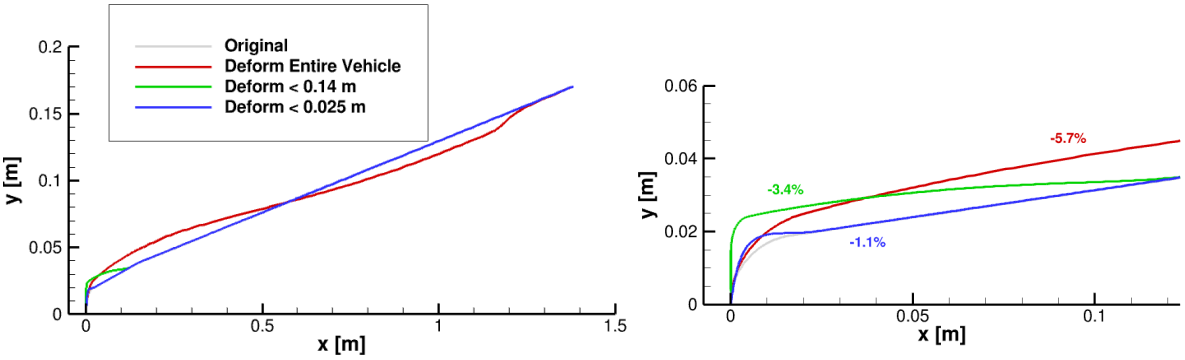

(a) Full body.

(b) Zoomed in.

Figure 5.12: Optimized shapes for varied regions defined by the design variables at TP1.

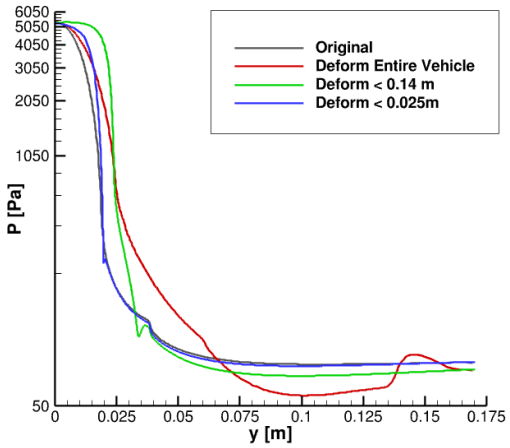

Figure 5.13: Wall pressure comparisons between original geometry and the optimized shapes at TP1.

Due to the benefit of reducing the drag, the entire vehicle is used in the subsequent optimizations. To attain a more comprehensive perspective on how a flight vehicle can be improved using these techniques, optimizations are performed at three trajectory points: TP1, TP5 and TP10. For all three cases, the number of design variables is set to 8 and constraints on the pressure, temperature, heat flux, total surface heat transfer rate and volume of the vehicle are imposed. The drag reduction for the optimized configurations is shown in Table 5.2, along with which constraints are active for the optimized geometry. The convergence for the optimization process for each of the trajectory points is shown in Figure 5.14, which reflects that the first step contributes the majority of the drag reduction. Additionally, the optimal configurations for each of the trajectory points are shown in Figure 5.15.

Table 5.2: Drag reduction and constraints for optimized geometries.

TP	**Drag Reduction**	**Pressure Constraint**	**Temperature Constraint**	**Heat flux Constraint**	**Heating Constraint**	**Volume Constraint**
1	5.7%	Inactive	Inactive	Inactive	Inactive	Active
5	1.5%	Active	Active	Active	Inactive	Inactive
10	2.6%	Active	Active	Active	Inactive	Inactive

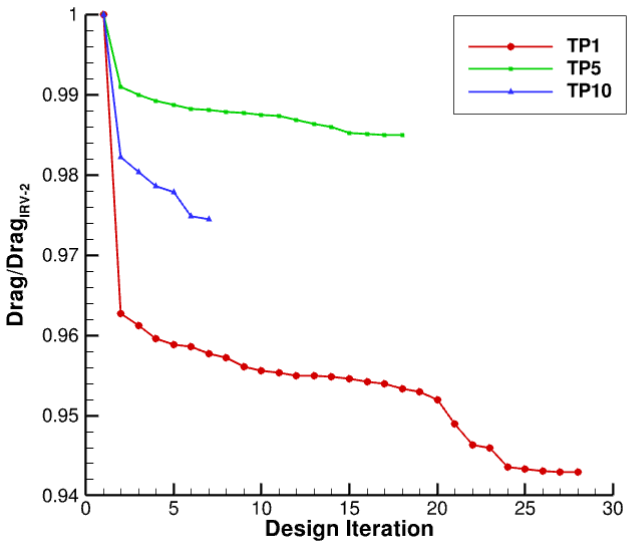

Figure 5.14: Objective function convergence rates for each trajectory point.

The first trajectory point's optimized configuration produces the largest reduction in drag in terms of percent difference. Additionally, the constraint that is active at the optimum is the volume constraint, whereas the other two optimizations are bounded by the temperature, heat flux and pressure constraints. The difference can also be observed in the geometry shown in Figure 5.15a, where both TP5 and TP10 produce similar optimized shapes towards the front of the vehicle. In contrast, the optimal solution for TP1 produces a blunter nose and front section until an axial location of approximately 0.3 m. Behind that section, the shape curves inward, allowing for expansion to occur until the end of the vehicle, as shown in Figure 5.15b. The fifth trajectory point follows a similar trend, although with less curvature and expansion, hence less drag reduction. Both TP1 and TP5 have a non-linearity near the base of the vehicle, which causes the gas to compress and increase the drag in that area. Conversely, the tenth trajectory point produces an optimal geometry where the aft smoothly transitions to the end of the vehicle with minimal curvature, avoiding additional compression. The reduction in drag is primarily due to the initial expansion from the front of the vehicle.

To understand what is driving these differences, especially with the first trajectory point,

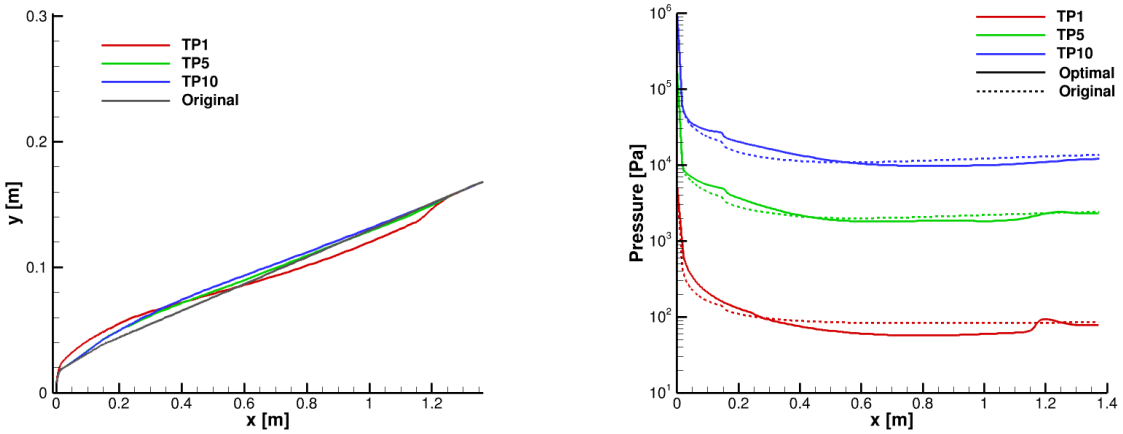

(a) Geometry (y-axis is not to scale).

(b) Wall pressures.

Figure 5.15: Optimized designs for the three trajectory points.

the initial gradient of the drag coefficient is studied for the original geometry for each trajectory point. The gradient for each design variable is shown in Figure 5.16, including both the implicit and explicit derivatives. A positive value of dC_D/ds indicates an increase in drag coefficient. To build intuition on how the design variables affect the shape of the vehicle, Figure 5.17 portrays a scaled version of perturbing each design variable, individually. The odd design variables (which are perturbations in the y-direction) show a distinct change to the shape, expanding normal to the wall. The even design variables (which are perturbations in the x-direction) show minimal changes to the original shape. This is reflected in Figure 5.16, where the even design variables produce gradients close to zero, indicating little to no influence on the drag. Conversely, the odd design variables show a much greater impact. Focusing on the odd design variables, the total derivatives for each trajectory point show a monotonically increasing influence on the drag coefficient. However, TP5 and TP10 both show negative gradients for all the odd design variables, while TP1 shows positive gradients for the fifth and seventh design variables.

The cause for the apparent difference in the gradients is better understood by focusing on the implicit and explicit derivatives. The trend for the explicit derivative does not change for each

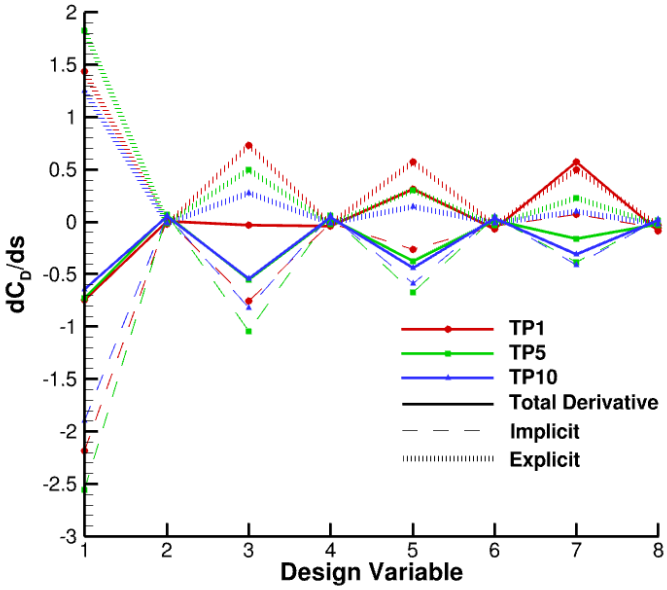

Figure 5.16: Initial gradient for the IRV-2 geometry for each design variable.

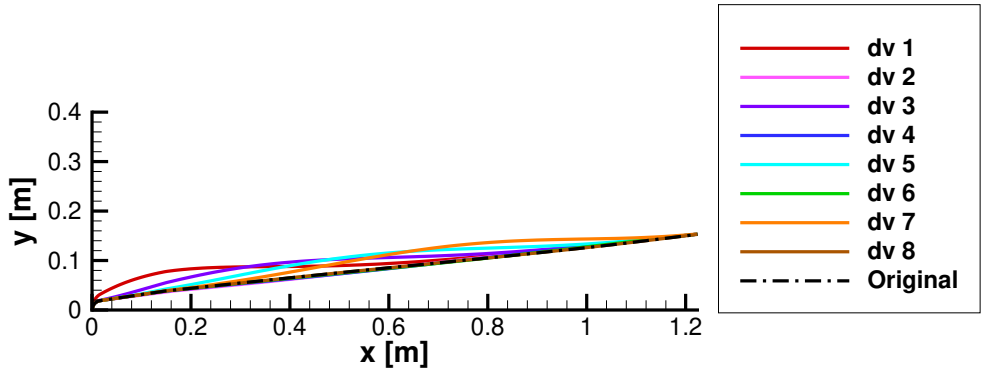

Figure 5.17: Deformations due to individual perturbations of each design variable for the IRV-2 geometry.

trajectory point. Instead, the difference appears in the implicit derivative, or the influence of the shape on the pressure and shear forces. Focusing on the seventh design variable, which is shown as cyan in Figure 5.17, a bump in the flow is observed on the aft section of the vehicle. For both TP5 and TP10 the implicit derivative is negative, indicating the expansion after the bump is able

to reduce the drag on the vehicle. Conversely, TP1 shows a positive gradient for the seventh design variable, indicating that the expansion after the bump is not able to offset the drag contribution caused by the compression of the bump. To better explore this phenomenon, the wall pressure sensitivity to the seventh design variable is shown in Figure 5.18. The shape of the wall sensitivity indicates that the bump from perturbing the design variable causes an initial compression around x = 0.6 m, followed by an expansion around 1.1 m that decreases the pressure. For TP5 and TP10, the magnitudes for the compression and expansion are roughly equivalent, but the aft section will be applied to a larger area, hence the decrease in drag. For TP1, which is at a higher altitude, there is not enough expansion to counteract the compression, and this leads to an overall positive drag gradient. This initial difference for the gradient of the drag coefficient is indicative of how the optimal shape forms, where an emphasis on blunting the front of the vehicle (defined by the lower numbered design variables) is seen for TP1, followed by a curve inwards in the aft section of the vehicle (defined by the higher numbered design variables).

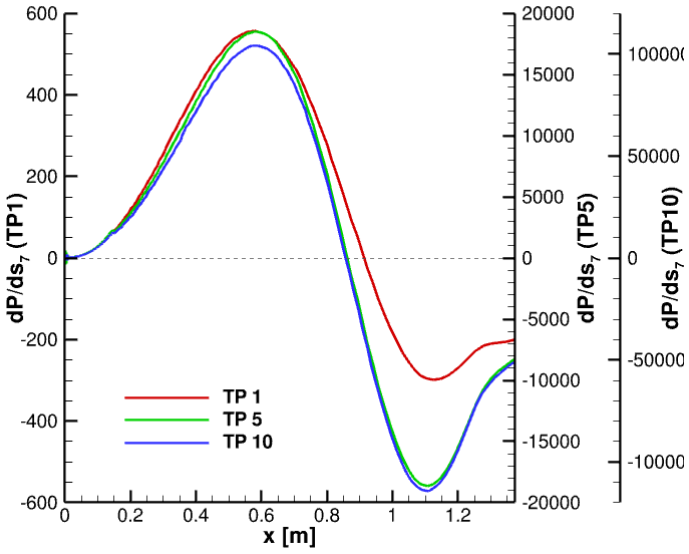

Figure 5.18: Wall pressure sensitivity to the 7th design variable for the IRV-2 case for each trajectory point.

The three optimal geometries are also evaluated at the first ten trajectory points for the

IRV-2 mission. A plot of the drag coefficient as well as the percent difference from the nominal geometry at each point is shown in Figure 5.19. The results for the optimal geometry found at TP1 show that it produces the largest reduction in drag for the first two trajectory points, but causes a large *increase* to the drag at the later points, reaching almost a 10% increase. The results for the optimal geometry found at TP5 consistently produce reductions in drag across the ten points, although only a 1% to 2% difference. Finally, the results for the optimal geometry at TP10 show a slight increase in drag at the early trajectory points, but a slight reduction for the remaining points. Depending on the mission requirements and trajectory, one of these optimal geometries may be more beneficial than the original geometry.

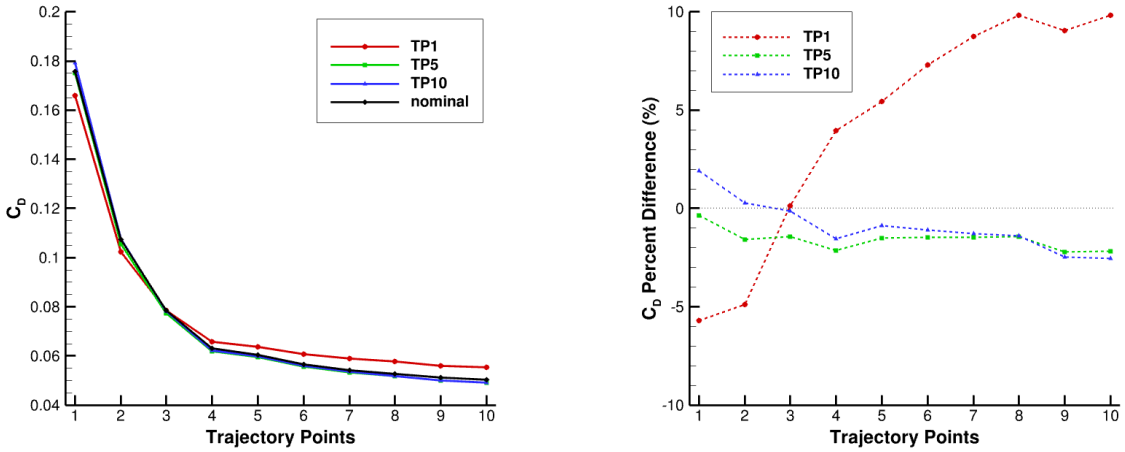

(a) Drag coefficient.

(b) Drag coefficient percent difference from nominal.

Figure 5.19: Drag performance for each design at different trajectory points.

Continuing the comparison of each design along the trajectory, Figure 5.20 shows the pressure and heat flux constraints, which are bounded as ≤ 1 at the design point. The optimal configuration designed at TP1 initially shows a pressure ratio below one, indicating that the constraint is not active, as stated previously in Table 5.2. However, the pressure for the next four trajectory points is larger than the original IRV-2 configuration. In terms of heat flux, the design from TP1 is also inactive and remains less than the original heat flux throughout the mission. Both designs from

TP5 and TP10 follow similar behaviors to each other, hovering around the same pressure and heat flux as the original geometry for the majority of the trajectory. Both the pressure and heat flux constraints are active at the design point for each optimized design.

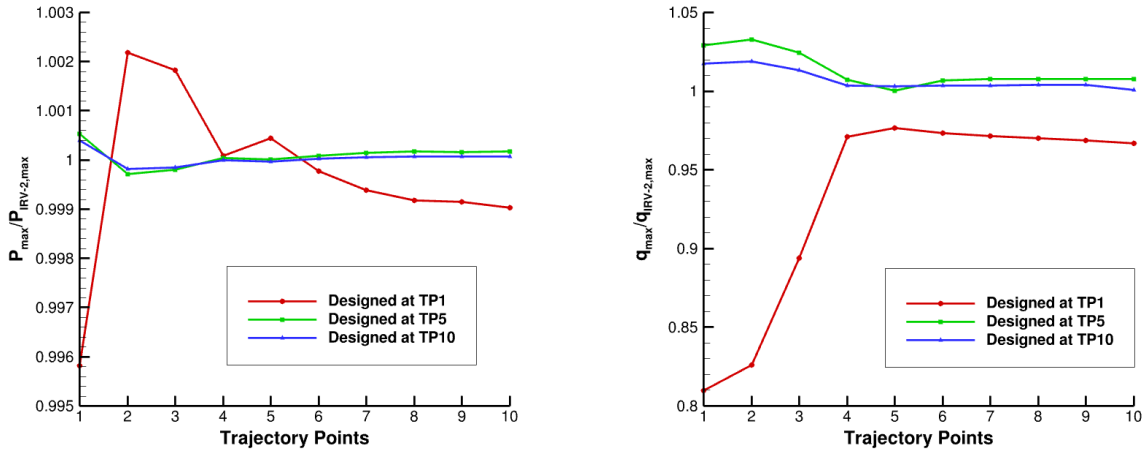

(a) Pressure.

(b) Heat flux.

Figure 5.20: Maximum wall value for each design at different trajectory points.

The maximum pressure and heat flux, which are being used for the constraint, are found around the stagnation point for all of these designs. Therefore, the stagnation line provides insight on the behavior of these constraints. Figure 5.21 juxtaposes the stagnation line for each design at three different trajectory points. For all the trajectory points, the design from TP1 shows a larger shock-standoff distance than the original geometry. By increasing the shock-standoff distance, the heating on the nose of the vehicle is reduced, which causes the decrease to heat flux observed in Figure 5.20b. This can be observed in Figure 5.21a, where the slope of the temperature at the wall is reduced compared to the original geometry, which reflects a reduction in heat flux. Meanwhile, the designs for TP5 and TP10 produce stagnation lines that are very similar to those of the original geometry, which is reflected in the heat flux constraint hovering around unity.

All of this work can be found in Refs. [55] and [54]. The three final optimizations at the different trajectory points result in very different shapes and behaviors of the vehicle. If the

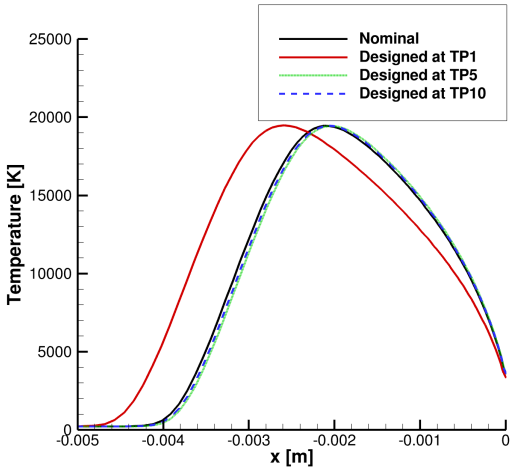

(a) TP1.

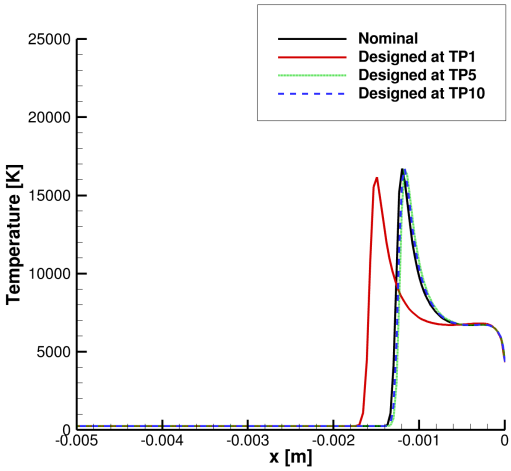

(b) TP5.

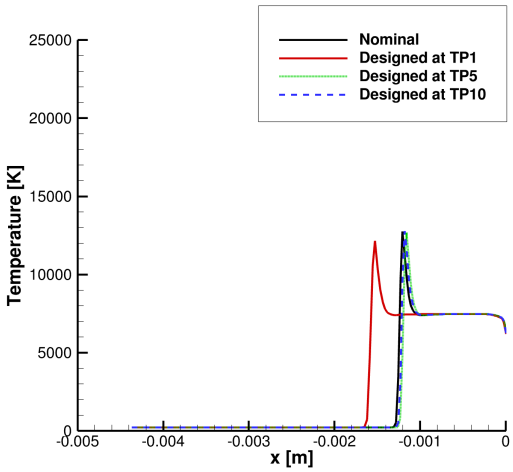

(c) TP10.

Figure 5.21: Stagnation line profile for each design at three trajectory points.

intention of optimization is to reduce the drag overall, the optimized configuration at TP5 provides the most consistent reduction in drag. However, if the heat flux has a hard constraint due to the material properties, the optimization performed at TP1 may be the safer option. Overall, because this is a re-entry vehicle experiencing many conditions over the flight, an additional analysis of interest could be including multiple trajectory points into one optimization. This could very easily be performed with the sensitivity analysis and optimization developed in this work.

5.3 Scramjet Inlet Optimization

As shown in Chapter 3, aero-thermo-elastic interactions can have a large influence on the flowfield of a scramjet inlet. These deformations can impact the temperatures and pressures experienced by the inlet, combustion chamber and nozzle structures and are therefore critical to design choices. The deformation also impacts the engine and aerodynamic performance of the vehicle and are therefore critical to mission requirements. Evaluating the aero-thermo-elastic effects of the inlet after design may cause unwanted redesigns for some conditions. Because of this, including FTSI earlier in the design process, although requiring more computation expense, may decrease the chances of catastrophic failures or the need for redesigns later on. This section focuses on including FTSI in the design process through shape optimization. The chapter juxtaposes designing with and without aero-thermo-elastic effects to present the benefits of including deformation earlier on.

5.3.1 Test Setup

The optimization in this section builds on the scramjet sensitivity analyses begun in Chapter 4. The UNSW experimental configuration is used as the baseline for the OML of the inlet, with a slight adjustment to reflect a more realistic vehicle: the outer surface of the forebody is redefined as a flat surface after the leading edge, rather than keeping the geometry as a cantilevered plate from the experiment, as shown in Figure 5.22. The entire internal structure is defined as a solid block of TZM-C03, which properties are shown in Figure 3.49. The flight conditions applied for the optimization are the same as Section 3.6 and described in Table 3.10. Because this geometry is a solid block, rather than a cantilevered plate studied in Section 4.2 or a hollow structure studied in Section 3.3, less deformation occurs, as shown in Figure 4.15. The reason a solid structure, which is less realistic due to weight and packaging concerns, is chosen for this analysis is to eliminate issues of surface contacts. Including considerations for contact in the optimization process is instead a recommendation for future work, described in Section 6.3. Due to the relatively low speed of Mach 5.85, the aerothermodynamic model does not include thermal or chemical nonequilibrium.

However, turbulence is included using the Spalart-Allmaras RANS model.

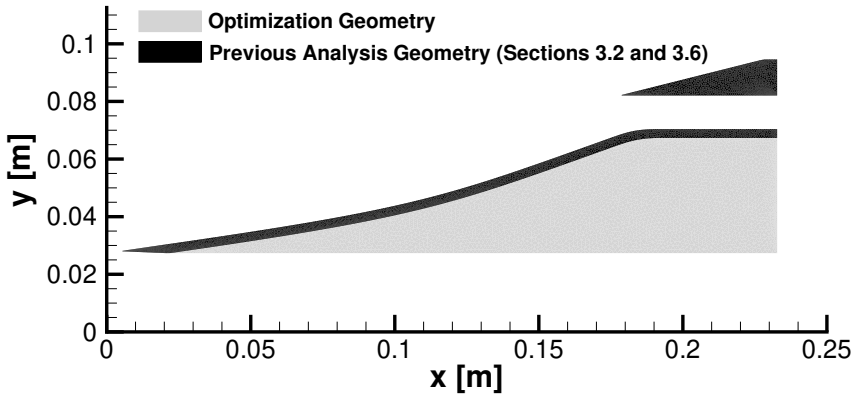

Figure 5.22: New geometry for the optimization compared with the original UNSW design.

One aim of this study is to compare ignoring or including FTSI in the optimization process. Therefore, cases ignoring FTSI begin with aero-only optimizations and are followed by an aero-thermo-elastic analysis on the optimized geometry. These aero-thermo-elastic results can then be compared to the optimized geometry from the aero-thermo-elastic optimization. As shown in Figure 4.16, there is a large difference between the aero-thermo-elastic forward analysis and the aero-only case for a cold-wall condition. Part of this difference is due to the thermal boundary condition assumption, which can influence the shock train even when deformation is not included (as shown in Section 3.6). To reduce this influence and ideally produce an aero-only analysis that is more similar to the expected performance of an FTSI analysis, determining the thermal boundary condition is essential. Therefore, a secondary goal for this study is to determine if certain wall thermal boundary conditions can assist in producing a more accurate optimization for a single aero-only discipline.

To begin, an aero-thermo-elastic forward analysis on the original geometry is performed, followed by three aero-only cases. One of these aero-only cases is a cold wall condition (300 K), one is a hot wall condition (800 K), and one is a variable temperature condition. The variable temperature

condition (labeled *Aero-only: FTSI Temp.*) uses the resultant boundary condition (BC) from the FTSI case and applies the same wall temperature profile to the undeformed configuration using RBF interpolation. The comparison of wall pressures for the four different cases is shown in Figure 5.23a, which indicate a large difference between the wall pressure for the cold wall case and the three others. Figure 5.23b depicts the mass flux weighted pressure at the outlet, since the low-fidelity engine models uses that quantity to determine the 1D pressure. Again, the largest difference is produced between the cold wall case and the three others, further highlighted in the performance parameters in Table 5.3.

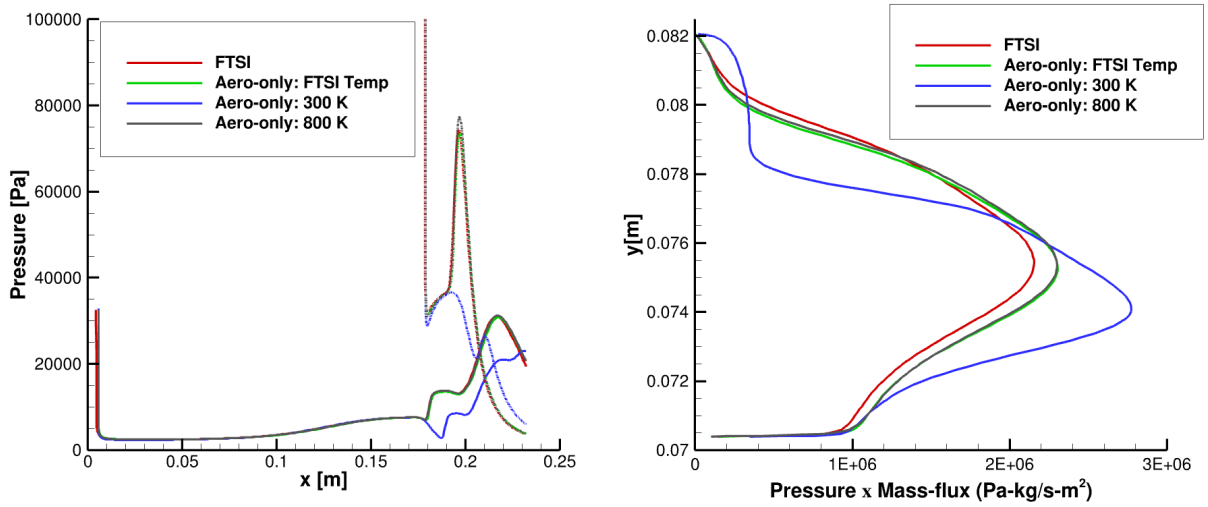

(a) Wall pressure.

(b) Outlet mass-flux weighted pressure.

Figure 5.23: Pressure distributions for the FTSI case compared with three aero-only cases for the original UNSW geometry.

Table 5.3: Initial geometry quantities of interest.

	FTSI	**Aero-only: FTSI Temp.**	**Aero-only: 300 K**	**Aero-only: 800 K**
PR	20.37	20.97	21.36	21.10
π_C	0.3266	0.3314	0.3537	0.3281
L/D	1.072	1.015	0.5360	1.013

In order to better understand this flowfield, a contour plot of the Mach number for the aero-only 800 K is shown in Figure 5.24. In addition, the separation region is outlined for the four different cases. The region for the FTSI case, the aero-only variable temperature case and the 800 K isothermal case all show nearly identical regions of separation, whereas the cold wall, 300 K case shows a smaller region in terms of both the length and height. These comparisons are reflected in the wall pressure profiles in Figure 5.23a. All the hot wall cases represent more realistic temperatures a scramjet structure might experience. The similarity between these three cases indicates that this is a relatively stable flowfield condition for the given configuration (including the material and structural choices). Therefore, even the small deformation does not show a drastic change in the resulting flowfield. Because of this more drastic difference with the cold wall case, only the hot wall condition and FTSI temperature condition are used for the aero-only optimization cases.

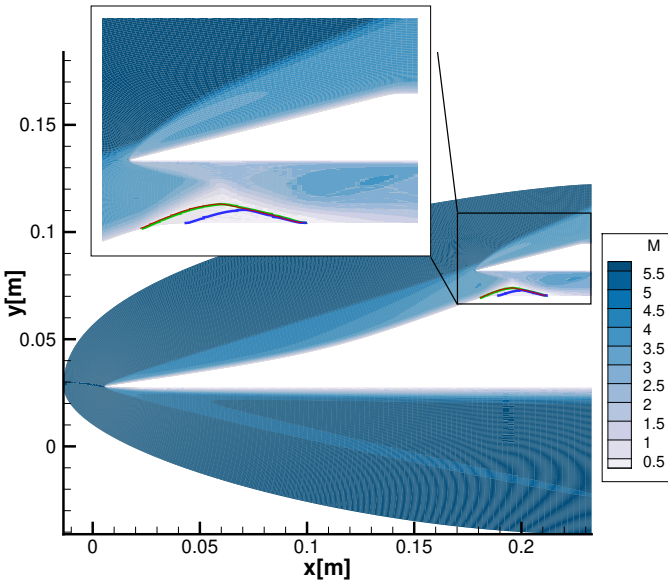

Figure 5.24: Mach number contours of the original geometry. The solid green, dotted black, solid blue and thin red lines represent the separation region for the aero-only FTSI Temp. case, aero-only 800 K isothermal case, aero-only 800 K isothermal case and the FTSI case, respectively.

For all three optimization cases, the setup is as follows: the reciprocal of the normalized compression ratio is chosen for the objective ($20/PR$); the lift-to-drag ratio is chosen as a constraint and must be greater than or equal to 1; the total pressure recovery is chosen as a second constraint

and must be greater than or equal to 30%. Additionally, box constraints on each of the design variables are all set at ([-0.05,0.05]) and are only implemented to restrict the optimization from beginning the line search with design variables that are too large; these constraints are not active for the optimized solutions. Four design variables are chosen to control the compression surface, as shown in Section 4.4. Python's SLSQP algorithm is used to drive the optimization, similar to the IRV-2 optimization case explained previously. Because the coupled optimization is much more computationally expensive, previous results from the Coupling Approach Section in Chapter 3 are integrated into the process. After the gradient is computed, the optimization algorithm searches along the descent direction until a reduced objective is found. This stage requires several coupled forward analyses; however, by incorporating the result that a sandwich coupling is within 20% of the fully two-way coupled FTSI solution, if the second iteration of the CFD produces an objective greater than 20% of the current best, the algorithm continues on to the next search evaluation rather than finishing the fully coupled analysis. Once the optimizations are completed, the final flowfield is coupled to the engine model described in Section 2.1.3 using a kerosene fuel.

5.3.2 Optimization Results

Although there is the capability to include aero-thermo-elastic results in the optimization process, it's important to understand if and when this analysis is needed in the case of a scramjet inlet since FTSI computations are costly. Hence, the aero-only cases are used as a comparison as a less costly solution. The results of the optimized geometries are shown in Table 5.4, including the objective percent-difference, the compression ratio, the total pressure recovery and the lift-to-drag ratio. For the aero-only cases, these are evaluated with and without aero-thermo-elastic considerations for the optimized geometries. For the FTSI optimization case, the objective is only able to be decreased by about 1% while including the deformations. Meanwhile, the aero-only case that uses the FTSI temperature as a boundary condition reduces the objective by 4.5% during the aero-only optimization. However, when deformation is included, the objective is reduced further by almost 23%; this reduction does come at the cost of violating both the π_c and L/D constraints.

Similarly, the isothermal aero-only case reduces the objective by 3.7% when ignoring FTSI, but increases the objective by 1.4% when deformation is included compared to the original isothermal aero-only case. Despite that, the case still provides more compression than the optimized FTSI case. Although the FTSI case produces the smallest reduction in the objective, by including the aero-thermo-elastic effects in the process, no unexpected violations occurred.

Table 5.4: Optimized geometry quantities of interest.

		FTSI	**Aero-only: FTSI Temp**	**Aero-only: 800 K**
Obj. % Difference	$(aero)$	-	-4.504%	-3.743%
	$(FTSI)$	-1.047%	-22.95%	+1.355%
PR	$(aero)$	-	21.96	21.92
	$(FTSI)$	20.59	27.22	20.82
$\boldsymbol{\pi_C}$	$(aero)$	-	0.3019	0.3019
	$(FTSI)$	0.3023	0.2880	0.2902
L/D	$(aero)$	-	1.001	1.000
	$(FTSI)$	1.067	0.7486	1.074

Narrowing in on the optimization process, Figure 5.25 shows the history of the convergence for each of the analyses. As shown in Table 5.3, the FTSI case begins with the smallest compression ratio (i.e. the largest objective), followed by the aero-only case using the FTSI temperature BC. The trajectory of the FTSI optimization shows the shallowest slope, taking much smaller steps for each iteration than the aero-only cases. The aero-only case using an isothermal BC begins with the steepest descent, but transitions to a more gradual descent for many iterations. Finally, the aero-only case using the FTSI temperature BC shows a consistent descent until reaching convergence. The trajectory of the three all look very distinct, despite a somewhat similar beginning, as shown in Figure 5.23. The filled in circles in Figure 5.25 indicate when FTSI is included; for the aero-only cases, deformation is only considered as a post-processing step for the optimized configuration. For each aero-only case, only the final point shows the FTSI result, both of which show a drastic divergence in the aero-only objective. Interestingly, the two have different reactions when FTSI is included, which will be explored further.

The optimized geometries are shown in Figures 5.26a and 5.26b, ignoring and including

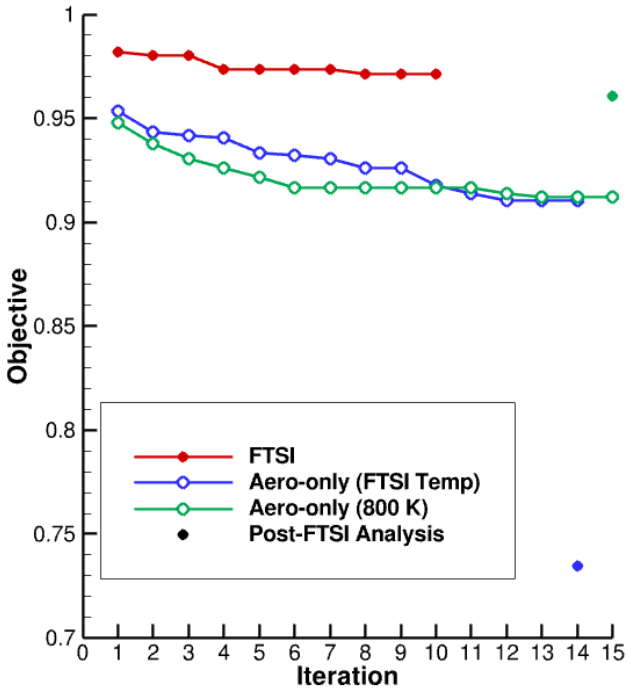

Figure 5.25: Objective function convergence rates for each optimization case.

deformation, respectively. For the undeformed configurations in Figure 5.26a, both aero-only cases show a shallower slope directly after the leading edge compared to the original geometry. The isothermal case begins more shallow, but quickly increases to a steeper slope approximately 2 cm from the leading edge. The two converge on the same tangent approximately 15 cm from the leading edge, with a steeper slope than the original geometry. Once aero-thermo-elastic effects are accounted for, the general trends of the aero-only optimization case remain. However, the two show slightly more deformation than the original geometry in both the x- and y-directions. The FTSI optimization case only shows minor differences with the original geometry, which is why only a minor reduction in the objective is observed.

When analyzing the original geometry, the effects of FTSI are not extremely drastic, as observed in Table 5.3. However, for the optimized geometry with the FTSI temperature BC, the difference between the compression ratio is very significant with and without deformation. This indicates a more sensitive geometry, which is likely undesirable for a realistic configuration.

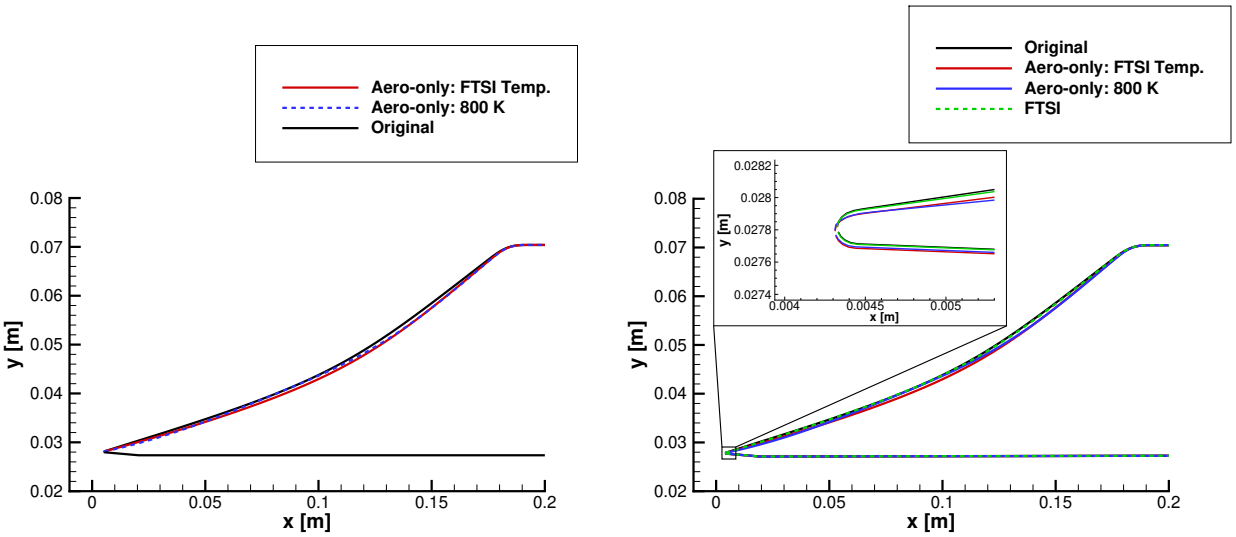

(a) Undeformed.

(b) Deformed.

Figure 5.26: Optimized geometries for the forebody region.

Studying Figure 5.26b, the overall deformation at the leading edge of the forebody is very similar to the original deformation, indicating that this is not an issue of more deformation. Instead, it shows that the impacts are very case specific based on the flowfield. To understand what in particular is causing this sensitivity, the wall pressures and Mach number contours are shown in Figures 5.27b and 5.28. Additionally, the geometry of the aero-only case with the FTSI temperature BC is shown in Figure 5.27a. Beginning with the aero-only cases, the original geometry has a steeper gradient directly behind the leading edge and, therefore, begins with a higher pressure along the forebody surface. Further along the forebody, the optimized geometry has a steeper gradient, which then causes the pressure to increase even higher than the original geometry before the internal closure of the inlet. Continuing along the wall pressure for the forebody, the sudden increase in pressure at approximately x = 0.18 m indicates where the separation region begins. The rise in pressure for the optimized geometry begins ahead of the original geometry, indicating the separation region has moved forward. Similarly, both the final rise in pressure along the forebody and the rise along the cowl after the leading edge both show a shift forward for the optimized case. These rises in pressure are indications of oblique shock interactions with the wall of the shock train. Finally,

when observing the pressure at the end of the inlet along the cowl, the optimized pressure deviates from the original pressure, showing a stark increase earlier on. Once again, this increase in pressure is a shock-boundary layer interaction, indicating that the shock train has shifted forward enough to have an additional interaction. This final increase in pressure contributes to the increase in compression ratio observed in the objective. When considering FTSI, the inclusion of deformation exacerbates all of these differences. For the final rise in pressure along the cowl, the S-shape and shift forward of the final rise indicate a separation has occurred along the forebody wall, which will be discussed more below.

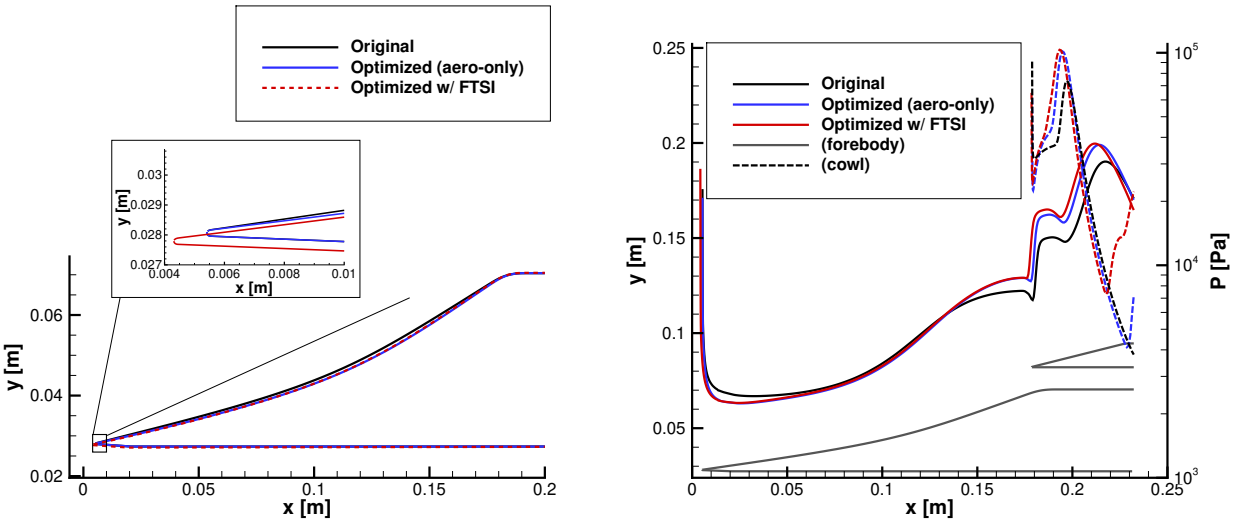

(a) Forebody geometry.

(b) Wall pressure.

Figure 5.27: Comparisons with and without deformation for the optimized aero-only case using the FTSI temperature BC.

Many of the features described above can be further examined with the Mach number contour plot of the optimized geometry, along with the designated separation regions, in Figure 5.28. Important distinctions between the aero-only and aero-thermo-elastic results are observed. The aero-only case not only has the separation region appearing along the shoulder of the forebody, but there is also slight separation seen on the cowl. In the original geometry, no separation is observed on the cowl, as seen in Figure 5.24. Once the coupled forward analysis is applied to this optimized

geometry, the separation region on the forebody shifts forward, indicating a more favorable pressure gradient around the shoulder. Additionally, the separation directly across on the cowl grows and a new separation region at the end of the inlet along the cowl is created. The new creation of separation regions indicates that the flowfield is less stable/robust to deformation than the original geometry. Because of this instability, large changes in the inlet performance are observed when FTSI is considered despite the deformation being fairly similar to the original deformation, as shown in Figure 5.27a.

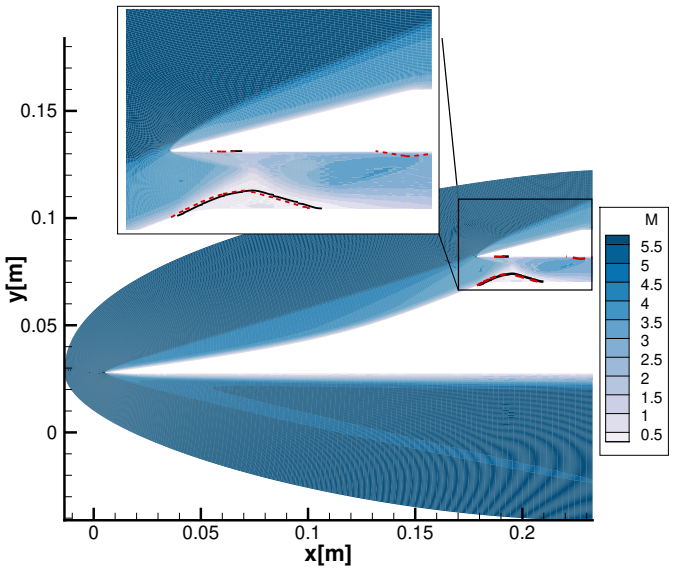

Figure 5.28: Mach number contours of the optimized geometry for the aero-only case using the FTSI temperature BC. The solid black and dashed red lines indicate the separation regions for the aero-only solution and the FTSI solution, respectively.

The post-optimality sensitivities can be used to better understand this instability compared with the original geometry. Figure 5.29 shows the objective's sensitivity to each design variable. The original geometry sensitivity is shown in red, with the aero-only case shown with a solid line and open circle and the FTSI case shown with a dashed line and closed circle. The two show very similar sensitivities, indicating that the deformation does not cause a large change to the sensitivity of the configuration to shape change. Conversely, the optimized geometry, shown in blue, shows a drastic change between the aero-only case and the FTSI case, despite the fact that the total deformation

amount is comparable to the original geometry. This indicates a more sensitive flowfield that has large reactions to a change in shape. Also note that the sensitivities of the optimized geometry are close to zero, indicating a local minimum, although the active constraints bar the sensitivity from reaching zero exactly.

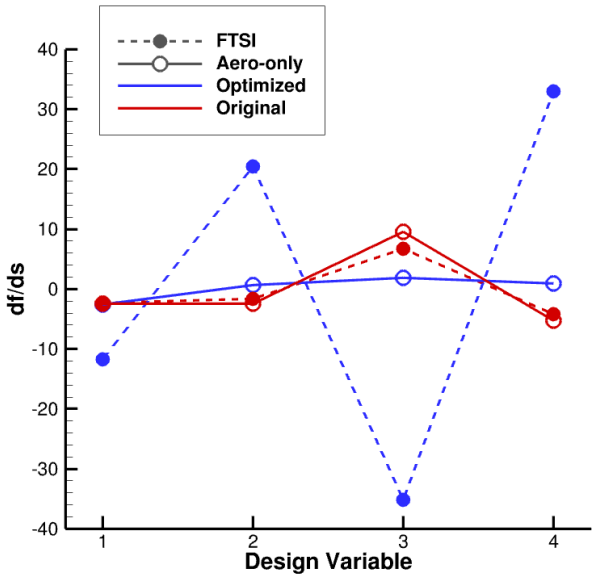

Figure 5.29: Sensitivity of the objective to each design variable for the original and optimized configurations with and without FTSI for the aero-only FTSI temperature BC case.

Transitioning now to the aero-only case with an isothermal wall BC, the flowfield reacts oppositely when the shape is perturbed by deformations. Instead of further increasing the compression ratio, there is a loss of pressure. Similar to the previous case, the separation bubble along the shoulder is shifted forward for the optimized case, as seen in Figure 5.30. However, when focusing on the wall pressure at the end of the cowl, the inclusion of FTSI does not cause such a drastic shift in the final increase in pressure, nor does it produce an S-shaped curve. This indicates that a separation region has not occurred in that area, which is further captured in Figure 5.31. A small separation region is shown along the cowl just aft of the leading edge, but one does not appear along the exit as is shown in Figure 5.28. The inclusion of FTSI does show a larger increase in pressure at the exit along the cowl, but also a reduction in pressure along the forebody, which

overall produces a smaller compression ratio.

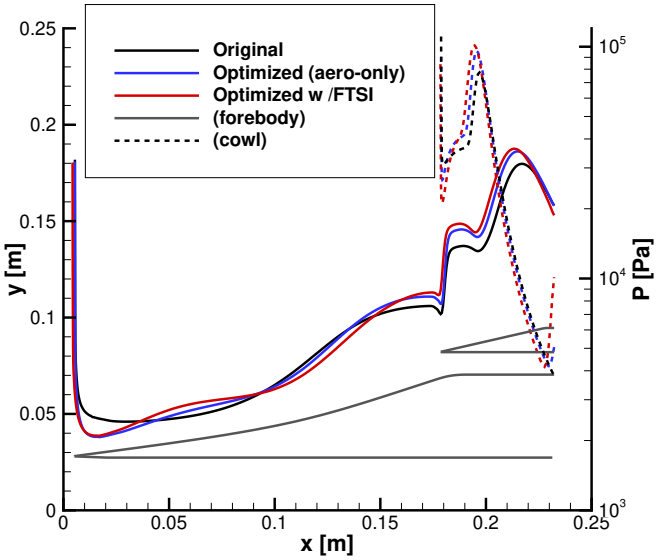

Figure 5.30: Comparisons of the wall pressure with and without deformation for the optimized aero-only case using the isothermal wall BC.

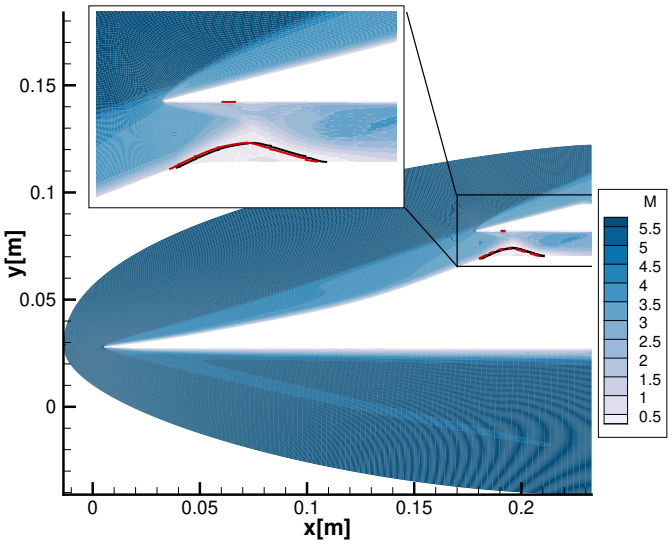

Figure 5.31: Mach number contours of the optimized geometry for the aero-only case using the isothermal BC of 800 K. The solid black and dashed red lines indicate the separation regions for the aero-only solution and the FTSI solution, respectively.

The difference between the two aero-only cases can be further observed through the differences in their post-optimality sensitivities. Figure 5.32 again shows the sensitivity of the objective for each design variable, this time for the aero-only isothermal case. The original sensitivities remain the same in red, but optimized sensitivities in blue show a very different story than Figure 5.29. All but one sensitivity for the FTSI inclusion case remains relatively similar to the aero-only optimized case. However, there is a large sensitivity to the third design variable. Interestingly, the sensitivity is the opposite sign to the aero-only FTSI temperature BC case. Overall, the sensitivities increased with FTSI more than the original case indicating that the optimized configuration is more sensitive to shape change. Once again, the optimized aero-only case shows sensitivities close to zero, indicating a local minimum has been reached.

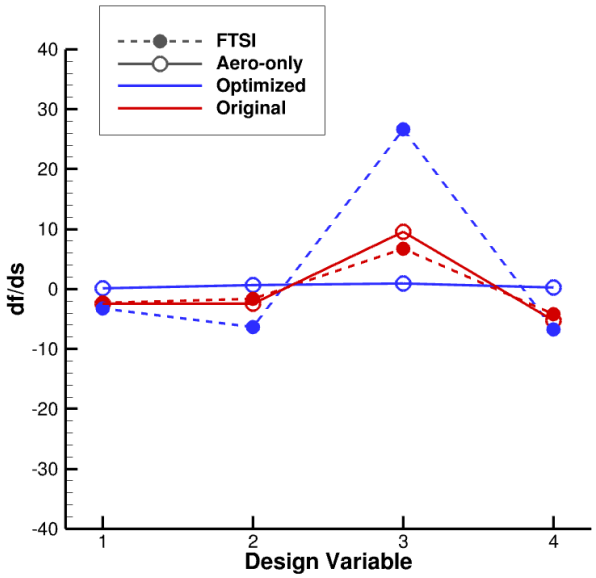

Figure 5.32: Sensitivity of the objective to each design variable for the original and optimized configurations with and without FTSI for the aero-only isothermal case.

All of these geometry changes, whether through shape optimization or the aero-thermo-elastic deformation, have implications on the subsequent engine processes. Therefore, it's important to understand the flow conditions that are provided to the combustion chamber. The pressure and temperature outflow conditions are presented in Figure 5.33 for each of the optimized cases.

Beginning with the pressure in Figure 5.33a, all the cases except one show a similar shape along the outflow with varying magnitudes. The exception is when aero-thermo-elastic deformations are applied to the optimized aero-only configuration with the FTSI temperature BC. A much higher pressure appears on the cowl due to the separation occurring there, causing the major increase to the exit pressure and compression ratio when compared to its optimized configuration neglecting deformation. Meanwhile, the aero-only optimized configuration using the isothermal BC shows an increase in pressure along the cowl and a decrease in pressure along the forebody when deformation is applied, further shown in Figure 5.30. Overall, when a mass-flux weighted average of the pressure is considered, deformation causes a lower pressure at the exit. Finally, the FTSI optimized case still shows the lowest mass-flux weighted average of pressure.

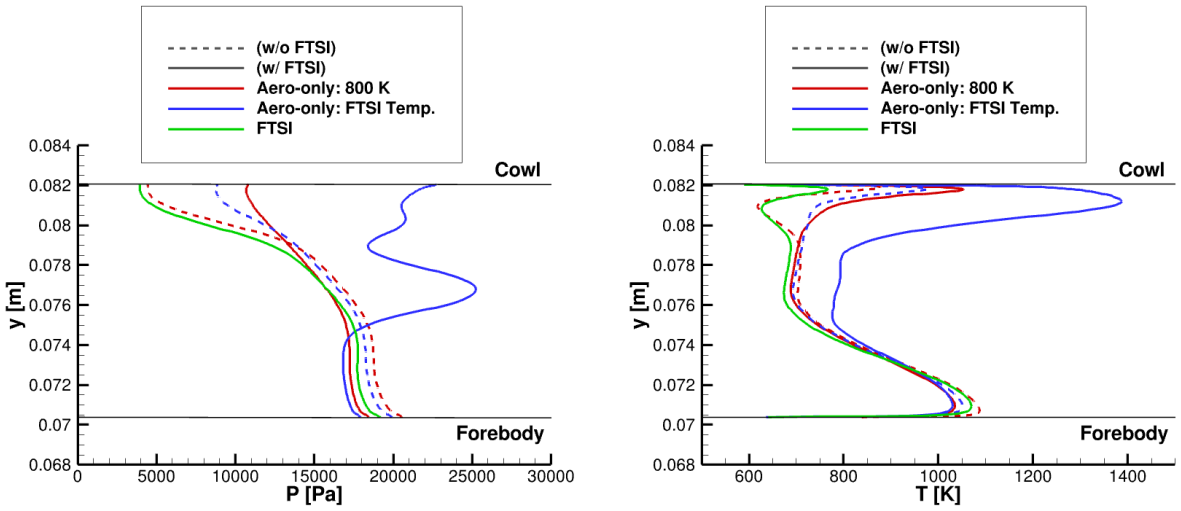

(a) Pressure.

(b) Temperature.

Figure 5.33: Comparisons of the outflow properties for each optimized case including and excluding aero-thermo-elastic effects in dashed and solids lines, respectively.

Focusing on the temperature profiles in Figure 5.33b, once again, the FTSI optimized case shows the lowest overall temperature for the outflow. Additionally, a large increase in temperature occurs near the cowl when FTSI is applied to the aero-only optimized cases. For the aero-only case using an FTSI temperature BC, the increase in maximum temperature rises almost 400 K with

deformation. The high temperatures and pressures fed into the combustion chamber can cause material degradation issues if not properly accounted for. Therefore, even though this configuration produces the largest compression ratio, this comes at a cost of higher temperatures and pressures experienced by the engine.

The repercussions to the engine can be further studied using Table 5.5, which includes the engine and combustion chamber properties when the optimized results are coupled with the low-fidelity engine model. Each of the three optimized configurations are presented, including both the aero-only and aero-thermo-elastic results when applicable. As a baseline comparison, the original geometry is also included. Typically, the compression ratio is used as an indicator for the performance of the inlet, but the inclusion of the engine model allows the specific impulse to be compared as a performance metric for the entire engine. Despite the increases to the compression ratio, the specific impulses for the aero-only results show that the optimized configurations are less effective than the original geometry. This result indicates that the compression ratio on its own may not be enough of an indicator for the performance of the engine and highlights the need to consider coupling the engine model into the optimization process, which will be discussed further in the Recommendations for Future Work Section. Once aero-thermo-elastic effects are considered, the optimized FTSI case shows only a marginal increase in the specific impulse. However, the aero-only optimized case involving the isothermal boundary condition shows a lower specific impulse than the original when FTSI is considered. The aero-only optimized case using the FTSI temperature BC shows the largest increase to the specific impulse, corresponding to the largest increase in compression ratio. The same trends are seen with the total engine efficiency, and interestingly the efficiency of the aero-only optimized case with the FTSI temperature BC increased.

The second set of quantities in Table 5.5 includes the maximum temperatures and pressures of the gas experienced within the combustion chamber, which builds upon the conditions shown in Figure 5.33. Overall, the temperatures all remain slightly below 2000 K, with only about a 20 K variation. The one exception is again when an aero-thermo-elastic analysis is applied to the optimized configuration of the aero-only case using the FTSI temperature BC. Here, an increase

Table 5.5: Engine properties for the original and optimized geometries.

		Original	**FTSI** (optimized)	**Aero-only: FTSI Temp** (optimized)	**Aero-only: 800 K** (optimized)
$\mathbf{I}_{SP}$ (s)	(*aero*)	1695	-	1689	1694
	(*FTSI*)	1684	1685	1793	1668
$\boldsymbol{\eta}_o$	(*aero*)	0.5737	-	0.5712	0.5723
	(*FTSI*)	0.5729	0.5729	0.5903	0.5700
$\mathbf{T}_{max,combust}$ (K)	(*aero*)	1957	-	1977	1978
	(*FTSI*)	1953	1958	2064	1974
$\mathbf{P}_{max,combust}$ (Pa)	(*aero*)	37610	-	39100	39430
	(*FTSI*)	36550	36920	50630	37360

of 100 K is observed for the quasi-1D analysis. For the maximum pressures experienced within the combustion chamber, the optimized configurations only considering aero-only show an increase in pressure around 4-5%. When FTSI is considered, the FTSI optimized and the isothermal aero-only optimized configurations show a slight increase of pressure around 1-2%. However, when deformation is applied for the aero-only FTSI temperature BC case, the pressure increases by 39%, which could surpass the structural limits of the combustion chamber.

The original reason behind including the two aero-only cases was to discover if a more accurate thermal boundary condition could produce comparable results to an FTSI optimization, while being less computationally limiting. However, including the FTSI temperature boundary condition case for the aero-only optimization did not provide an optimized geometry more similar to the FTSI optimization case. Instead, it resulted in the most unstable of the three cases, indicating that the deformation and two-way coupled aero-thermal analysis are critical to informing the design decisions, especially for a stable configuration. Although the aero-only cases provided a large increase to the compression ratio, both also violated the constraints which could cause critical engine failures. Only incorporating FTSI in the optimization led to a well-constrained solution that maximized the compression ratio.

However, only a roughly 1% reduction in objective is observed for the FTSI optimization; even the aero-only cases only saw a few percentage points of reduction. This study looked at

minor changes to the forebody of a vehicle to optimize the compression ratio. Much more drastic shape changes could instead be included, such as allowing the length of the forebody to change or deforming the cowl. Additionally, since the constraints/requirements drove the final optimized configurations, more information on those requirements could produce more optimal configurations. The current work is mostly a proof-of-concept for the FTSI and aerothermodynamic optimization processes for the scramjet inlet, but many more studies can be used to further define the optimal configuration.

5.4 Chapter Summary

Two optimization examples were performed in this chapter on a re-entry vehicle and a scramjet inlet. The optimization approach, including the semi-analytical sensitivity analyses, was built for shape optimization of hypersonic structures. Both a single discipline approach, only using the aerothermodynamic forward and sensitivity analyses, or a multi-disciplinary approach, coupling the aero-thermo-elastic forward and sensitivity analyses, can be included in the optimization algorithm. This chapter began with an aerothermodynamic-only optimization of the reentry vehicle, IRV-2, which had been evaluated and optimized in previous literature.

For this work, the methodology was applied to the IRV-2 geometry to reduce drag. The IRV-2 geometry has been optimized in the past using alternative methods, focusing on the nose region. Here, the current analysis was extended to the full vehicle to understand the impacts of the shape on the overall drag coefficient. This research has indicated that the reduction in drag at the nose was insufficient to reduce the drag of the entire vehicle and instead increased the drag due to increased pressure along the wall. Including the entire vehicle in the optimization process, while still only deforming the nose, allowed for a decrease in drag by 1.1% (accounting for the various constraints, such as the pressure, temperature, total heat rate and heat flux). Additionally, by defining the design variables to control the shape along the entire length of the vehicle, the drag reduction improved by 5.7% while including an additional volume constraint.

Two more optimal geometries were determined for trajectory points TP5 and TP10, which

were able to reduce the drag by 1.5% and 2.6% at their flight conditions, respectively. The optimal shape varied for the three different trajectory points, indicating that the altitude and density influence the minimum drag shape. These three optimal designs were then evaluated at the first ten trajectory points of the IRV-2 flight. The optimal design for TP1 performed the best at the high altitude cases but increased the drag substantially at the lower altitudes. In contrast, the optimal design for TP5 was consistently able to reduce the drag for all the trajectory points, although the reduction was only around 1% to 2%. Specific mission requirements are necessary to help define which design would be most optimal. Additionally, the aerodynamic and aerothermodynamic constraints were defined by the IRV-2 vehicle's original values. If these constraints can be relaxed to account for the material property restrictions, additional reductions in drag may be possible.

The second optimization focused on a scramjet inlet and included both aero-only optimizations and an FTSI optimization in order to compare the influence of deformation on the vehicle and on the optimization process. The objective was to minimize the normalized reciprocal of the compression ratio (i.e. maximize the compression ratio) while constraining the total pressure recovery and the lift-to-drag ratio. The FTSI optimized configuration increased the compression ratio the least, but when deformations were applied to the aero-only cases, the results violated the constraints. For the aero-only cases, one applied an isothermal boundary condition of 800 K and the other applied the resultant temperature profile from the FTSI analysis on the original geometry. Although the two showed promising results when deformation was neglected, the quantities of interest changed considerably when deformation was applied. This indicated that the optimized geometries were more sensitive to shape changes. This result was further confirmed when the post-optimal sensitivities were compared for the objective against the original sensitivities. The aero-only optimized sensitivities were close to zero, indicating a local minimum, but many of the sensitivities were an order of magnitude greater than that of the original geometry when including deformation. The stark increase highlights an unstable solution prone to large changes and showcases that these sensitivities can help decide the robustness of the designs to shape change.

Chapter 6

Conclusion

This concluding chapter summarizes the research presented in this dissertation. Highlights of both the tool development and applied results are discussed, followed by a record of the major contributions of this work. The next section details the recommendations for the potential future directions of this work, specific to aero-thermo-elastic analysis and optimization. Finally, a summary of the publications involving this work are presented for further reference.

6.1 Dissertation Summary

A brief summary of each of chapter is presented in this section. Beginning with Chapter 2, two main tools were developed: an aero-thermo-elastic forward analysis and an aero-thermo-elastic optimization procedure. The forward analysis tightly coupled a hypersonic CFD solver with a thermo-structural FEM solver and then subsequently integrated with a low-fidelity engine model. Together, this forward analysis assesses the deformation and temperature of the structure, the impact on the inlet flowfield, and the repercussions to the engine. This FTSI analysis can be used on its own, or incorporated into the optimization toolkit. Within the multidisciplinary optimization process, several sub-development tasks were presented, including an aerothermodynamic sensitivity analysis and a coupled FTSI sensitivity analysis, which are both integral for gradient-based optimizations. The aerothermodynamic sensitivity analysis computes the gradient through the direct method, accounting for thermal nonequilibrium, chemistry, and turbulence. This analysis was then tightly coupled to a thermo-structural sensitivity analysis computed through a forward

finite-difference. These tools were then applied to hypersonic applications such as a scramjet and a reentry vehicle.

Chapter 3 employed the aero-thermo-elastic analysis for many studies involving scramjet inlets. One of the studies assessed the influence of heating/thermal expansion on the deformation process, which led to using thermal management as an effective tool to reduce unwanted repercussions from the deformation. A second study narrowed in on how the level of coupling between the disciplines affected the accuracy and efficiency of the solution; various situations were assessed such as material, thermal management and flight conditions. A third study determined the influence deformation had on the sensitivity of different performance metrics – lift, drag, specific impulse, etc. – to varying flight conditions often expected across a mission. When compared to an isothermal, rigid structure, the behavior of the lift coefficient was more strongly influenced by the dynamic pressure for the deforming case. Finally, a fourth study altered the aero-thermo-elastic analysis to include thermal transients to assess the deformation and its impact on the engine as the vehicle reaches thermal equilibrium; the temperature variations across the transient affected the size of the separation region and the specific impulse of the engine. In addition to these studies, a validation effort was performed against an experiment involving an inlet surface that could dynamically deform due to structural loading (no thermal strains were captured in the experiment).

Chapter 4 transitioned to applications of the aero-only and aero-thermo-elastic sensitivity analyses. The overarching goal of this chapter was to benchmark the newly developed model before incorporating it into an optimization. The chapter began by studying a canonical hypersonic geometry: a nose-cone. Good agreement was found between the semi-analytical sensitivity analysis and a forward finite-difference, which was used as a benchmark for the flowfield, wall, and objective variable sensitivities. With confidence in the aero-only sensitivity analysis, the FTSI sensitivity analysis was assessed against another forward finite-difference, this time using coupled forward analyses. Again, great agreement was seen for the nose-cone case. Next, the same effort was applied to a scramjet inlet geometry. The nose-cone case included thermal nonequilibrium, chemistry and laminar flow for its flight conditions; the scramjet inlet ignored thermal nonequilibrium and

chemistry, but included turbulence modeling using RANS. Because of both the RANS capabilities and the complicated flow structures seen in inlets, the sensitivity analysis again needed to be benchmarked before use. Both the aero-only and aero-thermo-elastic sensitivity analyses showed good agreement with the finite-difference solutions for the wall quantities, inlet outflow quantities and an objective function. Therefore, the sensitivity analysis could be incorporated into the optimization procedure with confidence for the canonical nose-cone and scramjet inlet cases.

Chapter 5 discussed optimizations of two different hypersonic systems: a reentry vehicle and a scramjet inlet. Beginning with the reentry vehicle, the optimization methodology was applied to the IRV-2 geometry to reduce drag while considering constraints such as temperature, pressure, heat flux and volume. Alternate studies had been performed previously on the nose region of the geometry, so the main contribution of this study was extending the new sensitivity analysis to the full, 1.4 meter long vehicle. The initial study concluded that performing the optimization only on the nose did not result in a reduction in drag when the entire vehicle was considered; instead a slight increase in drag was predicted. Subsequently, three optimizations were performed on the entire vehicle with certain regions controlled by the design variables. The first optimization controlled the nose-region, as defined by the previous study by Eyi et al. [35], and produced a geometry that reduced the drag by 1.1%; this indicates that including the entire geometry is necessary even when perturbing only the nose region. The largest reduction in drag from these studies was when the entire vehicle could deform. The next study optimized the entire geometry at three separate trajectory points and then evaluated the optimized solutions along the first ten trajectory points. The optimized geometry at the first trajectory point, corresponding to the highest altitude, had the most drastic reduction in drag for the higher altitude trajectory cases, but then dramatically increased the drag at lower altitudes. Conversely, the optimized geometry for the mid-altitude trajectory point produced a slighter, but consistent reduction in drag along all the evaluated points.

The final half of Chapter 5 was the culmination of all the previous efforts: an aero-thermo-elastic optimization of a scramjet inlet leveraging the coupled forward and sensitivity analyses.

Along with this multiphysics optimization problem, two additional single-discipline aerothermodynamics optimizations were performed as a comparison. The first aero-only optimization used an isothermal boundary condition of 800 K. The second was an attempt to bridge the gap between an aero-thermal model and an aero-only model by applying the original wall temperature solution from the FTSI forward analysis without any deformation. A variation of the compression ratio was used as the objective and the lift-to-drag ratio and the TPR were used as the constraints. Overall, the coupled FTSI optimization was able to increase the compression ratio by over 1%, while the aero-only isothermal and FTSI temperature cases were able to increase the compression ratio by 3.9% and 4.7%, respectively. However, once an aero-thermo-elastic forward analysis was performed on the optimized geometries for the aero-only cases, both violated constraints and showed major differences in the objectives. In particular, the aero-only case using the FTSI temperature BC had a 24% change in the objective when deformations were included due to additional separation occurring. This indicated that certain configurations were more sensitive to shape changes and could be easily driven to unstable solutions. Therefore, although the FTSI optimization case only increased the compression ratio by 1%, it did provide a robust and stable configuration while including the deformation.

6.2 Research Contributions

The major contributions related to fluid-thermal-structural interactions and optimization for hypersonic applications are summarized as follows:

- **Development of a high-fidelity aero-thermo-elastic model:** A fully coupled CFD-FEM framework was developed and demonstrated to model fundamental interactions between the flowfield, material and structure of a hypersonic vehicle. Initial validation efforts showed good agreement to the limited test data available for the case of a scramjet inlet.
- **Use of thermal management as a mitigation strategy:** The aero-thermo-elastic model provided insights for the root cause of the deformations for a generic scramjet inlet.

Thermal expansion contributes significantly to deformation and certain undesirable outcomes; as a consequence, thermal management is a potential option for controlling/reducing the undesired distortion of the inlet.

- **Influence of the deformation on the performance of a scramjet engine:** A scramjet is highly dependent on the geometry to perform the proper compression, so even the relatively minor deformations occurring influences the composition of compressed air at the exit of the inlet. The modified flow will then alter the isolator, combustion chamber and nozzle flow conditions, resulting in different pressures and temperatures than originally designed for and changes to the engine performance.

- **Development of a hypersonic aerothermodynamic sensitivity analysis:** A reliable and efficient hypersonic aerothermodynamic sensitivity analysis was developed, demonstrated and benchmarked for hypersonic applications (including chemistry, thermal non-equilibrium, and turbulence). There are few codes currently capable of providing a semi-analytical sensitivity analysis with full hypersonic capabilities such as this.

- **Development of a hypersonic aero-thermo-elastic sensitivity analysis:** A fully coupled hypersonic aero-thermo-elastic sensitivity analysis was developed, demonstrated and benchmarked for hypersonic applications. This tool was demonstrated on a nose-cone geometry and a scramjet inlet, but there are several other multidisciplinary applications this can be used for.

- **Shape optimization of a re-entry vehicle:** The aerothermodynamic sensitivity analysis was integrated into an optimization procedure to reduce the drag of the IRV-2 vehicle. The area controlled by shape optimization and the effect of flight conditions on the optimized geometry were studied and compared along a trajectory.

- **Multidisciplinary optimization of a scramjet inlet:** The aerothermodynamic and aero-thermo-elastic sensitivity analyses were integrated into optimization procedures to

maximize the compression ratio of an inlet while maintaining the TPR and lift-to-drag ratio. The aerothermodynamic-only optimization cases were able to increase the compression ratio more than the aero-thermo-elastic optimization, but were shown to be much more sensitive to FTSI when included.

6.3 Recommendations for Future Work

Aero-thermo-elastic Analysis Specific

A recurring question asked about the thermal management mitigation strategy is how this can realistically be designed. One potential area of future work is fleshing out this thermal management system more by adding detail to the model, rather than using an isothermal boundary condition. For example, stacking different materials and different thicknesses could help produce a more evenly distributed temperature and expansion across the top and bottom panels. Otherwise, using an active cooling technique such as routing a small portion of fuel through channels to be used as a coolant can be analyzed to determine if the fuel consumption is possible. A more detailed model is necessary to accurately account for either thermal management system. In addition, the current scramjet inlet internal geometry is very simple, where a "bulkhead" is only considered at the end of the geometry as a fixed boundary conditions. Updating the geometry to include multiple bulkheads into the thermal and structural analyses may help to limit both the thermal expansion and structural deformation.

In addition to modifying the performance of the engine and vehicle, the intense thermal and structural loads experienced by the vehicle along the trajectory will have implications on the structural limits of the vehicle. Due to the large variations in temperatures that are experienced by the vehicle throughout each stage of the flight, the material properties will vary greatly, as shown in Figure 3.22. In order to avoid damage, or even failures, the path-dependent aero-thermo-elastic effects should be evaluated. The flight path of the vehicle and its thermal and structural loading will also affect the cycle fatigue, which will likely need to be studied further both experimentally and numerically.

Overall, simple, low-fidelity models were used to assess the aerodynamic and engine performance *trends* when deformation was included. There are several modifications that can be used to better predict the expected *quantities*, as well. For example, using a mass-flux weighted average to ascertain 1D flow properties from a 2D CFD solution has been used for this dissertation, as well as previously in Refs. [7] and [67]. However, a more accurate representation could include the Conserved Mass/Momentum/Energy Method described in Baurle and Gaffney [7], which captures continuous stream thrusts and forces. In terms of the aerodynamic coefficients, the inlet was considered as an individual structure for this analysis and the entire wetted surface was included in the computation; however, when considering the entire engine, much of the lift and drag from the inlet is instead considered within the streamtube of the engine and is not included in the aerodynamic quantities. Updating the aerodynamic model to ignore the forces within the streamtube of the engine and include the spillage would provide a more comprehensive depiction of the inlet integrated into the engine.

Another area for development is a tighter coupling between the high-fidelity CFD and the low-fidelity combustion model. Currently, only a one-way coupling is implemented, with the flow of information going from inlet/isolator to the combustor. Although the majority of the flow *should* remain supersonic, at least the boundary layer will be subsonic. For subsonic conditions, a more realistic boundary condition at the exit of the isolator is necessary, which can come from the pressure computed in the low-fidelity combustion model. Therefore, a two-way coupling between the inlet/isolator and the combustor models would allow for a more accurate model for the boundary layer. In non-ideal cases, which can arise due to the inlet deforming, a normal shock can appear in the isolator when there are high enough pressures in the combustion chamber. This can lead to an unsteady behavior, unstart, which could force the normal shock all the way through the inlet. Modeling this possible behavior is integral to the engine performance and therefore should be captured by tightly coupling the two models currently in place.

For many applications, such as assessing an entire flight envelope or mission profile, the use of high-fidelity FTSI such as that developed in this thesis is intractable. Instead, reduced order mod-

eling (ROM) or surrogate modeling can be developed from high-fidelity data points. These ROMs would allow for a faster examination of the parameter space, making it easier to integrate into the design phase of a hypersonic vehicle. Additionally, ROMs would make uncertainty quantification possible. Many of the parameters in CFD are considered uncertain for hypersonic applications. Moreover, the high-temperature materials necessary for the harsh thermal environments have many uncertainties associated with the material properties, especially at such high temperatures. Therefore, quantifying the uncertainty in the aero-thermo-elastic deformations and its impacts on the vehicle and engine performance is critical for vehicle development and testing. Due to the computational cost of these aero-thermo-elastic analyses, uncertainty quantification is not possible with the current implementation, but could be with ROMs. The local sensitivity analysis developed in this dissertation using a semi-analytical method can also be incorporated into training the ROMs [121].

Additionally, the current work focuses on a static analysis for the structural response under cruise conditions. However, dynamic interactions are likely to occur, whether that includes the entire structure or just a panel. The current model is not able to capture these dynamic behaviors, such as flutter and vibrations. Additional work assessing the timescales and coupling between a fluid model and a dynamic structural response is needed to model such phenomena. Previous research on geometries such as panels has concluded that the structural timescale is two orders of magnitude smaller than the thermal response time, allowing for a quasi-static thermal-structural assumption where thermally-induced vibrations are not expected [26]. In cases where this assumption cannot be made, the computational cost for solving unsteady and dynamic coupled problems often leads to researchers using either low-fidelity or ROMs [86]. Extensive research in this area for the aero-thermo-elastic effects have used Kriging surrogate modeling [86, 93], but the application of such on a scramjet has yet to be performed.

Optimization Specific

The low-fidelity engine model is currently not included in the scramjet optimization procedure, but would not be difficult to integrate into the framework. The engine model is dependent on

the outflow conditions of the inlet/isolator; the flow conditions and sensitivities necessary for the model are already output in the current procedure. Therefore, the only additional work is performing a directional finite-difference to compute engine specific quantities of interest once the FTSI coupled sensitivity analysis is complete. If the engine model is integrated into the iterative forward analysis, as described in the previous section, then the finite-difference engine model sensitivity analysis would need to be incorporated into the coupled process in Equations 2.32 and 2.33. Note that a finite-difference is tractable for this portion of the analysis due to the low computational cost. By incorporating the engine model into the optimization, the specific impulse or thrust can be directly optimized and additional constraints can be placed on the expected temperatures and pressures in the combustion chamber and nozzle.

Often times engine and/or vehicle requirements work in opposition, such as decreasing the drag of the vehicle while increasing capture area of the engine. A useful tool to weigh both in the design space is creating a Pareto front of optimized geometries. A Pareto front is a set of solutions for a multi-objective problem, where each solution represents the best case for a certain weight between the objectives. The majority of the dissertation included building tools and structure capable of solving complex optimization problems. With the tools in place, exploring Pareto solutions for various objectives is possible. Examples of objectives could include specific impulse, drag, engine efficiency, weight, deformation, and many more. A Pareto front can also be used to optimize for multiple points along a mission profile such as ascent, cruise and descent.

Within the Coupled Sensitivity Analysis Chapter, the sensitivity for the thermo-structural analysis is performed through a finite-difference. However, MORIS is built specifically for topology and shape optimizations, therefore several of the necessary features are already in place. There is a current effort to integrate the built-in sensitivity analysis and optimization features of MORIS. The current limitations are the shape definition of the surface mesh and a direct method integration. Previously, MORIS' optimization focused on level-set function definitions for the surface, but an externally provided surface is necessary for this work. An additional update is that MORIS use adjoint sensitivities and therefore needs to reformat the definitions of the sensitivities in order to

integrate with the coupled sensitivity procedure outlined in this work. One of the benefits of using the inherent sensitivity and optimization structure within MORIS rather than a finite-difference is that complicated features of shape and topology optimization are already implemented. The scramjet optimization in Chapter 5 was performed on a completely solid structure, unlike the panel-like geometries in Chapter 3. The reason for this decision was to avoid overlapping structures during the optimization process. Preliminary work is already implemented in MORIS to account for these contacts, and can therefore be integrated into an optimization process. In addition, a topology optimization can be performed with minimal changes to the current process. By leveraging both shape and topology optimizations, it is possible to design more realistic structural components while considering deformation.

One limiting factor for topology optimization or for 3D analysis is that the direct method expense is a function of the number of design variables. To alleviate this concern, an adjoint analysis can be implemented instead of the direct method within LeMANS so the computation expense of the gradient is dependent on the number of objective functions rather than the number of design variables. This endeavor would require restructuring the expressions from Equations 2.32 and 2.33 using an adjoint construction:

$$\begin{bmatrix} \frac{\partial \mathcal{R}_F}{\partial q_F} & \frac{\partial \mathcal{R}_F}{\partial q_T} & \frac{\partial \mathcal{R}_F}{\partial q_S} \\ \frac{\partial \mathcal{R}_T}{\partial q_F} & \frac{\partial \mathcal{R}_T}{\partial q_T} & 0 \\ \frac{\partial \mathcal{R}_S}{\partial q_F} & \frac{\partial \mathcal{R}_S}{\partial q_T} & \frac{\partial \mathcal{R}_S}{\partial q_S} \end{bmatrix}^T \begin{bmatrix} \psi_F \\ \psi_T \\ \psi_S \end{bmatrix} = \begin{bmatrix} \frac{d\boldsymbol{f}}{d\mathbf{q}_F} \\ \frac{d\boldsymbol{f}}{d\mathbf{q}_T} \\ \frac{d\boldsymbol{f}}{d\mathbf{q}_S} \end{bmatrix} \tag{6.1}$$

$$\frac{df}{d\mathbf{s}} = \frac{df}{d\mathbf{s}}^{expl} - [\psi_F^T, \psi_T^T, \psi_S^T] \begin{bmatrix} \frac{\partial \mathcal{R}_F}{\partial \mathbf{s}} \\ \frac{\partial \mathcal{R}_T}{\partial \mathbf{s}} \\ \frac{\partial \mathcal{R}_S}{\partial \mathbf{s}} \end{bmatrix} \tag{6.2}$$

where ψ is the adjoint *vector*. Therefore, solving the adjoint expression in LeMANS would require additional restructuring, as well. An adjoint analysis is already performed in MORIS, so extensive changes to the code would not be necessary except for the communication with LeMANS. Therefore, although a new implementation in LeMANS is necessary, the flexibility for the number of design

variables would allow for many more complex optimizations.

A repercussion of using a gradient-based technique is that the search looks for a *local* minimum. Therefore, if the design space is riddled with several local minima, it is unclear how close the solution is to a global minimum. To assist in design space exploration, a multi-fidelity optimization process could be leveraged. For example, a lower-fidelity aero-thermo-elastic model can be used in conjunction with a global search algorithm. Once a minimum is discovered, the high-fidelity technique described in this work can be used to further search that area. This does not necessarily ensure a global minimum, since the low-fidelity model could be missing key aspects of the physics drivers; however it does allow for a larger exploration of the design space.

Section 5.3.2 touches on the use of the post-optimality sensitivities to inform the robustness of a solution to shape change, but these sensitivities could shed light to further help with design decisions. Post-optimality derivatives represent the "change in the optimized disciplinary constraints with respect to a change in the system design variables" [84]. Certain design trends, such as how design variables or other parameters effect the objective and constraints, can be discerned from this type of study [6]. By doing so, it is possible to derive which design variables or constraints are the most influential in driving the design. The results could indicate to engineers how to best alleviate driving constraints. If there is some flexibility in those constraints, a post-optimality study could quantify the benefit of relaxing or tightening said constraint. For example, if the material temperature constraint is currently the biggest limiter for the optimization, it might be worthwhile to explore different materials in order to relax that constraint. Conversely, if the stress within the body does not have a large influence, it may be beneficial to tighten that constraint and possibly use a lighter material or structure for mass savings. This type of study can also be used to drive a robust design that is impervious to known uncertainties.

6.4 List of Publications

6.4.1 Journal Articles

I. J. Horing, I. D. Boyd, and K. K. Maute, "Fully Coupled Analysis of Aero-thermo-elastic Deformation of a Scramjet Inlet," *Journal of Propulsion and Power*, 40.2 (2024), pp. 181-193.

II. J. Horing, K. K. Maute, and I. D. Boyd, "Aerothermodynamic Sensitivity Analysis and Optimization of a Re-entry Vehicle," *Journal of Thermophysics and Heat Transfer*, Articles in Advance.

6.4.2 Conference Papers

I. J. Horing, I. D. Boyd, and K. K. Maute, "Influence of Active Cooling on Engine Performance by Mitigating Aero-thermo-elastic Deformation of a Hypersonic Inlet," *25th AIAA International Space Planes and Hypersonic Systems and Technologies Conference*, AIAA 2023-3037, May 2023. (**Best Paper**)

II. J. Horing, K. K. Maute, and I. D. Boyd, "Effects of a Deforming Scramjet Inlet on Hypersonic Vehicle Performance and Sensitivity," *JANNAF 40th Air-breathing Propulsion*, December 2023.

III. J. Horing, K. K. Maute, and I. D. Boyd, "Coupling Strategies for Analyzing Fluid-Thermal-Structural Interactions of a Hypersonic Inlet" *AIAA SCITECH 2024 Forum*, AIAA 2024-0110, January 2024.

IV. J. Horing, I. D. Boyd, and K. K. Maute, "Influence of Freestream Fluctuations on the Aero-thermo-elastic Behavior of a Hypersonic Inlet," *26th Conference of the International Society for Air Breathing Engines*, ISABE-2024-056, September 2024.

V. J. Horing, K. K. Maute, and I. D. Boyd, "Aerothermodynamic Sensitivity Analysis and Optimization of a Re-entry Vehicle," *AIAA SCITECH 2025 Forum*, AIAA 2025-1560, January 2025.

VI. J. Horing, K. K. Maute, and I. D. Boyd, "Aero-thermo-elastic Sensitivity Analysis of a Scramjet Inlet," *AIAA SCITECH 2025 Forum*, AIAA 2025-1720, January 2025.

Bibliography

[1] NIST/SEMATECH e-Handbook of Statistical Methods. https://doi.org/10.18434/M32189, 2012.

[2] Gridpro til manual. Distributed by Program Development Company, White Plains, NY, USA, 2023.

[3] A Sh Agazhanov, DA Samoshkin, and Yu M Kozlovskii. Thermophysical properties of inconel 718 alloy. In Journal of Physics: Conference Series, volume 1382, page 012175. IOP Publishing, 2019.

[4] John David Anderson. Modern Compressible Flow: With Historical Perspective, volume 12. McGraw-Hill New York, 1990.

[5] Pedro PB Araújo, Marcos VS Pereira, George S Marinho, João FA Martos, and Paulo GP Toro. Optimization of scramjet inlet based on temperature and mach number of supersonic combustion. Aerospace science and technology, 116:106864, 2021.

[6] Jasbir Singh Arora. Introduction to Optimum Design. Elsevier, 2004.

[7] R. A. Baurle and R. L. Gaffney. Extraction of one-dimensional flow properties from multidimensional data sets. Journal of Propulsion and Power, 24(4):704–714, 2008.

[8] I. E. Beckwith, M. Cooper, J. J. Gallagher, and J. J. Jones. Heat-Transer Measurements On a Concave Hemispherical Nose Shape With Unsteady-Flow Effects At Mach Numbers of 1.98 and 4.95. Technical Report NACA-RM-L58D25A, NASA Langley Research Center, July 1958.

[9] Sudip Bhattrai, Liam P McQuellin, Gaetano Currao, Andrew Neely, and David R Buttsworth. Experimental study of the aeroelastic response and performance of a hypersonic intake. In 23rd AIAA International Space Planes and Hypersonic Systems and Technologies Conference, 2020.

[10] Sudip Bhattrai, Liam P. McQuellin, Gaetano M. D. Currao, Andrew J. Neely, and David R. Buttsworth. Experimental Study of Aeroelastic Response and Performance of a Hypersonic Intake Ramp. Journal of Propulsion and Power, 38(1):158–170, January 2022.

[11] Sudip Bhattrai, Andrew Neely, Gaetano Currao, and Liam McQuellin. Impact of aeroelasticity on hypersonic intake performance. In International Society of Airbreathing Engines, pages ISABE–2019–24059, 09 2019.

[12] Frederick Billig, Robert Baurle, Chung-Jen Tam, and Stephen Wornom. Design and analysis of streamline traced hypersonic inlets. In 9th international space planes and hypersonic systems and technologies conference, page 4974, 1999.

[13] Byrenn Birch, David Buttsworth, and Fabian Zander. Measurements of freestream density fluctuations in a hypersonic wind tunnel. Experiments in Fluids, 61(7):158, 2020.

[14] F G Blottner, M Johnson, and M Ellis. Chemically reacting viscous flow program for program multi-component gas mixtures. Sandia National Lab., SC-RR-70-754, 1 1971.

[15] Kevin G Bowcutt. Multidisciplinary optimization of airbreathing hypersonic vehicles. Journal of Propulsion and Power, 17(6):1184–1190, 2001.

[16] Kevin G Bowcutt. Physics drivers of hypersonic vehicle design. In 22nd AIAA International Space Planes and Hypersonics Systems and Technologies Conference, page 5373, 2018.

[17] Kevin G Bowcutt, John D Anderson, and Diego Capriotti. Viscous optimized hypersonic waveriders. In 25th AIAA Aerospace Sciences Meeting, page 272, 1987.

[18] Matthew Brown, Neil Mudford, Andrew Neely, and Tapabrata Ray. Robust design optimization of two-dimensional scramjet inlets. In 14th AIAA/AHI Space Planes and Hypersonic Systems and Technologies Conference, page 8140, 2006.

[19] Paige B. Burbank and Robert L. Stallings. Heat-Transfer and Pressure Measurements on a Flat-Face Cylinder at a Mach Number Range of 2.49 to 4.44. Technical Report NASA-TM-X-19, NASA Langley Research Center, August 1959.

[20] Adolf Busemann. Die achsensymmetrische kegelige überschallströmung. Luftfahrtforschung, 19(4):137–144, 1942.

[21] David R Buttsworth. Ludwieg tunnel facility with free piston compression heating for supersonic and hypersonic testing. In Proceedings of the 9th Australian Space Science Conference (ASSC 2009). University of Southern Queensland, 2010.

[22] Graham V Candler, Heath B Johnson, Ioannis Nompelis, Vladimyr M Gidzak, Pramod K Subbareddy, and Michael Barnhardt. Development of the us3d code for advanced compressible and reacting flow simulations. In 53rd AIAA Aerospace Sciences Meeting, page 1893, 2015.

[23] Ismail Celik, U Ghia, P.J. Roache, Chris Freitas, H Coloman, and Peter Raad. Procedure for Estimation and Reporting of Uncertainty Due to Discretization in CFD Applications. Journal of Fluids Engineering, 130(078001), July 2008.

[24] Sean R Copeland, Francisco Palacios, and Juan J Alonso. Adjoint-based aerothermodynamic shape design of hypersonic vehicles in non-equilibrium flows. In 52nd Aerospace Sciences Meeting, page 0513, 2014.

[25] Adam Culler and Jack McNamara. Coupled Flow-Thermal-Structural Analysis for Response Prediction of Hypersonic Vehicle Skin Panels. In 51st AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Structures, Structural Dynamics, and Materials and Co-located Conferences. American Institute of Aeronautics and Astronautics, April 2010.

[26] Adam Culler and Jack McNamara. Fluid-Thermal-Structural Modeling and Analysis of Hypersonic Structures under Combined Loading. In 52nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, Structures, Structural Dynamics, and Materials and Co-located Conferences. American Institute of Aeronautics and Astronautics, April 2011.

[27] Adam J. Culler and Jack J. McNamara. Studies on Fluid-Thermal-Structural Coupling for Aerothermoelasticity in Hypersonic Flow. AIAA Journal, 48(8):1721–1738, August 2010.

[28] Gaetano MD Currao, Sudip Bhattrai, and Andrew J Neely. Fluid-thermal-structure interactions on a hypersonic two-dimensional intake. In ISABE Conference, 2019.

[29] Kyle A Damm, Rowan J Gollan, Peter A Jacobs, Michael K Smart, Seonguk Lee, Eunsa Kim, and Chongam Kim. Discrete adjoint optimization of a hypersonic inlet. AIAA Journal, 58(6):2621–2634, 2020.

[30] João VMB de Siqueira, Mauricio AP Rosa, and Guilherme B Ribeiro. Three-dimensional cfd investigation of a scramjet inlet under different freestream conditions. Thermal Science and Engineering Progress, 27:101051, 2022.

[31] Pramote Dechaumphai, Earl A. Thornton, and Allan R. Wieting. Flow-thermal-structural study of aerodynamically heated leading edges. Journal of Spacecraft and Rockets, 26(4):201–209, July 1989.

[32] Nicholas J DiGregorio, Thomas K West, and Seongim Choi. Robust design under uncertainty of hypersonic inlets. In AIAA AVIATION 2022 Forum, page 3862, 2022.

[33] Barry Edney. Anomalous heat transfer and pressure distributions on blunt bodies at hypersonic speeds in the presence of an impinging shock. Technical report, Flygtekniska Forsoksanstalten, Stockholm (Sweden), 1 1968.

[34] Walter C. Engelund. Hyper-x aerodynamics: The x-43a airframe-integrated scramjet propulsion flight-test experiments. Journal of Spacecraft and Rockets, 38(6):801–802, 2001.

[35] Sinan Eyi, Kyle M. Hanquist, and Iain D. Boyd. Aerothermodynamic Design Optimization of Hypersonic Vehicles. Journal of Thermophysics and Heat Transfer, 33(2):392–406, April 2019. Publisher: American Institute of Aeronautics and Astronautics.

[36] Sinan Eyi, Kyle M. Hanquist, and Iain D. Boyd. Shape Optimization of Reentry Vehicles to Minimize Heat Loading. Journal of Thermophysics and Heat Transfer, 33(3):785–796, July 2019. Publisher: American Institute of Aeronautics and Astronautics.

[37] PM Geuzaine. An implicit finite volume method for compressible turbulent flows on unstructured meshes. 2003.

[38] Frederick W Gibson and John S Mixson. Flutter investigation at a mach number of 7.2 of models of the horizontal-tail and vertical-tail surfaces of the x-15 airplane. Technical report, 1959.

[39] Peter A Gnoffo. Conservation Equations and Physical Models for Hypersonic Air Flows in Thermal and Chemical Nonequilibrium, volume 2867. National Aeronautics and Space Administration, Office of Management, Scientific and Technical Information Division, 1989.

[40] Peter A Gnoffo. An upwind-biased, point-implicit relaxation algorithm for viscous, compressible perfect-gas flows. Technical report, No. NASA-TP-2953., 1990.

[41] David G. Goodwin, Raymond L. Speth, Harry K. Moffat, and Bryan W. Weber. Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes. `https://www.cantera.org`, 2021. Version 2.5.1.

[42] Dai Guangyue, Lei Zeng, Hongyin Jia, Lei Liu, Bo Qiu, and Yewei Gui. Preliminary study on the influence of aerothermoelastic deformation on 2-d hypersonic inlet performance. In 21st AIAA International Space Planes and Hypersonics Technologies Conference, Xiamen, China, March 2017. American Institute of Aeronautics and Astronautics.

[43] Jürgen Häberle and Ali Gülhan. Investigation of two-dimensional scramjet inlet flowfield at mach 7. Journal of Propulsion and Power, 24(3):446–459, 2008.

[44] Joseph Hank, James Murphy, and Richard Mutzman. The x-51a scramjet engine flight demonstration program. In 15th AIAA International Space Planes and Hypersonic Systems and Technologies Conference, page 2540, 2008.

[45] Joseph Hank, James Murphy, and Richard Mutzman. The x-51a scramjet engine flight demonstration program. In 15th AIAA International Space Planes and Hypersonic Systems and Technologies Conference, AIAA-2008-2540, April 2008.

[46] William H Heiser, David T Pratt, and Daniel H Daley. Hypersonic Airbreathing Propulsion. AIAA, Washington, DC, 1994.

[47] Iñigo Hernando, Mario Alfredo Renderos, Magdalena Cortina, Jose Exequiel Ruiz, Jon Iñaki Arrizubieta, and Aitzol Lamikiz. Inconel 718 laser welding simulation tool based on a moving heat source and phase change. Procedia CIRP, 74:674–678, 2018.

[48] Oliver Hohn and Ali Gülhan. Experimental investigation on the influence of yaw angle on the inlet performance at mach 7. In 48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, page 938, 2010.

[49] Michael E. Holloway, Ross S. Chaudhry, and Iain D. Boyd. Assessment of hypersonic double-cone experiments for validation of thermochemistry models. Journal of Spacecraft and Rockets, 59(2):389–400, 2022.

[50] R. N. Hopko and H. K. Strass. Some Experimental Heating Data on Convex and Concave Hemispherical Nose Shapes and Hemispherical Depressions on a 30 Deg Blunted Nose Cone. Technical Report NACA-RM-L58A17A, NASA Langley Research Center, March 1958.

[51] Jennifer Horing, Iain D Boyd, and Kurt K Maute. Influence of active cooling on engine performance by mitigating aero-thermo-elastic deformation of a hypersonic inlet. In 25th AIAA International Space Planes and Hypersonic Systems and Technologies Conference, AIAA-2023-3037, May 2023.

[52] Jennifer Horing, Iain D Boyd, and Kurt K Maute. Influence of freestream fluctuations on the aero-thermo-elastic behavior of a scramjet inlet. In International Society of Airbreathing Engines 2024, ISABE-2024-056, May 2023.

[53] Jennifer Horing, Kurt Maute, and Iain D Boyd. Aero-thermo-elastic sensitivity analysis of a scramjet inlet. In AIAA SCITECH 2025 Forum, page 1720, 2025.

[54] Jennifer Horing, Kurt Maute, and Iain D Boyd. Aerothermodynamic sensitivity analysis and optimization of a hypersonic re-entry vehicle. Journal of Thermophysics and Heat Transfer, 0(0):1–13, 2025.

[55] Jennifer Horing, Kurt Maute, and Iain D Boyd. Aerothermodynamic sensitivity analysis and optimization of a hypersonic re-entry vehicle. In AIAA SCITECH 2025 Forum, page 1560, 2025.

[56] Jennifer Horing, Kurt K Maute, and Iain D Boyd. Effects of a deforming scramjet inlet on hypersonic vehicle performance and sensitivity. In JANNAF PIB/CS/APS/EPSS/ESHS Joint Meetings, 2023.

[57] Jennifer Horing, Kurt K Maute, and Iain D Boyd. Coupling strategies for analyzing fluid-thermal-structural interactions of a hypersonic inlet. In AIAA SCITECH 2024 Forum, page 0110, 2024.

[58] Jennifer A. Horing, Iain D. Boyd, and Kurt K. Maute. Fully coupled analysis of aerothermoelastic deformation of a scramjet inlet. Journal of Propulsion and Power, 40(2):181–193, 2024.

[59] Kevin Jackson, Mark Gruber, and Todd Barhorst. The hifire flight 2 experiment: an overview and status update. In 45th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, page 5029, 2009.

[60] Kevin R. Jackson, Mark R. Gruber, and Salvatore Buccellato. Mach 6–8+ hydrocarbon-fueled scramjet flight experiment: The hifire flight 2 project. Journal of Propulsion and Power, 31(1):36–53, 2015.

[61] Soudeh Kamali. Development of a High-Fidelity Aero-Thermo-Elastic Analysis and Design Capability. University of Wyoming, 2021.

[62] Soudeh Kamali, Dimitri Mavriplis, and Evan Anderson. Development and validation of an aerothermoelastic analysis capability. AIAA Journal, 60(1):287–301, 2022.

[63] Soudeh Kamali, Dimitri J Mavriplis, and Evan M Anderson. Sensitivity analysis for aero-thermo-elastic problems using the discrete adjoint approach. In AIAA Aviation 2020 Forum, page 3138, 2020.

[64] Arthur Kantrowitz et al. Preliminary Investigation of Supersonic Diffusers. National Advisory Committee for Aeronautics, 1945.

[65] Scott D Kasen. Thermal management at hypersonic leading edges. University of Virginia, 2013.

[66] Heather L Kline, Thomas D Economon, and Juan J Alonso. Mulit-objective optimization of a hypersonic inlet using generalized outflow boundary conditions in the continuous adjoint method. In 54th AIAA Aerospace Sciences Meeting, page 0912, 2016.

[67] Heather L. Kline, Francisco Palacios, and Juan J. Alonso. Sensitivity of the Performance of a 3-Dimensional Hypersonic Inlet to Shape Deformations. In 19th AIAA International Space Planes and Hypersonic Systems and Technologies Conference, Atlanta, GA, June 2014.

[68] HL Kline and JJ Alonso. Adjoint of generalized outflow-based functionals applied to hypersonic inlet design. AIAA Journal, 55(11):3903–3915, 2017.

[69] DA Knoll, PR McHugh, and David Elliot Keyes. Newton-krylov methods for low-mach-number compressible combustion. AIAA Journal, 34(5):961–967, 1996.

[70] Robert H Korkegi. Comparison of shock-induced two-and three-dimensional incipient turbulent separation. AIAA Journal, 13(4):534–535, 1975.

[71] David Kuntz, Basil Hassan, and Donald Potter. An iterative approach for coupling fluid/thermal predictions of ablating hypersonic vehicles. In 33rd Thermophysics Conference, page 3460. American Institute of Aeronautics and Astronautics.

[72] David W Kuntz, Basil Hassan, and Donald L Potter. Predictions of ablating hypersonic vehicles using an iterative coupled fluid/thermal approach. Journal of Thermophysics and Heat Transfer, 15(2):129–139, 2001.

[73] Nicolas Lamorte and Peretz Friedmann. Hypersonic aeroelastic stability boundary computations using radial basis functions for mesh deformation. In 18th AIAA/3AF International Space Planes and Hypersonic Systems and Technologies Conference, page 5943, 2012.

[74] Nicolas Lamorte and Peretz P Friedmann. Hypersonic aeroelastic and aerothermoelastic studies using computational fluid dynamics. AIAA Journal, 52(9):2062–2078, 2014.

[75] Nicolas Lamorte, Peretz P. Friedmann, Derek J. Dalle, Sean M. Torrez, and James F. Driscoll. Uncertainty Propagation in Integrated Airframe–Propulsion System Analysis for Hypersonic Vehicles. Journal of Propulsion and Power, 31(1):54–68, January 2015.

[76] J.-H. Lee. Basic governing equations for the flight regimes of aeroassisted orbital transfer vehicles. In 19th Thermophysics Conference, AIAA-1984-1729, June 1984.

[77] Liu Lei, Dai Guangyue, Zeng Lei, Wang Zhenfeng, and Gui Yewei. Experimental model design and preliminary numerical verification of fluid–thermal–structural coupling problem. AIAA Journal, 57(4):1715–1724, 2019.

[78] Jack Levine and Charles B. Rumsey. Heat-Transfer Measurements on a 5.5- Inch-Diameter Hemispherical Concave Nose in Free Flight at Mach Numbers up to 6.6. Technical Report NASA-MEMO-10-21-58L, NASA Langley Research Center, December 1958.

[79] Siyi Li, Lei Liu, Shenshen Liu, Xiaofeng Yang, Ziyi Wang, and Yanxia Du. Analysis of intake performance of forebody and inlet on pre-deformation. In Journal of Physics: Conference Series, volume 2458, page 012028. IOP Publishing, 2023.

[80] Walter Thomas Maier. A Discrete Adjoint Framework for Turbulent Hypersonic Flows in Thermochemical Nonequilibrium. Stanford University, 2023.

[81] Laurie Marshall, Catherine Bahm, Griffin Corpening, and Robert Sherrill. Overview with results and lessons learned of the x-43a mach 10 flight. In AIAA/CIRA 13th International Space Planes and Hypersonics Systems and Technologies Conference, page 3336, 2005.

[82] Alexandre Martin, Leonardo C Scalabrin, and Iain D Boyd. High performance modeling of atmospheric re-entry vehicles. Journal of Physics: Conference Series, 341(1):012002, feb 2012.

[83] Alexandre Martin, Leonardo C. Scalabrin, and Iain D. Boyd. High performance modeling of atmospheric re-entry vehicles. Journal of Physics: Conference Series, 341(1):012002, February 2012.

[84] Joaquim RRA Martins and Andrew Ning. Engineering Design Optimization. Cambridge University Press, 2021.

[85] Dimitri J Mavriplis. Discrete adjoint-based approach for optimization problems on three-dimensional unstructured meshes. AIAA Journal, 45(4):741–750, 2007.

[86] Jack J. McNamara and Peretz P. Friedmann. Aeroelastic and Aerothermoelastic Analysis in Hypersonic Flow: Past, Present, and Future. AIAA Journal, 49(6):1089–1122, June 2011.

[87] David J Munk, Dries Verstraete, and Gareth A Vio. Effect of fluid-thermal–structural interactions on the topology optimization of a hypersonic transport aircraft wing. Journal of Fluids and Structures, 75:45–76, 2017.

[88] Aaron G Neville and Graham V Candler. Computational-fluid-dynamics-based axisymmetric aeroshell shape optimization in hypersonic entry conditions. Journal of Spacecraft and Rockets, 52(1):76–88, 2015.

[89] L Noël, M Schmidt, Christian Messe, JA Evans, and Kurt Maute. Adaptive level set topology optimization using hierarchical b-splines. Structural and Multidisciplinary Optimization, 62:1669–1699, 2020.

[90] Oğuz K Onay and Sinan Eyi. Ablation analyses of optimized nose tips for hypersonic vehicles. Journal of Thermophysics and Heat Transfer, 34(1):78–89, 2020.

[91] Chul Park. Assessment of two-temperature kinetic model for ionizing air. Journal of Thermophysics and Heat Transfer, 3(3):233–244, 1989.

[92] Logan P. Riley and Andrew Zakrajsek. Thermal-Mechanical Deflection Effects on Aero-Propulsive Performance of a Reference High-Speed Air-Breathing Vehicle. In AIAA Propulsion and Energy 2021 Forum, AIAA Propulsion and Energy Forum. American Institute of Aeronautics and Astronautics, July 2021.

[93] Zachary B Riley and Jack J McNamara. Interaction between compliant structures and boundary-layer transition in hypersonic flow. AIAA Journal, 55(8):2645–2663, 2017.

[94] Zachary B Riley, Ricardo A Perez, Gregory W Bartram, S Michael Spottswood, Benjamin P Smarslok, and Timothy J Beberniss. Aerothermoelastic experimental design for the aedc/vkf tunnel c: Challenges associated with measuring the response of flexible panels in high-temperature, high-speed wind tunnels. Journal of Sound and Vibration, 441:96–105, 2019.

[95] Patrick E Rodi. Reduced edney type-iv cowl shock-on-lip heating by leading edge geometry optimization. In 54th AIAA Aerospace Sciences Meeting, page 0911, 2016.

[96] Gianluigi Rozza, Anwar Koshakji, Alfio Quarteroni, et al. Free form deformation techniques applied to 3d shape optimization problems. Communications in Applied and Industrial Mathematics, 4:1–26, 2013.

[97] Yousef Saad. Iterative Methods for Sparse Linear Systems. Other Titles in Applied Mathematics. Society for Industrial and Applied Mathematics, January 2003.

[98] Leonardo Scalabrin and Iain Boyd. Numerical simulations of the fire-ii convective and radiative heating rates. In 39th AIAA Thermophysics Conference, 2007, AIAA Paper 2007-4044.

[99] Leonardo C Scalabrin. Numerical Simulation of Weakly Ionized Hypersonic Flow Over Reentry Capsules. PhD thesis, University of Michigan, 2007.

[100] Mathias Schmidt, Lise Noël, Keenan Doble, John A Evans, and Kurt Maute. Extended isogeometric analysis of multi-material and multi-physics problems using hierarchical b-splines. Computational Mechanics, 71(6):1179–1203, 2023.

[101] Christopher Seager and Ramesh K Agarwal. Hypersonic blunt-body shape optimization for reducing drag and heat transfer. Journal of Thermophysics and Heat Transfer, 31(1):48–55, 2017.

[102] M.M. Selim and R.P. Koomullil. Mesh Deformation Approaches – A Survey. Journal of Physical Mathematics, 7(2):1 – 9, 2016.

[103] Ascher H Shapiro. The Dynamics and Thermodynamics of Compressible Fluid Flow. New York: Ronald Press, 1953.

[104] Vilas J Shinde, Datta V Gaitonde, and Jack J McNamara. Supersonic turbulent boundary-layer separation control using a morphing surface. AIAA Journal, 59(3):912–926, 2021.

[105] Jeffrey P Slotnick, Abdollah Khodadoust, Juan Alonso, David Darmofal, William Gropp, Elizabeth Lurie, and Dimitri J Mavriplis. Cfd vision 2030 study: A path to revolutionary computational aerosciences. Technical report, No. NF1676L-18332., 2014.

[106] Michael Smart. Scramjets. The Aeronautical Journal, 111(1124):605–619, 2007.

[107] Michael K Smart. How much compression should a scramjet inlet do? AIAA Journal, 50(3):610–619, 2012.

[108] MK Smart. Design of three-dimensional hypersonic inlets with rectangular-to-elliptical shape transition. Journal of Propulsion and Power, 15(3):408–416, 1999.

[109] MK Smart. Optimization of two-dimensional scramjet inlets. Journal of Aircraft, 36(2):430–433, 1999.

[110] S Michael Spottswood, Timothy J Beberniss, Thomas G Eason, Ricardo A Perez, Jeffrey M Donbar, David A Ehrhardt, and Zachary B Riley. Exploring the response of a thin, flexible panel to shock-turbulent boundary-layer interactions. Journal of Sound and Vibration, 443:74–89, 2019.

[111] Thomas H Squire and Jochen Marschall. Material property requirements for analysis and design of uhtc components in hypersonic applications. Journal of the European Ceramic Society, 30(11):2239–2251, 2010.

[112] Wei-Yi Su, Hang An, and Mou-Yuan Wang. Fluid–thermal–structure interaction of hypersonic inlets under different aspect ratios. AIAA Journal, pages 1–13, 2023.

[113] Kenneth Sutton and Peter Gnoffo. Multi-component diffusion with application to computational aerothermodynamics. In 7th AIAA/ASME Joint Thermophysics and Heat Transfer Conference, AIAA-1998-2575, June 1998.

[114] John E Theisinger and Robert D Braun. Multi-objective hypersonic entry aeroshell shape optimization. Journal of Spacecraft and Rockets, 46(5):957–966, 2009.

[115] Milton O Thompson. At the edge of space: the X-15 flight program. Smithsonian Institution, 2013.

[116] Earl A Thornton and Pramote Dechaumphai. Coupled flow, thermal, and structural analysis of aerodynamically heated panels. Journal of aircraft, 25(11):1052–1059, 1988.

[117] David M Van Wie. Hypersonics: Past, present, and potential future. Johns Hopkins APL Technical Digest, 35(4):335–341, 2021.

[118] Paulo B Vasconcelos, Liam P McQuellin, Krishna M Talluru, and Andrew J Neely. High-speed fluid-structure interactions on a compliant panel under shock impingement. AIAA Journal, 61(3):1077–1094, 2023.

[119] Walter G Vincenti, Charles H Kruger Jr, and T Teichmann. Introduction to Physical Gas Dynamics. American Institute of Physics, New York, 1966.

[120] Antonio Viviani, Andrea Aprovitola, Luigi Iuspa, and Giuseppe Pezzella. Aeroshape design of reusable re-entry vehicles by multidisciplinary optimization and computational fluid dynamics. Aerospace Science and Technology, 105:106029, 2020.

[121] Gary Weickum, MS Eldred, and Kurt Maute. A multi-point reduced-order modeling approach of transient structural dynamics with application to robust design optimization. Structural and Multidisciplinary Optimization, 38:599–611, 2009.

[122] ME White. The national aerospace plane program and the apl role. Johns Hopkins APL Technical Digest, 13(1):218–229, 1992.

[123] PW White and J Hoffman. Earth global reference atmospheric model (earth-gram): User guide. Technical report, No. NASA/TM-20230014404, 2023.

[124] Charles R Wilke. A viscosity equation for gas mixtures. The Journal of Chemical Physics, 18(4):517–519, 1950.

[125] Michael J Wright, Graham V Candler, and Deepak Bose. Data-parallel line relaxation method for the navier-stokes equations. AIAA Journal, 36(9):1603–1609, 1998.

[126] Michael James Wright. A Family of Data-Parallel Relaxation Methods for the Navier-Stokes Equations. University of Minnesota, 1997.

[127] Kun Ye, Zhengyin Ye, Chunna Li, and Jie Wu. Effects of the aerothermoelastic deformation on the performance of the three-dimensional hypersonic inlet. Aerospace Science and Technology, 84:747–762, 2019.