ABSTRACT

Title of Thesis: OPTIMIZATION AND TRAJECTORY ANALYSIS OF MORPHING WAVERIDERS

John Jessen
Master of Science, 2025

Thesis Directed by: Professor Christoph Brehm
Department of Aerospace Engineering

Professor Johan Larsson
Department of Mechanical Engineering

A waverider is a flight vehicle that creates a shock wave attached to the leading edge, which leads to high lift-to-drag ratio at a design Mach number. A morphing waverider changes the lower surface in order to keep the shock wave attached to the leading edge at a range of Mach numbers. The classic waverider design method creates the vehicle shape with sharp leading edges from a known flow field. This is unrealistic for vehicle design, first because real-world leading edges are rounded, and second because this method limits the design space of possible geometries to a small subset that generate flow fields that can be calculated analytically.

This objective of this fundamental research was to bypass the current limitations facing round-leading-edge waverider optimization by including a reduced-order blunt-leading-edge model and a Computational Fluid Dynamics (CFD) algorithm in the optimization process. A new morphing methodology was created that could apply additional off-design features to the leading

edge and lower surface of planar shock-derived waveriders. An analytical model for off-design waveriders was developed and used to optimize morphing waverider geometries alongside the CHAMPS+ CFD tool. The impact of morphing on waverider range was evaluated using a direct collocation trajectory optimization algorithm.

While the analytical model and CFD tool predicted similar performance trends, the CFD-based method was better suited to capture the complex features experienced by round-leading-edge morphing waveriders. For both methods, the optimized morphing waverider geometries were predicted to give lift-to-drag ratio and range increases over non-morphing waveriders. To provide additional insights into the effects of morphing on round-leading-edge waverider performance, further morphing optimization iterations, as well as the addition of constraints which take into account real-world low-cost access to space considerations, are recommended.

OPTIMIZATION AND TRAJECTORY ANALYSIS OF MORPHING WAVERIDERS

by

John Jessen

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Master of Science
2025

Advisory Committee:
Professor Christoph Brehm, Chair/Advisor
Professor Johan Larsson, Co-Advisor
Professor James Baeder
Dr. Kevin Bowcutt

2025

Acknowledgments

The views expressed in this article are those of the author and do not necessarily reflect the official policy or position of the Air Force, the Department of Defense or the U.S. Government.

I would like to extend my deepest gratitude to those who have supported my thesis research. Firstly, I would like to thank my advisors, whose guidance and mentorship have been essential to the completion of this work.

I am deeply grateful to those who took part in my advisory committee, for their guidance on computational fluid dynamics topics and expertise on waverider design and space flight.

I also would like to express my sincere thanks to the students in the computational research laboratory. As my friends, mentors, coworkers, and code-debuggers, my experience at the University of Maryland would not have been the same without them.

My gratitude also extends to my advisors and the faculty at the U.S. Air Force Academy, whose encouragement and support have provided the foundation for my academic journey.

Finally, to my family, for everything——your love, patience, and belief in me mean more than words can express.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

The 21st Century has witnessed remarkable advances in low-cost space access. The civilian spaceflight industry aims to make commercial suborbitable travel affordable, while significant strides have been made by others in developing cost-effective, reusable launch vehicles. As reusability becomes increasingly common in modern spacecraft and the cost of launching vehicles into orbit continues to decrease, attention is turning once again to the design of high-performing, efficient civilian reentry vehicles. One such vehicle design is the morphing waverider, which could achieve high performance by adjusting its lower surface to generate a shockwave in such a way that it produces high lift-to-drag ratios across a gliding trajectory.

1.1 Reentry Flight

Numerous challenges face orbital reentry vehicles. These include thin shock layers that closely follow the leading edges of vehicles, and complex boundary layers, where entropy layers, interaction between the boundary layer and inviscid flow, and the presence of high-frequency disturbances all affect boundary layer thickness and the transition to turbulent flow [1].

The difficulty of modeling these features significantly limited the performance of early reentry vehicles. Examples of such vehicles include the Apollo Capsule and the Space Shuttle Orbiter, both of which incorporated conservative design features, such as large blunt surfaces.

While these features helped mitigate high-temperature effects and ensured these vehicles could survive in this poorly understood regime, they also considerably reduced performance. As a result, high-performing, low cost reentry vehicles were not realized for human space flight during the 20th Century.

Now, however, improvements in computing power and Computational Fluid Dynamics (CFD) capabilities have led to an increased understanding of the physical phenomena present during reentry. Additionally, the development of materials capable of withstanding extremely high temperatures, such as ultra-high temperature ceramics, have enabled these vehicles to utilize progressively smaller leading edge rounding. These technological advancements have sparked renewed interest in the waverider, a class of vehicle with theoretical performance far superior to those of traditional reentry vehicles.

1.2 The Waverider

The modern waverider was proposed in 1959 by Terence Nonweiler as a reentry vehicle that could achieve high performance [2,3]. The waverider forms he described were caret-shaped, with sharp, triangular planforms when viewed from above. These designs produced a planar, attached shock at sharp leading edges. The upper surface aligned with the freestream flow, while the lower surface followed the disturbed flow resulting from the planar shock at the leading edge. The rear surface was tangent to the freestream. Waverider performance was estimated by integrating the shockwave-generated pressure at the surface of vehicle. Using planar shock relations [1], Nonweiler predicted significantly higher lift-to-drag ratios for these vehicles than was standard in other reentry vehicle geometries of the time. For this reason, his approach would

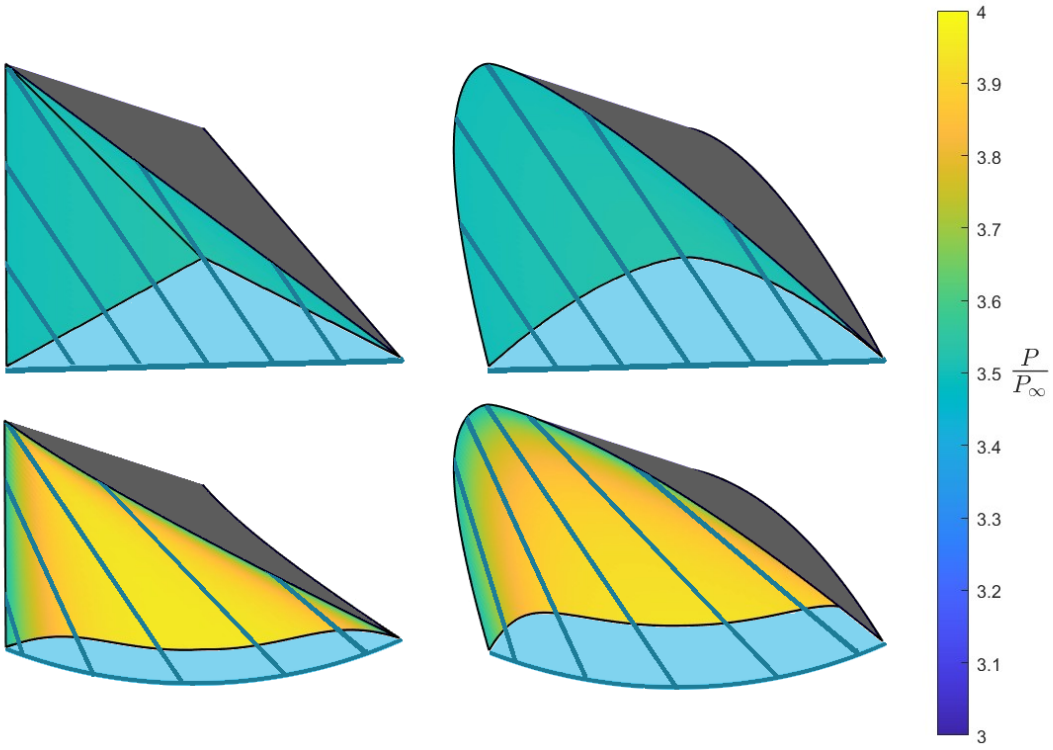

Figure 1.1: Caret and power-law leading edge waverider geometries at Mach 6 and 20 km altitude with lower surface pressure distributions and shock outlines; upper: planar shock generating body; lower: conical shock generating body

serve as the foundation for waverider design and performance evaluation.

This formulation was extended to conical shock waves by Jones [4], who used the Taylor-Maccoll conical shock wave solution [5] to calculate the lower surface streamlines and the pressure distribution. This study also used a caret-type leading edge and demonstrated that the more general shock solution offered increased performance over purely planar waverider designs. Subsequent research on waveriders led to the creation of a power-law formulation for the leading edge [6], which allowed for additional design flexibility. This formulation was later simplified for strictly planar-derived waveriders [7].

The emergence of high-speed computing enabled the study of waveriders that used more complex generating shock shapes which required more complex solutions to their flow fields.

This included cones with elliptical cross-sections or with some inclination [8], and osculating cones [9], which used cones of varying shock angle across the leading edge to create complex lower-surface shock surfaces.

1.2.1 Waverider Optimization

With high-speed computing also came the desire to determine the optimal waverider geometry for any given flight condition. This effort began in 1987 with Bowcutt et al., who investigated conical-shock generating vehicles with an upper expansion surface [10], and was further developed by Corda et al., who used a generalized axisymmetric shock-generating body [11]. These studies are of particular importance—first, because they were the earliest to account for the effect of viscous forces from the boundary layer on performance, and second, because they demonstrated that even after adding viscous effects, optimized waveriders produced lift-to-drag ratios significantly higher than those of other contemporary reentry vehicles.

Although waveriders optimized for a single flight condition achieved impressive performance, CFD studies revealed significant performance decreases when these waveriders were flown in off-design conditions. This included situations where leading edge rounding was applied [12], or where the vehicles were flown in flight conditions they were not optimized for [13]. The primary cause of this performance drop in off-design conditions was often pressure bleed from a shock wave that no longer conformed to the leading edge, as shown in Fig. 2.1. This presents a significant challenge for real-world gliding waveriders, which must operate across a wide range of flight conditions and also feature a round leading edge.

While understanding the effects of leading edge rounding is of the utmost importance for

vehicle design, its true impact on waverider optimization has only been investigated a handful of times by researchers. Among the most significant are studies by Guo et al., which used a complex analytical model that predicted some effects of leading edge rounding to optimize rounded conical shock waveriders, including viscous effects [17]. This study was similar to previous efforts that used Newtonian theory to approximate round-leading-edge drag [18], but also used symbolic regression, which uses optimization algorithms to fit equations to complex physical phenomena, in order to develop a more advanced model that accounts for stream-wise pressure bleed effects. These studies are of great importance to this thesis, first, because they demonstrated an initial methodology to account for off-design effects in waverider design, and second, because of their finding that geometric optimization which incorporates leading edge blunting results in vastly different geometries and significantly improved performance, compared to theoretical 'on-design' waveriders. Despite the impressive capabilities of this model, it is important to note that it only accounts for some of the effects of leading edge rounding on waverider performance, and its results have only been compared to computational fluid dynamics results geometries with minimal pressure bleed. For this reason, it remains likely that computational fluid dynamics is the best method for accurately optimizing round-leading-edge waveriders.

1.2.2 Morphing Waveriders

Multiple design methods have been explored to mitigate the off-design effects encountered when flying at multiple flight conditions. As discussed by Zhao et al., these include "star-body," "tandem," and "parallel" waverider types, among others, all of which involve combining the features of multiple waveriders optimized for various flight conditions to enhance overall glide

performance [14]. The challenges with these design approaches are twofold: first, they typically only include a small number of flight conditions in the optimization process, and second, they are highly subjective approaches which often result in only modest increases in overall performance.

One particularly promising waverider design for operating across a range of flight conditions is the 'morphing' waverider designed by Jesse Maxwell [15, 16]. This design features a movable lower surface that flexes as the Mach number varies, ensuring that the generated shock wave remains attached to the sharp leading edge. Using this approach, Maxwell found that a morphing lower surface significantly improved vehicle lift-to-drag ratio, range, and even maneuverability for various mission types and atmospheres, including Earth or Mars reentry glides.

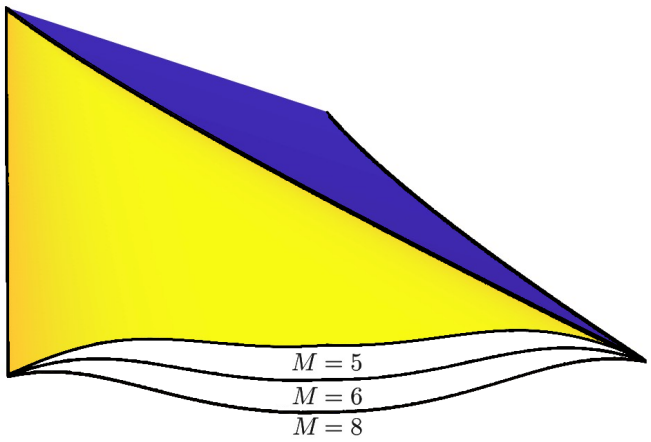

Figure 1.2: Morphing conical-shock derived caret waverider at various Mach numbers

While previous morphing waverider theory offers significant improvements in performance, the use of the analytical Taylor-Maccoll solution to generate waverider geometries required several underlying assumptions that may limit understanding of the true performance achievable using morphing surfaces. First, the requirement of a known conical shock wave shape for waverider generation required that all morphing waverider configurations generated an attached conical

shock wave. Not only does this assumption require significant changes in lower surface shape in-flight, as shown in Fig. 1.2, it artificially limits the geometry design space. The sharp-leading edge analytical method also implicitly assumed that waverider geometries designed such that the shock wave is attached for a sharp leading edge vehicle are optimal, even when leading edge rounding is applied. This is counter to the results found by Guo et al.

These results, along with studies that improved waverider performance after applying arbitrary polynomial perturbations across the lower surface [19], seemed to indicate that morphing waveriders that undergo optimization that include the effects of rounding could have improved performance over those that do not. By implementing CFD simulations to estimate performance, the restrictions imposed by traditional analytical methods are lifted. This study seeks to transition from an improved round-leading-edge analytical model to a CFD-based approach to generate morphing waveriders with improved performance.

1.2.3 Trajectory Analysis

Once the lift and drag forces for a waverider have been determined, it is often necessary to characterize the range and flight path it may achieve by approximating its physical flight path. One method is to use a reduced-order equilibrium-glide model, such as the one used by Maxwell [15], to approximate the general features of a reentry trajectory without accounting for dynamic effects.

Other approaches have employed more advanced techniques, such as guidance and navigation logic [20], to direct vehicle trajectories. Angle-of-attack control, paired with complex genetic algorithm optimization, was used to optimize the trajectory of finned flight systems [21], as well

as waverider-based reentry vehicles [22]. Another approach includes nonlinear direct collocation optimization methods, which model both the system controls and physical states as polynomial splines [23]. These methods are particularly well suited for accounting for many constraints and flight profiles with multiple distinct phases, including rocket launch, or hazard avoidance. This technique was utilized by Rizvi et al. [24], as well as Yu and Nai-Gang [25] to identify optimal trajectories for waverider vehicles with distinct rocket launch, waverider glide, and terminal phases. These studies demonstrated that the high lift-to-drag ratio performance of waverider-class vehicles enables long-range mission capabilities, and that trajectory optimization algorithms can offer significant insight into waverider performance.

1.3 Problem Statement

While the waverider optimization studies mentioned have made significant strides in improving waverider performance using a variety of methods, surface morphing—arguably the most promising of these methods—has not yet been explored using geometry optimization that accounts for off-design effects, such as pressure bleed from rounded leading edges. Additionally, trajectory analyses of morphing waveriders have not been conducted using advanced trajectory optimization tools, and there remains a limited understanding of the true performance improvements offered by morphing technologies. For this reason, this study aims to define optimal morphing waverider vehicles with rounded leading edges and assess the lift-to-drag ratio and trajectory impacts of these vehicles compared to a standard non-morphing waverider.

1.3.1 Objectives

The goal of this thesis is to explore the use of morphing technologies to improve the performance of waveriders for civilian transportation applications. Specifically, this research focuses on improving the range and lift-to-drag ratio of waveriders using analytical methods, computational fluid dynamics, and trajectory optimization. The key objectives of this study are:

1. To create a reduced-order performance evaluation tool for round-leading-edge waveriders using a mixture of analytical relations, and to compare its accuracy to CFD;

2. To identify rounded leading edge waverider geometries that offer improved lift-to-drag ratios, starting from theoretical optimum 'on-design' geometries, by applying arbitrary shape perturbations;

3. To use a trajectory optimization tool to generate morphing geometries that outperform an initial non-morphing design;

4. And to compare the performance impacts of analytical methods versus CFD tools, as well as the associated trends.

1.3.2 Thesis Outline

The following chapter provides the methodology used to generate a waverider and analyze its performance. Mathematical definitions for a round-leading-edge, planar shock-derived waverider with leading edge and lower surface perturbations are given, along with the two methods selected for approximating surface forces and lift-to-drag ratio: a lower-order analytical model,

and a CFD solver. The geometry optimization algorithm is described and validated, followed by details of the waverider dynamics and the trajectory analysis tool.

The third chapter describes the methodology used to optimize fixed and morphing waveriders. The analytical performance model is used to optimize a traditional non-morphing vehicle, and its resulting trajectory is simulated. By analyzing the flight conditions along the trajectory, control points are identified where morphing may occur in order to best improve performance. Using these points, three different morphing vehicles are created, each with different constraints placed on the ability of their lower surfaces or leading edge to adjust during morphing. Finally, additional trajectories are generated for the morphing vehicles, and the lift-to-drag ratio and range performance of each is compared with those of the initial vehicle.

The fourth chapter applies the established methodology to optimize a series of morphing waveriders using the computational fluid dynamics tool. The performance of these vehicles is compared to the initial fixed vehicle, as well as to the optimal vehicles created using the analytical model in the previous chapter. The computational geometries are re-simulated using the analytical model, and the differences between the two methods are discussed.

Finally, important findings are summarized and discussed in the final chapter. These include primarily that CFD tools were well suited for waverider geometry optimization, and also that notable performance increases were predicted as a result of surface morphing. Key weaknesses of the methodology used are discussed, and recommendations for future studies are provided.

Chapter 2: Approach

This chapter details the complete mathematical definition of the waverider-class vehicle, the geometry optimization algorithm, and the performance evaluation of various waverider designs used for the present study. Two distinct techniques are employed to estimate the physical properties on the vehicle's surface: first, an analytical model utilizing tangent-wedge shock relations, Modified Newtonian pressure coefficient theory, and a rear surface pressure correction; and second, a high-efficiency inviscid Euler computational fluid dynamics code. Both methods incorporate Eckert's reference temperature method to estimate skin friction.

This section also provides an overview of the geometry optimization tool, including its setup, constraints, and validation. Finally, the trajectory analysis tool used to approximate vehicle performance is described, highlighting the waverider equations of motion and the trajectory constraints for the glide phase.

2.1 Waverider Construction

Planar shock-derived waveriders are created by first defining a two-dimensional oblique shock based on the desired shock angle and flight condition, followed by defining an outer profile (or planform) that intersected this shock wave. The vehicle surfaces are then created by tracing the upper surface backward from the leading edge, parallel to the freestream direction, and by

tracing the lower surface streamlines aft of the oblique shock. This ensures that the waverider generates the defined oblique shock along its leading edge and captures the high pressure from the shock wave across its lower surface. For a visualization, see Fig. 2.4.

Additional transformations are then considered which perturb the leading edge and lower surface to 'off-design' conditions, such that the shock generated by the leading edge is not necessarily perfectly captured across the lower surface. These operations are performed with the goal of counteracting the effects of leading edge blunting, which is required for physical waverider designs and already has an off-design effect by forming a normal shock at some standoff distance forward of the leading edge. These blunting effects can be seen below in Fig. 2.1, which shows the front-view flow field for an on-design waverider before and after leading edge blunting is applied. When the leading edge is sharp, the on-design waverider captures the entire shock wave across its lower surface, which typically leads to high performance at this flight condition. However, once leading edge blunting is applied, significant pressure bleed from the lower surface occurs. This phenomena tends to considerably reduce vehicle performance [18], and can be easily seen in Fig. 2.17, where significant pressure loss occurs after applying blunting to an on-design sharp leading edge waverider.

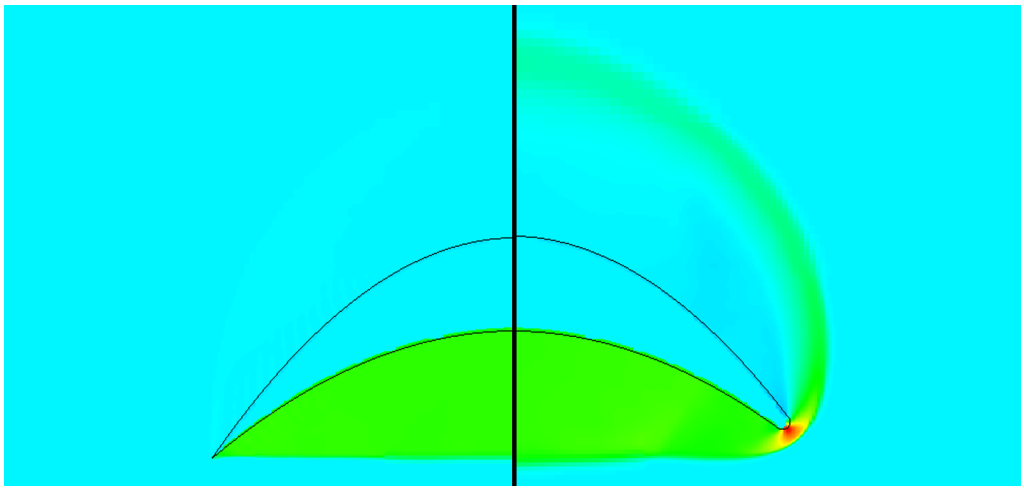

Figure 2.1: Effects of large radius leading edge blunting on an on-design waverider; front-view pressure field

2.1.1 Initial Surface Generation

The waverider leading edge parameterization used in this work is first defined using a power law derivation used by Ryan Starkey [7] to generate planar-shock waveriders. This geometric form was developed in three dimensions, with the coordinate system centered at the waverider's nose-tip. In this system, the waverider extends lengthwise in the positive x-direction, with width extending from the centerline in the y-direction and height extending in the negative z-direction, with positive z defined as pointing upwards. The leading edge is first defined in the x- and y-dimensions using the power-law parameter A.

$$y_{LE}(x) = Ax^n \tag{2.1}$$

This is sometimes rearranged as:

$$x_{LE}(y) = \left(\frac{y}{A}\right)^{1/n} \tag{2.2}$$

with bounds $0 \leq x \leq L$ and $-w/2 \leq y \leq w/2$. A is a sizing constant defined by the vehicle length, w, length, L, and power-law exponent, n. Exponent n can be varied from 0 to 1, and in effect controls the sharpness or bluntness of the upper planform at the nosetip, as shown by Fig. 2.2 in the x-y plane.

$$A = \frac{w}{2L^n} \tag{2.3}$$

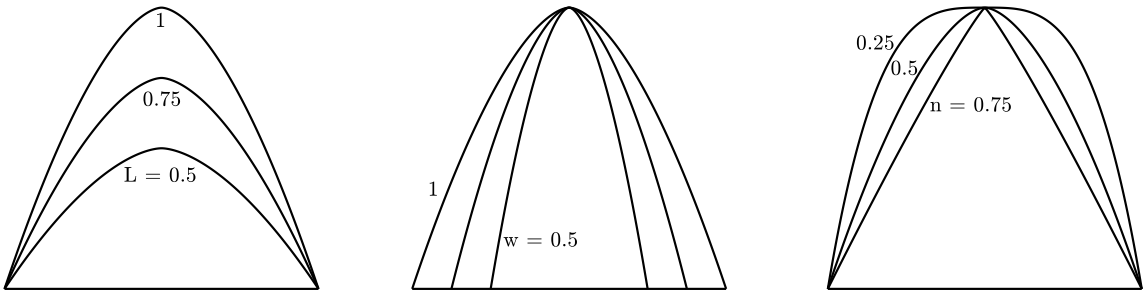

Figure 2.2: Various leading edge shapes; upper planform

The curvature in the x-z plane is derived slightly differently than the standard power-law formulation. Instead of being defined solely by the shape of an oblique shock, an additional curve factor is introduced. This factor applies an off-design curvature to the leading edge:

$$z_{LE}(x) = -x\tan\beta\left(\frac{x}{L}\varepsilon - \varepsilon + 1\right) \tag{2.4}$$

The leading edge is therefore controlled by modifying the desired shock angle, β, and the curve factor ε. The curve factor was bounded between -1 and 1, where $\varepsilon = 0$ represents the standard linear formulation. More negative or positive values of ε introduce concavity or convexity, respectively, across the leading edge. This adjustment typically results in vehicle geometries resembling those of conical-shock derived waveriders and is intended to allow for

additional compression or expansion at the leading edge. Fig. 2.3 illustrates the effect of changing these parameters in the x-z plane:

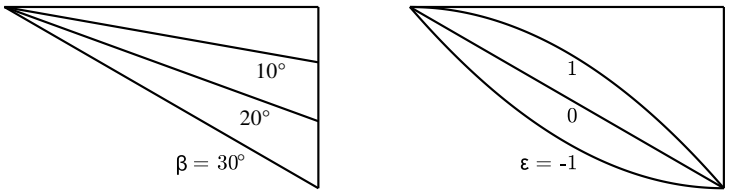

Figure 2.3: Various waverider leading edge shapes; side-view

After defining the leading edge curve, a sharp-leading edge waverider is created by constructing the upper freestream and lower compression surfaces. An initial on-design compression surface is generated using a constant wedge angle, θ. This value is derived using the trigonometric planar shock relation, assuming a very large design Mach number:

$$\theta = \tan^{-1}\left(\frac{2\cot\beta\sin^2\beta}{\gamma + \cos(2\beta)}\right) \tag{2.5}$$

The upper and lower surface z-coordinates are then bounded by the leading and trailing edges, with $0 \leq x \leq L$ and $-y_{LE}(x) \leq y \leq y_{LE}(x)$. The upper surface is created by tracing backwards in the freestream direction, while the lower surface is constructed using the planar shock wedge angle. This process is shown in detail in Fig. 2.4

$$z_{upper}(y) = z_{LE}(x_{LE}(y)) \tag{2.6}$$

$$z_{lower}(x, y) = z_{upper}(y) - [x - x_{LE}(y)]\tan\theta \tag{2.7}$$

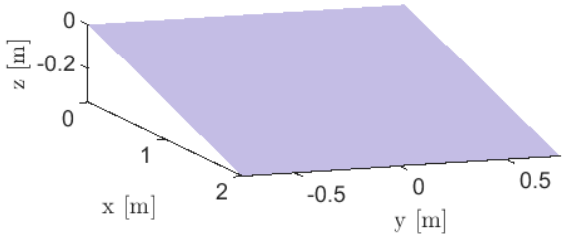

(a) Oblique shock defined by β

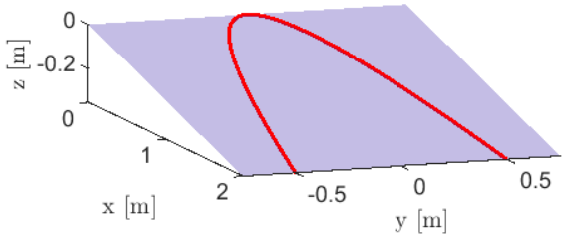

(b) leading edge shape defined by L, w, n, ε

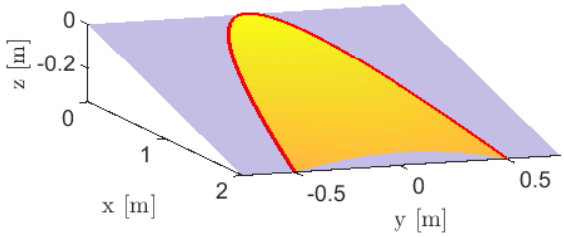

(c) On-design lower surface defined by θ

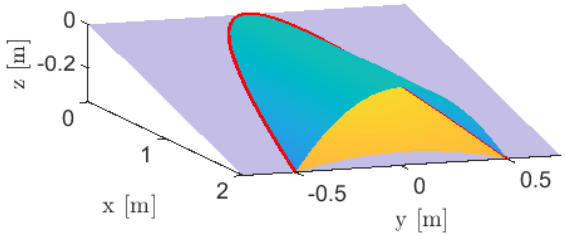

(d) Upper surface defined by freestream

Figure 2.4: Waverider surface generation; $L = 2$, $w = 1$, $n = 0.5$, $\beta = 10°$, $\varepsilon = 0$

2.1.2 Chebyshev Polynomial Perturbation

Next, arbitrary curved perturbations are applied across the lower surface to allow for changes in compression along the span. This is achieved by applying a Chebyshev factor to the wedge angle along the span, using three control points: p_1, p_2, and p_3. The function is determined using Chebyshev polynomials of the first kind [26]:

$$\begin{aligned} T_0(\eta) &= 1 \\ T_1(\eta) &= \eta \\ T_{n+1}(\eta) &= 2\eta T_n(\eta) - T_{n-1}(\eta) \end{aligned} \tag{2.8}$$

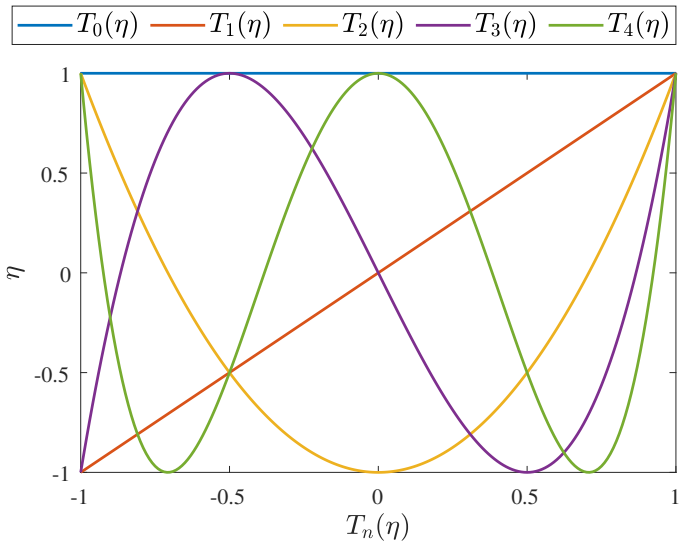

Figure 2.5: Chebyshev polynomials of the first kind

Since Chebyshev polynomials are defined in the dimension $-1 \leq \eta \leq 1$, a coordinate transformation is required to apply them across the span $0 \leq y \leq w/2$ and to mirror them across the centroid. The full fourth-order Chebyshev polynomial, with this axis transformation, is:

$$\begin{aligned} \eta(y) &= \frac{4|y|}{w} - 1 \\ T_n(y) &= T_n[\eta(y)] \\ T(y) &= \sum_{n=0}^{4} a_n T_n(y) \end{aligned} \tag{2.9}$$

To solve for the coefficients a_n, a system of equations is solved using the control points p_1, p_2, and p_3. The solution applies a zero slope condition at the vehicle centroid, $y = 0$, and enforces the leading edge boundary condition at $y = w/2$:

$$\begin{bmatrix} 0 \\ p_1 \\ p_2 \\ p_3 \\ 0 \end{bmatrix} = \begin{bmatrix} T_0(w/2) & T_1(w/2) & T_2(w/2) & T_3(w/2) & T_4(w/2) \\ T_0(w/3) & T_1(w/3) & T_2(w/3) & T_3(w/3) & T_4(w/3) \\ T_0(w/6) & T_1(w/6) & T_2(w/6) & T_3(w/6) & T_4(w/6) \\ T_0(0) & T_1(0) & T_2(0) & T_3(0) & T_4(0) \\ T_0'(0) & T_1'(0) & T_2'(0) & T_3'(0) & T_4'(0) \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \\ a_4 \end{bmatrix} \tag{2.10}$$

Finally, the lower surface is perturbed by applying the Chebyshev polynomial scale factor from Eq. 2.9 to the formulation in Eq. 2.7, resulting in Eq. 2.11. As shown in Fig. 2.6, this allows the wedge angle to be scaled linearly up or down, or to vary across the vehicle span.

$$z_{lower}(x, y) = z_{upper}(y) - [x - x_{LE}(y)] \tan\left(T(y)\theta\right) \tag{2.11}$$

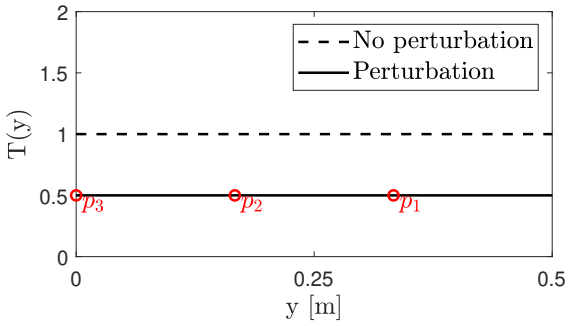

(a) Chebyshev $T(y)$; $p_1 = p_2 = p_3 = 0.5$

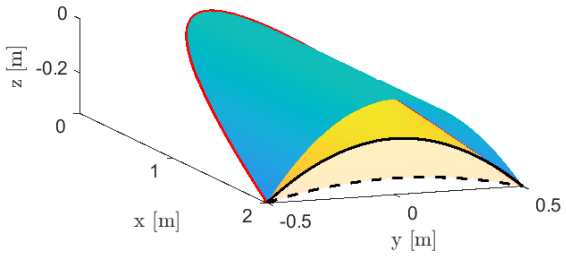

(b) Waverider; $p_1 = p_2 = p_3 = 0.5$

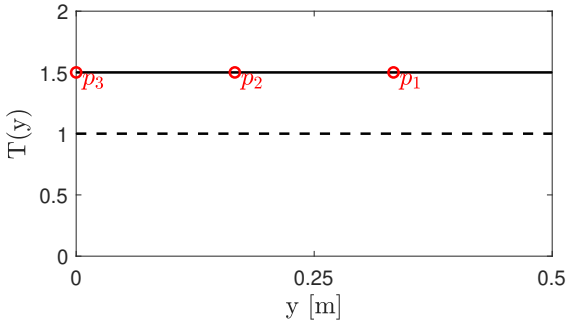

(c) Chebyshev $T(y)$; $p_1 = p_2 = p_3 = 1.5$

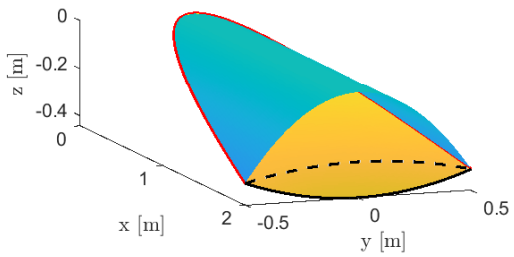

(d) Waverider; $p_1 = p_2 = p_3 = 1.5$

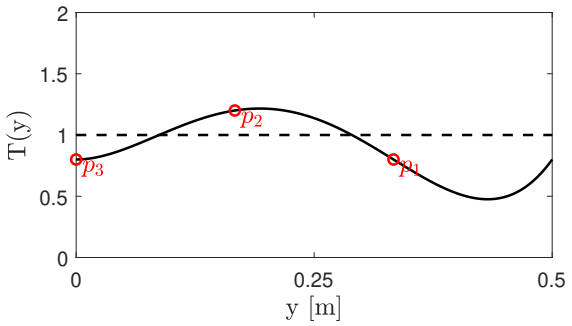

(e) Chebyshev $T(y)$; $p_1 = 1.2$, $p_2 = 0.8$, $p_3 = 1.2$

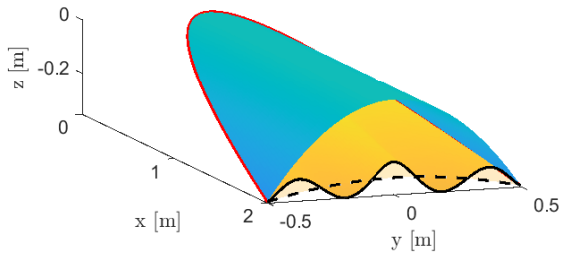

(f) Waverider; $p_1 = 1.2$, $p_2 = 0.8$, $p_3 = 1.2$

Figure 2.6: Waverider with Chebyshev perturbations; $L = 2$, $w = 1$, $n = 0.5$, $\beta = 10°$, $\varepsilon = 0$

2.1.3 Blunted leading edge

At this point, the waverider has an infinitely sharp corner at the leading edge. While this feature would, in theory, provide a perfectly attached shock at the leading edge and result in very high lift-to-drag ratios, it makes construction difficult and likely would not survive the extremely high temperatures that would result from the attached shock in a realistic flight environment. Therefore, a rounded leading edge of some radius, R, is required. The present work rounds the leading edge using the "adding material" method proposed by Tincher and Burnett [27]. While traditional leading edge blunting involves cutting material away from the leading edge to create a fillet, adding material methods may provide improved performance by reducing pressure leakage at the leading edge and preserving the original vehicle shape [15]. To construct the new leading edge, the original upper surface is removed, and a two-dimensional round leading edge is extruded from the lower surface. Next, a new upper surface is traced back from the rounded leading edge in the freestream direction. Finally, a rear surface tangent to the freestream is added, enclosing the vehicle geometry. This final process can be seen in Fig. 2.7.

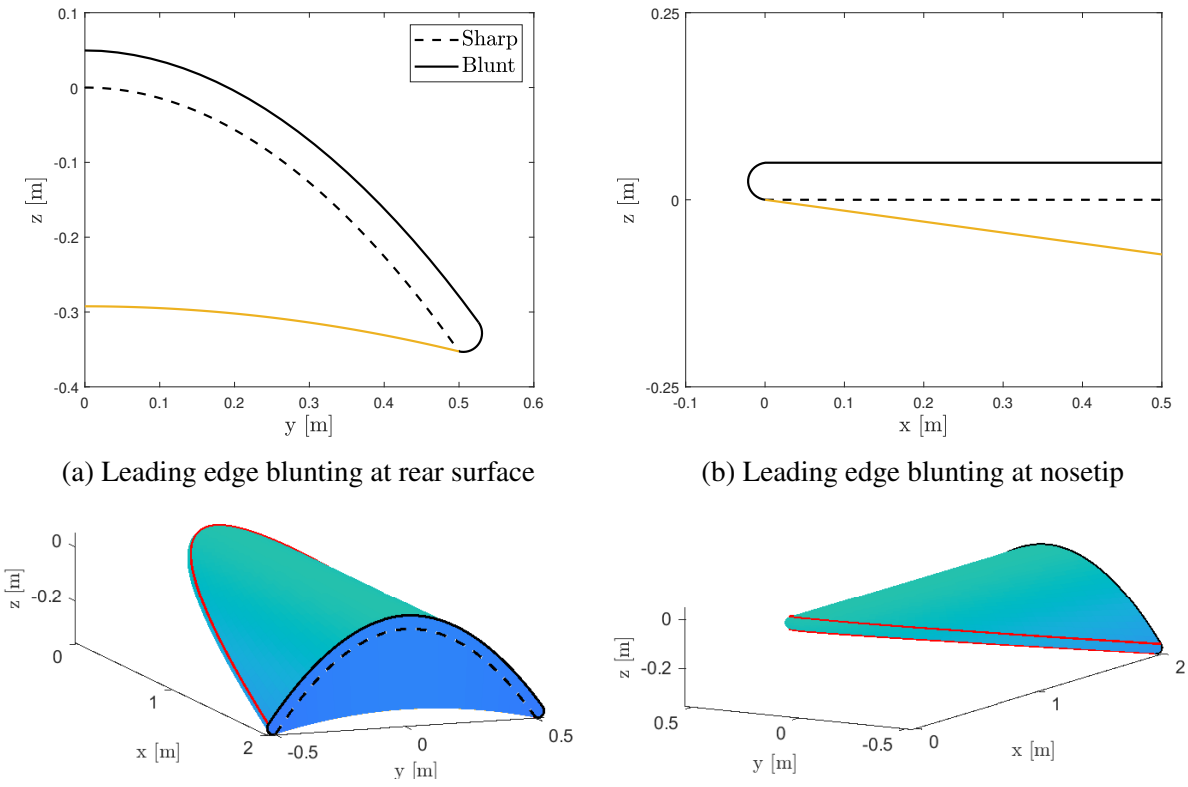

(a) Leading edge blunting at rear surface

(b) Leading edge blunting at nosetip

(c) Waverider rear surface

(d) Waverider leading edge

Figure 2.7: Waverider leading edge blunting full geometry; $L = 2$, $w = 1$, $n = 0.5$, $\beta = 10°$, $\varepsilon = 0$, $p_1 = 1$, $p_2 = 1$, $p_3 = 1$, $R = 0.025$

In summary, the waverider design process outlined provides the flexibility to create a wide range of waverider geometries. These geometries can feature sharp or blunt leading edges and can resemble traditional on-design configurations or include perturbations (ε, p_1, p_2, and p_3) for off-design configurations. Each geometry is defined as a function $f(L, w, n, \beta, \varepsilon, p_1, p_2, p_3, R)$ and is created using the following steps:

1. Define the leading edge using parameters L, w, n, β, and ε along with Eqs. 2.1 through 2.4;

2. Solve the system of equations for the lower surface Chebyshev polynomial $T(y)$ using

parameters p_1, p_2, and p_3, and using Eq. 2.10;

3. Define the lower surface using Eqs. 2.5 and 2.11;

4. If the nose radius is zero, use Eq. 2.6 to define the upper surface. Otherwise, apply the method outlined by Tincher and Burnett to build a round leading edge outward from the lower surface using trigonometric relations. Then, apply a curve-fit to the upper bounds of the new leading edge and create the new upper surface of the waverider with Eq. 2.6;

5. Define the rear surface and enclose the waverider geometry.

2.2 Analytical Model

To estimate the performance of different waverider configurations in a computationally efficient manner, a reduced-order analytical model for waverider aerodynamics is developed. While the aerodynamics of sharp leading edge planar waveriders can be approximated accurately using planar shock and expansion wave theory, these methods do not capture the complex flow characteristics of vehicles with blunt leading edges or those with lower surfaces that are not aligned with the shock wave flow field. For this reason, simple flow relations are paired with pressure corrections for blunt leading edge waveriders. These relations, in addition to skin friction relations, provide a straightforward, albeit lower-order, estimate of waverider performance at a fraction of the cost of CFD.

2.2.1 Inviscid Pressure Models

Different pressure approximations, as described by Anderson [1], are used depending on the relative angle, θ, between the vehicle surface and the freestream. A tangent-wedge method is applied to approximate the pressure, P, on the windward surfaces. This method assumes discontinuous, adiabatic shock waves in a perfect gas. On the leeward surfaces, Prandtl-Meyer expansion wave relations are used, which assume isentropic expansion. Past a critical expansion angle, however, a proposed base pressure model proposed by Maxwell [15] is used. Each relation is described as follows:

$$\frac{P}{P_\infty} = \begin{cases} 1 + \dfrac{2\gamma}{\gamma+1}(M_\infty^2 \sin^2\beta - 1); & \theta \geq 0 \\ \left(1 + \dfrac{\gamma-1}{2} M_\infty \theta\right)^{\frac{2\gamma}{\gamma-1}}; & \dfrac{2}{\gamma-1} M_\infty \left[\left(\dfrac{1}{M_\infty}\right)^{\frac{\gamma-1}{2\gamma}} - 1\right] < \theta < 0 \\ \dfrac{1}{M_\infty}; & \theta \leq -\dfrac{2}{\gamma-1} M_\infty \left[\left(\dfrac{1}{M_\infty}\right)^{\frac{\gamma-1}{2\gamma}} - 1\right] \end{cases} \tag{2.12}$$

where shock angle, β, is approximated by assuming large M_∞ and small angles θ and β [1]:

$$\beta = \theta \left[\frac{\gamma+1}{4} + \sqrt{\left(\frac{\gamma+1}{4}\right)^2 + \frac{1}{M_\infty^2 \theta}}\right] \tag{2.13}$$

At the blunted leading edge, however, oblique shock relations provide a relatively poor pressure approximation. Instead, Modified Newtonian theory [1] is used:

$$c_{p_{max}} = \frac{2}{\gamma M_\infty^2} \left\{ \left[\frac{(\gamma+1)^2 M_\infty^2}{4\gamma M_\infty^2 - 2(\gamma-1)}\right]^{\frac{\gamma}{\gamma-1}} \left[\frac{1 - \gamma + 2\gamma M_\infty^2}{\gamma+1}\right] - 1 \right\} \tag{2.14}$$

$$c_p = c_{p_{max}} \sin^2\theta \tag{2.15}$$

From the definition of pressure coefficient:

$$\frac{P}{P_\infty} = \frac{1}{2} c_p \gamma M_\infty^2 + 1 \tag{2.16}$$

To further improve the pressure approximation for blunt leading edge waveriders, an analytical pressure increment correction formulated by Guo et al. [17] is used. This model, developed using symbolic regression and designed specifically for blunt leading edge waverider geometries, is intended to better account for the effects of stagnation heating on pressure as well as pressure bleed to surfaces aft of the leading edge. It is important to note that while this equation is a function of freestream Mach number, M_∞, angle of attack, α, sweep angle, Λ, distance from the leading edge, x, and blunting radius, R, Guo et al. found that the upper surface pressure bleed effects were essentially independent of angle of attack, always setting $\alpha = 0^\circ$ for the upper surface pressure increment calculations.

$$\frac{\Delta P}{P_\infty} = \frac{(0.18\cos\Lambda + 0.37\sin\alpha) M_\infty^2}{2.14 + x/R} + \frac{-0.85\cos\Lambda \cdot \sin^2\alpha \cdot M_\infty^2}{\exp(0.41\sin^2\alpha \cdot x/R)} + \frac{-0.046 M_\infty^2}{\exp(1.34\cos\Lambda \cdot x/R)} \tag{2.17}$$

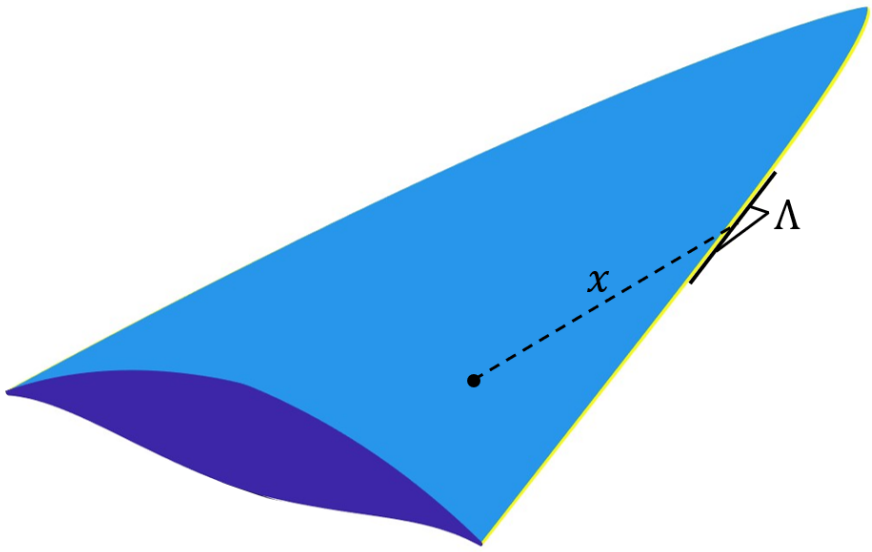

Figure 2.8: Leading edge pressure increment dimensions of interest for a sample point on the waverider surface

2.2.2 Skin Friction

An additional skin friction estimate is required to approximate the viscous drag a waverider would experience in flight. While the characteristics of the boundary layer are extremely complex and difficult to accurately predict, the large scale of vehicles tested in this study, paired with the requirement for high computational efficiency, leads to the desire for skin friction approximations that are simple to implement. For this study, the Eckert reference temperature formulation [28] is utilized, which provides a simple but relatively accurate estimate for the Reynolds number using the local surface temperature and Mach number:

$$T^* = T\left[1 + 0.032M^2 + 0.58\left(\frac{T_w}{T} - 1\right)\right] \tag{2.18}$$

where T_w is wall temperature. The reference temperature is then used to calculate reference pressure using ideal gas relations, the viscosity using Sutherland's Law [29], and finally the reference Reynolds number:

$$\begin{aligned} \rho^* &= \frac{P}{R_{air} T^*} \\ \mu^* &= \mu_{ref} \left(\frac{T^*}{T_{ref}} \right)^{3/2} \frac{T_{ref} + S}{T^* + S} \\ Re_x^* &= \frac{\rho^* u_e x_{LE}}{\mu^*} \end{aligned} \tag{2.19}$$

Where $R_{air} = 287$ J/kg·K, $\mu_{ref} = 1.716 \times 10^{-5}$ kg/m·s, $T_{ref} = 273.15$ K, and $S_{ref} = 110.4$ K. Similar to pressure, the tangent-wedge and Prandtl-Meyer expansion relations are applied to estimate the surface temperature and Mach number, as required by Eq. 2.18. For sufficiently large expansion angles, the impact of skin friction is assumed to be negligible. Thus, a rear surface approximation is not required for temperature or Mach number:

$$\frac{T}{T_\infty} = \begin{cases} \left(\frac{P}{P_\infty} \right) \frac{2 + (\gamma - 1) M^2 \sin^2 \beta}{(\gamma + 1) M^2 \sin^2 \beta}; & \theta \geq 0 \\ \left(1 + \frac{\gamma - 1}{2} M_\infty \theta \right)^2; & -\frac{2}{(\gamma - 1) M_\infty} < \theta < 0 \end{cases} \tag{2.20}$$

$$M = \begin{cases} \frac{(\gamma - 1) M_\infty^2 \sin^2 \beta + 1}{2\gamma M_\infty^2 \sin^2 \beta - \gamma + 1}; & \theta \geq 0 \\ M_\infty \left(1 + \frac{\gamma - 1}{2} M_\infty \theta \right)^{-1}; & -\frac{2}{(\gamma - 1) M_\infty} < \theta < 0 \end{cases} \tag{2.21}$$

The transition from laminar to turbulent flow is an additional characteristic of boundary layers that is extremely difficult to predict and can have a significant impact on aircraft performance. Here, a low-order approximation is used during vehicle optimization to increase

computational efficiency. The local skin friction coefficient for laminar and turbulent compressible flow is calculated using compressible flat plate relations given in Anderson [1]:

$$c_{f,lam} = \frac{0.664}{\sqrt{Re_x^*}} \tag{2.22}$$

$$c_{f,tur} = \frac{0.0592}{(Re_x^*)^{0.2}} \tag{2.23}$$

where laminar to turbulent transition is predicted using a formulation used by Bowcutt et al. for waverider analysis [10]. This method produces the expected transition effects: transition occurs at a consistent distance from the leading edge and also occurs earlier on the lower surface, where compression occurs. Fig. 2.9 shows the estimated transition location for a sample vehicle at Mach 7 and 30 km altitude.

$$\log_{10}(Re_T) = 6.421 \cdot \exp(1.209 \times 10^{-4} M_e^{2.641}) \tag{2.24}$$

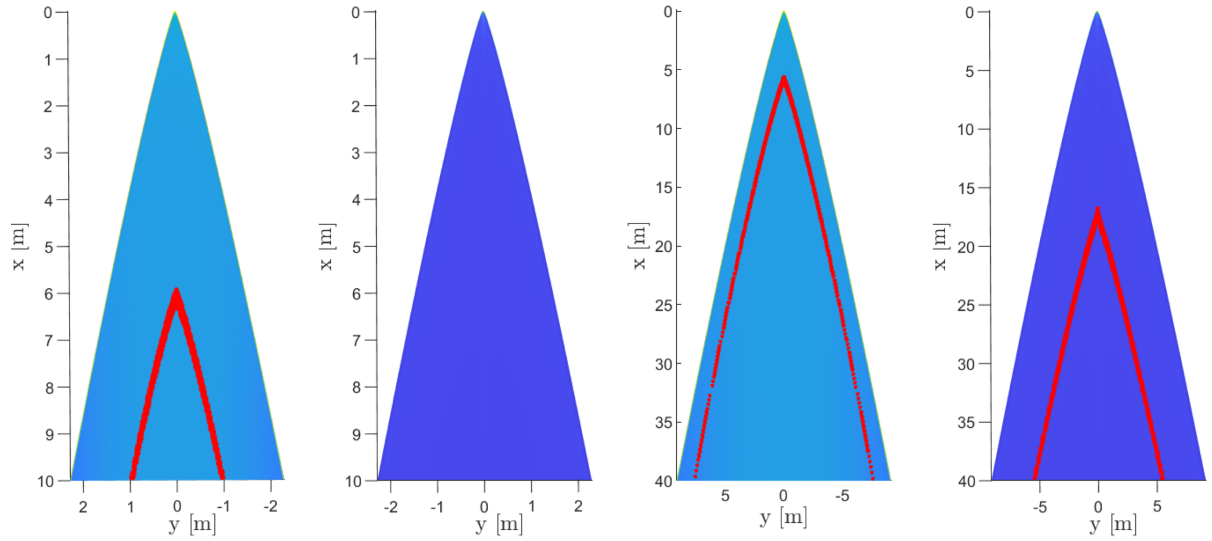

(a) Lower, upper surfaces; $L = 10$ m

(b) Lower, upper surfaces; $L = 40$ m

Figure 2.9: Laminar to turbulent transition location using Eckert reference, Bowcutt et al. transition estimate; Mach 7, 30km

2.2.3 Grid Convergence

To obtain a performance metric for any waverider geometry, the inviscid pressure and skin friction forces are discretely integrated over a vehicle mesh with N triangular cells, as shown in Fig. 2.10. In this study, lift and drag are the primary performance measures and are calculated as follows:

$$\begin{aligned} \vec{F}_{inv} &= -\sum_{i=1}^{N} p_i \vec{A}_i \\ \vec{F}_{visc} &= -\sum_{i=1}^{N} c_f A_i q_i \frac{\vec{V}_i}{||\vec{V}_i||} \\ \text{Lift} &= F_{inv,3} + F_{visc,3} \\ \text{Drag} &= F_{inv,1} + F_{visc,1} \end{aligned} \tag{2.25}$$

where A_i is the local cell area, q_i is the local dynamic pressure of the freestream, and V_i is local velocity.

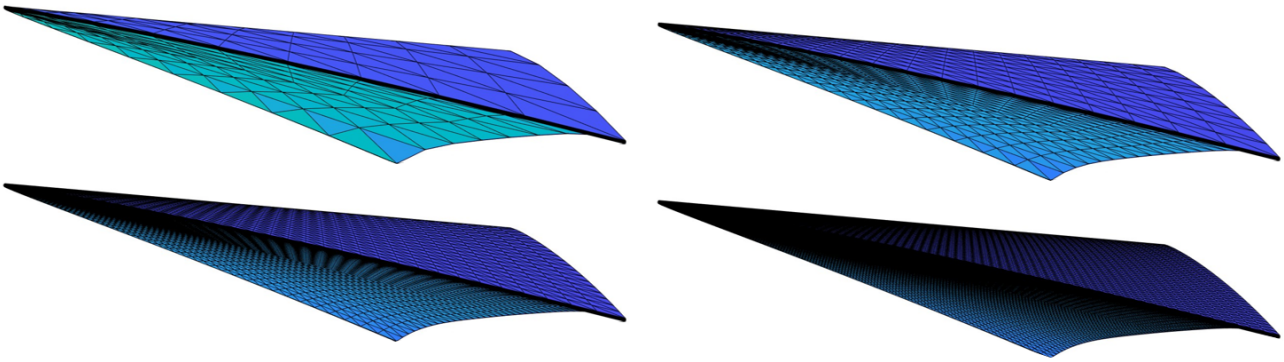

Figure 2.10: Grid refinement geometries at Mach 7 and 30km; 1132, 3032, 9232, 31232 cells

A grid convergence study of the analytical lift-to-drag ratio is performed at Mach 7 and 30 km altitude. A mesh size with an error of 0.1% or less is desired. As shown in Fig. 2.11, the optimal mesh size for this geometry and flight condition is 31,232 cells.

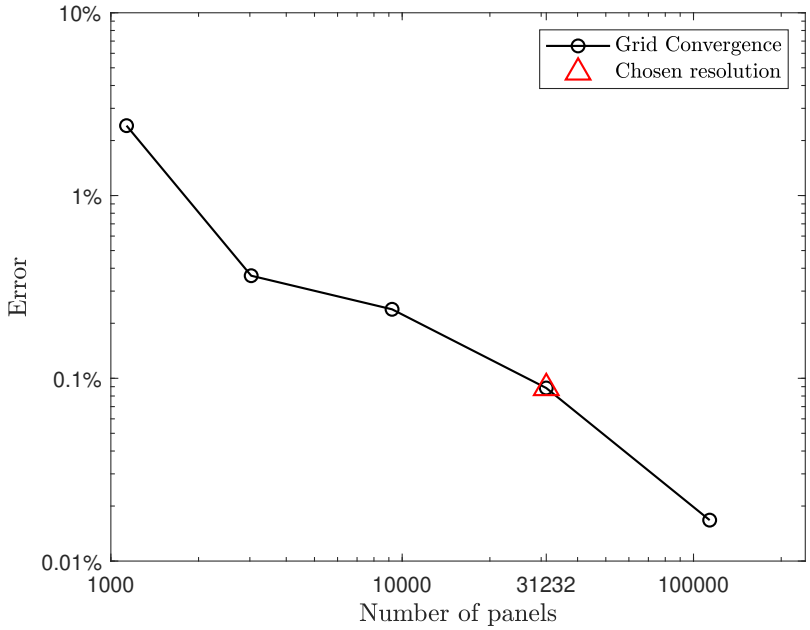

Figure 2.11: Analytical model L/D grid refinement study at Mach 7 and 30km

2.3 Computational Fluid Dynamics

While the previously described analytical model provides a good initial approximation of waverider performance, the included equations do not account for more complex flow characteristics experienced by a blunted leading edge waverider in-flight. Among these features, the most significant is pressure loss on the lower compression surface due to the normal shock standoff.

To address this issue, the GPU-based Cartesian High-order Adaptive Multi-Physics Solver (CHAMPS+) CFD tool [30] is used. This tool is based on the academic CHAMPS code [31], which has a CPU integration. For this study, CHAMPS+ is used to simulate waverider aerodynamics by solving the inviscid, compressible Navier-Stokes equations. This provides force data that more accurately approximates a waverider's actual flight performance by calculating the shock shape generated by an off-design geometry. A key feature of the CHAMPS+ code is its GPU integration, which allows simulations to run at extremely high speeds. Although the time required for simulations is still an order of magnitude longer than the analytical model—from a few milliseconds per analytical calculation to a few minutes per simulation—its highly parallelized structure offers significant speed improvements over traditional CPU-based codes, making it feasible to incorporate this higher-fidelity tool into the geometry optimization process.

2.3.1 CHAMPS+ Computational Fluid Dynamics Tool

The CHAMPS+ solver is ideal for geometry optimization because it does not require manual volume grid generation, instead using an immersed-boundary method to automatically refine a Cartesian volume grid for any arbitrary vehicle geometry. By applying grid refinement zones to the leading edge of the waverider geometries, computational grids that effectively model

the important flow features of a blunted leading edge waverider can be generated rapidly and automatically by the CHAMPS+ algorithm for each geometry, as shown in Fig. 2.12. For the cases used in the present work, the solver was configured with a 5th-order WENO scheme to allow for both shock capturing and improved accuracy. It is important to note that this method still requires the use of the Eckert reference temperature formulation (described in Section 2.2.2) to approximate the effects of skin friction on vehicle performance, and Maxwell's base pressure approximation from Eq. 2.12. In effect, the CHAMPS+ CFD tool provides more accurate estimates of inviscid surface pressure, temperature, and Mach number and replaces analytical model Eqs. 2.12 (with the exception of the base pressure model) through 2.17, 2.20, and 2.21, respectively.

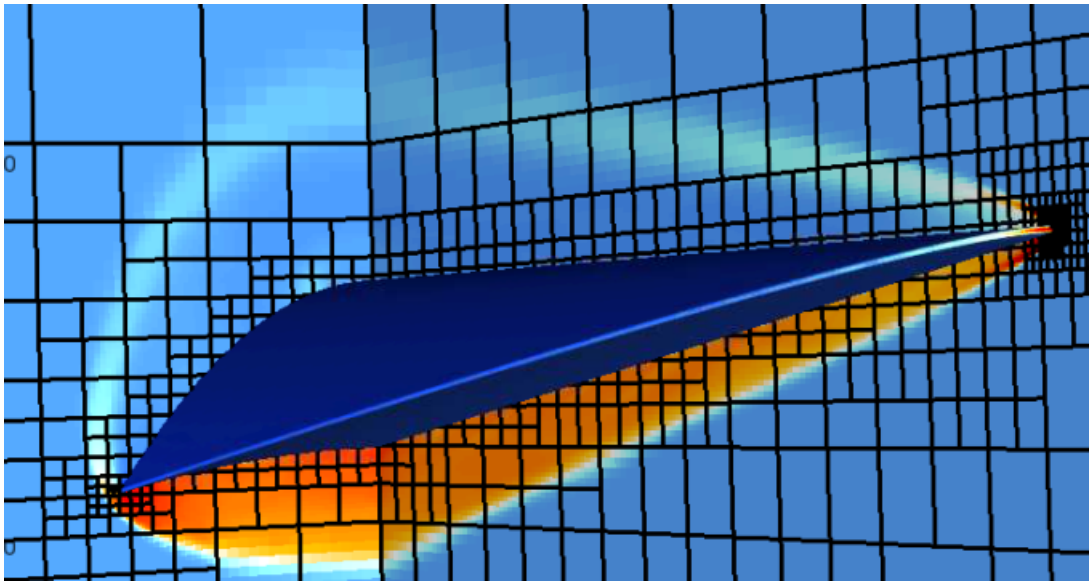

Figure 2.12: Waverider block structure with near-body and leading edge superimposed over pressure flow field; 3.2M grid cells

2.3.2 Grid and Convergence

Multiple refinement zones are used to accurately capture the important high-gradient flow regions around the waverider vehicle efficiently. Since the most important data of interest are the surface variables such as pressure, temperature, and velocity, high grid refinement is applied near the vehicle surface—especially near the leading edge—while accuracy is sacrificed in regions further away from the waverider to increase simulation speed.

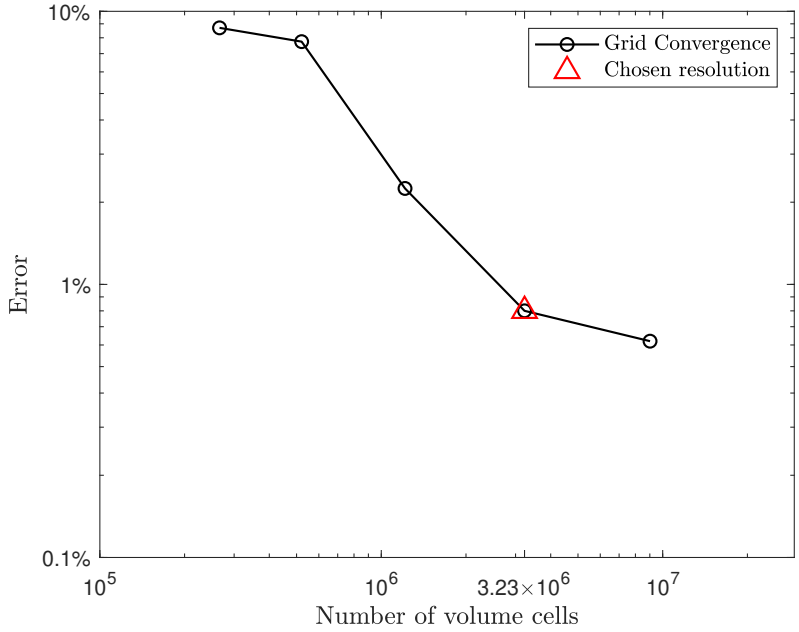

Figure 2.13: Computational fluid dynamics L/D grid refinement study at Mach 7 and 30 km

To determine the optimal volume grid resolution, a grid convergence study is conducted. To achieve a good balance between accuracy and compute time, the coarsest grid with an error of 1% or less is chosen. This results in a medium-fine grid with 3.2 million cells using the sample geometry shown in Fig. 2.10. Fig. 2.13 shows the convergence study for this resolution, using a waverider with a 40 m length and a nose radius of $R = 0.1$ m at a flight condition of Mach 7

and 30 km altitude. Fig. 2.14 shows the resulting grid refinement and flow features at the leading edge for this resolution:

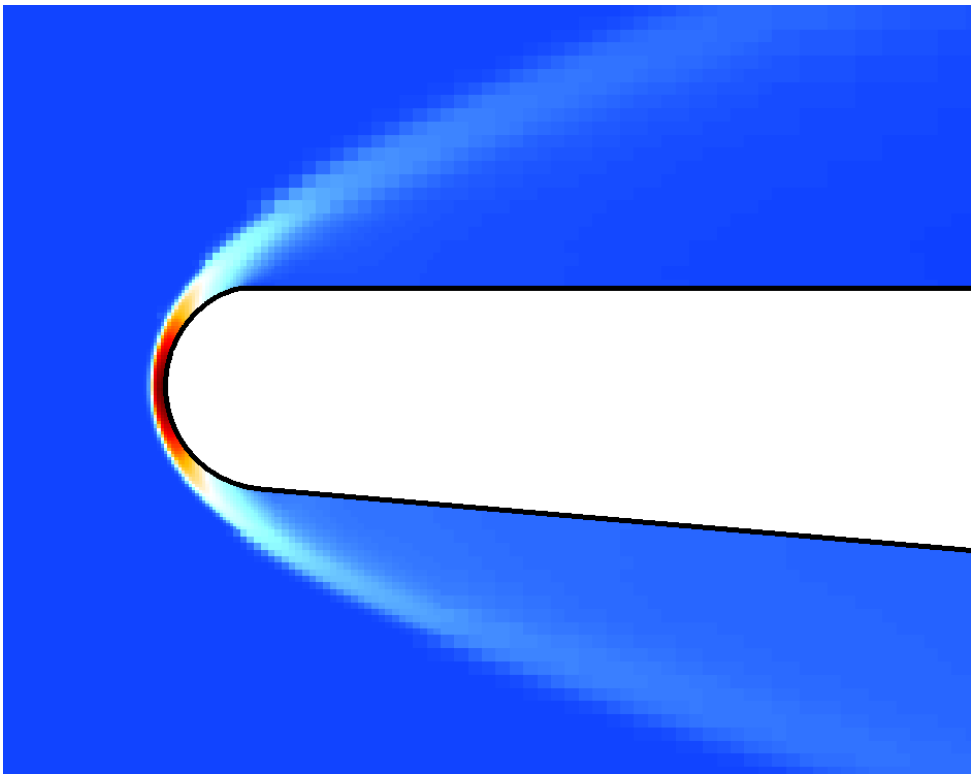

Figure 2.14: Sample waverider leading edge grid refinement and pressure flow field in x-z plane, $R = 0.1$ m

2.3.3 Comparison to Analytical Model

Many optimization runs are performed using both the CFD and analytical models for a wide variety of waverider geometries, with different lower surfaces, leading edge planforms, and blunting radii. To compare the differences between these methods, two sample geometries are selected: one predicted to have high performance based on the analytical model, and the other which shows high performance according to the CFD results. These geometries, labeled '1' and '2' respectively, are shown below in Fig. 2.15:

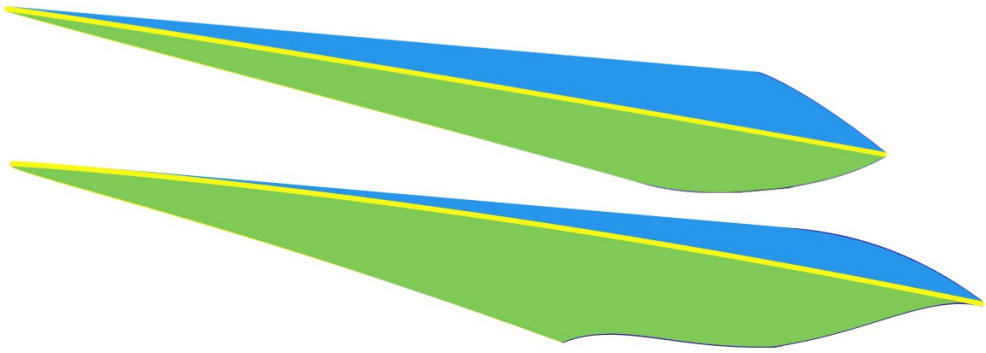

Figure 2.15: Waveriders used for model comparison; analytical-optimized Geometry 1 (top) and CFD-optimized Geometry 2 (bottom)

An angle of attack sweep from $-6^\circ \leq \alpha \leq 6^\circ$ is performed at Mach 8 and 20 km altitude for bluntness ratios varying from sharp ($R/L = 0\%$) to significantly rounded ($R/L = 1\%$). As shown in Fig. 2.17, the analytical model accurately predicts the overall trends for both geometries across varying angles of attack and bluntness ratios. Furthermore, while it closely matches the CFD lift-to-drag predictions for Geometry 2, it noticeably over-predicts the magnitude for Geometry 1. This discrepancy is likely due to lateral pressure bleed, caused by excessive local wedge angle or thickness, which is not accounted for in the analytical model. As a result, waverider geometries that would experience a high degree of lateral pressure bleed in real flight have over-predicted performance by the analytical model. This phenomenon is shown in greater detail in Fig. 2.16, which further shows that Geometry 1 experiences significantly more pressure bleed than Geometry 2.

Additionally, Fig. 2.18 compares the surface pressure distribution for both geometries calculated using the analytical model and CFD. It can be seen clearly that the pressure on the leading edge of geometry 1 was uniformly overestimated, likely due to the fact that significant pressure loss around the leading edge is not modeled in these equations. Geometry 2, on the other

hand, underestimates pressure in locations where high pressure occurred in the CFD, notably the wingtips, and overestimated pressure where it was lower in the CFD. This indicates again that even when the analytical model accurately predicts lift-to-drag ratio performance for round-leading edge geometries, this is probably due to the analytical model closely matching the average pressure distribution across the waverider surface, and not accurately predicting the actual effects of leading edge rounding on the flow. This trend is in addition to a tendency to overestimate pressure on the blunt leading edge where sweep angle is low, such as near the nosetip, and underestimate it where sweep is high, which again is likely due to complex shock shape effects which are not accounted for in the analytical equations.

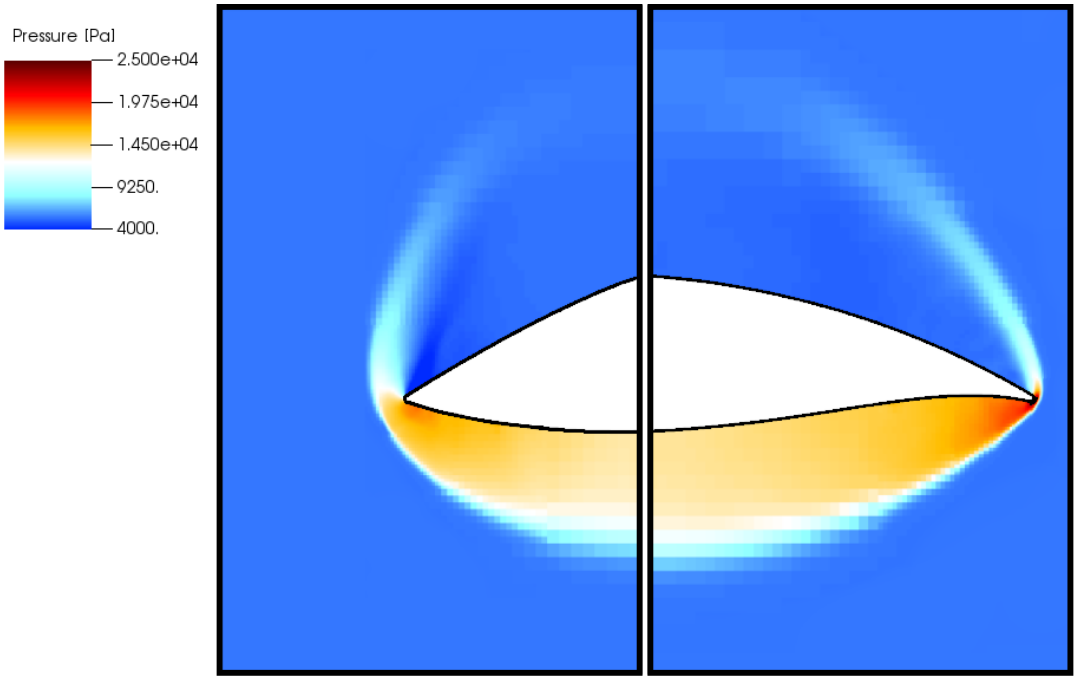

Figure 2.16: Computational fluid dynamics pressure bleed comparison; geometry 1 & 2 at Mach 8, 20 km altitude

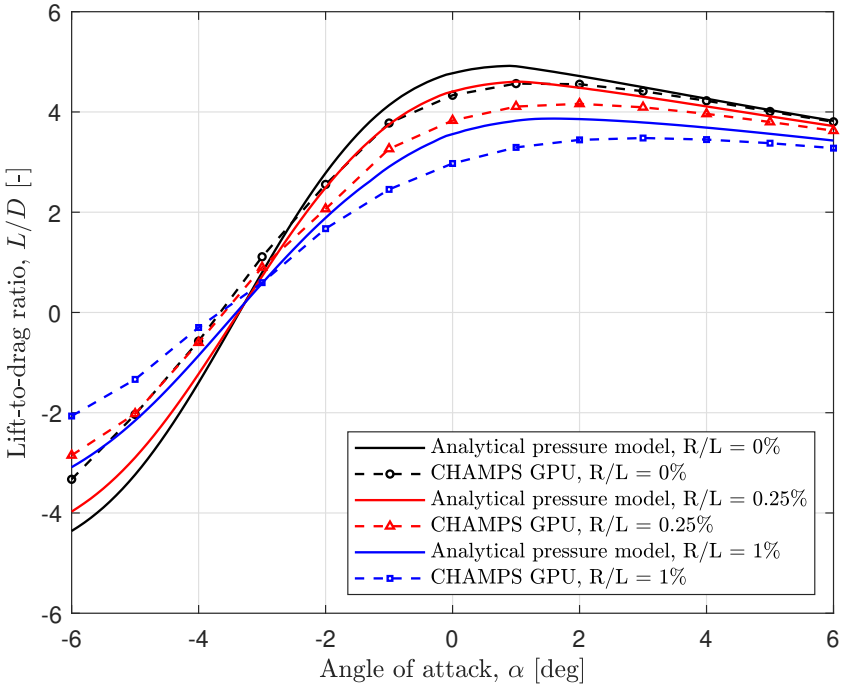

(a) Geometry 1

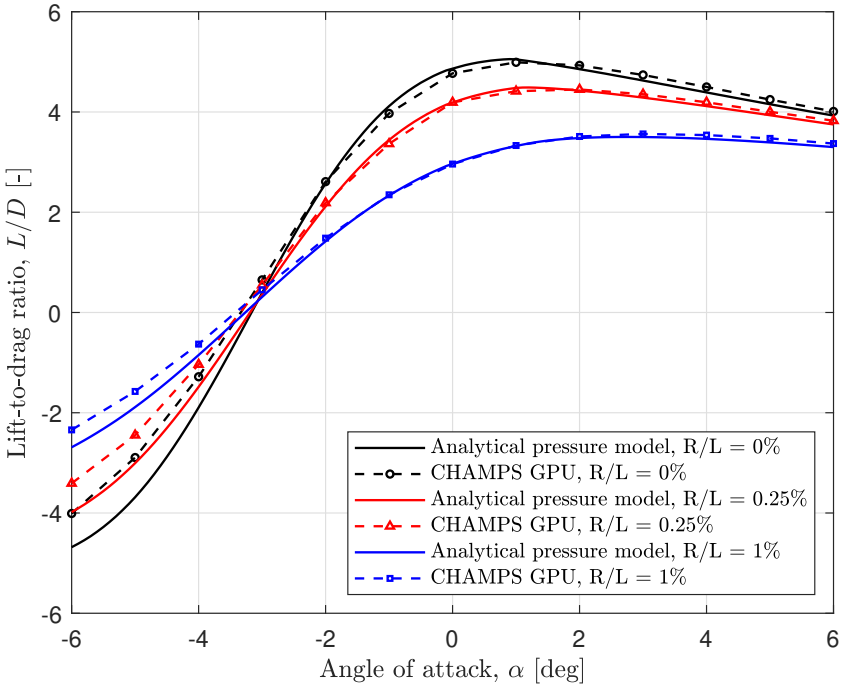

(b) Geometry 2

Figure 2.17: Lift-to-drag ratio versus angle of attack for analytical model and computational fluid dynamics; Mach 8, 20 km

These results seem to suggest that the pressure increment model from Guo et al. (Eq. 2.17) was created using the results from waveriders with low amounts of pressure bleed at the leading edge. While the symbolic regression model and overall pressure approximation technique used predict the performance of round-leading-edge waveriders with remarkable accuracy when the amount of pressure bleed is small, the decrease in performance for highly off-design waveriders highlights the weaknesses of analytical models at predicting performance in increasingly complex flow fields. This again seems to indicate that without the inclusion of higher-fidelity methods in the optimization process, analytical optimization techniques could converge to designs that take advantage of weak points in an analytical model. For this reason, two distinct morphing vehicle optimizations will be conducted and subsequently compared in this study, one solely using the analytical model, and one using CFD.

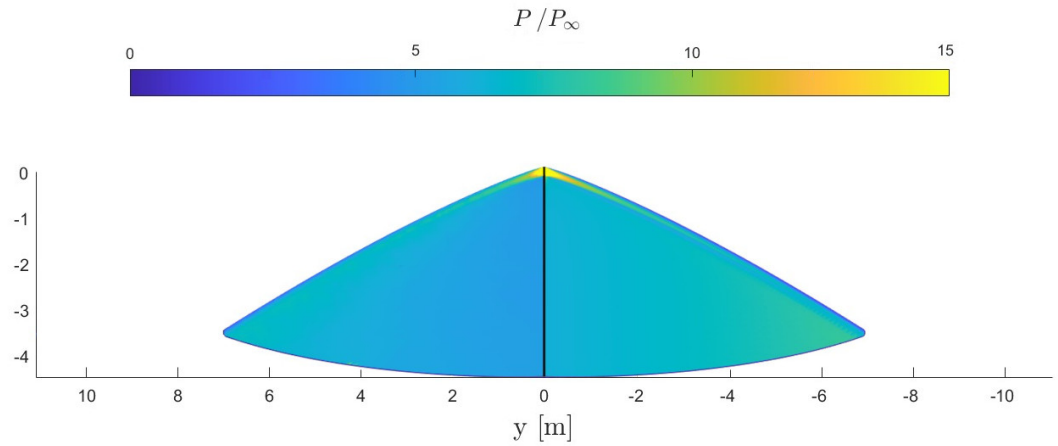

(a) Geometry 1 pressure distribution; left: CFD, right: analytical model

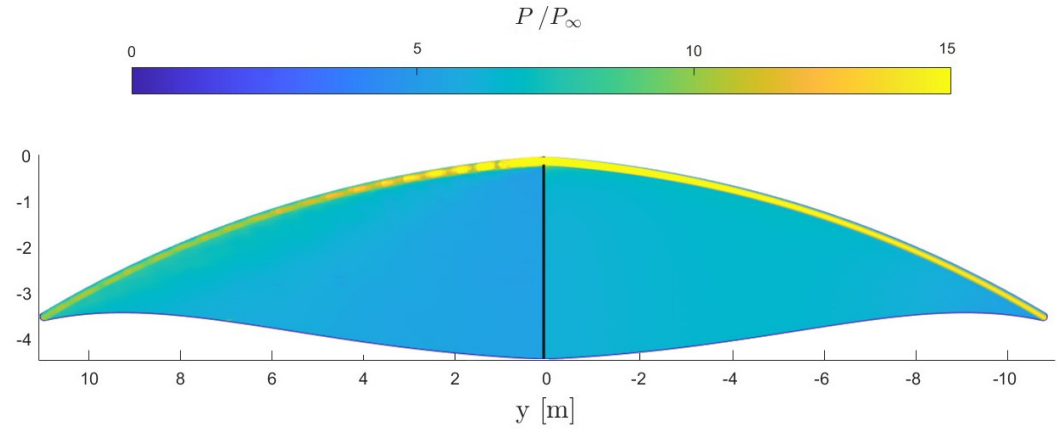

(b) Geometry 2 pressure distribution; left: CFD, right: analytical model

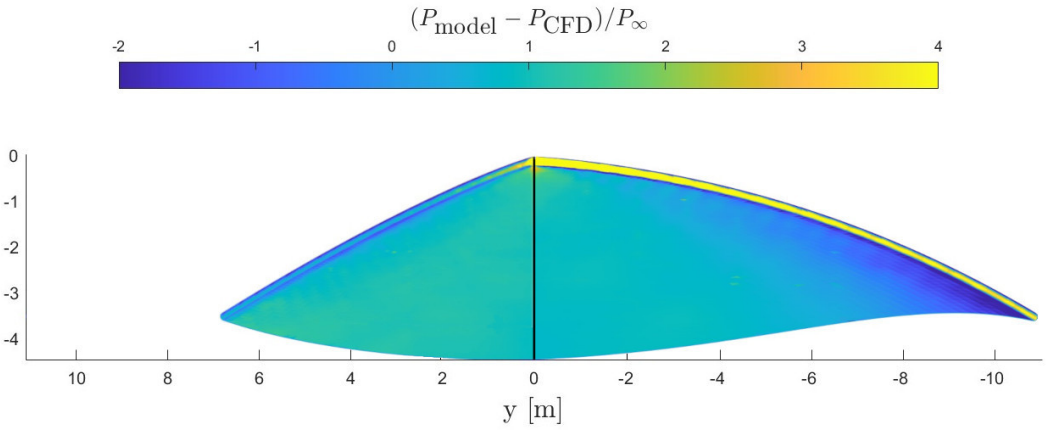

(c) Differences in analytical model and CFD pressure; left: geometry 1, right: geometry 2

Figure 2.18: Surface pressure distribution for CFD versus analytical model; Mach 8, 20 km altitude

2.4 Geometry Optimization

In this study, the goal is to design a morphing waverider that can adjust its geometry to maximize its lift-to-drag ratio throughout a mission glide trajectory. To identify the optimal waverider geometry for a given flight condition, a pattern search optimization algorithm [32] is employed using MATLAB's built-in optimization toolbox. Pattern search is particularly useful due to its derivative-free nature, allowing it to efficiently navigate the complex and often discontinuous design space of the waverider geometry formulation used in this study.

2.4.1 Geometric Constraints

For the geometry optimization process, it is necessary to provide realistic upper and lower bounds on the design variables that defined the geometry function $f(L, w, n, \beta, \varepsilon, p_1, p_2, p_3, R)$ to ensure that the optimization problem is well-defined for the search algorithm. For all optimization runs, only geometries of constant length and blunting radius are compared. Additional constraints are applied based on their mathematical definition, such as for n and ε, or simply to prevent the construction of obviously suboptimal waveriders—particularly geometries that are extremely wide or thick. The variable bounds applied during this optimization stage can be seen below:

$$0.1L \leq w \leq 1.5L \qquad 0.1 \leq n \leq 0.9 \qquad 1^\circ \leq \beta \leq 20^\circ \qquad -1 \leq \varepsilon \leq 1$$

$$0.5 \leq p_1 \leq 3 \qquad 0.5 \leq p_2 \leq 3 \qquad 0.5 \leq p_3 \leq 3$$

To assist in validating the algorithm and impose real-world limitations on waverider geom-

etry, a restriction on volumetric efficiency, η_{vol} is enforced:

$$\eta_{\text{vol}} = \frac{V^{2/3}}{S} \tag{2.26}$$

Where V is the vehicle's volume, and S is its planform area. Vehicles with high volumetric efficiency tend to be wider and thicker, offering more usable volume, while the opposite is generally true for vehicles with low volumetric efficiency.

2.4.2 Optimization and Validation

Since both the analytical model and the CFD algorithm require inputs for Mach number, altitude, and angle of attack, the optimization design vector consists of these values with the waverider geometric parameters. For each optimization sweep, Mach number, altitude, and angle of attack are fixed, leading to the design vector:

$$\vec{X} = \{L, w, n, \beta, \varepsilon, p_1, p_2, p_3, R, M, h, \alpha\} \tag{2.27}$$

As discussed earlier, the lift-to-drag ratio is the primary function of interest for this study, yielding the objective function:

$$J = -L/D(\vec{X}) \tag{2.28}$$

For geometry optimization, the pattern search algorithm is used. This non-gradient-based technique evaluates the objective function in discrete steps, selecting the step with the minimum function value iteratively [32]. This algorithm is chosen because of its gradient-free implementation and its suitability for engineering problems with complex, discontinuous objective

functions that may have multiple local minima [33]. While the lift-to-drag ratio objective function (Eq. 2.28) is generally convex, discontinuities arise by applying geometric constraints, such as excluding geometries with intersecting upper and lower surfaces. These issues, combined with local minima that sometimes occur at extreme flight conditions make pattern search a more appropriate choice than other methods, such as fmincon. Additionally, the relatively expensive CFD function evaluations make true global optimization methods, like genetic algorithms, impractical due to the high number of function evaluations they require.

To validate this scheme, multiple convergence tests are conducted using the analytical model's viscous lift-to-drag ratio calculations. A discrete sweep of a limited design space is performed to identify the true minimum, and the results are compared with the pattern search converged solution. Fig. 2.19 shows the results of a study with both one and two geometric control points. Only vehicle width is varied for the one-dimensional case, and width and design shock angle are varied for the two-dimensional case. As shown, pattern search performs well in converging to the true minimum in both cases.

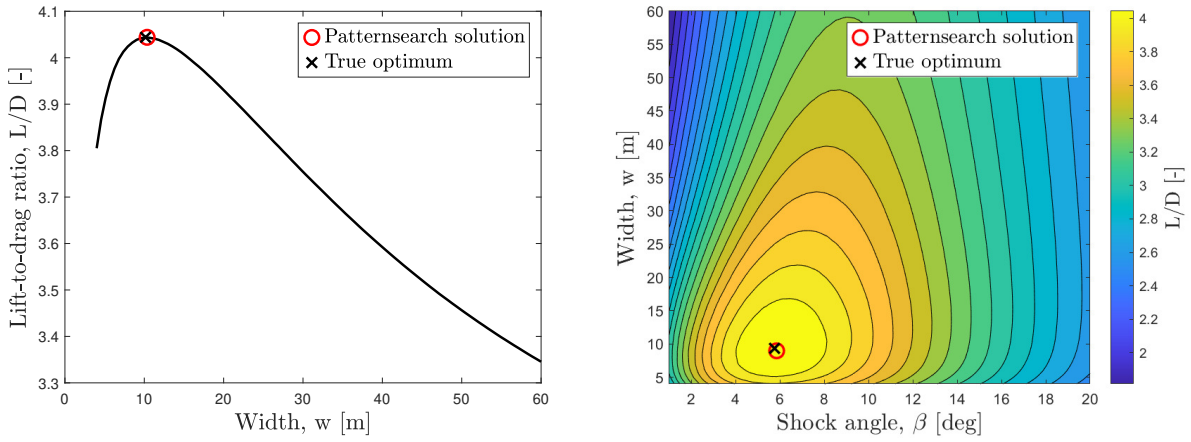

(a) 1-D (b) 2-D

Figure 2.19: Pattern search lift-to-drag ratio optimization validation

The final validation study is conducted with all geometric variables bounded as described in Section 2.4.1, using a waverider with $L = 40$ m and $R = 0.1$ m at Mach 7.5 and 30 km altitude. Here, waverider geometry is optimized to create a Pareto front of the design space, i.e., to maximize the lift-to-drag ratio across a range of volumetric efficiencies, $0 \leq \eta_{\text{vol}} \leq 1$. These results are then compared to those from a brute-force attempt, where one million geometries are simulated across the given geometric range. As shown in Fig. 2.20, the pattern-search-converged Pareto front not only visually matches the Pareto front of viscous performance versus volumetric efficiency found in other studies [7, 34], but also converges to geometries with performance equal to or better than those found using the brute-force sweep. This indicates again that the pattern search algorithm is well-suited for optimizing the chosen waverider geometry design space.

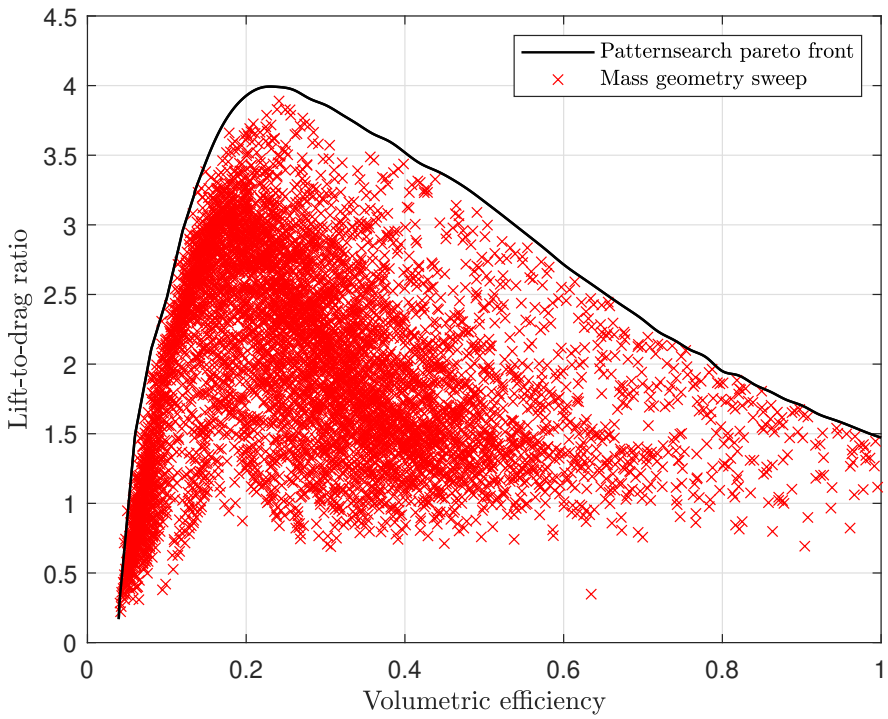

Figure 2.20: Pattern search optimized Pareto front of lift-to-drag ratio and volumetric efficiency using the analytical model; Mach 7.5, 30 km

2.5 Trajectory Analysis

While lift-to-drag ratio is the primary measure of waverider performance in this study, a trajectory optimization tool is also used to visualize the general effects morphing may have on a waverider's glide path and range. To identify waverider trajectories, the direct collocation functionality of the Imperial College London Optimal Control Software (ICLOCS2) is used [35]. This tool discretizes the waverider trajectory into a system of discrete points, or collocation nodes. By applying the waverider dynamics and constraints across these nodes, an optimization problem can be solved that minimizes the desired objective while satisfying the physical dynamics of waverider flight. This algorithm is particularly effective for the waverider gliding trajectory, where many constraints are involved, and the flight paths are typically complex.

The full waverider flight profile is broken up into three distinct trajectory optimization stages. The most important of these is the glide phase, where range is maximized for fixed initial and final flight conditions by controlling angle of attack. For visualization purposes, a boost phase, in which a rocket launch delivers the waverider to its starting flight condition of the glide with an additional thrust control, and a land phase, in which the waverider approaches a specified terminal condition, are simulated. It is important to emphasize that the launch and landing phases are not considered in the performance analysis process.

2.5.1 Equations of Motion

Waverider dynamics are defined using a simplified three-degree-of-freedom system of aircraft equations of motion, where angle of attack control affects waverider pitch and motion in the x- and z-directions. It is assumed that changes in the waverider's angle of attack attitude are

instantaneous as controlled, and no control surface movement or associated aerodynamic effects are simulated. The chosen three-degree-of-freedom formulation has been used in similar gliding aircraft trajectory optimization problems, including space shuttle reentry [23] and rocket-assisted gliding waverider trajectories [24, 25], and accounts for Earth's curvature and its associated centrifugal effects. Vehicle motion is described using the following states:

$$\begin{aligned}
\dot{h} &= v\sin\gamma \\
\dot{\phi} &= \frac{v}{h+R_e}\cos\gamma \\
\dot{v} &= \frac{T-D(h,v,\alpha)}{m} - g\sin\gamma \\
\dot{\gamma} &= \frac{L(h,v,\alpha)}{mv} + \cos\gamma\left(\frac{v}{h+R_e} - \frac{g}{v}\right)
\end{aligned} \tag{2.29}$$

where h is altitude above ground level, ϕ is longitude, v is velocity, γ is flight path angle, and R_e is Earth's radius. The primary forces considered during the gliding trajectory are lift, L, and drag, D, which are stored as discrete aerotables as a function of altitude, velocity, and angle of attack, α. Also considered are gravity, g, calculated using Earth's gravitational constant $\mu = 3.986 \times 10^{14}$ $\mathrm{m}^3\ \mathrm{s}^{-2}$, and thrust, T, for the rocket launch trajectories.

$$g = \frac{\mu}{(h+R_e)^2} \tag{2.30}$$

Additional constraints are included, most notably stagnation point temperature, T_w, and for the rocket boost phase, fuel mass, m_f. Stagnation point temperature is iterated using explicit Euler time integration and discrete aerotables based on Fay-Riddell heating relations. This data is defined as a function of altitude, velocity, and leading edge radius, and is calculated using Eqs.

2.33 through 2.35. Fuel mass flux is calculated using engine thrust and specific impulse.

$$\dot{T}_w = \frac{T_w(h, v, r) - T_w(h + \dot{h}\Delta t, v + \dot{v}\Delta t, r)}{\Delta t} \tag{2.31}$$

$$\dot{m}_f = -\frac{T}{I_{sp} g} \tag{2.32}$$

2.5.2 Flight Conditions and Constraints

Stagnation point heat flux is calculated at all points throughout the trajectory using a simplified Fay-Riddell approximation for axisymmetric flows [1]:

$$q_w = 0.763 Pr^{-0.6} \sqrt{\rho_e \mu_e \frac{dU_e}{dx}} (h_{aw} - h_w) \tag{2.33}$$

where Pr is the Prandtl number, h_{aw} is the adiabatic wall enthalpy, h_w is the wall enthalpy, and $\mathrm{d}U_e/\mathrm{d}x$ is the stagnation point velocity gradient. It is assumed that the adiabatic wall enthalpy is equivalent to the stagnation enthalpy, $h_{aw} = h_0$, and $\mathrm{d}U_e/\mathrm{d}x$ is calculated using the Newtonian approximation and assuming a small $\mathrm{d}x$ near the stagnation point:

$$\frac{\mathrm{d}U_e}{\mathrm{d}x} = \frac{U_\infty}{R} \sqrt{\frac{2\rho_\infty}{\rho_e}} \tag{2.34}$$

Finally, stagnation point temperature is calculated assuming radiative and convective equilibrium and iterating stagnation heating and temperature until the Stefan-Boltzmann law is satisfied:

$$q_w = \varepsilon \sigma T_w^4 \tag{2.35}$$

where σ is the Stefan-Boltzmann constant and ε is the waverider material emissivity at the leading edge, for this study assumed to be $\varepsilon = 0.9$. For all phases of flight, a conservative stagnation point temperature limit of 2700 K is used to represent a ceramic material with modest heat resistance.

2.5.2.1 Boost Phase

The boost phase of the trajectory is initiated at ground level and completed once all usable fuel in the rocket has been expended and a specified altitude, in this case 35 km altitude and Mach 10, or 3086.5 m/s, is reached. For a 40 m length waverider, rocket mass is based loosely on a design sized in-between existing launch vehicles with open-source performance values, with a total mass of 850,000 kg including 800,000 kg of usable fuel. The full boost phase constraints and boundary conditions are as follows:

$$
\begin{array}{ccc}
h_{0,b} = 0 \text{ km} & h_{0,b} \leq h_b \leq h_{f,b} & h_{f,b} = 35 \text{ km}, \\
\phi_{0,b} = 0^\circ & \phi_b \geq \phi_{0,b} \, , & \\
v_{0,b} = 0 \text{ m/s} & v_{0,b} \leq v_b \leq v_{f,b} & v_{f,b} = 3086.5 \text{ m/s}, \\
\gamma_{0,b} = 90^\circ & \gamma_{f,b} \leq \gamma_b \leq \gamma_{0,b} & \gamma_{f,b} = 0^\circ \, , \\
T_{w,0,b} = 288 \text{ K} & T_{w,b} \leq 2700 \text{ K}, & \\
m_{0,b} = 850,000 \text{ kg} + m_{WR} & m_{f,b} \leq m_b \leq m_{0,b} & m_{f,b} = 50,000 \text{ kg} + m_{\text{WR}}
\end{array}
$$

2.5.2.2 Glide Phase

The glide phase begins at specified release conditions and ends at a chosen glide ceiling, designed to test the optimizer's ability to avoid no-fly conditions and maintain high-performance flight. For the majority of simulated trajectories, a glide ceiling of 15 km altitude and 1000 m/s

is selected.

$$h_{0,g} = 35 \text{ km} \qquad h_{0,g} \leq h_g \leq 100,000 \text{ km} \qquad h_{f,g} = 15 \text{ km},$$

$$\phi_{0,g} = \phi_{f,b} \qquad \phi_g \geq \phi_{0,g} \,,$$

$$v_{0,g} = 3086.5 \text{ m/s} \qquad 1000 \text{ m/s} \leq v_g \leq 10,000 \text{ m/s},$$

$$\gamma_{0,g} = 0^\circ \qquad -89^\circ \leq \gamma_g \leq 89^\circ,$$

$$T_{w,0,g} = T_{w,f,b} \qquad T_{w,g} \leq 2700 \text{ K}$$

2.5.2.3 Land Phase

Finally, constraints are applied to the land phase to ensure the vehicle completes the mission at a slight nose-down attitude and a low-supersonic velocity. This is done because all analytical and GPU simulations are performed at supersonic speeds, and it is assumed that the lifting-body characteristics of the waverider vehicle would require a landing approach similar to the space shuttle. The full constraints are:

$$h_{0,l} = 15 \text{ km} \qquad h_{f,l} \leq h_l \leq h_{0,l} \qquad h_{f,l} = 0 \text{ km},$$

$$\phi_{0,l} = \phi_{f,g} \qquad \phi_l \geq \phi_{0,l} \,,$$

$$v_{0,l} = v_{f,g} \qquad 0 \text{ m/s} \leq v_l \leq 10,000 \text{ m/s} \qquad 300 \text{ m/s} \leq v_{f,l} \leq 350 \text{ m/s},$$

$$\gamma_{0,l} = \gamma_{f,g} \qquad -89^\circ \leq \gamma_l \leq 89^\circ \qquad -5^\circ \leq \gamma_{f,l} \leq 0^\circ,$$

$$T_{w,0,l} = T_{w,f,g} \qquad T_{w,l} \leq 2700 \text{ K}$$

2.5.3 Optimization

The objective of the waverider trajectory optimization is to maximize the range of the glide phase of flight. This means the objective function is simply described by the initial and final glide latitudes:

$$J_g = -(\theta_{f,g} - \theta_{0,g}) \tag{2.36}$$

In order to converge to a realistic solution for the boost and land phases of flight, the optimization problem is formulated to minimize flight time:

$$J_b = t_{f,b} \tag{2.37}$$

$$J_l = t_{f,l} - t_{0,l} \tag{2.38}$$

This set of optimization problems tends to work extremely well with the ICLOCS2 algorithm, which quickly yields full trajectories that closely resemble boost-glide vehicle flight dynamics. Fig. 2.21 shows a sample trajectory predicted for a general low-performance, on-design caret waverider:

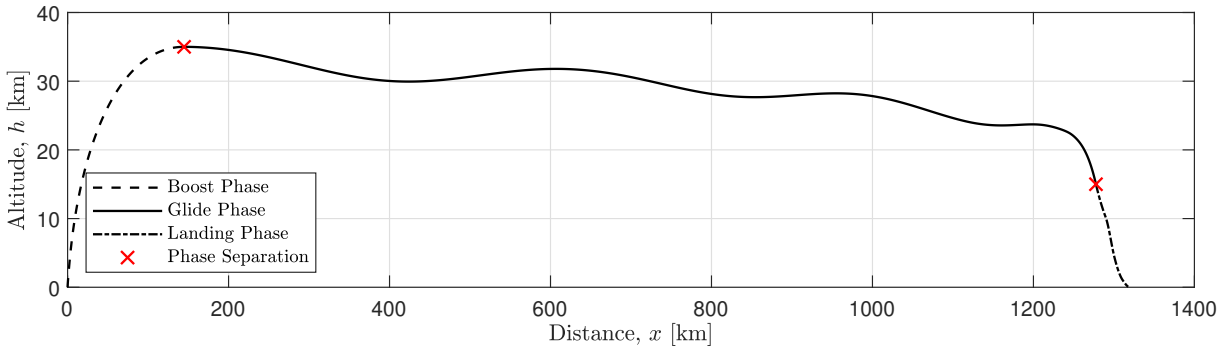

Figure 2.21: Sample waverider trajectory with boost, glide, and land phases; analytical pressure estimates ($L = 40$m, $w = 16.7$m, $n = 0.86$, $\beta = 13.2°$, $\varepsilon = 0$, $p_1 = 0.8$ $p_2 = 0.8$, $p_3 = 0.8$)

Chapter 3: Problem Formation Using Analytical Model

In this study, a 'fixed' waverider is a vehicle that maintains a constant geometry throughout a boost-glide mission, similar to conventional waveriders. In contrast, 'morphing' waveriders incorporate an additional control mechanism which applies geometric changes to the entire lower surface—or sometimes even the leading edge—at various points during the trajectory to increase mission range. Fig 3.1 shows an example of a morphing waverider geometry at three different lower surface configurations and a fixed leading edge.

As stated earlier, the goal of this study is to identify a methodology to select a family of waverider geometries that a vehicle can morph between during a trajectory, and to identify the potential range impacts of this hypothetical control surface using analytical and CFD methods. This chapter details the design process used to establish a baseline vehicle and its trajectory, select flight conditions along the trajectory where morphing would occur, and optimize at these conditions to create a set of geometries the morphing vehicle could switch between throughout the trajectory to constantly maximize lift-to-drag ratio. Finally, the performance and range impacts of morphing are evaluated. The lower-order analytical model is used initially, followed by higher fidelity optimization using CFD in Chapter 4.

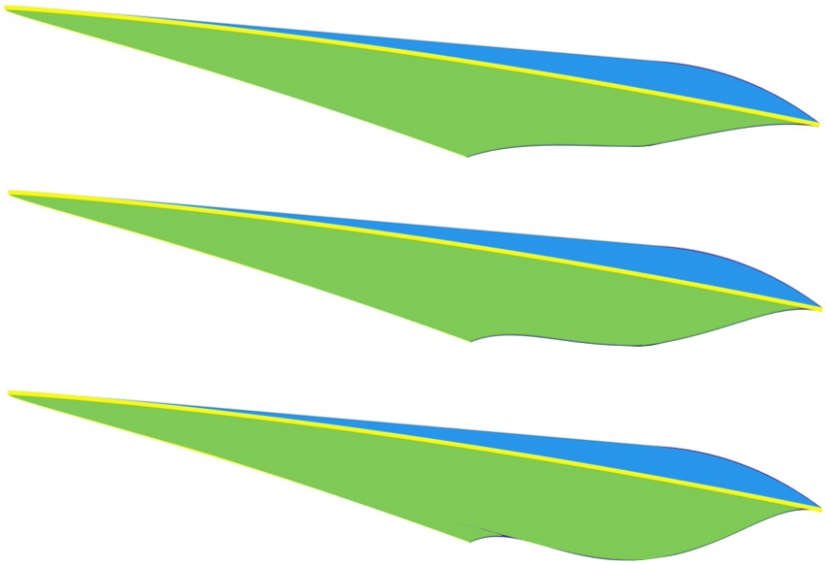

Figure 3.1: Morphing waverider with three lower surface configurations and a fixed leading edge

3.1 Fixed Waverider Design

For this study, an initial fixed waverider design is optimized and used as a baseline for comparing morphing waverider performance. This geometry represents a traditional reentry glide vehicle, designed with a fixed fuselage and traditional control surfaces. Generating the glide trajectory for this baseline vehicle provides multiple flight conditions for which the morphing vehicles will be optimized. Theoretically, these multiple optimization steps could improve the average lift-to-drag ratio across the trajectory and improve mission range. This section details the selection of an initial fixed waverider geometry and the resulting trajectory.

3.1.1 Geometry optimization

The fixed waverider geometry is optimized for maximum lift-to-drag ratio per Eq. 2.28. Optimization is performed using geometric constraints $L = 40$ m and $R = 0.1$ m, representing

a roughly space shuttle-sized transport vehicle with a rounding ratio $R/L = 0.25\%$, flying at the midpoint between the initial and terminal flight conditions of the glide phase: Mach 6.85, $h = 25$ km altitude, and $\alpha = 0°$. The resulting optimal waverider achieves a lift-to-drag ratio of 3.97 at this condition. Its geometry is shown in Fig. 3.2 and is defined by variables:

$$w = 12.81 \text{ m}, \;\; n = 0.90, \;\; \beta = 5.42°, \;\; \varepsilon = -0.35, \;\; p_1 = 1.47, \;\; p_2 = 1.54, \;\; p_3 = 1.57$$

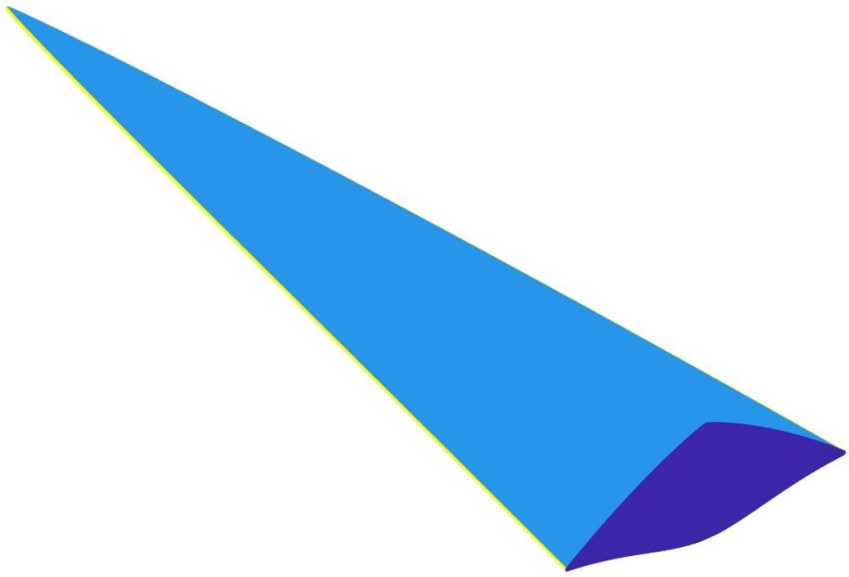

Figure 3.2: Analytical-model-optimized fixed waverider; $L = 40$ m, $R = 0.1$ m, $M = 6.85$, $h = 25$ km, and $\alpha = 0°$

To generate the glide trajectory for this geometry, aerotables of lift and drag forces are created, with values calculated across Mach, altitude, and angle of attack. The lift-to-drag ratio data is displayed for the fixed waverider across Mach number and altitude at $\alpha = 0°$ in Fig. 3.3a, and across Mach number and angle of attack at $h = 25$ km in Fig. 3.3b. As can be seen, lift-to-drag ratio tends to increase as Mach number increases. As altitude increases, lift-to-drag ratio increases as the flight condition where optimization occurred approaches, then decreases as

a result of the decreasing air density. Furthermore, peak performance occurs at a positive angle of attack at each flight condition. These characteristics are typical of optimized waveriders, showing that the analytical model provides a good approximation of the overall performance trends that align well with literature results.

However, a slight discontinuity in slope can be observed near zero angle of attack for the low Mach flight conditions in 3.3b. This likely arises from the pressure approximation in Eq. 2.12, which is also discontinuous in slope across flow angle $\theta = 0°$ and highlights additional weaknesses of this model. Still, in most cases, this discontinuity is minimal and has a negligible impact on optimization.

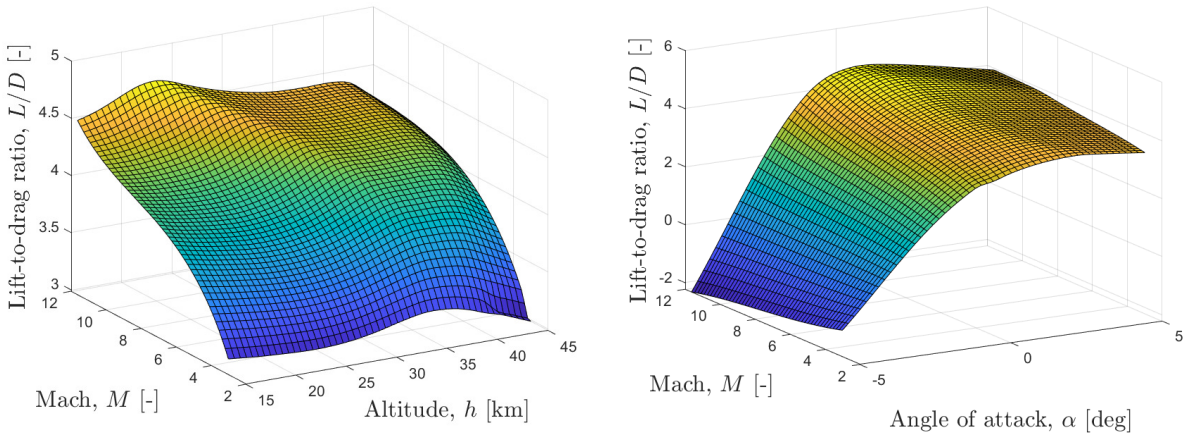

(a) Fixed at $\alpha = 0°$

(b) Fixed at $h = 25$ km

Figure 3.3: Fixed waverider lift-to-drag ratio performance across altitude, Mach number, and angle of attack

3.1.2 Fixed Vehicle Trajectory using Analytical Model

This section discusses the trajectory optimization results for the initial fixed waverider, optimized using the analytical model. Time histories of system dynamics, stagnation temperature, flight distance, angle of attack control, and lift-to-drag ratio performance are shown in Fig. 3.4.

Fig. 3.4a shows the altitude of the fixed waverider over time. As shown, the glide lasts 16 minutes and 52 seconds, with the vehicle exhibiting an oscillatory motion throughout. This motion closely resembles the phugoid dynamic stability mode [36] which often results from the trajectory optimization of waverider vehicles [24, 25]. On one hand it could be that this oscillatory motion results in part from the initial condition selection. Since the fixed vehicle was not optimized to be in a trimmed condition at Mach 10 and 35 km altitude, it initially pitches down, losing altitude and increasing velocity. As this happens, the thicker air and faster speed increase vehicle lift, which causes a pitching-up motion. Phugoid motion, characterized by oscillations resulting from flying outside of a trimmed condition, explains some of the vehicle's motion. However, since this is an optimized trajectory, it follows that this oscillatory motion is not solely due to the vehicle being outside of its trim condition, but could also be a preferable to using control to dampen out this motion.

This behavior is also observable, though less distinctly, in the Mach number time history shown in Fig. 3.4b. The Mach number increases in regions corresponding to decreases in altitude, and vice versa, again typical of phugoid motion. However, the Mach number also shows a nearly linear decrease throughout the trajectory. The oscillatory behavior is further illustrated in Fig. 3.4c, which depicts the trajectory flight path angle. Near the end of the trajectory, the vehicle sharply pitches down, likely to avoid the imposed lower velocity limit. It is also important to note that the observed trajectory results in an estimated leading edge stagnation temperature well below the set limit of 2700 K, as depicted in Fig. 3.4d. Finally, the range of the fixed waverider is 2087.2 km, as seen in Fig. 3.4e.

Fig. 3.4f shows the angle of attack flown compared to the angle of attack for maximum lift-to-drag ratio throughout the trajectory. As can be seen, the vehicle attempts to maintain

maximum lift-to-drag ratio throughout nearly the entire glide. This again likely indicates that it is more important for the optimal trajectory to control to maintain this high performance attitude rather than to dampen dynamic stability effects. This behavior continues until the vehicle pitches over during the terminal phase of the glide. As a result, the time-averaged lift-to-drag ratio is 3.94, shown in 3.4g.

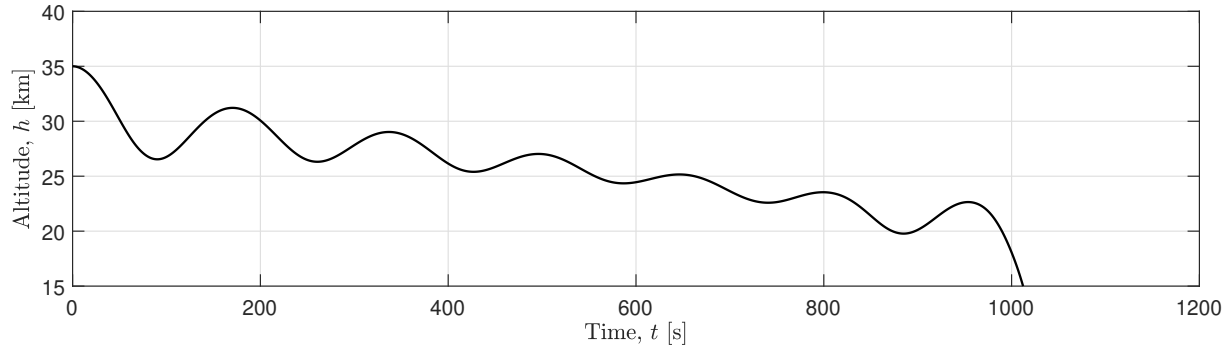

(a) Altitude time history

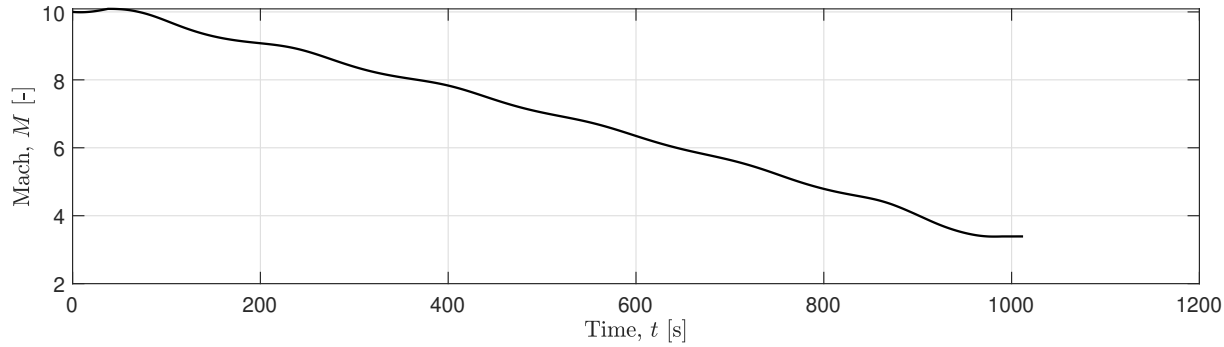

(b) Mach number time history

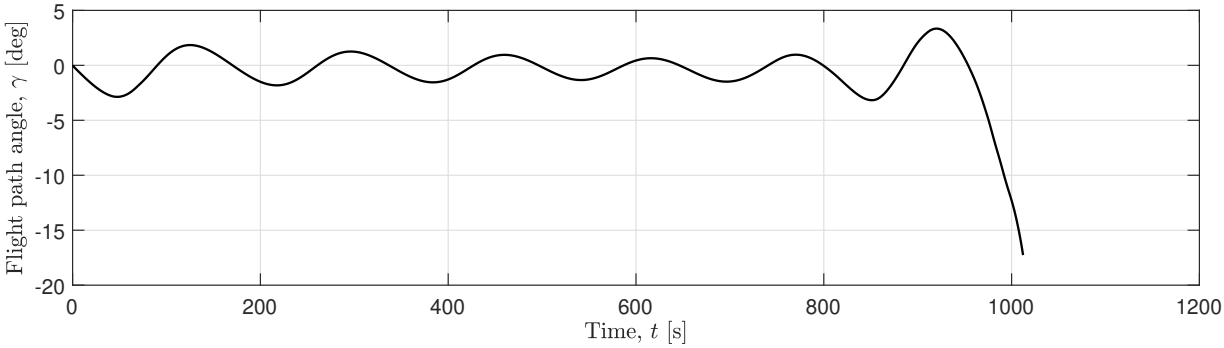

(c) Flight path angle time history

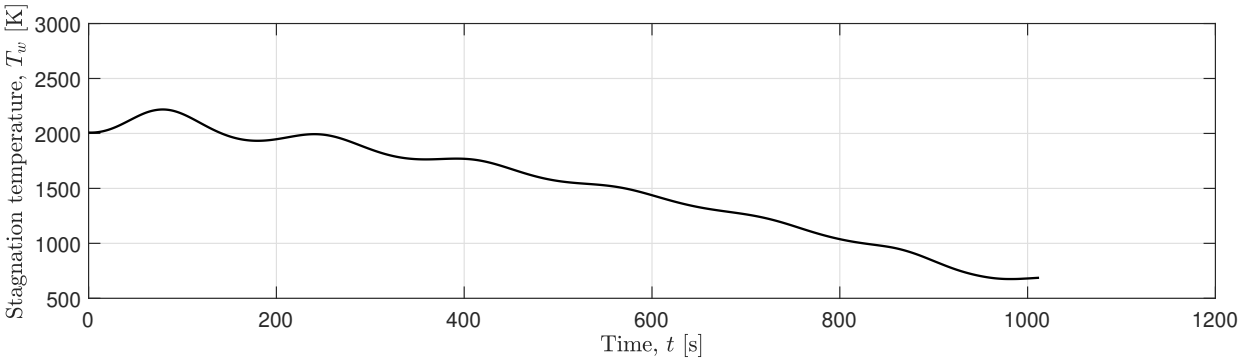

(d) Stagnation temperature time history

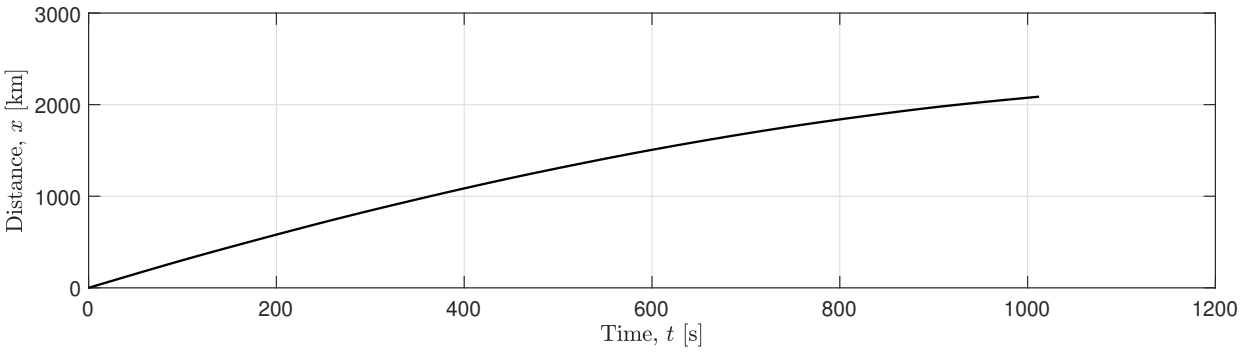

(e) Distance time history

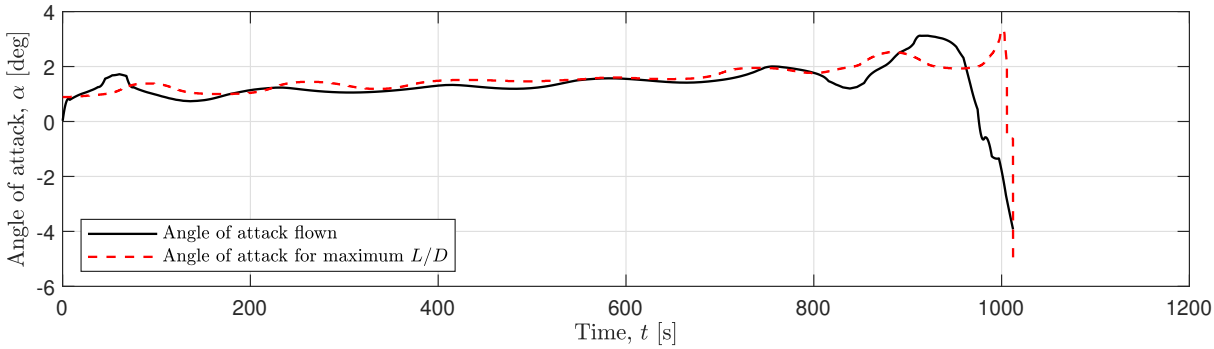

(f) Angle of attack time history

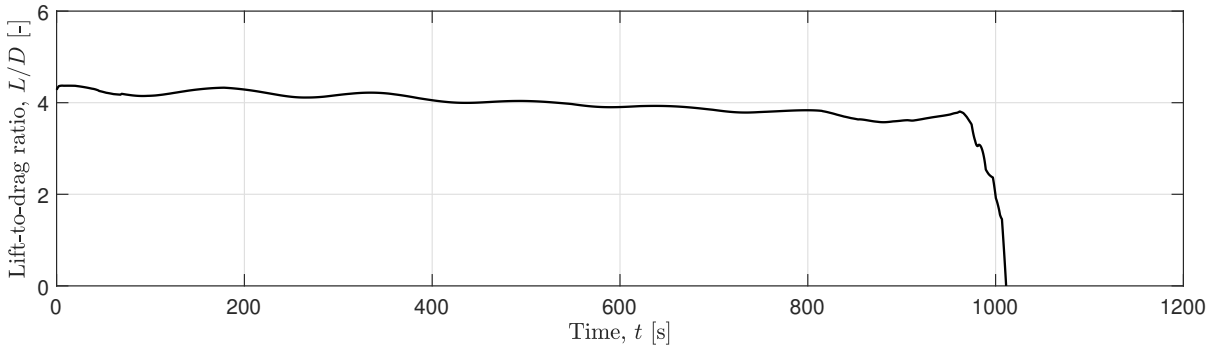

(g) Lift-to-drag ratio time history

Figure 3.4: Fixed waverider time histories; Mach 10, 35 km release attitude, analytical pressure estimates

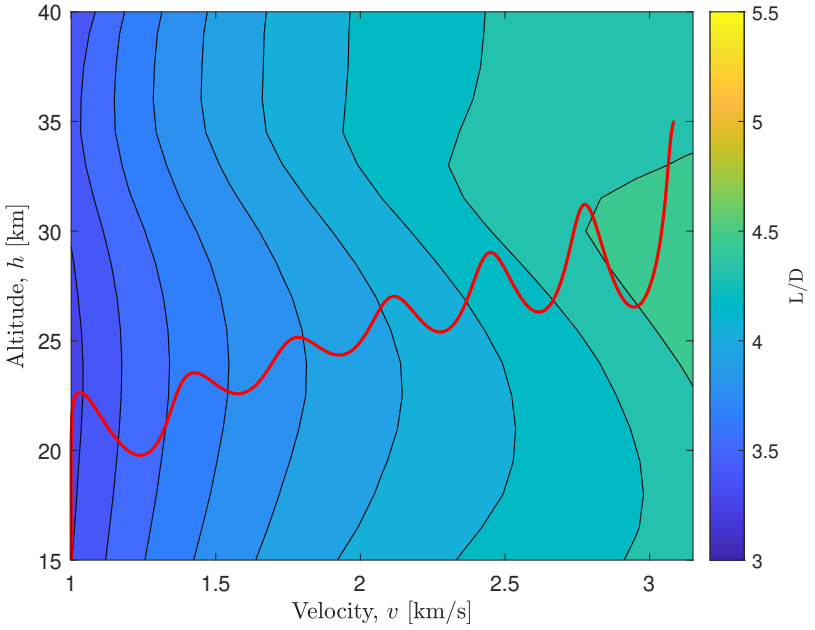

Figure 3.5: Fixed waverider L/D aerotables and trajectory; $\alpha = 1°$; analytical pressure estimates ($L = 40$ m, $w = 12.81$ m, $n = 0.90$, $\beta = 5.42°$, $\varepsilon = -0.35$, $p_1 = 1.47$ $p_2 = 1.54$, $p_3 = 1.57$)

Fig. 3.5 shows the trajectory of the fixed waverider overlaid with its lift-to-drag ratio aerotables at a $1°$ angle of attack. At nearly every velocity point in the trajectory, the lift-to-drag ratio increases with altitude. This suggests that the oscillatory behavior of the waverider's glide path may contribute to the increased range by capitalizing on higher lift-to-drag performance at the higher altitudes.

3.2 Morphing Waverider Design

To define a set of geometries for a morphing waverider, four additional flight conditions are selected from the initial trajectory and a new geometry is optimized at each. These conditions are selected at equidistance velocity points along the trajectory both up and down the trajectory from the initial optimization point. Fig. 3.6 displays the lift-to-drag aerotables and fixed trajectory,

highlighting the new optimization points. Table 3.1 shows these conditions in greater detail.

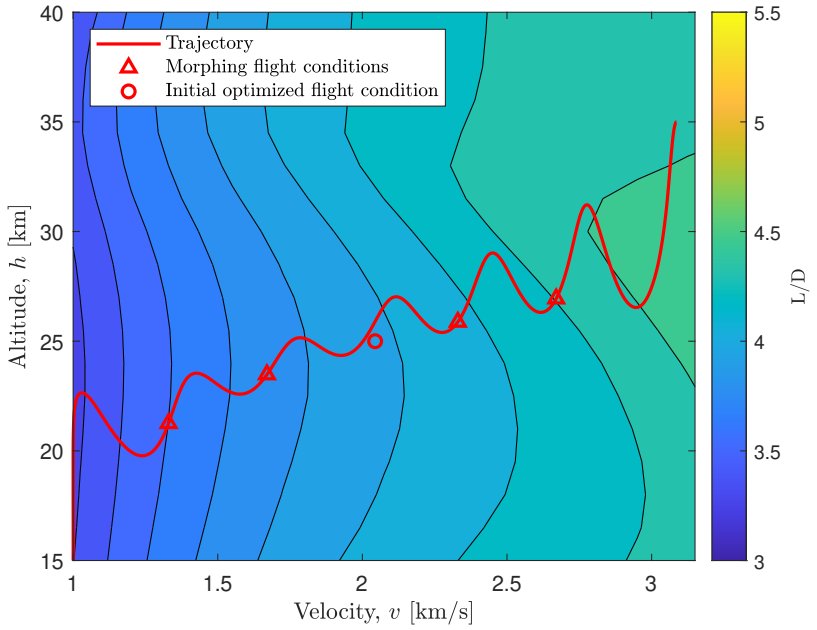

Figure 3.6: Fixed waverider L/D aerotables and trajectory with morphing points; $\alpha = 1°$; analytical pressure estimates ($L = 40$ m, $w = 12.81$ m, $n = 0.90$, $\beta = 5.42°$, $\varepsilon = -0.35$, $p_1 = 1.47$ $p_2 = 1.54$, $p_3 = 1.57$)

Table 3.1: Analytical waverider optimization flight conditions

Flight Condition	M [-]	h [km]	α [deg]
Fixed	6.85	25.00	0.00
1	8.91	26.76	1.24
2	7.80	25.70	1.51
3	5.62	23.12	1.54
4	4.49	21.29	1.44

Three methodologies for morphing geometry optimization are explored. The first keeps the leading edge fixed, only varying parameters p_1, p_2, and p_3. The second allows limited flexibility in the leading edge, where parameters w, n, β, and ε can be varied from the initial shape by up to 20% of the full range shown in 2.4.1, while p_1, p_2, and p_3 are adjusted as normal. The third

methodology, representing a 'perfect' waverider, allows all parameters to vary freely to maximize performance at every flight condition, providing the theoretical upper bound of waverider performance achievable using morphing.

3.2.1 Morphing with Fixed leading edge

The converged fixed leading edge geometries for each flight condition are shown in Table 3.2 and Fig. 3.7. These additional four geometries all have decreased Chebyshev control point values, leading to thinner wedge angles compared to the initial fixed geometry. This is likely due to the increased angle of attack at these flight conditions. Moreover, these points further decrease with Mach number, which could be attributed to the modified Newtonian assumption at the leading edge, resulting in reduced drag as Mach decreases and, consequently, less lower surface lift required for maximum lift-to-drag ratio.

Table 3.2: Analytical waverider fixed leading edge morphing geometric parameters

Flight Condition	w [m]	n [-]	β [deg]	ε [-]	p_1 [-]	p_2 [-]	p_3 [-]
Fixed	12.81	0.90	5.42	-0.35	1.47	1.54	1.57
1	12.81	0.90	5.42	-0.35	1.04	1.16	1.20
2	12.81	0.90	5.42	-0.35	1.09	1.16	1.12
3	12.81	0.90	5.42	-0.35	0.96	0.97	0.96
4	12.81	0.90	5.42	-0.35	0.88	0.88	0.82

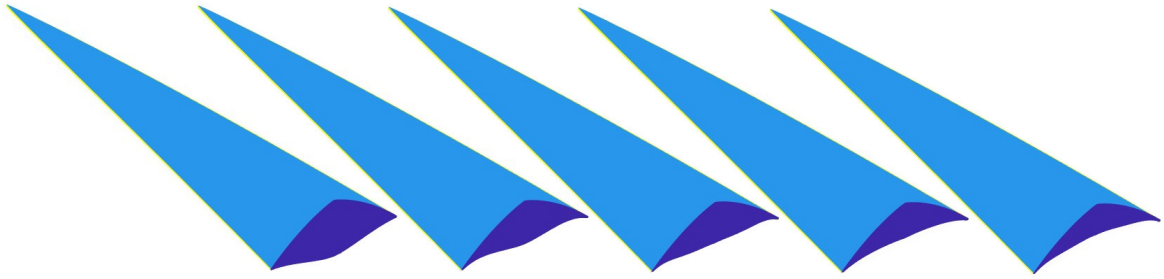

Figure 3.7: Analytical waverider fixed leading edge morphing geometries; left to right: fixed geometry, morphing geometries 1-4

3.2.2 Morphing with flexible leading edge

The flexible leading edge morphing geometries for each flight condition are shown in Table 3.3 and Fig. 3.8. Similar to the fixed leading edge geometries, the Chebyshev control points decrease with Mach number throughout the trajectory. While no distinct trend is observed in width, w, and curve factor, ε, shock angle, β decreases slightly throughout the trajectory. This again suggests that thinner vehicles are more desirable at lower Mach numbers. A planform factor, n, at the upper limit of 0.9 is optimal for all vehicles, which indicates that the waverider glide profile favors sharper planform shapes, likely because these geometries reduce leading edge pressure drag and especially span-wise viscous effects.

Table 3.3: Analytical waverider flexible leading edge morphing geometric parameters

Flight Condition	w [m]	n [-]	β [deg]	ε [-]	p_1 [-]	p_2 [-]	p_3 [-]
Fixed	12.81	0.90	5.42	-0.35	1.47	1.54	1.57
1	12.55	0.90	5.02	-0.42	1.12	1.23	1.20
2	12.70	0.90	5.01	-0.51	1.11	1.16	1.14
3	12.87	0.90	5.00	-0.47	0.98	0.98	0.96
4	11.56	0.90	5.00	-0.43	0.92	0.91	0.88

Figure 3.8: Analytical waverider constrained leading edge morphing geometries; left to right: fixed geometry, morphing geometries 1-4

3.2.3 Perfect Morphing

The unconstrained morphing geometries for each flight condition are shown in Table 3.4 and Fig. 3.9. As with the previous results, thinner vehicles with smaller profiles are preferred as Mach number decreases. However, this trend is primarily driven by a decrease in width and a decrease in shock angle, rather than by changes in the Chebyshev control points. This is likely due to the increased flexibility in the leading edge parameters for thin vehicles and small β, resulting from the oblique shock formulation in Eq. 2.5 and the thickness-multiplying effects of the Chebyshev points. Similar to the flexible leading edge morphing, no distinct trend is observed in leading edge curve factor, and planform n-factor is constant at 0.9.

Table 3.4: Analytical waverider perfect morphing geometric parameters

Flight Condition	w [m]	n [-]	β [deg]	ε [-]	p_1 [-]	p_2 [-]	p_3 [-]
Fixed	12.81	0.90	5.42	-0.35	1.47	1.54	1.57
1	12.66	0.90	4.74	-0.28	1.17	1.30	1.30
2	12.56	0.90	4.51	-0.40	1.17	1.23	1.20
3	10.59	0.90	3.16	-0.06	1.39	1.35	1.33
4	8.56	0.90	2.55	-0.09	1.42	1.41	1.43

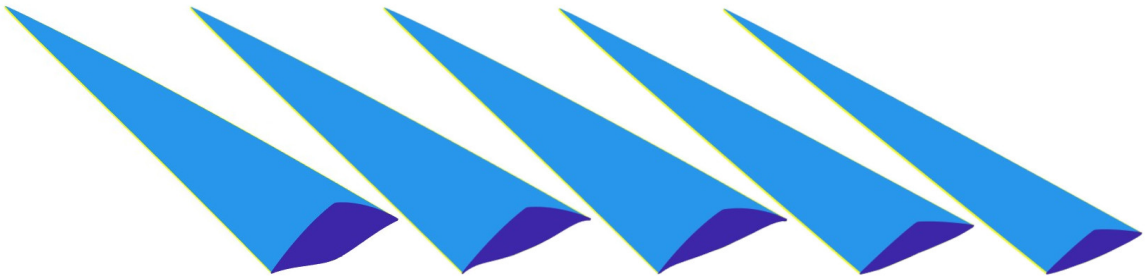

Figure 3.9: Analytical waverider perfect morphing geometries; left to right: fixed geometry, morphing geometries 1-4

Fig. 3.10 shows the front-view profiles of all three morphing waverider geometry configurations. Of particular interest are the 'perfect' waverider geometries, shown in Fig. 3.10c. These geometries have more compact profiles, suggesting that a more optimal set of morphing geometries with leading edge constraints could be created if the initial fixed geometry has a profile somewhere between that of the perfect morphing geometries. This could likely be achieved by re-optimizing the geometry at a slightly positive angle of attack, similar to the approach taken for the subsequent morphing geometries. Since this method was not initially applied, however, The constrained leading edge geometries in 3.10a and 3.10b differ noticeably from their theoretical best-performing counterparts.

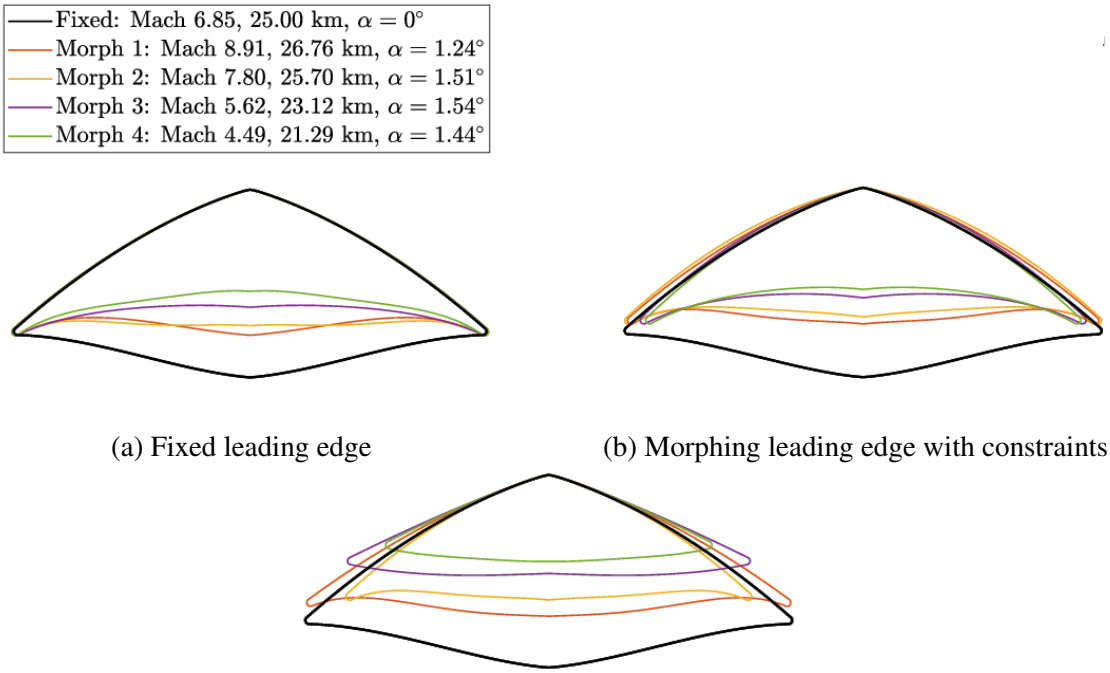

(a) Fixed leading edge

(b) Morphing leading edge with constraints

(c) Perfect morphing

Figure 3.10: Analytical morphing waverider geometries; front profile comparison

3.3 Analytical Model Results

Now that these sets of morphing geometries are optimized, new aerotables are created for each to be used during trajectory optimization. At every flight condition in the lift-to-drag aerotable, the values from the configuration with the highest lift-to-drag ratio are used. This is shown for the fixed-leading edge morphing vehicle at $\alpha = 1°$ in Fig. 3.11. This figure is extremely important, as it reveals many valuable insights about the chosen morphing geometry selection process. First, it appears that optimizing geometries at various points throughout the initial trajectory had the intended effect: as a trajectory progresses, the morphing waverider geometries are used near the conditions they were optimized for. That is, with the exception of the

initial fixed geometry configuration. As shown, this configuration is always outperformed by the morphing geometries, which suggests that a better method could have been used in selecting this geometry. It is likely that if this geometry were re-optimized at a slightly positive angle of attack, the overall performance of the resulting morphing waverider would be improved. Additionally, it appears that a geometry optimized for low-Mach number, low-altitude flight is the best performing at the initial phase of flight. This shows that the inclusion of additional optimization flight conditions could further improve morphing waverider performance.

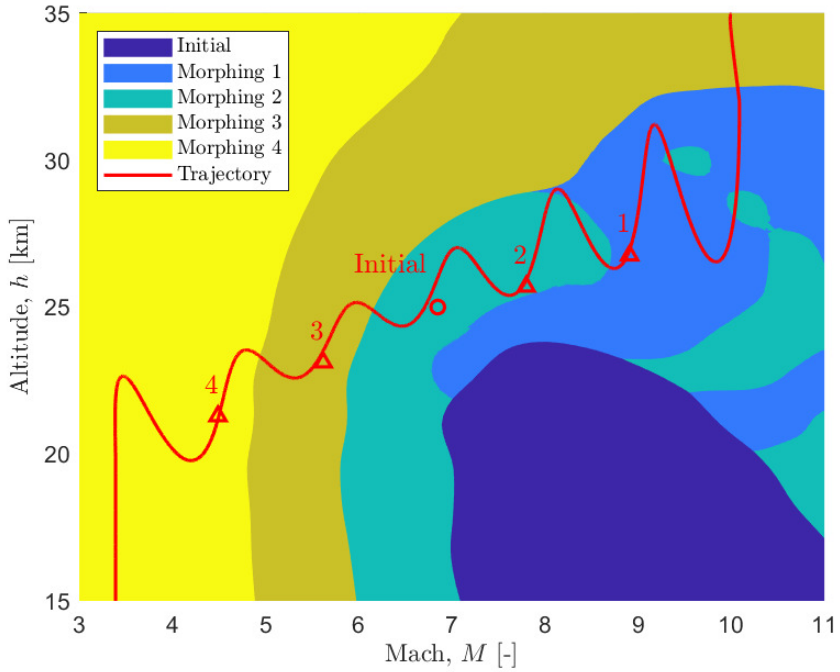

Figure 3.11: Highest-performing fixed-leading edge morphing waverider configuration across flight conditions at $\alpha = 1°$; initial fixed waverider trajectory superimposed

3.3.1 Aerotable Comparison

Fig. 3.12 compares the lift-to-drag ratio of the initial fixed geometry with the three morphing geometries, assuming each vehicle followed the trajectory of the initial fixed geometry. Under

this assumption, only a small increase in lift-to-drag ratio is observed using morphing, with the fixed and flexible leading edge vehicles showing nearly identical increases in performance. The perfect waverider provides the best performance throughout the trajectory, particularly during the initial oscillations and the final dive. The overall lift-to-drag ratio data for each vehicle is presented in Table 3.5.

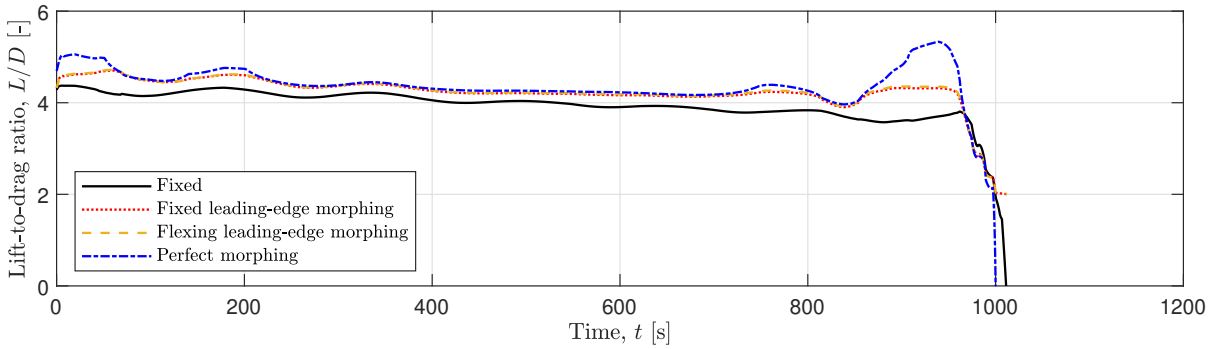

Figure 3.12: Fixed versus morphing waverider L/D performance over time assuming constant trajectory; Mach 10, 35 km release attitude, analytical pressure estimates

Table 3.5: Analytical waverider optimization flight conditions

Vehicle type	L/D_{avg} [-]	L/D_{max} [-]	L/D_{min} [-]
Fixed	3.94	4.37	-0.34
Fixed leading edge morphing	4.23	4.70	2.00
Flexible leading edge morphing	4.24	4.72	2.00
Perfect morphing	4.27	5.33	-5.71

While the constant-trajectory assumption provides useful insights into how morphing improved lift-to-drag ratio at various flight conditions along the initial trajectory, it is insufficient for approximating the resulting increase in range. This is because vehicles with different lift-to-drag ratio data will inevitably travel at different conditions across their trajectory. Therefore, additional trajectory optimization is performed for all three morphing geometries, as shown in

Fig. 3.13. The results reveal that the fixed and flexible leading edge morphing waveriders show very similar increases in lift-to-drag ratio throughout the altitude and velocity range, especially as altitude increases. These vehicles also exhibit altitude and velocity oscillations with noticeably lower amplitude. In contrast, the perfect morphing vehicle sees significantly larger increases in lift-to-drag ratio. Consequently, the optimal trajectory has much larger amplitude oscillations, likely due to both the increased lift throughout the trajectory, and also because improved performance could always be reached by increasing altitude.

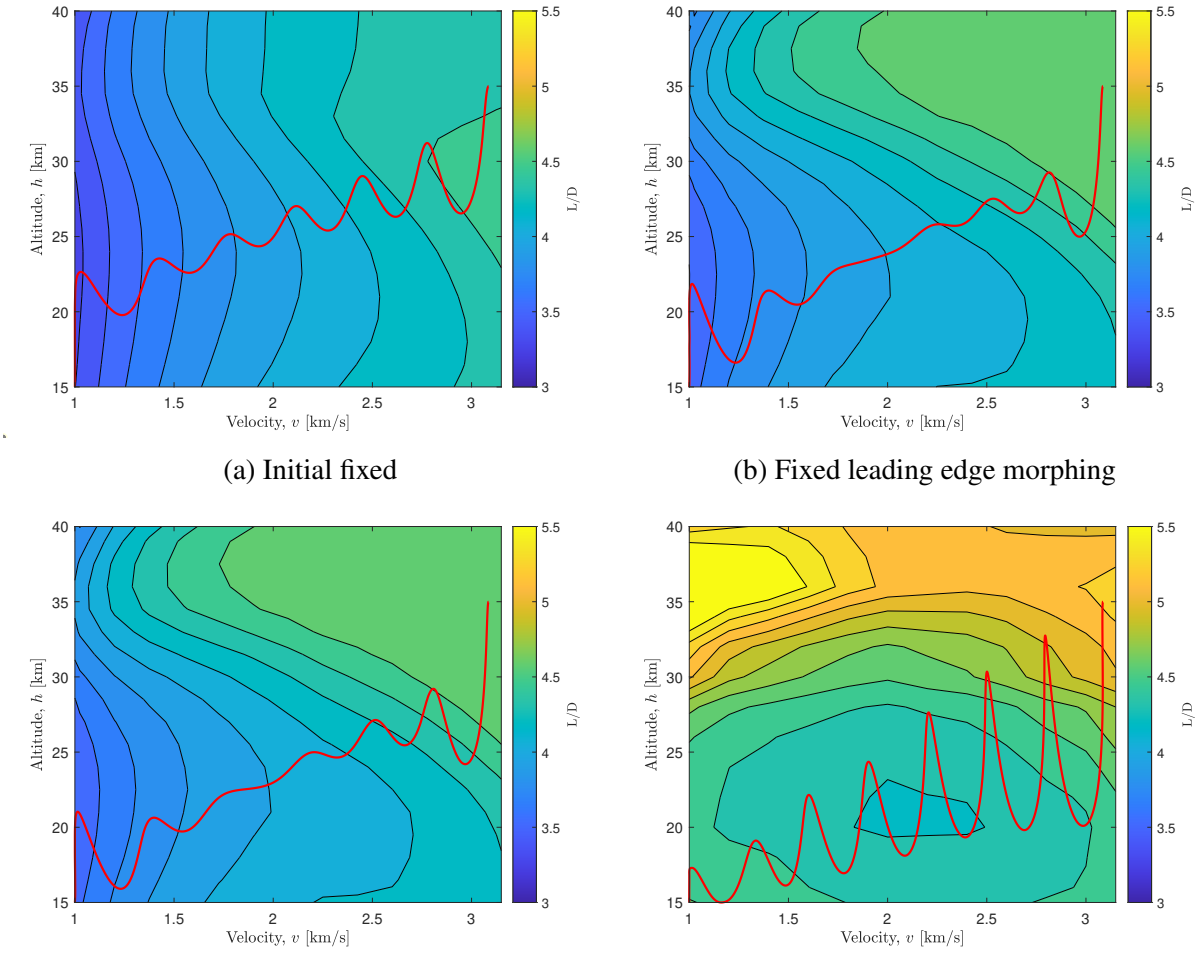

(a) Initial fixed

(b) Fixed leading edge morphing

(c) Flexible leading edge morphing

(d) Perfect morphing

Figure 3.13: Fixed and morphing waverider L/D aerotables and trajectory; $\alpha = 1°$; analytical pressure estimates

3.3.2 Morphing Trajectory

In this section the trajectory optimization results for the analytical model morphing geometries are compared with the initial fixed-geometry waverider results from Section 3.1.2. Overall trajectory statistics are provided in Table 3.6. While the constant-trajectory assumption in Table 3.5 predicted average lift-to-drag ratio increases of 7.4%, 7.6%, and 8.4% for the fixed leading edge, flexible leading edge, and perfect morphing waveriders, respectively, the actual trajectory analysis reveals significantly larger improvements: 12.2%, 14.5%, and 28.7%. These increases in lift-to-drag ratio result in noticeable increases in trajectory range, maximum Mach number, stagnation temperature, and flight time, as shown in Table 3.7. For these vehicles, each parameter increases in tandem with lift-to-drag ratio; the perfect morphing waverider performs best, followed by the flexible leading edge, then fixed leading edge morphing vehicles.

Table 3.6: Trajectory results for fixed and morphing waveriders: analytical model

Vehicle Type	Range [km]	M_{max} [-]	$T_{w_{\text{max}}}$ [K]	L/D_{avg} [-]	T_f [min]
Fixed	2087.20	10.09	2218.00	3.94	16.87
Fixed morphing	2256.04	10.15	2276.98	4.42	18.32
Flexible morphing	2266.40	10.17	2316.91	4.51	18.38
Perfect morphing	2458.01	10.32	2486.87	5.07	20.22

Table 3.7: Performance changes for morphing waveriders: analytical model

Vehicle Type	Range	M_{max}	$T_{w_{\text{max}}}$	L/D_{avg}	T_f
Fixed morphing	+8.1%	+0.6%	+3.1%	+12.2%	+8.6%
Flexible morphing	+8.6%	+0.8%	+4.5%	+14.5%	+9.0%
Perfect morphing	+17.8%	+2.3%	+12.1%	+28.7%	+19.9%

The trajectories for all three morphing vehicles are compared with the initial fixed waverider in this section. Fig. 3.14a shows the altitude over time. Similar to the results in Fig. 3.13, the perfect morphing vehicle exhibits higher amplitude oscillations in the phugoid compared to the initial fixed vehicle, while the morphing vehicles with constraints on the leading edge demonstrate comparably less oscillatory behavior. Additionally, the morphing vehicles all achieve increased flight times. As shown in Fig. 3.14b, the morphing vehicles were able to maintain higher Mach numbers for longer periods of time. The higher-amplitude oscillations contribute to increased maximum Mach numbers, as well as larger amplitude flight path angle time histories, seen in Fig. 3.14c. More notably, the morphing vehicles experience higher stagnation point temperatures, as shown in Fig. 3.4d. Although the increase in heating is slight for the morphing vehicles leading edge constraints, this suggests that while morphing can increase waverider range, it could also necessitate the use of more heat-resistant materials at the leading edge. Finally, as shown in Table 3.7, morphing noticeably increases the range compared to the fixed design.

The control time history of the morphing waveriders differs significantly from that of the fixed waverider. As seen in Fig. 3.4f, the morphing vehicles fly at a higher angle of attack, with the perfect morphing vehicle even reaching the 5° control limit. These increases in angle of attack control directly correlate with increases in overall lift-to-drag ratio, as shown in Fig. 3.4g. This suggests that the angle of attack a geometry is optimized at has a significant effect on the maximum lift-to drag ratio angle of attack. Additionally, the largest gains in lift-to-drag ratio for the morphing vehicles occur near the start and end of the glide. This is likely because the initial fixed vehicle was optimized at Mach 6.85, 25 km, and $\alpha = 0^\circ$, a flight condition that all vehicles approach near the midpoint of their respective trajectories. Consequently, as the vehicles' actual flight conditions diverge from the initial optimization point, the potential benefits of morphing

increase. This indicates that morphing waveriders could provide the greatest lift-to-drag ratio and range improvements for vehicles operating across a broad spectrum of flight conditions.

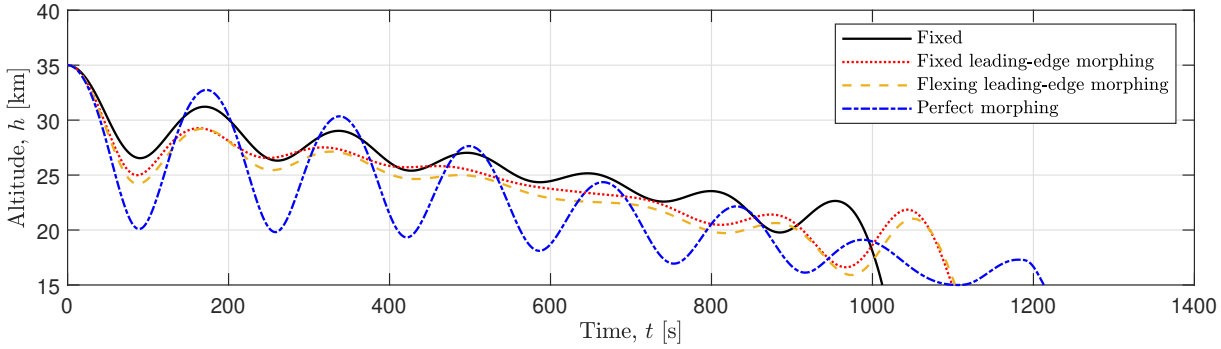

(a) Altitude time history

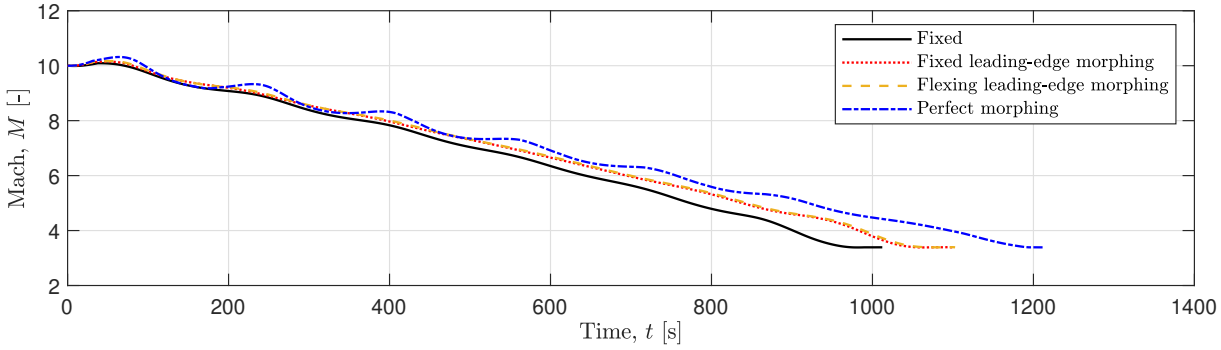

(b) Mach number time history

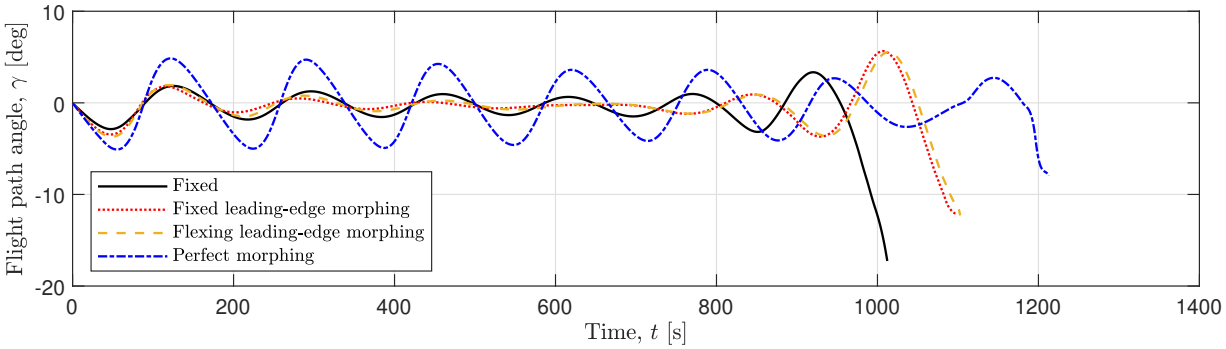

(c) Flight path angle time history

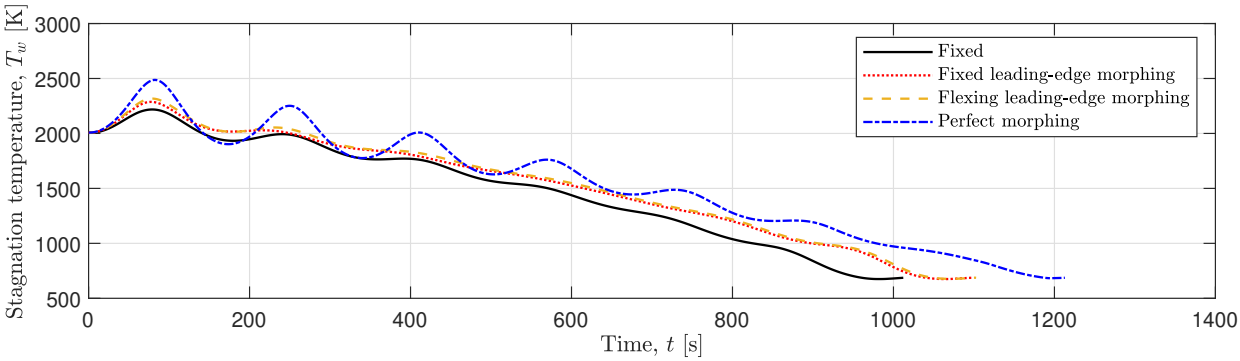

(d) Stagnation temperature time history

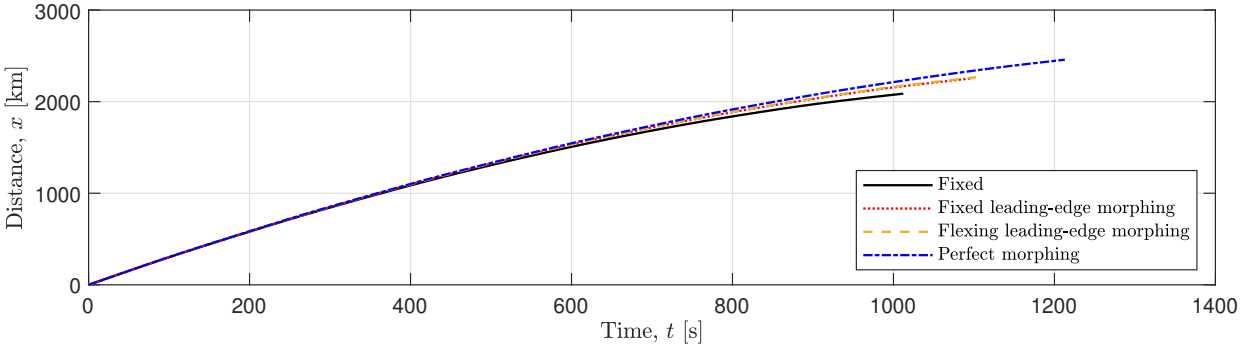

(e) Distance time history

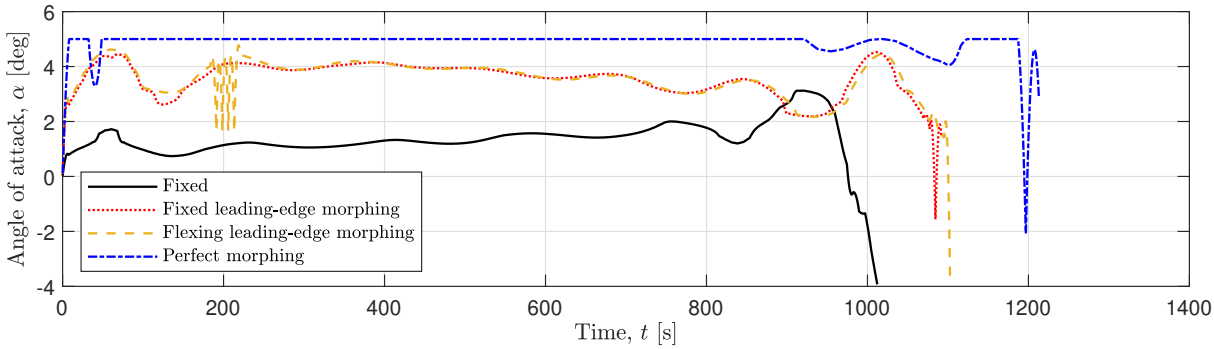

(f) Angle of attack time history

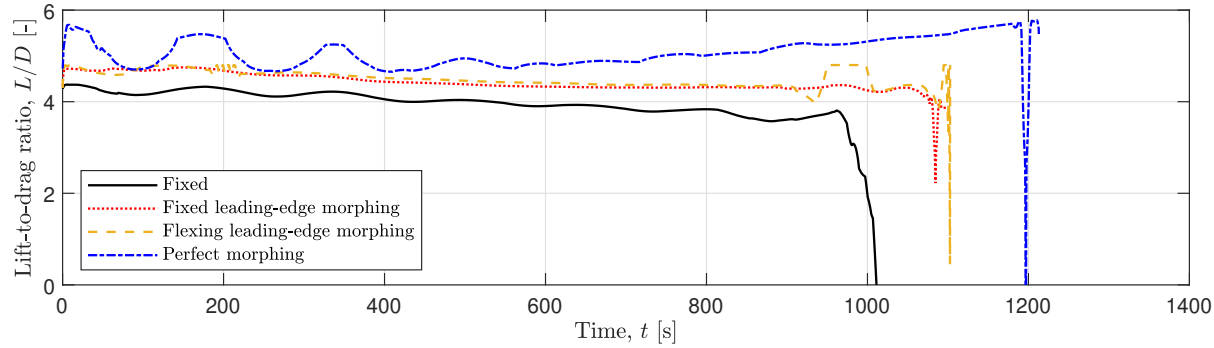

(g) Lift-to-drag ratio time history

Figure 3.14: Fixed and morphing waverider time histories; Mach 10, 35 km release attitude, analytical pressure estimates

Chapter 4: Results using Computational Fluid Dynamics

Now that a methodology for optimizing and evaluating the trajectories of morphing waverider geometries has been developed using the analytical model, this section repeats the process using CFD to calculate performance. The geometry of a new, non-morphing waverider is optimized and its trajectory is simulated to aid in the subsequent optimization of morphing geometries. The impacts of morphing are then calculated using CFD. Finally, the impact of CFD-based optimization is evaluated by simulating the morphing geometry optimized using the analytical model in CFD, and comparing its performance to that of the morphing geometry optimized using CFD.

4.1 Morphing Waverider Design using Computational Fluid Dynamics

First, an initial fixed waverider geometry is optimized at the same flight condition as in Chapter 3: Mach 6.85, 25 km altitude, and 0° angle of attack. The optimized geometry, which achieves a lift-to-drag ratio of 3.82, is shown isometrically in Fig. 4.1, and its front-view flow field is shown in Fig. 4.2. It is defined by:

$$w = 19.06 \text{ m}, \;\; n = 0.90, \;\; \beta = 9.00^\circ, \;\; \varepsilon = -0.56, \;\; p_1 = 0.98, \;\; p_2 = 1.02, \;\; p_3 = 0.97$$

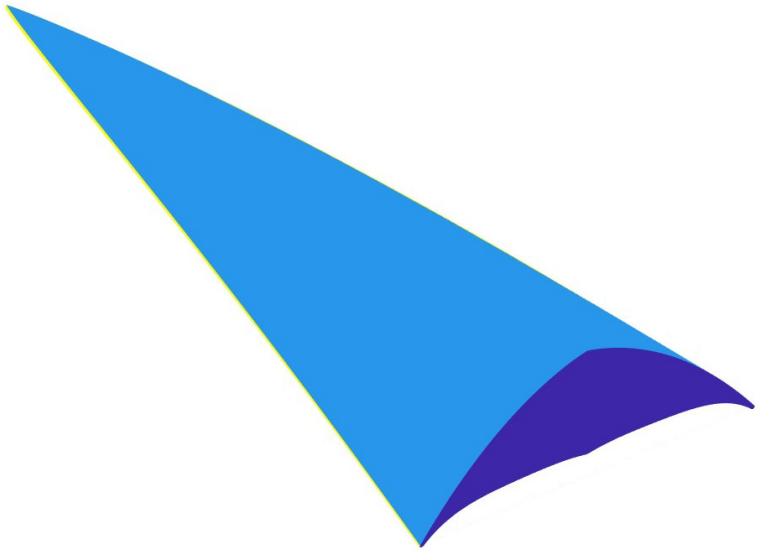

Figure 4.1: CFD-optimized, fixed waverider using; $L = 40$ m, $R = 0.1$ m, $M = 6.85$, $h = 25$ km, and $\alpha = 0^\circ$

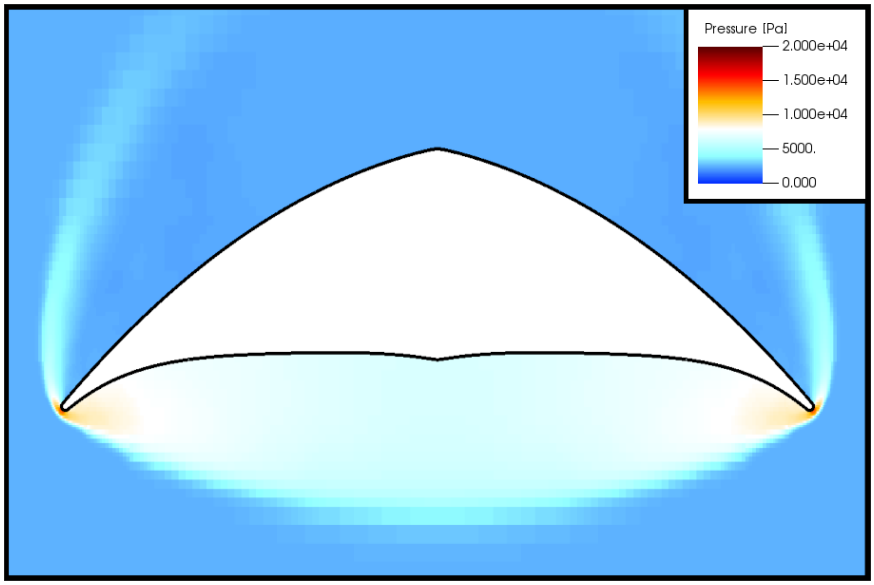

Figure 4.2: Front-view pressure flow field for CFD-based fixed waverider

Following the steps of the previous chapter, discrete lift and drag aerotables are generated for this geometry across a range of flight conditions using CFD simulations. These aerotables are used to generate a gliding trajectory from 35 km to 15 km altitude, starting at Mach 10, to identify the flight conditions where additional morphing optimization will take place. Fig.

4.3 shows this trajectory superimposed over the lift-to-drag ratio aerotable data, along with the selected morphing flight conditions. As shown, trends of the lift-to-drag ratio similar to the analytical model (Fig. 3.6) are observed, particularly the fact that higher lift-to-drag ratio can generally be achieved at high altitude points in the oscillations in the early stages of the trajectory. It is important to note that aerotable resolution was decreased due to the expensive function evaluations of the CFD, even with a GPU-based implementation. This explains the jagged contours in the performance contours.

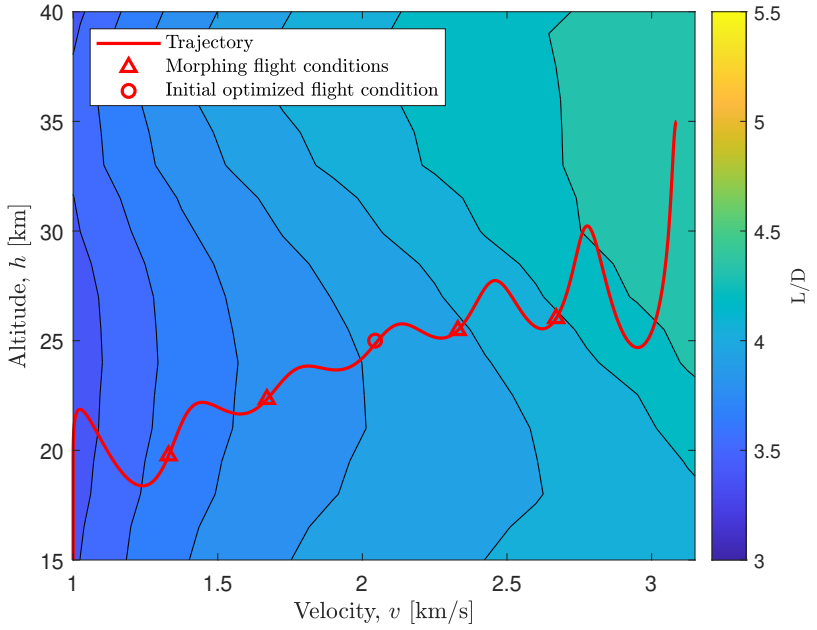

Figure 4.3: Fixed waverider L/D aerotables and trajectory with morphing points; $\alpha = 1°$; CFD pressure estimates ($L = 40$ m, $w = 19.06$ m, $n = 0.90$, $\beta = 9.00°$, $\varepsilon = -0.56$, $p_1 = 0.98$ $p_2 = 1.02$, $p_3 = 0.97$)

The flight conditions selected for morphing optimization are shown again in Table 4.1. It is notable that slightly lower altitudes are selected for morphing than when using the analytical model, which is an initial indication that lower overall altitudes are flown in trajectories that use CFD data. Additionally, the flight conditions have noticeably larger angles of attack, from 1°

using the analytical model, to about 2°. This could suggest that the CFD-optimized waverider achieves its maximum lift-to-drag ratio at higher angles of attack, and again likely indicates that the choice to optimize the initial fixed vehicle at 0° could impact the overall performance of the resulting morphing geometries.

Table 4.1: CFD waverider optimization flight conditions

Flight Condition	M [-]	h [km]	α [deg]
Fixed	6.85	25.00	0.00
1	8.93	26.02	2.30
2	7.80	25.47	2.45
3	5.63	22.33	2.35
4	4.51	19.76	2.00

4.1.1 Morphing with Fixed leading edge

The optimized morphing geometries with leading edge fixed are displayed in Table 4.2 and Fig. 4.4. Similar to the analytical model optimization, the Chebyshev control points decrease as Mach number decreases. A key difference is that all geometries have Chebyshev values that are significantly lower than for the fixed geometry starting point, and only have slight variation between flight conditions. Again, this could be because the fixed geometry was optimized for $\alpha = 0^\circ$, while its actual trajectory was flown at a slightly positive angle of attack. Because the chosen initial optimization flight condition may not have been well-suited for the resulting trajectory, it is likely that the fixed-leading-edge constraint was a serious restriction during morphing waverider optimization. Additionally, because the angle of attack difference suggests that there likely is a better performing fixed waverider than the one chosen, future work should include the angle of

attack in the optimization procedure. This would allow the advantages of surface morphing to be better leveraged.

Table 4.2: CFD waverider fixed leading edge morphing geometric parameters

Flight Condition	w [m]	n [-]	β [deg]	ε [-]	p_1 [-]	p_2 [-]	p_3 [-]
Fixed	19.06	0.90	9.00	-0.56	0.98	1.02	0.97
1	19.06	0.90	9.00	-0.56	0.65	0.65	0.62
2	19.06	0.90	9.00	-0.56	0.64	0.64	0.61
3	19.06	0.90	9.00	-0.56	0.55	0.56	0.53
4	19.06	0.90	9.00	-0.56	0.50	0.52	0.51

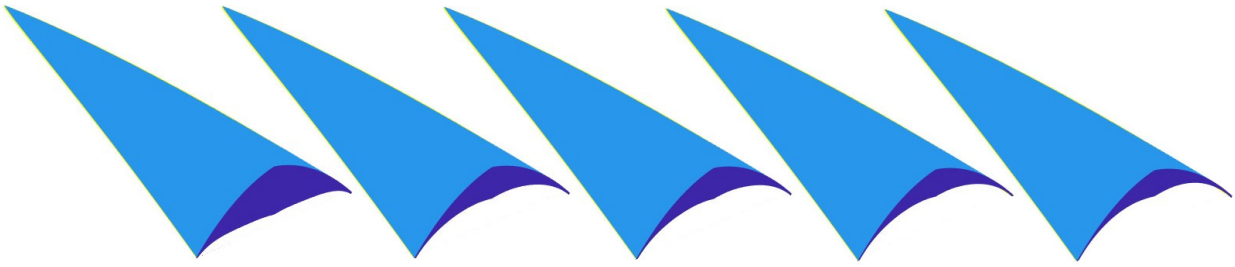

Figure 4.4: CFD waverider fixed leading edge morphing geometries; left to right: fixed geometry, morphing geometries 1-4

The front-view flow fields for all fixed-leading-edge morphing geometries are shown in Fig. 4.5. Each optimized geometry generates a shock that closely follows the leading edge, minimizing pressure bleed to the upper surface. However, it also clear that the leading edge shape of the initial fixed geometry is too wide at the higher mach numbers, leading to the lower surface shock to collapse towards the lower surface near the wingtips. The effects of angle of attack are also shown—significant expansion is observed on the upper surface of each waverider. This effect is especially obvious for geometry 3 due to the pressure range shown. Additionally, it can be seen that the shock produced was conical-like in shape.

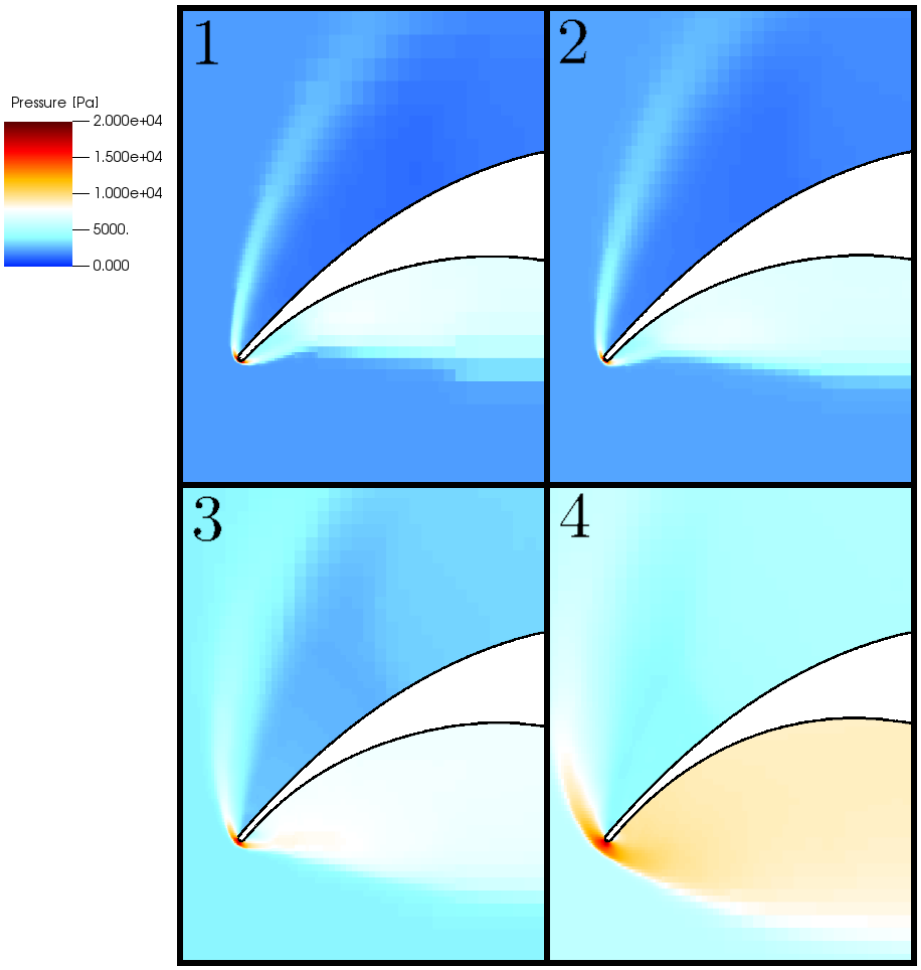

Figure 4.5: Front-view pressure flow field for CFD-based fixed leading edge morphing waverider

4.1.2 Perfect Morphing

The unconstrained morphing geometries optimized using CFD are shown in Table 4.3 and Fig. 4.6. While these vehicles are significantly thinner in the z-direction than the initial vehicle, the combination of design shock angle, β, and Chebyshev control points results in vehicles with decreasing thickness with Mach number. Additionally, these vehicles increase their width as Mach number increases, which could be the result of optimal geometries seeking to extend their wingtips close to the generated shock wave, which likely increases in span as Mach number decreases, following shock relations. The power law coefficient, n, stays constant for each vehicle,

again indicating that vehicles with more triangular planforms perform best limit leading edge pressure drag and streamwise viscous effects. Finally, it is important to note that no trend is seen in the curve factor, ε. This is likely due to the fact that as β decreases, the sensitivity of the leading edge shape to changes in curve factor decrease, as shown in Eq. 2.4.

Table 4.3: CFD waverider perfect morphing geometric parameters

Flight Condition	w [m]	n [-]	β [deg]	ε [-]	p_1 [-]	p_2 [-]	p_3 [-]
Fixed	19.06	0.90	9.00	-0.56	0.98	1.02	0.97
1	17.00	0.90	2.50	0.13	1.20	1.36	1.44
2	17.34	0.90	3.50	-0.43	0.90	0.98	0.97
3	17.65	0.90	1.75	-0.01	0.87	1.02	1.16
4	18.06	0.90	1.49	0.57	0.66	0.86	0.97

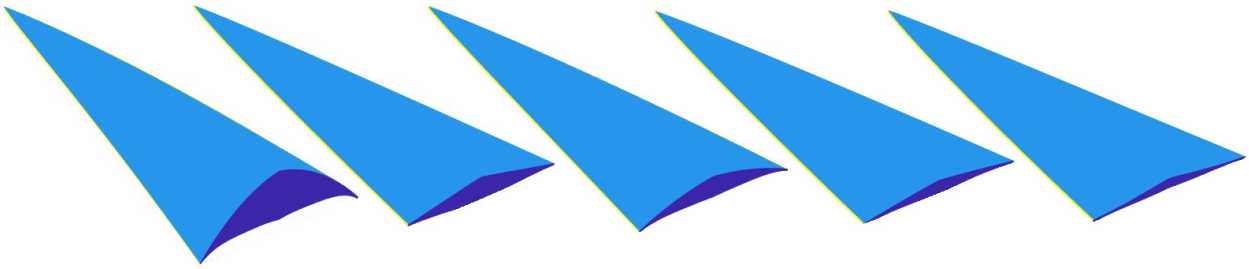

Figure 4.6: CFD waverider perfect morphing geometries; left to right: fixed geometry, morphing geometries 1-4

The perfect waverider flow fields are shown in Fig. 4.7. The lack of leading-edge constraints allows the optimizer to converge towards geometries that closely capture the shock waves across the body at the higher Mach numbers, without the complex shock effects seen at these conditions for the fixed-leading-edge morphing waverider. However, it is also evident that the lower Mach number geometries did not extend the width to completely capture the shock—instead taking advantage of suction at the wingtips due to the large inviscid expansion that the geome-

tries generate. This could show an advantage of using CFD-based waverider optimization: these tools capture off-design effects that improve waverider performance, while traditional analytical models, which assume an attached shock wave, do not.

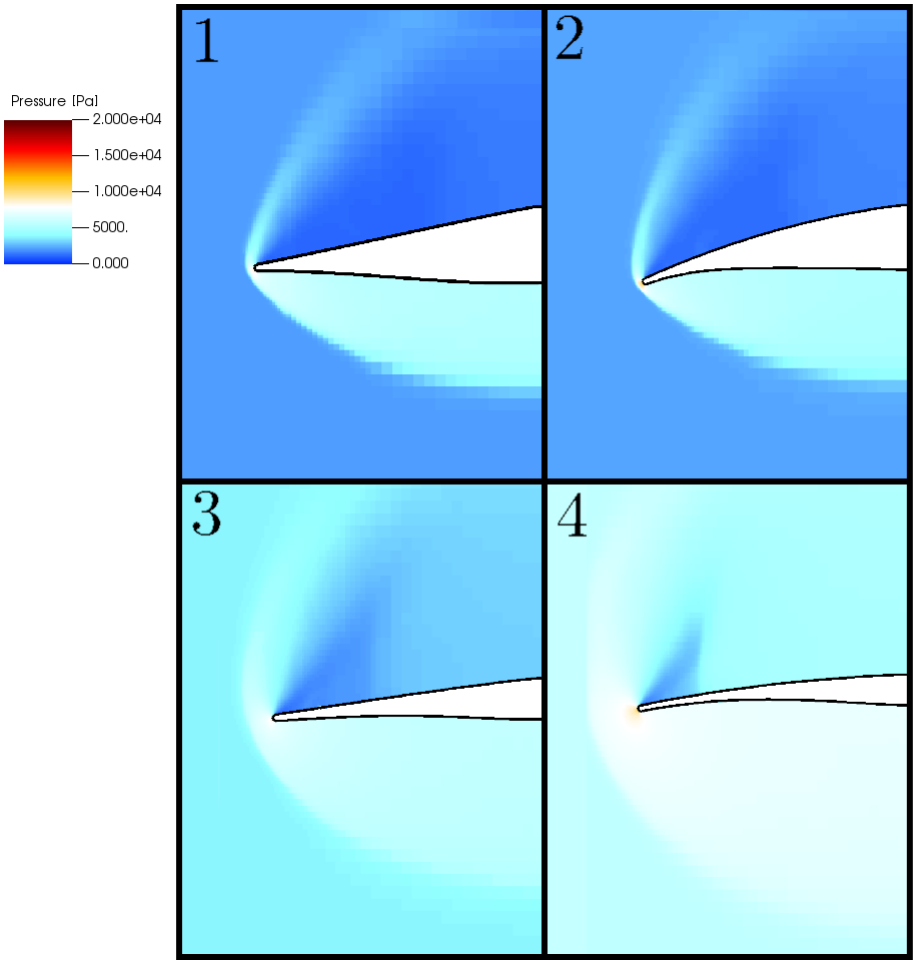

Figure 4.7: Front-view pressure flow field for CFD-based perfect morphing waverider

Fig. 4.8 shows a more in-depth view of the CFD-optimized morphing waverider front-view profiles. It again shows that the converged fixed-leading-edge morphing waverider geometries (4.8a) experience relatively little variation in lower surface shape, which could be due to the restrictions imposed by optimizing the initial vehicle at zero degrees angle of attack. However, it is interesting to note that the perfect morphing geometries also maintained similar profiles

throughout the trajectory (4.8b), which is a significant departure from the results of the analytical model. This may suggest that by allowing angle of attack to vary during the optimization process, a morphing waverider with leading edge constraints, or even a non-morphing waverider, could achieve comparable performance to the unrestrained 'perfect' configuration using CFD estimates. Also of interest are the differences between the analytical model (Fig. 3.10) and CFD-predicted optimal morphing shapes. The CFD-optimized geometries in every case are wider, likely to capture the shock wave across the lower surfaces, and thicker, probably because they fly at a larger angle of attack, which already has a similar effect to increasing vehicle thickness.

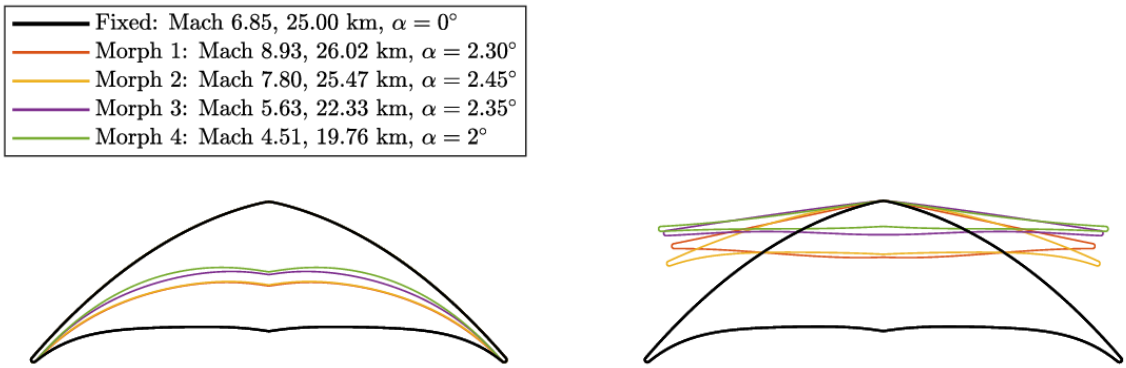

(a) Fixed leading edge

(b) Perfect morphing

Figure 4.8: CFD morphing waverider geometries; front profile comparison

4.2 Trajectory Analysis

This section discusses the trajectories resulting from the initial fixed and morphing waverider geometries designed using the CFD-based optimization. Table 4.4 shows the trajectory results of each vehicle and Table 4.5 shows the predicted impact of morphing using the flight condition selection methodology. As can be seen, surface morphing is predicted to give significant per-

formance improvements, but also result in higher stagnation point heating, which is similar to the predictions of the analytical model. While these comparisons are made to a fixed vehicle, which was likely optimized using incorrect angle-of-attack assumptions, it is notable that both sets of CFD-optimized morphing geometries achieved significantly better performance than the analytical model predicted for waveriders it optimized (Table 3.6). The impacts of optimizing unrestrained 'perfect' morphing for example, increased from a range of 2458 to 2725 km, and from an average lift-to-drag ratio of 5.07 to 5.63. This is notable, and emphasizes that while analytical models may predict overall trends accurately, they may at the same time give overall results that do not closely match CFD.

Table 4.4: Trajectory results for fixed and morphing waveriders: CFD

Vehicle Type	Range [km]	M_{max} [-]	$T_{w_{\text{max}}}$ [K]	L/D_{avg} [-]	T_f [min]
Fixed	2072.09	10.14	2288.63	3.93	16.81
Fixed morphing	2406.02	10.23	2397.40	4.80	19.72
Perfect morphing	2728.46	10.37	2550.25	5.64	22.49

Table 4.5: Performance changes for morphing waveriders: CFD

Vehicle Type	Range	M_{max}	$T_{w_{\text{max}}}$	L/D_{avg}	T_f
Fixed morphing	+16.1%	+0.9%	+4.7%	+22.1%	+17.3%
Perfect morphing	+31.7%	+2.3%	+11.4%	+43.5%	+33.8%

The trajectory results of each vehicle are compared below. One important takeaway is that the CFD results produced similar overall trends to those predicted by the analytical model. This includes the presences of a phugoid-like oscillation in both the altitude and Mach number time histories, with amplitude increased as additional morphing was allowed. These altitude and

Mach number results are shown most clearly in Fig. 4.9a. This effect is also shown in the Mach number time histories of Fig. 4.9b, which also shows that the morphing vehicles maintained high velocity longer. The oscillation is again shown in the flight path angle data of Fig. 4.9c. Morphing vehicles are predicted to experience higher stagnation point heating for longer periods of time, again showing the costs that come with improved reentry vehicle performance. Fig. 4.9e clearly shows that the morphing waveriders achieved significantly longer ranges and flight times than the fixed vehicle.

The angle of attack time histories are shown in Fig. 4.9f. Again, similar trends to the analytical model are observed, including positive angle of attack behavior throughout the trajectory, with an increased pitching up behavior near the terminal minutes of the glide. This behavior is likely a result of the final boundary conditions and represents the vehicle expending the maximum amount of energy to increase the range of the gliding stage. The morphing vehicles also flew at higher angles of attack than the non-morphing vehicles, approaching the angle of attack limit just like the analytical morphing waveriders did. Interestingly, the 'perfect' morphing waverider did not consistently maintain a higher angle of attack than the fixed leading edge waverider. This could be an effect of the restrictive angle of attack limitations, or could indicate that these vehicles had similar angles of attack for maximum lift-to-drag ratio. Finally, this lift-to-drag ratio data is shown in Fig. 4.9g, which clearly shows that the as the morphing restrictions were reduced, lift-to-drag ratio performance increased significantly.

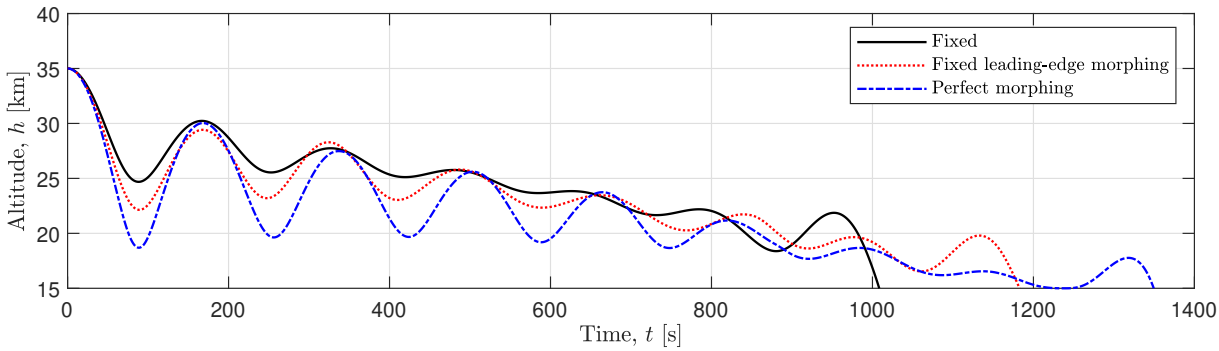

(a) Altitude time history

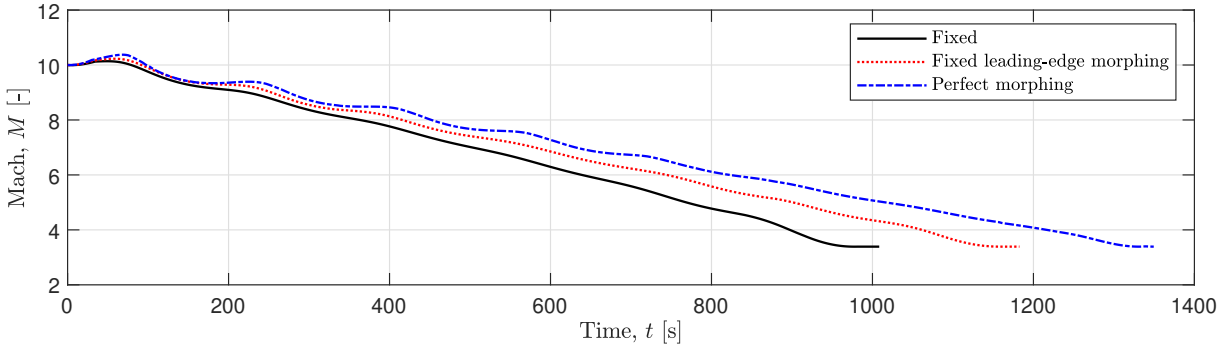

(b) Mach number time history

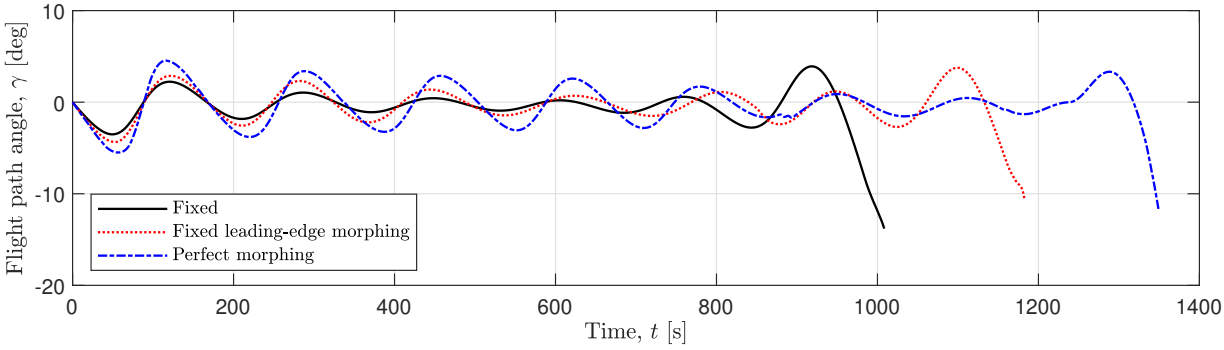

(c) Flight path angle time history

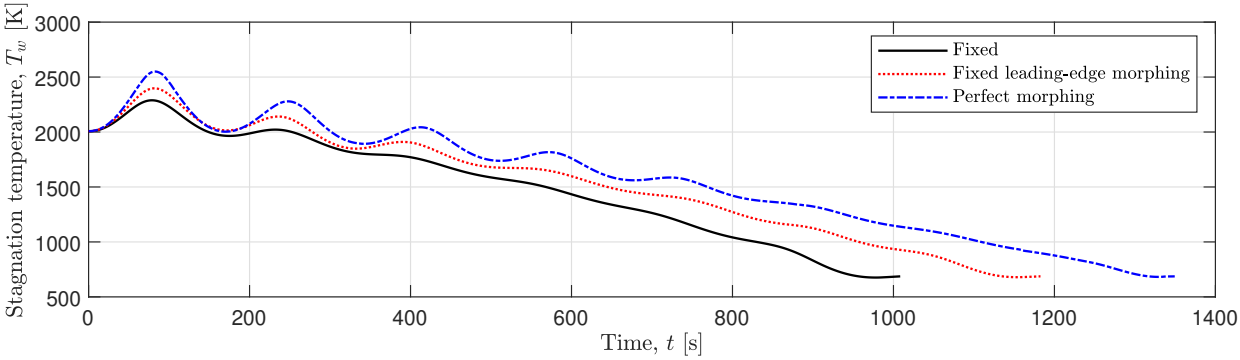

(d) Stagnation temperature time history

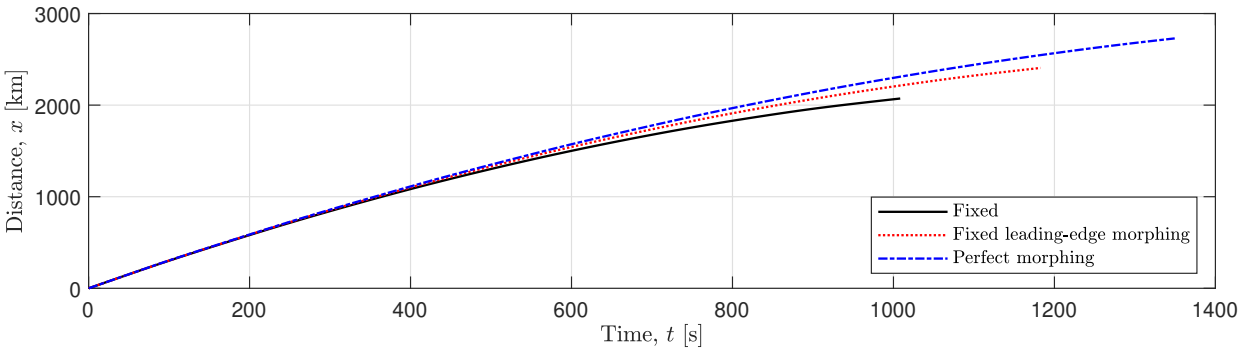

(e) Distance time history

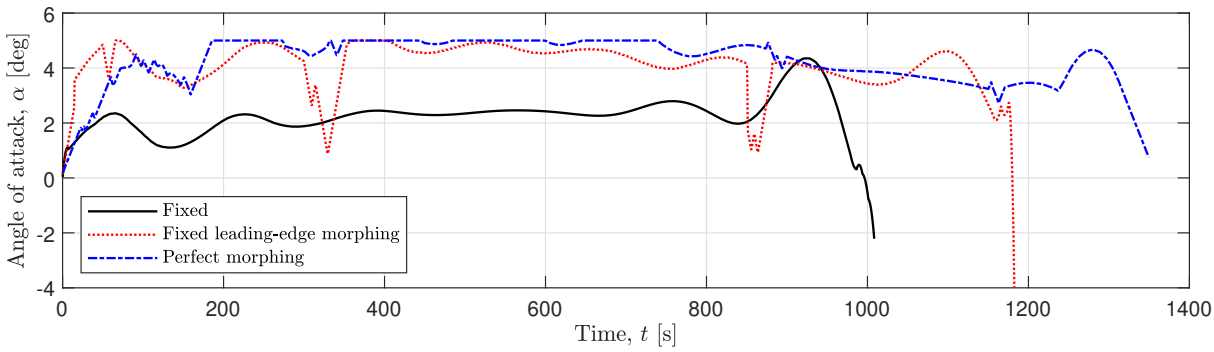

(f) Angle of attack time history

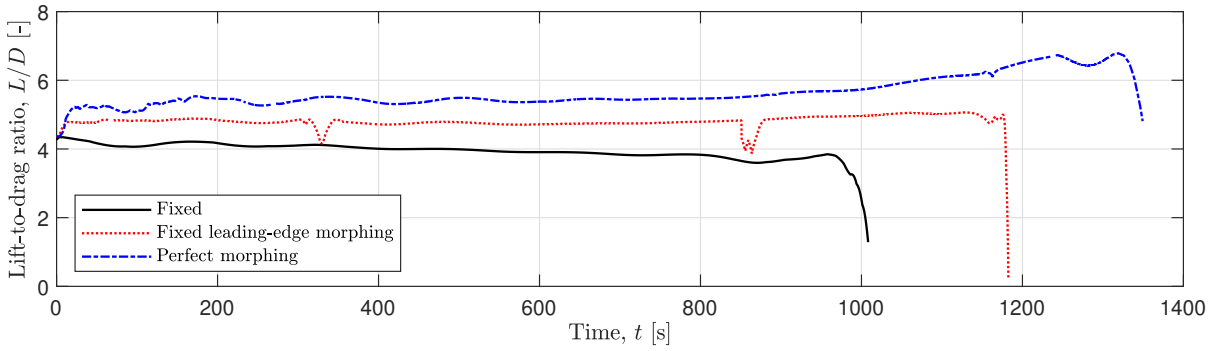

(g) Lift-to-drag ratio time history

Figure 4.9: Fixed and morphing waverider time histories; Mach 10, 35 km release attitude, CFD pressure estimates

4.2.1 Comparison to Analytical Model

To offer final comparisons between the analytical and CFD methods, the optimal geometries found using CFD are re-simulated using the analytical model, and vice-versa. For simplicity,

waverider trajectories generated using aerotables created using the CFD or analytical model will be called 'CFD-' or 'analytical-model-simulated,' while waverider geometries created using each method will be called 'CFD-' or 'analytical-model-optimized.' First, Figs. 4.10a and 4.10b show the CFD-optimized altitude and distance time histories simulated using both models. Evidently, the different models predict different trajectory behaviors for identical vehicles. While trends in performance are similar, the analytical model simulated trajectories with higher overall altitudes and shorter ranges than the CFD does. This indicates that even though the analytical model was shown to accurately predict the lift-to-drag ratio performance at some flight conditions for waveriders optimized using CFD in Fig. 2.17, this trend does not extend to the analytical-model-simulated trajectories in this study. This trend has serious implications for the safety predictions of the analytical model—Fig. 4.10c shows that the analytical model under-predicts stagnation point temperatures, which could lead to structural failure if a flight tested vehicle was designed with materials based on these results.

The analytical-model-optimized perfect morphing waverider trajectory is resimulated using CFD and compared to the CFD-optimized and simulated perfect morphing vehicle. The trajectory results are shown in this section, and important statistics are highlighted in Table 4.6. As can be seen, CFD optimization provides improved performance in every area of interest, even including a decrease in stagnation heating. Fig. 4.11a shows the altitude time histories for each vehicle. Here, the difference in flight time is noticeable, and it can be seen that the analytical-model-optimized vehicle has difficulties maintaining altitude throughout its trajectory. The analytical-model-optimized vehicle also experienced decreases in Mach number at a greater rate, as shown in Fig. 4.11b. The stagnation temperature time histories are shown in Fig. 4.11c. The analytical-model-optimized waverider experienced higher stagnation point heating in the first moments of

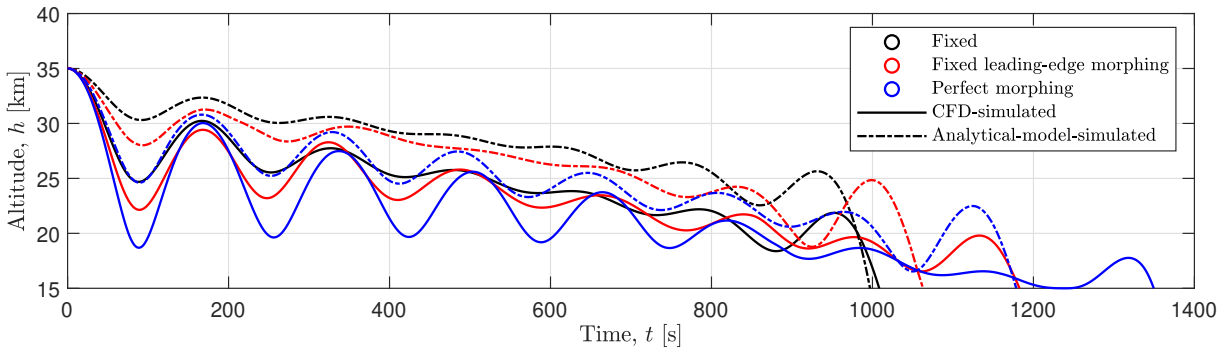

(a) Altitude time history

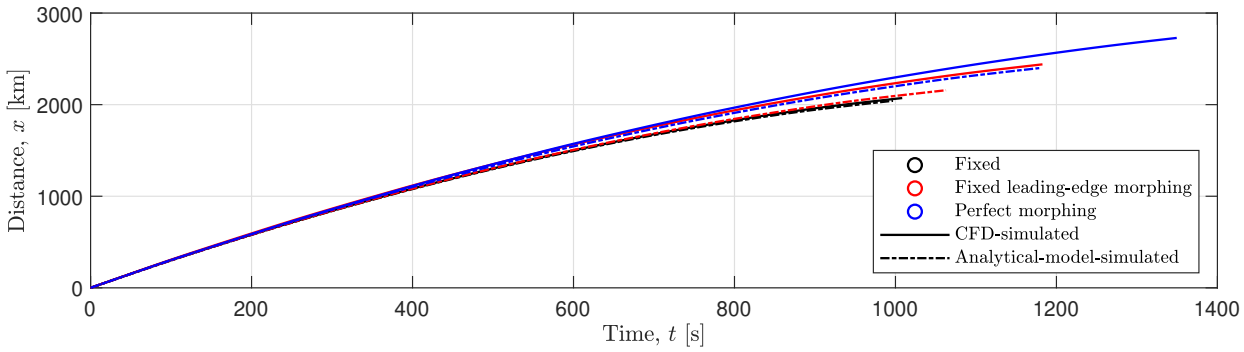

(b) Distance time history

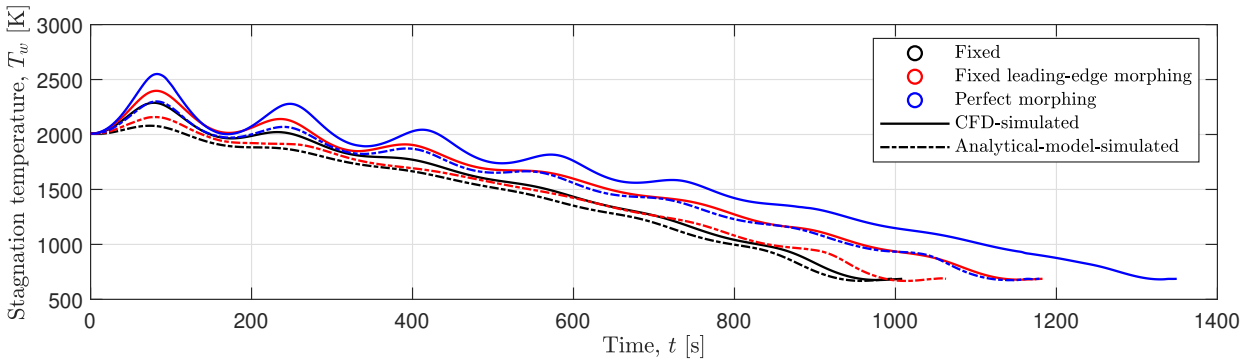

(c) Stagnation temperature time history

Figure 4.10: Analytical-model and CFD-simulated time histories for waverider geometries optimized using CFD

the glide, likely due to the lower altitude it reached in the first oscillation. The significant range impact can be observed in Fig. 4.11d. Fig. 4.11e shows the lift-to-drag ratio results of each vehicle. As shown, the analytical-model-optimized vehicle actually experienced similar lift-to-drag ratio performance for most of the trajectory, only experiencing a significant decrease near

the terminal stage of the glide where altitude and Mach number were low. This could suggest that the analytical-model-optimized geometries for these flight conditions perform particularly poorly in CFD.

Table 4.6: Performance comparison of analytical-model versus CFD-optimized perfect morphing waverider geometries; CFD-simulated trajectory results for both

Vehicle Type	Range [km]	M_{max} [-]	$T_{w_{\text{max}}}$ [K]	L/D_{avg} [-]	T_f [min]
Analytical Model	2261.58	10.37	2631.31	5.22	17.43
CFD	2724.90	10.38	2533.82	5.63	22.44
Performance impacts (%)	+20.5%	+0.1%	-3.7%	+7.85%	+28.7%

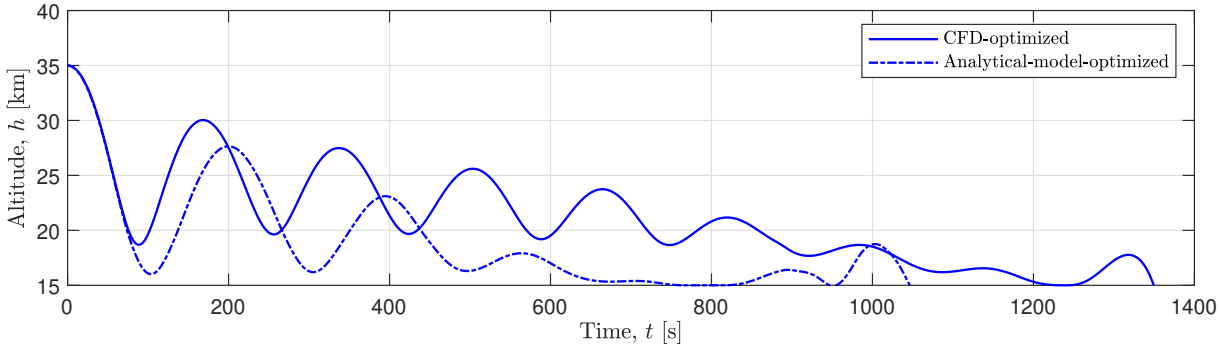

(a) Altitude time history

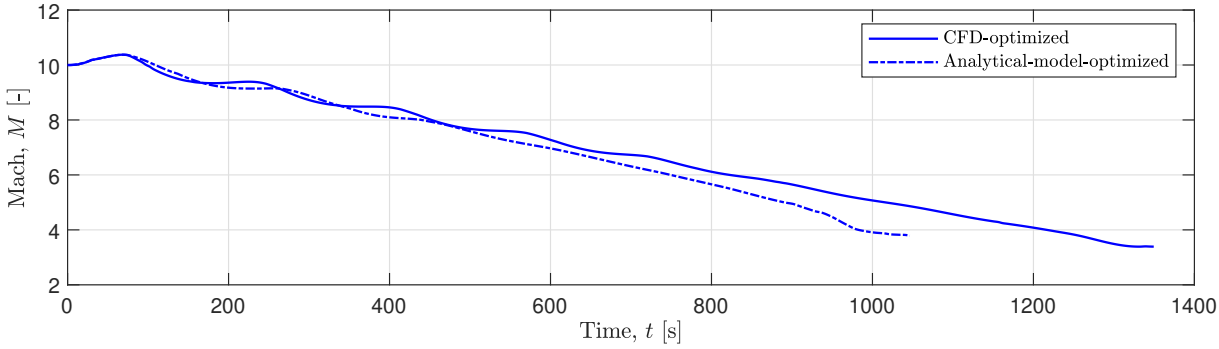

(b) Mach number time history

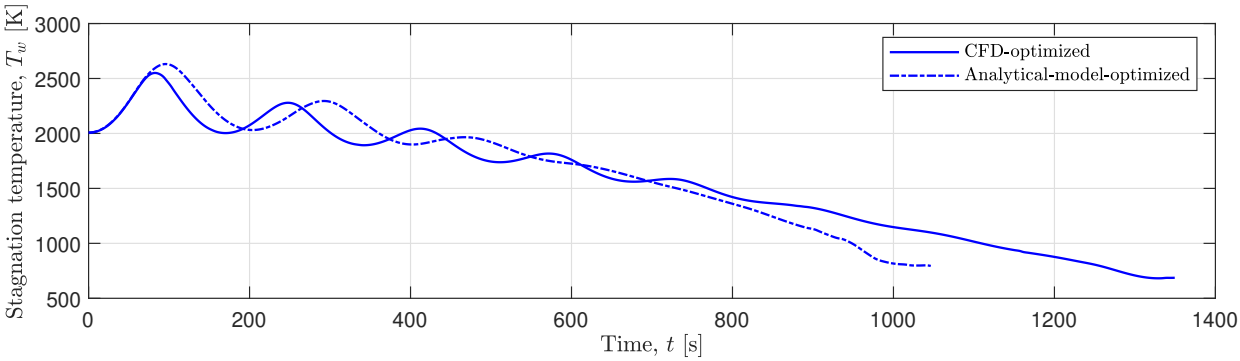

(c) Stagnation temperature time history

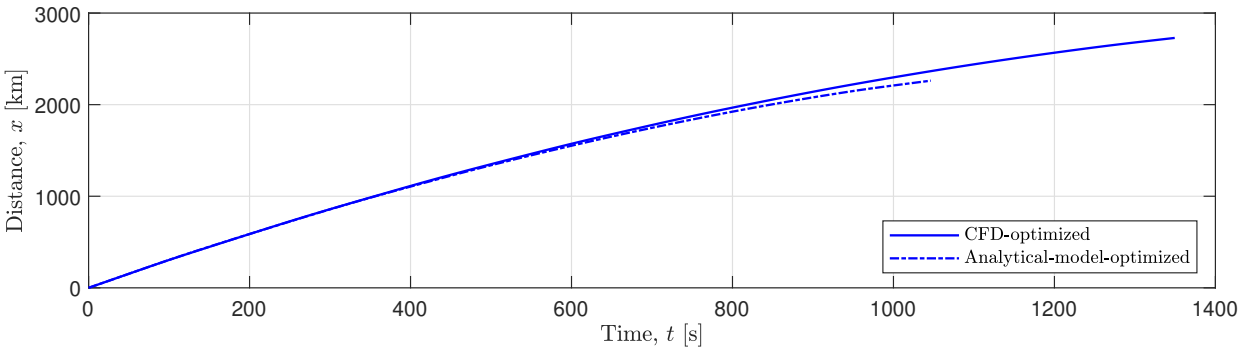

(d) Distance time history

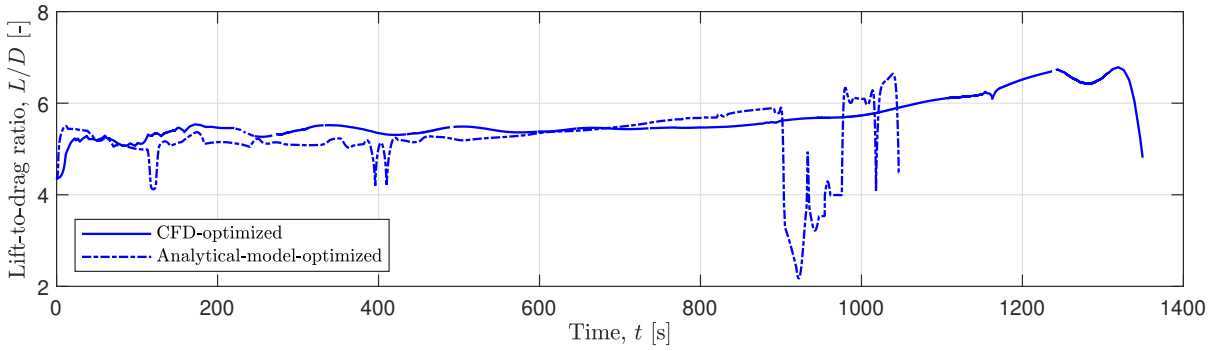

(e) Lift-to-drag ratio time history

Figure 4.11: CFD-simulated time histories for perfect morphing waverider geometries optimized using both the CFD and analytical model

These differences in performance between the analytical-optimized geometries are highlighted in Fig. 4.12, which shows the CFD-simulated flow fields of the analytical-optimized geometries. While each geometry experiences more pressure bleed than its CFD-optimized counterpart in Fig. 4.7, the low-speed geometries 3 and 4, in particular, experience extreme

pressure bleed. These flow fields, in conjunction with the trajectory data, highlight the weaknesses of the analytical model in predicting optimal waverider geometries at off-design conditions. While high-performing geometries can sometimes be optimized because of the analytical model's strengths at predicting pressure distributions for on-design vehicles, the failure of the model to predict features such as shock wave shape and pressure bleed can lead to morphing vehicles that perform poorly at some flight conditions.

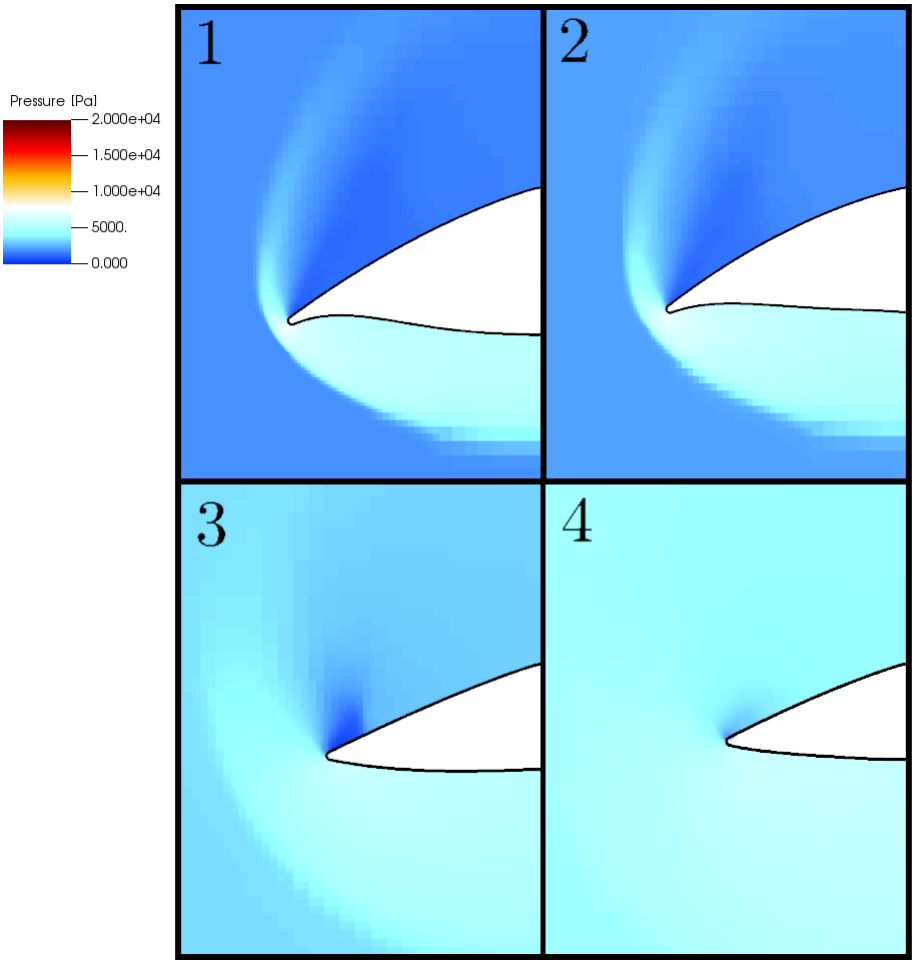

Figure 4.12: Front-view pressure flow field for analytical model-based perfect morphing waverider

4.3 Summary

The design methods used in this study to optimize morphing waveriders show potential for use in the optimization of vehicles with more restrictive design restraints. At single flight conditions, the developed analytical model provides estimates of waverider performance that are similar to those from CFD. However, these trends do not extend to waveriders flying across multiple flight conditions or in significantly off-design situations, likely due to the analytical model's weaknesses in predicting shock wave shape and pressure bleed effects. This highlights the need for CFD-based optimization techniques for morphing waveriders. By implementing CFD into the optimization process, vehicles were created that achieved high performance and experienced minimal pressure bleed. When used to generate families of morphing geometries, both the CFD and analytical model produced vehicles that were predicted to increase performance over the initial fixed waverider. However, the initial choice of flight conditions used to optimize these vehicles may not have been optimal, potentially requiring additional morphing optimization iterations to produce morphing and fixed vehicles with the best performance, and also to better understand the true impacts of morphing on realistic high-performing waveriders.

Chapter 5: Conclusions

The purpose of this study was to design and analyze various methods for evaluating waverider performance and to use these tools to optimize the lift-to-drag ratio and range of morphing waveriders for civilian transportation applications. A planar-shock-based waverider geometry definition was created that allowed for leading edge rounding and the application of arbitrary perturbations to vehicle surfaces. This tool effectively enabled the creation of intentionally off-design geometries and the exploration of how these effects could be combined to improve the performance of round-leading-edge planar waveriders. An inviscid analytical pressure model was developed using tangent-wedge oblique shock and expansion relations, paired with Modified Newtonian theory and symbolic regression equations to model a blunted leading edge. The GPU-based CHAMPS+ computational fluid dynamics tool was then employed to efficiently solve the inviscid Navier-Stokes equations and provide a higher-order approximation of inviscid performance. The Eckert reference temperature method was utilized to predict viscous drag for both models.

Once these two methods were developed, they were paired with a gradient-free Pattern search optimization algorithm, which effectively identified waverider geometries that achieved optimal lift-to-drag ratio performance at a desired flight condition. It was found that these methods predicted significantly different optimal geometries under identical input conditions, and

these differences were discussed. These prediction methods were also utilized to develop discrete lift and drag aerotables across a range of flight conditions, enabling the use of the ICLOCS2 direct collocation algorithm to generate optimal glide trajectories.

Groups of morphing waverider geometries were identified by selecting several flight conditions along an initial waverider trajectory and optimizing additional geometries at these new conditions. To generate a single set of aerotables for a morphing vehicle with multiple configurations, aerotables were first created for each geometry and then combined, with the lift and drag values of the local highest-performing geometry saved at every flight condition. This process resulted in trajectories that simulated the transition between different morphing geometries as the trajectory progressed. Multiple morphing geometries, with varying constraints, were created using both the analytical and CFD methods, and the results were compared. Finally, the optimal morphing geometry generated using the analytical model was simulated in CFD to compare the impacts of each method. In general, the key results of this study are as follows.

1. The reduced-order analytical model accurately predicted overall trends in waverider performance by utilizing a combination of simple pressure approximation methods, and provided highly accurate lift-to-drag ratio predictions for waveriders with minimal leading edge pressure bleed;

2. The analytical and CFD-based prediction tools converged to designs with off-design features that produced conical-like shock shapes. The weaknesses of the analytical model at predicting the effects of lateral pressure bleed led it to converge to vehicles with high pressure losses in CFD simulations;

3. The simple method chosen to identify and combine families of morphing waverider ge-

ometries using an initial trajectory resulted in significant predicted range improvements over a fixed waverider when using morphing, from 8% (fixing the leading edge) to 17% (unconstrained morphing) when using the analytical method for optimization and performance evaluation, and from 17% (fixing the leading edge) to 31% (unconstrained morphing) using CFD. This suggested that performance improvements could still be achieved after applying additional constraints on areas such as volume, structure, and heating;

4. The weaknesses of the analytical model were compounded when used to generate vehicle trajectories: when compared to analytical-model optimized geometries evaluated using CFD, geometries optimized and evaluated in CFD improved range by +20% and lift-to-drag ration by +7%.

5.1 Recommendations for Future Work

The wide range of topics covered in this study has led to the exploration of several aerodynamic and vehicle design concepts that would benefit from further analysis. Such analysis would allow for an improved characterization of the benefits of lower surface morphing in the context of real-world flight performance. It is first recommended that a variety of generating shock shapes be explored for the creation of the initial on-design waverider geometry. The planar shock-derived waverider is the simplest of waverider shapes and can even be considered a special case in the design space of conical shock-derived vehicles, as the radius of curvature becomes very large. Generating shock shapes with more degrees of freedom have already been shown to result in better-performing waveriders [10]. Therefore, utilizing the perturbation method developed in this study with improved shock shapes is likely to provide a deeper understanding of the true

effects of these perturbations and lead to better-performing morphing vehicles in general.

Next, both performance estimation tools used in this study could likely be improved to better approximate the true flow features observed during reentry. For the analytical method, it is possible that improved panel-method techniques or aerodynamic relations such as the pressure increment model from Guo et al. [17], could be developed to more accurately predict the effects of leading edge rounding on the lower surface pressure distribution, as well as rear surface expansion effects. If successful, these advancements could enable analytical tools to converge to waverider geometries that achieve close-to-optimal performance in CFD simulations. Further improvements may also be possible using higher-fidelity viscous drag estimates. As the capabilities of GPU-based computational fluid dynamics continue to improve, higher fidelity viscous models like Reynolds-Averaged Navier-Stokes (RANS) or Large-Eddy Simulation (LES) could be integrated into the optimization loop.

Additionally, while the methodology used to generate a family of morphing waveriders was predicted to provide large performance increases over a non-morphing vehicles, the decision to optimize the initial fixed geometry at zero degrees angle of attack resulted in suboptimal starting vehicles. Morphing waverider design would likely benefit from an iterative process that converges to morphing flight conditions that better align with the trajectory flown by the resulting geometry. In this study, it is likely that fixed and morphing vehicles with improved performance could be designed if the initial fixed vehicle were optimized at the angle of attack for maximum lift-to-drag ratio. Furthermore, the selection of morphing flight conditions at equidistant points throughout the trajectory was somewhat arbitrary, and primarily driven by time constraints. Additional studies may identify alternative methodologies that lead to higher-performance morphing waveriders.

It would also be valuable to consider the effects of more realistic constraints on geometry generation and trajectory analysis. This could include constraints typically considered during the vehicle development process, such as usable volume and structural limitations. Realistic mechanical constraints also warrant further attention. Additionally, the impacts of restrictive constraints on morphing waverider range were not fully explored. These constraints include lower maximum allowable stagnation temperatures, n-factor limitations, and altitude restrictions based on modern-day restrictions for overland supersonic flight. The rocket launch and terminal phases of flight are also essential components of a real-world civilian reentry glide vehicle, and morphing waverider optimization would benefit by including these flight stages.

Finally, the high speed of GPU-based computational fluid dynamics tools has the potential to enable the optimization of waverider geometries with more complex design factors included into their definition. These could include the addition of fuselage sections with specified passenger or cargo space, the modeling of deflecting control surfaces, or even engine integration. Researchers have already demonstrated that realistic reentry vehicle design features can be integrated automatically into waverider geometry definitions [22]. Combined with the findings in this study, which show that high-efficiency CFD can be incorporated into the waverider optimization process, it is clear that vehicles generating complex flowfields can be optimized automatically, provided the correct implementation can be used. Overall, it is recommended that the process outlined in this paper be repeated, incorporating restrictive and realistic constraints on the vehicle geometry and mission parameters.

Bibliography

[1] Anderson, J. D., Jr., *Hypersonic and High-Temperature Gas Dynamics*, 3rd ed., AIAA, Reston, VA, 2019.

[2] Nonweiler, T.R.F., "Aerodynamic Problems of Manned Space Vehicles," *Journal of the Royal Aeronautical Society*, Vol. 63, pg. 521-528, 1959.

[3] Nonweiler, T.R.F., "Delta Waves of Shape Amenable to Exact Shock Wave Theory", *Journal of Royal Aeronautical Society*, Vol. 67, pg. 39, 1963.

[4] Jones, J.G., "A Method for Designing Lifting Configurations for High Supersonic Speeds Using the Flow Fields of Non-Lifting Cones," Royal Aircraft Establishment, Report No. Aero. 2674, Mar 1963.

[5] Taylor, G.I., Maccoll, J.W., "The Air Pressure on a Cone Moving at High Speed", *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol. 139, No. 838, pp. 278-311, 1933.

[6] Pike, J., "A Design Method for Aircraft Basic Shapes with Fully Attached Shock Waves Using Known Axisymmetric Flow Fields," R.A.E. Technical Report No. 66069, 1966.

[7] Starkey, R. P., *A Parametric Study of L/D and Volumetric Efficiency Trade-Offs for Waverider Based Vehicles*, M.S. Thesis. University of Maryland, College Park, 1998.

[8] Rasmussen, M., "Waverider Configurations Derived from Inclined Circular and Elliptic Cones," *Journal of Spacecraft and Rockets*, Vol. 17, No. 6 , pp. 537-545, 1980.

[9] Sobieczky, H., Dougherty, F.C., Jones, K., "Hypersonic Waverider Design from Given Shock Waves," Proc. 1st Int. Hypersonic Waverider Symposium, Univ. of Maryland, 1991.

[10] Bowcutt, K. G., Anderson, J., D., Capriotti, D., "Viscous Optimized Hypersonic Waveriders," *AIAA 25th Aerospace Sciences Meeting*, AIAA Paper 87-0272, Jan 1987.
doi: 10.2514/6.1987-272

[11] Corda, S., Anderson, J.D., Jr., "Viscous Optimized Waveriders Designed from Axisymmetric Flow Fields," AIAA Paper 88-0369, 1988.

[12] Takashima, N., and Lewis, M. J., "Navier-Stokes Computation Of a Viscous Optimized Waverider," *Journal of Spacecraft and Rockets*, Vol. 31, No. 3, pp. 383-391 May-Jun 1994.

[13] He, X., Rasmussen, M. L., "Computational Analysis of Off-Design Waveriders", *Journal of Aircraft*, Vol. 31, No. 2 , pp. 345-353. 1994.
10.2514/3.46493

[14] Zhao, Z., Huang, Wei., Yan, L., and Yang, Y., "An Overview of Research on Wide-Speed Range waverider Configuration," *Progress in Aerospace Sciences*, Vol. 113, 2020.
10.1016/j.paerosci.2020.100606

[15] Maxwell, J., *Morphing Waveriders for Atmospheric Entry*, Ph.D Dissertation. University of Maryland, College Park, 2019.

[16] Maxwell, J., and Pheonix, A., "Morphable Hypersonic Waverider and Trajectory Optimized for Atmospheric Entry," AIAA SPACE and Astronautics Forum and Exposition, Sep 2017.
doi: 10.2514/6.2017-5357

[17] Guo, S., Liu, W., Zhang, C., and Liu, Y., "Aerodynamic Optimization of Hypersonic Blunted Waveriders Based on Symbolic Regression," *Aerospace Science and Technology*, Vol. 144, Jan 2024.
doi: 10.1016/j.ast.2023.108801

[18] Santos, W., "Bluntness Effects on Lift-to-Drag Ratio of Leading Edges for Hypersonic Waverider Configurations," *18th AIAA/3AF International Space Planes and Hypersonic*

Systems and Technologies Conference, 2011.
doi: 10.2514/6.2012-5802

[19] Liu, W., Zhang, C., Wang, F., and Ye, Z., "Design and Optimization Method for Hypersonic Quasi-Waverider," *AIAA Journal*, Vol. 58, No. 5, pp. 2132-2146 May 2020.
doi: 10.2514/1.J059087

[20] Faruqi, F., "Integrated Navigation, Guidance, and Control of Missile Systems: 2-D Dynamic Models," Australian Government Dept. of Defense, Defense Science and Technology Group, Sep 2013.

[21] James, B., *Aero Database Development and Two-Dimensional Hypersonic Trajectory Optimization for the High-Speed Army Reference Vehicle*, M.S. Thesis. University of Maryland, College Park, 2023.

[22] Lobbia, M. A., "Multidisciplinary Design Optimization of Waverider-Derived Crew Reentry Vehicles," *Journal of Spacecraft and Rockets*, Vol. 54, No. 1, pp. 233-245 Jan-Feb 1994.
10.2514/1.A33253

[23] Betts, J. T., *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*, 2nd ed., Society for Industrial and Applied Mathematics, Philidelphia, PA, 2010.

[24] Rizvi, S. T., LinShu, H., DaJun, X., and Shah, S., "Trajectory Optimisation for a Rocket-Assisted Hypersonic Boost-Glide Vehicle," *The Aeronautical Journal* , Vol. 121, Mar 2017.
doi: 10.1017/aer.2017.11

[25] Yu, L., and Nai-gang, C., "Optimal Attack Trajectory for Hypersonic Boost-Glide Missile in Maximum Reachable Domain," *2009 IEEE International Conference on Mechatronics and Automation*, Aug. 2009.
doi: 10.1109/ICMA.2009.5246695

[26] Boyd, F., *Chebyshev and Fourier Spectral Methods*, 2nd ed., Dover, Mineola, NY, 2001.

[27] Tincher, D. J., and Burnett, D., W., "Hypersonic Waverider Test Vehicle: A Logical Next Step," *Journal of Spacecraft and Rockets*, Vol. 31, No. 3, pp. 392-399, May-Jun 1994.
doi: 10.2514/3.26451

[28] Eckert, E. R. G., "Survey of Boundary Layer Heat Transfer at High Velocities and High Temperatures," WADC Tech. Rep., pp. 59-624, Apr 1960.

[29] Sutherland, W., "The Viscosity of Gases and Molecular Force," *Philosophical Magazine*, S. 5, pp. 507-531, 1893.

[30] van Noordt, W., McQuaid, J., Larsson, J., and Brehm, C., "GPU-Accelerated Simulations of Turbulent Flows With the Immersed Boundary Method," *AIAA AVIATION FORUM AND ASCEND 2024*, Jul 2024.
doi: 10.2514/6.2024-3536

[31] van Noordt, W., Ganju, S., and Brehm, C., "An Immersed Boundary Method for Wall-Modeled Large-Eddy Simulation of Turbulent High-Mach-Number Flows," *Journal of Computational Physics*, Vol. 470, Dec 2022.
doi: 10.1016/j.jcp.2022.111583

[32] Audet, C., and Dennis Jr., J. E., "Analysis of Generalized Pattern Searches," *SIAM Journal on Optimization*, Vol. 13, No. 3, pp. 889-903, 2002.
doi: 10.1137/S1052623499352024

[33] Abramson, M. A., "Mixed Variable Optimization of a Load-Bearing Thermal Insulation System Using a Filter Pattern Search Algorithm," *Optimization and Engineering*, Vol. 5, pp. 157–177, Jun 2004.
doi: 10.1023/B:OPTE.0000033373.79886.54

[34] Kim, B. S., Rasmussen, M. L., and Jischke, M. C., "Optimization of Waverider Configurations Generated from Axisymmetric Conical Flows," *Journal of Spacecraft and Rockets*, Vol. 20, No. 5, pp. 461-469, 1983.

[35] Nie, Y. Faqir. O., and Kerrigan, E. C., "ICLOCS2: Try this Optimal Control Problem Solver Before you Try the Rest," *2018 UKACC 12th International Conference on Control (CONTROL)*, pp. 336-336, 2018.
doi: 10.1109/CONTROL.2018.8516795

[36] Yechout, T. R., *Introduction to Aircraft Flight Mechanics*, 2nd ed., AIAA, Reston, VA, 2014.