UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Aeronautics, Astronautics and Computational Engineering

Computational Engineering & Design Group

ON DEVELOPING EFFICIENT PARAMETRIC GEOMETRY MODELS FOR WAVERIDER-BASED HYPERSONIC AIRCRAFT

by

Konstantinos Kontogiannis

Thesis for the degree of Doctor of Philosophy

November 2017

UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Computational Engineering & Design Group

Thesis for the degree of Doctor of Philosophy

ON DEVELOPING EFFICIENT PARAMETRIC GEOMETRY MODELS FOR WAVERIDER-BASED HYPERSONIC AIRCRAFT

Konstantinos Kontogiannis

ABSTRACT

This work is focused on tackling the high dimensionality and complex nature of waverider-based high speed aircraft design through the development of effective and efficiently parameterized parametric geometry models. The first part of the work is focused on the parameterization and handling of waverider forebody geometries. Different design approaches and a novel design method are presented, each offering direct control of different aspects of the geometry. This can be utilized to directly implement any design constraints or to enable straightforward interfacing with additional geometry components. The new three-dimensional leading edge waverider design method that is proposed is a step away from inverse and one towards direct waverider design. A series of requirements for designing valid three-dimensional leading edge curves are also highlighted. A method to compare different parameterization schemes in order to avoid over or under-parameterizing the geometries and assist in deciding on the number of degrees of freedom and control points for the design-driving curves of the inverse design methods is also presented. This enables the designer to make better educated decisions during the parametric model development phase and when parameterizing hypersonic-design-specific components for which we have limited experience and detailed data in the literature. Complementing the waverider forebody component of the aircraft are a series of blunt leading edge shape formulations. Their effectiveness and efficiency compared to other blunting approaches is highlighted. They are suitable for generating blunt shapes for any wedge-like geometry and they can also be used for inlet cowls, sidewalls, control surfaces, etc. They also offer second order continuity at the interface between the blunt part and the original geometry, which can have a favourable effect on the receptivity and turbulent transition mechanism. Finally, a parametric geometry model development framework consisting of a revised aerodynamic design process that involves design loops to better tune the parametric model, and a geometry engine developed to enable these early design loops, is presented. This is complimented by a number of implementation specific findings and proposed features such as an interactive GUI with real-time updates and dynamically controlled resolution of the generated geometries.

Table of Contents

List of Tables

List of Figures

DECLARATION OF AUTHORSHIP

I, ...

declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

....

....

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;

2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;

3. Where I have consulted the published work of others, this is always clearly attributed;

4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;

5. I have acknowledged all main sources of help;

6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;

7. Parts of this work have been published as:

    i.   Kontogiannis, K., Sobester, A., & Taylor, N. J. (2015). *On the Conceptual Design of Waverider Forebody Geometries*. 53rd AIAA Aerospace Sciences Meeting, Kissimmee, FL.

    ii. Kontogiannis, K., Cerminara, A., Taylor, N. J., Sobester, A., & Sandham, N. (2015). *Parametric Geometry Models for Hypersonic Aircraft Components: Blunt Leading Edges*. 20th AIAA International Space Planes and Hypersonic Systems and Technologies, Glasgow, UK.

    iii. Kontogiannis, K., Taylor, N., & Sóbester, A. (2016). *Efficient Parameterization of Waverider Geometry Models for Hypersonic Aircraft: Integrated External Inlet Compression*. 54th AIAA Aerospace Sciences Meeting, San Diego, CA.

    iv. Kontogiannis, K., Sóbester, A., & Taylor, N. (2016). *Efficient Parameterization of Waverider Geometries*. Journal of Aircraft, 54(3), 890–901.

    v.   Kontogiannis, K., Sóbester, A., & Taylor, N. (2017). *Waverider Design Based on Three-Dimensional Leading Edge Shapes*. Journal of Aircraft, 54(5), 2010–2012.

Signed: ....

Date: ....

Acknowledgements

First and foremost, I would like to thank my advisor Dr András Sóbester for his guidance throughout my studies. He pointed me in the right direction every time I needed to and offered a significant amount of his time reviewing research material and making sure my studies were progressing smoothly. This research has only been possible because of his help.

I would also like to thank Dr Nigel Taylor, Capability Leader Aerodynamic Tools & Methods in MBDA UK, for his input throughout this research and the industry perspective he offered. The route this research followed was greatly influenced by our collaboration. Additionally, I would like to thank my second advisor Dr Neil Sandham for his valuable advice whenever I sought it and for also reviewing and providing feedback on various pieces of research material. I would also like to offer my thanks to Adriano Cerminara, who recently finished his PhD studies, for supporting part of the research material in this thesis.

Most importantly, I wish to thank my family who has supported and encouraged me throughout my studies abroad, and finally my girlfriend, Fiona, who has always been there for me and has so far put up with me in my time here.

Definitions and Abbreviations

English letters

C_L lift coefficient

C_D drag coefficient

D_{local} distance of upper surface and shockwave on an osculating plane

d   parametric curve knot multiplication number; thickness of leading edge section

G^i geometric continuity of order i

h   flight altitude

k   coefficient of thermal conductivity

k_i i^{th} knot or break point of parametric curve

L_{local} local length of waverider on an osculating plane

L/D   lift to drag ratio

M   Mach number

n   order of parametric curve

P_i coordinate array of control point *i*

p   static pressure

p_w mean wall static pressure

p'_w wall pressure fluctuation amplitude

p'_{inf} freestream pressure fluctuation amplitude

p_{inf} freestream static pressure

p_{local} local static pressure

q   wall heat flux

q'_w wall heat flux fluctuation amplitude

R   radius

R_c radius of curvature

R_{gas} gas constant

Re   flow Reynolds number

r, θ   radial and angular coordinates of polar coordinate system

T   temperature

$\vec{T}$ local tangent vector of curve

T_A, T_B	Ferguson spline tangent vectors
t	curve parameter ranging from 0 to 1
u	flow velocity; curve/surface parameter
u_e	flow velocity at the edge of the boundary layer
$\overrightarrow{u_{inf}}$	freestream velocity vector
v	curve/surface parameter
V_r	radial velocity component
V_θ	angular velocity component
U_∞	freestream flow velocity
var_i	the i^{th} shape defining variable for the blunt leading edge
V_r	the radial component of velocity in spherical coordinates
V_θ	the angular component of velocity in spherical coordinates
V/A	the ratio of the volume to the wetted/surface area of a 3D geometry
w_i	weight of the i control point in rational parametric curves
x, y, z	Cartesian coordinates
x', y', z'	first derivative of curve coordinates
x", y", z"	second derivative of curve coordinates

Greek letters

α	osculating plane angle to vertical axis
β	effective shock angle, different from shock angle for curved shock waves
γ	specific heat ratio; angle between shockwave profile and upper surface
δ	flow deflection behind shock; boundary layer thickness
ε	emissivity
θ	angle between upper surface profile and horizontal axis
θ_c	cone half angle
θ_s	shock angle
λ	effective leading edge sweep angle
μ	fluid viscosity
ρ	radius of curvature; fluid density
ρ'	density fluctuation
φ	angle between straight line segments of discretized curve representation

Subscripts

inf	freestream properties
i, j, k	indices
w	properties at wall

Abbreviations

CAD	Computer Aided Design
CAE	Computer Aided Engineering
CFD	Computational Fluid Dynamics
FCC	Flow Capture Curve
FFD	Free-Form Deformation
ICC	Inlet Capture Curve
ICFA	Internal Conical Flow A
LSPC	Lower Surface Profile Curve
LSPD	Lower Surface Profile Definition
MOC	Method of Characteristics
NACA	National Advisory Committee for Aeronautics
NASA	National Aeronautics and Space Administration
PLED	Planform Leading Edge Definition
SSTO	Single Stage To Orbit
TSTO	Two Stage To Orbit
RANS	Reynolds Averaged Navier-Stokes
RBLE	Rational Bézier Leading Edge
SWPC	Shock Wave Profile Curve
USPC	Upper Surface Profile Curve
USPD	Upper Surface Profile Definition

Chapter 1:   Introduction

Over the last couple of decades engineering design has shifted from the drawing board to computer monitors, and also shifted from direct modelling to the development of parametric geometry models. One of the factors that have greatly contributed to that is the growing use of computational design optimization techniques, especially in complex engineering design areas such as aeronautics. The present work is focused on the development of such parametric geometry models for the definition of the outer mould lines of hypersonic aircraft.

With one of mankind's greatest dreams of being able to fly being fulfilled in the early 20th century, it was predetermined that what would follow would be the desire to fly further and faster, and despite all the strategic reasons that might have pushed countries to investing in research that would enable that, our pure desire of wanting to go faster cannot be denied. Aviation has come a long way since surpassing one of the first barriers, the speed of sound, and has already made the first steps in order to take the next leap.

Hypersonic flight has been an area of interest and research since the early Cold War era, when even jet engines were in their early stages, and as with many other aerospace technological advances of that time it was mostly driven by defence and military oriented projects. However, in the following years research efforts were redirected to the space race, with reconnaissance satellites seeming like a more viable option than a hypersonic reconnaissance aircraft, and also to the development of intercontinental ballistic missiles. Since the turn of the century, sustained hypersonic flight has been increasingly attracting research interest worldwide.

Flight at hypersonic speeds, meaning from Mach 5 and above, presents many challenges due to the more complex aerothermodynamic phenomena involved. This is reflected in the small number of high speed aircraft examples available. The most successful operational manned aircraft to date is the X-15, which was rocket engine powered and capable of Mach 6.7. The only air

breathing aircraft that comes close is the SR-71 reconnaissance plane which was capable of Mach 3.5 sustained flight, and although that was not in the hypersonic region, it incorporated various key technologies such as dual mode turbojet-ramjet engines and the need for cooling due to aerodynamic heating. The space shuttle was another prime example of a high-hypersonic velocity aircraft design. As far as civil aviation is concerned, where low-cost is much more important than sheer speed, the only two notable examples of aircraft that have seen operation are the Concorde and the Tupolev Tu-144 in the late 1960s, operating at much lower supersonic speeds of around Mach 2. In more recent years and as scramjet engine technologies are making significant progress, sustained flight at hypersonic speeds has been attracting increasing interest for various applications, ranging from small/unmanned and large/transport hypersonic cruise craft to single-stage or two-stage to orbit launch vehicles. There are currently several research programs on hypersonic flight and development of scramjet engines run by a number of institutes around the globe.

With the on-going development of scramjet engines many obstacles and inefficiencies that were related to the atmospheric use of rockets engines, which were the only option to reach such velocities in the past, can be overcome. An air-breathing hypersonic vehicle only has to carry the fuel on-board and not the oxidizer. It can also be a lifting body as opposed to most rocket-powered vehicles that counter gravity with thrust alone, thus making a horizontal take-off and landing configuration possible. Despite being under development for more than a decade, scramjets are currently the most promising technology for sustained hypersonic flight. The use of scramjet engines is one of the factors that greatly influences aircraft design. Given the operating envelope of altitudes and Mach numbers where scramjet operation is viable, as well as the principle on which they function, the design of such engines requires much larger intakes and exhaust nozzles compared to conventional engine designs of subsonic and low-supersonic aircraft. This leads to engines that form a very substantial part of the vehicle and need to be highly integrated with the aircraft structure from the early stages of the design process. Various methods of airframe-engine integration have been proposed in concept studies over the years, and as this field is in its early stages, there is no specific configuration that can be considered to be more advantageous over the others yet.

Amongst the most promising concepts for designing lifting body geometries, and especially the forebody, which is the forward part of the airframe that greatly influences the overall characteristics of the flow around the vehicle and the engine intake, are waveriders. Initially conceived by Nonweiler (1959), they are aircraft shapes where the generated shockwave in the oncoming flow is attached to the leading edge of the body. The high pressure area behind the shock

is separated or 'trapped' on the underside of the vehicle resulting in an increase in lift, hence the name 'waverider'.

Hypersonic aircraft and waverider design methods in the past were inevitably focused on generally low-fidelity approaches, such as empirical relations and inverse design methods based on analytical solutions of inviscid flowfields, which needed to be strongly supported by experiments. With the advances in computational resources, computational fluid dynamics and modelling of aerothermodynamic phenomena in recent years, much more accurate computations of the performance characteristics of entire hypersonic vehicles have become possible. Moreover, the use of multidisciplinary design search and optimization methods is becoming a viable option for more complex applications. In order to take full advantage of high-fidelity simulations coupled with computational design-search and optimization techniques, efficient and effective parametric geometry models are required. Developments in computer-aided design (CAD) have made significant progress on this aspect over the last years, many times driven by the aerospace industry itself. Hypersonic aircraft design is inherently a multidisciplinary problem where the aircraft components are also much more intertwined, thus requiring that many aspects of the design be addressed simultaneously in a coupled approach. This, along with the fact that there is limited design experience in the field and larger design spaces need to be explored, further emphasizes the need for effectively and efficiently parameterized geometry. Along with modern computational geometry tools, several of the aforementioned inverse design methods are still being used, especially in the early design stages, as they can provide adequate approximations with minimal computational cost. These methods can also provide some good starting points or baseline shapes upon which more complex geometries can be built, and dramatically reduce the dimensionality of larger scale design optimization studies. However, details on their parameterization are much more scarce in comparison to the material available on conventional aircraft design.

Nowadays, developing a computational parametric geometry model has essentially replaced hand drawing, however, the ratio of formally backed scientific design decisions to ones based on the intuition, experience and imagination of an aerodynamicist has probably not shifted either way during that transition. Among the goals of this work is to provide some methods to further justify decisions in the development stages of the parametric geometry and also to contribute in some areas by adding to the pool of available tools for hypersonic aircraft design.

The aim of the current work is to develop efficient parametric geometry models to aid the design of hypersonic aircraft, with the majority of findings also applicable to supersonic design. This will be accomplished through the development of a hypersonic aircraft parametric geometry engine and through additional investigations in areas that will be identified during that process. One of the

more flexible inverse design methods, the osculating cones method (Sobieczky, 1990), will be mainly used to obtain the baseline waverider shapes. These designs are essentially inviscid but using such inverse design methods is a very efficient way to obtain 'waveriding' lifting body geometries. Aspects of those baseline geometries will be further altered, manipulated and parameterized using existing and novel approaches where deemed necessary in order to construct efficient and flexible geometry models that will generate aircraft shapes able cope with all aspects of hypersonic flight as well as take advantage of modern computational design techniques. The geometry engine will be developed in order to assist a proposed aerodynamic design process that aims to tackle a number of additional challenges faced by hypersonic aircraft design.

The question this thesis aims to answer is: How can the parametric geometry model development phase for a waverider-based hypersonic aircraft, which is an area where we have limited collective experience, be improved and better rationalized in order to ensure efficiency and effectiveness of the model prior to high fidelity design studies?

To answer that question, the following research objectives have been identified:

1. Investigation of different parameterization approaches offered to the numerous waverider design methods that have been proposed in order to efficiently meet design requirements and assist the highly coupled design of hypersonic aircraft.

2. The parameterization of a geometry model generally involves finding the optimum trade-off between flexibility (breadth of design space coverage), robustness and dimensionality (number of design variables). In the case of hypersonic waverider forebody geometries, this thesis is to propose methods to estimate the correct balance prior to resorting to expensive computational simulations.

3. Explore the possibility of developing novel design methods for waveriders and for their blunt leading edge to enhance our design capabilities and potentially increase the efficiency and effectiveness of parametric models.

4. Develop a geometry engine that enables the development of efficient parametric geometry models by utilizing the proposed methods, and can be utilized in a realistic aerodynamic design process.

Chapter 2: Background

2.1   Hypersonic Aircraft Design

Hypersonic speed in aerodynamics is one that is highly supersonic. Since the 1970s, flight speeds of Mach 5 and above have been considered to be in the hypersonic regime, with NASA classifying speeds between Mach $10 - 25$ as being in the high hypersonic regime and those even higher as re-entry speeds. What differentiates it from the supersonic regime are a number of effects that start being noticeable at speeds higher than Mach 5. The combination of all of these effects, as illustrated in Figure 2.1, introduce additional challenges to hypersonic aircraft designers.
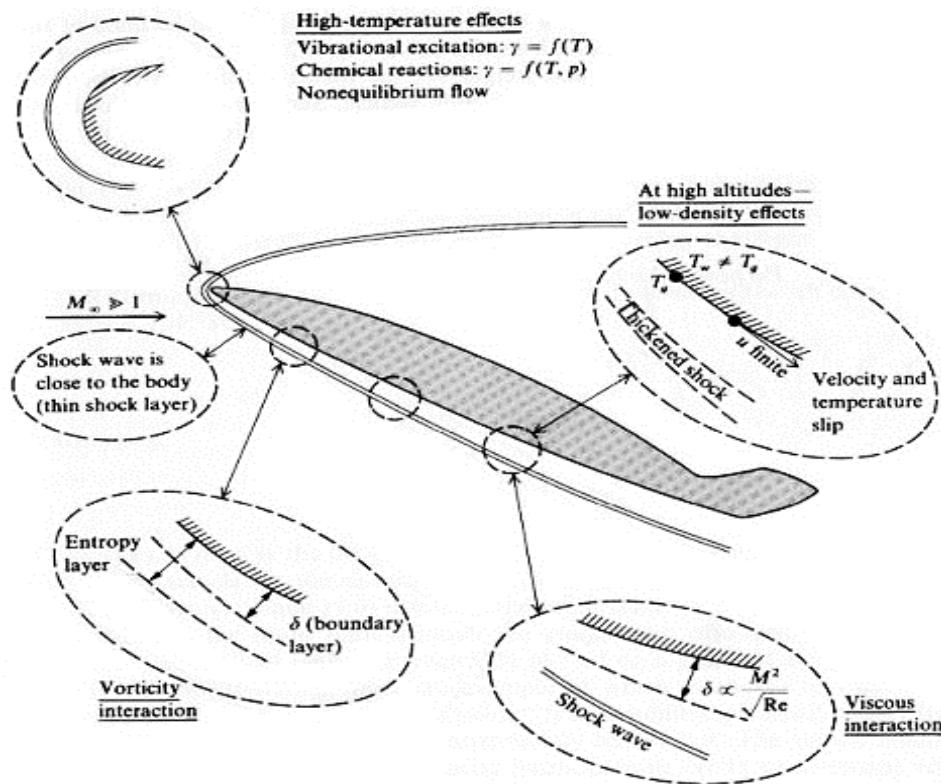
Figure 2.1.   Flow characteristics affecting hypersonic vehicles (Anderson, 2006).

The flow characteristics affecting hypersonic aircraft, as described by Anderson (2006), can be summarized as:

- In high Mach numbers the shock lies close to the body and interacts with the boundary layer in what is called the shock layer.
- The gases experience a dramatic temperature increase behind the shock and close to the body, aerodynamic heating effects substantially dictate the aircraft design and require the use of blunt leading edges.
- Blunt leading edges generate oblique shocks with an entropy gradient causing a phenomenon called vorticity interaction.
- The boundary layer in hypersonic flow, unlike supersonic and subsonic flows, gets thicker as Mach number increases due to the high temperatures and subsequent increase of viscosity and decrease of density. This enhances the viscous interactions in the shock layer.
- In sufficiently high velocities or re-entry speeds, the temperature rises enough to excite the vibrational energies of the atoms and cause dissociation or even ionization. Radiation from the hot gases also starts being significant.
- In very high altitudes, the continuum assumption for air starts to break down due to very low densities.

These characteristics are highly intertwined and their combination renders hypersonic vehicle aerodynamic design a very challenging endeavour and greatly differentiates it from supersonic designs. All these effects might not be present simultaneously depending on the flight envelope, as seen in Figure 2.2. The primary design considerations for a hypersonic vehicle can therefore differ significantly depending on the application. The present work is mostly focused on hypersonic air-breathing cruise vehicles, though a significant part of it applies to all classes of hypersonic aircraft. It might have been almost impossible to take all these effects and their interactions into account when utilizing design methods based on empirical relations and assumptions that had to be made a few decades ago. However, accurate numerical simulations for the majority of these complex phenomena are currently possible. This, combined with the limitations of supersonic and hypersonic wind tunnels and testing facilities, emphasizes the importance of computational design techniques in this field. All the aforementioned phenomena contribute to the multidisciplinary nature and the complexity of the design evaluation methods and simulations required. This generally leads to increased computational costs in the design process. A number of the hypersonic flow characteristics mentioned earlier, such as the complications of rarefied air and re-entry velocities are not within the scope of the present work. However, the developed parametric geometries are to be utilized to enable complex design studies that tackle

most of the aforementioned characteristics, even though they are not explicitly accounted for during the parametric model development phase.

Aerodynamic heating is a major design-driving concern for all hypersonic aircraft. In the case of the X-51 waverider demonstrator the thermal considerations were something that even affected the flight duration of the mission (Lane, 2007). The temperature and heat flux in various areas of the airframe need to be considered. Sharp leading edges that would be ideal from an aerodynamic perspective cannot be used and have to be blunted. A hot-structure and thermal protection system design philosophy which affects both the aerodynamic and structural design of the aircraft needs to be followed. Cooling and thermal protection systems can also have a significant contribution to the mass of the aircraft. Parts of the structure, such as leading edges and the engine, might need to be cooled by circulating cryogenic fuel before burning it. The XB-70 Valkyrie, developed in the 1960s, was an example of such an application with parts of the structure cooled by circulating fuel. Moreover, thermal protection system features such as blunt leading edges tend to have a detrimental effect on aerodynamic performance characteristics such as lift to drag ratio and quality of the inlet flow. Decisions on the materials used also need to be made early in the design process due to the influence their selection can have on the overall design. These multidisciplinary considerations can significantly increase the complexity and cost of early design stages and also increase the computational cost of design evaluation methods.
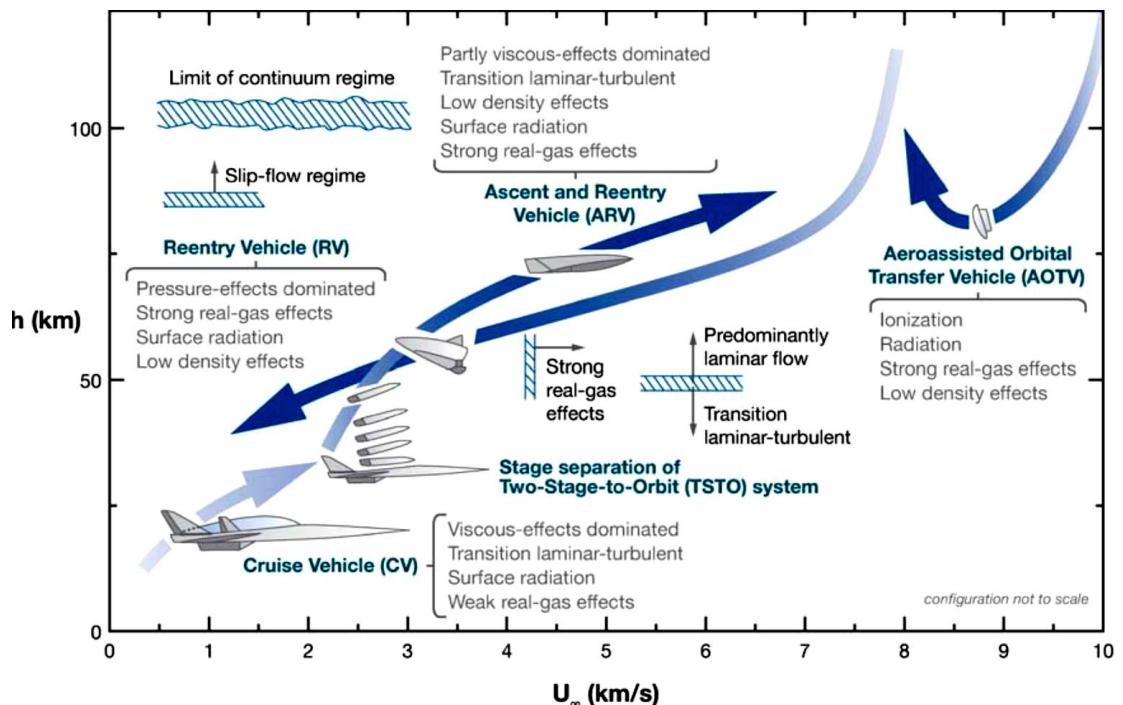
Figure 2.2.   Four distinct classes of hypersonic aircraft with the major hypersonic aerodynamic effects they experience (Bertin, 1994).

Another significant difference between hypersonic and subsonic aircraft design is the level of integration between the airframe and the engine, with two major reasons behind this. The first is that due to the high velocities and altitudes that constitute the air-breathing hypersonic flight corridor, the thrust requirement and relative size of the engine components increases. The engine flow-path forms a substantial part of the airframe, with inlet compression usually starting at the leading edge of the forebody and the nozzle ending at the aft end. This can be seen, for example, in the X-43 experimental aircraft, illustrated in Figure 2.3. The second reason is that the engine inlet and additional components must be close to the body and within the shock layer formed by the forebody to avoid shock interactions. Shock interference and impingement are difficult to predict and can even have catastrophic effects as was observed in one of the X-15 test flights (Watts, 1968). This requires a coupled design approach that generally increases the dimensionality and further contributes to the multidisciplinary nature of the design problem. The propulsion system, for example, can have a strong influence on lift and pitching moments, rendering multidisciplinary optimization approaches necessary for effectively configuring hypersonic vehicles (Bowcutt, 2001). The design process proposed by McClinton (2008), which was adopted for the design of the LAPCAT 2 MR2 aircraft (Langener *et al.*, 2012), incorporated nose-to-tail propulsion analysis in all design stages, a unique requirement of designing hypersonic air-breathing systems. Additionally, although still under active research and development, the operation of hypersonic aircraft with scramjet engines in particular relies in very tight margins. Unlike, for example, the very early and roughly put together jet engines that, from their early years, surpassed the performance of the most advanced at the time propeller propulsion systems, getting a scramjet engine to consistently work and produce adequate amounts of thrust has so far required significant development effort. McClinton (2008) emphasizes the need for coupled propulsion system and airframe optimization due to the small excess thrust generated by scramjet engines, which makes balancing the thrust and drag for a hypersonic aircraft much more challenging than for a conventional aircraft.

Figure 2.3.   The X-43 hypersonic aircraft (NASA graphic).

The highly coupled design of hypersonic aircraft is further enhanced by the fact that, throughout aircraft design, as flight velocities increase the wing sweep also increases. This trend

reaches a point where the wings are blended with the body, with blended wing-body lifting shapes being the optimal configurations for generating sufficient amounts of lift and favourable lift to drag ratios at hypersonic velocities. Hypersonic aircraft also show significant differences to subsonic and low-supersonic designs when it comes to their stability and controllability characteristics, with substantial research effort being focused in this area (Fidan, 2003; Xu, 2015). This is both due to inherent instabilities that require fly-by-wire systems for control at hypersonic velocities and due to the large range of velocities that a hypersonic vehicle that needs to land/take-off will experience. From a geometry perspective, so far traditional rudders and elevators/ailerons have been used in many hypersonic vehicle designs. The space shuttle also utilized a body flap under the engine nozzles for better pitch control during re-entry.

Hypersonic aircraft design is still in its early stages, with only a few experimental aircraft having taken flight. As far as the geometries and configurations of hypersonic aircraft go, one can observe a wide variety of them being proposed in the literature. A few examples of concept hypersonic aircraft and spaceplanes can be seen in Figures 2.4-2.6. The overall challenges and issues hypersonic aircraft design faces are much more complex as Sziroczak (2016) describes in a recent review, ranging from operational issues to aerothermodynamic, propulsion and structural challenges. The issues discussed here, however, are the main items that affect the outer mould line definition and aerodynamic aspect of the design process.
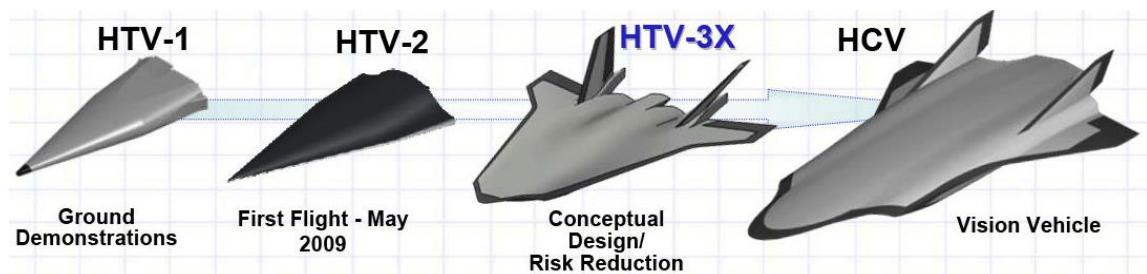
Figure 2.4. Conceptual geometries of the Falcon project (DARPA Falcon project).
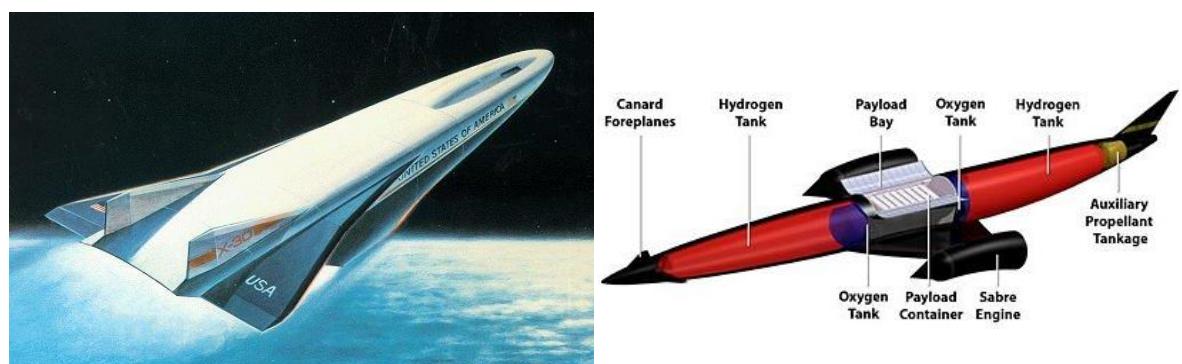
Figure 2.5. Artist's concept of X-30 NASP (left), and the Skylon SSTO concept (right, Reaction Engines Ltd.).
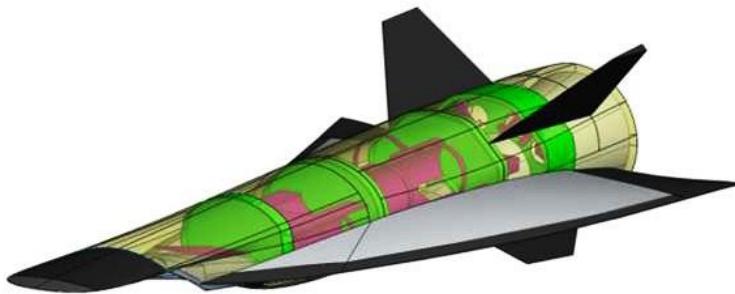
Figure 2.6.   HyShot VII concept (HIFiRE 8).

2.2   Waveriders

Among the most promising aircraft shapes for hypersonic flight are waveriders. They are blended wing-body geometries configured to keep the shockwave generated by the oncoming hypersonic flow attached to the entire leading edge of the lifting body, as illustrated in Figure 2.7. This way the high-pressure region of the shockwave is 'trapped' under the body while being separated from the flow in the upper surface region, resulting in increased lift characteristics that lead to potential lift to drag ratios that cannot be matched by conventional designs.
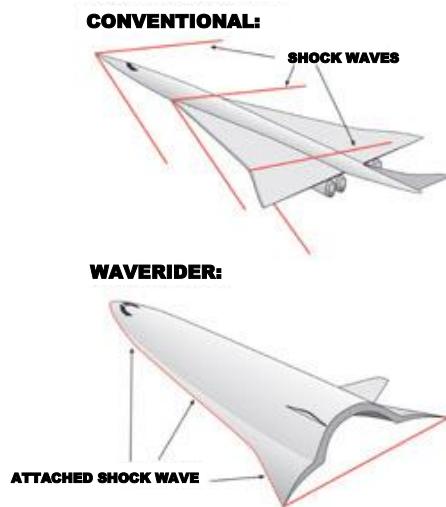
Figure 2.7.   Conventional versus waverider configuration, Lobbia (2004).

As a concept they were introduced by Nonweiler (1959) as a novel high lift geometry/configuration for winged re-entry vehicles in order to tackle the heating issues that space vehicles of the time were experiencing during re-entry. This would be achieved by using vehicles able to generate lift at high altitudes where air density is very low, enabling them to perform part of the deceleration and re-entry process there. The resulting peak heating rates would be much lower than those of ballistic blunt shaped vehicles. Nonweiler derived various waverider shapes by adopting an inverse design method utilizing two-dimensional inviscid oblique shock analytical

solutions for supersonic flow past a wedge. His initial designs were delta/caret shaped configurations based on planar shock flowfields. In the design method, a leading edge is first prescribed on the shockwave and from there the lower surface is constructed by stream-tracing the flow right behind the shock. The upper surface is designed to be parallel to the oncoming flow so that no shockwave would form there and the pressure would be relatively low, as seen in Figure 2.8. A study later done by Squire (1971) confirmed that such 'caret' (from the symbol '^') winged shapes generate substantially more lift when compared to flat winged designs.
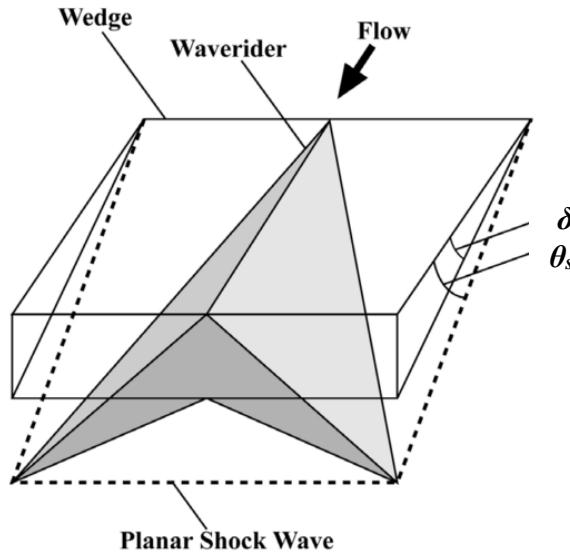
Figure 2.8.   Nonweiler's caret shaped waverider based on a planar shock (Lobbia, 2004).

2.2.1   Base Flowfields for Inverse Design

Stream tracing a shock containing, inviscid flowfield is the most common inverse design approach that has been followed to generate waverider geometries, illustrated in Figure 2.9. The leading edge shape on the selected shockwave is usually defined by designing the projection of the leading edge on the base plane. The base plane is the plane normal to the freestream direction at the aft end of the waverider geometry. The wave-riding lower surface of the geometry is generated by tracing the streamlines of the flowfield from the shockwave to the base plane. The basic flowfield and shape of the shockwave is of major importance to the shape and performance of the waverider. Various flowfields and flowfield generating bodies have been considered by researchers over the years in order to improve the characteristics of waverider geometries. Such inverse design methods can be computationally inexpensive and also guarantee that the shock will lie on the leading edge of the blended wing-body shapes.
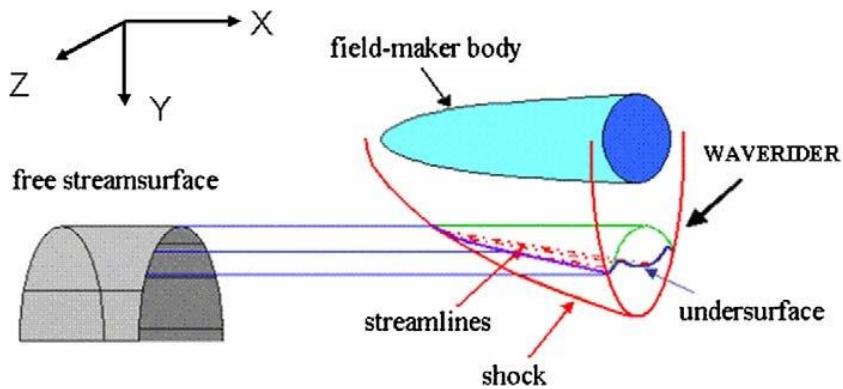
Figure 2.9.   Waverider inverse design (Mangin, 2006).

Two dimensional and three dimensional axisymmetric flowfields can be rapidly calculated with analytical methods, the method of characteristics and other computationally inexpensive approaches. This has rendered them ideal for use with inverse waverider design methods. Waveriders were initially introduced as 'caret' or V-shaped wings by Nonweiler. Other types of waveriders based on two dimensional wedge flowfields have been constructed using, for example, power-law curves for defining the projection of the leading edge on the base plane (Mazhul, 2004). Mazhul (2010) has also investigated waverider design using two dimensional isentropic compression curved wedge flowfields.

Conical shock based supersonic flowfields for waverider design were initially proposed by Jones (1963), illustrated in Figure 2.10. They can be rapidly obtained by solving the Taylor-Maccoll equations (1933) for supersonic inviscid flow around a cone, and have thus been particularly popular. The cone-derived waveriders primarily demonstrated superior volumetric efficiency when compared to wedge derived ones, low volumetric efficiencies being a significant issue for the latter. Corda (1988) later utilized minimum drag pointed power-law axisymmetric bodies at zero angle of attack to generate the basic flowfields for waverider design. The resulting viscous optimized power-law derived waveriders were able to match or outperform cone-derived ones in lift-to-drag comparisons. Axisymmetric bodies of various shapes have since been used to derive waveriders. Mangin (2006) used blunt power-law bodies also achieving slightly higher lift-to-drag ratios than cone-derived designs. Curved cones that further isentropically compress the flow and von Karman ogives (Ding, 2015) have also been considered. Goonko (2000) considered another type of axisymmetric flowfield for the inverse design of waverider forebodies: converging supersonic flow inside constricting ducts, illustrated in Figure 2.11. He reported that the geometries exhibited higher lift coefficients than those derived from planar and conical flowfields although their volumetric efficiency might have suffered.
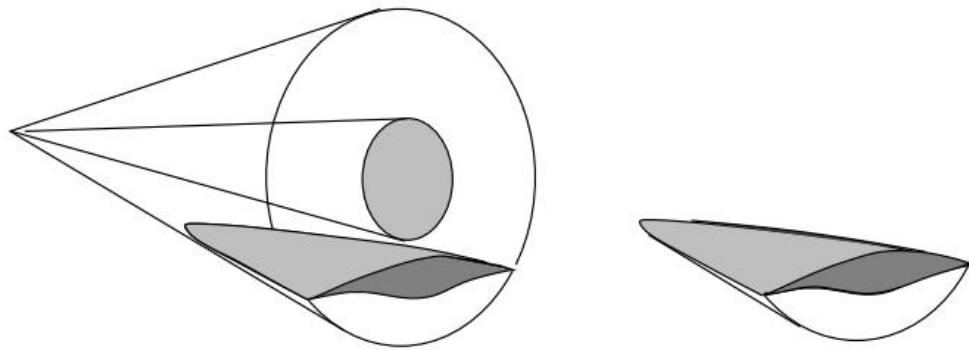
Figure 2.10. Conical-derived waverider design overview (Kinsey, 1998).
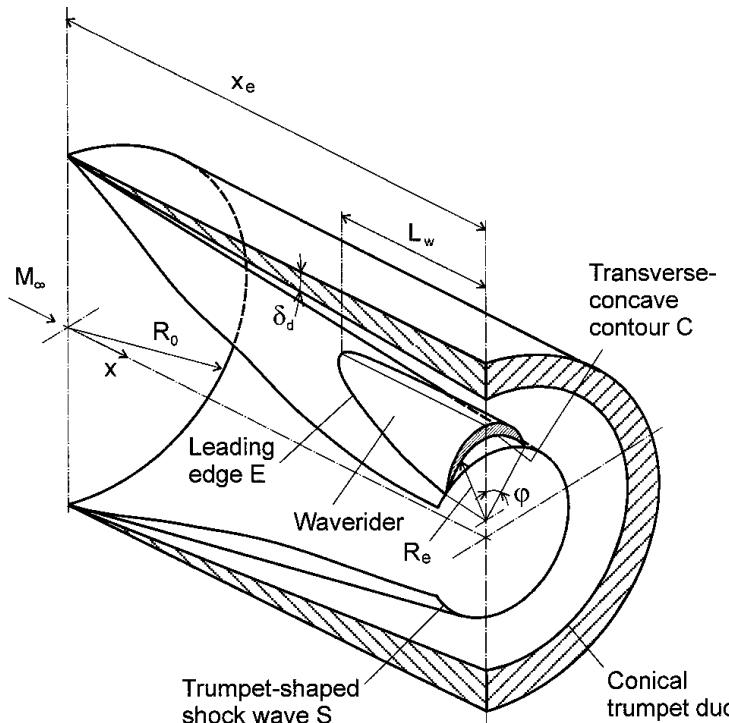
Figure 2.11. Converging flow derived waverider forebody design method by Goonko (2000).

More complex three-dimensional flowfields were first used by Rasmussen (1980) who employed elliptic and inclined circular cones to generate waverider geometries utilizing the hypersonic small-disturbance theory. These were an improvement over simple cone-derived designs and exhibited better volumetric efficiencies and reduced wave drag. Rasmussen (1982) and Jischke (1983) went on to experimentally investigate the aerodynamic features and performance characteristics of such geometries. Takashima and Lewis (1994; 1995) from the University of Maryland, used combined wedge-cone flowfield generating bodies for waverider design, seen in Figure 2.12. These geometries demonstrated a good balance of volumetric efficiencies and lift-to-drag ratios. Additionally, they offered a region of uniform wedge-like flow for more straightforward

engine inlet integration while maintaining the advantages of cone-derived waveriders. Advances in computational resources rendered numerically calculating the supersonic flowfield with Euler CFD simulations affordable, allowing any shock generating body to be considered while keeping the same principles of inverse design seen in Figure 2.9. Given the extent to which the final geometry depends on the shockwave generating body in such design methods, a very wide variety of shapes have been considered with a good example being the work of Cui *et al.* (2007) that incorporated very complex bodies such as quadrate cones, floriated cones and cross cones, illustrated in Figure 2.13. In their study they still obtained the best aerodynamic efficiency from an elliptical cone's flowfield. More sophisticated chemically reacting and inviscid flowfields have also been considered for waverider design. According to Anderson (1992), assuming a frozen chemistry for all but re-entry Mach numbers in waverider design is reasonable.
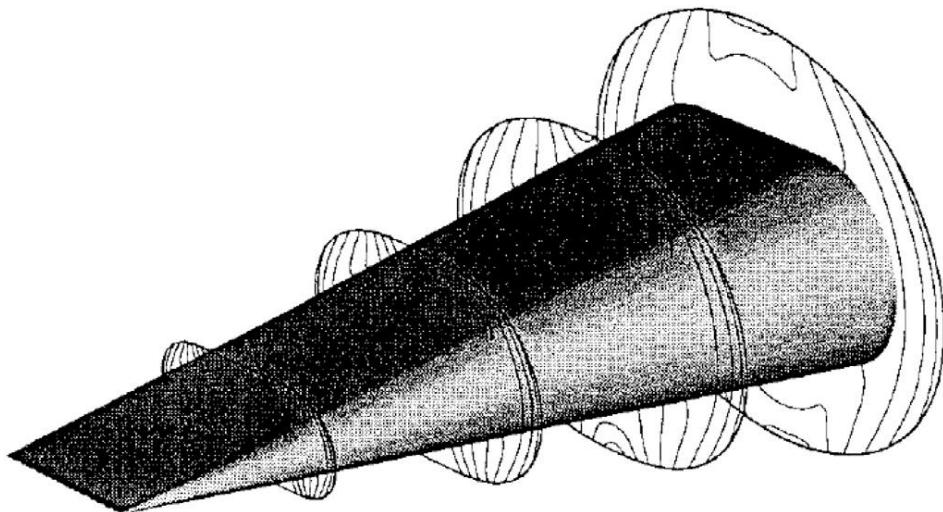
Figure 2.12. Wedge-cone flowfield generating body for waverider design (Takashima & Lewis, 1994).
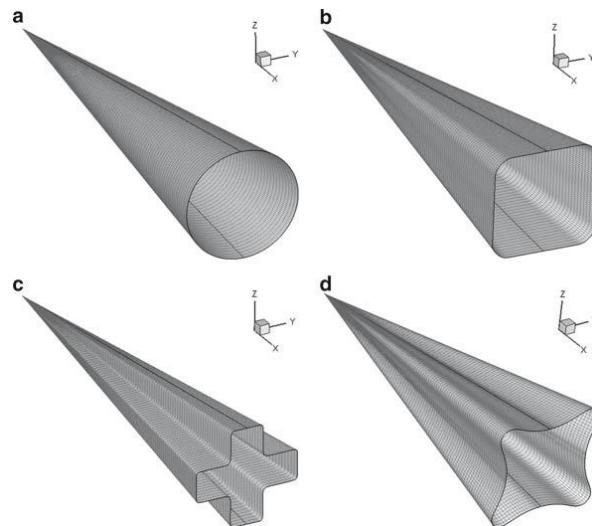
Figure 2.13. Conical shock generating bodies considered by Cui *et al.* (2007), a) circle cone, b) quadrate cone, c) cross cone, and d) floriated cone.

Inverse design methods also have the advantage of providing a fully defined flowfield for inlet integration purposes. Wedge derived waveriders are characterized by a uniform two-dimensional flowfield on the underside that is favorable for an engine intake, however their low volumetric efficiency and the fact that other shapes could outperform such designs renders them non-practical. Designs based on axisymmetric shock generating bodies such as cones exhibit a higher volumetric efficiency and improved performance characteristics. However, the flow on the underside of such a waverider is three-dimensional making the potential integration of an air breathing engine inlet more difficult. A more flexible approach that tackles this issue and gives much more control over the design is the osculating cones method proposed by Sobieczky (1990). In this method the three dimensional flowfield is the result of an ensemble of axisymmetric flowfields applied on osculating planes (or slices) in the waverider's general spanwise direction. The shock of this inverse design method can be more freely shaped. Uniform two-dimensional flow can be obtained by having no curvature of the shock in selected areas, while the favorable characteristics of using an axisymmetric shock generating body can be utilized in others. This method is a departure from the previous methods that focused on defining the shape of a single body to generate the flowfield and, although not an exact method of inverse design, its flexibility and efficiency have made it very popular and it has been widely used. The method is explained in detail in Appendix A. Chauffour and Lewis (2004; 2005) from the University of Maryland, studied a more accurate formulation that could account for the effect of pressure gradients on the cross flow direction. The resulting geometries were still fairly close to the initial method, further verifying the assumption that crossflow due to pressure gradients between the osculating planes is limited. The method has also been experimentally validated (Mill, 1998).

Further extensions to the osculating cones method that allow more than just conical flowfields to be used on each osculating plane are the osculating axisymmetric flow method for waverider design (Sobieczky, 1997) and the osculating flowfield method (Rodi, 2005; 2007). For the first method, the flowfield used on the osculating planes of the design can be different to a cone-based one and calculated with the MOC. The osculating flowfield method offers even more flexibility, allowing the flowfield to change between the osculating planes of the design. Rodi utilized power-law axisymmetric bodies instead of cones for calculating the flow on each osculating plane. The local shape could be controlled to more efficiently compress the flow on the planes where the engine intake would be located while the flowfield of the remaining osculating planes could be configured for higher aerodynamic efficiency and performance. The potential advantages of the design method can be summarized as: (a) improved streamwise lift distribution offering reduced trim drag, (b) improved volumetric efficiency for equivalent aerodynamic performance, c) able to shift the center of mass position more forward than with the osculating cones method

expanding the design space from a stability perspective, (d) potential reduced viscous drag through delayed boundary-layer transition, and (e) an improved aft body close-out geometry. You (2009) proposed a further extension to the osculating flowfield method that also allowed converging flow and concave regions of the shock shape on the osculating planes of the design, something that could be manipulated to integrate inward turning inlets. The osculating cones method and its extensions are able to cover a large design space and are also able to replicate or, at least, get very close to geometries provided by the majority of inverse design methods discussed earlier. This has rendered them very popular for waverider design.

2.2.2   Waverider Adoption

Early work on hypersonic vehicle design, and especially waveriders, was faced with a number of insurmountable, for the time, problems when it became apparent that viscous effects played an important role and needed to be considered early in the design process. This was one of the reasons that contributed in subduing the initial hype and expectations for this class of geometries until Bowcutt *et al.* (1987) introduced a family of viscous optimized waveriders, seen in Figure 2.14. A semi-empirical approach was utilized to include skin friction forces in the calculations. For a given Mach number a family of optimum waverider geometries could be obtained by varying the shock angle. A comparison of the optimum designs for each shock angle for a Mach number of 6 is presented in Figure 2.15. The higher L/D was observed for the waverider designed with a shock angle of 12°. Shock angles higher than the optimum exhibit a larger portion of wave drag while for lower shock angles the portion of skin friction drag greatly increases. This optimization technique could be applied to optimize any aspect of the design.
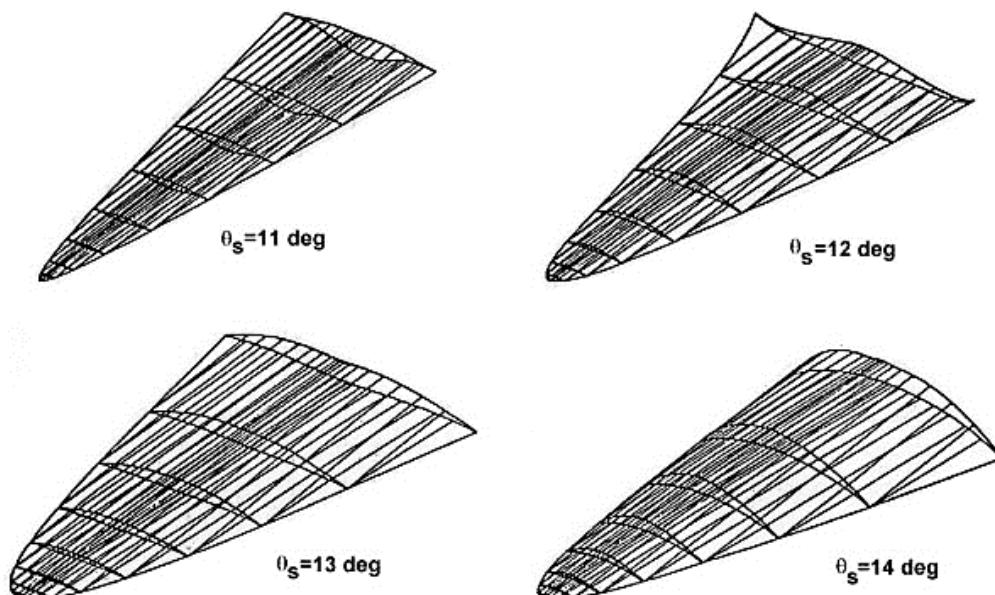
Figure 2.14. Family of Mach 6 viscous optimized waveriders, Bowcutt (1987).
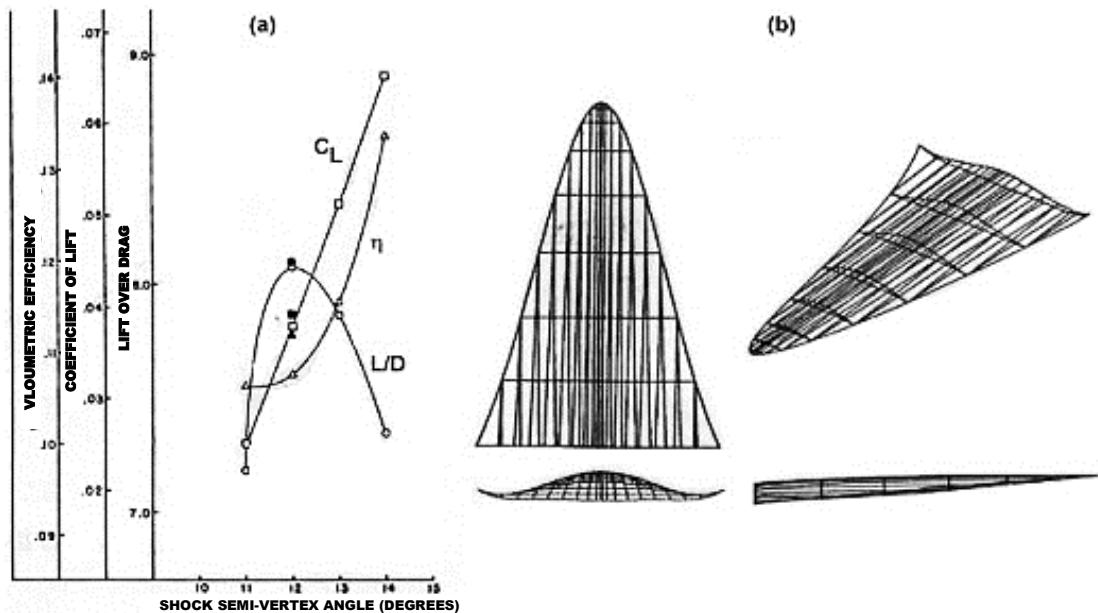
Figure 2.15. (a) Optimum waverider L/D for different shock angles. (b) The overall best Mach 6 design for a shock angle of 12^0, Bowcutt (1987).

The results were very meaningful when considering that the geometries seemed to perform better than the empirical L/D barrier that was proposed by Kuchemann (1978) for conventional aircraft. Following further work by Corda and Anderson (1988) on viscous optimized axisymmetric flowfield derived geometries, Anderson redefined the empirical L/D barrier for waveriders, as seen in Figure 2.16.
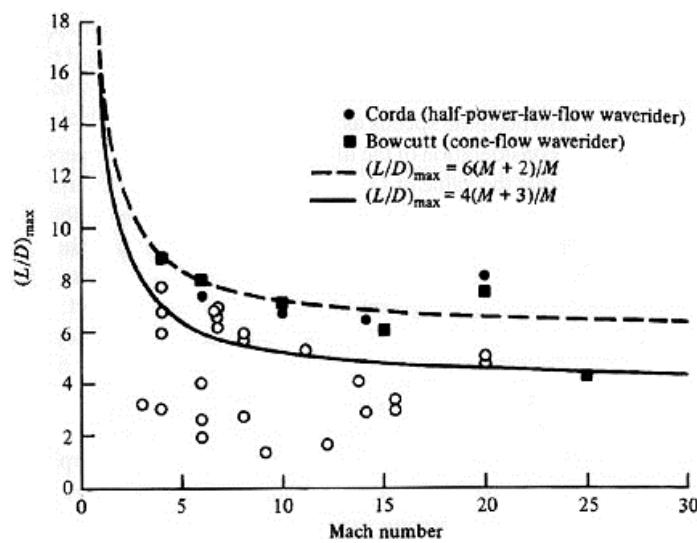
Figure 2.16. The empirical L/D barrier for conventional designs and waveriders (dashed) along with example design points (open: conventional, solid: waveriders), Anderson (2000).

Due to the nature of the design method these geometries were optimized for a single design point, and it was believed that any deviation from the design Mach number and angle of attack could substantially affect the aerodynamic performance. A number of investigations by He and

Rasmussen (1994) and Eggers *et al.* (1995) showed that the off-design aerodynamic characteristics of waverider geometries does not always substantially degrade, while the L/D of certain waverider shapes can even increase for lower than the design point Mach numbers and slightly off-design angles of attack. Rodi (2012) has also investigated waverider shapes designed to maximize vortex lift at various Mach numbers by controlling the sweep of the leading edge. The off-design performance of waverider shapes is therefore something that can also be optimized.

There is a great amount of literature focused on purely optimizing the conceptual design stage aerodynamic performance of idealized waverider geometries, such as in the works of Corda (1988), Lin (1995) and Mangin (2006). Such optimization studies were many times conducted to showcase the capabilities of proposed design methods. However, as the birth of a realistic and viable hypersonic air-breathing vehicle seems to draw near, more realistic approaches on vehicle design are becoming of greater interest.

All of the aforementioned design methods and their results are for ideal waveriders with sharp leading edges that can keep the shock attached and the high pressure flow at the lower surface isolated, preventing leakage to the upper surface. However, due to the nature of the aerodynamic heating effects in hypersonic flows, it would be impossible to adopt a sharp leading edge design. A sharp object made from the most exotic of materials available today would still be unable to withstand sustained hypersonic flow without degrading. With the blunting of leading edges that is required, drag generally increases, and in the case of waveriders a substantial reduction in aerodynamic efficiency due to additional loss of lift can be observed. This is due to the fact that the resulting shockwave will not lie were the ideal inverse design method predicts and there will be leakage from the high-pressure flow on the underside to the upper part (Blosser *et al.*, 1995). Moreover, blunt leading edges result in thicker boundary layers that affect the engine inlet's operation, as well as the overall performance of the vehicle. The experimental results of Gillum and Lewis (1997) on a blunt leading edge Mach 14 waverider showed a maximum L/D value of 3.7 while the theoretical calculations predicted one of 4.61. This was attributed to increased drag and the loss of lift due to the aforementioned leakage from the underside. Among the shapes considered for blunt leading edges, apart from circular arcs, power law curves have been used (O'Brien and Lewis, 1998 and 1999; Santos, 2005 and 2010) while more flexible shapes utilizing fourth order Bézier curves have also been proposed by Rodi (2013) and further investigated by Hinman (2017). Additionally, in receptivity and turbulent transition numerical studies and experiments, ellipses and super-ellipses have also been considered for the leading edges of flat plates. Super-ellipses are modified ellipses that maintain second order continuity at the interface with the flat plate geometry (Lin *et al.*, 1992), which is shown to affect the boundary layer's receptivity to freestream

disturbances. The literature on more sophisticated shapes for the design of blunt leading edges is otherwise limited.

Another issue that rendered most of the early designs unrealistic was the low volumetric efficiency of waverider shapes, because there was not enough space for all of the subsystems required to operate the vehicle to be fitted. The size of many of these subsystems has, however, dramatically decreased over the years. Aerodynamic efficiency, volumetric efficiency and thermal management form three of the primary design objectives in waverider-based hypersonic aircraft design. In the aerodynamic design process most of these objectives are competing and, for example, trying to improve volumetric efficiency will, in general, result in worse aerodynamic performance. As with many other engineering applications, the optimal design will come down to a trade-off between a number of such competing parameters.

A wide variety of conceptual designs and configurations based on waverider geometries have been proposed, ranging from hypersonic cruise missiles and aircraft, to aero-gravity-assist capable vehicles (Lewis, 1991). The number of real world applications on the other hand has been fairly limited up to now. One of the first aircraft to utilize a waverider-like concept was the XB-70 Valkyrie in the late 1950s. Its role would be that of a deep penetration strategic bomber capable of sustained Mach 3 flight at 70,000 feet. The lift characteristics of the vehicle were greatly enhanced by a wing design that utilized compression lift, although it was not inversely designed. The X-51 'Waverider' is another experimental aircraft first flown in 2010, mainly to test scramjet engines and various aspects of hypersonic flight. Designed for speeds of Mach 5+, it incorporated waverider concepts in its design and can be considered as a glimpse into the future of hypersonic flight.

In building a hypersonic aircraft parametric geometry model, inversely designed waverider geometries can be considered as the main component that needs to be further manipulated and interfaced with additional components to complete the aircraft shape. Its role can range from forming a substantial part of the geometry, similar to the design of the XB-70 or the LAPCAT MR2 concept, to being a small part of it as seen in the X-51. To complete the waverider based aircraft geometries more conventional computational geometry tools are usually employed, which will be briefly reviewed in the next section.

Figure 2.17. The XB-70 Valkyrie in flight (NASA photo, 1968).
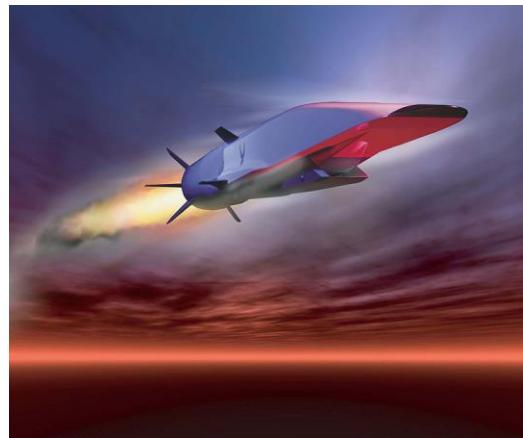
Figure 2.18. Artist's impression of the X-51 Waverider in flight (USAF graphic).

2.3   Computational Geometry for Engineering Applications

One could say that geometry is the very essence of aerodynamic design. It is where one starts and what one is looking for in the end. In the last few decades the tools available to a designer have drastically changed. Drawing boards have been replaced with computer screens where designers can draw and navigate around full three-dimensional geometries, something that could soon be replaced with designers being immersed in virtual or augmented reality environments with their creations within hand's reach.

A wide range of techniques and mathematical formulations have been developed to represent geometry in the computer graphics and CAD fields. Computational geometry is no longer considered to be a tool to merely 'draw' 3D objects, but is essential in enabling computer aided engineering (CAE) and manufacturing (CAM). Moreover, with the widespread applicability of computational design optimization methods that started with the works of Hicks and Henne (1978), focus has turned to parametric geometry representations where the geometric shape is controlled by a number of design variables. Such parametric geometry models can be utilized to automate the process of improving a design. It is, therefore, natural that developments in computational geometry tools have often been driven by the aerospace and automotive industries themselves. Parametric geometries have matured a lot since the National Advisory Committee for Aeronautics' (NACA) first parametric wing shapes in the 1930s, and in this section we will review a number of commonly employed methods to control the shape of curves and surfaces.

The desirable features of a good geometric representation technique from an aircraft design perspective as listed by Kulfan (2006), are as follows:

- Able to produce smooth and realistic shapes.
- Efficient and numerically stable mathematical formulation.

- Able to represent a variety of shapes without using an excessive number of design variables.
- Allow direct control of key design parameters.
- Provide easy control for designing and editing the shape of a curve.
- Have an intuitive geometric interpretation.
- Be systematic and consistent in the way different geometries are represented.
- Robustness.

We can classify methods that parametrically control the shape in two distinct categories. First, there are schemes that fully define geometries and how they are controlled through their formulation. These are complete on their own and are not built upon other baseline shapes. The second category of schemes are ones that do utilize an existing baseline geometry and add control over its shape by perturbing it. A couple of examples of such approaches are free-form deformation methods and methods that add bump functions to add flexibility to an existing geometry. Finally, there is a class of optimization heuristics that deal with the numerical geometry directly. The concept of design variables is fundamentally different in this case as every node in the computational mesh can be treated as one. An example is illustrated in Figure 2.19 for a structural optimization case. The equivalent in aerodynamic optimization are adjoint CFD based approaches (Giles, 2000) that can numerically calculate the sensitivity of an objective function to movements of the individual surface grid cells. Adjoint optimization techniques for aerodynamic design in particular are better suited at driving a design to the local optimum once a larger design space has been explored.
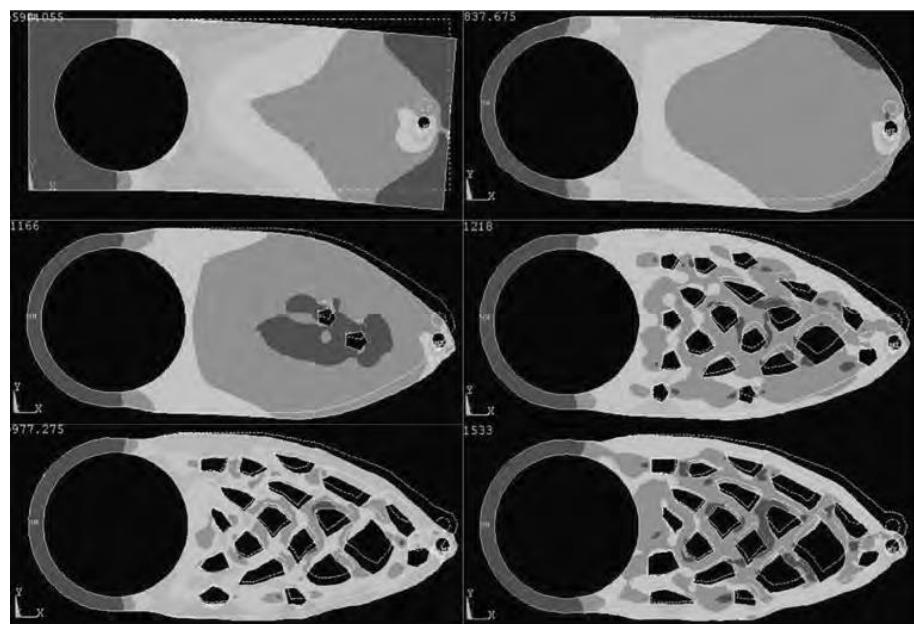
Figure 2.19. Geometrical flexibility at the numerical grid level. Snapshots of the progress (left-to-right, top-to-bottom order) of a structural optimization run, Sóbester (2014).

2.3.1 Bézier curves and splines

Popularized by Pierre Bézier in 1962 through his work as a design engineer with the Renault car company, Bézier curves are parametric curves widely used in computer graphics and related fields due to their intuitive formulation. They are based on Bernstein polynomials and possess properties that make their use very attractive to designers.

Bézier curves are defined by a set of control points, $\mathbf{P_0}$ to $\mathbf{P_n}$, with n being their order. The curve is the path traced by the function $\mathbf{B}(t)$, with t being the parameter of the curve ranging from 0 to 1, equation (2.1). The polygon that results from connecting the control points with lines in the respective order is called the curve's control polygon, shown in Figure 2.20. The basis functions (Bernstein polynomials) for a cubic Bézier curve are shown in Figure 2.21.
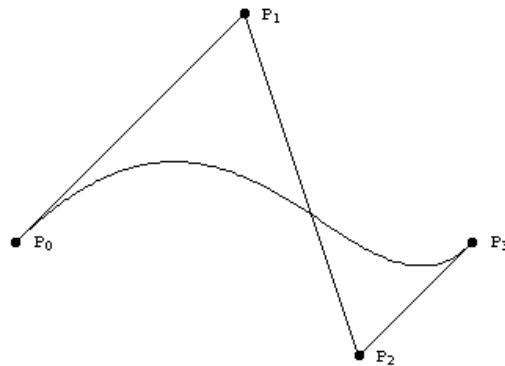
Figure 2.20. Cubic Bézier curve with the control points and control polygon.

$$\mathbf{B}(t) = \sum_{i=0}^{n} \binom{n}{i}(1-t)^{n-i}t^i \mathbf{P}_i \, , \qquad t \in [0,1] \, , \tag{2.1}$$

with $\binom{n}{i}$ the binomial coefficient and $(1-t)^{n-i}t^i$ the Bernstein polynomials.
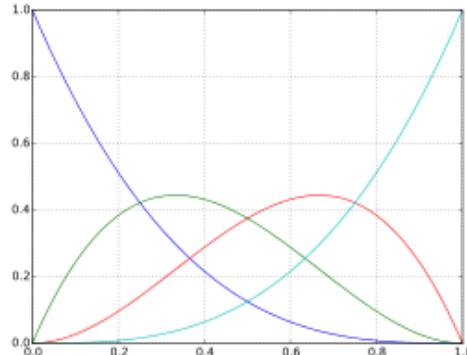
Figure 2.21. The basis functions for a cubic Bézier curve: $y_0 = (1-t)^3$ (blue), $y_1 = 3(1-t)^2\,t$ (green), $y_2 = 3(1-t)\,t^2$ (red), and $y_3 = t^3$ (cyan).

Some of their properties are:

- The end points of the curve are the first and last control points, although the intermediate points generally do not lie on the curve.
- The slope at the start and end of the curve matches that of the first and last sections of the control polygon.
- A curve can be split into more curves that will also be Bézier curves.
- A Bézier curve lies within the convex hull defined by its control polygon.
- The variation diminishing property applies to them.
- Affine transformations applied to the control points are also applied to the curve.
- A curve of degree *n* can be represented exactly by a curve of higher degree.
- The curve's derivatives can be easily obtained.

A formulation that gives more control over the curve with the addition of adjustable weights is the rational Bézier curve, equation (2.2). The value of each weight essentially controls the amount of influence of the respective control point on the curve. These curves are able to represent a wider variety of shapes and exactly represent conic sections.

$$B(t) = \frac{\sum_{i=0}^{n} \binom{n}{i} (1-t)^{n-i} t^i w_i P_i}{\sum_{i=0}^{n} \binom{n}{i} (1-t)^{n-i} t^i w_i} \tag{2.2}$$
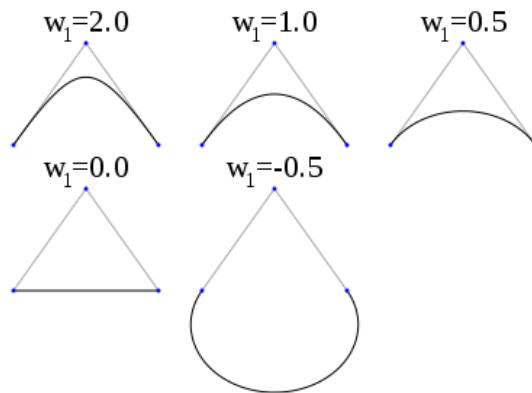
Figure 2.22. Rational quadratic Bézier curve for different values of w_1.

While quadratic and cubic Bézier curves are common, higher order curves are more rarely used. Instead, to represent more complex shapes, lower order curves can be linked together in a piecewise manner in what is called Bézier splines. The process is very straightforward, given two successive curves with control points $P1_i$, *i*=1, 2, … *n1* and $P2_i$, *i*=1, 2, … *n2*, the end point of the first has to coincide with the start point of the second in order for them to meet, so $P1_{n1}=P2_0$.

Smoothness at the interface between the two is achieved with tangency of the respective last and first sections of the control polygon.

The Ferguson spline (Ferguson, 1964) is a similar formulation and essentially a reparameterization of a Bézier curve. What mainly differentiates it is that, while the end control points control the two ends of the curve similarly to the Bézier curve, the end tangents are controlled directly by specifying a tangent vector's direction and magnitude, as seen in Figure 2.23. The Bézier curve tangents at the end points are controlled with the placement of the second and ($n − 1$) control points.
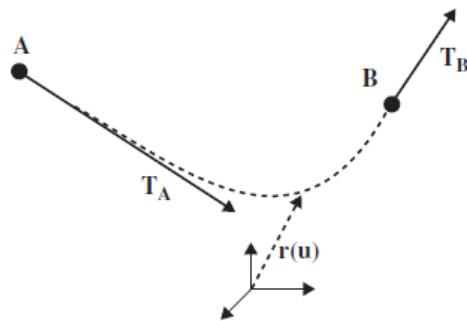
Figure 2.23. Ferguson's spline with control points and tangency vectors.

2.3.2   B-splines and Non-Uniform Rational B-splines

B-splines, with 'B' standing for 'basis', are piecewise polynomial functions made of ($n + 1$) − d segments with ($n + 1$) the number of control points P_i, and d the degree of each segment. Initially investigated by Lobachevsky as early as the nineteenth century, they were named by Schoenberg (de Boor, 1978) who used them for statistical data smoothing. The two main differences from Bézier splines is that the domain in which they are defined is subdivided by knots, and the B-spline basis functions are non-zero on a number of adjacent intervals and not the entire domain of the curve, unlike the Bézier curve's basis functions. This enables local control of the curve, which is not possible with Bézier curves.

The curve is defined over a domain $k_0 \leq u \leq k_{n+1+d}$, with the points where $u = k_i$ known as the knots or break points that define where the segments start and end. The vector $[k_0, k_1, … , k_{n+1+d}]$ is the spline's knot vector, with values of the knots being in ascending order. Unlike Bézier splines, the ends of a B-spline do not coincide with the first and last control points, unless there are multiple control points or knots. A B-spline will pass through a knot multiplied d times. When all internal knots are distinct and equally spaced, the curve is a uniform B-spline that is continuous and also has continuous derivatives up to a $d − 1$ degree. If internal knots are multiplied, the continuity of

the curve is reduced by 1 for every additional knot. B-splines can be defined as weighed sums of their control points similarly to Bézier curves:

$$S(u) = \sum_{i=0}^{n} \boldsymbol{P}_i B_{i,d}(u) \,, \tag{2.3}$$

with the basis functions $B_{i,d}$ recursively calculated:

$$B_{i,1}(u) = \begin{cases} 1, & k_i \leq x < k_{i+1} \\ 0, & otherwise \end{cases} ,$$

$$\tag{2.4}$$

$$B_{i,d}(u) = \frac{u - k_i}{k_{i+d-1} - k_i} B_{i,d-1}(u) + \frac{k_{i+d} - u}{k_{i+d} - k_{i+1}} B_{i+1,d-1}(u) \,.$$

By varying the knot vector in a non-uniform manner and also adding adjustable weights as was done for the rational Bézier curves, a much more versatile family of parametric curves is obtained. Those are the non-uniform rational basis splines (NURBS). They are essentially a superset of all the previously described parametric curves that is being widely used in computer graphics and design. They can be expressed in a similar way as a Bézier curve or B-spline:

$$N(u) = \frac{\sum_{i=0}^{n} w_i \boldsymbol{P}_i B_{i,d}(u)}{\sum_{i=0}^{n} w_i B_{i,d}(u)} \,, \tag{2.5}$$

with $B_{i,d}$ the basis functions described earlier.

Such parametric curves are extensively used throughout the current work, initially for parameterizing the design driving curves of the waverider inverse design method and then for designing additional aircraft components and surfaces.

2.3.3   Surfaces

Extending the parametric curves that were described earlier to surfaces can be accomplished in a number of ways, all of which are fairly straightforward. Surfaces are generally described as two variable equations, *S(u,v)*, with each of the two parameters usually used to scan the surface in a

different direction. The simplest of those are lofted surfaces and surfaces generated by translating a curve along a second one.

Lofted surfaces are essentially a weighted sum of two curves and can be described with:

$$L(u,v) = (1-u)B_1(v) + uB_2(v) \, , \tag{2.6}$$

with B_1 and B_2 the two parametric curves that are connected with straight lines to generate the surface.

The surface obtained by translating a curve along another intersecting one is given by:

$$T(u,v) = B_1(u) + B_2(v) - P_{0,0} \, , \tag{2.7}$$

with $a_{0,0}$ the point of intersection of the two curves.

Coons surfaces are a combination of the two previous approaches. They are defined by four boundary curves and the surface is essentially the sum of two lofts between opposite curves, minus the weighted sum of the corner points (BL_{ab}).

$$C(u,v) = L_a(u,v) + L_b(u,v) - BL_{ab}(u,v) \, , \tag{2.8}$$

$$BL_{ab}(u,v) = (1-u, u) \begin{pmatrix} P_{0,0} & P_{0,m} \\ P_{n,0} & P_{n,m} \end{pmatrix} \begin{pmatrix} 1-v \\ v \end{pmatrix} \, , \tag{2.9}$$

with L_a and L_b calculated as before.

Utilizing the parametric curve formulations that were described earlier, it is possible to generate parametric surface representations that offer much more control over how the surface is shaped in-between the boundary curves. These are mainly Bézier, B-spline and NURBS surfaces and

each of them can be obtained as a tensor product of two respective curves. A Bézier surface for example is given by:

$$S_B(u,v) = \sum_{i=0}^{n} \sum_{i=0}^{m} B_i^n(u) B_j^m(v) \boldsymbol{P}_{i,j} .$$

(2.10)

The control points, $\boldsymbol{P}_{i,j}$, now form a rectangular grid. The control points at the edges define curves identical to the curves used as a base for the parametric surface, while the intermediate control points are able to shape the internal parts of it, as illustrated in Figure 2.24.
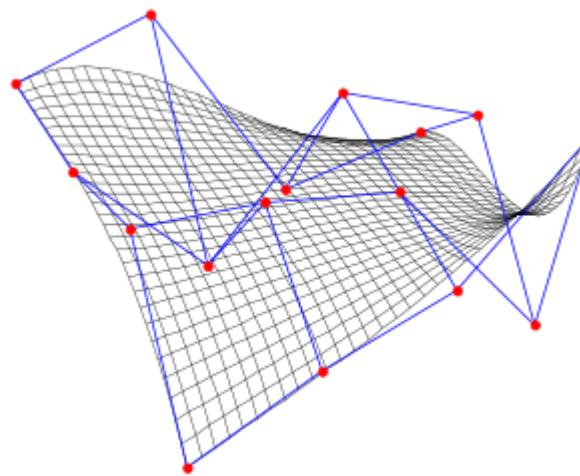
Figure 2.24. Example Bézier surface with its control points and grid/polygon.

While it is preferable to use the inverse design methods to obtain the forebody geometry for a hypersonic aircraft, as will be discussed later, many of the additional geometry components, such as the blunt leading edge, will utilize these tools to complete the aircraft geometry.

2.3.4   Free-Form Deformation

First introduced by Sederberg and Parry (1986) and later extended by Coquillart (1990), free-form deformation (FFD) is a geometric technique that offers control over the shape of existing geometries by manipulating locations of points related to it. The general idea is that an object is enclosed and mapped within a hull that can be easily deformed. The hull's deformations can then be translated to deformations of the object. The deformation of the hull is based on three-dimensional analogs of parametric curves such as the ones described earlier. With this approach, flexibility can be added to existing geometries while it is also an approach that can deal directly with numerical representations of surfaces (such as numerical surface grids), rendering it very powerful.

The Bézier curve derived FFD, using a grid of control points $\textbf{\textit{P}}_{i,j,k}$, is given by:

$$X(u,v,s) = \sum_{i=0}^{m} \sum_{j=0}^{n} \sum_{k=0}^{p} \textbf{\textit{P}}_{i,j,k} f_i(u) g_j(v) h_k(s) \, , \qquad (2.11)$$

with $f_i(u)$, $g_j(v)$, and $h_k(s)$ the Bernstein polynomials.
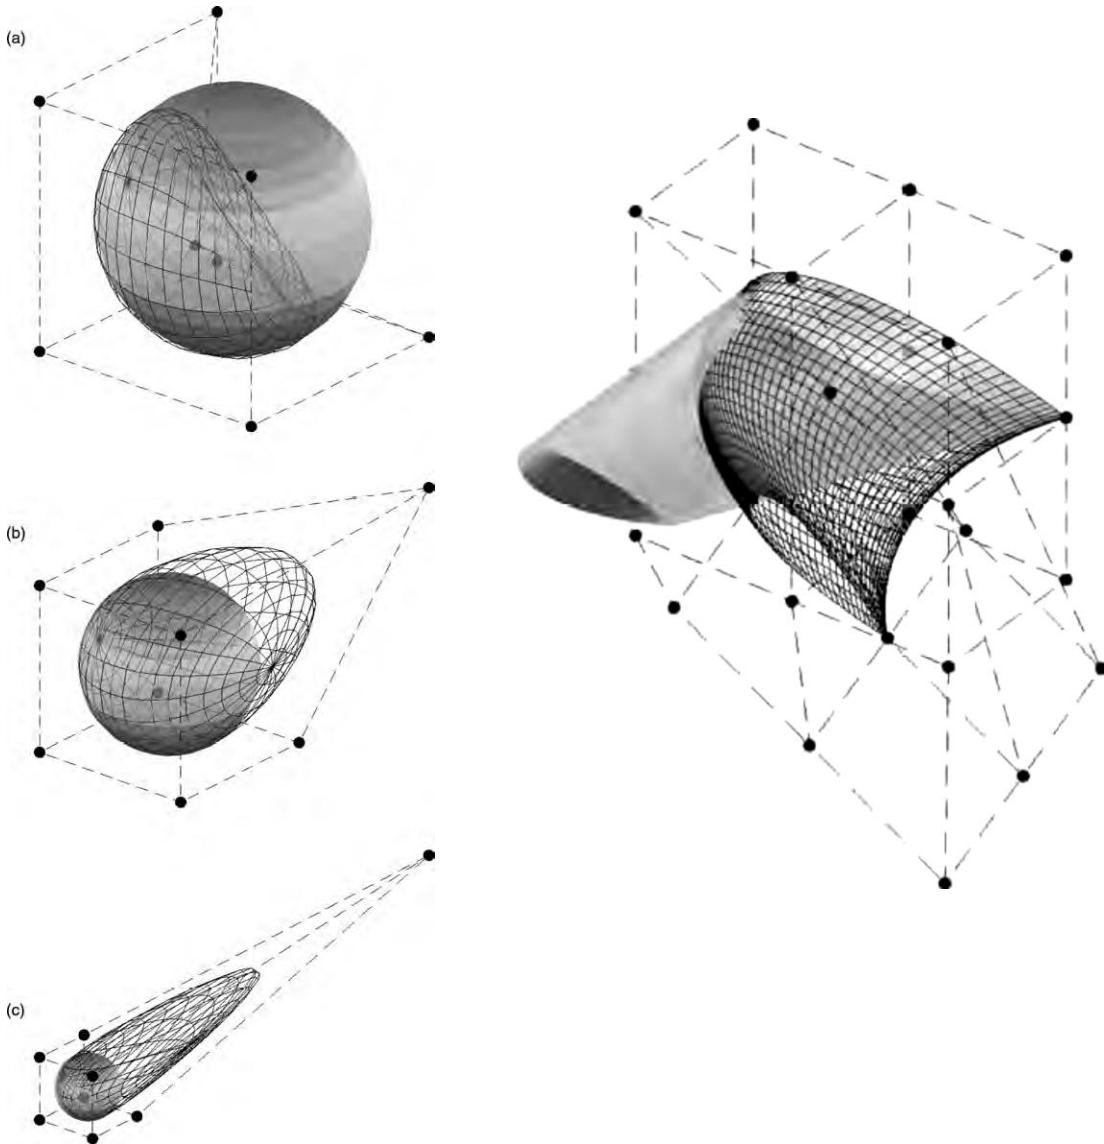
Figure 2.25. Free-form deformation of a sphere using a 2x2x2 gird (left), and of a generic wing shape using a 2x3x4 grid (right), Sóbester (2014).

As far as the current work goes, free-form deformation can be utilized if we wish to further manipulate the inversely designed waverider geometries, especially their upper surface, in order to further improve performance characteristics, for the packaging of internal subsystems and also for interfacing it with additional geometry components.

28

2.4 Parameterization of Waverider-based Hypersonic Aircraft Configurations

The fundamental differences between hypersonic, subsonic and supersonic aircraft designs that were described earlier, affect the classification and integration of different geometry components that compose the aircraft. These can be summarized in the following: As flight speed increases increased sweep angles of the wings are preferable. This trend reaches a point where the wings are fully blended with the fuselage and the aircraft looks more like a lifting body. The size of engine intakes and nozzles increases due to altitude and Mach number requirements, resulting in fully integrated engines that make up a large portion of the aircraft. Hypersonic aircraft design is also limited by the few past proven examples and, relatively, little accumulated design experience. Design methods and parametric geometry formulations for conventional/subsonic aircraft have been extensively studied, and one can find very detailed information in the literature on how to design a parametric fuselage, wing, engine nacelle and the interfacing geometry components for the above. Modern military fighter aircraft design is where these distinct geometries start being blended together, that, however, is many times due to additional design requirements such as super-maneuverability and radar signature. On the other hand, the equivalent literature on hypersonic design and geometry parameterization is much shallower, especially for hypersonic-design-specific geometry components.

The majority of aircraft shapes that have been proposed for hypersonic flight have waverider or waverider-like forebodies and blended wing-body airframes. The waverider component can either be only a small part of the entire airframe, like in the case of the X-51 in Figure 2.18, or a substantial part of it, as seen in the work of Lobbia (2004) in Figure 2.26.
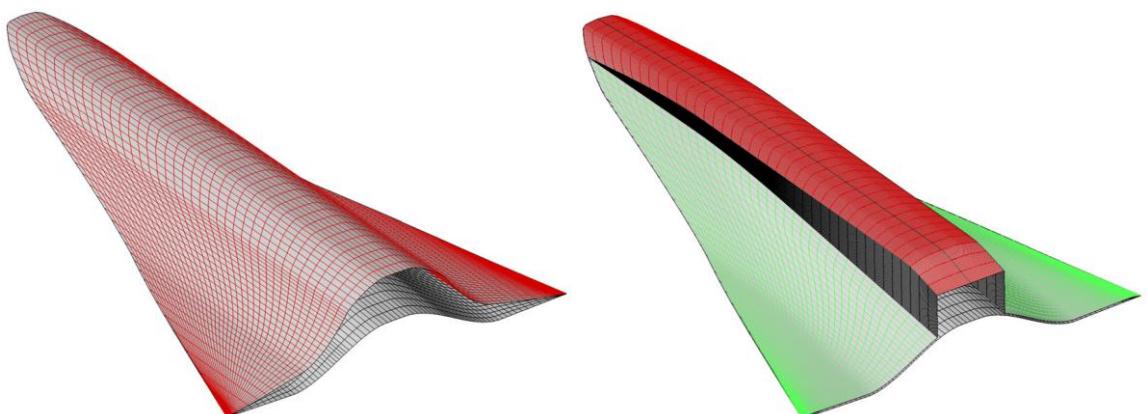
Figure 2.26. Idealized waverider (left) and modified waverider geometry used for the entire airframe (right), Lobbia (2004).

All the methods mentioned in Section 2.2 are useful tools in designing the baseline lifting body shapes suitable for hypersonic flight, where fuselage and wings are no longer distinct. They are however only part of the toolbox, while they are not a standalone tool on their own either, as they require parametric curve formulations for their underlying design-driving curves. To elaborate, most of the inverse design methods mentioned earlier utilize a parametric curve to define the shape of the leading edge on the shockwave from where stream tracing begins. The more sophisticated methods also have one more curve defining the shape of the shockwave or a method to control the shape of the flowfield generating body. These design-driving curves have traditionally been designed on the base plane at the aft end of the geometry. While a broad range of geometry formulations are suitable for this purpose, and different formulations have been used from study to study, rarely have those decisions been analyzed. Especially when it comes to the degrees of freedom and amount of flexibility given to those design-driving curves.

The waverider shapes of Figure 2.26 are parameterized as seen in Figure 2.27, with 4 control points per design-driving curve of the osculating cones method, totaling around 9 degrees of freedom. The geometry model of the entire transport aircraft utilized 24 design variables. In the work of Tian *et al.* (2013), a hypersonic aircraft model parameterized with a total of 227 design variables with the inlet capture curve of the waverider component utilizing 21 control points is employed, part of the seemingly over-parameterized geometry is seen in Figure 2.28. Additional NURBS patches are used for the upper and lower surface resulting in a complete aircraft shape. The geometry is then optimized utilizing a low-fidelity panel method for design evaluations. Such a parametric model should be used in very late design stages, otherwise one can only hope that the initial arrangement of the geometry results in a 'good enough' candidate, and design search algorithms are limited to the simplest of evaluation methods or adjoint optimization techniques.
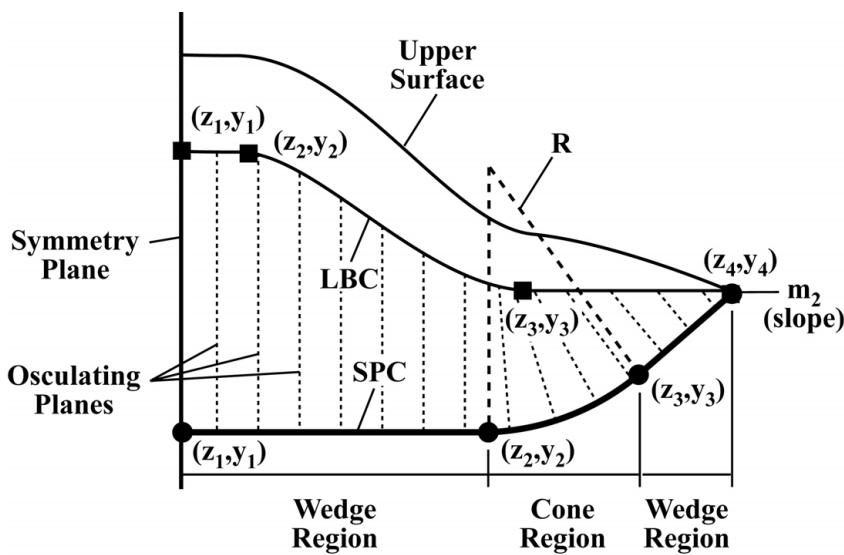
Figure 2.27. Waverider parameterization used by Lobbia (2004).

The two aforementioned examples from the open literature represent two ends of the spectrum, with a great number of studies utilizing similar parametric geometries to conduct optimization investigations into various performance characteristics of waverider based aircraft. Two additional examples are the initial viscous optimized waveriders of Bowcutt (1987) and optimization with off-design considerations by Takashima (1999). Research projects in hypersonic flight that are defense oriented do not have such details on the geometric parameterization openly available. Overall, we observe that a number of the decisions behind the parameterization of the waverider component are usually based on intuition and experience, and this is an area where the current investigations will focus, by attempting to further justify such choices as far as handling the waverider component of the aircraft goes.
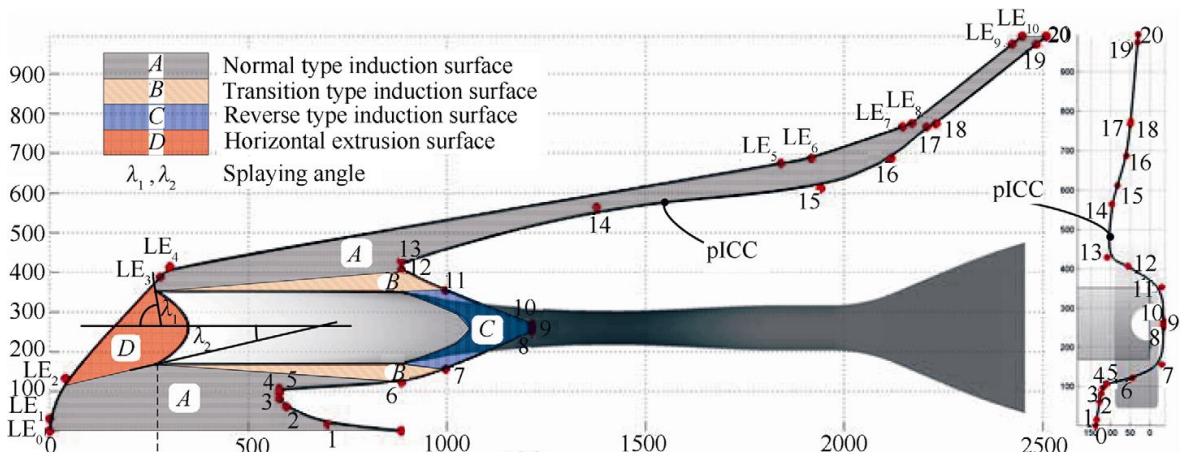
Figure 2.28. Bottom and rear view of waverider parameterization of the inlet integrated waverider by Tien (2013), with control points for the main feature edges of the geometry (LE: leading edge, pICC: inlet capture curve).

Apart from the waverider component, the rest of the toolbox consists of methods to design additional components required to reach more realistic aircraft shapes, usually by employing more conventional parametric geometry formulations like the ones described in section 2.3. Those can be blunt leading edges, control surfaces, scramjet engine components, interfacing geometries, further shaping of the upper surface of the waverider etc. One of the geometry features that is most closely integrated with the waverider forebody is the blunt leading edge. A relatively small number of studies have investigated potential advantages to using more sophisticated shapes for blunt leading edges, specifically for hypersonic aircraft (O'Brien, 1999; Santos, 2005; Rodi, 2013; Hinman, 2017). Improvements in drag and thermal load or in the shock stand-off distance and subsequent leakage to the upper surface could be observed. These become significant mostly for smaller aircraft where the thickness of the blunt leading edge is a scale that is comparable to the rest of the aircraft.

On the other hand, a great number of studies have looked into designing and integrating scramjet engine inlets onto waveriders and hypersonic aircraft in general. Research in this field goes hand in hand with the continued effort to develop operational scramjet engines as there are many challenges in designing inlets able to cope with operation at off design conditions and unstart issues while also being efficient. The X-43 and X-51, seen in Figure 2.18, both employed conventional ramp compression inlets on the underside of the body and conventional control surfaces, a straightforward approach that has been utilized in various design studies. This approach was also utilized by Takashima (1999) with the parameterization of the two-dimensional engine flow-path seen in Figure 2.29. Other concepts, such as the MR2 hypersonic cruiser configuration from the LAPCAT II program (Steelant and Langener, 2008) incorporate dorsal mounted inward turning inlet design that directly faces the freestream, seen in Figure 2.30. Inward-turning inlets utilize axisymmetric flowfields that compress the flow internally, towards the axis of symmetry. Overall, the propulsion system can have a strong impact on lift and pitching moment further complicating the interaction, and matching the two is a complex task requiring significant optimization effort as described by Bowcutt (2001).
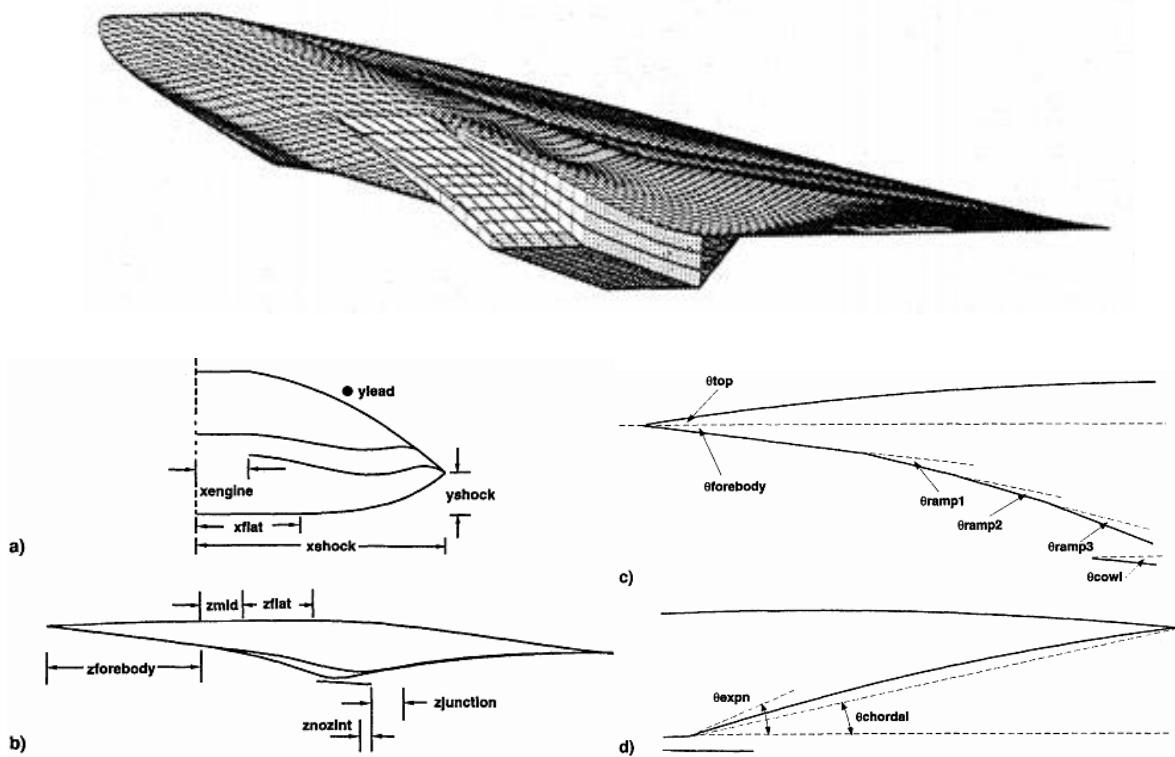
Figure 2.29. Two-dimensional ramp inlet integrated on waverider geometry (top) and parameterization of the scramjet flow-path (bottom), (Takashima, 1999).

Figure 2.30. The MR2 hypersonic cruiser (LAPCAT II).

The parameterization of each of the main components of the aircraft is more coupled than in conventional aircraft and an integrated design approach needs to be followed. Parametric geometry tools developed for such purposes are usually very specialized to the specific configuration that has been chosen, and different configurations affect the coupled parameterization of geometry components that need to work together. A tool that will enable multiple configurations will need to be flexible in the way the geometry is parameterized and interfaced with additional aircraft components.

Finally, a few studies have also focused on developing interactive waverider and hypersonic aircraft design tools, with a graphical user interface for ease of use and optimization capabilities built in. One of the first examples is Center's (1993) Waverider Interactive Parameter Adjustment Routine (WIPAR), seen in Figure 2.31. The software was one of the first waverider design tools to enable rapid generation, modification and optimization of waverider geometries in a graphical environment, and was further used in subsequent studies (Jones, 2002). Szema (2010) also developed a design tool for generating waverider-based hypersonic aircraft geometries with integrated streamtraced inward turning inlets. The waverider geometries were parameterized using a NURBS curve with 20 control points to define the leading edge on the designed shock shape from a planform perspective, and another NURBS curve defining the shape of the aft end of the waverider geometry since not all streamlines need to be traced all the way to the base plane or aft end of the selected shock. The inlet integration method is limited to the method patented by Elvin (2007) with the Busemann-type inlets facing the freestream and integrated along the leading edge of the waverider forebody, illustrated in Figure 2.32.
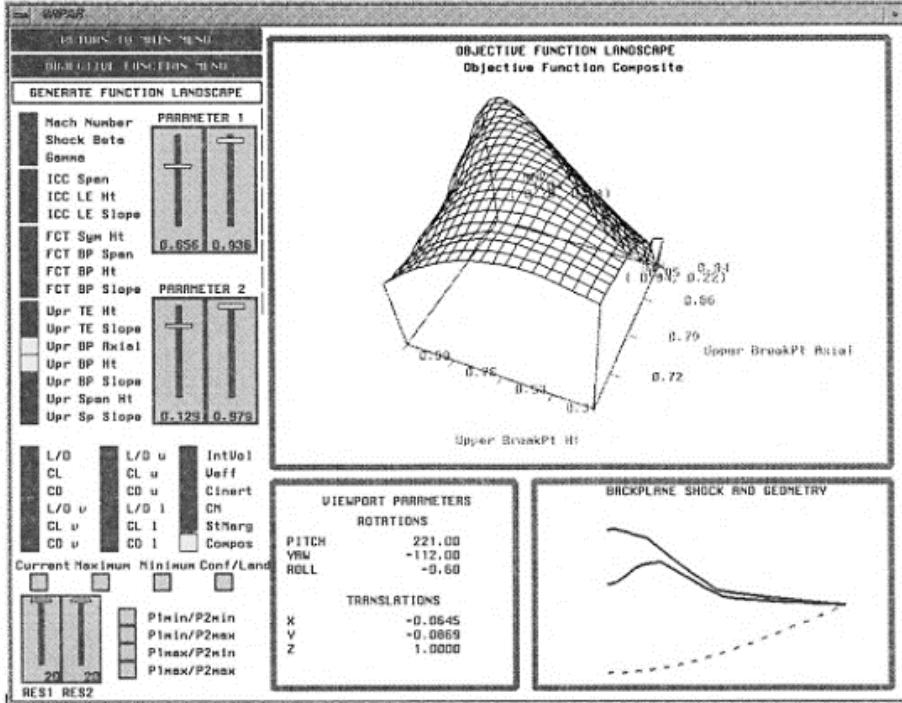
Figure 2.31. The WIPAR interactive waverider design tool (Sobieczky, 2004).
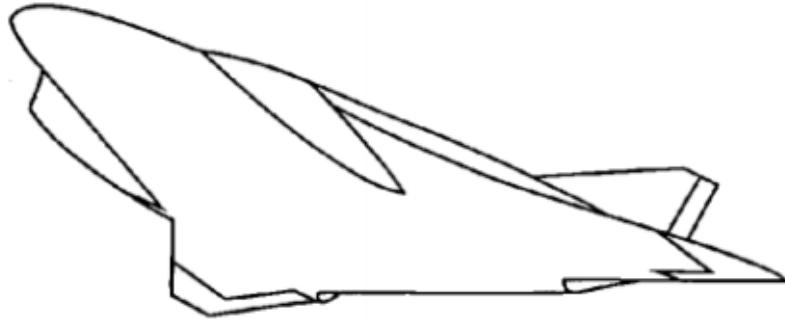
Figure 2.32. Two inward turning inlets integrated along the leading edge of the aircraft (Elvin, 2007).

2.5   Motivation & Approach

The goal for this work is to investigate and establish guidelines on how best to effectively and efficiently parameterize and handle the geometries defining the outer mould lines of hypersonic aircraft, and especially the waverider blended lifting body shapes. Parametric geometry models with those characteristics are required to make higher fidelity aerodynamic design optimization studies, which utilize our constantly evolving numerical and computational analysis capabilities, more affordable. Additionally, the material in this thesis attempts to assist in the transition from the theoretical waverider design methods and concepts to realistic design processes, by understanding and developing parameterization techniques and tools for waverider-

based aircraft. Most of the work is focused on air-breathing hypersonic cruise vehicles, although a substantial part of it also applies for all classes of hypersonic aircraft. The work has been undertaken as part of GHandI (Geometry Handling and Integration), a UK Aerospace Technology Institute project, in collaboration with MBDA UK.

With waveriders being amongst the most promising types of geometries for hypersonic flight, various performance aspects of waverider based hypersonic aircraft geometries have been the subject of optimization studies. Other studies have focused on further developing the geometries by proposing propulsion system integration approaches. However, detailed information and investigations on their parameterization and handling is something that has generally seen limited scientific research effort. Moreover, conceptual design studies and conceptual geometries for hypersonic aircraft still show a relatively large discrepancy in proposed configurations when compared to other aeronautical fields. This requires design studies which explore a larger design space and multiple configurations. It is likely that in a decade or two we will have a very specific image and configuration in our minds when we think of a 'hypersonic air-breathing aircraft' and their parameterization will be more straightforward. Right now, however, we are at the stage where that 'image' is being shaped by researchers around the world, rendering this a very intriguing field to work in. The ultimate goal and hope for this work is for it to be one more brush stroke in painting that image.

The approach that will be followed is that of going through the development of a parametric geometry engine for the design of waverider-based hypersonic aircraft, while addressing the research objectives raised in the introduction section. These will be addressed through: 1) developing new types of parametric geometry formulations and design methods for a number of the aircraft's components, 2) providing a deeper analysis of existing design techniques and approaches with emphasis on parametric model efficiency and robustness, 3) attempting to rationalize the decisions made in developing the parametric geometries, and especially their flexibility and parameterization approach, and 4) implementation-specific considerations in order to enable a revised aerodynamic design process for waverider-based hypersonic aircraft.

The osculating cones waverider generation method will be used as the main tool to obtain the baseline shapes, as it is a method that efficiently enables a reasonable amount of flexibility and can also be seen as a superset of a number of simpler inverse design methods. Focus will be given on the parameterization and requirements of the method's design-driving curves with efficiency and robustness in mind, something that has not been systematically and sufficiently covered in relevant literature, although the method has been around for quite some time. Two aspects of efficiently parameterizing the inverse design method are investigated. First is altering the method

itself in order to enable the geometry to more efficiently meet potential design requirements and readily interface with additional aircraft components, addressing the first research objective set in the introduction. Second is looking into the required flexibility and degrees of freedom of the design driving curves in order to enable various meaningful shapes without over-parameterizing the geometry model, addressing the second research objective. A novel waverider design method that is based on the osculating cones method was also developed. It forms a clear departure from having to design a shock containing flowfield beforehand and allows full control over the three-dimensional shape of the leading edge.

Among the additional geometry features that are better suited to be built into the forebody geometry generation tool is the blunt leading edge. This was identified as an area where a mathematically formal, robust and efficient geometry representation, better suited for more detailed investigations, could be introduced. Forebody and inlet design for hypersonic aircraft go hand in hand, therefore, an inlet design tool that will enable seamless integration and interfacing with the generated forebody geometries will also be developed. This will increase the robustness and enable the geometry engine to seamlessly support the development of a variety of aircraft configurations. The geometry engine focuses on generating the aircraft components that are not straightforward to directly design in modern CAD software, such as the baseline waverider shapes and streamtraced inlets. An example on how the geometries generated by the geometry engine can be further built upon and completed will also be presented, and functionality to export and complete the geometries in the majority of CAD software packages will be available. Finally, we will have a look into how the geometry engine fits into a conceptual aerodynamic design process and how it can be utilized in order to address a number of the additional challenges observed in hypersonic aircraft design.

Chapter 3:   Waverider Forebody Parameterization

The high lift-to-drag ratio that waverider geometries demonstrate at high velocities renders them ideal for hypersonic applications. A waverider shape, as generated by an inverse design method, is however incapable of tackling the envelope of aerothermodynamic challenges presented by hypersonic flight. It does however provide a sensible starting point or baseline geometry, which can be further manipulated and built upon to reach more realistic airframe shapes. Just as utilizing a distribution of circular cross sections to build the fuselage of an airliner is one of the starting points for commercial aircraft design, inverse waverider design methods can be considered as the equivalent for obtaining lifting, blended wing-body airframes that can form the base of a hypersonic aircraft geometry model.

Using parametric surface formulations to directly design such shapes is not a well-posed problem, and would result in a dramatic increase in the dimensionality of the design space. An example of a waverider-like geometry constructed with Bézier patches can be seen in Figure 3.1. The number of design variables, even after applying a number of relations between the control points to restrict the geometry to waverider-like shapes, is at least double than what is needed to adequately parameterize an inverse design algorithm. Additionally, a subsequent optimization algorithm that would utilize this parametric geometry would have a 'hard time' even identifying feasible waverider geometries, let alone finding the best. Inverse design methods have become the standard for waverider design due to the much higher efficiency in generating such shapes, where, in addition to the control given by the parametric design-driving curves, hypersonic flowfield information is woven into the fabric of the geometries.
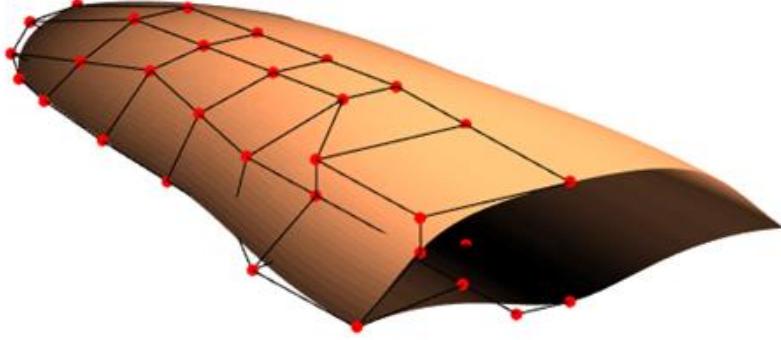
Figure 3.1.   Waverider-like geometry generated using 4 cubic Bézier patches, control points and
             polygon visible.

For the purposes of this work the osculating cones method will be primarily used as it provides a reasonable amount of flexibility in parameterizing the waverider shape. The method, as well as a brief analysis on the requirements and limitations of its design-driving curves, is described in Appendix A (the reader is advised to read the appendix at this point if unfamiliar with the method). Most of the findings however, can also be applied to the method's extensions, the osculating axisymmetric flow and osculating flowfield methods, as well as the majority of inverse design algorithms. In this chapter we turn our focus to the parameterization of the waverider shapes, whose role in the entire airframe can range from being just a small part of the forebody, as is the case with the X-51, to being a substantial part of the entire airframe, as seen in the work of Takashima (1999) and Lobbia (2004), both shown in section 2.4. A sample geometry of a simple waverider-based hypersonic aircraft parametric model that is closer to a more realistic aircraft and incorporates an osculating cones generated waverider forebody is shown in Figure 3.2 for context – in this chapter we focus on handling the geometry of the forebody (highlighted in the figure), and addressing how to efficiently parameterize it both as a standalone component and for an application like this.
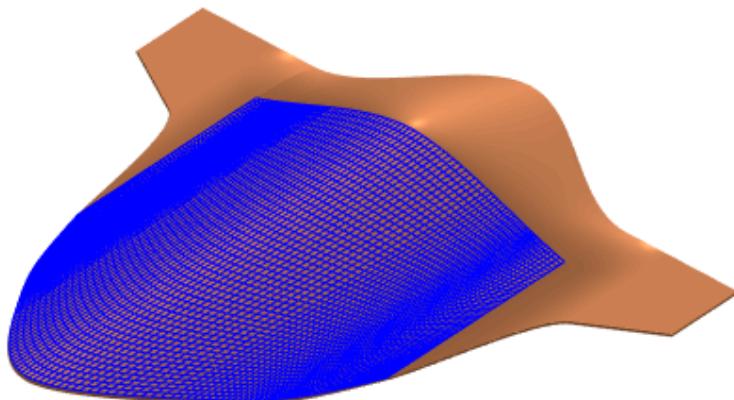
Figure 3.2.   Simple waverider-based hypersonic aircraft geometry example, with the wireframe of
             the osculating cones generated part of the geometry.

3.1 Different Approaches to Inverse Design Methods

Most waverider inverse design methods involve obtaining the shape of the leading edge on a previously selected shock shape. This can be accomplished in a number of different ways, each providing direct control of a different aspect of the geometry. Moreover, this can be exploited when there are specific design requirements or constraints, to enable the geometry to meet them more efficiently. In this section three of these different approaches to the inverse design method are documented, each one essentially providing a different origin for the stream tracing process. These three variations are applicable to all inverse design algorithms that utilize similar stream tracing techniques on calculated supersonic flowfields, from planar to axisymmetric as well as the aforementioned osculating flowfield or osculating axisymmetric flow methods. The distinct feature edges of the geometry whose shapes can be directly controlled with each approach are illustrated in Figure 3.3. The shockwave shape is defined by a shockwave profile curve on the base plane in all cases (see Appendix A for details on the osculating cones design method as many of its features and characteristics are mentioned in this section). Finally, a new hybrid approach that can combine the advantages of the different methods is also described.
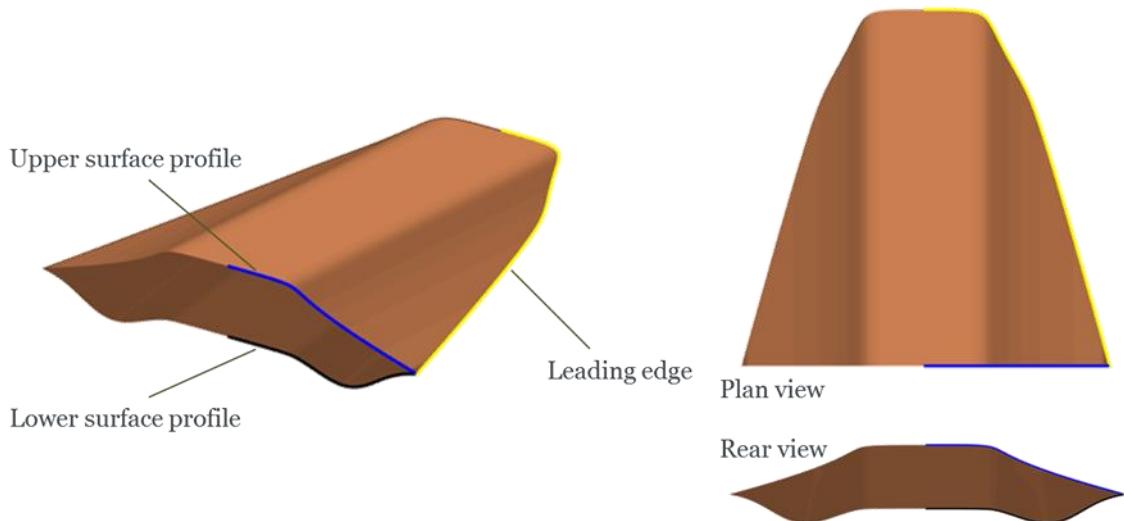
Figure 3.3.   The geometry's feature edges that can be directly controlled by the different approaches to the design method.

3.1.1 Upper Surface Profile Definition (USPD)

The most commonly used approach involves prescribing the upper surface profile on the base plane, along with the shockwave profile. This is the approach adopted in the original osculating cones method paper by Sobieczky (1990). The leading edge is essentially defined by projecting the upper surface profile, or inlet capture curve, upstream onto the shock surface. The implementation

of this method is fairly straightforward, as after picking an osculating plane and locating its point of intersection with the upper surface profile curve, the leading edge point can be trigonometrically located since the shock angle, θ_s, is known. The limitation for the upper surface profile curve is that on each osculating plane, the trace of the upper surface profile curve on the plane must be between the shockwave profile trace and the centre of the local osculating cone. If it is not, there is no feasible way to trace a streamline and the design process breaks down. This method provides direct control of the shape of the leading edge as viewed from the front/back, utilizes two profile curves that are drawn on the same plane, and also gives good insight into the structure of the osculating plane-based geometry generation. Examples of waverider forebodies generated with this method are presented in Figure 3.4.
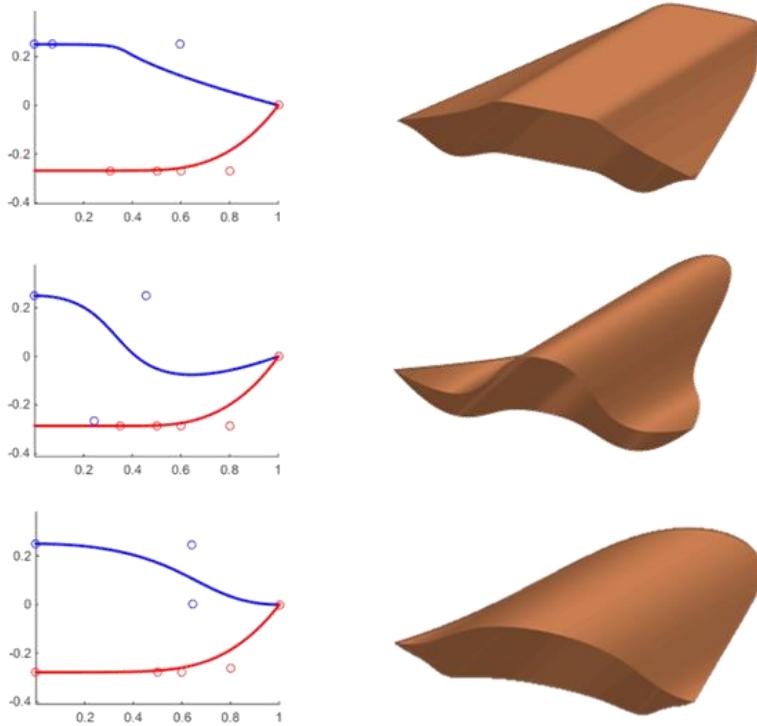
Figure 3.4.   USPD designed example geometries with the respective upper surface and shockwave profile curves on the base plane (with their control points).

3.1.2   Planform Leading Edge Definition (PLED)

Directly designing the planform shape of the geometry is another approach that is sensible from an engineering point of view. The approach presented here gives direct control of the shape of the leading edge from a planform perspective, a method hinted as being used in a presentation by Grandine (2006) (other detailed occurrences could not be identified in the literature).

The upper surface profile curve at the base plane is, in this case, replaced by a leading edge profile curve on a horizontal plane. The leading edge shape is then determined by vertically

projecting that leading edge curve onto the shock surface. More specifically, for the osculating cones method, the leading edge point on each osculating plane is located by finding the point of intersection between the shockwave trace and the vertical projection of the leading edge profile curve on that osculating plane, as seen in Figure 3.5. The lower surface is then generated by tracing the streamlines from that point to the base plane as usual. The equivalent constraint in this method is the need for the leading edge point on each osculating plane to be behind the tip of the local osculating cone, otherwise stream tracing becomes infeasible. Two example geometries generated with this method are shown in Figure 3.6.

With this method, the shape of the leading edge remains identical to the leading edge profile curve from a planform point of view, but the effective sweep angle of the leading edge relative to the freestream, which is a combination of the sweep and dihedral angles, can vary. The work of Rodi (2011) on geometrical relationships for osculating cones and osculating flowfield waveriders provides a set of useful mathematical relationships for controlling the effective sweep angle of the final geometry when using the USPD approach. Through Rodi's geometrical relations, the effective sweep angle of the leading edge relative to the freestream is directly controlled, while the plan view sweep angle can vary –essentially the opposite of the approach presented here. Both approaches can be useful depending on the design requirements that are present.
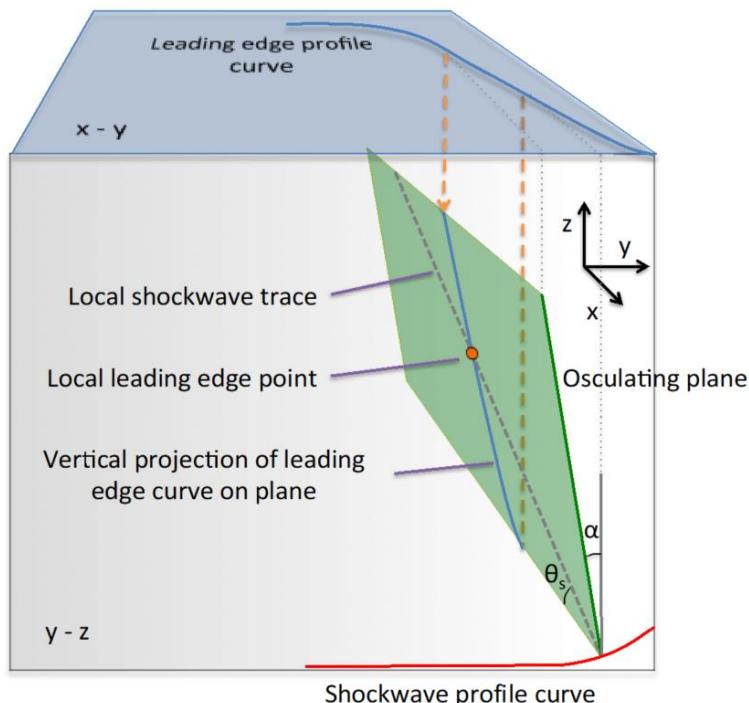
Figure 3.5.   3D geometrical relations for locating the leading edge point on each osculating plane for a plan view definition of the leading edge shape.
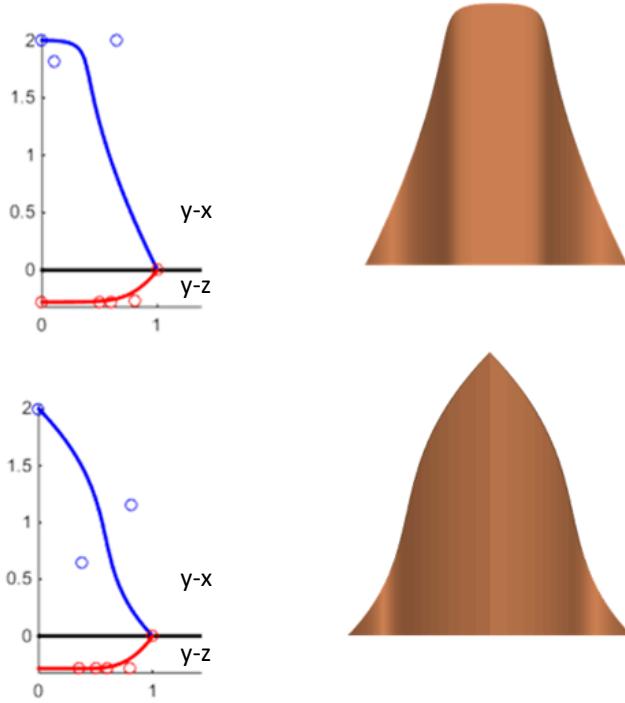
Figure 3.6.   PLED designed geometries with the leading edge profile curve from a plan view (y-x plane) and the shockwave profile curve on the base plane (y-z plane).

3.1.3   Lower Surface Profile Definition (LSPD)

The third approach to the osculating cones method incorporates the direct definition of the lower surface profile at the base plane, and is essentially an inverse of the USPD method. This method has been used in a number of studies (e.g. Lobbia, 2004), as it allows direct control of the underside of the waverider for easier engine integration. In this approach, a lower surface profile curve is used and the lower surface is first constructed by tracing the streamlines upstream up to the shock, where the leading edge profile is determined. The upper surface can then be constructed by moving in the freestream direction, from the leading edge to the base plane.

The limitation in this method is that the trace of the lower surface profile curve on each plane must not be inside the local osculating cone, as that would render stream tracing impossible. This approach gives more direct control of the lower surface shape at the base plane, where, as mentioned, there might be specific needs for integrating additional aircraft components like an engine inlet. Two example geometries with the corresponding design driving curves are illustrated in Figure 3.7.
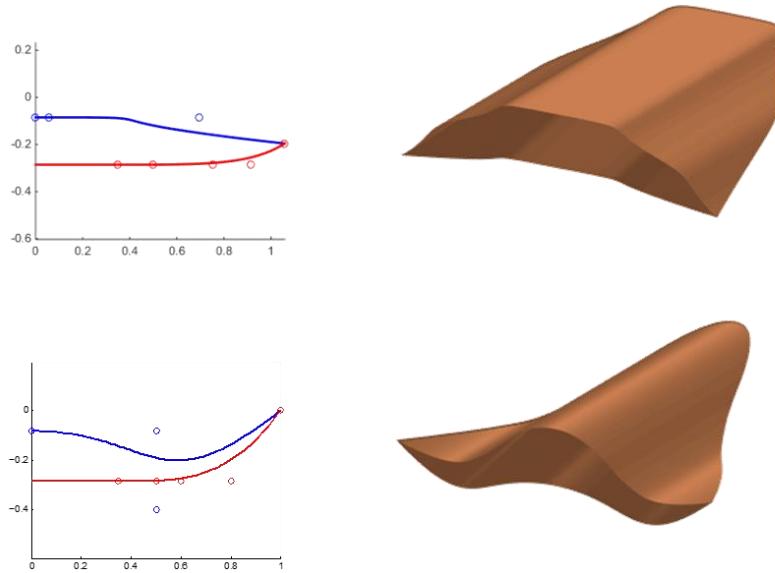
Figure 3.7.   LSPD designed geometries with the lower surface and shockwave profile curves on the base plane (with their control points).

3.1.4   Hybrid Design Approach

Encountering multiple design requirements or constraints is fairly common, particularly when the aircraft configuration becomes more realistic and several components also need to be interfaced. This can be efficiently addressed with the hybrid design approach that is presented in this section. With this approach it is possible to blend two of the aforementioned methods in order to improve the efficiency and robustness with which the parametric geometry can satisfy multiple design requirements simultaneously.

This can be accomplished by splitting the geometry generation process at a desired spanwise location, indicated by a chosen osculating plane, and applying a different design approach for each side. In this case we have two stream tracing origin defining curves, each covering only part of the full geometry. They can be designed as usual with the only difference being that the start or end point of one of the two profile curves on the interfacing plane is constrained at a position dictated by the other curve in order to satisfy continuity. Additionally, a tangency condition can be applied if smoothness of the upper and lower surfaces is desired at that interface. Finally, the geometry generation can be split over more than two sections, but this should only be needed in rare circumstances in standard design practice when complex design constraints are present.

Figure 3.8 shows a geometry designed with a hybrid LSPD and USPD approach, with the two design driving curves highlighted. By utilizing a hybrid approach, different parts of the geometry can be more readily designed to cater for specific design requirements. For example, it would make sense to use the LSPD approach for the central part of the geometry for engine inlet integration

purposes and design the rest of the geometry with the PLED method in order to directly control the sweep angle to satisfy heating constraints or enable vortex lift. Being able to directly apply design constraints to one of the design driving curves can increase parametric efficiency of the geometry model; in the case of multiple sets of constraints, however, the impact of doing so via the hybrid approach proposed here will be greater still.

USPD designed side

LSPD designed central part
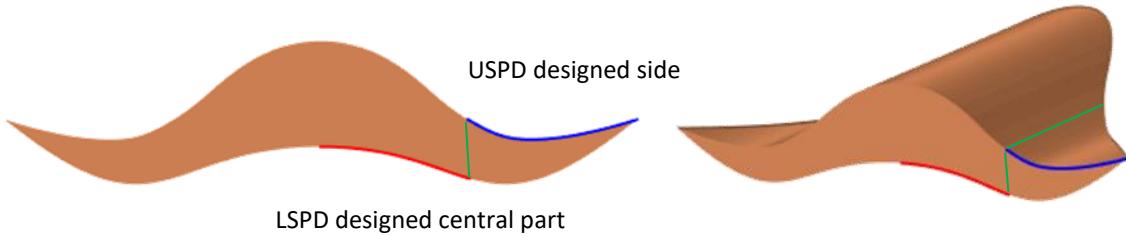
Figure 3.8.   Rear and perspective view of LSPD/USPD designed hybrid geometry with interfacing osculating plane highlighted.

The choice of the appropriate method can have a strong impact on the efficiency of the design process. To indirectly apply certain types of constraints on the feature edges of the geometry, more complex parameterizations and curve fitting methods are required. To quantify the benefits, in complexity and resources, of directly enforcing such design requirements, we conducted a series of investigations where directly designed geometries are recreated with indirect approaches (Appendix C). In general, to adequately match the desired shape the control needed over the design-driving curves usually had to increase.

Resulting upper surface profile

Flat lower surface profile

Shockwave profile curve
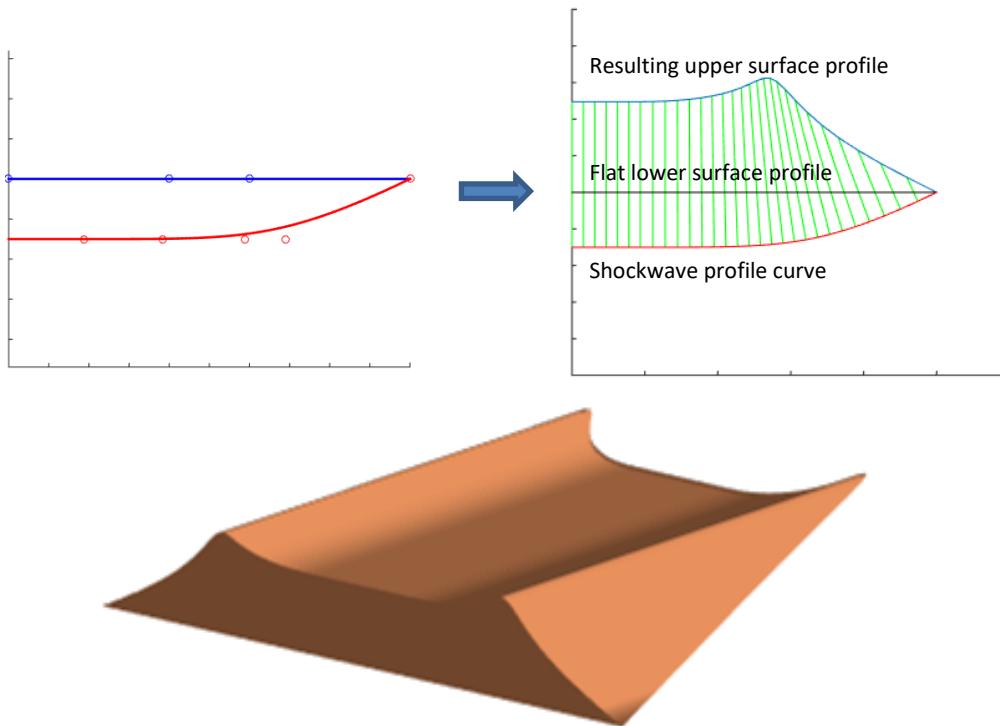
Figure 3.9.   The design-driving curves for a flat lower surface profile and a few of the osculating planes (top), with the resulting LSPD designed waverider geometry.

The geometry of one of those investigations can be seen in Figure 3.9, which depicts a waverider with a flat lower surface designed using the LSPD method by defining a straight line for the lower surface profile. Generating this geometry using the USPD method requires the use of a parametric curve able to match the resulting upper surface profile of the geometry, which in turn requires a number of additional design variables and optimization effort. The best least squares fit for the resulting upper surface profile curve using a quadratic, cubic and fourth order Bézier curve can be seen in Figure 3.10.
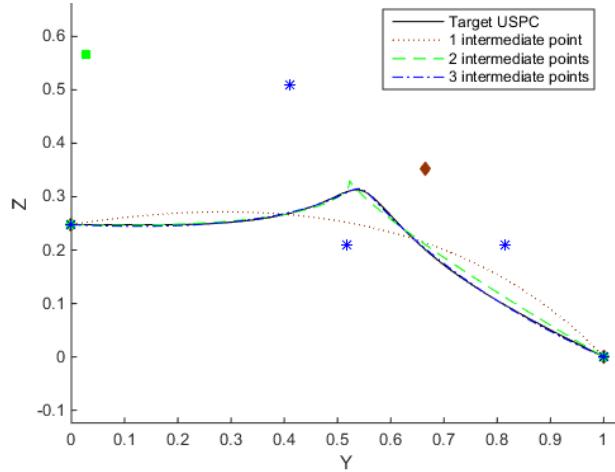
Figure 3.10. Target USPC and best fitting Bézier curves of increasing flexibility.

The quadratic Bézier curve with only 1 intermediate control point is dismissed as it is clearly not flexible enough to approximate the desired shape. The cubic one with two intermediate control points shows a great improvement and the fourth order curve seems to lie right on top of the target upper surface profile curve. An additional fifth order curve was fitted and the deviations of the parameterized curve from the target USPC as well as the deviation of the resulting lower surface trailing edge of the waverider geometry from the intended flat one were calculated, seen in Figure 3.11.
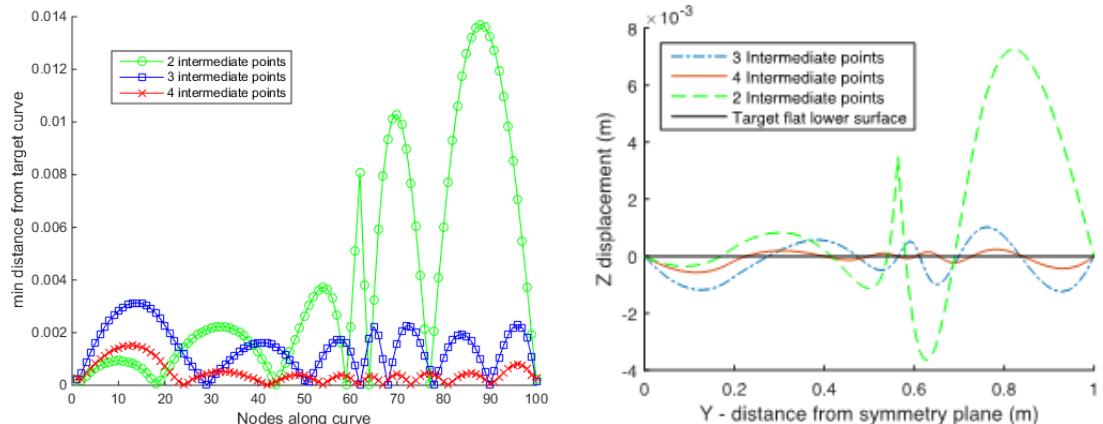
Figure 3.11. Deviation of parameterized curve from target USPC (left) and z-deviation of the resulting lower surface of the waverider (right).

This is a process we could afford if we were only optimizing a standalone waverider geometry, but the additional cost might be prohibitive in the context of a full airframe optimization where the waverider is just one of the components that needs to be interfaced with additional geometry. As seen in this case and the additional investigations presented in Appendix C, a geometry cannot be exactly matched in the majority of cases. Therefore, when indirectly trying to match a specific shape, additional considerations on how to handle the small deviations would be required. This would involve the definition of an acceptable tolerance, with the geometry having to be snapped or morphed to the exact desired shape once that is satisfied. These can all be avoided by directly enforcing any requirements on the parameterization of the respective design-driving curve. Eliminating the need for curve-fitting techniques and subsequent snapping of the final geometry to the exact desired shape also increases the robustness of the parametric geometry model.

A wide range of design requirements can be directly catered for by using the different approaches and their hybrids. These can be geometrical constraints to interface the geometry with additional aircraft components, such as a flat lower surface at the centre of the geometry for a compression ramp, a circular arc for an axisymmetric inward turning type inlet, or even flat lower surface profiles at the wing tips to hinge control surfaces. The leading edge sweep can also be directly designed to control heating or to enable vortex lift at lower speeds. In general, one of the design approaches is often better suited than the others for a particular set of requirements. With the hybrid approach the parametric geometry can also cater for multiple requirements, which are expected especially when integrating additional aircraft components. An example parametric aircraft configuration utilizing such a hybrid design approach will be presented later, in Chapter 5. These techniques provide tools to handle waverider geometries more efficiently, and proper utilization of the methods discussed here is of significant importance as the parametric geometry model gets more realistic and complex. The increased robustness offered by proper utilization of these parameterization approaches also enhances the capabilities of a geometry engine that can seamlessly utilize them. This is increasingly important for a geometry engine that is not tied to a single configuration but aims to support the development of a variety of aircraft configurations and forebody-inlet integration techniques.

3.2   Waverider Design Based on 3D Leading Edge Shapes

Inverse waverider design has traditionally relied on first generating or selecting a supersonic shock containing flowfield. The geometry is then 'carved' out of that flowfield by designing the leading edge on the shock surface and tracing the streamlines downstream of the shock to generate the 'waveriding' lower surface of the geometry. Essentially all inverse design methods, from the

simplest planar shock based designs initially proposed by Nonweiler (1959), to ones based on arbitrary shock generating bodies or the osculating cones method and its extensions, have followed this approach. A novel design approach is presented here, which differs in that the three dimensional shape of the leading edge does not have to be designed on a predefined shockwave shape. Instead, the method can explore the potential of building a waverider out of any leading edge shape; although there are limitations and any arbitrary 3D curve will not necessarily be able to provide a waverider shape.

The proposed method calculates the appropriate shape of the shockwave to fit a given 3D leading edge. Due to the shape of the shockwave being unknown and not a design-driving parameter, the approach followed in the osculating cones method and its extensions has to be utilized since they allow arbitrary shaping of the shockwave based on simple and easy to calculate axisymmetric flowfields. The proposed method can also be viewed as a further extension, or a different way to parameterize the osculating cones and similar (osculating axisymmetric flow and osculating flowfield) geometry generation methods, but without the need to design a shockwave profile.

3.2.1 Method Description

The starting point of the proposed design method is a 3D curve that fully defines the leading edge of the waverider geometry. The forward end of the curve is usually at the symmetry plane and the aft end defines where the base plane of the design method lies, as shown from different perspectives in Figure 3.12. In order to construct the appropriate shockwave shape we utilize the geometrical relationships for osculating cones and osculating flowfield waveriders derived by Rodi (2011). These relate the effective leading edge sweep, the effective angle of the shockwave, and the angle between the leading edge curve and shockwave profile from a base plane perspective, shown in Figure 3.13. The effective leading edge sweep, λ, is the sweep angle of the leading edge relative to the freestream (usually different from the sweep angle viewed from a planform perspective due to dihedral). It is calculated as:

$$\lambda = 90 - acos\left(\frac{\vec{T} \cdot \overrightarrow{u_{inf}}}{\left|\overrightarrow{u_{inf}}\right|}\right), \tag{3.1}$$

with $\vec{T}$ the leading edge unit tangent vector and $\overrightarrow{u_{inf}}$ the freestream velocity vector.

The effective angle of the shockwave, β, is equal to the shock angle for the osculating cones method, but has to be calculated when the shockwave is curved and not of equal strength along the length of the geometry for the osculating flowfield and osculating axisymmetric flow methods.

This involves the calculation of the angle of the ray between the foremost and aft-most points of the curved shock. The angle and shape of the shock, depending on the method followed, is chosen in the beginning along with the flow conditions. Given the effective leading edge sweep, the effective shock angle and an adaptation of Rodi's equation we can calculate the angle γ, between the leading edge projection on the base plane and the shockwave. We can then calculate the angle of the osculating planes to the vertical axis, α, using equation (3.2):

$$\sin(\gamma) = \tan(\lambda)\tan(\beta) \qquad \text{(Rodi, 2011)}$$

$$\gamma = \text{asin}(tan(\lambda)\,tan(\beta))$$

$$\alpha = \gamma - \theta \tag{3.2}$$

Given the angles, α, for all osculating planes along the leading edge we only need to calculate the point where the shockwave is on each osculating plane. The distance between the shockwave profile and the leading edge's projection on the base plane can be calculated given the effective shock angle, β, and the local length of the geometry on each osculating plane, which is fully defined by the three-dimensional leading edge. That distance is equal to:

$$D_{local} = L_{local}\tan(\beta) \tag{3.3}$$
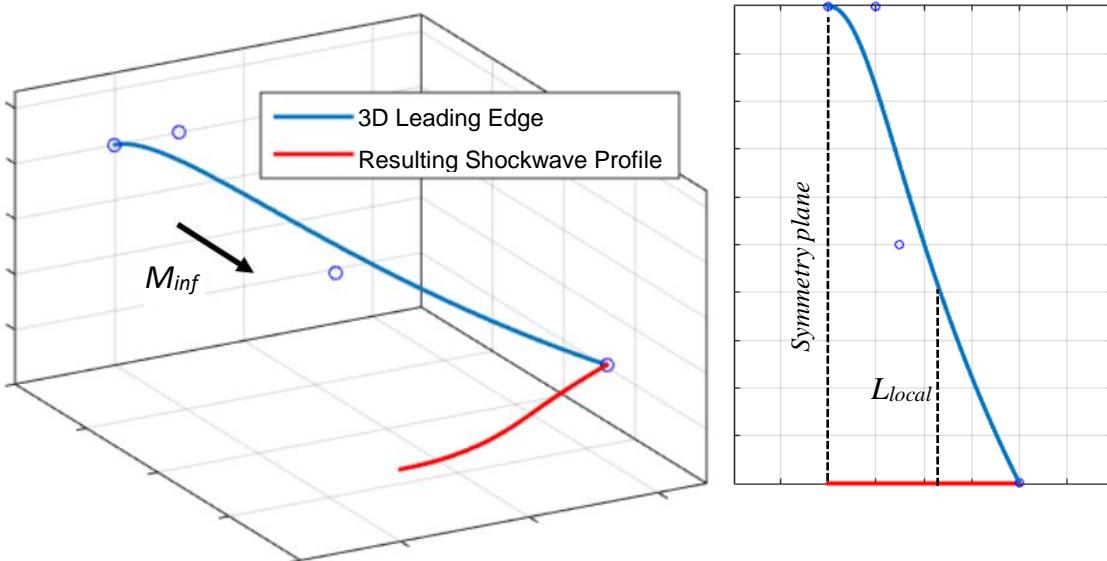
Figure 3.12. Perspective and planform view of the 3D leading edge curve and its resulting shockwave profile on the base plane.
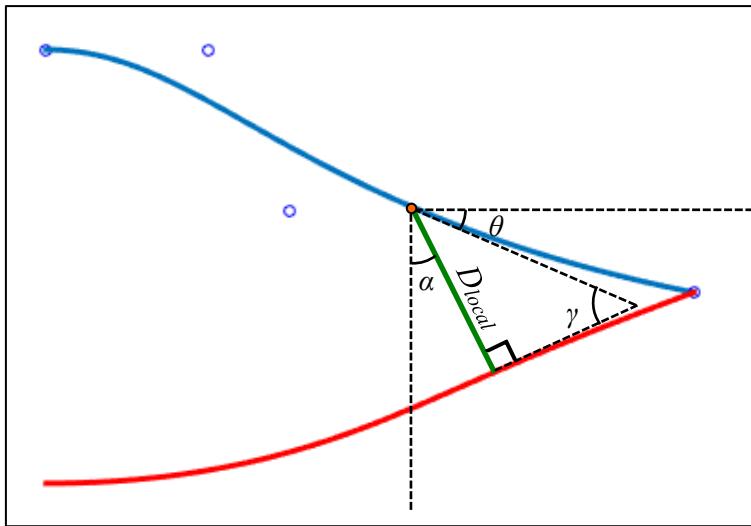
Figure 3.13. Base plane view of the leading edge with the geometrical features used for the calculation of the osculating plane angles and resulting shockwave profile.

Knowing the distance of the shockwave from the base plane projection of the leading edge on each plane we can now fully reconstruct the shockwave profile curve of the osculating cones method. The radius of curvature along the obtained shockwave can be calculated either numerically or through a series of complex chain rules relating it to the derivatives of the 3D leading edge curve; here we use the former. To increase the accuracy of the numerical calculation we can utilize additional 'dummy' shockwave profile curve points that are not used past this point and are placed in pairs on either side and close to the points where the radius of curvature is evaluated in order to enable a much more accurate three point calculation. The geometry generation process can, from this point on, be completed as described by the osculating cones method by tracing streamlines from the leading edge to the base plane to generate the lower surface of the waverider. The waverider generated using the 3D leading edge curve of Figure 3.12 and Figure 3.13, a shock angle of 15^0, and freestream Mach number equal to 6 can be seen in Figure 3.14.

It is, of course, impossible to generate waverider geometries out of any arbitrary 3D shape that instantiates a parametric curve. It is not very straightforward to formulate constraints that can be applied to the three-dimensional leading edge shape in order to always satisfy the requirements for generating a waverider geometry. However, the computational cost of constructing the potential shockwave profile from any 3D curve is negligible and we can instantly know if a specific curve can be matched with a viable waverider shockwave shape. Among the parametric curves tested we found that having second order continuity is preferable. Another obvious constraint is that the effective sweep should not go beyond the freestream Mach angle. Moreover, we can expand the capabilities of the osculating cones method, as well as the range of supported leading edge shapes, by enabling handling of concave shockwave shapes, as will be explained in the next

section. Once this was applied, we found that the majority of leading edge shapes that should intuitively be able to provide a waverider shape would work with the method, rendering it much less restrictive than originally anticipated.
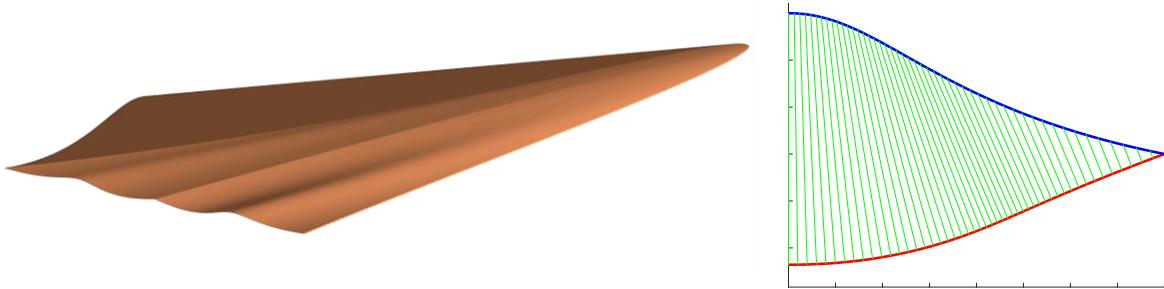
Figure 3.14. Resulting 3D leading edge designed waverider with converging ICFA flow used for part of the geometry generation (left). Osculating planes and profile curves on the base plane (right).

3.2.2   Additional Considerations

Not being in direct control of the shockwave shape, since it is derived in order to fit the three-dimensional leading edge, leads to cases where the shockwave profile will also be concave. Normally, the osculating cones method can only handle convex shockwave profiles. Therefore, we can expand the range of leading edge geometries that meet the requirements of the method if we enable handling of shockwave shapes with concave regions. A concave shockwave shape also points to a converging flow, and the type of converging flow that best fits our purposes is the ICFA flow (Internal Conical Flow A) as described by Moelder (1967). Converging flowfields have been previously used for waverider design by Goonko (2000). This type of flowfield enables stream tracing on osculating planes where the shockwave shape is concave by utilizing the Taylor-Maccoll equations. It does however come with its own limitation. The ICFA flowfield terminates at a singularity beyond which stream tracing is not possible. This effectively places a limit on the minimum concave radius of curvature of the shockwave profile curve, which depends on two parameters: the angle in spherical coordinates where that singularity occurs for the given flow conditions and the local length of the waverider geometry on the osculating plane, as shown in Figure 3.15. Getting around this by allowing the lower surface generation to continue past the singularity following the last given direction of the ICFA streamline can be considered, since the shape of the shock will not be affected as long as we do not reach regions where the shock is getting close to the local axis. Finally, while the ICFA flowfield can work with the osculating cones method since they both feature straight shocks, when the proposed 3D leading edge design method is used with the osculating axisymmetric flow or osculating flowfield approaches, other types of converging flowfields will need to be numerically constructed.
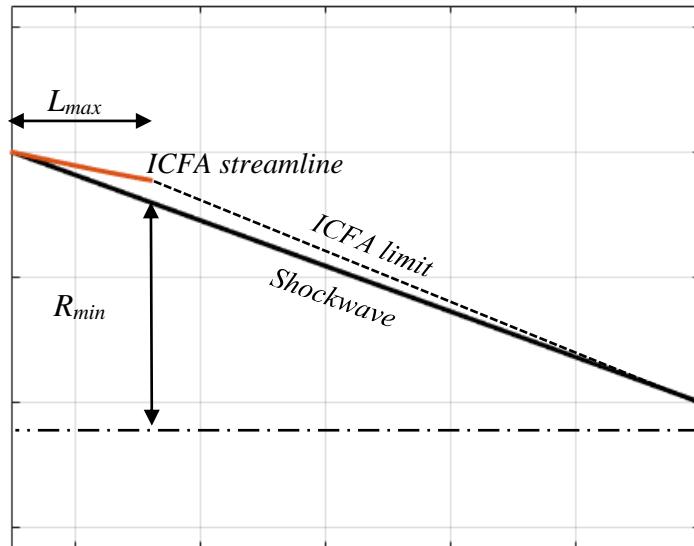
Figure 3.15. ICFA flowfield limit with resulting maximum local length and minimum local radius of curvature.

One advantage of previous design methods is that by directly designing the shockwave in the region of the engine inlet, its integration with the forebody becomes very straightforward. While the method presented here does not inherently provide direct control of the shockwave shape, there are two potential ways to obtain the desired flow in the inlet region. If, for example, a uniform flow and planar shock shape is required at a certain region, this can be enabled by keeping the angle of the osculating planes, α, constant. This, in turn, means that a constraint needs to be placed on the three-dimensional leading edge curve relating its effective sweep angle, λ, with its inclination from a base plane perspective, θ (y-z plane if x is the freestream direction), where:

$$\theta = \operatorname{asin}(tan(\lambda)\,tan(\beta)) - \alpha \tag{3.4}$$

This effectively constraints the leading edge, or part of it, to lie on a single plane that locally describes the desired planar shock. Depending on the type of parametric curve used, this can be complex to implement. Another way to obtain the desired shockwave shape in part of the underside of the geometry while maintaining any advantages provided by the three-dimensional definition of the leading edge curve for the rest of it, is to use a hybrid design approach as explained in 3.1.4. We can use the osculating cones method as originally described and specify an upper surface profile or inlet capture curve and shockwave profile shape only for the region of the engine inlet for example, and use the 3D leading edge method for the rest of the geometry. Equation (3.4) will only need to be applied at the interface of the two methods and essentially provides a tangency

condition for one end of the 3D leading edge curve; this is much easier to implement than applying the condition over a segment of the curve.

For geometries that are fully defined by a three-dimensional leading edge, the angle of the osculating plane at the symmetry needs to be zero, otherwise the lower surface will either be discontinuous or self-intersecting. This places a constraint on the tangency at one end of the three-dimensional leading edge that lies on the symmetry plane, equation (3.5). This translates to the tangent of the curve at the symmetry plane being restricted to lie on a plane that has an angle equal to the shock angle with the freestream and is perpendicular to the symmetry plane, illustrated in Figure 3.16. For smooth leading edge geometries at the symmetry plane this is always satisfied as the tangent is parallel to the y axis and will lie on the plane that describes the shock no matter its angle. If a pointed leading edge waverider is to be designed, this constraint can be applied by restricting the second control point of a Bézier curve to that plane, since the first intermediate control point determines the tangency at the respective end of the curve. If the entire curve is designed on such a plane, with an angle relative to the freestream equal to the shock angle, a wedge flow planar shock waverider is obtained. An example of this can be seen in Figure 3.17.

$$\sin \theta = tan(\lambda)\, tan(\beta) \tag{3.5}$$
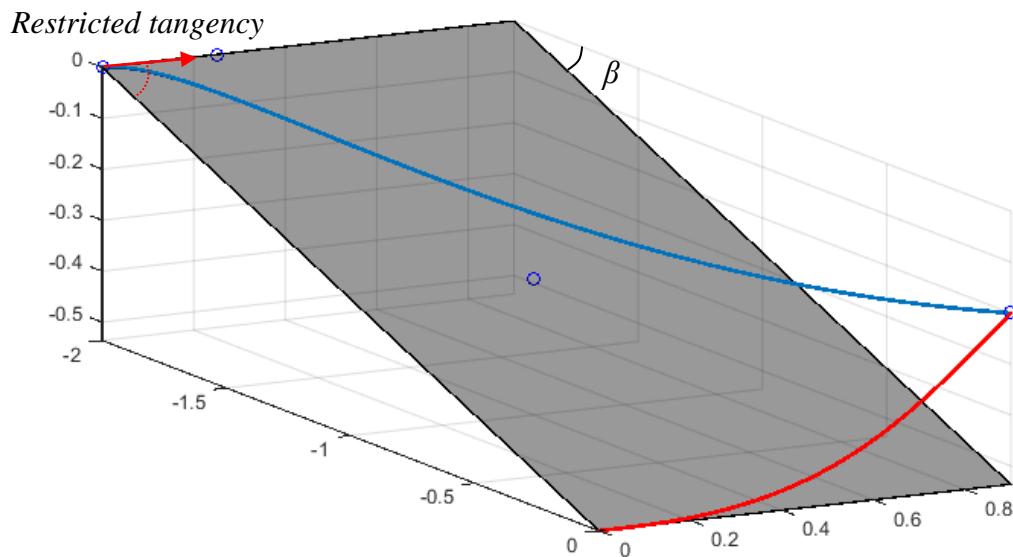
Figure 3.16. Three-dimensional leading edge with its control points and tangency-restricting plane.

Symmetry plane view of 3D leading edge

3D perspective view

Resulting planar shock
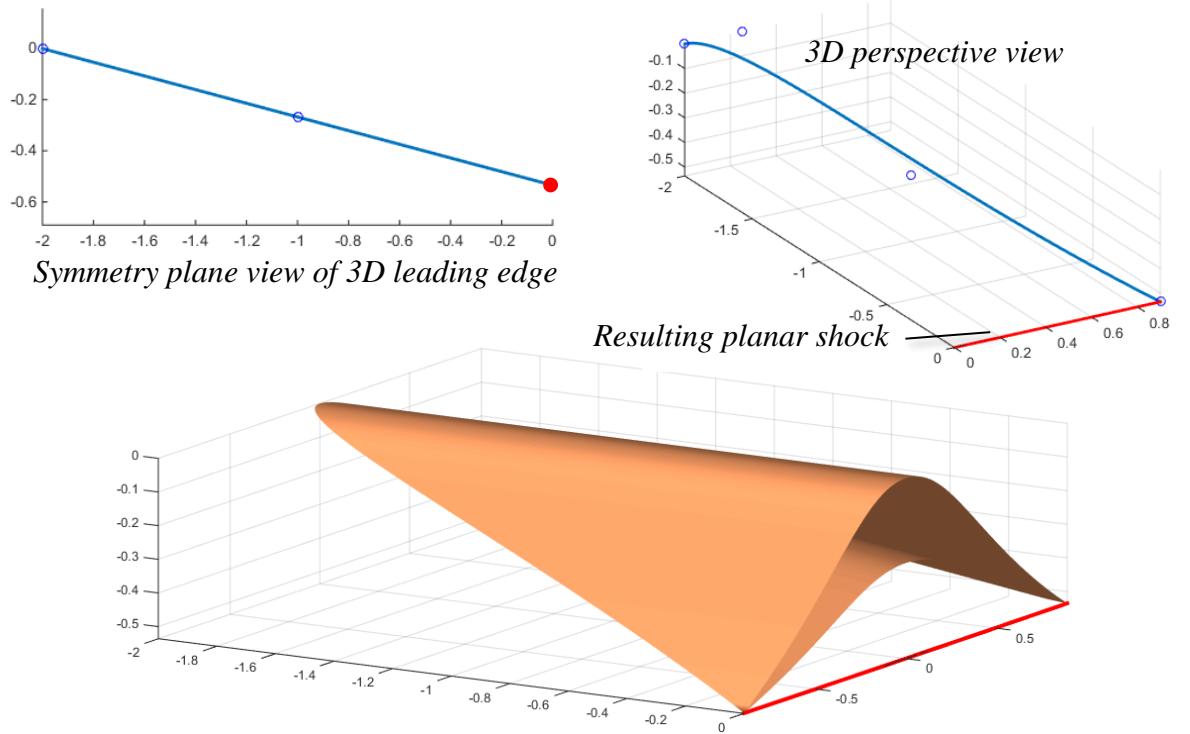
Figure 3.17. Planar wedge shock waverider designed by having the entire leading edge on the tangency restricting plane.

As was mentioned earlier, parametric curves with high degrees of geometric continuity are generally preferred for parameterizing the 3D leading edge. When designing the waverider in a piecewise manner, either by a hybrid design approach or by expressing the three-dimensional leading edge with two or more parametric curves, the geometric continuity at the interfaces needs to be considered. The geometric continuity is related to the continuity and smoothness of the resulting surfaces. For the original osculating cones method, continuity limitations apply for the shockwave profile curve, as is explained in Appendix A. The upper surface profile curve, on the other hand, can be much more freely designed. In this case, both curves have been replaced by one three-dimensional leading edge curve. The equivalent limitations can be translated for the leading edge curve by looking at its continuity from different perspectives. For this we consider the local projection of the leading edge on two planes, the first is a plane that is tangent to the local shockwave shape and perpendicular to the local osculating plane, and the second a plane that is perpendicular both to the first plane and to the local osculating plane. These planes change along the curve, Figure 3.18 illustrates both planes for the starting point of the three dimensional leading edge. Second order continuity (G^2) of the local leading edge projection on the second, perpendicular to the local shock, plane needs to be maintained. Third order (G^3) is required to obtain smooth surfaces without edges, similar to what applies for the shockwave profile curve in the osculating cones method. That is not required for the local projection on the shock-tangent plane, where only geometric continuity is required (G^0).
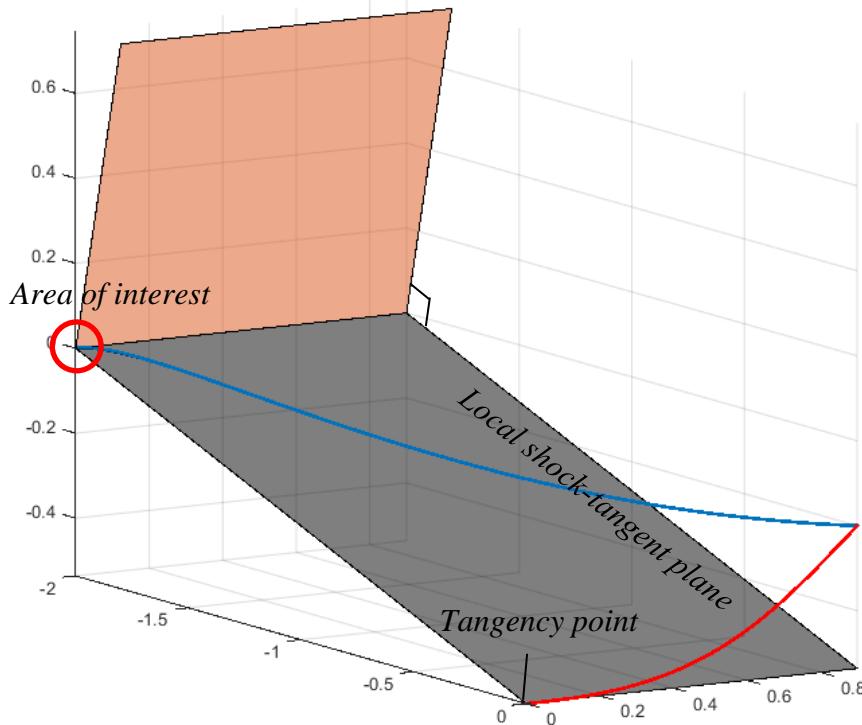
Figure 3.18. Plane tangent to shockwave profile curve (grey) and the plane perpendicular to that and the local osculating plane (orange). The local osculating plane for this part of the curve is the plane of symmetry.

Again, this is not a constraint that can be applied across a parametric curve but one that should be considered at interfacing points. Let us consider a hybrid designed waverider where the central part of the geometry is designed with a planar wedge shock for flow uniformity using the USPD method, followed by a Bézier curve describing a three-dimensional leading edge. The first control point will be restricted to the interface. The USPD designed region features a shock with no curvature. In order to obtain zero curvature of the shock at the start of the 3D leading edge controlled section, the second and third control points need to be on the local shock-tangent plane. In this case this plane is the wedge shock of the USPD designed region. This will result in a continuous lower surface. In order to also guarantee a smooth lower surface at the interface the fourth point will also need to be on that plane. The 3D leading edge does not need to have an increased level of continuity at the interface, since the points are free to move on the shock-tangent plane. However, having the points on the shock-tangent plane satisfies the continuity requirement from the perspective of the plane perpendicular to it, as was explained earlier.

The effect of this can also be seen on the symmetry plane of 3D leading edge designed waveriders. By placing the second and third control points on the tangency-restricting plane the start of the shock on the symmetry plane is planar with no curvature. Without the fourth control point on the tangency-restricting plane as well, the curvature of the shock starts rapidly increasing

as we move in the spanwise direction on either end of the symmetry plane, which results in an edge. By placing the fourth control point on the tangency-restricting plane the curvature decrease and increase as we move through the symmetry plane is smooth, resulting in a smooth lower surface at the symmetry plane. Two examples illustrating this can be seen in Figure 3.19. Overall, taking these requirements into account when parameterizing the method will increase the number of designed leading edge shapes that result in valid waverider geometries.

2 intermediate points on tangency-restricting plane

3 intermediate points on tangency-restricting plane
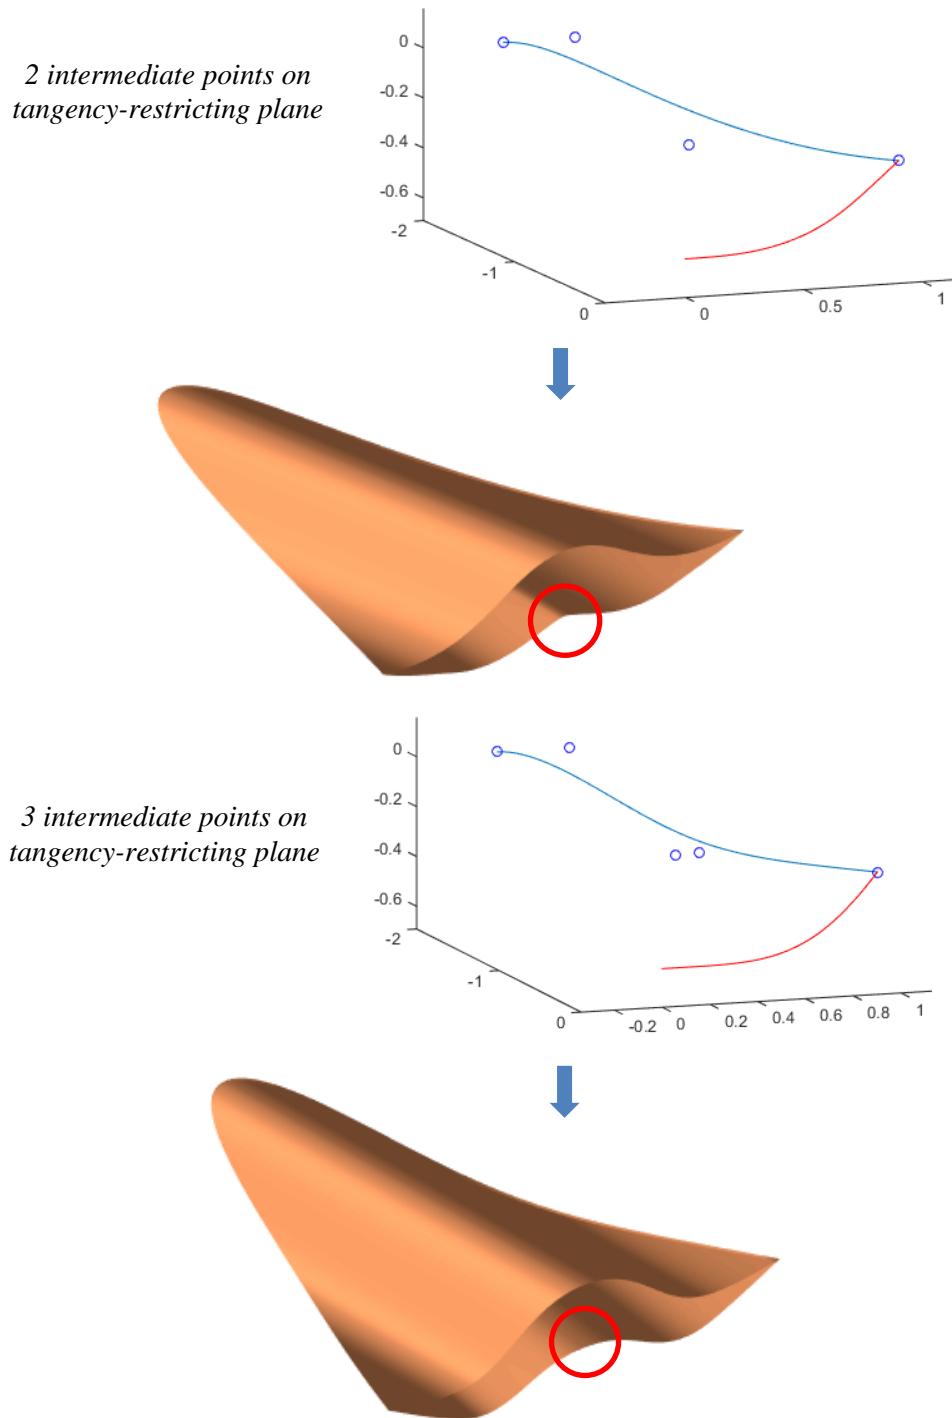
Figure 3.19. 3D leading edge designed waveriders with different levels of continuity at the symmetry plane.

The method presented in this section enables generation of waverider geometries from fully defined three-dimensional leading edge shapes. Although it utilizes principles of existing waverider generation methods, it is a clear departure from the notion of having to specify a shockwave shape and then draw the leading edge on the shock surface. Apart from enabling any 3D leading edge shape to be considered for generating waverider geometries, it can also potentially increase the efficiency of parametric waverider geometry models as only one curve needs to be parameterized, albeit a three-dimensional one. This means that, given specific design requirements for the shape of the leading edge, it can be more efficient to use this method to parameterize the geometry prior to a design optimization study, with the alternative being the use of two parametric planar curves. It can also be seen as a step away from inverse design and one towards direct waverider design.

As far as the aerodynamic performance of the geometries obtained goes, the method can match what the original osculating cones method is capable of. We expect that by utilizing this novel approach with the osculating axisymmetric flow and osculating flowfield extensions of the osculating cones method the designs will be able to match the performance of geometries designed with those, since the underlying design concept is the same. Another area that is worth further investigating is the potential expansion of the design space from an aircraft stability perspective, offered by enabling concave regions of the shock which can lead to greater flexibility in moving the centre of aerodynamic forces.

3.2.3   CFD Validation

The proposed design method utilizes part of the osculating cones waverider generation algorithm, and assumes that the effect of cross-flow pressure gradients in the ensemble of osculating flowfields is minimal. It is, therefore, expected that the validity of the resulting flowfields will be equivalent to that of the osculating cones method, which has been confirmed through multiple studies. There are two aspects to further validating the inviscid aspect of the proposed design method, which in this case will be done through qualitative comparisons of the predicted flow features with CFD calculated data. The first is investigating the flowfield and shock shape that is generated by the geometries to verify that there is indeed a leading edge attached shock of the predicted shape. The second is confirming the predicted behaviour of the flowfield in regions where the ICFA flow is utilized. The latter can be accomplished again by comparing the predicted and the calculated shockwave shape in regions where it is has a concave shape.

The inviscid flowfield around the forebody geometries can be numerically calculated with Euler simulations. Since a major part of the validation will be comparing the shape of the shockwave calculated through CFD with the predicted one, the computational grid has been further refined in

the region of the predicted shockwave for shock capturing purposes. The forebody geometry investigated features a shock shape with a slightly concave region where the ICFA flowfield was utilized within its limitations. It is a 3D leading edge designed Mach 6 waverider, fully defined by the shape of its leading edge. An overview of the geometry and pressure iso-lines plotted on three different slices can be seen in Figure 3.20. A pressure contour of the base plane, where the predicted shockwave shape is also drawn, is shown in Figure 3.21. The use of the term 'predicted' shockwave shape points to the fact that the shockwave profile is not directly designed, but is the result of the specified three-dimensional leading edge shape and it is generated to fit that.
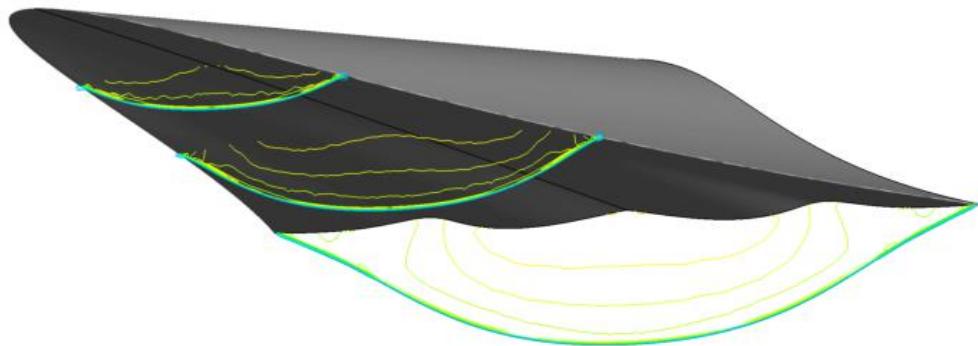
Figure 3.20. Pressure iso-lines in three normal to the freestream slices for qualitative analysis.
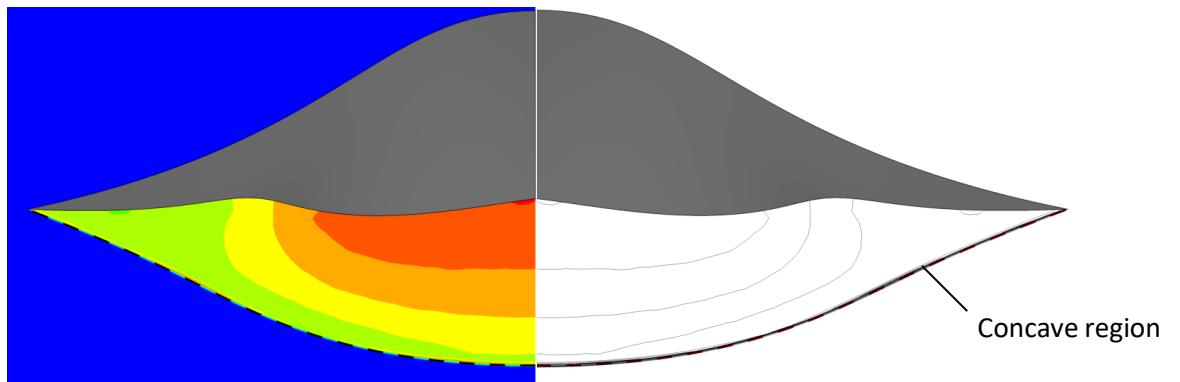
Figure 3.21. Euler calculated pressure contour and iso-lines at the base plane with predicted shock shape (dashed line).

Through this qualitative comparison we can confirm that the geometry recreates a leading edge attached shock that matches the shape of the design method predicted shock. The concave region of the shock also matches, further verifying the use of ICFA flowfields in the osculating cones algorithm. Additional confidence in the ability of the method to generate geometries that perform as predicted is given by the comparison of lift to drag ratios computed through CFD and through the design-predicted flowfield characteristics. The predicted L/D for the previous geometry was

7.2625 while the one calculated through the Euler simulation was 7.2661, a negligible difference. Finally, a comparison of dimensionless pressure (p_{local}/p_{inf}) contours of the lower surface of the waverider can be seen in Figure 3.22.
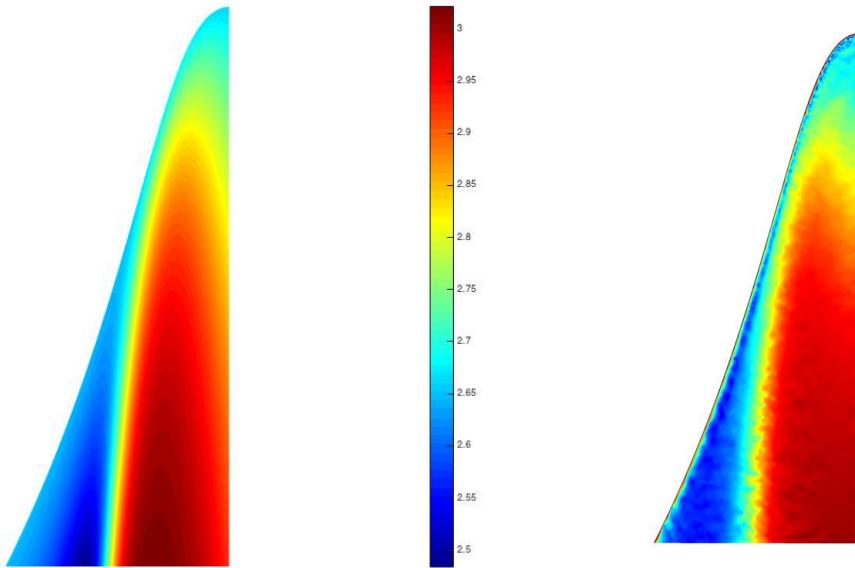
Figure 3.22. Dimensionless pressure contour of the lower surface, design method predicted (left) and Euler CFD (right).

3.3   Flexibility of the Design-Driving Curves

In the previous sections we looked at different approaches to inverse design methods in order to enable the parametric geometry to meet potential design constraints and requirements more efficiently. This was accomplished by essentially changing the stream tracing origin of the design algorithm in order to directly control a specific part of the geometry. In this section we are looking into the parameterization and geometric flexibility of the design-driving curves that are used to define that origin. Limited design experience of such vehicles means that no guidelines are available to help the engineer steer a course between over-parameterising and over-constraining (that is, equipping the geometry with an insufficient number of degrees of freedom) a model. One approach to justifying geometry parameterization options before moving to higher fidelity and larger scale design studies of more realistic aircraft geometries, is to perform low fidelity design space explorations in order to assess the parameterization itself by comparing different parameterization options. In a way, this is a process through which 'design experience' is generated in order to support decisions made in the development of the parametric geometry. Inversely designed waverider geometries are ideal for such studies since their standalone aerodynamic performance characteristics can be rapidly evaluated with sufficient accuracy.

3.3.1 Methodology

We will be conducting low-fidelity design space explorations and investigating the effect of varying the flexibility and degrees of freedom of the design, in order to obtain more insight on where the limits of over and under-parameterizing the geometry model may lie. This is an area where the development of the parametric geometry model starts blending into the preliminary design stages of the concept aircraft. The main feature that sets such a study apart from the early design stages for a given mission, is that the geometry is still incomplete and this assessment concerns only one component of the final geometry. We are also 'freezing' a number of parameters to focus on the parameterization of the part being investigated.

Essentially, the methodology presented here and the conclusions drawn from it are to be used when deciding a parameterization scheme both for a standalone waverider geometry as well as when developing a more complete aircraft parametric geometry model, to assist with part of its parameterization – the waverider component. The reason for such investigations is that, in the literature, case studies with parametric geometries using from 4 to 20 design variables for the design-driving curves of the waverider component are common – the core of our argument here is that a rational basis can be developed for the choice of such numbers. In other words, there is no guarantee that the behaviour of the high fidelity flow simulations will yield an objective landscape that behaves in the same way as our low cost surrogate here, but what we are proposing here is the closest affordable approximation. Capturing some of the salient features of the real landscape at a fraction of the cost is likely to yield a better parameterisation than the judgement calls having to be made by engineers at present.

Conducting a parameterization flexibility study on more realistic aircraft configurations that require higher fidelity analysis tools is computationally prohibitive; therefore, it is done at the highest fidelity and design detail where we can still afford it. Aerodynamic performance characteristics of the baseline sharp leading edge waverider forebodies can be quickly estimated by using the flowfield characteristics available through the inverse design process. For the osculating cones method for example, flowfield metrics can be obtained through the Taylor-McColl equations for the lower surface and freestream values can be used for the parallel to the freestream upper surface. A boundary layer integral method can also be utilized to estimate viscous skin friction effects (Appendix D). A very simple evaluation method was implemented, assuming laminar flow for the entire forebody and no interaction between the viscous and inviscid parts. Inviscid pressure forces and viscous skin friction forces are calculated for the wetted surface of the geometry (excluding the base – on which freestream static pressure was assumed to be acting throughout this analysis) and then integrated to obtain the resulting force on the body. The calculation method

is further explained in section 5.3.2. As the two example comparisons presented in Table 3.1 show, the inviscid L/D results obtained in this way are within ~0.5% of the results obtained from inviscid Euler computations, while the viscous results are within ~1.5% of laminar RANS (Reynolds Averaged Navier-Stokes) computations with freestream conditions equivalent to flight at 75,000ft for the first and 100,000ft for the second case. ANSYS Fluent was used for the comparisons, with convergence of the computational grid achieved for around 4 million cells.

Table 3.1.   Two of the inviscid and viscous L/D calculations comparisons with CFD.

Geometry	Euler CFD	Inviscid calculated	RANS laminar CFD	Viscous calculated
Mach 6	5.34	5.358	4.82	4.847
Mach 8	9.85	9.824	7.36	7.268

Since our objective is to develop the parametric geometry model, what matters most is relative changes of the characteristics between different designs and not the absolute values, something that this simplified approach can adequately provide. One can of course opt to use more complex evaluation tools of equivalent computational cost to account for more factors, such as local inclination methods for a shaped upper surface, turbulent boundary layers, etc.; as well as different parametric geometry formulations or even focus on more objectives such as the on-design moments/stability.

What we are looking to compare is the sets of optimum solutions, or Pareto fronts, derived following a simple optimization process utilizing the different parameterization schemes. By comparing the Pareto fronts we don't have to focus on a single design point in cases with competing optimization objectives. This approach allows us to compare the parameterization schemes over a wider area of the design space. It is also impossible to identify an optimum balance between competing objectives (such as L/D and volumetric efficiency) so early in the design stage, when the geometry model is still being developed. When comparing parameterizations with the same number of degrees of freedom the comparison is fairly straightforward, the Pareto front of the parameterization scheme that provides a better way to parameterize the geometry will usually dominate the Pareto front of the other scheme, thus being more effective. There might be cases where one parameterization scheme performs better than the other only in parts of the entire

design space, something that we will be able to identify through the Pareto front comparison. When comparing parameterizations with different degrees of freedom, however, the level of convergence and populations of the Pareto fronts are also metrics to be considered. This is especially true when a fixed computational resources approach is followed for all optimization runs, which is an approach that better represents a realistic design scenario. By following a fixed computational resources approach we are able to also compare the efficiency of different parameterizations and how readily accessible the 'good' parts of the design space are during an optimization process. Through these comparisons we should be able to identify under-parameterized geometry models –where their performance is restricted due to insufficient flexibility or flexibility applied in a meaningless manner, and over-parameterized geometry models –where the amount of flexibility provided is greater than what the optimization algorithm and available resources can handle in order to locate the best designs the parameterization can offer.

In the sections that follow we present two case studies, each providing different conclusions on the parameterization options evaluated. In the first study we examine the effects of varying the flexibility of the upper surface profile curve when using the USPD method for waverider design, while in the second we compare different parameterizations of the shockwave profile curve.

3.3.2 Case 1

For the first case, designs were optimized with four different curve parameterization schemes of the upper surface profile curve, using one, two and three movable intermediate control points (meaning three, four and five control point Bézier curves, respectively). The shockwave profile curve and shock angle are fixed for all cases, as are the length and width of the geometry, restricting the end control points. The only further constraints we enforce are those that ensure physically and geometrically plausible waverider shapes. These constraints allow us to focus on the part of the geometry that is being evaluated – the upper surface profile curve. The parameterization schemes are as follows:

• 1 intermediate point – able to move in two directions, 2 design variables.

• 2 intermediate points – both able to move in two directions, 4 design variables.

• 3 intermediate points – all able to move in two directions, 6 design variables.

• 2 intermediate y-restricted points – equally spaced along the span and able to move in the vertical direction only, 2 design variables.

A set of simulations were conducted at a Mach number of 8 and Reynolds number of 8.316×10^6 based on the length of the forebody. We are seeking to maximize two objectives: the lift to drag ratio (L/D) and the volumetric efficiency (V/A, volume of the forebody geometry/wetted area excluding the base plane). These two goals generally compete in waverider design, so a germane comparison between the parameterization schemes is an analysis of the relative positions of their respective Pareto fronts in V/A versus L/D space.

In order to generate the fronts, we ran a basic genetic algorithm (Goldberg, 1989) search for each parameterization. A fixed computational cost approach with 3000 calculations for each parameterization was followed. The resulting Pareto fronts and some of the geometries along them are presented in Figure 3.23. The distribution of evaluated designs in the multi-objective landscape from which the Pareto fronts were derived for each case is shown in Figure 3.24. Adding more degrees of freedom could further advance the overall Pareto front, but such advances need to be viewed in the context of the cost of the increments in degrees of freedom and expansion of the design space. It is already visible that with more degrees of freedom the convergence of the Pareto front degrades.
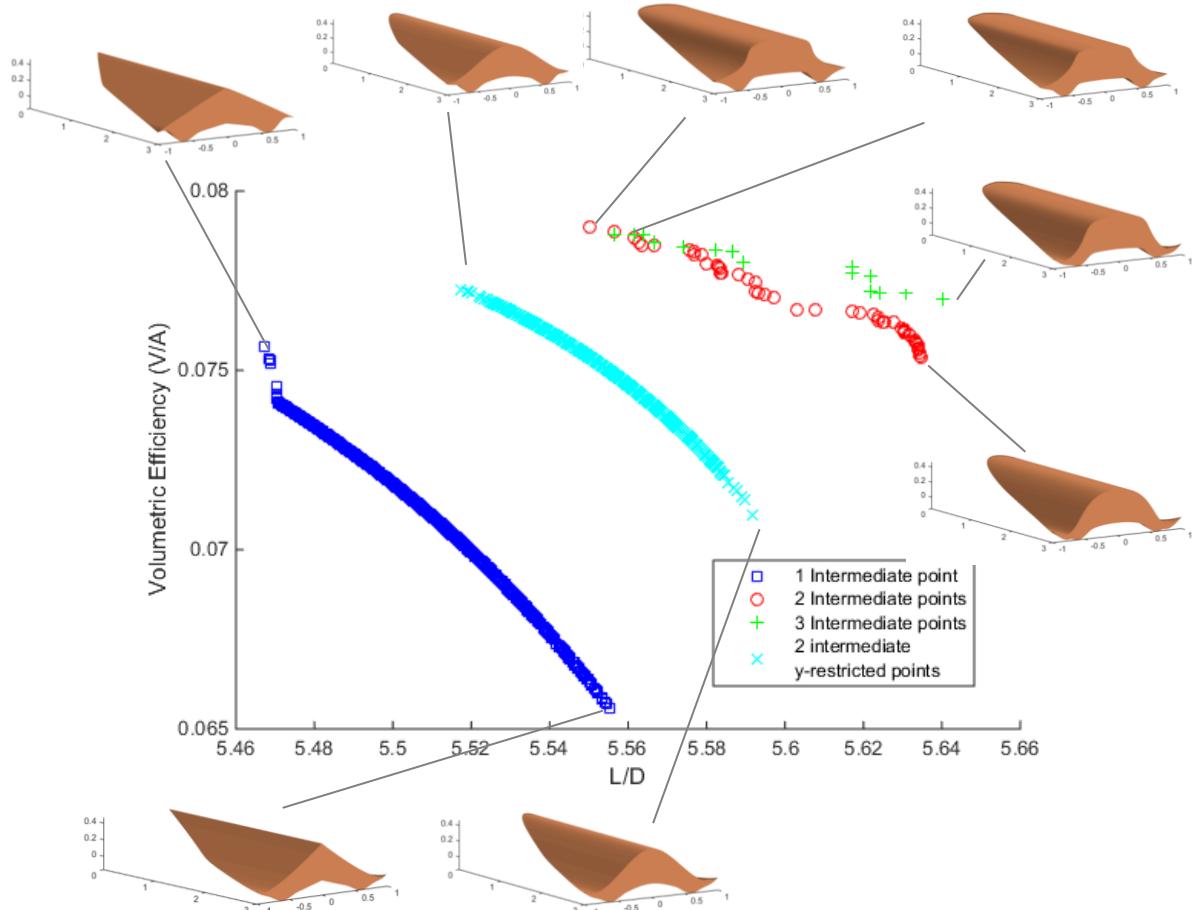
Figure 3.23. Pareto fronts for different parameterization options of the upper surface profile curve.

The Pareto fronts identified by the two design variable optimizations (the first and fourth cases) consist of about 900 designs each, the second case with four design variables has 50 and the third one with six design variables only has 15. The latter Pareto fronts are not resolved to the same degree as the two design variable cases due to the convergence (or end) criterion being a simple computational budget limitation. A much greater number of calculations and computational resources would be needed to achieve the same level of Pareto front convergence for the more complex parameterizations. What also needs to be considered is that if we cannot reach the optimum geometries with this number of evaluations when focusing on this part of the geometry, it will be even harder to do so when the geometry has been integrated in a more complex geometry model and is being optimized along with additional components. Although, admittedly, the generic genetic algorithm is not the most efficient tool we can use.
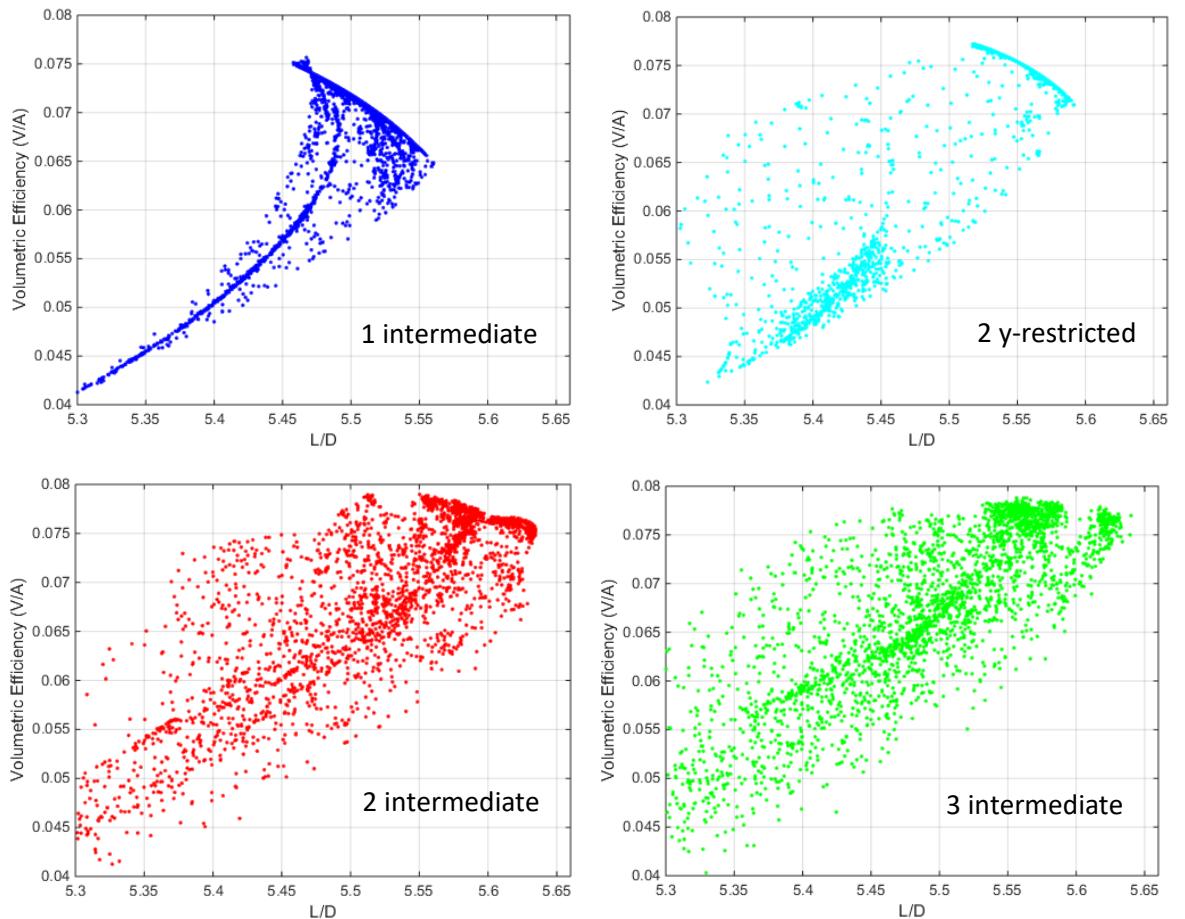
Figure 3.24. Design objective landscape for the four optimization cases.

From the Pareto front landscape of Figure 3.23 we can draw a number of conclusions that can assist us when deciding on the parameterization of the inversely designed waverider component of an aircraft geometry model. The option of using a quadratic Bézier curve with a single intermediate control point yields, predictably, the poorest performance, as with the same number of design variables and using two intermediate control points with one degree of freedom each, a

much better set of optimal solutions can be obtained. This also points to an advantage of enabling double curvature shapes for the profile curve. Allowing the two intermediate control points to move in two directions yields a further improvement that would potentially be desirable should the increased dimensionality of the design space and subsequent increase in the cost of more advanced design studies be acceptable. The same applies to using three intermediate control points, although the improvements are marginal at this stage. Finally, further information can be obtained from the distribution of the ~3000 evaluated designs on the design objective landscape for each case, seen in Figure 3.24. In this figure we observe that even for the two higher dimensionality cases whose Pareto fronts were not as densely populated, the region close to the fronts is still well populated. This results in a good number of designs cluttered in that region that still outperform the Pareto fronts of the first two cases.

Given a fixed computational budget one could say that the most flexible parameterization is the best because its Pareto front, even without being fully converged, is able to totally dominate the others. However, when the waverider component is to be optimized along with additional aircraft components, consistently being close to the Pareto front may be just as important. This means that it is not always straightforward to identify a clear winner among the different parameterizations tested, but the information obtained from such investigations provides more insight and useful information in order to further support the decisions made during the development of the parametric geometry.

3.3.3 Case 2

Evaluations such as the one above can be applied to almost all aspects of the parametric geometry. In the following example we consider the impact of adding more flexibility to the shockwave profile curve. As shown in Figure 3.25, we are using a shockwave profile consisting of a straight line for the central region to obtain uniform flow, and a 4th order Bézier with the three intermediate points in line in order to achieve third order continuity at the interface (point 0) with the straight line segment.

In this case we are looking to evaluate how much control we need over positioning those three intermediate points. Point 0 is fixed to 35% of the total span as a design constraint for the area where uniform flow is required. The span and length of the waverider are also fixed, therefore point 4 is only allowed to move up and down, and the maximum distance between the shockwave and upper surface curves is restricted. The 2 intermediate points of the upper surface profile curve are equally spaced along the span and only allowed to move up and down. This is done because, with changes to the shape of the shockwave profile curve, the limitations that apply for the upper

surface profile curve as mentioned in Section III change, and we wish to allow the geometry to exploit any changes in the shockwave shape. The angle of the shock is also restricted to 15 degrees. The three different parameterizations of the shockwave profile curve that were evaluated are:

- Equally spaced - One variable controlling the position of point 3 along the span, with points 1 and 2 equally spaced between 0 and 3.
- 1 independent - One variable controlling the position of point 3 along the span independently, and one more controlling the position of point 2, with only point 1 equally spaced between 0 and 2.
- All independent - Three variables controlling the position of each of the points, 1, 2, and 3.

In all cases the points are restricted to being in the same order, i.e. point 1 is always between 0 and 2, etc.

The L/D evaluation method used was the one described earlier with the flow conditions remaining the same. The objectives to be maximized are the lift to drag ratio (L/D) and volumetric efficiency (V/A). Again, 3000 evaluations were run for each parameterization using a simple genetic algorithm. The resulting Pareto fronts and two geometries close to the two ends can be seen in Figure 3.26. The distribution of designs in lift-to-drag vs volumetric efficiency landscape from which the Pareto fronts were derived for each case can be seen in Figure 3.27.
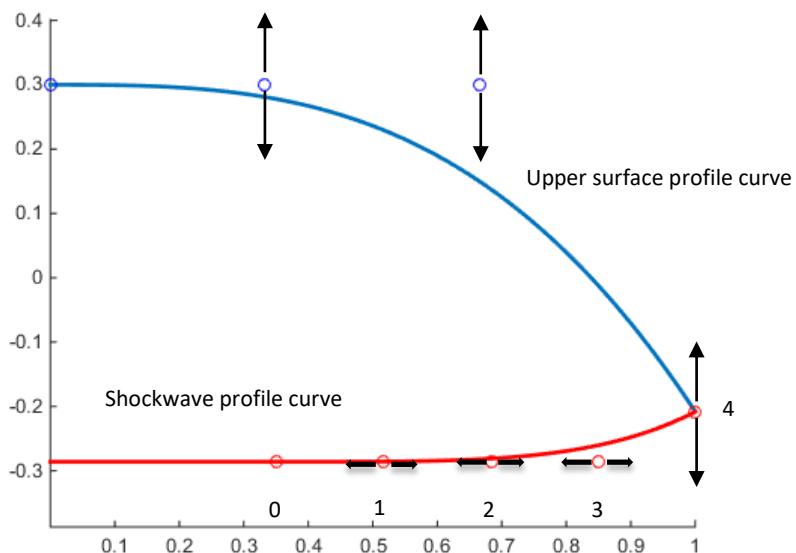
Figure 3.25. Shockwave profile curve control point schemes.
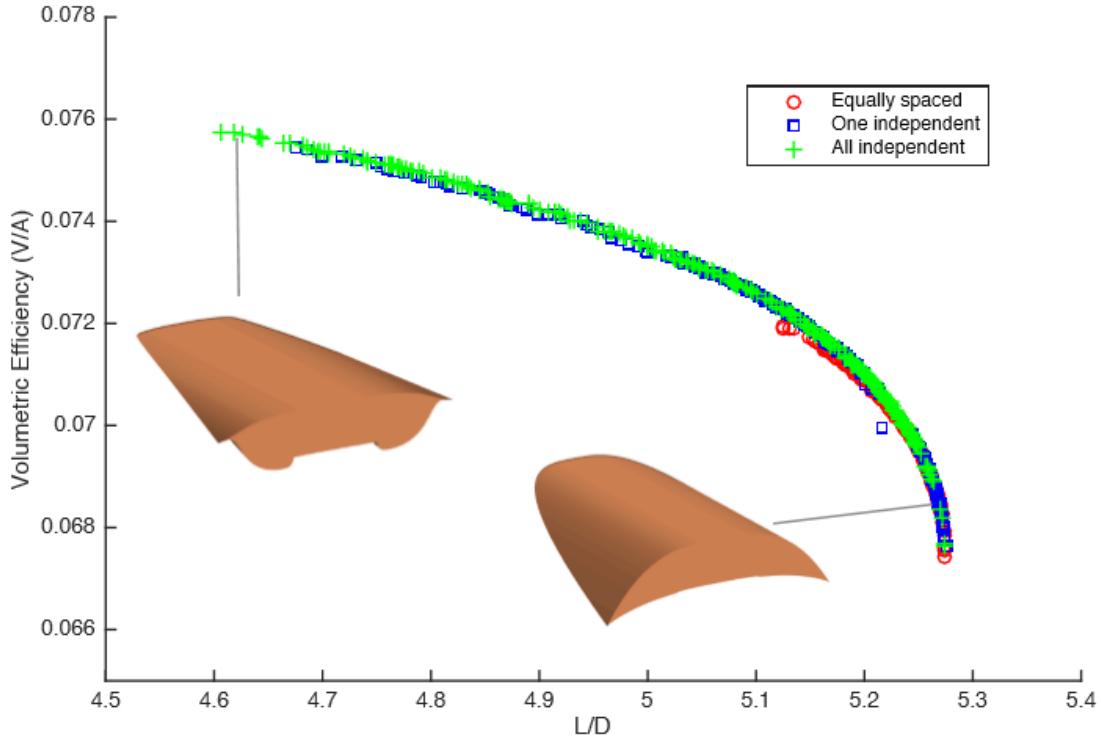
Figure 3.26. Pareto fronts for different parameterization schemes of the shockwave profile curve.
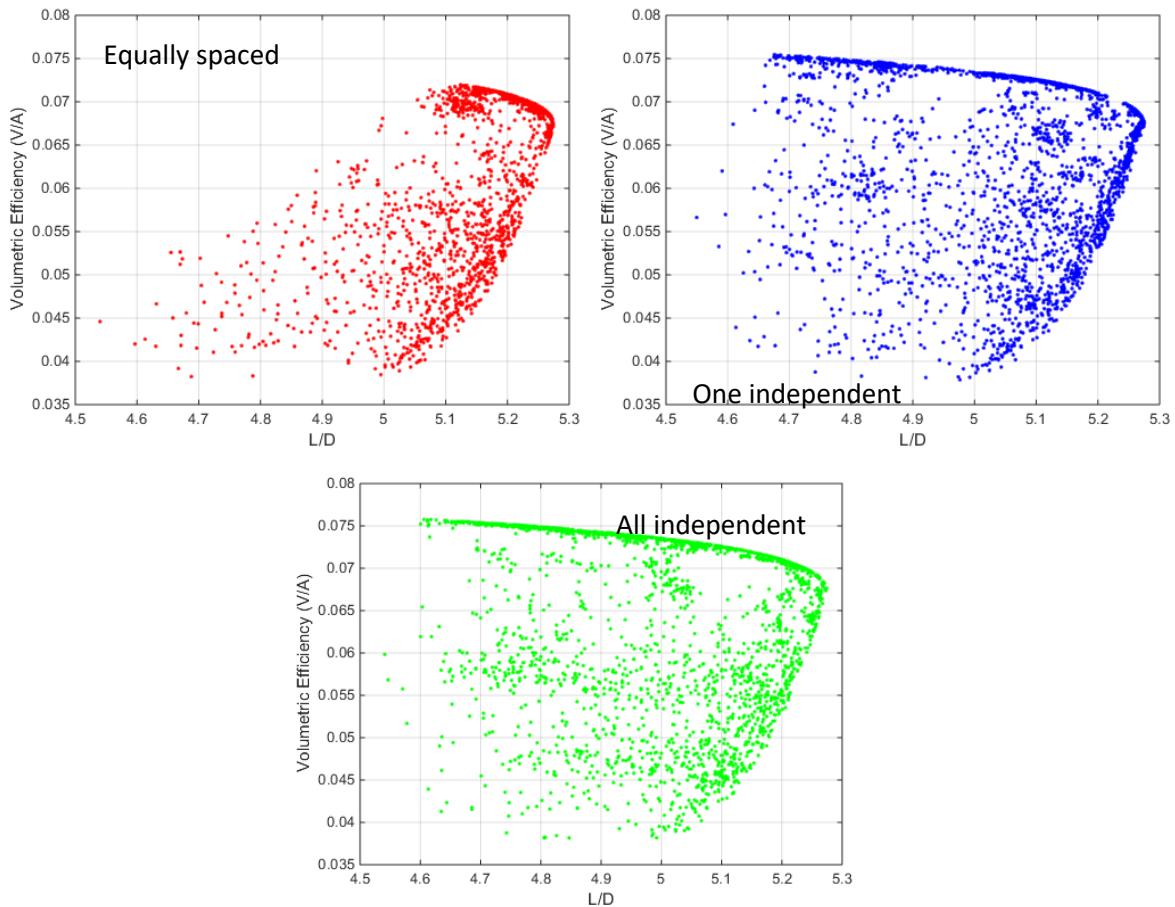
Figure 3.27. Design objective landscape for the three SWPC optimization cases.

In this case, we do not see significant advantages when increasing the flexibility of the geometry. We observe that allowing point 3 to move independently from the other two enables a wider range of Pareto optimum solutions and results in an expansion of the design space, however the improvements at the region of the first Pareto front (red circles) are very small. Having control over the position of all three points seems to yield no further advantages. Additionally, moving to the regions on the right side of the second and third case Pareto fronts can be better achieved by increasing the shock angle. Therefore, controlling all points with a single design variable would be an acceptable option, while controlling point 3 independently might also be considered, since it is also the point that controls the tangency at the end of the curve (at point 4). This is a good example of a case where a clear conclusion can be drawn: that there is no point in increasing the flexibility further. The one or two degrees of freedom that are 'saved' may not seem significant but as the dimensionality and complexity of the design studies increases, each one of those degrees of freedom exponentially adds to the difficulty of the problem.

It is important to, once again, note at this stage, that the goal of the design searches presented here is not the identification of a final optimal geometry. Instead, the objective here was to use an optimizer to exercise each parameterization, testing their flexibility in the 'good' region of the design space and testing the ability of each to access these regions. Additionally, the two cases presented are mostly meant as examples illustrating the proposed methodology. By conducting such parametric flexibility studies one can further justify the parameterization options that are used in subsequent design studies. Conclusions drawn from such studies are also meant to assist the development of more complete hypersonic aircraft parametric models, by providing information and a means to compare different parameterization options for the waverider component. While it is true that efficiently parameterizing each component through low fidelity studies will not necessarily translate to increased efficiency of the more complex parametric geometry, such studies can provide good indications on where the limits of over and under-parameterizing the geometry may lie. Therefore, such parametric flexibility studies, apart from offering a better understanding on handling such geometries, are also useful precursors of more complex design studies. They can be utilized to identify the necessary level of flexibility, or at least offer an educated guess in terms of the flexibility versus performance trade-off, before expensive resources are committed to more complex and higher fidelity design optimizations. More details on how parametric flexibility studies fit in an aerodynamic design process will be given in Chapter 5.

The current chapter has been focused on the parameterization and handling of waverider forebody geometries. We looked into different approaches to parameterizing inverse design methods and also proposed a method to assist in deciding the amount of flexibility required by their

design-driving curves, with the goal of increasing the efficiency and effectiveness of the parametric geometries. A novel waverider design method that does not need a prescribed shock shape and allows full control over the three-dimensional leading edge was also proposed. The geometry feature we turn our attention to next is the blunt leading edge, since it is the geometry feature that is most closely integrated with the forebody from a geometrical perspective.

Chapter 4:   Blunt Leading Edge Shape Parameterization

All inverse waverider design methods produce sharp leading edge, inviscid theory based geometries. One of the first issues that need to be addressed when departing from the simplicity of inviscid designs is that of the aerodynamic heating effects that, along with manufacturability constraints, preclude the use of sharp leading edges at hypersonic speeds. Two early examples of non-ballistic high-speed vehicle geometries that needed blunt leading edges are the Space Shuttle and the X-15 research aircraft. Waverider shapes however rely on the sharpness of their leading edge in order to keep the shock attached to it, and for this reason the performance of such aircraft is expected to degrade and be more sensitive to the use of finite thickness leading edges. It is, therefore, of great interest to investigate methods to optimally introduce bluntness to the leading edge of such vehicles.
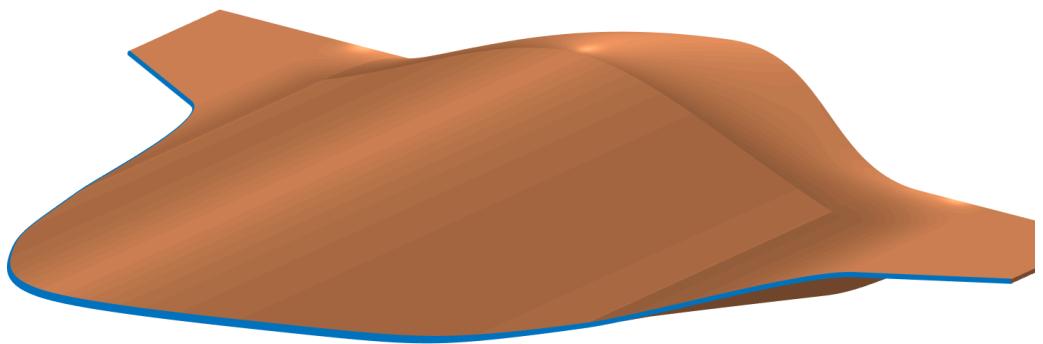
Figure 4.1.   Simple waverider-based hypersonic aircraft geometry example, with the blunt leading edge highlighted.

Over the years, not many approaches to introducing bluntness to leading edges of slender waverider bodies have been proposed. On the other hand, the shape of ballistic blunt re-entry vehicles is an area that has seen much more scientific interest. For hypersonic aircraft, most

common are circular leading edge shapes that have been widely used and studied due to their simplicity. There are, however, benefits in looking into more sophisticated shapes than a constant radius of curvature for the leading edge bluntness, and a number of studies have looked into different geometry parameterizations for this purpose. Power-law shapes have been considered in studies by O'Brien (1998; 1999) and Santos (2005) that explored potential advantages of having more control over the shape of the blunted part. As Santos reports, the shapes are not able to outperform circular ones when it comes to the combination of lowering both heating and drag, and those studies focused on the shock's stand-off distance. More recently, parametric curves able to achieve a wide variety of shapes have also been considered with promising results. Rodi (2013) incorporated Bézier curves with movable control points to control the shape of the blunted part, with initial analysis showing that for the same peak heating rate the optimized shapes showed drag values around 20% lower. Rodi (2015) went on to incorporate the optimized shapes on full 3D waverider forebodies with favourable to very small aerodynamic improvements depending on the relative size and drag contribution of the leading edge. The said Bézier curve-based parametric geometry consisted of three movable control points and required around five design variables for its parameterization. These leading edge shapes were further studied by Hinman (2017) who proposed optimization methods based on numerically efficient design evaluation tools in order to reduce the design optimization cost. Finally, ellipses and super-ellipses (modified ellipses able to achieve curvature continuity at the flat plate/wedge interface) have been utilized in receptivity studies to determine the effects of leading edge radius of curvature and curvature continuity (Lin *et al*., 1992) in turbulent transition mechanisms.

In this work, a family of novel parametric geometries for hypersonic aircraft leading edges has been developed, with the goal of providing an effective, efficient and robust geometry parameterization. It is a geometric formulation able to achieve a wide variety of shapes with desirable characteristics, while keeping the number of associated design variables low. In this chapter, we have a more detailed look into these rational Bézier curve based leading edge (RBLE) shapes, starting from their purely geometrical characteristics and how to effectively control them, to evaluating their aero-thermal performance characteristics. We back these up with CFD analyses of the geometries, where their performance is also compared to common circular leading edge shapes. Preliminary results on receptivity and turbulent transition mechanism investigations conducted to examine the effect of the increased geometric continuity of the proposed geometries are also presented. Finally, these 2D shapes are integrated onto full 3D waverider forebody geometries.

4.1 Leading Edge Geometry Parameterization

Sharp leading edges of waverider forebodies can generally be described as wedge-like geometries on the vehicle's cross sections. Such section profiles allow for a parametric curve as simple as a quadratic Bézier curve to be used to generate a curved blunt leading edge shape after truncating the original geometry. Using a quadratic Bézier curve, however, fixes the shape of the leading edge, rendering it non-parametric. The two end points would be placed at the interface of the truncated part on the original geometry, and the intermediate one is restricted in order to satisfy first order continuity at the interfaces of the curve and the original geometry. Two alternatives, similar to this approach, are either using higher order curves with more movable control points, as is the case in Rodi's (2013) work, or by using rational Bézier curves whose shape can be further adjusted with weights. The latter of the two approaches has been followed in this work and the formulations developed, with their advantages, are discussed next.

4.1.1 Quadratic Rational Bézier Leading Edge

This is perhaps the simplest alteration of the original wedge-like geometry. The original sharp wedge is truncated and part of it is replaced with a quadratic rational Bézier curve as seen in Figure 4.2. The parametric curve's end points (P0 and P2) are on the interface with the original truncated geometry, while the middle control point (P1) is fixed to satisfy continuity conditions at the interfaces since the curve is tangent to its control polygon at its end points. Varying the weight w_1 of the middle point then controls the shape, while the weights of the end points, w_0 and w_2, are fixed at a value of 1.
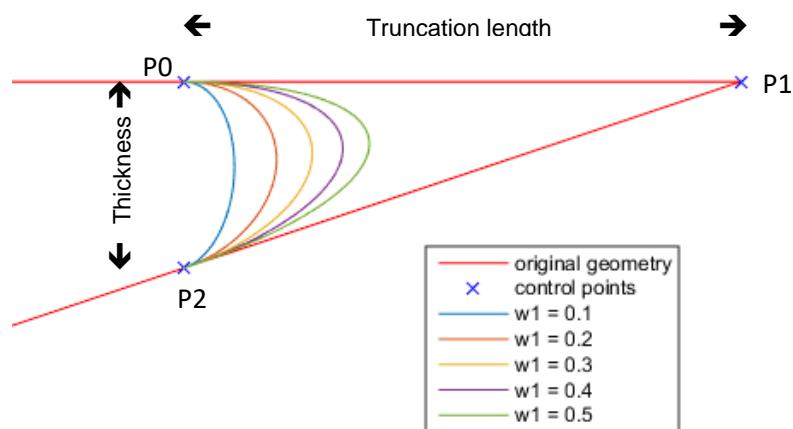
Figure 4.2. Rational quadratic Bézier curve applied to a 2D wedge-like sharp leading edge with different weights w_1.

The geometries that are expected to be of most interest and have aero-thermal characteristics that outperform the circular arc leading edges are those that have a larger radius of

curvature at the stagnation point and tighter ones closer to the edges, as this better distributes the heat load across the blunted part. An example of such a shape is the bluntest of the family shown in Figure 4.2, with a w_1 value of around 0.1.

These shapes can be controlled with two parameters, the truncation length (or the thickness at the interface) and the weight of the middle point w_1. The first can be a design variable itself since it directly controls the scale of the blunted region. However, varying the weight directly produces a non-linear variation of shapes for the range of values of w_1. This is best illustrated in Figure 4.3, where the weight is directly linked to the design variable that ranges from 0 to 1 in the first case, and equal to the square of the design variable's value in the second case. The second method offers a more intuitive variation of shapes across the design variable range, with slightly increased resolution at low values where geometries with better aero-thermal characteristics are expected.

Movement of the forward tip for every 0.1 increment of

var_1 is more linear after the reparameterization

$$w_1 = var_1 \qquad\qquad w_1 = var_1{}^2$$
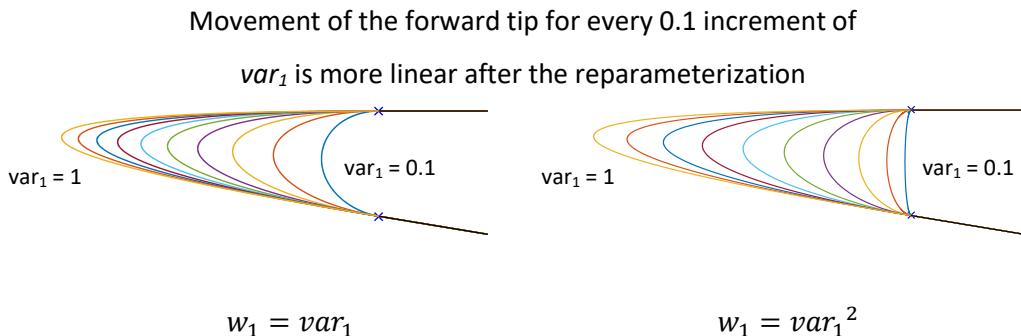
Figure 4.3.   Leading edge shape range for w_1 directly linked a design variable (left) and linked to the square of a design variable (right) ranging from 0.1 to 1.0 with equal spacing.

This reparameterization of the shape might seem insignificant for this simple case, but it becomes much more meaningful for the more complex parameterization schemes that follow.

4.1.2   Cubic Rational Bézier Leading Edge

With the quadratic RBLE shapes, a variety of what intuitively seems like meaningful shapes can be obtained. There could however be benefits in having more control over the shape of the blunted geometry by using parameterizations that provide more flexibility. A simple and elegant way to accomplish this without adding too much complexity with movable control points, is by duplicating the intermediate point. This way the order of the curve increases by one and the two intermediate control points are on top of each other. They are still fixed to satisfy continuity conditions but there are now two weights that can be controlled. The parameterization can be seen in Figure 4.4.
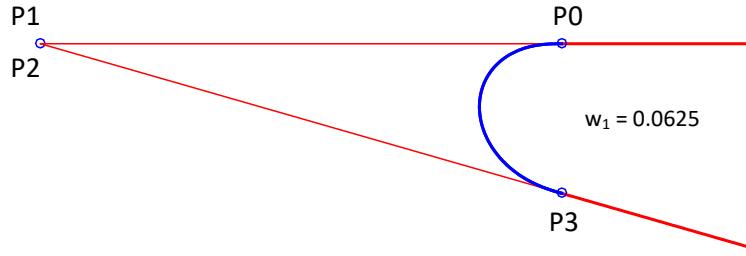
Figure 4.4.   Cubic rational Bézier curve leading edge.

Once again, for more intuitive control, a non-linear mapping between the weights and the design variables was implemented. This reparameterization was mainly developed through a trial and error approach. The square of the first variable defines the mean value of the two weights and the second one is used to balance them as follows:

$$w_1 = var_1{}^2 \left(1 + var_2\right), \tag{4.1}$$

$$w_2 = var_1{}^2 \left(1 - var_2\right),$$

with the variable ranges: $var_1 \in (0, 1]$, $var_2 \in [-1, 1]$.

The mean value of the two weights is therefore equal to the square of the first design variable which mostly has control over how far the 'tip' will extend. The second design variable balances the two weights without moving the tip of the geometry in the direction of the flow, but shifts the balance of the shape upwards or downwards. This is best demonstrated by the example geometries that are given in Figure 4.5 for various values of the design variables.
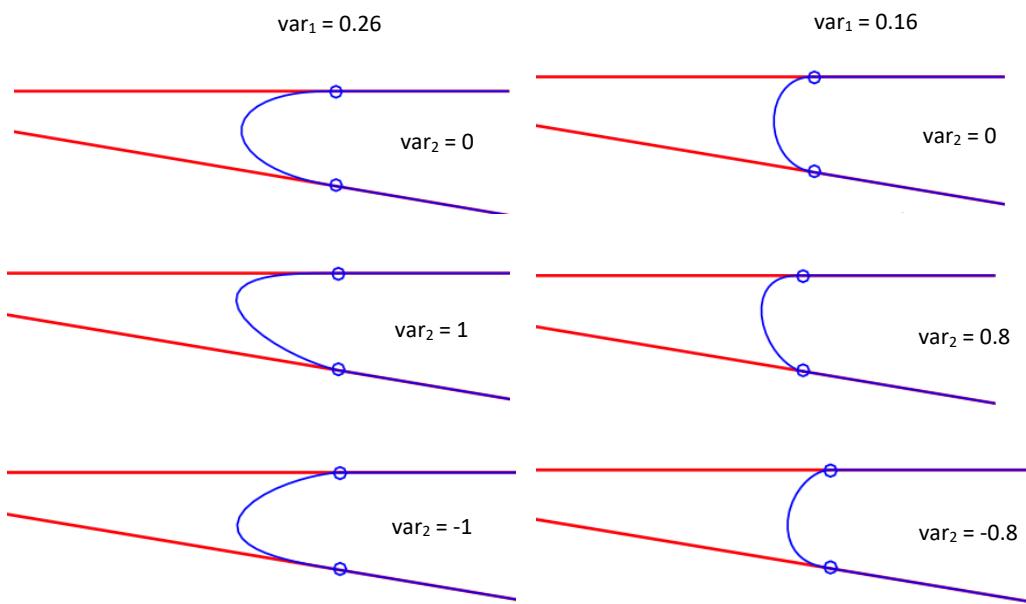
Figure 4.5.   Leading edge shape examples of the cubic rational Bézier parameterization.

Apart from the increased level of control over the shape compared to the first scheme that was described, another advantage is the increased level of geometric continuity that can be achieved at the interface between the original wedge-like geometry and the parametric curve. Almost all the approaches that have been considered so far, from the simple circular leading edge shapes to power law shapes and even parametric curve shapes such as the quadratic RBLE described earlier, only achieve first order geometric continuity (G^1) at the interface of the blunted part. To the human eye this seems smooth enough. However, the flow properties can be sensitive to the jump in the radius of curvature of the surface, potentially affecting the receptivity and turbulent transition mechanisms. The cubic RBLE shapes that are presented here, however, can achieve continuity of the radius of curvature (G^2) at the interface of the blunted part.

The radius of curvature at the endpoints of a rational Bézier curve can be calculated as follows (Sederberg, 2012):

$$\rho(t_0) = \frac{w_1^2}{w_0 w_2} \frac{n}{n-1} \frac{a^2}{h}. \tag{4.2}$$
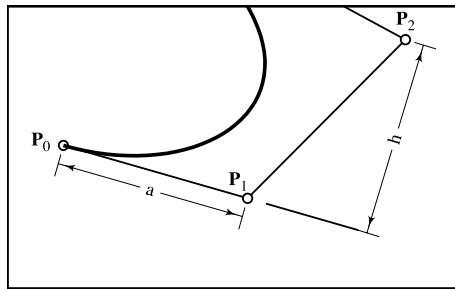
Figure 4.6.   Features for calculation of radius of curvature at end point of rational Bézier curve.

When the initial sharp leading edge is a wedge shape consisting of straight lines of infinite radius of curvature, the parametric shape described earlier matches that infinite radius of curvature at the end points, as $h = 0$ for both ends of the curve. The numerically calculated radius of curvature at stations along the curve of a cubic RBLE geometry is shown in Figure 5.
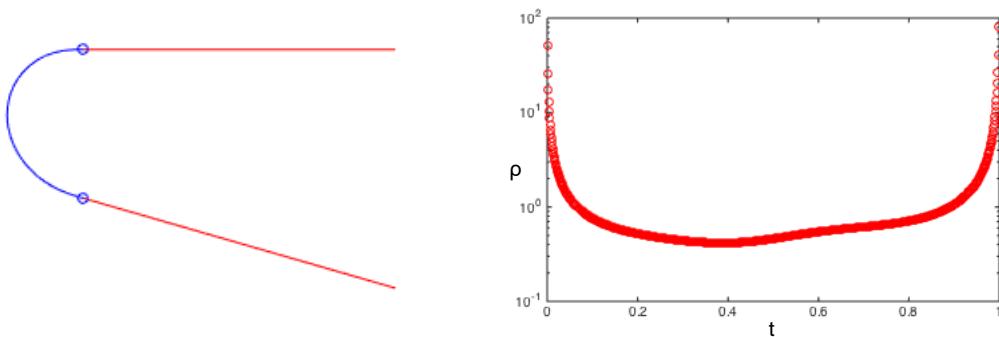
Figure 4.7.   Radius of curvature along cubic rational Bézier leading edge.

It is also possible to match relatively large (compared to the thickness of the blunted region) radii of curvature of the original geometry at the interface, even when the radius of curvature is finite. That is usually the case for the lower surface of waverider geometries or after morphing and further shaping of the upper surface. If, for example, the radius of curvature is finite only on the lower surface, one of the control points can be moved to a position with $h \neq 0$ with respect to one end of the curve. An example is shown in Figure 4.8, where P2 can move along the line defined by P1 and P3 to obtain the desired radius of curvature at P0 while maintaining an infinite radius of curvature at P3. The parameterization at this point becomes more complex as the control point needs to move according to the values of the weights to keep the radius of curvature at P0 constant.
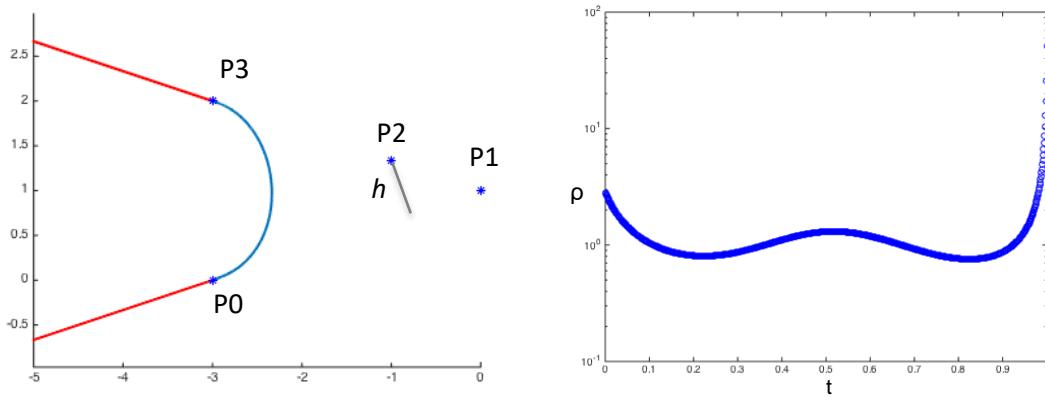
Figure 4.8.   Cubic rational Bézier leading edge with a finite radius of curvature on one end (P0).

4.1.3   Fourth Order Rational Bézier Leading Edge

To achieve even greater flexibility, a third intermediate control point can also be placed on top of the two of the previous method. The order of the curve is now four, and the number of adjustable weights for the parameterization increases to three. The addition of a third adjustable weight opens up a wider range of potential shapes that can now be generated. The weights w_1, w_2 and w_3, control the top, middle and lower part of the shape respectively (when the control points are ordered as seen in Figure 4.9). Balancing them to obtain meaningful shapes for a range of design variable values, though, is a little more delicate than for the previous two approaches. The relationships between the design variables and weights were chosen here as follows:

$$w_1 = var_1{}^2 \left(1 + \frac{1}{3} var_2 \right),$$

$$w_3 = var_1{}^2 \left(1 - \frac{1}{3} var_2 \right), \tag{4.3}$$

$$w_2 = -var_1{}^2 + 0.1 \, var_3 \,,$$

with the variable ranges: $var_1 \in (0, 1]$, $var_2 \in [-1, 1]$, $var_3 \in [0, 1]$.
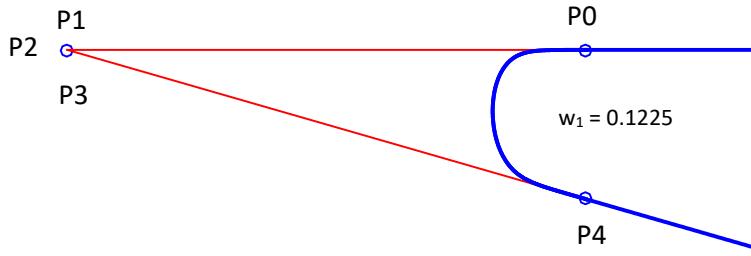
Figure 4.9.   Fourth order rational Bézier curve leading edge.

With the above reparameterization each of the variables is given a more specific role. The first and the second ones, similar to the cubic RBLE case, control the mean value of the weights at the sides, w_1 and w_3, and the balance between the two respectively. The weight of the middle point w_2 is balanced with the third design variable and controls how blunt the middle part of the shape is, with its lowest value being roughly the limit where the geometry starts becoming concave. The end result is a shape parameterization where each design variable has a specific role and the characteristics it gives to the final geometry are retained when the other design variables are altered. This is best illustrated in Figure 4.10 where a number of example geometries are given for different design variable values.
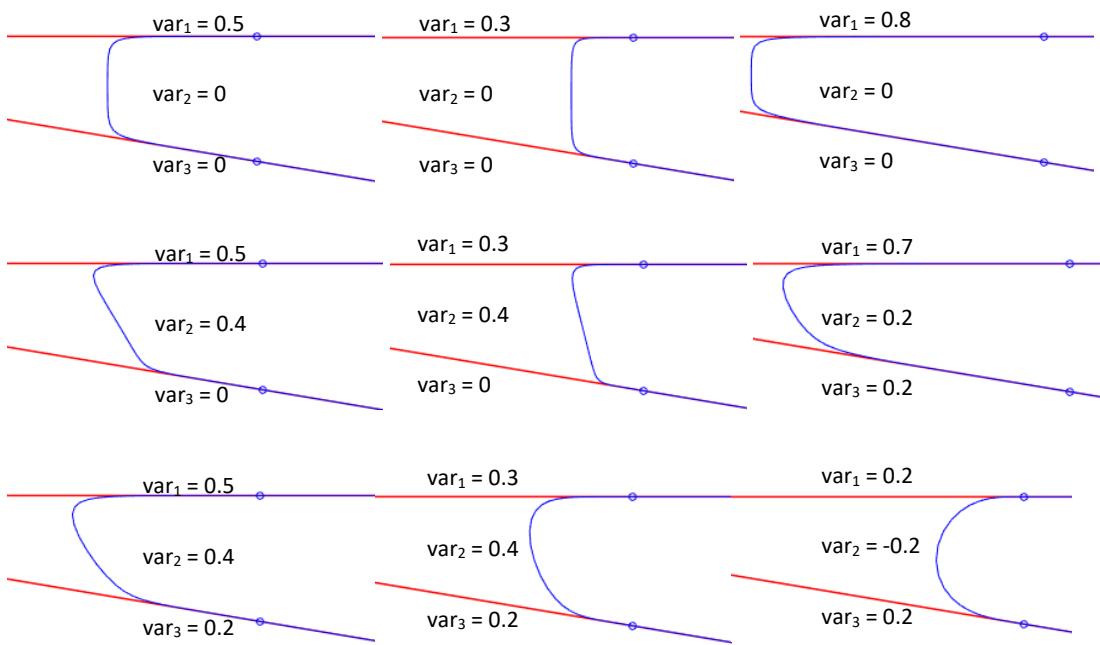
Figure 4.10. Leading edge shape examples of the fourth order RBLE parameterization.

The fourth order RBLEs, when used on a wedge geometry consisting of straight lines, theoretically provide third order (G^3) geometric continuity at the interface, though in practice this is of limited significance in both geometrical and flow dynamics terms. Effectively though, it can provide an even smoother transition to zero curvature or infinite radius of curvature when compared with the cubic rational Bézier parameterization of Section 4.1.2, as seen in Figure 4.11.
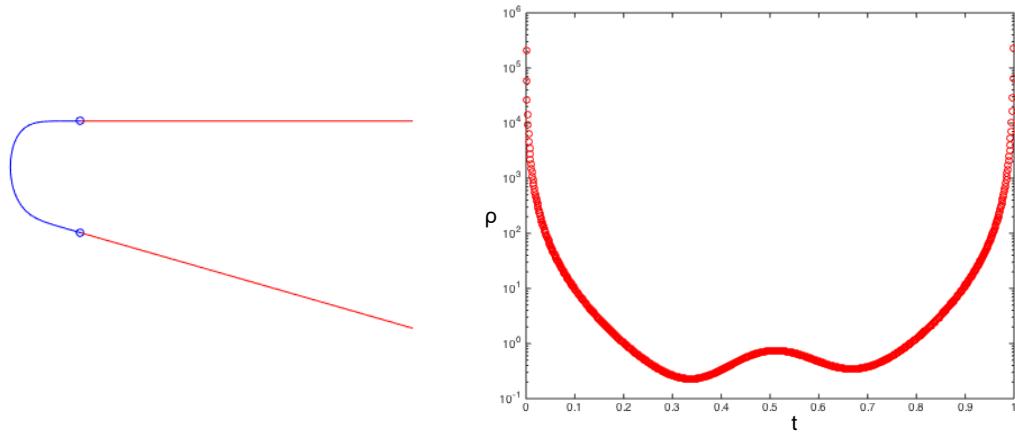
Figure 4.11. Radius of curvature across fourth order rational Bézier leading edge.

The variety of shapes obtained with this last approach is significantly wider, while we also see that, even for a fixed value of the initial truncation length or thickness at the interface, the parameterization is able to drive itself to shapes of seemingly lower thickness. It should be noted that because balancing the weights is more delicate in this case, it is harder to match finite radii of curvature at the end points as was done for the previous approach. It can be attempted by moving P2 accordingly, but careful reparameterization of the weights and design variables is required, something that will not be further investigated at this stage.

The extent to which it is beneficial to have this much control over the shape of the leading edge will be examined in more detail in the following sections. As far as manufacturability of such shapes goes, it is something that strongly depends on the scale of the bluntness needed. Finally, from now on each shape will be described only by the values of the design variables as they were defined in this section, and not the values of the weights.

4.2   2D CFD Analysis of Shapes

The ultimate goal of being able to look into a variety of different shapes for the blunt leading edges of hypersonic aircraft is to enable the design to get the best out of the tradeoff between the aerodynamic performance characteristics and the limits of the materials and thermal protection system. Leading edges with increased bluntness are required to cope with the increasing aerodynamic heating of higher velocities. As the bluntness increases so does the drag, while the lift can decrease due to the high-pressure flow at the underside of waverider shapes no longer being completely 'trapped', with the shock standing further away from the tip of the leading edge. The shape of the leading edge essentially affects the flowfield around the entire aircraft. As the bluntness-induced drag and peak temperatures at the leading edge are something that can be examined locally, a series of 2D CFD simulations to study the local aero-thermal properties of the shapes described earlier were initially conducted.

4.2.1   Cold Wall Simulations

ANSYS Fluent was used for the simulations and a number of different case setups were considered. The first set of cases was set up with a cold wall (T_w=300K) and flow conditions matching Mach 8 flight at 100,000ft (p_{inf} = 1090Pa, T_{inf} = 227K). Similar investigations were done for Mach 6 and varying dynamic pressures as well, however, the majority of the results presented here are for the Mach 8 case mentioned earlier unless otherwise stated. For all cases laminar flow was assumed around this part of the leading edge and therefore no turbulence model was used. The two values of interest were the drag and the peak heat flux. Grid independence was achieved with a 500x120 (along wall x wall-normal) initial structured C-type mesh, and three levels of adaptive mesh refinement in areas of high pressure gradient for better shock capturing. The peak *y+* value across the surface of the leading edges was around 2. From this point we will consider a circular leading edge with radius R to be equivalent to the parametric shapes described earlier with thickness 2R for convenience, as a thickness cannot be defined the same way for circular leading edges due to the two interface points with the wedge not being in the same streamwise position.

In Figure 4.12 we can see how the drag coefficient and peak heat flux vary for a 1cm thick quadratic RBLE, for a range of 0.15 to 0.3 of the shape variable. The equivalent circular leading edge data point is also plotted. The wedge angle for this and most of the geometries that are presented was 9.5°, which is the deflection angle that results in a ~15° shock angle under the given flow conditions. All drag coefficient values are for the geometries extended downstream up to the point where the wedge thickness is 1.67cm.

In Figure 4.12 we observe that, first of all, the parametric shapes can come close to the circular leading edge both in geometry and performance characteristics. As the geometry gets more blunt the drag increases while the peak heat flux at the stagnation region decreases and can reach values up to ~20% lower than the equivalent circular leading edge. For geometries that become too blunt the heat flux around the sharply curved part close to the interface overcomes that of the stagnation point and therefore the peak heat flux increases again. The way to utilize these geometries if, for example, the initial goal was to replace the equivalent 0.5cm radius circular leading edge, is to adjust their thickness as well. This will enable the shape to have a lower drag coefficient for the same peak heat flux, something that will be demonstrated later on. It is also worth noting that, for cold wall cases of the same wall temperature, we observe the same pattern between the characteristics of the geometries for different flow conditions; this was also observed by Rodi (2015). An example of the distribution obtained for a Mach 6, 75,000ft altitude case is shown in Figure 4.13.
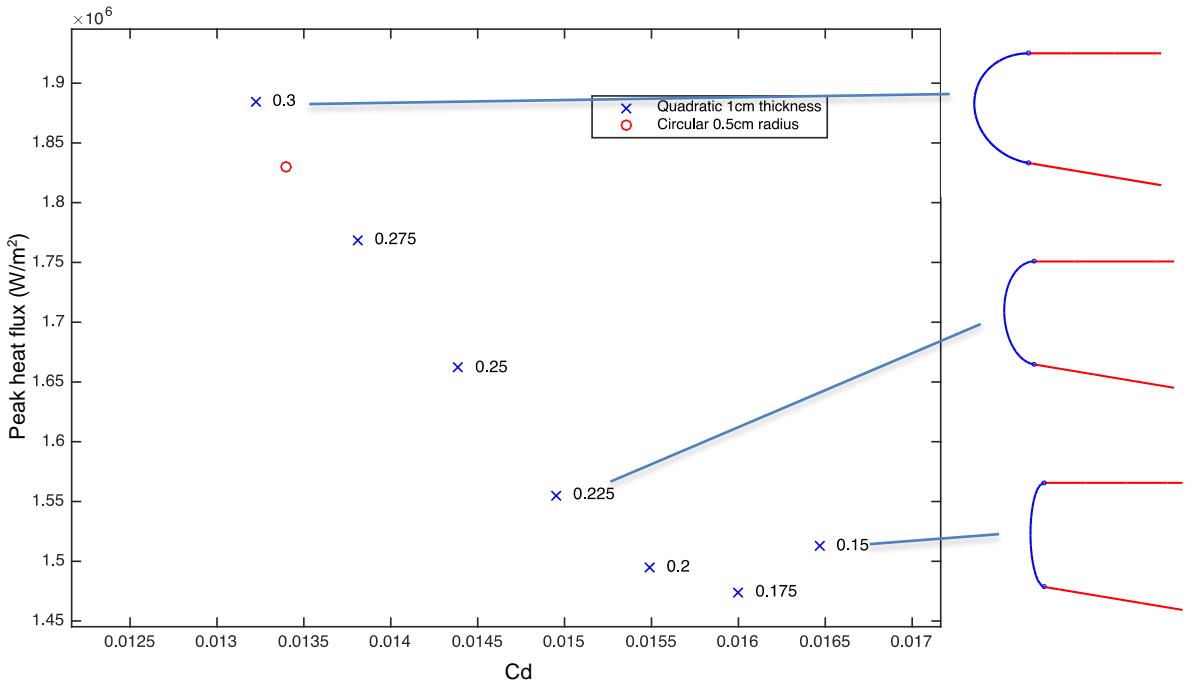
Figure 4.12. Drag coefficient and peak heat flux of a quadratic rational Bézier curve of 1cm thickness, T_w=300K, M_{inf}=8, p_{inf}=1090Pa, T_{inf}=227K.
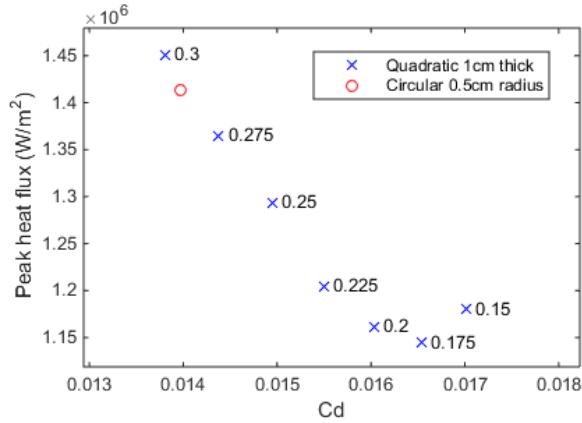
Figure 4.13. Drag coefficient and peak heat flux for T_w=300K, M_{inf}=6, p_{inf}=3498Pa, T_{inf}=219.5K.

To more quickly obtain the sets of optimal geometries (Pareto fronts) of shapes controlled by multiple design variables such as the cubic and 4[th] order RBLE parameterizations, meta-models were constructed using the kriging approach (Jones, 1998). Kriging is an interpolation method with interpolated values modeled using a Gaussian process. The meta-models were based on a latin hypercube-obtained (McKay, 1979) initial sample of 20 CFD calculations for the 2 design variable controlled cubic RBLEs and 40 for the 3 design variable controlled 4[th] order RBLEs. The meta-models can then be used to do a more exhaustive search of the design space and construct the Pareto fronts.
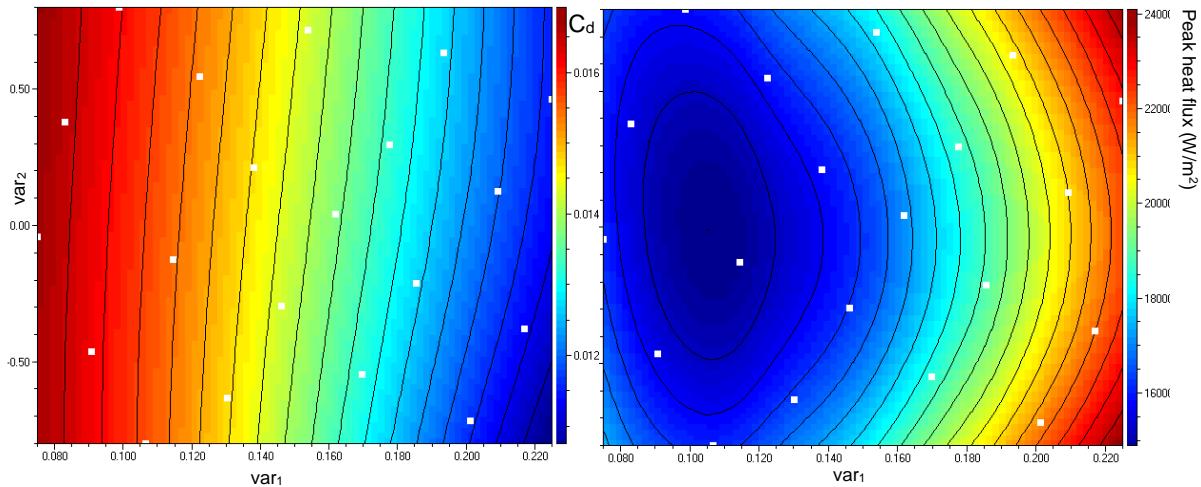
Figure 4.14. Contours of the kriging meta-model for drag and peak heat flux of the cubic RBLE with the initial sample as white squares, generated using QstatLab.

The simple form of the objective function landscape means that values obtained using the meta-model with the given initial sample gives very good approximations of the values obtained from CFD calculations, although their accuracy can be further improved with more data points. Contours of the meta-models can be seen in Figure 4.14, along with the initial sample points that were calculated. Utilizing the kriging meta-models we are then able to more densely populate the Pareto fronts of the 2 and 3 design variable controlled parameterizations using a brute force approach, since evaluation of the kriging meta-models requires negligible computational effort. The result can be seen in Figure 4.15.
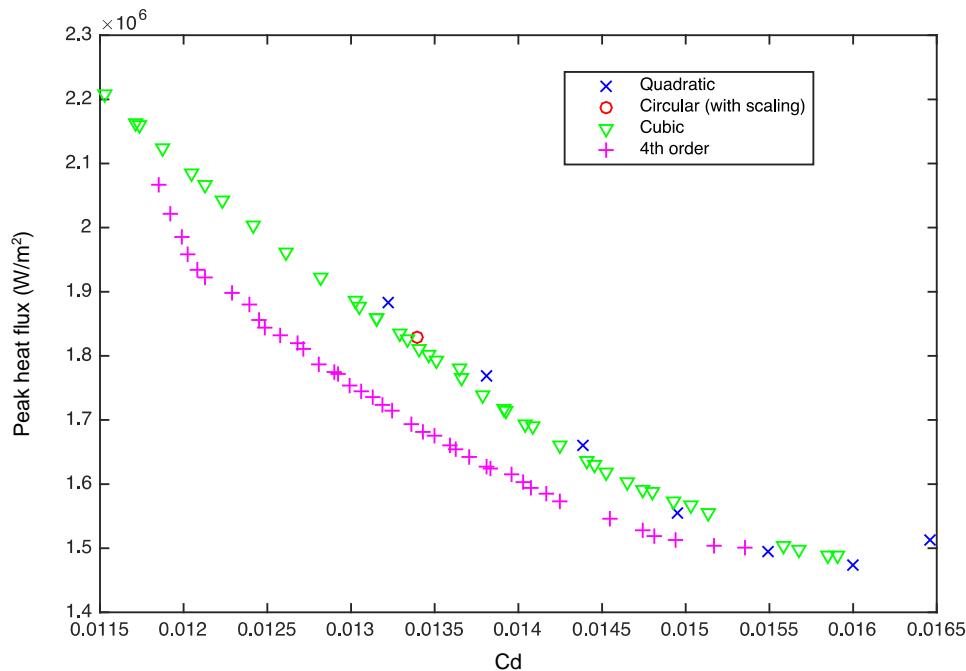
Figure 4.15. Sets of optimal solutions for the three different RBLE parameterizations, T_w=300K, M_{inf}=8, p_{inf}=1090Pa, T_{inf}=227K.

In Figure 4.15 we observe that while the solutions given by the more complex 4th order RBLE shapes have a clear advantage over the quadratic ones, the cubic ones seem to be performing very similarly to the latter. Additionally, for the cubic RBLE all the optimal solutions seem to consist of a fixed value of the second design variable (the one that shifts the balance to the upper or lower part), with variations of the first variable populating the Pareto front. Therefore, in case the cubic RBLE second design variable is fixed, the complexity of the parameterization would not be higher than the quadratic one, while, as was explained in Section II, it is a formulation that can provide G^2 continuity of the geometry.

Now, to provide a more direct comparison between the parameterized shapes and the simple circular leading edge, the drag and peak heat flux of circular leading edges of varying radii and quadratic RBLE shapes of variable thickness will be compared. In this case the quadratic RBLE is now controlled with two design variables, the thickness and one shape parameter (var_1), therefore a similar use of meta-models was followed for this comparison as well. The set of optimal solutions can be seen in Figure 4.16, where it is also compared with circular leading edges. It is worth noting that the Pareto front was dominated by geometries whose shape parameter (var_1) had a value of around 0.182, with only the thickness varying along the distribution.
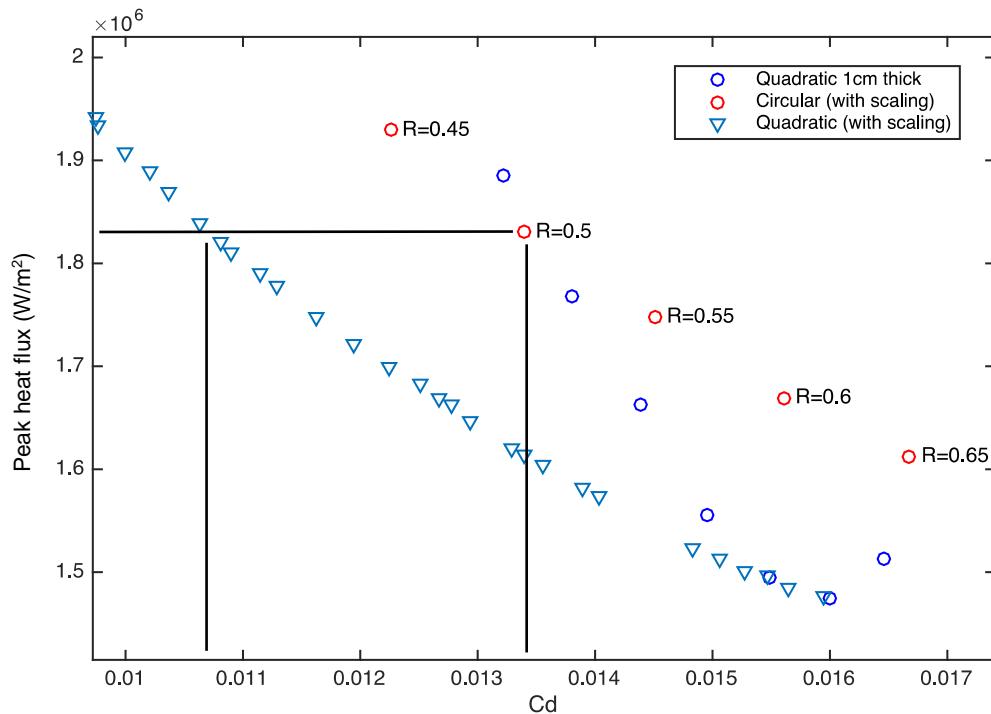
Figure 4.16. Circular leading edge and quadratic RBLE comparison, T_w=300K, M_{inf}=8, p_{inf}=1090Pa, T_{inf}=227K.

What we also see in Figure 4.16, is that a circular leading edge with a radius of 0.5cm can be replaced with a thinner quadratic RBLE with the same peak heat flux that will have a drag coefficient

around 20% lower. This RBLE would be ~30% thinner, i.e. 0.7cm thick in this case. With an equivalent process, the same circular leading edge can also be replaced with an RBLE that demonstrates the same drag coefficient while having a ~12% lower peak heat flux. Finally, in Figure 4.17 we see a comparison of the heat flux distribution around the blunt part of the geometry of the quadratic RBLE, with *var₁=0.182*, that dominates the aforementioned Pareto front and the equivalent circular leading edge.
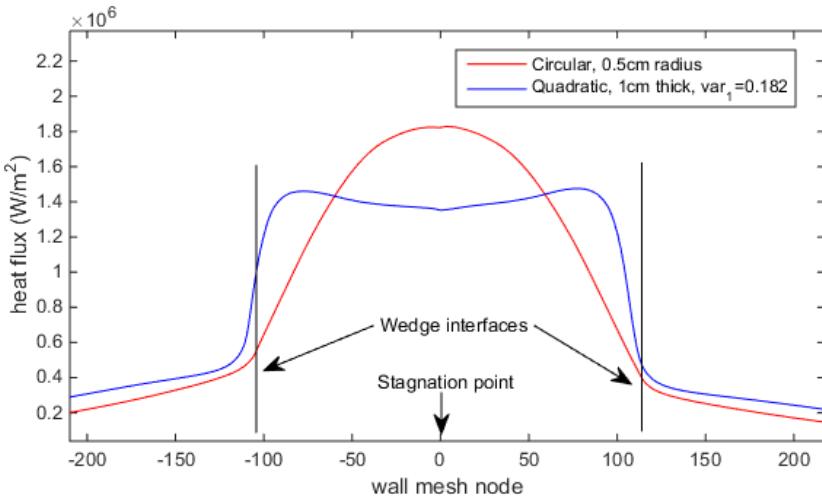
Figure 4.17. Heat flux distribution for the dominant quadratic RBLE and the equivalent circular leading edge, T_w=300K, M_{inf}=8, p_{inf}=1090Pa, T_{inf}=227K.

4.2.2   Equilibrium Temperature Simulations

While peak heat flux calculations have been used in studies since in many cases they can be compared with experimental data, they are not the best representation of a thermal protection system meant for sustained flight. The next series of simulations were set up to calculate equilibrium temperatures along the wall, using conditions that are more representative of the operation of a thermal protection system intended for hypersonic cruising. A first set of cases was run with a radiative equilibrium temperature condition at the wall, and a second set with conductive heat transfer within the solid leading edge tip as well. The flow conditions in the results presented were again set to match Mach 8 flight at 100,000ft. The emissivity of the wall is set to 0.9, which is a relatively high value but would be desirable for a thermal protection system, and any radiation originating from the hot gases around the geometry is neglected, so only the wall radiates energy. For the second set of cases the same characteristics were utilized but head conduction through the solid leading edge was also modeled. A relatively high thermal conductivity of 200W/mK was assumed for the heat conduction calculations within the solid leading edge. The values of interest are now the drag coefficient and the peak wall temperature.
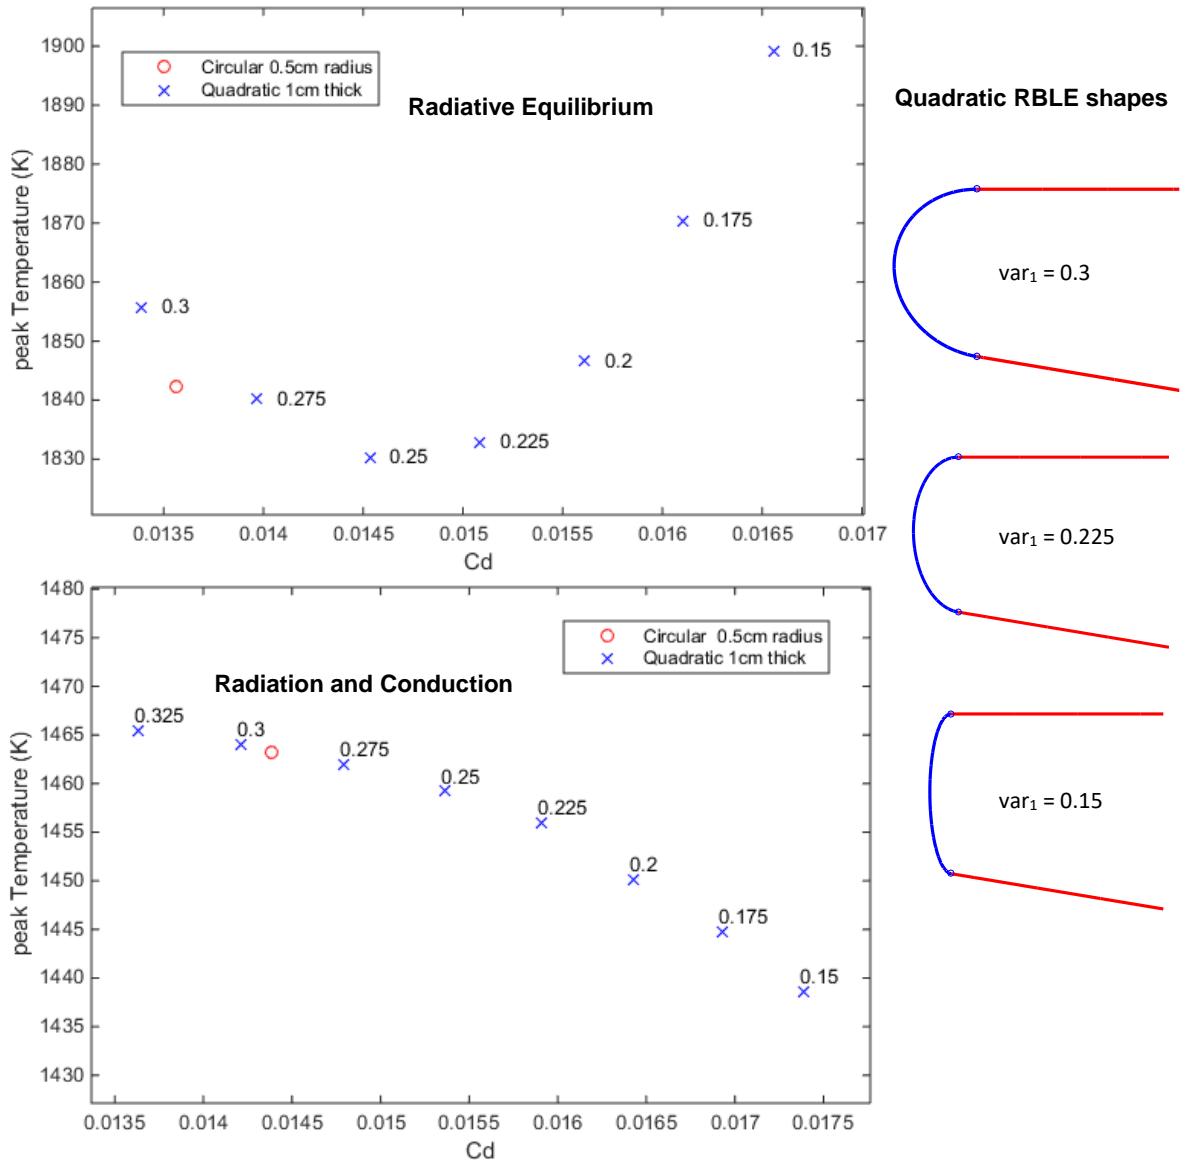
Figure 4.18. Peak temperature and drag coefficient for radiative equilibrium temperature conditions (top) and radiative and convective equilibrium conditions (bottom) for different values of var_1, $\varepsilon=0.9$, $k=200W/mK$, $M_{inf}=8$, $p_{inf}=1090Pa$, $T_{inf}=227K$.

In Figure 4.18 we can see the drag coefficient and peak temperatures that were calculated for the different shapes for both case setups. The results of Figure 4.18 are for a set of quadratic RBLE shapes for values of the design variable (var_1) ranging from 0.15 to 0.325 and a thickness of 1cm, also compared with the equivalent (0.5cm radius) circular leading edge.

A first observation is that the distributions shift significantly when radiation and thermal conduction is accounted for, while different distributions were also observed for different values of the thermal conductivity of the leading edge material. Even for the radiative equilibrium temperature set of cases, that are equivalent to a case with 0 thermal conductivity, the shape that shows the lowest peak temperature is different from the one where the lowest peak heat flux was

observed (see Figure 4.12). Moreover, as expected, the range of temperatures around the leading edge becomes narrower with increasing thermal conductivity, which significantly drives down the peak temperatures, illustrated in Figure 4.19.
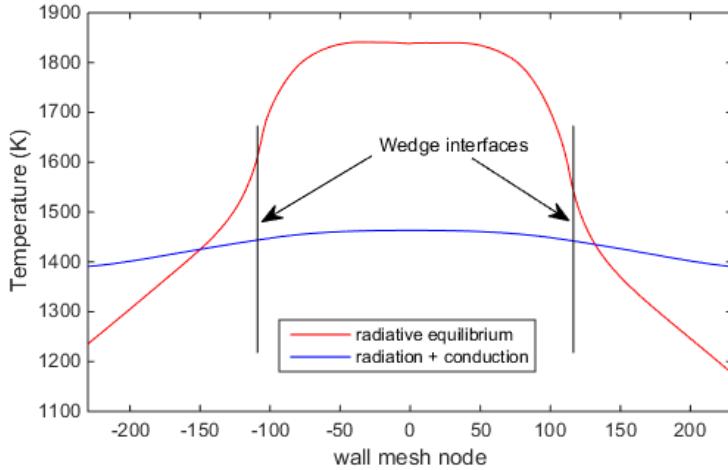
Figure 4.19. Temperature distributions around circular leading edge for radiative equilibrium conditions and radiation with heat conduction, $\varepsilon=0.9$, k=200W/mK, M_{inf}=8, p_{inf}=1090Pa, T_{inf}=227K.

In Figure 4.20 we see the heat flux and temperature around the surface of the blunt leading edge, and also the temperature distribution within the solid leading edge, colored on the computational mesh used for the thermal conduction calculations. We observe that around the blunted part of the leading edge there is a positive heat flux, meaning heat being transferred to the leading edge, while as we move downstream the heat flux is negative with heat being taken away from the leading edge. This is due to the temperature in that area being higher that the radiative equilibrium temperature due to the conduction of heat from the hotter blunt region. The result of this is lower peak temperatures at the blunted region of the shape, and higher temperatures downstream.
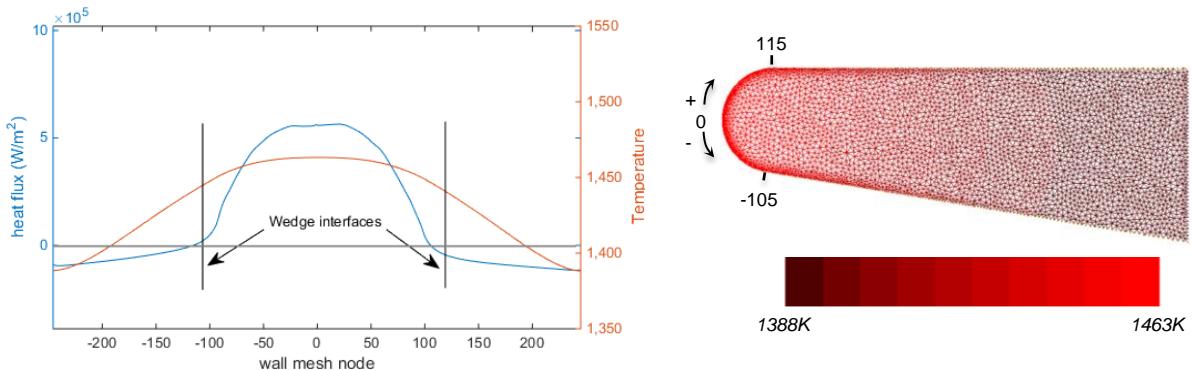
Figure 4.20. Heat flux and temperature around circular leading edge (left) and temperature distribution within the solid (right, mesh visible), $\varepsilon=0.9$, k=200W/mK, M_{inf}=8, p_{inf}=1090Pa, T_{inf}=227K.

Apart from being interesting from an aerodynamics perspective, these observations indicate that the optimal shapes strongly depend on the specific conditions of each case, with the differences in the distributions seen in Figure 4.18 and Figure 4.12 supporting that. To elaborate: even if for cold wall conditions a Mach number independence has been observed when it comes to the optimum shapes, it is the nature (active, passive etc.) and specific characteristics of the thermal protection system that will direct any design process on a case by case basis. That is where such shapes, efficiently parameterized with a low number of design variables and able to achieve a wide range of meaningful shapes, assist in making higher fidelity and multi-physics design studies affordable. The shape parameterizations introduced seem to be able to match or outperform equivalent shape parameterizations found in relevant literature, while also utilizing a lower number of design variables. The increased efficiency renders them ideal for use in more complex hypersonic design cases where more sophisticated blunt leading edge shapes are to be explored.

4.3   Preliminary Receptivity – Transition Investigation

The boundary layer turbulent transition can have a profound effect on the design and performance of a hypersonic vehicle. With the leading edge being where the boundary layer starts developing, even small changes in its shape can potentially affect the transition mechanism. In this section, some preliminary results of investigations into the effect of using a geometry with increased geometric continuity at the wedge interface are presented. These simulations were run by Adriano Cerminara, a fellow PhD student at the University of Southampton, using the direct numerical simulation code he was developing. Following the author's wish to investigate the receptivity of the proposed geometries that offer second order continuity, the author generated the geometries and surface meshes. The DNS runs were then conducted by Adrinano Cerminara to compliment the main parametric geometry investigations presented in the paper that is the basis of this chapter (Kontogiannis *et al.*, 2015). The numerical simulation, aimed at studying the leading edge receptivity process to fast acoustic waves, was carried out using a fourth-order accurate direct numerical simulation code (Cerminara and Sandham, 2015). The flow conditions are hypersonic flow at Mach 7.3 over a 20^o half-angle blunt wedge and unit Reynolds number of 4.4x10^6, with acoustic waves being inserted in the domain with 10 frequencies ranging from 50 kHz to 500 kHz and amplitude equal to 1.0x10^{-4} relative to the freestream density.

A 0.1mm radius circular leading edge shape is compared with an equivalent fourth order rational Bézier curve leading edge shape. The latter leading edge shape, as described in Section 4.1, provides second order geometric continuity of the wall geometry at the wedge junction. The specific geometry was designed to be roughly close to the geometry of a circular leading edge in the area around the stagnation point, attempting to keep an almost constant radius of curvature in that

section, and a smoother transition to infinite radius of curvature at the interface with the wedge. The main reason for this is that the investigation is intended to study the effect of the increased geometric continuity and not just a different geometry. The two geometries can be seen in Figure 4.21, while the radius of curvature along each curve is given in Figure 4.22.
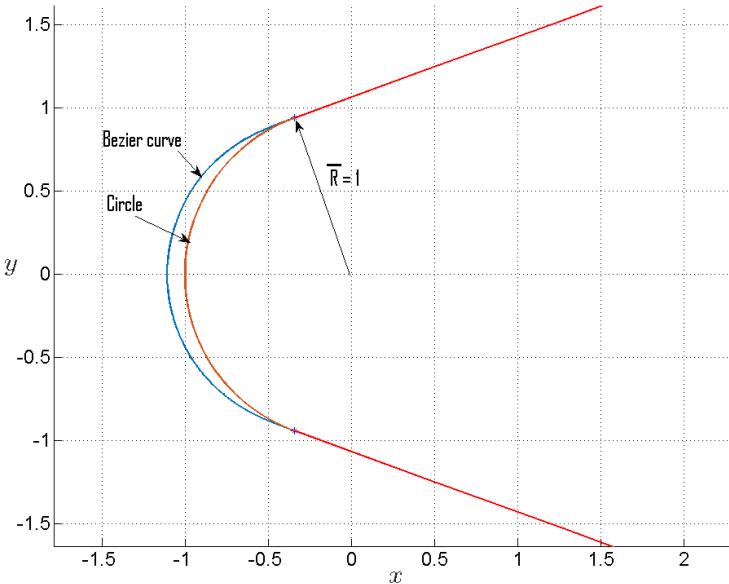
Figure 4.21. Sketch of the two leading edge geometries in non-dimensional coordinates: circle (red line), and rational Bézier curve (blue line).
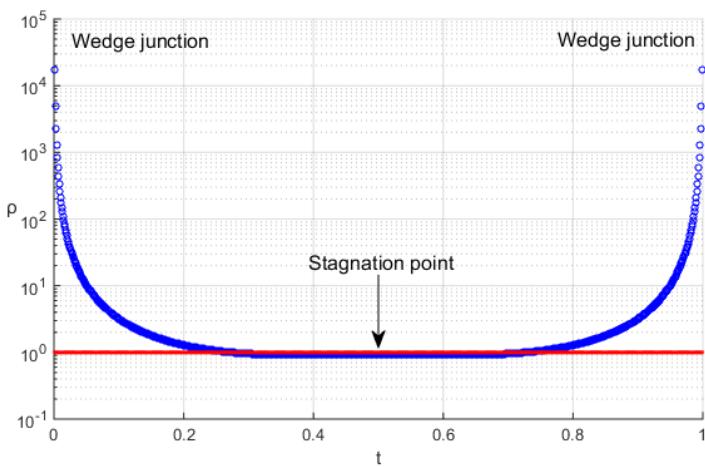
Figure 4.22. Numerically calculated radius of curvature along the leading edge curves: circle (red), rational Bézier curve (blue).

The smoother transition of the radius of curvature from the curved part to the wedge, results in a smoother pressure distribution around the interface as can be seen in Figure 4.23. A snapshot of the density fluctuation field that gives some further insight into the receptivity process is shown in Figure 4.24, this pattern is qualitatively in good agreement with results found in the literature (e.g. Kara *et al.*, 2007).
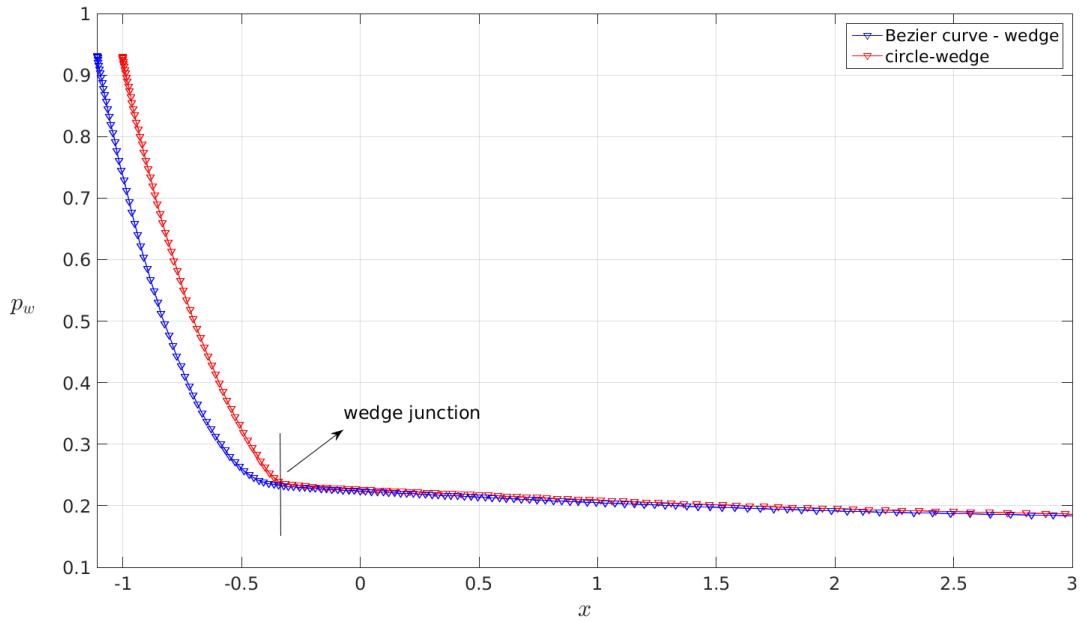
Figure 4.23. Mean pressure profiles along the wall (stagnation point of RBLE shape is ~0.1 R upstream of the circular one).
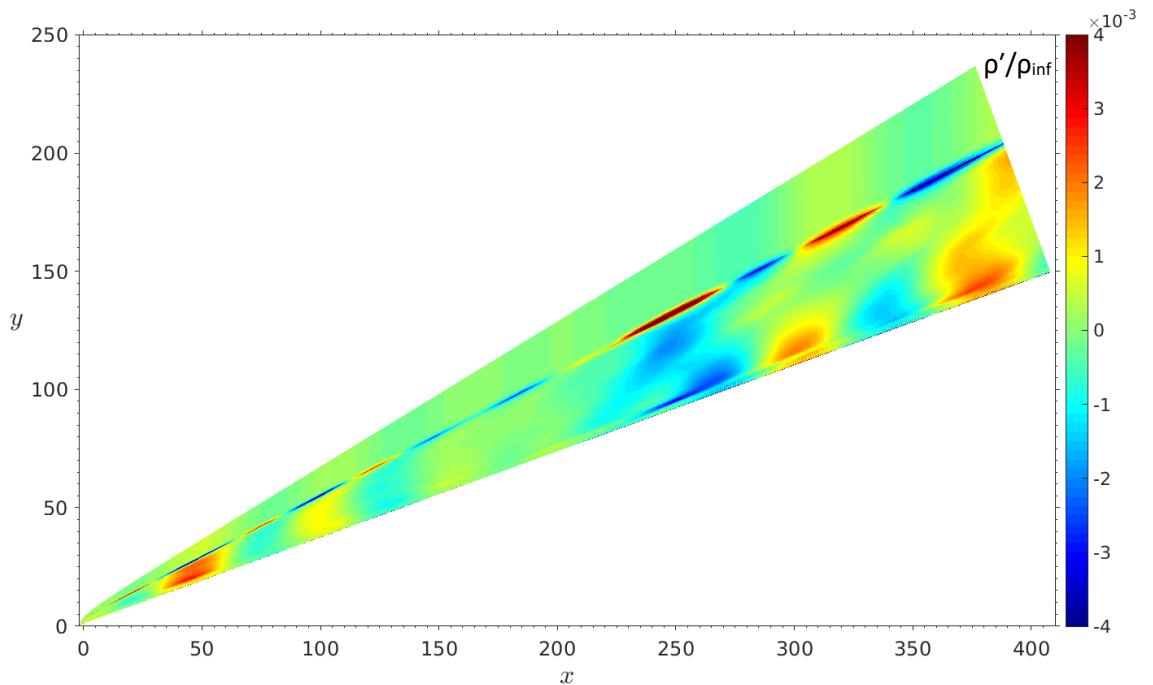
Figure 4.24. Instantaneous density fluctuation field (ρ'/ρ_{inf}).

Finally, Figure 4.25 shows a comparison between the two cases at three different points along the wedge, corresponding to the x non-dimensional coordinate values x=297.3, x=319.4 and x=347.5. Here, the normalized pressure and heat flux fluctuation amplitude spectra, computed with a Fast Fourier Transform approach, are presented on the three different positions on the wall. In particular, the pressure fluctuation amplitudes are given at all three positions, while the heat flux

fluctuation amplitude spectrum is given only at the position x=347.5 (see Cerminara and Sandham, 2015, for more details).
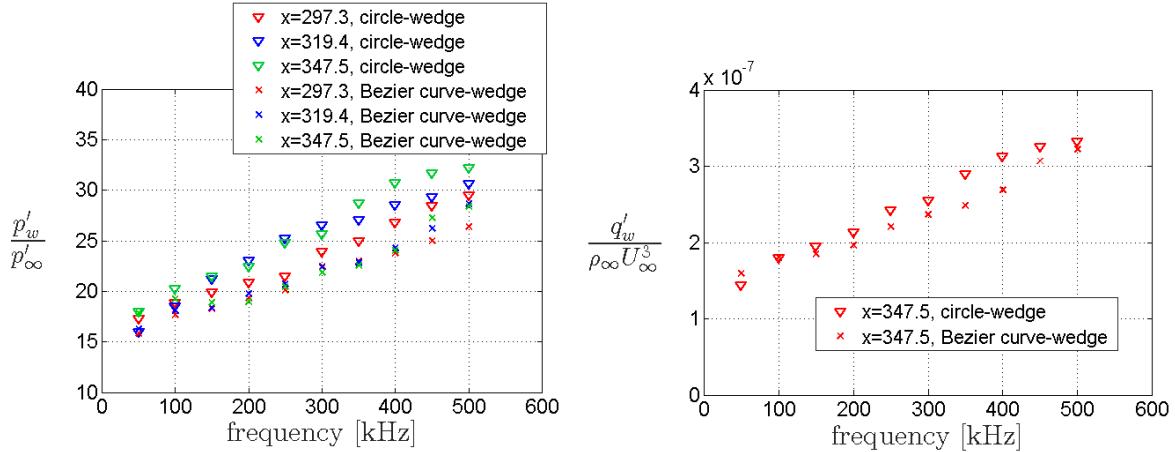
Figure 4.25. Pressure (left) and heat flux (right) fluctuation amplitude spectra at three different points along the wall.

The results in Figure 4.25 reveal that the fluctuation trend at the different frequencies is qualitatively similar between the two cases, for both the pressure and the heat flux, with the fluctuation amplitude increasing at increasing frequencies. However, in the case of the rational Bézier curve leading edge the receptivity is in general lower than the case with a circular leading edge. This difference in fluctuation amplitude between the two leading edge geometries seems, in turn, to increase with the frequency and with the x distance along the wall, for both pressure and heat flux spectra, up to a maximum reached around 350 kHz - 400 kHz . In particular, the maximum difference is obtained at 400 kHz and x=347.5 (where the pressure and heat flux fluctuation amplitudes for the circle-wedge case are respectively about 20% and 14% higher than for the rational Bézier curve leading edge case).

The difference in the receptivity levels between the two different leading edge geometries may be due to the slightly stronger shock formed in front of the circular leading edge (with slightly higher radius of curvature at the stagnation point), which produces a higher amplification of the acoustic waves behind the shock in the nose region, and hence higher wall pressure fluctuation levels transmitted downstream. However, part of this difference may be also a consequence of the curvature continuity in the case of the rational Bézier leading edge, which results in a smoother variation of the wall pressure at the wedge junction (as shown in Figure 4.23), and might produce differences in the mechanism through which the external waves are internalised into the boundary layer.

In conclusion, from the results shown above, the fourth order rational Bézier leading edge geometry that provides G^2 continuity, seems to have a slightly stabilizing effect on the wall receptivity to fast acoustic waves, in particular at the higher frequencies. However, in order to quantify the effect of the different shock curvature and strength due to the small difference in stagnation point radius, and the effect of the curvature continuity, further investigation is needed. Moreover, the study needs to be complemented with the receptivity to slow acoustic waves, which lead to the generation of the dominant unstable mode inside the boundary layer in hypersonic flows, as evident by results available in the literature (e.g. Kara *et al.*, 2007; Malik and Balakumar, 2007).

Overall, the investigation presented was conducted to provide some initial evidence to the claim that more sophisticated shapes and increased levels of geometric continuity at the blunt leading edge-wedge interface can affect the turbulent transition mechanism.

4.4 Integration on 3D Waverider Geometries

There are a number of different approaches when it comes to integrating the 2D geometry formulations described earlier on 3D waverider forebody geometries. In this last section the method which was implemented in the geometry engine will be presented together with some general remarks on the subject, as evaluating and quantifying the merits of each different approach requires complex and computationally expensive case studies.

It is common to truncate the sharp leading edge geometry perpendicular to the leading edge in order to accommodate the bluntness. There are, however, benefits in truncating the geometry and positioning each RBLE section along the osculating planes in the case of an osculating cones/osculating flowfield generated waverider forebody or along the planes on which the streamlines were traced in axisymmetric or wedge shock-based inverse design methods. In this way, maintaining geometric continuity at the interface of the blunt leading edge shape with the upper and lower surface of the waverider becomes more straightforward, especially when 2^{nd} order geometric continuity (which can be achieved with the parameterization described in this work) is desirable. Moreover, even if the original geometry was not G^2 continuous in the spanwise direction due, for example, to the shape of the leading edge, G^2 continuity will be maintained along the general direction of the streamlines (were general direction hints at the fact that bluntness in swept parts will generate cross-flow to adjacent osculating planes). Both approaches are illustrated in Figure 4.26. To obtain the desired thickness when truncating the geometry with the second method (II), the truncation length has to be adjusted according to the local sweep angle and inclination of the osculating planes when truncating along them.
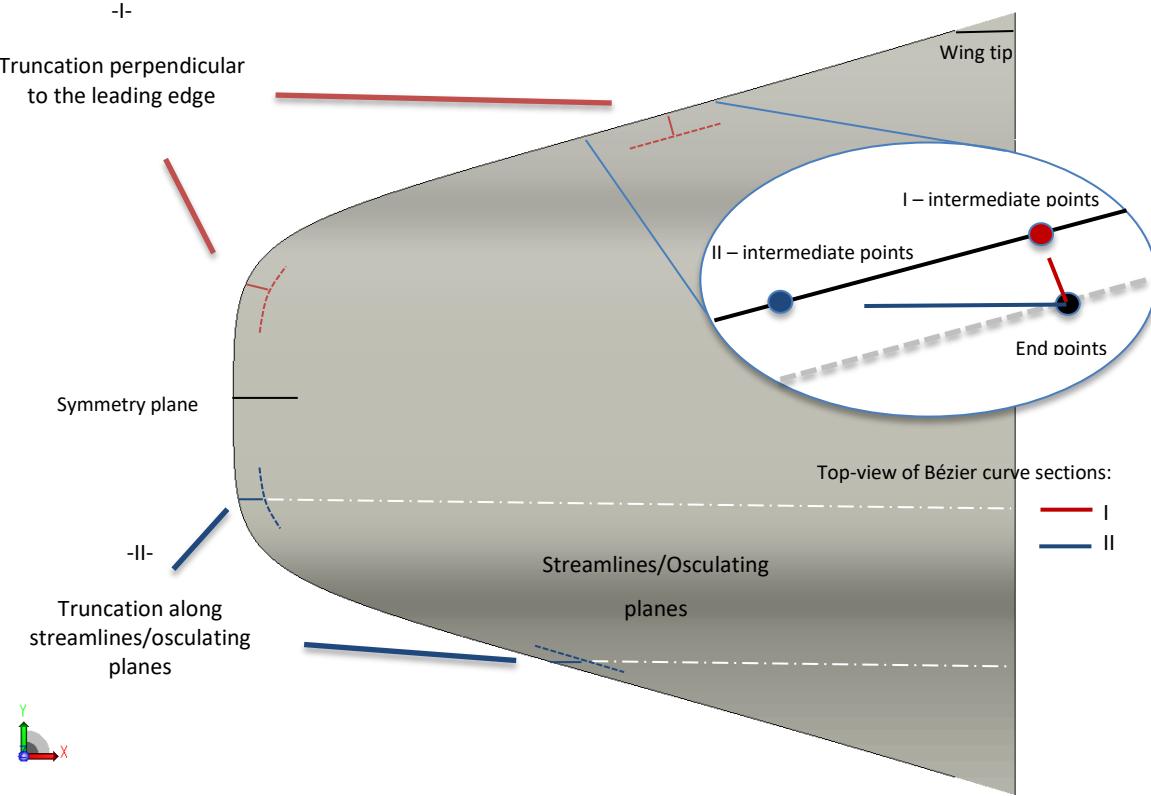
Figure 4.26. Top view of a sharp leading edge waverider forebody with illustrations of truncation options to accommodate the bluntness.

At this stage, apart from the thickness and shape of each section, it is also meaningful to control how those characteristics are distributed along the leading edge, from the symmetry plane at the front to the wing tip. While the 2D evaluations will be valid for the symmetry plane section, swept sections are subject to reduced aerodynamic heating. More specifically, the stagnation point heating of circular swept leading edges can be related to the non-swept one using the empirical equation (Hayes, 1992), which is valid for laminar flow:

$$\frac{q_{swept}}{q_{non-swept}} = (\cos \lambda)^{1.1}.$$

(4.4)

Therefore, in order to optimally utilize the thermal protection system to its limits, the thickness can be reduced as the local sweep increases. The relation for the heating rate of two circular leading edges with different radii is (Bertin, 1994):

$$\frac{q_{R1}}{q_{R2}} = \sqrt{\frac{R2}{R1}}.$$

(4.5)

Using these empirical relations and substituting the radius with the thickness of the RBLE shapes, the thickness across the leading edge can be distributed in order to obtain roughly the same peak heating rate as the symmetry section, equation (4.6). This will result in lower drag while

utilizing the thermal protection system to its limits across the entire leading edge. This approach was also followed by Rodi (2015).

$$d_{swept} = d_{symmetry}(\cos \lambda)^{2.2} \qquad (4.6)$$

When this method is directly applied to geometries that have areas along the leading edge with sudden variations in sweep angle, the planform shape of the original geometry can be significantly altered. Areas where the leading edge is almost non-swept and the sweep starts increasing rapidly are most sensitive. Figure 4.27 includes an example of this, with the sweep starting to increase significantly around *y*=0.2. The resulting geometry ends up with two 'bumps' that are faced head on with the oncoming flow while their thickness is much lower than the nominal zero-sweep thickness.
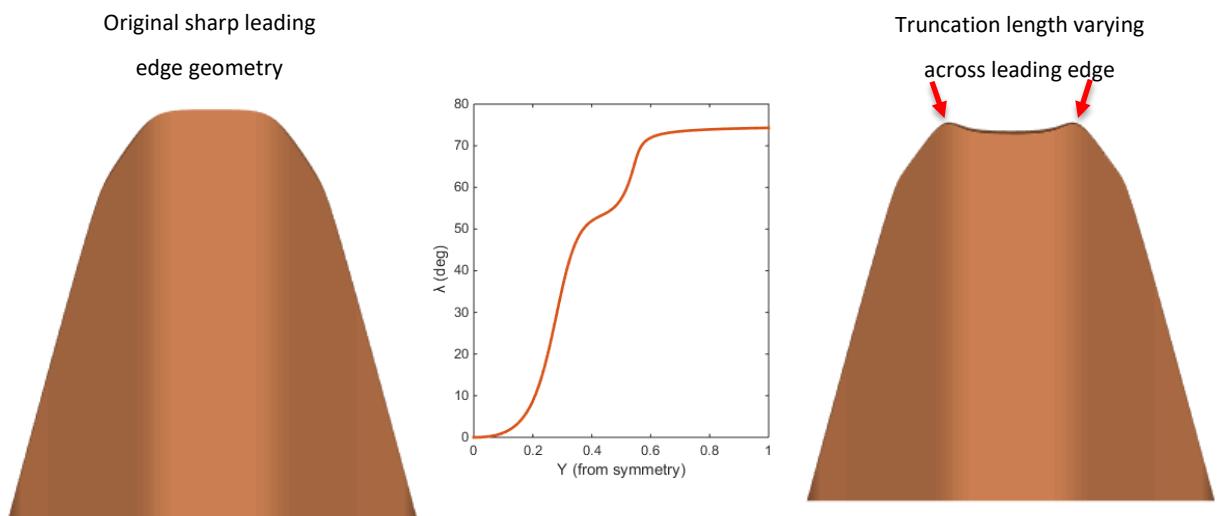
Figure 4.27. Planform deformation when varying truncation length across leading edge according to original geometry's sweep.

To counter this issue at the front of the geometry a limiter can be applied that will not allow the geometry to be truncated less than what is needed in order prevent it from protruding in front of the sections before it, as we move from the symmetry plane to the wing tip. The result of applying this limiter can be seen in Figure 4.28, together with the modified truncation distribution.
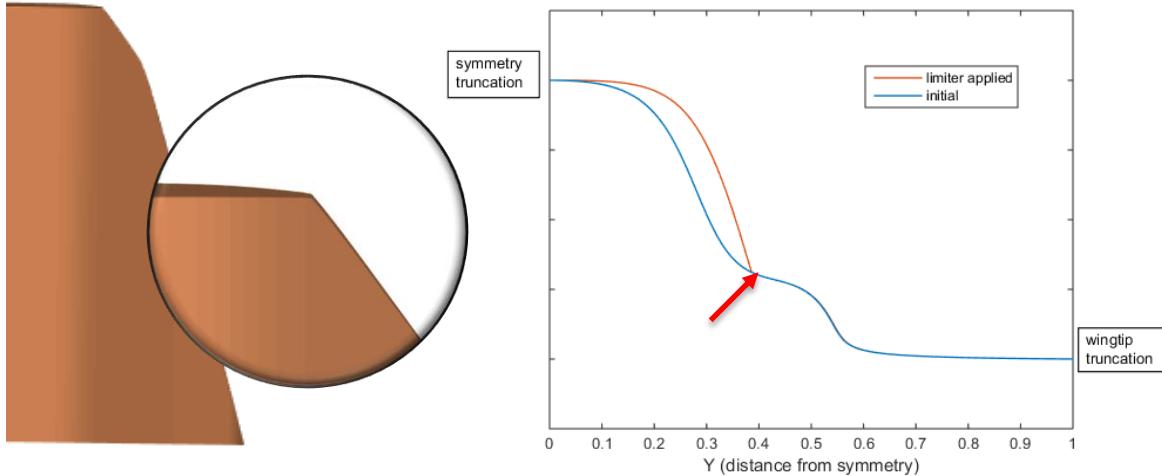
Figure 4.28. Truncation length distribution and resulting geometry with limiter applied.

The limiter does solve the previous issue, however it creates an edge at the point where it switches off. To remedy this, a smoothing of the truncation length distribution in that area can be performed, using, in this case for example, a locally weighted linear regression method (MATLAB User Guide) provided by MATLAB. The results of this can be seen in Figure 4.29.
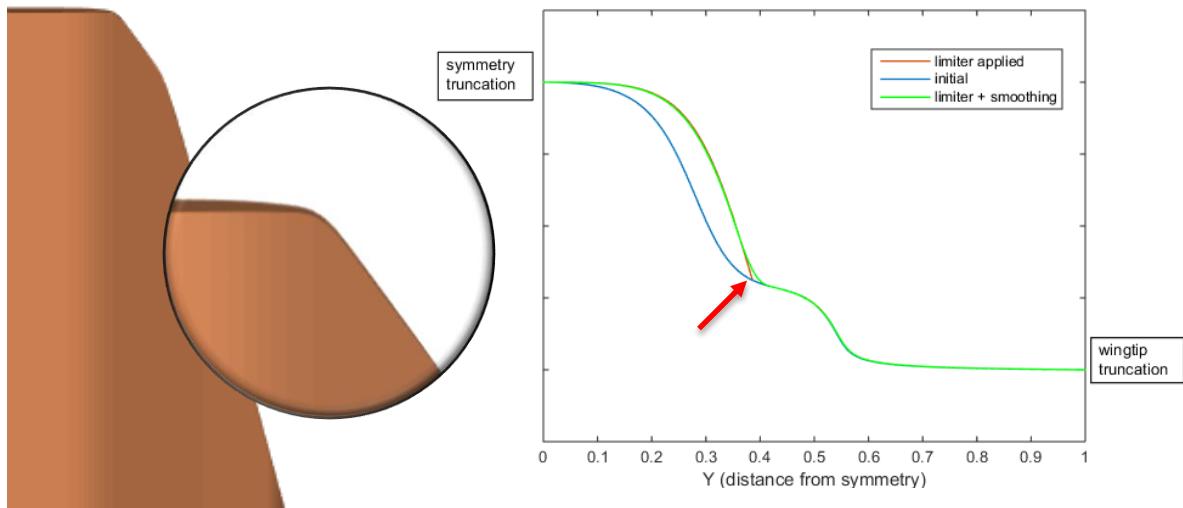
Figure 4.29. Truncation length distribution and resulting geometry with limiter and smoothing applied.

An alternative to the latter is running an iterative method that will keep correcting the thickness distribution as the planform shape of the geometry and effective sweep changes, until a converged shape is reached. However, a more elegant solution is instead offered by not just truncating the geometry, but also raising the upper surface in order to accommodate the bluntness of the leading edge. The two approaches can be blended as seen in Figure 4.30. Although this introduces an additional design parameter (the displacement of the upper surface) that also has to be distributed along the leading edge, it also opens up a number of additional options. One of those

options is the ability to generate the required thickness to accommodate the blunt leading edge without affecting the planform shape of the waverider. This can be achieved by balancing the truncation and displacement in order to keep the tip of the geometry at each section in the same place from a planform perspective. This way, the planform shape and sweep angle remain unchanged, while the shape of the upper surface changes. Additionally, the small wedge angles that usually characterize waverider designs result in deformations of the upper surface that are much smaller than the deformations of the planform would be if only truncation was used. This approach was also incorporated in our geometry engine as a more elegant solution to the ones discussed earlier and can be applied when the leading edge curve is smooth. It should not be applied to leading edges that are not smooth since a discontinuity in the sweep angle across the leading edge will result in a discontinuity of the upper surface due to its displacement and truncation being related to its local sweep angle.
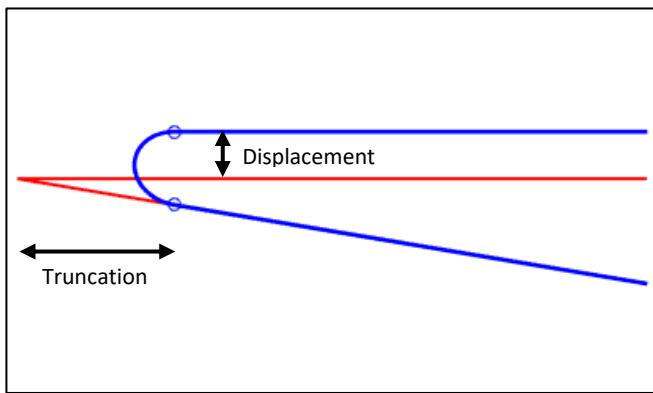
Figure 4.30. Blended displacement of upper surface and truncation of original geometry to accommodate the bluntness.

What also needs to be considered is that the equations described earlier for distributing the thickness of a blunt leading edge according to the sweep do not take into account any effects that the variations of sweep along the leading edge can have on aerodynamic heating, as they generally apply for straight, constant sweep leading edge segments. Additionally, the flow along the attachment line of swept leading edges can transition after a certain length which will result in higher heating rates. It is, therefore, expected that the optimal integrated 3D leading edge geometries can only be reached by utilizing controlled and directly optimized distributions. The previous analysis does however provide a reasonable basis for a self-designing geometry, which avoids the use of one or two more design variables that would have to control that custom distribution.

Finally, it is worth noting that apart from the blunt leading edge's thickness, it can also be beneficial to vary the shape parameters that control the RBLE geometries across the leading edge. Aerodynamic heating considerations apart, this can potentially further limit any losses in lift due to

leakage of the high pressure flow from the underside of the waverider around the leading edges by adjusting the standoff distance of the shock. Additionally, we may find that for swept leading edges a different set of optimal shapes is obtained when it comes to drag and peak heating rates or temperatures. To quantify any potential advantages though, a series of full 3D CFD simulations would be required. The parametric shapes presented in this chapter should enable tackling such complex investigations in an efficient and robust manner. As mentioned earlier, they are able to match or outperform other proposed leading edge shape parameterizations in the literature, while generally utilizing fewer design variables. Ultimately, the additional effort of designing and manufacturing a more sophisticated blunt leading edge shape over a circular one can outweigh the potential benefits. This strongly depends on the relative size of the blunt leading edge, as also seen in Rodi's (2015) work.

We have so far focused on the parameterization of the waverider forebodies and the shape of their blunt leading edge, which is the most tightly integrated additional geometry feature. For the rest of this thesis we turn our focus on how these baseline waverider shapes and the parameterization methods discussed so far fit in an aerodynamic design process and how they can be incorporated in a geometry engine that can be utilized to develop parametric geometries for more realistic design studies.

Chapter 5: A Geometry Tool for the Waverider-based Hypersonic Aircraft Design Process

The approach followed from the early stages of the investigations was to start developing a parametric geometry engine for the design of waverider-based hypersonic aircraft. The material presented in the previous chapters has focused on handling and efficiently parameterizing the inversely designed forebody of the aircraft. In this chapter, a geometry tool that enables a revised design methodology to tackle a number of challenges presented by conceptual and preliminary hypersonic aircraft design is presented. This revised methodology is based on traditional aircraft design processes and aims to tackle two of the main issues early hypersonic aircraft design faces; the increased complexity due to the coupled design approach of the main components of the aircraft from the conceptual design stages, and the very limited availability of data and design experience in the field.

In the sections that follow, the main differences and additional challenges of hypersonic aircraft design when compared to conventional/commercial aircraft design, are initially described. A revised aerodynamic design process, aiming to tackle a number of those challenges, is then presented and a set of requirements for the parametric geometry tool to enable that are identified. Afterwards, having identified requirements of the outer mould line geometry generation tool from an overall design process perspective, the structure and features of the geometry engine that has been developed to fit that role are presented. A number of specific design, parametric efficiency and user interaction features, as well as features satisfying concerns that have been raised through our collaboration with the industry are also presented. Finally, we go through the development of an example parametric configuration of a complete conceptual aircraft geometry that highlights the robustness and versatility offered by the geometry engine.

Efficiency of the parametric geometry model has been a key driving factor during the development of the parametric geometry engine. The majority of parametric efficiency considerations have focused on the inversely designed waverider part of the geometry and its blunt leading edge. Additionally, apart from efficient parameterization, one of our main objectives is to develop a robust geometry engine that is not limited to a single aircraft configuration and engine integration approach, but one that can support the development of various configurations. Other than the waverider design, the geometry engine consists of a number of tools that further manipulate the forebody or design additional components. We first developed a tool for designing streamtraced inlet geometries that can be coupled with the forebody. Most of the additional geometry components required to complete the aircraft, such as control surfaces and interfacing geometries, are usually designed with more conventional parametric geometry tools available in all modern CAD software packages. Methods to rapidly evaluate the performance of those additional geometry components in order to enable parametric flexibility studies are generally less accurate or computationally expensive (e.g. local inclination methods and finite volume CFD methods). We also developed a tool to design such additional components, mostly as an example on how the geometries can be completed, while functionality to export the forebody geometries to commercial CAD software for further development has also been implemented. The main focus of the geometry engine, however, is the forebody inverse design capability, which is not readily available through direct geometry generation methods utilized by the majority of CAD software packages.

5.1   Waverider-based Hypersonic Aircraft Design Process

5.1.1   Differences to Conventional Aircraft Design

Hypersonic aircraft design presents a number of additional challenges when compared to conventional subsonic and even supersonic aircraft, both from an overall design process and an aerodynamic design process perspective that this work is focused on. The overall aircraft design process can generally be divided in three phases, the conceptual design phase, the preliminary design phase and the final or detail design phase, as seen in Figure 5.1. In the conceptual design phase, once a set of requirements and mission objectives are set, the aircraft configuration is determined and various sizing and performance characteristics are estimated. In the preliminary design phase the aerodynamic shape and airframe structures are further developed with the assistance of CFD and wind tunnel testing, until a finalized design is reached. If at this point the requirements set at the beginning of the design process are met, the design can progress to the final design stages. In general, sufficient market interest is also required at this stage. In the final or detail design phase all the components and subsystems of the aircraft need to be fully defined and

the entire manufacturing process needs to be developed. It is the most expensive and time consuming phase of the design process.
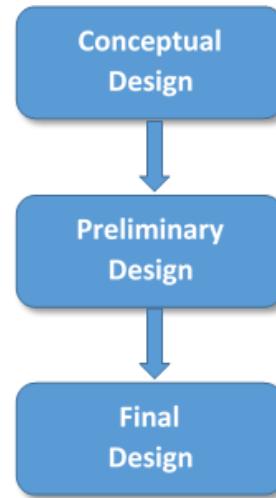
Figure 5.1.   The three aircraft design phases.

The conceptual design phase can also be divided in three more stages: the general concept definition, the configuration definition and the configuration refinement stage. The configuration definition stage is primarily driven by specific design requirements that might have been set with respect to the role or functionality of the aircraft and by referring to empirical data and past examples. The designer is, at this stage, able to identify the types of configurations that are proven to be able to satisfy the majority of the requirements that have been set, through large databases of already designed and proven aircraft. Data and detailed studies for standalone aircraft components are also readily available. These can include detailed data for existing propulsion systems that could satisfy the given set of requirements, airfoil shapes, wing design and fuselage integration methods depending on the loading and desired flight speed, etc. Additionally, data from various trade off investigations and simple empirical relations can be combined to make very well educated decisions and enable designs unlike any before seen, albeit similar. Once a configuration and preliminary sizing has been determined, especially in the case of transport aircraft and commercial airliners, a relatively detailed geometry of the entire aircraft can be obtained with little effort, using available conceptual design tools.

When it comes to hypersonic aircraft design, however, the availability of such data is severely limited. This is also reflected in the diversity of configurations that are being proposed in conceptual design studies. Examples of such configuration decisions that are made at this stage can include choosing the type of engine inlet and how it is integrated with the forebody or the use of a waverider forebody and how much of the entire aircraft the waverider component covers. In many cases different and diverse options will need to be investigated and a much larger design space will

need to be explored. The configuration might also need to be significantly revised during the configuration refinement stage, while the level of confidence in the configuration selected will still hardly reach the confidence provided by imitating a proven past example in conventional aircraft design. Limited past collective experience in this design field is therefore one of the hurdles that greatly affect the early conceptual design stages of hypersonic aircraft.

The additional challenges of hypersonic aircraft design that were described in Section 2.1, all contribute to increasing the complexity of the hypersonic aircraft design process when compared to conventional aircraft design. The integrated and coupled design nature of the primary components of a hypersonic aircraft is one of those factors. In this aspect, hypersonic aircraft design is a little more similar to supersonic and military aircraft design, however the level of integration and the point in the design process where it becomes important still stand out. Usually, in conventional aircraft design, each of the primary components of the aircraft can be designed independently to a large degree. The interaction between the propulsion system, wings, fuselage and control surfaces is not very important in the early design stages, while there is also empirical data available to predict their interactions and direct their integration. Detailed investigations on their interaction and integration that require CFD and wind tunnel testing are usually left for the preliminary phase of the design process. Designing interfacing geometries such as the wing-body fairing and jet engine pylons can also be an isolated design processes performed at subsequent design stages. For hypersonic aircraft, on the other hand, the major components of the aircraft are usually designed integrated as a whole from the conceptual design phase, as was described in Section 2.1. The design can still be broken down to a certain extent, for example as long as the forebody is required to generate an area of uniform pre-compressed flow on-design, an engine inlet to fit that could then be designed. Even in this simplified case though, once angle of attack and sideslip need to be accounted for, the resulting more complex interaction is something that needs to be accounted for while designing each component. This coupled design approach and the higher level of interaction between the main geometry components of a hypersonic aircraft generally increase the dimensionality of the design problem. Increased number of design parameters that need to be considered simultaneously along with the need to explore a larger design space greatly increase the effort and computational cost required by design studies.

The pronounced multidisciplinary nature of hypersonic aircraft design also becomes important in much earlier stages than in conventional aircraft design. This is primarily due to aerodynamic heating, which is one of the primary design-driving factors from the early design stages. Such multidisciplinary considerations can significantly increase the complexity and cost of the early design stages as they further contribute to the multi-dimensionality of the design problem and also increase the computational cost of design evaluation methods.

The main differences and additional challenges of early hypersonic aircraft design compared to conventional aircraft design can be summarized in the following:

- Limited design experience in the field and very few past proven examples and empirical data to rely on. This generally requires the exploration of a more vast design space.

- Integrated design approach that increases the dimensionality of design studies from the early design stages.

- More pronounced multidisciplinary nature, and especially aerodynamic heating effects, that strongly influence the design and add to the multi-dimensionality and computational resources required for design evaluation.

- The effectiveness of hypersonic aircraft is dictated by tighter margins between lift, drag and thrust, this in turn requires more detailed simulations that can provide higher levels of confidence earlier in the design process.

- Well defined geometries are required early in the design process in order to enable coupled analysis of the different aircraft components.

With computational design search and optimization having become standard practice in aerodynamic design, all of the above essentially lead to increased computational cost. The part of it that stems from the higher fidelity and multidisciplinary simulations required cannot be avoided. After all, a drive to increase the fidelity and contribution of CFD in earlier design stages while also taking multidisciplinary aspects into consideration has been present for aircraft design in general. This drive is even stronger in the hypersonic design field. What can be improved, however, is the design process and the efficiency of computational design studies.

5.1.2   Overview of a Generic Aerodynamic Design Process

Apart from the points raised in the previous section, the conceptual design process of hypersonic aircraft does not greatly differ from that of conventional aircraft. A generic conceptual design process for the definition of the outer mould lines of an aircraft geometry involving computational design search and optimization is described in the flowchart of Figure 5.2.

The objectives and design requirements for a hypersonic mission, similar to a conventional aircraft, can include range, payload mass and volume, flight speed, etc. It is likely that these objectives are a little more flexible than in conventional aircraft design, and are adjusted during the conceptual design phase due to the amount of innovation required to realize them.

The available empirical data used at this stage consists of past proven aircraft configurations that have fulfilled similar missions, previous concept studies, design trade-off studies of aircraft or

individual components, empirical relations and information that can provide adequate information for selecting the initial configuration of the aircraft. This was identified as one of the more challenging areas when it comes to hypersonic aircraft design.
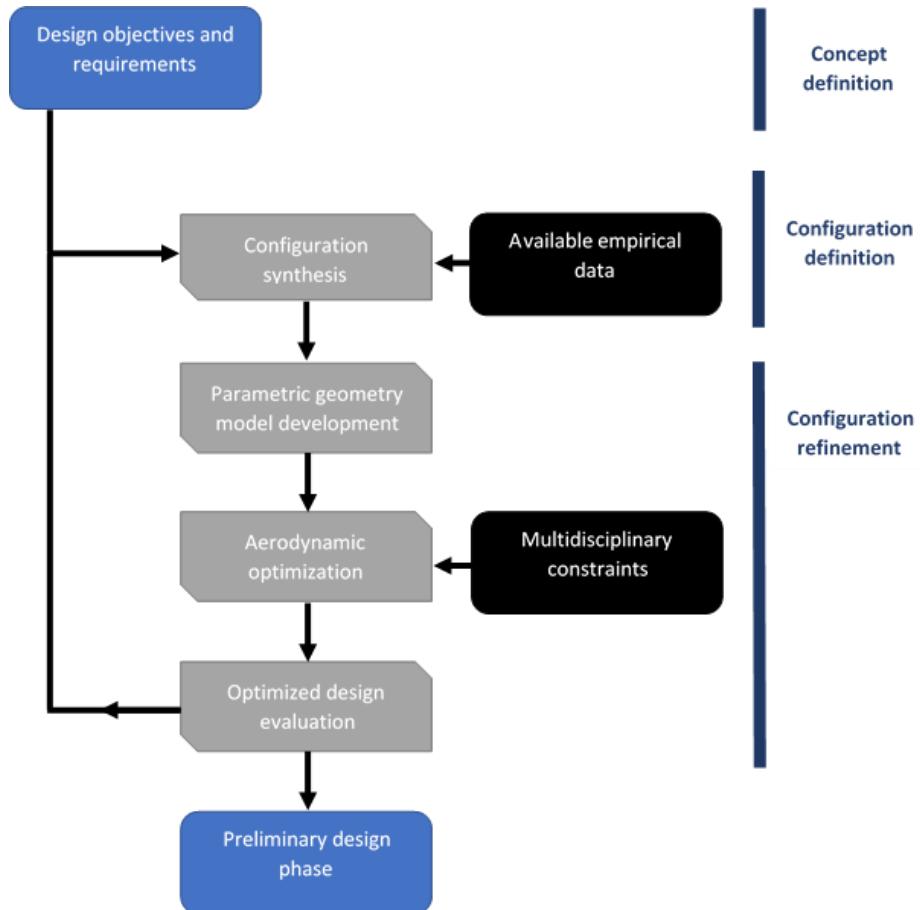
Figure 5.2.   A generic concept aircraft aerodynamic design process.

In the configuration synthesis phase the layout and integration approach of the main components of the aircraft is decided. Different configurations are generally designs that cannot be covered by a single parametric geometry model, or at least the parametric geometry would be unable to transition from one configuration to another smoothly. This stage involves decisions on very basic features as well as more detailed ones. These can be: using a waverider lifting body or not, using a ramp based inlet compression system or an inward turning type inlet, the type of control surfaces, or even something with greater impact such as adopting a single or two stage approach. Once a configuration and basic layout has been decided a parametric geometry model can generally be developed.

The parametric geometry model development phase is required to make future updates to the design easier to implement as well as to enable automatic computational optimization techniques. This phase has been the main focus of this thesis, while it is a process that is usually not thoroughly explained and rationalized in the literature. Developing an effective and at the same

time efficient parametric geometry can greatly affect the outcome of the subsequent aerodynamic optimization. Details on parameterization and geometry handling practices also being areas were hypersonic aircraft design is lacking, as was explained in section 2.4.

The aerodynamic optimization phase, which can also be seen as a multidisciplinary problem since it is subject to multidisciplinary constraints, is the process through which the parametric geometry model is fine tuned. This process can be repeated in the conceptual design phase with varying, usually increasing, levels of fidelity and complexity. The bulk of the computational cost at this stage lies in the design evaluations through flowfield computations which are repeated for several designs. Other than the cost of each design evaluation, the other major factor that dictates the computational resources required to converge to an optimal solution, is the dimensionality of the design problem. The choice of optimization algorithm can also have a certain effect on the cost of a design study, albeit not as profound as long as sensible decisions are made in its selection and setup. The exception to the so called 'curse of dimensionality' are adjoint based optimization methods. These methods however are restricted to locating local optimum solutions and are therefore better suited to later design stages.

The outcome of the optimization or design study is finally evaluated and the process repeats with either fundamental updates to the configuration, updates in the parametric geometry model, and/or increases in the fidelity and complexity of the design problem. Once the requirements set in the concept definition phase are met and the geometry is also compatible with other multidisciplinary constraints, the design process can move on to the preliminary design phase where the design is further refined. As far as hypersonic aircraft design goes, one could say that certain design stages that would be better suited to the preliminary design phase for conventional aircraft, are performed in the conceptual design phase here. This is based on the complexity of the required evaluation methods, with 3D CFD simulations being commonly utilized in the conceptual design phase, and the coupled design approach of the components. It is also possible that the boundaries of conceptual and preliminary design phases are more difficult to define in this case due to the relative infancy of the field.

5.1.3 The Revised Hypersonic Aerodynamic Design Process

To improve the efficiency of the design process we are focusing on the development of the parametric geometry model and in rapidly generating additional data in the form of trade off studies that cover a wide design space. These design loops are highlighted in the flowchart of the revised design process shown in Figure 5.3. The goal of this approach is to improve the effectiveness and efficiency of the parametric geometries before moving on to full scale design studies and

aerodynamic optimization processes. The highlighted design loops are essentially a cost effective way to tune the parametric geometry model to a degree. The low fidelity evaluation techniques and segregated approach of the parametric flexibility studies, as will be explained later, do not guarantee an optimally parameterized full aircraft geometry, but provide useful information and insight on the parameterization of non-standard geometry components whose parameterization methods are not as mature.
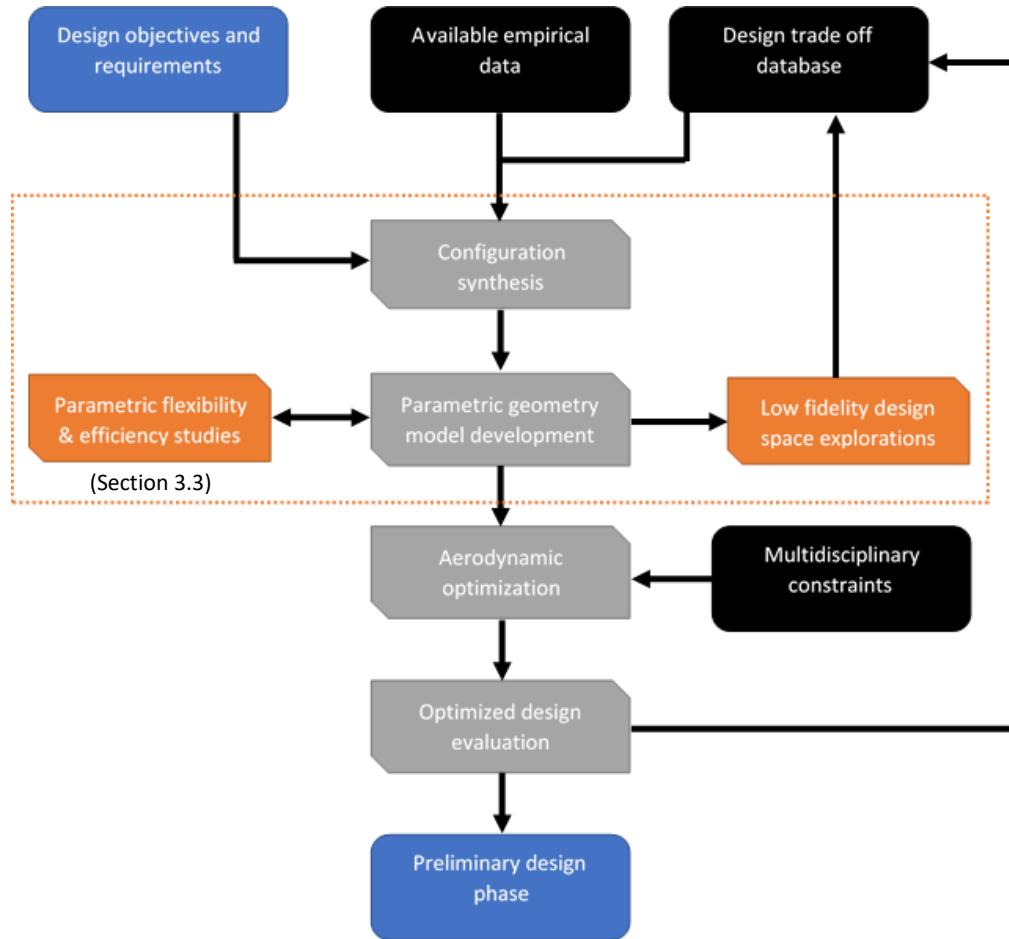
Figure 5.3.   Revised hypersonic aircraft aerodynamic design process.

The parametric flexibility and efficiency studies, as described in Section 3.3, are utilized to compare different parameterization schemes. They can be performed to study the parameterization of geometry components that can be rapidly evaluated on their own; waverider forebodies are a very good example of this. Parametric flexibility studies of other integrated components can also be performed by freezing the rest of the design. As was explained in the aforementioned section of Chapter 3, efficiently parameterizing each component through such processes will not necessarily translate to equivalent efficiency of the complete parametric geometry model being developed. It does, however, enable the designer to make better educated decisions. Apart from the waverider forebodies and blunt leading edges that were examined in the previous chapters, similar parametric flexibility studies can be performed on a number of aircraft

components. The main requirement is the applicability of rapid evaluation methods. These can range from simple CFD simulations, as was seen for the blunt leading edge shapes, to local inclination methods coupled with semi-empirical skin friction calculations for more arbitrary geometries that are used to complete the airframe. It is, of course, up to the experience and intuition of the designer to decide for which components such studies make sense. The intention of this process is to cover the lack of detailed information on the parameterization of hypersonic aircraft components, something that is generally available for conventional aircraft design.

To enable parametric flexibility and efficiency studies simple parametric geometry models need to be developed, focused on specific parts of the geometry and with varying degrees of flexibility and formulations. The range of these is also a decision the designer needs to make. The data generated should not be focused on single design points, as at this stage of the design process clear goals cannot be set for each geometry component before they have been integrated. Instead, a multi-objective approach is preferable, usually best presented as sets of Pareto-optimum fronts. Comparing Pareto fronts of two to three objectives is a fairly straightforward way to compare the capabilities of different geometry parameterizations. Parameterization schemes that totally dominate others or ones that are dominant only in parts of the design space can be identified. Additionally, by following a fixed computational resources approach, as was explained in Section 3.3, a good indication of their efficiency in reaching 'good' regions of the design space can also be highlighted. Moreover, the range of the design parameters can be better tuned through this process to exclude invalid or very poor performing regions of the design space.
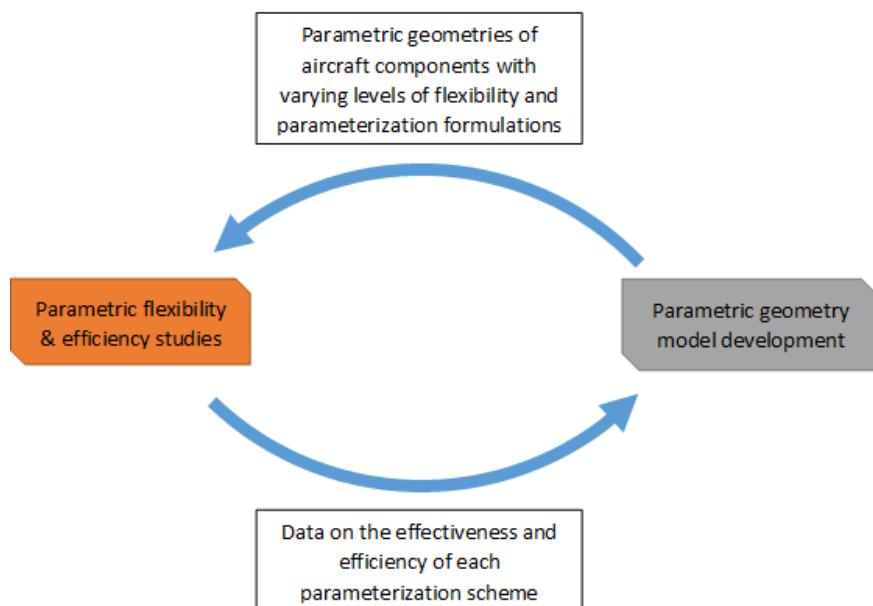
Figure 5.4. Synergy between parametric geometry development and geometry flexibility studies.

This cost effective design loop primarily achieves a potential decrease in the dimensionality of the design space without restricting the capabilities of the parametric geometries. This primarily

addresses the higher computational cost issue of design studies by potentially reducing the amount of design evaluations required to reach satisfying and/or converged solutions. Finally, it also covers for part of the limited design experience issue by identifying good practices in parameterizing hypersonic aircraft specific components. A schematic of this design loop is shown in Figure 5.4.

The second design loop involved in the development of the configuration and parametric geometry model is used to further make up for the limited design experience and availability of empirical data in the hypersonic design field. The low fidelity design space explorations, shown in the design loop of Figure 5.5, are essentially multi-objective design optimization studies that utilize fairly simplified parametric geometries and low fidelity evaluation methods. Complete airframe configurations as well as the effect of certain types of components (e.g. types of inlets) can be investigated. Unlike the aerodynamic optimization process that follows, these design studies are focused on covering a wider and more diverse area of the design space, one that cannot be covered by a single parametric geometry model. Additionally, they should not be focused on a single design point and specific goals but should also be approached as multi-objective design explorations. This way a design trade off database can be built, utilizing Pareto front information that provide a good indication of the capabilities of the different configurations over a wider area of the design space. Through this process the available empirical data can be further enhanced in order to improve the configuration synthesis and parametric geometry development. One could see this process as building up 'design experience' without committing vast amounts of resources. It should be noted that data obtained through parametric flexibility studies may also be suitable for this purpose.
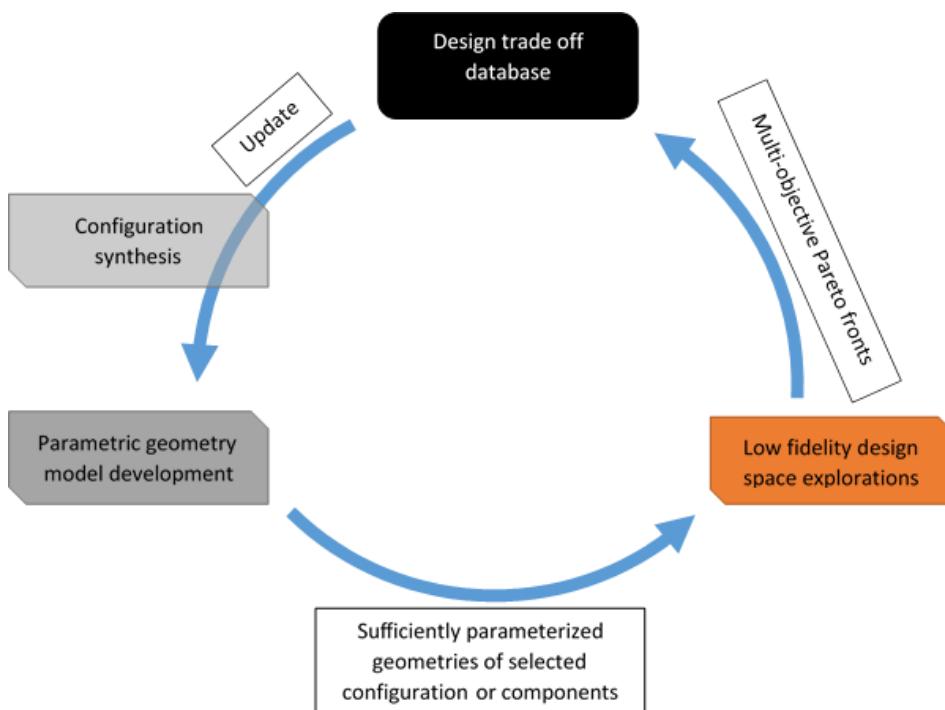
Figure 5.5.   Role of low fidelity design space explorations.

To enable such design space explorations the airframe parametric geometry models generally need to be more complete than for the parametric flexibility studies. Low fidelity rapid evaluation methods are also required, which again renders waverider-based hypersonic aircraft configurations very suitable. For example, available flowfield data from inverse design methods combined with local inclination methods and semi-empirical skin friction calculations can be utilized for performance evaluation of these more complete aircraft configurations, a technique that has been utilized in various conceptual design studies in the literature. The objectives of these studies can range from lift, lift to drag ratio, volumetric efficiency and mass estimates, to stability characteristics and thermal management metrics as long as the evaluation method remains relatively 'cheap'. An example for such a study would be an investigation with an inward turning versus a ramp compression inlet integrated on a forebody geometry. The forebody design should not be frozen in order to take advantage of the changes introduced by the different inlet configurations.

On the one hand, similar design loops would generally be performed anyway in the generic aircraft design process, as it was shown in Figure 5.2. This would be accomplished through the entire aerodynamic optimization loop that is usually performed with low fidelity tools during the first design iterations. On the other hand it is important to both separate this phase in which the configuration is still very fluid and also to obtain multi-objective design trade off data which can also be useful if the mission objectives need to be adjusted later on or even for different missions and design problems. Overall, both of the design loops that were described earlier aim to provide a basis for making rationalized and better justified decisions throughout the parametric geometry model development. The primary focus of this work, however, is the geometry tool that can enable such processes, and we turn our attention to that next.

5.2   The Parametric Geometry Generation Tool

The previous section of this chapter has set the stage for the hypersonic aircraft parametric geometry engine that we set out to develop. The complete geometry engine consists of different tools that are illustrated in Figure 5.6. This work has been primarily focused on the parameterization of the waverider forebody geometries with most of the software development effort also focused on that part of the geometry engine. Tools to parameterise and generate the remaining, directly designed components of the aircraft can be found in the majority of CAD software. We did, however, also develop a tool to generate the aft end and integrate control surface geometries, mostly as an example on how the geometries can be completed. The forebody and engine inlet designs are highly coupled, and the tool for generating streamtraced inlets to fit the generated waverider shapes also needed to be developed. The ultimate goal of this development process is

to provide the design tools that, along with traditional CAD techniques, can be used to completely define the outer mould lines of a concept hypersonic aircraft.
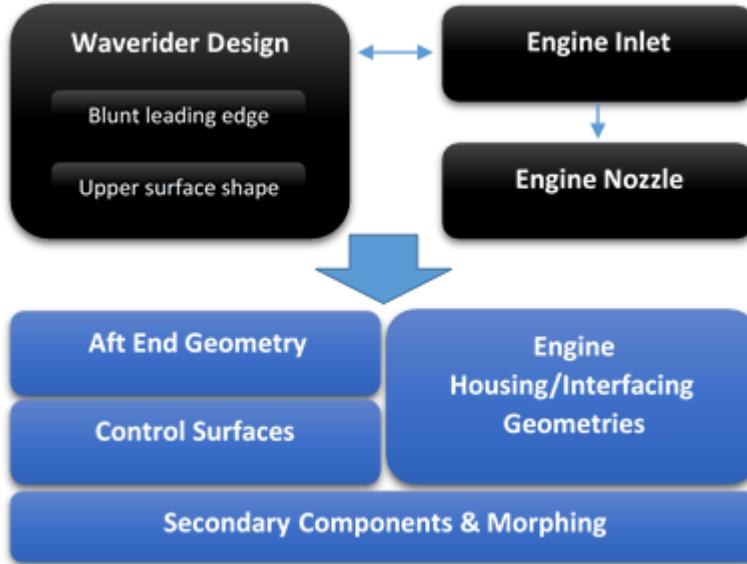
Figure 5.6.   Components of the complete geometry engine, usually inversely designed (black) and directly designed (blue).

5.2.1   Requirements

Enabling the revised design process that was described in section 5.1.3 introduces a number of requirements that are explained in the following section. A number of additional requirements stem through our collaboration with the industry. These requirements of the geometry generation tool can be summarized as:

➢ Ease of use to enable the development of multiple parametric geometry configurations required by the revised design process.

➢ Built-in evaluation tool for inversely designed geometries that can also be used for parametric flexibility studies.

➢ Flexibility in the way the waverider component is parameterized in order to readily interface with additional geometry components and support different configurations and inlet integration approaches.

➢ Functionality to easily introduce different/custom design-driving curve formulations for the waverider forebodies, other than Bézier curves.

➢ Interfacing between the forebody and engine inlet geometry tools in order to support various configurations and inlet integration approaches.

➢ Control over the accuracy and fidelity of the output geometries (for exporting to CAD software, for computational grid generation, and as good geometry handling practice in general).

➢ Robustness control and the ability to identify regions in the design space that lead to invalid geometries.

Performing parametric flexibility studies and low-fidelity design space explorations requires rapid prototyping of multiple parametric geometries with varying amounts of flexibility and different configurations. By ease of use of the geometry generation tool we mean that the designer/user should be able to readily make decisions on the ranges of the design variables and also have a good idea of the flexibility offered by the amount of control points or degrees of freedom used in every case. This is mostly a process based on experience and intuition but since it will need to be repeated several times, the geometry generation tool can be adjusted to potentially enhance this stage of parametric model development.

The aerodynamic performance of inversely designed waverider geometries is something that can be rapidly evaluated utilizing the flowfield information used for their design. For those types of geometries it makes a lot of sense to perform these aerodynamic performance calculations during their generation since the additional computational cost is almost negligible. Waverider geometries also happen to be among the hypersonic-design-specific geometry components we are most unfamiliar with when it comes to their detailed parameterization. Our capability to rapidly evaluate their performance renders them ideal for parametric flexibility studies, which is why a waverider performance evaluation algorithm was built into the waverider geometry generation tool.

Apart from supporting parametric flexibility studies, the geometry generation tool that has been developed incorporates all the parametric efficiency oriented design techniques presented in Chapters 3 and 4. As far as the waverider forebodies go, the different design methods presented in 3.1 and 3.2 enable the geometry to readily interface with additional geometry components and can also be utilized to apply potential design requirements more efficiently. Design of blunt leading edges is also an integral part of the waverider forebody generation process and is built into the same geometry generation tool. The proposed parametric geometries and application methods presented in Chapter 4 have been implemented.

While the two parametric flexibility case studies of Chapter 3 primarily used Bézier curves, utilizing different types of parametric curve formulations and even parametric flexibility studies across different formulations should be straightforward to implement. Additionally, the forebody and streamtraced inlet design tools should seamlessly interface with each other to generate matching geometries. A forebody can be designed to accommodate a specific inlet or a

streamtraced inlet can be designed to fit a specific forebody. In general, similar parametric geometry tools that have been utilized in the literature have been specialized to one selected configuration. The intention for this geometry engine is to be robust enough in order to enable the development of parametric geometry models for multiple configurations with minimal additional software development effort.

Finally, the user should also have control over the accuracy and fidelity of the geometries output by the geometry engine. Apart from being good geometry handling practice in general, this becomes more important when the geometries are to be imported in CAD software to be further built upon. A method to quantify that accuracy is necessary and any control approach needs to be dynamic in order to account for the fact that these are parametric geometries. Such features also contribute to a more polished geometry generation tool that can be utilized throughout the conceptual design phase and potentially during preliminary design as well.

5.2.2   Overview

The main pillar of the geometry engine is the waverider forebody generation tool, which is able to generate blunt leading edge waveriders utilizing the parameterization approaches presented in the previous chapters. This is complemented by a set of tools that we developed for further advancing the design of waverider forebodies to hypersonic aircraft. Those include: tools to generate streamtraced inlets, a single-expansion-ramp nozzle (SERN) and a set of simplified tools for completing the aircraft geometry, following the structure presented in Figure 5.6. To address the first of the requirements set earlier regarding 'ease of use', we also developed an accompanying interactive graphical user interface (GUI) providing more intuitive control. All aforementioned geometry generation tools have been programed in MATLAB, spanning a total of around 10,000 lines of commented code.

The main script of the geometry engine is responsible for generating the waverider forebody and applying a blunt leading edge. Design variables and parameters are read from a *.mat file (MATLAB format) which can be generated and edited either manually, by an automated optimization algorithm or with the developed GUI. A chart illustrating the primary functions and inputs/outputs of the forebody generation algorithm is seen in Figure 5.7. The selected configuration and engine inlet integration method affects the parameterization of the forebody, which is coupled to that of the inlet. This is achieved by tailoring the design-driving curve functions and the way the forebody is parameterized (utilizing the different approaches presented in Chapter 3) to the requirements of the coupling method that is being followed. Once a sharp leading edge geometry has been generated, it can be evaluated with the performance evaluation script. A blunt

leading edge can also be designed, finalizing the baseline forebody geometry. This geometry can then be output in two formats, a *.dat proprietary format compatible with software used by MBDA UK, or a standard *.stl (stereolithography) format compatible with most major CAD and grid generation software. More details on these and a number of additional features of the forebody generation algorithm and the geometry engine are presented in the following sections.
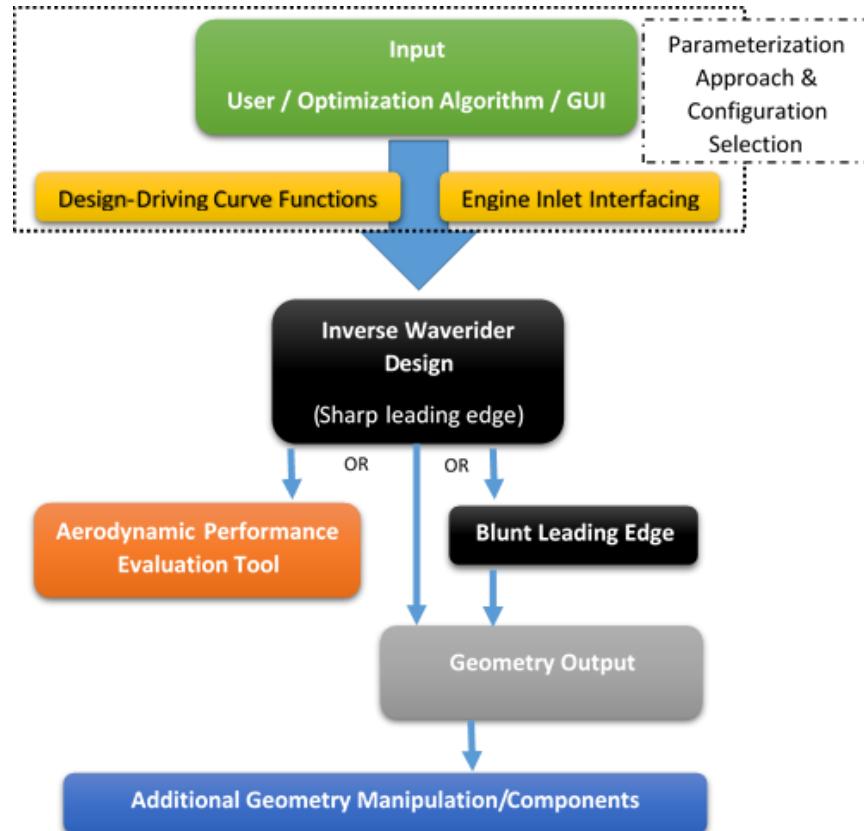
Figure 5.7.   Waverider Geometry Generation Tool Overview.

5.3   Specialized Features

To address the requirements set earlier a number of features that are described in this section were developed. Examples on how to design and parameterize the additional geometry components to further build upon the waverider forebody are given in 5.4.

5.3.1   Interactive Graphical User Interface (GUI)

To address the ease of use requirement, we developed an interactive graphical user interface, able to utilize the majority of functions provided by the geometry engine, in MATLAB. Among its novel features is the ability to manipulate the waverider geometries in real-time with control points that can be dragged with the mouse cursor, rendering around 4 updates per second on a standard laptop. This real-time manipulation of the geometries better relates to human

understanding and we also found that it favourably adjusts the learning curve for engineers new to waverider design. The GUI is a tool that enables more intuitive user interaction, and can potentially enhance and speed up the process of setting up design studies and generating waverider forebodies to interface with specific additional aircraft components. It also enables the development of multiple parameterization schemes that can be manually explored and visualized prior to parametric flexibility studies. This section contains a brief explanation of its main features and functionality. An overview of the GUI is shown in Figure 5.8.
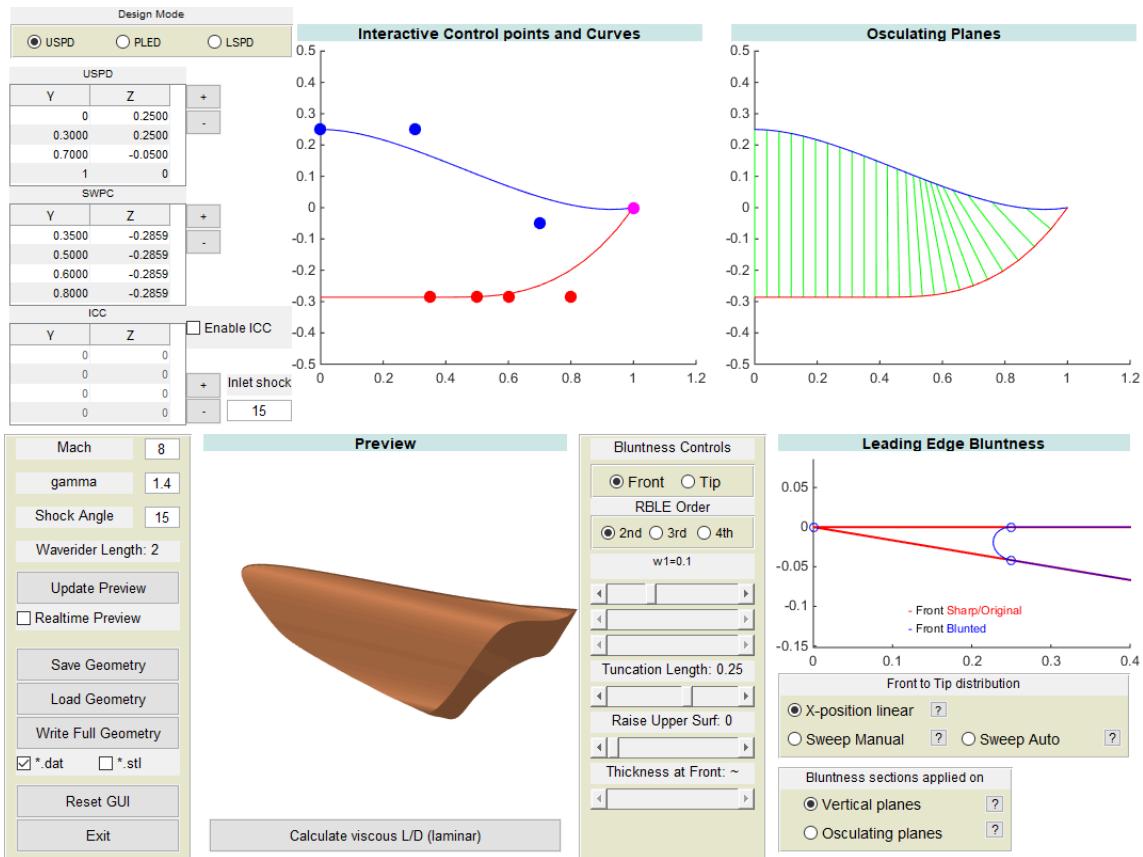
Figure 5.8.   Overview of the interactive GUI.

On the top left panel, Figure 5.9, the user can switch between different design modes and also configure the control point coordinates for the two design-driving curves. The option to design an additional external compression surface on the underside of the waverider and in the part where the flow is uniform is also available. A geometry like that can form the external compression part of a mixed compression inlet and replace the first compression ramp of a traditionally designed planar compression inlet while avoiding the use of sidewalls, its design is described in Appendix E. The figure displaying the control points and design-driving curves is interactive, the control points can be manually dragged with the mouse cursor, with the figure updating in real time, and also be controlled by entering numerical values in the relevant tables. The figure on the top right showing the osculating planes of the design also features real-time updates. This, apart from providing more

insight into the geometry generation method, especially to a user that is beginning to get accustomed to the osculating flowfield nature of the inverse design algorithm, also enables the user to easily identify the limits where the design-driving curves result in invalid geometries by breaching the limitations mentioned in 3.1. This way, picking sensible upper and lower limits for the design parameters is more straightforward. Finally, control points can be added and removed with the '+' and '-' buttons, enabling the design of more complex shapes by increasing the flexibility of the design driving curves. When Bézier curves are being used, a degree elevation algorithm is utilized in order to increase the flexibility without affecting the current shape.
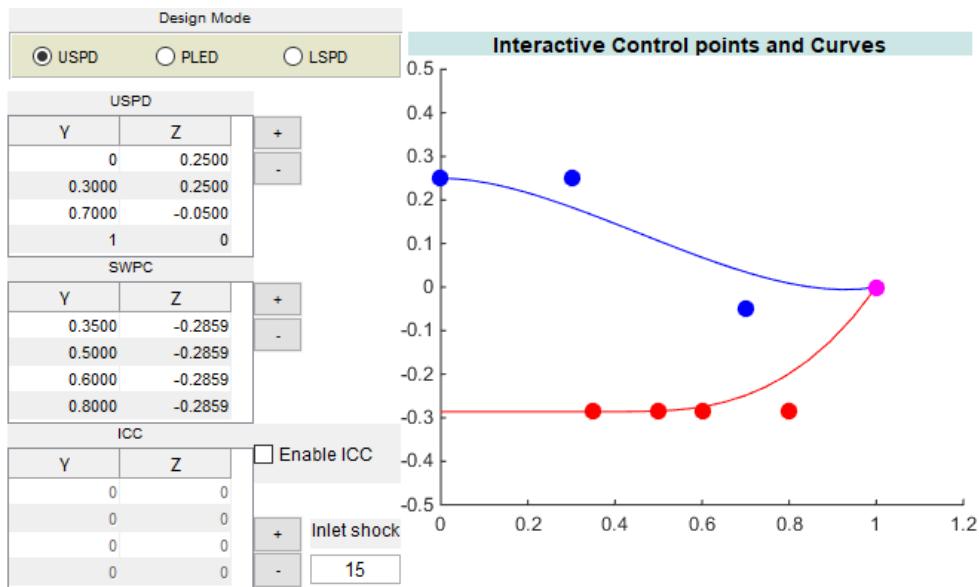
Figure 5.9. Top left panel of the GUI, design mode selection and design-driving curve control.

The main controls are located on the bottom left of the GUI, seen in Figure 5.10, next to the preview window. The text input in the top three boxes define the flowfield features of the cone in supersonic flow that is used in the osculating cones method, they are: the freestream Mach number, the specific heat ratio of the gas and the conical shock angle. The length of the forebody is displayed underneath. The remaining controls are:

- **Update Preview/Real-time Preview**: With the real-time preview checkbox unchecked, any changes to the design driving curves are reflected in the three-dimensional preview geometry after the 'Update Preview' button is clicked. When 'Real-time Preview' is checked, the geometry in the preview window updates automatically and in real-time as the control points are dragged or control point coordinates are input, rendering around 4 updates per second. To enable that, the ODE solution is evaluated once and interpolated with a spline function, which is in turn used to construct the streamtraced lower surface. Additionally, the preview geometry is of relatively low resolution. This functionality allows the user to instantly identify how changes to the design-driving curves are reflected in the

final geometry, and also enables them to more easily identify design variable limits for the control points. Additionally, as was mentioned earlier, we found that it is a much more intuitive way to introduce the design method to designers/engineers without prior knowledge of the inverse design algorithm.

- **Save/Load Geometry**: Waverider geometries generated by the GUI and used by the geometry engine are fully described in ~.mat (MATLAB format) files that can be saved from or loaded into the GUI. Therefore, geometries from optimization runs that were designed through an automated process can also be loaded, inspected and modified as well.

- **Write Full Geometry**: This function saves the currently designed waverider to a ~.mat file and passes it on to the main script of the engine which in turn writes the final, high resolution geometry in the specified format (proprietary .dat and/or standard .stl).

- **Calculate viscous L/D**: This function again saves the currently designed waverider variables to a ~.mat file and passes it on to the aerodynamic performance evaluation tool mentioned in 5.3.2. The geometry needs to be a sharp leading edge waverider.
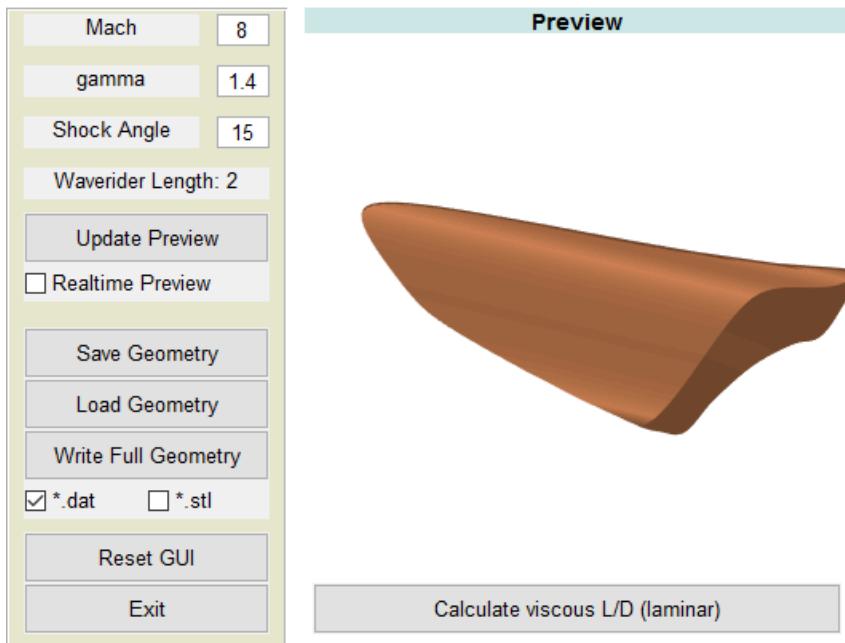
Figure 5.10. Main controls and preview window of the GUI.

Finally, the bottom right panel contains the blunt leading edge shape controls, seen in Figure 5.11. The three RBLE geometry definition methods described in chapter 4 have been implemented. The user can control the shape of the leading edge, and the truncation and/or displacement of the upper surface that is needed in order to accommodate it. These can be defined at the front, which is on the symmetry plane of the geometry and at the wing-tips, with the parameters being distributed in between using three different approaches. The automated method discussed in 4.4 that attempts to design a geometry with equal peak heating rate across the leading edge by relating

the effective sweep angle to the thickness, is the third, 'Sweep Auto', method. When this is selected only a thickness at the front needs to be defined and the truncation and displacement are blended in order to keep the planform shape unaffected. In the 'Sweep Manual' mode the selected values between the front and the tip are linearly distributed according to the local sweep angle. A third approach, 'X-position linear', linearly distributes the values between the front and tip according to their streamwise position. This was implemented for non-smooth leading edges where the sweep angle along the leading edge is not continuous. The bluntness sections can be either applied on the osculating planes of the inverse design algorithm or on vertical planes, the second option being preferable when a combination of large dihedral angles and angles, α, of the osculating planes towards the sides of the waverider can cause significant skewness of the blunt shape.
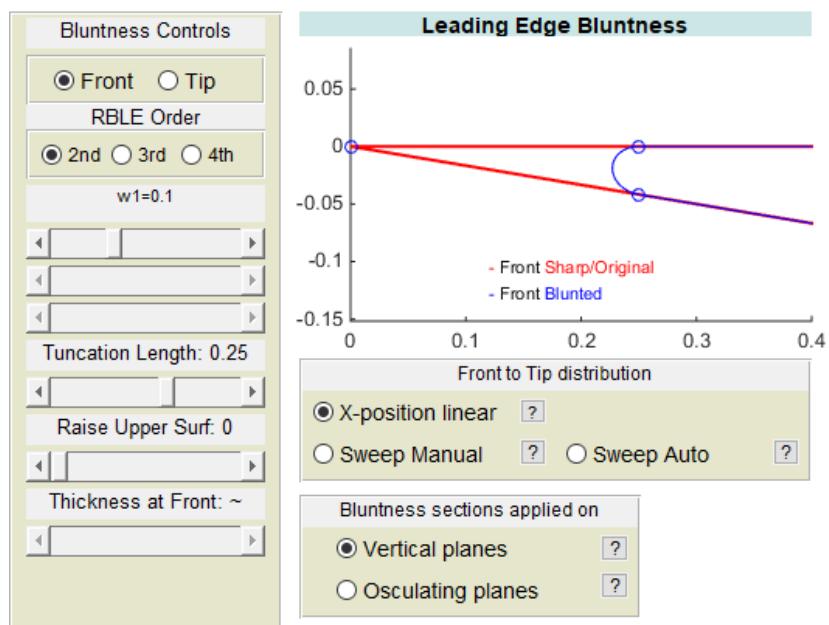
Figure 5.11. Blunt leading edge shape controls of the GUI.

5.3.2   Aerodynamic Performance Evaluation Tool

The on-design aerodynamic performance characteristics of waverider geometries with sharp leading edges can be rapidly evaluated with a reasonable amount of accuracy, as was discussed in section 3.3. The aerodynamic performance evaluation tool used for the design and parameterization flexibility studies, utilizes the flowfield characteristics of the inverse design algorithm along with a semi-empirical boundary layer calculation method in order to estimate the viscous lift and drag characteristics of the geometries. The volume and wetted area of the forebody geometries are also calculated, and it is possible to obtain on-design pitching moment or centre of aerodynamic forces data as well, the latter being important mainly when the waverider forms a substantial part of the entire airframe.

For the force calculations, Mach number values are initially obtained from the Taylor-Maccoll (A.1) conical flowfield solution using equation (5.1) for the lower surface of the waverider and freestream values are used for the upper surface which is parallel to the freestream.

$$M_{local} = \sqrt{\frac{2}{(\gamma - 1)\left(\frac{1}{V_r^2 + V_\theta^2} - 1\right)}} . \tag{5.1}$$

After setting a freestream static pressure and temperature the flow characteristics can be dimensionalized, enabling the calculation of absolute velocities and pressures on the surfaces of the waverider. The characteristics behind the forebody shock are first calculated with oblique shock relations:

$$M_{n\ inf} = M_{inf} \sin\theta_s ,$$

$$M_{n2}^2 = \frac{1 + \left(\frac{\gamma - 1}{2}\right)M_{n1}^2}{\gamma M_{n1}^2 - \frac{\gamma - 1}{2}} ,$$

$$M_2 = \frac{M_{n2}}{\sin(\theta_s - \delta)} , \tag{5.2}$$

$$\frac{p_2}{p_{inf}} = \frac{2\gamma M_{inf}^2 \sin^2\theta_s - (\gamma - 1)}{\gamma + 1} , \tag{5.4}$$

$$\frac{T_2}{T_{inf}} = \frac{[2\gamma M_{inf}^2 \sin^2\theta_s - (\gamma - 1)][(\gamma - 1)M_{inf}^2 \sin^2\theta_s + 2]}{(\gamma + 1)^2 M_{inf}^2 \sin^2\theta_s} , \tag{5.3}$$

with n indicating the Mach number normal to the shock and 2 indicating the values behind the oblique shock.

Isentropic flow relations are then used to calculate values along the streamlines describing the lower surface of the waverider:

$$\frac{p_{local}}{p_2} = \frac{1 + \frac{\gamma - 1}{2}M_2^2}{1 + \frac{\gamma - 1}{2}M_{local}^2} , \tag{5.5}$$

$$\frac{T_{local}}{T_2} = \left(\frac{1 + \frac{\gamma - 1}{2}M_2^2}{1 + \frac{\gamma - 1}{2}M_{local}^2}\right)^{\frac{\gamma}{\gamma - 1}} . \tag{5.6}$$

A boundary layer calculation method can then be applied along each streamline to calculate the viscous forces. Through this process the forces are first calculated on the nodes describing the discretized surface patches of the geometry. The total forces acting on the geometry are then calculated by summing the force of each triangular element/face of the discretized surface patches. No interaction between the boundary layer and inviscid flow was assumed, therefore we do not take into account the displacement thickness of the boundary layer that affects the effective inviscid shape of the forebody. For the semi-empirical boundary layer calculations the local velocity outside the boundary layer is calculated using equation (5.7):

$$u_e = M_{local} * \left(\gamma R_{gas} T_{local}\right)^{0.5} \tag{5.7}$$

For large waverider forebodies and high Reynolds numbers, a reference temperature method that assumes turbulent flow for the entire body is a simple and adequate way to calculate viscous forces. White's (1974) reference temperature method has shown good correlation with experimental results, and has been used for waverider applications before (Lobbia, 2004). For smaller waverider forebodies where the flow can be laminar for the entire waverider geometry or a significant part of it, the integral boundary layer method described by Drela and Giles (1987) was implemented. It is a two-equation integral formulation based on dissipation closure and compatible with both laminar and turbulent flows. Both are described in Appendix D.

The aerodynamic performance evaluation tool can be executed from within the GUI to evaluate the geometry currently being designed or as a standalone in order to evaluate any sharp leading edge waverider described in a ~.mat file. A series of validation cases showing good agreement with CFD calculations have been examined, with two examples shown in Table 3.1 earlier. Nonetheless, accuracy of the results is of lesser importance compared to the sensitivity of the tool in identifying variations between different geometries when it comes to parameterization flexibility studies. On that end, the current tool provides a quick and robust method to calculate and compare viscous aerodynamic performance characteristics of simplified waverider geometries. Its runtime on a modern laptop is around 10 seconds, and can be further reduced by utilizing lower resolution geometries.

5.3.3   Design-Driving Curves

While Bézier curves have been used for the majority of this work, other types of parametric curves can be readily incorporated. This can be accomplished by pointing the main script to a

different subroutine/function where the parametric curve is defined. The MATLAB syntax and I/O structure are as follows, explained using the current Bézier curve function as an example:

```
[coords,der1,der2] = BezierCurve(P,t)
```

Input: *P* is an array containing control point locations (x1 y1 z1; x2 y2 z2; ...; xn yn zn), which can be 3D or 2D, and *t* is a column vector with entries from 0 to 1 where the parametric function is evaluated.

Output: when the function is called with a single output argument the coordinates of the curve for the given values of *t* are output in the first output array (*coords*), x(t), y(t), z(t). When additional output arguments are present in the function call, the second output argument contains first derivative information x'(t), y'(t), z'(t), and the third output argument contains second derivative information of the parametric function, x"(t), y"(t), z"(t). The latter are used for the calculation of the radius of curvature of the shockwave profile curve and are not required for the curve defining the stream tracing origin.

The function can consist of more than one parametric geometry formulations describing different sections of the curve. For example a circular arc can be described for values of *t* between [0 , 0.3) and a spline for [0.3 , 1]. *P* does not necessarily need to contain only control point information but can also contain radii, weights etc. as long as their use is consistent between the function and the function calls. This allows any parametric curve formulation to be utilized or constructed in order to fit specific shapes, which in turn enables applying specific design requirements that might be present directly to one of the design driving curves. This can increase the efficiency of the parametric geometry model as was explained in Section 3.1, and it can also enable parametric flexibility studies comparing more than one type of parametric curve. Additionally, it enables more straightforward interfacing between the forebody and engine inlet design tools, as is explained in the next section.

5.3.4   Engine Inlet Interfacing

A separate engine inlet design tool that can be used to generate compression flowfields and streamtraced inlet geometries for planar compression inlets, axisymmetric spike compression inlets and Busemann-type inward turning inlets was also developed. The inlet geometries are generated up to the beginning of the isolator prior to the combustion chamber. The main waverider forebody generation tool is also able to generate an integrated external compression surface without sidewalls that can replace the first compression ramp of planar compression inlets. A brief description of the three types of inlets and the external compression geometry is given in Appendix E.

The capabilities of the inlet design tool include: external or mixed planar compression inlets consisting of discrete shocks or isentropic compression regions, external or mixed compression axisymmetric cone and spike based inlets, and internal compression Busemann-type inlets. The Busemann-type inlets can also be truncated and fitted with an ICFA axisymmetric flow region that generates a straight shock in the beginning. The exact method to integrate the two types of flowfields is described by Otto (2015). An example geometry for each inlet type can be seen in Figure 5.12. The flowfield generation is fully parameterized and arbitrary capture shapes can be utilized to generate three-dimensional streamtraced inlets. These inlet geometries assume uniform oncoming flow and can therefore be integrated in parts of the waverider forebody were that condition is met (usually regions of the underside designed with that requirement), or they need to be facing the freestream prior to any significant forebody interference.
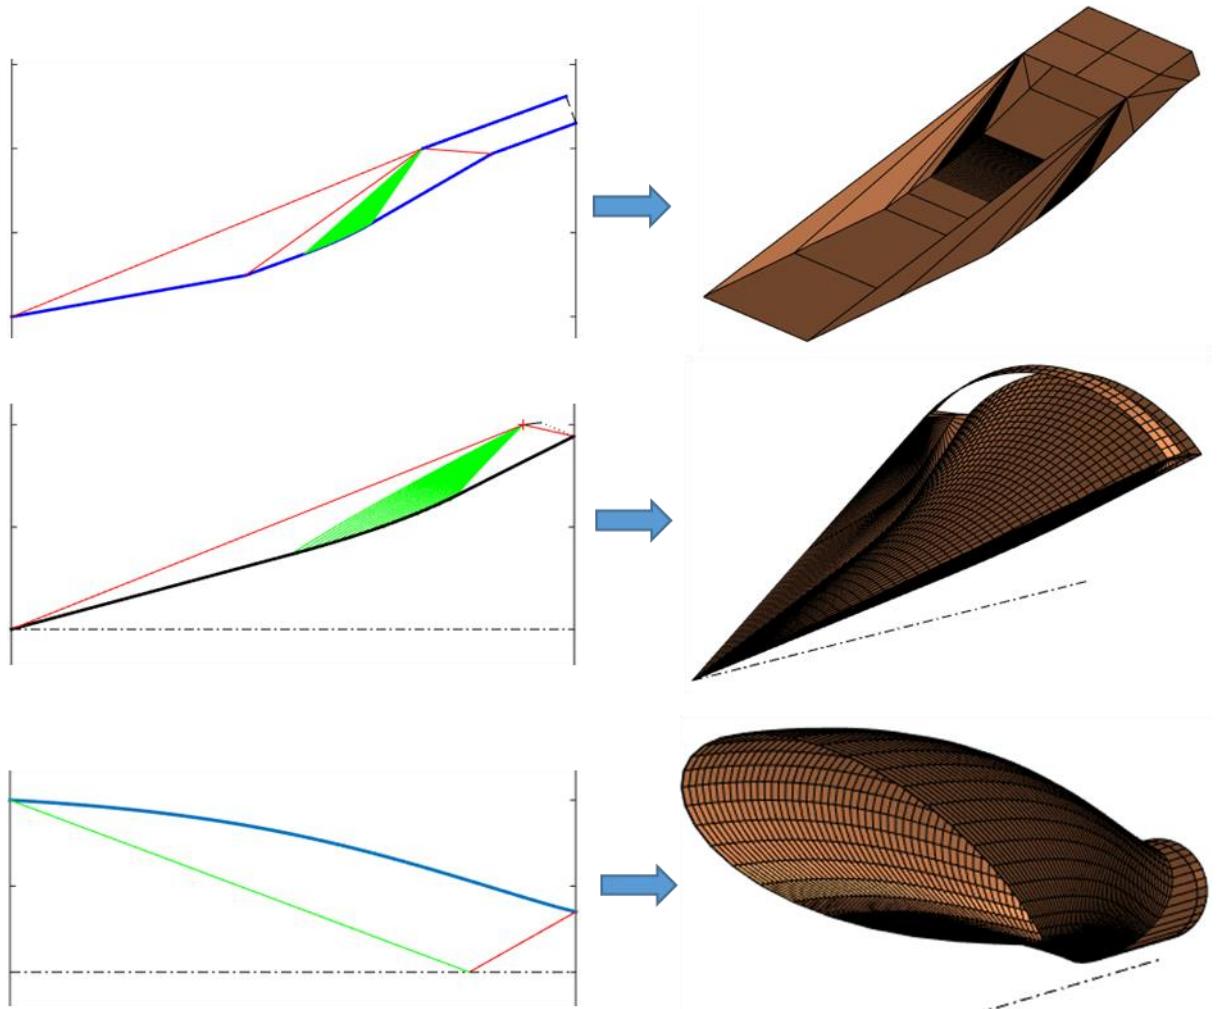
Figure 5.12. The three types of base flowfields and resulting streamtraced inlets, shock waves (red) and isentropic compression Mach waves (green) shown.

The primary reason for developing this tool is to enable a high level of design integration between the forebody and inlet. To that end, the design can be both forebody-driven, with the inlet

designed to match the respective lower surface, for example, of the waverider, or inlet-driven, where the inlet is fully parameterized and the forebody shape adapts to match it. In the first case, an example of which can be seen on the left of Figure 5.13, the waverider forebody is fully controlled by its parametric design-driving curves and part of the capture shape of the inlet depends on the shape of the underside of the forebody. Since this needs to be a uniform flow region of the forebody, the shape of the leading edge of the waverider translates to the shape of the capture shape of the inward turning inlet. The rest of the capture shape is directly designed although certain constraints might still apply for it. For example, given the location of the focus point of an inward turning inlet, the resulting inlet geometry needs be behind the shock generated by the waverider forebody. In the second case, which can be seen on the right of the same figure, the inlet is fully defined and part of the leading edge of the waverider forebody is fully controlled by its capture shape. In this case the LSPD or a hybrid waverider generation method need to be used, with part of the design-driving curve fully defined. Considering the sensitivity and tight margins under which scramjet engines are designed and operate, encountering very specific requirements as far as the shape of the inlet goes is possible, and such an inlet-driven approach would aim to cater for that.
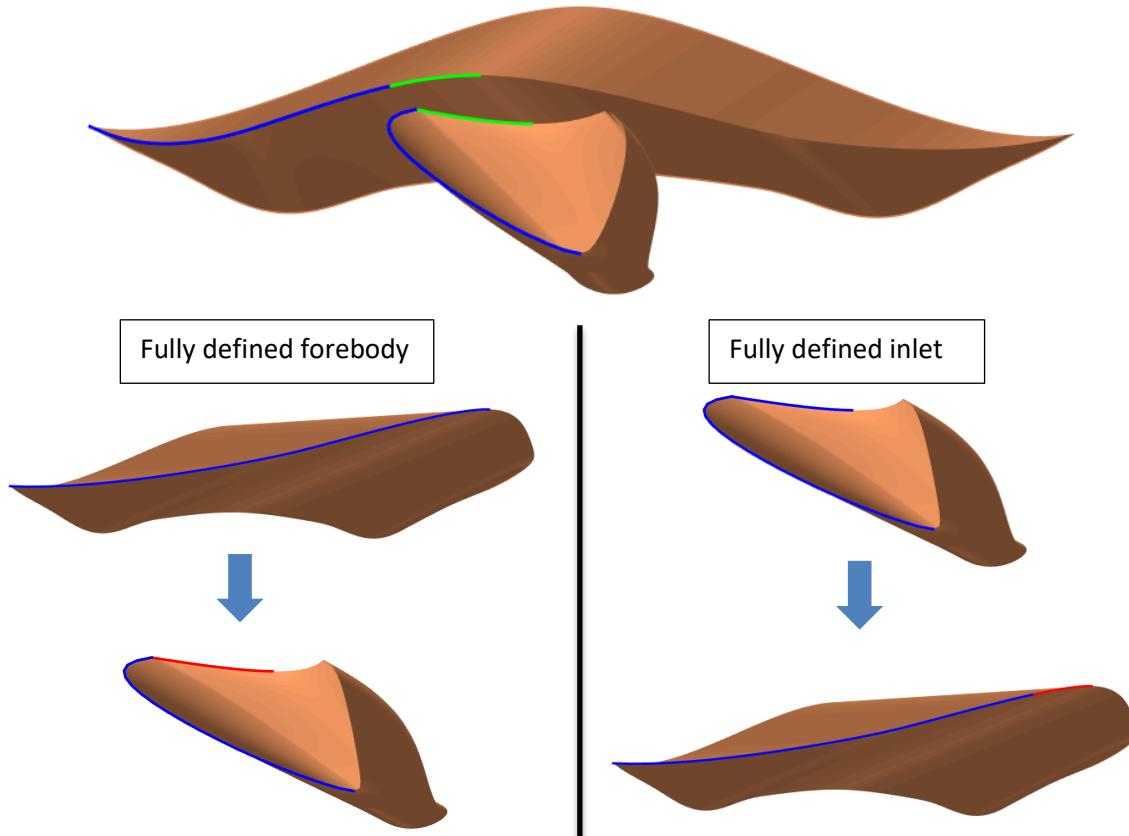
Figure 5.13. Waverider forebody and attached inward turning inlet internal flow-path without housing geometry (top). Integrated parameterization approaches (bottom) with dependent part of the design-driving curves in red and directly designed in blue.

The inlet design tool along with the different parameterization approaches presented in section 3.1 enable the geometry engine to support the development of different aircraft configurations with various inlet integration approaches. The geometry engine is, therefore, not targeted at a single configuration and inlet integration method as has been observed in a number of similar design tools in the literature. It has been developed to be modular and robust in order to enable a wide range of configurations with minimal additional software development effort from the designer, making exploration of different configurations and configuration changes in the early design stages much easier to facilitate.

5.3.5 Dynamic Geometric Fidelity

The geometric fidelity and accuracy of the output geometries has been an area of interest in our collaboration with the industry throughout this project. While one aspect of the geometry parameterization is providing enough flexibility to the design driving curves in order to enable certain shapes, the other aspect is the accuracy and required resolution of the geometry in order to accurately represent the underlying formulations.

The typical process that a computational geometry goes through prior to a CFD simulation, for example, can be seen in Figure 5.14. It should be noted that while a number of parametric geometry formulations can be directly used by grid generation or even CAM software, in this case it would be a much more complex task due to the nature of the inverse design algorithm that involves streamtracing a flowfield that is usually described by differential equations. Therefore, at this stage, we are interested in the information that can potentially be lost when moving from the first stage to the second one, the latter being the output of the geometry engine. Generally, the accuracy of CFD results will not be meaningfully affected unless the quality of the geometries significantly degrades. However, importing the geometries to other CAD software to further develop them is a process more sensitive to geometry tolerances.

A number of geometries were initially investigated, with more detailed information found in Appendix B. We are however developing a parametric geometry model, and while measuring a set of geometries at this stage can give us a certain level of confidence to using 100x100 surface resolutions for example, it is not a robust enough approach. Doing the same analysis for a large variety of generated shapes would be one method to statistically increase our level of confidence. There is, however, another more robust and efficient approach to this.
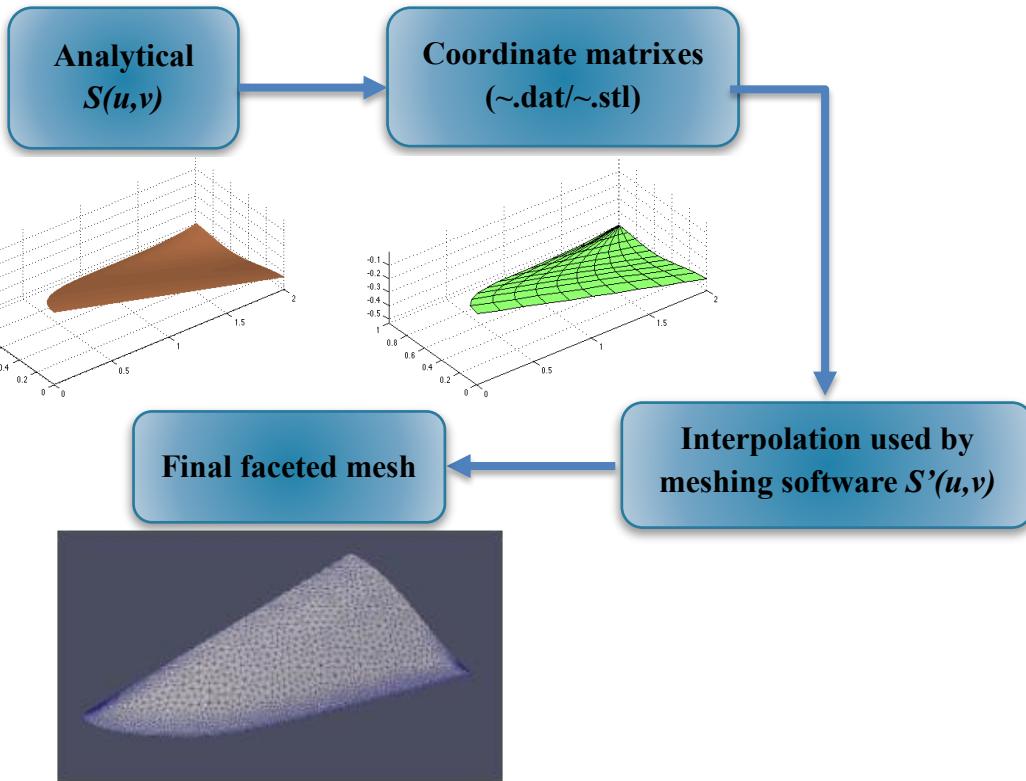
Figure 5.14. Example of stages a computational geometry goes through before a CFD analysis.

The desired resolution of the surfaces is determined at the start of the waverider generation algorithm. This is expressed as the number of osculating planes in the spanwise direction and the number of points describing each streamline in the streamwise direction. In addition to that, a maximum allowable angle deviation for subsequent line segments of the discretized feature edges of the faceted geometry can also be manually set. This way, as the osculating planes are being drawn and the leading edge points calculated, if the angle between two successive line segments is over the defined limit, additional osculating planes are drawn in-between. The approach is illustrated in Figure 5.15. If the angle is still too large after the subdivision, the process is repeated by adding a point/plane in the previous gap. This way most resolution transitions between consecutive line segments are limited to double/half and infinite loops that could otherwise be encountered under certain conditions are also avoided. This dynamic method of controlling the fidelity of the generated geometries was implemented in the geometry engine for the main feature edges of the waverider forebody as well as the streamtraced inlet generation tools. It is not necessary to apply this in the streamwise direction along each streamline since sudden changes in their curvature are not expected under any flow conditions given the nature of the flowfields utilized. An example geometry where the resolution has been adjusted in the area where the leading edge was more tightly curved can be seen in Figure 5.16.
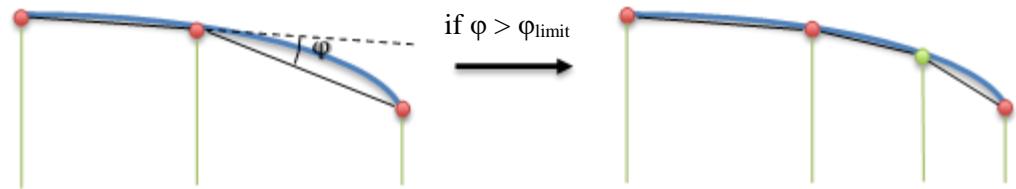
Figure 5.15. Method of drawing additional osculating planes when angle between line segments is above the limit.
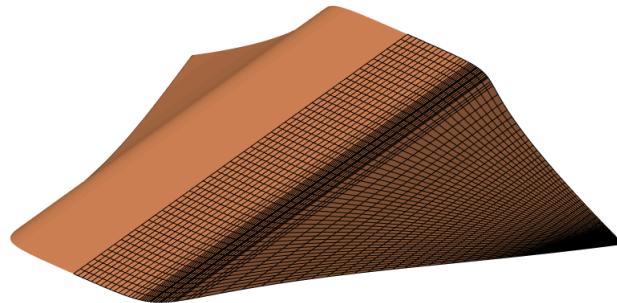
Figure 5.16. Dynamic resolution adjusted waverider forebody, initial surface resolution: 50x50, angle limit between line segments: 5^0.

5.4   Completing the Aircraft Geometry

The majority of the investigations and software development throughout this work concerns the parameterization of the waverider forebody geometries, as that is the primary component that is not readily available through conventional CAD approaches. Additional aircraft components can be designed utilizing conventional parametric geometry formulations. Parameterization options of these directly designed components required to complete a full waverider-based aircraft geometry are not as easy to justify with methods as the one presented in Chapter 3. Performing parameterization flexibility studies starts becoming more complex as it is not possible to isolate the design of geometry components behind the forebody when its design is not finalized. The design experience and intuition of the designer starts being more important at this stage, and there is a much wider variety of approaches one can follow to complete the aircraft geometry as well. The geometry model from this stage on can also be completed with the aid of commercially available CAD software. For these reasons, further manipulation of the waverider forebody geometries and the integration of additional geometry components has been implemented in a separate tool, mostly as an example on how the forebody shapes can be further integrated in a full aircraft geometry engine.
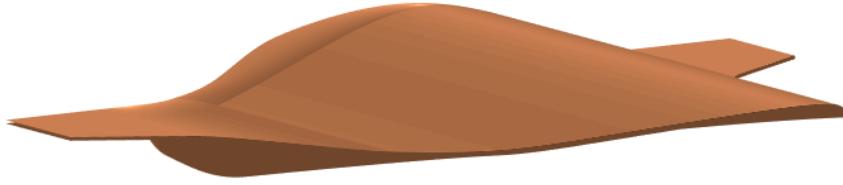
Figure 5.17. Simplified waverider-based aircraft geometry.

5.4.1 Upper Surface Manipulation

The upper surface of the waverider can be manipulated without affecting the 'waveriding' properties of the lower surface as long as the leading edge shape is also not affected. This can contribute in improving the aerodynamic characteristics of the geometry, the integration of additional geometry components such as a canopy, and also to alter the distribution of available volume for the packaging of internal subsystems. There is a plethora of available options for changing the shape of the upper surface, and a straightforward method to do that is a simple form of free-form deformation that will be described here.

A parametric surface patch, such as a Bézier or NURBS, can be mapped to the rectangular coordinate matrix of the upper surface, and then movements of its control points can be translated to perturbations of the upper surface. In our example, displacement in the vertical (z) direction was mapped so that the vertical displacement of each point (i,j) of the coordinate matrix of the upper surface can be obtained by calculating the $z(u,v)$ of the parametric patch, with:

$$u = \frac{i-1}{Ni-1},$$

$$v = \frac{j-1}{Nj-1},$$

(5.8)

Ni and *Nj* the dimensions of the upper surface patch coordinate matrix.

The control points on the leading edge are restricted (blue in Figure 5.18), the remaining control points can freely move (red), while movements of the control points on the symmetry plane (green) should be related to movements of the others in order to maintain surface continuity and smoothness. Additionally, the upper surface should not be allowed to intersect the lower surface. With all control points at *z=0* the upper surface remains unchanged. An example of a waverider forebody with a deformed upper surface can be seen in Figure 5.19.

Bump functions and other types of geometry manipulation techniques can be applied in a similar manner. Additionally, the aforementioned mapping and deformation can be applied only on part of the upper surface if needed. It is also possible to take such changes of the upper surface shape into account with a simple aerodynamic performance evaluation tool. This could be achieved by applying a local inclination-based calculation method as the one utilized by Lobbia (2004), which could enable parameterization flexibility studies for this part of the geometry model as well.
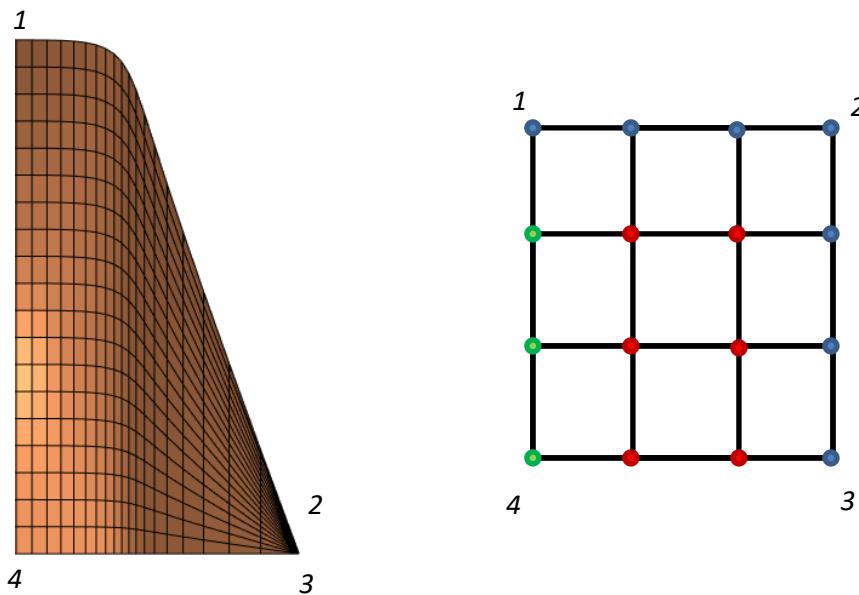
Figure 5.18. Parametric surface patch mapped to the upper surface of the waverider.
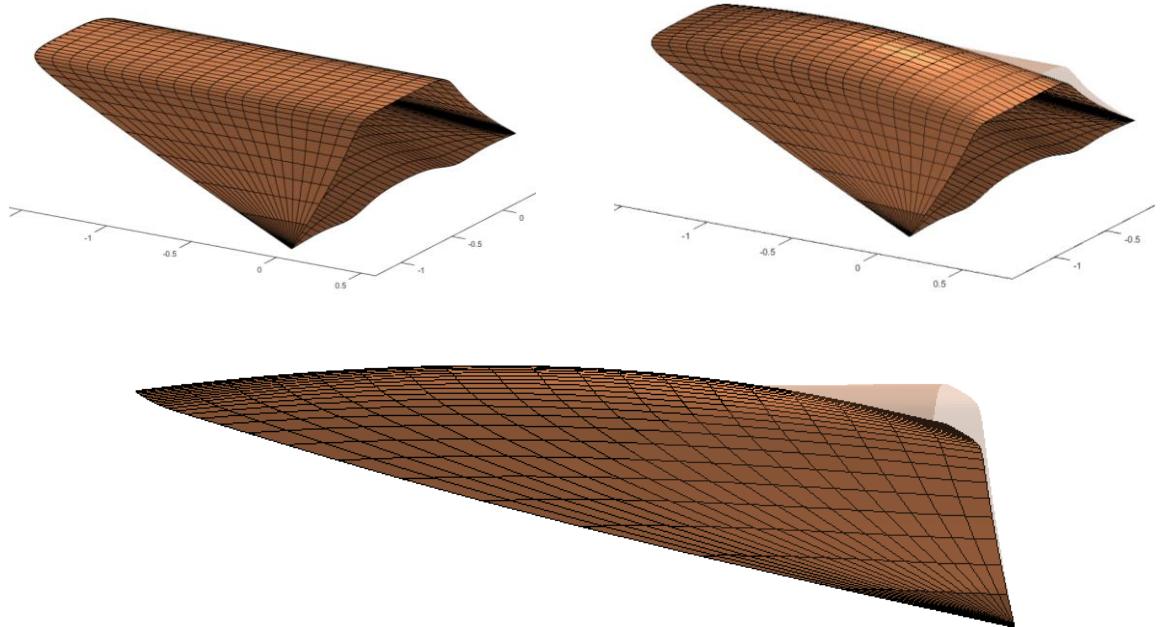
Figure 5.19. Original waverider forebody (top-left) and resulting waverider forebody with deformed upper surface.

5.4.2 Aft Section of the Geometry

Whether the inversely designed waverider geometry forms the majority of the entire airframe or is just a small part of it at the front, its aft end needs to be completed or interfaced with additional geometry. If the waverider does form the majority of the airframe then one option is to manipulate the upper surface in such a manner that it ends up being a watertight waverider geometry, while leaving space for a nozzle at the back. Examples of this can be seen in the work of Lobbia (2004) where the central part of the airframe has volume to hold the payload and has the engine nozzle at the aft end, while the upper surface at the sides collapses to the lower surface in order to form a delta wing type airframe.

However, when additional geometry is required, a simple method to extend the airframe and design the aft end is by adding cross sections of controlled shape and completing the geometry utilizing a lofting approach. By allowing the designer to use any number of cross sections a wide range of parametric geometries can be generated, from very simple ones utilizing just one cross section at the end of the geometry to much more sophisticated and complex ones. The method implemented is based on B-splines for lofting. The B-splines are drawn in the general streamwise direction, from the end of the waverider geometry to the aft end of the airframe which is defined by a final cross section curve. At least tree control sections are used with the second control section being restricted in order to satisfy tangency conditions at the interface of the two surface patches. The smooth transition at the interface can be controlled with a single parameter that determines the distance of the tangency control section where the second control points of the B-splines lie, as seen in Figure 5.20.
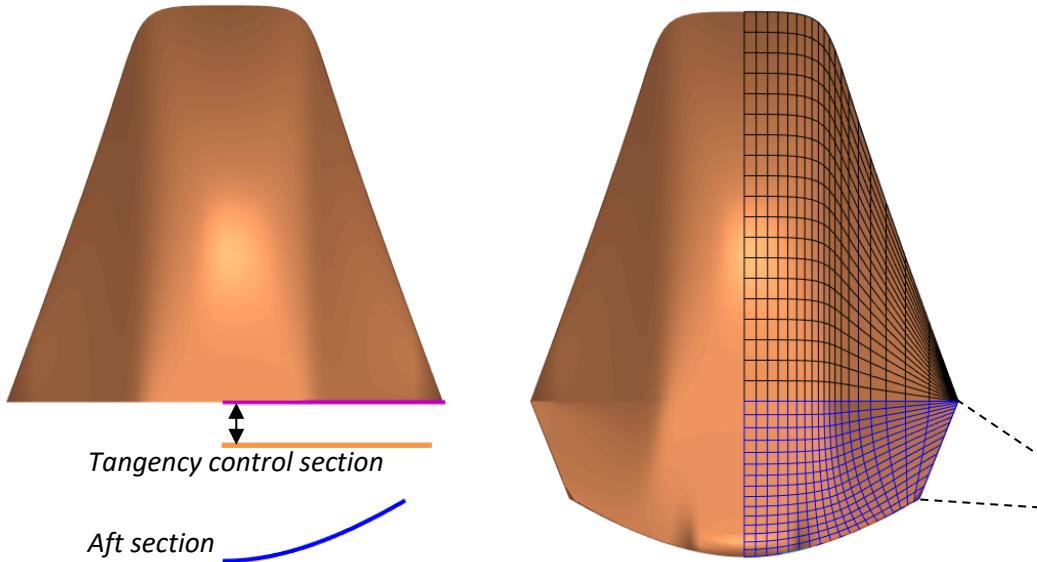
Figure 5.20. Plan view of lofting process. Waverider interface (purple), tangency control line (orange) and aft cross section (blue). Potential winglet integration illustrated.

Both upper and lower surfaces can end on a single curve -resulting in a watertight geometry, or a shape that matches the engine exhaust nozzle, as seen in Figure 5.21. More control can be provided by using additional controllable cross sections between the interface with the waverider geometry and the aft end shape, something that becomes more meaningful for airframes much larger than the waverider forebody. The issue, however, is that with each additional cross section the number of design variables increases significantly. The minimum number of design variables required for this type of geometry are design variables controlling the 3D curve or cross section at the aft end of the geometry and a tangency parameter at the interface which can potentially be fixed in a design study. It is also worth noting that even the lower surface of the baseline waverider can be truncated and does not have to be used as a whole. This is as long as the lower surface close to the leading edge is not changed and no drastic changes to the slender shape of the body are made.

Smoothness maintained at interface
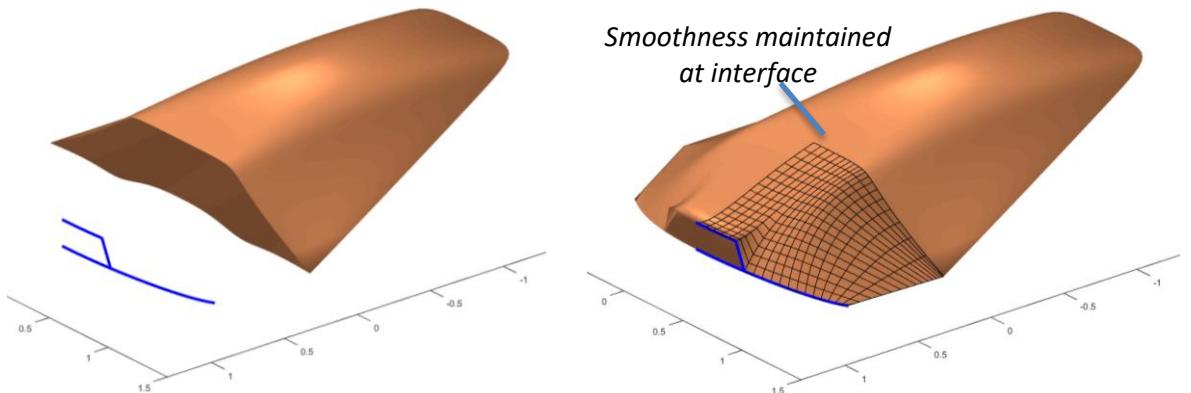
Figure 5.21. Lofted surface at the aft end of the geometry, controlled with a single cross section at the end.

5.4.3   Winglets and Interfacing Geometry at Sides

Utilizing a similar approach to the latter, additional geometry such as winglets can also be attached to the sides of the waverider geometry. In the current implementation a given wing geometry can be placed in relation to the forebody and an interfacing geometry to connect the two can be generated with the aforementioned B-spline based lofting approach. The waverider geometry can be truncated at a chosen osculating plane and only the central part is then used, as seen in Figure 5.22. In this case tangency at both ends is required and is controlled with two tangency parameters that control the distance of the tangency control sections from the two ends of the interfacing surface patch, as seen in Figure 5.22. Apart from controlling the tangency at each end, in certain cases the tangency parameters need to be tuned to avoid self-intersecting geometries. The design parameters involved in this case are the location of the winglet and the two tangency parameters which can potentially be fixed in a design study. The winglet might also need

be designed separately, something that is not covered here. An example geometry of a blunt leading edge waverider with winglets and upper and lower interfacing geometries can be seen in Figure 5.23.

Tangency parameters control the distance of the tangency maintaining control sections from the interface.

Interface with aft sections
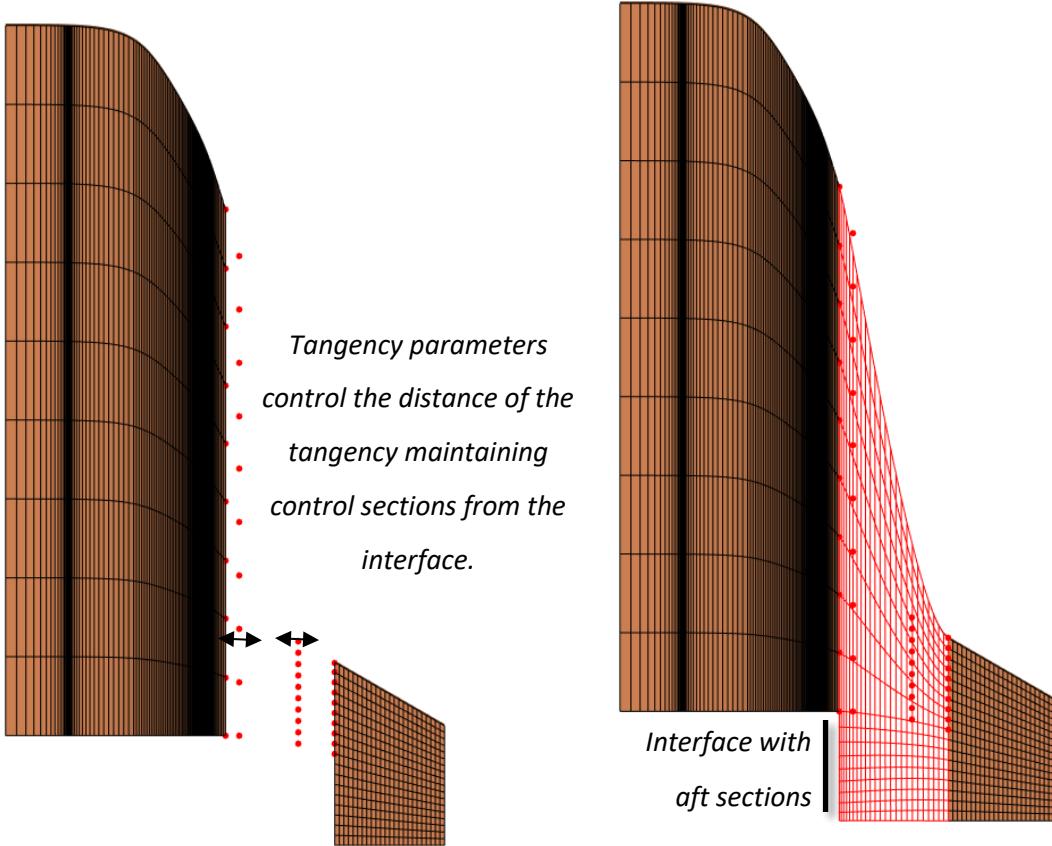
Figure 5.22. Plan view of the lofting process for the winglet interface.
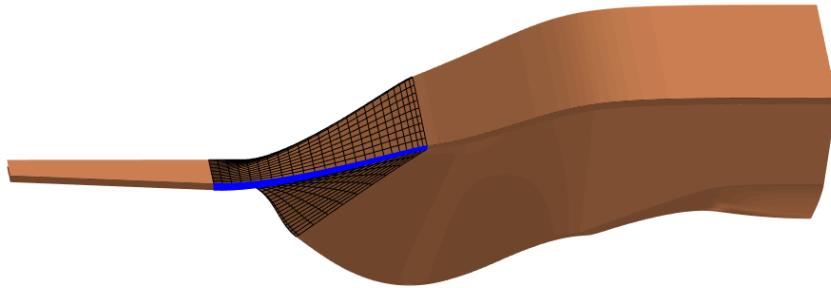
Figure 5.23. Front view of the interfacing geometry (black-wired), with the blunt part highlighted (blue).

Additional geometry components can be designed in a similar manner or even with the use of tools available in commercially available CAD software. The parametric geometry options discussed in this section are meant as examples of approaches that can provide increasing amounts of flexibility based on the requirements of the designer. It should be noted that the example geometries in this section were arbitrarily chosen to showcase the parametric geometry structure and do not necessarily reflect functional airframe shapes.

5.5 Parametric Configuration Example

In this last section of this thesis we will have a look into how the geometry engine can be utilized to develop complete –from a conceptual design perspective– parametric aircraft geometries based on waverider forebodies. A configuration that employs a less commonly used approach of engine inlet integration will be described. It is a dorsal mounted inward turning inlet approach, similar to that of the MR2 concept (Figure 2.30). The most common and straightforward integration approach found in the literature are underbelly mounted planar compression inlet configurations. The purpose of this section is to showcase the robustness and versatility of the geometry engine that has been developed, especially in handling waverider forebody geometries.

5.5.1 Dorsal Mounted Inward Turning Inlet Configuration

For this example we will follow an inlet-driven design approach, where the capture shape of the inlet is directly designed and the waverider forebody is then parameterized to accommodate it. In a realistic design scenario such an approach would be preferable when the performance of the inlet is a sensitive area of the design. This applies for inward turning inlets specifically due to their starting characteristics being greatly affected by their capture shape. The core of the parametric geometry model for such a configuration is therefore the inward turning inlet. The inlet geometry is controlled by the parameters defining its base flowfield and a parametric capture shape. The streamtraced inlet generation process can be seen in Figure 5.24.

Flowfield Definition:
- M_{inf}
- Free-standing Conical Shock Angle
- Busemann Truncation Angle

Base Flowfield

Capture Shape
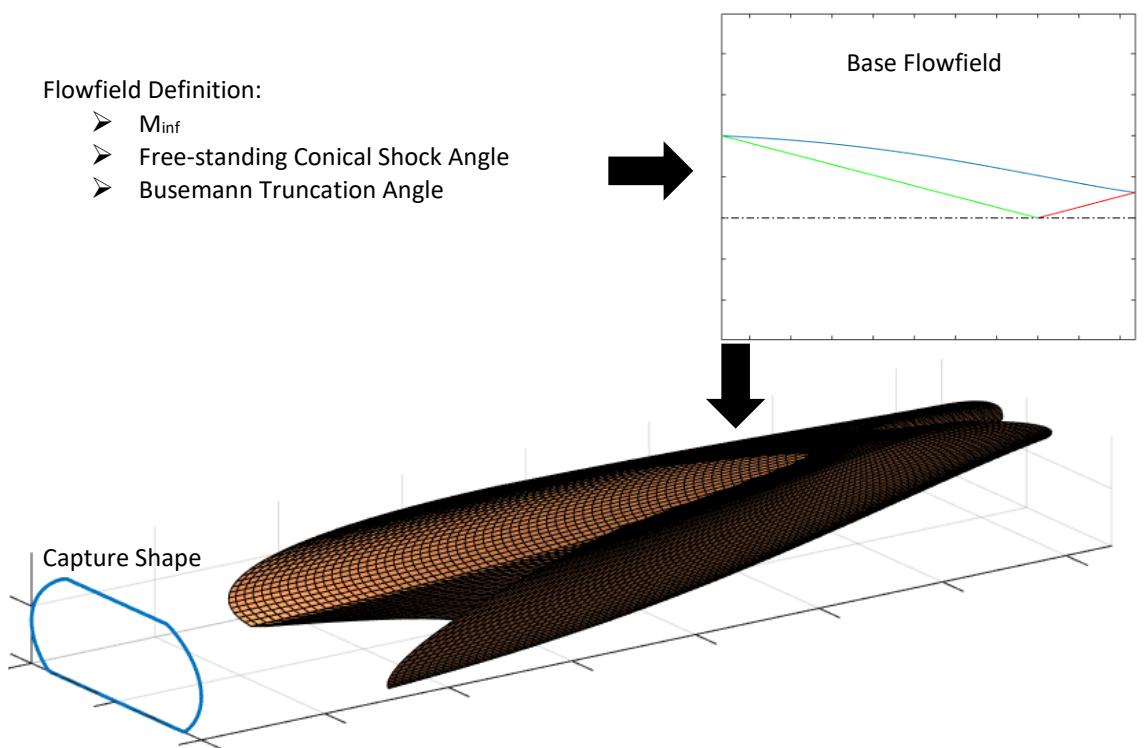
Figure 5.24. Streamtraced inward turning inlet parameterization.

The rest of the engine flow-path for this model is simplified, containing a straight duct that serves as the isolator and combustion chamber followed by a conical-type nozzle that is controlled by placing an end section at the aft end of the aircraft to expand the exhaust gasses. The MR2 concept that incorporated a similar design, utilized a planar expansion to first transition to a circular section for the first part of a similar nozzle. The engine flow-path can be seen in Figure 5.25.
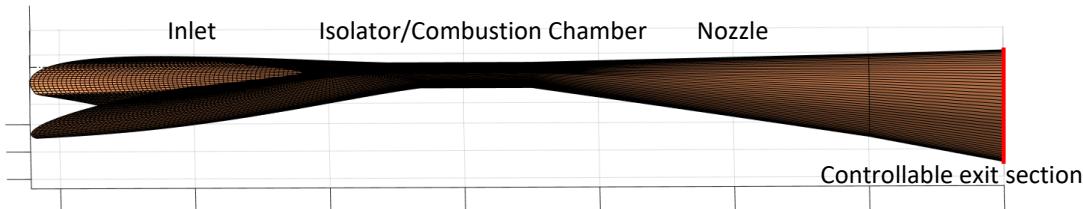
Figure 5.25. Engine flow-path.

In order to match the inlet with the waverider forebody, part of the capture shape needs to also form part of the upper surface profile curve of the forebody. This can either be accomplished by restricting the design of the central part of the upper surface profile curve when utilizing the USPD method, or by using a hybrid design approach where the unrestricted part of the geometry can be parameterized with a different design method. In this case the second approach was chosen with the rest of the stream tracing origin curve being defined with the PLED approach which gives direct control over the planform shape of the lifting body. Part of the capture shape is therefore used for the USPD part of the parameterization and a parametric curve that controls the rest of the geometry utilizing the PLED approach. The starting point of the planform designed leading edge curve is constrained to the end point of the USPD inlet-constrained section.
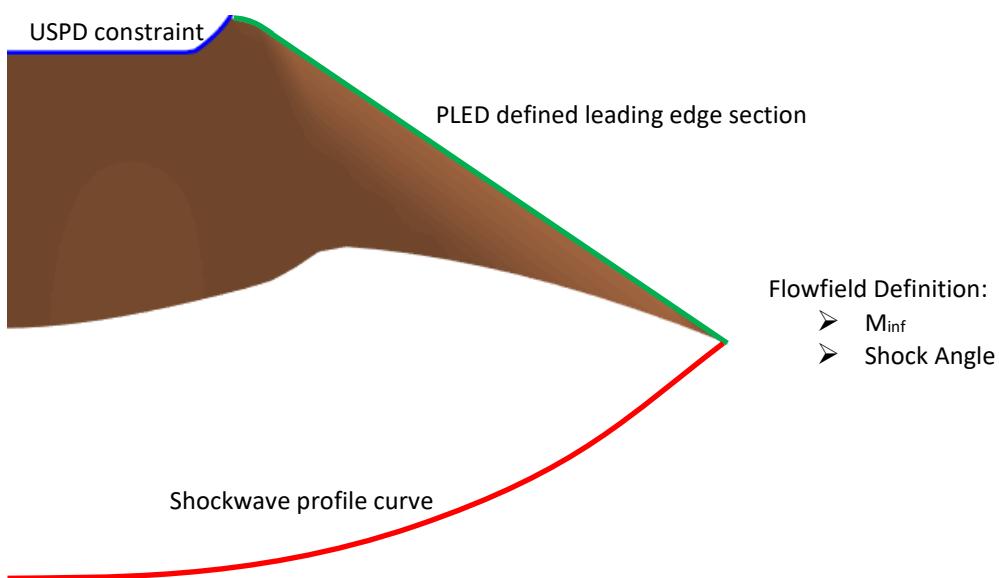
Figure 5.26. Front view of half-symmetric waverider and design-driving curves.

Side view

Front view

Top view

Freedom of inlet placement
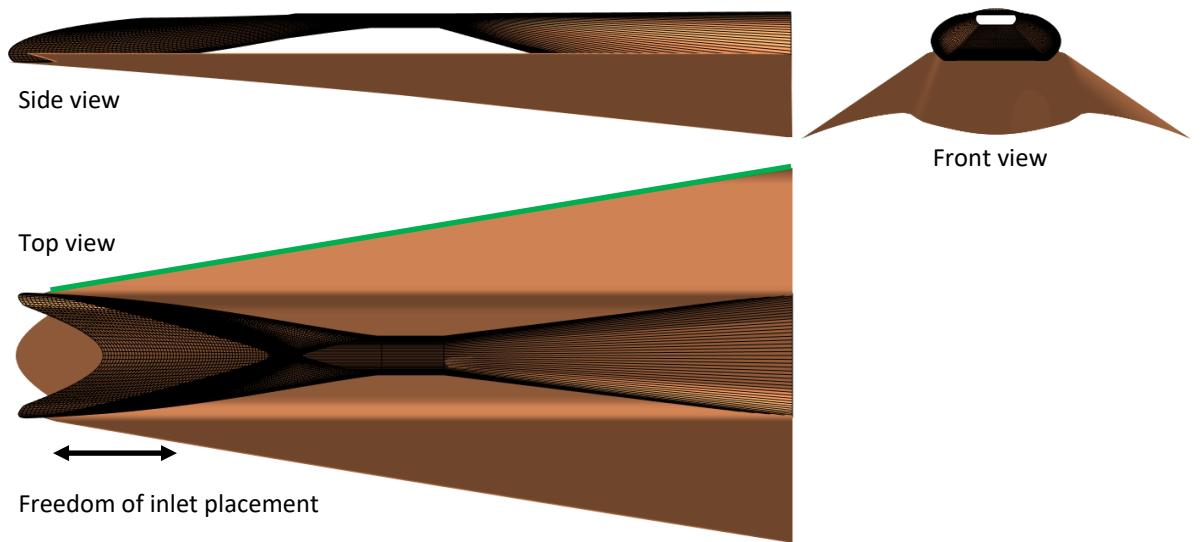
Figure 5.27. Perspective views of the matched waverider and engine flow-path.

With the inlet capture shape matching the shape of the upper surface profile at the central section of the waverider their matching is straightforward, seen in Figure 5.27. The only additional parameter that dictates their integration is the streamwise position that the inlet is located, since the flow over the top of the waverider forebody is uniform and there is a certain freedom in placing it. From this point on, the geometry is completed with more traditional CAD approaches, like the ones that were described earlier in section 5.4. In Figure 5.28 the additional surfaces used for housing the engine flow-path and closing the aft end of the geometry can be seen.

The engine flow-path is covered with a lofted surface (Figure 5.28, green). The upper surface of the original waverider geometry that is inside this lofted surface is not needed. Both upper and lower surfaces of the original waverider geometry do not need to be used as a whole and can be truncated and replaced or morphed. The 'waveriding potential' of the geometry will not be affected as long as the lower surface around the leading edge region is not altered and the body remains slender. For this configuration part of the original waverider lower surface is replaced with an interfacing surface patch and the upper surface is deformed in order to transition the geometry to the desired aft section shape. The aft section is defined by: the end of the nozzle housing geometry, a part of the lower surface profile that was not truncated, and additional directly designed geometry. The lower surface patch that has been added (red) is used to smoothly transition the lower surface to the designed aft section shape and the upper surface is deformed (purple) to meet the lower surface forming an edge that closes the geometry. This edge of the aft section shape can be designed to accommodate additional control surfaces. Finally, two vertical stabilizers that have been directly designed in CAD software are added on the engine housing surface to complete the geometry model, which can be seen in Figure 5.29.
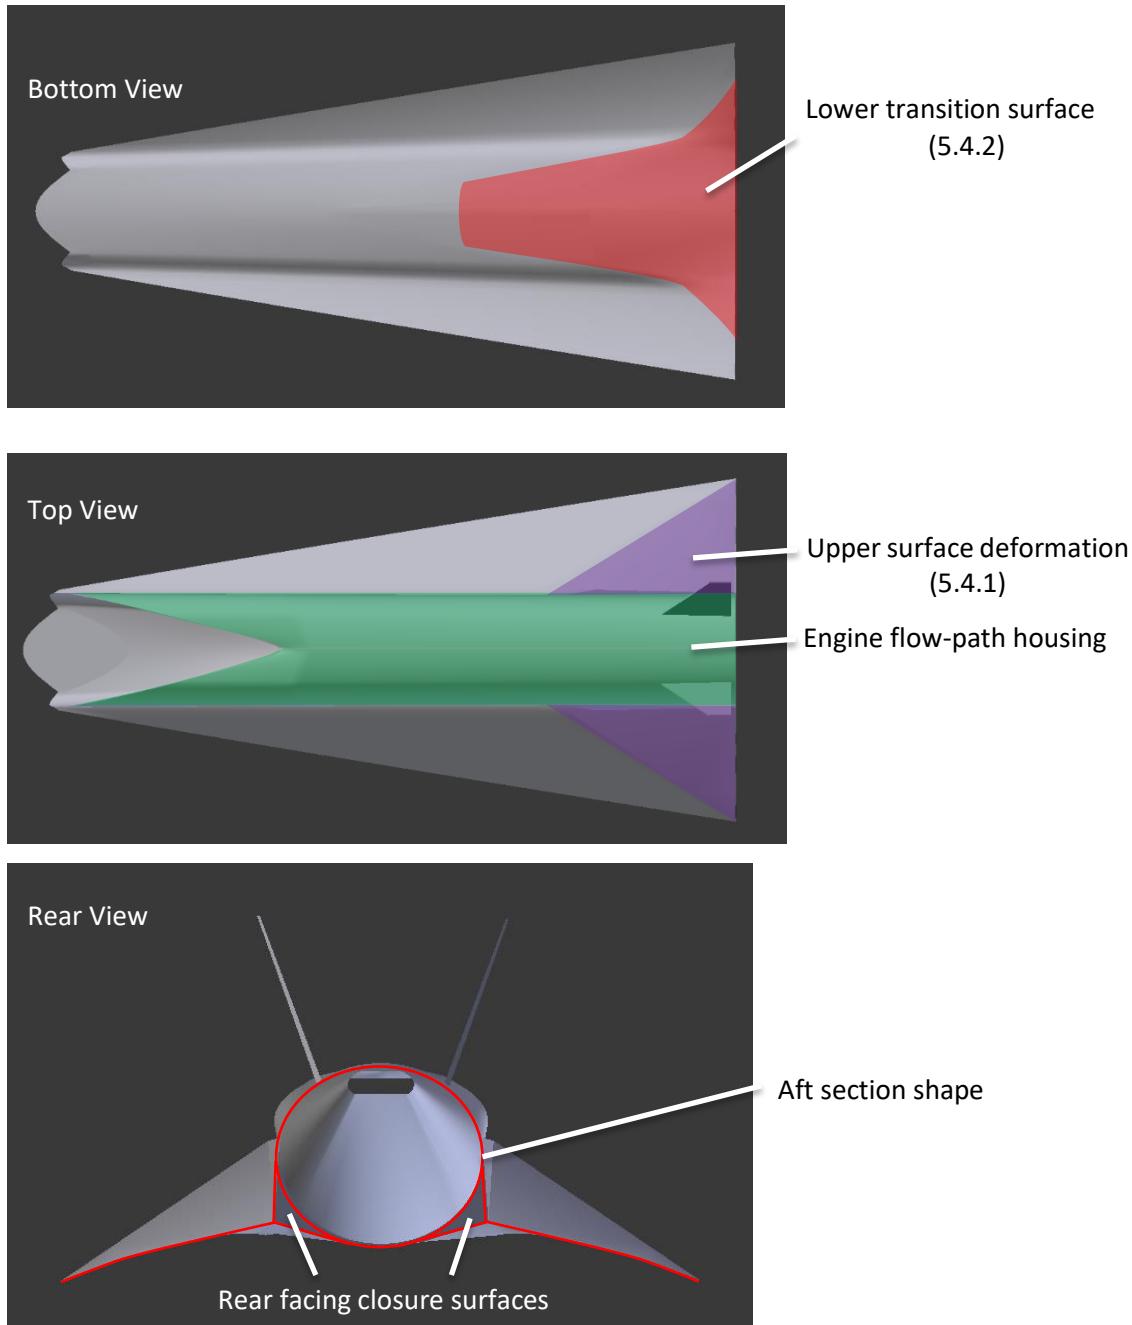
Figure 5.28. Engine flow-path housing and aft section completion.

The design is parameterized with a set of design-driving curves and several primary and secondary design variables. The parametric curves and shapes that are used to define the geometry are:

➢ Inlet capture shape
➢ Waverider planform leading edge curve
➢ Waverider shockwave profile curve
➢ Aft section closure shape
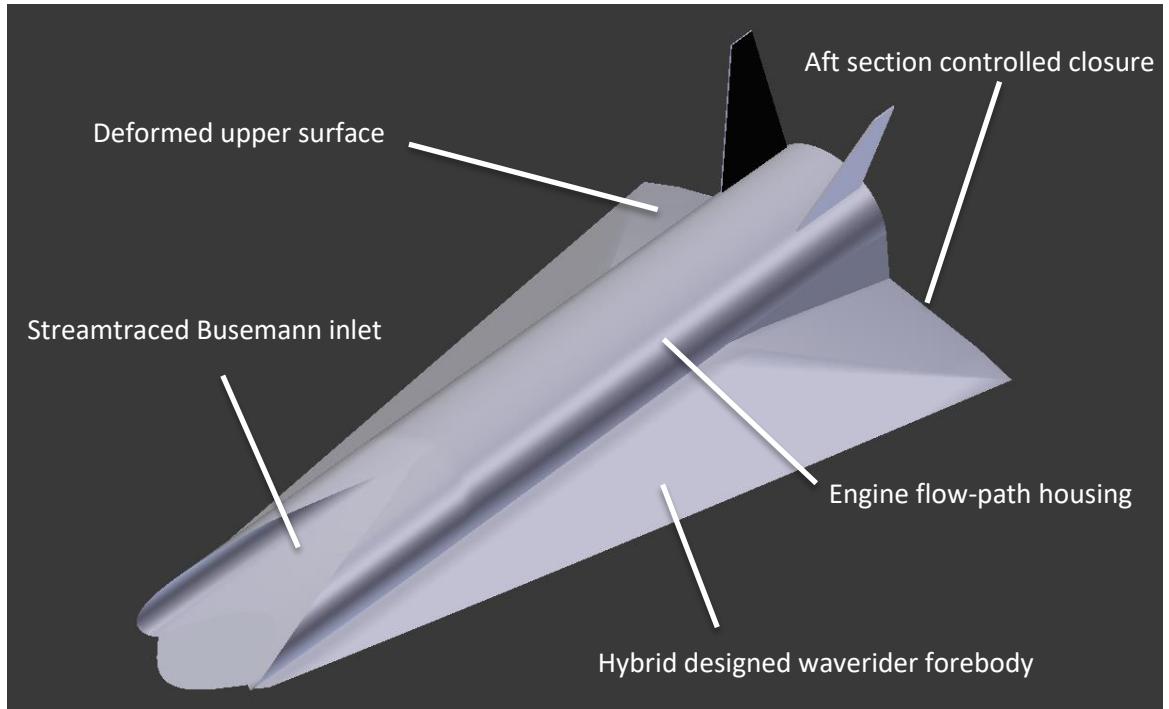
Figure 5.29. Perspective view of the waverider-based hypersonic aircraft model with a dorsal mounted inward turning inlet.

Among the primary design variables that dictate the design are:

➢ Freestream Mach number

➢ Shock angle of the waverider forebody

➢ Shock angle of the conical free standing shock of the inlet

➢ Busemann inlet truncation angle

➢ Inlet streamwise placement variable

➢ Isolator-combustor length

➢ Nozzle length (together with the isolator, combustor and inlet length they also define the length of the waverider)

➢ Nozzle area ratio

➢ Control surface shape characteristics and placement

The secondary design variables include parameters that control the transition and lofted surfaces and how smooth the transition from one shape to the other is, parameters that define the points where the original waverider geometry ends, etc. These are generally design variables that, especially in the early design iterations, would be frozen to values that work for a range of geometries or their parameterization would not need to be independent but could be controlled by other features of the geometry model.

By allowing around five degrees of freedom to each design-driving curve and shape, the total number of primary design variables will be around 30. Parameterization flexibility studies can, at this stage, assist in the parameterization of these design-driving curves and shapes. Depending on how much control we wish to have over secondary features the number of design variables will further increase, especially when considering that this applies for a very simple engine flow-path and for a waverider with a sharp leading edge. For such high dimensionality design problems each degree of freedom that can be saved without restricting the capabilities and potential of the parametric geometry model can have a significant effect on subsequent design optimization studies, something that puts the high dimensionality statements and the parametric efficiency considerations of the previous chapters into perspective.

To parameterize this hypersonic aircraft configuration the additional software and code development required by the user was limited to: an algorithm for truncating the upper and lower surfaces of the original waverider geometry and linking the generated surfaces, writing the desired design-driving curve and shape formulations and integrating them in the geometry engine, and generating any geometry that wasn't already included in the capabilities of the geometry engine such as engine flow-path sections and control surfaces. More diverse waverider-based hypersonic aircraft configurations can be developed with a similar process given the ability of the geometry engine to develop coupled parametric geometry models utilizing the approaches that were discussed in Chapter 3. The versatility and ease of development of multiple configurations with varying degrees of geometric flexibility satisfy the primary requirements set for enabling and supporting the revised aerodynamic process that was described in the beginning of this chapter. Finally, apart from demonstrating a design tool, the purpose of this chapter has been to show how the proposed design methods and geometry parameterization approaches can be practically implemented and utilized to develop parametric geometry models for more realistic design studies.

Chapter 6:   Conclusions

Hypersonic aircraft design is a very active research field with numerous design concepts and ideas emerging from a number of institutes and research facilities around the world. The multidisciplinary nature of this field is reflected in the plethora of active research areas, ranging from design of airframes and engines to thermal protection systems, understanding of hypersonic fluid dynamics features and flight control systems.

While some of the geometry generation tools and concepts for hypersonic applications have been around for the better part of a century, difficulties in a number of the other research fields have slowed down their application and in some cases their evolution as well. Inverse design methods for generating waverider forebody geometries have been utilized in a great number of studies, with optimizations conducted on all aspects of their aerodynamic performance. However, details on handling and parameterizing them were identified as an areas where current literature was lacking. Both the design algorithms themselves and the parameterization of the underlying design-driving curves that most of them utilize, could be further investigated to improve the effectiveness and efficiency of parametric aircraft geometries based on them. The importance of efficiently parameterized geometry models is highlighted by the multidisciplinary nature of the problem and the limited collective design experience in this design field. This is in comparison to the transport aircraft design field, where much more design experience has been accumulated. For example, when designing an airliner the engineer knows that an efficient and robust method for parameterizing the fuselage in the early stages is to use circular cross sections controlled with a radius, and there is an understanding of the benefits and drawbacks of using something more sophisticated.

Chapter 6

6.1 Contributions

The focus of the present work has been parametric geometry model development for the outer mould lines of waverider-based hypersonic aircraft. This can be divided in three distinct areas: 1) investigations into efficiently handling and parameterizing existing waverider design methods, 2) introducing new design tools and methods for hypersonic aircraft design, and 3) implementation of specific considerations to aid with the effective parameterization of hypersonic aircraft geometries, which can range from very specific features of the design tools to specialized design processes.

The main contributions of the current thesis can be summarized as:

- Different approaches to parameterizing waverider inverse design methods in order to enable the geometry to more efficiently meet potential design requirements.
- New waverider design method based on 3D leading edges, which is a clear departure from the notion of specifying the flowfield/shockwave shape beforehand.
- Parameterization flexibility studies of the design driving curves as a method to avoid over/under-parameterizing the geometry.
- New parametric shapes for blunt leading edges, more efficiently parameterized and with potential advantages in turbulent transition.
- A revised design methodology implementing parameterization flexibility studies and low fidelity design space explorations to tackle a number of hypersonic design challenges.
- A geometry engine supporting the parametric model efficiency features and design processes with a few added novel features, such as real-time geometry manipulation.

With the majority of the research being conducted while developing the parametric geometry engine, more findings were made on a number of areas. These secondary contributions include:

- Quantification of the increase in complexity when indirectly applying design constraints that could otherwise be directly applied by utilizing the appropriate waverider parameterization approach (USPD, LSPD, PLET or hybrid).
- A set of requirements for parameterizing valid leading edge curves for the 3D leading edge waverider design method.
- Two approaches for adjusting the deformation of the planform shape of the waverider when a blunt leading edge with varying thickness across the span is applied.
- An interactive GUI that can potentially speed up the development of multiple parametric geometry models and increase the efficiency of the human element in the design process.
- Dynamic geometric fidelity of the parametric waverider forebody geometries to ensure high quality geometries that can also be completed in external CAD software.

- A versatile and flexible coupled parameterization of the waverider forebody and different types of inlets in order to improve the efficiency and robustness of the parametric aircraft configurations.

- Examples of how the aft sections of the aircraft that are interfaced with the waverider forebody can be designed and parameterized.

The three approaches to parameterizing the waverider inverse design method allow the direct enforcement of potential design requirements and constraints on the feature edges of the geometry. Interfacing the geometry with additional geometry components can also benefit from that. Overall, utilizing the appropriate parameterization approach or a hybrid one can increase the efficiency and robustness of the parametric geometry and the design process. The first, USPD, approach is the one most commonly used, while the second, LSPD, has been used in studies involving integrated engine inlets. There seemed to be no detailed description of the PLED approach in the literature although it might have been used on occasion. Apart from systematically documenting the three approaches, we attempted to quantify the potential benefits of using the one that best fits potential design requirements or constraints. Utilizing hybrids of the approaches can assist in more complex design problems where the geometry is subject to constraints and also needs to be interfaced with additional components.

The novel waverider design method based on three-dimensional leading edge shapes that was presented can also be seen as an additional method to parameterize the osculating cones and similar design methods. However, apart from the parametric model efficiency benefits similar to the three aforementioned approaches, it is also a clear departure from the notion of designing the shock generating flowfield first. It can, therefore, also be seen as a step away from inverse design and one towards direct design of waverider geometries. Any arbitrary three-dimensional leading edge curve is not fit for functioning as the leading edge of a waverider. However, testing the capacity of a curve to be matched with a waverider compatible shockwave required negligible time and computational effort. High levels of geometric continuity of the leading edge curve are generally preferred and a number of constraints that can be applied on the parametric curves to improve handling of the geometries were also explored. The range of curves that provide valid shock containing flowfields was further expanded by enabling concave shock axisymmetric flowfields on the osculating planes of the design method. The best match of such a flowfield for the osculating cones method is the ICFA flowfield, which is also described by straight shocks. For the osculating axisymmetric flow and osculating flowfield methods, inward turning flowfields with appropriate shock shapes will need to be numerically constructed. The geometries generated by this design method cover what the osculating cones method can provide with the added design space offered by the use of concave shock shapes. The design method was validated with inviscid

CFD simulations and through qualitative and quantitative comparisons between the predicted and calculated flowfield characteristics.

The detailed parameterization of the design-driving curves utilized by inverse design methods was also identified as an area were a method to further support the decisions made during the development of the parametric geometry model would be useful. Towards that goal, an approach involving parametric flexibility studies was presented. Through such studies different parameterization schemes with varying degrees of freedom can be compared. The limits of over and under-parameterizing a geometry can also be identified. The designer can then make better educated guesses when 'putting together' the parametric geometry model, a process which generally also relies on experience and intuition.

An integral part of the forebody geometries moving forward is the blunt leading edge. Apart from circular leading edges the use of more sophisticated shapes has lately seen scientific interest. A leading edge parameterization scheme with three approaches of increasing flexibility and complexity was developed and presented. The geometries seem to match or outperform similar approaches while being more efficient (utilizing fewer design parameters). Through the investigations conducted, we also highlighted the need for efficiently parameterized geometries that need to be adapted on a case by case basis since the characteristics and requirements of the thermal protection system can greatly affect the ideal leading edge shapes. We also conducted a preliminary investigation on how the receptivity and transition mechanism can be affected by increased levels of geometric continuity provided by the proposed geometries compared to circular leading edge shapes, with potentially favourable results. Additionally, we presented different options and addressed a number of issues regarding their integration on three-dimensional forebody geometries, including a semi-empirical approach to distributing the thickness across the leading edge of the waverider and two methods to control the deformation of the planform that results from varying the thickness of the blunt leading edge across the span.

Finally, we presented a framework for the development of waverider-based hypersonic aircraft parametric geometry models. This framework consists of, 1) a revised aerodynamic design process with additional design loops aimed at aiding the development of the aircraft configurations and parametric geometry models, and 2) a geometry engine to enable this revised process which also incorporates all the aforementioned waverider parametric efficiency considerations and design approaches. Initially, we described the revised aerodynamic design process that incorporates parametric flexibility studies and low-fidelity design space explorations to assist the development of parametric geometry models. We then identified a set of requirements of the geometry engine in order to enable and assist these design loops. These included the ability to quickly develop

multiple parametric geometry models for parametric flexibility studies and the ability to seamlessly interface and develop coupled parametric models of the main geometry components, such as the forebody and inlet.

The geometry engine was focused on the inversely designed and specific to hypersonic aircraft design components that cannot be readily designed and parameterized using modern CAD software. The detailed parameterization of these hypersonic-design-specific aircraft components is the area where the additional design loops, which are proposed as part of the revised aerodynamic design process, are mostly required. Additionally, the geometry engine was not focused at a single configuration but is modular and robust, supporting the development of multiple configurations and airframe-engine integration approaches through a versatile and coupled forebody and inlet parameterization. We also highlighted a number of implementation specific considerations and features that further increase the utility of the geometry engine. We developed an interactive GUI as one of the methods to generate the input for the geometry generation, which improved and sped up the process of developing multiple parametric geometries and was also found to greatly assist in introducing the waverider design concept to designers unfamiliar with it. The main geometry parameterization tools are complimented by conventional CAD approaches, and the geometries can also be exported in a standard format to be further built upon and completed in third-party CAD software packages. Finally, an example hypersonic aircraft configuration was developed to highlight the versatility of the developed geometry engine.

6.2   Further Work

The current work has primarily focused on the development of the parametric geometry models. However, to really quantify the effectiveness of all the aforementioned considerations and design methods, they need to be used in realistic and detailed design studies. The computational resources required for such a study will be significant as performing the additional design loops to develop an effective and efficiently parameterized geometry only makes sense when the following aerodynamic optimization involves computational requirements an order of magnitude or two higher.

Regarding the novel waverider design method based on three-dimensional leading edge shapes that was proposed, there are two areas that would be worth further investigating. The first is assessing the advantages of enabling concave shock regions through the use of ICFA flowfields, which can, for example, lead to a potential expansion of the design space as far as aircraft stability goes. The second is developing a detailed methodology for applying this design approach to the osculating axisymmetric flow and osculating flowfield extensions of the osculating cones waverider

design method. The proposed blunt leading edge shape formulations are ready to be used in any design where more sophisticated leading edge shapes are required, and most of their advantages have been sufficiently showcased. The effect of increased degrees of surface continuity on receptivity and the turbulence transition mechanism is, of course, an area where much more detailed investigations are required, especially since turbulent transition can have significant effects on the increased skin friction and heating affecting hypersonic aircraft. Use of more sophisticated leading edge shapes on inlet cowl lips and side walls is also a subject that would be interesting to further investigate.

Ultimately, the geometry engine and proposed parametric geometry model development framework is, at this point, a collection of tools to develop parametric geometries for use in aerodynamic optimization studies. These tools have been developed to tackle a number of hypersonic aircraft design challenges, utilizing novel features and design methodologies. Additionally, the material in this thesis is aimed at assisting the transition from the various theoretical design methods and scientific studies to geometries that can be handled effectively and efficiently in realistic design processes. Finally, the geometry engine was also developed to bring all the material presented in this work together, while some of the investigations were triggered through the process of developing the parametric geometry engine itself.

It is the hope of the author that, the detailed documentation and novel design approaches for handling waverider geometries, the leading edge shape formulations and the proposed design methodology for developing effective and efficient parametric geometry models, will prove to be valuable or even useful information to the reader, whether it is for further research on the subject or for designing a hypersonic aircraft.

Appendices

Appendix A **Osculating Cones Waverider Design**

The osculating cones waverider design method is essentially a superset of the wedge and conical shockwave inverse design methods. In both those methods a leading edge that lies on the shockwave generated by a cone or wedge is first defined and then the lower surface of the waverider is designed by tracing the streamlines of the cone or wedge flowfield, downstream from that leading edge. This is based on the fact that for inviscid flows a streamsurface can be replaced with a solid boundary of the same shape and will generate the same flowfield, hence the inverse design aspect of the method. The upper surface on the other hand can be arbitrarily designed although it does affect all aspects of the vehicle's performance. It is usually manipulated to improve aerodynamic characteristics, to enable the packaging of internal subsystems and to interface additional geometry components such as control surfaces.

The osculating cones weaverider design method allows much more flexibility in the design as the shockwave is not generated by a single rigid body. Instead, the flowfield is divided in osculating planes in the general span wise direction (Figure A.1), and with the assumption that the cross-flow between the planes is minimal and can be ignored, the flowfield of each plane can be based on slightly different bodies. In this case a conical flowfield solution is scaled on each plane in order to match the curvature of the desired shock profile. The shock and cone angle for all the planes is the same as shocks of equal strength are required to prevent strong pressure gradients and cross-flow between planes. If a shock profile with a constant radius of curvature is used along all of the osculating planes, a simple conical shock waverider is generated. If one with no curvature is used then the flowfield is locally two-dimensional with the result being identical to a planar wedge shock derived one. By being able to blend those two however, many more design options open up. The general configuration that has been of interest is using a section with no curvature close to the center of the body in order to get a uniform two-dimensional flow, whilst using a curved shock shape for the planes further out on the sides of the waverider. This way the centre area's flowfield is more suitable for an engine intake while the overall design retains some of the good characteristics of cone based designs such as better volumetric efficiency. This flexibility of the design method in controlling the shockwave and the leading edge shapes offers an expanded design space for waverider design. More details on the design algorithm can be found in the following.

A.1   Method Description

In this waverider design method, the geometry of the waverider is controlled by a number of design parameters and two design-driving curves. Starting from the intended operating flight altitude and velocity of the vehicle, the flight Mach number and specific heat ratio γ of the medium

Appendix A

(air) can be obtained. The next feature of the design that is chosen is the shock angle θ_s. It can influence the design substantially as it directly affects the strength of the shock and can therefore have a great effect on all aspects of the vehicle's performance. Finally, the shockwave shape and the leading edge shape are controlled by two curves at the base plane (aft section) of the waverider geometry, the shockwave profile curve (SWPC) or inlet capture curve (ICC) and the upper surface profile curve (USPC) or flow capture curve (FCC), as seen in Figure A.1. The shockwave profile curve alone, given a fixed shock angle, fully defines the three-dimensional shape of the shockwave on which the design will be based. The upper surface profile curve, through controlling the shape of the upper surface or the upper trailing edge of the waverider, also controls the shape of the leading edge from a freestream perspective since the top surface is parallel to the freestream. It is the curve that defines the origin for the streamtracing process. The length of the waverider is directly linked to the shock angle and maximum distance between the two curves, which is usually at the osculating planes around the symmetry plane of the waverider.
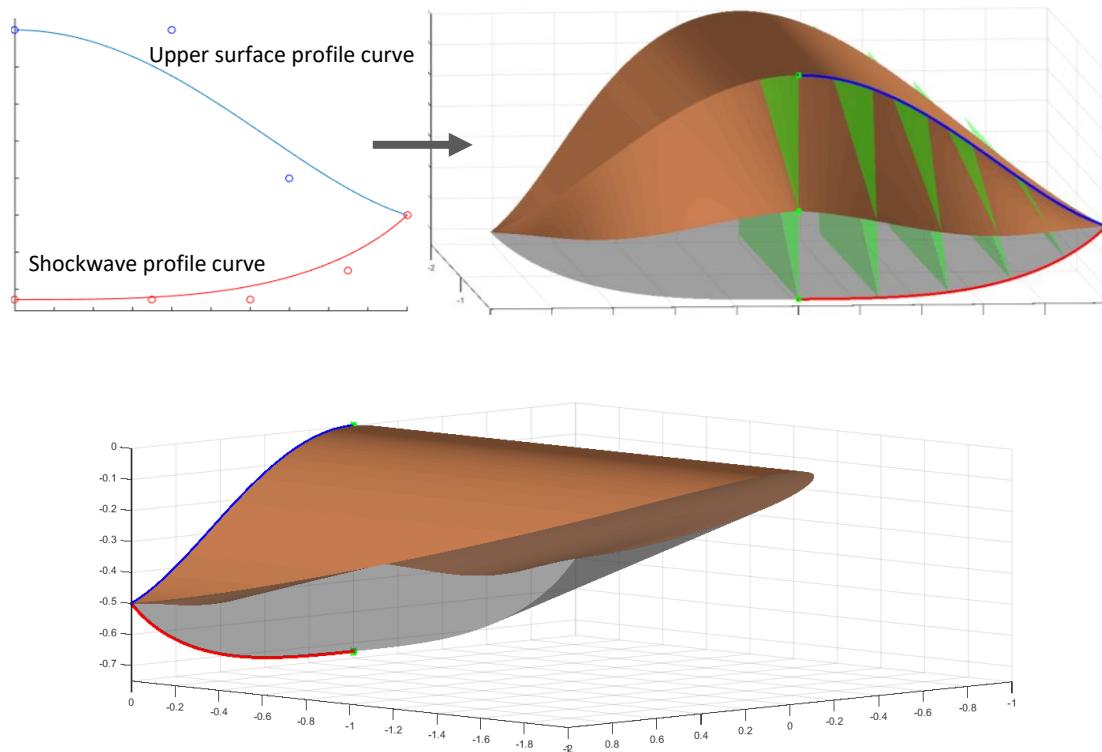
Figure A.1.   Design-driving curves on base plane (top-left) and perspective views of osculating cones waverider, with the shock surface (grey), five of the osculating planes (green) and shockwave and upper surface profile curves.

Using the shock angle, freestream Mach number and specific heat ratio, the generic cone flowfield can be calculated by numerically solving the Taylor-Maccoll equations (A.1):

$$\frac{\gamma - 1}{2}\left[V_{\max}^2 - V_r^2 - \left(\frac{dV_r}{d\theta}\right)^2\right]\left[2V_r + \frac{dV_r}{d\theta}\cot\theta + \frac{d^2V_r}{d\theta^2}\right] - \frac{dV_r}{d\theta}\left[V_r\frac{dV_r}{d\theta} + \frac{dV_r}{d\theta}\left(\frac{d^2V_r}{d\theta^2}\right)\right] = 0\,,$$

$$V_\theta = V_r' = \frac{dV_r}{d\theta}. \tag{A.1}$$

These equations are expressed in a polar coordinate system with θ the angle of the ray from the cone axis and V_r and V_θ the radial and angular velocity components as seen in Figure A.2. The solution process starts by using the oblique shock relations (A.2) to calculate the initial deflection of the flow, δ, and Mach number, M_2, immediately after the shock using the freestream Mach number and shock angle. The velocity vector is then transformed to the polar coordinate system and the Taylor-Maccoll differential equations are integrated by advancing θ from $θ_s$ at the shock wave until V_θ=0 which is the normal velocity condition at the surface of the cone. The cone half angle $θ_c$ where the integration stops is therefore also obtained. In case a cone half angle is to be set in the beginning instead of the shock angle, an iterative process is required where a shock angle is initially guessed and depending on the resulting cone half angle it is adjusted until the desired and calculated values of $θ_c$ match. The velocity field is only a function of θ as the velocity does not depend on the radius and is the same along each ray, it is a self-similar solution.

$$\tan\delta = 2\cot\theta_s \frac{M_{inf}^2 \sin^2\theta_s - 1}{M_{inf}^2(\gamma + \cos 2\theta_s) + 2}\,,$$

$$M_2 = \frac{1}{\sin(\theta_s - \delta)}\sqrt{\frac{1 + \frac{\gamma - 1}{2}M_{inf}^2\sin^2\theta_s}{\gamma M_{inf}^2\sin^2\theta_s - \frac{\gamma - 1}{2}}}. \tag{A.2}$$
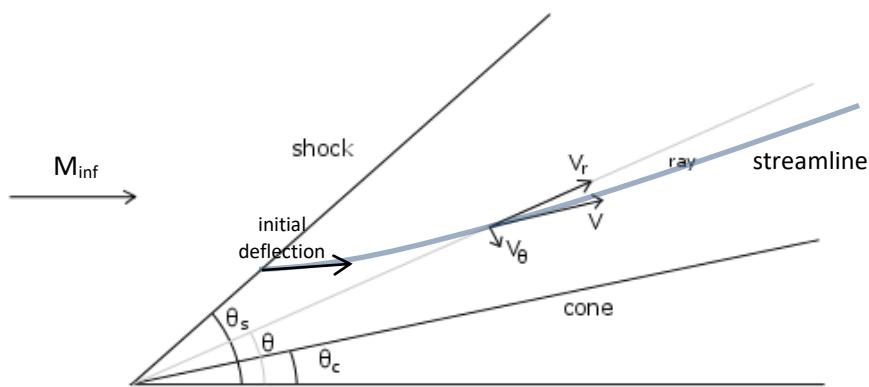

Figure A.2.   Axisymmetric flowfield of a cone in supersonic flow.

After the two profile curves have been defined, the inclination α and radius of curvature R_c along the shockwave profile curve is calculated. The inclination is used to define the angle of the

osculating planes that need to be perpendicular to the curve, while the local radius of curvature defines the scaling of the osculating cone and its flowfield on each plane, illustrated in Figure A.3.
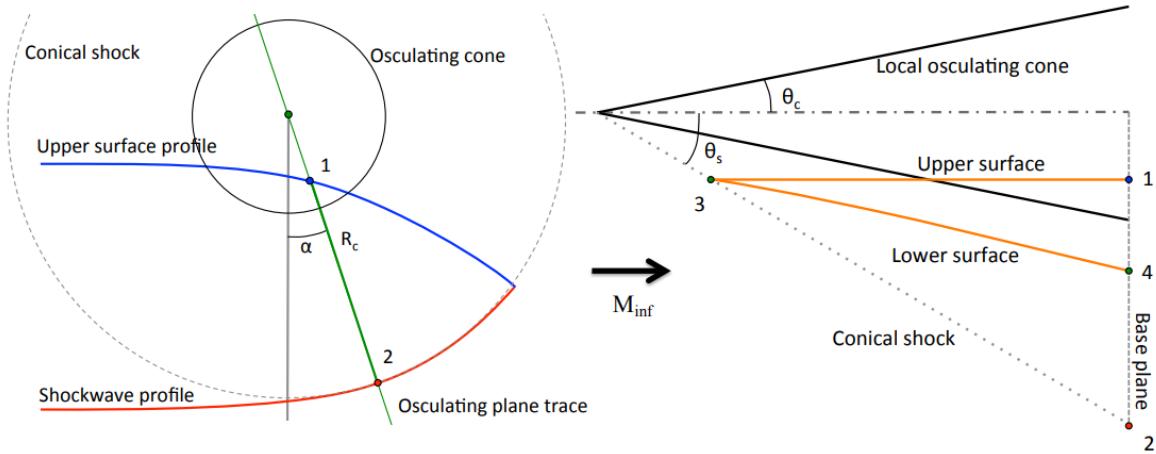

Figure A.3.   Single osculating plane features on the base plane (left) and waverider design on each plane (right). The curve between 3 and 4 describes a streamline of the flowfield.

For each osculating plane of the design the traces[1] of the upper surface profile curve and shockwave profile curve are first located, points 1 and 2 in Figure A.3. A limitation that applies here is that if the local radius of curvature is smaller than the distance between those two points, the local cone axis will be under point 1, in which case stream tracing becomes impossible as the leading edge is beyond the tip of the local flowfield generating cone. The process continues by calculating the leading edge points (point 3 on Figure A.3) for each osculating plane, which are directly upstream of the upper surface profile curve trace point and lie on the shockwave. The upper surface of the waverider can then be fully defined, as it is chosen to be parallel to the flow and is outlined by the leading edge and the upper surface profile curve.

To construct the lower surface of the waverider the streamlines on each osculating plane are traced from the leading edge points downstream until the base plane at the back is reached (point 4 on Figure A.3). Knowing the local radius of curvature and the shock angle, the tip of the local osculating cone can be positioned on the plane. It is then possible to determine the position of each point as we march downstream in the local cone's polar coordinate system, mainly the angle θ that is necessary for the calculations. By numerically solving the Taylor-Maccoll equations the radial and angular components of the velocity can be calculated for that angle θ. The velocity vector is then transformed back to the Cartesian coordinate system and from a given point we can compute the location of the next one on the same streamline by marching in the direction of the vector until the

---

[1] The word 'trace' is used for the point/path of intersection between a curve/surface and a plane. This usage is not to be confused with the terminology of the flowfield stream-tracing process.

base plane is reached, where the process stops. Once this is done for all osculating planes the lower surface is fully defined and the waverider geometry is finalized.

A.2   Design-Driving Curve Limitations

The first limitation that applies to the upper surface profile curve is that the upper surface point on each osculating plane must not be further from the shock than the local radius of curvature. Breaching this limitation would also mean that the adjacent osculating planes would intersect between the shockwave and upper surface profile curves. To remedy this, we would need either a different upper surface profile curve design or a larger radius of curvature for the relevant part of the shockwave profile curve. In section 3.1, the equivalent limitations that apply for the different parameterization approaches that utilize other streamtracing origin defining curves are also described. There is one more limitation should be applied to the upper surface profile curve: it should have no more than one point on each osculating plane. Otherwise the geometry will either have self-intersecting parts or discontinuities. The streamtracing origin defining curve can otherwise be freely designed. Its level of continuity, smoothness and any sharp corners directly translate to the continuity and edges of the waverider geometry in the spanwise direction.

The design is, however, more sensitive to the shape of the shockwave profile curve. This is due to the fact that the cone flowfield on each osculating plane is scaled according to the radius of curvature of the shockwave shape. The radius of curvature is linked to the first and second derivatives of the parametric curve with the following equation:

$$R_c(t) = \left| \frac{(x'(t)^2 + y'(t)^2)^{3/2}}{x'(t)y''(t) - x''(t)y'(t)} \right| \tag{A.3}$$

The higher observed sensitivity of the design stems from the fact that changes to the second order derivative are difficult for the human perception to discern, and in this case such changes end up having a significant effect on the final design. The radius of curvature affects the scaling of the local flowfield and position of the local osculating cone, also affecting the curvature of the streamline that is being traced. Given a smooth leading edge, in order to obtain smooth surfaces in the spanwise direction third order continuity of the shockwave profile curve is required, making this one of the very few occurrences in CAD where such high level of geometric continuity is required of a parametric curve. This is not an issue when a single infinitely differentiable function is used for the entire curve, but something that needs to be considered when piecewise functions are used. For example, a fourth order spline would be required to satisfy this condition.

A common approach to parameterizing the shockwave profile curve is using a straight line segment describing the central part of the flowfield and a parametric curve for the rest of it. This

results in a uniform flowfield favouring the integration of an engine inlet on the central part of the underbelly of the aircraft, while the benefits of more tightly curved axisymmetric flowfields can be utilized for the rest of the design. In the following examples we use a straight line segment and a Bézier curve to design such a shockwave shape. In order to achieve third order continuity at their interface, the first three intermediate points of the parametric curve need to be on a straight line, each one adding an additional degree to the level of continuity at the interface with the straight line segment. Having only the first control point in line, which controls the tangency and achieves first order continuity would result in a discontinuity of the lower surface. Having two points in line and second order continuity would result in a sharp edge appearing on the lower surface. Both examples can be seen in Figure A.4.
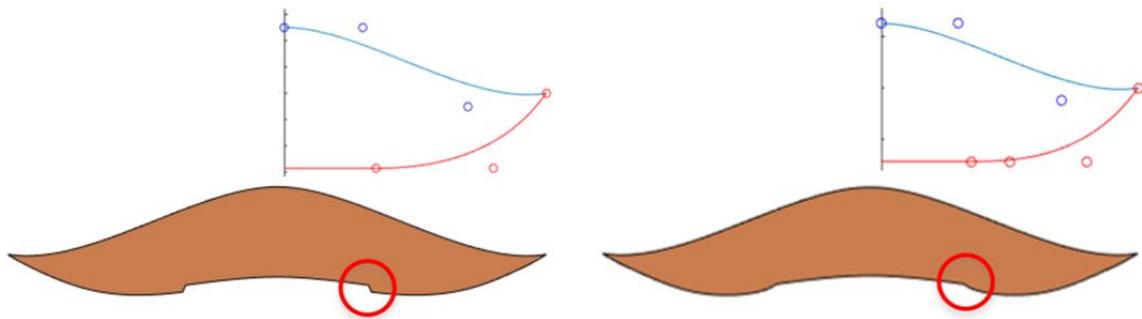
Figure A.4. First order (left) and second order (right) continuity of the shockwave profile curve and resulting discontinuity and sharp edge in the respective waverider geometries. Design-driving curves and their control points visible.

The presence of a sharp edge or not is not a major concern from a conceptual design aerodynamic performance perspective, it is however a concern from a CAD perspective. Knowing when to expect a sharp edge on the geometry and adjusting the topology of the surface patches to accommodate it is a good geometry handling practice, especially when the geometries are to be further built upon with conventional CAD techniques.

Appendix B **Geometry Accuracy**

In this section a number of investigations and considerations into the accuracy of the output geometries of the parametric geometry engine are presented. Figure A.5 shows a typical process through which a geometry goes prior to a CFD simulation and at this stage we are interested in errors arising from the first stage of moving from the analytically described surfaces to discretized coordinate matrixes.
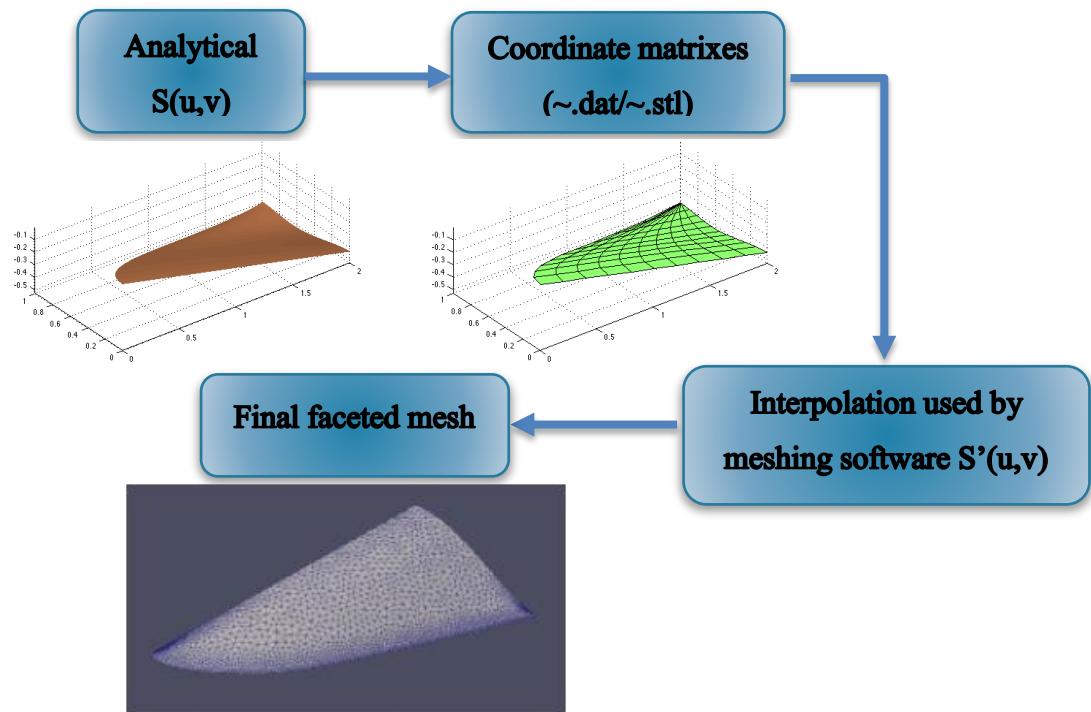
Figure A.5.   Process through which a geometry goes from CAD to CFD.

Loss of information and errors at this stage originate from two sources, the accuracy of the cloud of points that describe the coordinate matrixes, and the loss of information due to the discretization of the surfaces to faceted surface patches. The accuracy of the points comes down to the accuracy of the streamtracing algorithm that is used, since each line of points in the streamwise direction represents a streamline on an osculating plane of the design method. This in turn depends on the accuracy of the Taylor-Maccoll differential equation solution which is provided by MATLAB's *ode45* ordinary differential equation solver and the downstream marching method that is used. Tighter tolerances of *ode4*5 yielded insignificant differences when compared to the size of the steps of the downstream marching algorithm, which is where more focus was given.

For streamtracing we initially employed a simple Euler marching method, which utilized the velocity direction given by the differential equation solution evaluation on the current location to march to the next point on the streamline with first order accuracy. Using this method, the relative

error of the end point of 10 streamlines along the span of a waverider geometry was initially calculated, with point 1 being on the symmetry plane. The relative error was calculated as the deviation of the last point of the streamline on the base plane, from a geometry that was generated with 10,000 steps for the same length. Geometries generated with 10, 100 and 1000 steps are compared as seen in Figure A.6, with the expected one order of accuracy increments with each ten-fold increase of the resolution. To improve this, a predictor-corrector second order marching method was employed and the same results using that can be seen in Figure A.7. This second order method was a considerable improvement both as far as the accuracy and the number of steps required to achieve acceptable tolerances go.
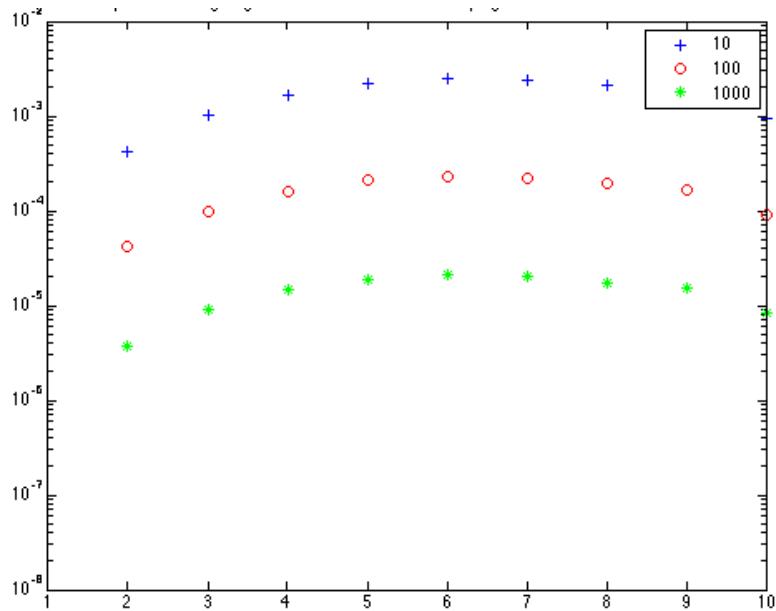

Figure A.6.   Relative deviation of 10 lower surface points along the span of a waverider using a simple Euler (first order accurate) method for stream tracing.
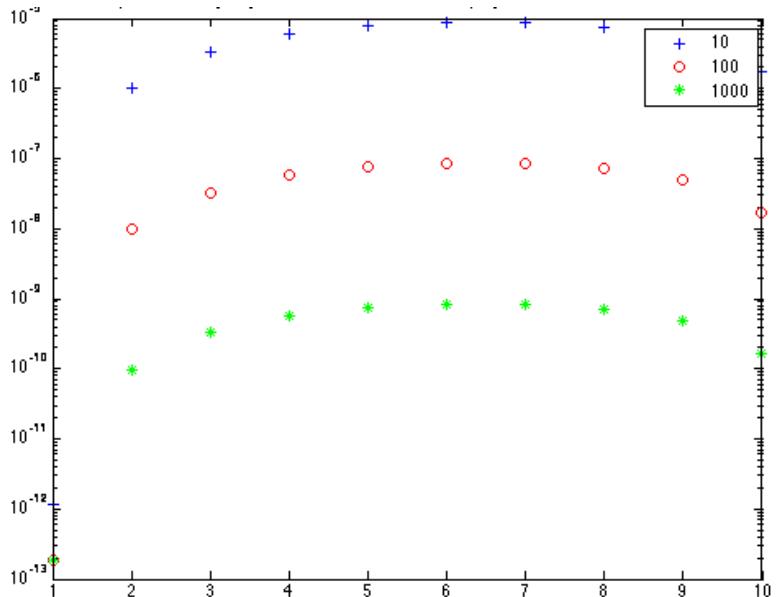

Figure A.7.   Relative deviation of 10 lower surface points along the span of a waverider using the predictor-corrector method for stream tracing.

The second source of error rises from discretizing the surfaces and describing them in a faceted format. In solid modelling, the accuracy of a faceted surface representation is usually defined using two metrics: 1) a maximum distance deviation between the flat faces and the corresponding original geometry, and 2) an angle deviation between edge segments that are used to represent the original surface's curves. To measure the distance deviation that results from discretizing the waverider surface, we can calculate the distance of a surface patch or face, to a number of directly calculated points $S(u_i, v_i)$, or essentially a higher resolution geometry. In Figure A.8, a single patch that is defined by 4 points in the 100x100 resolution, is compared to 100 directly calculated points that are represented by that patch. A contour of the deviations calculated for the entire waverider lower surface, with 100 equivalent measurements per patch, is seen in Figure A.9 and Figure A.10.
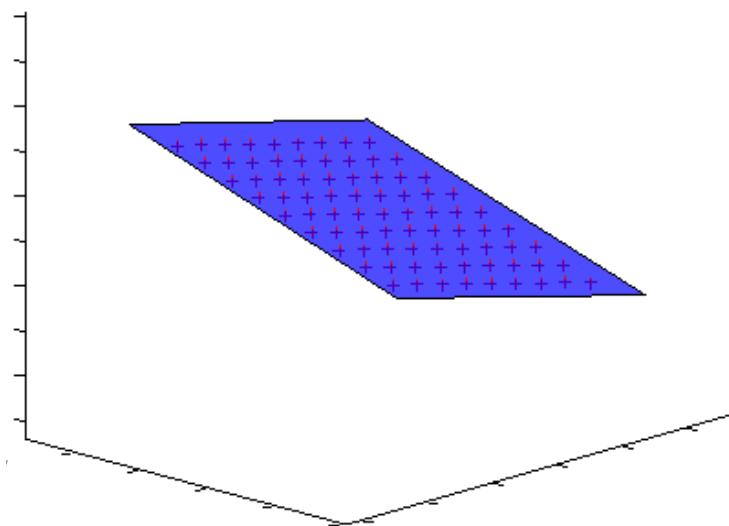
Figure A.8.   A single patch of a 100x100 resolution geometry and the 100 corresponding points whose distance is measured.

The resulting maximum deviation that was calculated was 6.4×10^{-5} for this waverider geometry with a half-span of one unit and a two unit length, the mean deviation was 4.7×10^{-5}. The maximum deviation is, as expected, located in parts of the geometry that are highly curved.
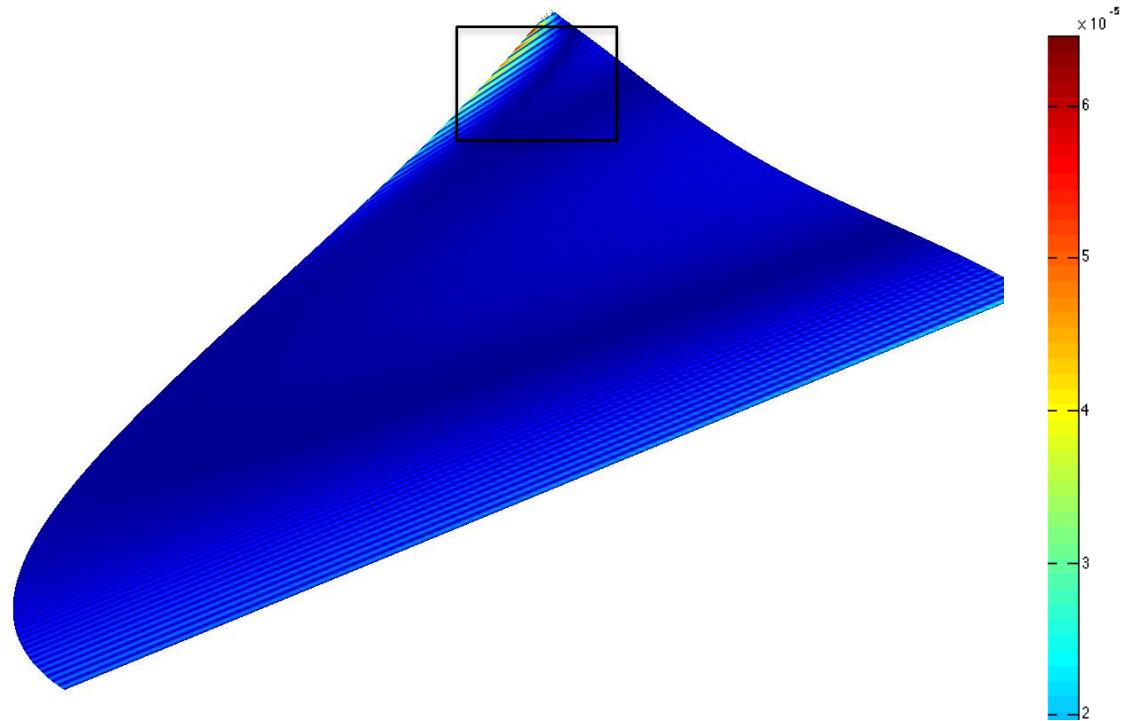
Figure A.9. Deviation of faceted 100x100 lower half-symmetric surface of a waverider from directly calculated points, 100 measurements per patch.
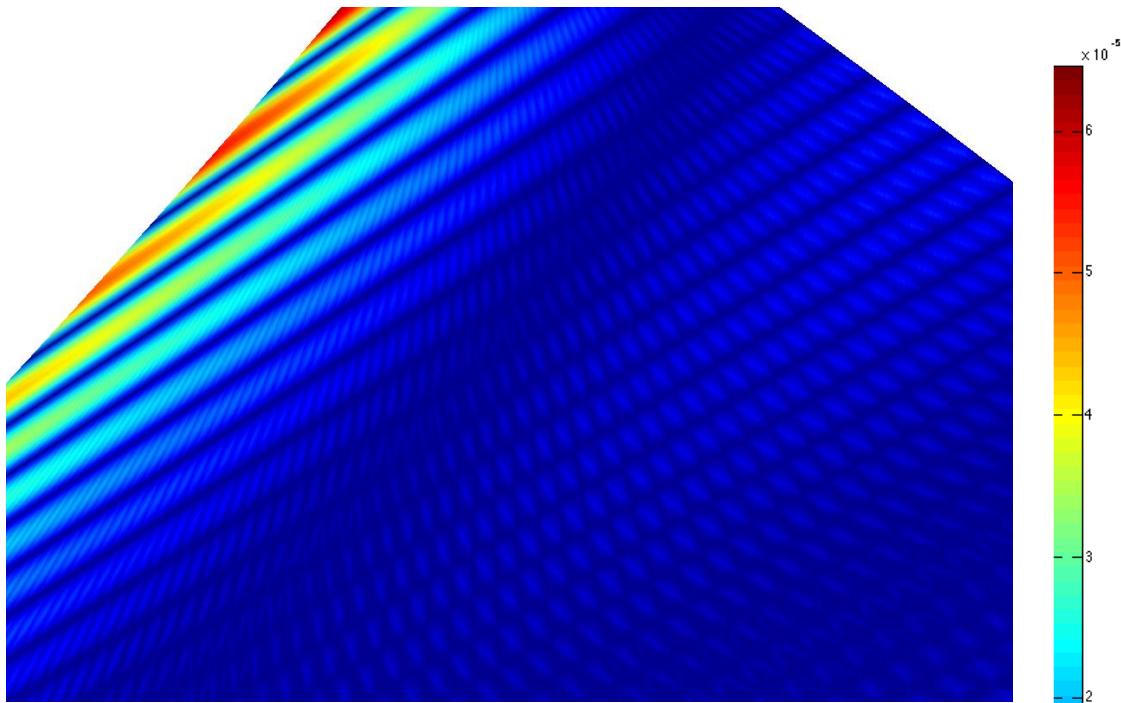
Figure A.10. Deviation of faceted 100x100 lower surface of a waverider from directly calculated points near the wing-tip, 100 measurements per patch.

Measuring the second type of deviation, the angle deviation, is most important for the geometry's feature edges. In waverider forebody geometries it is mainly the leading edge curve that is of interest, as the upper surface and lower surface profiles are both related to that and also

they are not the actual end of the geometry, but just an interface for the aft part of the aircraft. There are two approaches to quantifying how well discrete line segments represent curved edges. The first one is directly calculating the angle deviation and distance between the original curve and its discretized representation. A second, more meaningful and comprehensive approach that is usually employed in solid and surface modelling software, is by measuring the angle between each successive line segment of the discretized curve. This angle, along with a known accuracy of the nodes at the end points of each line segment, contain all the information needed to quantify how smoothly and accurately a curved edge is represented when discretised. The angle between the line segments representing the leading edge of the geometry of Figure A.9 is calculated and shown in Figure A.11, something that is simple and straightforward to read.
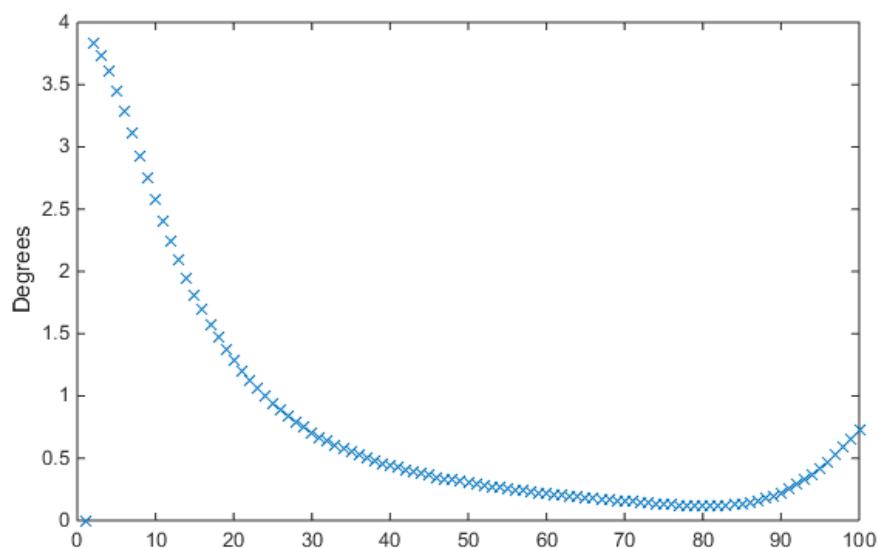
Figure A.11. Angle between 100 successive line segments representing the leading edge curve of a waverider forebody geometry.

At this stage it was determined that while a certain level of confidence regarding the surface resolution that should be used can be obtained through such investigations, given the parametric aspect of the geometry engine a more robust approach was needed. This is explained in section 5.3.5.

Appendix C **Matching Geometries Between the Design Approaches**

This section contains data complementing the content in section 3.1. We have investigated the ability of each of the three different design approaches to replicate waverider geometries originally designed with another approach. This, first of all, serves as a cross-validation of the three different approaches (USPD, PLED, and LSPD). In a number of cases, it also highlights the benefits of using a specific design approach when there are design constraints that can be directly applied on it, by displaying the increase in complexity required in order to indirectly satisfy them. The investigations also provide an indication into how geometry perturbations, errors and tolerances are transferred between the different profile curves, while it also gives us some more insight into the required flexibility of the design driving curves in order to enable certain shapes that are, for example, readily available using a different design approach.

To replicate a geometry generated by one of the design methods with another one, the design-driving curve has to match the respective shape of the geometry. For example, a waverider forebody geometry given by the PLED method has a certain upper surface profile on the base plane, which also depends on the shape of the shockwave. If, for the same shockwave shape, that upper surface profile is replicated by the parameterized curve of the USPD method, the same geometry can be obtained. A least squares fitting approach was followed to match the respective design-driving curves with the target shapes on each case. Additionally, parameterizations with varying flexibility were tested, in the form of additional intermediate control points of the parametric Bézier curves. All the geometries are designed with a non-dimensional half-span of one.

The LSPD generated geometry that was also presented in section 3.1.4, was generated by defining a flat lower surface profile at the base plane of the waverider forebody. This is a case that highlights the advantages of opting for the most suitable design method when there are design constraints that can be directly applied to one of the design-driving curves, obtaining a specific lower surface shape in this example.

153
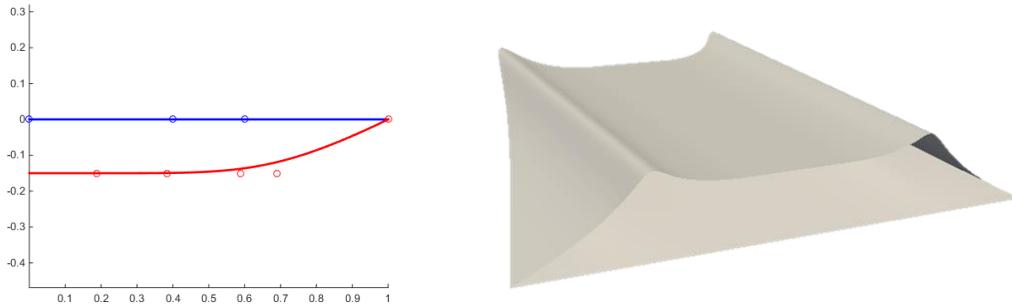
Figure A.12. LSPD designed geometry with flat lower surface.

Initially, we use the USPD method and attempt to match the resulting geometry's upper surface profile with parametric curves of varying flexibility by differentiating the number of intermediate control points. The end points are fixed in all cases.
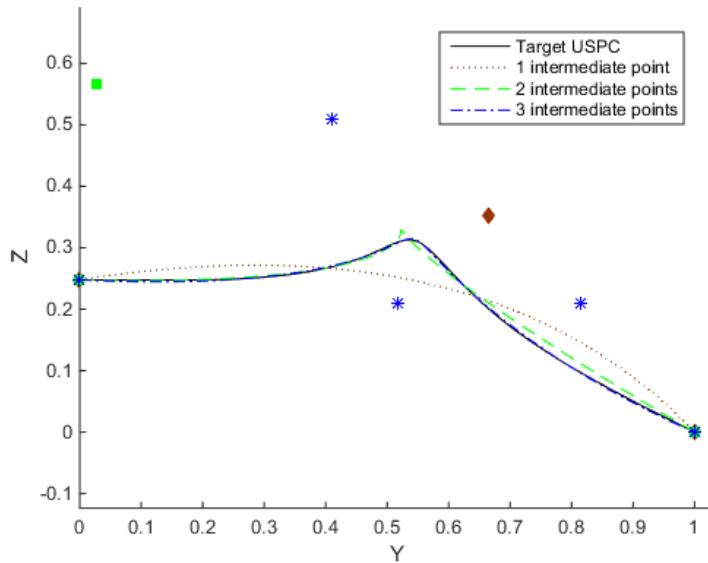
Figure A.13. Matching the upper surface profile to replicate the geometry with the USPD method.
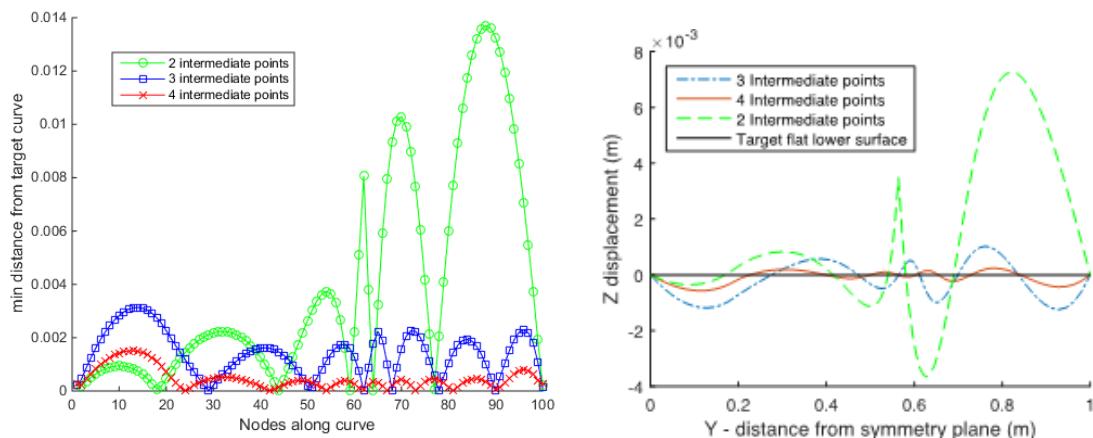
Figure A.14. Deviations of the USPD design driving curve (left) and deviations of the resulting lower surface profile (right).
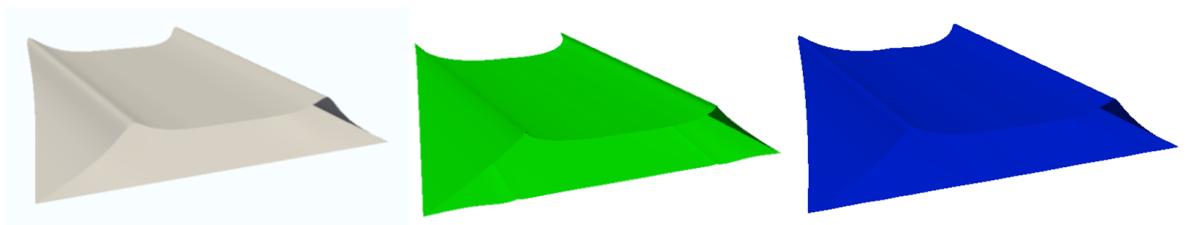
Figure A.15. Original geometry (left), 2 intermediate point approximation (middle), and 3 intermediate point approximation (right).

Using the PLED method we attempt to match the leading edge planform shape with the design-driving curve, again using increasing levels of flexibility by varying the number of intermediate control points.
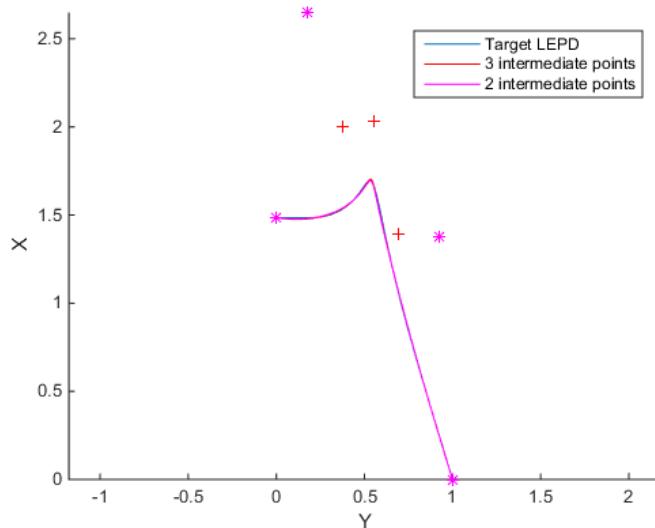
Figure A.16. Matching the leading edge profile to replicate the geometry with the PLED method.
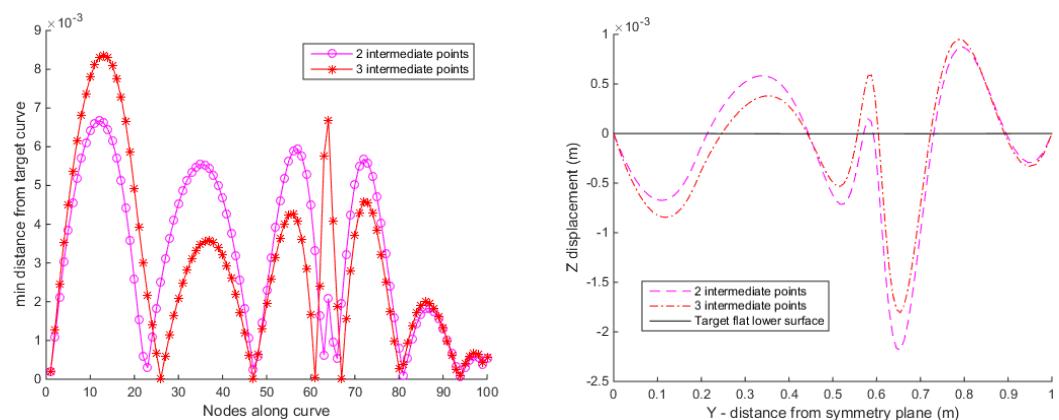
Figure A.17. Deviations of the PLED design driving curve (left) and deviations of the resulting lower surface profile (right).
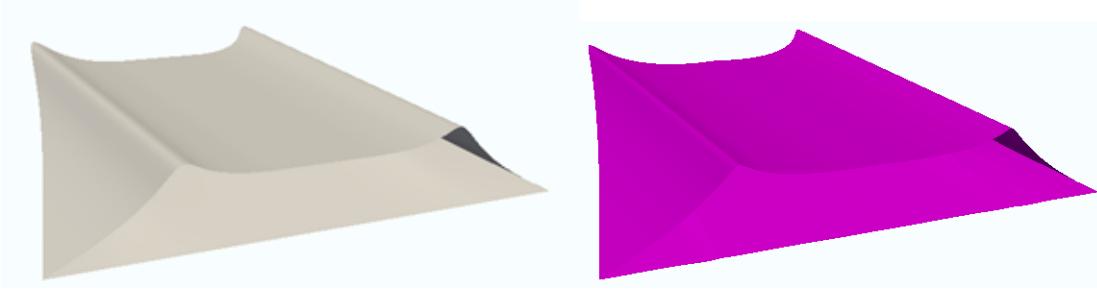
Figure A.18. Original geometry (left) and 2 intermediate point approximation (right).

Moving on to the PLED designed geometry, which was manually generated with the interactive GUI using the PLED method. The geometry was designed to roughly resemble some of the concept geometries found in the literature while the shock shape in this case is also different from the one used for the previous comparison.
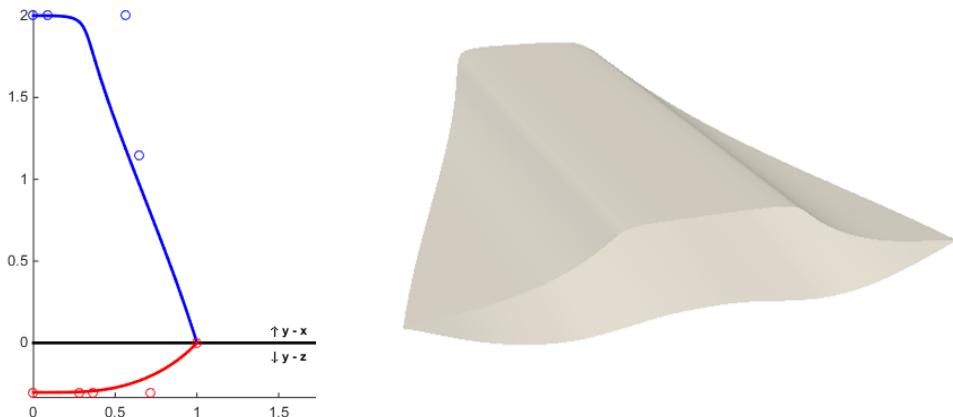
Figure A.19. PLED manually designed geometry.
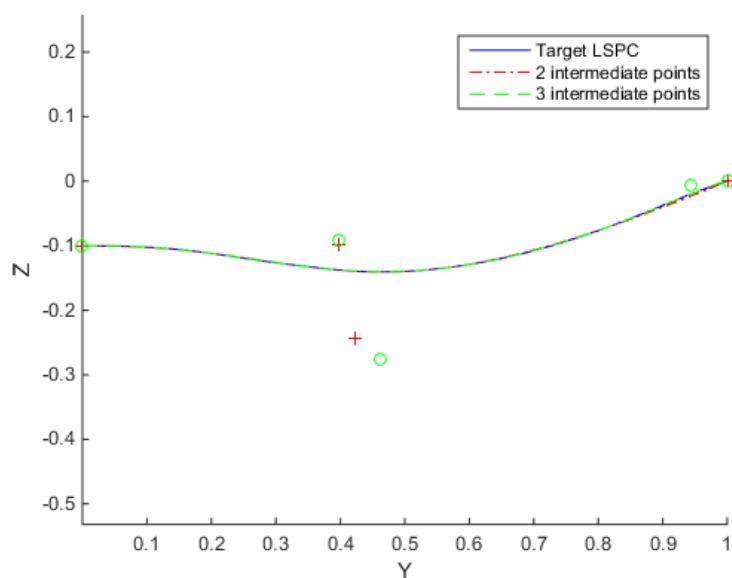
Figure A.20. Matching the lower surface profile to replicate the geometry with the LSPD method.
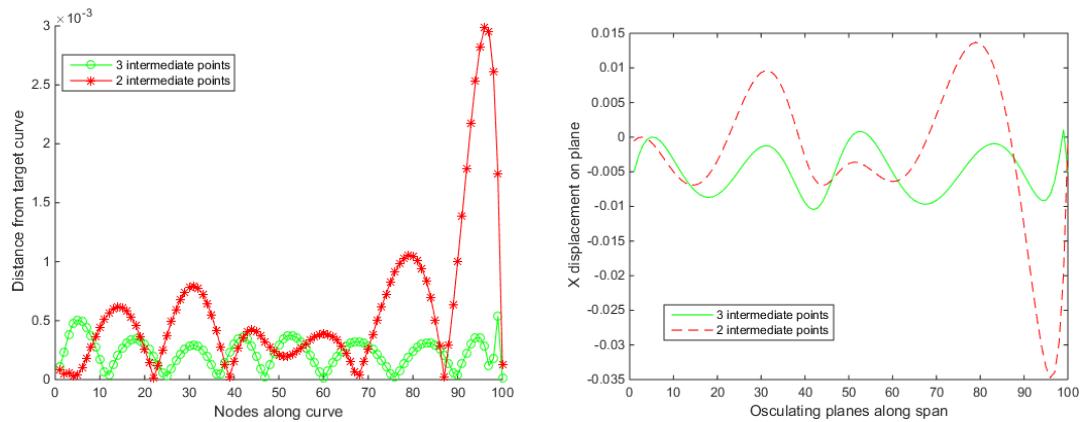
Figure A.21. Deviations of the LSPD design driving curve (left) and x-displacement of the resulting leading edge points on each osculating plane (right).
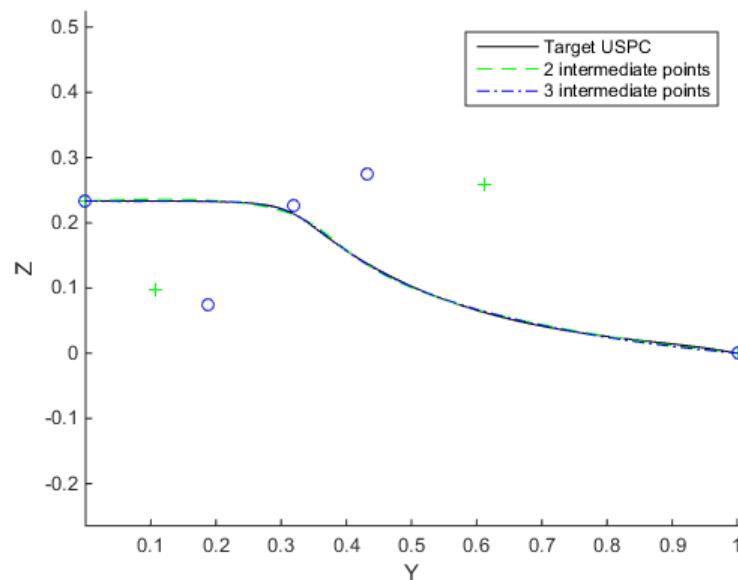

Figure A.22. Matching the upper surface profile to replicate the geometry with the USPD method.
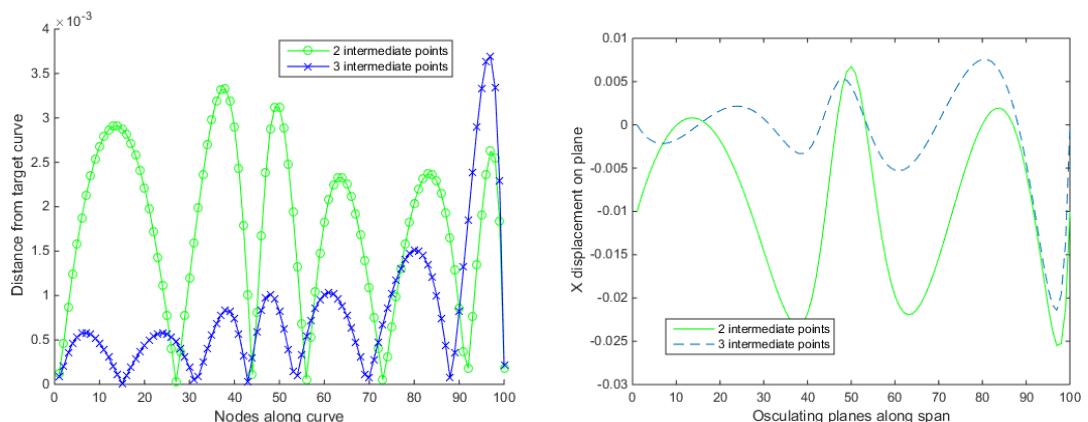
Figure A.23. Deviations of the USPD design driving curve (left) and x-displacement of the resulting leading edge points on each osculating plane (right).

Appendix C

The differences in the geometries in these cases are much smaller than in the case of the flat lower surface waverider.
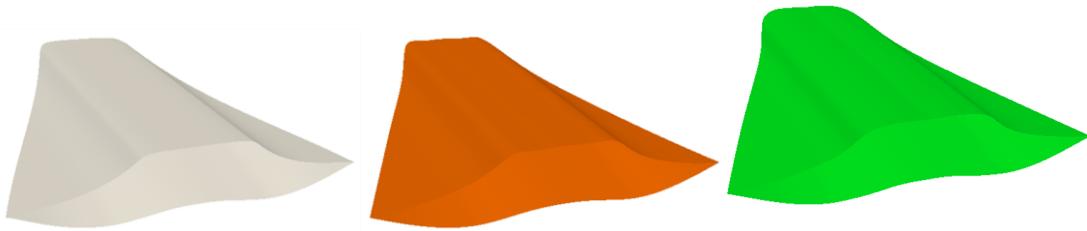
Figure A.24. Original PLED generated geometry (left), LSPD generated 2 intermediate control point approximation (middle), and USPD generated 2 intermediate control point approximation (right).

The final geometry that was investigated was one of the Pareto-optimum USPD designed geometries of section 3.3, displayed bellow. In this case, as the design was very close to the constraints of the design algorithm for certain parts of the curve (mentioned in 3.1.1), the curve fitting process would in certain cases cause the design-driving curve to violate those limitations. The reason for this is that geometries close to the limitation take full advantage of the cone flowfield, which results in higher volumetric efficiencies. The geometry could still be replicated properly as long as the design-driving curve was flexible enough to have very small deviations from the target profile.
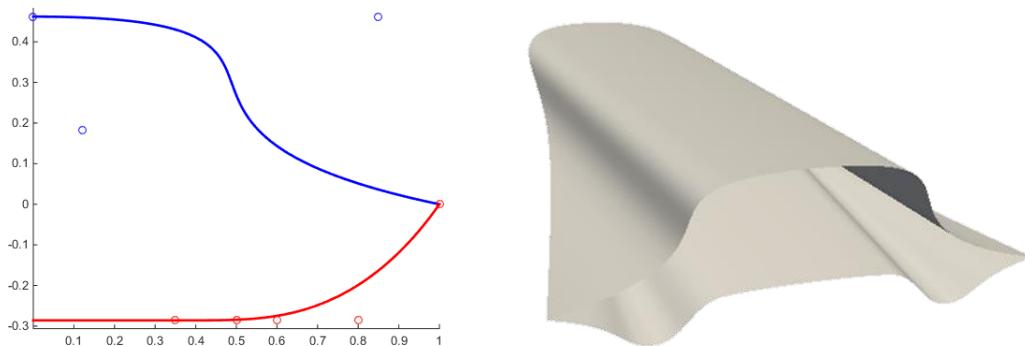
Figure A.25. One of the USPD generated aerodynamic and volumetric efficiency optimum geometries.
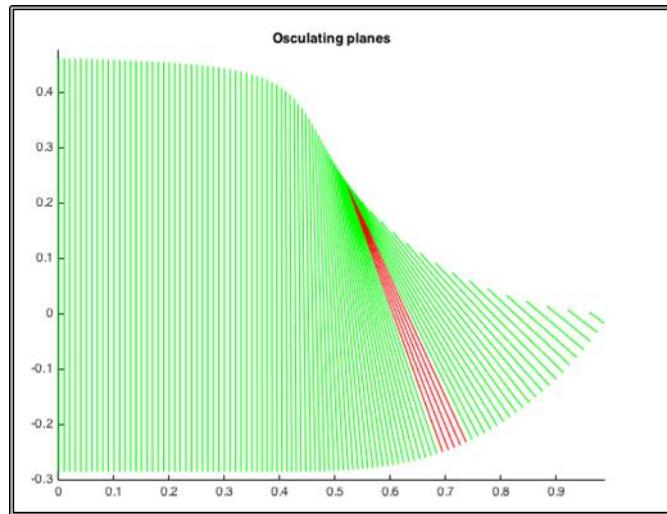
Figure A.26. Example failed design due to the osculating planes intersecting under the upper surface profile curve.
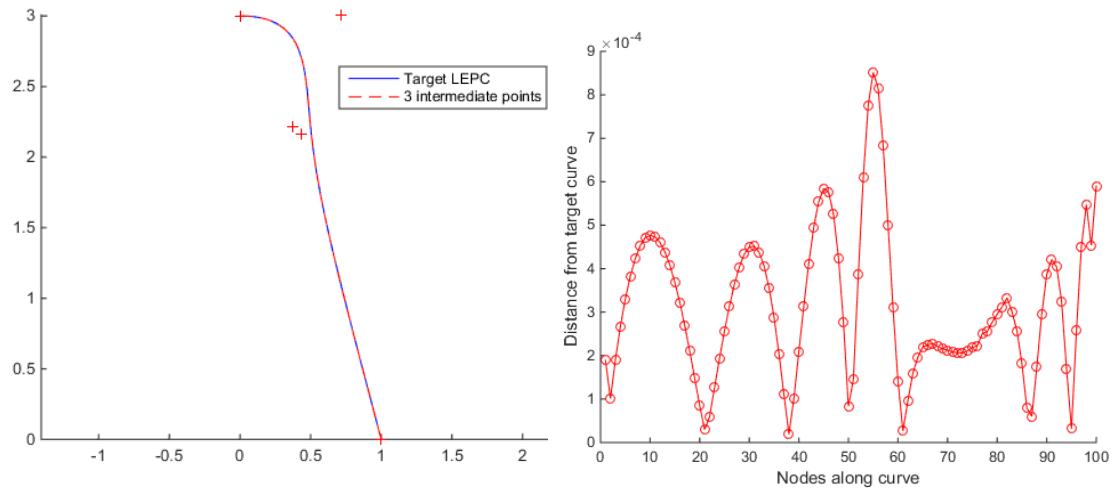
Figure A.27. Matching the leading edge profile to replicate the geometry with the PLED method (left), deviations of the PLED 3 intermediate control point design-driving curve from the target profile (right).
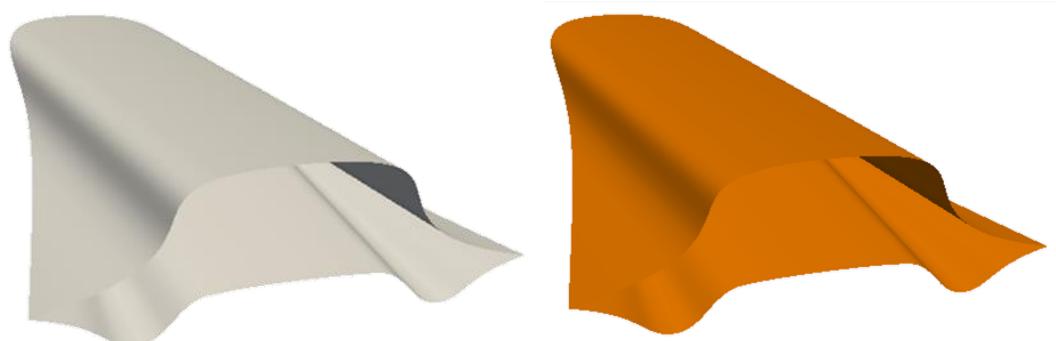
Figure A.28. Original USPD obtained waverider geometry (left) and PLED generated approximation (right).

Appendix C

The above, first of all, confirms that the three variations of the design method function as expected. Smaller deviations of the design driving curve from the target profile lead to smaller deviations in all aspects of the geometry. Additionally, we get an indication of the increased efficiency of the parametric geometry that is offered by directly applying potential design constraints or requirements to a design-driving curve when the appropriate design approach is used. Attempting to do so with an indirect method generally increases the complexity of the parametric geometry and also introduces additional optimization effort, while an exact match is not possible and one will have to deal with tolerances or resort to slight modifications of the geometry after it has been generated.

Finally, the volumetric efficiency and lift to drag ratios of all the geometries (calculated with the aerodynamic performance evaluation tool described in 5.3.2) are presented in the tables bellow. As expected, the differences in aerodynamic performance are insignificant.

Table A.1. Performance characteristics comparison for the PLED designed flat lower surface waverider.

	PLED (Original)	USPD (2 points)	USPD (3 points)	LSPD (2 points)	LSPD (3 points)
V/A	0.061456	0.061578 (+0.2%)	0.061499 (+0.1%)	0.061501 (+0.1%)	0.061473 (+0.03%)
L/D	5.0463	L/D = 5.0475 (+0.02%)	5.0479 (+0.03%)	5.0489 (+0.04%)	5.0464 (+0.002%)

Table A.2. Performance characteristics comparison for the LSPD manually designed waverider.

	LSPD (Original)	USPD (2 points)	USPD (3 points)	USPD (4 points)	PLED (2 points)	PLED (3 points)
V/A	0.051945	0.051522 (-0.8%)	0.051855 (-0.17%)	0.051919 (-0.05%)	0.051849 (-0.18%)	0.051792 (-0.29%)
L/D	5.1230	5.1175 (-0.10%)	5.1225 (-0.01%)	5.1228 (-0.004%)	5.1228 (-0.004%)	5.1221 (-0.012%)

Table A.3.   Performance characteristics for the USPD optimum-designed waverider.

USPD (Original)	PLED (3 points)
V/A = 0.078848	V/A = 0.078831 (-0.02%)
L/D = 5.5563	L/D = 5.5560 (-0.005%)

Appendix D **Viscous Skin Friction Calculation Methods**

Pressure forces for waverider geometries can be accurately obtained through the design algorithm. For more complex shapes other than inversely designed waveriders, local inclination and simple Newtonian methods can also be used with sufficient accuracy. The calculated pressure forces and velocities on the surface of the geometry, can then be utilized by semi-empirical skin friction calculation methods in order to calculate the viscous aerodynamic performance.

D.1   Reference Temperature Method

Reference temperature methods are simple methods to calculate skin friction forces for hypersonic flow over flat plates. For more sophisticated shapes they can be applied by splitting the wetted surface in small panels and assuming each one behaves like a flat plate. White's reference temperature method (White, 1974) in particular has shown good correlation with experiments for turbulent hypersonic flow. It is therefore suitable for use with relatively large geometries where the majority of the flow is expected to be turbulent.

The friction coefficient can be calculated as:

$$C_f = \frac{0.455}{S^2 \, ln^2 \left(0.06 \frac{Re_e}{S} \frac{\mu_e}{\mu_w} \sqrt{\frac{T_e}{T_w}} \right)} \tag{A.4}$$

With, *e* the subscript indicating values at the edge of the boundary layer, *w* the subscript indicating values at the wall, and Re_e being the local Reynolds number calculated according to the distance travelled along the current streamline.

Sutherland's law can be used to calculate the viscosity at varying temperatures:

$$\mu = \mu_0 \left(\frac{T}{T_0} \right)^{1.5} (T_0 + 110)/(T + 110) \tag{A.5}$$

With μ_0 = 1.7894x10^{-5} kg/(m s) and T_0 = 288.16K.

S is calculated as follows:

$$S = \frac{\sqrt{\frac{T_{aw}}{T_e} - 1}}{\sin^{-1}(\frac{A}{C}) + \sin^{-1}(\frac{B}{C})}$$

$$A = \frac{T_{aw} + T_w}{T_e} - 2$$

Appendix D

$$B = \frac{T_{aw} - T_w}{T_e}$$

$$C = \left[\left(\frac{T_{aw} + T_w}{T_e}\right)^2 - 4\frac{T_w}{T_e}\right]^{0.5} \tag{A.6}$$

The subscript *aw* indicates for adiabatic wall, and the adiabatic wall temperature can be calculated using the recovery factor which is related to the Prandtl number:

$$r = \sqrt[3]{Pr} \tag{A.7}$$

And equation:

$$r = \frac{T_{aw} - T_w}{T_{te} - T_e} \tag{A.8}$$

With T_{te} the total temperature at the edge of the boundary layer and *Pr* can be assumed to be equal to 0.715.

D.2   Boundary Layer Integral Method

Boundary layer integral methods are generally more computationally demanding than reference temperature methods, however, they allow more accurate calculations of skin friction forces. The method proposed by Drela and Giles (1987), offers compatibility between the laminar and turbulent formulations and is suitable for use in applications where transition is important. It is a two equation boundary layer integral method based on dissipation closure, with a transition prediction formulation. Such two-equation methods can also describe thin separated regions, although for hypersonic aircraft this might only be of interest for flight in off-design low Mach numbers.

Their formulation utilizes the following integral momentum and kinetic energy shape parameter equations:

$$\frac{d\theta}{d\xi} + (2 + H - M_e^2)\frac{\theta}{u_e}\frac{du_e}{d\xi} = \frac{C_f}{2} \tag{A.9}$$

$$\theta\frac{dH^*}{d\xi} - [2H^{**} + H^*(1 - H)]\frac{\theta}{u_e}\frac{du_e}{d\xi} = 2C_D - H^*\frac{C_f}{2}$$

Where: C_D the dissipation coefficient, C_f the skin friction coefficient, H the shape parameter, H* the kinetic energy shape parameter, H** the density shape parameter, θ the momentum thickness,

M_e the Mach number at the edge of the boundary layer, u_e the velocity at the edge of the boundary layer, and ξ the streamwise coordinate of the thin sear layer.

The following dependencies are needed for closure:

$$H^* = H^*(H_k, M_e, Re_\theta), \qquad H^{**} = H^{**}(H_k, M_e)$$

$$C_f = C_f(H_k, M_e, Re_\theta), \qquad C_D = C_D(H_k, M_e, Re_\theta)$$

With Re_θ the momentum thickness Reynolds number and H_k the kinematic shape parameter defined with constant density across the boundary layer, derived by Whitfield:

$$H_k = \frac{H - 0.29 M_e^2}{1 + 1,113 M_e} \tag{A.10}$$

Drela and Giles utilize a Flakner-Skan one parameter profile family for laminar closure and the skin friction and velocity profile formulation of Swaford (1983) for turbulent closure. An e^9 type transition prediction formulation is also derived and incorporated into the viscous formulation (Ingen, 1956). For the current work, only the laminar closure was utilized for the parametric flexibility studies of Chapter 3. The relationships that are used for laminar closure are the following:

$$H^* = 1.515 + 0.076 \frac{(4 - H_k)^2}{H_k}, \qquad H_k < 4$$

$$H^* = 1.515 + 0.040 \frac{(H_k - 4)^2}{H_k}, \qquad H_k > 4$$

$$Re_\theta \frac{C_f}{2} = -0.067 + 0.01977 \frac{(7.4 - H_k)^2}{H_k - 1}, \qquad H_k < 7.4$$

$$Re_\theta \frac{C_f}{2} = -0.067 + 0.022 \left(1 - \frac{1.4}{H_k - 6}\right)^2, \qquad H_k > 7.4 \tag{A.11}$$

$$Re_\theta \frac{2C_d}{H^*} = 0.207 + 0.00205(4 - H_k)^{5.5}, \qquad H_k < 4$$

$$Re_\theta \frac{2C_d}{H^*} = 0.207 - 0.003 \frac{(H_k - 4)^2}{(1 + 0.02(H_k - 4)^2)}, \qquad H_k > 4$$

$$H^{**} = \left(\frac{0.064}{H_k - 8} + 0.251\right) M_e^2$$

Appendix E Streamtraced Inlet Design and Forebody External Compression

E.1 Streamtraced Inlet Design

Inlets for supersonic applications are usually based on three types of flowfields, two-dimensional planar compression flowfields, axisymmetric outward-turning or cone/spike flowfields, and inward-turning or Busemann-type flowfields. Streamtraced inlet geometries can be generated by tracing a throat or capture area outline through those flowfields. This is based on the principle that a stream-surface will replicate the inviscid flowfield it encompasses. Streamtraced inlets are most commonly generated for inward-turning flowfields, while two-dimensional planar compression ramps or compression cones and spikes are often used as a whole.

The design of two-dimensional planar compression inlets is dictated by the oblique shock relations:

$$\tan \delta = 2 \cot \theta_s \frac{M_1^2 \sin^2 \theta_s - 1}{M_1^2 (\gamma + \cos 2\theta_s) + 2},$$

$$M_2 = \frac{1}{\sin(\theta_s - \delta)} \sqrt{\frac{1 + \frac{\gamma - 1}{2} M_1^2 \sin^2 \theta_s}{\gamma M_1^2 \sin^2 \theta_s - \frac{\gamma - 1}{2}}}, \qquad \text{(A.12)}$$

$$\frac{p_2}{p_1} = \frac{2\gamma M_1^2 \sin^2 \theta_s - (\gamma - 1)}{\gamma + 1},$$

with δ the deflection angle of the flow behind a shock angled at θ_s, and M_1 the Mach number before the shock. M_2 is the Mach number after the shock and p_1 and p_2 the pressure before and after the shock respectively.

Apart from discrete shocks and flow deflections, planar compression flowfields can also compress the flow by smooth isentropic turning, which is described by the Prandtl-Meyer function, $v(M)$. The flow turning δ is related to the change in Mach number by:

$$\delta = v(M_2) - v(M_1) \qquad \text{(A.13)}$$

The external compression part of planar compression inlets is usually designed by focusing the shocks and isentropic compression Mach lines at a focus point. A cowl lip is usually placed in relation to that focus point and the oblique shock relations can be further utilized to describe internal compression flow, with internal compression usually performed by deflecting the flow in the opposite direction, as in the mixed compression flowfield seen at the top-left of Figure 5.12.

Both outward-turning cone based and inward-turning Busemann flowfields are described by the Taylor-Maccoll equations (A.1). The conical flowfield calculation was described for the osculating cones method in Appendix A. Based on cones are also inlet spike geometries that accommodate additional isentropic compression of the flow. The surface of these flared cone geometries can be designed with the method of characteristics in order to focus the isentropic compression Mach lines at a focus point, an example can be seen at the middle of Figure 5.12.

Inward-turning Busemann-type axisymmetric flowfields, seen at the bottom-left of Figure 5.12, are generated by first defining the free-standing conical shock angle and Mach number at the downstream end of the flowfield. The Taylor-Maccoll equations, expressed in spherical coordinates with the origin being the tip of the free-standing shock, are then integrated with an upstream direction and increasing ray angles from the shock. The integration ends when the flow is parallel to the freestream direction, indicating the first Mach line of the flowfield. A resulting freestream Mach number is then calculated. In order to design an inlet for a specific freestream Mach number an iterative process is usually followed. Both Busemann-type and cone flowfields are self-similar solutions, with the velocity being constant along the rays of the spherical coordinate system.

E.2   Forebody External Compression

Planar compression inlet ramps require sidewalls for the external compression region in order to contain the high pressure flow. However, the flow close to the corners between the compression ramp and the sidewall is subject to strong viscous effects. These effects, along with the boundary layer developed on the sidewalls, adversely affect the uniformity of the flow at the cowl lip section. The leading edges of the sidewalls also require additional aerodynamic heating management. Two types of external compression geometries that do not require sidewalls and can potentially overcome those disadvantages were developed and implemented in our geometry engine. Their design methodology is described in the following section.

The first method further utilizes the osculating cones design method to generate an additional compression surface on the underside of the waverider forebody. This can be accomplished by placing the additional osculating cones designed geometry in the uniform flow region of an osculating cones designed forebody geometry. Instead of freestream values, the design algorithm for this compression geometry utilizes the flow conditions behind the forebody-generated shock, which needs to be planar in that region. The axis of the local axisymmetric osculating flowfields is tilted relative to the freestream at an angle equal to the deflection angle, δ, of the planar shock of the forebody. The lower surface shape of the waverider at that part of the geometry becomes the upper surface profile for the inlet compression geometry. An inlet capture curve (ICC) defining the shock shape of this external compression surface can then be drawn using

a parametric curve. The first end point on the symmetry plane should be placed within the region of the forebody-generated shockwave where the flow is uniform. Placing the other end point of the ICC at the edge of the uniform flowfield maximizes the capture area of the inlet. A shock angle is chosen for the ICC defined shock and the osculating cones method from then on is followed as usual to generate the external compression geometry of the inlet. The design-driving curve parameterization of both the waverider forebody and the osculating cones designed external compression surface can be seen in Figure A.29.
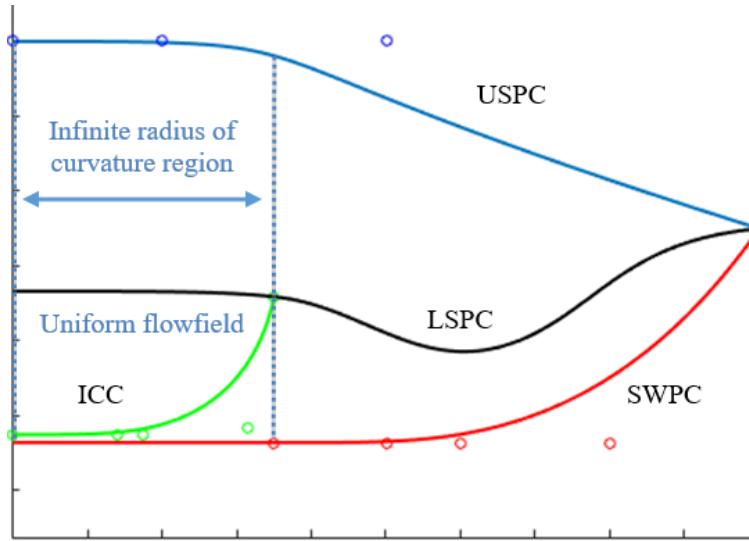
Figure A.29. Osculating cones waverider and external compression geometry design-driving curves.

The resulting waverider geometry, for a freestream Mach number of 8 and both forebody and inlet shock angles set to 15^0, can be seen in Figure A.30. A qualitative analysis of the results is provided by the CFD results shown in Figure A.31, with the forebody and external compression shocks attached to the respective leading edges. A comparison between Mach number contours of the Euler CFD and the flowfield predicted by the design method can be seen in Figure A.32. The contours are plotted on the base plane where the design-driving curves of the geometry were drawn. Overall, we observe a very good agreement of the Mach number values and shock shape, further validating the design method.

Designing the rest of the inlet from this point on becomes more complicated as the resulting flowfield is not completely uniform due to the nature of the conical shock flowfields. One solution would be to have the inlet cowl only capture the central part of the externally compressed flow, which can be designed to be sufficiently uniform by accordingly designing the ICC, and follow a planar design. This way the adverse effects of using sidewalls are traded with increased pressure drag stemming from the flow that is further compressed but not captured. Inverse design and rapid geometry generation methods seem to otherwise come to a dead-end, with potentially complex

Appendix E

CFD calculations and direct design techniques required to move a detailed inlet design forward. Utilizing the osculating flowfield method for this external compression geometry may, however, allow better control of the uniformity of the resulting flow and open up additional options. The second type of external compression geometry without sidewalls that is explained in the following section, offers a compression flowfield better compatible with planar compression inlets.

Compression geometry interface

Inlet Capture Curve
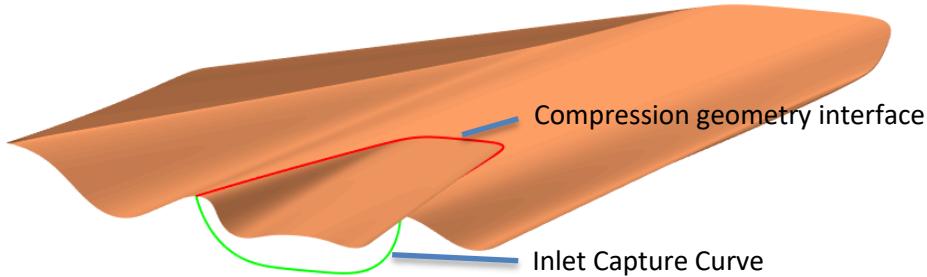
Figure A.30. Osculating cones method generated external compression surface integrated on waverider forebody.
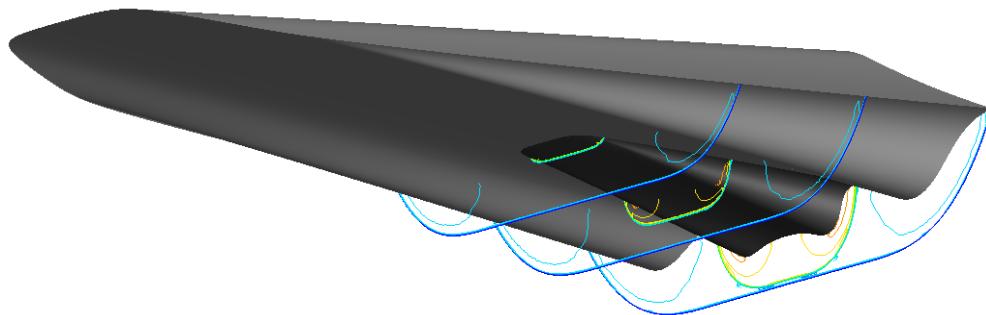

Figure A.31. Pressure iso-lines on three freestream-normal slices, Euler CFD results.
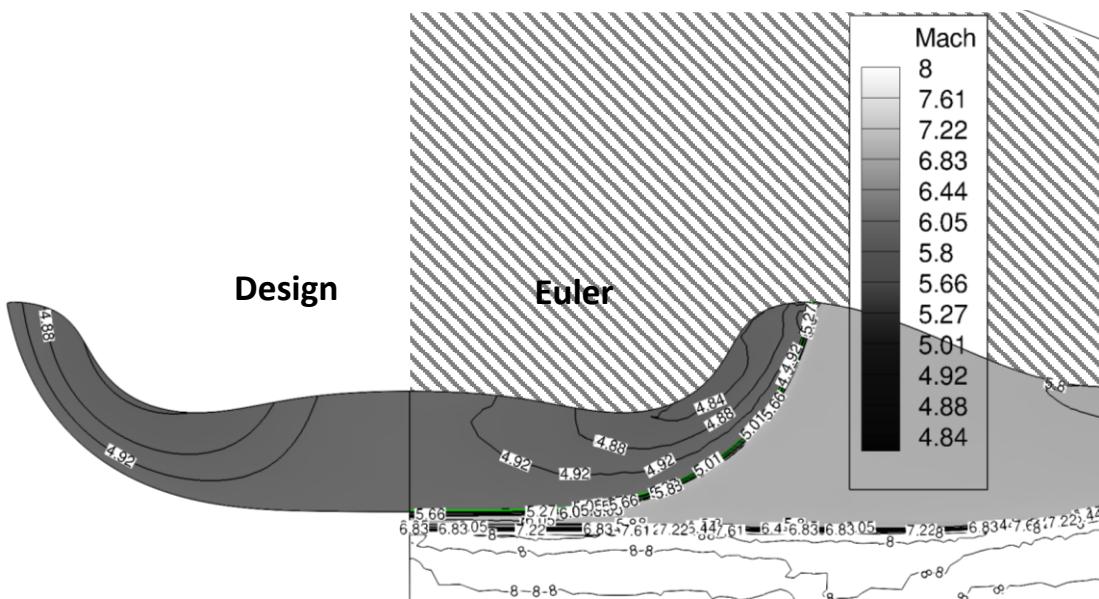
Figure A.32. Comparison between design predicted and Euler calculated Mach number contours at the base plane on the aft end of a waverider with an external compression surface attached. Predicted shock shape plotted in CFD results as well (green).

The second method appears to be similar to the approach followed to design the external compression ramp of the X-51 Waverider experimental aircraft, and was also very recently investigated by Jiang (2017). In this method, a simple planar compression ramp without sidewalls is used, placed within a uniform flow region of a waverider forebody similar to a simple planar compression inlet. However, it starts wider than the required width of the inlet cowl that defines the captured flow, and gets thinner at an angle equal to or slightly larger than the Mach angle of the compressed flow, as seen in Figure A.33 and Figure A.34. Flow uniformity is lost within the Mach cones starting at the truncated edge of the ramp. By truncating the ramp at an angle close to the Mach angle the non-uniform flow that has been compressed and will not be captured can be minimized. This leaves a uniform flow region at the central section of the inlet ramp. An approach like this might not be suitable for supersonic applications due to the large Mach angles of lower Mach numbers, which result in more discarded flow. It is however better suited for hypersonic speeds due to the small Mach angles of the flow.
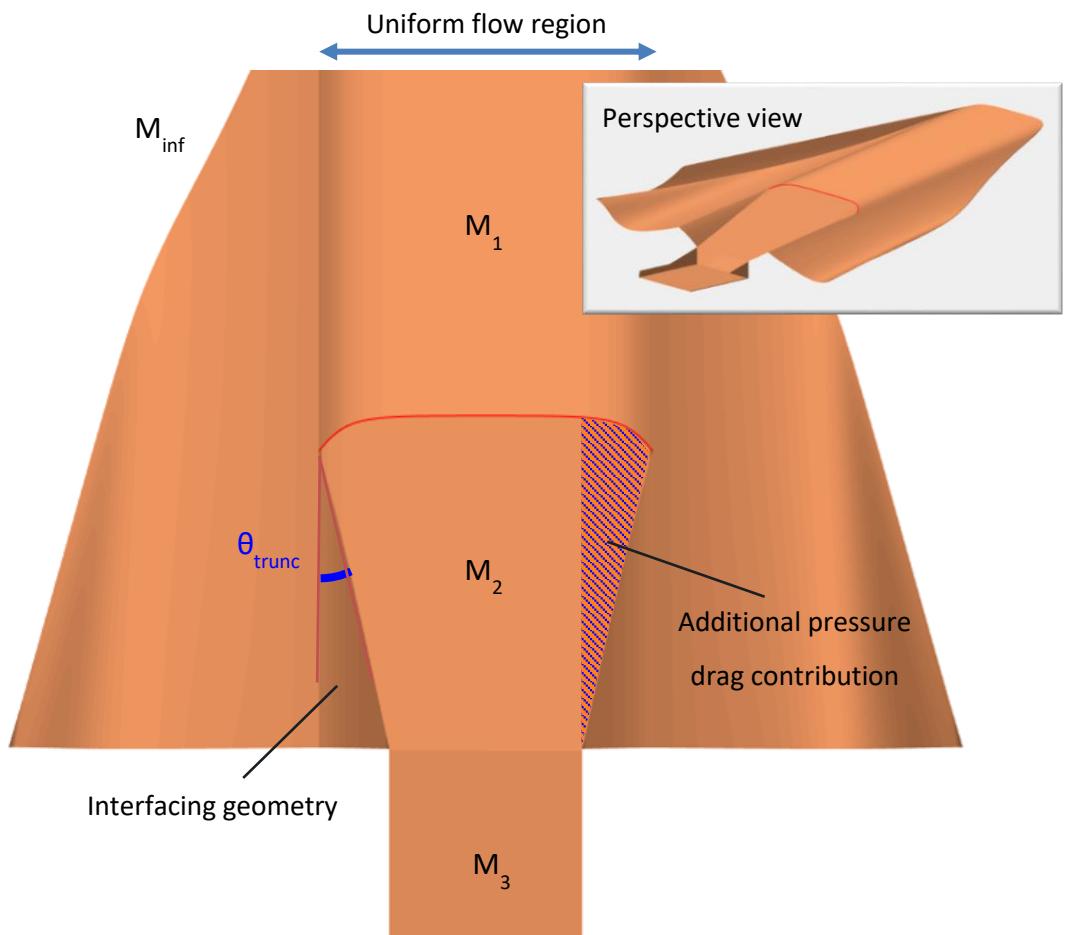
Figure A.33. Bottom view of truncated external compression ramp on the underside of waverider geometry. M_1 the Mach number behind the forebody shock, M_2 the Mach number behind the external compression shock, M_3 the Mach number behind the internal compression shock originating at the cowl lip, and θ_{trunc} the truncation angle.
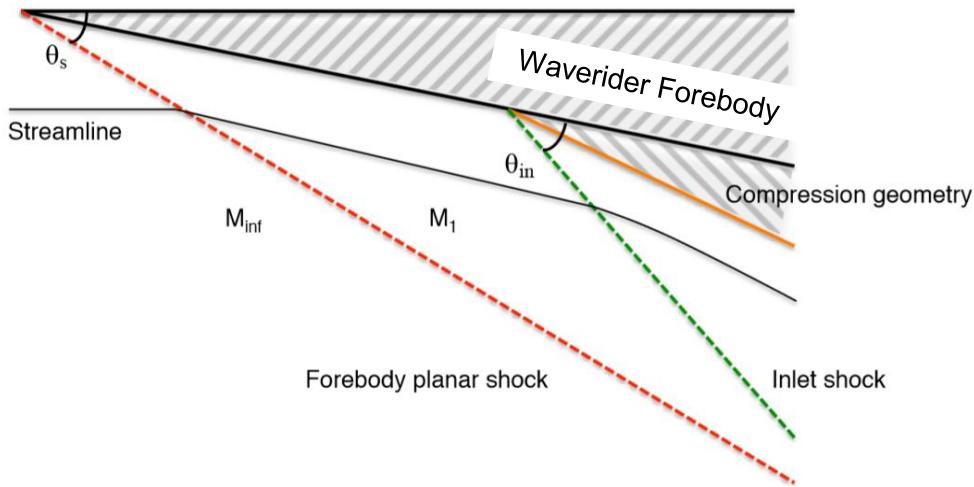
Figure A.34. Symmetry plane slice of waverider and external compression geometry, with θ_s the forebody shock angle, θ_{in} the inlet external compression shock angle.

This type of compression surface has an additional contribution to pressure/wave drag due to the part of the flow that is compressed and then discarded. It is for this reason that we try to minimize the flow that is not captured by truncating the ramp at the smallest angle possible. On the other hand, skin friction drag is reduced when compared to a ramp with side walls due to the reduced wetted area as well as the absence of corner flow in the external compression part of the inlet. The increased pressure of the flow in the region that is not captured by the inlet cowl also has a small contribution to lift. CFD calculations for this geometry verify the uniformity of the flow that is captured by the inlet cowl, as seen in Figure A.35 for the Mach 8 designed waverider with integrated compression surfaces that was also shown in Figure A.33.
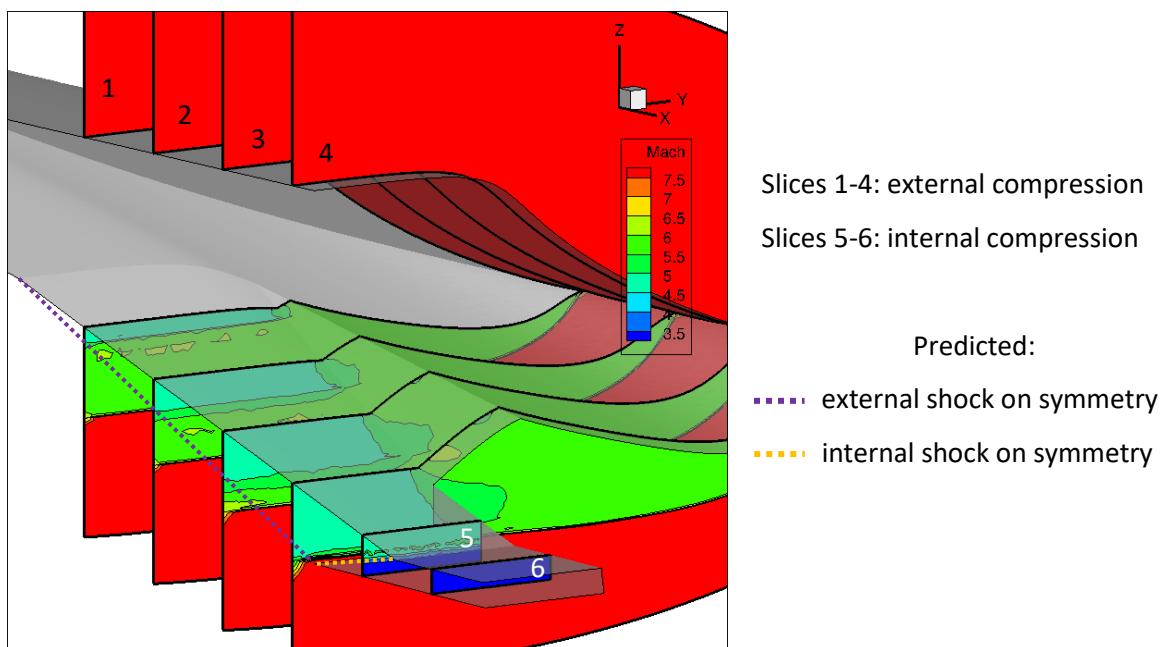
Figure A.35. Mach number contour on four slices around half-symmetric waverider forebody with attached truncated external compression surface and internal compression section.

Glossary of Terms

Base plane: Plane normal to the freestream direction at the aft end of waverider geometries. In the majority of design methods the design-driving curves are drawn on this plane.

Binomial coefficient: The binomial coefficient $\binom{n}{k}$ expresses the number of ways of picking k unordered outcomes from n possibilities and is given by: $\frac{n!}{(n-k)!\,k!}$

Curse of dimensionality: The various complexities and phenomena that arise when analysing data in high-dimensional spaces. For engineering design, as the number of design variables and dimensionality of a design problem increases, the volume of the design space increases so rapidly that the available data become very sparse.

Osculating circle/cone/plane: Osculating circle of a curve at a certain point, is the circle that is locally tangent at that point and also has the same radius of curvature. Osculating cones are cones whose trace on the base plane forms osculating circles of the shockwave profile curve also drawn on the base plane. Osculating planes are the planes defined by the radial line segment between the osculating circle centre and the osculating circle-curve tangency point, and are also parallel to the freestream direction.

Variation diminishing property: In mathematics, the variation diminishing property of certain mathematical objects involves diminishing the number of changes in sign. The variation diminishing property of Bézier curves is that they are smoother than the polygon formed by their control points. If a line is drawn through the curve, the number of intersections with the curve will be less than or equal to the number of intersections with the control polygon.

Bibliography

Anderson, J. D. (2006), "Hypersonic and High Temperature Gas Dynamics", 2nd edition, AIAA Education Series.

Anderson, J. D. (2000), "Introduction to Flight", 4th ed. Boston: McGraw-Hill, pp. 622-626 and 681-707.

Bertin, J.J. (1994), "Hypersonic Aerothermodynamics", AIAA Education Series, New York, New York.

Blosser, M. L, Blankson I. M, Schwoerke, S, Brunson, D, and Hagseth, P (1995), "Wing leading edge design concepts for airbreathing hypersonic waveriders", Journal of Aircraft, vol. 32, pp. 307-312.

Bowcutt, K. G. (2001), "Multidisciplinary Optimization of Airbreathing Hypersonic Vehicles", Journal of Propulsion and Power, 17(6), pp. 1184-1190.

Bowcutt, K. G., Anderson, J. D., JR, and Capriotti, D. P. (1987), "Viscous Optimised Hypersonic Waveriders", AIAA, AIAA-87-0272.

Center, K. B. (1993), "An Interactive Approach to the Design and Optimization of Practical Hypersonic Waveriders", Ph.D. Dissertation, Department of Aerospace Engineering Sciences, University of Colorado.

Center, K. B., Sobieczky, H., and Dougherty, F. C. (1991), "Interactive Design of Hypersonic Waverider Geometries", AIAA Paper 91-1697, AIAA 22nd Fluid Dynamics, Plasma Dynamics and Lasers Conference, Honolulu, HI.

Cerminara, A., and Sandham, N., D. (2015), "Leading-edge Receptivity to Acoustic Waves for High-Speed Flows over a Blunt Wedge", AIAA Paper 2015-3078, 45th Fluid Dynamics Conference, Dallas, TX.

Chang, J. (1991), "A study of viscous interaction effects on hypersonic waveriders", PhD thesis, University of Maryland, College Park.

Chauffour, M. (2004), "Shock-Based Waverider Design with Pressure Corrections, and Computational Simulations", MSc Thesis, University of Maryland.

Chauffour, M.-L., and Lewis, M. (2004), "Corrected Waverider Design for Inlet Applications", 40th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, Fort Lauderdale, FL.

Cockrell Jr., and C. E. (1994), "Interpretation of waverider performance data using computational fluid dynamics", Journal of Aircraft, vol. 31, pp. 1095-1100.

Coquillart, S. (1990), "Extended free-form deformation: a sculpturing tool for 3D geometric modelling", ACM SIGGRAPH Computer Graphics, 24 (4), pp. 187–196.

Bibliography

Corda, S. (1988), "Viscous Optimized Hypersonic Waveriders Designed from Flows over Cones and Minimum Drag Bodies", Ph.D. Thesis, University of Maryland.

Corda, S., and Anderson, J. D. (1988), "Viscous Optimized Waveriders Designed from Axisymmetric Flow Fields", AIAA Paper 88-0369, 26th Aerospace Sciences Meeting, Aerospace Sciences Meetings, Reno, NV.

Cox, C., Lewis, C., *et al.* (1996). "The LoFLYTE Program", Space Plane and Hypersonic Systems and Technology Conference, Norfolk, VA.

Cui, K., Yang, and G. (2007), "The effect of conical flowfields on the performance of waveriders at Mach 6", Chinese Science Bulletin, 52, pp. 57-64.

De Boor, C. (1978), "A Practical Guide to Splines", Springer-Verlag.

Ding, F., Shen, C., Liu, J., and Huang, W. (2015), "Comparison between novel waverider generated from flow past a pointed von Karman ogive and conventional cone-derived waverider", Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering, 229(14), pp. 2620-2633.

Drela, M., and Giles, M. B. (1987), "Visouc-inviscid analysis of transonic and low Raynolds number airfoils", AIAA Journal, Vol 25, No 10, pp. 1347-1355.

Eggers, T., Strohmeyer, D., Nickel, H., and Radespiel, R. (1995), "Aerodynamic off-design behavior of integrated waveriders from take-off up to hypersonic flight", AIAA Paper 1995-6091, 6th International Aerospace Planes and Hypersonics Technologies, Chattanooga, TN.

Eggers, A. J. , Ashley, H., Springer, G. S., Bowles, J. V., and Ardema, M. D. (1990), "Hypersonic Waverider Configurations from the 1950's to the 1990's", Proceedings of the First International Hypersonic Waverider Symposium, University of Maryland, College Park, MD.

Elvin, J.D. (2007), "Integrated inward turning inlets and nozzles for hypersonic air vehicles," US Patent US2007/0187550.

Farin, G. (1993), "Curves and surfaces for computer aided geometric design. A practical guide", Acad. Press.

Fidan, B., Mirmirani, M., and Ioannou, P. (2003), "Flight Dynamics and Control of Air-Breathing Hypersonic Vehicles: Review and New Directions", 12th AIAA International Space Planes and Hypersonic Systems and Technologies, Norfolk, VA.

Giles, M.B., and Pierce, N.A. (2000), "An introduction to the adjoint approach to design", Flow, turbulence and combustion, 65(3-4), pp. 393-415.

Goldberg, G.E. (1989), "Genetic Algorithms in Search, Optimization and Machine Learning", Addison-Wesley, Reading, MA.

Goonko, Y.P., Mazhul, I.I., Markelov, G.N. (2000), "Convergent-flow derived waveriders", Journal of Aircraft, 37(4), pp. 647-654.

Grandine, T. A., Bowcutt, K. G., and Smith, T. R. (2006), "Parametric Hypersonic Design with Osculating Cone Waveriders," Boeing Presentation, FSP Meeting, Vorau, Styria, Austria.

Hayes, J.R., and Neumann, R.D. (1992), "Introduction to the Aerodynamic Heating Analysis of Supersonic Missiles", AIAA Progress in Astronautics and Astronautics, Volume 142, Mendenhall M.R. (Ed.).

He, X., and Rasmussen, M. L. (1994), "Computational analysis of off-design waveriders", Journal of Aircraft, 31(2), pp. 345-353.

Hicks, R.M. and Henne, P.A. (1978), "Wing design by numerical optimization", Journal of Aircraft, vol. 3, pp. 407-412.

Hinman, W. S., Johansen, C. T., & Rodi, P. E. (2017), "Optimization and analysis of hypersonic leading edge geometries", Aerospace Science and Technology, 70(Supplement C), pp. 549-558.

Ingen, J. L. (1956), "A Suggested Semi-Empirical Method for the Calculation of the Boundary Layer Transition Region", Delft University of Technology, Department of Aerospace Engineering, Report VTH-74.

Jiang, C., Hu, S., Gao, Z., Lee, C.-H., and Xue, H. (2017), "Mach Line Cutting of Compression Surfaces for Two-Dimensional Planar Inlets", AIAA Journal, 55(9), pp. 3219-3226.

Jischke, M. C., Rasmussen, M. L., and Daniel, D. C. (1983), "Experimental surface pressures on cone-derived waveriders for freestream M = 3-5", Journal of Spacecraft and Rockets, 20(6), pp. 539-545.

Jones, D.R., Schonlau, M. and Welch W.J. (1998), "Efficient global optimization of expensive black-box functions", Journal of Global Optimization, 13, pp. 455-492.

Jones, J.G. (1963), "A method for Designing Lifting Configurations for High Supersonic Speeds Using the Flow Fields of Non-Lifting Cones", Royal Aircraft Establishment, Report 2674.

Jones, K. D., and Center, K. B. (2002), "Waverider Design Methods for Non-Conical Shock Geometries", AIAA Paper 2002-3204, 3rd Theoretical Fluid Mechanics Meeting, St. Louis, MO.

Jones, K. D., Sobieczky, H., Seebass, A. R., and Dougherty, F. C. (1995), "Waverider Design for Generalized Shock Geometries", Journal of Spacecraft and Rockets, Vol. 32, No. 6, pp. 957-963.

Kara, K., Balakumar, P., and Kandil, O. A. (2007), "Receptivity of Hypersonic Boundary Layers Due to Acoustic Disturbances over Blunt Cone", AIAA Paper 2007-945, 45th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV.

Bibliography

Kinsey, D. D. (1995), "Toward the Direct Design of Waverider Airframes", PhD Thesis, University of Michigan.

Kontogiannis, K., Cerminara, A., Taylor, N.J., Sóbester, A., and Sandham, N. (2015), "Parametric Geometry Models for Hypersonic Aircraft Components: Blunt Leading Edges", AIAA Paper 2015-3580, 20th AIAA International Space Planes and Hypersonic Systems and Technologies Conference, Glasgow, UK.

Kuchemann, D. (1978), "The Aerodynamic Design of Aircraft", Pergamon, Oxford, 448-510.

Kulfan, B. M. (2006), "Fundamental" Parametric Geometry Representations for Aircraft Component Shapes", AIAA 2006-6948, Proceedings of the 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Portsmouth, VA.

Lane, J. (2007), "Design processes and criteria for the X-51A flight vehicle airframe", Boeing Integrated Defence Systems, Huntington Beach, CA, USA.

Langener, T., Steelant, J., Roncioni, P., Natale, P., and Marini, M. (2012), "Preliminary Performance Analysis of the LAPCAT-MR2 by means of Nose-to-Tail Computations", AIAA Paper 2012-5872, 18th AIAA/3AF International Space Planes and Hypersonic Systems and Technologies Conference, Tours, France.

Lewis, M. J., and Chauffour, M.-L. (2005), "Shock-Based Waverider Design with Pressure Gradient Corrections and Computational Simulations", Journal of Aircraft, 42(5), pp. 1350-1352.

Lewis, M. J., and McRonald, A. D. (1992), "The Design of Hypersonic Waveriders for Aero-Assisted Interplanetary Trajectories", AIAA Paper 91-0053, 29th Aerospace Sciences Meeting, Reno, NV.

Lin, N., Reed, H., and Saric, W. S. (1992), "Effect of Leading-Edge Geometry on Boundary-Layer Receptivity to Freestream Sound", Instability, Transition, and Turbulence, pp. 421-440.

Lin, S., and Luo, Y. (1995), "Optimization of Waverider Configurations Generated from Inclined Circular and Elliptic Cones", AIAA Paper 95-6089, International Space Planes and Hypersonic Systems and Technologies, Chattanooga, TN.

Liu, C., Peng, B., Chen, B., and Ji, C. (2016), "Rapid Design and Optimization of Waverider from 3D Flow", 16th AIAA Aviation Technology, Integration, and Operations Conference, Washington, D.C.

Lobbia, M. (2004), "A Framework for the Design and Optimization of Waverider-Derived Hypersonic Transport Configurations", PhD Thesis, University of Tokyo.

Malik, M. R., and Balakumar, P. (2007), "Acoustic receptivity of Mach 4.5 boundary layer with leading-edge bluntness", Theoretical and Computational Fluid Dynamics, 21 (5), pp. 323-342.

Mangin, B., Benay, R., Chanetz, B., and Chpoun, A. (2006), "Optimization of viscous waveriders derived from axisymmetric power-law blunt body flows", Journal of Spacecraft and Rockets, vol. 43, pp. 990-998.

Mazhul, I. I. (2010), "Off-design regimes of flow past waveriders based on isentropic compression flows", Fluid Dynamics, 45(2), pp. 271-280.

Mazhul, I. I., and Rakhimov, R. D. (2004), "Hypersonic Power-Law Shaped Waveriders in Off-Design Regimes" Journal of Aircraft, 41(4), pp. 839-845.

McClinton, C.R. (2008), "High Speed/Hypersonic Aircraft Propulsion Technology Development", In Advances on Propulsion Technology for High-Speed Aircraft, Educational Notes RTO-EN-AVT-150, Paper 1, Neuilly-sur-Seine, France: RTO.

McKay, M.D., Beckman, R.J., and Conover, W.J. (1979), "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code", Technometrics, American Statistical Association. 21 (2): pp. 239-245.

Mill, R.W., Argrow, B.M., Center, K.B., Brauckmann, G.J., and Rhode, M.N. (1998), "Experimental Verification of the Osculating Cones Method for Two Waverider Forebodies At Mach 4 and 6", AIAA Paper 98-0682, 36th Aerospace Sciences Meeting, Reno, NV.

Molder, S. (1967), "Internal, axisymmetric, conical flow," AIAA Journal, vol. 5, pp. 1252-1255.

Nonweiler, T. R. F. (1963), "Delta Wings of Shape Amenable to Exact Shockwave Theory", Journal of the Royal Aeronautical Society, Vol. 67, p. 39.

Nonweiler, T.R.F. (1959), "Aerodynamic Problems of Manned Space Vehicles", Journal of the Royal Aeronautical Society, Vol. 63, pp. 521-528.

O'Brien, T.F., and Lewis, M.J. (1999), "Power Law Shapes for Leading-Edge Blunting with Minimal Shock Standoff", AIAA Journal of Spacecraft and Rockets, Vol., 36, No. 5, September-October 1999, pp. 653-658.

O'Brien, T.F., and Lewis, M.J. (1998), "Power Law Leading Edges for Waveriders Designed with Shock Attachment," AIAA Paper 1998-0600, 36th AIAA Aerospace Sciences Meeting and Exhibit, Aerospace Sciences Meeting, Reno, NV.

Otto, S. E., Trefny, C. J., & Slater, J. W. (2016), "Inward-Turning Streamline-Traced Inlet Design Method for Low-Boom, Low-Drag Applications", Journal of Propulsion and Power, 32(5), pp. 1178-1189.

Palacios, F., Colonno, M.R., Aranake, A.C., Campos, A., Copeland, S.R., Economon, T.D., Lonkar, A.K., Lukaczyk, T.W., Taylor T.W.R., and Alonso, J.J (2013), "Stanford University Unstructured (SU2): An open-source integrated computational environment for multi-physics simulation and design", AIAA Paper 2013-0287, 51st AIAA Aerospace Sciences Meeting and Exhibit, Grapevine, TX.

QstatLab, v6, Design and Optimization Software, available at: http://qstatlab.co.uk.

Bibliography

Rasmussen, M. L., Jischke, M. C., and Daniel, D. C. (1982), "Experimental Forces and Moments on Cone-Derived Waveriders for M = 3 to 5", Journal of Spacecraft and Rockets, 19(6), pp. 592-598.

Rasmussen M L. (1980), "Waverider configurations derived from inclined circular and elliptic cones", J Spacecraft Rockets, 17(5): pp. 537-545.

Rodi, P.E. (2015), "Integration of Optimized Leading Edge Geometries onto Waverider Configurations", AIAA Paper 2015-1700, 53rd AIAA Aerospace Sciences Meeting, Kissimmee, FL.

Rodi, P. E. (2013), "Optimization of Bezier Curves for High Speed Leading Edge Geometries", AIAA Paper 2013-1004, 51st AIAA Aerospace Sciences Meeting, Grapevine, TX.

Rodi, P. (2012), "Vortex Lift Waverider Configurations", AIAA Paper 2012-1238, 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, Nashville, TN.

Rodi, P. E. (2011), "Geometrical Relationships for Osculating Cones and Osculating Flowfield Waveriders", AIAA Paper 2011-1188, 49th Aerospace Sciences Meeting, Orlando, FL.

Rodi, P., and Genovesi, D. (2007), "Engineering-Based Performance Comparisons Between Osculating Cone and Osculating Flowfield Waveriders", 37th AIAA Fluid Dynamics Conference and Exhibit, Miami, FL.

Rodi, P. E. (2005), "The Osculating Flowfield Method of Waverider Geometry Generation", AIAA 2005-511, 44th AIAA Aerospace Sciences Meeting, Reno, NV.

Santos, W. (2011), "Bluntness Effects on Lift-to-Drag Ratio of Leading Edges for Hypersonic Waverider Configurations", AIAA Paper 2011-2321, 17th AIAA International Space Planes and Hypersonic Systems and Technologies Conference, San Francisco, CA.

Santos, W.F.N., and Lewis, M.J. (2005), "Aerothermodynamic Performance Analysis of Hypersonic Flow on Power Law Leading Edges," AIAA Journal of Spacecraft and Rockets, Vol. 42, No. 4, July-August 2005, pp. 588-597.

Sederberg, T. W. (2012), "Computer Aided Geometric Design", Course Notes, Brigham Young University.

Sederberg, T. W., and Parry, S. R. (1986), "Free-form Deformation of Solid Geometric Objects", ACM SIGGRAPH Computer Graphics, 20 (4), pp. 151-160.

Sóbester, A., and Forrester, A.I.J. (2014), "Aircraft Aerodynamic Design: Geometry and Optimization", Wiley.

Sobieczky, H. (2014), "New Design Concepts for High Speed Air Transport", Springer, Vienna.

Sobieczky, H., Zores, B., Wang, Z.,and Qian, Y.J. (1997), "High Speed Flow Design Using Osculating Axisymmetric Flows", Proc. 3rd Pacific Int. Conf. on Aerospace Science and Technology, pp. 182-187.

Sobieczky, H., Dougherty, F., and Jones, K.D. (1990), "Hypersonic Waverider Design from Given Shock Waves", First International Hypersonic Waverider Symposium, University of Maryland, College Park, MD.

Squire, L. G (1971), "A comparison of the lift of flat delta wings and waveriders at high angles of incidence and high Mach numbers", Ingenieur Archiv, vol. 40, pp. 339-352.

Steelant, J., and Langener, T. (2014), "The LAPCAT-MR2 Hypersonic Cruiser Concept", ICAS-2014-0428, 29th Congress of the International Council of the Aeronautical Sciences, St. Petersburg.

Swafford, T. W. (1983), "Analytical approximation of two-dimensional separated turbulent boundary-layer velocity profiles", AIAA Journal, 21(6), pp. 923-926.

Szema, K.-Y., Liu, Z., and Munipalli, R. (2010), "An Efficient GUI Design Tool for High-Speed Airbreathing Propulsion Integration", 28th AIAA Applied Aerodynamics Conference, Chicago, IL.

Sziroczak, D., & Smith, H. (2016), "A review of design issues specific to hypersonic flight vehicles", Progress in Aerospace Sciences, 84 (Supplement C).

Takashima, N., & Lewis, M. (1995), "Wedge-cone waverider configuration for engine-airframe interaction", Journal of Aircraft, 32(5), pp. 1142-1144.

Takashima, N., & Lewis, M. (1994), "Waverider configurations based on non-axisymmetric flow fields for engine-airframe integration", 32nd Aerospace Sciences Meeting and Exhibit, Reno, NV.

Taylor, G. I., Maccoll, J. W. (1933), "The Air Pressure on a Cone Moving at High Speeds", Proc. Roy. Soc. London A, 139, pp. 279-311.

Tian, C., Li, N., Gong, G., Su, Z. (2013), "A parameterized geometry design method for inward turning inlet compatible waverider", Chinese Journal of Aeronautics, Volume 26, Issue 5, October 2013, pp. 1135-1146.

Tincher, D., and Burnett, D. (1994), "Hypersonic Waverider Test Vehicle: A Logical Next Step", Journal of Spacecraft and Rockets, 31, pp. 392-399.

Watts, J. D. (1968), "Flight experience with shock impingement and interference heating on the X-15-2 research airplane", NASA technical memorandum, Flight Research Center, Edwards, CA, USA.

White, F. M. (1974), "Viscous Fluid Flow", McGraw-Hill Co., New York.

Xu, B., and Shi, Z. (2015), "An overview on flight dynamics and control approaches for hypersonic vehicles", Science China Information Sciences; 58: 1–19.

Bibliography

You, Y.C., Zhu, C.X., and Guo, J.L. (2009), "Dual Waverider Concept for the Integration of Hypersonic Inward-turning Inlet and Airframe Forebody", AIAA Paper 2009-7421, 16th AIAA/DLR/DGLR International Space Planes and Hypersonic Systems and Technologies Conference, Bremen, Germany.