

Efficient Parameterization of Waverider Geometries

Konstantinos Kontogiannis[†], András Sóbester[‡]
University of Southampton, Southampton, UK

Nigel Taylor[§]
MBDA, Filton, UK

This paper summarizes the results of investigations into the development of parametric waverider geometry models, with emphasis on their efficiency, in terms of their ability to cover a large feasible design space with a sufficiently small number of design variables to avoid the ‘curse of dimensionality’. The work presented here is focused on the parameterization of idealized waverider forebody geometries that provide the baseline shapes upon which more sophisticated and realistic hypersonic aircraft geometries can be built. We consider three different aspects of rationalizing the decisions behind the parametric geometry models developed utilizing the osculating cones method. Initially, we discuss three different approaches to the design method itself. Each approach provides direct control over different aspects of the geometry for which very specific shapes would be more complex to obtain indirectly, thus enabling the geometry to more efficiently meet any related design constraints. We then look into a number of requirements and limitations that affect the available options for the parametric design-driving curves of the inverse design method. Finally, we estimate the performance advantages that open up with increasing flexibility of the design-driving curves in the context of a design optimization study. This allows us to reduce the risk of over-parameterizing the geometry model, while still enabling a variety of meaningful shapes. While we mainly used the osculating cones method here, most of the findings also apply to other similar inverse design algorithms.

Nomenclature

L/D	=	lift over drag ratio
M_{inf}	=	freestream Mach number
R_c	=	radius of curvature
V/A	=	volumetric efficiency, volume/wetted area
α	=	local plane inclination angle
θ_c	=	cone half angle
θ_s	=	shockwave angle

I. Introduction

W^AVERIDER geometries are considered amongst the most promising in making hypersonic flight viable due to the high lift-to-drag ratios they exhibit at high speeds. This is accomplished by keeping the shockwave generated by the oncoming hypersonic flow attached to the entire leading edge of the lifting body, effectively trapping the high-pressure region of the flow at the underside of the body, greatly increasing the lift. The initial enthusiasm for the waverider concept following its introduction in the late 1950s [1] was suppressed by concerns for the hypersonic viscous flow effects, off-design performance of the geometries, and the erosion of the high L/D capabilities of the idealized shapes following modifications required to enable practical aerospace vehicles. A number of studies on more realistic, viscous optimized waveriders [2, 3], along with advanced analysis tools and better understanding of hypersonic flow, later brought them back to the forefront.

[†] Ph.D. Student, University of Southampton, Boldrewood Campus, Southampton, UK, SO16 7QF.

[‡] Associate Professor, University of Southampton, Boldrewood Campus, Southampton, UK, SO16 7QF; SMAIAA.

[§] Capability Leader Aerodynamic Tools & Methods, MBDA, PO Box 5, Filton, UK, BS34 7QW; SMAIAA.

Since the introduction of the idea, various approaches to designing such geometries have been proposed [4-6], the majority of which are inverse design methods where the bodies are generated utilizing calculated supersonic flowfields and stream tracing techniques. Early inverse design methods were based on planar and axisymmetric shockwave flowfields, both of which were limited by the poor balance of L/D ratios, volumetric efficiency and ease of engine inlet integration. While a variety of shock and flowfield generating bodies have been considered over the years, a more flexible design method that departed from the notion of having a rigid shock generating body is the osculating cones method [7, 8]. The ability to shape the shockwave more freely enables significant improvements in performance characteristics, although it also results in a more complex design problem as the number of design parameters increases, leading to an exponential increase in the size of the design space. Further evolutions of the method are the osculating axisymmetry [9] and osculating flowfield [10] methods, where the shock generating body for each osculating plane is not restricted to a cone, allowing for greater flexibility in the design. Although the geometries provided by inverse waverider design methods are not capable of dealing with all aspects of hypersonic flight and they cannot function as complete airframes by themselves, they do provide a very efficient way to obtain the baseline shapes upon which more sophisticated designs can be based.

As a number of key technologies advanced and the concept of an air-breathing hypersonic aircraft started to look increasingly feasible, various concept geometries and configurations have been emerging in studies. A large group of studies has focused on optimizing certain performance aspects of waverider shapes or optimizing waverider based hypersonic aircraft. These have ranged from simple shape optimizations to complex frameworks for hypersonic aircraft conceptual design studies, a couple of examples being Ref. [11, 12]. Both ends of this spectrum inherently introduce or utilize parametric geometry concepts. However, so far, at least in the open literature, little attention has been given to the development stages of the parametric formulations that define the geometry itself and drive the design algorithms. In other words, while, in many cases, advanced and modern geometry formulations and geometric representations have been incorporated, the decisions behind their use have generally not been thoroughly and systematically investigated and rationalized during the development of the respective parametric geometry models. Thus, at present, the boundaries between over-restricting and over-parameterizing a hypersonic aircraft parametric geometry model are not clear. It is one of the goals of the present work to demonstrate a methodology to evaluate and compare the parametric geometry formulations themselves, from a conceptual design perspective.

Apart from the complexities deriving from hypersonic flow features, another distinctive characteristic of hypersonic air-breathing aircraft designs is the degree to which the engine, inlet and nozzle must be integrated with the rest of the body, as their size relative to the aircraft increases in order to satisfy the needs of high-altitude hypersonic flight. The result of this is a very complex design problem that often needs to be tackled as a whole, both due to the complexity stemming from hypersonic flight characteristics as well as the need to explore a wider design space due to the limited range of available past examples and limited hypersonic airframe design experience in general. Therefore, to enable modern, multi-disciplinary computational design optimization techniques to operate effectively and efficiently, it is very important that the geometry parameterization is as efficient as possible.

The purpose of the present work is to give more insight and provide a basis upon which the decisions involved in the development stages of a waverider-based, parametric geometry model can be rationalized, as well as have a detailed look into a number of geometry parameterization options for the design of waverider forebody geometries generated with inverse design methods, such as the osculating cones method. The material presented is focused on handling of idealized waverider shapes, which can be seen both as standalone geometries or as a component that is further manipulated and embedded in a more complex aircraft model. In the latter case, the scope of the present study is focused on efficiently parameterizing the inversely designed component, and also enabling the geometry to more readily interface with other components of the aircraft. First, following a brief description of the osculating cones method, we discuss the reasoning behind the decision to use inverse design algorithms. We then consider three different aspects of rationalizing the decisions behind the parametric geometry models developed, with many of the findings also applicable to other similar inverse design algorithms. At the first stage of our investigations we see that the osculating cones method can be altered to meet a number of potential design requirements more efficiently. To that end, we discuss three different variations of the method, each providing more direct control to a different aspect of the final geometry, as well as a new hybrid approach able to combine those. We then examine a number of specific requirements and limitations that apply to the design-driving curves of the inverse design method. Finally, the amount of flexibility required by those curves is assessed in order to ensure that the parameterization is capable of producing a variety of meaningful shapes without requiring an excessive number of degrees of freedom. This is accomplished in the context of low fidelity design optimization studies, by estimating the performance advantages that open up with increasing the flexibility of the design driving curves (thus essentially expanding the design space). It is essentially a problem formulation for comparing and quantifying the effectiveness of different parameterization options.

II. Osculating Cones Method

The osculating cones method is an inverse design algorithm proposed by Sobieczky in 1990 [7]. The method utilizes the axisymmetric flowfield of a cone in supersonic flow, whose properties can be calculated by numerically integrating the Taylor-McColl equations [13], in a more arbitrary, non-axisymmetric manner. The design is driven by two curves, a curve controlling the shockwave shape and a curve from which the shape of the leading edge on the shockwave can be determined. Both curves are drawn on the base plane, which is the plane at the back of the geometry and perpendicular to the freestream. A conical flowfield solution of a selected shock angle is then applied to osculating planes that are perpendicular to the shockwave profile curve, while being scaled to match the curvature of the local shock profile, as seen in Figure 1. The ensemble of osculating planes describes a three dimensional flowfield that can be stream-traced to generate the ‘waveriding’ lower surface of the geometry. The resulting three-dimensional flowfield is not an exact solution but a very good approximation of the actual flowfield that the geometry will generate, as small pressure gradients across different planes can cause limited cross-flow. Applicability of the design method has been computationally validated and experimentally investigated [14]. A geometry example with its corresponding design-driving curves is shown in Figure 2.

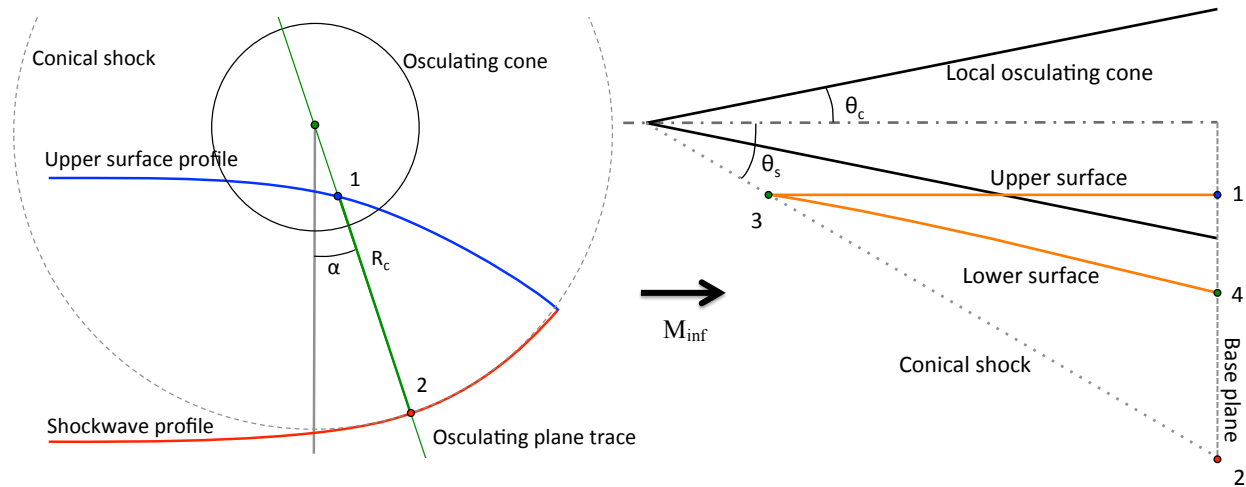


Figure 1. Single osculating plane features on the base plane (left) and waverider design on each plane (right).




Figure 2. Design-driving curves with their control points and resulting osculating cones waverider with the shock surface (grey), a few of the osculating planes (green) and shockwave (red) and upper surface (blue) profile curves illustrated.

At this point, it would also be fitting to discuss why we opted to use an inverse design method, such as the osculating cones method, to obtain the baseline waverider geometries, and then further manipulate those and build upon them to complete the hypersonic aircraft shapes. One may reasonably ask why we could not simply define a fully controllable, parametric wireframe covered with Non-Uniform Rational B-Spline (NURBS) patches. The answer lies in the superior efficiency of the inverse design algorithm, which only requires the parameterization of

two design-driving curves, and a shock angle. While direct waverider design is not a well-posed problem, the inverse design algorithm on the other hand, is able to generate plausible waverider shapes with a small fraction of the degrees of freedom of a geometry consisting of parametric surface patches such as the one seen on the left of Figure 3.

It is possible that a sufficiently flexible NURBS patch could potentially provide geometries that outperform the inversely designed ones, optimizing to locate those geometries would, however, pose a significant computational challenge, rendering it impractical. Even then, one could argue that by further perturbing the inversely designed geometries one could enable the same levels of performance, at a lower cost still, with an example of an inversely design geometry with free-form deformation applied to the upper surface shown on the right of Figure 3. Therefore, the superior efficiency of the inverse design methods is what has made them the ‘standard’ in waverider design.




Figure 3. Waverider-like geometry generated using 4 cubic Bézier patches, with its control points and polygon visible (left), osculating cones waverider with free-form deformation applied on the upper surface (right).

Finally, 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 4 for context – for the remainder of this paper we will be focusing on handling the geometry of the forebody (highlighted in the figure), addressing how to efficiently parameterize it both as a standalone component and for an application like this.




Figure 4. Simple waverider-based hypersonic aircraft geometry example, with the wireframe of the osculating cones generated part of the geometry.

III. Different Approaches to the Design Method

Most waverider inverse design methods involve obtaining the shape of the leading edge on a previously selected shock shape. That 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 we document three of those different approaches to the design method, 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 axisymmetry methods. The distinct feature edges of the geometry whose shapes can be directly controlled with each approach are illustrated in Figure 5. The shockwave shape is defined by a shockwave profile curve on the base plane in all cases. Finally, a new hybrid approach that can combine the advantages of the different methods is also described.




Figure 5. The geometry’s feature edges that can be directly controlled by the different approaches to the design method.

A. 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 Sobieczky⁷. 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 center 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 6.

** We use the word ‘trace’ throughout this paper 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.

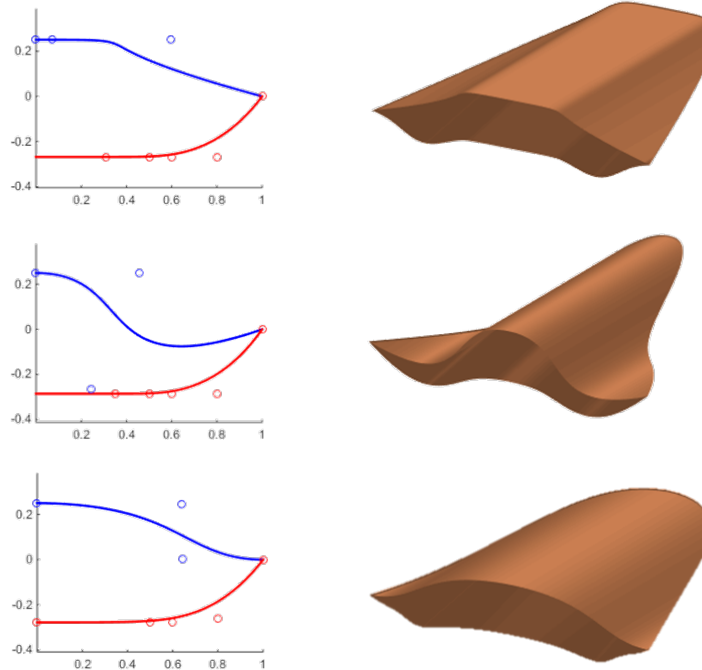


Figure 6. USPD designed example geometries with the upper surface and shockwave profile curves on the base plane (with their control points).

B. Planform Leading Edge Definition (PLED)

Directly designing the planform shape of the geometry is another approach that makes sense 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 [15] (we were unable to identify detailed occurrences 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 the same plane, as seen in Figure 7. 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 8.

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 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 [16] on geometrical relationships for osculating cones and osculating flowfield waveriders provides a set of useful mathematical relationships for 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.




Figure 7. 3D geometrical relations for locating the leading edge point on each osculating plane for a plan view definition of the leading edge shape.




Figure 8. PLED designed geometries with the leading edge profile curves from a plan view (y-x plane) and the shockwave profile curve on the base plane (y-z plane) (with their control points).

C. 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. Ref. [11]), 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 also 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 an engine inlet or a desired lower surface dihedral angle towards the tips of the lifting body. Two example geometries with the corresponding design driving curves are illustrated in Figure 9.

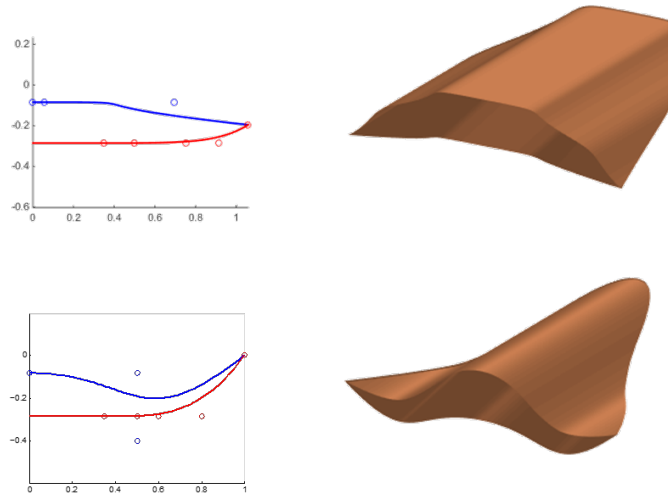


Figure 9. LSPD designed geometries with the lower surface and shockwave profile curves on the base plane (with their control points).

D. Hybrid design approach

Finally, as encountering multiple design requirements or constraints is fairly common, it is possible to blend two of the aforementioned methods in order to improve the efficiency with which the parametric geometry can satisfy them 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 also be applied if smoothness of the upper and lower surfaces is desired at that interface. Moreover, the geometry generation can be split over more than two sections, but this should only be needed in rare circumstances in standard design practice.

Figure 10 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 and/or low speed flight characteristics. Using the LSPD method only for the outer part of the geometry would also make sense to directly control the dihedral angle of the wings, for example to accommodate a control surface. As we mentioned earlier, being able to directly apply design constraints to one of the design driving curves increases parametric efficiency; in the case of multiple sets of constraints, the impact of doing so via the hybrid approach proposed here will be greater still.




Figure 10. Rear and perspective view of LSPD/USPD designed hybrid geometry.

The choice of the appropriate method has a strong impact on design process efficiency. Figure 11 illustrates this: it 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 would require the use of a parametric curve to match the resulting upper surface profile of the geometry, which would require a number of additional design variables and optimization effort, as we will see in section IV/B. We could afford this if we were only optimizing a standalone waverider geometry, but the additional cost might be prohibitive in the context of a full airframe optimization.

A wide range of design requirements can be catered for 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 center of the geometry for a compression ramp, a circular arc for an inward turning type inlet, or even flat profiles at the wing tips to hinge control surfaces. The leading edge sweep can 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 cater for multiple requirements, to be expected especially when integrating additional aircraft components. These techniques provide tools to handle waverider geometries much more efficiently, and proper utilization of the methods discussed here is, also, of significant importance as the parametric geometry model gets more realistic and complex.

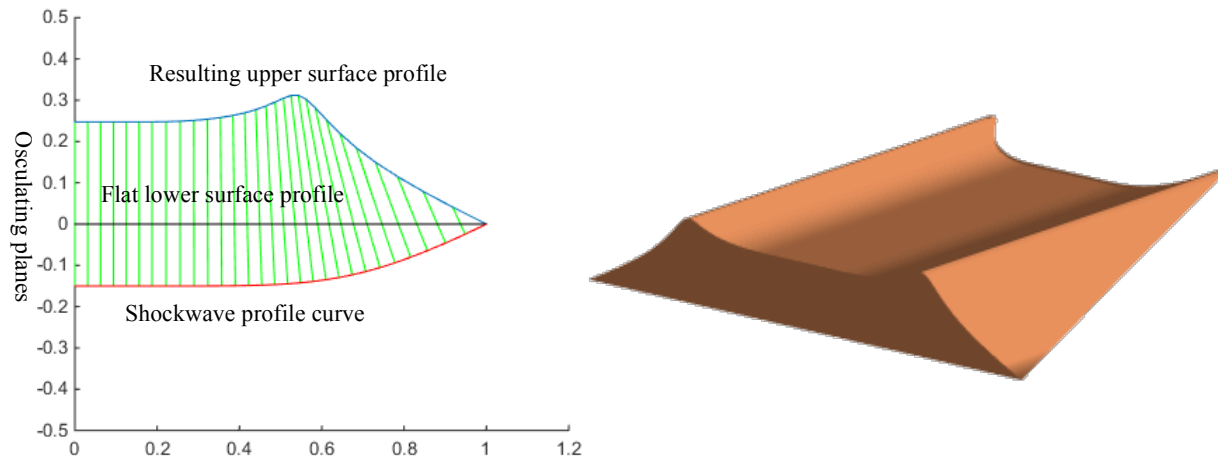


Figure 11. The design-driving curves for a flat lower surface profile and a few of the osculating planes (left), with the resulting waverider geometry (right).

IV. Parameterizing the design-driving curves

A. Limitations of the design-driving curves

With the osculating cones method, the design is very sensitive to the shape of the shockwave profile curve, mainly due to the fact that the local osculating cone flowfield on each plane is scaled according to the radius of curvature of the shockwave profile. Different scaling of the local flowfield affects the location of the local shock-generating cone on each plane and the curvature of the streamline. Given a smooth leading edge, a smooth second order derivative (third order continuity) of the shockwave profile curve is needed in order to obtain a smooth lower surface. This is not an issue when a single infinitely differentiable function is used for the entire curve, but when piecewise functions are used, it needs to be considered. For example, a fourth order spline would be required to satisfy this condition.

In what follows we use a line segment for the part of the shockwave under the center of the forebody where two-dimensional uniform flow – favoring the placement of an engine inlet – is desirable. The rest of the shockwave profile curve is defined using a Bézier curve. In order to achieve third order continuity at the interface between the line segment and the Bézier curve, the first three intermediate control points of the curve have to be on a line. Merely making the shockwave profile curve smooth (first order continuity) with only the first point at the same vertical coordinate would result in a discontinuity of the resulting lower surface, while with second order continuity the lower surface would simply not be smooth and an edge would be present at the resulting lower surface interface. Both cases are illustrated in Figure 12, while Figure 9 earlier featured two geometries where third order continuity is enforced, resulting in a smooth lower and upper surface.




Figure 12. First order (left) and second order (right) continuity at the interface between the straight line and the Bézier curve of the shockwave profile result in discontinuities and sharp edges in the lower surface.

The curve defining the origin of the stream tracing algorithm and leading edge or lower surface profile can be much more flexible in all three aforementioned approaches and any curve parameterization could be considered. Even a series of straight line segments can be used as long as the limitations mentioned in section III are not breached. The only additional restriction that should be applied is that the curve that defines the origin of the stream tracing algorithm should not allow more than one solution in locating the leading edge point on each osculating plane, otherwise surface discontinuities will occur.

As long as these limitations are applied, any curve formulation can be used, and this leaves the designer with numerous options, especially when trying to determine the amount of flexibility that should be given to the design-driving curves once a formulation has been selected. This is the design process step we turn our attention to next.

B. Matching geometries between the methods

The investigation presented in this section was conducted as a means to obtain further insight into a number of subjects. We are reverse engineering a waverider geometry with a different approach to the design method than the one it was originally generated with. Apart from being a simple cross-validation between the three design approaches presented earlier, this investigation serves three more purposes. First, we are looking into how deviations or inaccuracies are translated between the profile curves of the geometry. Second, it is an attempt to quantify the additional effort and increase in dimensionality caused by using an indirect approach to enforce an otherwise straightforward restriction. Finally, it is also a method to obtain some initial insight into the minimum required flexibility of the design-driving curves, in order to enable shapes that are readily available through the other approaches.

To replicate a geometry with a different design method than the one it was created with, the design-driving curve has to match the respective shape of the geometry. The example presented here is the geometry of Figure 11, where a waverider forebody with a flat lower surface profile is initially designed using the LSPD method. This geometry has a more complex upper surface profile in contrast with the very simple line segment used for the design-driving curve, and, although it is a peculiar case, it serves the purpose of providing a straightforward example case. To obtain a flat lower surface with the USPD method, the upper surface profile curve (USPC) would need to match that upper surface shape of the LSPD generated geometry. The best representations of that target upper surface shape using Bézier curves with increasing number of control points are presented in Figure 13. The end points are fixed, while the locations of the intermediate points are calculated to provide a least squares fit to the target curve. For the comparisons we assume a half span of 1m for the geometry.




Figure 13. Matching the upper surface profile of the flat lower surface design with 3, 4 and 5 control point Bézier curves.

Using one intermediate control point is clearly not going to serve the purpose, but two or three points seem to improve approximation accuracy quite significantly. A curve using an additional fourth intermediate control point was also used and the deviations of each optimal representation from the target curve are shown in Figure 14, together with the deviations of the resulting lower surface profiles at the back of the forebody geometries. Figure 15 shows the final geometry using the two intermediate point approximation.

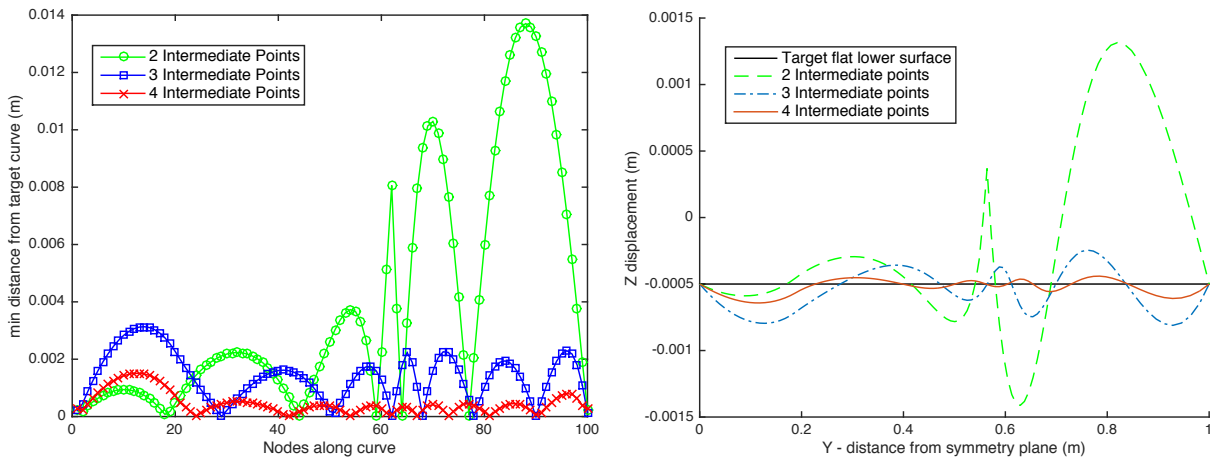


Figure 14. Deviation of the parameterized curves from the target profile curve (left) and deviation of the resulting lower surface profile (right).

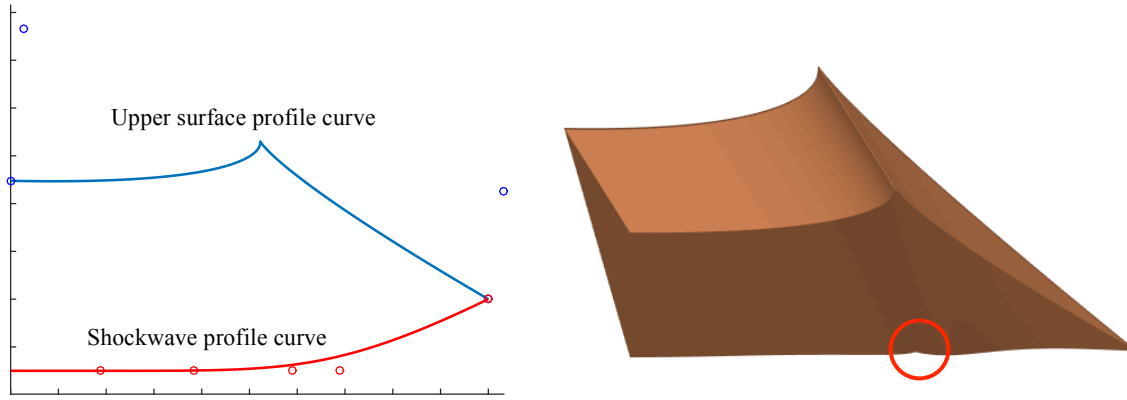


Figure 15. Flat lower surface profile geometry approximation with the upper surface profile definition method and a third order Bézier curve (2 intermediate control points).

While in Figure 15 the imperfections are clearly visible in the final geometry, using one or two extra control points reduces them to the order of $\sim 1\text{mm}$ for a geometry with a half-span of 1m . Therefore, to enable the geometry to achieve such shapes, at least a cubic Bézier curve would be needed with two intermediate control points able to move in both directions, i.e. four design variables more than if the lower surface profile definition method was used. Figure 14 gives us a good indication of how deviations or errors are translated from one feature edge to the other, something that can be useful from a CAD perspective when deciding on the tolerances and resolution of the design driving curves.

V. Parameterization Flexibility Studies

Another 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 this last section, we are 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 it may lie. This is an area where the development of the parametric geometry model starts blending into the preliminary design stages. 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, while 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. In the literature, case studies with parametric geometries that utilize 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 needs to be developed for the choice of such numbers.

Conducting a parameterization efficiency 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 values obtained by solving the Taylor-McColl equations for the lower surface and using freestream values for the upper surface, with a boundary layer integral method to estimate viscous skin friction effects. A very simple 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. As the two example comparisons presented in Table 1 show, the inviscid L/D results obtained in this way are within $\sim 0.5\%$ of the results obtained from inviscid Euler computations, while the viscous results are within $\sim 1.5\%$ of laminar RANS (Reynolds averaged Navier-Stokes) computations. ANSYS Fluent [17] was used for the comparisons, with convergence of the computational grid achieved for around 4 million cells. 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.

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, while in the second we compare different parameterizations of the shockwave profile curve.

A. Case 1.

Designs were optimized with four different curve parameterizations 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 parameterizations 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.

Here we present the results of a set of simulations conducted at a Mach number of 8 and Reynolds number of 8.316×10^6 based on the length of the forebody geometry. 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 conflict 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. Adding more degrees of freedom could further advance the 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. This needs to be assessed in the case of a high fidelity study, where the parametric geometry model in development is going to be used, and in the context of the computational cost of each simulation and the overall computational budget.

In order to generate the fronts, we ran a genetic algorithm [18] 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 16.

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 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.

From the Pareto front landscape of Figure 16 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 performances, 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.

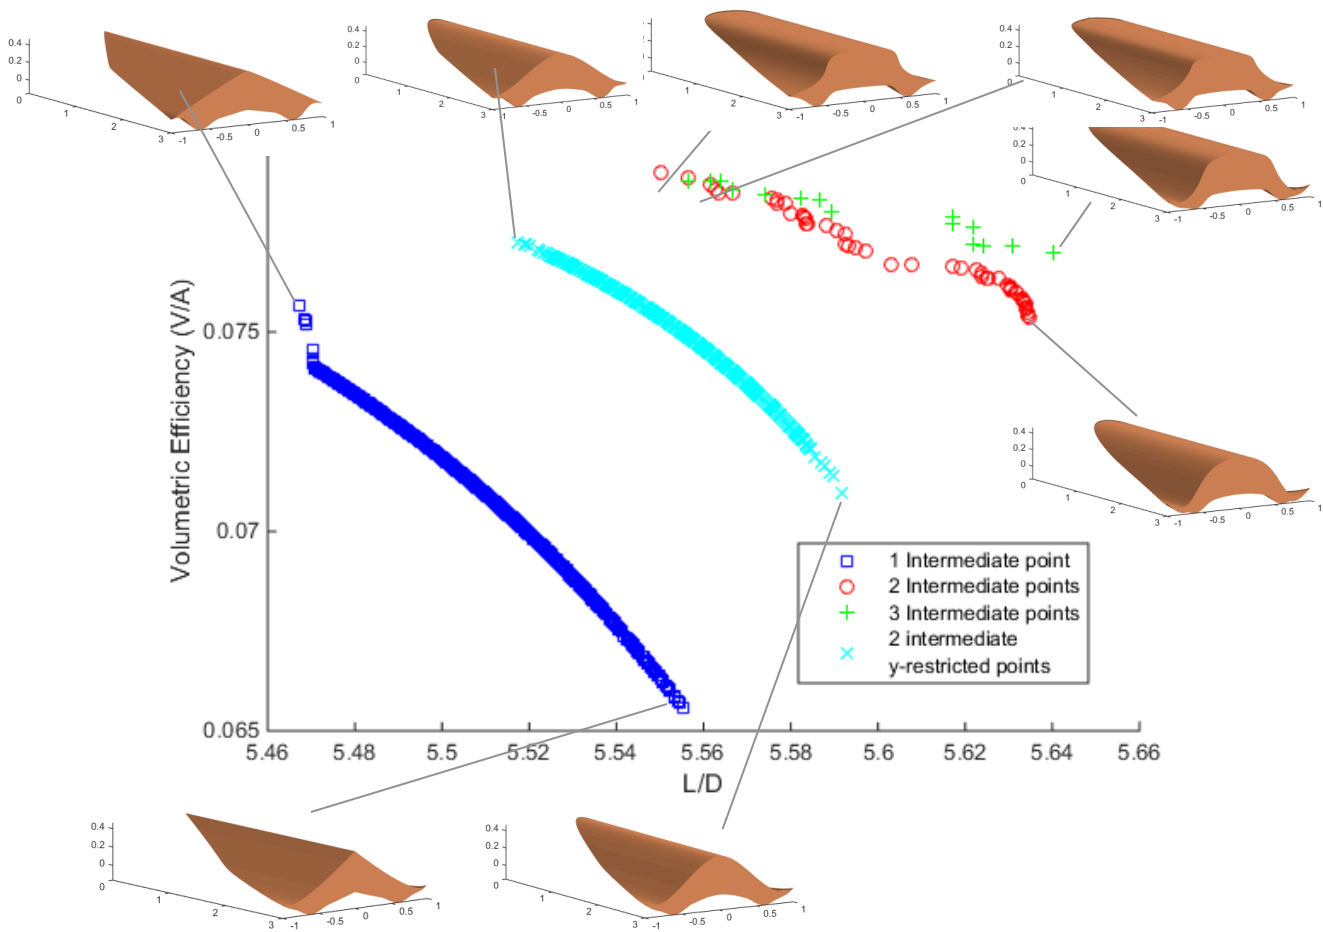
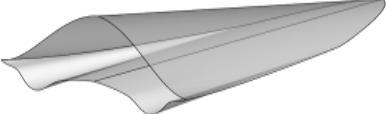
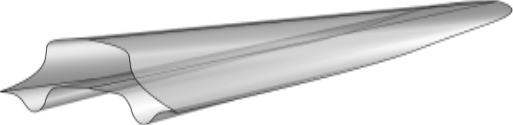


Figure 16. Pareto fronts for different parameterizations of upper surface profile curve.

Table 1. Two of the inviscid and viscous L/D calculation comparisons with CFD

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

B. 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 seen in Figure 17 and also mentioned earlier in Section IV, 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 again 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 18.

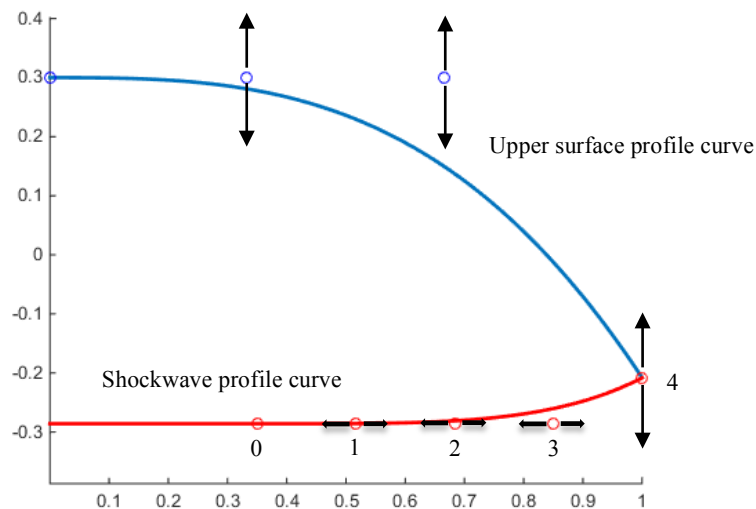


Figure 17. Shockwave profile curve control points under investigation.

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, 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. 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).

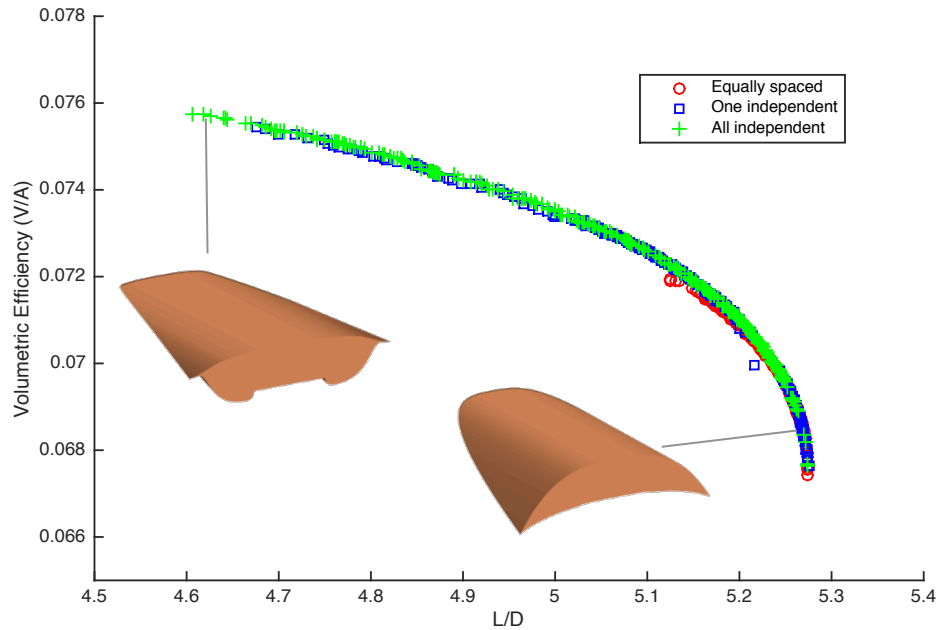


Figure 18. Pareto fronts for different parameterizations of the shockwave profile curve.

It is important to 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 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, flexibility trade studies, such as the one described here, apart from offering a better understanding on handling such geometries, are also useful precursors of more complex design studies, as they 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.

VI. Conclusions

In this paper we presented a series of investigations focused at the efficient handling and parameterization of waverider forebody shapes generated with inverse design methods. We documented three different variations, as well as a new hybrid approach, of the osculating cones waverider forebody design method that offer direct control over different aspects of the final geometry, and we highlighted the importance of opting for the approach that is best suited to any specific design requirements that might be present. Additionally, we reviewed and analyzed the requirements and limitations that apply to the parametric design-driving curves of the osculating cones waverider design method. It is apparent that at least third order continuity of the shockwave profile curve is required to obtain smooth surfaces, while the curve defining the leading edge shape on the shockwave can be much more flexible. To provide further insight and also a basis upon which we can evaluate different parameterization schemes in the parametric geometry model development stages, we presented a methodology to evaluate and compare different parameterization options through a series of multi-objective optimization studies. Two example cases were presented where we attempted to determine the effect of increased flexibility of the design-driving curves on two performance characteristics of interest. Essentially, by conducting such investigations, the decisions made in the process of developing the parametric geometry model can be further rationalized, which is of significant importance

due to the limited experience available in the hypersonic design field. Most of the material presented is also relevant to other similar inverse waverider design techniques, such as the osculating axisymmetry and osculating flowfield methods or even wedge and conical shock based ones. Overall, the work presented can provide useful guidelines to anyone developing or using a waverider-based hypersonic aircraft parametric geometry engine.

Acknowledgments

The authors would like to acknowledge the financial support provided by MBDA UK and Innovate UK: the work reported herein has been undertaken as part of GHandI (TSB 101372), a UK Aerospace Technology Institute project.

Data supporting the computational results of section V are openly available from the University of Southampton repository at <http://dx.doi.org/10.5258/SOTON/403521>.

This paper has been accepted for publication in the AIAA Journal of Aircraft.

References

- [1] Nonweiler, T. R. F., “Aerodynamic Problems of Manned Space Vehicles”, *Journal of the Royal Aeronautical Society*, Vol. 63, 1959, pp. 521–528.
- [2] Bowcutt, K.G., Anderson, J.D., Capriotti, D., “Viscous Optimized Hypersonic Waveriders”, AIAA Paper 87-0272, *24th Aerospace Sciences Meeting*, Reno, NV, 1987.
doi:10.2514/6.1987-272
- [3] Corda, S., and Anderson, J. D., “Viscous Optimized Hypersonic Waveriders Designed from Axisymmetric Flowfields”, AIAA Paper 88-0369, Jan. 1988.
doi:10.2514/6.1988-369
- [4] Eggers, A.J., Ashley, H., Springer, G.S., Bowles, J.V., and Ardema, M.D., “Hypersonic Waverider Configuration from the 1950’s to the 1990’s”, AIAA Paper 93-0774, *31st Aerospace Sciences Meeting*, Reno, NV, 1993.
- [5] Jones, J. G., “A Method for Designing Lifting Configurations for High Supersonic Speeds Using the Flow Fields of Non-Lifting Cones”, Royal Aircraft Establishment, Report 2674, Mar. 1963.
- [6] Rasmussen, M., “Waverider Configurations Derived from Inclined Circular and Elliptic Cones”, *Journal of Spacecraft and Rockets*, Vol. 17, No. 6, 1980, pp. 450-458.
doi:10.2514/3.57771
- [7] Sobieczky, H., Dougherty, F., and Jones, K.D., “Hypersonic Waverider Design from Given Shock Waves”, *First International Hypersonic Waverider Symposium*, University of Maryland, College Park, MD, 1990.
- [8] Center, K.B., Sobieczky, H. and Dougherty, F.C., “Interactive Design of Hypersonic Waverider Geometries”, AIAA Paper 91-1697, *AIAA 22nd Fluid Dynamics, Plasma Dynamics and Lasers Conference*, Honolulu, HI, June 1991.
doi:10.2514/6.1991-1697
- [9] Sobieczky, H., Zores, B., Wang, Z., and Qian, Y.J., “High Speed Flow Design Using Osculating Axisymmetric Flows”, *Proc. 3rd Pacific Int. Conf. on Aerospace Science and Technology*, 1997, pp. 182-187.
- [10] Rodi, P.E., “The Osculating Flowfield Method of Waverider Geometry Generation”, AIAA Paper 2005-0511, *44th AIAA Aerospace Sciences Meeting*, Reno, NV, 2005.
doi:10.2514/6.2005-511
- [11] Takashima, N., “Optimization of Waverider-Based Hypersonic Vehicle Designs”, Ph.D. Dissertation, Department of Aerospace Engineering, University of Maryland, College Park, May 1997.
- [12] Lobbia, M., “A Framework for the Design and Optimization of Waverider-Derived Hypersonic Transport Configurations”, PhD Thesis, University of Tokyo, Japan, 2004.
- [13] Taylor, G. I., Maccoll, J. W., “The Air Pressure on a Cone Moving at High Speeds”, *Proc. Roy. Soc. London A*, 139, pp 279-311, 1933.
- [14] Mill, R.W., Argrow, B.M., Center, K.B., Brauckmann, G.J., and Rhode, M.N., “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, 1998.
doi:10.2514/6.1998-682
- [15] Grandine, T.A., Bowcutt, K.G., and Smith, T.R., “Parametric Hypersonic Design with Osculating Cone Waveriders”, PowerPoint presentation, Boeing, 2006.
- [16] Rodi, P.E., “Geometrical Relationships for Osculating Cones and Osculating Flowfield Waveriders”, AIAA Paper 2011-1188, *49th Aerospace Sciences Meeting*, Orlando, FL, 2011.
doi:10.2514/6.2011-1188
- [17] ANSYS FLUENT, Software Package, Release 14.0, ANSYS, Inc.
- [18] Goldberg, G.E., “Genetic Algorithms in Search, Optimization and Machine Learning”, Addison-Wesley, Reading, MA, 1989.
doi:10.5860/CHOICE.27-0936