University of
Strathclyde
Glasgow

# Strathprints Institutional Repository

http://strathprints.strath.ac.uk/

# Towards Robust Aero-Thermodynamic Predictions for Re-Usable Single-Stage to Orbit Vehicles

Romain Wuilbercq,[*] Abdul Ahmad,[†] Thomas Scanlon[‡] and Richard Brown[§]

*Centre for Future Air-Space Transportation Technology,*
*Department of Mechanical and Aerospace Engineering,*
*University of Strathclyde, Glasgow G1 1XJ, Scotland*

**Re-usable single stage to orbit launch vehicles promise to reduce the cost of access to space, but their success will be particularly reliant on accurate and robust modelling of their aero-thermodynamic characteristics. For preliminary design and optimization studies, relatively simple numerical prediction techniques must perforce be used, but it is important that the uncertainty that is inherent in the predictions of these models be understood. Predictions of surface pressure and heat transfer obtained using a new reduced-order model that is based on the Newtonian flow assumption and the Reynolds analogy for heating are compared against those of a more physically-sophisticated Direct Simulation Monte Carlo method in order to determine the ability of the model to capture the aero-thermodynamics of vehicles with very complex configuration even when run at low enough resolution to be practical in the context of design optimization studies. Attention is focused on the high-altitude regime where lifting re-usable Single-Stage to Orbit configurations will experience their greatest thermal load during re-entry, but where non-continuum effects within the gas of the atmosphere might be important. It is shown that the reduced-order model is capable of reproducing the results of the more complex Monte Carlo formalism with surprising fidelity, but that residual uncertainties exist, particularly in the behaviour of the heating models and in the applicability of the continuum assumption given the onset of finite slip velocity on surface of vehicle. The results suggest thus that, if used with care, reduced-order models such as those described here can be used very effectively in the design and optimization of space-access vehicles with very complex configuration, as long as their predictions are adequately supported by the use of more sophisticated computational techniques.**

## Introduction

THE principal hope for achieving Single-Stage to Orbit (SSTO) with a practically-useful payload mass fraction is seen to be the development of hybrid propulsion technology that exploits, at least at low to intermediate altitudes, the very high specific impulse that is achievable with air-breathing jet engines[1]. Since the ratio of thrust to weight of all practically-conceivable hybrid propulsion systems (at least in air-breathing jet mode) is less than unity, this will necessitate almost universally a lifting configuration for the vehicle. Optimisation to include the low-speed, low altitude flight regime will almost inevitably result in a configuration that is at least partially aircraft-like, and thus a vehicle that is geometrically much more complex than the space vehicles that have been designed in the past (see Fig. 1).

It is well worth bearing in mind that the relatively simple geometry of most spacecraft that have been designed up to now - even those with a lifting configuration such as the Space Shuttle - was, in many cases, borne out of a well-justified scepticism regarding the ability of contemporary analysis techniques to characterise properly the aerodynamics of vehicles with any greater geometric complexity. Indeed, both ascent and descent pose a set of demanding requirements on the accuracy with which the aerodynamic performance of any SSTO vehicle needs to be characterised during its design, but for rather different reasons.

Mis-prediction of aerodynamic drag and lift during ascent to orbit has a direct impact on the sizing of an SSTO vehicle for a given payload capacity in much the same way as for conventional expendable launcher technology, except of course that the impact of uncertainty in the aerodynamic characterisation of the vehicle will have a much
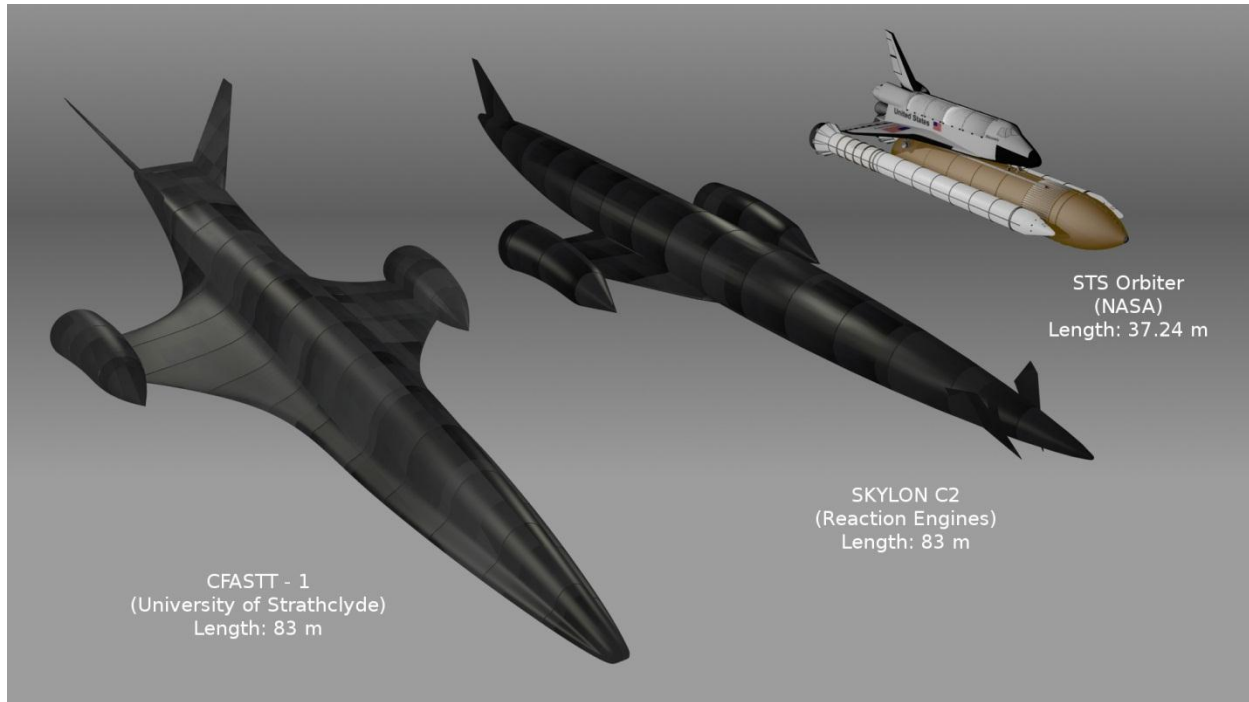
---

[*] Postgraduate Research Student
[†] Postgraduate Research Student.
[‡] Senior Lecturer
[§] Professor, AIAA Senior Member.

**Figure 1. The configuration of a historical re-usable craft (Shuttle, at right) compared to proposed re-usable SSTO vehicles (Skylon, middle, and CFASTT-1, at left).**

more severe effect on the viability of most SSTO vehicles given their characteristically marginal payload performance.

For re-usable SSTO vehicles, however, the very requirement that the vehicle be able to re-enter the earth's atmosphere and be recovered intact for re-launch poses a particularly significant structural mass penalty in terms of the thermal protection system (TPS) that is required during descent to dissipate into the atmosphere the energy gained by the craft during ascent. To reduce the overall mass of the craft to the extent that a worthwhile payload mass fraction is achieved, there is strong design pressure to operate close to the boundaries of existing materials technology. Re-usable SSTO vehicles have limited luxury, for instance, to exploit ablative systems or other simple TPS technologies requiring extensive post-flight refurbishment without destroying the economics of the vehicle or its operational flexibility. One of the principal lessons of the Shuttle programme was indeed how sensitive the economics of an ostensibly re-usable vehicle could be to the robustness and maintainability of its TPS[2].

Mis-prediction of the aerodynamic heating during re-entry is thus a particularly serious problem in the design of a re-usable SSTO vehicle given that any predictive uncertainty must be accommodated by over-design of the TPS, or, in the case where the

heating of the vehicle is over-predicted to the point where indications are that the thermal load cannot be accommodated by a passive system, complex and expensive palliative measures, for example a switch to an active TPS (where for instance the surface of the vehicle is cooled by recirculation of fuel), may be introduced unnecessarily. All of these measures add to the structural mass and reduce the payload that the vehicle is able to carry into orbit.

The aero-thermodynamics of re-entry into the earth's atmosphere is extremely complex. Amongst other effects, the detailed geometry of the shock waves produced by the vehicle, the interaction between the flow behind the shock waves and the viscous flow near the surface of the vehicle, and the presence of real-gas effects such as excitation of internal modes of energy storage within the gas molecules and chemical reactions with the surface of the vehicle all have a strong effect on its aerodynamic loading and heating. Much progress has been made towards developing computational methods that can incorporate all these effects, but even with the advent of massively parallel computation, these methods are extremely time-intensive, often taking several days to produce aero-thermodynamic data for a single point within the operational envelope of the vehicle.

For this reason, preliminary design, and particularly design optimization, must still rely to a significant

degree on the use of much simpler models to characterise the aero-thermodynamics of any prospective vehicle. Practical methods for preliminary design and optimisation are generally limited to those based on very simple models for the behaviour of the flow around the vehicle, for instance those based on a Newtonian flow assumption[3] or the approximation of the vehicle as a composite of simple, non-interacting geometric components for which the flow can be calculated using simple shock theory - albeit perhaps with some empirical embellishment. Indeed, a number of industry-standard methods for employing these techniques currently exist[4,5]. Historically this approach has been very successful[**] notwithstanding notable exceptions such as the well-known Shuttle pitch moment anomaly[6].

The concern, however, is that the inevitable increase in configurational complexity that will accompany the practical embodiment of future SSTO vehicles will force designers to retreat from the conservatism with which these industry-standard models have been applied in the past. It seems thus prudent to re-visit the assumption that these simple models will remain valid for use in the preliminary design of future space-access vehicles, particularly those that have the geometric complexity that is associated with some proposed SSTO designs. In that vein, this paper examines the properties of a newly constructed reduced-order model that has been optimized for use in the multi-disciplinary design and optimization (MDO) of space access vehicles[7]. In this particular application, speed of computational execution is paramount, and thus the behaviour of the method, when applied to the analysis of the aerodynamic and thermal loads on SSTO vehicles with reasonably complex configuration, will be explored both when the technique is applied at high enough resolution for its predictions to be comparable to the output of CFD calculations, as well as when the resolution of the model is downgraded to the point where its application in the context of MDO becomes practically feasible given current limitations on computational power.

It is important to realize though that, in the context of MDO, the purpose of a reduced-order model is not necessarily to reproduce, to the last word in fidelity, the predictions of the sophisticated model for which it is intended to be a surrogate. Instead, the purpose

of the model is simply to yield results that are representatitive enough, or that embody the predictions of the more sophisticated technique to high enough fidelity, so that the optimization process is led efficiently into the same basin of attraction within the configuration space of the system as to which the use of the more sophisticated computational technique would lead. In this way, the more sophisticated technique can then be relied upon in the final stages of the optimization process to further refine the search for the optimum. It will be shown that, if used with care and within the limits of their applicability, reduced-order models of the type described in this paper can yield valuable insight into both the quantitative and qualitative aspects of the aero-thermodynamic loading on complex SSTO configurations, as well as, by inference, form the basis of a viable computational strategy for the multi-disciplinary design and optimization of such vehicles.

In the present analysis, attention will be focused on the high-altitude re-entry regime preceding the onset of maximal heating. The trajectory that the vehicle follows through this regime needs to be tightly constrained so that the capacity of its thermal protection system is not ultimately exceeded during its descent. The aero-thermodynamic properties of the vehicle are difficult to model accurately in this regime, however, because at the highest altitudes they are best characterised as being due to the free flow of molecules over the vehicle while at the lower boundary of the regime they are more effectively treated as being due to a continuum gas flow. By kind acquiescence of Reaction Engines Limited, their Skylon vehicle[8] will be used to illustrate the pertinent physics. Bearing in mind though the limitations of the computational techniques that are employed in the present study, any inference from the results presented here as to the actual performance of the Skylon vehicle should be made with extreme care.

Additional results are presented for a variation on the SSTO configuration employed by Skylon, which has been developed at Strathclyde University and called CFASTT-1. This configuration makes some concessions to ease of manufacture and structural efficiency in order to explore a somewhat different aerodynamic design philosophy to that of Skylon. The engine configuration of the original Skylon is retained, but the rounded fuselage shape of the vehicle is replaced by a flat-bottomed profile, and lateral strakes along the length of the fuselage provide additional lift at high angle of attack. In addition, the loading on the strakes, together with that on the sculpted underside of the rear fuselage, supplant the foreplanes of the original design in

---

[**] Bearing in mind the earlier statement regarding the limitations on geometric complexity that prudent designers have been forced to adopt in the light of the limitations of their methods.

yielding the nose-up pitching moment that is required to maintain a proper re-entry attitude. Longitudinal control under low-speed flight conditions is provided by a V-tail arrangement that, under re-entry conditions, lies in the shadow of the fuselage and wings in order to protect it from excessive heating.

## I. Computational Model

The reduced-order model used to produce the results presented below is a simple panel-based implementation of modified Newtonian theory[9] that has been written principally by the first author of this paper and called HyFlow. The method contains many similarities to, and in fact is inspired to some extent by, existing reduced-order models such as SHABP[4] and CBAERO[10], but contains some innovations to make it amenable to very rapid computations of the aero-thermal loading on arbitrary vehicle configurations in order to support a larger ongoing programme involving the multi-disciplinary design and optimization of hypersonic vehicles. The method calculates the pressure coefficient $C_P$ on the surface of the vehicle, described in practice in terms of its tessellation into a set of flat, triangular panels, according to

$$C_P = C_{P\max} \sin^2 \delta \qquad (1)$$

where $\delta$ is the local inclination of the surface to the free stream and $C_{P\max}$ is an estimate of the stagnation pressure downstream of the shock layer that usually precedes the vehicle during hypersonic flight. The viscous contribution to the loads on the vehicle is then computed by first determining the streamline pattern over the surface of the vehicle. The aim of this procedure is to connect each panel describing the surface of the vehicle to a point on the surface at which the flow is stagnant, and thence to determine the local Reynolds number distribution over the surface of the vehicle. The streamline distribution is determined using a simple tracking routine that propagates the lines in piecewise-continuous fashion from panel to panel while ensuring that they always remain tangent to the flow on the surface of the vehicle. The stagnation features within the flow are determined through the simple expedient of detecting those panels on which the tracking procedure fails to find a viable upstream direction in which to continue propagation of the streamlines. The running length back along the streamline from the relevant stagnation point to the panel in question is then used to determine the local Reynolds number. The Meador-Smart reference enthalpy method[11] is then used to estimate the skin friction coefficient, and thus the viscous force on the panel. Transition from laminar to turbulent flow is modelled using a simple

algebraic criterion[3]. The Reynolds analogy is then applied to compute the local rate of heat transfer to the surface. A Mangler correction[3,10] is used to scale the two-dimensional estimates of the skin friction and heat transfer for each panel that are obtained using this approach into values that are appropriate to a three-dimensional body. It should be remarked though that the choice of scaling factors strongly influences the quantitative predictions of the heating rate that are produced by this approach. Unfortunately there appears to be no particularly substantive basis for the selection of these factors in practice and this should be borne in mind when assessing the results that are presented later in this paper.

The approach to calculating the thermal flux to the surface of the vehicle via the Reynolds analogy is valid where the slip velocities on the surface of the vehicle are small, in other words under conditions that are close to continuum, but can be called into question when the Knudsen number characterizing the flow over the vehicle becomes large. For that reason, a future extension to the present approach is being considered where a bridging-function type approach (akin to that of Chrusciel[12], or Ivanov et al.[13], for instance) might be adopted to account better for the rarefied nature of the flow at high altitude. Significant success with this approach has been achieved when applied to simple shapes (for example to blunt re-entry vehicles), and some of the results presented at the end of this paper are intended to shed some light on the likely success of this intended approach when applied to complex vehicle geometries. Given though the present state of development of the model, and even though the focus within this paper is on predicting the aero-thermodynamics of SSTO vehicles under high-altitude re-entry conditions, cases will only be considered where the flow can be characterized as being reasonably close to continuum.

For purposes of verification, the predictions of HyFlow are compared to equivalent results for the aero-thermal loading on the surface of the vehicle that have been produced using a Direct-Simulation Monte Carlo (DSMC) approach. The code used to produce these results has been written within the framework of the OpenFOAM[14] open-source C++ CFD toolbox and is known as dsmcFoam. The DSMC approach is particularly well-suited to modelling the pressure and thermal loads on the vehicle in the transition regime between continuum and free molecular flow, and the main features of dsmcFoam include the capability to perform both steady and transient DSMC calculations for gases

4

consisting of multiple chemical species, to model arbitrary 2D or 3D geometries using unstructured polyhedral meshes and the ability to exploit unlimited parallel processing. The version of dsmcFoam used to produce the results presented below determines the energetics of intermolecular collisions for polyatomic species using the variable hard sphere (VHS) model[15], and, where included within simulations, the phenomenological Larsen-Borgnakke model is used to distribute post-collisional energy between the translational and rotational modes of the particles[15]. A series of successful benchmark trials have been carried out which have validated the dsmcFoam code for non-reacting flows[14], and developments currently in hand include the implementation of a comprehensive capability using Bird's quantum chemistry model[16] to model the chemical reactions that might take place in the flow under certain flight conditions.

For the DSMC calculations presented here, the atmosphere surrounding the vehicle was assumed to be a binary mixture composed of 23% $O_2$ and 77% $N_2$ by mass. The sensitivity of the DSMC predictions to various assumptions in its formulation, in particular the accuracy with which the thermodynamic characteristics of the atmosphere were modelled, was first assessed by conducting two sets of simulations. The first set of simulations modelled the gas molecules as possessing only their translational mode of energy storage. A second set of calculations additionally accounted for the rotational inertia of the gas molecules, thus allowing a second mode of energy storage and the possibility of a non-equilibrium distribution of energy between the two modes under suitable circumstances. Of course, the rotational degrees of freedom are just the first in a chain of additional modes of energy storage that may be activated under particular flight conditions. For instance, under the real atmospheric conditions that pertain during re-entry, vibrational excitation of the gas molecules as well as dissociation of both oxygen and nitrogen are likely also to contribute to the energetics of the gas.

The flight case that was modelled in the present study corresponds to a re-entry condition for the Skylon vehicle, with the craft at an altitude of 114km travelling at Mach 22.7 and an incidence of 66°. At this altitude the mean free path of the molecules in the atmosphere is about 1.5m, yielding a Knudsen number, based on the length of the vehicle, of the
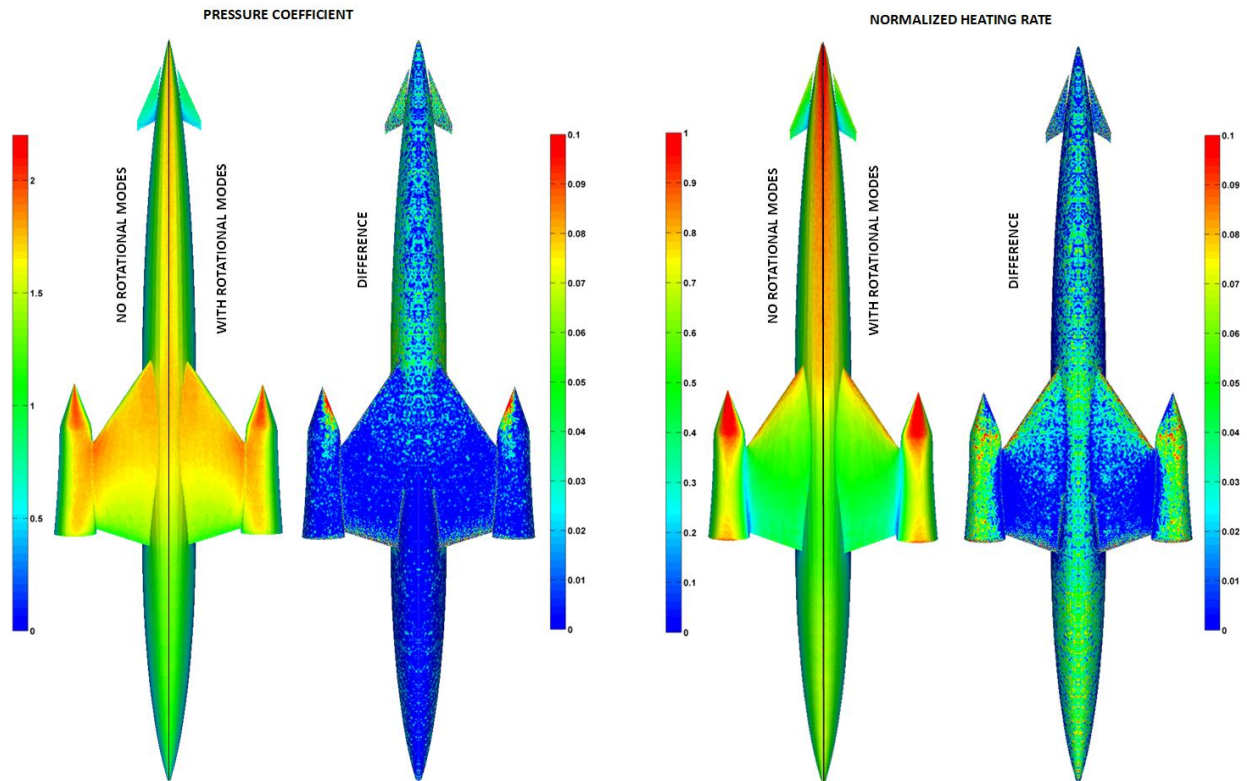


**Figure 2. DSMC-predicted pressure coefficient (at left) and normalized heat transfer rate (at right) on the lower surface of the Skylon re-usable SSTO vehicle flying at Mach 22.7 at 114km altitude. Calculations with and without gas rotational degree of freedom compared.**

order of 0.01. Under these conditions the flow is ostensibly close to the continuum boundary of the transition from continuum to rarefied flow (for example see the calculations by Wilmoth, Blanchard and Moss[17] that suggest that the overall loads on the vehicle would only be very moderately affected by rarefaction effects at this Knudsen number) but is well within the modelling capabilities of the DSMC approach. It is important to remark though that care needs to be taken with such a broad characterization of the flow as that given above, given the possible presence of local non-continuum effects on parts of the vehicle surface, particularly in the neighbourhood of shock interactions and in regions with large surface curvature.The effect of the choice of gas model on the predictions of surface pressure and local heating rate for the Skylon vehicle at the 114km test condition is shown in Fig. 2. The principal differences between the DSMC predictions of the pressure distribution on the vehicle with the two different gas models are confined to the regions of shock interaction on the inner faces of the nacelles, and within a plume of weakly non-equilibrium flow that extends up the sides of the fuselage just forward of the wing[18]. Given the lateral symmetry of the

vehicle, these differences yield a variation in the predicted lift and lift-to-drag ratio of the vehicle in this flight condition of only about one percent, however. The pitching moment about the centre of mass of the vehicle (bearing in mind that, because the vehicle is trimmed, the predicted value is small in any case) is somewhat more strongly affected, with a difference of 14% being observed between the two predictions. As with the traditionally-accepted explanation for the Shuttle pitch moment anomaly[6] it is highly likely that the larger discrepancy in the case of the pitching moment arises from the effect of small differences distributed over the surface of the vehicle being compounded by their distance from the centre of gravity of the vehicle.

Figure 2 also shows the distribution of heating rate over the lower surface of the vehicle, normalised to unity at the point of maximum heating rate, as predicted by the DSMC approach with the two different gas models and the vehicle at 114km altitude. The undersides of the fuselage and nacelles bear the majority of the thermal flux from the gas to the surface of the vehicle, but results presented later suggest that this is for a more interesting reason than
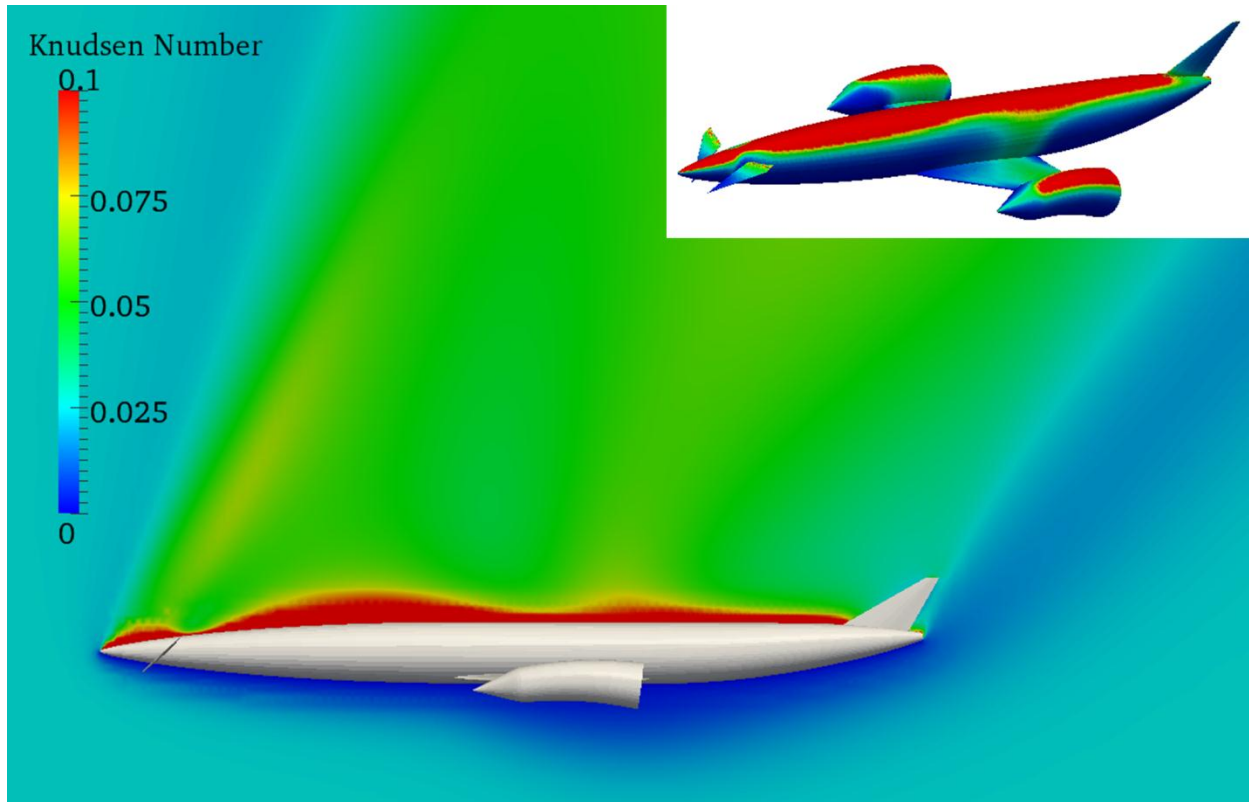


**Figure 3. Variation of molecular mean free path as predicted by dsmcFoam (a) on the slice through the flow that contains the plane of longitudinal symmetry of the Skylon and (b) on the surface of the same vehicle (top right) at 114km altitude.**

American Institute of Aeronautics and Astronautics

simply because the radius of curvature of these parts of the vehicle is small relative to that, for instance, of the underside of the wing - where the heating load is significantly lower. Allowing the gas additional rotational degrees of freedom modifies subtly the predicted distribution of heating rate to the surface of the vehicle, principally on the underside of the nacelles and rear fuselage, but also along the leading edges of the wings. Although the data presented in Fig. 2 show significant statistical scatter, the difference between the predictions using the two different gas models is of the order of 10%.

Despite the existence on parts of the vehicle of some significant differences between the predictions of pressure and heating rate that are obtained using the two different gas models, these results suggest that the principal sources of uncertainty within the DSMC predictions of the flow over complex SSTO configurations such as Skylon are geometrically localized and can be reasonably well understood in terms of their physical origin. Indeed, thus, the data presented here tends to confirm the robustness of the aero-thermodynamic predictions obtained using

dsmcFoam under the given flight conditions, and thus their suitability, if treated with the requisite caution, as a standard against which the predictions of the reduced-order HyFlow model can be assessed.

Before proceeding further, it is important to justify more carefully the use of the essentially continuum-based approach of the reduced-order, HyFlow model when applied to the analysis of the aero-thermal loads on the surface of the vehicle at the flight conditions modelled within this paper. This is important since, as described above, these conditions lie very close to the domain of flight where the onset of non-continuum effects might well be expected. To support the validity of the approach adopted here, Fig. 3 shows the variation of molecular mean free path, as predicted by dsmcFoam, on the slice through the flow that contains the plane of longitudinal symmetry of Skylon when operating at the 114km flight condition described above. The figure shows the vehicle to be preceded by a very diffuse shock structure that compresses the gas and reduces the molecular mean free path near the windward surface of the vehicle quite considerably even compared to
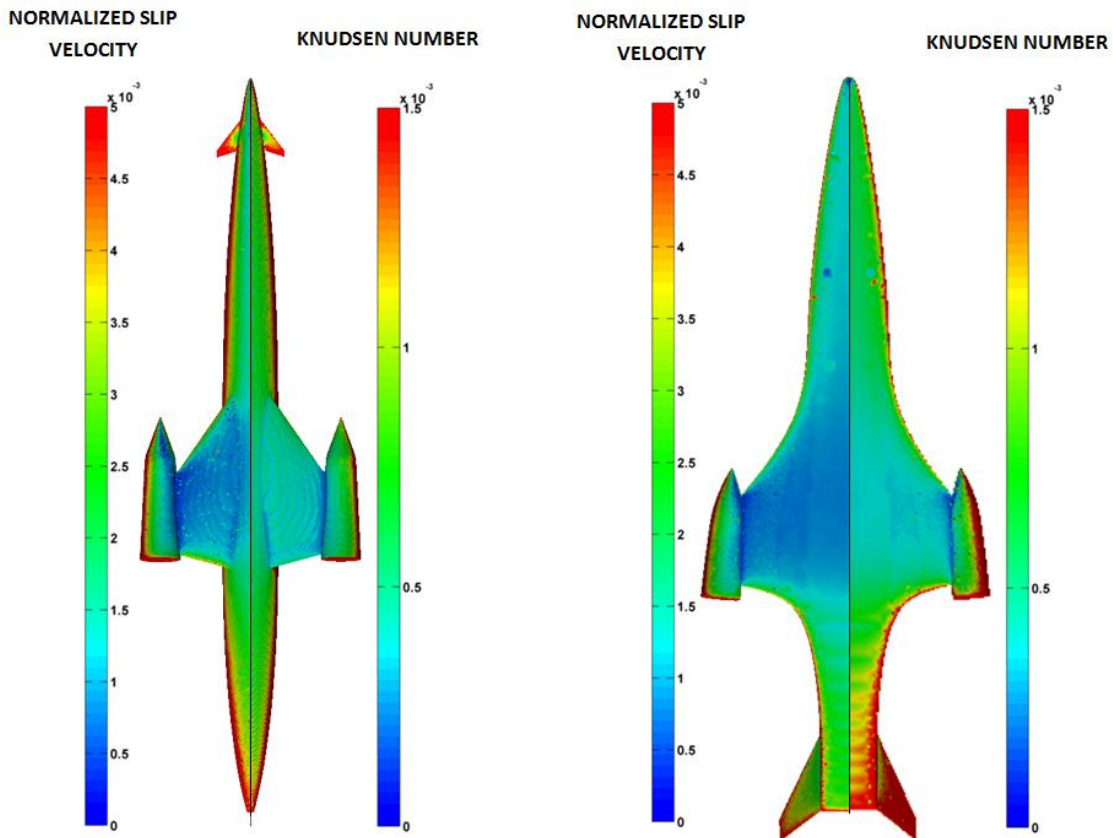


**Figure 4. Variation of slip velocity and molecular mean free path on the windward surfaces of Skylon (left) and CFASTT-1 (right) at 114km altitude as predicted by dsmcFoam.**

its free-stream value, which, as remarked earlier, at the 114km altitude of the simulation is approximately 1.5m. The extended plume of rarefied flow that is produced downwind of Skylon as a result of the acceleration of the flow from the windward to the leeward side of the vehicle is also very clearly apparent in the figure. Figure 3 also shows the overall distribution of molecular mean free path on the surface of the vehicle (scaled by vehicle length to obtain a measure of the global Knudsen number), and confirms that the principal regions of flow rarefaction are confined to the upper surface of the vehicle. Indeed, the mean free path of the gas molecules on the windward side of the vehicle is comparatively small. Figure 4 shows in more detail the distribution of the molecular mean free path on the windward surface of Skylon and compares this to the DSMC-calculated slip velocity on the underside of the vehicle. The appearance of velocity slip at solid surfaces is generally held to be the first direct indicator of the onset of rarefied flow and the

associated non-equilibrium within the gas that leads eventually to the breakdown of the Reynolds analogy and the direct link between the heat transfer to the surface of the vehicle and the shear stress. The predicted distributions of slip velocity and mean free path shown in Fig. 4 show that velocity slip becomes significant only in those regions of the surface where the aero-thermal loads on the vehicle are moderate in any case (see Fig. 6 later), and that the molecular mean free path is small enough almost everywhere on the underside of the vehicle (bearing in mind the reservations expressed above regarding local rarefaction effects) for the flow on the windward side of the vehicle to be treated, with some confidence, as a continuum. These observations taken together suggest thus that the use of the modified Newtonian approach for calculating the pressure distribution on the surface of the vehicle, as well as the boundary layer approach and Reynolds analogy for calculating the heat transfer to the surface of the vehicle, at least for the underside of the vehicle where those
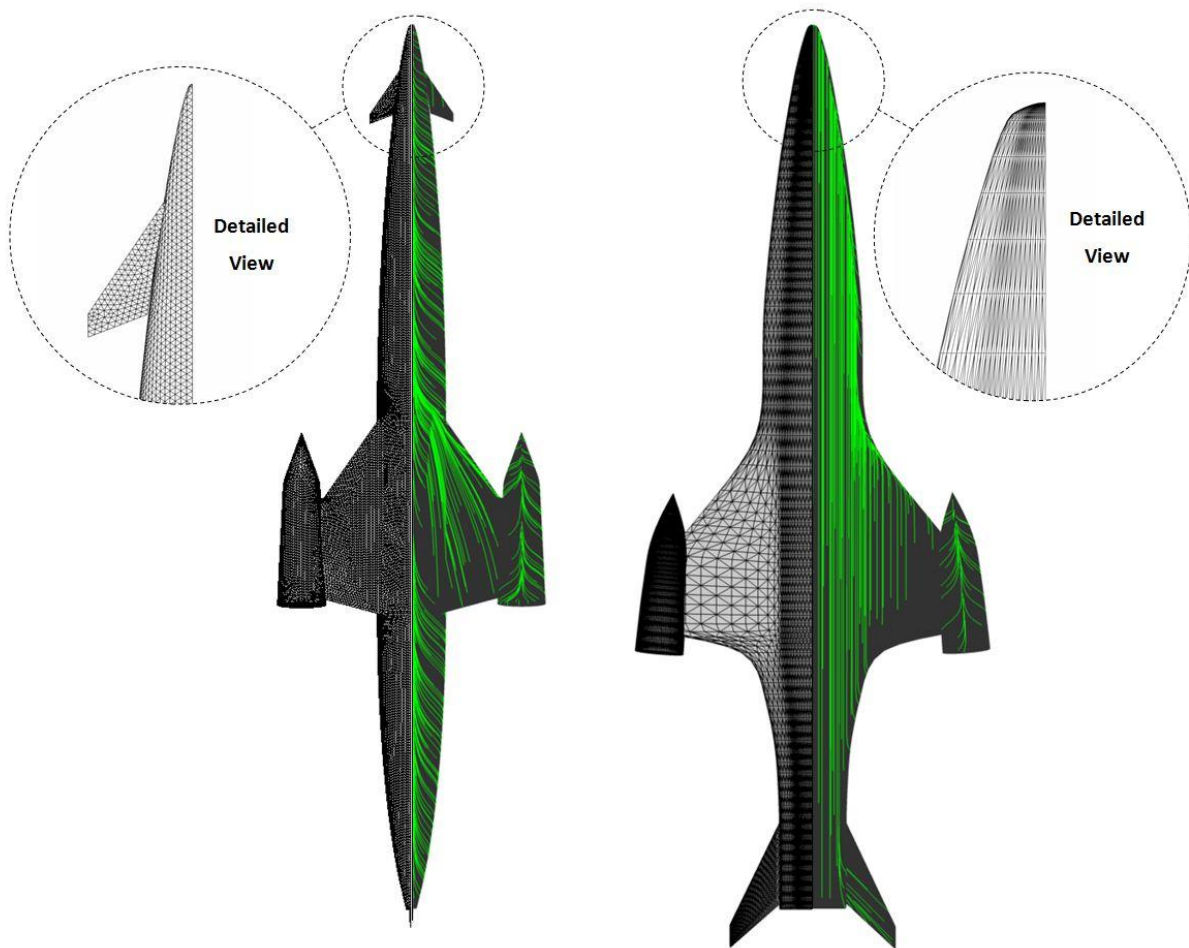


**Figure 5. Surface discretization and HyFlow-calculated streamline pattern on the lower surface of both Skylon (at left) and CFASTT-1 (at right).**

properties have greatest bearing on the overall aero-thermal loads on the vehicle, should be entirely appropriate at the flight conditions that are analysed in this paper.

## II. Reduced-Order Model

High-resolution calculations with HyFlow were compared with DSMC predictions for both the Skylon and the CFASTT-1 vehicles at the 114km flight condition. This was done in order to explore the capability of the reduced-order approach to capture accurately the aero-thermodynamics of a complex aerospace vehicle when run at high enough resolution for its output to be comparable with that obtained from more powerful, conventional numerical techniques. Indeed, in these calculations, the surface of the vehicle was discretized finely enough for the resolution of the HyFlow computations to be comparable to that of equivalent, properly constituted  DSMC calculations that were run with sufficient cell and particle density to resolve accurately the flow surrounding the vehicle.

In the reduced-order simulations, the surface of the Skylon vehicle was discretized using about 80,000 triangular panels and that of the CFASTT-1 vehicle using about 150,000 panels. In order to yield results with as little bias as possible, attention was paid to

achieving as close as possible to a uniform distribution of equilateral triangular panels over the surface of each of the vehicles, rather than, for instance, concentrating panels in the regions of the highest gradients of heat transfer to the surface. Figure 5 shows the distribution of panels on the surface of both vehicles that was adopted for this high-resolution validation study. The figure also shows the pattern of streamlines that was obtained on the surface of each vehicle using the procedure employed within HyFlow.

At the 114km flight condition, the flow on the lower surface of the fuselage of Skylon is characterized by a single stagnation line that runs the entire length of the centreline of the body. The surface streamlines diverge from this stagnation feature, subsequently wrapping around the fuselage and migrating the surface flow towards the lee side of the vehicle. The flow then leaves the surface along a separation line that runs the length of the top of the fuselage. The qualitative nature of the flow over the engine nacelles is similar to that over the fuselage. Bearing in mind the importance of the running length of the streamlines in governing the local Reynolds number and thus the heating of the vehicle, the flow pattern that is found on the lower surface of the Skylon wing at this flight condition is extremely interesting. The streamlines over this surface originate predominantly
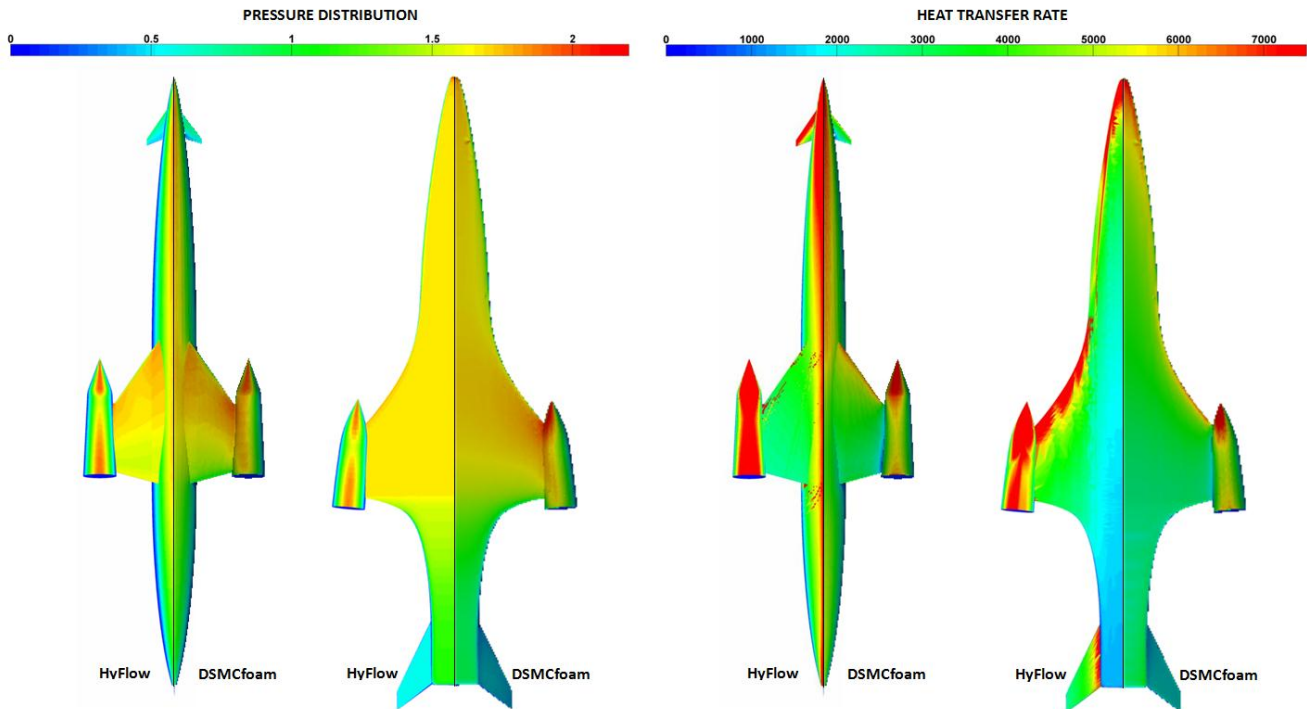


**Figure 6. Comparison of HyFlow and dsmcFoam-predicted pressure distributions on the Skylon and CFASTT-1 vehicles (left) and heat transfer rates (right).**

in the stagnation line along the centerline of the fuselage, and thus are significantly longer than if they had originated, for instance, in a stagnation line lying along the leading edge of the wing itself. Indeed, Fig. 5 reveals the streamlines on the forward part of the wing to lie almost parallel to the wing leading edge, and this rather fortuitous property of the Skylon geometry appears to have a marked effect on the heating of the lower surface of the wing, as shall be shown later in this paper. In contrast, the foreplanes, operating, as they do at this flight condition, at a smaller angle of attack to the oncoming flow than the wing, are characterized by stagnation lines running along their leading edge, and by very short streamlines that run almost parallel to the chord of these surfaces.

#### Table 1   DSMC-predicted vs HyFlow-predicted aerodynamic loads

**SKYLON**

|  | dsmcFoam | HyFlow |
|---|---|---|
| **Lift [N]** | 229.87 | 254.4847 |
| **Drag [N]** | 1274.56 | 1243.4314 |
| **Cl** | 0.6364 | 0.7046 |
| **Cd** | 3.5288 | 3.4426 |
| **L/D** | 0.18 | 0.2047 |
| **Pitching Moment [N.m]** | 5031.86 | 5433.82 |

**CFASTT-1**

|  | dsmcFoam | HyFlow |
|---|---|---|
| **Lift [N]** | 485.23 | 624.1963 |
| **Drag [N]** | 1878 | 1998.2 |
| **Cl** | 1.34 | 1.7282 |
| **Cd** | 5.2 | 5.5323 |
| **L/D** | 0.26 | 0.3124 |
| **Pitching Moment [N.m]** | 11864.85 | 12061 |

The streamline pattern on CFASTT-1 is significantly different to that on Skylon. The presence of the flat-bottomed fuselage and the strakes along the sides of the forward fuselage yields a flow over the windward surface of the vehicle that is closely aligned with the free-stream direction. This results in a set of elongated streamlines that run practically the entire length of the underside of the vehicle. These streamlines originate in a crescent-shaped stagnation feature that follows closely the geometry of the leading edge of the strakes on the forward fuselage and replaces the single elongated stagnation feature that characterizes the flow on the lower surface of Skylon. Most interesting is that the strong spanwise flow on the underside of the wings of Skylon is replaced on CFASTT-1 by a set of streamlines that run essentially parallel to the chord of the wing and originate in a stagnation line just underneath its leading edge. The consequences of this difference in streamline pattern for the predicted heating of the lower surface of the wing will be discussed later in this paper.

Figure 6 compares the pressure distributions on the lower surface of the Skylon vehicle as predicted by HyFlow and by dsmcFoam. The figure also shows similar results for the CFASTT-1 vehicle. HyFlow does not capture the slightly elevated pressure in the junction between the nacelle and wing of both vehicle configurations and indeed simplifies the pressure distribution over the nacelle itself. Nevertheless, for Skylon, HyFlow very accurately reproduces the DSMC-predicted pressure field on the wing as well as that on the fuselage. For CFASTT-1, HyFlow again somewhat simplifies the pressure field in the region where the wing intersects the rear fuselage, but in both cases the similarity in the predicted pressure distributions that are generated by the two computational techniques is remarkable. The close agreement between the predicted pressure fields translates into very close agreement in the predictions by the two approaches of the overall loads on the vehicle. Table 1 shows that for Skylon and CFASTT-1 the difference between the drag coefficients predicted using HyFlow and dsmcFoam is about 2.5 and 6.4 percent respectively, and the lift coefficients 11 and 29 percent respectively, yielding a difference in predicted lift to drag ratio between the two methods of about 14 and 20 percent respectively. The discrepancy between the predictions of the pitching moment that are obtained using the two different methods is approximately 8 percent for Skylon and approximately 2 percent for CFASTT-1. Accurate prediction of the overall loads on the vehicle is an obvious pre-requisite for use of the method in trajectory design and optimization, and, although the calculations presented here are certainly at too high a resolution for practical or routine use in this context - given present limitations on computational power- the results presented here provide a very powerful

basic confirmation of the capabilities of HyFlow in this regard.

Figure 6 also compares the distribution of local heat flux to the lower surface of the Skylon vehicle as predicted by HyFlow and by dsmcFoam. The figure also shows similar results for the CFASTT-1 vehicle . As is to be expected, the maximal heating on the vehicles is to be found close to the regions of flow stagnation near the attachment features that were identified using the streamline mapping procedure described earlier. Both HyFlow and dsmcFoam are in broad agreement regarding the regions of maximal heating on Skylon, although HyFlow tends to overpredict the extent of the region of greatest heating on the underside of the engine nacelles. In contrast, the extent of the region of maximal heating on the lower forward fuselage of Skylon, as predicted by HyFlow, agrees well with the predictions of the DSMC method. For the CFASTT-1 configuration, the agreement between the HyFlow and DSMC predictions of the heating distribution on the

underside of the fuselage is fair. Both methods agree that the maximal heating is to be found along an arc-shaped trajectory corresponding to the location of the attachment line just inboard of the leading edge of the fuselage strakes. The major disagreement between the two computational approaches is that HyFlow predicts a region of intense surface heating along the leading edges of the underside of the wing of CFASTT-1 that is very much less marked in the predictions of the DSMC method.

The interesting observation, however, is that no such region of intensified heating near the wing leading edge is predicted using either HyFlow or the DSMC approach for Skylon. Some insight into the origins of this disagreement can be obtained by considering that, in the HyFlow formalism, regions of more intense heating are associated with lower local Reynolds number or, equivalently, with those parts of the vehicle surface that are separated by only a short running length along the streamlines from their associated stagnation features. On Skylon, where the
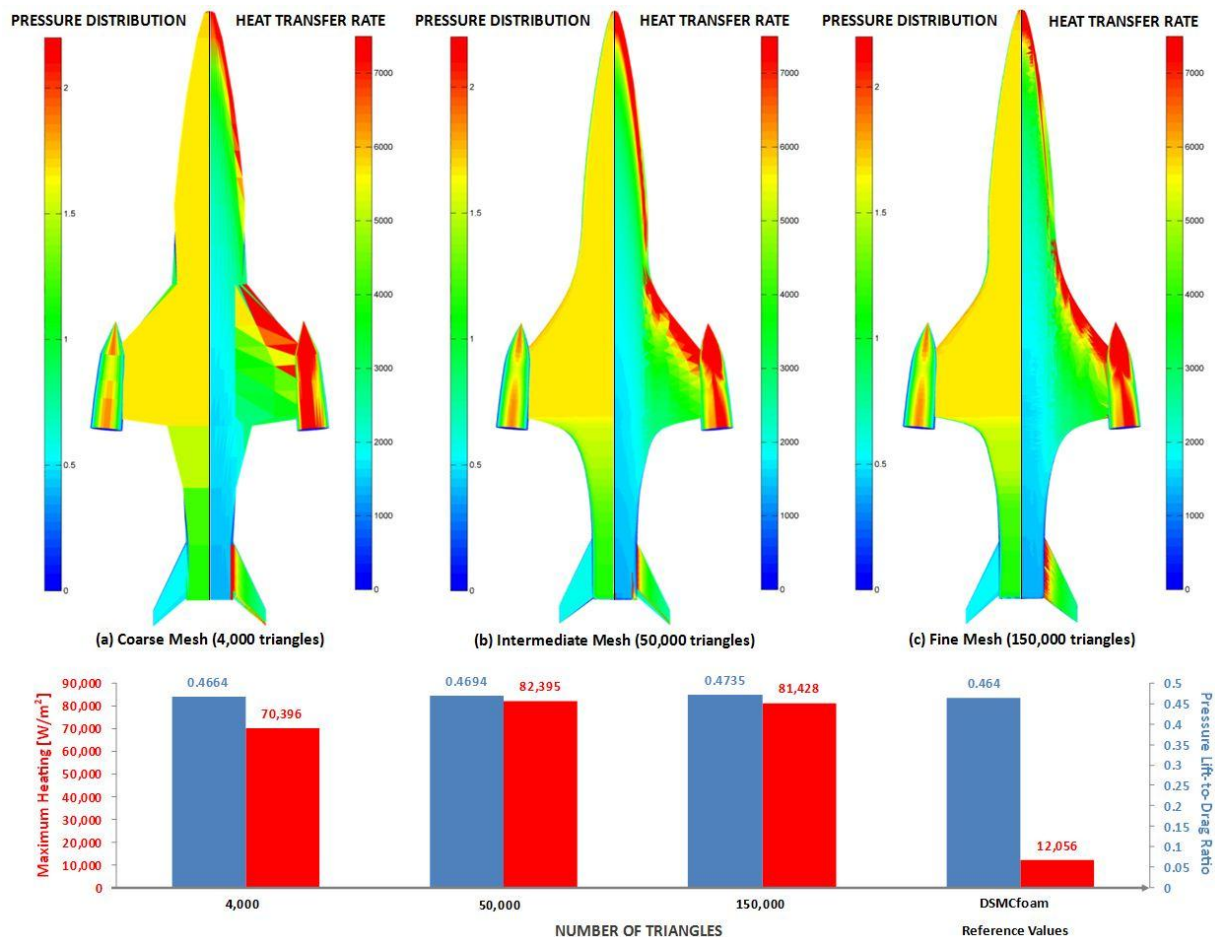


**Figure 7. HyFlow-predicted pressure coefficient and heat transfer rate on the lower surface of CFASTT-1 re-usable SSTO vehicle flying at Mach 22.7 at 114km altitude for (a) Coarse mesh, (b) Intermediate mesh and (c) Fine mesh.**

streamlines near the leading edge of the wing originate, for the most part, quite some considerable distance away on the fuselage, the local Reynolds numbers on the forward part of the underside of the wing are relatively large, and thus the wing would be expected to be subject to lower heating in this region than would CFASTT-1. This configuration is of course subject to much higher heating because the streamlines on the lower surface of the wing of this vehicle originate on the leading edge of the wing itself, and then run aft in an essentially chordwise direction, resulting in comparatively low local Reynolds numbers on the forward regions of the under-surface of the wing.

In the context of trajectory design and optimization, the heating on the vehicle manifests itself usually as a constraint on the allowable performance of the system, and it is usually a global measure of the heat load on the vehicle rather than the distribution of heating over its surface (e.g. maximal heating or integrated heat load) that is of greatest relevance in this context. By suitable selection of the Mangler scaling factors (see earlier), the maximal heating rate on the surface of Skylon, as predicted by HyFlow at the 114km operating condition, can be brought into reasonably close agreement with the predictions of the DSMC approach (where the maximum heating is 16,230 $W/m^2$). As a cross-check of the approach, when the same factor is applied to the calculation of the maximum heating rate on the CFASTT-1 configuration, the HyFlow-predicted result is quite significantly larger than the value predicted by dsmcFoam (see Fig. 7). The uncertainty that is introduced by the presence of these (essentially tuneable) scaling factors in the expressions used in HyFlow to compute the skin friction renders a method for calculating the heat flux to the surface of the vehicle that is not entirely satisfactory or robust, but is arguably workable as long as suitable correlation is available against a check case that has been produced using a higher-fidelity computational approach (e.g. the DSMC model used here).

The difference in the times taken to execute the two different computational approaches is highly relevant given the intended practical application of the reduced-order HyFlow method to vehicle design and optimisation. The DSMC calculations took in the order of a week when run on 96 parallel processors while a single pressure distribution could be calculated on a desktop PC in a fraction of a second using HyFlow. The calculation of the streamline patterns is the most expensive element of the HyFlow computation, and for the Skylon geometry with 80,000 panels as described above, the additional overhead associated with streamline generation causes the calculation of the heating distribution on the vehicle to take an additional 28 hours on the same desktop machine. The present algorithm within HyFlow calculates an individual streamline for each individual surface panel, however, and this approach is somewhat inefficient and wasteful. Later developments of the method will most likely employ some form of interpolation between nearby panels to reduce the computational overhead associated with streamline generation. Obviously these computational times render both methods in any case too expensive to be used in design-level optimization studies where perhaps thousands of evaluations of the aero-thermodynamics of the vehicle under various flight conditions are required. The advantage of the reduced-order approach is the potential, however, to run at very much lower resolution than is possible with the DSMC approach and still to yield pertinent results. Indeed, in the following section of the paper, the ability of the HyFlow model still to capture the salient features of the aerodynamics and heating of the CFASTT-1 vehicle when run at significantly degraded resolution is explored.

## III.  Simulations at Reduced Resolution

For application in the context of design optimization, speed of execution is paramount. Since the computational effort required to obtain a solution using HyFlow scales almost directly with the number of panels used to discretize the surface of the vehicle, significant reductions in execution time can be achieved by running the model at reduced resolution. An important characteristic of any reduced-order technique when following this approach is then, of course, that the aero-thermodynamic predictions of the model be as insensitive as possible to changes in the quality of the discretization of the surface geometry of the vehicle. This is particularly important when the number of panels used to discretize the surface of the vehicle is reduced to a point where execution times result in a practically feasible implementation within the context of design optimization – bearing in mind of course the effects of future improvements in available computer power in weakening this argument.

Figure 7 compares the distributions of pressure and heating rate over the lower surface of the CFASTT-1 vehicle when the surface of the vehicle is discretized using 4,000, 50,000 and 150,000 panels respectively. The finest discretization is the same as that used in the studies described above, the intermediate discretization is typical of that which might be used for quick design studies, whereas the coarsest
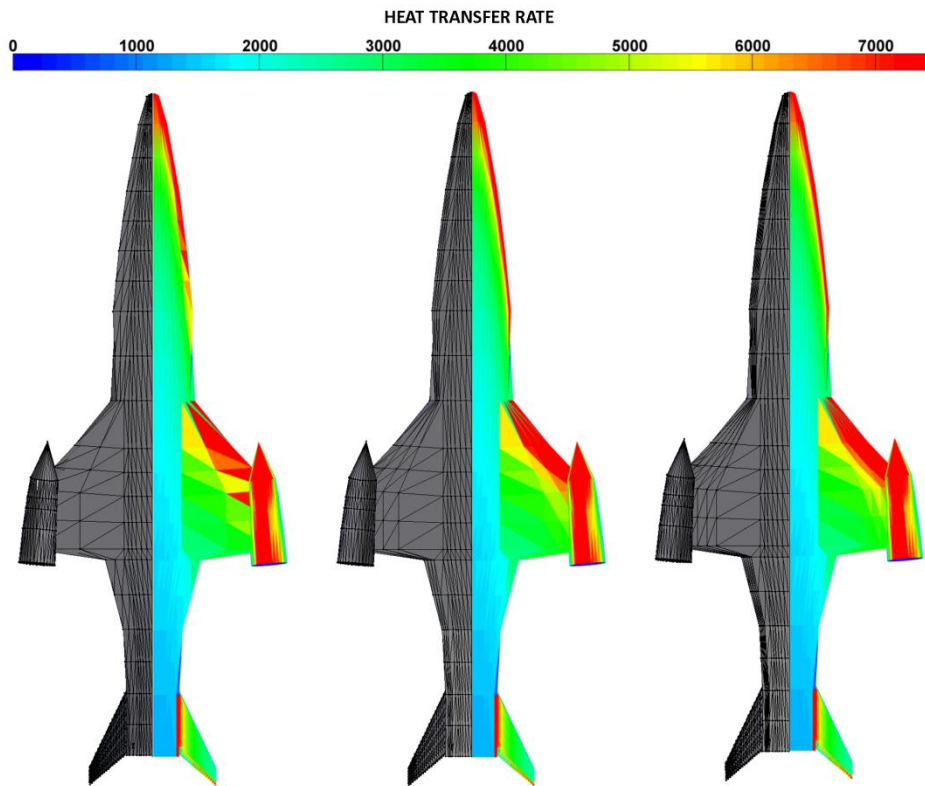
**Figure 8. HyFlow-predicted heat transfer rate on the lower surface of CFASTT-1 showing the effect of redistributing the panels on the surface of the vehicle on the region of intense heating.**

discretization is representative of what might be used as the basis of the aero-thermodynamic model within a multi-disciplinary optimization tool, for instance for the determination of optimal re-entry trajectories. It can be seen that even as the number of panels used to discretize the vehicle is reduced drastically, the salient features of both the pressure and heat flux distributions on the surface of the vehicle are retained.

Particularly important though is that the overall loads on the vehicle, as well as any suitably global quantification of the heating rate, be robust to such degradation in the resolution of the model for it to be of any use in its intended context. The data presented at the bottom of Fig. 7 suggests that reducing the resolution from 150,000 to 4,000 panels has an effect of about 1.5% on the overall loads on the vehicle and about 13.5% percent on the predicted maximum heating rate, but the time taken to compute a solution is in fact reduced by a factor of 15,000. It is also possible to improve the output of the model, and thus to maintain its accuracy to some extent when run at very low resolution, by careful optimization of the distribution of the panels on the surface of the

vehicle. Figure 8 shows the effect of keeping approximately the same number of panels but distributing them intelligently on the surface of the vehicle in order to capture properly the salient characteristics of its aero-thermodynamics. In particular the resolution of the region of intense heating on the forward part of the CFASTT-1 wing can be improved significantly by redistributing panels onto this part of the vehicle's surface from those areas of the vehicle that are subject to a more benign heating environment. By doing this, the predictions of peak heating rate obtained at the lowest resolution of the model can be made to match quite closely those obtained when the model is run at much higher resolution.

## IV.  Application to Design

The advantage of the reduced-order model is of course the ease with which parametric variations in the properties of the system can be explored. The geometry of the lower forward fuselage of the CFASTT-1 configuration was designed partially to explore the possibility of producing a more practical design, as far as the implementation of a thermal

protection system is concerned, than the Skylon vehicle. The underside of the forward fuselage of Skylon is subject to relatively intense heating in a broad swath along almost its entire length, thus, in the current implementation of the design, requiring extensive thermal protection over a large proportion of its surface. In addition its foreplanes are exposed to a particularly intense heat flux, which its parent company intend to address through the provision of active cooling. The CFASTT-1configuration on the other hand does not require foreplanes to trim into its re-entry attitude, and is subject to equivalently intense heating only in a narrow, crescent shaped arc that follows closely the geometry of the leading edge of the strakes on the forward fuselage (see Fig. 6). Inboard of this arc, as remarked earlier, the heat flux is considerably reduced as a consequence of the streamline geometry that is peculiar to this configuration, and, particularly near the rear of the fuselage where the streamwise Reynolds numbers become comparatively large, the heating environment is very benign compared to that of Skylon. This should allow the thermal protection system for this

around the leading edges of the strakes rather than distributed over the entire underside of the vehicle as might be necessary with Skylon.

At the 114km flight condition, the flow over the windward surface of the vehicle should ostensibly be completely laminar. It is well known though that surface roughness can increase the effective Reynolds number of the flow near the surface and promote premature transition to turbulent flow. In practice, however, the degree to which the surface is roughened, or how the roughness of the surface influences the properties of the flow, is not always easy to determine *a priori* (the Shuttle was particularly susceptible to this syndrome) and from the design point of view it is better to create a system which is robust to the effects of such unknowns. To illustrate the power of reduced order modelling in examining the sensitivity of a given design to such uncertainties in characterizing the aero-thermodynamics of the vehicle, the effects on the surface heat flux of varying the effective Reynolds number on the lower surface of the Skylon and CFASTT-1 vehicles was investigated.
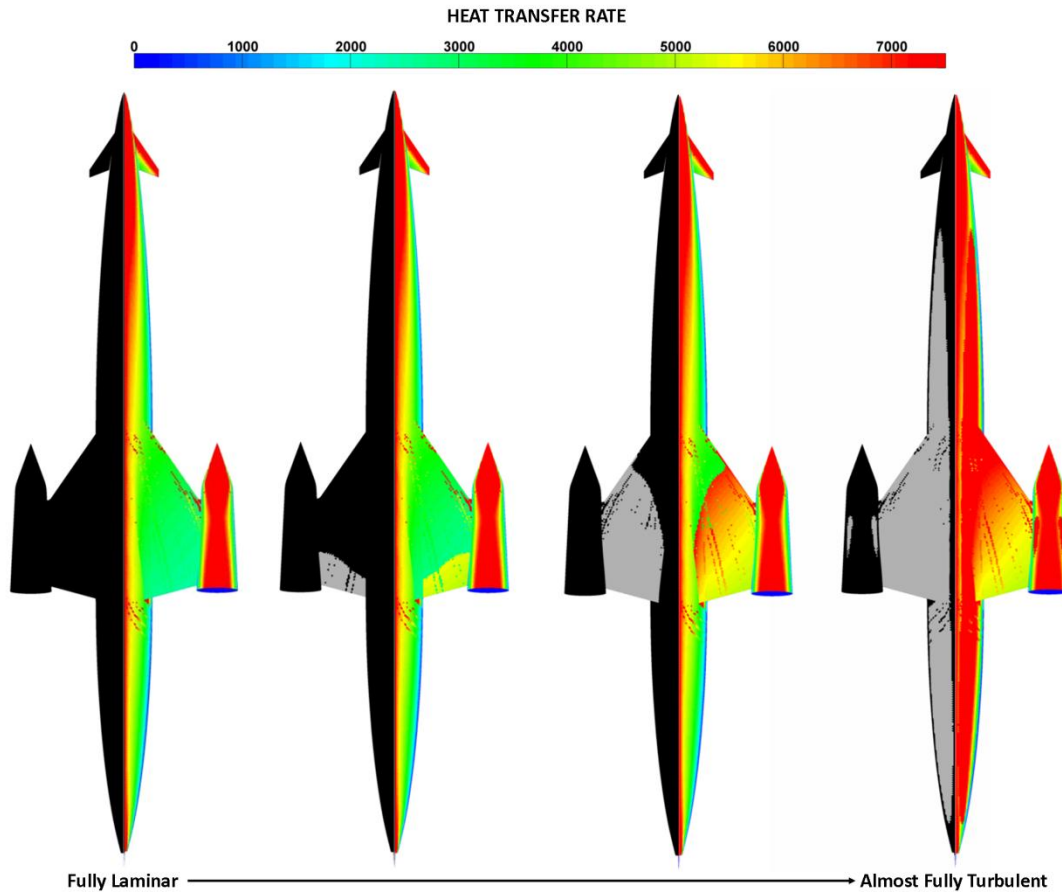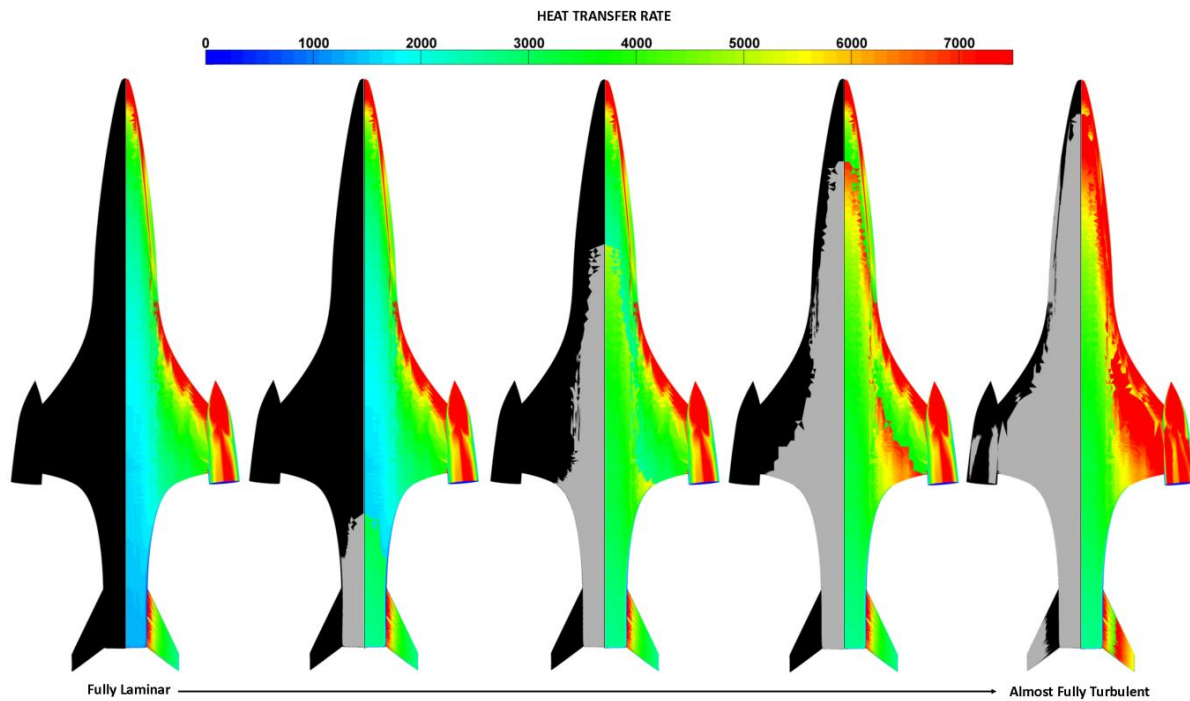


**Figure 9. The effect of transition Reynolds number on the extent of laminar (black) and turbulent flow (grey) on the underside of Skylon and the subsequent effect on heat transfer rate as predicted using HyFlow.**

**Figure 10. The effect of transition Reynolds number on the extent of laminar (black) and turbulent flow (grey) on the underside of CFASTT-1 and the subsequent effect on heat transfer rate as predicted using HyFlow.**

Figure 9 shows the effect of uncertainty in the effective Reynolds number on the predicted extent of turbulent flow on Skylon, and thus on the distribution of heating rate over the surface of the vehicle. The left half-figures show the division of the lower surface of the vehicle into regions of laminar (black) and turbulent (grey) flow for various feasible values of the effective Reynolds number, ranging from that yielding fully laminar flow over the surface to that yielding an almost completely turbulent flow.

The right half-figures then portray the effect of the existence of the regions of turbulent flow on the predicted heating rate on the surface of Skylon. Transition is seen to occur first near the trailing edges of the wings. As the effective Reynolds number is increased, these regions of turbulent flow expand to cover most of the under-surface of the wing. As the effective Reynolds number is increased further, a threshold is reached where transition is suddenly encountered on the fuselage, and this quickly results in large areas of the bottom and sides of the fuselage being exposed to a fully turbulent flow. Where the Reynolds analogy is appropriate, the existence of turbulent flow is generally associated with an increased heat flux to the surface of the vehicle. This is clearly evident in Fig. 9 where it can be seen how the wing, and eventually the fuselage of Skylon, is exposed to an increasingly intense heating regime as the effective Reynolds number is increased. The implications for the conservatism that needs to be incorporated into the design of the thermal protection system for a vehicle with this configuration, given that the TPS might be required to accommodate any uncertainty in the predicted extent of turbulent flow on the surface of the vehicle, are thus quite striking. This is particularly the case given that this particular geometry contains no inherent features that limit or contain the growth and spread of regions of turbulent flow across its surface.

Figure 10 shows similar results for CFASTT-1. Given the chordwise pattern of streamlines over the under-surface of the wing, turbulent flow and the associated elevation in local heating rate is delayed on these parts of the vehicle, at least compared to Skylon, until a much higher effective Reynolds number is reached. In addition, those regions of the fuselage that are subjected to the greatest increase in heating rate as a result of the onset of turbulent flow are confined to those parts of the geometry where the heating rates are already relatively benign thanks to the particular streamline pattern that is produced on this configuration. The implications for the design of the thermal protection system for this vehicle have yet to be explored fully, but the flat under-surface

together with the strakes on the forward fuselage would appear to produce a design that is relatively more robust than Skylon to any uncertainties in the extent of turbulent flow over the underside of the vehicle.

## V. Limitations of the Approach

As remarked earlier, the 114km altitude of the flight condition that forms the basis for all the analyses presented above lies at the lower boundary of the regime in which the rarefied nature of the ambient gas, and thus the possibility of non-equilibrium within the flow over the vehicle, needs to be properly accounted for. Indeed, even though the conditions of the calculations presented here might be judged according to a global measure of the Knudsen number to still lie within the continuum regime, localized regions of non-equilibrium may still be encountered within the flow, for instance in the neighbourhood of shock interactions and near to those parts of the surface where the local curvature is large. One method of characterizing the importance

of non-equilibrium effects within the flow is to measure the slip velocity at the surface of the vehicle. In a fully continuum flow, the slip velocity is zero as colliding gas particles accommodate their momentum fully to the presence of the surface of the vehicle. As the gas becomes more rarefied, an ever smaller proportion of the tangential momentum of the gas is transferred to the surface of the vehicle during collisions and the slip velocity becomes finite. Of particular relevance to the present study is that, where the slip velocity is non-zero, the Reynolds analogy that is currently used within HyFlow to calculate the heat transfer to the surface of the vehicle can be called into question.

Several methods for accounting for the onset of non-continuum effects in reduced order models for the aero-thermodynamics of aerospace vehicles have been proposed, the most interesting of which being those that rely on an algebraic bridging function to merge local surface properties from those encountered in the continuum regime to those that
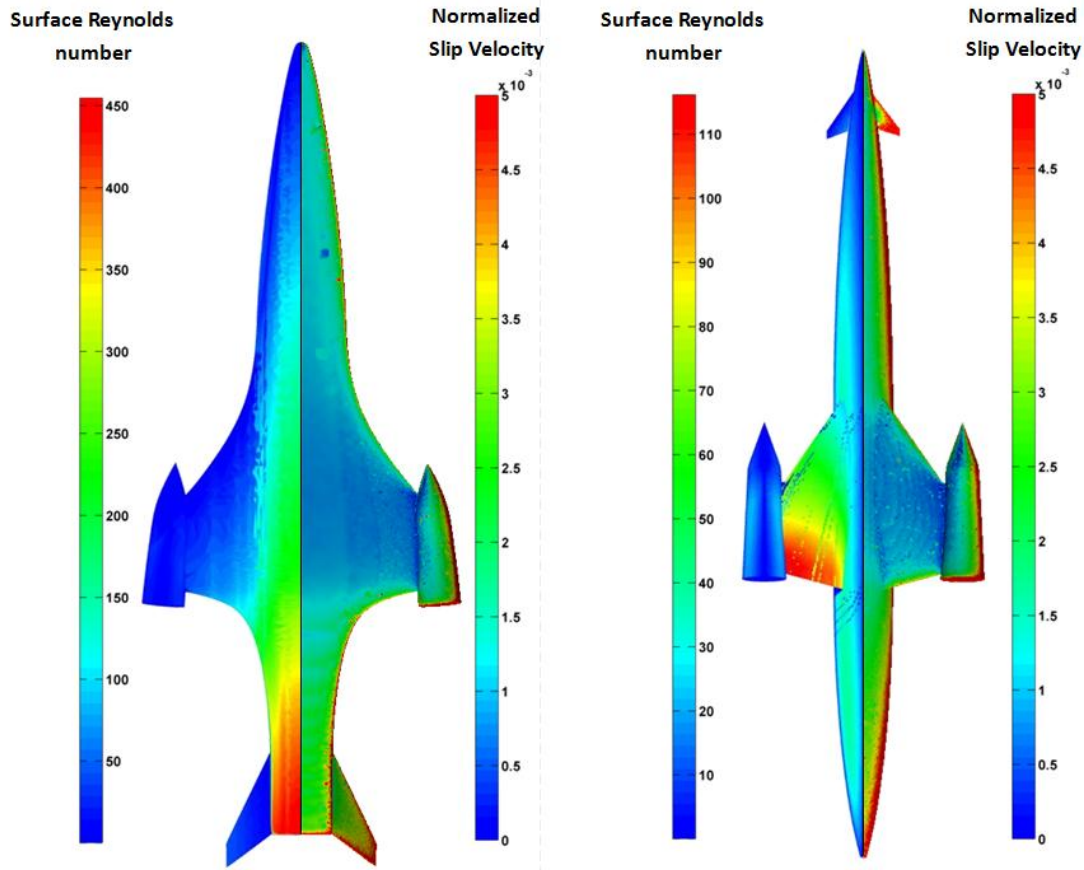


**Figure 11. Reynolds Number as predicted by HyFlow, compared to the normalized slip velocity as predicted by dsmcFoam, on the lower surface of CFASTT-1 (left) and Skylon (right).**

would be experienced were the vehicle be subjected to free molecular flow, in other words to a regime in which the oncoming flow is so dilute that collisions within the gas can be ignored. These methods have been applied very successfully to analyse the aero-thermodynamics of simple shapes (for example to blunt re-entry vehicles) but the ability of the approach to handle more complex vehicle geometries remains to be tested.

These approaches traditionally rely on the local Reynolds number being a suitable gauge of the local importance of rarefied flow effects on the surface of the vehicle[13,14]. In this vein, Fig. 11 compares the distribution of slip velocity, as calculated using dsmcFoam, to the distribution of local Reynolds number, as calculated using HyFlow, on the lower surfaces of the Skylon and CFASTT-1 vehicles. Although under the flight conditions presented here the slip velocity over the underside of both vehicles is relatively small, both cases show rather mixed correlation between the slip velocity and the local Reynolds number. For instance the large Reynolds numbers under the tail of CFASTT-1 do seem to correlate well with an increase in slip velocity there, and likewise the high Reynolds numbers on the aft surface of the Skylon wing do seem to be associated with elevated slip velocities along its trailing edge, but the extent of the correlation in these two rather selective instances needs to be treated rather circumspectly, it seems, given that, for instance, the rapid increase in the slip velocity around the sides of the forward fuselage of Skylon does not appear to be associated with particularly high local Reynolds numbers. The results presented here suggest that some significant work may still need to be done before the techniques described here can be used reliably to extend the low-order modeling of the aero-thermodynamic characteristics of vehicles with anything other than rather simple configuration into the rarefied flow regime.

## VI. Conclusions

Re-usable single stage to orbit launch vehicles promise to reduce the cost of access to space, but their success will be particularly reliant on accurate and robust modelling of their aero-thermodynamic characteristics. For preliminary design and and optimization studies, relatively simple numerical prediction techniques must often be used in place of more sophisticated, accurate but time-consuming analysis tools. It is important though that these simplified techniques be capable of reproducing the behaviour of the models for which they are intended to be a surrogate, at least to sufficient fidelity for

them to lead the design process into the same basin of attraction of the configuration space of the system. It is vital thus that the uncertainty that is inherent in the predictions of these simplified models be well understood before they are used in practical applications.

In this paper, a new reduced-order model, called HyFlow, that is based on the Newtonian flow assumption and the Reynolds analogy for surface heating, is described. Its predictions of surface pressure and heat transfer are compared against those of a more physically-sophisticated method based on the Direct Simulation Monte Carlo technique in order to determine the ability of the model to capture the aero-thermodynamics of vehicles with the inherently very complex configuration that will be characteristic of future SSTO vehicles.

Attention is focused on the lower boundary of the high-altitude regime where lifting re-usable SSTO configurations will experience their greatest thermal load during re-entry, but where non-continuum effects within the gas of the atmosphere start to become important. The flight condition of the case analysed in this paper, although on the borderline of applicability of the assumption of continuum flow over the vehicle surface, is such that the flow is subject to only a small degree of slip on parts of the underside of the vehicle and hence can be treated to all practical intent as a continuum. Extension of the formalism presented here further into the transition regime is an important goal of the research that underpins this paper, but preliminary results presented here suggest that currently-accepted approaches, at least those which attempt to correlate the onset of non-continuum with the local Reynolds number on the surface of the vehicle, may not be entirely appropriate to the class of vehicles that are of interest in this study. The development of bridging techniques that are appropriate to the inherently very complex geometry of currently proposed SSTO vehicles is indeed likely to be a very fruitful area of future research.

Despite these reservations, predictions using the HyFlow reduced-order model presented in this paper of the pressure distribution on the lower surface of the two different vehicles that were analysed in the present study match very well those of the DSMC approach, and the principal discrepancies seem to be confined to those regions where the geometry of the vehicle, for example in the wing-nacelle junction, may lead to particle reflections from the surface that enhance the importance of the off-body molecular

collisions that are not captured by the modified Newtonian approach. Although the case analysed here was not particularly affected by the presence of shock interactions within the flow field, it is expected that some work will also have to be done in future to address any effects on the pressure distribution on the surface of the vehicle that result from such interactions, given that their origin is also in the flow away from the vehicle surface.

The predictions of surface heating using the approach adopted in this paper are reasonably satisfactory, especially in terms of the location and extent of the regions of maximal heating on the surface of the vehicle, but the presence of a number of essentially arbitrary correction factors within the formalism leads to an approach which is not entirely robust, particularly given that the most suitable choice of these factors seems to be rather dependent on the details of the geometry of the vehicle. When these factors can be set by reference to the predictions of the more sophisticated analysis, however, the reduced order method is capable of reproducing the results of the more complex method, and thus of acting as its surrogate, with quite surprising fidelity.

Within the context of multi-disciplinary design and optimization, the speed of computational execution of the method is paramount. The predictions of the reduced-order model appear to be very robust to large reductions in the resolution of the geometry of the surface of the vehicle, and by reducing the number of panels into which the surface of the vehicle is discretized, several orders of magnitude improvement in execution time can be achieved with minimal degradation in the fidelity of the method. It appears too that judicious choice of the way in which the geometry of the surface of the vehicle is tesellated can sustain the accuracy of the method even when the discretization of the surface of the vehicle is made coarse enough for the method to be practical in the context of design or trajectory optimization studies.

Finally, the power of reduced order modelling in being able to perform fast studies of the effect of varying certain key parameters on the aero-thermodynamic behaviour of the system is illustrated by comparing the effect of changing the effective Reynolds number of the flow on the distribution of heating on the bottom surface of the two different vehicles that were analysed in this study. The results suggest the advantages of certain geometric features of the vehicle as far as maximal heating is concerned, but it should be borne in mind that these inferences, no matter how suggestive or interesting, are of course subject to confirmation using more sophisticated analytical techniques.

Taken as a whole, the results presented in this work suggest thus that, if used with care, reduced-order models such as those described in this paper can be used very effectively in the design and optimization of the next generation of future re-usable SSTO vehicles, even those with very complex configuration as currently proposed, as long as their predictions are adequately supported by the use of more sophisticated computational techniques both to confirm their validity, as well as, within a few limited instances, to set appropriate numerical values for some of the essential parameters that lie at the heart of their formulation.

## Acknowledgements

## References

[1]Varvill, R., and Bond, A., "A Comparison of Propulsion Concepts for SSTO Reuseable Launchers," *Journal of the British Interplanetary Society*, Vol. 56, pp. 108-117, 2003.

[2]Heppenheimer, T.A., *Development of the Space Shuttle 1972-1981: History of the Space Shuttle Volume 2*, Smithsonian Books, 480 pp., 2002.

[3]Anderson, J.D., *Hypersonic and High-Temperature Gas Dynamics*, McGraw-Hill Series in Aeronautical and Aerospace Engineering, 672 pp., 1988.

[4]Gentry, A.E., Smyth, D.N., Oliver, W.R., *The Mark IV Supersonic-Hypersonic Arbitrary-Body Program. Volume I. User's Manual*, Douglas Aircraft Corporation, Long Beach, CA., 269 pp., 1973.

[5]Kinney, D.J., "Aero-Thermodynamics for Conceptual Design," AIAA-2004-31, *42nd AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada, January 2004.

[6]Hirschel, E.H., and Weiland, C., *Selected Aerothermodynamic Design Problems of Hypersonic Flight Vehicles*, Springer, 518 pp., 2009.

[7]Pescetelli, F., Minisci, E., Maddock, C., Taylor, I., Brown, R.E., "Ascent Trajectory Optimisation for a Single-Stage-To-Orbit Vehicle with Hybrid Propulsion," *18th AIAA/3AF International Space*

*Planes and Hypersonic Systems and Technology Conference*, Tours, France, September 2012.

[8]Varvill, R., and Bond, A., "The Skylon Spaceplane: Progress to Realisation," *Journal of the British Interplanetary Society*, Vol. 61, pp. 412-418, 2008

[9]Lees, L., "Hypersonic Flow," *Fifth International Aeronautical Conference*, Los Angeles, CA, 1955.

[10]Kinney, D.J., "Aerothermal Anchoring of CBAERO Using High Fidelity CFD", *45th AIAA Aerospace Sciences Meeting*, Reno, NV, January 2007.

[11]Meador, W.E., Smart, M.K, "Reference Enthalpy Method Developed from Solutions of the Boundary-Layer Equations," *AIAA Journal* Vol. 43, No. 1, pp. 135-139, January 2005.

[12]Chrusciel, G.T., "Rarefied Flow Heat Transfer Model for Slender Bodies at Large Angle of Attack," *23rd International Symposium on Rarefied Gas Dynamics*, Whistler, BC, July 2002.

[13]Vashchenkov, P., Ivanov, M., and Krylov, A., "Numerical Simulations of High-Altitude Aerothermodynamics of a Promising Spacecraft Model," *27th International Symposium on rarefied Gas Dynamics*, Pacific Grove, CA, July 2010.

[14]Scanlon, T.J., Roohi, E., White, C., Darbandi, M. and Reese, J.M., "An open-source, parallel, DSMC code for rarefied gas flows in arbitrary geometries," *Computers and Fluids*, Vol. 39, pp. 2078-2089, 2010.

[15]Bird, G.A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Clarendon Press, 484 pp., 1994.

[16]Bird, G.A., "Chemical Reactions in DSMC", *27th International Symposium on Rarefied Gas Dynamics*, Pacific Grove, CA, July 2010.

[17]Wilmoth, R.G., Blanchard, R.C., Moss, J.N., "Rarefied Transitional Bridging of Blunt Body Aerodynamics," *21st International Symposium on Rarefied Gas Dynamics*, Marseille, France, July 1998.

[18]Ahmad, A., Maddock, C., Scanlon, T., and Brown, R., "Prediction of the Aerodynamic Performance of Re-Usable Single-Stage to Orbit Vehicles," *Proceedings of the First International Space Access Conference*, Paris, France, September 2011.