Noname manuscript No.
(will be inserted by the editor)

Design optimisation using computational fluid dynamics applied to a land–based supersonic vehicle, the BLOODHOUND SSC

B. Evans · T. Morton · L. Sheridan · O. Hassan · K. Morgan · J.W. Jones · M. Chapman · R. Ayers · I. Niven

Received: date / Accepted: date

Abstract This paper details the computational design optimisation strategy employed to achieve an engineering solution to the problem of excessive supersonic lift at the rear of the BLOODHOUND SSC (SuperSonic Car) during its design. The optimisation problem is described first, followed by details of the computational fluid dynamics procedure employed to study the aerodynamic performance of the vehicle and the design optimisation strategy utilising Design of Experiments. The 'optimised' design resulting from this study is presented in the final section and contrasted with the 'unoptimised' baseline geometry. The final vehicle geometry presented in this paper is, at the time of writing, being built and is due to be tested in 2013 in an attempt to increase the World Land Speed Record from 763 mph to 1,000 mph.

Keywords CFD · Design of Experiments · design optimisation · BLOODHOUND SSC

The authors gratefully acknowledge the financial support provided for this work by the UK Engineering and Physical Sciences Research Council in the form of research grant EP/F032617.

B. Evans, O. Hassan, K. Morgan, J. Jones
College of Engineering
Swansea University
E-mail: b.j.evans@swansea.ac.uk

T. Morton, L. Sheridan
MathWorks

M. Chapman, R. Ayers, I. Niven
BLOODHOUND Ltd.

1 Introduction

1.1 The World Land Speed Record

The World Land Speed Record was first set at a modest 39 mph at Achères in France in 1898. The Record has been broken around sixty times and the current (supersonic) Record of 763.035 mph was set by Andy Green in THRUST SSC (SuperSonic Car) at the Black Rock Desert, Nevada in 1997. A significant feature of the THRUST SSC project was the manner in which computational fluid dynamics (CFD) was applied to guide the aerodynamic design process [1–3].

The BLOODHOUND project was officially launched in October 2008 with a primary engineering objective of designing, building and running a car to achieve a new World Land Speed Record of 1,000 mph. This engineering objective is coupled with an educational objective to promote science, technology, engineering and mathematics to schoolchildren in the UK via the BLOODHOUND education programme. Once again, CFD has been chosen as the primary tool to guide the aerodynamic design of the vehicle. Efficient post–processing of CFD data has proved to be invaluable in allowing the design cycle to progress as effectively as possible. Post-processing of raw CFD outputs has allowed the team of engineers involved to interrogate the aerodynamic behaviour of each configuration of the vehicle as its design has evolved. This paper outlines the methodology and results of combining CFD simulations with design optimisation to solve a specific aerodynamic problem relating to the design of BLOODHOUND SSC: defining the rear wheel, rear suspension and rear car body geometry as shown in Figure 1.

The rules and regulations associated with the World Land Speed Record are provided by the Federation In-

ternationale de l'Automobile [4]. These include the engineering restrictions placed upon the vehicle itself and the details of the timing procedures for speed measurement.

Fig. 1 Rear vehicle geometry of BLOODHOUND SSC

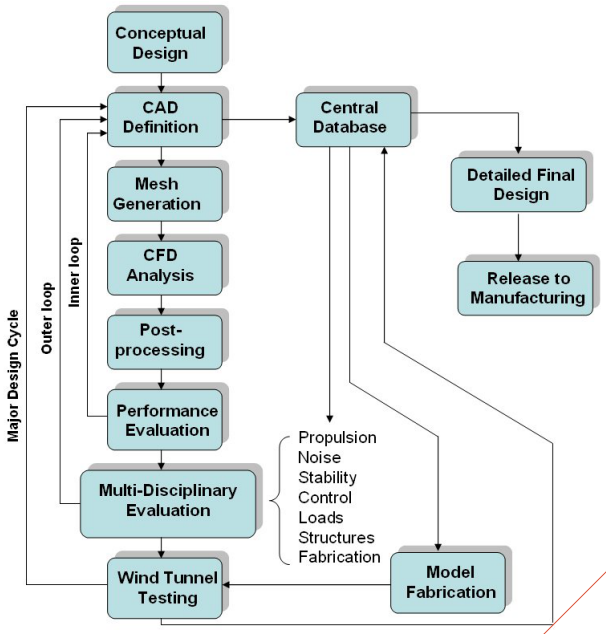

Fig. 2 A typical multi–disciplinary aerospace design cycle

1.2 Background to Computational Aerodynamic Optimisation

Over the past three decades, computational fluid dynamics (CFD) has revolutionised the way in which the aerospace industry tackles problems of aerodynamic design. In particular, unstructured mesh methods [5–7] now allow grids on complex three–dimensional geometries to be generated in a matter of hours that might once have taken several months using multiblock techniques for quasi–structured meshes [8,9]. In light of this, CFD has become an integral part of the typical aerodynamic design cycle, as indicated in Figure 2 [10]. This flow chart indicates the emphasis now placed on CFD within the inner and outer design loops. On the BLOODHOUND project, for practicality, financial restrictions and time constraints, the *Major Design Cycle* loop is essentially the actual vehicle runway and desert testing. This is where validation of the CFD modelling used in the inner loops will take place. It may well be the case that vehicle testing will require significant redesign and entering back into the inner design loops (note that, at the time of writing the BLOODHOUND is in the build phase with anticipated build completion December 2012 followed by vehicle testing in 2013).

It is only relatively recently, i.e., the last 10 years, that significant effort has been invested into research concerning coupling CFD modelling with computational optimisation [12] that has led to practical applications of industrial interest. Aerodynamic optimisation problems are computationally expensive and often highly ill–conditioned. The computational difficulty of solving aerodynamic optimisation problems requires significant research investment in this field. Jameson introduced the adjoint method to computational aerodynamic design by applying control theory to fluid mechanics which has successfully been applied to the design of a concept race plane for the annual Reno air races and the optimisation of the British–Aerospace MDO datum wing [13]. The adjoint methodology is typically the preferred approach for design problems with a high number of design variables. However, the implementation of the adjoint method for the design of complex configurations is still not fully matured. As this study was for a time–critical aspect of the design of the BLOODHOUND SSC, the authors agreed that a more trusted and dependable optimisation approach in which experience was available would be adopted. This led the authors to select a Design of Experiments (DoE) approach to the solution of the problem.

In general, the variety and complexity of aerodynamic optimisation problems facing the aerospace industry at present requires the tools of multi–disciplinary design optimisation. Designers must trade analysis fidelity with analysis time. The tools of design optimisation can save time at several stages: design of experiment tools can predict the most efficient points to

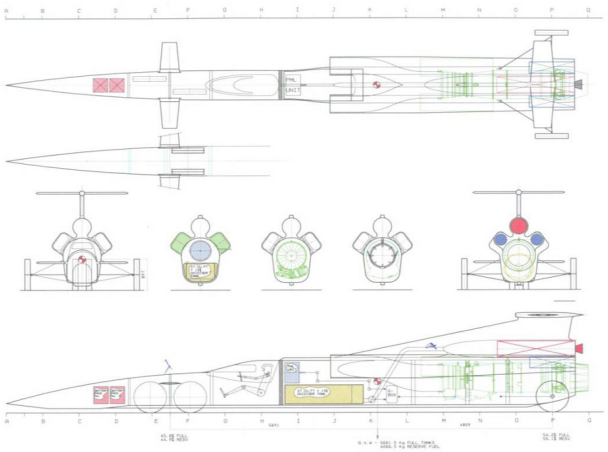

Fig. 3 Drawing illustrating the initial concept layout of BLOODHOUND SSC

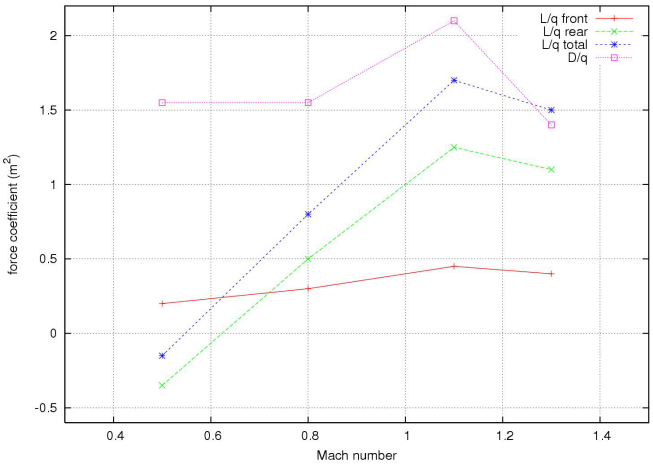

Fig. 5 Lift and drag coefficients against Mach number for 'pre–optimised' config9

collect data with a CFD model; reduced–order statistical models of the CFD model results can be evaluated much faster than the full CFD model, and a wide range of optimisation algorithms can find optimal configurations using those statistical models. The appropriate use of these computational tools has a positive impact on the man hours required to complete loops in the design cycle shown in Figure 2, particularly the inner loop. As an example of the impact of these methods in industry, most of the external geometry on the Boeing 787 airplane was designed using aerodynamic optimisation tools [11].

Finally, it is also worth noting that multi–disciplinary optimisation is now becoming a goal of the aerospace industry [14]. This field is very much in its infancy. Only simple coupled aerodynamic, structural and electomagnetic problems have been tackled but, in principle, these techniques have much potential for impacting the outer design loop depicted in Figure 2 [10].

2 The Optimisation Problem

The initial concept for the 1,000 mph BLOODHOUND vehicle is shown in Figure 3. With reference to Figure 2, this is the conceptual design providing the starting point for the design process.

The design iterations taking the external shape of the vehicle from the conceptual design through to a realisable engineering solution is shown in Figure 4. Each configuration ('config') is a product of the outer loop of the design cycle in Figure 2. The external vehicle shape iterations, feeding into the wider engineering design cycle, were driven primarily by CFD simulations of the full vehicle. Typically, each full aerodynamic vehicle design iteration, i.e., the inner loop in Figure 2 required 2 months for meshing, and sufficient simulation and post–processing to guide the designers on to the subsequent configuration [15]. This is due to the large number of steady state solutions required for each new configuration over a range of Mach numbers at design and off–design conditions (e.g. in yaw / crosswind). A more thorough description of the CFD approach will be provided in section 3.

The final aerodynamic challenge to be solved at the 'config10' stage was to minimise significant supersonic lift at the rear of the vehicle whilst also reducing the overall vehicle supersonic drag. The target supersonic drag coefficient (D/q) to be achieved was 1.4 m^2 although, obviously, reduction below this value would have been beneficial. Note that both lift and drag coefficients referred to in this paper are simply total lift or drag divided by dynamic pressure, q (in Pa), and, therefore, have units m^2. The lift and drag coefficients as a function of Mach (M) number of the 'pre–optimised' config9 vehicle geometry are shown in Figure 5. The overall lift coefficient is also shown in terms of its distribution between the front wheels (L/q front) and the rear wheels (L/q rear). It is clear that the majority of the contribution the vehicle's increase in lift at supersonic speed is driven by aerodynamic forces at the rear of the vehicle. Also, the drag variation indicates that the peak vehicle drag is significantly higher than the 1.4 m^2 supersonic performance target.

The config9 rear geometry and $M_\infty = 1.3$ flow field are shown in Figure 6. These plots indicate that the source of the rear supersonic lift is the strong shock wave generated by the rear wheels and suspension structure and, in particular, the interaction of this shock wave with the main vehicle body. The positions of the upstream bow shock generated by each of the rear wheels

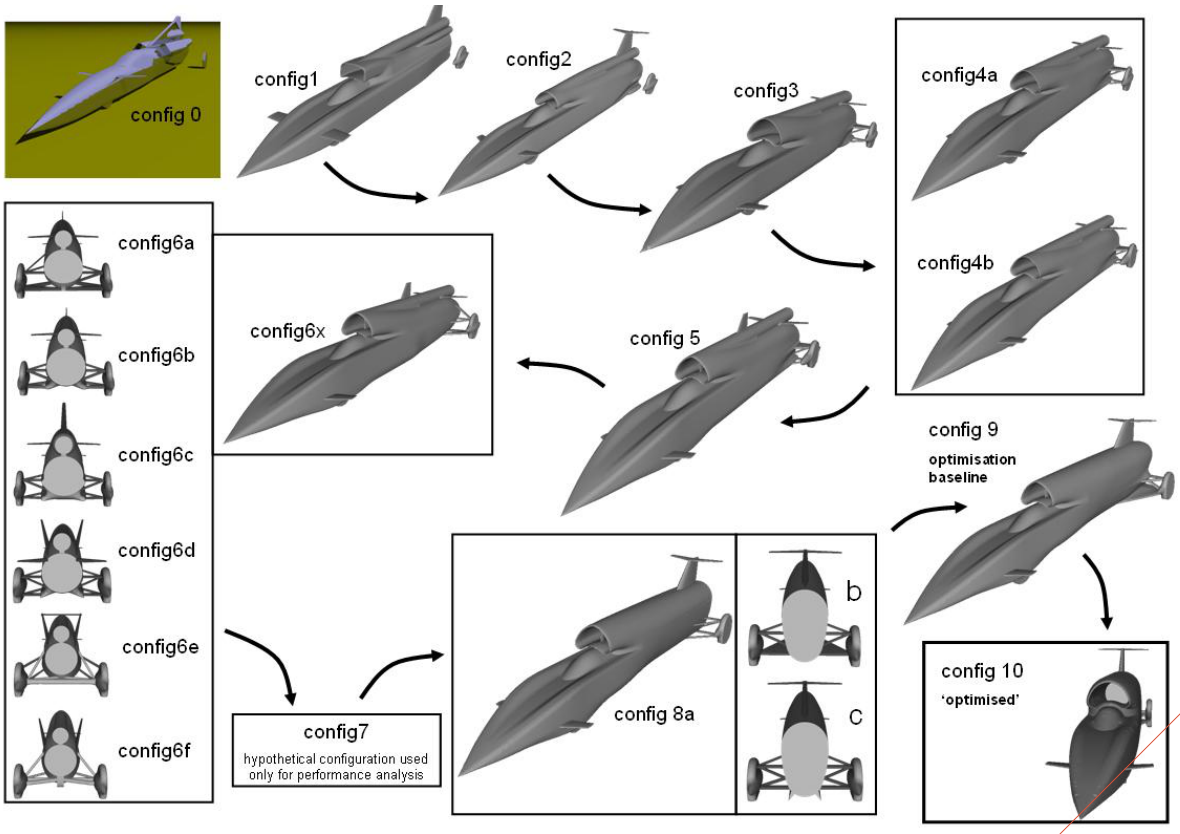

Fig. 4 BLOODHOUND design iterations from 'config 0' (2007) to 'config10' (2010)

is evident on either side of the vehicle in Figure 6. Aft of the shock wave is a large region of subsonic, and high pressure, flow. This flow then accelerates back to supersonic speed with a sonic condition at the 'pinch point' in the geometry between the car and rear wheel fairing. This over–accelerated flow is then decelerated back to the freestream speed by a terminal shock aft of the car as the flow merges with the rest of the wake. The effects of this complex flowfield on the underside of the car is evident in Figure 7 which shows the variation in the config9 underside pressure coefficient distribution between $M_\infty = 0.5$ and $M_\infty = 1.3$. The transition of the rear portion of the vehicle underside from behaving as a traditional 'downforce generating' diffuser to a supersonic lifting device is evident.

3 Computational Fluid Dynamics Approach

An overview of the Computational Fluid Dynamics (CFD) process from geometry specification through to post–processing and data analysis is given in Figure 8. The CFD procedure detailed in this section is a development of the scheme detailed in [7] specifically for the particularly demanding aerodynamics associated with a supersonic vehicle, with rotating wheels travelling across a moving ground plane. This complexity results in inherently non–linear flow features such as shockwave–bounday layer interactions and jet–wake interactions that are difficult to capture accurately and efficiently utilising a stable CFD scheme. The novel application also required some unique boundary condition applications such as supersonic moving surfaces (wheels /ground-plane).

3.1 Governing Equations

For the complex, turbulent flow throughout the relevant Mach number range for BLOODHOUND SSC ($M_\infty < 1.4$) , the appropriate governing equations are the compressible Navier–Stokes equations. Relative to a cartesian $Ox_1x_2x_3$ coordinate system, the steady state compressible Favre–averaged Navier–Stokes equations [16] are expressed in the integral form as,

$$\int_\Gamma \mathbf{F}^j(\mathbf{U}) n_j d\Gamma = \int_\Gamma \mathbf{G}^j(\mathbf{U}) n_j d\Gamma \qquad j = 1, 2, 3 (1)$$

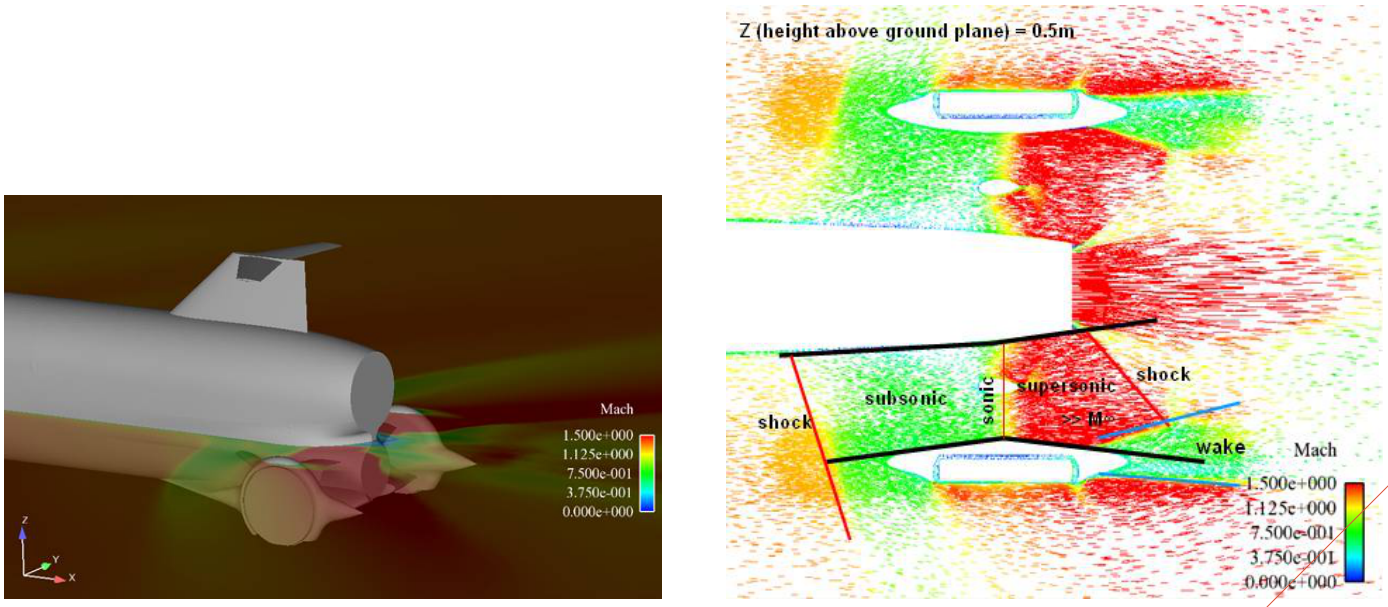

(a) (b)

Fig. 6 The rear supersonic shock system at M_∞, showing :(a) config9 rear geometry and Mach distribution on a plane at z=0.9m; (b) velocity vector field on a horizontal plane at z=0.5m, where z is the height above the ground plane

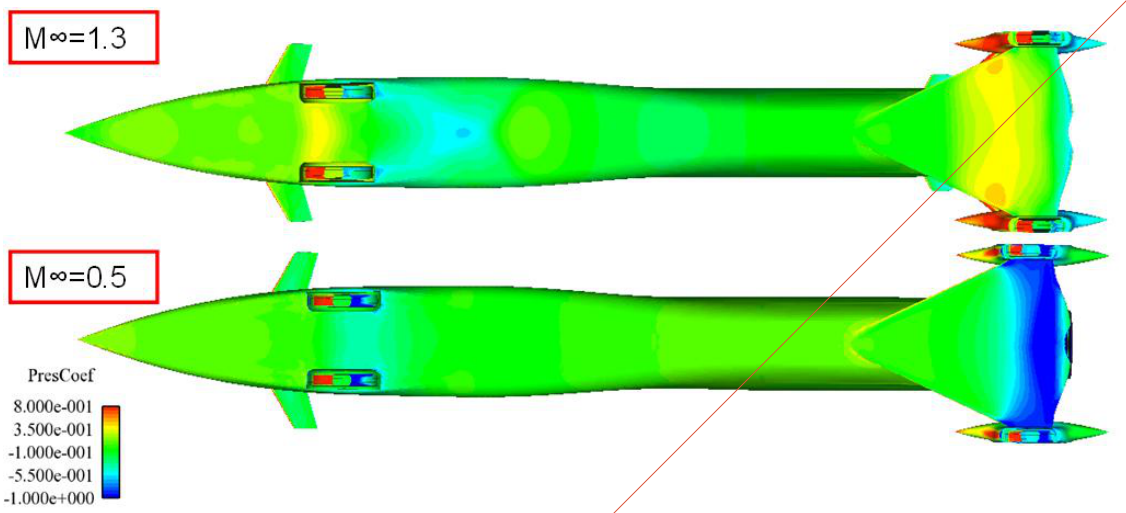

Fig. 7 Underside pressure coefficient distribution at $M_\infty = 0.5$ and $M_\infty = 1.3$ for config9

where Γ denotes the closed surface bounding a three dimensional domain Ω and the summation convention has been adopted. The unknown vector $\mathbf{U}$ is defined by

$$\mathbf{U} = \begin{bmatrix} \varrho \\ \varrho u_1 \\ \varrho u_2 \\ \varrho u_3 \\ \varrho \epsilon \end{bmatrix} \tag{2}$$

and the inviscid and viscous flux vectors are given as

$$\mathbf{F}^j = \begin{bmatrix} \varrho u_j \\ \varrho u_1 u_j + p\delta_{1j} \\ \varrho u_2 u_j + p\delta_{2j} \\ \varrho u_3 u_j + p\delta_{3j} \\ u_j\left(\varrho\epsilon + p\right) \end{bmatrix} \qquad \mathbf{G}^j = \begin{bmatrix} 0 \\ \tau_{1j} \\ \tau_{2j} \\ \tau_{3j} \\ u_k\tau_{kj} - q_j \end{bmatrix} \tag{3}$$

respectively. The unit outward normal vector to Γ is $\mathbf{n} = (n_1, n_2, n_3)$ and δ_{kj} denotes the Kronecker delta. In these equations, ϱ is the averaged fluid density, u_i is the component of the averaged fluid velocity in the direction x_i, p is the averaged fluid pressure and ϵ is the

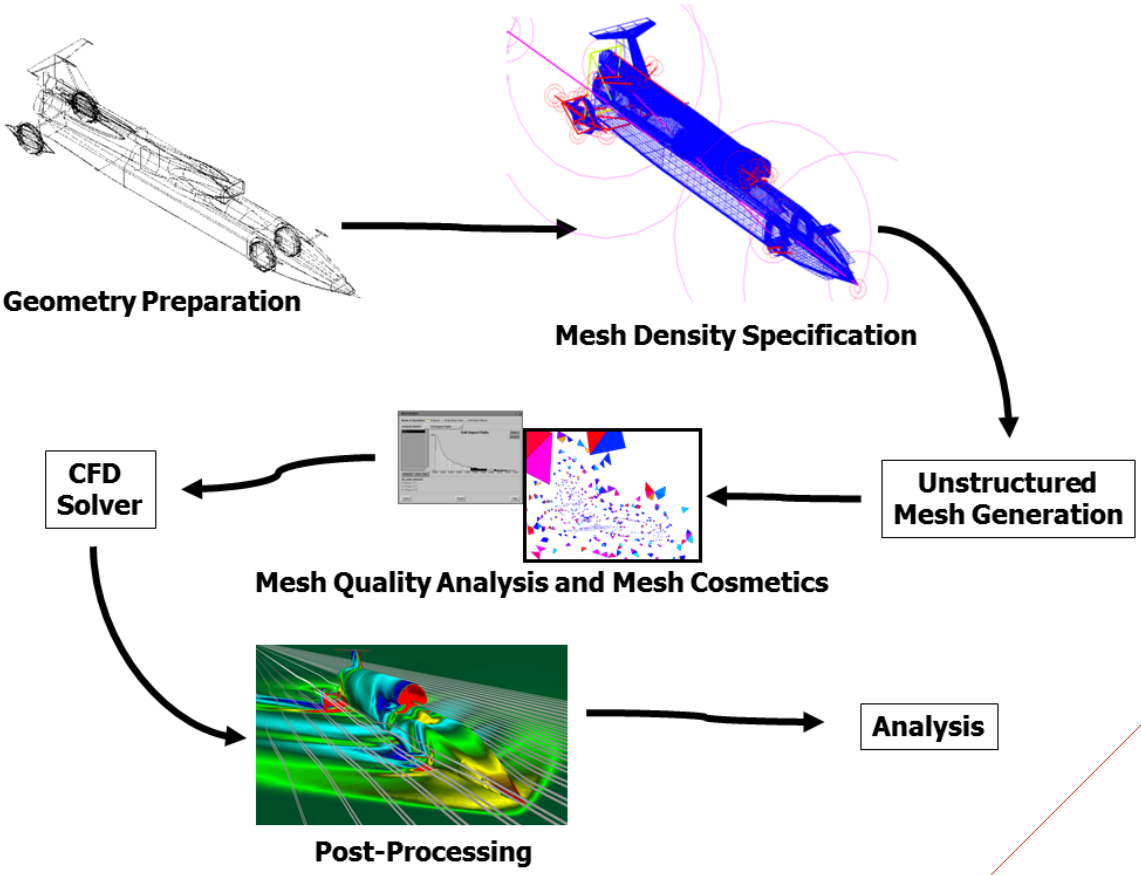

Fig. 8 Overview of the CFD process

averaged specific total energy of the fluid. The averaged deviatoric stress tensor is defined by

$$\tau_{ij} = -\frac{2}{3}\mu\frac{\partial u_k}{\partial x_k}\delta_{ij} + \mu\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right) \tag{4}$$

and the averaged heat flux is

$$q_j = -k\frac{\partial T}{\partial x_j} \tag{5}$$

Here, μ denotes the sum of the laminar and turbulent viscosities, k is the sum of the laminar and turbulent thermal conductivities and T is the averaged absolute temperature. The air is assumed to be calorically perfect and the averaged state equations

$$p = \varrho \mathrm{R}T \qquad \epsilon = c_v T + \frac{1}{2}u_k u_k \tag{6}$$

are employed, where R is the gas constant, $c_v = c_p - \mathrm{R}$ is the specific heat at constant volume, c_p is the specific heat at constant pressure and $c_p/c_v = 1.4$. The laminar and turbulent Prandtl numbers are assumed to be constant. The laminar viscosity varies with temperature according to Sutherland's law [17]. For the simulations to be shown here, the variation of the kinematic turbulent viscosity is obtained from the one–equation turbulence model of Spalart and Allmaras [18].

3.2 Domain Discretisation

For any practical aerodynamic simulation, a CAD system is normally employed to provide the definition of the boundary geometry. This is the 'CAD Definition' stage with reference to Figure 2. The geometry definition obtained in this way frequently requires treatment in order to produce a watertight surface description that is suitable for the processes of surface and volume mesh generation. This treatment process will involve the merging of surfaces and intersection points and the smoothing of underlying surface definitions. In the work presented here, this process was accomplished by using the CADfix [19] commercial software package.

Before the start of the discretisation process, the desired element size distribution is defined by the user in terms of a mesh control function [20,21]. The first step in the discretisation process is the triangulation of the computational boundaries and this is accomplished using an advancing front approach [22]. The advancing layers method [5] is used next to generate stretched tetrahedral elements adjacent to the boundary surface components that represent solid walls. The height of each layer, and the number of layers, is specified by the

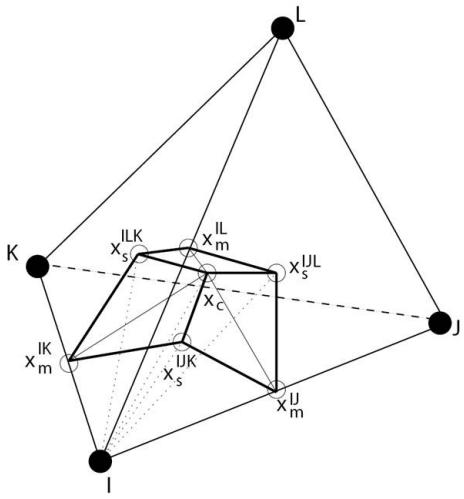

Fig. 9 Illustration of the portion of the element Ω_I of the dual mesh, associated with node I, that is contained inside a tetrahedral element

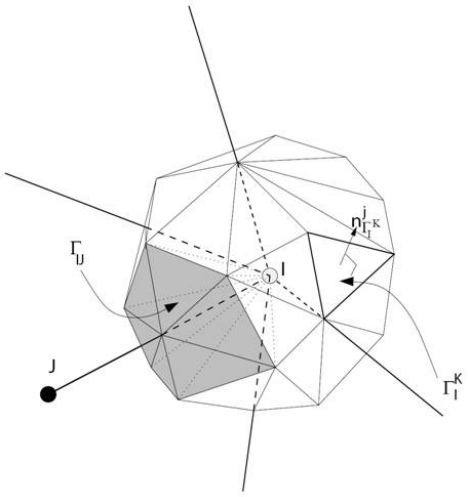

Fig. 10 The volume Ω_I of the dual mesh surrounding an interior node I

user in an attempt to ensure that the expected boundary layer profile can be adequately represented. Following the use of the advancing layers method, the remainder of the computational domain is discretised using a standard Delaunay isotropic tetrahedral mesh generation procedure [6]. As a final step, appropriate elements of the tetrahedral mesh generated by the advancing layers method are merged to produce a consistent hybrid mesh of tetrahedra, pyramids and prisms [23]. In this study a typical mesh for the full car CFD model was of the order 100 million hybrid elements with approximately 25 layers in the boundary layer portion of the mesh. For the parametric model, this was reduced to approximately 15 million hybrid elements and 20 layers in the boundary layer mesh to achieve, what was deemed to be, sufficient resolution.

The median dual [24,5] is employed within that portion of the mesh consisting of isotropic tetrahedral elements. This dual is constructed by connecting edge midpoints, element centroids and face centroids in the domain mesh, in such a way that only one node is present in each dual mesh element. With this strategy, each node I of the domain mesh is associated with a volume Ω_I of the dual mesh. The boundary surface of the volume Ω_I is denoted by $\partial\Omega_I$. Each edge of the domain mesh is associated with a segment of the dual mesh interface between the nodes connected to the edge. This segment is a surface constructed from triangular facets, where each facet is connected to the midpoint of the edge, a neighboring element centroid and the centroid of an element face connected to the edge, as illustrated in Figure 9. The midpoint of the edge between node I and J is denoted by $\mathbf{x}_m^{IJ}$, the centroid of the face with vertices I, J and K is denoted by $\mathbf{x}_s^{IJK}$ and the element centroid is designated by $\mathbf{x}_c$. The bold lines on the dual mesh in this figure illustrate the boundaries between the edges with which the dual mesh segment is associated. With this dual mesh definition, the volume Ω_I can be viewed as being constructed in terms of a set of tetrahedra, as illustrated for a typical interior node I in Figure 10. The surface of the dual mesh cell surrounding node I is defined in terms of the closed set of planar triangular facets Γ_I^K, where each facet only touches a single edge of the domain mesh. The set of facets touching the edge between nodes I and J is denoted by Γ_{IJ}.

In general, the median dual approach cannot be used for the hybrid elements generated by the procedure described above, as the elements produced by merging tetrahedra in this way may be warped so severely that a vertex can lie outside the corresponding median dual cell. This may occur in regions of high geometry curvature or at the interface between the hybrid and isotropic meshes. To overcome this problem, the information contained in the original tetrahedral mesh is used to ensure that the topology of the control volume cells is valid [25].

3.3 Equation Discretisation

In the cell vertex finite volume method, equation (1) is applied to each cell of the dual mesh in turn. To enable the numerical integration of the inviscid fluxes, a set of coefficients is calculated for each edge by using the surface segments of the dual mesh cell that are associated

with the edge. For an internal edge connecting nodes I and J, these coefficients are written as

$$C_{IJ}^{j} = \sum_{K \in \Gamma_{IJ}} A_{\Gamma_I^K} n_{\Gamma_I^K}^{j} \tag{7}$$

for $j = 1, 2, 3$, where $A_{\Gamma_I^K}$ is the area of the facet Γ_I^K and $n_{\Gamma_I^K}^{j}$ is the component, in direction x_j, of the unit outward normal vector to the facet from the viewpoint of node I, as illustrated in Figure 10. Additional coefficients need to be introduced to enable integrals over the computational boundary to be approximated [23].

The integral of the inviscid fluxes over the surface of the dual mesh cell associated with an edge is approximated by assuming the flux to be constant over the surface and equal to its computed value at the midpoint of the edge. The surface integral of the inviscid flux over the complete dual mesh surface for a typical internal node I is then approximated as

$$\int_{\Gamma_I} \mathbf{F}^j n^j d\Gamma \approx \sum_{J \in \Lambda_I} \frac{C_{IJ}^{j}}{2} \left(\mathbf{F}_I^j + \mathbf{F}_J^j \right) \tag{8}$$

where Λ_I denotes the set of nodes connected to node I by an edge in the domain mesh.

The integral in equation (1) containing the viscous fluxes is approximated in a similar manner as

$$\int_{\Gamma_I} \mathbf{G}^j n^j d\Gamma \approx \sum_{J \in \Lambda_I} \frac{C_{IJ}^{j}}{2} \left(\mathbf{G}_I^j + \mathbf{G}_J^j \right) \tag{9}$$

However, gradients of the flow variables need to be determined before this formula can be employed. The evaluation of the gradient of a function may be performed in several ways within the finite volume framework. For edge based implementations, a common approach is to calculate the nodal values of the derivatives of the flow variables by using a finite volume method and then to use these values directly in equation (9). While this is convenient computationally, the result is a five point discretisation stencil, as compared to the traditional three point stencil of many finite element and finite difference schemes. The five point stencil effectively doubles the mesh spacing for the evaluation of the viscous terms and may allow checkerboarding of the solution and a consequent destabilization [26].

A different edge based approach can be developed which results in a more compact stencil [27]. The viscous terms contribute mainly in the boundary layers, which are characterized by high gradients normal to solid surfaces and, except for localized regions such as around shocks, relatively small gradients tangential to the surface. The introduction of the quasi regular meshes, with grid lines parallel and normal to the wall, ensures that the high gradients are captured by the compact stencil and that the five point stencil terms are marginalized. It follows that, with the form of meshes that are employed in the boundary layer regions, this treatment of the viscous terms reduces to the familiar three point scheme used for structured meshes.

3.4 Stabilisation

This discretisation procedure results in a stencil for the convective terms that is central difference in character, which means that the addition of stabilizing dissipation is necessary before the solution of practical flows may be attempted. This is achieved by replacing the physical inviscid flux function on the right hand side of equation (8) by a consistent numerical flux. Here, the HLLC numerical flux function [28] is adopted, which is a modification of the original HLL scheme [29]. The function is constructed by assuming that the solution to the Riemann problem, between states $\mathbf{U}_I$ and $\mathbf{U}_J$, is represented by a contact wave and two acoustic waves separating four constant states. The acoustic waves may be either shocks or expansion fans. The method employs an exact resolution of the Riemann problem, while averaging the wave speeds in an appropriate manner. The implementation employed involves an approximation to the acoustic waves that improves the transition from subsonic to supersonic speeds [28]. This results in a scheme that is appropriate for the range of Mach number encountered by the BLOODHOUND vehicle.

3.5 Boundary Conditions

To complete the specification of the problem, boundary conditions must be defined over the entire boundary of the computational domain. With regards to the simulations necessary to provide data for the optimisation problem detailed in this paper, this includes farfield boundaries, viscous walls, jet engine / rocket exhausts and symmetry planes. Note that the jet / rocket exhaust boundary condition was only used in full vehicle simulations, and was replaced by a solid 'sting' in the parametric geometry model. This will be discussed further in section 4.

3.5.1 Inflow and Outflow Boundaries

At basic inflow and outflow boundaries, a characteristic treatment [30] is employed to determine the number, and type, of conditions that require specification at any given point. At inflow for the Spalart–Allmaras turbulence model, the turbulent viscosity is arbitrarily set to a value equal to ten percent of the laminar viscosity value. This assumption is in line with CFD studies using

similar schemes in the literature [7–9, 15] The turbulent viscosity value at outflow is obtained by the extrapolation of values from the interior of the computational domain.

3.5.2 Viscous Walls

For viscous flow, the no slip condition

$$\mathbf{u} = \mathbf{u}^w \tag{10}$$

is strongly applied, where $\mathbf{u}^w$ is a specified wall velocity. When simulating a vehicle with rotating wheels, the vector $\mathbf{u}^w$, at a point with position vector $\mathbf{r}$ on the wheel relative to a user–specified position on the axis of rotation, is computed as

$$\mathbf{u}^w = \boldsymbol{\omega} \times \mathbf{d} \tag{11}$$

where, as illustrated in Figure 11(a),

$$\mathbf{d} = \mathbf{r} - \frac{(\boldsymbol{\omega}.\mathbf{r})}{|\boldsymbol{\omega}|}\boldsymbol{\omega} \tag{12}$$

is the position vector of the point relative to the centre of the wheel and $\boldsymbol{\omega}$ is the user–specified wheel rotation vector. All the simulations are performed with the vehicle at rest and this requires that the ground moves with the speed of the vehicle but in the opposite direction. The combined effect of the rotating wheel and rolling ground velocity boundary condition field is illustrated, for a single wheel, in Figure 11(b).

3.5.3 Jet engine / Rocket exhaust

Supersonic inflow conditions, provided by the engine and rocket manufacturers, are applied at the engine and rocket exhausts. The turbulent viscosity of the Spalart–Allmaras turbulence model is prescribed and is, again, set at a value equal to ten percent of the laminar viscosity value.

3.5.4 Symmetry Plane

At symmetry planes, the flux of the unknown vector $\mathbf{U}$ normal to the surface is assumed to be zero at all points on the plane. That is

$$\mathbf{F}^j n^j = 0 \tag{13}$$

where n^j is a vector normal to the surface. This implies that the fluid velocity everywhere on a symmetry plane is tangential to it:

$$u^j n^j = 0 \tag{14}$$

3.6 Computational Performance

To ensure a reasonable turnaround time for each simulation, the computational performance of the solution algorithm is improved by the use of parallel processing. The parallel implementation involves the subdivision of the original domain into a number of sub–domains, such that each sub–domain comprises a distinct set of mesh edges, along with the corresponding nodes that form the ends of each edge. This has the effect that nodes at the interface between two sub–domains are duplicated. The solver is executed, in parallel, utilising one process per sub–domain. At the start of a time step, the interface nodes obtain contributions from the interface edges. These partially updated contributions are transmitted to the corresponding interface nodes in the neighbouring sub–domains. A loop over the interior edges is followed by the receiving of the interface node contributions and the subsequent updating of all interior nodal values. This procedure is implemented in such a way as to allow computation and communication to be performed concurrently, where permitted by the parallel computer's hardware.

In order to optimise performance, and achieve scalability on a large number of processors, the chosen domain decomposition strategy must produce sub–domains of a balanced size and with a minimum number of cut edges. This ensures that each processor has to perform an equal amount of work and that the amount of communication between processors is minimised. This has been achieved by utilising the METIS family of partitioning algorithms [31]. In order to further improve computational performance, each sub–domain node list is renumbered, to maximise the use of cache memory. Computations were performed on a 64 processor (384 core) PC cluster running Intel Xeon E7450 processors with Mellanox DDR Infiniband for communication. Typical single Mach number steady state simulations on the parametric geometry model running on 60 cores would take approximately 12 hours to converge to 3 significant figure accuracy in the force coefficients.

4 Construction of the Parametric Geometry Model

Having studied the complexity of the flow field and geometry driving the aerodynamics at the rear of the vehicle, it was determined that the most efficient and cost–effective approach to solve the aerodynamic problem would be to parameterise the geometry and incorporate a design of experiments (DoE) optimisation methodology to guide the design. The first step in the process was

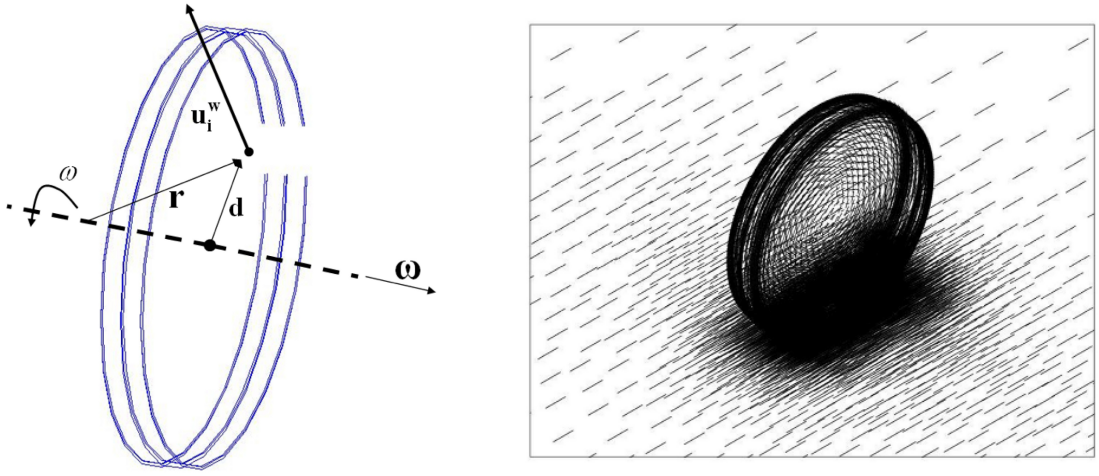

(a) (b)

Fig. 11 Simulation of a rotating wheel, showing :(a) interpretation of the boundary condition implementation; (b) the velocity vector field on the rotating wheel and on the rolling ground.

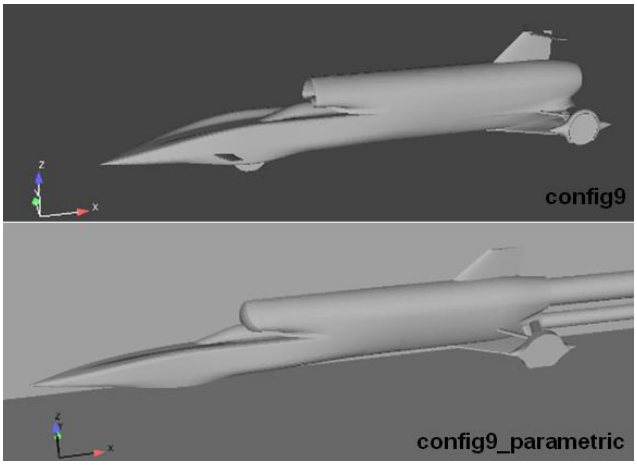

Fig. 12 Comparison of the config9 full car geometry and parametric model equivalent

Model	Mach∞	L/q total (m^2)	L/q front (m^2)	L/q rear (m^2)	D/q (m^2)	Drag LoA above ground plane (m)
config9 – full car	**1.3**	**1.39**	**0.369**	**1.021**	**1.88**	**0.63**
config9 - parametric	**1.3**	0.65 * 2 **= 1.3**	-0.119 * 2 **= -0.238**	0.769 * 2 **= 1.54**	0.93 * 2 **= 1.86**	**0.77**

Table 1 Comparison of the config9 full car geometry and parametric model lift and drag results

to simplify the geometry in order for it to be parameterised. It was decided to significantly simplify the front half of the vehicle by blocking the air intake and removing the front wheels and winglets, reduce the model from a full car to a half car by the use of a symmetry plane in the CFD model, simplify the rear wheel and suspension geometry and replace the jet and rocket exhausts with a pair of solid 'stings'. A comparison of the full config9 geometry with the simplified, parameterised equivalent is shown in Figure 12.

4.1 Validation of the Parameterised Model

Before the DoE process could begin, it was necessary to confirm that the simplified, parameterised model was representing the flow physics in the region of the rear wheels and suspension in a sufficiently acceptable manner. Since the optmisation was to be focussed on the supersonic ($M_\infty = 1.3$) behaviour of the geometry, a pair of test cases, involving the full car CFD model and a parameterised model, at $M_\infty = 1.3$ were compared. The lift and drag results from this pair of simulations is summarised in Table 1. Note that since the parametric geometry model is 'half–vehicle' the force coefficients have been multiplied by 2 in order to compare them with the full vehicle equivalents. It is evident that the parameterised model does not accurately reproduce the results generated by the full car CFD analysis, especially at the front of the car where key features of the geometry are not modelled. However, on comparing the flow visualisations for the two simulations it was deemed that the parameterised model simulation was capturing the key features of the flow physics, i.e., shock positions and strengths, expansions, separations etc, sufficiently well so that improvements to the parametric model using optimisation would translate back to the full geometry at the end of the process.

4.2 'Screening' the Variables for the DoE Study

The parameterised vehicle model consisted of 17 variable parameters including specifications for the wheel positions, the suspension strut positions and shaping of the rear car body and base. These are summarised in Figure 13. An initial study was run in order to understand which were the most powerful parameters in determining the aerodynamic behaviour of the vehicle. Each of the 17 parameters was run as a CFD model at $M_\infty = 1.3$ at their minimum and maximum positions in the design space about some baseline configuration. These minima and maxima in the design space parameters were dictated by practical engineering constraints on how far the geometry could be manipulated. In the majority of cases these were manufacturing or packaging constraints. In the remainder of cases the bounds were a function of stressing limits and predicted aero–elastic behaviour of the structure.

This initial study took approximately 1 week of CFD analysis time. At the end of the initial study, the rear wheel track was determined to be the most important parameter controlling both lift and drag by a considerable margin, with the minimum track width being optimal in the reduction of both lift and drag. This minimum rear track was dictated by limitations resulting from the requirement to achieve a safe roll stability margin for the vehicle. It was, therefore, decided to fix the rear wheel track at the mimimum value acceptable from a vehicle dynamic stability perspective (1.76m) and focus on the next 5 most powerful parameters in the DoE study. Concentrating on these 5 parameters minimised unnecessary computational expense in the optimisation process. These 5 'screened' parameters are highlighted in bold in Figure 13 along with the code letters used to reference them throughout the DoE study in brackets.

5 The DoE Process

The design optimisation procedure utilised in this study consists of:

1. Designing an experiment
2. Collecting data at the design points
3. Fitting a statistical model to the data
4. Running an optimisation using the model to find the optimal solution

Each of these items will be discussed in the following subsections:-

5.1 Design of Experiments

Design of Experiments (DoE) is the process for selecting the best design points within a specified design space. The DoE methodology can help choose the most effective design points at which to gather data in order to get the most information in the shortest possible time. DoE selects the design points to collect the most statistically useful data, and just enough of it to fit the models. When collecting the necessary data at each design point is time–consuming, as is the case when the design point data relies on a CFD simulation, as described in section 3, the most efficient design can save a great deal of time and money. A good choice of design points was critically important to the BLOODHOUND project because running a CFD simulation to achieve data for a single design point in this study typically took in the region of 12 hours of wallclock time. Lift and drag data was collected from the CFD model at the design points chosen by the DoE, then a statistical model was fitted to the design point data. The statistical model was then used to find optimal solutions. DoE tools were used to choose a set of experimental points that allowed estimation of the model with the maximum confidence while minimising computational expense (i.e., number of design points).

5.2 Choosing a Design

Prior knowledge about the expected shape of the response, based on aerodynamic first–principles, was used to help choose the best design. Engineering knowledge of theoretical relationships between lift, drag and the 5 selected design parameters was used to predict the model shape to best fit the data. It was predicted that a quadratic model could be used to give an appropriate balance between accuracy and computational expense of approximation models for the relationships between lift and drag coefficients with respect to the five design parameters.

Both space–filling and optimal designs were considered and then rejected. Space–filling designs are useful if the response surface shape is likely to be complex or if there is little or no information about the underlying effects of design parameters on responses. There was no reason to believe that the response surface would be complex, therefore space–filling designs were rejected. Optimal designs are good for systems with well–understood constraints, as they concentrate points near the edges of the design space. Optimal designs were rejected because there was no need to explore boundaries. At the time of the study, the BLOODHOUND project required fast, practical answers due

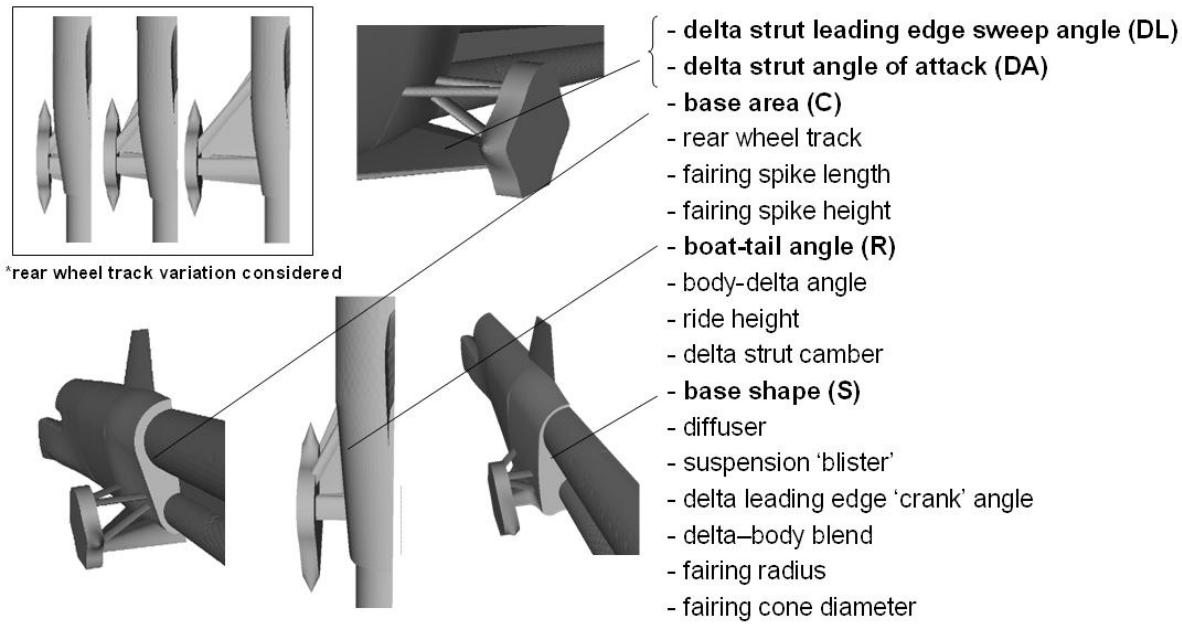

Fig. 13 Summary of the paramameters in the optimisation study

to time constraints on the overall design process. Given this demand, the best choice was a simple classical DoE. Classical designs are very well researched and are suitable for simple design spaces, such as the hypercube or sphere. Prior engineering knowledge predicted that the response could be fitted well by a quadratic model. The Central Composite Design (CCD) was chosen because it is a classical design commonly used for quadratics. The design points within the 5–dimensional space to be studied included a factorial design (the corners of a hypercube) together with a centre point and star points in the middle of each face of the cube. This is depicted, in the 3–dimensional case, in Figure 14.

The Central Composite Design specified 43 sampling points, each requiring a single CFD run at $M_\infty = 1.3$ in order to compute lift and drag coefficients. These 43 CFD runs took approximately 1 month to complete including meshing, solver time and post–processing.

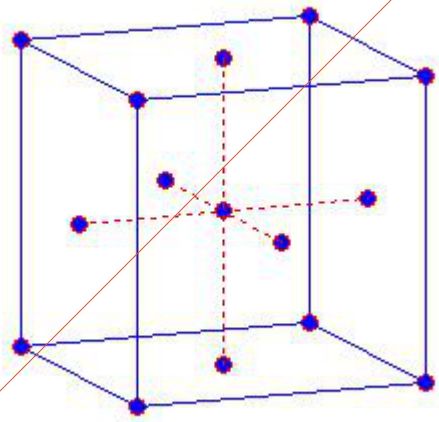

Fig. 14 Sampling points in a Central Composite Design space

5.3 Fitting Statistical Models to the CFD Data

The data from the CFD runs was used to fit a selection of response surface models and choose the best. As prior knowledge predicted a quadratic model, a quadratic model was fitted first. The stepwise quadratic process was also attempted. Stepwise is a process which eliminates the least useful model terms, one by one, in an attempt to achieve the simplest model that fits the data accurately [32]. The more complex Radial Basis Function (RBF) model was also used to investigate any possible complexity in the response surfaces. The RBF model looked like a plane, and as a result of this observation a linear model was fitted. The stepwise process did not succeed in simplifying the quadratic model all the way to a linear model because it only 'sees' one term at a time. All the alternative models were compared using graphical tools and statistical measurements. Measuring the root mean square error (RMSE) between the models and the data indicated that the simplest (linear) model produced a satisfactory fit with a RMSE of $0.027m^2$. A summary of the number of parameters in each model and the RMSE in fitting the CFD lift data is shown in Table 2.

Model	Number of Parameters	RMSE (m^2)
Full Quadratic	21	0.0297
Quadratic with Stepwise	8	0.0255
Radial Basis Function	5	0.0263
Full Linear	6	0.027
Linear with Stepwise	5	0.0266

Table 2 Comparison of the statistical models used to fit the CFD data

The simplicity of the linear model is demonstrated in the resulting lift formula shown in equation 15.

$$\begin{aligned} L/q = & \quad f(C, DA, R, S) \\ = & \ 0.55634 + 0.021225C + 0.024637DA \\ & -0.0070771R + 0.02416S \end{aligned} \tag{15}$$

Comparing models is a tradeoff process of accuracy versus model complexity. Less parameters are preferred to give more confidence that the noise in the CFD data is not being replicated in the model. In this case the simplest statistical model (the linear model) was judged the best because it provided the best tradeoff in describing the response accurately enough with the minimum model complexity. The residuals and 'predicted versus observed' plots are shown for the linear model in Figure 15 and a comparison of the linear and quadratic surfaces is shown in 16.

The y–axis on the top plot in Figure 15 shows the residuals, i.e., the difference between the value of rear L/q predicted by the linear model and the value for rear L/q predicted by the CFD simulation. The x–axis is the observation number (a variable that indexes through the CFD simulation predictions).

The lower plot shows the value of rear L/q predicted by the linear model on the y–axis, plotted against the value of the rear L/q predicted by the CFD simulations. The black line represents the $y = x$ line. Data points that are on this line have the property that the prediction made by the linear model exactly matches the prediction from the CFD simulation. The small black bars around the data points are the 95% confidence intervals for the linear model.

5.4 Using Optimisation with the Statistical Models

The initial optimisation study focussed entirely on lift minimisation, the objective function being simply the rear L/q. The optimal solution was sought using the best models describing the relationship between the 5 DoE parameters and the outputs, lift and drag. Due to

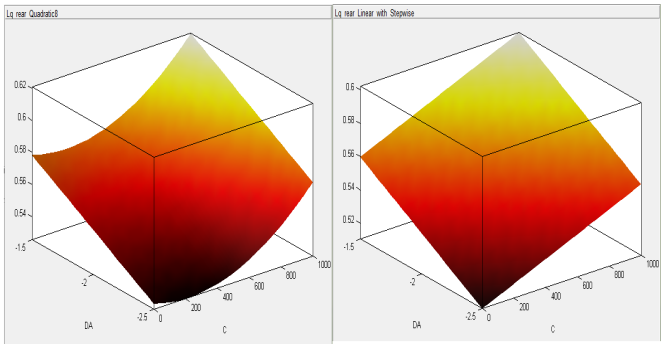

Fig. 16 Comparison of quadratic (left) and linear (right) response surfaces

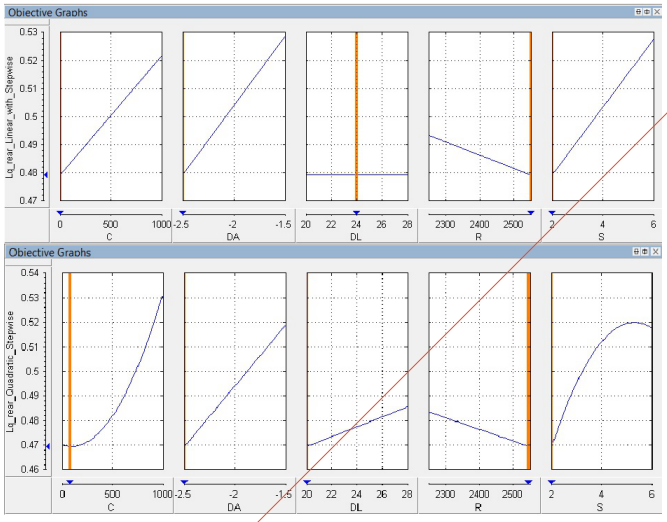

Fig. 17 Rear lift variation against the 5 optimisation parameters for linear (top) and quadratic (bottom) models

the simple nature of the linear model (see equation 15), the optimal setting for the 5 parameters was straightforward to identify. Note that other than the limits on the range of the DoE parameters, visible in Figure 17, there were no other constraints. The gradient–based optimisation algorithm 'fmincon' in MATLAB [33] was also used as a check using the quadratic stepwise model. Plots of lift against each of the five parameters along with indications of the optimal position in the design space (orange bars) using both linear and quadratic stepwise models is shown in Figure 17. Note that the letter codes for the variables are detailed in Figure 13. There is clearly very little difference between the optimal solutions predicted by the linear and quadratic stepwise models. Because of this, the linear model was adopted for the remainder of the optimisation investigations.

Figure 18 summarises the variations in lift and drag coefficients with the 5 DoE parameters. Clearly the linear model provides a better data fit for the lift data than for the drag data. But it is also clear that, with

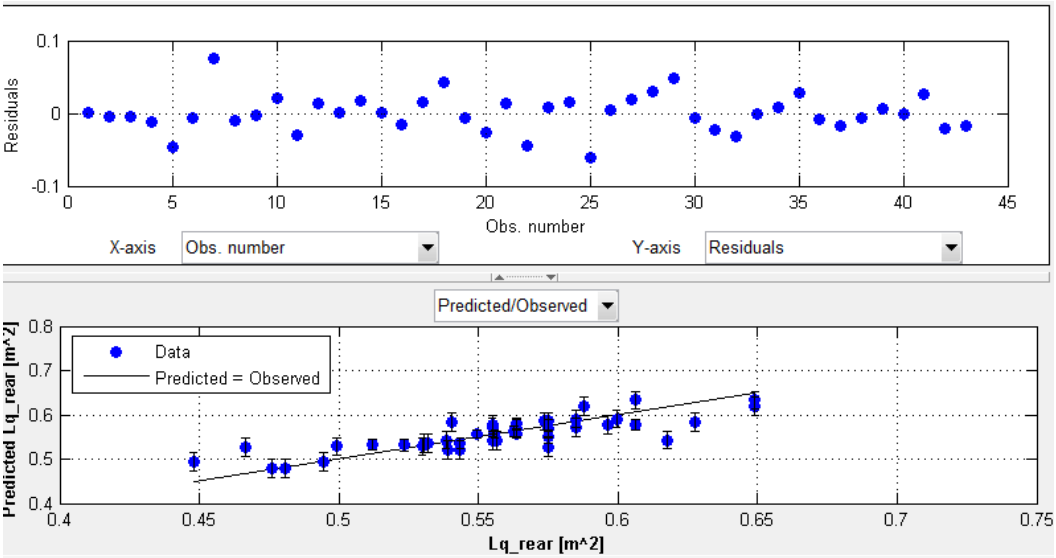

Fig. 15 Residuals and 'predicted versus observed' plots for the linear model

Model	Mach∞	L/q total (m^2)	L/q front (m^2)	L/q rear (m^2)	D/q (m^2)	Drag LoA above ground plane (m)
config9 – full car	**1.3**	**1.39**	**0.369**	**1.021**	**1.88**	**0.63**
config9 - parametric	**1.3**	0.65 * 2 **= 1.3**	-0.119 * 2 **= -0.238**	0.769 * 2 **= 1.54**	0.93 * 2 **= 1.86**	**0.77**
optimised parametric	**1.3**	0.15 * 2 **= 0.30**	-0.107 * 2 **= -0.214**	0.257 * 2 **= 0.514**	0.79 * 2 **= 1.58**	**0.87**

Table 3 Comparison of the config9 full car geometry, config9 parametric model and optimised parametric model lift and drag results

the exception of the boat–tail angle parameter (R), the mimimum drag parameter positions also coincide with mimimum lift parameter positions making it straightforward to select the best parameter values to minimise both lift and drag.

The plots in Figure 19 highlight the relative significances of the parameters and terms in the model in controlling both lift and drag (the top 30 are shown in each case). It is evident that delta angle of attack, base shape and base area are the most important parameters controlling lift, while base area and a combination of delta angle of attack, boat–tail angle and base shape are the most important parameters controlling drag.

A summary of the improvements in lift and drag coefficients between the baseline config9 parametric model and the optimised parametric model are shown in Table 3.

6 Post–optimisation Performance

Having used the statistical model to determine the optimal position within the 5–dimensional design space, given that the rear wheel track had been fixed at its minimum acceptable level, these parametric geomety changes were then translated back to a full car CFD model to compare with the pre–optimised config9 CFD results. Running simulations across the Mach range from $M_\infty = 0.5$ through to $M_\infty = 1.3$ resulted in the lift and drag variations detailed in Figure 20. Comparing this with the pre–optimised equivalent in Figure 5 it can be seen that the peak supersonic lift coefficient has been reduced from $1.7m^2$ to $0.4m^2$, and the peak drag coefficient reduced from $2.1m^2$ to $1.4m^2$ at $M_\infty = 1.1$. Assuming a dynamic pressure, $q = 83,900Pa$ at $M_\infty = 1.1$, this translates to a lift force and drag force reduction of 109kN and a 59kN, respectively.

Figure 21 shows the final full car geometry and $M_\infty = 1.3$ pressure coefficient distribution on the upper and lower surfaces of the vehicle. It is clear, upon comparing this supersonic underside pressure distribution with Figure 7, that the optimised geometry has significantly reduced the high pressures resulting from the rear wheel shock system. The rear wheel track reduction has also been an in important factor in minimizing the peak drag.

An overview of the improvement in the vehicle aerodynamic performance as a result of the optimisation across the Mach range from $M_\infty = 0.5$ to $M_\infty = 1.3$ is provided in Figure 22. It is evident that an optimisation focussed only on aerodynamic performance at

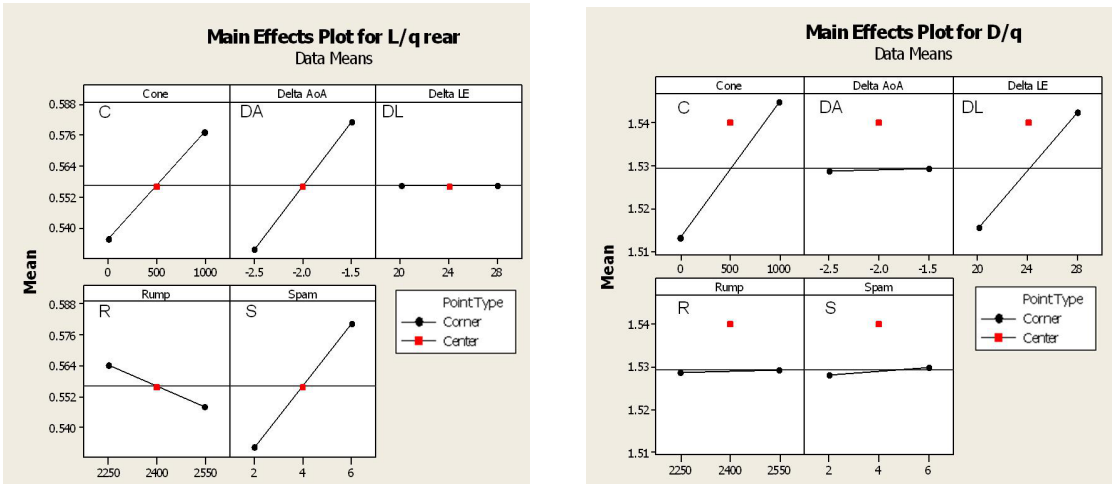

(a) (b)

Fig. 18 Summary of DoE results showing :(a) Rear L/q variation with DoE parameters; (b) D/q variation with DoE parameters

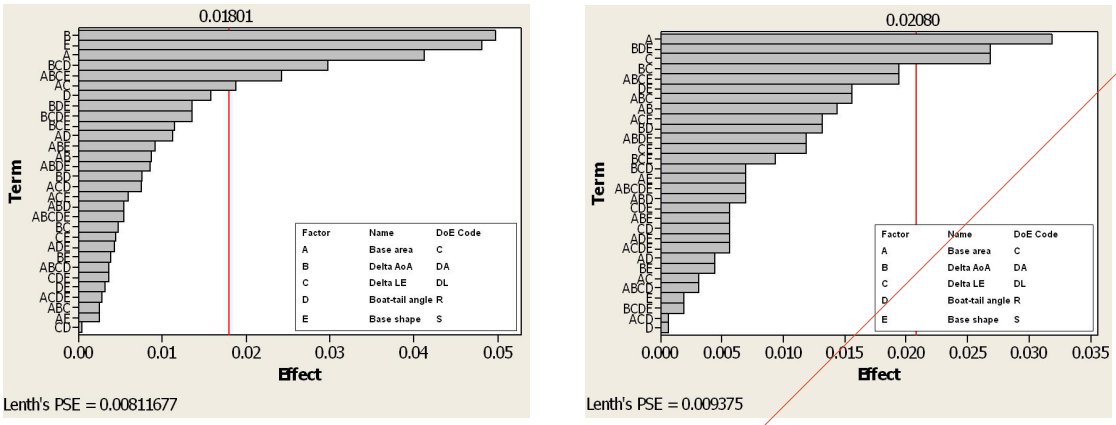

(a) (b)

Fig. 19 Plots of DoE parameter interactions for :(a) Rear L/q ; (b) D/q

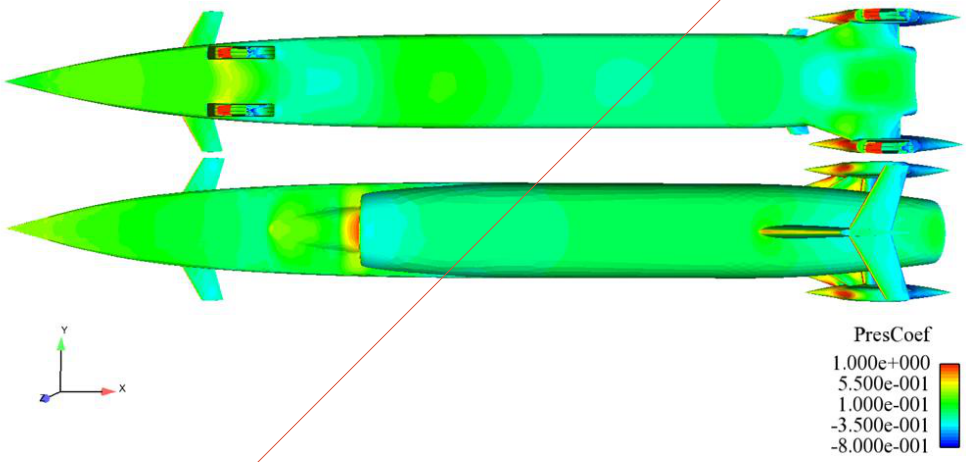

Fig. 21 'Post-optimisation' config10 upper and lower $M_\infty = 1.3$ pressure coefficient distribution

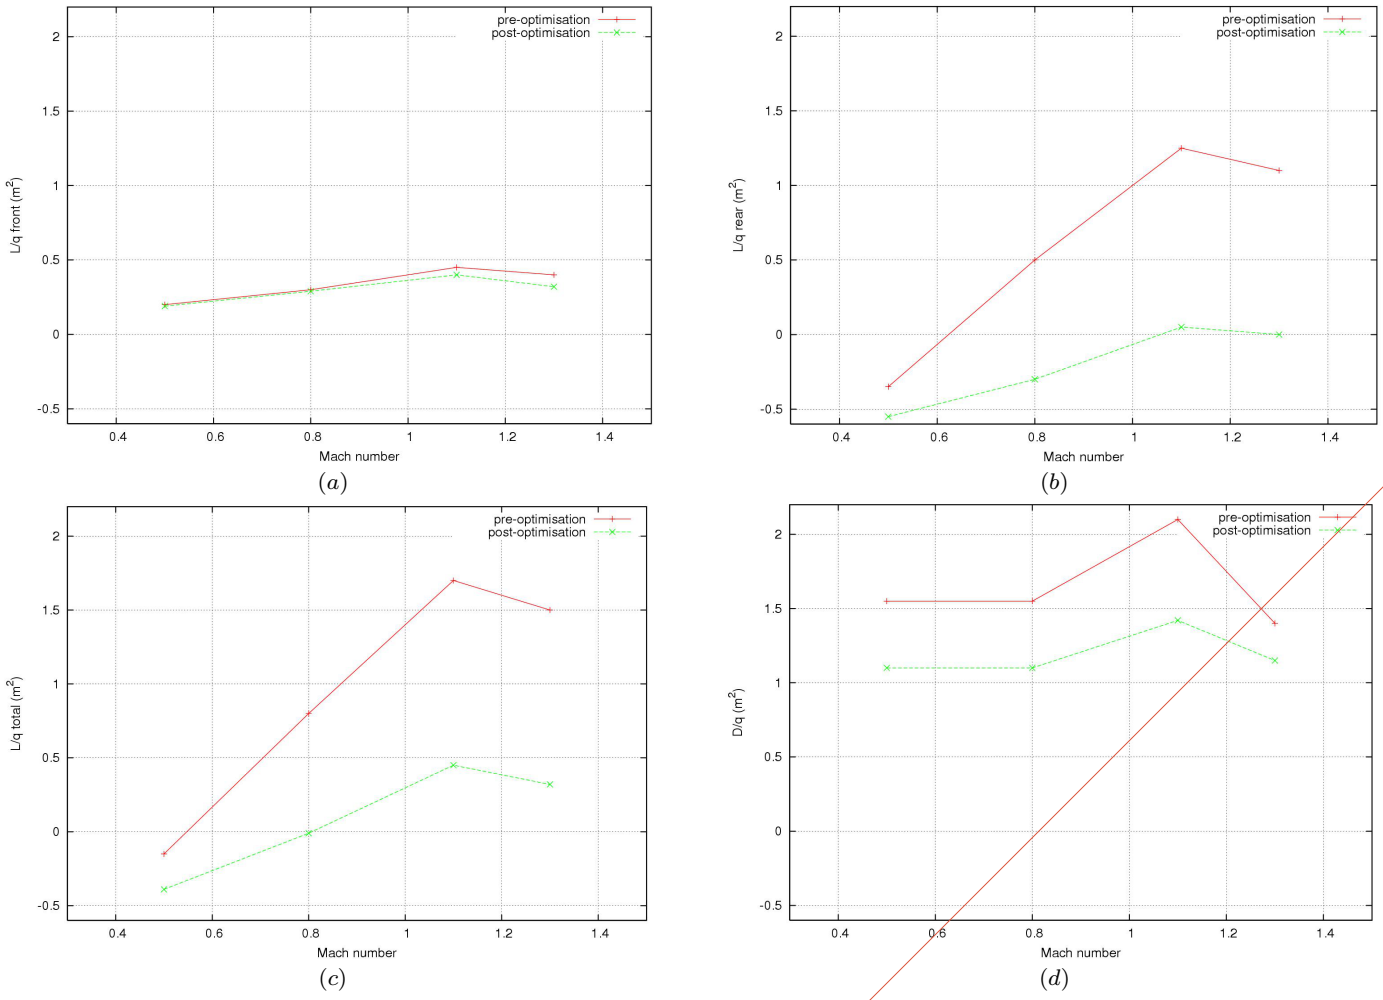

Fig. 22 Comparisons of pre-optimised and post-optimised lift and drag across the Mach range :(a) Front L/q ; (b) Rear L/q; (c) Total L/q; (d) D/q

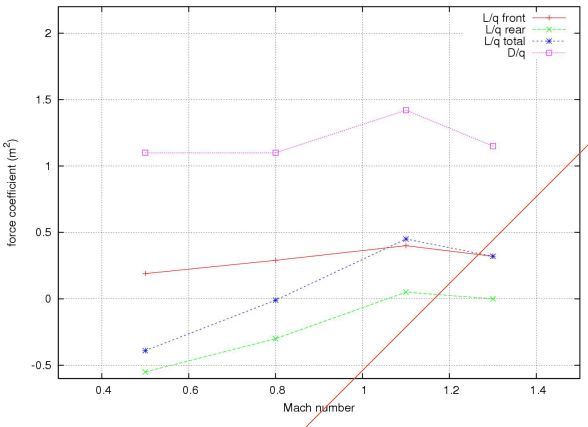

Fig. 20 Lift and drag coefficients against Mach number for 'post-optimisation' config10

$M_\infty = 1.3$ has provided benefits in terms of both lift and drag across the whole Mach range of interest.

7 Final Remarks

This paper has detailed the computational design optimisation strategy employed to achieve an engineering solution to the problem of excessive supersonic lift at the rear of the BLOODHOUND SSC (SuperSonic Car) during its design. The method utilised involved coupling data from CFD simulations with a parametric design optimisation utilising Design of Experiments. The final vehicle geometry predicted as optimal in this paper is, at the time of writing, being built and is due to be tested in 2013 in an attempt to increase the World Land Speed Record to 1,000 mph.

Acknowledgements The authors gratefully acknowledge the financial support provided for this work by the UK Engineering and Physical Sciences Research Council in the form of research grant EP/F032617. The work presented was a collaborative effort between Swansea University, BLOODHOUND

Ltd and MathWorks with computational support provided by INTEL.

References

1. R. Noble, *Thrust*, Transworld, London, 1998.
2. *The Story of ThrustSSC*, Corgi Books, London, 1998.
3. K. Morgan, O. Hassan, N. P. Weatherill, Why didn't the supersonic car fly?, *Mathematics Today, Institute of Mathematics and Applications*, **35**:110, 1999.
4. http://www.fia.com/en-GB/sport/records/Pages/Introduction.aspx
5. O. Hassan, K. Morgan, E. J. Probert, J. Peraire, Unstructured tetrahedral mesh generation for three–dimensional viscous flows, *International Journal for Numererical Methods in Engineering* **39**:549–567, 1996.
6. N. P. Weatherill, O. Hassan, Efficient three–dimensional Delaunay triangulation with automatic boundary point creation and imposed boundary constraints, *International Journal for Numererical Methods in Engineering* **37**:2005–2039, 1994.
7. K. Morgan, J. Peraire, Unstructured grid finite element methods for fluid mechanics, *Reports on Progress in Physics* **61**:569–638, 1998.
8. N. P. Weatherill, C. R. Forsey, Grid generation and flow calculation for complex aircraft geometries using a multi–block scheme, *AIAA Paper*, **85**, 1985.
9. S. Allright, Multiblock topology specification and grid generation for complete aircraft configurations, *Applications of Mesh Generation to Complex 3D Configurations (Conf. Proc. No 464)*, (Paris: AGARD), 11.1-11.11
10. A. Jameson, J.C. Vassberg, Computational Fluid Dynamics: Its Current and Future Impact, *AIAA Paper 2001-0538*, 2001.
11. http://www.boeing.com/commercial/787family/
12. E. Arian, A. Iolo, Challenges in Aerodynamic Optimization, *Computational Fluid Dynamics Review*, World Scientific, pp 447–467, 2010.
13. A. Jameson, Aerodynamic Design Via Control Theory, ICASE report NO. 88–64, *Journal of Scientific Computing* **3**:233–269, 1988.
14. d. Lee, K. Srivinas, L.F. Gonzalez, Robust Multidisciplinary Design Optimisation Using CFD and Advanced Evolutionary Algorithms, *Computational Fluid Dynamics Review*, World Scientific, pp 469–491, 2010.
15. B. J. Evans et al, Computational Fluid Dynamics Applied to the Aerodynamic Design of a Land–Based Supersonic Vehicle, *Numerical Methods for Partial Differential Equations* **27**:141–159, 2011 (MAFELAP 2009 Special Issue).
16. D. C. Wilcox, *Turbulence Modelling for CFD*, DCW Industries Inc, 1998.
17. F. M. White, *Viscous Fluid Flow*, Third Edition, McGraw Hill, Boston, 2006.
18. P. R. Spalart, S. R. Allmaras, A one–equation turbulent model for aerodynamic flows, *AIAA Paper 92–0439*, 1992.
19. www.cadfix.com
20. J. Peraire, M. Vahdati, K. Morgan, O. C. Zienkiewicz, Adaptive remeshing for compressible flow simulations, *Journal of Computational Physics* **72**:449–466, 1987.
21. J. Peiró, J. Peraire, K. Morgan, *FELISA system reference manual. Part 1–basic theory*, University of Wales Swansea Report C/R/821/94, 1994.
22. J. Peiró, J. Peraire, K. Morgan, The generation of triangular meshes on surfaces, in C. Creasy and C. Craggs (editors), *Applied Surface Modelling*, Ellis–Horwood, Chichester, 25–33, 1989.
23. K. A. Sørensen, *A multigrid accelerated procedure for the solution of compressible fluid flows on unstructured hybrid meshes*, PhD Thesis, University of Wales, Swansea, 2002.
24. J. Peraire, J. Peiró, K. Morgan, Multigrid solution of the 3D compressible Euler equations on unstructured tetrahedral grids, *International Journal for Numerical Methods in Engineering* **36**:1029–1044, 1993.
25. K. A. Sørensen, O. Hassan, K. Morgan, N. P. Weatherill, A multigrid accelerated hybrid unstructured mesh method for 3D compressible turbulent flow, *Computational Mechanics* **31**:101–114, 2003.
26. P. I. Crumpton, P. Moinier, M. B. Giles, An unstructured algorithm for high Reynolds number flows on highly stretched grids, in C. Taylor, J. T. Cross (editors), *Numerical Methods in Laminar and Turbulent Flow*, Pineridge Press, Swansea, 561–572, 1997.
27. K. A. Sørensen, O. Hassan, K. Morgan and N. P. Weatherill. A multigrid accelerated hybrid unstructured mesh method for 3D compressible turbulent flow. *Computational Mechanics* **31**, 101–114, (2003).
28. P. Batten, N. Clarke, C. Lambert, D. M. Causon, On the choice of wavespeeds for the HLLC Riemann solver, *SIAM Journal of Scientific Computation* **18**:1553–1570, 1997.
29. A. Harten, P. D. Lax, B. Van Leer, On upstream differencing and Godunov–type schemes for hyperbolic conservation laws, *SIAM Review* **25**:35–61, 1983.
30. W. J. Usab, E. M. Murman, Embedded mesh solutions of the Euler equations using a multiple–grid method, in W. G. Habashi (editor), *Advances in Computational Transonics*, Pineridge Press, Swansea, 447–472, 1985.
31. G. Karypis, V. Kumar, *METIS4.0: Unstructured graph partitioning and sparse and sparse matrix ordering system*, Technical Report, Department of Computer Science, University of Minnesota, 1998.
32. Hocking, R. R., *The Analysis and Selection of Variables in Linear Regression*, Biometrics, **32**, 1976.
33. Nocedal, J. and S. J. Wright. Numerical Optimization, Second Edition. Springer Series in Operations Research, Springer Verlag, 2006