Computational Fluid Dynamics Applied to the Aerodynamic Design of a Land-Based Supersonic Vehicle

B. J. Evans, O. Hassan, J. W. Jones, K. Morgan, L. Remaki
Civil and Computational Engineering Centre, School of Engineering, Swansea University, Swansea SA2 8PP, Wales, United Kingdom

Received 1 July 2010; accepted 12 July 2010
Published online 1 October 2010 in Wiley Online Library (wileyonlinelibrary.com).
DOI 10.1002/num.20644

The BLOODHOUND SSC project was publicly announced in October 2008, with a primary engineering objective of designing, constructing and running a vehicle capable of achieving a speed of 1000 mph on land. The aerodynamic design of this vehicle is to be accomplished using computational simulation only and this paper describes the development and application of the approach adopted. The computational model employs a cell vertex finite volume algorithm for the solution of compressible viscous flow problems on unstructured hybrid meshes. A one equation turbulence model is adopted and the solution of the steady flow equations is obtained by explicit relaxation. For the combination of high Mach number, complex geometry and complex boundary conditions, involving rotating surfaces and a rolling ground, a consistent HLLC numerical flux function is adopted to ensure a stable procedure. To illustrate the impact of the approach upon the final configuration, a number of simulations undertaken to aid the aerodynamic design are described.

Keywords: BLOODHOUND SSC; finite volume method; HLLC numerical flux function; hybrid unstructured mesh; parallelization; Spalart–Allmaras turbulence model; steady 3D compressible aerodynamic flow

I. INTRODUCTION

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 at the Black Rock Desert, Nevada in 1997. A significant feature of the THRUST SSC project was the manner in which computational fluid dynamics was applied to guide the aerodynamic design process [1, 2].

The BLOODHOUND SSC project was publicly announced in October 2008 and involves the development of a vehicle which will attempt to take the World Land Speed Record to 1000 mph. Again, it has been decided that computational fluid dynamics will be used to guide the aerodynamic design of the vehicle. The computational approach that has been adopted will be outlined in this article and consists of a finite volume method [3, 4] for simulating turbulent high speed

Correspondence to: K. Morgan, Civil and Computational Engineering Centre, School of Engineering, Swansea University, Swansea SA2 8PP, Wales, United Kingdom (e-mail: k.morgan@swansea.ac.uk)
Contract grant sponsor: UK Engineering and Physical Sciences Research Council; contract grant number: EP/F032617

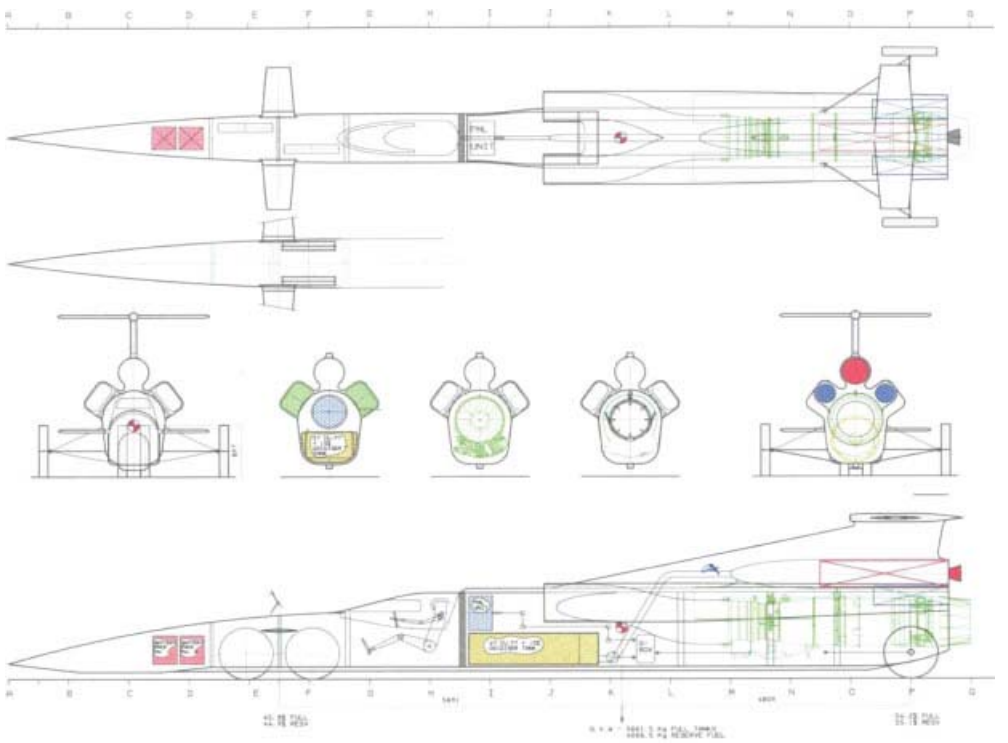

FIG. 1. Drawing illustrating the initial concept layout of BLOODHOUND SSC. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

compressible flows on hybrid unstructured meshes. Boundary layer regions are discretized using stretched prisms and pyramids formed from merging tetrahedra generated by an advancing layers method [5]. Outside the boundary layers, isotropic tetrahedral elements are generated [6]. The result is a hybrid mesh approach that enables the efficient meshing of complex configurations. Stabilization of the equation system [7] is achieved by using a consistent HLLC flux function [8,9], to enable the simulation of the moving vehicle, and methods for modeling high speed rotating wheels and rolling surfaces are implemented. The stabilized equation system is solved by explicit iteration, with the computational performance enhanced by means of effective parallelization [10].

The initial concept for the BLOODHOUND vehicle is shown in Fig. 1. The vehicle body is long and slender with two widely spaced rear wheels outboard of the main body providing roll stability. The vehicle thrust is provided by a turbofan jet engine, shown in green, and a hybrid rocket, shown in red, which are capable of delivering 90 and 120 kN of thrust, respectively. The impact of the use of the computational techniques during the design evolution of the BLOODHOUND SSC is illustrated by considering the aerodynamic behavior of the front wheels of the vehicle and their configuration relative to the nose.

II. PROBLEM FORMULATION

Relative to a cartesian $Ox_1x_2x_3$ coordinate system, the steady state compressible Favre–averaged Navier Stokes equations [11] governing airflow over a moving vehicle are expressed in the integral form

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

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

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

and the inviscid and viscous flux vectors are given as

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

respectively. The unit outward normal vector to Γ is $\boldsymbol{n} = (n_1, n_2, n_3)$ 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 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 [12]. The variation of the kinematic turbulent viscosity is obtained from the one–equation turbulence model of Spalart and Allmaras [13].

III. DOMAIN DISCRETIZATION

For any practical aerodynamic simulation, a CAD system is normally employed to provide the definition of the boundary geometry. The definition obtained in this way frequently requires treatment, to produce a watertight surface description that is suitable for the processes of surface and

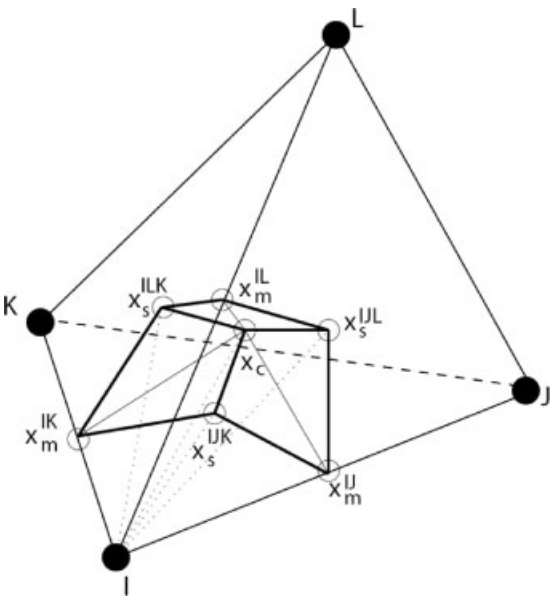

FIG. 2. Illustration of the portion of the element Ω_I of the dual mesh, associated with node I, that is contained inside a tetrahedral element.

volume mesh generation. This treatment process will involve the merging of surfaces and intersection points and the smoothing of underlying surface definitions. In our work, this process is accomplished by using the CADfix [14] commercial software package.

Before the start of the discretization process, the desired element size distribution is defined by the user in terms of a mesh control function [15, 16]. The first step in the discretization process is the triangulation of the computational boundaries and this is accomplished using an advancing front approach [17]. 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 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 discretized 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 [10].

IV. EQUATION DISCRETIZATION

Nodes are located at the vertices of the generated mesh and the spatial discretization of Eq. (1) is accomplished using a cell centered finite volume procedure. This requires the construction of a dual mesh, in which each cell of the dual is associated with a single node of the hybrid mesh.

A. Dual Mesh Construction

The median dual is employed within that portion of the mesh consisting of isotropic tetrahedral elements. This dual is constructed by connecting edge midpoints, element and face centroids in

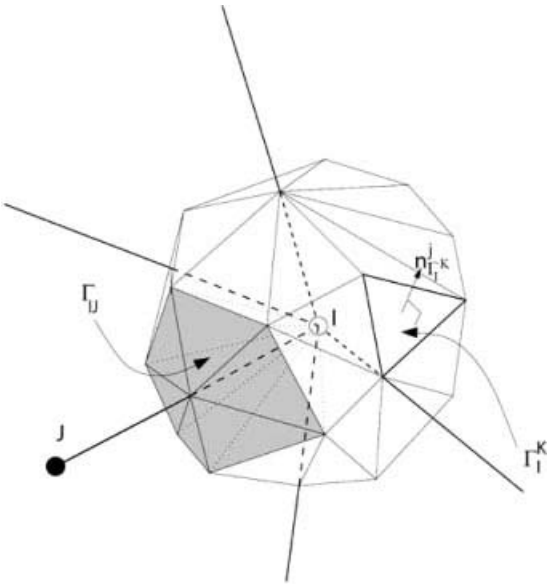

FIG. 3. The volume Ω_I of the dual mesh surrounding an interior node I.

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 Fig. 2. The midpoint of the edge between node I and J is denoted by $\boldsymbol{x}_m^{IJ}$, the centroid of the face with vertices I, J, and K is denoted by $\boldsymbol{x}_s^{IJK}$ and the element centroid is designated by $\boldsymbol{x}_c$. 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 Fig. 3. 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 [18].

B. Treatment of the Inviscid Fluxes

In the cell vertex finite volume method, Eq. (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 Fig. 3. Additional coefficients need to be introduced to enable integrals over the computational boundary to be approximated [10].

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

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

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

C. Treatment of the Viscous Fluxes

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

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

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 Eq. (9). While this is convenient computationally, the result is a five point discretization 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 [19].

A different edge based approach can be developed which results in a more compact stencil [20]. To achieve this, the derivative of a flow variable at an edge midpoint is represented in terms of its components parallel and perpendicular to the edge. The component perpendicular to the edge is calculated in the normal manner, while the component parallel to the edge is approximated by using a two point finite difference scheme [19]. 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.

D. Stabilization

This discretization 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 Eq. (8) by a consistent numerical flux. Here, the HLLC numerical flux function [8] is adopted, which is a modification of the original HLL scheme [9]. The function is constructed by assuming that the solution to the Riemann problem, between states $\boldsymbol{U}_I$ and $\boldsymbol{U}_J$, 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 [8].

E. Boundary Conditions

To complete the specification of the problem, boundary conditions must be defined over the entire boundary of the computational domain.

Inflow and Outflow Boundaries. At basic inflow and outflow boundaries, a characteristic treatment [21] is employed to determine the number and the 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 10% of the laminar viscosity value. The turbulent viscosity value at outflow is obtained by the extrapolation of values from the interior of the computational domain.

Viscous Walls. For viscous flow, the no slip condition

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

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

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

where, as illustrated in Fig. 4(a),

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

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 Fig. 4(b).

Engine Inflow. When modeling a jet engine, the internal portion of the engine is considered to lie outside the computational domain. This means that an outflow boundary condition is required

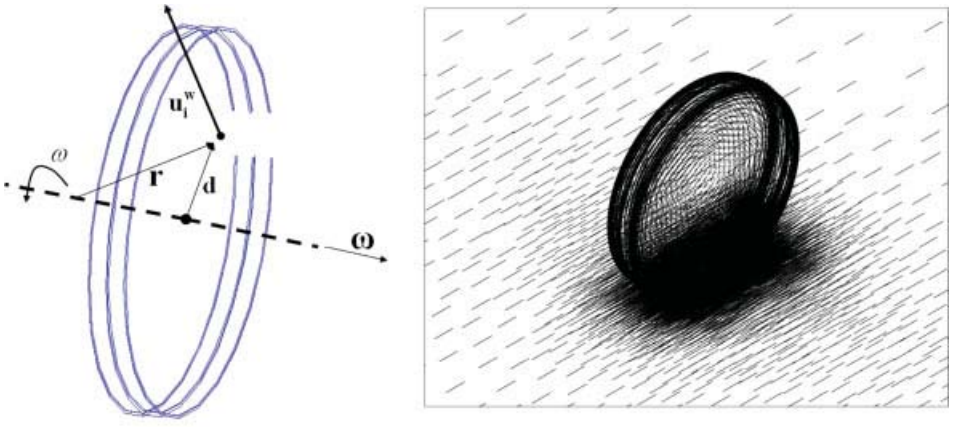

FIG. 4. Simulation of a rotating wheel, showing: (a) interpretation of the boundary condition implementation; (b) the velocity vector field on the rotating wheel and on the rolling ground. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

at the engine compressor face. The most useful description of the demand placed on an aerodynamic flow by a jet engine face is the specification of a non–dimensional mass flow function [22], which is defined as

$$W^* = \frac{\dot{m}\sqrt{c_p T_0}}{p_0} \tag{13}$$

where $\dot{m}$ is the mass flow rate, T_0 is the total temperature and p_0 is the total pressure. This mass flow function depends upon the throttle setting of the engine and the free stream Mach number. The incorporation of the total pressure and total temperature into this definition enables the flow model to automatically account for the irreversible aerodynamic losses resulting from shock systems and flow separation in the intake duct when computing the flow across the engine face. This is not the case when a simple mass flow rate is prescribed at the engine face. The mapping between the throttle setting, the Mach number and the mass flow function is information that can be provided by the manufacturer of the engine. At each stage of the solution process, the requirements of the mass flow function are satisfied using an iterative technique, based upon locally modifying the velocity vector field at the engine face.

Jet Engine/Rocket Outflow. 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 arbitrarily set at a value equal to 10% of the laminar viscosity value.

F. Discrete Equation

Following the discretization of the inviscid and the viscous fluxes, the discrete form of the governing Eq. (1) can be written as

$$\boldsymbol{R}(\boldsymbol{U}) = \mathbf{0} \tag{14}$$

The discrete equation for the Spalart–Allmaras model [13] is constructed in a similar fashion, with the convective term discretized using a first order upwind method [7]. The complete equation system is solved by explicit relaxation using a three stage Runge Kutta approach with local time stepping [10].

G. Improving 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 subdomains, such that each subdomain 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 subdomains are duplicated. The solver is executed, in parallel, utilizing one process per subdomain. 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 neighboring subdomains. 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.

To optimize performance, and achieve scalability on a large number of processors, the chosen domain decomposition strategy must produce subdomains 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 minimized. This has been achieved by utilizing the METIS family of partitioning algorithms [23]. To further improve computational performance, each subdomain node list is renumbered, to maximize the use of cache memory. All computations were performed on the School of Engineering PC cluster of AMD opteron processors at Swansea University.

V. APPLICATION TO BLOODHOUND SSC

The design of a manned land-based vehicle that will travel at speeds significantly exceeding the speed of sound for a sustained period of time is extremely challenging. Unsurprisingly, this design objective poses some serious aerodynamic problems. Of primary importance in understanding the aerodynamic characteristics of such a vehicle is the accurate prediction of shock wave positions, including shock reflections, within the flow domain. For example, the 900 mm diameter wheels on BLOODHOUND SSC are designed to rotate at 10,000 rpm at the vehicle's top speed of 1050 mph. This implies that the top of the wheel will be moving forward relative to a stationary observer at 2100 mph, which is equivalent to a Mach number of around 2.8. In addition, the rules of the Land Speed Record imply that the vehicle must be steered using its wheels alone, which highlights the importance of understanding the aerodynamic characteristics of the wheels across their range of steering angle.

A. Wheel Design

BLOODHOUND SSC will have solid metal wheels (aluminium or titanium), as radial accelerations of the order 50,000g will be experienced at the wheel rim at top speed. It is anticipated that improved traction, between these solid wheels and the compliant alkali playa desert surface across which the vehicle is expected to run, can be achieved by using circumferential keels on

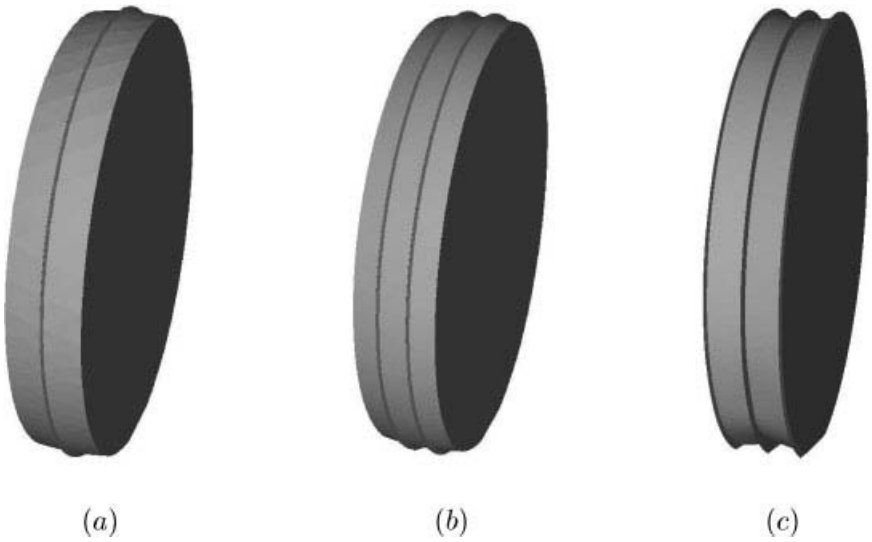

FIG. 5. Possible wheel design options, showing: (a) the single keel; (b) the double keel; (c) the triple keel.

the wheel surface. Such keels should help generate the lateral loads essential for stability and steering control. The aerodynamic implications of the number of keels used, and their location, was analysed using computational modeling.

Three wheel designs with different keel configurations were considered. The first design employed a single keel running around the centreline of the wheel; the second design used two keels running around the wheel profile; the third design had a single keel running around the wheel centreline together with a keel on each wheel edge. These three geometric wheel design options are illustrated in Fig. 5.

The interaction between the aerodynamic flow and the ground surface in the vicinity of the wheel contact patch is deemed to be of critical importance. This is due to the fact that the rear wheels will have to run across a surface that has been disturbed by the front wheels. Therefore, it is important for the wheel design to be such that the pressure field generated in the vicinity of the contact patch is as benign as possible. Analysis of the aerodynamic effects of the wheel design were carried out using local models, consisting of the wheel and rolling ground alone. The hybrid meshes employed were typically of the order of 1.8 M elements, 320 K nodes, and 2.5 M edges. The simulations were performed using 8 CPUs of the Swansea C2EC cluster, with convergence requiring 6 K cycles for supersonic problems and 20 K cycles for subsonic problems. At a typical computation rate of 13.8 cycles per minute, this equates to wallclock run times of around 7 and 24 hrs for supersonic and subsonic problems, respectively. Both lift and drag coefficients were converged to an accuracy of three significant figures. Figure 6 shows a detail of the computed Mach number distribution close to the ground surface for two different keel configurations at a free stream Mach number $M_\infty = 0.5$. Figure 7(a) shows a detail of the predicted pressure distribution, for the double keel wheel, on the vertical plane running through the wheel centre when $M_\infty = 1.0$. The pressure coefficient distribution on the top (red) and bottom (green) of the wheel rim is shown, in Fig. 7(b), demonstrating that at these high speeds the wheel is generating significant lift. Figure 8 indicates the impact of the wheel keel configuration on the variation of the lift and drag with Mach number for the wheel. It is apparent that the number of keels adopted has a significant influence on the aerodynamic properties of the wheel. For example, at $M_\infty = 1.1$,

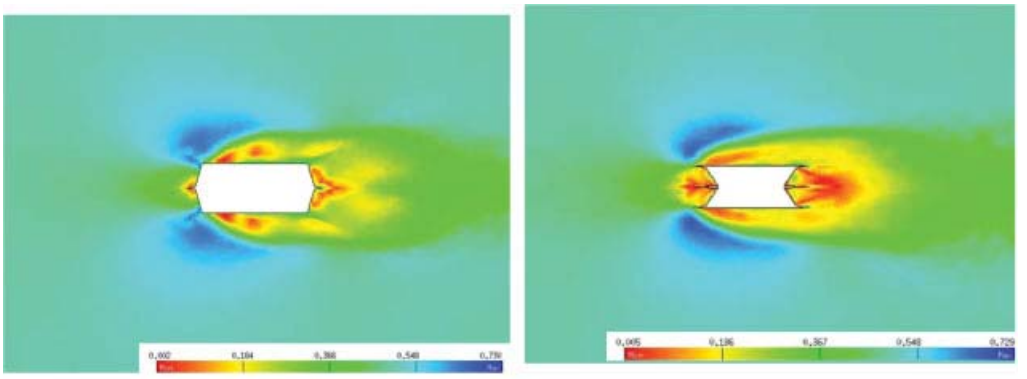

(a) (b)

FIG. 6. Mach number contours on a plane 0.01m above the surface around the wheel contact patch, for two different wheel profiles at $M_\infty = 0.5$: (a) a single keel wheel; (b) a triple keel wheel. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

there is a difference of half a ton per wheel in the lift forces generated by the single and triple keel options, while the difference in the drag forces being generated is of the order 2 kN per wheel. These differences are primarily due to the compression levels achieved just upstream of the wheel, at the contact patch, because of the different keel configurations.

It was decided that from a purely aerodynamic point of view, the single keel behaved more favorably than the double keel wheel, and the double keel wheel more favorably than the triple keel wheel. At the time of writing, the wheel design choice had been narrowed down to the use of either a single or a double keel. The final decision will be made after testing these two wheel profiles on the desert, to determine their dynamic response to the alkali playa surface.

B. Front Wheel Configuration

Sketches indicating the initial concept for the external shape and packaging of BLOODHOUND SSC are shown in Fig. 1. From these sketches, it can be observed that two front wheel configuration

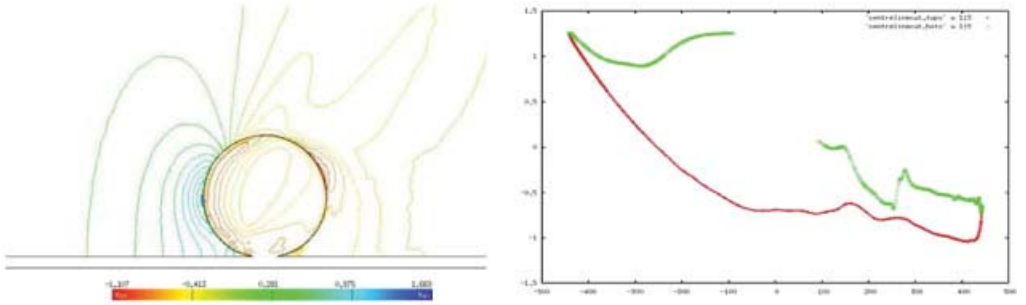

(a) (b)

FIG. 7. A typical pressure coefficient distribution at a free stream Mach number $M_\infty = 1.0$ (a) on the wheel surface and vertical plane through the wheel centre; (b) on the wheel rim, with top in red and bottom in green. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

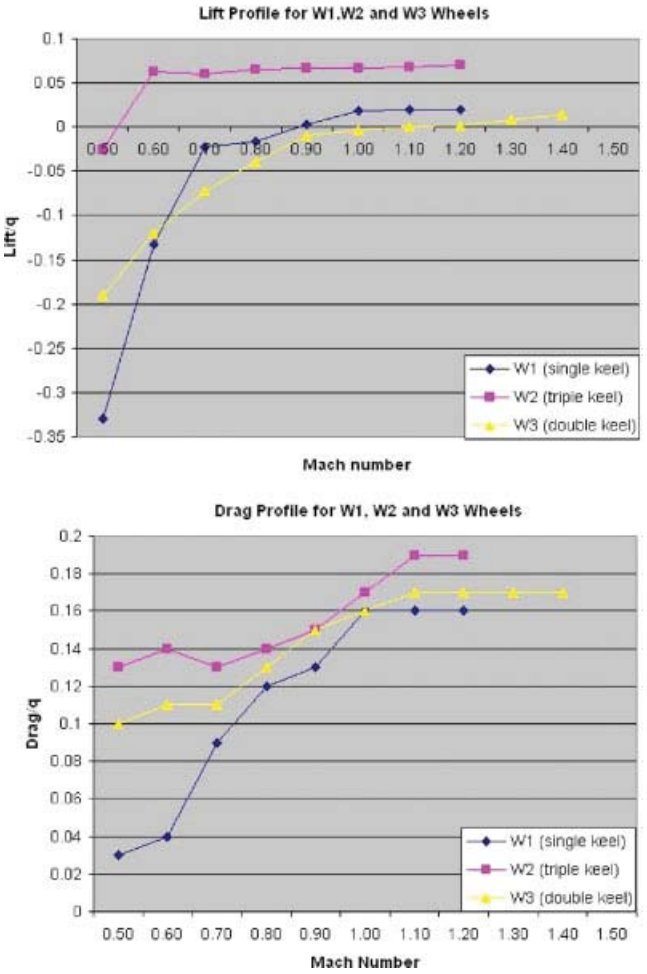

FIG. 8. Variation in the lift and drag, normalized by dynamic pressure, q, with free stream Mach number M_∞, for each wheel design considered. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

options were under consideration. One option has the wheels located in a staggered asymmetric configuration, while the other option displays a parallel symmetric configuration. The advantage of the staggered configuration is that, in this case, the packaging of the suspension allows for a more slender nose, and hence lower aerodynamic drag, particularly at supersonic speeds, when wave drag can be expected to dominate. However, it is desirable that the vehicle should respond symmetrically both to yaw and steering input. Initially, the response of the staggered configuration was considered, with the wheel fairing options detailed in Fig. 9, to investigate if symmetric behavior was achievable. The simulations were carried out using a local model, consisting of the vehicle nose, cockpit and the front wheel pair rotating within their wheel wells. The hybrid meshes typically were of the order of 11.3 M elements, 1.9 M nodes, and 14.8 M edges. The simulations were performed using 10 CPUs of the Swansea C2EC cluster, with convergence requiring 5 K cycles for supersonic problems and 15 K cycles for subsonic problems. Typically, 2.4 cycles

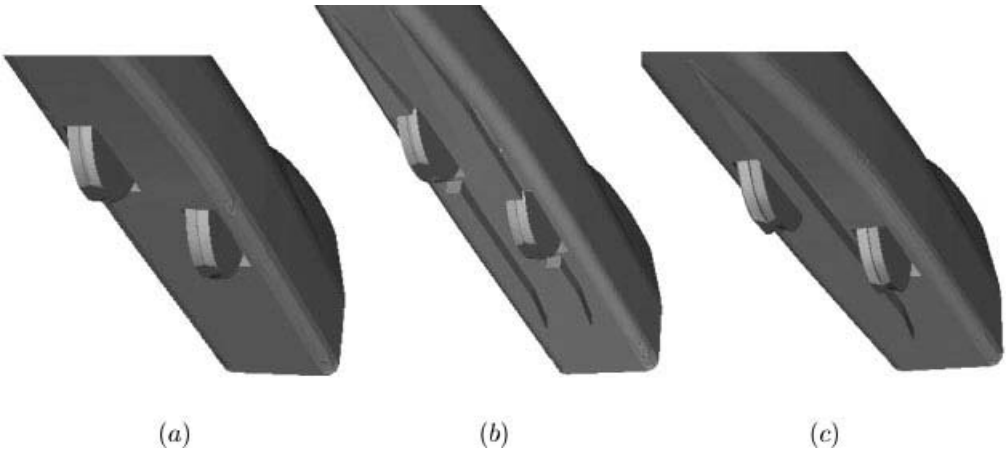

FIG. 9. Staggered wheel geometry options: (a) no fairings; (b) fairing option 1; (c) fairing option 2.

were computed per minute, which equates to a wall clock runtime of around 24 and 100 hs for supersonic and subsonic problems, respectively. Both lift and drag coefficients were converged to an accuracy of three significant figures.

Figure 10 shows a detail of a typical subsonic flow pattern on the ground surface in the vicinity of the contact patch for the staggered wheel layout and Fig. 11 presents the computed yaw and steer lateral force response for the various staggered wheel fairing options with both the single and triple keel wheel geometries. It is apparent that a symmetric aerodynamic response between -2° and $+2^\circ$ has not been achieved. It is believed that this is largely due to the interaction between the trailing wheel and the wake of the leading wheel. Based on these simulations, the

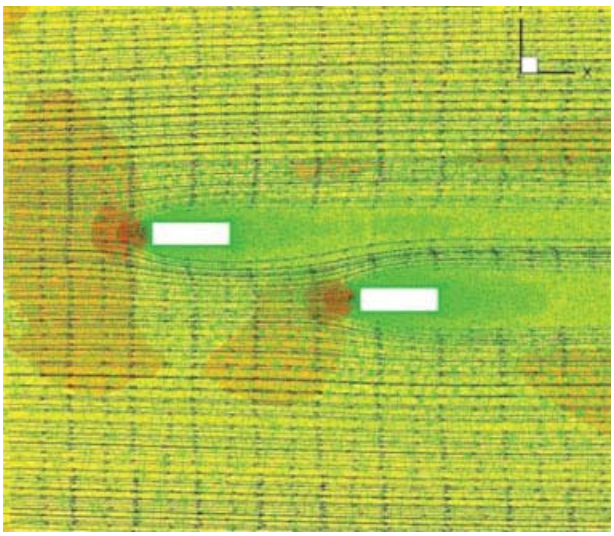

FIG. 10. Subsonic simulation with a staggered front wheel configuration, showing streamlines and Mach contours. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

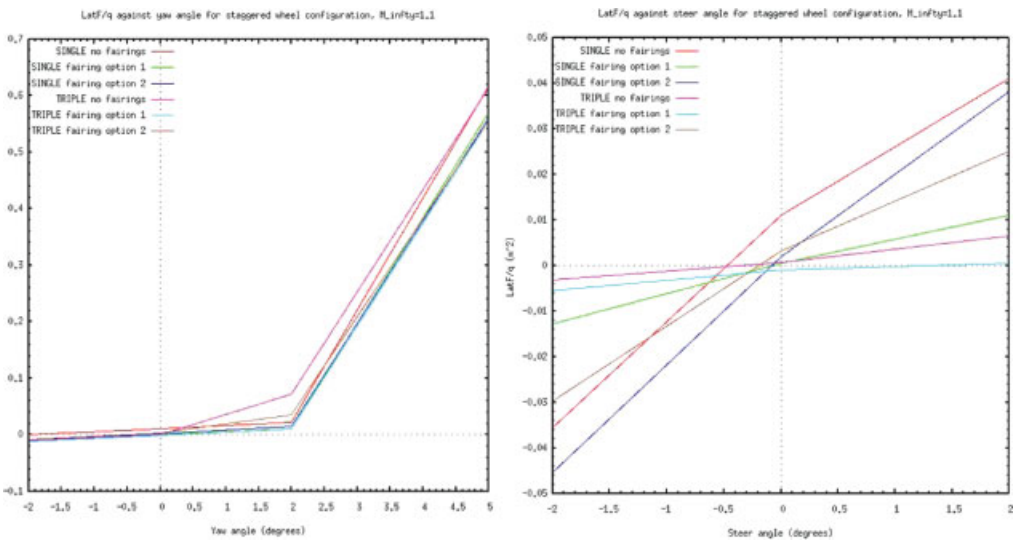

FIG. 11. Yaw and steer characteristic with staggered wheel configuration at $M_\infty = 1.1$. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

parallel wheel configuration shown in Figure 12 was adopted for continued design development. Figure 13 indicates the streamline pattern and local Mach number distribution, created with the parallel configuration, on a plane 0.01 m above the ground, close to the contact patch for a single keel at $M_\infty = 1.4$. Figure 14 indicates the lateral force response to steer and yaw angle for a single keel parallel configuration, with and without fairings, at $M_\infty = 1.1$. It is evident that an aerodynamic response to yaw and steer angle that is much closer to symmetric has been achieved. The parallel wheel configuration, was therefore, adopted in preference to the staggered wheel option, at the expense of a small drag penalty.

FIG. 12. Parallel wheel geometry with fairings. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

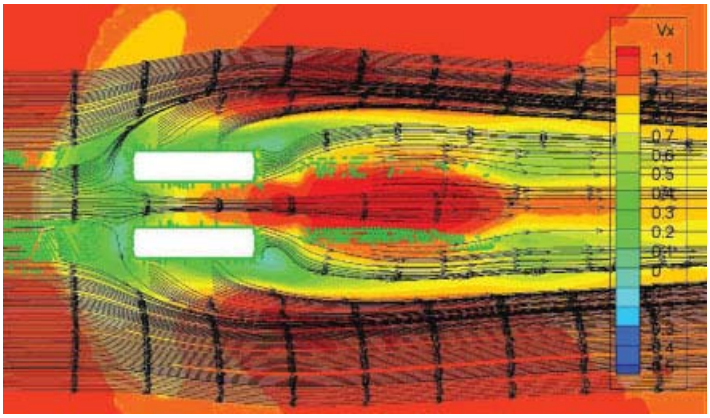

FIG. 13. Simulation with a parallel front wheel configuration showing streamlines and Mach contours at $M_\infty = 1.4$. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

A final consideration, when studying the flow pattern over the BLOODHOUND nose and front wheel configuration, is the predicted particle movement in the vicinity of the wheels. BLOODHOUND SSC will be running across an alkali playa desert surface and the wheels will have significant interaction with the surface. This wheel/surface interaction, combined with the influence of the vehicle's pressure field on the surface, will lead to dust and particle entrainment into the flow field. Comparisons of particle traces, such as those shown in Fig. 15, give an indication of the extent to which this particle entrainment will affect the wheel well surfaces and surfaces downstream of the front wheels. This, of course, assumes that the entrained particles behave as though they are massless.

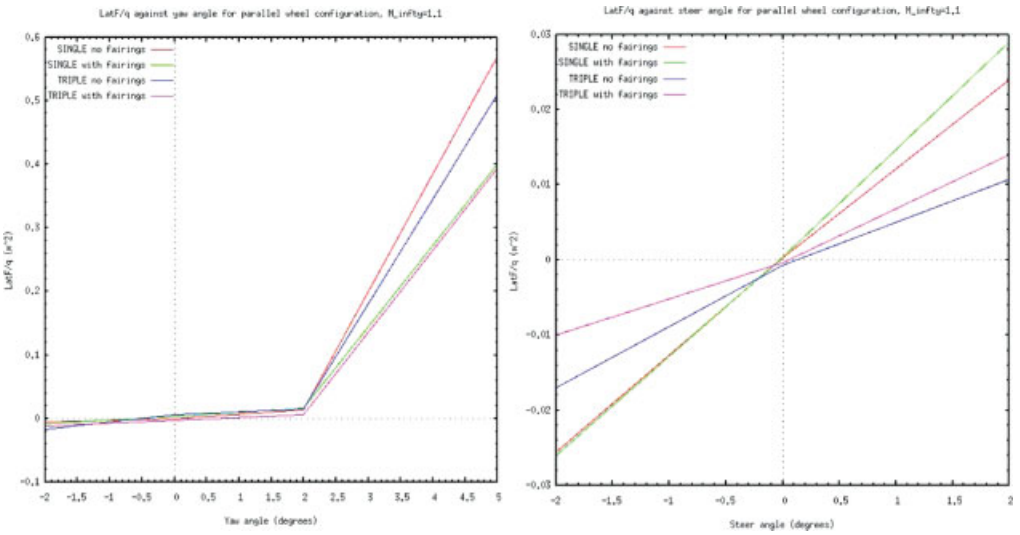

FIG. 14. Yaw and steer characteristic for a parallel single keel wheel configuration at $M_\infty = 1.1$. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

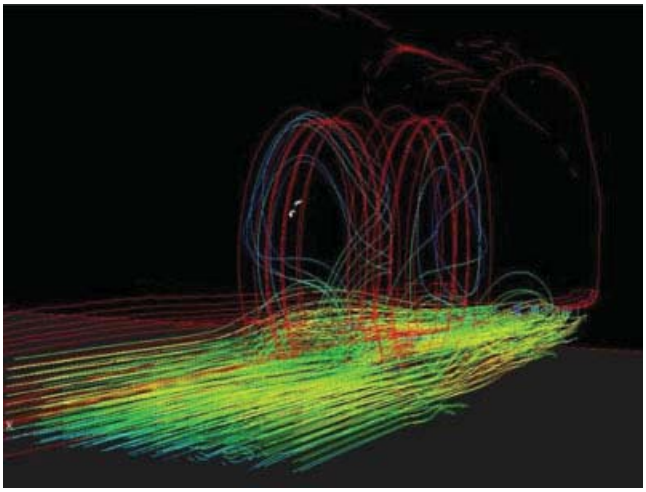

FIG. 15. Simulation with a parallel front wheel configuration showing velocity–colored particle traces at $M_\infty = 1.1$. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

C. Further Detailed Local Modelling

Other local model simulations have also been undertaken and have influenced the design of the full vehicle geometry. These include a local analysis of the winglet and tail designs, to ensure the optimum size is selected to satisfy vertical load control, in the case of the winglets, and satisfactory static margin and, hence, yaw stability across the Mach number range, in the case of the tail. Extensive jet engine intake duct modeling was undertaken to ensure that satisfactory flow is delivered to the engine compressor face across the Mach number range. This modeling guided the decision to convert from a twin intake design, see Fig. 1, to the current single intake design, which will be seen later in Fig. 17. Details such as the duct throat area and area distribution were also designed in this fashion. Finally, a parametric optimization study was undertaken focusing on the rear of the vehicle, including the rear body shape, suspension and wheel fairing. The aim of this study was to find a rear geometry configuration that satisfied both lift and drag targets for the vehicle rear across the Mach number range. Details of this study will be provided in a forthcoming paper (Evans et al. in preparation).

D. Full Vehicle Simulations

The understanding gained from the local CFD simulations was used to guide the construction of geometries for full vehicle simulations. These simulations enable the aerodynamic behavior of the complete vehicle to be predicted.

Full vehicle simulation meshes were typically undertaken using meshes of the order of 35 M elements, 10 M nodes and 46 M edges. A section through a typical hybrid mesh used for these simulations is shown in Fig. 16, where the transition from the quasi–structured boundary layer mesh to the isotropic Delaunay mesh is clearly apparent. The simulations were performed using 120 CPUs of the Swansea C2EC cluster, with convergence requiring 8 K cycles for supersonic problems and 20 K cycles for subsonic problems. With 6.5 cycles being computed each minute, this equates to wall clock run times of around 21 and 50 hs for supersonic and subsonic problems, respectively. Lift and drag coefficients were converged to an accuracy of two significant figures

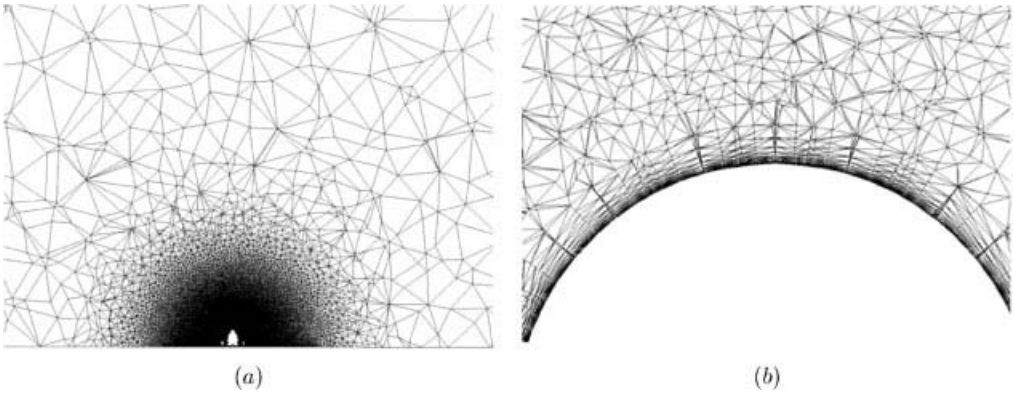

FIG. 16. Details of a cut through a typical hybrid mesh used for the simulation of flow over the complete BLOODHOUND SSC vehicle.

for these simulations, as this level of accuracy was deemed to be sufficient for identifying trends to guide the vehicle design.

Figure 17 shows the computed pressure coefficient distribution over the surface of BLOODHOUND configuration 9, from September 2009, at a free stream Mach number of 1.3. Also

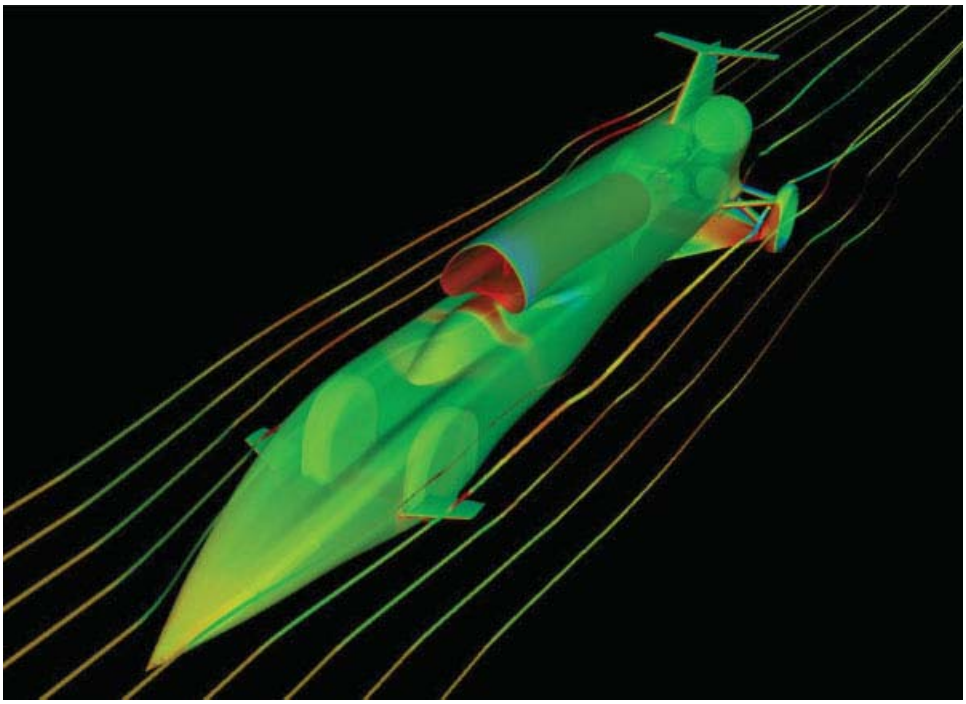

FIG. 17. Pressure coefficient distribution and velocity–coloured stream ribbons for a simulation of flow over the complete BLOODHOUND vehicle. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

displayed in this figure are a number of velocity–colored stream ribbons. Note the high pressure field above the driver's canopy, as a result of the twin–shock system formed upstream of the intake duct, the significant compression as the flow is decelerated within the intake duct and the shock system formed by the rear wheels causing a high pressure field on the top of the suspension strut. It is the presence of this large planform strut that generates the significant down force at the rear of the vehicle that is necessary to stop the vehicle from lifting off the desert surface at high speed.

VI. CONCLUSIONS

A fully viscous finite volume procedure, using the HLLC convective flux function and the Spalart–Allmaras turbulence model, on unstructured hybrid meshes has been employed to study the aerodynamic behavior of a supersonic Land Speed Record vehicle with rotating wheels and rolling ground. Local flow simulations, focusing on such aspects as the front wheel configuration, combined with full vehicle simulations have guided the aerodynamic design process of the BLOODHOUND supersonic car. At the time of writing (January 2010), the design phase of the vehicle is reaching its conclusion and the project is moving into the vehicle build phase.

References

1. R. Noble, Thrust, Transworld, London, 1998.
2. The Story of ThrustSSC, Corgi Books, London, 1998.
3. J. Peraire, J. Peiró, and K. Morgan, Multigrid solution of the 3D compressible Euler equations on unstructured tetrahedral grids, Int J Numer Methods Eng 36 (1993), 1029–1044.
4. D. J. Mavriplis and V. Venkatakrishnan, A 3D agglomeration multigrid solver for the Reynolds–averaged Navier–Stokes equations on unstructured meshes, Int J Numer Methods Fluids 23 (1996), 527–544.
5. O. Hassan, K. Morgan, E. J. Probert, and J. Peraire, Unstructured tetrahedral mesh generation for three–dimensional viscous flows, Int J Numer Methods Eng 39 (1996), 549–567.
6. N. P. Weatherill and O. Hassan, Efficient three–dimensional Delaunay triangulation with automatic boundary point creation and imposed boundary constraints, Int J Numer Methods Eng 37 (1994), 2005–2039.
7. K. Morgan and J. Peraire, Unstructured grid finite element methods for fluid mechanics, Reports Prog Phys 61 (1998), 569–638.
8. P. Batten, N. Clarke, C. Lambert, and D. M. Causon, On the choice of wavespeeds for the HLLC Riemann solver, SIAM J Scientific Comput 18 (1997), 1553–1570.
9. A. Harten, P. D. Lax, and B. Van Leer, On upstream differencing and Godunov–type schemes for hyperbolic conservation laws, SIAM Review 25 (1983), 35–61.
10. K. A. Sørensen, A multigrid accelerated procedure for the solution of compressible fluid flows on unstructured hybrid meshes, Ph.D. Thesis, University of Wales, Swansea, 2002.
11. D. C. Wilcox, Turbulence Modelling for CFD, DCW Industries Inc., 1998.
12. F. M. White, Viscous Fluid Flow, 3rd Ed., McGraw Hill, Boston, 2006.
13. P. R. Spalart and S. R. Allmaras, A one–equation turbulent model for aerodynamic flows, AIAA Paper 92–0439, 1992.
14. www.cadfix.com. Accessed 27 August 2010.
15. J. Peraire, M. Vahdati, K. Morgan, and O. C. Zienkiewicz, Adaptive remeshing for compressible flow simulations, J Comput Phys 72 (1987), 449–466.

16. J. Peiró, J. Peraire, and K. Morgan, FELISA system reference manual. Part 1–basic theory, University of Wales Swansea Report C/R/821/94, 1994.

17. J. Peiró, J. Peraire, and K. Morgan, The generation of triangular meshes on surfaces, C. Creasy and C. Craggs, editors, Appl Surf Model, Ellis–Horwood, Chichester, 1989, pp. 25–33.

18. K. A. Sørensen, O. Hassan, K. Morgan, and N. P. Weatherill, A multigrid accelerated hybrid unstructured mesh method for 3D compressible turbulent flow, Comput Mech 31 (2003), 101–114.

19. P. I. Crumpton, P. Moinier, and M. B. Giles, An unstructured algorithm for high Reynolds number flows on highly stretched grids, C. Taylor and J. T. Cross, editors, Numerical Methods in Laminar and Turbulent Flow, Pineridge Press, Swansea, 1997, pp. 561–572.

20. K. A. Sørensen, O. Hassan, K. Morgan, and N. P. Weatherill, A multigrid accelerated hybrid unstructured mesh method for 3D compressible turbulent flow, Comput Mech 31 (2003), 101–114.

21. W. J. Usab and E. M. Murman, Embedded mesh solutions of the Euler equations using a multiple–grid method, W. G. Habashi, editor, Advances in Computational Transonics, Pineridge Press, Swansea, 1985, pp. 447–472.

22. J. D. Anderson, Fundamentals of aerodynamics, McGraw–Hill, New York, 2001.

23. G. Karypis and V. Kumar, METIS4.0: Unstructured graph partitioning and sparse and sparse matrix ordering system, Technical Report, Department of Computer Science, University of Minnesota, 1998.