AIAA 2003–3531

Numerical Analysis and Design of Upwind Sails

Sriram and Antony Jameson
Dept of Aeronautics and Astronautics
Stanford University, Stanford, CA

Margot G Gerritsen
Department of Petroleum Engineering
Stanford University, Stanford, CA

21st AIAA Applied Aerodynamics Conference
June 23–26, 2003/Orlando, FL

Numerical Analysis and Design of Upwind Sails

Sriram and Antony Jameson *
Dept of Aeronautics and Astronautics
Stanford University, Stanford, CA

Margot G Gerritsen †
Department of Petroleum Engineering
Stanford University, Stanford, CA

The aim of this study is to develop and validate numerical methods that solve the inviscid field equations (Euler) to simulate and design upwind sails. The three dimensional Euler equations for compressible flow are modified using the idea of artificial compressibility and discretized on unstructured tetrahedral grids to provide estimates of lift and drag for upwind sail configurations. Convergence acceleration techniques like multigrid and residual averaging are used along with parallel computing platforms to enable these simulations to be performed in a few minutes. To account for the elastic nature of the sail cloth, this flow solver is coupled to NASTRAN to provide estimates of the deflections caused by the pressure loading. The results from the aeroelastic simulations show that the major effect of the sail elasticity was in altering the pressure distribution around the leading edge of the head and the main sail. Adjoint based design methods have also been developed, and used to induce changes to the camber distribution of the main sail. The design process resulted in a camber distribution that allows smooth entry of the flow through the leading edge of the main sail thereby, reducing the leading edge suction peaks that would be detrimental to the growth of the boundary layer.

1. Introduction

Races like the Americas Cup have seen significant improvements in the design of both the hull and the sails over the last two decades. Competing syndicates are constantly pushing the aerodynamic and structural limits of the designs as improvements of less than 0.5 % in the speed of the boat results in savings of about 25-35 seconds which is typically greater than the margin of victory for these races. For the windward leg of the race, a good measure of the performance of a design is the distance that the boat travels directly to windward in a given time. This performance index (called the speed-made-good by the boat) is dependent on both the speed of the boat and her true sailing course which in turn are dependent on the aerodynamic and hydrodynamic forces produced by the sails and the hull. In general, the performance of the boat can be improved by reducing the resistance of the hull and the drag of the sails. However changes in the aerodynamic and hydrodynamic forces alter the equilibrium of the sail boat which then has to adjust its speed and sailing course to maintain equilibrium. Hence, the design of sails boats has to be carried out in an environment where the analysis and design procedures for the sails and the hull are tightly coupled to realize a meaningful overall design.

Traditionally, designers have used Velocity Prediction Programs (VPP)to design a boat with good overall performance. These programs solve for the motion of the boat using force and moment balance relations while ensuring that the resulting motion is stable. These programs typically compute the aerodynamic and hydrodynamic forces using potential flow solvers, which are computationally inexpensive and easy to implement, enabling the designer to evaluate a wide array of designs. However, there exists large regions of rotational flow and significant viscous interaction, where the assumptions of potential flow are not valid.

Alternate descriptions of the flow field that account for the rotational nature of the flow can provide more accurate estimates of the induced drag (which accounts for 60 % of the total drag) for the upwind leg of the race. Modeling the viscous effects during the upwind and downwind leg of the race will help in the design of sails and hulls with reduced viscous drag. Simulations of the evolution of the free-surface of the sea and its interaction with the hull can help provide a virtual tool that can supplement tow tank testing. The development of robust and accurate tools to predict the above mentioned physical phenomena could lead to improved designs that can benefit the racing community in particular and the sailing world in general.

Computational techniques that solve the Euler or the Navier-Stokes equations are increasingly being used by competing syndicates in races like the Americas Cup. For sail configurations, this desire stems from a need to understand the influence of the mast on the boundary layer and pressure distribution on the main sail, the effect of camber and planform variations of the sails on the driving and heeling force produced by them and the interaction of the boundary layer profile of the air over the surface of the water and the gap between the boom and the deck on the performance

*Professor, Stanford University
†Professor, Stanford University

of the sail. Traditionally, experimental methods along with potential flow solvers have been widely used to quantify these effects. While these approaches are invaluable either for validation purposes or during the early stages of design, the potential advantages of high fidelity computational methods makes them attractive candidates during the later stages of the design process.

The aim of this study is to develop and validate numerical methods that solve the inviscid Euler field equations to simulate and design upwind sails. The three dimensional Euler equations for compressible flow are modified using the idea of artificial compressibility, and discretized on unstructured tetrahedral grids to provide estimates of lift and drag for upwind sail configurations. Convergence acceleration techniques like multigrid and residual averaging are used along with parallel computing platforms to enable these simulations to be performed in a few minutes. To account for the elastic nature of the sail cloth, this flow solver is coupled to NASTRAN to provide estimates of the deflections caused by the pressure loading. The results of this aeroelastic simulation, illustrate that the major effect of the sail elasticity, is in altering the pressure distribution around the leading edge of the head and the main sail.

We also address the issue of how to calculate improved sail shapes using an optimization procedure based on the adjoint method. This is used to induce changes to the camber distribution of the main sail with the goal of reducing the leading edge suction peaks that could be detrimental to the growth of the boundary layer. The design process results in an camber distribution that allows smooth entry of the flow through the leading edge of the main sail, thereby reducing the leading edge suction peaks.

2. Analysis with CFD

2.1 Finite Volume Discretization of the Flow Equations

A vast repertoire of computational codes have been developed by Antony Jameson to analyze aerodynamic configurations in transonic flight.[1,2] These codes model the fluid as a compressible fluid and a variety of numerical techniques have been developed to efficiently solve the governing equations of a compressible fluid with embedded supersonic regions.

In the limit of truly incompressible flow, or zero Mach number, alternate methods must be used to compute the flow field to preserve the accuracy, robustness and convergence properties of the flow solution. The fundamental problem in transitioning from a compressible fluid model to an incompressible one is the loss of the evolution equation for the density. Since the density is constant, a time dependent constraint must be imposed on the continuity equations to ensure a divergence-free velocity field. In addition, the eigenvalues resulting from the system of conventional hyperbolic Euler equations for compressible flows become infinite in the limit of incompressible flow. This is due to the fact that incompressible flows exhibit infinite sound speeds. Hence, the use of compressible flow solvers in the incompressible flow limit introduces widely varying eigen-speeds, resulting in extremely stiff equations.

To overcome this difficulty, the present work uses the artificial compressibility method, an approach first proposed by Chorin in 1967[3] as a method to solve viscous flows. Artificial compressibility methods introduce a psuedotemporal equation for the pressure through the continuity equation. This approach removes the troublesome sound waves associated with compressible flow formulations as the Mach number approaches zero. The eigenvalues of the original system are now replaced with an artificial set that renders the new set of equations well-conditioned for numerical computation. When combined with multigrid acceleration procedures, artificial compressibility proves to be particularly effective[4]. Converged solutions of incompressible flows over a main sail can be obtained in about 75-100 multigrid cycles.

Using the idea of artificial compressibility, the equations of motion of an incompressible, inviscid fluid can be cast in the following form.

$$\frac{\partial w}{\partial t} + P\left\{\frac{\partial F}{\partial x} + \frac{\partial G}{\partial y} + \frac{\partial H}{dz}\right\} = 0. \tag{1}$$

Here, the dependent variables w, the inviscid flux vectors f, g and h and the pre-conditioning matrix P are described by

$$w = \begin{Bmatrix} p \\ u \\ v \\ w \end{Bmatrix},\ F = \begin{Bmatrix} u \\ u^2+p \\ uv \\ uw \end{Bmatrix},\ G = \begin{Bmatrix} v \\ vu \\ v^2+p \\ vw \end{Bmatrix}, \tag{2}$$

$$H = \begin{Bmatrix} w \\ wu \\ wv \\ w^2+p \end{Bmatrix},\ P = \begin{bmatrix} \Gamma^2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{3}$$

This system of equations has no physical meaning until a steady state is reached. At steady state, the time dependent pressure term drops from the continuity equation resulting in the true steady state equations for an incompressible flow. Further, Γ can be selected to accelerate the time decay to steady state. Using the finite volume approach, the governing equations can be cast in the integral form for each computational volume in the domain as follows,

Conservation of Mass

$$\frac{d}{dt}\int_V p dV + \int_S \Gamma^2 (u \cdot n)\, dS = 0 \tag{4}$$

Conservation of Momentum

$$\frac{d}{dt}\int_V udV + \int_S u(u\cdot n)dS = -\int_S pndS \qquad (5)$$

Spatial discretization of equation (4) and (5) leads to the following equation for each point of interest in the computational mesh.

$$\frac{d}{dt}V_i w_i + \sum_k F_k.n_k S_k = 0 \qquad (6)$$

where p is the pressure, u is the velocity vector, n is the unit normal at the surface of the control volume, V and S are the volume and surface area of the control volume respectively, F is the flux through the control volume and the summation of the fluxes is over the control volume that surrounds each node of the mesh.

In the present study, a cell-vertex scheme is used for the implementation of the finite volume scheme on unstructured tetrahedra. Dual meshes constructed from planes bisecting each edge of the mesh are used to accumulate the fluxes at each node. Boundary conditions are then enforced along the triangular faces that lie on the boundary to account for the one-sided control volumes for the nodes on the boundary. The rest of the discussion in this section outlines the details of the implementation of the spatial discretization operators when used with artificial compressibility methods, the evaluation of the numerical diffusion terms and the multigrid algorithm.

In Equation (3), Γ is called the artificial compressibility parameter in the light of the analogy that may be drawn between the above equations and the equations of motion for a compressible fluid whose equation of state is given by $p = \Gamma^2\rho$. Thus, ρ is an artificial density and Γ may be referred to as an artificial speed of sound. When the temporal derivatives tend to zero, the set of equations satisfy precisely the incompressible Euler equations, with the consequence that the correct pressure may be established using the artificial compressibility formulation. The pre-conditioning matrix, P, may be viewed as a device to create a well posed system of hyperbolic equations that are to be integrated to steady state along lines similar to well established compressible flow Finite Volume formulations. In addition, the artificial compressibility parameter may be viewed as a relaxation parameter for the pressure iteration.

The eigenvalues of the system of equations in equation (1), are given by

$$\lambda_1 = U, \lambda_2 = U, \lambda_3 = U + a, \lambda_4 = U - a$$

where,

$$a^2 = U^2 + \Gamma^2(\psi^2 + \eta^2 + \xi^2)$$

and

$$U = u\psi + v\eta + w\xi$$

The terms ψ, η, ξ represent the slopes of the characteristic system of waves, are arbitrary and defined.

The choice of Γ is crucial in determining the convergence and stability properties of the numerical scheme. Typically, the convergence and stability rate of the scheme is dictated by the slowest waves and the stability of the scheme by the fastest. In the limit of large Γ, the difference in wave speeds can be large. Although this situation would presumably lead to a more accurate solution through the penalty effect in the pressure equation, very small time steps would be required to ensure stability. Conversely, for small Γ, the difference in the maximum and minimum wave speeds may be significantly reduced, but at the expense of accuracy. Thus a compromise between the extremes is achieved by choosing Γ to be

$$\Gamma^2 = C(u^2 + v^2 + w^2)$$

where C is a constant of the order of unity. In regions of high velocity and low pressure where suction occurs, Γ is large to improve accuracy, and in regions of low velocity, Γ is correspondingly reduced.

Under these assumptions on the choice of the preconditioner, P, and the application of the Finite Volume method for a cell-vertex scheme results in a set of ordinary differential equation for each node of the computational mesh,

$$\frac{d}{dt}(V_i w) + PQ_i = 0. \qquad (7)$$

where V_i is the volume around each node and Q_i represents the flux through the faces of the control volume. To prevent, odd-even decoupling at adjacent nodes which may lead to oscillatory solutions, a dissipation term is added to the flux calculation to modify the above equation

$$\frac{d}{dt}(V_i w) + P[Q_i - D_i] = 0. \qquad (8)$$

or

$$\frac{d}{dt}(V_i w) + R_i = 0. \qquad (9)$$

where R_i is the residual at each node in the computational mesh.

The resulting system of equations are integrated in time using an explicit multistage scheme with coefficients that maximize the stability region of the time-stepping scheme. To further accelerate convergence to steady state, local time-stepping and residual averaging techniques are introduced. Detailed numerical analysis of the spatial discretization[6,7] and the time stepping scheme[8] can be obtained from previous studies listed in the citations.

2.2 Dissipation

Numerous research efforts during the eighties led to to the development of a mathematical frame-work to

add numerical diffusion to the discretized equations. While the focus of these studies were on deriving stable and accurate techniques to resolve shocks in the flowfield, the mathematical frame-work can be inherited for incompressible flows that use the artificial compressibility method with some modifications that limit the amount of numerical diffusion.

Local Extremum Diminishing (LED) schemes that guarantee that new extrema are not generated during the evolution of the solution have proven to be robust and efficient. These schemes limit the reconstructed solution and fluxes at cell interfaces by using limiters that can be constructed from gradient information from a stencil of points around each computational point. A variety of choices exist for the form of the limiters, and the JST scheme [8] can be reformulated in the class of SLIP schemes for a particular choice of the limiter. Over the years the JST scheme has proven to be robust and accurate for a variety of aerodynamic flows. It can be represented as

$$d_{j+\frac{1}{2}} = \epsilon^2_{j+\frac{1}{2}} \Delta w_{j+\frac{1}{2}} - \epsilon^4_{j+\frac{1}{2}} \left(\Delta w_{j+\frac{3}{2}} - 2\Delta w_{j+\frac{1}{2}} + \Delta w_{j-\frac{1}{2}} \right)$$

When used for incompressible flows, the first term on the R.H.S. can be dropped and the remainder used as the numerical diffusion term that prevents odd-even coupling. All the computations in this study has been performed using the JST scheme as the basis for numerical dissipation.

2.3 Multigrid

Multigrid techniques are widely used to accelerate the convergence of a system of equations to steady state. A general framework for the development of full-approximation multigrid methods for non-linear equations can be outlined as follows.

Consider,

$$\mathcal{L}u = \mathcal{F}$$

discretized on a mesh with spacing h as

$$\mathcal{L}_h v_h = \mathcal{F}_h$$

This can be rewritten as

$$\mathcal{L}_h(v_h + \delta v_h) = \mathcal{F}_h$$

where δv represents correction to the present estimate v_h or

$$\mathcal{L}_h \delta v_h + R_h = 0 \tag{10}$$

where R_h is the residual. on the coarse grid, the above equation can be replaced by

$$\mathcal{L}_{2h} \delta v_{2h} + I^h_{2h} R_h = 0 \tag{11}$$

where I^h_{2h} represents the aggregation or restriction operator. The correction to the present fine grid solution can be represented as

$$v^{new}_h = v_h + I^{2h}_h \delta v_h$$

where I^{2h}_h represents an interpolation operator. We can add and subtract the following from equation (11)

$$\mathcal{L}_{2h} v_{2h} - \mathcal{F}_{2h} = R_{2h}$$

to get

$$\mathcal{L}_{2h}(v_{2h} + \delta v_{2h}) - \mathcal{F}_{2h} + I^h_{2h} R_h = 0$$

This leads to the full approximation scheme (FAS)

$$\mathcal{L}_{2h}(v^+_{2h}) - \mathcal{F}_{2h} + I^h_{2h} R_h - R_{2h} = 0$$

Then

$$v^+_h = v_h + I^h_{2h}(v^+_{2h} - v_{2h})$$

In the present work a series of non-nested meshes are used for the multigrid cycle and the solution and residual from each mesh are aggregated to the coarser mesh while interpolating the correction from the coarser to the fine mesh. Detailed descriptions of the multigrid scheme can be obtained from[9].

For unstructured grids, the nature of the grids to be used in the multigrid cycle is a question of ongoing debate. In the present study a series of non-nested meshes was generated by a grid generator (MESHPLANE). The initial solution from a particular mesh is advanced in pseudo-time to obtain new estimates of the flow variables. On transfer to the next level of the multigrid, the solution for the coarse grid mesh points is interpolated from the four nodes of the fine mesh tetrahedron that contains this node. Further, the accumulated residual at each fine mesh point is distributed to each node of the tetrahedron in the coarse mesh that that encloses the fine mesh node. The interpolating factors for each node are computed from weights which are based on the volume included by a given node and opposite face of the tetrahedron (figure 1). This reduces to a second order interpolation scheme on equilateral tetrahedra and has been found to be sufficient for the present calculations. The estimate of the residuals from the fine mesh is used to advance the solution on the coarse mesh where the larger scale errors can be more efficiently accomplished on the coarse mesh. Further levels in the multigrid cycle involve the same operations are before, thereby using grids that are coarser and coarser to convect the error terms out of the computational domain faster. The ascending side of the multigrid cycle estimates a correction from each grid which is then interpolated to the next finer mesh in the sequence (figure 2). The corrections from the coarser mesh are transferred using

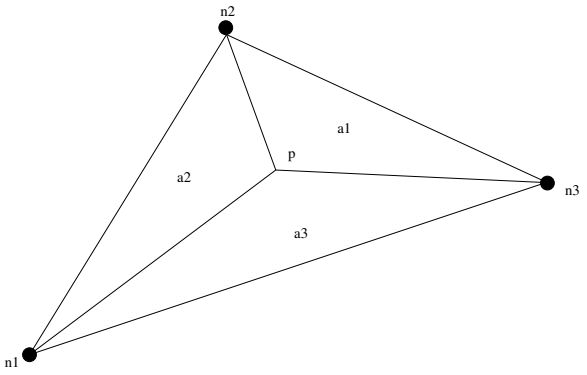

Fig. 1 Interpolation coefficients for use in the multigrid cycle

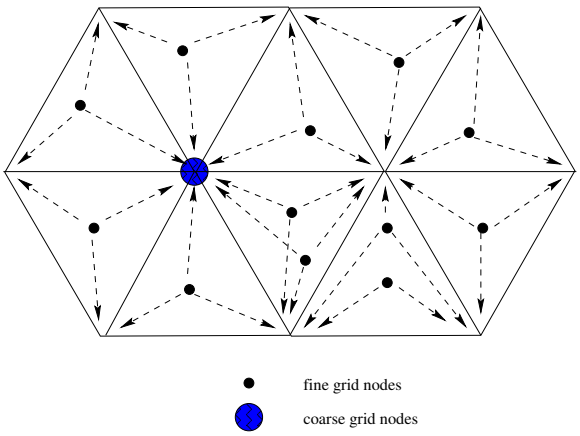

Fig. 2 Transfer of solution, residuals and corrections between the fine and coarse mesh

similar interpolating factors as for the aggregation operations. Multigrid cycles which progress in the shape of a W have generally been found to provide faster convergence to steady state than the simple V cycle.

3. Parallel Implementation of the Unstructured Multigrid Flow Solver

To exploit the availability of modern parallel computing platforms, a combination of the AIRPLANE[1] and FLO77[10] computational programs, was parallelized. Due to the unstructured nature of the computational grid, a wide variety of possible data structures to implement the underlying numerical algorithms exist. The following sections outline the choice of data structures and algorithms that were made to parallelize the flow solver.

3.1 Domain Decomposition, Load Balancing

A modified coordinate bisection method was used to recursively divide a given computational mesh into sub-domains (figure 3). The sub-domains were created such that they contained approximately the same number of computational nodes. No effort was made to optimize the domain decomposition process so as to minimize the number of edges that are shared between

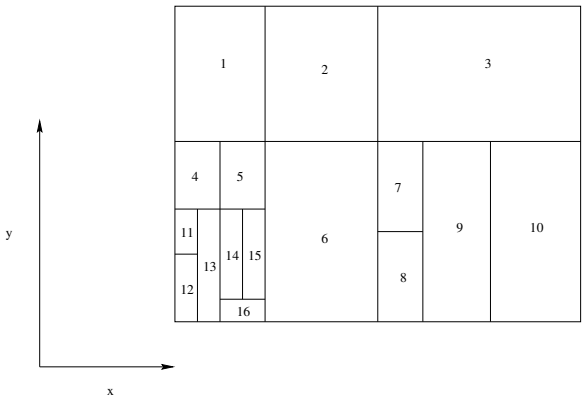

Fig. 3 Domain decomposition of a rectangular region using a modified bisection method

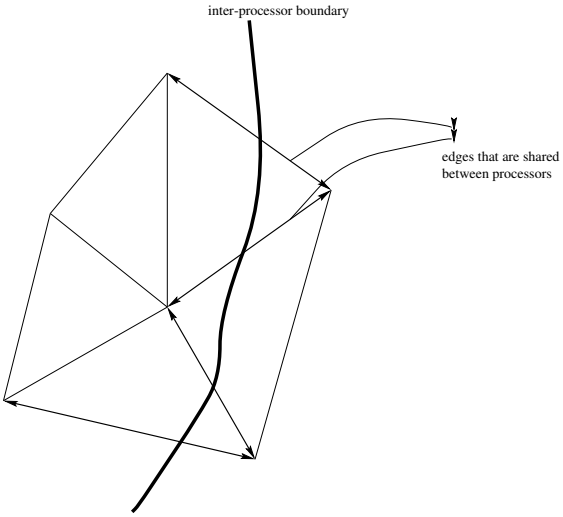

Fig. 4 Halo nodes and the distribution of edges along processor boundaries

the sub-domains. The sub-domains are distributed among the available processors in a manner that minimizes a combination of the computational cost associated with the domains in each processor and the cost of communication among the processors for a given distribution. This methodology was found to result in well balanced load distributions.

As the flow solver uses an edge-based data structure to accumulate the fluxes at each vertex, the edges surrounding the nodes that lie within a partition are accumulated. If an edge connects a pair of nodes that lie across processor boundaries, this edge is duplicated in the two processors and 'halo' nodes are constructed for both processors. This idea is illustrated further in figure 4.

3.2 Parallel Implementation of the Multigrid Algorithm

The use of multigrid techniques for flow analysis necessitates the need to exchange information between fine and coarse grid points and vice-versa. While the

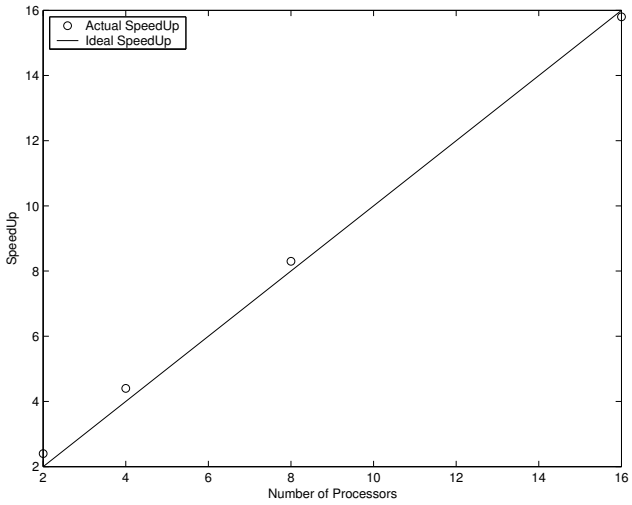

Fig. 5 Speedup from the parallel implementation

residuals from the fine grid points need to be accumulated at the coarse grid points, the corrections from the coarse grid points need to be transferred back to the fine grid points. Hence, efficient point location methods are used to determine the interpolation coefficients associated with each computational node in the multigrid cycle. The required search methodology has been implemented using an octree-based search routine. Each grid in the multigrid cycle is recursively divided into octants that contain a certain number of points. Using this data structure, a given point is identified within an octant and the closest node closest in this octant is determined. The tetrahedra connected to this node are checked to see if they contain the search point. Once such a tetrahedron is identified, interpolation coefficients to aggregate the residuals and to interpolate the corrections are constructed. Further, to reduce the communication cost among processors during the transfer of information between the fine and coarse grids, sub-domains on the coarser grids are constructed and distributed so as to conform to the division and distribution of the fine mesh. Typical computational times to build the octree and to build the interpolation tables are in the range of a few minutes for a sequence of meshes containing a million nodes. The efficiency of the parallel implementation is shown in figure 5.

4. Aeroelastic Computations to Predict the Flying Shape of Upwind Sails

The incompressible flow solver described in the previous sections has been used to obtain the characteristics of sail configurations of high performance yachts for races like the Americas Cup. Modifications to the inlet boundary condition were incorporated so as to simulate the boundary layer profile over the sea. Figure 7 shows a representative head and main sail combination used in the Americas Cup. Due to the fast turn-around times of the flow solver it was possible to use this computational package to obtain aerodynamic characteristics of the sail configurations for a variety of wind conditions, and also study the effect of varying twist, camber and sail trim. Further, in order to predict the flying shape of the sail geometries, this flow solver was coupled to NASTRAN. The aeroelastic package uses an iterative algorithm that transfers the pressure loading obtained from the flow solver to the structural model, and uses the deflections from the structural analysis to modify the computational mesh for the fluid. This process is iteratively carried out until the deflections are below a particular threshold.

4.1 Structural Model

In the structural model of the sail the sail cloth is discretized into quadrilateral membrane finite elements with four nodes (after neglecting the presence of batten pockets). These elements withstand all external forces through tension and are incapable of resisting bending moments. The translational and rotational degrees of freedom along the foot of the main sail are suppressed. Along the mast, the translational degrees of freedom are inhibited while allowing for rotational motion. For the head sail, the point of attachment of the foot to the rig is constrained. The leech of the main and head sail are allowed to move freely to induce a geometric twist due to the aerodynamic loading. The mast is assumed to be rigid during the structural and aeroelastic calculations. The presence of battens and tension cables and other structural elements of the sail rig is neglected from this analysis.

The linear system of equations relating the displacements to the force field is advanced to a steady state by an iterative process that incrementally adds the load while obtaining a converged displacement field for each step. This procedure was introduced in order to allow for large deflections of the sail geometry. Wrinkling of the structure, which is an important consideration especially around the leading edge (luff) and at the sail tip, is not anticipated by this model, but the use of a numerical algorithm to track large deformations allows for wrinkling models to be included at a later stage. Figures 13 and 14 pictorially depict the boundary conditions used for the structural model.

4.2 Aeroelastic Coupling Procedure

The pressure loading from the flow solver is fed to the structural analysis to estimate the deflected shape of the sail. To enable the transfer of loads and displacements to be conservative, the fluid mesh on the surface and the structural mesh are identical, eliminating the need for interpolation. The deflected shape of the sail is used to deform the computational mesh. The popular 'spring-analogy' method has been used to track the mesh deformations. While this method restricts the allowable deflections and may impair the quality of the deformed mesh, it provides a simple tool to track

mesh deformations. The deformed mesh is then used to compute a new pressure loading for the sail. This iterative process offers no formal guarantee of convergence, but it usually predicts the deflected shape to reasonable accuracy in a few steps (typically 5 for sail geometries).

4.3 Results of Aeroelastic Simulations

The deflected shapes of the head and the main sail are shown in figures 19 and 20. It can be seen from these plots that the lower sections of the head and the main sail do not undergo appreciable deformation. The largest deflections occur in the mid-sections of the main sail. As the point of attachment of the main sail to the mast and the leading edge of the head sail were not allowed to move, the aerodynamic loading changes the twist of the sail geometry. This has a favorable influence on the pressure distribution, especially on the head sail (figures 17 and 18). The pressure distribution over the head and sail after the aeroelastic simulation highlights the need to perform aeroelastic analysis to obtain the flying shape of these sail geometries. While the lift and the drag distribution of the deformed shape is not significantly different from the undeformed shape, the pressure distribution over the sail sections show that the twist and the camber distribution of the head and the main sail have been altered to provide a smooth entry for the flow over the leading edge of both components.

5. Aerodynamic Shape Optimization

With the availability of high performance computing platforms and robust numerical methods to simulate fluid flows, it is possible to shift attention to automated design procedures which combine CFD with gradient-based optimization techniques. Typically, in gradient-based optimization techniques, a control function to be optimized (the sail shape, for example) is parameterized with a set of design variables and a suitable cost function to be minimized is defined. For aerodynamic problems, the cost function is typically lift, drag or a specified target pressure distribution. Then, a constraint, the governing equations can be introduced in order to express the dependence between the cost function and the control function. The sensitivity derivatives of the cost function with respect to the design variables are calculated in order to get a direction of improvement. Finally, a step is taken in this direction and the procedure is repeated until convergence is achieved. Finding a fast and accurate way of calculating the necessary gradient information is essential to developing an effective design method since this can be the most time consuming portion of the design process. This is particularly true in problems which involve a very large number of design variables as is the case in typical three dimensional shape optimization.

The control theory approach[11–13] has dramatic computational cost advantages over the finite-difference method of calculating gradients. With this approach the necessary gradients are obtained through the solution of an adjoint system of equations of the governing equations of interest. The adjoint method is extremely efficient since the computational expense incurred in the calculation of the complete gradient is effectively independent of the number of design variables.

In this study, a continuous adjoint formulation has been used to derive the adjoint system of equations. Hence, the adjoint equations are derived directly from the governing equations and then discretized. Hence, this approach has the advantage over the discrete adjoint formulation in that the resulting adjoint equations are independent of the form of discretized flow equations. The adjoint system of equations have a similar form to the governing equations of the flow and hence the numerical methods developed for the flow equations[1, 8, 14] can be reused for the adjoint equations. The gradient is derived directly from the adjoint solution and the surface motion independent of the mesh modification. This is critical for this design methodology to work on unstructured meshes. If the gradient depends on the form of the mesh modification, then the field integral in the gradient calculation has to be recomputed for mesh modifications corresponding to each design variable. To reduce the computational cost with this approach,[15–17] the number of design variables would have to reduced by parameterizing the geometry. However, this reduced set of design variables would be incapable of recovering all possible shape variations. Using the gradients computed with this new formulation, a steepest descent method is used to improve an existing design.

5.1 The General Formulation of the Adjoint Approach to Optimal Design

For flow about an airfoil, or wing, the aerodynamic properties which define the cost function are functions of the flow-field variables, w, and the physical location of the boundary, which may be represented by the function, $\mathcal{F}$, say. Then

$$I = I(w, \mathcal{F}),$$

and a change in $\mathcal{F}$ results in a change

$$\delta I = \frac{\partial I^T}{\partial w}\delta w + \frac{\partial I^T}{\partial \mathcal{F}}\delta \mathcal{F}, \tag{12}$$

in the cost function. Using control theory, the governing equations of the flow field are introduced as a constraint in such a way that the final expression for the gradient does not require re-evaluation of the flow-field. In order to achieve this, δw must be eliminated from equation (12). Suppose that the governing equation R which expresses the dependence of w and $\mathcal{F}$ within the flow field domain D can be written as

$$R(w, \mathcal{F}) = 0 \tag{13}$$

Then δw is determined from the equation

$$\delta R = \left[\frac{\partial R}{\partial w}\right] \delta w + \left[\frac{\partial R}{\partial \mathcal{F}}\right] \delta \mathcal{F} = 0 \tag{14}$$

Next, introducing a Lagrange Multiplier ψ, we have

$$\delta I = \frac{\partial I^T}{\partial w}\delta w + \frac{\partial I^T}{\partial \mathcal{F}}\delta \mathcal{F} - \psi^T \left(\left[\frac{\partial R}{\partial w}\right] \delta w + \left[\frac{\partial R}{\partial \mathcal{F}}\right] \delta \mathcal{F} \right)$$

$$\delta I = \left(\frac{\partial I^T}{\partial w} - \psi^T \left[\frac{\partial R}{\partial w}\right] \partial w \right) \delta w + \left(\frac{\partial I^T}{\partial \mathcal{F}} - \psi^T \left[\frac{\partial R}{\partial \mathcal{F}}\right] \delta \mathcal{F} \right)$$

Choosing ψ to satisfy the adjoint equation

$$\left[\frac{\partial R}{\partial w}\right]^T \psi = \frac{\partial I}{\partial w} \tag{15}$$

the first term is eliminated and we find that

$$\delta I = \mathcal{G} \delta \mathcal{F} \tag{16}$$

where

$$\mathcal{G} = \frac{\partial I^T}{\partial \mathcal{F}} - \psi^T \left[\frac{\partial R}{\partial \mathcal{F}}\right] \tag{17}$$

This process allows for elimination of the terms that depend on the flow solution with the result that the gradient with respect with an arbitrary number of design variables can be determined without the need for additional flow field evaluations.

After taking a step in the negative gradient direction, the gradient is recalculated and the process repeated to follow the path of steepest descent until a minimum is reached. In order to avoid violating constraints, such as the minimum acceptable wing thickness, the gradient can be projected into an allowable subspace within which the constraints are satisfied. In this way one can devise procedures which must necessarily converge at least to a local minimum and which can be accelerated by the use of more sophisticated descent methods such as conjugate gradient or quasi-Newton algorithms. There is a possibility of more than one local minimum, but in any case this method will lead to an improvement over the original design.

5.2 Design Using the Euler Equations

The application of control theory to aerodynamic design problems is illustrated in this section for the case of three-dimensional wing design using the compressible Euler equations as the mathematical model. It proves convenient to denote the Cartesian coordinates and velocity components by x_1, x_2, x_3 and u_1, u_2, u_3, and to use the convention that summation over $i = 1$ to 3 is implied by a repeated index i. Then, the three-dimensional Euler equations may be written as

$$\frac{\partial w}{\partial t} + \frac{\partial f_i}{\partial x_i} = 0 \quad \text{in } D, \tag{18}$$

where

$$w = \begin{Bmatrix} \rho \\ \rho u_1 \\ \rho u_2 \\ \rho u_3 \\ \rho E \end{Bmatrix}, \quad f_i = \begin{Bmatrix} \rho u_i \\ \rho u_i u_1 + p\delta_{i1} \\ \rho u_i u_2 + p\delta_{i2} \\ \rho u_i u_3 + p\delta_{i3} \\ \rho u_i H \end{Bmatrix} \tag{19}$$

and δ_{ij} is the Kronecker delta function. Also,

$$p = (\gamma - 1)\,\rho \left\{ E - \frac{1}{2}\left(u_i^2\right) \right\}, \tag{20}$$

and

$$\rho H = \rho E + p \tag{21}$$

where γ is the ratio of the specific heats.

Consider a transformation to coordinates ξ_1, ξ_2, ξ_3 where

$$K_{ij} = \left[\frac{\partial x_i}{\partial \xi_j}\right], \quad J = \det(K), \quad K_{ij}^{-1} = \left[\frac{\partial \xi_i}{\partial x_j}\right],$$

and

$$S = JK^{-1}.$$

The elements of S are the cofactors of K, and in a finite volume discretization they are just the face areas of the computational cells projected in the x_1, x_2, and x_3 directions. Using the permutation tensor ϵ_{ijk} we can express the elements of S as

$$S_{ij} = \frac{1}{2}\epsilon_{jpq}\epsilon_{irs}\frac{\partial x_p}{\partial \xi_r}\frac{\partial x_q}{\partial \xi_s}. \tag{22}$$

Then

$$\begin{aligned} \frac{\partial}{\partial \xi_i} S_{ij} &= \frac{1}{2}\epsilon_{jpq}\epsilon_{irs}\left(\frac{\partial^2 x_p}{\partial \xi_r \partial \xi_i}\frac{\partial x_q}{\partial \xi_s} + \frac{\partial x_p}{\partial \xi_r}\frac{\partial^2 x_q}{\partial \xi_s \partial \xi_i} \right) \\ &= 0. \end{aligned} \tag{23}$$

Now, multiplying equation (18) by J and applying the chain rule,

$$J\frac{\partial w}{\partial t} + R(w) = 0 \tag{24}$$

where

$$R(w) = S_{ij}\frac{\partial f_j}{\partial \xi_i} = \frac{\partial}{\partial \xi_i}\left(S_{ij} f_j\right), \tag{25}$$

using (23). We can write the transformed fluxes in terms of the scaled contravariant velocity components

$$U_i = S_{ij} u_j$$

as

$$F_i = S_{ij} f_j = \begin{bmatrix} \rho U_i \\ \rho U_i u_1 + S_{i1} p \\ \rho U_i u_2 + S_{i2} p \\ \rho U_i u_3 + S_{i3} p \\ \rho U_i H \end{bmatrix}.$$

Assume now that the new computational coordinate system conforms to the wing in such a way that the wing surface B_W is represented by $\xi_2 = 0$. Then the flow is determined as the steady state solution of equation (24) subject to the flow tangency condition

$$U_2 = 0 \quad \text{on } B_W. \tag{26}$$

At the far field boundary B_F, conditions are specified for incoming waves, as in the two-dimensional case, while outgoing waves are determined by the solution.

The weak form of the Euler equations for steady flow can be written as

$$\int_{\mathcal{D}} \frac{\partial \phi^T}{\partial \xi_i} F_i d\mathcal{D} = \int_{\mathcal{B}} n_i \phi^T F_i d\mathcal{B}, \tag{27}$$

where the test vector ϕ is an arbitrary differentiable function and n_i is the outward normal at the boundary. If a differentiable solution w is obtained to this equation, it can be integrated by parts to give

$$\int_{\mathcal{D}} \phi^T \frac{\partial F_i}{\partial \xi_i} d\mathcal{D} = 0$$

and since this is true for any ϕ, the differential form can be recovered. If the solution is discontinuous (27) may be integrated by parts separately on either side of the discontinuity to recover the shock jump conditions.

Suppose now that it is desired to control the surface pressure by varying the wing shape. For this purpose, it is convenient to retain a fixed computational domain. Then variations in the shape result in corresponding variations in the mapping derivatives defined by K. As an example, consider the case of an inverse problem, where we introduce the cost function

$$I = \frac{1}{2} \iint_{B_W} (p - p_d)^2 \, d\xi_1 d\xi_3,$$

where p_d is the desired pressure. The design problem is now treated as a control problem where the control function is the wing shape, which is to be chosen to minimize I subject to the constraints defined by the flow equations (24). A variation in the shape will cause a variation δp in the pressure and consequently a variation in the cost function

$$\delta I = \iint_{B_W} (p - p_d) \, \delta p \;\; d\xi_1 d\xi_3. \tag{28}$$

Since p depends on w through the equation of state (20–21), the variation δp can be determined from the variation δw. Define the Jacobian matrices

$$A_i = \frac{\partial f_i}{\partial w}, \quad C_i = S_{ij} A_j. \tag{29}$$

The weak form of the equation for δw in the steady state becomes

$$\int_{\mathcal{D}} \frac{\partial \phi^T}{\partial \xi_i} \delta F_i d\mathcal{D} = \int_{\mathcal{B}} (n_i \phi^T \delta F_i) d\mathcal{B},$$

where

$$\delta F_i = C_i \delta w + \delta S_{ij} f_j,$$

which should hold for any differential test function ϕ. This equation may be added to the variation in the cost function, which may now be written as

$$\begin{aligned} \delta I = & \iint_{B_W} (p - p_d) \, \delta p \;\; d\xi_1 d\xi_3 \\ & - \int_{\mathcal{D}} \left(\frac{\partial \phi^T}{\partial \xi_i} \delta F_i \right) d\mathcal{D} \\ & + \int_{\mathcal{B}} \left(n_i \phi^T \delta F_i \right) d\mathcal{B}. \end{aligned} \tag{30}$$

On the wing surface B_W, $n_1 = n_3 = 0$. Thus, it follows from equation (26) that

$$\delta F_2 = \begin{bmatrix} 0 \\ S_{21} \delta p \\ S_{22} \delta p \\ S_{23} \delta p \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ \delta S_{21} p \\ \delta S_{22} p \\ \delta S_{23} p \\ 0 \end{bmatrix}. \tag{31}$$

Since the weak equation for δw should hold for an arbitrary choice of the test vector ϕ, we are free to choose ϕ to simplify the resulting expressions. Therefore we set $\phi = \psi$, where the co-state vector ψ is the solution of the adjoint equation

$$\frac{\partial \psi}{\partial t} - C_i^T \frac{\partial \psi}{\partial \xi_i} = 0 \quad \text{in } D. \tag{32}$$

At the outer boundary incoming characteristics for ψ correspond to outgoing characteristics for δw. Consequently one can choose boundary conditions for ψ such that

$$n_i \psi^T C_i \delta w = 0.$$

Then, if the coordinate transformation is such that δS is negligible in the far field, the only remaining boundary term is

$$- \iint_{B_W} \psi^T \delta F_2 \; d\xi_1 d\xi_3.$$

Thus, by letting ψ satisfy the boundary condition,

$$S_{21}\psi_2 + S_{22}\psi_3 + S_{23}\psi_4 = (p - p_d) \quad \text{on } B_W, \tag{33}$$

we find finally that

$$\begin{aligned} \delta I \quad = & - \int_{\mathcal{D}} \frac{\partial \psi^T}{\partial \xi_i} \delta S_{ij} f_j d\mathcal{D} \\ & - \iint_{B_W} (\delta S_{21}\psi_2 + \delta S_{22}\psi_3 + \delta S_{23}\psi_4) \, p \, d\xi_1 d\xi_3. \end{aligned} \tag{34}$$

Here the expression for the cost variation depends on the mesh variations throughout the domain which appear in the field integral. However, the true gradient for a shape variation should not depend on the way in which the mesh is deformed, but only on the true flow solution. In the next section, we show how the field integral can be eliminated to produce a reduced gradient formula which depends only on the boundary movement.

5.3 Adjoint Equations for the Euler Equations Modified by the Artificial Compressibility Method

Although the adjoint equation represents a linear set of partial differential equations for the adjoint variables, they are of the same form of the flow equations. The numerical solution procedures developed for the flow equations are applied to the adjoint system with the appropriate boundary conditions. The adjoint co-state flux terms are modified to account for the introduction of the artificial compressibility terms in the governing flow equations. The methodology followed here is derived from the work of Luigi Martinelli and G. Cowles.[20] The adjoint field equations can be expressed as a time dependent system of the form,

$$\frac{\partial \psi}{\partial t} - [A_i]^T \frac{\partial \psi}{\partial x_i} = 0 \tag{35}$$

where

$$\psi = \left\{ \begin{array}{c} p \\ \phi_1 \\ \phi_2 \\ \phi_3 \end{array} \right\} \tag{36}$$

Hence, this system can be integrated to steady state using a pre-conditioner similar to that used in the method of artificial compressibility. The adjoint 'continuity' equation is augmented by a time derivative of the adjoint pressure p.

$$\frac{\partial p}{\partial t} - \Gamma^2 \frac{\partial \phi_i}{\partial x_i} = 0 \tag{37}$$

The form of Γ is identical to that used for the flow equations since the magnitude of the eigenvalues of the flux Jacobians for the two systems are identical. Together with equation (37), the adjoint system is discretized and solved in a manner that is consistent with that used for the flow equation.

5.4 Reduced Gradient Formulations

Continuous adjoint formulations have generally used a form of the gradient that depends on the manner in which the mesh is modified for perturbations in each design variable. To represent all possible shapes the control surface should be regarded as a free surface. If the surface mesh points are used to define the surface, this leaves the designer with a thousands of design variables. On an unstructured mesh evaluating the gradient by perturbing each design variable in turn, would be prohibitively expensive because of the need to determine corresponding perturbations of the entire mesh. This would inhibit the use of this design tool in any meaningful design process.

In order to avoid this difficulty an alternate formulation to the gradient calculation is followed in this study. This idea was developed by Jameson and Sangho Kim[18] and was validated for two and three dimensional problems with structured grids. However, as it is possible to devise mesh modification routines that are computationally cheap on structured grids, the major benefit of this alternate gradient formulation is for general three dimensional unstructured grids. To complete the formulation of the control theory approach to shape optimization, the gradient formulations are outlined next. The formulation for the reduced gradients in the continuous limit is presented in the context of transformation between the physical domain and the computational domain, and is easily extended to unstructured grid methods where these transformations are not explicitly used.

The evaluation of the field integral in equation (34) requires the evaluation of the metric variations δS_{ij} throughout the domain. However, the true gradient should not depend on the way the mesh is modified.

Consider the case of a mesh variation with a fixed boundary. Then,

$$\delta I = 0$$

but there is a variation in the transformed flux,

$$\delta F_i = C_i \delta w + \delta S_{ij} f_j.$$

Here the true solution is unchanged. Thus, the variation δw is due to the mesh movement δx at fixed boundary configuration. Therefore

$$\delta w = \nabla w \cdot \delta x = \frac{\partial w}{\partial x_j} \delta x_j \, (= \delta w^*)$$

and since

$$\frac{\partial}{\partial \xi_i} \delta F_i = 0,$$

it follows that

$$\frac{\partial}{\partial \xi_i} \left(\delta S_{ij} f_j \right) = - \frac{\partial}{\partial \xi_i} \left(C_i \delta w^* \right). \tag{38}$$

It is verified by Jameson and Sangho Kim[18] that this relation holds in the general case with boundary movement. Now

$$\begin{aligned} \int_{\mathcal{D}} \psi^T \delta R d\mathcal{D} &= \int_{\mathcal{D}} \psi^T \frac{\partial}{\partial \xi_i} C_i \left(\delta w - \delta w^* \right) d\mathcal{D} \\ &= \int_{\mathcal{B}} \psi^T C_i \left(\delta w - \delta w^* \right) d\mathcal{B} \\ &- \int_{\mathcal{D}} \frac{\partial \psi^T}{\partial \xi_i} C_i \left(\delta w - \delta w^* \right) d\mathcal{D}. \end{aligned} \tag{39}$$

Here on the wall boundary

$$C_2\delta w = \delta F_2 - \delta S_{2j} f_j. \tag{40}$$

Thus, by choosing ψ to satisfy the adjoint equation and the adjoint boundary condition, we have finally the reduced gradient formulation that

$$\begin{aligned}\delta I \quad &= \int_{\mathcal{B}_\mathcal{W}} \psi^T \left(\delta S_{2j} f_j + C_2 \delta w^*\right) d\xi_1 d\xi_3 \\ &- \iint_{B_W} (\delta S_{21}\psi_2 + \delta S_{22}\psi_3 + \ \delta S_{23}\psi_4)\, p\, d\xi_1 d\xi_3 \end{aligned} \tag{41}$$

5.5 The need for a Sobolev inner product in the definition of the gradient

Another key issue for successful implementation of the continuous adjoint method is the choice of an appropriate inner product for the definition of the gradient. It turns out that there is an enormous benefit from the use of a modified Sobolev gradient, which enables the generation of a sequence of smooth shapes.

When the metric perturbations, δS_{ij}, are related to the surface motion, equation (41) finally yields the cost variation in the form of an inner product defined over the surface

$$\delta I = (\mathcal{G}, \delta\mathcal{F}) = \int_B \mathcal{G}\delta\mathcal{F} d\xi_B$$

where $\delta\mathcal{F}$ denotes the surface displacement. Then the update

$$\mathcal{F}^{n+1} = \mathcal{F}^n - \lambda\mathcal{G}^n$$

would result in an improvement

$$\delta I = -\lambda\left(\mathcal{G}, \mathcal{G}\right)) \leq 0$$

It turns out, however, that the gradient is generally a function in a lower smoothness class than the initial shape, with the result that there is a progressive loss of smoothness with each iteration.

This can be corrected by introducing a modified gradient which corresponds to a weighted Sobolev inner product, of the form

$$\langle u, v\rangle = \int (uv + \epsilon u' v')d\xi$$

This is equivalent to replacing $\mathcal{G}$ by $\overline{\mathcal{G}}$ where in one dimension

$$\overline{\mathcal{G}} - \frac{\partial}{\partial\xi}\epsilon\frac{\partial\overline{\mathcal{G}}}{\partial\xi} = \mathcal{G},$$

with $\overline{\mathcal{G}} = 0$ at the end points and making a shape change

$$\delta\mathcal{F} = -\lambda\overline{\mathcal{G}}$$

This both preserves the smoothness of the redesigned shape, and acts as a preconditioner which reduces the number of design steps needed to reach an optimum solution.

6. Mesh Deformation

The modifications to the shape of the boundary are transferred to the volume mesh using the spring method. This approach has been found to be adequate for the computations performed in this study.

The spring method can be mathematically conceptualized as solving the following equation

$$\frac{\partial\Delta x_i}{\partial t} + \sum_{j=1}^{N} K_{ij}(\Delta x_i - \Delta x_j) = 0$$

where the K_{ij} is the stiffness of the edge connecting node i to node j and its value is inversely proportional to the length of this edge, Δx_i is the displacement of node i and Δx_j is the displacement of node j, the opposite end of the edge. The position of static equilibrium of the mesh is computed using a Jacobi iteration with known initial values for the surface displacements.

7. Results

The overall design process is illustrated in figure 6.

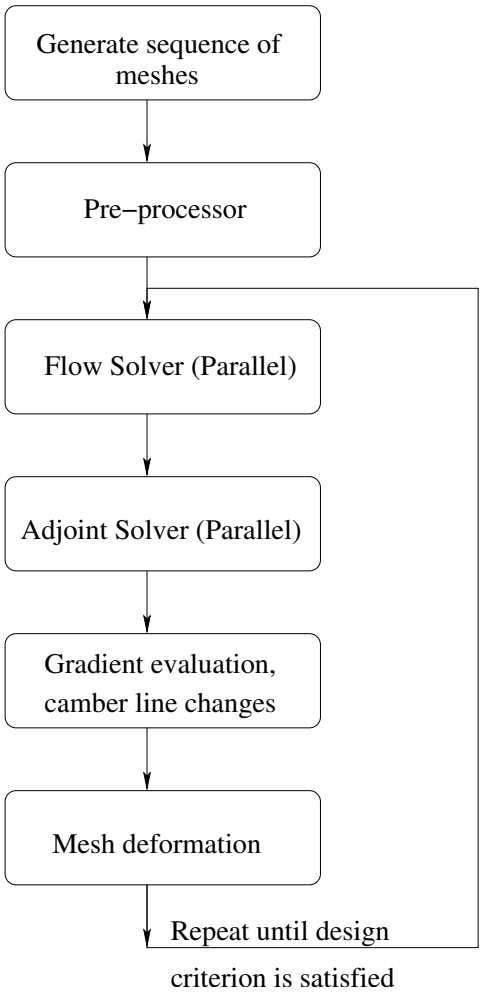

Fig. 6 Flow chart of the overall design process

7.1 Shape Optimization of Airfoils in Incompressible Flow

In order to validate the design procedure, two dimensional problems were studied first. An inverse design problem that recovers the pressure distribution over the Onera M6 airfoil was used to validate the gradient formulations and the adjoint solution. A three level multigrid cycle was used for obtain steady state solutions for the flow and the adjoint equations. The grids were generated using a conformal mapping technique. The initial airfoil shape had a NACA 0012 profile and the initial pressure distribution is shown in figure 21. Around 40 design cycles were required to recover the target pressure and shape by the design process (figure 21).

7.2 Inverse Design of Wings in Incompressible Flow

To validate the design process for three dimensional flows, a test problem similar to the two dimensional case was used. The initial wing had the planform of the Onera M6 but had NACA 0012 airfoil sections. The target pressure distribution corresponded to the steady state pressure distribution over the Onera M6 wing. Three levels of multigrid were used to obtain steady state flow and adjoint solutions. The meshes were generated using an automated grid generator and interpolation coefficients were accumulated in a preprocessing step. The parallel implementation of the flow and adjoint solvers were used to reduce the computational time of the design process. Modifications to the shape of the wing were transmitted to the interior mesh using the spring deformation method which worked well for this problem.

Figures 22,23,24,25 show that the target pressure distribution has been recovered in about 50 design cycles. These computations took under 30 minutes requiring 8 processors of an SGI Origin 300.

7.3 Inverse Design for Sail Geometries

The results of the flow and aeroelastic simulations, show that the interaction of the head sail with the main reduces the development of sharp pressure gradients around the leading edge. This interaction is crucial to the performance of the main sail as it allows the main sail to be set at a higher angle to the center-line of the boat. These results also show that the region above the head sail has large suction peaks which is a cause of concern. The aerodynamic shape optimization procedure validated in the previous sections was used to redesign the main sail, with an aim of reducing the pressure gradient around the luff of the main sail. An inverse design procedure was employed and the target pressure distribution was obtained by smoothing the pressure distribution on the main sail obtained from the aeroelastic analysis. Figures 26,27,28 and 29 show that a significant portion of the leading edge of the main sail has been redesigned to allow for smooth entry of the flow. The associated reduction in sharp suction peaks should have a favorable affect on the growth of the boundary layer over the upper surface. The change to the sections are shown in figure

8. Conclusions

Rotational inviscid flow solutions to sail configurations can now obtained in a cheap and efficient manner. These can be used to substitute potential flow models to obtain better estimates of the induced drag. Adjoint based design methods can be used to redesign sail shapes to smooth out suction peaks and improve the sail performance. The use of gradient formulations that depend only of the surface mesh allows adjoint based methods to be used for unstructured grids in a computationally efficient manner. Hence, it is now possible to devise completely automated shape optimization procedures for sail configurations. Further exploiting the flexibility of unstructured grids, it is possible to devise a hierarchy of optimization procedures that can be used to tackle planform optimization, sail settings, hull/keel shapes and other components of the overall design process. This holds great promise for the design of improved sails.

9. Ongoing and future directions of Research

Our current efforts are focussed on developing computational methods for viscous simulations. Preliminary results from a two equation turbulence model have provided good estimates of the friction drag for airplane configurations. Boundary layer methods coupled to inviscid flow solutions are also being investigated to provide a computationally cheap method to estimate the viscous effects. Redesign of hull and keel shapes are within sight, and it is hoped that in the near future, a combined analysis and design tool for high performance boats will be realized.

References

[1] A. Jameson and T.J. Baker, Improvements to the Aircraft Euler Method, *AIAA Paper 87-0353*, 25^{th} AIAA Aerospace Sciences Meeting, Reno, January, 1987.

[2] J. Reuther, A. Jameson, J. Farmer, L. Martinelli and D. Saunders, Aerodynamic Shape optimization of complex aircraft configurations via an adjoint method, *AIAA Paper 96-0094*, 34^{th} AIAA Aerospace Sciences Meeting, Reno, January, 1996.

[3] A. Chorin, A Numerical Method for Solving the Incompressible Viscous Flow problem, *Journal of Computational Physics* Vol. 2, pp 12-26, 1967.

[4] J. Farmer, L. Martinelli and A. Jameson, A Fast Multigrid Method for Solving Incompressible Hydrodynamic Problems with Free Surfaces, *AIAA Paper 93-0767*, 31^{st} AIAA Aerospace Sciences Meeting, Reno, January, 1993.

[5] A. Rizzi and L. Eriksson, Computation of inviscid incompressible flow with rotation, *Journal of Fluid Mechanics* Vol. 153, pp 275-312, 1985.

[6] A. Jameson, Analysis and Design of numerical schemes for gas dynamics II, Artificial diffusion and discrete shock structure, *International Journal of Computational Fluid Dynamics*, Vol. 5, pp. 1-38, 1995.

[7]A. Jameson, Multigrid algorithms for compressible flow calculations, In W. Hackbusch and U. Trottenberg, editors, *Lecture Notes in Mathematics, Vol. 1228*, pages 166–201. Proceedings of the 2nd European Conference on Multigrid Methods, Cologne, 1985, Springer-Verlag, 1986.

[8]A. Jameson, W. Schmidt and E. Turkel, Numerical Solution of the Euler equations by finite volume methods using Runge-Kutta time stepping schemes, *AIAA Paper 81-1259*, June, 1981.

[9]A. Jameson, Multigrid algorithms for compressible flow calculations, *Proceedings of the 2nd European Conference on Multigrid Methods*, Cologne, 1985.

[10]A. Jameson, Unpublished notes and computational techniques for the solution of the unsteady Euler equations on unstructured grids using multigrid methods.

[11]J.L. Lions, Optimal Control of Systems Governed by Partial Differential Equations, *Springer-Verlag*, New York, 1971. Translated by S.K. Mitter.

[12]O. Pironneau, Optimal Shape Design for Elliptic Systems, *Springer-Verlag*, New York, 1984.

[13]A. Jameson, Optimum Aerodynamic Design Using Control Theory, *Computational Fluid Dynamics Review 1995* Wiley, 1995.

[14]T.J. Barth, Apects of unstructured grids and finite volume solvers for the Euler and Navier-Stokes equations, *AIAA Paper 91-0237*, 29^{th} AIAA Aerospace Sciences Meeting, Reno, January, 1994.

[15]K. Anderson and V. Venkatakrishnan, Aerodynamic Design Optimization on Unstructured grids using a continuous adjoint formulation, *AIAA Paper 97-0643*, 34^{th} AIAA Aerospace Sciences Meeting, Reno, January, 1997.

[16]J. Elliot and J. Peraire, Aerodynamic design using unstructured meshes, *AIAA Paper 96-1941*, 33^{rd} AIAA Aerospace Sciences Meeting, Reno January, 1996.

[17]S. E. Cliff, S.D. Thomas, T. J. Baker, A. Jameson and R. M. Hicks, Aerodynamic Shape optimization using unstructured grid method, *AIAA Paper 02-5550*, 9^{th} AIAA Symposium on Multidisciplinary Analysis and Optimization, Atlanta, September, 2002.

[18]A. Jameson and Sangho Kim, Reduction of the Adjoint Gradient Formula in the Continuous Limit, *AIAA Paper*, 41^{st} AIAA Aerospace Sciences Meeting, Reno January, 2003.

[19]A. Jameson, L. Martinelli and J. Vassberg, Using CFD for Aerodynamics - A critical Assesment, *Proceedings of ICAS 2002*, September 8-13, 2002, Toronto, Canada.

[20]G. Cowles and Luigi Martinelli, A Control-Theory Based Method for Shape Design in Incompressible Viscous Flow using RANS, *AIAA Fluids 2000-2544*, June 19-22, Denver, CO.

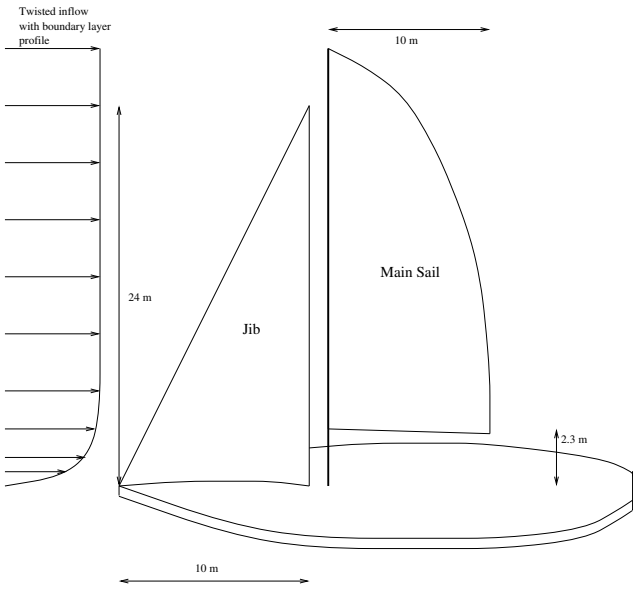

Fig. 7 Sail geometry

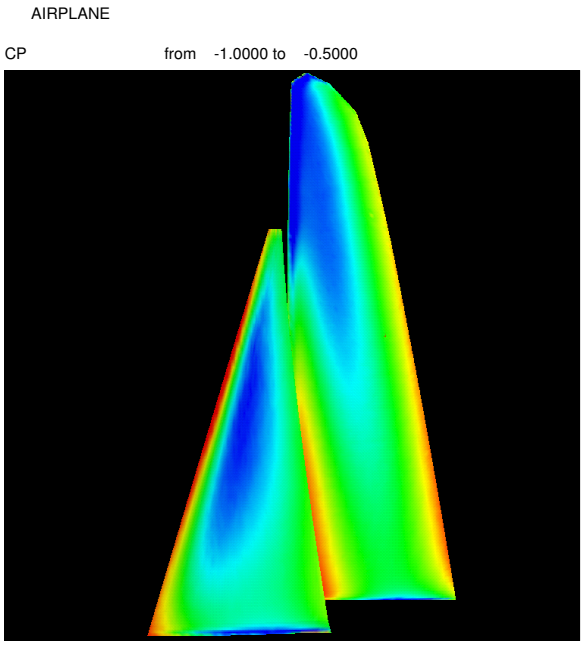

Fig. 9 Pressure contours on the leeward side

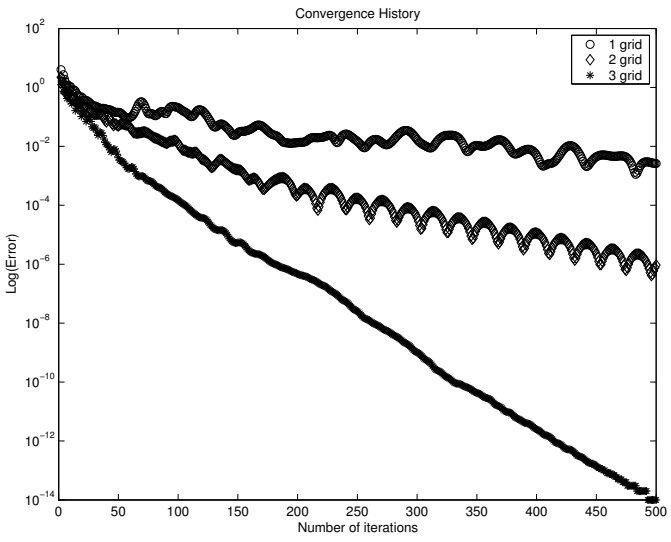

Fig. 8 Convergence history for sail geometries with artificial compressibility

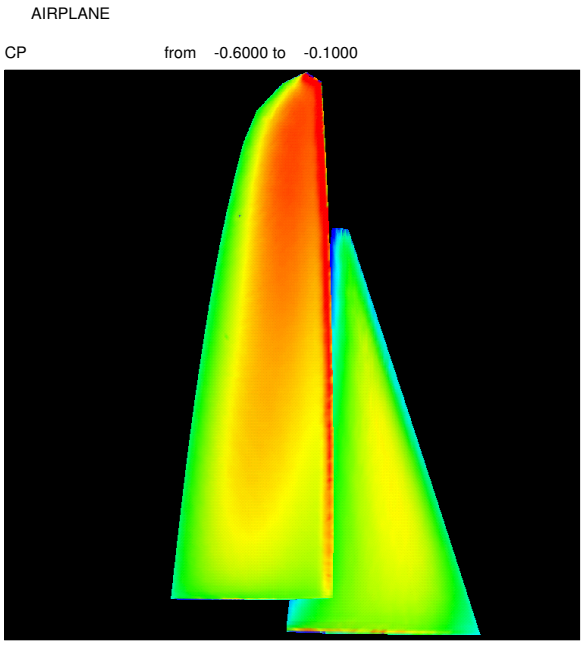

Fig. 10 Pressure contours on the windward side

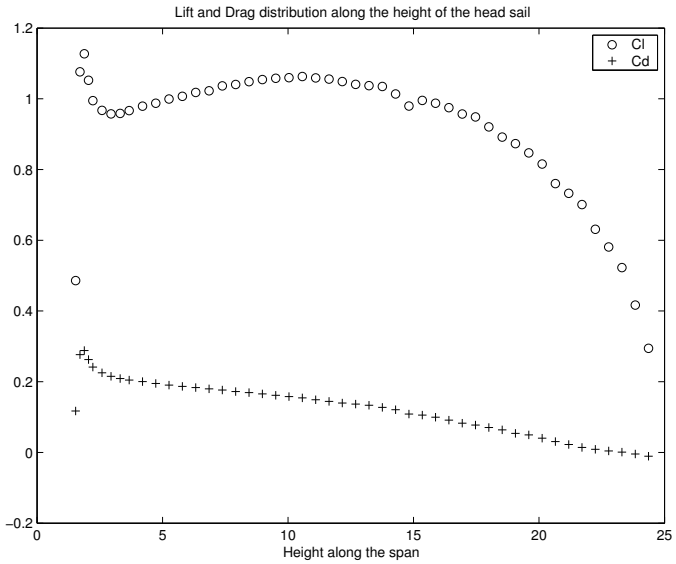

Fig. 11 Spanwise force distributions on the head sail

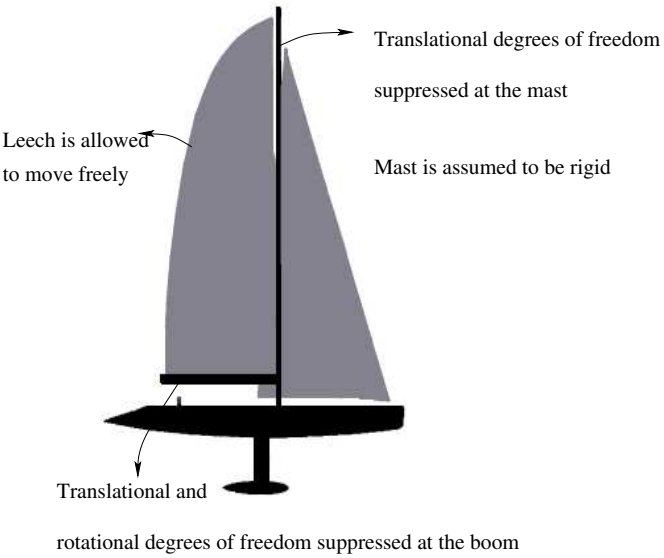

Fig. 13 Boundary conditions for the main sail

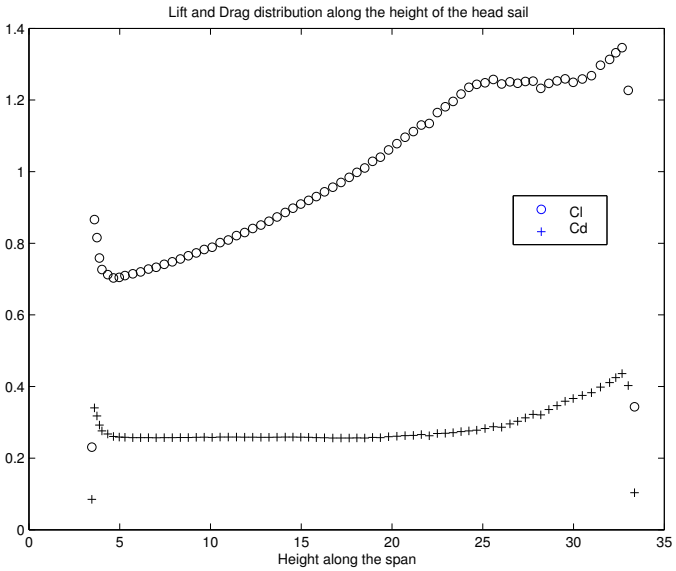

Fig. 12 Spanwise force distributions on the main sail

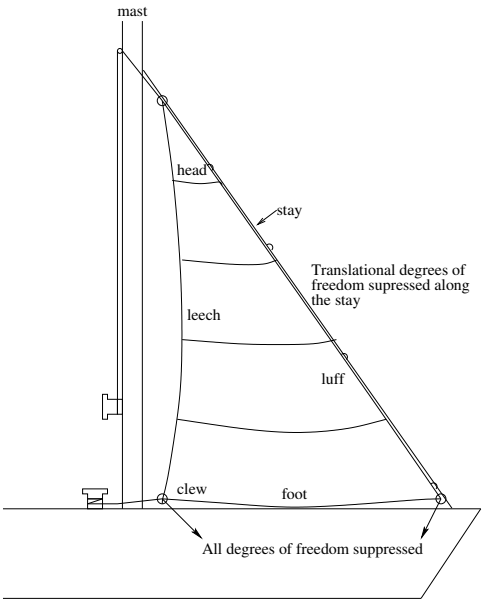

Fig. 14 Boundary conditions for the head sail

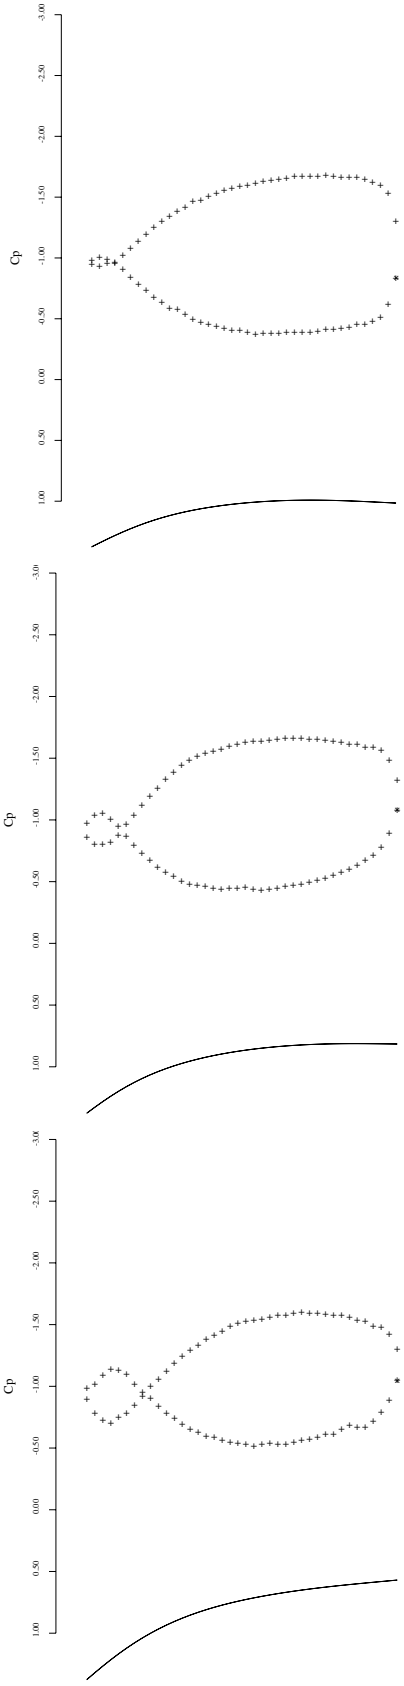

Fig. 15 Pressure distributions along sections at 1, 25 and 85 percent of the height of head sail

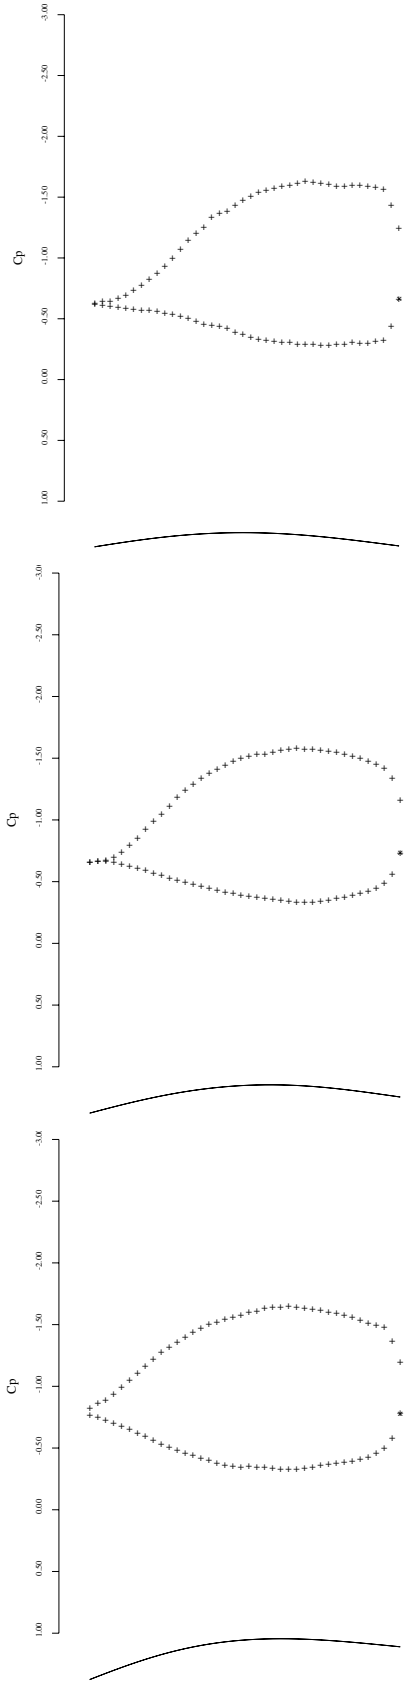

Fig. 16 Pressure distributions along sections at 1, 25 and 85 percent of the height of the main sail

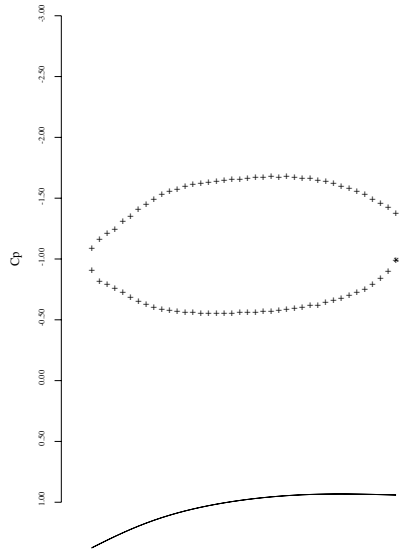

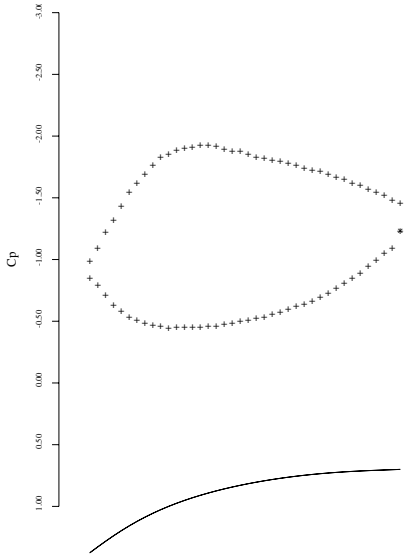

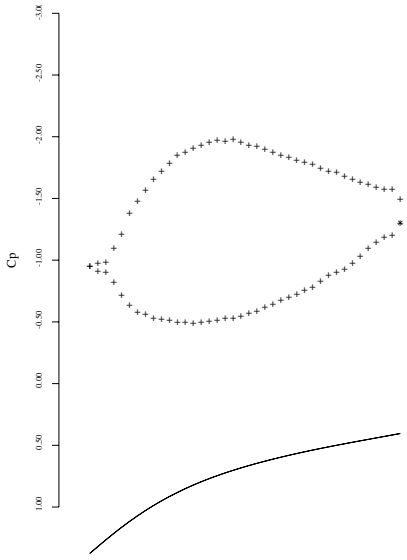

Fig. 17 Pressure distributions along sections at 1, 25 and 85 percent of the height of head sail after aeroelastic analysis

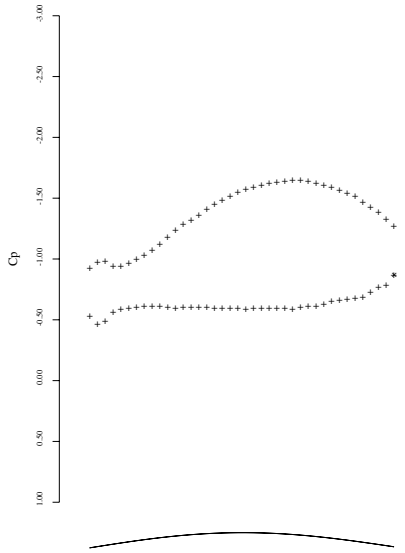

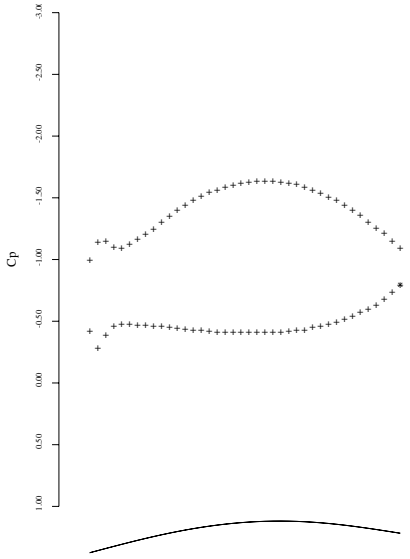

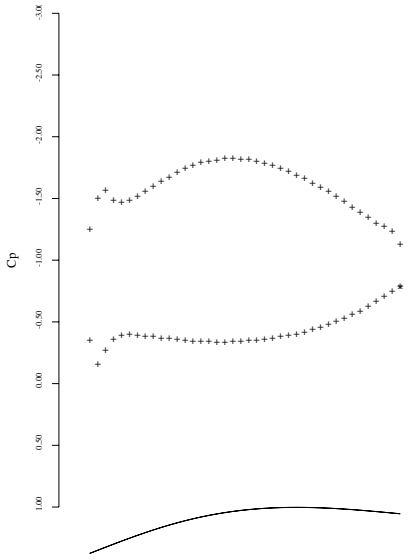

Fig. 18 Pressure distributions along sections at 1, 25 and 85 percent of the height of main sail after aeroelastic analysis

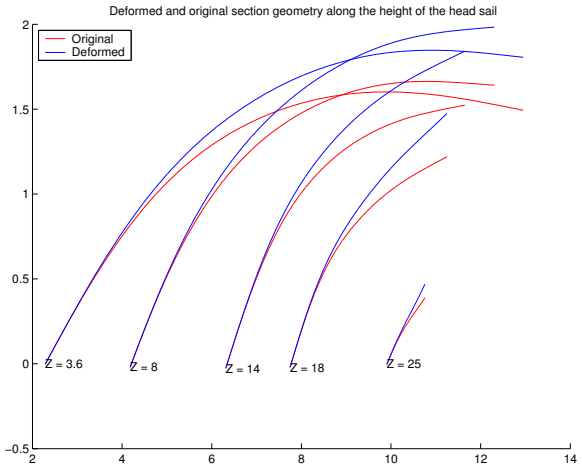

Fig. 19 Original and deformed sail sections for the head sail

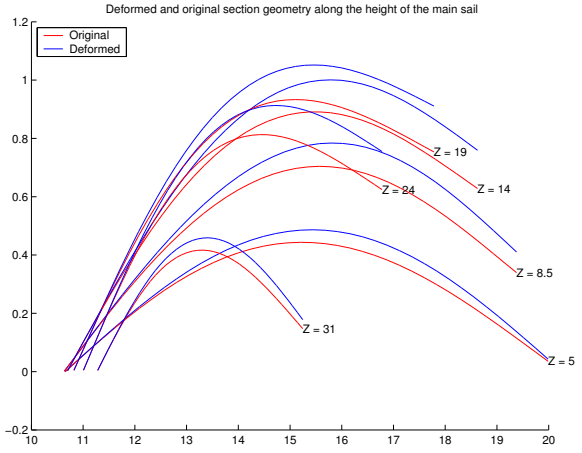

Fig. 20 Original and deformed sail sections for the main sail

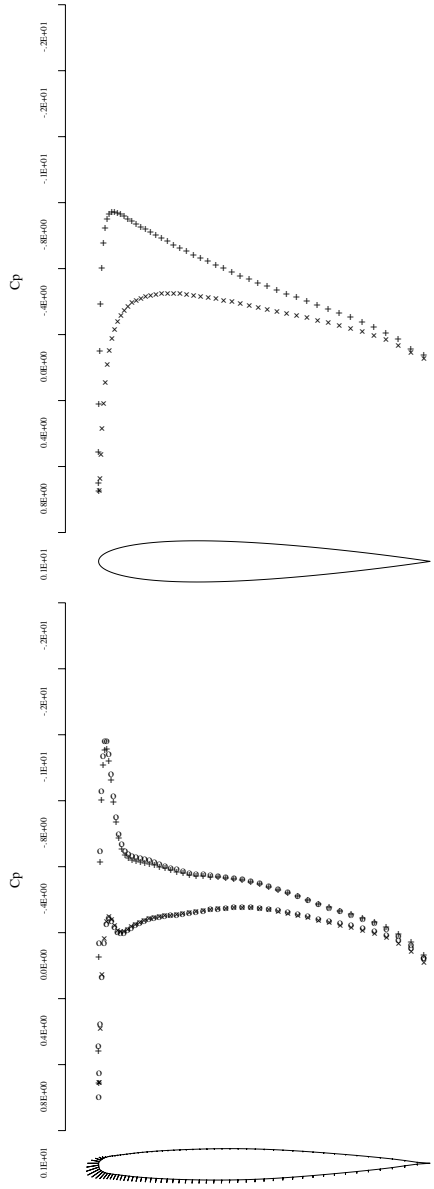

Fig. 21 Initial and final pressure distribution, o is the target pressure distribution, x is the computed pressure distribution for the redesigned airfoil

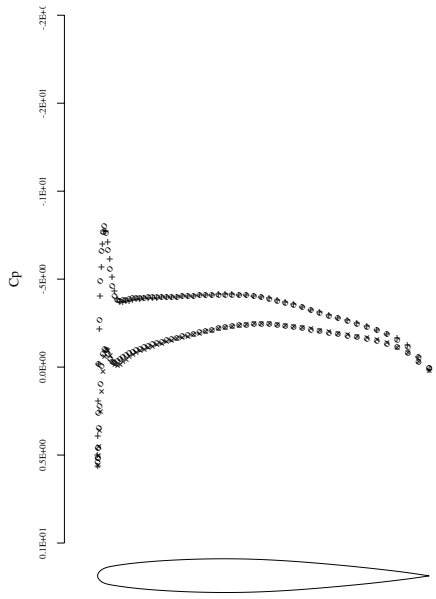

Fig. 22 Attained(+,x) and target(o) pressure distributions at 0 % of the wing span

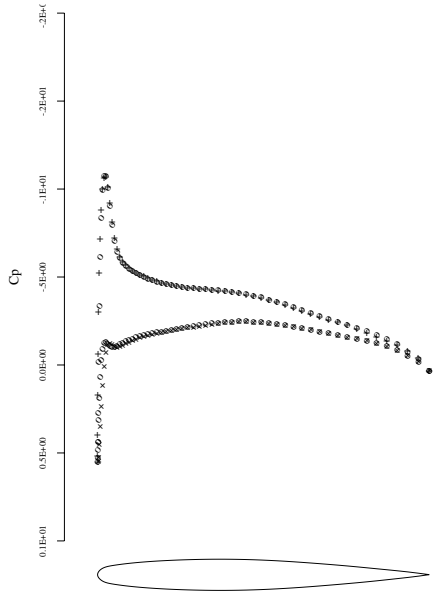

Fig. 23 Attained(+,x) and target(o) pressure distributions at 25 % of the wing span

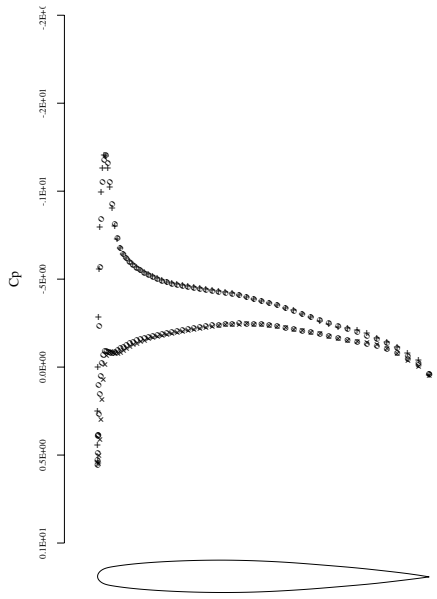

Fig. 24 Attained(+,x) and target(o) pressure distributions at 75 % of the wing span

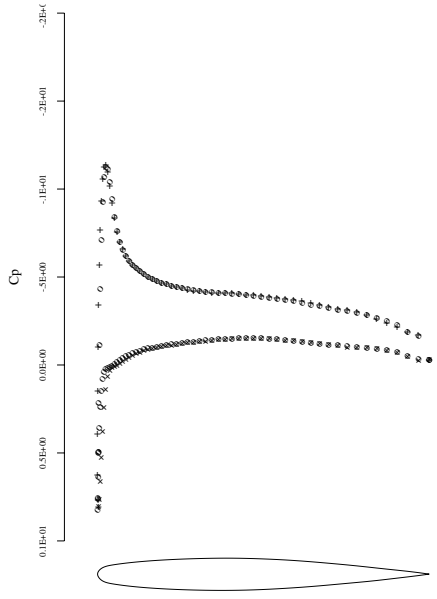

Fig. 25 Attained(+,x) and target(o) pressure distributions at 100 % of the wing span

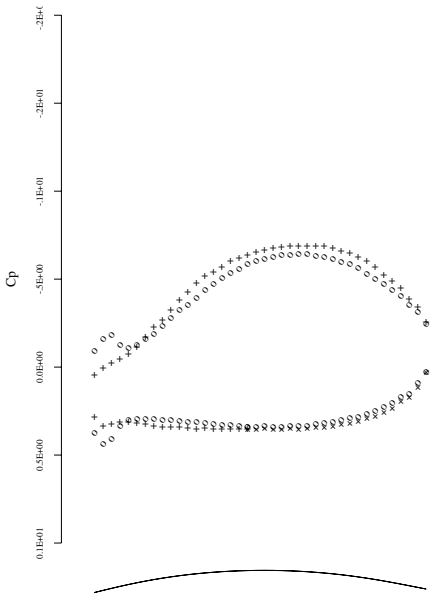

Fig. 26 Initial (o) and final pressure distribution at 15% height on the main sail

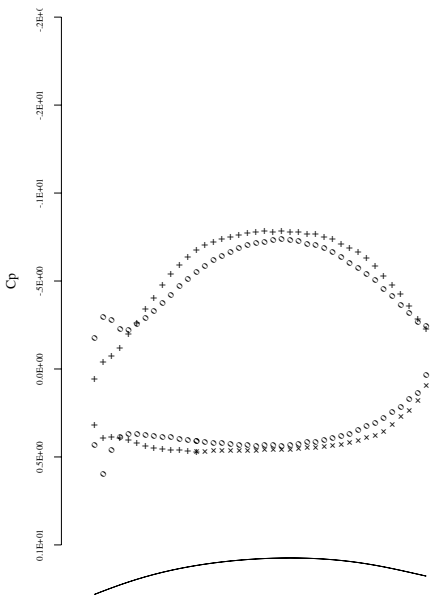

Fig. 27 Initial (o) and final pressure distribution at 32% height on the main sail

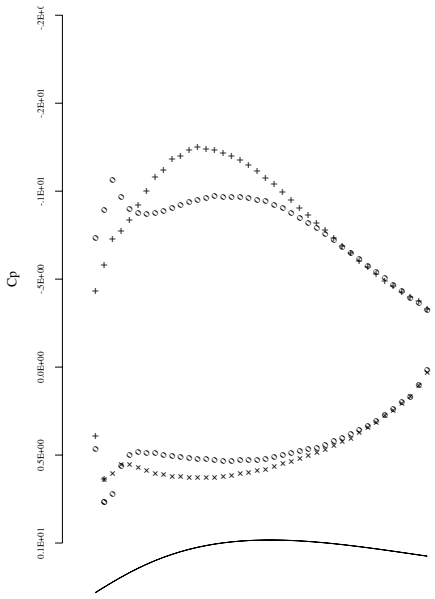

Fig. 28 Initial (o) and final pressure distribution at 75% height on the main sail

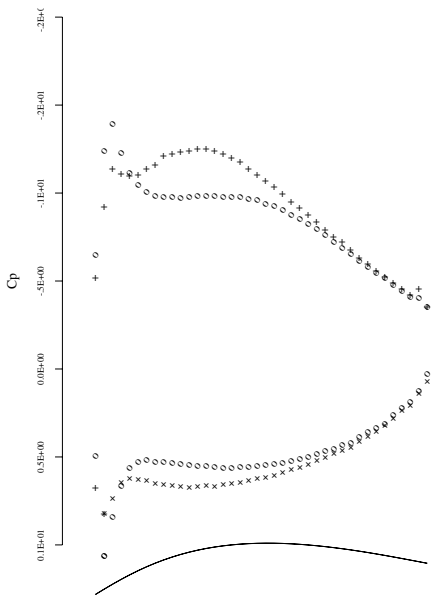

Fig. 29 Initial (o) and final pressure distribution at 85% height on the main sail

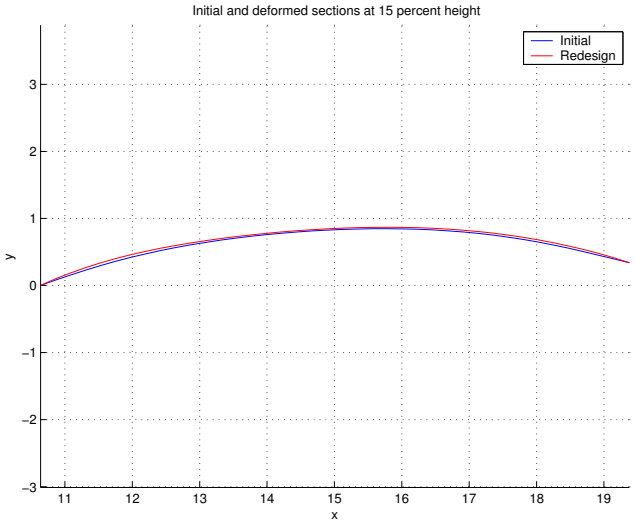

Fig. 30 Initial and redesigned camber line at 15% of height

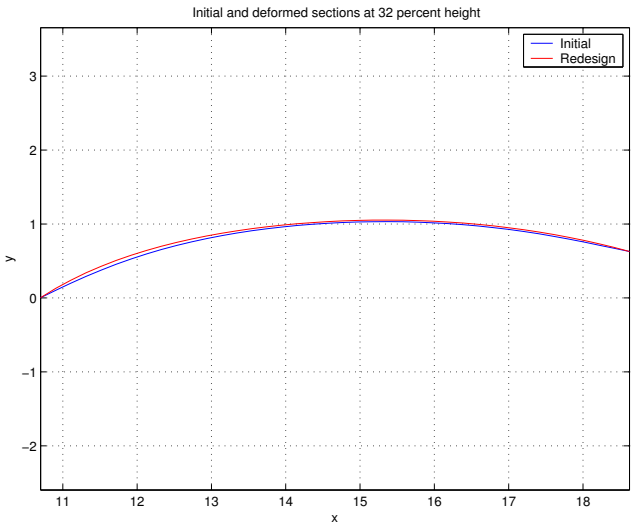

Fig. 31 Initial and redesigned camber line at 32% of height

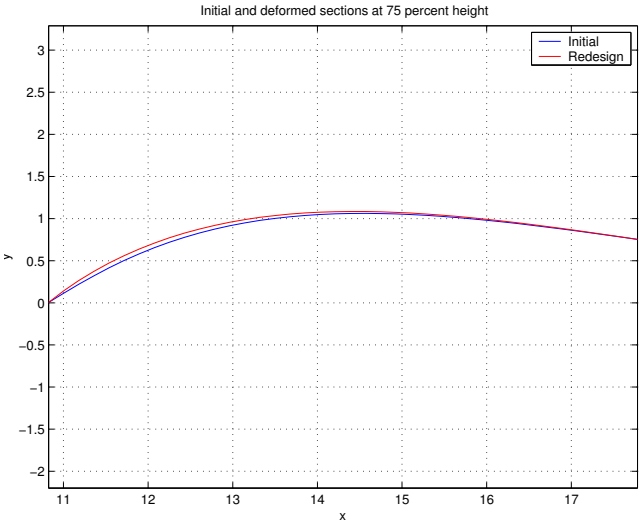

Fig. 32 Initial and redesigned camber line at 75 % of height

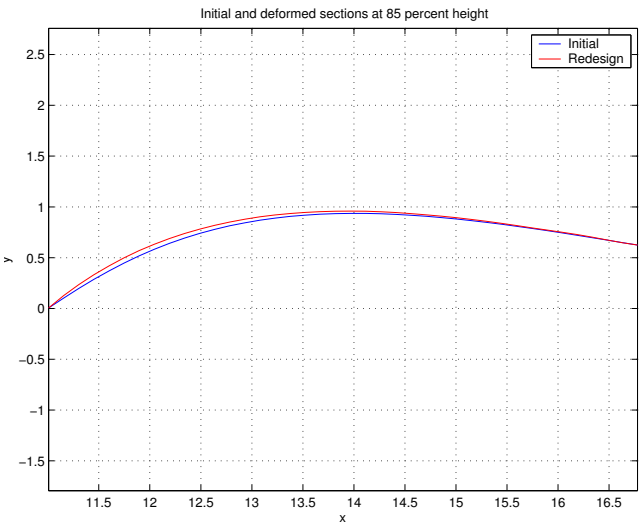

Fig. 33 Initial and redesigned camber line at 85 %of height