AERODYNAMIC DESIGN USING THE EULER ADJOINT APPROACH

M. R. Cross
BAE SYSTEMS Airbus UK
murray.cross@baesystems.com

Keywords: *optimisation, adjoint, aerodynamic design*

Abstract

This paper describes and demonstrates aerodynamic optimisation using the Euler adjoint approach as developed by Jameson on wing/body geometries. This method consists of an Euler flow solver which has been coupled to an integral boundary layer method, an Euler adjoint solver and a simple optimiser. Also described is a new technique for extending the method to more complex configurations, which retains the speed benefit of the structured mesh adjoint technique and combines it with the geometric flexiblity of a viscous-coupled unstructured flow solver. The capability is demonstrated on the W4 wing/body geometry and extended to W4 wing/body/nacelle configuration.

1 Introduction

Optimisation is becoming increasingly important in the field of aircraft design. The complexity of today's aircraft and the competitive pressures of international industry are driving the use of automatic computer optimisation methods within aircraft design. The aerodynamic design process is a typical example where optimisation methods can yield significant benefits, and this paper discusses and presents results showing some of the capability of optimisation using the Euler adjoint approach.

Historically, CFD optimisations have always been computationally expensive, with the CPU time proportional to the number of design variables defining the problem when using a gradient-based method, and possibly even more expensive for stochastic approaches such as genetic algorithms. The difficulty has been in calculating the sensitivity of the cost function with respect to the N design variables, often this has involved N+1 CFD calculations using the traditional finite difference method. When defining the shape of a wing for example, the number of parameters describing the section and planform shapes can be of order 100 or more. This could lead to order 100 CFD calculations for one set of sensitivities, which in most cases is prohibitively expensive.

Recently, new approaches to aerodynamic optimisation have been developed. The approach described here is the continuous Euler adjoint method, as developed by Jameson. This method calculates aerodynamic sensitivities by solving both the Euler flow and Euler adjoint equations. This approach means that with one flow and one adjoint solution all the sensitivities can be calculated independent of the number of parameters used to define the surface shape. Therefore one design iteration for one design point is roughly the cost of two flow solutions. Having derived the sensitivities the optimisation method uses a gradient based method to modify the geometry to achieve a step improvement in the objective function.

When modeling the airflow around a wing, capturing the viscous effects is very important if an accurate representation of the forces and flow features is required. In commercial aircraft design the design tolerances demand accurate prediction of the aircraft performance and

therefore viscous effects must be included in any optimisation process. The two main routes for including the viscous effects are moving to a Navier-Stokes adjoint formulation [4], or coupling the Euler solver, via transpiration boundary conditions, to a boundary layer solver [7]. The latter approach has been employed here having the advantage of rapid solution turn-around, with the disadvantage that the Euler adjoint has no visibility of the skin friction drag.

For practical design use, the capability must exist to optimise a wing at one or more design conditions. When optimising at a single point, the code can produce a wing with a local minimum in the cost function at that particular condition, but the off design performance is often worse than the original wing. For this reason in most circumstances, three design conditions are used. When more than one design condition is used, Euler and adjoint solutions are calculated for each design point and the sensitivities combined to give a new wing shape.

2 Method

2.1 Euler and Adjoint Solver

The code described in this paper as originally developed by Jameson, is a finite volume cell centred structured mesh single block code, which uses the Runge-Kutta method for integrating the fluxes in time, to obtain a steady-state solution. The Euler and adjoint equations are solved in a similar manner. Local time stepping and multigrid are used to accelerate convergence, with enthalpy damping also being used in the Euler solution. Stability is achieved using 2^{nd} and 4^{th} order artificial dissipation, with a 2^{nd} derivative of the pressure as a switch to apply the 2^{nd} order dissipation in the shock region. The dissipation can be applied either in an isotropic or an anisotropic manner depending on the required accuracy and the required convergence speed. Developments have been made at Airbus UK, to enable the Euler solver to be run either purely as an inviscid code, or in conjunction with a boundary layer code with a transpiration boundary conditions applied to the wing surface.

The quasi-simultaneous coupling technique is used to link the Euler solver with an integral boundary layer code which is based on the lag-entrainment method.

The adjoint boundary conditions are set according to the type of optimisation required. The Euler adjoint boundary conditions are set by a function of pressure defined over the wing surface. This limits the setting of the objective function to integrated pressure values such as drag, lift, pitching and bending moments, or pressure difference between current Euler predicted values and target values. One of the possible disadvantages of the adjoint approach for improving the performance of an aircraft is the reliance on integrated pressure drag. Integrated pressure values are not always very accurate as a measure of drag as small errors in the flow solution can lead to large errors in the level of drag, but generally it has been found that the incremental values are in the right direction even if the absolute drag values are not exact.

2.3 Parameterisation

The code uses each grid node on the wing surface as a design parameter. This leads to an enormous number of parameters, but this does not affect the optimisation speed due to the nature of the adjoint method. After calculating the sensitivities, the point is moved in the direction along the normal grid line which reduces the cost function. The change at the surface is propagated outward into the domain at a reducing rate. As the parameterisation method does not inherently preserve surface smoothness, the modifications are heavily smoothed with implicit and explicit smoothing to reduce the risk of discontinuities.

2.4 Optimiser

A simple optimisation procedure is used to update the geometry. The gradient as calculated from the adjoint and Euler variables is smoothed and under-relaxed. Each surface point is then moved a small distance in the direction which reduces the cost function.

There are two main categories of objective function. One is minimisation (or maximisation) of integrated pressure quantities such as drag, and the other is minimisation of differences between current and target surface pressures, giving an inverse design capability. Drag, pitching moment and wing bending moment can all feature in the objective function. The addition of the moment objective functions gives the designer more control on the type of pressure distributions produced by the code. Minimising wing bending moment brings the spanwise centre of lift inboard which can reduce the wing structural weight. Maximising wing pitching moment, or minimising the nose down pitching moment can reduce the tail force required to trim an aircraft which also can improve overall aircraft performance. Another important role of the pitching moment objective function, is that the gradient of the recompression region towards the rear of the aerofoil sections can be controlled. This is important to prevent the optimiser from excessively increasing the rear loading at the expense of increased profile drag.

3. W4 Testcase

The testcase used throughout this paper is the W4 wing geometry with a circular cross-section semi-infinite body, at a Mach number of 0.78 [2]. The structured mesh used in this paper is reasonably coarse at roughly 164,000 cells.

Both drag optimisations and inverse designs will be shown on the W4 configuration, highlighting the benefits of this optimisation approach. This paper demonstrates the ability of the code to reduce wing drag, while maintaining wing thickness and lift coefficients. The wing thickness constraint needs to be present to avoid the optimiser reducing the wing thickness, spar depth and hence the fuel volume. Single design points are used in all the testcases to demonstrate the design capability, although in real design situations three design points are generally used.

All the coefficients quoted in this paper will be integrated pressure values, except for wave drag figures which come from a Lock wave drag calculation method [1].

4 Wing/Body Results

4.1 W4 wing/body Drag Minimisation

A single point drag minimisation case is shown here with a target CL=0.5. The design history for 20 iterations is shown in Figure 1.

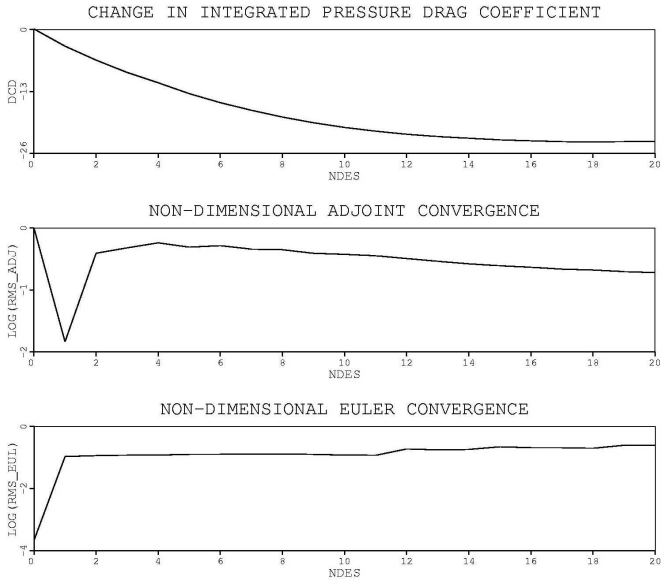

Figure 1 – Design History for wing/body Cd minimisation

The top graph shows the drag reducing steadily by roughly 23 drag counts. The other two plots show the log of the reduction in the first residual of the adjoint and Euler calculation for each design cycle. The level of convergence for the first iteration of each is greater than the

subsequent cycles as more iterations are used to start off the solution. The solutions for subsequent cycles are restarted from the previous design cycle.

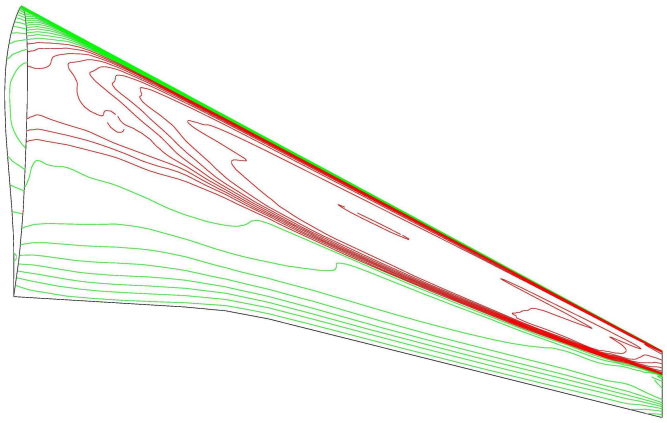

Figure 2 – Mach contours original geometry

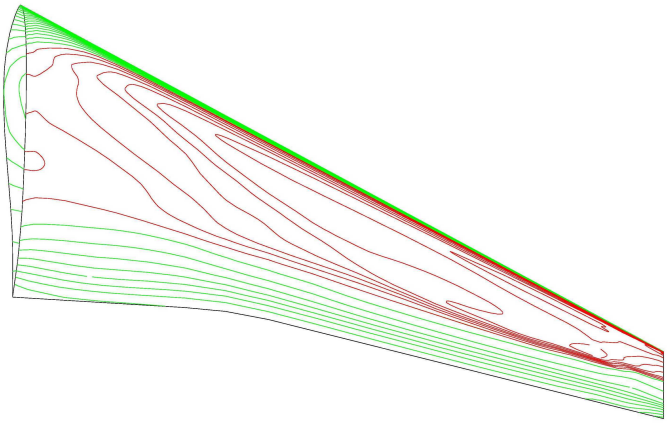

Figure 3 – Mach contours optimised geometry

The change in shock pattern can be seen from the Mach number contours in Figures 2 and 3, the red colour indicating supersonic, and the green subsonic flow. The majority of the drag benefit is coming from reducing the strength of the shock wave, which was achieved with increased mid-chord thickness and camber.

More details can be seen when looking at sectional cuts. Figure 4 shows cuts at four stations across the wing, plotting pressure coefficients for the original and final geometries.

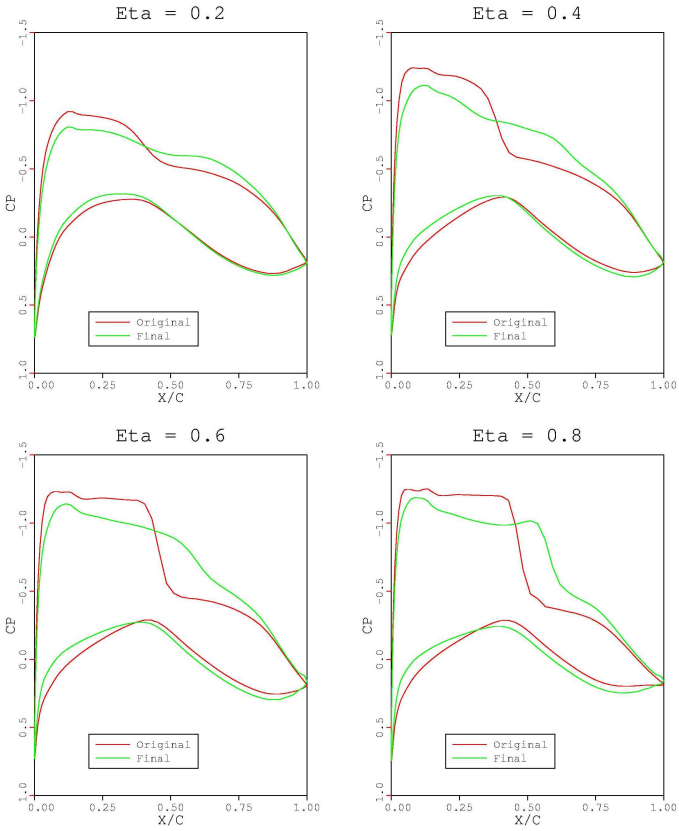

Figure 4 – Sectional Cp's across wing

4.2 W4 Drag Minimisation with fixed loading

The method includes the capability to modify the local twist across the wing in order to achieve a desired loading distribution. Loading distribution on the wing significantly affects the required wing strength (and therefore wing weight), together with vortex drag. When changing the loading, a change in section design is required across the wing to reoptimise the pressures. By combining the drag minimisation with the target loading capability, useful sensitivity figures can be extracted to feed into configuration optimisation. Useful trades can

therefore be made between wing weight and wing drag.

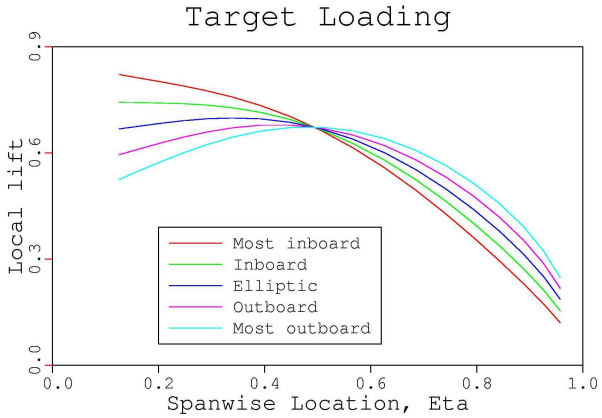

Figure 5 – Target loading distributions

Five cases were run with the target loadings given in Figure 5 and the resultant sectional Cp's can be seen in Figure 6. In each case the geometry has been redesigned to reduce the shock strength while maintaining the prescribed loading. The advantage of this approach is that a consistent methodology is used for each of the loading distributions, so the final trades between loading and drag should be consistent.

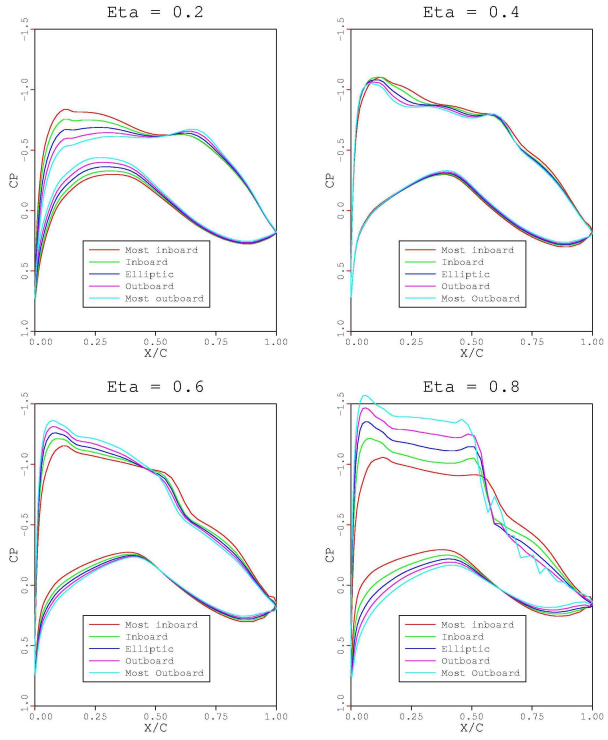

Figure 6 – Cp distributions for 5 span loadings

The high loadings and strong shock on the most outboard loaded wing are causing oscillations in the viscous-coupled Euler solution, which are not physically realistic.

4.3 W4 Wing/Body Inverse Design

The final pressures from the first example of drag minimisation (paragraph 4.1) were given as the target for an inverse design case. Starting with the same geometry and same conditions as previously the geometry was modified to achieve this target. As seen in Figure 7 the match was good after 50 iterations, and the difference between the target and actual pressures was still reducing.

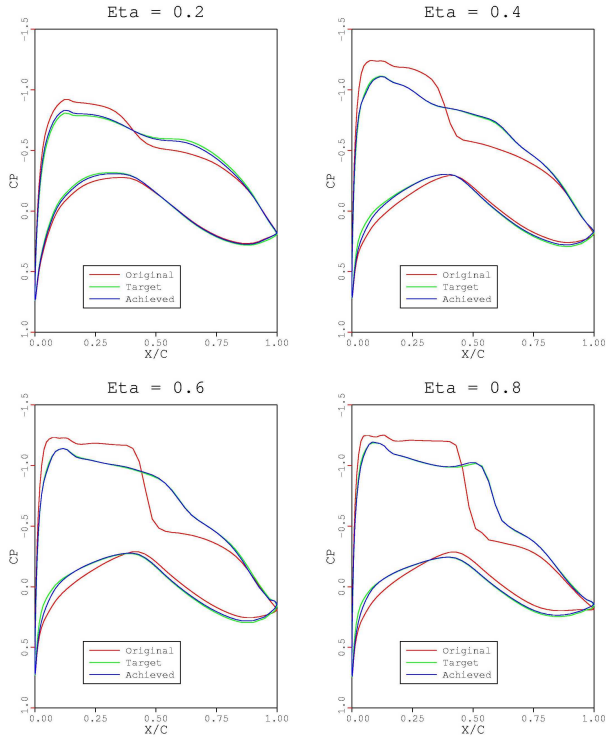

Figure 7 Initial, Target and Final Cp's

The match between achieved and target pressures is very good in this case because the target pressures came from a real geometry. In most real design cases a desired pressure distribution is given which is not fully achievable either because it is physically unrealistic or because the geometry is constrained to be smooth chordwise and spanwise. In these cases the nearest achievable pressure match is produced.

5. Extension of Method to Complex Configurations

One of the limitations of the method previously described is the geometry modeling aspect which only allows a single block structured mesh. This limits the geometric flexibility to wing/body type geometries. An extension of the method has been developed which shows promise in extending the adjoint method to more complex configurations while retaining many of the speed advantages of the single block structured mesh.

The adjoint boundary conditions are set by supplying pressures from the Euler solution. The single block viscous-coupled Euler solver is now replaced in this new approach with an unstructured viscous-coupled Euler solver. The flow solution is interpolated onto the structured adjoint mesh for solving the sensitivies. The aerodynamic sensitivities are solved on the structured mesh and these sensitivities are fed into the optimiser to drive the geometry change. This geometry change needs to be fed back into the unstructured Euler solution. The main drawback in this approach would be the CPU required to regenerate the unstructured mesh. In order to reduce the CPU and complexity of regenerating the unstructured mesh for each design iteration, transpirations are used to simulate the geometry modification in the Euler solution. This cycle is then repeated for as many iterations as required.

The method to simulate a change in geometry using transpirations works in a similar way to the viscous-coupling concept. The inviscid streamlines are deviated by applying a normal velocity at each point on the wing surface. This gives a reasonable approximation to small changes in geometry as demonstrated in the wing/body testcase below.

6. Results for Complex Configurations

6.1 W4 Unstructured Flow Solver Wing/Body Drag Minimisation

Two meshes are used in this example. The structured mesh as used previously, and an unstructured mesh based on the same geometry. This case was only run for 5 design iterations to minimise the drag at design CL=0.46. The

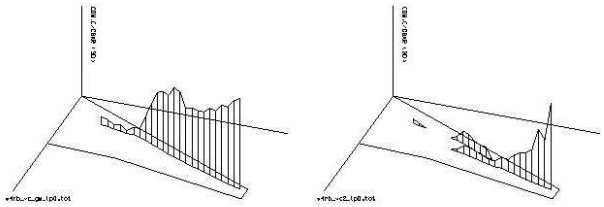

Figure 8 – Wave drag distribution

wave drag was reduced from 10 counts to 4 as can be seen in Figure 8 where the height of the bars represent the level of wave drag.

The reduction in the strength of the shocks can also be seen in the pressure coefficient plots in Figure 9. The new geometry was remeshed and run at the same conditions to show the accuracy of the transpiration geometry approximation. This result is also plotted on Figure 9, and except for a slight change in shock location outboard, shows a reasonable match.

The changes in the geometry are similar to the earlier testcase, with increased thickness and camber in the shock region and reduced incidence to achieve the design CL. These changes significantly reduce the wave drag of the wing in the region of the design point, but there will be serious degradation of the performance at higher Mach numbers. Optimisation at higher Mach numbers often has the opposite effect of reducing the camber to attempt to reduce the pre-shock Mach number. This is why it is important to use multiple design points in a real design situation even

though the drag benefit at the primary design point would be reduced.

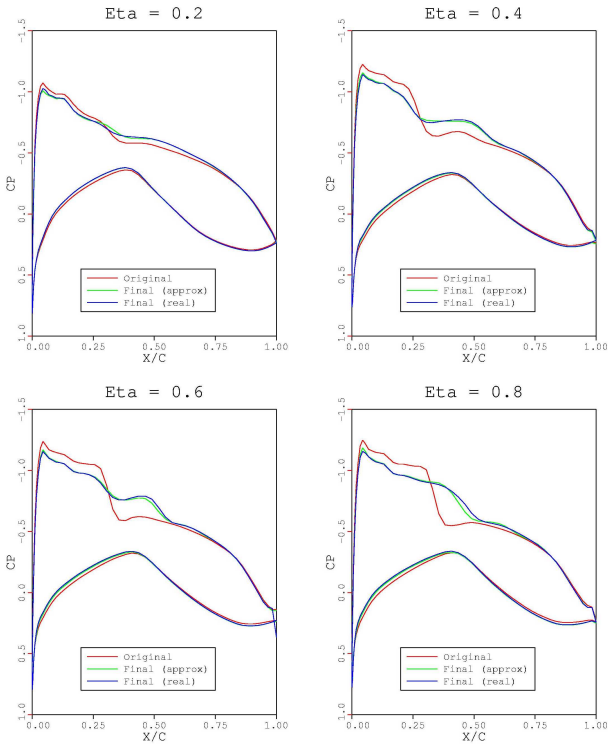

6.2 W4 Wing/Body/Nacelle Drag Minimisation

When designing a wing the influence of the powerplant must be taken into account if the final wing is to be of optimal performance. This testcase attempts to modify the wing in order to reduce wing drag while simulating the blockage effect of a through-flow nacelle.

The geometry modeled by the structured and unstructured meshes were different in this case. The geometry modeled by the unstructured mesh contained a long cowl nacelle (not attached to the wing) and consisted of nearly 2 million cells.

Regions existed on the structured mesh where interpolation could not be performed due to the presence of the nacelle in the unstructured mesh. These values were detected and repaired using an averaging technique.

The process described here calculates approximate sensitivities because the shape of the solution domain for the Euler and the adjoint calculation is different. The boundary conditions which drive the adjoint solution are valid, however the propagation of the adjoint fluxes into the field will not be exact as the nacelle's influence is not being modeled. This approximation is thought to be reasonable, however it would not be possible to include the nacelle drag in the objective function.

The code was run in the same way as the wing/body testcase at design CL=0.53, for 10 design iterations and this was achieved in under 15hrs on a HP N-Class processor.

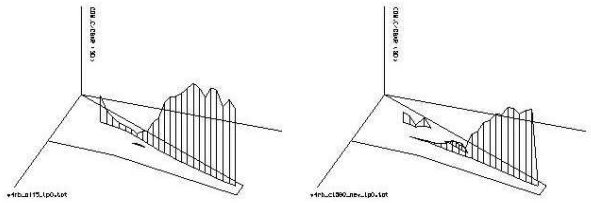

Figure 10 – Wave drag distribution

The wave drag was reduced from 24 counts to 14 (see Figure 10), again by increasing thickness and camber mid chord. The single strong shock front is being weakened and a λ type shock pattern is emerging.

The original geometry pressure contours can be seen in Figure 11, and Figure 12 shows the modified pressures.

In general the presence of a nacelle can significantly change the nature of the shock structure on the upper surface of the wing, especially at higher Mach numbers. This is why it is important to model these effects when optimising the geometry.

Figure 11 – Initial Cp contours

Figure 12 – Final Cp contours

7. Conclusion

A practical three-dimensional Euler adjoint optimisation method has been described and results presented on the W4 wing. Drag minimisation and inverse design have been demonstrated on the wing/body geometry using a structured mesh viscous-coupled Euler solver.

A method for extending the capability of the adjoint approach to more complex geometries has been proposed with promising results. The main advantage of this method being that the pressures from a more realistic geometry are driving the adjoint equations so that a wing geometry can be optimised in the presence of a component such as a nacelle. This approach also has the advantage of providing a rapid optimisation capability by combining the rapid structured mesh adjoint solution and the geometric flexibility of an unstructured mesh flow solver.

8. References

[1] ESDU 87003. A method of determining the wave drag and its spanwise distribution on a finite wing in transonic flow, ESDU International, April 1997.

[2] Fulker. Pressure distributions on research wing W4 mounted on an axisymmetric body, AGARD-AR-303, volume II.

[3] Jameson, Alonso, Reuther, Martinelli, Vassberg. Aerodynamic shape optimization techniques based on control theory, AIAA98-2538, June 1998.

[4] Jameson. Optimum aerodynamic design using control theory, Computational Fluid Dynamics Review, June 1995.

[5] Lorentzen. Surface parametrisation in the context of aerodynamic shape optimization with an adjoint solver, ECCOMAS 1998.

[6] Reuther, Jameson, Farmer, Martinelli, Saunders. Aerodynamic shape optimization of complex aircraft configurations via adjoint formulation, AIAA96-0094, January 1996.

[7] Szmelter. Aerodynamic wing optimisation, AIAA 99-0550, January 1999.