

Chapter 7

Optimisation tests and results in two dimensions

Once the four components i.e. the optimiser, the grid perturbation program, the CFD code MERLIN and the adjoint solver described in the last four chapters have been tested separately, they are put together to form the optimisation chain. Before applying this suite of codes to the three-dimensional optimisation of a BWB, the ultimate goal of this work, it has been tested on simple two-dimensional problems. This is what is described in this chapter. The purpose of this testing was to understand how the method worked and what could be done to make it more efficient. The reasons for choosing only a two-dimensional problem for these investigations are exactly the same as for the verification of the adjoint solver in Section 6.7 i.e. a quick turn-around and the possibility to study the effects of convergence of the flow and adjoint solvers on the optimisation process.

In particular a study was conducted on the influence of convergence levels, grid size and physical models on the accuracy of the gradient. This is explained in the first part of this chapter. The second part describes two-dimensional aerofoil drag minimisations using a direct optimisation method based on Sequential Quadratic Programming. This looks into the influence of the same parameters on the overall optimisation and not just on a single gradient. This part also investigates different ways of handling the constraints. Finally two-dimensional optimisation results using the variable-fidelity method are presented.

7.1 Influence of approximation on the accuracy of the gradient

This section conducts an investigation into the sensitivity of the gradient obtained by the adjoint method with respect to three classes of parameters: convergence levels of both the flow and adjoint solvers, grid refinement and nature of the physical model used in the adjoint solver. All these tests are conducted on the RAE2822 aerofoil with the same conditions as in Section 6.7 i.e. Mach number $M_\infty = 0.725$, incidence $\alpha = 2.54^\circ$, Reynolds




Figure 7.1: Influence of the level of convergence of the flow solver on the value of lift and drag coefficients.

$$\text{Accuracy of } C_D \text{ at } j \text{ orders} = \frac{|C_{D,L}^{12 \text{ orders}} - C_{D,L}^{j \text{ orders}}|}{|C_{D,L}^{12 \text{ orders}}|}$$

number $Re = 6.5 \times 10^6$ based on chord and freestream temperature $T_\infty = 283.0 \text{ K}$. Unless otherwise stated, the grid is the same $129 \times 33 \times 2$ grid and both the flow and adjoint solvers are run in their fully turbulent mode, the Baldwin-Lomax turbulence model being employed. The upper and lower surfaces of the aerofoil are parameterised using 10 Bézier parameters each. Since the end points are fixed for both surfaces, this reduces the number of design variables to 16. The angle of incidence of the aerofoil counts for the 17th design variable.

7.1.1 Influence of the levels of convergence

7.1.1.1 Flow solver on objective function

This paragraph looks at the influence of the level of convergence of the flow solver MERLIN on the value of the objective function (here lift and drag coefficients). This is done by stopping the flow calculation at different stages of convergence and by comparing each time, the output aerodynamic coefficients with these obtained with a 12 order converged flow solution. This postulates that the 12 order converged solution gives accurate aerodynamic coefficients based on the given physical model, discretisation and grid. The accuracy of the drag coefficient at the j^{th} order of convergence is for example given by

$$\text{accuracy of } C_D \text{ at } j \text{ orders} = \frac{|C_{D,12 \text{ orders}} - C_{D,j \text{ orders}}|}{|C_{D,12 \text{ orders}}|}$$

The convergence level is measured by the total residual R_{total} normalised by the freestream total residual. This investigation was carried out for both the lift and drag coefficients and

its results can be seen in Figure 7.1. The accuracy is represented using both a linear scale and a logarithmic scale. For engineering accuracy the linear scale is sufficient and shows that there is no need to go beyond 4 orders of convergence to get an accurate value of C_L and C_D . The logarithmic scale shows that the accuracy is still improved by continuing to converge the solution. It also shows that the behaviour of the lift and drag coefficients is very similar. The conclusion of this investigation is that a modest level of convergence of the flow solver is required to get accurate values of its output aerodynamic coefficients.

7.1.1.2 Flow solver on gradient


The influence of the level of convergence of the flow solver MERLIN on the accuracy of the gradient is investigated here. This is inspired by a similar study by Nadarajah in Reference [56]. Reference [45] looks at the same problem but presented in a different way.

The present investigation is done by stopping the flow solver at different stages of convergence as in the previous paragraph. The obtained flow solution is then used by the adjoint solver to calculate the sensitivity derivatives of either lift or drag coefficient. To avoid any effect of the level of convergence of the adjoint solver, it is converged each time to 12 orders. Each gradient is composed of 17 components corresponding to the 17 design variables. The accuracy of each component is assessed by comparison to the same component obtained when both the flow and adjoint solvers are converged to 12 orders. For the k^{th} component of the gradient for drag, the accuracy at the j^{th} order of convergence of the flow solver is measured as

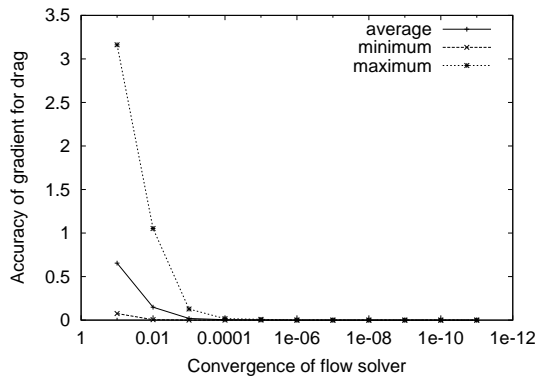
$$\text{accuracy of } k^{\text{th}} \text{ gradient at } j^{\text{th}} \text{ order} = \frac{\left| \frac{dC_D}{d\beta_k} \Big|_{12 \text{ orders}} - \frac{dC_D}{d\beta_k} \Big|_{j \text{ orders}} \right|}{\left| \frac{dC_D}{d\beta_k} \Big|_{12 \text{ orders}}} \quad (7.1)$$

This is presented in Figure 7.2(a) for all 17 design variables. Since the details for all the curves is not very important, the envelope of these curves (minimum and maximum) and the average over the 17 components are added to Figure 7.2(a). In order to clarify the pictures, only these three curves are presented for the rest of the results. The same simplified graph is shown in linear scale in Figure 7.2(b) and in logarithmic scale in Figure 7.2(c). Figure 7.3 shows the results for the sensitivity derivatives of lift coefficient.


The behaviour of the lift and drag coefficients is very similar. Both gradients show some sensitivity to the convergence of the flow solver and the more the flow solution is converged, the better it is for the gradient. The gradient appears to be one order of magnitude more sensitive to the convergence of the flow solver than the objective function itself. This implies that convergence levels of 5 orders or more are now required to get an accurate gradient.



(a) All the 17 design variables. Logarithmic scale



(b) Linear scale



(c) Logarithmic scale

Figure 7.2: Influence of the level of convergence of the flow solver on the value of the sensitivity derivatives of drag coefficient. Accuracy of k^{th} gradient at j^{th} order =

$$\frac{\left| \frac{dC_D}{d\beta_k} \right|_{12 \text{ orders}} - \left| \frac{dC_D}{d\beta_k} \right|_{j \text{ orders}}}{\left| \frac{dC_D}{d\beta_k} \right|_{12 \text{ orders}}}$$




Figure 7.3: Influence of the level of convergence of the flow solver on the value of the sensitivity derivatives of lift coefficient. Accuracy of k^{th} gradient at j^{th} order =

$$\frac{\left| \frac{dC_L}{d\beta_k} \text{ 12 orders} - \frac{dC_L}{d\beta_k} \text{ } j \text{ orders} \right|}{\left| \frac{dC_L}{d\beta_k} \text{ 12 orders} \right|}$$

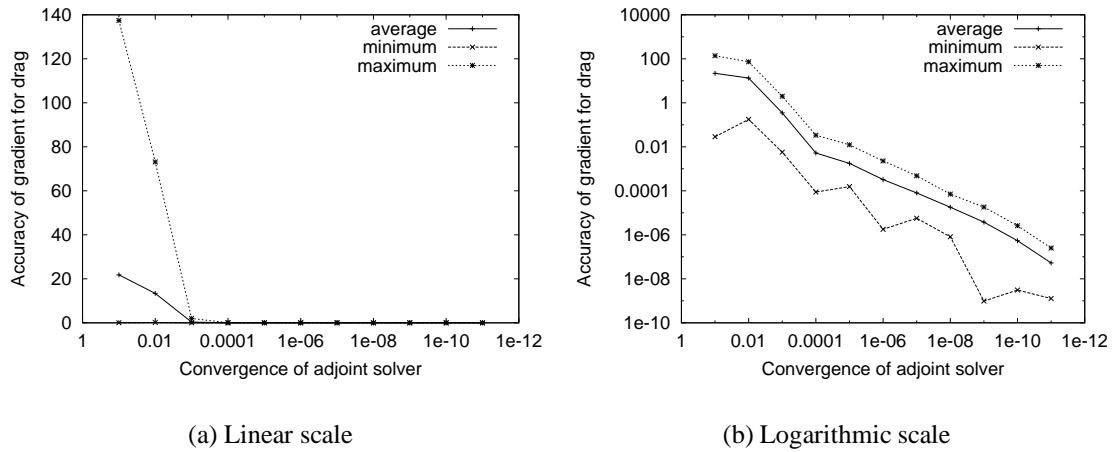


Figure 7.4: Influence of the level of convergence of the adjoint solver on the value of the sensitivity derivatives of drag coefficient. Accuracy of k^{th} gradient at j^{th} order =

$$\frac{\left| \frac{dC_D}{d\beta_k} \text{ 12 orders} - \frac{dC_D}{d\beta_k} \text{ } j \text{ orders} \right|}{\left| \frac{dC_D}{d\beta_k} \text{ 12 orders} \right|}$$

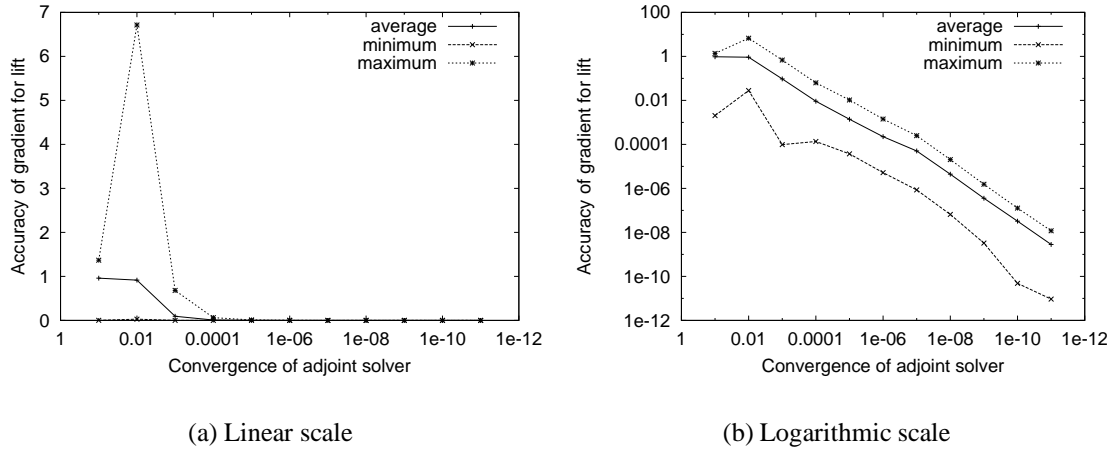


Figure 7.5: Influence of the level of convergence of the adjoint solver on the value of the sensitivity derivatives of lift coefficient. Accuracy of k^{th} gradient at j^{th} order =

$$\frac{\left| \frac{dC_L}{d\beta_k} \text{ 12 orders} - \frac{dC_L}{d\beta_k} \text{ } j \text{ orders} \right|}{\left| \frac{dC_L}{d\beta_k} \text{ 12 orders} \right|}$$

7.1.1.3 Adjoint solver on gradient

In this paragraph, the influence of the level of convergence of the adjoint on the gradient is investigated. A similar study was carried out by Nadarajah.^[56] In this work, the flow solver is always converged to 12 orders to remove its influence while the adjoint solver is converged to different levels. Each time the 17 components of the gradient are calculated and compared to the reference provided with 12 order converged flow and adjoint solutions. Equation (7.1) is still used to measure the accuracy of the gradient but this time the level of convergence refers to the adjoint solver rather than the flow solver.

Figures 7.4 and 7.5 show the results of this study for drag and lift coefficient gradients respectively. The gradient appears to be much more sensitive to the convergence of the adjoint solver than to the convergence of the flow solver, especially for the drag coefficient. This is in contradiction to the findings in Reference [56] where 1 order of convergence of the adjoint solver is sufficient to get a gradient accurate to 4 significant digits. Here this accuracy is obtained after more than 4 orders of convergence for the least sensitive design variable and up to 8 orders for the most sensitive one. This clearly shows that the convergence of the adjoint solver is very important here and that it needs to be converged to 5 or 6 orders in order to get a reasonably accurate gradient. Several reasons might explain the difference with Nadarajah's work: first, Nadarajah does not explain how he measures the convergence of the adjoint solver. Is it based on the total residual or on the residual of only the first component of the adjoint vector? And how is this residual non-dimensionalised? Secondly the reference gradient is not clearly identifiable




Figure 7.6: Influence of the grid size on the value of the sensitivity derivatives of drag, lift and pitching moment coefficients. $\text{Ratio} = \frac{\text{sensitivity obtained on } 65 \times 17 \text{ grid}}{\text{sensitivity obtained on } 129 \times 33 \text{ grid}}$.

in Reference [56]. Finally, and this is perhaps the main reason, Nadarajah's adjoint solver does not linearise the turbulent viscosity and considers it as a constant. It is possible that the turbulent viscosity plays an important role here and the way in which it is linearised in the present adjoint solver might require a good convergence level to provide accurate gradients.

This concludes the section on the influence of convergence levels on the objective function and its gradient. The main conclusion is that good convergence of both the flow and adjoint solvers is required to obtain reasonably accurate gradients. It should not be forgotten that the influence of the adjoint solver and of the flow solver have been considered in this section independently from each other, one of the two codes being well converged each time. In practice to save computing time, both will be converged to less stringent requirements and this is likely to have a cumulative negative effect on the accuracy of the gradient. The next section briefly looks at the effect of grid size on the gradient.

7.1.2 Influence of the grid size

The aim of this short section is not to investigate in detail the influence of grid refinement on the adjoint. The idea was that since the adjoint vector and hence the gradient are quite costly to calculate, a possible way of reducing computing time might be to calculate them on a coarser grid than the one used for the flow solution. Hence the accuracy of a gradient obtained on a coarse grid is investigated. This is how it is done: the 129×33 C-type grid around the RAE2822 aerofoil is the basis grid. Note that this is already a coarse grid for Navier-Stokes calculations. A flow solution converged to 12 orders is calculated on this grid as well as a gradient converged to 12 orders too. This latter serves as a reference.

A very coarse grid is made by removing every other points in each direction to obtain a 65×17 grid. The 12 order converged flow solution is transferred to this coarse grid without making any flow computation: an average involving the four (2-D problem) flow values surrounding each subgrid cell-centre is performed to get the relevant flow value for this subgrid cell. The adjoint solver is then run on the coarse grid using this transferred flow solution. The obtained gradient is compared to the reference gradient component by component.

The results are shown in Figure 7.6 that presents for each design variable the ratio

$$\text{ratio for drag sensitivity } k = \frac{\frac{dC_D}{d\beta_k}_{65 \times 17}}{\frac{dC_D}{d\beta_k}_{129 \times 33}}$$

This is done for drag, lift and pitching moment coefficients. Obviously if the grid size had no influence, the ratio would be 1 for all the variables. A line indicating this value of 1 is plotted in Figure 7.6 to enable the comparison. It is clear from this Figure, that calculating the gradient from a subgrid of the grid used for the flow solution is not appropriate at all since some of the sensitivity derivatives obtained in this way are several times higher or smaller than what they should be, some even being of the wrong sign, which could be damaging in an optimisation process. Note that the ratio for the 9th design variable (that corresponds to the active design variable the closest to the leading edge on the lower surface) for the pitching moment could not be represented on the graph due to its very large magnitude. The subgrid is certainly much too coarse to give a proper representation of the flow field to the adjoint solver. The idea of calculating the adjoint on a subgrid to save computing time was thus abandoned through this investigation.

7.1.3 Influence of the physical model

As in the previous section, the idea behind the investigation in this section is to reduce the cost of computing the gradient. This is done by changing the physical model employed in the adjoint solver. The flow solution is still fully turbulent but the physical model in the adjoint solver is downgraded either to turbulent without taking into account the linearisation of the turbulent viscosity ($\mu_t = \text{constant}$) or to viscous laminar. The reference for the gradient is still with a fully turbulent adjoint. These simplifications surely save some computing time, but is it at the cost of accuracy?

The answer to this question lies in Figures 7.7 and 7.8 that show for each design variable k the ratios

$$\text{ratio for drag sensitivity } k = \frac{\frac{dC_D}{d\beta_k}_{\text{turbulent adjoint with } \mu_t = \text{constant}}}{\frac{dC_D}{d\beta_k}_{\text{turbulent adjoint}}}$$




Figure 7.7: Influence of the physical model (here turbulent adjoint but with $\mu_t = constant$) on the value of the sensitivity derivatives of drag, lift and pitching moment coefficients. Ratio = $\frac{\text{sensitivity obtained with a turbulent adjoint with } \mu_t = constant}{\text{sensitivity obtained with a turbulent adjoint}}$.

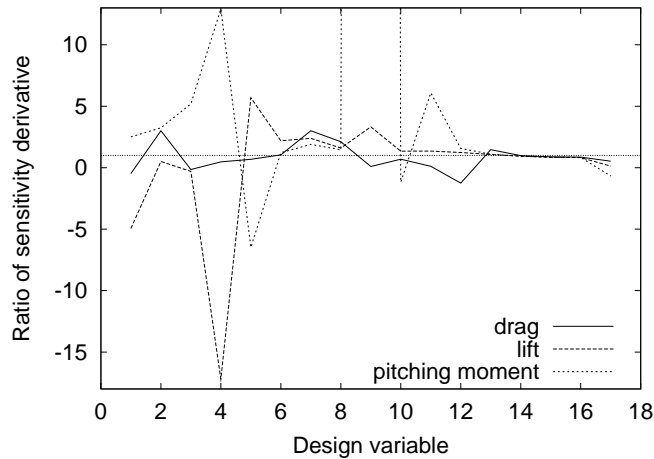


Figure 7.8: Influence of the physical model (here viscous laminar adjoint) on the value of the sensitivity derivatives of drag, lift and pitching moment coefficients. Ratio = $\frac{\text{sensitivity obtained with a viscous laminar adjoint}}{\text{sensitivity obtained with a turbulent adjoint}}$.

and

$$\text{ratio for drag sensitivity } k = \frac{\frac{dC_D}{d\beta_k \text{ viscous laminar adjoint}}}{\frac{dC_D}{d\beta_k \text{ turbulent adjoint}}}$$

respectively. This is done for drag, lift and pitching moment coefficients. The results prove again that this kind of approximation is not valid in order to calculate the gradient since some of the design variables are several orders of magnitude wrong, some also having the wrong sign. In both graphs, the ratio for the 9th design variable could not be plotted due to its large magnitude. It has to be noted however that the approximation of discarding only the linearisation of the turbulent viscosity is less damaging than discarding the whole turbulent viscosity as expected.

This completes the first part of the chapter that investigated the influence of possible approximations on the accuracy of the gradient and the objective function. The aim was to reduce as much as possible the computing time needed to calculate both of these in an optimisation process. The main conclusion is that no obvious simplification was found to accelerate the computations and on the contrary, good convergence levels of both the flow and adjoint solvers are required to obtain reasonably accurate gradients. Remember that this is only considering gradients for Navier-Stokes flows, maybe the conclusions would be different for inviscid or laminar flows. In addition the study is made in a particular case for a given aerofoil geometry, computational grid and flow conditions. This only shows that for this case, approximations are not valid but it might be possible to find other instances where they work.

However this part only looked at the gradients taken in isolation which might be biased since the gradients are only one of the components of the optimisation. It is more likely though that if each computation of the gradient is accurate, then the optimisation process will have no problem but it might be possible that the optimiser can accommodate less accurate gradients and this could possibly result in lower computing time. That is why the next part of this chapter investigates the influence of approximations no longer only on the gradients alone but on the whole optimisation process. This is done on a two-dimensional drag minimisation problem using a direct optimisation method based on Sequential Quadratic Programming.

7.2 Two-dimensional optimisation using a direct SQP method

Before looking at the effects of approximation on optimisation, this part introduces the optimisation problem to be solved that will be used as a test case. It then details the scaling of the design variables and the constraints and this will serve as a presentation of the

reference optimisation results that will be compared to the results obtained with approximations that will then be examined. Finally the problem of how to handle the constraints will be studied.

7.2.1 The optimisation problem

The optimisation problem that serves as a test case in this part is the constrained drag minimisation of an aerofoil. The original aerofoil section is taken from the outer wing of a Blended Wing-Body. Because of the wing sweep angle, it has to be transformed to two dimensions before doing the optimisation. Likewise the three-dimensional flow conditions are transformed to two dimensions to obtain: Mach number $M_\infty = 0.725$, unit Reynolds number $Re = 5.691 \times 10^6 / m$ and freestream temperature $T_\infty = 218.8 K$. The local lift coefficient at which the aerofoil operates, is also transformed into two dimensions and serves as a target C_L at which the optimisation has to be carried out. A geometrical constraint on the interior volume (in fact, it is an area since the problem is only two-dimensional) of the aerofoil is added to prevent the optimiser from thinning it too much in order to reduce drag. The optimisation problem is as follows.

$$\begin{aligned} \text{Minimise} \quad & C_D \\ \text{Subject to:} \quad & C_L \geq 0.652 \\ & V_0 \leq V \leq 2 \times V_0 \end{aligned} \tag{7.2}$$

where V is the interior volume and V_0 the interior volume of the baseline aerofoil equal to 5.99×10^{-2} . Note that the constraints have been set to inequality constraints even if equality constraints are desired. The reason for this is that some tests were carried out using equality constraints but the optimiser was spending most of its time trying to satisfy these demanding constraints rather than doing some optimisation. Transforming them into inequalities gives much more freedom to the optimiser and since the problem is to minimise drag, the optimum design should be very close to the minimum admissible C_L and volume and thus be very close to satisfying equality constraints.

The upper and lower surfaces of the aerofoil are free to deform and the perturbation added to the baseline curves is parameterised using 10 Bézier parameters for each surface. Since the leading and trailing edge points are fixed, this makes 16 design variables. In addition the angle of incidence is considered as a design variable.

Indeed the constraint on lift in this work is handled as a hard constraint and we let the optimiser decide how it combines changes in angle of incidence and in geometry to match the required C_L . This is contrary to what is usually done in aerodynamic optimisation where the angle of incidence is not governed by the optimiser but is changed by the flow solver to match the required lift.^[7,45,57,83,201] In such approaches the constraint on lift does not appear as a constraint in the optimisation, it is only an internal constraint to the flow solver like the angle of incidence is an internal free variable. The approach taken here where the optimiser controls the incidence in addition to shape variables, is thought

to be sounder since it lets the optimiser decide what is best for the optimisation problem, either a change of shape or a change of incidence or a combination of both, since it is able to see their effects both on lift and drag independently. In the traditional approach, it only sees the value of drag and has no knowledge of the effects of design changes on lift. The method considered here has however a cost since it requires the calculation of the gradient for lift that is not needed in the traditional approach but this is certainly compensated by the cost of each flow calculation that is done at constant angle of incidence rather than constant C_L and hence should be faster.

The geometrical constraint on interior volume is also handled as a hard constraint but this does not pose any problem. The volume is easily calculated analytically from the position of the grid points defining the geometry. Likewise the gradient of the volume with respect to the design variables is calculated analytically from the grid sensitivities at a very low cost.

Bounds are imposed on the design variables. Constraining lower bounds are imposed on the design variables parameterising the upper surface curve in order to avoid as much as possible surface crossings and grid deformation problems. For instance, the first design variable close to the leading edge is obliged to be always positive to prevent points close to the leading edge from passing under the y coordinate of the leading edge point, which would create grid generation problems. The other design variables for the upper surface can have moderate negative values in order to change the curvature of the surface if necessary but this is limited to avoid crossing the lower surface curve and generating a non-physical aerofoil. In the same way, design variables for the lower surface are limited by an upper bound allowing only moderate positive values. Note that due to the range still allowed to the design variables, it is still possible to get surface crossing, a more explicit geometrical constraint would have to be set up to avoid it completely. Avoiding it by putting bounds on the design variables can also be achieved but this would seriously limit the possible shape deformations. In addition it was found that mild surface crossing do not cause problems to the grid deformation code or to the CFD solver and that the optimiser does not stay very long with designs of this kind since the constraint on interior volume and the decrease of drag are difficult to achieve.

The baseline geometry and the CFD grid around it are shown in Figure 7.9. The grid is a C-type grid composed of $137 \times 35 \times 2$ points. The optimisation is started with the angle of incidence set to $\alpha = 2^\circ$ that makes the initial design point infeasible for the constraint on lift. Unless stated otherwise, the flow solver is run each time until the total residual has reached 10^{-6} or until 5000 implicit iterations have been completed, whichever comes first. The adjoint solver is stopped when the residual has also reached 10^{-6} or when 3000 implicit iterations have been performed, whichever comes first as before. The number of maximum iterations chosen should be sufficient to reach the 6 order convergence requirement when starting from freestream conditions. In addition both the flow and adjoint solvers use their restart capabilities to start each calculation from the previous solution in




Figure 7.9: Baseline geometry and the CFD grid around it.

order to reduce as much as possible the computing time. All the optimisations are started from converged flow and adjoint solutions on the baseline geometry. All the calculations are run on an Alpha workstation with a single 667 MHz processor to enable comparisons.

7.2.2 Scaling and reference optimisation

As pointed out in Reference [113], scaling can refer to a lot of things in optimisation. Here only the scaling of the design variables and the scaling of the objective function and the constraints will be discussed.

The scaling of the design variables is important as the example given by Vanderplaats^[112] shows: a two-dimensional design space that has the shape of a narrow valley for which a gradient-based optimiser has a lot of difficulties to operate, can be transformed into a perfectly circular bowl-type design space where the optimiser would find the minimum in a few iterations, simply by scaling appropriately the design variables. This has nothing to do with the physics of the optimisation problem to solve, it is just a mathematical transformation that should improve the conditioning of the problem. The main idea for scaling the design variables is that they should all have approximately the same magnitude. This is usually achieved by dividing each design variable by its initial physical value to start with all design variables set to 1. This is not possible with the present parameterisation since it is a perturbation that is represented by the variables and hence their initial value is zero. The scaling of the design variables has been obtained by imposing that an increment of 1 of each design variable taken in isolation produces a maximum displacement of 2% of the aerofoil chord. This gives approximately the same weight to all the design variables and enables shape parameters and angle of incidence that are variables of different nature, to be mixed together.

Scaling the objective functions and the constraints is necessary when using a penalty term for the constraint added to the objective function. It is also a good practice to give a similar weight to all the constraints when the problem contains more than one constraint. Finally it is also desirable that the magnitude of the objective function be of the order of unity.^[113] There is not any simple method to achieve all of this at the same time. Here we follow the approach given in Reference [112] for scaling the constraints when using penalty terms, that is based on the rate of change of the functions rather than on their magnitude. The method is slightly modified but is essentially the same: the objective and constraint functions are all divided by the norm of their initial gradient so that at the beginning of the optimisation, they all have a gradient with the same norm equal to 1. The L_2 norm is used to calculate the norm of the gradients.

This scaling is applied consistently to all the optimisation problems solved in this thesis even when a penalty term is not employed. This is to maintain consistency and to enable comparison. Note that this scaling, as well as the scaling of the variables, comes in addition to possible internal scaling done by the optimisation routines employed in this work.

		C_L	C_D	Improvement in C_D
Initial		0.6520	0.01894	-
Optimised	scaled	0.6520	0.01377	-27.3%
	not scaled	0.6520	0.01379	-27.2%

Table 7.1: Aerodynamic coefficients of the optimised aerofoils at the target C_L .



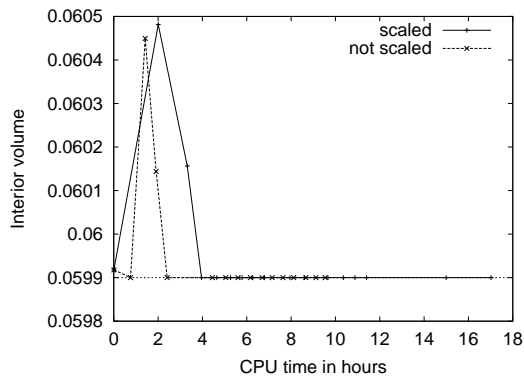
To assess the effect of the scaling of the objective function and the constraints, the two-dimensional test case optimisation problem is solved with and without scaling. With scaling, it becomes

$$\begin{aligned}
 \text{Minimise} \quad & \frac{C_D}{\left\| \frac{dC_D}{d\beta^0} \right\|} = 0.02292 \\
 \text{Subject to:} \quad & \frac{C_L}{\left\| \frac{dC_L}{d\beta^0} \right\|} \geq \frac{0.652}{\left\| \frac{dC_L}{d\beta^0} \right\|} = 0.4754 \\
 & \frac{V_0}{\left\| \frac{dV_0}{d\beta^0} \right\|} \leq \frac{V}{\left\| \frac{dV_0}{d\beta^0} \right\|} \leq \frac{2 \times V_0}{\left\| \frac{dV_0}{d\beta^0} \right\|} = 0.02872
 \end{aligned} \tag{7.3}$$

Without scaling it is solved as written in Equation (7.2). The optimisation with scaling will serve as a reference optimisation to which other methods will be compared. The NAG routine E04UCF is the SQP optimiser used.

The results of these two optimisations are presented in Figure 7.10 that shows the evolution of interesting parameters during optimisation. The same format will be used to present other optimisation results. Figure 7.10(a) presents the evolution of the drag coefficient C_D that is the objective function to minimise, against the CPU time measured in hours for the optimisations with and without scaling. Both methods are very similar and manage to reduce drag. This is even more impressive when looking at the values of C_D obtained at the correct lift in Table 7.1: a 27% improvement is achieved compared to the initial geometry. The method with scaling takes a long time though to reach its best point but the improvement over previous points is marginal after 10 hours of calculation at which point both methods have reached the same drag improvement, the method without scaling being slightly quicker.

Figure 7.10(b) shows the evolution of the lift coefficient during optimisation. This is still plotted against CPU time. There is a constraint on C_L and its boundary is shown by a dashed line at 0.652. As already explained, the initial point is infeasible for this constraint since it starts at $C_L = 0.576$. Both methods manage to satisfy the constraint at the end of

(a) Objective function C_D (b) Lift coefficient C_L 

(c) Interior volume

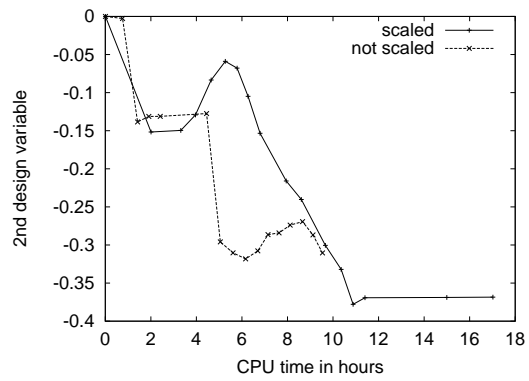


(d) 2nd design variable(e) 6th design variable(f) 17th design variable: increment in angle of incidence

Figure 7.10: Evolution of different parameters during the aerofoil optimisations comparing the method with scaling that serves as reference to the method without scaling.


the optimisation and it can be seen that the NAG routine E04UCF is the optimiser since it allows for infeasible points during optimisation. The \times and $+$ indicate optimisation cycles so it is clear that the method without scaling reaches the constraint boundary much quicker in terms of computing time as well as number of iterations than the method with scaling.

The evolution of the interior volume, the other constraint of the problem, is plotted in Figure 7.10(c). The boundary for this constraint is $V_0 = 0.0599$, the baseline geometry being just feasible obviously. Both methods satisfy well this geometrical constraint and as expected, the minimum drag is achieved at the minimum allowable interior volume.

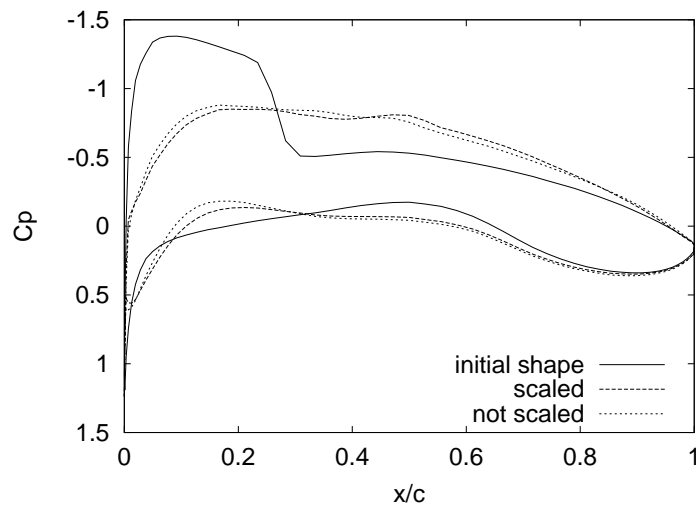
Figures 7.10(d) to 7.10(f) show the evolution of three design variables against time. The 2nd design variable has an influence on the upper surface of the aerofoil very close to the leading edge while the influence of the 6th design variable is felt further aft but still on the upper surface. There is no particular reason for choosing these two shape variables, they are just an example of what other variables might do as well. The 17th design variable is the increment in angle of incidence. It is plotted since it is the only design variable not controlling shape, so its behaviour might be interesting. From Figures 7.10(d) to 7.10(f) it can be seen that as already explained, the optimisation starts with a value of zero for all the design variables. Note that the values plotted are scaled variables, not the true value of the Bézier parameters or the increment in angle of incidence. For this latter however, the scaling coefficient is 1.146 so the value displayed is close to the true increment in degrees. The methods with and without scaling are very similar again, reaching almost the same design points. The method without scaling seems to have been cut short of reaching a plateau though, while the method with scaling does not evolve after 11 hours of calculations. This confirms the impression given by Figure 7.10(a) that the minimum is already reached after 10 ~ 11 hours. It is also interesting to notice that after 6 hours of computation, the design variables are still evolving a lot although drag, lift and interior volume are almost unchanged.

The optimum found by both methods is shown in Figure 7.11. Figure 7.11(a) presents the two optimal shapes in addition to the baseline geometry. The modifications of this latter are quite important, the optimised aerofoil having more camber. The optimal shapes are very close to each other as expected, the differences observed in the design variables being translated nevertheless in identifiable shape differences.

The chordwise pressure distribution on these shapes is plotted in Figure 7.11(b). The pressure distribution for the baseline geometry has been obtained for a C_L of 0.652 and not for the initial C_L while the other two pressure distributions are extracted from the results of the optimisation whatever the final C_L . As seen from Figure 7.10(b), this should be very close to 0.652 anyway. A very strong shock wave can be observed at 25% chord on the upper surface of the baseline aerofoil. This shock wave has been almost completely eliminated by the optimisation even if there is still a very small pressure bump



(a) Aerofoil shapes



(b) Chordwise pressure distributions

Figure 7.11: Result of the aerofoil optimisations comparing the method with scaling that serves as reference to the method without scaling.

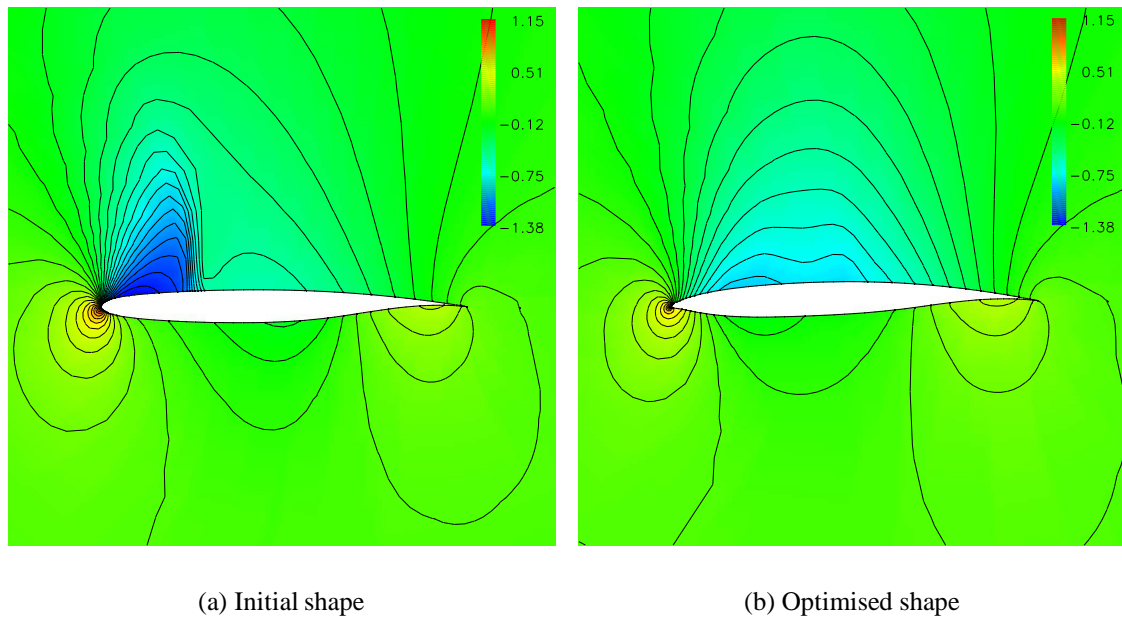


Figure 7.12: Contours of pressure coefficient on the initial shape and on the optimised shape obtained by using the scaled optimisation method.

on the aerofoil optimised by the method with scaling. This elimination of the shock wave was expected and it proves that the drag minimisation method is working since by doing so, the optimisation gets rid of the wave drag. Since the optimised shapes are slightly different, the pressure distributions are also so but they are still very close.

The elimination of the shock wave can also be seen in Figure 7.12 that shows the flow field pressure contours on the baseline aerofoil and the aerofoil obtained with the method with scaling. The flow field on this latter shape is represented here since as already explained, the method with scaling is used as a reference for the rest of the two-dimensional optimisations presented afterwards. This is the only figure representing the optimised flow field that will be shown since, when the optimised designs are close to each other, it is not possible to differentiate them on this kind of picture. The shock wave at 25% chord on the initial aerofoil can clearly be seen in Figure 7.12(a). The absence of it can equally be seen on the optimised shape. Since the scale of colours for the pressure coefficient is the same for both figures, the reduction of the suction on the upper surface of the optimised aerofoil is visible as it is on the pressure distribution of Figure 7.11(b).

This is the end of the presentation of the results of the first two-dimensional optimisations that compared the method with scaling of the objective function and the constraints to the method without scaling. As explained previously, the method with scaling is chosen as a reference for the rest of the results examined in this chapter and all the other optimisations

will be scaled as well. From these results it seems however that the method without scaling is slightly quicker hence better. The method with scaling is nevertheless kept as a reference to enable comparison with later optimisations involving penalty terms where some scaling is necessary.

This section, after explaining the scaling of the design variables and the scaling of the objective function and the constraints, proved that the optimisation chain set up in this study is working and that large drag reduction can be achieved, at least in two dimensions. However it also showed that the optimisation is time consuming: 10 hours of CPU time might be alright but this is only for a two-dimensional relatively coarse Navier-Stokes grid. Since the aim of this thesis is to apply optimisation to a three-dimensional Blended Wing-Body, computational times involved will become huge with the current method. Hence some testing aiming at reducing the cost of optimisation on this simple two-dimensional problem was carried out as the next section explains.

7.2.3 Influence of approximation on optimisation


This section reconsiders some ideas encountered in the first part of this chapter in Section 7.1 but is looking at the overall optimisation rather than at a single gradient. The aim as just explained, is to reduce computing time while maintaining a good level of optimisation compared to the method of reference presented in the previous section. Throughout this section the optimiser used is the NAG routine E04UCF and the optimisation problem solved is problem (7.3). The first test is studying the influence of the level of convergence of the adjoint solver on the optimisation.

7.2.3.1 Influence of the level of convergence of the adjoint solver

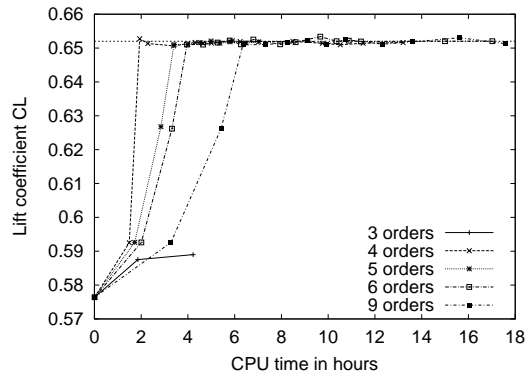
The reference optimisation converges the adjoint solver to 10^{-6} each time, this subsection looks at what is happening for different levels of convergence. Four additional optimisations are carried out with an adjoint solver converged to 10^{-3} , 10^{-4} , 10^{-5} and 10^{-9} respectively. The comparison of these optimisations is made in Figure 7.13.

From all the graphs in Figure 7.13, it appears clearly that the optimisation where the adjoint convergence is limited to 3 orders fails prematurely. The very coarse approximation of the gradient introduced by this poor convergence certainly causes trouble for the optimiser which is unable to find a feasible descent direction.

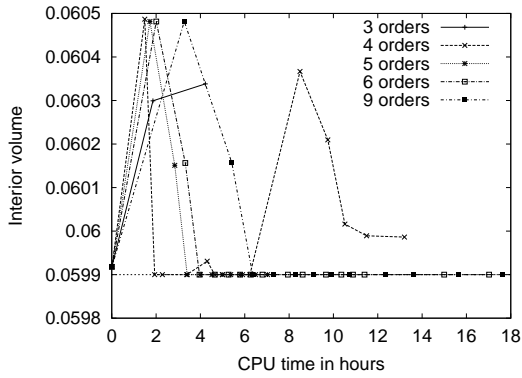
The optimisation with a 4-order converged adjoint encounters some difficulties as well most noticeably seen in Figure 7.13(c) where the boundary for the interior volume is quickly reached but the optimiser bounces back again and never manages to reduce drag by thinning the aerofoil. The approximation is certainly too strong for the optimiser and when considering drag in Figure 7.13(a) the method is never competitive.



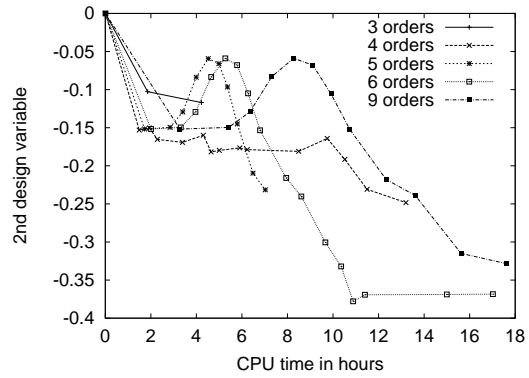
(a) Objective function C_D



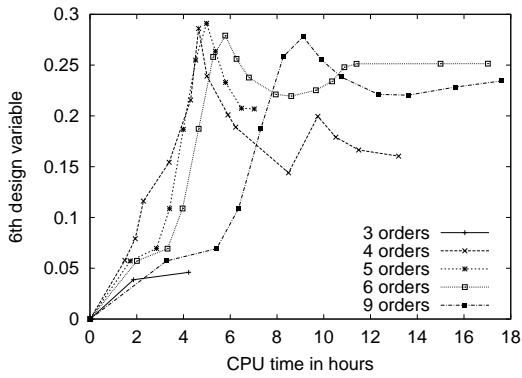
(b) Lift coefficient C_L




(c) Interior volume



(d) 2nd design variable



(e) 6th design variable



(f) 17th design variable: increment in angle of incidence

Figure 7.13: Evolution of different parameters during the aerofoil optimisations comparing the influence of the level of convergence of the adjoint solver.

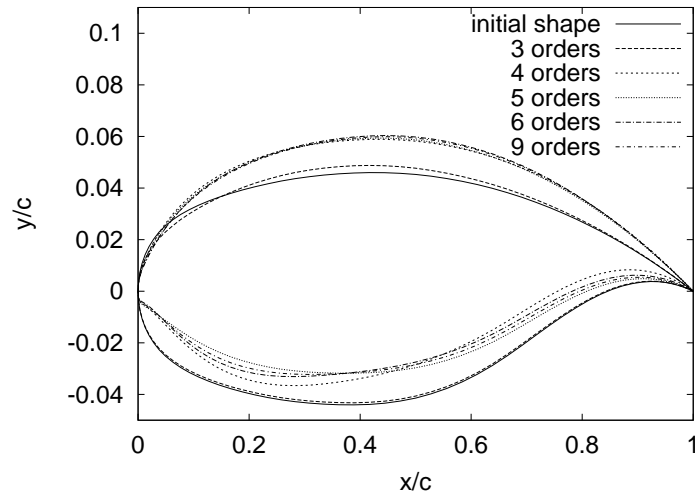
		C_L	C_D	Improvement in C_D
Initial		0.6520	0.01894	-
Optimised	3 orders	0.6520	0.01554	-18.0%
	4 orders	0.6520	0.01391	-26.6%
	5 orders	0.6520	0.01387	-26.8%
	6 orders	0.6520	0.01377	-27.3%
	9 orders	0.6520	0.01380	-27.1%

Table 7.2: Aerodynamic coefficients of the optimised aerofoils for the optimisations comparing the influence of the level of convergence of the adjoint solver, at the target C_L .


The other three optimisations are very similar, only computing time differentiating them. Indeed the symbols on the graphs of Figure 7.13 correspond to optimisation iterations and show that the three optimisations are passing by the same values at the end of each cycle. The only difference is the CPU time that shifts the points horizontally. There is a consistency in this shifting since the optimisation where the adjoint is converged to 5 orders is the quickest, followed by the reference optimisation with 6 orders and the optimisation with 9 orders. The latter is very expensive. There are however some slight differences in the path of the optimisation with a 5-order converged adjoint noticed for example at the peak value after 5 hours in Figure 7.13(e), that might explain the abrupt interruption of the optimisation. 5 orders might just be the limit of consistency of the gradients. The 9-order method is very costly and would not be used in practice. It proves however that the reference 6-order method has the correct level of approximation when considering the adjoint solver.

Table 7.2 looks in more details at the optimised aerofoils. Despite being not successful, the optimisation with a 3-order converged adjoint manages to reduce drag by 18% over the initial geometry. The other optimisations are relatively close to each other, those with a well-converged adjoint achieving slightly better than the other two as would be expected. It is a pity that the 4-order method is not time efficient because it does reduce the drag considerably. The 5-order method is interesting in the light of these results: it does not minimise drag as well as the methods with well-converged adjoints, but its improvements are already high and it is faster. Hence considering a compromise between computing time and accuracy, it is difficult to choose which method is best between the 5-order and the 6-order converged adjoint.

Although it is not performing well when looking at the objective function or constraints, the method with a 4-order converged adjoint ends up with design variables close to the optimum when considering Figures 7.13(d) to 7.13(f). This is translated in the final shapes and pressure distributions of Figure 7.14: omitting the 3-order shape that is poorly optimised, the other four optima are very close to each other. The main differences in shape are seen on the lower surface of the aerofoil and are not easily explained. This



(a) Aerofoil shapes



(b) Chordwise pressure distributions

Figure 7.14: Result of the aerofoil optimisations comparing the influence of the level of convergence of the adjoint solver.

obviously has some effect on the pressure distribution for the lower side as seen in Figure 7.14(b). Although shape differences are very small on the upper surface, they still have some impact on the C_p distribution: the further the adjoint solver is converged, the better it is (since the 4-order method shows some remains of the shock wave that are gradually smeared out as the method uses a better converged adjoint).


This terminates this subsection investigating the effects of the level of convergence of the adjoint solver on the overall optimisation. It showed that poor convergence is damaging, causing the optimisation to fail but it also gave indications that slightly reducing the level of convergence might be a way of reducing computing time. The method with a 5-order converged adjoint is indeed not far from being the best: its optimum is very close to the optimum of more accurate methods and it is reached more rapidly. It is a very good compromise between reaching the true minimum and reaching this minimum as quickly as possible. This should be particularly interesting for three-dimensional optimisation.

7.2.3.2 Influence of the grid size

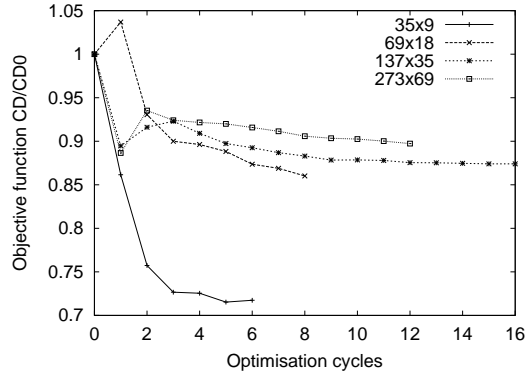
This subsection examines the influence of the grid refinement on the overall optimisation. Here the grid is the same for the flow solver and for the adjoint solver. The reference optimisation uses a 137×35 grid shown in Figure 7.9. A finer grid is made by doubling the number of points in each direction to obtain a 273×69 grid and two subgrids are also created by removing one point out of two in each direction from the reference grid to obtain a 69×18 grid and a 35×9 grid. This latter grid is extremely coarse and is employed only for numerical investigation because it is expected that little of the flow physics can be captured by such a grid. These four grids are used for optimisation, the results being presented in Figure 7.15.

In Figure 7.15(a), the evolution of the drag coefficient for the four optimisations is plotted against computing time. Clearly the optimisation on the fine grid takes a very long time compared to the optimisation on the coarsest grid, which makes the graph difficult to read. That is why the remaining of the results are plotted against optimisation cycles rather than computing time to eliminate this scale effect. Since it is the only graph showing CPU time, the fact that the optimisation on the fine grid does not start at 0 hour has to be explained here: the flow solver could not indeed reach the required convergence level of 10^{-6} and stalled before that. Since this level of convergence is reached later during the optimisation on modified geometries and since this same requirement is used in the other optimisations, it was applied to the initial baseline and the flow solver wasted a long time reaching the maximum number of iterations.


The other important point to notice in Figure 7.15(a) is that the optimisations are starting from different values of drag. They indeed only share the initial angle of incidence of $\alpha = 2^\circ$ and since the grids are very different, they produce different initial values of drag but also of lift and interior volume as seen in Figures 7.15(c) and 7.15(d) respectively.




(a) Objective function C_D vs CPU time




(b) Ratio C_D/C_{D0} vs optimisation cycles




(c) Lift coefficient C_L



(d) Interior volume



(e) 2nd design variable



(f) 6th design variable

Figure 7.15: Evolution of different parameters during the aerofoil optimisations comparing the influence of the grid size.


(g) 17th design variable: increment in angle of incidence

Figure 7.15: Evolution of different parameters during the aerofoil optimisations comparing the influence of the grid size. (Concluded)

The difference in drag between the coarsest and the finest grid is almost an order of magnitude but there is some consistency because as the grid is refined, C_D converges to a common value. Notice that the reference grid 137×35 gives a drag value close to the fine grid, which justifies its choice as a reference.

This difference in drag between the grids makes it very difficult again to see any drag improvement in Figure 7.15(a). That is why the evolution of the drag coefficient is plotted again but as a ratio $\frac{C_D}{C_{D0}}$ of the drag coefficient over the initial drag coefficient in Figure 7.15(b). This shows, together with the evolution of C_L and of the interior volume in Figures 7.15(c) and 7.15(d) respectively, that the four optimisations are successful in the sense that they reduce drag while satisfying the constraints on their lower boundaries. Moreover the improvement in drag is increasing as the mesh becomes coarser. This seems to come from the poor accuracy of the discretisation on the coarse grids and does not have any physical backing unfortunately.

The aim of performing such testing was to see if the optimisation would follow the same path whatever the grid refinement or even reach the same optimum in terms of design variables even if the physics of the problem is not resolved accurately or even not at all. Since the coarse grid optimisations are very fast, this would have been very interesting. Unfortunately when looking at the evolution of the design variables in Figures 7.15(e) to 7.15(g), this does not seem to happen: the optimisations on the three finest grids follow more or less the same path for the 6th and 17th variables but the evolution of the 2nd design variable contradicts that. The conclusion is hence that there is no obvious interest in applying optimisation directly to a coarser subgrid of the initial grid in the hope of following the same optimisation path at a reduced cost.


		C_L	C_D	Improvement in C_D
Initial		0.6520	0.01894	-
Optimised	35×9	0.6520	0.01609	-15.0%
	69×18	0.6520	0.01418	-25.1%
	137×35	0.6520	0.01377	-27.3%
	273×69	0.6520	0.01402	-26.0%

Table 7.3: Aerodynamic coefficients of the optimised aerofoils on the 137×35 grid for the optimisations comparing the influence of the grid size, at the target C_L .


Since the four optimisations reach different optimal points, the resulting aerofoil shapes are very different on both surfaces as can be seen in Figure 7.16. Note that the optimal design variables of each optimisation have been applied to the reference 137×35 grid to obtain these shapes. The pressure distributions calculated on these optimised aerofoils are also different. They have been obtained on a 137×35 grid with the incidence found by each optimisation hence they should not correspond exactly to the target C_L of 0.652. The peculiar shapes of the leading edge of the aerofoils obtained with the optimisations on the two coarsest grids, result in strange pressure behaviours in that region. These two shapes also show some remains of a shock wave on the upper surface while the optimisations on the two fine grids have completely eliminated it.

These optimisations take a new dimension when their final design changes are applied to the reference grid and the performance of the resulting aerofoils tested on this grid as shown in Table 7.3. The results presented so far for these optimisations showed that these performed on coarse grids were very quick and successful but it is difficult to relate them to the reference optimisation on a medium grid in a way that could make this latter faster. Checking the design changes on the reference grid partially enables this. Table 7.3 shows that the coarse grid optimisations make substantial improvements, not far from what the fine grid optimisations achieve. Hence instead of performing a costly fine grid optimisation, one might be tempted to compromise i.e. to accept that the design obtained may not be the true optimum (but will not be far from it), as long as this design is obtained very quickly.

This concludes the subsection that investigated the influence of the grid size on the overall optimisation. No obvious ways of reducing the computing time of fine grid optimisations was found but performing optimisations on coarse grids is still interesting: it was shown in this case that they substantially improve the design when checked on a finer grid, at a reduced cost. They can hence be an alternative to expensive fine grid optimisations.



(a) Aerofoil shapes



(b) Chordwise pressure distributions

Figure 7.16: Result of the aerofoil optimisations comparing the influence of the grid size.

		C_L	C_D	Improvement in C_D
Initial		0.6520	0.01894	-
Optimised	laminar	0.6520	0.01407	-25.7%
	turbulent with $\mu_t = \text{const}$	0.6520	0.01402	-26.0%
	fully turbulent	0.6520	0.01377	-27.3%


Table 7.4: Aerodynamic coefficients of the optimised aerofoils for the optimisations comparing the influence of the physical model of the adjoint solver, at the target C_L .

7.2.3.3 Influence of the physical model of the adjoint solver


The idea of Section 7.1.3 to look into using an adjoint solver that is not fully turbulent is reconsidered here but this time at the global level of the whole optimisation and not just a single gradient. The problem still involves Navier-Stokes flows so the flow solver is turbulent and the grid is still the reference 137×35 coarse Navier-Stokes grid. The only changes are the physical model employed in the adjoint solver. Four cases are considered: an inviscid adjoint, a viscous laminar adjoint, a turbulent adjoint that discards the linearisation of the turbulent viscosity that we call turbulent with $\mu_t = \text{constant}$ and the reference case with a fully turbulent adjoint. The evolution of the parameters during these optimisations is shown in Figure 7.17.

In Figure 7.17(a), the history of the drag coefficient is displayed. It clearly shows that the optimisation with an inviscid adjoint failed (it did not converge) and this is confirmed by the rest of the graphs. There are two reasons for this: first, the adjoint solver never manages to converge properly, being in an inviscid mode while the flow solution and the grid are turbulent. Hence the accuracy of the gradient that should already be bad only by changing the physical model, is worsened by the poor convergence. Secondly the inconsistency of the gradient causes a lot of trouble to the optimiser that requires a high number of flow and adjoint solutions per iteration. The combination of these two problems results in a huge computing time as seen in Figure 7.17(a) with no success in the optimisation. Since it is clear that this optimisation does not bring any improvement, its results are not shown entirely on the remaining of the graphs of Figure 7.17 and instead they zoom in on the other optimisations by reducing the range of the CPU time axis. For more clarity the evolution of C_D is repeated in Figure 7.17(b).


The optimisations with a laminar adjoint and a turbulent adjoint with $\mu_t = \text{constant}$ could be interesting when looking at C_D and the interior volume since they both reduce drag and satisfy the constraint on volume. However as can be seen in Figure 7.17(c), they do not satisfy the constraint on lift, which reduces their interest. As is confirmed by Table 7.4, their improvement in drag is altered when the target lift coefficient is matched. They are however a bit faster than the fully turbulent optimisation as expected because each adjoint is quicker to calculate. The inconsistency in the gradient due to the use of




(a) Objective function C_D




(b) Zoom on the objective function C_D discarding the results with the inviscid adjoint




(c) Lift coefficient C_L



(d) Interior volume



(e) 2nd design variable



(f) 6th design variable

Figure 7.17: Evolution of different parameters during the aerofoil optimisations comparing the influence of the physical model of the adjoint solver.


(g) 17th design variable: increment in angle of incidence


Figure 7.17: Evolution of different parameters during the aerofoil optimisations comparing the influence of the physical model of the adjoint solver. (Concluded)

a different physical model might be causing some trouble to the optimiser that spends a lot of time per iteration in the case of the laminar adjoint optimisation. The combination of these factors make these optimisations with a simplified physical model in the adjoint solver, no better than the Navier-Stokes optimisation.


The path followed by these optimisations in the design space shown in Figures 7.17(e) to 7.17(g), is interesting however. They are indeed close to what the consistent optimisation is doing at least at the beginning of the optimisation because they then seem to have been abruptly terminated before reaching the same optimum as the fully turbulent optimisation. This puts the optimisation with $\mu_t = \text{constant}$ ahead of the consistent optimisation during most of the process but without a priori knowledge of the behaviour of the fully turbulent optimisation, it is impossible to know when it would be needed to switch from a turbulent adjoint with $\mu_t = \text{constant}$ to a fully turbulent adjoint to carry on the optimisation. Little can be gained from the optimisation with a laminar adjoint since it is more or less in phase with the consistent optimisation as can be seen in Figure 7.17(g) for the increment in angle of incidence.

Since the optimisations stop at different design points as can be seen from the design variables, the resulting shapes shown in Figure 7.18 are different although very similar on the upper surface. The pressure distributions calculated on these shapes depict some similarity: the shock wave has been eliminated in all the cases. However these pressure distributions are obtained at the angles of incidence determined by the optimisations, hence they do not match the required lift coefficient of 0.652 as Figure 7.17(c) already showed. Thus a weak shock wave might reappear if they are all put at the correct incidence.

This is the end of this subsection that investigated the use of different physical models for



(a) Aerofoil shapes



(b) Chordwise pressure distributions

Figure 7.18: Result of the aerofoil optimisations comparing the influence of the physical model of the adjoint solver.

the adjoint solver while the flow solution is still fully turbulent. Once again no obvious way of reducing the computing time results from this investigation. The behaviour of the optimisations with a simplified adjoint could possibly be exploited however at the beginning of the optimisation to accelerate the process but a change from a simplified adjoint to a more consistent adjoint would be needed at some point and finding this moment might not be a trivial task and would require a lot more investigation. This idea is not developed further in this work. Instead a comparison is made between two consistent optimisations using either an Euler mode or a Navier-Stokes mode.

7.2.3.4 Comparison of Euler vs Navier-Stokes optimisation

Using an inviscid adjoint with a turbulent flow solution did not prove successful as seen in the previous subsection. A test was made with both a flow solution and an adjoint solution in Euler mode on the reference Navier-Stokes grid but this did not work better. However by changing the grid to an Euler grid i.e. by eliminating most of the grid stretching in the normal direction in the neighbourhood of the surface, the optimisation in inviscid mode worked well. This is what is presented in this subsection and the comparison is made with the fully turbulent reference optimisation. Both optimisations are consistent in the sense that the flow and adjoint solvers use the same physical model and the CFD grid is adapted to this model. To enable a fair comparison between the models, the CFD grid has the same number of points in both cases i.e. 137×35 . The results of these two optimisations are compared in Figure 7.19.

From Figure 7.19(a) to 7.19(c), it is clear that the Euler optimisation is successful and very fast: the decrease in drag coefficient is substantial and the constraints on lift and interior volume are well satisfied. It also appears that after 4 hours of calculation most of the optimisation is done, the rest of the time being spent on marginal improvements. Note that the optimisation stopped because it had reached the maximum number of optimisation cycles set to 50 for all the optimisations presented in this part, and not because an optimum was found. It is also clear from these results that each Euler optimisation cycle is much faster than a Navier-Stokes cycle. The use of the Euler mode changes also very much the magnitude of drag that is only composed of pressure drag in this case, and the value of lift, which makes the initial point feasible.

The evolution of the design variables is also interesting: after 4 hours of computing time, little seems to happen for the 2nd design variable in Figure 7.19(d) while the 6th and 17th design variables are still evolving until the end. The other important aspect is that the optimisations are following very different paths in the design space, ending up in well distinct points. The values of the 6th design variable for the final points do not even have the same sign.

The final shapes are thus different and can be seen in Figure 7.20. The pressure distributions calculated on them show that in both cases the initial strong shock wave has been

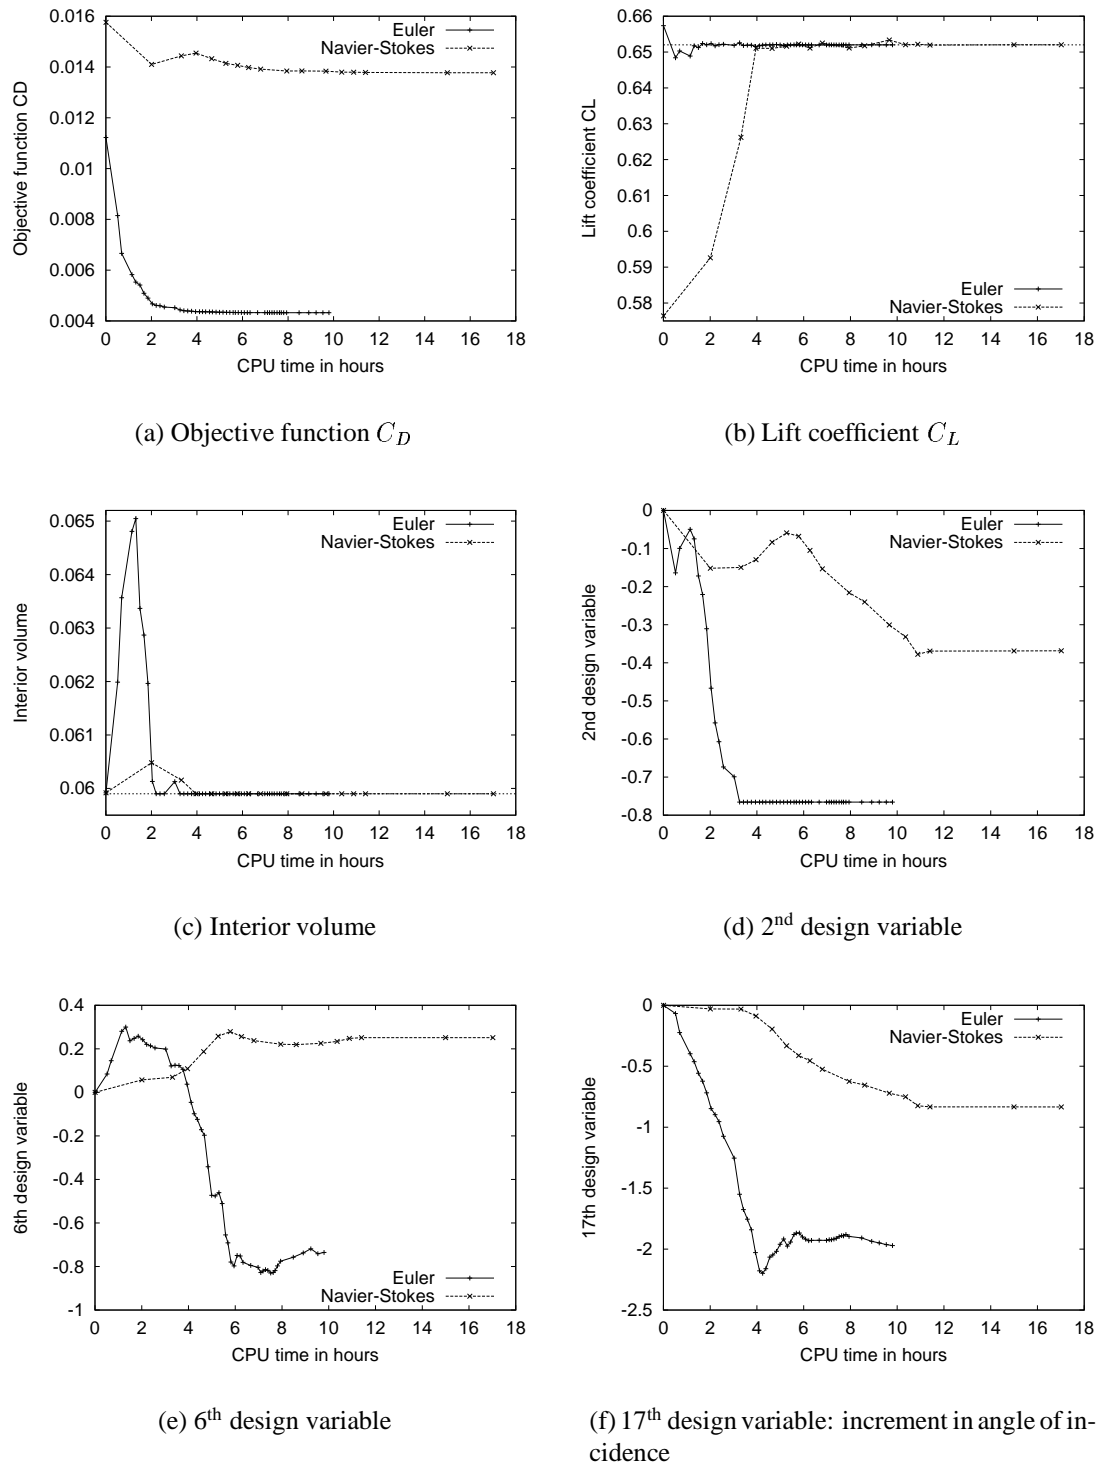




Figure 7.19: Evolution of different parameters during the aerofoil optimisations comparing an Euler to a Navier-Stokes optimisation.



(a) Aerofoil shapes



(b) Chordwise pressure distributions

Figure 7.20: Result of the aerofoil optimisations comparing an Euler to a Navier-Stokes optimisation.

		C_L	C_D	Improvement in C_D
Initial		0.6520	0.01894	-
Optimised	Euler	0.6519	0.01406	-25.8%
	Navier-Stokes	0.6520	0.01377	-27.3%

Table 7.5: Aerodynamic coefficients of the Euler and Navier-Stokes optimised aerofoils calculated from a Navier-Stokes solution on the reference grid at the target C_L .

eliminated. The Euler optimised aerofoil is more rear-loaded than the Navier-Stokes one. Note that the pressure distributions referring to the Euler optimisation are taken from an inviscid solution that is why two initial pressure distributions are plotted since otherwise the initial shape and target C_L are the same.

It is interesting to see what the performance of the Euler optimised aerofoil would be in the Navier-Stokes mode. This is what Table 7.5 presents. The Euler optimisation does remarkably well since it already manages to reduce the drag by almost 26% from the baseline geometry. Since this is obtained significantly faster than the Navier-Stokes optimum, there is some scope for reducing computing time, although as ever not in a straightforward way.

The conclusion of this subsection is that a consistent Euler optimisation is successful on the same reference problem. Its main advantage is that it is quick. The drawback however is that it does not share any behaviour with a consistent turbulent optimisation hence there is no easy way of combining the two, as starting with Euler and later on switching to Navier-Stokes, to provide an efficient optimisation. This test has nevertheless two positive consequences: it shows that an Euler optimisation is relevant in order to optimise a Navier-Stokes problem although the optimum might not perform as well as a Navier-Stokes optimum; and it paves the way for the use of the variable-fidelity method that will combine a low-fidelity Euler optimisation and a high-fidelity Navier-Stokes optimisation. This will be described in the last part of this chapter.

In the meantime, this paragraph also concludes this section presenting some basic testing on the two-dimensional optimisation test case employed in this study, done with the aim of reducing computing time. A slight reduction in the convergence requirements of the adjoint solver that should not change very much the result of the optimisation, seems to be the easiest way of slightly reducing computing time. The other investigations - using coarser grids, changing the physical model of the adjoint solver alone or of the entire optimisation - do not provide immediate ways of reducing computing time. However most of them lead to successful optimisations that manage to improve significantly the aerofoil performance when checked by a Navier-Stokes solution on a Navier-Stokes grid. Some of these optimisations even managed to achieve this quicker than the Navier-Stokes optimisation and are thus of interest. Some of these features will be used in the variable-fidelity

method described at the end of this chapter and others for the optimisations of a Blended Wing-Body in the next chapter. In the next section, different ways of handling the constraint on the lift coefficient are investigated.

7.2.4 Handling the constraints




In all the two-dimensional optimisation results presented so far, the optimiser is the NAG routine E04UCF. Since it is a standard SQP optimiser, it only satisfies the constraints when reaching an optimum and accepts infeasible points. The first part of this section looks at the use of the feasible optimiser FFSQP for this same two-dimensional problem. So far as well, the constraint on lift has been treated as a hard (separate) constraint. This means that the gradient of the lift coefficient is needed by the optimiser in addition to the gradient of the objective function. It works well but it is time consuming since it requires the solution of two adjoint equations for each optimisation cycle. Another possible way is to incorporate the constraint on lift into the objective function through the use of a penalty term. This would only require one adjoint solution per optimisation cycle. This is investigated in the second part of this section.

7.2.4.1 FFSQP vs E04UCF

As just mentioned, this subsection compares the use of the optimisation routine FFSQP to the routine E04UCF for the two-dimensional test case of this chapter. They have both been presented in detail in Section 3.3. The main difference between the two is that FFSQP is a feasible method that requires all the constraints to be satisfied at each iteration while E04UCF is a standard method in which the constraints are only satisfied close to the optimum. This difference is obvious when examining the results of the optimisations in Figure 7.21. The optimisation problem is the scaled problem (7.3).

The behaviour of the feasible method is clear in Figure 7.21(b): starting from an unfeasible point for the constraint on lift, the first thing that the optimiser does is to look for a feasible point that is far away from the constraint boundary and then gradually comes back to this boundary from the interior of the feasible region. It works but it is dramatically slow when compared to the standard SQP method. The requirement of feasibility pushes the design away from the optimum, generating a very thick aerofoil (Figure 7.21(c)) that creates a lot of drag (Figure 7.21(a)). Starting from this very bad but feasible point, the optimiser then slowly improves the design. Even if the method is not efficient, it is successful and the optimum found is very close to the one found by the standard SQP method.

This is confirmed by Table 7.6 and is emphasised by the path followed by the optimisation in the design space. When considering Figures 7.21(d) to 7.21(f), it is clear that the two optimisations have globally the same behaviour and stop at optimal design points very close to each other. This is reassuring and proves that the optimisation method taken as a reference so far is correct.

(a) Objective function C_D (b) Lift coefficient C_L 

(c) Interior volume




(d) 2nd design variable(e) 6th design variable(f) 17th design variable: increment in angle of incidence

Figure 7.21: Evolution of different parameters during the aerofoil optimisations comparing the optimisation routine FFSQP (feasible SQP) to the NAG routine E04UCF (standard SQP).

		C_L	C_D	Improvement in C_D
Initial		0.6520	0.01894	-
Optimised	FFSQP	0.6520	0.01377	-27.3%
	E04UCF	0.6520	0.01377	-27.3%

Table 7.6: Aerodynamic coefficients of the optimised aerofoils for the optimisations comparing the optimisation routine FFSQP (feasible SQP) to the NAG routine E04UCF (standard SQP), at the target C_L .

Since the optimal points are very close, their corresponding shapes, shown in Figure 7.22, are very similar, the only important difference being in the leading edge region on the lower surface. Thus the pressure distributions are also very close.


The conclusion to this subsection is that the feasible optimisation method is successful on this two-dimensional test case but it is much slower than the standard method employed so far. This is due to the feasibility requirement itself that creates an initial design point well inside the feasible region, too far from its boundaries. As a consequence, the routine FFSQP will no longer be used in this work in the case of a direct optimisation because it is too costly. The other conclusion to draw from this subsection is that the reference optimisation considered so far finds the correct optimum since the feasible method stops on the same design point.

7.2.4.2 Penalty function vs hard constraints


This subsection investigates the use of a penalty term added to the objective function for the constraint on lift rather than employing a separate constraint. This could reduce the total computing time since it eliminates the calculation of the adjoint for lift needed to calculate the gradient of lift. Only the constraint on C_L is included in the penalty term and the geometrical constraint on volume is still kept as a separate constraint. Since the interior volume and its gradient are calculated analytically in a fraction of a second this does not pose any problem. The optimiser employed in this subsection is the NAG routine E04UCF.

The optimisation problem to be solved becomes:

$$\begin{aligned} \text{Minimise } F &= \frac{C_D}{\left\| \frac{dC_D}{d\beta^0} \right\|} = 0.02292 + r_p \frac{(C_L - 0.652)^2}{\left\| \frac{dC_L}{d\beta^0} \right\|^2} = 0.4754^2 \\ \text{Subject to: } \frac{V_0}{\left\| \frac{dV_0}{d\beta^0} \right\|} &\leq \frac{V}{\left\| \frac{dV}{d\beta^0} \right\|} \leq \frac{2 \times V_0}{\left\| \frac{dV_0}{d\beta^0} \right\|} = 0.02872 \end{aligned}$$



(a) Aerofoil shapes



(b) Chordwise pressure distributions


Figure 7.22: Result of the aerofoil optimisations comparing the optimisation routine FFSQP (feasible SQP) to the NAG routine E04UCF (standard SQP).

where r_p is a penalty coefficient to be chosen. Note that the functions are still scaled by the norm of their initial gradient as suggested in Reference [112].


The choice of the penalty coefficient r_p poses problem. Ideally and theoretically, a sequential unconstrained minimisation^[112,233] should be performed: starting with a low value of the penalty coefficient, a first optimisation is carried out, then the penalty coefficient is increased and the optimisation is restarted from the previous optimum point and so on until the value of r_p becomes very high. This is necessary because putting too much weight on the penalty term produces an ill-conditioned objective function that could be difficult to minimise directly. By progressively increasing the penalty coefficient, the design is getting closer to the optimum and satisfies the constraints better and better. However this means repeating a sequence of optimisations and thus becomes very costly. In aerodynamic optimisation, to the knowledge of this author, sequential unconstrained minimisation has not been applied and people^[9,38,47,62,125,145–147,164,190,234] that use a penalty function, do a single optimisation with a constant penalty coefficient. Usually the result is satisfactory and the constraints are satisfied. However none of them really explains the choice of the value of the penalty coefficient. Here, to test this choice of penalty coefficient, we propose to test three values of r_p namely 1, 10 and 100. Each time a single optimisation is performed starting from the same initial design as all the two-dimensional optimisations of this chapter. It could have been interesting to perform a sequential unconstrained minimisation with these 3 values of r_p by starting the optimisation with $r_p = 10$ from the final design of the optimisation with $r_p = 1$ and so on. This would certainly have improved the conditioning of each optimisation but was not performed because as explained above, it is unlikely to reduce overall computing time, which is the aim to these tests. The results of these penalty function optimisations are presented in Figure 7.23.

Figure 7.23(a) provides information that was not available for separate constraints, i.e. the value of the objective function F . This is given as a ratio F/F_0 since the three objective functions are of very different magnitude due to the changes in the value of r_p . A quick look at Figure 7.23 shows that the optimisation with $r_p = 100$ fails. This is certainly due to the ill-conditioning mentioned above for large values of the penalty coefficient. The other two penalty optimisations try to find a minimum but are not entirely successful. They indeed fail to satisfy the constraint on lift as can be seen in Figure 7.23(c), thus the decrease in drag found in Figure 7.23(b) is altered when calculated at the target lift coefficient as can be seen in Table 7.7. As expected though, the optimisation with $r_p = 10$ puts a greater weight on the constraint on lift so it better satisfies this constraint than the optimisation with $r_p=1$ and thus manages to find a better optimum.


Even if not fully successful, these two optimisations follow a path in the design space similar to the optimisation with separate constraints as Figures 7.23(e) to 7.23(g) show. It is interesting to notice that the optimisation with $r_p = 1$ is quicker than the reference optimisation while the one with $r_p = 10$ is twice as slow as this latter. This can only be




(a) Ratio of objective function over initial objective function F/F_0 , both including a penalty term




(b) Drag coefficient C_D




(c) Lift coefficient C_L



(d) Interior volume



(e) 2nd design variable



(f) 6th design variable

Figure 7.23: Evolution of different parameters during the aerofoil optimisations using a penalty term approach with scaling.


(g) 17th design variable: increment in angle of incidence

Figure 7.23: Evolution of different parameters during the aerofoil optimisations using a penalty term approach with scaling. (Concluded)

		C_L	C_D	Improvement in C_D
Initial		0.6520	0.01894	-
Optimised	penalty coefficient=1	0.6520	0.01393	-26.5%
	penalty coefficient=10	0.6521	0.01385	-26.9%
	separate constraint	0.6520	0.01377	-27.3%

Table 7.7: Aerodynamic coefficients of the optimised aerofoils for the optimisations using a penalty term approach with scaling, at the target C_L .

		C_L	C_D	Improvement in C_D
Initial		0.6520	0.01894	-
Optimised	penalty coefficient=1	0.6519	0.01385	-26.9%
	penalty coefficient=10	0.6520	0.01386	-26.8%
	penalty coefficient=100	0.6520	0.01442	-23.9%
	separate constraint	0.6520	0.01377	-27.3%

Table 7.8: Aerodynamic coefficients of the optimised aerofoils for the optimisations using a penalty term approach without scaling, at the target C_L .

due to the poor conditioning of the optimisation problem, since mathematically the two penalty optimisations are the same. Note that it is worrying that an optimisation method that does not need the evaluation of the adjoint for lift, is much slower than one that needs this calculation. Despite not satisfying the constraint on lift, the good performance found in Table 7.7 and the fact that it is much quicker than the reference optimisation, make the optimisation with $r_p = 1$ an interesting option that will not be used any further in this work though.


An alternative was investigated with the suppression of the scaling in the objective function. The optimisation problem becomes:

$$\begin{aligned} \text{Minimise} \quad & F = C_D + r_p (C_L - 0.652)^2 \\ \text{Subject to:} \quad & \frac{V_0}{\left\| \frac{dV_0}{d\beta^0} \right\|} \leq \frac{V}{\left\| \frac{dV}{d\beta^0} \right\|} \leq \frac{2 \times V_0}{\left\| \frac{dV_0}{d\beta^0} \right\|} = 0.02872 \end{aligned}$$


The scaling for the constraint on interior volume is kept though because the previous results do not show any significant problem for this constraint. Again the three values 1, 10 and 100 are tested for the penalty coefficient r_p . The results are displayed in Figure 7.24.

Surprisingly they are much better than these obtained by using some scaling in the objective function. The three optimisations find an optimum. Only the optimisation with $r_p = 1$ is not entirely satisfactory since it does not satisfy accurately either of the constraints (see Figures 7.24(c) and 7.24(d)). Hence the very good value of drag (see Figure 7.24(b)) for its optimum is reduced when the lift coefficient is matched as Table 7.8 shows. The optimisation with $r_p = 100$ does not find a very good optimum, although it satisfies the constraints. Too much weight seems to be put on the constraint for lift in the objective function, as a result the optimiser does not decrease drag as much as it should. The optimisation with $r_p = 10$ is the most satisfactory overall since it satisfies well the two constraints and finds an optimum with a low drag.


When looking at the design variables in Figures 7.24(e) to 7.24(g), this latter optimisation seems to follow the same path as the reference optimisation although it is taking a lot of




(a) Ratio of objective function over initial objective function F/F_0 , both including a penalty term




(b) Drag coefficient C_D




(c) Lift coefficient C_L



(d) Interior volume



(e) 2nd design variable



(f) 6th design variable

Figure 7.24: Evolution of different parameters during the aerofoil optimisations using a penalty term approach without scaling.


(g) 17th design variable: increment in angle of incidence


Figure 7.24: Evolution of different parameters during the aerofoil optimisations using a penalty term approach without scaling. (Concluded)

smaller steps to get to the optimum. The values of the design variables at the optimum are not far from the reference values. The method with $r_p = 1$ is also very close to this optimum despite not satisfying the constraints very well and follows also more or less the same path. This explains its good performance in Table 7.8. The optimisation with a penalty coefficient of 100 stops at a very different point.


This is noticeable on the resulting aerofoil shapes shown in Figure 7.25: the one corresponding to $r_p = 100$ is situated between the initial shape and the optimal shapes of the other three optimisations as could be guessed from the evolution of the design variables. The shapes resulting from the other two penalty optimisations are very close to each other and to the reference shape as expected from the design variables. This means that the non-satisfaction of the constraint on interior volume by the optimisation with $r_p = 1$ is not as significant as Figure 7.24(d) would suggest. The pressure distributions obtained on the optimal shapes lead to the same conclusion: the one corresponding to $r_p = 100$ still has some features of the initial pressure distribution although the shock wave has disappeared and the three other distributions are very close to each other.

Although the results of these penalty optimisations are satisfactory, they become less interesting when looking at computing time. Indeed none of these optimisations is quicker than the reference one with separate constraint (if we consider that it reaches its optimum after 11.5 hours), although their computing cost should be reduced by not calculating the adjoint for lift. This shows that adding the penalty term to the objective function changes the nature of this function. It is now more difficult to find its minimum. Hence the expected gains in computing time vanish and the penalty approach loses its attraction.

To complete the investigation, tests without scaling in the penalty function as well as in



(a) Aerofoil shapes



(b) Chordwise pressure distributions

Figure 7.25: Result of the aerofoil optimisations using a penalty term approach without scaling.

the geometrical constraint were carried out. Their results are not shown here because they were not as good as the previous series of tests with scaling in the constraint on interior volume. This confirmed our conclusion that the penalty approach has some limitations and is not better than a direct approach with separate constraints. That is why a direct optimisation with a penalty objective function will not be used any more in the remaining of this thesis.

This is the end of this section that investigated different ways of handling the constraints for the two-dimensional test case of this chapter. It showed that the best approach is using a standard SQP method with separate constraints and that a feasible method or a penalty approach are not as efficient. This also terminates the second part of this chapter that performed some direct optimisations using an SQP optimiser on a two-dimensional test case. The next section looks, as an alternative, at the use of the variable-fidelity method for this same problem.

7.3 Two-dimensional optimisation using the variable-fidelity method

The aim of the previous part of this chapter was to gain some understanding of how the optimisation chain developed in this work, behaves. It was tested on a two-dimensional problem and it was found that although the basic method works well, it is very time consuming. Hence some testing was carried out to try and reduce the computing time without compromising the result of the optimisation, with the idea that the same method will then be applied to three-dimensional problems for which computing time is critical. Unfortunately these simple tests did not provide an easy way to do this so other optimisation methods were considered. This is what is described in this last part of the chapter.

The first section quickly surveys the literature to find innovative optimisation methods, still gradient-based but that are much quicker than a standard direct optimisation. Among them is the variable-fidelity method already introduced in Section 3.4 and employed in this work. This approach is applied to the two-dimensional optimisation test case and its results are presented afterwards.

7.3.1 Alternative methods

A short presentation of these methods can be found in the introduction of Reference [235]. We follow the same order and add a few references. All the methods presented here note the fact that the traditional aerodynamic optimisation method also called sequential optimisation or loosely-coupled optimisation, that consists in associating a flow solver and an adjoint solver to calculate the gradient to a numerical optimiser are too costly since they solve at each iteration two large systems of equations (one for the flow solution, one for the

adjoint solution) before updating the design. The aim of these methods is to reduce significantly the computing time to end up with an overall optimisation process that is equivalent to the cost of a few flow solutions, “a few” meaning here between 1 and 10. They manage this by considering the optimisation no longer as a sequence of flow solution - adjoint solution - design update but as an integrated process. The idea is similar to what has already been presented for the derivation of the continuous adjoint equations in Section 6.2 and is that the optimum design point must satisfy 3 equations:

- the flow equation
- the adjoint equation
- the optimality equation which states that, at the optimum, the gradient of the objective function is zero for an unconstrained problem

The traditional optimisation solves the first two equations at each optimisation cycle in order to advance in the resolution of the third equation but this is a waste of time since the first two equations only need to be satisfied at the optimum and not for other design points. Hence the alternative methods propose to solve the three equations simultaneously to reach the optimum when all three are satisfied.

This is the main idea of the One-shot method.^[151,236] It is achieved by an extensive use of a multigrid acceleration technique applied to the three equations simultaneously. The design variables are classified depending on whether they bring smooth changes to the geometry and the flow field or whether they imply high-frequency localised changes. The latter are only used on fine grids while global smooth design variables are employed throughout the grid levels. The use of a multigrid technique requires the implementation of relaxation and interpolation schemes for the three equations to pass from one grid level to another. This is applied but only for potential flows in Reference [151] at the cost of less than one flow solution on the fine grid.

Sung and Kwon^[145,147] use a similar idea but do not apply multigrid to the optimality equation. They successively perform a few multigrid cycles on the flow equations, followed by some cycles on the adjoint equations before updating the design and performing again some multigrid cycles on the flow equations and so on. That is what they call a tightly-coupled algorithm. Reference [35] shows that it is possible to employ the One-shot method without any adjoint solver by computing the gradient by finite-difference. This is of course less efficient than using an adjoint method. The drawback of the One-shot method is that it is mathematically complicated and needs a flow and an adjoint solver with a multigrid capability which is not always available.

That is why other techniques have been explored that could be simpler and applicable with any flow and adjoint solvers. The pseudo-time method^[237] is one of them. As already explained the method starts by noting that the optimum lies at the intersection of three hypersurfaces: the state (flow solution), the costate (adjoint solution) and the design

(optimality) hypersurfaces. Traditional optimisation methods work at the intersection of the state and costate hypersurfaces and move along this intersection to get closer to the design hypersurface. On the contrary the pseudo-time method works on the design hypersurface and gets closer to the other two hypersurfaces. In practice, the flow and adjoint equations are advanced by a few iterations on the entire field. Then both are solved in a small neighbourhood around the boundary that is undergoing design changes and the design is updated from these solutions. Then the process is repeated until convergence of the flow and adjoint equations on the entire field. This method is applied to the inverse design of a nozzle in Reference [222] and the inverse design of an aerofoil in Reference [224].

Finally a method that combines the One-shot and the pseudo-time approaches is the progressive optimisation method.^[235] This employs two flow solvers: an accurate flow solver for the flow equations and an auxiliary dissipative flow solver only used to calculate smooth sensitivities with an adjoint solver. The aim is to avoid noise as much as possible in the calculation of the gradient. Like the pseudo-time method, it uses partially converged flow and adjoint solutions and updates frequently the design. Like the One-shot method, it employs different levels of grid refinement. It indeed starts with a very coarse grid, solves partially the flow and adjoint equations, updates the design and does again a few iterations for the flow and adjoint solutions and so on until the process is converged on this coarse grid. Then this is repeated on a finer level of grid and so on until the equations are converged on the finest grid level. In Reference [235] this is applied to inviscid inverse designs and proves to be very efficient while in Reference [234] it is used for viscous turbulent inverse and direct design of aerofoils. In this latter case the adjoint is inviscid and always solved on the coarsest grid.

The problem with these three methods is that they are demonstrated on inverse design problems and of course they succeed in matching the target pressure distribution. For a direct design problem, they will undoubtedly also find an optimum but nothing guarantees that they will find the same optimum as a traditional method directly applied to the finest grid. That is why the present author preferred to employ the variable-fidelity method already presented in Section 3.4 and which is the last of the alternative methods examined here. The variable-fidelity method indeed constantly checks its improvement with respect to the high-fidelity model. There is no guarantee either that the optimum found would be the same as with a traditional method but at least, the designer can follow the design changes and improvements during the optimisation and he is confident that these improvements are meaningful improvements for the high-fidelity problem he is interested in.

The variable-fidelity method is quite different from the other three methods just presented and does not rely on the simultaneous solution of the flow, adjoint and optimality equations. It is much closer to a traditional optimisation method although it is very efficient like the three other alternative methods. Its application to the two-dimensional optimisation test case of this chapter is examined next.

7.3.2 Variable-fidelity results

The comparison made in Section 7.2.3.4 between consistent Euler and Navier-Stokes optimisations is the basis for the application of the variable-fidelity method (see Section 3.4 for details about the method) to this two-dimensional problem. It indeed proved that an Euler optimisation worked well, managed to improve significantly the design when checked with a Navier-Stokes solution and thus could be used as a low-fidelity model. That is why the variable-fidelity approach examined here employs a consistent Euler optimisation for the low-fidelity model (flow and adjoint solvers both in Euler mode working on an Euler grid) and a consistent fully turbulent optimisation for the high-fidelity model (flow and adjoint solvers in viscous turbulent mode with the complete linearisation of the eddy viscosity for the adjoint solver, working on a Navier-Stokes grid).

However the Euler optimisation presented in Section 7.2.3.4 was quite slow to be repeated several times in a row as part of the variable-fidelity method. This came primarily from the grid size since the Euler grid had the same number of points as the Navier-Stokes grid. Since the low-fidelity model of the variable-fidelity approach is only supposed to approximate the high-fidelity model and since it is anyway corrected to behave like this latter, some crude approximations are allowed. A compromise has to be found between rapidity of execution and accuracy for the low-fidelity model. That is why a very coarse Euler grid of 65×17 nodes is employed here for the low-fidelity model. It is shown in Figure 7.26 and has to be compared to the Navier-Stokes 137×35 grid of Figure 7.9 used for the high-fidelity model. This takes advantage of the findings made in Section 7.2.3.2 where coarse grid optimisations were investigated. This showed that these optimisations are very fast but at the same time manage to improve the Navier-Stokes design. Of course combining this to an Euler optimisation to create a low-fidelity model will reduce the fidelity of this model.

The low-fidelity optimisation is the optimisation problem (7.3) transformed into a corrected low-fidelity problem according to equation (3.6). It is thus:

$$\begin{aligned}
 & \underset{\mathbf{s}}{\text{Minimise}} && \frac{\widetilde{C}_{D_{lo}}(\boldsymbol{\beta}^q + \mathbf{s})}{\left\| \frac{d\widetilde{C}_{D_{lo}}(\boldsymbol{\beta}^q)}{d\boldsymbol{\beta}^q} \right\|} \\
 & \text{Subject to:} && \frac{0.652 - C_{L_{hi}}(\boldsymbol{\beta}^q) - \nabla C_{L_{hi}}(\boldsymbol{\beta}^q)^t \cdot \mathbf{s}}{\left\| \frac{dC_{L_{hi}}(\boldsymbol{\beta}^q)}{d\boldsymbol{\beta}^q} \right\|} \leq 0 \\
 & && \frac{V_0 - V_{hi}(\boldsymbol{\beta}^q) - \nabla V_{hi}(\boldsymbol{\beta}^q)^t \cdot \mathbf{s}}{\left\| \frac{dV_{hi}(\boldsymbol{\beta}^q)}{d\boldsymbol{\beta}^q} \right\|} \leq 0 \tag{7.4}
 \end{aligned}$$

$$\frac{-2 \times V_0 + V_{hi}(\boldsymbol{\beta}^q) + \nabla V_{hi}(\boldsymbol{\beta}^q)^t \cdot \mathbf{s}}{\left\| \frac{dV_{hi}(\boldsymbol{\beta}^q)}{d\boldsymbol{\beta}^q} \right\|} \leq 0$$

$$\beta_k^l - \beta_k^q \leq s_k \leq \beta_k^u - \beta_k^q \quad k = 1, 17$$

$$\|\mathbf{s}\| \leq \Delta_q$$

The variable-fidelity method avoids one of the problem encountered in Section 7.2.3.2 when performing optimisation on coarse grids and in Section 7.2.3.4 when doing an Euler optimisation, namely that using a lower-fidelity model changes the magnitude of the aerodynamic coefficients and thus that it is not entirely relevant to apply the same constraints with these models as with the high-fidelity model. Indeed when looking at Figure 7.15(c) and Figure 7.19(b) for example, it is clear that imposing a constraint at 0.652 does not have the same meaning on the coarse grids as on the fine grid and in Euler mode as in Navier-Stokes mode. Their constraint is not adapted to the model. This could explain why the performance of the optimised aerofoils when checked at the target lift coefficient on the reference grid, are not as good as for the Navier-Stokes optimisation. Maybe the constraints should be scaled in some way. The variable-fidelity method avoids this since the low-fidelity model is corrected to match the value of the high-fidelity model, hence the constraints are always consistent.

The merit function Φ_{hi} that is used by the variable-fidelity method to assess the high-fidelity performance of new designs (see Section 3.4), is similar to what has been used for the direct optimisation with a penalty approach except that now, since it is the only number measuring the performance of the design, the constraint on volume is incorporated into it. The merit function is thus:

$$\Phi_{hi} = \frac{C_D}{\left\| \frac{dC_D}{d\boldsymbol{\beta}^0} \right\|} = 0.02292 + 10 \frac{(C_L - 0.652)^2}{\left\| \frac{dC_L}{d\boldsymbol{\beta}^0} \right\|^2} = 0.4754^2 + 10 \frac{(V - V_0)^2}{\left\| \frac{dV_0}{d\boldsymbol{\beta}^0} \right\|^2} = 0.02872^2$$

Contrary to the direct optimisations performed using a penalty approach in Section 7.2.4.2, this penalty function Φ_{hi} does not drive the optimisation since this latter is performed through the low-fidelity problem (7.4). The penalty function here is an a posteriori check of the design and hence the problems encountered in Section 7.2.4.2 should be avoided.

The optimiser employed is FFSQP as already explained in Section 3.4. Mild convergence requirements and a maximum number of 20 iterations are imposed for the low-fidelity optimisation because it does not need to be very accurate but has to be as fast as possible. The results of the application of the variable-fidelity method to the two-dimensional test-case are shown in Figure 7.27 and compared to the reference direct optimisation using E04UCF as optimiser.




Figure 7.26: Coarse Euler grid used for the low-fidelity model in the variable-fidelity method. (to be compared to Figure 7.9)

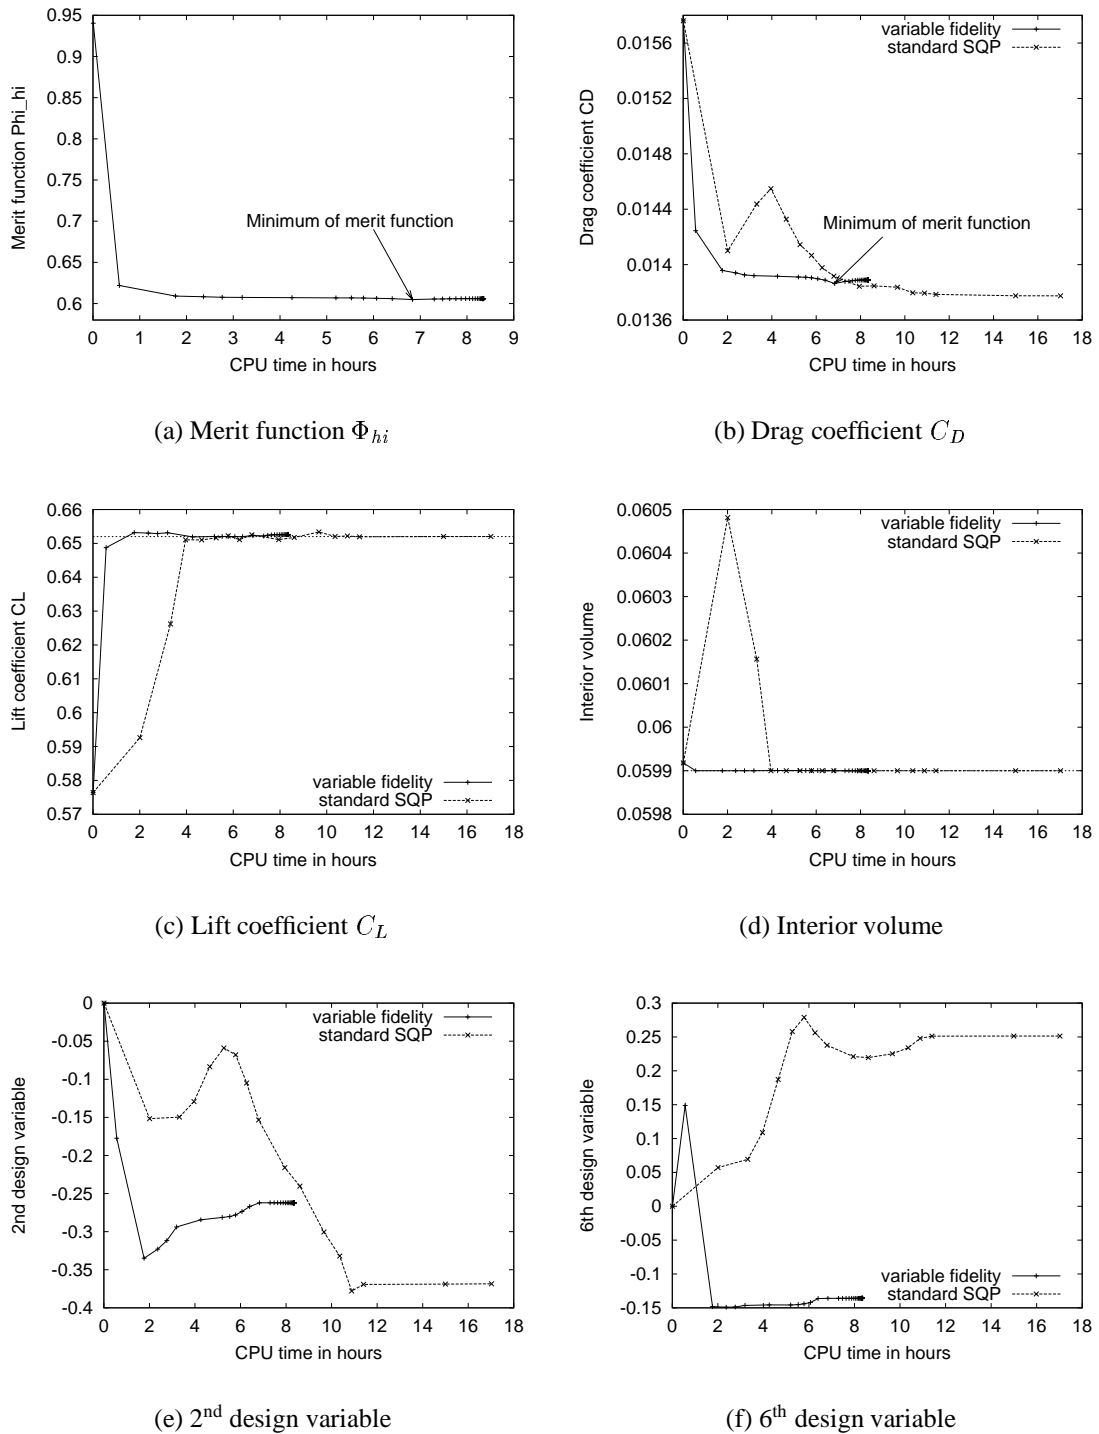


Figure 7.27: Evolution of different parameters during the aerofoil optimisations comparing the variable-fidelity method to the standard SQP method.


(g) 17th design variable: increment in angle of incidence

Figure 7.27: Evolution of different parameters during the aerofoil optimisations comparing the variable-fidelity method to the standard SQP method. (Concluded)

		C_L	C_D	Improvement in C_D
	Initial	0.6520	0.01894	-
Optimised	variable-fidelity	0.6520	0.01386	-26.8%
	standard SQP	0.6520	0.01377	-27.3%

Table 7.9: Aerodynamic coefficients of the optimised aerofoils for the optimisations comparing the variable-fidelity method to the standard SQP method, at the target C_L .


Figure 7.27(a) plots the evolution of the merit function Φ_{hi} . Unfortunately the initial value is very high due to the non-satisfaction of the constraint on lift, so the decrease of the merit function is not obvious after the initial drop. It does go down though until 7 hours of computation after which no improvement is made and the method stops when reaching the maximum number of cycles set to 50. From Figures 7.27(b) to 7.27(d), it is clear that the variable-fidelity method is successful in finding a valuable optimum. The constraint on lift is satisfied almost immediately and the interior volume is directly set to the target value. This is noticeable in Figure 7.27(b) since the drag coefficient falls quickly to very interesting values. The evolution after this drop is much slower. Note that the minimum in drag corresponds to the minimum of the merit function, which proves that the choice of this function is correct and that the method is working. The only problem is that the variable-fidelity method does not manage to find the same optimum as a standard SQP optimisation that stops on a better design, as Table 7.9 confirms.

This is clearly noticeable when looking at the design variables for the variable-fidelity method in Figures 7.27(e) to 7.27(g): they are following a very different path from the standard SQP method. Obviously this has an impact on the final shapes as shown in


Figure 7.28: the difference is particularly pronounced on the lower surface. Of course the pressure distributions are also different but the variable-fidelity method manages well to eliminate the shock wave. It is though keeping a small pressure bump on the upper surface exactly like the standard SQP approach.

The fact that a direct optimisation approach finds a better optimum than the variable-fidelity method although at a greater cost, is also encountered by Marduel *et al.*^[133] This slightly mars the otherwise excellent performance of the method: its handling of the constraints is indeed the best that has been seen in this chapter and the decrease in drag is very good, especially over the first two cycles. It unfortunately then slows down which makes the standard method almost as efficient to reach the same minimum. As a result the drop in computing time is not as great as expected. A reduction of computing time by a factor of 2 can be obtained over the standard SQP method if the variable-fidelity method is stopped after 5 \sim 6 iterations and would confirm the findings of Alexandrov,^[130–132] but not if the variable-fidelity method is left to find its own minimum, at least for this two-dimensional problem. The method is however very promising for three-dimensional optimisation if only for its very good handling of the constraints.

This concludes the presentation of the results obtained by the variable-fidelity method on the two-dimensional test case of this chapter. It also terminates this chapter that looked first of all at the accuracy of the gradient and how it was influenced by the convergence levels, the grid size or the physical model of the adjoint solver. This showed that a good convergence of both the flow and adjoint solvers is needed to obtain an accurate gradient and that computing the gradient on a subgrid or using a simpler physical model greatly compromise its accuracy. Then some results of aerofoil optimisations using a traditional approach with an SQP optimiser were presented. They showed that the optimisation chain set up in this study is working but also that it is quite costly, even for a two-dimensional problem. That is why some tests were carried out to find a simple way of reducing the cost of optimisation. Although some interesting points that will be used in the next chapter resulted from these investigations (a slight reduction of the convergence level of the adjoint solver is possible, coarse grid as well as Euler optimisations bring substantial improvements to the design), they did not provide an easy way of reducing the computing time of the Navier-Stokes optimisation. Hence the author had to investigate alternative methods and chose to concentrate on the variable-fidelity method already presented in Chapter 3. This latter approach showed promising results with a good handling of the constraints and a slight reduction of computing time. This was for a two-dimensional Navier-Stokes optimisation, the next chapter presents some results of the application of the optimisation chain to three-dimensional problems, in particular the optimisation of a Blended Wing-Body.



(a) Aerofoil shapes



(b) Chordwise pressure distributions

Figure 7.28: Result of the aerofoil optimisations comparing the variable-fidelity method to the standard SQP method.

This page has been left intentionally blank.

Chapter 8

Three-dimensional optimisations

The previous chapter described two-dimensional optimisations and the optimisation method worked well for this. However the goal of this work is to perform three-dimensional optimisations, particularly the optimisation of a Blended Wing-Body (BWB) aircraft. This is what is described in this chapter. It starts by explaining how the flow and adjoint solvers have been parallelised in order to tackle the computationally demanding three-dimensional optimisations. The first example of these optimisations is then presented and consists in the Navier-Stokes optimisation of the ONERA M6 wing. The optimisations of the BWB follow. Two categories of optimisations have been carried out on this geometry: Euler and Navier-Stokes optimisations, and within each of these, two optimisations are performed: with and without constraint on the pitching moment. The last part of this chapter discusses the performance of the optimisation method and possible ways to improve them. Likewise it assesses also the optimisations of the BWB.

8.1 Parallel computing

The previous chapter proved that reducing computing time for optimisation is not easy when modifying only the optimisation algorithm and the numerical models employed. The introduction of the variable-fidelity method brought some improvements but it is still expensive to run particularly for the three-dimensional problems of this chapter. Some further improvements are needed to be able to run three-dimensional optimisations in a reasonable time.

Parallel computing offers an easy way to reduce the wall-clock time spent on a calculation by distributing the work involved to several processors rather than on a single one for serial computing. It does not reduce the total computing time itself (it even slightly increases it) but it very much improves the computational time perceived by the user who has to wait to get an answer (wall-clock time). Two major approaches are possible to parallelise FORTRAN codes: Message Passing or Shared Memory parallelisation.

Message Passing uses libraries to distribute the data to the processors, to perform synchronisation between the processors and to manage the sending and receiving of information to and from different processors. All of this has to be programmed explicitly in the code. The two most popular libraries are PVM^[238] (Parallel Virtual Machine) and MPI^[239] (Message Passing Interface). The advantage of Message Passing is that it is portable and works on distributed memory machines such as clusters of PCs or workstations. However it involves a lot of programming to transform a serial code into a parallel code.

In contrast Shared Memory parallelisation can only be employed on shared memory machines which have several processors but a common memory accessible by all the processors. OpenMP^[240] is the language that is becoming a standard for Shared Memory parallelisation. It easily parallelises the loops of a serial code. Its main advantages are that the user does not have to worry about the communication between processors and that modifications to the serial code are minimal, the inconvenience being that it only works on shared memory machines with sufficient memory.

CFD is well suited for data parallelism also known as domain decomposition: instead of performing the computation on a large domain at once, the calculation is done on several smaller domains simultaneously. A multiblock CFD grid enables the computational domain to be easily divided into smaller blocks. The calculation on these blocks is done simultaneously on different processors at each iteration. For better efficiency, it is thus required to have a total number of blocks equal to a multiple of the number of processors available. Static load balancing is desirable to speed up the computation. It can be achieved by dividing the computational domain into blocks of exactly the same size so that the processors spend more or less the same time on each block.

Both the flow and adjoint solvers of this thesis have been parallelised using OpenMP. The aim was to obtain parallel codes as quickly as possible to allow for optimisation problems impossible on a single processor workstation. OpenMP was thus chosen for its simplicity of application as opposed to MPI and because shared memory machines were available to the author. Only two lines of codes had to be added to the serial codes to parallelise the main “DO loop” that is looping on the number of blocks.

Two machines were used for the parallel computations: an Alpha Unix server with 4 EV6.7 processors running at 667 MHz and 4 Gb of RAM, and an IBM SP parallel supercomputer composed of nodes of 16 Power3-II processors running at 375 MHz with 12 Gb of RAM per node. Both enable the use of OpenMP, though restricted to a single node of 16 processors for the IBM SP.

Only the flow and adjoint solvers were parallelised in the optimisation chain, the grid update program and the optimiser being kept as serial codes since their computing cost is very small. Due to constraints on the block topology of the CFD grid imposed by the use of the Baldwin-Lomax turbulence model and by the grid deformation code, perfect load

balancing could not be achieved and the parallel efficiency is not very high. A test was carried out with the adjoint solver and the fine 16-block Navier-Stokes grid employed in the optimisations of the BWB: on the 4 processors of the Alpha machine the speed up (=wall clock time of parallel code on 4 processors/wall clock time of serial code on single processor) of 2.6 giving a parallel efficiency (=speed up/number of processors) of no more than 66%; on one node of 16 processors of the IBM SP, the speed up is only 9.4 for a parallel efficiency of 59%.


Nevertheless the parallelisation of the flow and adjoint solvers achieved its goal i.e. to make it possible to tackle three-dimensional optimisations without suffering too much from huge computing times. A first example is given by the Navier-Stokes optimisation of the ONERA M6 wing presented in the next section.

8.2 Optimisation of the ONERA M6 wing


The ONERA M6 wing and the flow solution calculated with MERLIN around it have already been presented in Section 5.8. It is a good test case to test the three-dimensional optimisation methodology developed in this work, that is why the drag minimisation of the M6 wing is presented in this section. Similar optimisations of the M6 wing can be found in the literature either for inviscid flows^[18,136,145] or for viscous turbulent flows.^[147,162,190]

Here the drag minimisation is performed using the variable-fidelity method for turbulent flows. Following the two-dimensional example of Section 7.3, the low-fidelity model solves the Euler equations on a coarse Euler grid and the high-fidelity model solves the fully turbulent Navier-Stokes equations using the Baldwin-Lomax model on a Navier-Stokes grid. The two grids employed are shown in Figure 8.1: the low-fidelity grid is a $65 \times 17 \times 7$ grid divided into 4 blocks while the Navier-Stokes grid is the same $193 \times 49 \times 33$ grid as in Section 5.8 and is divided into 8 blocks.

Six wing master sections are modified during the optimisation and are shown in red in Figure 8.1. For consistency between the low- and high-fidelity models, these sections are common grid sections of the Euler and the Navier-Stokes grids. These sections are parameterised in the same way as the two-dimensional aerofoil of the previous chapter, i.e. with 8 active Bézier parameters controlling the shape modifications of each surface. In addition a design variable controls the increment in twist of each master section. The most outboard master section does not undergo shape modification in order to avoid the complexity of deforming the wing tip. It is only allowed to twist. This twist change is transmitted to all the sections that form the tip and that are situated outboard of this last master section. The master sections are connected linearly to form the wing. In total the optimisation hence uses 86 design variables ($5 \times 17 + 1$).



(a) Low-fidelity Euler grid



(b) High-fidelity Navier-Stokes grid

Figure 8.1: The two grid levels used for the optimisation of the ONERA M6 wing. The red sections of the high-fidelity grid are the master sections that are optimised and correspond to the sections of the low-fidelity grid.

The flow conditions are the same as in Section 5.8. In particular the overall angle of incidence is 3.06° but this will be modified during optimisation since the 6 master sections are free to twist independently from one another around their leading edge point.

The optimisation problem is to minimise drag at a constant lift while maintaining the internal volume of each master section. This can be written as:

$$\begin{aligned} \text{Minimise} \quad & C_D \\ \text{Subject to:} \quad & C_L \geq 0.2696 \\ & V_{0i} \leq V_i \leq 2 \times V_{0i}, \quad i = 1, \dots, 5 \end{aligned} \quad (8.1)$$

Note that no constraint on the interior volume of the 6th master section is needed since it only undergoes twist changes that do not modify the volume. Each design variable is scaled in the same way as in the previous chapter i.e. in order that an increment of 1 of each design variable taken in isolation produces a maximum deformation of 2% chord for each of the sections. The objective function and the constraints are scaled by the norm of their initial gradient in the variable-fidelity method. Problem (8.1) is transformed into a corrected low-fidelity model very similar to problem (7.4) with the addition of 8 linear constraints for the additional geometrical constraints on volume. The merit function Φ_{hi} employed in the optimisation to assess the performance of the design and corresponding to problem (8.1) is

$$\begin{aligned} \Phi_{hi} = & \frac{C_D}{\left\| \frac{dC_D}{d\beta^0} \right\| = 0.01143} + 10 \frac{(C_L - 0.2696)^2}{\left\| \frac{dC_L}{d\beta^0} \right\|^2 = 0.1477^2} + 10 \frac{(V_1 - 0.0313)^2}{\left\| \frac{dV_1}{d\beta^0} \right\|^2 = 0.01305^2} \\ & + 10 \frac{(V_2 - 0.0227)^2}{\left\| \frac{dV_2}{d\beta^0} \right\|^2 = 0.009506^2} + 10 \frac{(V_3 - 0.0173)^2}{\left\| \frac{dV_3}{d\beta^0} \right\|^2 = 0.007315^2} \\ & + 10 \frac{(V_4 - 0.0137)^2}{\left\| \frac{dV_4}{d\beta^0} \right\|^2 = 0.005812^2} + 10 \frac{(V_5 - 0.0114)^2}{\left\| \frac{dV_5}{d\beta^0} \right\|^2 = 0.004886^2} \end{aligned}$$

It is also very similar to the merit function used in two dimensions in Section 7.3 with additional terms for the constraints on interior volume.

Following the conclusions of the two-dimensional studies presented in the previous chapter, the flow and adjoint solvers are only required to converge to 10^{-5} in this optimisation. This should not impact too much on the accuracy while reducing computing time. To be more precise, each flow solution is stopped after 5000 implicit iterations or when the total residual has reached 10^{-5} , whichever comes first. The adjoint solutions are stopped after 3000 iterations or when the total residual has reached 10^{-5} , whichever comes first as before. Both the flow and adjoint computations restart from the previous solution. The




(a) Merit function Φ_{hi} (b) Drag coefficient C_D (c) Lift coefficient C_L

Figure 8.2: Evolution of different parameters during the variable-fidelity optimisation of the ONERA M6 wing.

optimisation is started from converged flow and adjoint solutions on the baseline geometry.

The optimisation history is shown in Figure 8.2. The computation is done on the 4 processor Alpha machine and the CPU time shown is the total computing time obtained by adding the time of the 4 processors. The merit function Φ_{hi} in Figure 8.2(a) quickly goes down and reaches a minimum after 265 hours and cannot then improve on that design. The drag coefficient in Figure 8.2(b) follows the same trend and the minimum of drag corresponds to the minimum of the merit function. Figure 8.2(c) shows that this minimum is very close to satisfying the constraint on lift, the difference with the target being very small despite the appearance of the graph.

The performance of the optimum point is summarised in Table 8.1 and compared to the

	C_L	$C_{D\ total}$	$C_{D\ press}$	$C_{D\ fric}$
Initial	0.2697	0.01736	0.01241	0.00495
Optimised	0.2694	0.01478	0.00973	0.00505
Target	0.2696	-	-	-

Table 8.1: Aerodynamic coefficients of the optimised ONERA M6 wing.

Origin of data	Initial			Optimised			
	C_L	C_D	L/D	C_L	C_D	Improvement in C_D	L/D
Present work	0.2697	0.01736	15.54	0.2694	0.01478	-14.9%	18.23
Nielsen and Anderson ^[162]	0.253	0.0168	15.06	0.253	0.0142	-15.5%	17.82
Lee <i>et al</i> ^[190]	0.2622	0.01751	14.97	0.2580	0.01586	-9.4%	16.27
Sung and Kwon ^[147]	N/A	N/A	N/A	-2.1%	N/A	-12.8%	N/A

Table 8.2: Comparison of aerodynamic coefficients obtained for optimised ONERA M6 wings.

baseline geometry. The drag coefficient loses 26 drag counts, that is a 15% drag improvement, while the lift is maintained. The drag breakdown shows that the improvement is coming from the pressure drag since the skin friction drag slightly increases. This is not surprising: since the internal volume is kept the same, there is not much hope of reducing the friction drag. The improvement in pressure drag is certainly coming from a reduction of the wave drag as the contour plots of Figure 8.3 suggest. This picture has to be compared with the pressure distribution on the baseline geometry in Figure 5.10. Clearly the lambda shock wave has been much reduced, except in the region where the two shocks merge and at the wing tip where limited shape modifications are allowed.

The chordwise pressure distributions of Figure 8.4 confirm this reduction in the strength of the shock wave. They are shown at the same spanwise locations where experimental data are available for the baseline geometry and that are plotted in Figure 5.11. At all the spanwise positions, the shock wave has been reduced by the optimisation, the sections at $\eta = 0.44$ and $\eta = 0.65$ almost managing to eliminate it completely.

Finally the six master sections are shown before and after optimisation in Figure 8.5. They are plotted in absolute coordinates that are scaled by the semi-span since the initial geometry is such that the semi-span is equal to 1. Both shape and twist deformations are shown but it clearly appears that the increment in twist is very small for the 6 sections. The shape changes are not large either but prove sufficient to greatly reduce the drag.

These results show that this optimisation is very successful and compares well with other Navier-Stokes optimisations found in the literature as can be seen in Table 8.2. The drag




Figure 8.3: Contours of pressure coefficient on the upper surface of the optimised ONERA M6 wing. (to be compared to Figure 5.10).

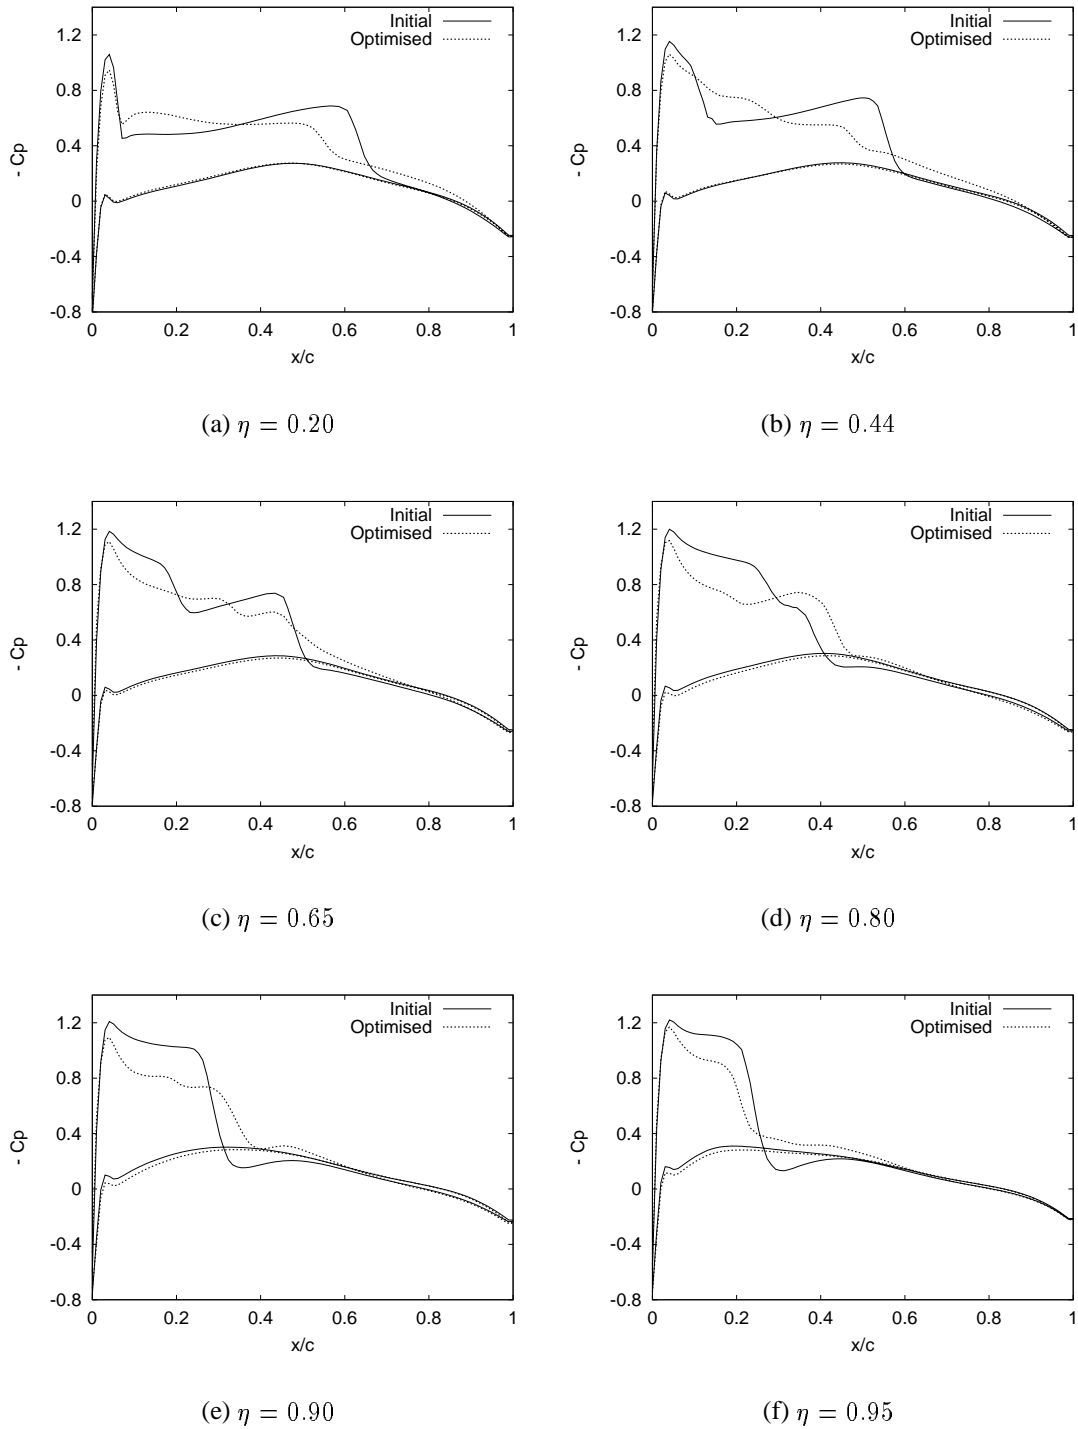



Figure 8.4: Chordwise C_p distributions for the optimised ONERA M6 wing.

(g) $\eta = 0.99$ Figure 8.4: Chordwise C_p distributions for the optimised ONERA M6 wing. (Concluded)

improvement is similar to that of Nielsen and Anderson,^[162] although their geometry changes are much more pronounced and they obtain very different pressure distributions. The present pressure distributions are closer to these of Lee *et al*^[190] and Sung and Kwon^[147] although their overall drag improvement is not as good.


The conclusion of this section is that the variable-fidelity method and the overall optimisation chain worked well on this three-dimensional Navier-Stokes optimisation of the ONERA M6 wing since a 15% drag improvement was achieved while maintaining lift. The goal of this test case, i.e. to demonstrate the applicability of the method to a three-dimensional problem, is achieved, it remains to apply this same procedure to the optimisation of a BWB as is described in the next section.

8.3 Optimisation of a BWB

8.3.1 Background about the BWB and about the present work

The Blended Wing-Body (BWB) is a tailless aircraft that integrates wing and fuselage. It is a revolutionary conceptual change from the classical design that has prevailed for the last 50 years i.e. a wing attached to a cylindrical fuselage and a tail to ensure the manoeuvrability of the aircraft. This concept is supposed to offer substantial aerodynamic performance improvement compared to the classical design. The BWB results from an evolution of the old concept of flying wings. This update of the flying wing is due to some work in the US over the last decade that attracted some attention due to its novelty.^[241–243]

Liebek in Reference [244] gives a good summary of the studies carried out on this subject in the US. It all started at the beginning of the 1990's by a conceptual study sponsored by



(a) Root section






(b) $\eta = 0.34$ (c) $\eta = 0.58$ (d) $\eta = 0.77$ (e) $\eta = 0.90$ (f) $\eta = 0.98$

Figure 8.5: Shape modification of the master sections of the ONERA M6 wing.

NASA. Then a feasibility program was started that involved NASA, McDonnell Douglas and several American universities, in particular Stanford University.^[20,245] This project involved some wind tunnel testings and the construction of a flying model. While this concerned a very large aircraft able to carry 800 passengers, Boeing then started its own study on a family of smaller BWBs for a maximum of 450 passengers once the concept had been proved to work.^[244] This study made use of the multidisciplinary conceptual design program WingMOD developed by Wakayama.^[246–248] The claimed performance of the resulting aircraft exceeds that of conventional airplanes.

The work in the US attracted the interest of other parties in the world:^[249] a group of MSc students under the supervision of Dr. Smith from the College of Aeronautics of Cranfield University designed a BWB in 1998,^[250,251] Airbus is also reported to study the concept as well as TsAGI in Russia.^[252] They are all considering a large aircraft capable of carrying 800 passengers. Smaller versions however have also been investigated at Tohoku University in Japan^[21] and at Cranfield University.^[253] Recently the possibility of a sonic BWB flying at a high subsonic Mach number has been considered.^[254]

The work described in this thesis on the BWB is part of the European Commission funded project entitled *MOB - A computational design engine incorporating multidisciplinary design and optimisation for Blended Wing Body configuration*. The aim of the project is to develop a computational design engine (CDE) i.e. an integrated suite of codes to perform the multidisciplinary design and optimisation of an aircraft. This CDE has to be accessed from different sites across Europe to enable people from different companies to work on the same project. The test case considered for this project to assess the CDE is a freighter BWB. The starting configuration is the BWB designed at Cranfield University in 1998.

A recent session at the 9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization in Atlanta was devoted to the MOB project and all its aspects were presented: an introduction of MOB,^[255] a description of the CDE^[256] and of the software architecture behind it,^[257] the data management system employed,^[258] the model generator based on the CAD software ICAD developed,^[259] the aerodynamic analyses and design performed on the BWB^[260] as well as the studies on aeroelasticity and flight mechanics carried out on this geometry.^[261]

The present study is part of the aerodynamic work done for the MOB project. As is briefly described in the next section, a first assessment of the baseline configuration was carried out before performing a twist inverse design of the BWB. The three-dimensional optimisation of the BWB with the adjoint method is described after that. The author only knows one other reference presenting a three-dimensional optimisation of a BWB and this was done in the MOB project: Weinerfelt^[262] used a continuous adjoint method to perform some Euler optimisation on the BWB. Hence the aerodynamic optimisation of the BWB that is presented in this thesis contributes significantly to the understanding of BWB aerodynamics

Number of points in normal direction	y^+	C_L	$C_{D\,total}$	$C_{D\,press}$	$C_{D\,fric}$
60	40	0.3357	0.02473	0.02436	0.00037
60	13	0.4089	0.02853	0.02488	0.00365
60	5	0.4136	0.03268	0.02504	0.00764
60	1	0.4164	0.03301	0.02538	0.00763
120	1	0.4214	0.03178	0.02411	0.00767

Table 8.3: Grid sensitivity study for the baseline BWB. $M_\infty = 0.85$, $\alpha = 3.0^\circ$.

and to the novelty of this study.

8.3.2 Preliminary work

This section quickly reviews the work that was carried out on the BWB geometry before doing any optimisation. This is fully reported in References [172, 260] and here only the aspects important for the optimisation are discussed.

The starting point was the assessment of the aerodynamic performance of the baseline BWB. To be sure that the aerodynamic coefficients given by the CFD code MERLIN for this configuration were accurate, a grid sensitivity study was carried out. This was necessary because the first grid created around the geometry produced a skin friction drag that was at least one order of magnitude smaller than what was expected. The baseline grid has $149 \times 60 \times 55$ points. Five grids are compared obtained only by changing the stretching in the normal direction and the number of points in that direction. The effect of the stretching on the flow solution is measured by an estimate of the non-dimensional distance y^+ at the first grid point away from the surface. The results of this grid sensitivity study are summarised in Table 8.3.

The first four grids have the same number of points in the normal direction and it is obvious that the grid stretching has an effect on the skin friction drag since the grid with $y^+ \approx 40$ totally underestimates its value. As the grid is further clustered towards the surface, the skin friction drag converges towards a common value. The same thing happens to the lift coefficient. The two grids with $y^+ \leq 5$ give similar results although the one with $y^+ \approx 1$ has less points away from the near wall region. To investigate this effect, the last grid was built. Its number of points is doubled in the normal direction while the stretching is kept the same as for the fourth grid. This results in a grid with nearly one million points that becomes very costly to run. The computation on this last grid shows that there is still some sensitivity in the aerodynamic coefficients but at least the grids with 60 points in the normal direction and $y^+ \leq 5$ give reasonable answers at a reduced computational cost. As a result of this grid sensitivity study, the grid with 60 points and $y^+ \approx 5$ was thought to be the best compromise, with some clustering towards

Twist distribution	C_L	$C_{D\ total}$	$C_{D\ press}$	$C_{D\ fric}$	$C_{D\ wave}$	L/D
baseline	0.4136	0.03268	0.02504	0.00764	0.00407	12.66
elliptic	0.4102	0.02837	0.02031	0.00806	0.00209	14.46
averaged	0.4090	0.02783	0.02008	0.00774	0.00180	14.70
triangular	0.4071	0.02866	0.02083	0.00783	0.00161	14.20

Table 8.4: Comparison of the performance of the three redesigned BWB geometries.

the wall but with still numerous points outside this near wall region giving reasonable numerical accuracy, all of this being obtained at an affordable computational cost. The normal stretching of this $y^+ \approx 5$ grid is reused for all the Navier-Stokes grids employed for the optimisations of the BWB.

The assessment of the baseline geometry revealed that at the design C_L of 0.41, a very strong shock wave was present on the upper surface of the outer wing and on the winglet. This resulted in a very high wave drag and a disappointing lift to drag ratio of $L/D = 12.66$. The study of the spanwise lift distribution showed that the outer wing was too highly loaded and that on the contrary, the fuselage was not generating enough lift. One of the theoretical advantage of the BWB is that the fuselage is capable of generating a fair amount of lift as opposed to the cylindrical fuselage of the conventional design but this was not happening here. It was then decided to try and shift the loading from the outboard to the inboard and to distribute it more evenly along the span in the hope that this would improve the aircraft efficiency.

This was achieved through a twist redesign using a panel code. The aircraft planform and the aerofoil sections were kept unchanged, only their twist angle was modified. An inverse design methodology was applied: a target spanwise lift distribution had to be matched by only changing the twist of the sections. Since there was no prior knowledge of what the best lift distribution was, three target distributions were tested: a triangular distribution, an elliptic distribution and an average between the triangular and elliptic distributions. Once the redesign had been completed with the panel code, the resulting geometries were tested using the Navier-Stokes solver MERLIN. The results of these computations are shown in Table 8.4. The wave drag values are extracted from the flow solutions following the method of Reference [263].

Table 8.4 shows that all of the redesigned geometries are better than the baseline one since the wave drag has been reduced by half or more. This directly impacts on the lift to drag ratio which is increased by up to 2 i.e. a 16% improvement. As expected, since the aerofoil sections are not changed, the skin friction drag varies very little. It has to be noticed that the redesign improves the pressure drag further than it reduces the wave drag, which contributes to the increase in aerodynamic performance. Since the geometry with a twist distribution average between elliptic and triangular gives the best L/D , it is




Figure 8.6: Baseline BWB geometry.

the design considered for further work on the BWB.

The twist inverse design proved that some substantial improvement could be achieved by slightly modifying the geometry. However its potential is limited since only the twist was changed and this was done manually. Some further degrees of freedom in the design have to be introduced, such as the aerofoil shapes for example, to further improve the performance of the BWB. This is what is done with the application of the optimisation methodology developed in this work to the optimisation of the BWB as is described next.

8.3.3 Baseline geometry and optimisation problem

The baseline geometry that serves as reference for the optimisation is shown in Figure 8.6 and is coming from the twist inverse design exercise but is not exactly the same geometry. This is due to some evolutions in the MOB project: the latest available MOB reference aircraft was considered to start the optimisation. The Navier-Stokes grid around this baseline geometry is shown in Figure 8.7. For obvious reasons of symmetry, only a semi-span

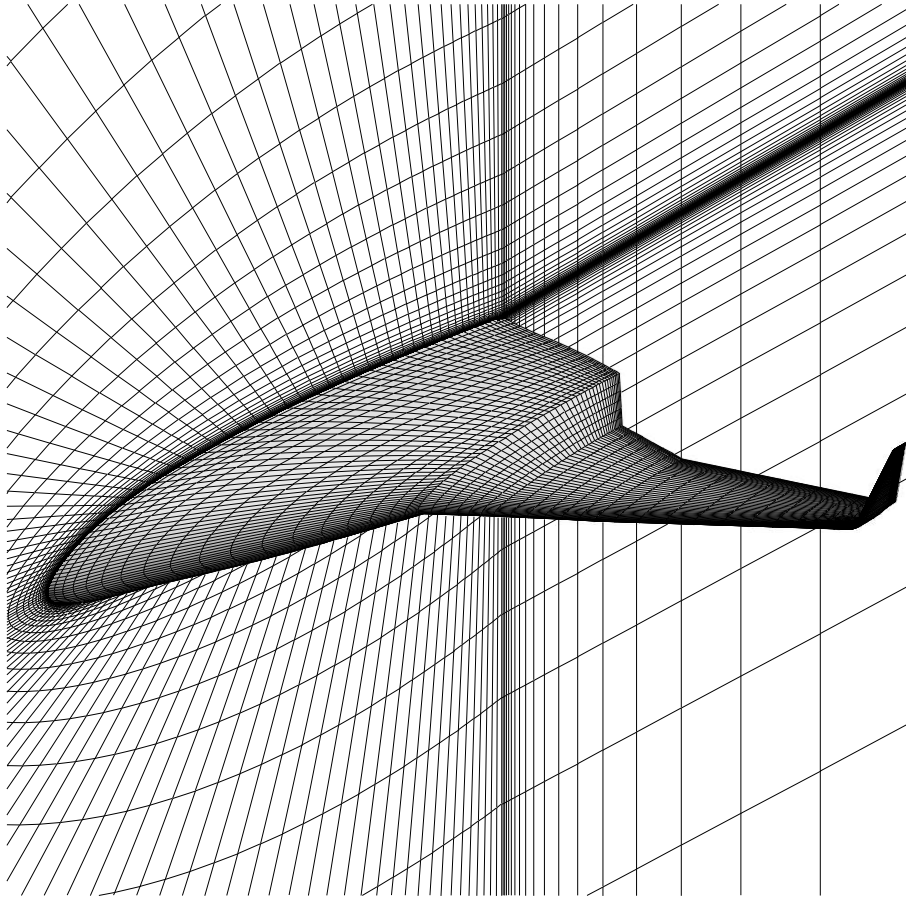


Figure 8.7: Fine Navier-Stokes grid around the baseline geometry.

model is employed. The grid has $161 \times 65 \times 45$ points divided into 16 blocks. It is considered throughout this optimisation as a fine Navier-Stokes grid and serves as a reference.

The flow conditions for the optimisation are: Mach number $M_\infty = 0.85$, unit Reynolds number $Re = 6.675 \times 10^6$ per metre and freestream temperature $T_\infty = 218.8 K$. At the design point, the lift coefficient is required to be 0.41 based on the trapezoidal area of 841.7 m^2 that is used to non-dimensionalise all the aerodynamic coefficients. The pitching moment is calculated around the centre of gravity situated 29.288 m behind the nose of the aircraft in the plane of symmetry. The reference length used for the pitching moment is the mean aerodynamic chord equal to 12.288 m. A positive value of the pitching moment corresponds to a nose down moment.

The flow solution at the design C_L is calculated using MERLIN. The resulting aerodynamic coefficients for the baseline geometry are given in Table 8.5. These are reference coefficients to which optimisations results will be compared. The wave drag here and for the optimised BWBs presented in the remaining of this chapter, only accounts for the drag

C_L	$C_{D\ total}$	$C_{D\ press}$	$C_{D\ fric}$	$C_{D\ wave}$	L/D	C_m
0.4101	0.02855	0.01885	0.00969	0.00101	14.37	0.07360

Table 8.5: Aerodynamic coefficients for the baseline BWB geometry at the design C_L .




Figure 8.8: Contours of pressure coefficient on the upper surface of the baseline BWB geometry at the design C_L .

generated by the wing and fuselage shock wave(s). The winglet wave drag is not included.

The corresponding contours of pressure coefficient on the upper surface of the BWB are shown in Figure 8.8. Despite the twist redesign, a strong shock wave is still present, that extends from the outer wing up to the fuselage. Since nothing has been done to the winglet in the inverse twist design, the original strong shock wave on its upper surface is still there and perturbs the flow at the tip of the outer wing.

Several drag minimisation optimisations of this baseline geometry are presented next. They all share the same parameterisation that is very similar to the one used for the optimisation of the ONERA M6 wing. 16 wing master sections are considered. The 15 most inboard are parameterised using 17 design variables each, like for the two-dimensional optimisations: 8 active Bézier parameters for the upper surface, 8 for the lower surface

and 1 for the twist increment of that section. The master section most outboard is only allowed to twist like for the M6 wing optimisation. This is also to avoid complication when deforming the wing tip where the winglet is attached. By only allowing twist changes, the winglet follows the changes of the last master section and is thus just rotated accordingly. Of course since the shape of the winglet is unchanged, it is unlikely that the strong shock wave sitting on it will disappear but making shape changes would require substantial improvements in the geometry modeller that are left as future work for the moment. In total this makes 256 active design variables ($15 \times 17 + 1$). They are all scaled as before i.e. so that taken in isolation, an increment of 1 in any of the design variable produces a maximum displacement of 2% of the chord of the appropriate master section.

The optimisation problem is again to minimise drag at a constant lift while maintaining the internal volume of each master section. An additional constraint on the pitching moment is or is not included depending on the optimisation. This can be written as:

$$\begin{aligned}
 \text{Minimise} \quad & C_D \\
 \text{Subject to:} \quad & C_L \geq 0.41 \\
 & (C_m)^2 \leq (0.001)^2 \quad \text{optional} \\
 & V_{0i} \leq V_i \leq 2 \times V_{0i}, \quad i = 1, \dots, 15
 \end{aligned} \tag{8.2}$$

The constraint on the pitching moment is necessary due to the nature of the BWB: it is a tailless aircraft. It has thus to be trimmed by deflecting the control surfaces situated all along the trailing edge of the aircraft. Some calculations done at NLR for the MOB project found that the trim drag i.e. the drag generated by the necessary deflection of the control surfaces to trim the aircraft, was very penalising. Hence it would be better if the BWB shape was trimmed naturally or required very small control surface deflections during the cruise condition. That is why this constraint on the pitching moment is included that aims at forcing the pitching moment to stay close to zero to get a trimmed aircraft. However this constraint is limiting: we will see that it reduces the potential improvements in L/D . In addition a consistent longitudinal stability analysis should include the effect of thrust but since the propulsion system is not considered in the MOB project, this effect has been neglected. Hence requiring a trimmed aircraft by only considering the aerodynamic forces is not very realistic. It is only if we suppose that the thrust acts on a line passing by the centre of gravity of the aircraft, the moment generated by the thrust being zero in that case. This is an assumption that we can make to justify the trim requirement. However in the general case, since the engines are supposed to be mounted on the upper surface of the fuselage, the thrust will generate a nose down pitching moment that will have to be compensated by a nose up pitching moment from the aerodynamic forces. Without further information, the requirement of a trimmed aircraft is restricting and possibly dubious. To demonstrate the feasibility of including a constraint on pitching moment, some optimisations will thus include this constraint, while to investigate the potential aerodynamic improvements, others will not.

All the optimisations of the BWB employ the variable-fidelity method and the optimisation problem (8.2) is transformed into a corrected-low fidelity optimisation similar to problem (7.4). Additional linear constraints are added for the interior volume and for the optional constraint on pitching moment. This latter, when employed, is written:


$$\frac{-(0.001)^2 + (C_{m\ hi}(\boldsymbol{\beta}^q))^2 + \nabla C_{m\ hi}^2(\boldsymbol{\beta}^q)^t \cdot \mathbf{s}}{\left\| \frac{dC_{m\ hi}^2(\boldsymbol{\beta}^q)}{d\boldsymbol{\beta}^q} \right\|} \leq 0$$

As for the optimisation of the ONERA M6 wing, the flow and adjoint solutions employed in the optimisations of the BWB are only required to converge to 5 orders. More specifically, the flow solutions stop after 5000 implicit iterations or when reaching 5 orders of convergence, whichever comes first, while the adjoint solutions are stopped after 3000 iterations or when reaching 5 orders of convergence, whichever comes first as before. Again all the computations restart from previous solutions to save some computing time and again the optimisations start from converged flow and adjoint solutions calculated on the baseline geometry with the grids and physical models employed in these optimisations.


Unfortunately, a Navier-Stokes optimisation using the fine grid for the high-fidelity model was not successful on this geometry for reasons that will be discussed in Section 8.4. Hence the author took advantage of the interesting findings made with the two-dimensional optimisations of the previous chapter to perform optimisations on the BWB with lower-fidelity models. This includes Euler optimisations on fine Euler grids and Navier-Stokes optimisations on coarse grids, both of which brought substantial improvement to the aerofoil design in 2D. This implies that, throughout these BWB optimisations, two points of view will be presented: an optimisation point of view that is concerned by the optimisation itself, analyses if it is successful or not and what improvements have been made; and the point of view of an aerodynamicist or designer whose aim is to improve the BWB geometry. In this latter case, whatever the optimisation method employed, the final optimised geometry is checked on a fine Navier-Stokes grid and compared to the initial reference BWB presented above to assess the true improvements. The next three sections details the results of these optimisations. The first two consider the Euler optimisations while the last one describes the Navier-Stokes optimisations on a coarse grid.

8.3.4 Euler optimisations of the BWB without constraint on C_m

Two Euler optimisations of the BWB without any constraint on the pitching moment were carried out. They used the variable-fidelity method despite the fact that both the low- and high-fidelity models solved the Euler equations. The low-fidelity model uses a coarse Euler grid with $81 \times 17 \times 22$ points divided into 16 blocks while the high-fidelity model employs a finer Euler grid with $81 \times 33 \times 39$ points also divided into 16 blocks. These two grids are shown in Figure 8.9. The 16 master sections can be seen in that picture.



(a) Low-fidelity Euler grid



(b) High-fidelity Euler grid

Figure 8.9: The two grid levels used for the Euler optimisations of the BWB. The red sections of the high-fidelity grid are the master sections that are optimised and correspond to the sections of the low-fidelity grid.

i	V_{0i}	$\left\ \frac{dV_{0i}}{d\beta^0} \right\ $
1	0.11578	2.8708×10^{-2}
2	0.10867	2.6857×10^{-2}
3	9.3004×10^{-2}	2.2742×10^{-2}
4	6.9747×10^{-2}	1.7098×10^{-2}
5	4.4219×10^{-2}	1.1502×10^{-2}
6	2.5742×10^{-2}	7.7115×10^{-3}
7	1.6391×10^{-2}	5.7425×10^{-3}
8	9.6732×10^{-3}	3.8391×10^{-3}
9	4.7623×10^{-3}	2.3559×10^{-3}
10	3.2007×10^{-3}	1.6716×10^{-3}
11	1.9455×10^{-3}	1.1047×10^{-3}
12	1.3716×10^{-3}	7.7117×10^{-4}
13	8.2362×10^{-4}	4.5765×10^{-4}
14	5.1351×10^{-4}	2.8591×10^{-4}
15	3.8734×10^{-4}	2.1882×10^{-4}

	Connections	
	linear	spline
$\left\ \frac{dC_D}{d\beta^0} \right\ $	0.005783	0.005731
$\left\ \frac{dC_L}{d\beta^0} \right\ $	0.10066	0.10250

Table 8.6: Value of the coefficients used in the evaluation of Φ_{hi} in Equation (8.3).

The difference between the two optimisations is coming from the connection between the master sections. Chronologically, the first optimisation was done with a linear connection between the master sections as for the M6 wing. It worked well but since the fuselage of the BWB is “fat” and curved, using linear connections to represent this part of the geometry creates slight slope discontinuities that do not look “nice” especially when transposed to a Navier-Stokes grid. It does not create flow discontinuities so it is perfectly valid for optimisation, the only problem being that the resulting optimal geometry would have to be smoothed in some way before being used by other disciplines. To avoid this problem, the grid update program described in Chapter 4 was modified to connect the sections by piece-wise cubic splines. This was done only on the fuselage since the inner and outer wings look very much like a conventional aircraft wing where linear connections are perfectly suited. The difference between the two types of connections is a posteriori relatively small but the change is still useful.

The merit function Φ_{hi} employed in the variable-fidelity method to assess the performance of the design and corresponding to problem (8.2) is, in both cases,

$$\Phi_{hi} = \frac{C_D}{\left\| \frac{dC_D}{d\beta^0} \right\|} + 10 \frac{(C_L - 0.41)^2}{\left\| \frac{dC_L}{d\beta^0} \right\|^2} + 10 \sum_{i=1}^{15} \frac{(V_i - V_{0i})^2}{\left\| \frac{dV_{0i}}{d\beta^0} \right\|^2} \quad (8.3)$$

with the value of the coefficients given in Table 8.6.




(a) Merit function Φ_{hi} (b) Drag coefficient C_D (c) Lift coefficient C_L

Figure 8.10: Evolution of different parameters during the Euler variable-fidelity optimisations of the BWB without any constraint on C_m .

		C_L	C_D	L/D	C_m
Linear connections	Initial	0.4100	0.02109	19.44	0.08957
	Optimised	0.4059	0.01719	23.61	0.10501
	Optimised at target C_L	0.4101	0.01738	23.59	0.10646
Spline connections	Initial	0.4101	0.02125	19.30	0.08973
	Optimised	0.4081	0.01688	24.17	0.13941
	Optimised at target C_L	0.4100	0.01697	24.16	0.14008

Table 8.7: Inviscid aerodynamic coefficients of the Euler optimised BWBs without any constraint on C_m .

The optimisation history is shown in Figure 8.10. The computation is done on the 4 processor Alpha machine and the CPU time shown is the total computing time. So the first point to notice is that these optimisations are relatively fast compared to the optimisation of the M6 wing that used the Navier-Stokes equations on a fine grid and the same Alpha machine. It is not surprising since only the Euler equations are employed here and the grids are quite small.

The optimisation history shows that both optimisations are successful: the merit function Φ_{hi} is reduced, which is reflected by a decrease in drag coefficient while the lift constraint is globally satisfied. Of the two, the optimisation with spline connections between the master sections is the best since it is faster and reaches a point that both has a lower drag and better satisfies the lift constraint. The aerodynamic coefficients of the optimised geometries compared to the initial ones are given in Table 8.7. This confirms this remark that the optimisation with spline connections achieves better performance: starting from a geometry with an L/D of 19.30, it finds an optimum with an L/D of 24.16. This is a reduction of 43 drag counts i.e. a 20% improvement. As stated earlier, these improvements have however an adverse effect on the pitching moment for both optimisations since the final value is greater than the initial one, dragging the design away from a trimmed aircraft. The fact that the pitching moment for the optimisation with linear connections is better than for the other optimisation might explain why the final design is not as good in term of L/D .

Figure 8.11 compares the pressure distribution of the initial geometry with the one of the optimised BWB. This is obtained in inviscid mode on the fine Euler grid. Both the initial and optimised geometries have spline connections. The pattern on the initial geometry is very similar to what is obtained with the Navier-Stokes equations in Figure 8.8: a very strong shock wave spans the whole wing up to the fuselage. In inviscid mode, as expected, this shock wave is situated further aft on the wing. The optimisation clearly manages to weaken this shock wave that has now disappeared on the inner wing. It seems it has however reinforced it at the wing tip where the interaction with the winglet shock, still present and very strong, is more pronounced.





Figure 8.11: Comparison of the contours of pressure coefficient on the upper surface of the initial BWB and of the Euler optimised BWB without any constraint on C_m . Euler calculations.

This overall picture is completed by the chordwise pressure distributions given at six spanwise locations shown in Figure 8.12. Though the pressure distribution at the root is not much affected, significant changes occur as early as 17% of the span i.e. midway through the fuselage, where the build-up of the shock wave has been totally eliminated by the optimisation. The station at $\eta = 0.40$ is a very strong example of this elimination of the shock wave achieved by both optimisations. As we move further outboard on the wing, a shock wave starts to form though and is even stronger than on the initial geometry (station at $\eta = 0.93$). The last station at 98% of the span i.e. very close to the junction between the wing and winglet shows the disturbed flow field in this region with strange pressure distributions. Overall both optimisations give similar pressure distributions despite differences in the global aerodynamic coefficients.

This similarity in pressure distributions is understood when looking at the shape modifications made at these stations and shown in Figure 8.13: the two optimisations produce shapes that are very close to each other. The sections are shown in their true coordinates that are non-dimensionalised by the chord of the root section. The sections are thus seen in their twisted position. Please note however that the scale on the vertical axis is very different from the one on the horizontal axis, which gives the impression that the outer wing sections are twisted at 45° while it is only 6 or 7° in reality. This scale issue also gives the false impression that design changes are small on the fuselage and large on the



(a) Root section


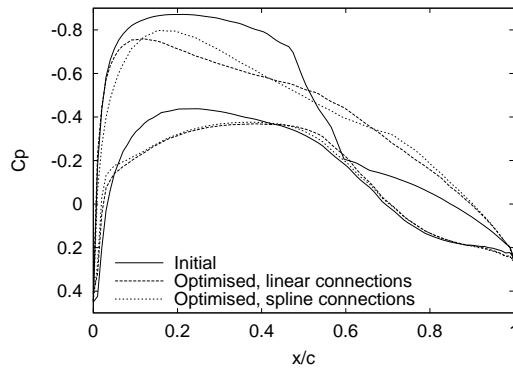




(b) 4th master section at $\eta = 0.17$ (c) 8th master section at $\eta = 0.40$ (d) 12th master section at $\eta = 0.71$ (e) 14th master section at $\eta = 0.93$ (f) 15th master section at $\eta = 0.98$

Figure 8.12: Chordwise C_p distributions for the Euler optimised BWBs without any constraint on C_m . Euler calculations.



(a) Root section






(b) 4th master section at $\eta = 0.17$ (c) 8th master section at $\eta = 0.40$ (d) 12th master section at $\eta = 0.71$ (e) 14th master section at $\eta = 0.93$ (f) 15th master section at $\eta = 0.98$

Figure 8.13: Shape modification of some master sections for the Euler optimised BWBs without any constraint on C_m .




Figure 8.14: Spanwise twist distribution in degrees for the Euler optimised BWBs without any constraint on C_m .

	C_L	C_{Dtotal}	C_{Dpress}	C_{Dfric}	C_{Dwave}	L/D	C_m
Initial reference	0.4101	0.02855	0.01885	0.00969	0.00101	14.37	0.07360
Linear connections	0.4100	0.02572	0.01581	0.00992	0.00015	15.94	0.08526
Spline connections	0.4101	0.02576	0.01581	0.00995	0.00042	15.92	0.11533

Table 8.8: Navier-Stokes check on the fine Navier-Stokes grid of the Euler optimised BWBs without any constraint on C_m .

outer wing, which is not true in real values. However presenting the shapes in this way offers the advantage of being able to see that the shapes and twist angles are modified, because scaled back to the aerofoil chord, these changes would be too small to be seen.

The twist angles are however not changed much as can be seen with the spanwise twist distributions of Figure 8.14. These changes are very small compared to the twist redesign exercise of References [172, 260] but the optimiser has the additional freedom of modifying the aerofoil shapes to get where it wants to go so the twist angles might not play such an important role. It is also possible that the twist design variables should be scaled differently to give them more importance if that is what is required.

Despite these small changes, the improvements in L/D are substantial as seen before, at least in Euler mode. What has been presented so far for these two optimisations shows that they are successful: starting from a given Euler grid and flow solution, the optimisations have managed to find much better design points while satisfying the constraints. But is it still the case when transferred to the fine Navier-Stokes grid? This is now the point of view of the aerodynamicist who wants to know if the optimisation has managed to improve his starting geometry. The answer is given in Table 8.8 where the design changes obtained by the Euler optimisations are applied to the fine Navier-Stokes grid of Figure 8.7 that is then used by MERLIN to assess the aircraft performance.




Figure 8.15: Contours of pressure coefficient on the upper surface of the Euler optimised BWB without any constraint on C_m . Navier-Stokes calculation on the fine Navier-Stokes grid (to be compared to Figure 8.8).

Table 8.8 shows that the improvements are also substantial in Navier-Stokes mode. The optimisation with linear connections is better this time but the difference with the spline connection optimisation is very small. An improvement of 28 drag counts is achieved i.e. a reduction of 10% of the total drag. Even if the lift to drag ratio is still poor compared to conventional aircraft, this improvement is quite important. Table 8.8 shows that it is obtained by a reduction of the pressure drag, part of it coming from a lower wave drag. Since the wetted area should not be changed very much by the optimisation, there is little hope of reducing the skin friction drag that even slightly increases. It is also reassuring to see that the behaviour observed in Euler mode for the pitching moment is well reflected by the Navier-Stokes checks: both optima have a higher than initial pitching moment but the optimum corresponding to the linear connections has the lowest value.

These good results are confirmed when looking at the pressure distribution for these Navier-Stokes checks shown in Figure 8.15 for the linear connection optimum. The colour bar has the same scale as Figure 8.8 enabling a direct comparison. Clearly the strength of the shock wave has been reduced and it no longer affects the fuselage, being stopped in the middle of the inner wing. As in the inviscid case, its strength is reinforced though at the wing tip.

This is confirmed by the chordwise pressure distributions of Figure 8.16. The optimisation with linear connections manages to eliminate completely the shock wave at the stations $\eta = 0.17$ and $\eta = 0.40$ as in the Euler mode. The other optimisation does not perform as well at 40% of the span with still a small shock wave. At $\eta = 0.71$ both optimisations have a shock wave but again the one with spline connections is stronger. These two stations show that the optimum with spline connections has a stronger shock wave than the other one, which explains the higher wave drag found in Table 8.8. As seen in Figure 8.15, the shock wave becomes very strong close to the wing tip (station at $\eta = 0.93$) for both optima and is much stronger than on the initial geometry.

The spanwise lift distribution and loading extracted from these Navier-Stokes flow solutions are shown in Figure 8.17. Here again changes are much smaller than for the twist redesign. The changes made however seem to go in the wrong direction since part of the load on the fuselage is shifted to the outer wing while the inverse effect was sought through the twist inverse design. This can explain though why the shock wave is weaker on the fuselage and inner wing while stronger on the outer wing.


The conclusion of these Euler optimisations without any constraint on the pitching moment is that they are relatively efficient and successful. Their main advantage is that the behaviour of the Euler optimised BWBs translates very well on the fine Navier-Stokes grid, making them a very good tool for optimisation at a reduced cost. The improvements gained on the baseline BWB are substantial. However these gains are obtained to the detriment of the pitching moment since they pull the geometries away from a trimmed aircraft. Having in mind the comments made in the previous section, an Euler optimisation with a trimmed aircraft is presented next.

8.3.5 Euler optimisation of the BWB with constraint on C_m

The variable-fidelity method is still employed for this optimisation in exactly the same way as in the previous section: the Euler mode is used for the low- and high-fidelity models and the same coarse and fine Euler grids (Figure 8.9) are employed. Since the method with spline connections between master sections was satisfactory, it is this one alone that is kept here.

The only difference with the optimisations of the previous section is the optimisation problem solved. It is now problem (8.2) with the constraint on pitching moment. This has an effect on the merit function Φ_{hi} employed in the variable-fidelity method, that becomes

$$\Phi_{hi} = \frac{C_D}{\left\| \frac{dC_D}{d\beta^0} \right\|} + 10 \frac{(C_L - 0.41)^2}{\left\| \frac{dC_L}{d\beta^0} \right\|^2} + 10 \frac{C_m^2}{\left\| \frac{dC_m^2}{d\beta^0} \right\|} + 10 \sum_{i=1}^{15} \frac{(V_i - V_{0i})^2}{\left\| \frac{dV_{0i}}{d\beta^0} \right\|^2} \quad (8.4)$$



(a) Root section






(b) 4th master section at $\eta = 0.17$ (c) 8th master section at $\eta = 0.40$ (d) 12th master section at $\eta = 0.71$ (e) 14th master section at $\eta = 0.93$ (f) 15th master section at $\eta = 0.98$

Figure 8.16: Chordwise C_p distributions for the Euler optimised BWBs without any constraint on C_m . Navier-Stokes calculations on the fine Navier-Stokes grid.

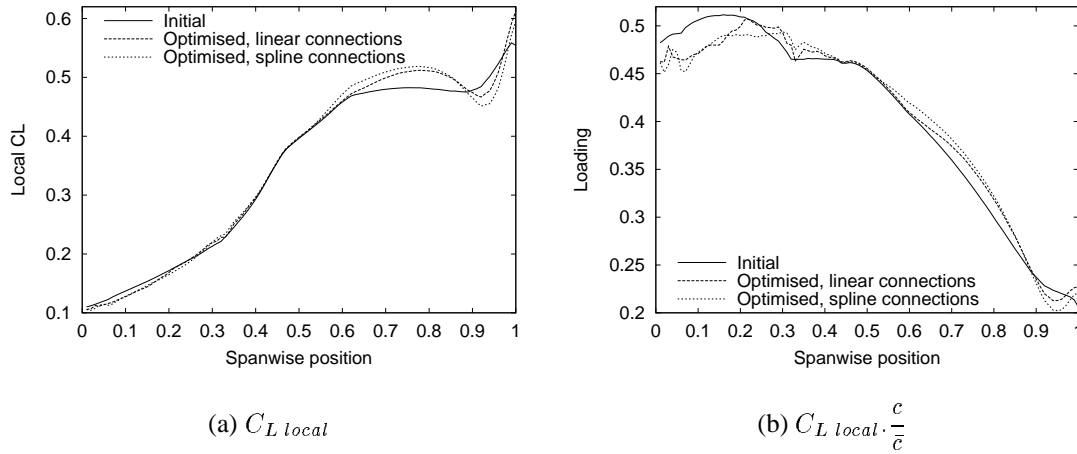


Figure 8.17: Spanwise lift distribution and loading for the Euler optimised BWBs without any constraint on C_m . Navier-Stokes calculations on the fine Navier-Stokes grid.

Note that the function C_m^2 and its gradient are used directly rather than C_m and its gradient. This should not change anything in the optimisation. The value of the different coefficients is still given by Table 8.6 with the addition of $\left\| \frac{dC_m^2}{d\beta^0} \right\| = 0.02937$. Note also that the formula for Φ_{hi} is slightly inconsistent with the constraint on pitching moment shown in problem (8.2): the upper limit of $(0.001)^2$ is used in the low-fidelity optimisation rather than zero, to provide an attainable target that could make the constraint feasible (it is important since FFSQP is the optimiser employed and thus searches for feasible constraints before anything). However what is really desired is not 0.001 but zero for the pitching moment, hence the way Φ_{hi} is set up should favour designs with a very low pitching moment even if it is relatively far from 0.001. Since this value is almost 2 orders of magnitude smaller than the initial pitching moment, this inconsistency should not however make any difference.

The optimisation is still performed on the 4 processor Alpha machine. The evolution of different parameters during this optimisation is shown in Figure 8.18. The first comment is that the process takes longer than before but this is not surprising since one more adjoint solution is required each time the full gradient is needed, due to the additional constraint on C_m . Figure 8.18 shows that the optimisation is successful: the merit function Φ_{hi} is well reduced; the drag is improved although not in a monotonic way; the constraint on lift is satisfied or almost at the optimum although the optimisation wandered away during the process; and the constraint on pitching moment is very close to being satisfied after starting from an initial point far away in the infeasible domain. It is remarkable that such a low pitching moment can be achieved while improving on drag and maintaining lift considering the poor initial point. A gradient-based method like the one employed here





(a) Merit function Φ_{hi} (b) Drag coefficient C_D (c) Lift coefficient C_L (d) Pitching moment coefficient C_m

Figure 8.18: Evolution of different parameters during the Euler variable-fidelity optimisation of the BWB with constraint on C_m .

	C_L	C_D	L/D	C_m
Initial	0.4101	0.02125	19.30	0.08973
Optimised	0.4088	0.01727	23.67	0.00292
Optimised at target C_L	0.4101	0.01733	23.66	0.00338

Table 8.9: Inviscid aerodynamic coefficients of the Euler optimised BWB with constraint on C_m .





Figure 8.19: Comparison of the contours of pressure coefficient on the upper surface of the initial BWB and of the Euler optimised BWB with constraint on C_m . Euler calculations.

could have reached a local minimum well before attaining this point.

The aerodynamic coefficients of the Euler optimum are shown in Table 8.9: an L/D of 23.66 is reached with a very low pitching moment. This compares favourably to the lift to drag ratio of 24.16 obtained by the optimisation without constraint on C_m but that had a pitching moment 40 times higher. This optimisation brings an 18% drag improvement (39 drag counts).

The pressure coefficients on the upper surface of the initial geometry are compared to these of the optimum in Figure 8.19. As previously, the strength of the shock wave is visibly reduced and only affects the region outboard of the mid-inner wing. This is further highlighted by the chordwise pressure distributions of Figure 8.20: the shock wave has been eliminated almost up to the station at $\eta = 0.71$. A careful comparison with the pressure distributions of Figure 8.12 reveals that the constrained aircraft carries



(a) Root section







(b) 4th master section at $\eta = 0.17$ (c) 8th master section at $\eta = 0.40$ (d) 12th master section at $\eta = 0.71$ (e) 14th master section at $\eta = 0.93$ (f) 15th master section at $\eta = 0.98$

Figure 8.20: Chordwise C_p distributions for the Euler optimised BWB with constraint on C_m . Euler calculations.



(a) Root section






(b) 4th master section at $\eta = 0.17$ (c) 8th master section at $\eta = 0.40$ (d) 12th master section at $\eta = 0.71$ (e) 14th master section at $\eta = 0.93$ (f) 15th master section at $\eta = 0.98$

Figure 8.21: Shape modification of some master sections for the Euler optimised BWB with constraint on C_m .




Figure 8.22: Spanwise twist distribution in degrees for the Euler optimised BWB with constraint on C_m .

less lift on the rear part of the sections (after 60% chord say) for all the stations except perhaps the one at $\eta = 0.71$. On the fuselage, it even carries negative lift at the rear of the sections (root section and station at $\eta = 0.17$). Furthermore a comparison of Figure 8.19 and Figure 8.11 shows that the shock wave on the constrained optimised BWB is situated further forward on the wing than for the unconstrained BWB. All of this gives an explanation as to how the low pitching moment is achieved: removing some lift on the rear part of the sections reduces the nose down pitching moment around the centre of gravity.

The shape deformations shown in Figure 8.21 also explain how this is achieved: on the fuselage, the twist is decreased and reverse camber seems to be introduced while on the wing ($0.40 \leq \eta \leq 0.71$), the rear part of the sections is thickened by transferring some volume from the lower surface leading edge region to the rear of the aerofoil on the upper surface. This was also done on the unconstrained BWB (Figure 8.13) but the phenomenon is amplified here. The twist changes shown in Figure 8.22 are bigger than previously and are closer in magnitude to the changes introduced by the twist inverse design. The trend is also very clear: as explained above, the twist is decreased on the fuselage (that stops at $\eta = 0.34$ just where the two twist curves cross each other) while it is increased on the wing.

From an optimisation point of view, this Euler optimisation with constraint on C_m proves to be very satisfactory: the introduction of the constraint on pitching moment reduces its magnitude very much as was expected and the good thing is that this does not degrade the very good drag improvement already achieved with an unconstrained problem. Let us consider now the designer point of view to see if these changes are transferable to a Navier-Stokes solution on a fine grid. The answer is yes as Table 8.10 proves: the changes made on the Euler optimised BWB, transferred to the fine Navier-Stokes grid, produce a geometry that is very efficient and has a very low pitching moment that is even slightly nose up now. Like for the BWB without any constraint on C_m , the wave drag has been

	C_L	$C_{D\ total}$	$C_{D\ press}$	$C_{D\ fric}$	$C_{D\ wave}$	L/D	C_m
Initial reference	0.4101	0.02855	0.01885	0.00969	0.00101	14.37	0.07360
Euler optimised	0.4100	0.02595	0.01592	0.01003	0.00023	15.80	-0.00401

Table 8.10: Navier-Stokes check on the fine Navier-Stokes grid of the Euler optimised BWB with constraint on C_m .

well reduced while the skin friction drag slightly increases. Overall, the drag is reduced by 26 drag counts i.e. 9%, all of this being achieved while reducing the pitching moment 18 times in magnitude.

The contours of pressure coefficient for this Navier-Stokes solution are shown in Figure 8.23. The most noticeable feature is that a double shock wave is present on the upper surface of the wing compared to a single one on the BWB with constraint on C_m . This should explain why the resulting wave drag is higher than for the linear connection optimum. Otherwise compared to the baseline geometry, the extent and strength of the shock wave have been much reduced as previously while the strong shock wave on the winglet is still there.

The corresponding chordwise pressure distributions are given in Figure 8.24. The negative lift carried at the rear of the fuselage sections in inviscid mode is found in viscous turbulent mode as well, as the root section and the station at $\eta = 0.17$ show. The shock wave is eliminated on the fuselage as well as further outboard (station at $\eta = 0.40$). The double shock wave is clearly visible at 71% of the span and extends up to 93% of the span.

The resulting spanwise lift distribution and loading are plotted in Figure 8.25. A fair amount of lift is lost at the wing tip probably due to the stronger shock wave (compare the station at $\eta = 0.98$ in Figures 8.24 and 8.16) at the tip that interacts with the winglet shock. Otherwise some lift is redistributed on the fuselage itself: it is removed from the centre part and shifted to the outboard of the fuselage.

The conclusion of this section is that this Euler optimisation of the BWB with constraint on pitching moment is successful since it achieves its aim of greatly reducing the pitching moment while reducing the drag at constant lift. The good thing is that once again the behaviour observed in Euler mode is very well reflected in Navier-Stokes mode on a fine grid. This has two consequences: first our belief that the Euler optimisation is a very efficient tool to perform optimisation that would be more costly on a higher-fidelity model is reinforced; secondly, the BWB geometry obtained is quite good and it will be difficult to find a much better one by starting from the same baseline and using a similar optimisation methodology. This is however what is tested in the next section: since substantial improvements have been achieved only by using a medium-fidelity physical model for these optimisations, it might be possible that by using a high-fidelity model i.e.




Figure 8.23: Contours of pressure coefficient on the upper surface of the Euler optimised BWB with constraint on C_m . Navier-Stokes calculation on the fine Navier-Stokes grid (to be compared to Figure 8.8).

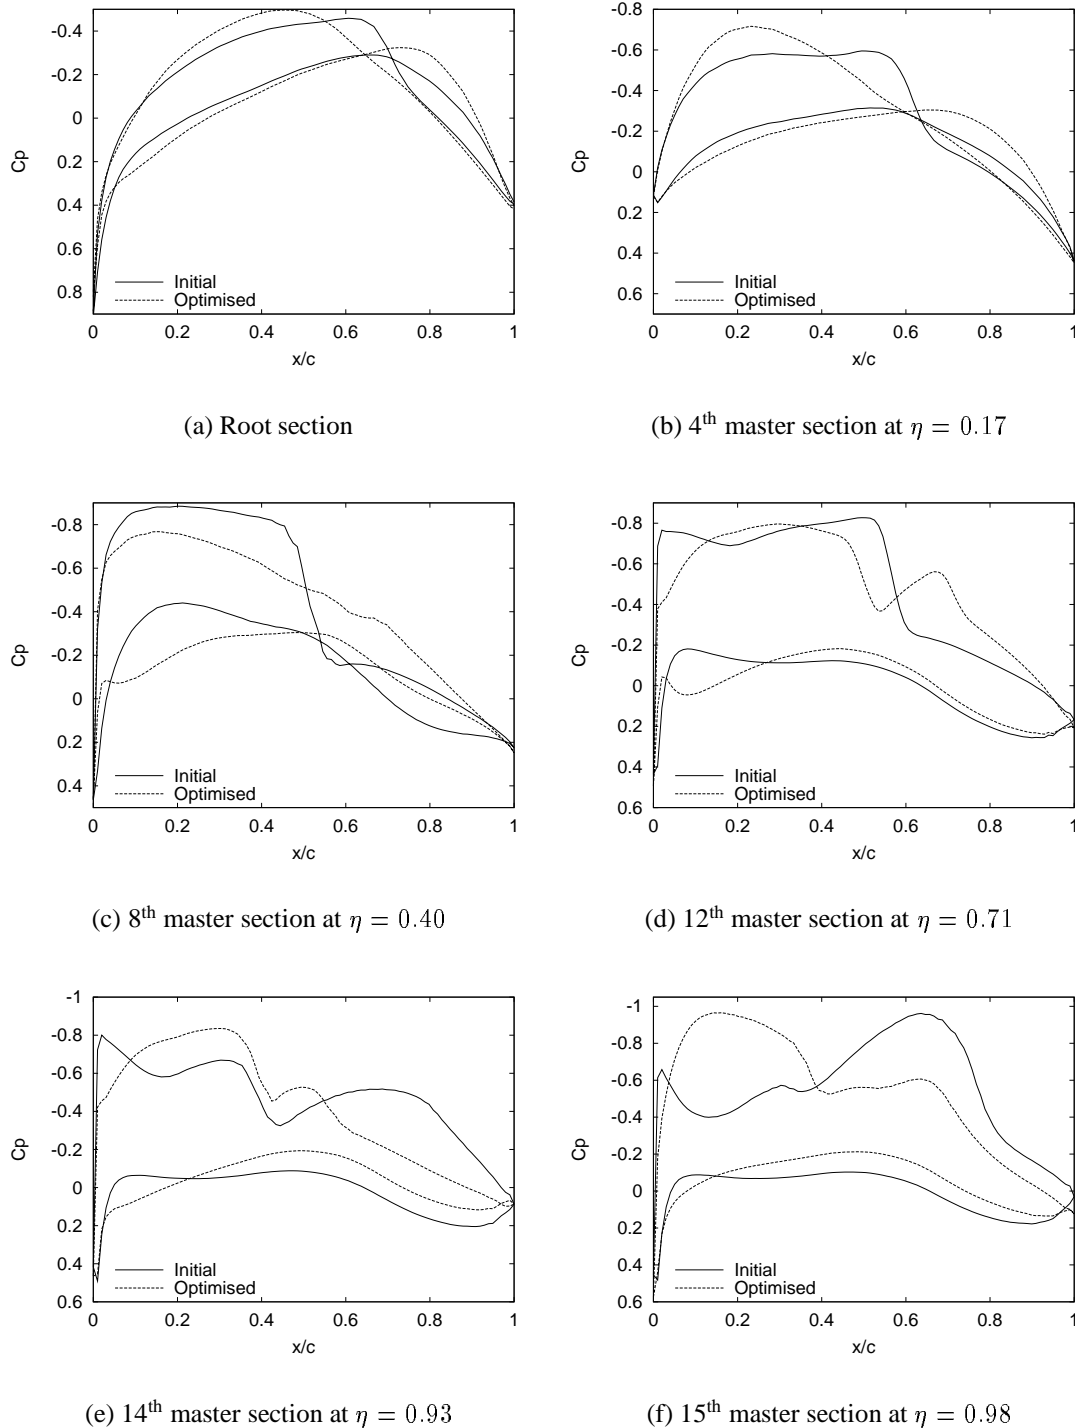


Figure 8.24: Chordwise C_p distributions for the Euler optimised BWB with constraint on C_m . Navier-Stokes calculations on the fine Navier-Stokes grid.

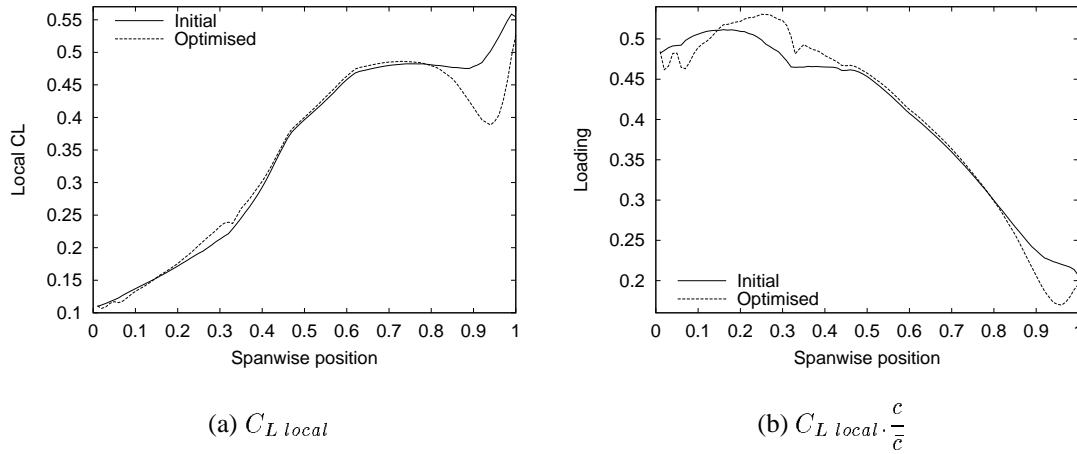


Figure 8.25: Spanwise lift distribution and loading for the Euler optimised BWB with constraint on C_m . Navier-Stokes calculations on the fine Navier-Stokes grid.

the turbulent Navier-Stokes equations, even greater performance could be attained. That is why Navier-Stokes optimisations of the BWB with and without constraint on the pitching moment have been carried out as is described next.

8.3.6 Navier-Stokes optimisations of the BWB on a coarse grid

Two optimisations of the BWB starting from the baseline geometry were carried out by using the Navier-Stokes equations on a coarse grid. These optimisations are with and without constraint on pitching moment to be able to compare them to the same optimisations performed in Euler mode and described in the previous two sections.

The variable-fidelity method is still employed: the low-fidelity model solves the Euler equation on the same coarse Euler grid shown in Figure 8.9(a); the high-fidelity model solves the Navier-Stokes equations on a coarse Navier-Stokes grid shown in Figure 8.26. This grid has the same number of points as the fine Euler grid employed previously (Figure 8.9(b)) i.e. $81 \times 33 \times 39$ points divided into 16 blocks, but the grid is much more clustered towards the surface of the wing. The grid stretching, at least for the first point away from the surface, is the same as for the fine Navier-Stokes grid of Figure 8.7 to provide an adequate grid for the Navier-Stokes equations. The surface grid on the BWB is the same as for the fine Euler grid and Figure 8.26 illustrates the master sections that correspond to the grid sections of the coarse Euler grid, in exactly the same way as previously. The parameterisation is also unchanged and uses the same 256 design variables.

The optimisation problem solved is still problem (8.2) and the merit function Φ_{hi} is




Figure 8.26: The coarse Navier-Stokes grid used by the high-fidelity model. The red sections are the master sections.

$\frac{dC_D}{d\beta^0}$	$\frac{dC_L}{d\beta^0}$	$\frac{dC_m^2}{d\beta^0}$
0.006460	0.07816	0.01369

Table 8.11: Value of the coefficients used in the evaluation of Φ_{hi} for the Navier-Stokes optimisations of the BWB.

either given by Equation (8.3) or Equation (8.4) depending on the presence or not of the constraint on pitching moment. The different values for the volumes and their gradient are unchanged and given in Table 8.6 while the norms of the gradient for the aerodynamic functions are given in Table 8.11.

The stopping criteria for the flow and adjoint solutions were slightly modified: for the low-fidelity model, everything is as before but the maximum number of iterations was reduced to 2500 for the high-fidelity Navier-Stokes flow solutions. This is because the flow solution never converged to 10^{-5} on this coarse Navier-Stokes grid and stalled after a reduction of three orders of magnitude of the total residual. This means that the flow calculation is always performed during the optimisation until the maximum number of iterations is reached even with the restart capability employed. Since the design changes are not too important around 1500 implicit iterations should be sufficient to get a non-changing flow solution. To be conservative, the maximum is set to 2500 iterations but this is already half of the 5000 iterations used for the Euler optimisations and thus should save a lot of computing time. These two optimisations were carried out on one node of 16 processors of the IBM SP parallel supercomputer. Since the time limit of the jobs run on this machine is of 12 wall-clock hours, the optimisations had to be restarted several times, which is not too much of a problem with the variable-fidelity method because each high-fidelity cycle could be done in isolation.

The evolution of the aerodynamic coefficients during the optimisation is plotted in Figure 8.27 for both optimisations with and without constraint on pitching moment. Since the magnitude of the merit function Φ_{hi} is greatly changed with the introduction or not of this constraint, it is represented in Figure 8.27(a) as a ratio of its evolving value over its initial value to be able to plot the two merit functions in a single graph. This graph shows that the optimisation is struggling to find a better design point especially in the unconstrained case. Nonetheless it manages to reduce the drag by a fair amount in this latter case. As expected it is less successful for the pitching moment constrained problem. The constraint on lift is well satisfied for the optimisation without constraint on C_m but it is more chaotic for the other optimisation. It is not surprising however that the minimum of Φ_{hi} is reached when the target lift is matched. The constraint on the pitching moment manages to reduce it but the optimal design is still far from the target imposed. This is the first obvious difference with the equivalent Euler optimisation that almost reached the feasibility boundary.

The aerodynamic performance of the optima is summarised in Table 8.12. The striking point is the very high drag for the three configurations that was already present in the graph 8.27(b) but that is emphasised here by the very low L/D ratios. It is not too surprising when remembering the two-dimensional optimisations on coarse grids of Section 7.2.3.2 that presented the same problem. The breakdown of drag into pressure drag and skin friction drag proves that both components are equally affected with values around 50% higher than what they should be. This shows that the grid quality is a major

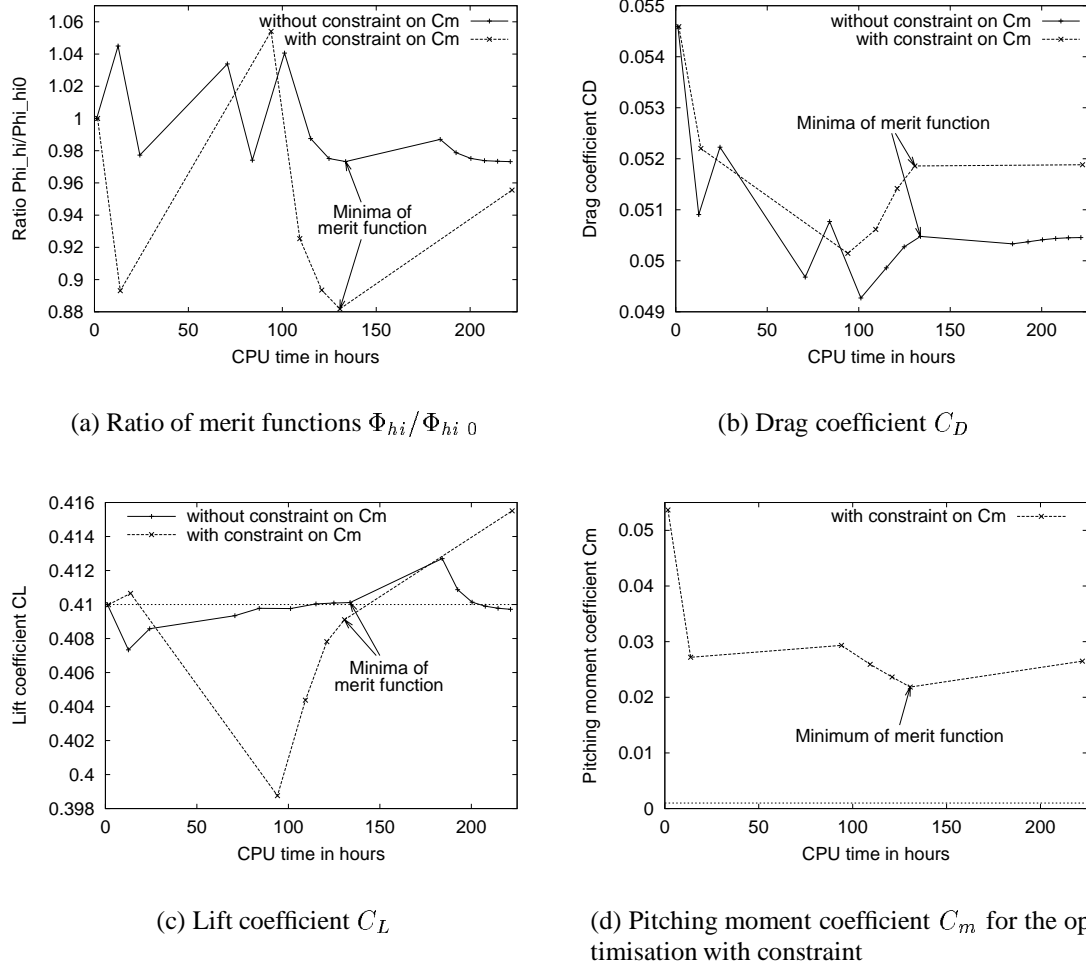


Figure 8.27: Evolution of different parameters during the Navier-Stokes variable-fidelity optimisations of the BWB on a coarse grid with and without constraint on C_m .

		C_L	$C_{D\ total}$	$C_{D\ press}$	$C_{D\ fric}$	L/D	C_m
Initial		0.4100	0.05457	0.03822	0.01635	7.51	0.05366
Without constraint on C_m	Optimised	0.4093	0.05049	0.03507	0.01542	8.11	0.07240
	Optimised at target C_L	0.4101	0.05053	0.03512	0.01542	8.11	0.07261
With constraint on C_m	Optimised	0.4091	0.05186	0.03662	0.01524	7.89	0.02184
	Optimised at target C_L	0.4100	0.05192	0.03668	0.01524	7.90	0.02207

Table 8.12: Aerodynamic coefficients obtained on the coarse grid for the Navier-Stokes optimised BWBs.





Figure 8.28: Comparison of the contours of pressure coefficient on the upper surface of the initial BWB and of the Navier-Stokes optimised BWB without any constraint on C_m . Coarse grid Navier-Stokes calculations.


factor in the solution accuracy and that totally wrong results can easily be obtained if no grid sensitivity study is performed. In the present case as for the two-dimensional optimisations, it is not very important since the aim is only to investigate the optimisation process, the final design point being then checked on a fine grid to assess its performance. Table 8.12 shows that the optimisations manage to reduce drag but the improvements are not as big as for the Euler optimisations: 7% for the optimisation without constraint on C_m and 5% for the other one. As mentioned above the pitching moment is not reduced as much as desired either.

The contours of pressure coefficient calculated on the coarse Navier-Stokes grid are shown in Figure 8.28. The optimised shape in this Figure corresponds to the optimisation without constraint on pitching moment. Surprisingly, despite what has just been said on the drag values, the pressure contours on this coarse grid for the baseline geometry are very close to what is obtained on the fine grid (Figure 8.8): the shock wave is strong and has the same extent, it even stops before the tip region like on the fine grid to leave some space for the strong suction at the rear of the wing tip due to the interaction with the winglet shock. The shock wave on the wing is smeared though but this is due to the coarsening of the surface grid. Even if the reduction in drag after optimisation is not as important as expected, the shock wave has been visibly weakened and its extent reduced.


Figure 8.29 shows the chordwise pressure distributions obtained on the coarse grid for




(a) Root section




(b) 4th master section at $\eta = 0.17$




(c) 8th master section at $\eta = 0.40$



(d) 12th master section at $\eta = 0.71$



(e) 14th master section at $\eta = 0.93$



(f) 15th master section at $\eta = 0.98$

Figure 8.29: Chordwise C_p distributions obtained on the coarse grid for the Navier-Stokes optimised BWBs.


	C_L	$C_{D\ total}$	$C_{D\ press}$	$C_{D\ fric}$	$C_{D\ wave}$	L/D	C_m
Initial reference	0.4101	0.02855	0.01885	0.00969	0.00101	14.37	0.07360
Without constraint on C_m	0.4101	0.02644	0.01663	0.00981	0.00020	15.51	0.09436
With constraint on C_m	0.4100	0.02706	0.01730	0.00978	0.00035	15.15	0.04043

Table 8.13: Navier-Stokes check on the fine Navier-Stokes grid of the optimised BWBs obtained by a Navier-Stokes optimisation on a coarse grid.


both optima and the initial point. At first sight, the optimisation hardly changes the pressure distributions from the initial distributions compared to what was happening for the Euler optimisations. A careful look though reveals that the shock wave is effectively reduced all along the span from the fuselage to the tip region and is even eliminated at 40% of the span. The optimised sections plotted in Figure 8.30 explain why the pressure distributions appear to change so little: it is because the design changes themselves are relatively small with the twist changes almost non-existent. The sections at $\eta = 0.93$ and $\eta = 0.98$ also show that the geometrical constraints on interior volume are not well satisfied since the volumes are greater than what is expected. We have not looked at these constraints so far because they were always well satisfied. The present problem is not explained since the shape parameterisation and geometrical constraints are the same as before.

As a consequence, the concluding comments from an optimisation point of view concerning these two coarse grid Navier-Stokes optimisations are mixed: on the one hand, the drag and pitching moment are reduced and the lift is maintained, all of this being reflected by a weakening of the shock wave as expected, but on the other hand the reduction in drag is not substantial, the constraint on pitching moment is still far from being satisfied and a problem with the constraint on interior volume appears. If the Euler optimisations had not been performed these coarse grid Navier-Stokes optimisations would have been qualified as successful but the insertion of the Euler optimisations in this thesis makes this conclusion less obvious.


The fine grid Navier-Stokes checks are nevertheless carried out on the optimal geometries, the resulting aerodynamic coefficients being given in Table 8.13. The good point is that some improvement has been made in the right direction with a decrease of 7% of the drag for the unconstrained optimum and 5% for the constrained one (the same reduction as on the coarse grid). This decrease is partly due to the weakening of the shock wave since the wave drag has been well reduced. These improvements are not as good as what was obtained with the Euler optimisations and has a consequence, the lift to drag ratios are relatively low. The reduction of the pitching moment by the introduction of the constraint on C_m is also going in the right direction but is nothing compared to what the Euler optimisation achieved. The only comforting comment is that these fine grid checks confirm once again the behaviour observed on a lower-fidelity model that is obtained in this case by coarsening the grid.




(a) Root section




(b) 4th master section at $\eta = 0.17$




(c) 8th master section at $\eta = 0.40$



(d) 12th master section at $\eta = 0.71$



(e) 14th master section at $\eta = 0.93$



(f) 15th master section at $\eta = 0.98$

Figure 8.30: Shape modification of some master sections for the Navier-Stokes optimised BWBs.




Figure 8.31: Contours of pressure coefficient on the upper surface of the Navier-Stokes optimised BWB without any constraint on C_m . Navier-Stokes calculation on the fine Navier-Stokes grid (to be compared to Figure 8.8).

The pressure contours of the optimised BWB without any constraint on pitching moment, calculated on the fine grid, are presented in Figure 8.31. Compared to the baseline picture of Figure 8.8, the strength of the shock wave has been reduced. Its extent has been reduced as well but not as much as the Euler optimisations did because the inner wing is still very much affected. A second shock wave is also forming just behind the main shock on the outer wing, looking like the double shock present on the Euler optimised BWB with constraint on C_m . The wing tip flow field does not seem to be affected by the design changes.

All of this is confirmed by the chordwise pressure distributions extracted from the fine grid solutions for these Navier-Stokes optimised geometries. These are plotted in Figure 8.32. For all the stations up to $\eta = 0.93$, it is clear that the optimisation has weakened the shock wave but this latter is still very much present where the Euler optimisation managed to completely eliminate it (stations at $\eta = 0.17$ and $\eta = 0.40$). The station at 70% of the span shows the double shock wave noticed on the contour plot but it is only present on the unconstrained BWB. Finally the last section at the tip shows that not very much is happening there as noticed above also.

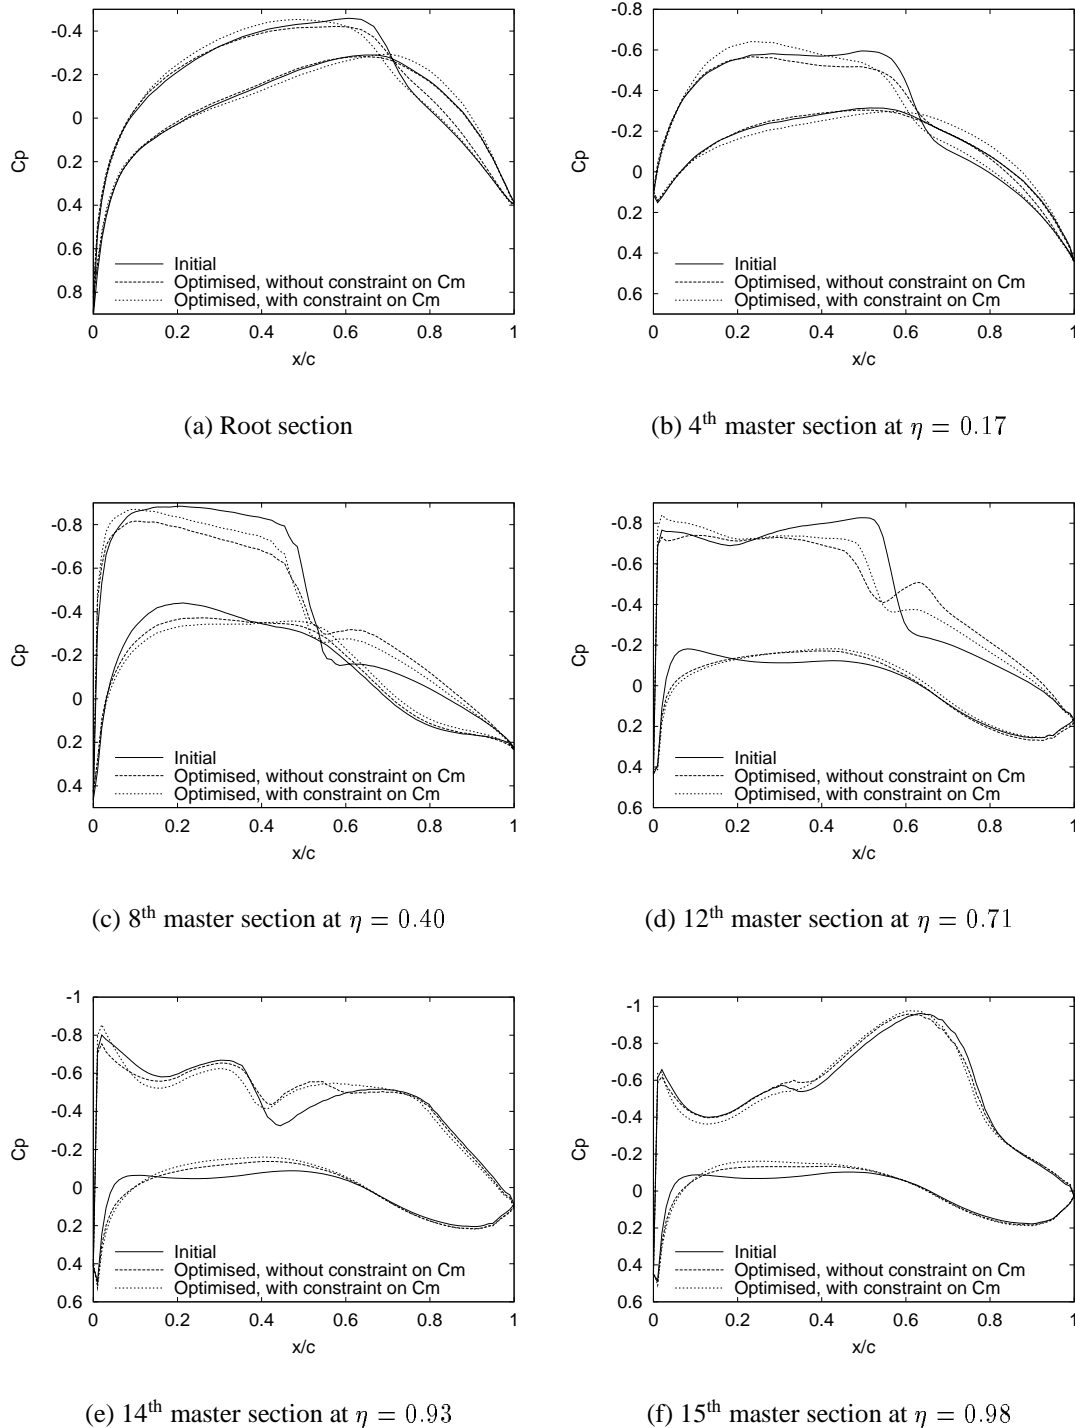


Figure 8.32: Chordwise C_p distributions for the Navier-Stokes optimised BWBs. Navier-Stokes calculations on the fine Navier-Stokes grid.

Since these Navier-Stokes optimisations did not improve on the baseline geometry as much as the Euler optimisations, there is no need to detail any further the aerodynamics of the optimised shapes. If any optimised geometry has to be picked up for further work, it will be one resulting from the Euler optimisations. This unfortunately sums up well this section on the Navier-Stokes optimisations of the BWB: they manage to improve the geometry but not as much as desired or expected.

This terminates this long part dedicated to the optimisation of a BWB. Overall the outcome is positive since the variable-fidelity method applied in Euler mode worked well: the decrease in drag was substantial, the constraint on lift well satisfied and the addition of a constraint on pitching moment for one of the optimisations was successful. The Euler optima when checked with the Navier-Stokes equations on a fine grid to be compared to the baseline geometry proved to be valuable geometries, reducing the total drag by 9% in the case of a trimmed aircraft. This demonstrated that the Euler optimisations are both relatively fast and a very good tool to perform aerodynamic optimisation on industry-type problems. The introduction of a higher-fidelity model with the Navier-Stokes optimisations on a coarse grid was however not so successful: they reduced the drag but not as much as expected and the constraint on pitching moment, when added, was not well satisfied. However these Euler and Navier-Stokes optimisations confirmed the findings made in the previous chapter i.e. that low-fidelity optimisations can improve substantially the performance of a design even when measured with a high-fidelity model, all of this being achieved at a reduced cost compared to a high-fidelity optimisation.

The other aspect that this part considered was the BWB itself. Here the outcome is not very satisfactory. The first theoretical advantage of the BWB is that of aerodynamic efficiency. Starting from an aerodynamically poor baseline geometry, we expected the optimisation to greatly improve on it and provide a much more efficient aircraft. The optimisation carried out its share of the work but did not do any miracle and the optimised geometry is still not very satisfactory. We will come back to this problem in the next section. To end this chapter, a section of discussion is indeed added that looks at the two aspects of this thesis i.e. the optimisation methodology with adjoint solver developed and the application of this methodology to the BWB.

8.4 Discussion

The aim of this section is to take a step back and reconsider with a global vision all the results presented in the last two chapters. Its goal is also to explain where the difficulties were encountered and why certain choices were made. This is to help people who would be tempted to develop a similar method, to understand our approach and possibly to avoid some of our mistakes by choosing other paths. This will thus introduce some ideas that will be described as future work in the next chapter of conclusions. As mentioned above,

this section deals with the two aspects of this work: the optimisation method itself and the BWB.

First of all let us come back to the choice of the variable-fidelity method to perform the three-dimensional optimisations. Chapter 7 showed that the variable-fidelity method was faster than a direct SQP method on the two-dimensional test case used. This is one of the reasons why it was chosen for the optimisations in the current chapter. However the SQP method was also tested on the BWB problem but did not work and thus has not been reported. The main problem encountered was robustness. The optimiser is fed with information in the form of the objective function, the constraints and their gradients. As soon as one of these informations was slightly inconsistent, the SQP method failed, often after wandering a bit trying to understand what was happening. Two main problems were the source of inconsistency: a divergence of the adjoint solver and some noise in the objective function. We will come back later to the problem of divergence of the adjoint solver. The numerical noise is mainly found for three-dimensional turbulent problems where as already explained in this thesis, the flow solution fails to converge properly and stalls generally after about 3 orders of convergence. Unfortunately the design changes during optimisation are often so small that when restarting the flow solution to assess a new design point, the flow solver is already in that region of stalled residual. Hence the aerodynamic coefficients it outputs might not be very accurate and might not reflect the design changes since the flow solution seems to have hardly evolved.

These problems come from the choice of the optimisation methodology that requires accurate information. From the two approaches identified by Giles,^[58,264] we chose the second one i.e. a quasi-Newton method that approximates the Hessian of the problem but thus needs accurate hence costly information. This approach is not robust to noise. Maybe a better choice would have been to try the other approach based on a steepest descent method that requires a lot of steps but for all these steps, the information does not need to be very accurate and hence can be obtained at a reduced cost. This is the approach followed by Jameson and his colleagues with their continuous adjoint solver. Their method although based on the steepest descent, relies heavily on the smoothing of the gradient that they perform to improve the conditioning of the optimisation process and that greatly accelerates the optimisation convergence.^[265]

This leads to the other problem encountered with an SQP method i.e. computational time. The variable-fidelity method brings some improvement on this matter but it still requires costly evaluations of the high-fidelity objective function and constraints and of their gradients. The alternative methods presented in Section 7.3 seem to be very cheap although their reported performance only relate to two-dimensional problems and to the knowledge of the author, have not been demonstrated for three-dimensional optimisations. The steepest descent-like method of Jameson has and due to the fact that it easily accommodates partially converged flow and adjoint solutions, it is fast, three-dimensional optimisations being performed in a few hours.^[6] A better use of parallel computing by increasing

the present efficiency on a 16-processor machine should also greatly improve the problem.

Finally another problem encountered with an SQP method is the generation of non-physical designs during optimisation. Putting geometrical constraints does not solve the problem entirely nor does using a feasible method to satisfy these constraints at each optimisation cycles. Indeed what happens is that inside one cycle, as soon as the optimiser knows the gradients, it is looking at points far away from where it is. Of course the evaluation of the constraints or of the objective function at these points, shows the optimiser that they do not bring any improvement and thus it reduces the length of the step it is taking in that direction. These evaluations have nevertheless to be performed in the first place and that is where problems occur. Even if the evaluations are carried out correctly at these points, it is a waste of time since the optimiser almost always reduces its step size to consider much closer points. This problem is partly alleviated in the variable-fidelity method by the use of the trust-region approach that forces the optimiser to look in a region close to the point where it is. In the low-fidelity optimisation since an SQP optimiser is still employed, the problem can reappear though, but generally the low-fidelity model is robust enough to cope and is fast enough to be able to afford these often unnecessary evaluations.

The introduction of the variable-fidelity method improved very much on all the problems encountered when using a direct high-fidelity SQP optimisation. Its first advantage is of course computational time that is reduced by the use of the low-fidelity model. But the low-fidelity model also brings robustness and is less noisy since in our case the Euler flow and adjoint solutions always converge well. As just explained, the use of the trust-region also contributes to the robustness and efficiency of the method.

Let us come back to the problem of the divergence of the adjoint solver mentioned above. What happens is that the adjoint solver sometimes diverges for no-apparent reasons since it works fine on very similar geometries. The problem is comparable to what happens when the CFL number of the flow solver is too large and provokes a divergence of the residual. Usually reducing the CFL number solves the problem for the flow solver. With the problem of divergence of the adjoint solver, this does not help at all. To partially fix this problem and still output a gradient in order to let the optimisation go on, the following was done in the adjoint solver: when the total residual reaches 100, which is clearly a sign of divergence, the calculation is restarted from the beginning and only 10 iterations are performed before calculating and outputting the gradient. Since the adjoint calculation is started from a previous solution, this normally ensures that the adjoint residual is low when outputting the gradient, which should be more or less a measure of its accuracy. However since only 10 iterations are performed from the previous solution, the gradient may be very close to what it was before and thus the optimisation will suffer. This problem was encountered only with the Navier-Stokes adjoint but both in two and three dimensions. In two dimensions, this was unfrequent and did not cause major problems. It did in three dimensions and is the main reason why Navier-Stokes optimisations of the BWB on a fine grid are not included in this thesis: they were attempted but did not prove successful,

essentially it is thought because this problem occurred. It also affected the coarse grid Navier-Stokes optimisations that have been presented and might explain why they did not work very well. Fortunately this did not affect the Navier-Stokes optimisation of the ONERA M6 wing that was successful. At the time of writing, we have been made aware of recent work^[266] attributing such non-convergence of the adjoint solver to the presence of complex conjugate pairs of eigenvalues of the adjoint lying outside the unit circle. These correspond to small amplitude limit cycles in the behaviour of the nonlinear solver, but cause blow-up of the linearised adjoint solver. These authors showed that a physically relevant adjoint can be obtained by using GMRES to stabilise the adjoint calculation.

It is indeed necessary to say that the optimisation methodology developed in this work, although not perfect, makes it possible to perform some valuable optimisations. The first example of this is the Euler optimisations of the BWB that worked very well, especially the one with constraint on pitching moment, and proved that this tool can be used in an industrial context to improve the design of an aircraft. However they are only Euler optimisations and a lot of time and effort would have been wasted in developing a Navier-Stokes adjoint solver if they were the only positive results. The Navier-Stokes optimisation of the ONERA M6 wing proves though that this was not a waste of time since it fully employed the capabilities of the method (fully turbulent adjoint and flow solvers, variable-fidelity method, parallelisation of the computations, multiblock grid) and proved successful. The fact that the M6 wing geometry and especially grid are simple might explain why the method worked very well there and not so well on the BWB geometry that has a more complicated grid mainly due to the presence of the winglet.

Let us now discuss the optimisations of the BWB. As already mentioned, they manage to improve substantially the baseline geometry but a lot still remains to be done. These optimisations only did what they were realistically able to do considering their possibilities: only the shape of the sections on the fuselage and the wing were allowed to change and yet within the constraint of keeping their internal volume and more importantly the aircraft planform. The implication of this is obvious with what is happening on the winglet: the strong shock wave present on the interior surface of the winglet on the baseline geometry (see Figure 8.8) is still as strong on the optimised geometries (see Figure 8.23 for example). This is not surprising since nothing was meant to correct this in the optimisation. As already explained, modifying the winglet implies finding an automatic way of deforming or regenerating the connection between the wing and the winglet, which is not a trivial task and the reason why this is not attempted in this work. The wave drag generated by the winglet shock wave is estimated to be of only 6 drag counts^[260] so its elimination will not solve all the problems of aerodynamic efficiency but it does contribute to bad performance: it not only affects the flow on the winglet itself but also interacts with the flow on the wing where some efficiency is certainly lost. Improving the winglet is definitely a requirement to get a better aircraft.

The main requirement though would be to optimise the BWB planform as well as volume

distribution. It is indeed expected that aerodynamic performance could be restored with such changes. However for such aircraft this no longer involves only the discipline of aerodynamics but all the disciplines involved in the design (flight mechanics and structural analysis) and truly becomes a multidisciplinary problem. Indeed the changes carried out in this thesis i.e. modifying the aerofoil profiles and twist should not greatly affect the structural part of the aircraft and its flight mechanics, at least to a first approximation, so this can be done in isolation as a pure aerodynamic optimisation. When changing the planform however, all the disciplines are affected and performing an aerodynamic shape optimisation would be meaningless without the consideration of other disciplines because the resulting shape may not be flyable. However it is felt that a great potential exists in optimising the planform and conversely, that not doing so limits very much the possible improvements.

One of the lessons learnt during this work is that you need to put stringent constraints on the optimisation problem otherwise the optimiser exploits all of the weaknesses in your defence and goes where you do not want it to go. It might not be obvious looking at the results presented in this work but this is the reason for introducing the constraints on internal volumes, the bounds on admissible shape displacements and the constraint on pitching moment for the BWB. The latter is a good example since when you are not constraining it, the optimiser minimises drag at constant lift as you want but also increases pitching moment, which you do not want. By constraining the problem, not much freedom is left to the optimiser to find its path and it is highly possible that the optimum it finds is not the true optimum in the design space available to the designer. Hence a unique optimisation is not enough and several have to be performed sequentially to improve the geometry step by step possibly by releasing some degrees of freedom and constraining others or by changing the objective function. Why not indeed optimise for minimum pitching moment at constant lift and drag and then come back to a drag minimisation problem? A good example of such a design methodology is given in Reference [9] where a high-speed civil transport aircraft is optimised by means of dozens of successive optimisations with a lot of intervention from the designer who does not just press a button and wait for the optimiser to do all the job. Such a methodology applied to the BWB would certainly find some improvements even if only as a pure aerodynamic optimisation.

Finally the optimisations of the BWB performed in this work are only single-point optimisations, which means that the optimisation only considers one flow condition. It is well-known that such approach is dangerous because the optimiser will improve very much the geometry at the flow conditions it is optimising for but the performance of the aircraft at off-design conditions is likely to get worse. This is the same problem as the one mentioned in the previous paragraph of not allowing too much freedom to the optimiser. A better approach is to use multipoint optimisation^[1,9,99,139,164] where several design conditions are considered at the same time, a composite objective function taking information from all the design points being minimised. Such an approach would not be too difficult to introduce in the present methodology but it becomes very computationally

intensive since the flow and adjoint equations have to be solved at all the design conditions for each design point. It is however the only approach that ensures the optimised aircraft is globally better and not just at one design condition. Ideally a multipoint optimisation of the BWB should thus be performed.

This terminates this section of discussion about the optimisation methodology and the problems encountered when developing it. It also discussed the optimisation of the BWB that is the second important contribution of this work. The main lesson is that time was an issue in this study because otherwise a lot of directions would need to be investigated to improve both the optimisation methodology and the design of the BWB. However within the timeframe, meaningful progress was achieved: concerning the optimisation chain, the key achievement is undeniably the successful Navier-Stokes optimisation of the ONERA M6 wing that employed the full capability of the method; regarding the BWB, the Euler optimisations managed to improve the geometry, certainly as realistically as possible having in mind the limitations mentioned above.

This is also the end of this long chapter. It first presented the parallelisation of the flow and adjoint solvers performed by using OpenMP. It then described the three-dimensional Navier-Stokes optimisation of the ONERA M6 wing. This was followed by an intensive part dedicated to the Euler and Navier-Stokes optimisations of the BWB. The present section of discussion ended this chapter. It is now time to end this thesis by presenting its conclusions and by giving some advice for future work.

This page has been left intentionally blank.

Chapter 9

Conclusions

9.1 Summary of achievements and findings

The first part of this thesis was dedicated to the description of the optimisation chain developed in this work. The major task was the development of the discrete adjoint solver to calculate efficiently the gradients used by the optimiser. This was successfully completed by verification through careful comparison between sensitivity derivatives calculated by the adjoint solver and those calculated by a finite-difference method. Optimisations on a two-dimensional problem proved that the chain was working but also that it was quite computationally demanding and that it was difficult to reduce easily this computing cost. This resulted in the adoption, for the rest of the thesis, of the variable-fidelity method of optimisation that combines a low-fidelity model with a high-fidelity model to accelerate the optimisation and which reduces overall computing time. This was found to be better than doing a direct high-fidelity optimisation with a traditional optimiser. To further reduce the computational cost, both the flow and adjoint solvers had to be parallelised using OpenMP.

With these improvements, three-dimensional Navier-Stokes optimisations were possible. The example that best illustrates the capabilities of the optimisation procedure developed, is the Navier-Stokes optimisation of the ONERA M6 wing. This optimisation was successful with a 15% drag improvement at constant lift. This showed that the first objective of this thesis, i.e. to develop an aerodynamic optimisation method based on a discrete adjoint solver for three-dimensional Navier-Stokes flows was achieved successfully.

The other objective of this work was to apply this capability to the optimisation of a Blended Wing-Body aircraft and this was done in the remaining of the thesis. Euler optimisations of a BWB geometry were performed using the variable-fidelity method. Substantial improvements were achieved in the Euler mode that were confirmed by Navier-Stokes calculations on a fine grid: an optimisation found a 9% drag improvement while lift was maintained and the aircraft was forced to be trimmed, which was not the case for the starting geometry. Navier-Stokes optimisations on the same geometry were less suc-

cessful though and did not improve the aircraft performance as much as the Euler optimisations. Even if some improvement was made on the starting BWB geometry, the main conclusion of these optimisations is that they did not manage to find an aerodynamically efficient shape that is the theoretical basis of the BWB concept. However limitations in the way the optimisation chain has been applied as explained in the last section of discussions of the previous chapter, might explain why limited results are obtained. This leads us to give some recommendations for future work as is described next.

9.2 Future work

Having started from a CFD code, the work described in this thesis managed to develop an optimisation method that proved to be successful on some problems but it is clearly far from perfect and lacks some generality. Some indications on how to make it a better tool are given here. Let us come back to Figure 1.1 and try to see what can be done for each component of the optimisation chain.

On the optimiser side, the main task would be to test different optimisation methodologies. This should involve testing a steepest descent-like method similar to what Jameson and his colleagues are using to see if this is more efficient. Developments of the variable-fidelity method should also be attempted following the example of Marduel.^[133] On the aerodynamic side, it could also be interesting to compare the present method where the angle of incidence is a design variable governed by the optimiser and lift is a separate constraint with a more traditional methodology where the constraint on lift is an internal constraint that is satisfied in the flow evaluation and does not appear in the optimisation problem. Which is the best? Extending the method to multipoint optimisation is certainly necessary for realistic usage.

The shape parameterisation and the grid deformation algorithm employed in this work performed well although with limitations. The winglet of the BWB could not be easily optimised for example. This is definitely a point of necessary improvement. Others could include the extension of the parameterisation to other methods and a truly multiblock grid deformation capability in order not to limit the possible grid topologies. The possibility of parameterising a fuselage in addition to a wing would be a valuable contribution.

The flow solver was relatively mature when starting this work so not much may be improved. However the parallelisation of the code using MPI would be useful to be able to use it in parallel on any platform. Developing a capability of performing flow solutions at constant lift coefficient rather than constant angle of incidence would also be necessary if an optimisation method that considers the constraint on lift as an internal constraint, was adopted. Deciding on the convergence of the flow solution based on the convergence of its corresponding aerodynamic coefficients instead of the convergence of the residual would also be a useful addition because the aerodynamic coefficients are the

flow solution outputs that are important to the optimisation. This could also bring some large computational savings since the flow solver would do just the necessary amount of iterations.

The main task concerning the adjoint solver is to investigate the problem of divergence described in Section 8.4. This is thought to be the main cause for the optimisation failures encountered in this work and that limited its extent. Solving this should make the method more robust and applicable to a wider range of problems. Like for the flow solver, the parallelisation of the adjoint code with MPI would be useful to broaden the type of platforms on which the optimisation chain can be run. Controlling the convergence of the adjoint solver by looking at the evolution of the gradient it calculates (the only quantity that matters for the optimisation), could also save some computing time. Finally implementing in the adjoint solver the two-equation turbulence models that already exist in the flow solver, would extend the applicability of the method to other types of flows. This would also eliminate the grid topology requirements imposed by the Baldwin-Lomax model although the cost of optimisation would be greatly increased in three-dimensions.

If some of the points mentioned above were performed, a more valuable optimisation of the BWB could be attempted. This would include optimising the winglet as a first task and then performing multipoint optimisations as already explained in Section 8.4. Also as mentioned in that section, modifying the planform of the BWB would be useful if some coupling with structural analysis and flight mechanics, even with simplified models of these, was developed in order to generate a flyable aircraft. This coupling would also be interesting for the current optimisations or at least an a posteriori check to see what impact the aerodynamic optimisation has on other disciplines.

9.3 Perspectives

The author believes that the two subjects of this thesis, i.e. aerodynamic optimisation and the Blended Wing-Body, are each promised a great future on their own. First the BWB: according to the Boeing study,^[244] this is the aircraft of the future due to the reduction in fuel burned it brings over conventional designs. Only public acceptance and economic/political reasons will decide if such an aircraft ever flies.

Aerodynamic optimisation is not limited by these factors and is highly likely to develop even more than it is today. The recent work by Martins, Reuther and Alonso^[57,267] pushes a new frontier: they are coupling a high-fidelity flow solver to a high-fidelity structural analysis code to perform the high-fidelity multidisciplinary optimisation of a flexible wing using an adjoint solver. Likewise, Nadarajah and Jameson^[268] have recently opened a new field by developing an adjoint-based shape optimisation method for unsteady flows. This is going one step further than the harmonic adjoint of References [10, 11] and is promised a great future for helicopter aerodynamics and turbomachinery. Finally looking further

ahead, when considering the future concept envisioned by NASA of a morphing wing aircraft^[269] that constantly modifies the shape of its wing according to what it wants to do, it is not impossible to think that one day, real-time aerodynamic shape optimisation will exist as an onboard system.

References

- [1] Jameson A. Re-engineering the design process through computation. *Journal of Aircraft*, Vol. 36, No. 1, pp. 36–50, January-February 1999.
- [2] Jameson A., Martinelli L., Alonso J.J., Vassberg J.C., and Reuther J. Simulation based aerodynamic design. *IEEE Aerospace Conference*. Big Sky, Montana, 18-25 March 2000.
- [3] Sobieszczanski-Sobieski J. and Haftka R.T. *Multidisciplinary aerospace design optimization: survey of recent developments*. AIAA Paper 96-0711, 1996.
- [4] Jameson A. Aerodynamic design via control theory. *Journal of Scientific Computing*, Vol. 3, No. 3, pp. 233–260, 1988.
- [5] Ahlstrom E., Gregg R., Vassberg J.C., and Jameson A. *G-force: the design of an unlimited class Reno racer*. AIAA Paper 2000-4341, 2000.
- [6] Jameson A. and Vassberg J.C. *Computational Fluid Dynamics for aerodynamic design: its current and future impact*. AIAA Paper 2001-0538, 2001.
- [7] Soemarwoto B.I., Laban M., Jameson A., Martins A.L., and Oskam B. *Adaptive aerodynamic optimization of regional jet aircraft*. AIAA Paper 2002-0260, 2002.
- [8] Alonso J.J., Kroo I., and Jameson A. *Advanced algorithms for design and optimization of quiet supersonic platform*. AIAA Paper 2002-0144, 2002.
- [9] Cliff S.E., Reuther J.J., Saunders D.A., and Hicks R.M. Single-point and multipoint aerodynamic shape optimization of high-speed civil transport. *Journal of Aircraft*, Vol. 38, No. 6, pp. 997–1005, November-December 2001.
- [10] Campobasso M.S., Duta M.C., and Giles M.B. Adjoint methods for turbomachinery design. *ISOABE Conference*. ISABE-2001-1055, 2001.
- [11] Giles M.B., Duta M.C., and Muller J.-D. *Adjoint code developments using the exact discrete approach*. AIAA Paper 2001-2596, 2001.
- [12] Cross M.R. Aerodynamic design using the Euler adjoint approach. *22nd International Congress of Aeronautical Sciences*. ICAS-2000-2.7.1., Harrogate, UK, 27 August-1 September 2000.

- [13] Szmelter J. *Aerodynamic wing optimisation*. AIAA Paper 99-0550, 1999.
- [14] Kleinveld S., Roge G., Dinh Q., and Daumas L. CAD-based wing body optimisation. *CEAS Aerospace Dynamics Research Conference*. Cambridge, UK, 10-12 June 2002.
- [15] Beux F., Dervieux A., Leclercq M.-P., and Stoufflet B. Techniques de contrôle optimal pour l'optimisation de forme en aérodynamique avec calcul exact du gradient. *Revue Scientifique et Technique de la Défense*, Vol. 25, pp. 159–166, 1994. In French.
- [16] Cormery M. Adjoint operator approach to aerodynamic shape optimisation for 3D transonic flows. *Proceedings of the 4th European Computational Fluid Dynamics Conference ECCOMAS 98*, Vol. 1, pp. 610–616. Athens, Greece, 7-11 September 1998.
- [17] Soemarwoto B.I. and Labrujere T.E. Airfoil design and optimization methods: recent progress at NLR. *International Journal for Numerical Methods in Fluids*, Vol. 30, No. 2, pp. 217–228, May 1999.
- [18] Weinerfelt P. and Enoksson O. Aerodynamic shape optimization and parallel computing applied to industrial problems. *Parallel Computational Fluid Dynamics 2000. Trends and Applications*. Trondheim, Norway, 22-25 May 2000.
- [19] Kim H.-J., Kim C., Rho O.-H., and Lee K.D. *Aerodynamic sensitivity analysis for Navier-Stokes equations*. AIAA Paper 99-0402, 1999.
- [20] Kroo I. *Aerodynamic concepts for future aircraft*. AIAA Paper 99-3524, 1999.
- [21] Pambagjo T.E., Nakahashi K., Obayashi S., and Matsushima K. *Aerodynamic design of a medium size Blended-Wing-Body airplane*. AIAA Paper 2001-0129, 2001.
- [22] Hicks R.M., Murman E.M., and Vanderplaats G.N. *An assessment of airfoil design by numerical optimization*. NASA TM X-3092, 1974.
- [23] Hicks R.M. and Vanderplaats G.N. *Application of numerical optimization to the design of low-speed airfoils*. NASA TM X-3213, 1975.
- [24] Vanderplaats G.N., Hicks R.M., and Murman E.M. Application of numerical optimization techniques to airfoil design. *Aerodynamic analyses requiring advanced computers*, Vol. II, pp. 749–768, Paper 25. NASA SP-347, Hampton, Virginia, 4-6 March 1975.
- [25] Beux F. and Dervieux A. Exact-gradient shape optimization of a 2-D Euler flow. *Finite Elements in Analysis and Design*, Vol. 12, pp. 281–302, 1992.

- [26] Hu H. Development of a sensitivity derivative version of the FPX CFD code via automatic differentiation. *Advances in Engineering Software*, Vol. 30, No. 4, pp. 273–279, April 1999.
- [27] Kim S., Alonso J.J., and Jameson A. *A gradient accuracy study for the adjoint-based Navier-Stokes design method*. AIAA Paper 99-0299, 1999.
- [28] Hou G.J.-W., Maraju V., Taylor A.C., Korivi V.M., and Newman P.A. Transonic turbulent airfoil design optimization with automatic differentiation in incremental iterative forms. *12th AIAA Computational Fluid Dynamics Conference*, Vol. 1, pp. 512–526, AIAA-95-1692-CP. San Diego, CA, 19-22 June 1995.
- [29] Green L.L., Newman P.A., and Haigler K.J. Sensitivity derivatives for advanced CFD algorithm and viscous modeling parameters via Automatic Differentiation. *Journal of Computational Physics*, Vol. 125, No. 2, pp. 313–324, May 1996.
- [30] Destarac D. and Reneaux J. Numerical optimization applied to transport aircraft aerodynamics. *La Recherche Aéronautique*, Vol. 2, pp. 39–55, English Edition, 1993.
- [31] Reneaux J. Numerical optimization method for airfoil design. *La Recherche Aéronautique*, Vol. 5, pp. 1–19, English Edition, 1984.
- [32] Eyi S., Chand K.K., Lee K.D., Rogers S.E., and Kwak D. *Multi-element high-lift design using the Navier-Stokes equations*. AIAA Paper 96-1943, 1996.
- [33] Eyi S., Lee K.D., Rogers S.E., and Kwak D. High-lift optimization using Navier-Stokes equations. *Journal of Aircraft*, Vol. 33, No. 3, pp. 499–504, May-June 1996.
- [34] Eyi S. and Lee K.D. Inverse airfoil design using the Navier-Stokes equations. *Engineering Optimization*, Vol. 28, No. 4, pp. 245–262, 1997.
- [35] Held C. and Dervieux A. One-shot airfoil optimization without adjoint. *Computers and Fluids*, Vol. 31, No. 8, pp. 1015–1049, August 2002.
- [36] Kaplan B. and Eyi S. *Inverse design of compressor cascades*. AIAA Paper 2001-0387, 2001.
- [37] Shebalin J.-P. and Tiwari S.N. *Noz-Op-2D: a CFD-based optimization system for axially symmetric rocket nozzles*. AIAA Paper 2001-1062, 2001.
- [38] Cliff S.E., Thomas S.D., Baker T.J., Jameson A., and Hicks R.M. Aerodynamic shape optimization using unstructured grid methods. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5550. Atlanta, GA, 4-6 September 2002.

- [39] Appel J.R. and Gunzburger M.D. Sensitivity calculation in flows with discontinuities. *14th AIAA Applied Aerodynamics Conference*, Vol. 2, pp. 721–729, AIAA–96–2471–CP. New Orleans, LA, 17–20 June 1996.
- [40] Baysal O. and Eleshaky M.E. *Aerodynamic design optimization using sensitivity analysis and computational fluid dynamics*. AIAA Paper 91-0471, 1991. Also available in *AIAA Journal*, Vol. 30, No. 3, pp. 718–725, March 1992.
- [41] Ajmani K. and Taylor A.C. *Discrete sensitivity derivatives of the Navier-Stokes equations with a parallel Krylov solver*. AIAA Paper 94-0091, 1994.
- [42] Korivi V.M., Newman P.A., and Taylor A.C. Aerodynamic optimization studies using a 3-D supersonic Euler code with efficient calculation of sensitivity derivatives. *5th AIAA/NASA/USAF/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 170–194, AIAA–94–4270–CP. Panama City Beach, FL, 7–9 September 1994. Also available in *Journal of Aircraft*, Vol. 35, No. 3, pp. 405–411, May–June 1998.
- [43] Kim C.S., Kim C., and Rho O.-H. *Aerodynamic sensitivity analysis for turbulent flows on Chimera overlaid grids*. AIAA Paper 2001-0267, 2001. Also available in *AIAA Journal*, Vol. 39, No. 5, pp. 838–845, May 2001.
- [44] Nadarajah S. and Jameson A. *A comparison of the continuous and discrete adjoint approach to automatic aerodynamic optimization*. AIAA Paper 2000-0667, 2000.
- [45] Kim S., Alonso J.J., and Jameson A. *Design optimization of high-lift configurations using a viscous continuous adjoint method*. AIAA Paper 2002-0844, 2002.
- [46] Kuruvila G., Narducci R., and Agrawal S. *Development and application of TLNS3D-adjoint: a practical tool for aerodynamic shape optimization*. AIAA Paper 2001-2400, 2001.
- [47] Nemec M. and Zingg D.W. *Towards efficient aerodynamic shape optimization based on the Navier-Stokes equations*. AIAA Paper 2001-2532, 2001. Also available in *AIAA Journal*, Vol. 40, No. 6, pp. 1146–1154, June 2002.
- [48] Nielsen E.J. and Anderson W.K. Aerodynamic design optimization on unstructured meshes using the Navier-Stokes equations. *AIAA Journal*, Vol. 37, No. 11, pp. 1411–1419, November 1999.
- [49] Korivi V.M., Taylor A.C., Newman P.A., Hou G.J.-W., and Jones H.E. An approximately factored incremental strategy for calculating consistent discrete aerodynamic sensitivity derivatives. *4th AIAA/USAF/NASA/OAI symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 465–478, AIAA–92–4746–CP. Cleveland, OH, 21–23 September 1992. Also available in *Journal of Computational Physics*, Vol. 113, pp. 336–346, 1994.

- [50] Carle A., Green L.L., Bischof C.H., and Newman P.A. *Applications of automatic differentiation in CFD*. AIAA Paper 94-2197, 1994.
- [51] Sherman L.L., Taylor A.C., Green L.L., Newman P.A., Hou G.J.-W., and Korivi V.M. First- and second-order aerodynamic sensitivity derivatives via automatic differentiation with incremental iterative methods. *5th AIAA/NASA/USAF/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 87–120, AIAA-94-4262-CP. Panama City Beach, FL, 7-9 September 1994. Also available in *Journal of Computational Physics*, Vol. 129, No. 2, pp. 307-331, 1996.
- [52] Anderson W.K., Newman J.C., Whitfield D.L., and Nielsen E.J. Sensitivity analysis for Navier-Stokes equations on unstructured meshes using complex variables. *AIAA Journal*, Vol. 39, No. 1, pp. 56–63, January 2001.
- [53] Svenningsen K.H., Madsen J.I., Hassing N.H., and Pauker W.H.G. Optimization of flow geometries applying quasianalytical sensitivity analysis. *Applied Mathematical Modelling*, Vol. 20, No. 3, pp. 214–224, March 1996.
- [54] Baysal O. and Eleshaky M.E. Aerodynamic sensitivity analysis methods for the compressible Euler equations. *Journal of Fluids Engineering*, Vol. 13, No. 4, pp. 681–688, December 1991.
- [55] Carle A., Fagan M., and Green L.L. Preliminary results from the application of automated adjoint code generation to CFL3D. *7th AIAA/USAF/NASA/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 2, pp. 807–817, AIAA-98-4807. St Louis, Missouri, 2-4 September 1998.
- [56] Nadarajah S. and Jameson A. *Studies of the continuous and discrete adjoint approaches to viscous automatic aerodynamic shape optimization*. AIAA Paper 2001-2530, 2001.
- [57] Martins J.R.R.A., Alonso J.J., and Reuther J.J. Complete configuration aero-structural optimization using a coupled sensitivity analysis method. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5402. Atlanta, GA, 4-6 September 2002.
- [58] Giles M.B. and Pierce N.A. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, Vol. 65, No. 3/4, pp. 393–415, 2000.
- [59] Newman J.C., Anderson W.K., and Whitfield D.L. *Multidisciplinary sensitivity derivatives using complex variables*. MSSU-COE-ERC-98-08, Mississippi State University, 1998.
- [60] Vatsa V.N. Computation of sensitivity derivatives of Navier-Stokes equations using complex variables. *5th symposium on the large-scale analysis and design and ISE*. Williamsburg, Virginia, 12-15 October 1999.

- [61] Barthelemy J.-F.M. and Hall L.E. Automatic differentiation as a tool in engineering design. *4th AIAA/USAF/NASA/OAI symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 424–432, AIAA-92-4743-CP. Cleveland, OH, 21-23 September 1992.
- [62] Mohammadi B. and Pironneau O. *Applied shape optimization for fluids*. Numerical mathematics and scientific computation, Oxford Science Publications, 2001.
- [63] Bischof C.H. and Griewank A. ADIFOR: a FORTRAN system for portable automatic differentiation. *4th AIAA/USAF/NASA/OAI symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 433–441, AIAA-92-4744-CP. Cleveland, OH, 21-23 September 1992.
- [64] *ADIFOR 2.0 Automatic Differentiation of Fortran*. Web page, available at <http://www.cs.rice.edu/~adifor/>. Last accessed on 03/09/2002.
- [65] Unger E.R. and Hall L.E. The use of automatic differentiation in an aircraft design problem. *5th AIAA/NASA/USAF/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 64–72, AIAA-94-4260-CP. Panama City Beach, FL, 7-9 September 1994.
- [66] Bischof C.H., Corliss G., Green L.L., Griewank A., Haigler K.J., and Newman P.A. Automatic differentiation of advanced CFD codes for multidisciplinary design. *Computing Systems in Engineering*, Vol. 3, No. 6, pp. 625–637, December 1992.
- [67] Bischof C.H., Green L.L., Haigler K.J., and Knauff T.L. Parallel calculation of sensitivity derivatives for aircraft design using automatic differentiation. *5th AIAA/NASA/USAF/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 73–86, AIAA-94-4261-CP. Panama City Beach, FL, 7-9 September 1994.
- [68] Taylor A.C., Oloso A.O., and J.C. Newman. CFL3D.ADII (version 2.0): an efficient, accurate, general-purpose code for flow shape-sensitivity analysis. *15th AIAA Applied Aerodynamics Conference*, Vol. 1, pp. 12–16, AIAA-97-2204. Atlanta, GA, 23-25 June 1997.
- [69] Oloso A.O. and Taylor A.C. Aerodynamic shape-sensitivity analysis and design optimization on the IBM-SP2. *15th AIAA Applied Aerodynamics Conference*, Vol. 1, pp. 481–488, AIAA-97-2273. Atlanta, GA, 23-25 June 1997.
- [70] Newman J.C., Hou G.J.-W., and Taylor A.C. *Observations regarding use of advanced CFD analysis, sensitivity analysis, and design codes in MDO*. NASA Contractor Report 198293, ICASE Report No 96-16, 1996.

- [71] Gumbert C.R., Hou G.J.-W., and Newman P.A. *Simultaneous Aerodynamic Analysis and Design Optimization (SAADO) for a 3-D rigid wing*. AIAA Paper 99-3296, 1999.
- [72] Gumbert C.R., Hou G.J.-W., and Newman P.A. *Simultaneous Aerodynamic Analysis and Design Optimization (SAADO) for a 3-D flexible wing*. AIAA Paper 2001-1107, 2001.
- [73] Gumbert C.R., Hou G.J.-W., and Newman P.A. *Simultaneous Aerodynamic and Structural Design Optimization (SASDO) for a 3-D wing*. AIAA Paper 2001-2527, 2001.
- [74] Newman P.A., Hou G.J.-W., Jones H.E., Taylor A.C., and Korivi V.M. Observations on computational methodologies for uses in large-scale, gradient-based, multidisciplinary design. *4th AIAA/USAF/NASA/OAI symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 531–542, AIAA-92-4753-CP. Cleveland, OH, 21-23 September 1992.
- [75] Faure C. *Odyssée index page*. Web page, available at <http://www-sop.inria.fr/safir/SAM/Odysee/odysee.html>, 1999. Last accessed on 03/09/2002.
- [76] Courty F., Dervieux A., Koobus B., and Hascoet L. *Application of a reverse symbolic differentiator to the low-storage assembly of an adjoint state equation. Contribution partner INRIA to Task3.1. Contribution to Deliverable D3-18*. AEROSHAPE Report, INRIA, France, 2001.
- [77] Taylor A.C., Green L.L., Newman P.A., and Putko M.M. *Some advanced concepts in discrete aerodynamic sensitivity analysis*. AIAA Paper 2001-2529, 2001.
- [78] Sobieszczanski-Sobieski J. The case for aerodynamic sensitivity analysis. *Sensitivity analysis in engineering*, NASA Conference Publication 2457, pp. 77–96. Langley Research Center, Hampton, Virginia, 25-26 September 1986.
- [79] Elliott J. and Peraire J. Practical 3D aerodynamic design and optimization using unstructured meshes. *6th AIAA/NASA/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 2, pp. 1819–1828, AIAA-96-4170-CP. Bellevue, WA, 4-6 September 1996. Also available in AIAA Journal, Vol. 35, No. 9, pp. 1479-1485, September 1997.
- [80] J.C. Newman. *Integrated multidisciplinary design optimization using discrete sensitivity analysis for geometrically complex aeroelastic configurations*. Ph.D. thesis, Virginia Polytechnic Institute and State University, 1997.
- [81] Newman J.C. and Taylor A.C. *Three-dimensional shape sensitivity analysis and design optimization using the Euler equations on unstructured grids*. AIAA-96-2464-CP, 1996.

- [82] Newman J.C., Taylor A.C., and Burgreen G.W. An unstructured grid approach to sensitivity analysis and shape optimization using the Euler equations. *12th AIAA Computational Fluid Dynamics Conference*, Vol. 1, pp. 1–10, AIAA–95–1646–CP. San Diego, CA, 19–22 June 1995.
- [83] Kim H.-J., Sasaki D., Obayashi S., and Nakahashi K. Aerodynamic optimization of supersonic transport wing using unstructured adjoint method. *AIAA Journal*, Vol. 39, No. 6, pp. 1011–1020, June 2001.
- [84] El-banna H.M. and Carlson L.A. Determination of aerodynamic sensitivity coefficients based on the transonic small perturbation formulation. *Journal of Aircraft*, Vol. 27, No. 6, pp. 507–515, June 1990.
- [85] El-banna H.M. and Carlson L.A. Aerodynamic sensitivity coefficients using the three-dimensional full potential equation. *Journal of Aircraft*, Vol. 31, No. 5, pp. 1071–1077, September–October 1994.
- [86] Burgreen G.W. and Baysal O. Three-dimensional aerodynamic shape optimization of supersonic delta wings. *5th AIAA/NASA/USAF/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 195–205, AIAA–94–4271–CP. Panama City Beach, FL, 7–9 September 1994.
- [87] Burgreen G.W. and Baysal O. Aerodynamic shape optimization using preconditioned conjugate gradient methods. *11th AIAA Computational Fluid Dynamics Conference*, Vol. 1, pp. 278–288, AIAA–93–3322–CP. Orlando, FL, 6–9 July 1993.
- [88] Pandya M.J. and Baysal O. Gradient-based aerodynamic shape optimization using alternating direction implicit method. *Journal of Aircraft*, Vol. 34, No. 3, pp. 346–352, May–June 1997.
- [89] Taylor A.C., Hou G.J.-W., and Korivi V.M. Methodology for calculating aerodynamic sensitivity derivatives. *AIAA Journal*, Vol. 30, No. 10, pp. 2411–2419, October 1992.
- [90] Hou G.J.-W., Taylor A.C., and Korivi V.M. Discrete shape sensitivity equations for aerodynamic problems. *International Journal for Numerical Methods in Engineering*, Vol. 37, No. 13, pp. 2251–2266, July 1994.
- [91] Eleshaky M.E. *A computational aerodynamic design optimization method using sensitivity analysis*. Ph.D. thesis, Old Dominion University, 1992.
- [92] Eleshaky M.E. and Baysal O. Aerodynamic shape optimization via sensitivity analysis on decomposed computational domains. *4th AIAA/USAF/NASA/OAI symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 98–109, AIAA–92–4698–CP. Cleveland, OH, 21–23 September 1992.

- [93] Eleshaky M.E. and Baysal O. *Design of a 3-D nacelle near flat plate wing using multiblock sensitivity analysis (ADOS)*. AIAA Paper 94-0160, 1994.
- [94] Eleshaky M.E. and Baysal O. Preconditioned domain decomposition scheme for three-dimensional aerodynamic sensitivity analysis. *AIAA Journal*, Vol. 32, No. 12, pp. 2489–2491, December 1994.
- [95] Lacasse J.M. and Baysal O. Shape optimization of single- and two-element airfoils on multiblock grids. *5th AIAA/NASA/USAF/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 223–231, AIAA-94-4273-CP. Panama City Beach, FL, 7-9 September 1994.
- [96] Taylor A.C., Hou G.J.-W., and Korivi V.M. *Sensitivity analysis, approximate analysis, and design optimization for internal and external viscous flows*. AIAA Paper 91-3083, 1991.
- [97] Arian E. and Ta'asan S. Analysis of the Hessian for aerodynamic optimization: inviscid flow. *Computers and Fluids*, Vol. 28, pp. 853–877, 1999.
- [98] Hou G.J.-W. and Sheen J. Numerical methods for second-order shape sensitivity analysis with applications to heat conduction problems. *International Journal for Numerical Methods in Engineering*, Vol. 36, No. 3, pp. 417–435, February 1993.
- [99] Melvin R.G., Huffman W.P., Young D.P., Johnson F.T., Hilmes C.L., and Bieterman M.B. Recent progress in aerodynamic design optimization. *International Journal for Numerical Methods in Fluids*, Vol. 30, No. 2, pp. 205–216, May 1999.
- [100] Arian E. and Vatsa V.N. *A preconditioning method for shape optimization governed by the Euler equations*. NASA/CR-1998-206926, ICASE Report No 98-14, 1998.
- [101] Knill D.L., Giunta A.A., Baker C.A., Grossman B., Mason W.H., Haftka R.T., and Watson L.T. Response surface models combining linear and Euler aerodynamics for supersonic transport design. *Journal of Aircraft*, Vol. 36, No. 1, pp. 75–86, January-February 1999.
- [102] Narducci R., Grossman B., Valorani M., Dadone A., and Haftka R.T. Optimization methods for non-smooth or noisy objective functions in fluid design problems. *12th AIAA Computational Fluid Dynamics Conference*, Vol. 1, pp. 21–32, AIAA-95-1648-CP. San Diego, CA, 19-22 June 1995.
- [103] Chung H.-S. and Alonso J.J. *Using gradients to construct response surface models for high-dimensional design optimization problems*. AIAA Paper 2001-0922, 2001.
- [104] Toomer C.A., Topliss M.E., and Hills D.P. Aerodynamic optimization using analytic descriptions of the design space. *Journal of Aircraft*, Vol. 35, No. 6, pp. 882–890, November-December 1998.

- [105] Topliss M.E., Toomer C.A., and Hills D.P. Rapid design space approximation for two-dimensional transonic aerofoil design. *Journal of Aircraft*, Vol. 33, No. 6, pp. 1101–1108, November-December 1996.
- [106] Papila N., Shyy W., Griffin L.W., and Dorney D.J. *Shape optimization of supersonic turbines using response surface and neural network methods*. AIAA Paper 2001-1065, 2001.
- [107] Carrier G., Knight D., Rasheed K., and Montazel X. *Multi-criteria design optimization of two-dimensional supersonic inlet*. AIAA Paper 2001-1064, 2001.
- [108] LingJun C. and Hai-Song A. *Conceptual/preliminary aircraft design using genetic algorithm and fuzzy mathematics*. AIAA Paper 99-0113, 1999.
- [109] LingJun C. and Qi-Peng C. *Genetic optimization for three-dimensional wing design method*. AIAA Paper 99-0400, 1999.
- [110] Oyama A., Obayashi S., Nakahashi K., and Nakamura T. Euler/Navier-Stokes optimization of supersonic wing design based on evolutionary algorithm. *AIAA Journal*, Vol. 37, No. 10, pp. 1327–1328, October 1999.
- [111] Wright W.A. and Holden C.M.E. Design optimisation of a simple 2-D aerofoil using stochastic search methods. *Conference proceedings adaptive computing in design and manufacture*. April 1998.
- [112] Vanderplaats G.N. *Numerical optimization techniques for engineering design: with applications*. McGraw-Hill Book Company, 1984.
- [113] Gill P.E., Murray W., and Wright M.H. *Practical optimization*. Academic Press, London, 1981.
- [114] Fletcher R. and Reeves C.M. Function minimization by conjugate gradients. *Computer Journal*, Vol. 7, No. 2, pp. 149–154, 1964.
- [115] Davidson W.C. *Variable metric methods for minimization*. A.E.C. Research and Development Report ANL-5990, Argonne National Laboratory, Argonne, Illinois, 1959.
- [116] Fletcher R. and Powell M.J.D. A rapidly convergent descent method for minimization. *Computer Journal*, Vol. 6, No. 2, pp. 163–168, 1963.
- [117] Broyden C.G. The convergence of a class of double-rank minimization algorithms. *Journal of the Institute of Mathematics and its Applications*, Vol. 6, pp. 76–90 and 222–231, 1970.
- [118] Fletcher R. A new approach to variable metric algorithms. *Computer Journal*, Vol. 13, No. 3, pp. 317–322, 1970.

- [119] Goldfarb D. A family of variable metric methods derived by variational means. *Mathematics of Computation*, Vol. 24, No. 109, pp. 23–26, January 1970.
- [120] Shanno D.F. Conditioning of quasi-Newton methods for function minimization. *Mathematics of Computation*, Vol. 24, No. 111, pp. 647–656, July 1970.
- [121] NAG (Numerical Algorithm Group). *NAG Fortran Library*. Web page, available at <http://www.nag.com/numeric/FL/FLdescription.asp>. Last accessed on 09/09/2002.
- [122] NAG (Numerical Algorithm Group). *E04UCF-NAG Fortran library routine documentation*.
- [123] AEMDesign. *What is FSQP*. Web page, available at <http://64.238.116.66/aemdesign/FSQPframe.htm>. Last accessed on 09/09/2002.
- [124] Zhou J.L., Tits A.L., and Lawrence C.T. *User's guide for FFSQP Version 3.7: a FORTRAN code for solving constrained nonlinear (Minimax) optimization problems, generating iterates satisfying all inequality and linear constraints*. Systems Research Center TR-92-107r2, University of Maryland, 1997.
- [125] Lambert P.-A. *Optimisation de formes en aérodynamique. Application à la conception des nacelles de moteurs civils*. Ph.D. thesis, Ecole Centrale Paris, France, 1995. In French.
- [126] Powell M.J.D. A fast algorithm for nonlinearly constrained optimization calculations. *Numerical Analysis, Dundee, 1977, Lecture Notes in Mathematics 630*, edited by Watson G.A., pp. 144–157. Springer Verlag, Berlin, 1978.
- [127] Lewis R.M. A trust region framework for managing approximation models in engineering optimization. *6th AIAA/NASA/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 2, pp. 1053–1055, AIAA-96-4101-CP. Bellevue, WA, 4-6 September 1996.
- [128] Alexandrov N.M. Robustness properties of a trust region framework for managing approximations in engineering optimization. *6th AIAA/NASA/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 2, pp. 1056–1059, AIAA-96-4102-CP. Bellevue, WA, 4-6 September 1996.
- [129] Alexandrov N.M., Dennis J.E., Lewis R.M., and Torczon V. *A trust region framework for managing the use of approximation models in optimization*. NASA/CR-201745, ICASE Report No 97-50, 1997.
- [130] Alexandrov N.M., Lewis R.M., Gumbert C.R., Green L.L., and Newman P.A. *Optimization with variable-fidelity models applied to wing design*. AIAA Paper 2000-0841, 2000.

- [131] Alexandrov N.M., Nielsen E.J., Lewis R.M., and Anderson W.K. First-order model management with variable-fidelity physics applied to multi-element airfoil optimization. *8th AIAA/USAF/NASA/ISSMO symposium on multidisciplinary analysis and optimization*. AIAA-2000-4886, Long Beach, CA, 6-8 September 2000.
- [132] Alexandrov N.M., Lewis R.M., Gumbert C.R., Green L.L., and Newman P.A. Approximation and model management in aerodynamic optimization with variable-fidelity models. *Journal of Aircraft*, Vol. 38, No. 6, pp. 1093–1101, November-December 2001.
- [133] Marduel X., Tribes C., and Trépanier J.-Y. Optimization using variable fidelity solvers: exploration of an approximation management framework for aerodynamic shape optimization. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5595. Atlanta, GA, 4-6 September 2002.
- [134] Samareh J.A. Survey of shape parameterization techniques for high-fidelity multidisciplinary shape optimization. *AIAA Journal*, Vol. 39, No. 5, pp. 877–884, May 2001.
- [135] Samareh J.A. Status and future of geometry modeling and grid generation for design and optimization. *Journal of Aircraft*, Vol. 36, No. 1, pp. 97–104, January-February 1999.
- [136] Samareh J.A. Novel multidisciplinary shape parameterization approach. *Journal of Aircraft*, Vol. 38, No. 6, pp. 1015–1024, November-December 2001. Also available as AIAA Paper 2000-4911.
- [137] Jameson A., Alonso J.J., Reuther J., Martinelli L., and Vassberg J.C. *Aerodynamic shape optimization techniques based on control theory*. AIAA Paper 98-2538, 1998.
- [138] Reuther J. and Jameson A. *Aerodynamic shape optimization of wing and wing-body configurations using control theory*. AIAA Paper 95-0123, 1995.
- [139] Reuther J., Jameson A., Alonso J.J., Rimlinger M.J., and Saunders D. *Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers*. AIAA Paper 97-0103, 1997. Also available in *AIAA Journal*, Vol. 36, No. 1, pp. 51-74, January-February 1999.
- [140] Jameson A. Optimum aerodynamic design using CFD and control theory. *12th AIAA Computational Fluid Dynamics Conference*, Vol. 2, pp. 926–949, AIAA-95-1729-CP. San Diego, CA, 19-22 June 1995.
- [141] Nadarajah S., Jameson A., and Alonso J.J. Sonic boom reduction using an adjoint method for wing-body configurations in supersonic flow. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5547. Atlanta, GA, 4-6 September 2002.

- [142] Dreyer L.L. and Martinelli L. *Hydrodynamic shape optimization of propulsor configurations using a continuous adjoint approach*. AIAA Paper 2001-2580, 2001.
- [143] Hicks R.M. and Henne P.A. Wing design by numerical optimization. *Journal of Aircraft*, Vol. 15, No. 7, pp. 407–412, July 1978.
- [144] Eyi S. and Lee K.D. *Effects of sensitivity analysis on airfoil design*. AIAA Paper 98-0909, 1998.
- [145] Sung C. and Kwon J.H. *An efficient aerodynamic design method using a tightly coupled algorithm*. AIAA Paper 2000-0783, 2000.
- [146] Sung C. and Kwon J.H. Accurate aerodynamic sensitivity analysis using adjoint equations. *AIAA Journal*, Vol. 38, No. 2, pp. 243–250, February 2000.
- [147] Sung C. and Kwon J.H. *Aerodynamic design optimization using the Navier-Stokes and adjoint equations*. AIAA Paper 2001-0266, 2001.
- [148] Jameson A. and Reuther J. Control theory based airfoil design using the Euler equations. *5th AIAA/NASA/USAF/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 206–222, AIAA-94-4272-CP. Panama City Beach, FL, 7-9 September 1994.
- [149] Reuther J. and Jameson A. *Control theory based airfoil design for potential flow and a finite volume discretization*. AIAA Paper 94-0499, 1994.
- [150] Reuther J., Jameson A., Farmer J., Martinelli L., and Saunders D. *Aerodynamic shape optimization of complex aircraft configurations via an adjoint formulation*. AIAA Paper 96-0094, 1996.
- [151] Ta'asan S., Kuruvila G., and Salas M.D. *Aerodynamic design and optimization in one shot*. AIAA Paper 92-0025, 1992.
- [152] Baysal O. and Ghayour K. Continuous adjoint sensitivities for general cost functionals on unstructured meshes in aerodynamic shape optimization. *7th AIAA/USAF/NASA/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 3, pp. 1483–1491, AIAA-98-4904. St Louis, Missouri, 2-4 September 1998. Also available in *AIAA Journal*, Vol. 39, No. 1, pp. 48-55, January 2001.
- [153] Burgreen G.W. *Three-dimensional aerodynamic shape optimization using discrete sensitivity analysis*. Ph.D. thesis, Old Dominion University, 1994.
- [154] Burgreen G.W. and Baysal O. Three-dimensional aerodynamic shape optimization using discrete sensitivity analysis. *AIAA Journal*, Vol. 34, No. 9, pp. 1761–1770, September 1996.

- [155] Burgreen G.W., Baysal O., and Eleshaky M.E. Improving the efficiency of aerodynamic shape optimization procedures. *4th AIAA/USAF/NASA/OAI symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 87–97, AIAA–92–4697–CP. Cleveland, OH, 21–23 September 1992.
- [156] Anderson W.K. and Bonhaus D.L. *Aerodynamic design on unstructured grids for turbulent flows*. NASA Technical Memorandum 112867, 1997.
- [157] Anderson W.K. and Venkatakrishnan V. *Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation*. AIAA Paper 97-0643, 1997.
- [158] Yong H., Zuobin C., and Gang L. Numerical optimization design for transonic airfoils. *4th Asian Computational Fluid Dynamics Conference*, pp. 595–602. Mi-nyang, China, 18–22 September 2000.
- [159] Hiernaux S. and Essers J.-A. Aerodynamic optimization using Navier-Stokes equations and optimal control theory. *A collection of the 14th AIAA Computational Fluid Dynamics Conference Technical Papers*, Vol. 1, pp. 419–428, AIAA–99–3297. Norfolk, Virginia, 26 June–1 July 1999.
- [160] Hiernaux S. and Essers J.-A. An optimal control theory based algorithm to solve 2D aerodynamic shape optimisation problems for inviscid and viscous flows. *RTO AVT Symposium on aerodynamic design and optimisation of flight vehicles in a concurrent multi-disciplinary environment*. RTO-MP-035-18, Ottawa, Canada, 18–21 October 1999.
- [161] Duvigneau R. and Visonneau M. *Shape optimization of incompressible and turbulent flows using the simplex method*. AIAA Paper 2001-2533, 2001.
- [162] Nielsen E.J. and Anderson W.K. *Recent improvements in aerodynamic design optimization on unstructured meshes*. AIAA Paper 2001-0596, 2001. Also available in AIAA Journal, Vol. 40, No. 6, pp. 1155–1163, June 2002.
- [163] Soto O. and Lohner R. *General methodologies for incompressible flow design problems*. AIAA Paper 2001-1061, 2001.
- [164] Nemec M. and Zingg D.W. Multi-point and multi-objective aerodynamic shape optimization. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5548. Atlanta, GA, 4–6 September 2002.
- [165] Reuther J., Alonso J.J., Vassberg J.C, Jameson A., and Martinelli L. *An efficient multiblock method for aerodynamic analysis and design on distributed memory systems*. AIAA-97-1893, 1997.
- [166] Jones W.T. and Samareh J.A. A grid generation system for multi-disciplinary design optimization. *12th AIAA Computational Fluid Dynamics Conference*, Vol. 1, pp. 474–482, AIAA–95–1689–CP. San Diego, CA, 19–22 June 1995.

- [167] Schmitt V. and Charpin F. Pressure distributions on the ONERA-M6-wing at transonic Mach numbers. *Experimental data base for computer program assessment*, AGARD-AR-138. AGARD, Neuilly-sur-Seine, France, 1979.
- [168] Pagaldipti N. and Chattopadhyay A. A discrete semi-analytical procedure for aerodynamic sensitivity analysis including grid sensitivity. *5th AIAA/NASA/USAF/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 161–169, AIAA-94-4268-CP. Panama City Beach, FL, 7-9 September 1994.
- [169] Perigo D. and Qin N. *A numerical investigation of the flows in and around plug nozzle configuration*. AIAA Paper 2001-2417, 2001.
- [170] D. Perigo. *A numerical investigation of the flows in and around clustered module plug nozzles*. Eng.D. thesis, College of Aeronautics, Cranfield University, UK, 2001.
- [171] Le Moigne A. and Qin N. A discrete adjoint method for aerodynamic sensitivities for Navier-Stokes flows. *CEAS Aerospace Dynamics Research Conference*. Cambridge, UK, 10-12 June 2002.
- [172] Qin N., Vavalle A., and Le Moigne A. A study of spanwise lift distribution for Blended Wing-Body aircraft. *CEAS Aerospace Dynamics Research Conference*. Cambridge, UK, 10-12 June 2002.
- [173] Osher S. and Solomon F. Upwind difference schemes for hyperbolic systems of conservation laws. *Mathematics of Computation*, Vol. 38, No. 158, pp. 339–374, April 1982.
- [174] Chakravarthy S.R. and Osher S. Numerical experiments with the Osher upwind scheme for the Euler equations. *AIAA Journal*, Vol. 21, No. 9, pp. 1241–1248, September 1983.
- [175] Van Leer B. Towards the ultimate conservative difference scheme. III - Upstream-centered finite-difference schemes for ideal compressible flow. IV - A new approach to numerical convection. *Journal of Computational Physics*, Vol. 23, pp. 263–299, March 1977.
- [176] Van Leer B. Towards the ultimate conservative difference scheme. V - A second-order sequel to Godunov's method (for ideal compressible flow). *Journal of Computational Physics*, Vol. 32, pp. 101–136, July 1979.
- [177] Baldwin B.S. and Lomax H. *Thin layer approximation and algebraic model for separated turbulent flows*. AIAA Paper 78-257, 1978.
- [178] Hirsch C. *Numerical computation of internal and external flows. Volume 2: Computational methods for inviscid and viscous flows*. John Wiley & Sons, 1988.

- [179] Hirsch C. *Numerical computation of internal and external flows. Volume 1: Fundamentals of numerical discretization*. John Wiley & Sons, 1988.
- [180] Toro E.F. *Riemann solvers and numerical methods for fluid dynamics*. Springer-Verlag, Berlin, 1997.
- [181] Spekreijse S.P. *Multigrid solution of the steady Euler equations*. Ph.D. thesis, Centrum Voor Wiskunde en Informatica, Amsterdam, 1987.
- [182] Qin N. *Lecture notes on computation of high speed flows*. College of Aeronautics, Cranfield University, UK, 1999.
- [183] Shaw S.T. *Numerical study of the unsteady aerodynamics of helicopter rotor aerofoils*. Ph.D. thesis, College of Aeronautics, Cranfield University, UK, 1999.
- [184] McNeil C.Y. *Efficient upwind algorithms for solution of the Euler and Navier-Stokes equations*. Ph.D. thesis, College of Aeronautics, Cranfield University, UK, 1995.
- [185] Wilcox D.C. *Turbulence modeling for CFD*. La Canada, California: DCW Industries, Inc, 1993.
- [186] Jameson A. and Baker T.J. Solution of the Euler equations for complex configurations. *6th AIAA Computational Fluid Dynamics Conference*, pp. 293–302, AIAA-83-1929. Danvers, Massachusetts, 13-15 July 1983.
- [187] Qin N., Ludlow D.K., and Shaw S.T. A matrix-free preconditioned Newton/GMRES method for Navier-Stokes solutions. *International Journal for Numerical Methods in Fluids*, Vol. 33, No. 3, pp. 223–248, June 2000.
- [188] Shaw S.T. and Qin N. *A preconditioned implicit Krylov-subspace method for the solution of the parabolised Navier-Stokes equations*. CoA Report 9706 NFP, College of Aeronautics, Cranfield University, UK, 1997.
- [189] Slater J.W. *ONERA M6 wing*. Web page, available at <http://www.grc.nasa.gov/www/wind/valid/m6wing/m6wing.html>, 2000. Last accessed on 22/07/2002.
- [190] Lee B.J., Kim C.S., Kim C., Rho O.-H., and Lee K.D. *Parallelized design optimization for transonic wings using aerodynamic sensitivity analysis*. AIAA Paper 2002-0264, 2002.
- [191] Vatsa V.N., Sanetrik M.D., and Parlette E.B. *Development of a flexible and efficient multigrid-based multiblock flow solver*. AIAA Paper 93-0677, 1993.
- [192] Yang J.-Y., Perng Y.-C., and Yen R.-H. Implicit weighted essentially nonoscillatory schemes for the compressible Navier-Stokes equations. *AIAA Journal*, Vol. 39, No. 11, pp. 2082–2090, November 2001.

- [193] Badcock K.J. and B.E. Richards. Implicit time-stepping methods for the Navier-Stokes equations. *AIAA Journal*, Vol. 34, No. 3, pp. 555–559, March 1996.
- [194] Deese J.E. and Agarwal R.K. Navier-Stokes calculations of transonic viscous flow about wing/body configurations. *Journal of Aircraft*, Vol. 25, No. 12, pp. 1106–1112, December 1988.
- [195] Nakahashi K., Sharov D., Kano S., and Kodera M. Applications of unstructured hybrid grid method to high-Reynolds number viscous flows. *International Journal for Numerical Methods in Fluids*, Vol. 31, No. 1, pp. 97–111, 1999.
- [196] Nakahashi K., Togashi F., and Sharov D. Intergrid-boundary definition method for overset unstructured grid approach. *AIAA Journal*, Vol. 38, No. 11, pp. 2077–2084, November 2000.
- [197] Radespiel R., Rossow C., and Swanson R.C. Efficient cell-vertex multigrid scheme for the three-dimensional Navier-Stokes equations. *AIAA Journal*, Vol. 28, No. 8, pp. 1464–1472, August 1990.
- [198] Vatsa V.N. Accurate numerical solutions for transonic viscous flow over finite wings. *Journal of Aircraft*, Vol. 24, No. 6, pp. 377–385, June 1987.
- [199] Vatsa V.N. and Wedan B.W. Development of a multigrid code for 3-D Navier-Stokes equations and its application to a grid-refinement study. *Computers and Fluids*, Vol. 18, No. 4, pp. 391–403, 1990.
- [200] Müller B. and Rizzi A. Navier-Stokes solution for transonic flow over wings. *Proceedings of the Seventh GAMM-Conference on Numerical Methods in Fluid Mechanics*, edited by Deville M., Vol. 20 of *Notes on Numerical Fluid Mechanics*, pp. 247–255. Vieweg, Louvain-La-Neuve, Belgium, 9-11 September 1987.
- [201] Nadarajah S., Jameson A., and Alonso J.J. *An adjoint method for the calculation of remote sensitivities in supersonic flow*. AIAA Paper 2002-0261, 2002.
- [202] Shenoy A.R. *Optimization techniques exploiting problem structure: applications to aerodynamic design*. Ph.D. thesis, Virginia Polytechnic Institute and State University, 1997.
- [203] Pironneau O. On optimum profiles in stokes flow. *Journal of Fluid Mechanics*, Vol. 59, No. 1, pp. 117–128, 1973.
- [204] Pironneau O. On optimum design in fluid mechanics. *Journal of Fluid Mechanics*, Vol. 64, No. 1, pp. 97–110, 1974.
- [205] Meric R.A. Finite element analysis of optimal heating of a slab with temperature dependent thermal conductivity. *International Journal of Heat and Mass Transfer*, Vol. 22, pp. 1347–1353, 1979.

- [206] Meric R.A. Boundary element methods for optimization of distributed parameter systems. *International Journal for Numerical Methods in Engineering*, Vol. 20, No. 7, pp. 1291–1306, July 1984.
- [207] Koda M. Sensitivity analysis of the atmospheric diffusion equation. *Atmospheric Environment*, Vol. 16, No. 11, pp. 2595–2601, 1982.
- [208] Koda M. Optimum design in fluid mechanical distributed-parameter systems. *Large Scale Systems*, Vol. 6, No. 3, pp. 279–291, June 1984.
- [209] Koda M. Sensitivity analysis of descriptor distributed parameter systems. *International Journal of Systems Science*, Vol. 19, No. 10, pp. 2103–2114, October 1988.
- [210] Jameson A. and Alonso J.J. *Automatic aerodynamic optimization on distributed memory architectures*. AIAA Paper 96-0409, 1996.
- [211] Jameson A. *Essential elements of computational algorithms for aerodynamic analysis and design*. NASA/CR-97-206268, ICASE Report No 97-68, 1997.
- [212] Jameson A., Pierce N.A., and Martinelli L. *Optimum aerodynamic design using the Navier-Stokes equations*. AIAA Paper 97-0101, 1997.
- [213] Ibrahim A.H. and Baysal O. *Design optimization using variational methods and CFD*. AIAA Paper 94-0093, 1994.
- [214] Ibrahim A.H., Hou G.J.-W., Tiwari S.N., and Smith R.E. *Variational methods in design optimization and sensitivity analysis for two-dimensional Euler equations*. AIAA Paper 97-0102, 1997.
- [215] Ibrahim A.H., Tiwari S.N., and Smith R.E. Stability analysis and aerodynamic design optimization of Euler equations using variational methods. *International Journal for Numerical Methods in Engineering*, Vol. 44, No. 11, pp. 1709–1726, April 1999.
- [216] Qiao Z., Qin X., and Yang X. Wing design using control theory. *4th Asian Computational Fluid Dynamics Conference*, pp. 653–663. Mianyang, China, 18-22 September 2000.
- [217] Qiao Z., Qin X., and Yang X. *Wing design by solving adjoint equations*. AIAA Paper 2002-0263, 2002.
- [218] Makino Y., Iwamya T., and Lei Z. *Fuselage shape optimization of a wing-body configuration with nacelles*. AIAA Paper 2001-2447, 2001.
- [219] Cabuk H and Modi V. Optimum plane diffusers in laminar flow. *Journal of Fluid Mechanics*, Vol. 237, pp. 373–393, April 1992.

- [220] Soemarwoto B.I. *The variational method for aerodynamic optimization using the Navier-Stokes equations*. NASA/CR-97-206277, ICASE Report No 97-71, 1997.
- [221] Ragab S.A. *Shape optimization in free surface potential flow using an adjoint formulation: surface ships*. AIAA Paper 2001-3042, 2001.
- [222] Iollo A., Kuruwila G., and Ta'asan S. *Pseudo-time method for optimal shape design using the Euler equations*. NASA CR-198205, ICASE Report No 95-59, 1995.
- [223] Iollo A. and Salas M.D. *Contribution to the optimal shape design of two-dimensional internal flows with embedded shocks*. NASA CR-195062, ICASE Report No 95-20, 1995.
- [224] Iollo A. and Salas M.D. Optimum transonic airfoils based on the Euler equations. *Computers and Fluids*, Vol. 28, No. 4-5, pp. 653–674, May 1999.
- [225] Iollo A., Salas M.D., and Ta'asan S. *Shape optimization governed by the Euler equations using an adjoint method*. NASA-CR-191555, ICASE Report No 93-78, 1993.
- [226] Arian E. and Salas M.D. *Admitting the inadmissible: adjoint formulation for incomplete cost functionals in aerodynamic optimization*. NASA/CR-97-206269, ICASE Report No 97-69, 1997.
- [227] Arian E. and Salas M.D. Adjoint formulation using auxiliary boundary equations (ABEs) demonstrated on the Stokes equations. *16th International conference on numerical methods in fluid dynamics*, pp. 355–360. Springer-Verlag, Berlin, 1998.
- [228] Giles M.B. and Pierce N.A. On the properties of solutions of the adjoint Euler equations. *6th ICFD Conference on Numerical Methods for Fluid Dynamics*. Oxford, UK, 1998.
- [229] Hiernaux S. and Essers J.-A. Aerodynamic shape optimization using continuous or discrete adjoint formulation. *Proceedings of the 4th European Computational Fluid Dynamics Conference ECCOMAS 98*, Vol. 1, pp. 598–603. Athens, Greece, 7-11 September 1998.
- [230] Qin N., Scriba K.W., and Richards B.E. Shock-shock, shock-vortex interaction and aerodynamic heating in hypersonic corner flow. *The Aeronautical Journal*, Vol. 95, No. 945, pp. 152–160, May 1991.
- [231] Spalart P.R. and Allmaras S.R. *A one-equation turbulence model for aerodynamic flows*. AIAA Paper 92-0439, 1992.
- [232] Newman J.C., Taylor A.C., Barnwell R.W., Newman P.A., and Hou G.J.-W. Overview of sensitivity analysis and shape optimization for complex aerodynamic configurations. *Journal of Aircraft*, Vol. 36, No. 1, pp. 87–96, January-February 1999.

- [233] Fiacco A.V. and McCormick G.P. *Nonlinear programming: sequential unconstrained minimization techniques*. John Wiley & Sons, Inc., 1968.
- [234] Dadone A. and Grossman B. *Fast convergence of viscous airfoil design problems*. AIAA Paper 2000-2547, 2000.
- [235] Dadone A. and Grossman B. Progressive optimization of inverse fluid dynamic design problems. *Computers and Fluids*, Vol. 29, No. 1, pp. 1–32, January 2000.
- [236] Ta'asan S. Multigrid one-shot methods and design strategy. *Inverse design and optimisation methods*, Lecture Series 1997-05. von Karman Institute for Fluid Dynamics, Rhode Saint Genese, Belgium, 21-25 April 1997.
- [237] Ta'asan S. *Pseudo-time methods for constrained optimization problems governed by PDE*. NASA CR-195081, ICASE Report No 95-32, 1995.
- [238] *PVM Parallel Virtual Machine*. Web page, available at http://www.epm.ornl.gov/pvm/pvm_home.html, 2002. Last accessed on 12/03/2003.
- [239] *The Message Passing Interface (MPI) standard*. Web page, available at <http://www-unix.mcs.anl.gov/mpi/>. Last accessed on 12/03/2003.
- [240] *OpenMP Simple, portable, scalable SMP programming*. Web page, available at <http://www.openmp.org/>. Last accessed on 12/03/2003.
- [241] NASA. *The Blended-Wing-Body. Super jumbo jet concept would carry 800 passengers*. Fact Sheet FS-1997-07-24-LaRC, 1997. Available at <http://oea.larc.nasa.gov/PAIS/pdf/FS-1997-07-24-LaRC.pdf>. Last accessed on 25/10/2002.
- [242] NASA. *The Blended Wing Body. A revolutionary concept in aircraft design*. Fact Sheet FS-2001-04-24-LaRC, 2001. Available at <http://oea.larc.nasa.gov/PAIS/BWB.html>. Last accessed on 25/10/2002.
- [243] Wilson J.R. New blend for an old wing design. *Aerospace America*, Vol. 38, No. 4, pp. 28–35, April 2000.
- [244] Liebeck R. *Design of the Blended-Wing-Body subsonic transport*. AIAA Paper 2002-0002, 2002.
- [245] *Blended Wing Body Project*. Web page, available at <http://aero.stanford.edu/BWBProject.html>. Last accessed on 25/10/2002.
- [246] Wakayama S. and Kroo I. The challenge and promise of Blended-Wing-Body optimization. *7th AIAA/USAF/NASA/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 1, pp. 239–250, AIAA-98-4736. St Louis, Missouri, 2-4 September 1998.

- [247] Wakayama S. Multidisciplinary design optimization of the Blended-Wing-Body. *7th AIAA/USAF/NASA/ISSMO symposium on multidisciplinary analysis and optimization*, Vol. 3, pp. 1771–1779, AIAA-98-4938. St Louis, Missouri, 2-4 September 1998.
- [248] Wakayama S. Blended-Wing-Body optimization problem setup. *8th AIAA/USAF/NASA/ISSMO symposium on multidisciplinary analysis and optimization*. AIAA-2000-4740, Long Beach, CA, 6-8 September 2000.
- [249] Whitlock J. *Blended Wing Body*. Web page, available at <http://www.geocities.com/witewings/bwb/index.html>, 2002. Last accessed on 25/10/2002.
- [250] *Blended Wing Body Airliner - the next generation of civil transport aircraft*. Web page, available at <http://www.wing.cranfield.ac.uk>. Last accessed on 25/10/2002.
- [251] Smith H. College of Aeronautics Blended Wing Body development programme. *22nd International Congress of Aeronautical Sciences*. ICAS-2000-1.1.4, Harrogate, UK, 27 August-1 September 2000.
- [252] Bolsunovsky A.L., Buzoverya N.P., Gurevich B.I., Denisov V.E., Dunaevsky A.I., Shkadov L.M., Sonin O.V., Udzhuhu A.J., and Zhurihin J.P. Flying wing - problems and decisions. *Aircraft Design*, Vol. 4, pp. 193–219, 2001.
- [253] Smith H. Aerospace Vehicle Design Group Project: BW-01 Blended Wing-Body advanced technology integration study. *Aerogram*, Vol. 10, No. 4, August 2002.
- [254] Gilmore R., Wakayama S., and Roman D. Optimization of high-subsonic Blended-Wing-Body configurations. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5666. Atlanta, GA, 4-6 September 2002.
- [255] Morris A.J. MOB a European distributed multi-disciplinary design and optimisation project. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5444. Atlanta, GA, 4-6 September 2002.
- [256] Laban M., Arendsen P., Rouwhorst W.F.J.A., and Vankan W.J. A computational design engine for multi-disciplinary optimisation with application to a Blended Wing Body configuration. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5446. Atlanta, GA, 4-6 September 2002.
- [257] Vankan W.J. and Laban M. A SPINeware based computational design engine for integrated multi-disciplinary aircraft design. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5445. Atlanta, GA, 4-6 September 2002.

- [258] Bartholomew P., Lockett H., and Gallop J. The development of the MOB data and product management system. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5497. Atlanta, GA, 4-6 September 2002.
- [259] La Rocca G., Krakkers L., and van Tooren M.J.L. Development of an ICAD generative model for Blended Wing Body aircraft design. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5447. Atlanta, GA, 4-6 September 2002.
- [260] Qin N. Aerodynamic studies for Blended Wing Body aircraft. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5448. Atlanta, GA, 4-6 September 2002.
- [261] Stettner M. and Voss R. Aeroelastics, flight mechanics and handling qualities of the MOB BWB configuration. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5449. Atlanta, GA, 4-6 September 2002.
- [262] Weinerfelt P. *3D aerodynamic shape optimization of the MOB geometry using a gradient based optimization method*. Progress report for Task 7 in the MOB project, 2001.
- [263] ESDU. *A method of determining the wave drag and its spanwise distribution on a finite wing in transonic flow*. ESDU 87003, 1995.
- [264] Giles M.B. Adjoint method for aeronautical design. *ECCOMAS Computational Fluid Dynamics Conference 2001*. Swansea, Wales, UK, 4-7 September 2001.
- [265] Jameson A. and Vassberg J.C. Studies of alternative numerical optimization methods applied to the Brachistochrone problem. *OptiCon'99 conference*. Newport Beach, CA, October 1999. Also available in *Computational Fluid Dynamics Journal*, Vol. 9, No. 3, pp. 281-296, October 2000.
- [266] Campobasso M.S. and Giles M.B. Effects of flow instabilities on the linear analysis of turbomachinery aeroelasticity. *AIAA Journal of Propulsion and Power*, Vol. 19, No. 2, pp. 250–259, March-April 2003.
- [267] Reuther J., Alonso J.J., Martins J.R.R.A., and Smith S.C. *A coupled aero-structural optimization method for complete aircraft configurations*. AIAA Paper 99-0187, 1999.
- [268] Nadarajah S. and Jameson A. Optimal control of unsteady flows using a time accurate method. *9th AIAA/ISSMO symposium on multidisciplinary analysis and optimization*, AIAA 2002-5436. Atlanta, GA, 4-6 September 2002.
- [269] Levine J. *The Dryden X-Press: The morphing aircraft*. Web page, available at http://www.dfrc.nasa.gov/PAO/X-Press/stories/043001/new_morph.html, 2001. Last accessed on 26/11/2002.

Appendix A

Linearisation of the boundary conditions

Let us denote $\mathbf{P} = (\rho \ u \ v \ w \ p)^t$ the vector of primitive variables and $\bar{\mathbf{P}} = (\rho \ \bar{U} \ \bar{V} \ \bar{W} \ p)^t$ the same vector but this time with the velocity calculated in the body fitted transformed coordinates. For all the boundary conditions detailed in Section 5.6.3, this appendix gives the matrix $\frac{\partial \mathbf{P}_1}{\partial \mathbf{P}_2}$ used in MERLIN and the additional matrices $\frac{\partial \mathbf{P}_0}{\partial \mathbf{P}_2}$ and $\frac{\partial \mathbf{P}_0}{\partial \mathbf{P}_3}$ used in the adjoint solver.

A.1 Calculation of $\frac{\partial \mathbf{P}_1}{\partial \mathbf{P}_2}$

A.1.1 Inviscid wall boundary condition

From equation (5.20) it is easy to calculate:

$$\frac{\partial \mathbf{P}_1}{\partial \mathbf{P}_2} = \frac{\partial \mathbf{P}_1}{\partial \bar{\mathbf{P}}_1} \cdot \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \cdot \frac{\partial \bar{\mathbf{P}}_2}{\partial \mathbf{P}_2} \quad (\text{A.1})$$

A.1.2 Viscous wall boundary condition

The linearisation of equation (5.21) gives

$$\frac{\partial \mathbf{P}_1}{\partial \mathbf{P}_2} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

A.1.3 Symmetry boundary condition

Since mathematically the symmetry boundary condition is identical to the inviscid wall boundary condition, the linearisation of this former is the same as equation (A.1).

A.1.4 Supersonic inflow boundary condition

In the case of supersonic inflow boundary, the flow value in the halo cells does not depend on what happens inside the domain hence

$$\frac{\partial \mathbf{P}_1}{\partial \mathbf{P}_2} = \mathbf{0} \quad (\text{A.2})$$

A.1.5 Supersonic outflow boundary condition

In subsection 5.6.3.5, it was explained that in the case of a supersonic outflow boundary, the flow value in cell 2 inside the domain is directly copied to the first halo cell 1 hence

$$\frac{\partial \mathbf{P}_1}{\partial \mathbf{P}_2} = \mathbf{I} \quad (\text{A.3})$$

where \mathbf{I} is the identity matrix.

A.1.6 Farfield boundary condition

Since the farfield boundary condition is the most complicated boundary condition presented in section 5.6.3, its linearisation is not as simple as for the other boundary conditions. However a careful application of the chain rule of differentiation, with the same

notation as in subsection 5.6.3.6, gives for a subsonic boundary:

$$\frac{\partial \mathbf{P}_1}{\partial \mathbf{P}_2} = \frac{\partial \mathbf{P}_1}{\partial \bar{\mathbf{P}}_1} \cdot \left[\begin{array}{cccccc} -\frac{\rho_b a_2}{2(\gamma-1)a_b \rho_2} + flag \cdot \frac{\gamma \rho_b s_2}{(\gamma-1)s_b \rho_2} & -\frac{\rho_b}{2a_b} & 0 & 0 & 0 & 0 \\ \frac{a_2}{2(\gamma-1)\rho_2} & \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & flag & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & flag & 0 \\ -\frac{\rho_b a_b a_2}{2(\gamma-1)\rho_2} + flag \cdot \frac{\gamma p_b s_2}{(\gamma-1)s_b \rho_2} & -\frac{\rho_b a_b}{2} & 0 & 0 & 0 & 0 \\ \frac{\rho_b a_2}{2(\gamma-1)a_b p_2} - flag \cdot \frac{\rho_b s_2}{(\gamma-1)s_b p_2} & \frac{a_2}{2(\gamma-1)p_2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{\rho_b a_b a_2}{2(\gamma-1)p_2} - flag \cdot \frac{p_b s_2}{(\gamma-1)s_b p_2} & -\frac{p_b s_2}{(\gamma-1)s_b p_2} & 0 & 0 & 0 & 0 \end{array} \right] \cdot \frac{\partial \bar{\mathbf{P}}_2}{\partial \mathbf{P}_2} \quad (\text{A.4})$$

where $flag = \begin{cases} 1 & \text{for outflow boundary} \\ 0 & \text{for inflow boundary} \end{cases}$.

For a supersonic boundary, either equation (A.2) or (A.3) is used.

A.1.7 Interface boundary condition

In the case of an interface boundary, the flow value in the halo cell of one block does not depend on what happens inside this same block hence the linearisation (A.2) has to be employed. This is only used in MERLIN, in the adjoint solver a more accurate treatment of this boundary condition is employed.

A.2 Calculation of $\frac{\partial \mathbf{P}_0}{\partial \mathbf{P}_2}$

According to the details given in section 5.6.3

$$\frac{\partial \mathbf{P}_0}{\partial \mathbf{P}_2} = \mathbf{0}$$

for the following boundary conditions:

- Inviscid wall boundary condition
- Viscous wall boundary condition
- Symmetry boundary condition

- Supersonic inflow boundary condition
- Supersonic outflow boundary condition

For a farfield boundary condition,

$$\frac{\partial \mathbf{P}_0}{\partial \mathbf{P}_2} = 2 \times \frac{\partial \mathbf{P}_1}{\partial \mathbf{P}_2} - \mathbf{I}$$

where $\frac{\partial \mathbf{P}_1}{\partial \mathbf{P}_2}$ is given by equation (A.4).

A.3 Calculation of $\frac{\partial \mathbf{P}_0}{\partial \mathbf{P}_3}$

For an inviscid wall boundary condition and a symmetry boundary condition, the linearisation of equation (5.20) gives

$$\frac{\partial \mathbf{P}_0}{\partial \mathbf{P}_3} = \frac{\partial \mathbf{P}_0}{\partial \bar{\mathbf{P}}_0} \cdot \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \cdot \frac{\partial \bar{\mathbf{P}}_3}{\partial \mathbf{P}_3}$$

For a viscous wall boundary conditions, equation (5.21) simply gives

$$\frac{\partial \mathbf{P}_0}{\partial \mathbf{P}_3} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

According to the details given in section 5.6.3

$$\frac{\partial \mathbf{P}_0}{\partial \mathbf{P}_3} = \mathbf{0}$$

for the following boundary conditions:

- Supersonic inflow boundary condition
- Supersonic outflow boundary condition
- Farfield boundary condition