Discrete Adjoint Method in *elsA* (part 2): Application to Aerodynamic Design Optimisation.

G. Carrier, S. Mouton, I. Salah El Din

ONERA, Applied Aerodynamics Department, 29 Avenue de la Division Leclerc, 92322 Châtillon, France

Abstract

The introduction of the Control Theory into the fields of Fluid Dynamics and Aerospace has been pioneered by Pironneau [8] and Jameson [5], opening the way for the development of adjoint CFD codes that enable efficient and accurate sensitivity calculations, valuable information in design optimisation. This paper presents three applications of aerodynamic design optimisation performed with the discrete adjoint technique, recently introduced into the *elsA* [1] software.

The first example deals with the enhancement of aerodynamic performance of a two-dimensional supercritical airfoil using local geometrical optimisation. A second application of the adjoint approach is then presented with a multipoint optimisation of a Supersonic Commercial Transport wing. Finally, the last application concerns the optimisation of the engine pylon of a large transport aircraft. This wide range of application demonstrates the maturity and advantages of the adjoint approach compared with more traditional approaches such as finite difference.

Keywords: Optimisation, Adjoint, Gradient, Aerodynamics, Design

1. Introduction

Once a revolution, automated design optimisation based on the control theory [8][5] is becoming a key point in aerodynamic shape design. Aerodynamic design optimisation relies on two main components which are the optimisation algorithm and the numerical evaluation based on Computational Fluid Dynamics (CFD) analysis. The former is meant to drive the optimum design search within the space of potential design variables, whereas the latter supplies the required characteristics (objective and constraints functions, gradients…) corresponding to any design configuration for which analysis is requested by the optimiser.

Accurate and robust CFD tools can be combined with a wide variety of optimisation methods. The optimisers can be categorised depending on the type of search that is performed within the design space: global or local. Among the most popular optimisers, Evolutionary Algorithms (EA) and gradient-based algorithms are examples of each category. An important practical obstacle to optimisation is the computational time needed for design analyses.

For aerodynamic optimisations the CFD based analyses can be prohibitively expensive despite a increasing computer power and codes efficiency. In the case of gradient-based optimisations the analysis process computational time depends strongly on the gradients evaluations efficiency. Scaling linearly with the number of design variables, the cost of gradients calculations by finite differences becomes rapidly a limitation. In this case the adjoint approach, pioneered by Jameson [5] and recently introduced in the *elsA* [1] solver by Peter [7] following the discrete approach, proves its usefulness. As a matter of fact, gradients can be obtained with this method at the price of a number of adjoint calculations equal to the number of functions to be differentiated (objective and constraints functions), which rarely exceeds the number of design variables in practical design problems.

The present paper describes three different aerodynamic design optimisations based on gradient method and using the adjoint approach, and carried out recently at ONERA:

- a mono objective multi-constraint optimisation involving a 2D airfoil trailing edge;
- a multipoint multi constraint supersonic wing design;
- a mono objective unconstrained 3D pylon design.

These three examples cover a wide range of problems 2D and 3D configurations in viscous and inviscid, at transonic and supersonic flow conditions.

2. Optimisation strategy and framework

The first two optimisation cases have been performed using an in-house optimisation framework with the Method of Feasible Directions (MFD). The CONMIN [10] optimiser is used directly or indirectly within DAKOTA [12] optimisation toolkit. The analysis module is built upon:

- parametric surface mesh deformation procedures;
- mesh deformation tools based on an analytical (integral) formulation associated with a transfinite interpolation technique;
- the *elsA* [1] software used to solve both the Euler or RANS flow equations and also the associated adjoint equations
- the ONERA drag and sensitivities postprocessor FFD41 [11];

Figure 1 gives an overview of such in-house optimisation

framework, as it has been used for supersonic aircraft wing optimisation.

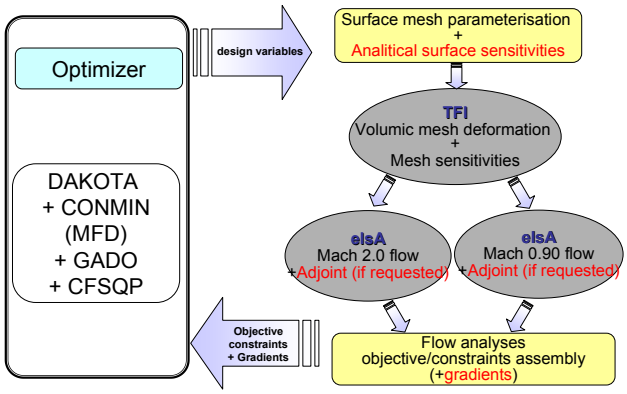

Figure 1: Optimisation loop architecture.

The Python [9] script language is used to assemble the different components. The flexibility of Python is useful to schedule the execution of the different components and organise data flow between them. For instance, the CFD analyses at the two cruise conditions are performed in parallel, using the Python threading library.

The third optimisation problem is achieved with the *Optalia* optimisation suite developed by Airbus, using DOT optimiser and the Modified MFD algorithm.

In all cases, flow analysis is performed using the ONERA's NEC SX8+ vector super calculator while the optimiser runs on an SGI ALTIX (Itanium processor) computer.

3. Thick Cambered Trailing Edge Optimisation

3.1. Background, problem formulation

Most of recent civil aircraft wings are based on efficient supercritical airfoil geometries. The introduction of thick trailing edges on such airfoils[3] aims at uncoupling the pressure distribution on upper and lower surface, hence allowing to increase the rear contribution to lift. For these configurations, smaller angles of attack are needed to maintain a target lift condition. As a consequence, the upper surface shock wave is weakened, resulting in a reduction of pressure drag. The optimisation of the lower surface thickness and curvature at the trailing edge is meant to ensure that the pressure drag gain remains larger than the viscous drag penalty generated by such a shape while maintaining a given lift level.

The supercritical airfoil geometry called upon for this study is based on the OAT15A model equipped with a 0.5% trailing edge thickness.

Considering the aim of the study, it seemed natural to choose a mono-objective multi-constraint optimisation problem. In this case the objective function (OF) is the total drag coefficient C_D and the constraints (CSTR) concern an acceptable lift level (C_L>C_{Lmin}) and a lower limit to the pitching moment (C_M>C_{Mmin}). The design variables are restricted to the angle of attack and two geometrical parameters, Δ and k, introduced in the following paragraph.

3.2. Parameterisation

The lower surface trailing edge follows a polynomial law allowing building up the thick cambered trailing edge:

$$\delta z = \Delta\left(\frac{x/C_0 - x_0/C_0}{1 - x_0/C_0}\right)^k \quad \forall x \in [x_0, C_0]$$

where δz is the thickness variation, Δ the additional thickness at the trailing edge compared to the reference case, k controlling the curvature of the deformation, x_0 the abscissa at the beginning of which the lower surface is modified, and C_0 the airfoil chord length.

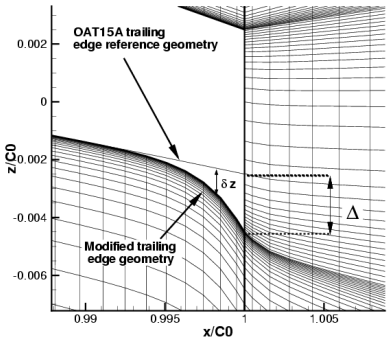

Figure 2: Thick cambered trailing edge parameterisation

The resulting geometrical aspect of a given non zero set of values is schemed in Figure 2.

3.3. Flow field analysis

The airfoil length is normalised and the domain is split into two sub domains. The boundaries are situated at more than 80 times the airfoil length. Domain number one is semi-circular with an 80 airfoil length radius and stops at the trailing edge. C-type meshing is combined to an H-type one to describe the whole domain, finally leading to a mesh having 24,050 cells. The meshes are refined adequately near the airfoil wall to capture the boundary layer. Refinements are also made in the wake for accurate aerodynamic coefficients computations (Figure 3).

The aerodynamic problem is solved using a RANS 2D steady flow approach with a perfect gas assumption and a Spallart Allmaras turbulence model. A backward-Euler integration scheme is coupled with an upwind Roe scheme extended to the second order thanks to the Van Albada limiter. A v-cycle multigrid approach is chosen to accelerate convergence and the system is solved using a 4-cycles scalar LU-SSOR implicit resolution. For stability and robustness reasons, a Harten entropy correction is applied both to the Euler resolution and to the uncoupled turbulent equations. The number of iterations is chosen depending on the convergence level requested and the corresponding CPU time. The whole study is carried out at a Mach number of 0.73 and a Reynolds number of 3.3 10^6.

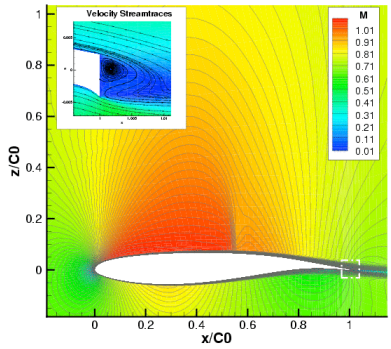

Figure 3 : Supercritical airfoil in transonic flow - Iso Mach contours and trailing edge streamlines

3.4. Adjoint state resolution

The reduced complexity of this problem allowed to perform to a non exhaustive but detailed study of the accuracy and/or consistency of the gradients computed with the adjoint approach. The sensitivity of the Adjoint state resolution to various numerical parameters has been observed.

An acceptable range of the *Harten entropy correction coefficient* was determined. Within this range the *convergence level of the analysis* computation stage have proved to play an important role on adjoint state resolution convergence.

The variation of the number of iterations of the analysis stage has put forward a problem of lack of stability. This black point has been overcome by using *artificial viscosity* similar to the classical Jameson scheme. Very small values second and fourth order coefficients have an immediate effect in delaying divergence significantly, as shown in Figure 4. Here again a range was determined for each parameters, so that acceptable convergence level of adjoint state resolution could be reached.

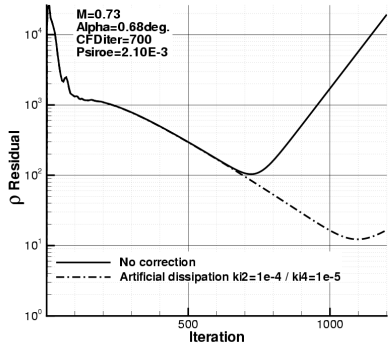

Figure 4 : Artificial viscosity influence on adjoint state resolution convergence

Consistency with first order finite difference results was successfully checked. Gradients were computed for a given set of geometrical parameters and angle of attack. The same highly converged analysis flow field was used in both cases. Every gradients component with respect to each design variable was evaluated. The finite difference computations were achieved using preliminarily determined discretisation steps for each design variable. Calculations lead to discrepancies ranging from *2 to 15% on gradients components*, that is satisfactory in comparison to the errors that may be introduced by the finite difference method. The pitching moment gradient with respect to the angle of attack is not taken into account as, by definition, the coefficient undergoes no variations.

Finally, the adjoint resolution gradient estimation sensitivity to artificial viscosity modification has shown that even for extreme but acceptable values, apart from $\partial C_M/\partial\alpha$, fluctuations *do not exceed 10%*.

3.5. Optimisation results

Optimisation configurations

The initial set of design variables configuration set is chosen identical to the one used for previous studies: the thickness is small with little curvature. The optimisation process has been achieved with both finite difference and adjoint state resolution computed gradients approaches. The resulting optima were compared and their consistency checked.

For the finite difference computations a fixed step approach has been chosen. This step is fixed at the previous studies values: $\delta\alpha$=0.005°, δk =0.005 and $\delta\Delta$=0.005C_0. Regarding adjoint gradient computation the design variables steps are set referring to the study on the adjoint state resolution convergence, which means: $\delta\alpha=10^{-3}$°, $\delta k=10^{-2}$ and $\delta\Delta=10^{-5}$m. Every presented configuration is computed using the same solver version (*elsA* v3.1.13).

In the following optimisation computations, when not otherwise mentioned, the initial constraint thickness parameter value is fixed at 0.1 and the minimum at 4.10^{-3}. Relative convergence levels are set to 10^{-6} and the convergence criterion is defined as three successive iterations undergoing variations lower than the convergence level.

Impact of optimisation on airfoil performance

The first numerical optimisations on thick cambered trailing edge problem were achieved using finite difference gradient computations and an adapted CONMIN optimiser. The optima were found for two lift conditions: C_L=0.7 and C_L=0.75. Table 1 and table 2 summarise the resulting geometries and respective aerodynamic performance.

	k	Δ	α	C_D	C_M
OAT15A C_L=0.7	0.	0.	1.057	0.01386	-0.1306
INIT C_L=0.7	1.5	0.0005	0.929	0.01373	-0.1368
DF Opt C_L=0.7	2.545	0.00113	0.7	0.01363	-0.1466

Table 1 :Finite difference optimisation for $C_L \geq 0.7$ lift constraint

	k	Δ	α	C_D	C_M
OAT15A C_L=0.75	0.	0.	1.057	0.01386	-0.1306
INIT C_L=0.75	1.5	0.0005	1.166	0.01478	-0.1368
DF Opt C_L=0.75	6.207	0.00162	0.664	0.01427	-0.1579

Table 2 : Finite difference optimisation for $C_L \geq 0.75$ lift constraint

Both optimisations lead to a modification of the trailing edge geometry by a thickening and an increase of the curvature resulting in a large decrease of the angle of attack necessary to satisfy the lift constraint. The corresponding $C_L(C_D)$ polars

show that the optima are more efficient for the given lift conditions than the reference OAT15A configuration. The optimisation is all the more efficient for non adapted computation conditions (C_L=0.75).

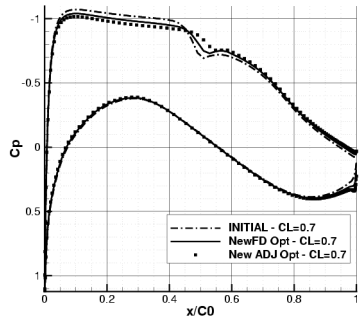

Figure 5 : Pressure coefficient distribution of initial and optimised configurations at C_L=0.7

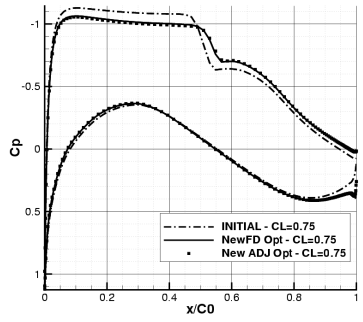

Figure 6 : Pressure coefficient distribution of initial and optimised configurations at C_L=0.75

During the optimisation, the leading edge expansion is reduced as well as the upper surface shock wave amplitude. The latter also appears downstream compared to the shock wave position for the original OAT15A (Figure 5 and Figure 6). Finally, the rear load is increased, in order to satisfy the lift constraint.

Optimisation using Adjoint state resolution

The same optimisations have been achieved by substituting to the finite difference approach the adjoint state resolution. The resulting optimum configurations and the corresponding aerodynamic performances are summarised in Table 1 and Table 2.

CL=0.7	*k*	*Δ*	*α*	C_D	C_M
ADJ	3.17	0.00158	0.536	0.01364	-0.1536
FD	2.545	0.00113	0.7	0.01363	-0.1466
%ADJ / FD	24.3	39	23	**0.07**	**4.8**

Table 1 : Adjoint approach optimisation for CL≥0.7 lift constraint

CL=0.75	*k*	*Δ*	*α*	C_D	C_M
ADJ	3.3	0.00204	0.623	0.01426	-0.1595
FD	6.2	0.00162	0.664	0.01428	-0.1579
%ADJ / FD	47	26	8	**0.14**	**1.1**

Table 2: Adjoint approach optimisation for CL≥0.75 lift constraint

Finite difference and Adjoint approaches give very close results. The greatest discrepancy observed (lower than 5%) concerns the pitching moment coefficient. The differences observed remain local, at the trailing edge level, due to different geometrical design variables.

The behaviour of the optimisation process observed throughout the evolution of the objective and the constraints is coherent with the initial request as the drag always decreases from one iteration to the next and the constraints remain satisfied. The results of the behaviour of the optimisation relatively to the design variables and the aerodynamic coefficients for the two lift constraints are given in Figure 7 and Figure 8.

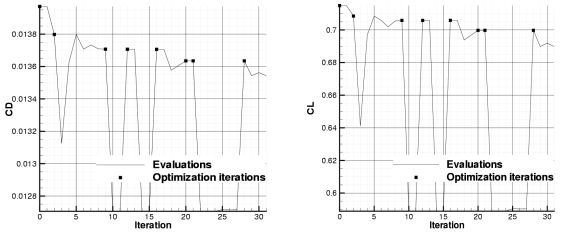

Figure 7 : Drag and lift evolutions throughout the optimisation process for C_L≥0.7 constraint

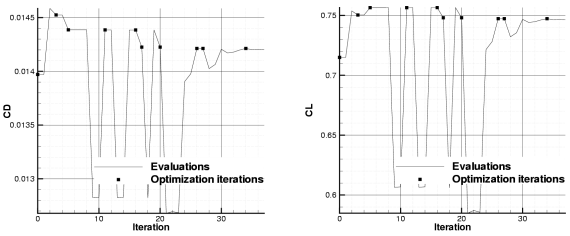

Figure 8: Drag and lift evolutions throughout the optimisation process for CL≥0.75 constraint

The analysis of the evolution of the various parameters throughout the optimisation process is similar to the previous one: important variations for the first iterations followed by an adjustment of the values to minimise the objective function without violating the constraints. The main difference with the finite difference optimisation is that the number of iterations is shortened, and hence the optimum reached sooner.

The quality of the optimum has been assessed by modifying the constraint thickness and the convergence level of the optimiser. This is motivated by the fact that the optimum does not lead to a satisfactory lift constraint: 0.747 instead of 0.75 and 0.698 instead of 0.7. What appears clearly is that the angle of attack undergoes little, if any, modifications. The reduction of the minimum constraint thickness from 10^{-3} to 10^{-5} has an impact on the geometrical design variables, which local influence helps to reach the constraints limits more accurately. Considering the actual optimisation module, the only limitation to the use of very thin constraint thickness is that the convergence criteria may never be reached and the optimisation would have to be stopped manually.

3.6. Discussion & Prospect

The aim of this study was to validate a new gradient evaluation technique based on adjoint state resolution for an elementary 2D case, the OAT15A supercritical airfoil. The results obtained during previous studies were used as a base to compare with. The preliminary study has shown that the behaviour of the gradient evaluation methods were consistent.

The adjoint and finite difference based optimisations lead to solutions with equivalent aerodynamic performances even though the optimum geometrical parameters may differ in a significant way as their influence is local once the upper surface and lower surface are uncoupled. The results of optimisations encourage the use of the adjoint approach for more complex cases. The solution to this kind of optimisation problem leads to various local optima and is highly dependent on the configuration and the parameterisation of the process. Enhancement such as the Kuhn-Tucker conditions, to check the convergence is effectively met, will be taken into consideration in future optimisations.

4. Wing optimisation of a Supersonic Commercial Transport aircraft

In this second application of optimisation, the *elsA* adjoint capabilities is used to optimise the wing shape of a generic Supersonic Commercial Transport (SCT) aircraft, the flow being modelled with the three-dimensional compressible Euler equations.

4.1. Background, problem formulation

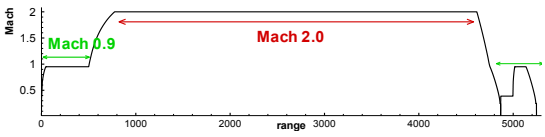

Figure 9 : Realistic mission profile of a supersonic transport aircraft.

For different reasons, a realistic mission profile for a supersonic commercial transport aircraft will include both transonic and supersonic cruise legs, as depicted in Figure 9. Such an aircraft will thus be asked to offer good aerodynamic performances in both cruise conditions, yielding a multipoint (MP) aerodynamic optimisation problem, which can be formulated as:

> *"Find the wing shape of a supersonic transport aircraft providing optimum trade-off of the aerodynamic performance over both supersonic and transonic cruise conditions, while full-filling the constraints imposed at each mission point."*

Several approaches are possible to tackle this MP optimisation problem, including the use of stochastic algorithms that intend to provide a set of optimal trade-off between both objectives. In this work, the simplest approach has been chosen that consists in optimising a single composite objective function defined as a linear combination of aerodynamic performance in both cruise conditions. The resulting mono-objective optimisation problem, defined by:

Minimize: $k_1 \cdot CD_{M=2.0}(\alpha) + k_2 \cdot CD_{M=0.9}(\alpha)$

s.c. • $CL_{M=2.0}(\alpha) \geq 0.11$

• $CL_{M=0.9}(\alpha) \geq 0.21$

• Geometrical constraints

with $\alpha=(\alpha_1, \ldots \alpha_N)$: wing shape design variables + A.o.A. at each flight points

can be solved with a conventional gradient-based method and the adjoint technique can be applied to efficiently provide the necessary sensitivity information of the different functions.

The results presented in the following have been obtained for equally weighted objectives ($k_1 = k_2 = 0.5$), and with constraints on the minimum lift coefficients of $C_L^{min}{}_1$=0.11 and $C_L^{min}{}_2$=0.21, respectively for the supersonic and transonic cruise conditions.

4.2. Parameterisation

During the optimisation, the reference wing geometry (presented in Figure 10) is deformed through a parametric analytical method which allows to freely modify the twist and camber of the wing, for a fixed wing planform. The deformations can be specified in several spanwise control sections and linearly interpolated in between. For the present optimisations two control sections have been used, at mid-span and wing tip, respectively. All deformations being cancelled at the wing root. Therefore, the wing deformations are controlled with six design variables: one for the twist and two for the camber in each of the two control sections. The displacement vectors field corresponding to these six deformation "modes" are presented in Figure 11.

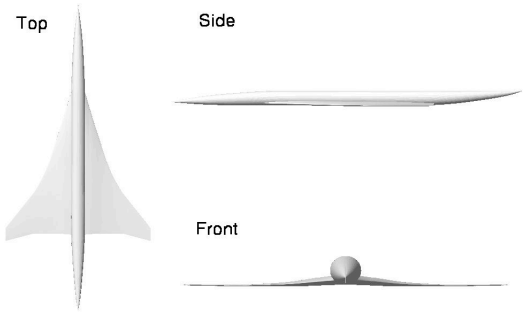

Figure 10 : Three-views of the reference SCT aircraft.

Two aerodynamic parameters for the angles of incidence at supersonic and transonic flight conditions are added to the six parameters controlling the wing geometry, resulting in a total number of 8 design optimisation parameters.

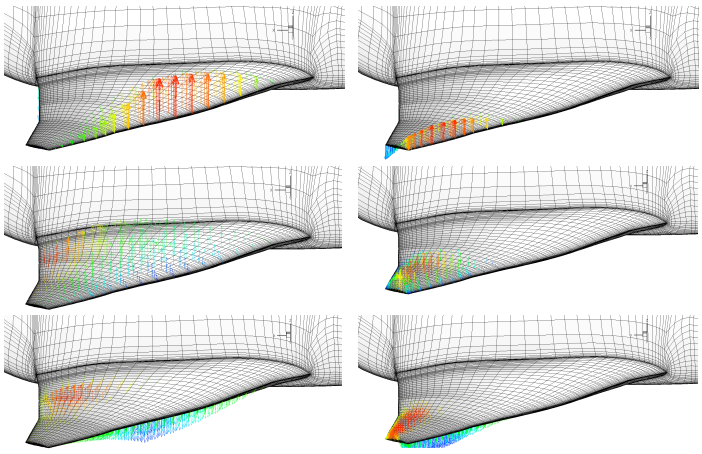

Figure 11 : Wing shape deformations allowed by the wing parameterisation with 6 parameters (twist and camber); Effects of the three mid-span section parameters (left) and those for the tip section (right).

4.3. Flow field analysis

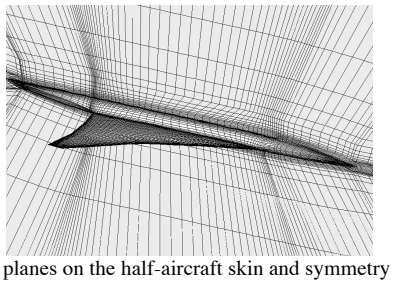

Figure 12 : Mesh planes on the half-aircraft skin and symmetry plane

The analysis of each aircraft geometry requires two CFD calculations, for the transonic (Mach number of 0.9) and supersonic (Mach number of 2.0) flight conditions. These calculations are performed with the ONERA multiblocks structured CFD software *elsA* used to solve the three dimensional compressible Euler equations using a CFD mesh of approximately 600, 000 cells (Figure 12).

The spatial discretisation of the inviscid fluxes are based on the Roe-Harten scheme (entropy correction constant of 0.05). Convergence to the steady flow solution is achieved iteratively using the combination of the backward Euler scheme associated with an scalar implicit stage (LU-SSOR). A 2-level multigrid strategy is employed to accelerate the convergence of the calculation at transonic regime. The convergence of residuals and forces obtained for the reference geometry are given in Figure 13. A sufficient convergence is obtained after 300 iterations at Mach 2.0 and 600 iterations at Mach 0.9.

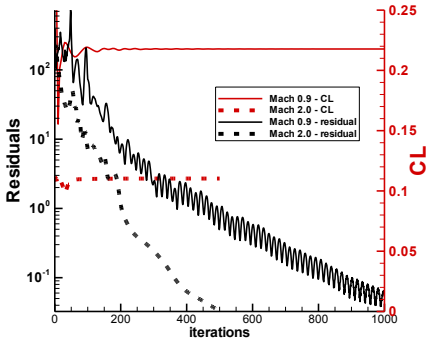

Figure 13 : Convergence of residuals and lift coefficient for the Euler calculations at Mach 2.0 and 0.9 on the reference geometry.

Note that, if an inviscid flow modelling yields sufficient accuracy on the aerodynamic coefficients evaluations and their sensitivities (in the context of the shape optimisation, under the assumption of a fixed wetted area) at supersonic speed, it is more questionable for the transonic performance analysis, especially with the mesh used in this study.

4.4. Adjoint state resolution

Each time the optimiser requires a gradient (sensitivities) information, this gradient is calculated with the adjoint method [7]. The objective function formulation involving the flow solutions at both supersonic and transonic conditions, calculating its gradient will also requires adjoint states to be calculated, for each flow condition. The *elsA* adjoint solver is used to calculate these different adjoint states. Provided the Euler flow solution is sufficiently converged, no additional dissipation was required to converge the adjoint states calculations. An example of adjoint residual convergence at both flow conditions is presented in Figure 14, for the reference geometry. Within 300 iterations, the residuals decreased by about 4 orders of magnitude.

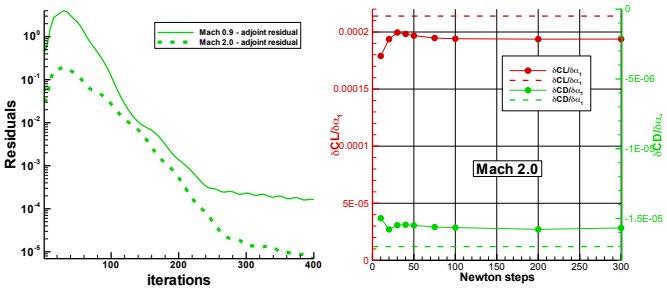

Figure 14 : Convergence of the (pressure drag and lift) adjoint residuals at Mach 2.0. for the reference configuration (left) and of the corresponding gradients (right, one component).

More interesting is the convergence of the resulting gradient with the adjoint Newton resolution iterations (Figure 14), which shows that the gradient is adequately converged within less than 100 iterations. The finite-difference (with a step of 10^{-4}) evaluation of the gradient is indicated by a horizontal dashed line.

4.5. Optimisation results

The optimisation has been conducted with the DAKOTA/CONMIN optimiser. The history of this optimisation run is given in Figure 15. The initial design has an aerodynamic drag (invicid flow) of 81 drag counts (d.c.) at Mach 2.0 and of 93 d.c. at Mach 0.9. This multipoint optimisation resulted in a design with an aerodynamic drag of 70.1 d.c. at Mach 2.0 and of 85.7 d.c. at Mach 0.9, corresponding to improved performance at both flight conditions: 13% drag reduction at Mach 2.0 and 8% at Mach 0.9.

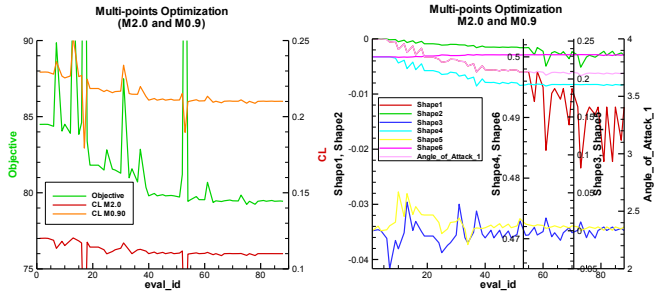

Figure 15 : Optimisation history of the multipoint optimisation for Mach 2.0 and 0.9 performed with the gradient optimiser DAKOTA/CONMIN; Convergence of the objective and constraints functions (left) and of some of the design variables.

Figure 16 shows the modification of the pressure distribution acting on the upper side of the wing resulting from the MP optimisation by comparing the final design to the initial design, which corresponds to a non-cambered untwisted wing.

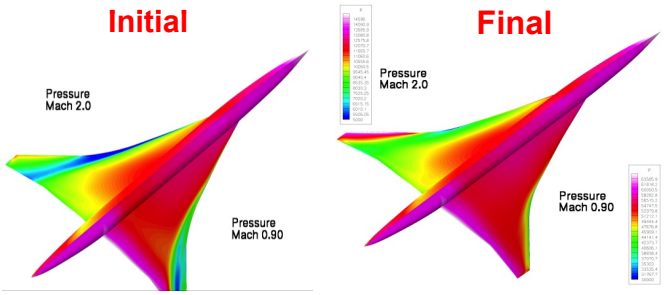

Figure 16 : Comparison of the pressure distribution at Mach 2.0 and Mach 0.9, on the suction side of the baseline and multipoint optimised configurations.

5. Optimisation of engine pylon on a large transport aircraft

5.1. Background, problem formulation

The larger size of modern aircraft engines with high by-pass ratio leads to increasing difficulties regarding engine integration under the wing. The features of the transonic flow on the whole wing are modified by the propulsive system, causing drag penalties. To minimise these penalties, the shape of the pylon attaching the engine to the wing lower surface must be carefully designed.

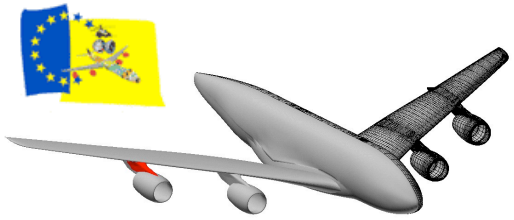

Figure 17: Geometry and mesh of a large civil aircraft and pylon to optimise in red

This work, performed within the EU project *VIVACE* (ST1.2.7.), focuses on the drag minimisation of a large civil aircraft (Figure 2) thanks to alterations of the shape of the outboard pylon.

Variations of lift coefficient C_{Lp} with respect to the vector of design parameters α is assumed to be small enough to allow keeping the angle of attack constant during the optimisation process. These variations are taken into account into the objective formulation as a negative weighting term to drag gains. Based on designers' experience, only pressure drag is expected to vary much from one pylon shape to another of similar wetted area, which allows us to consider pressure drag C_{Dp} only. Finally, the objective function J to minimise writes as:

$$J(\alpha) = C_{Dp}(\alpha) - \left(\frac{dC_{Dp}}{dC_{Lp}}\right)_{\alpha=0} \left(C_{Lp}(\alpha) - C_{Lp}(\alpha=0)\right)$$

No non-linear constraint is imposed during the optimisation. Design variables are properly bounded both to satisfy geometrical constraints and to avoid deformations exceeding the flow solver capabilities.

A gradient method is used to solve the optimisation problem. It uses the method of feasible directions with a one-dimensional search of local minimum at each iteration step. The gradient of the objective function is provided to the optimiser and computed thanks to the adjoint method.

5.2. Parameterisation

Based on previously mentioned existing tools developed at Airbus France, it was chosen to apply deformation on a pre-existing mesh. The pylon surface mesh is modified thanks to 17 bumps spread over the pylon surface. Along a longitudinal mesh line (nearly parallel to the flow), they are shaped according to Hicks-Henne bump definition [4]; and in the transverse direction, they have on each side either a linear or cubic spline shape. The bumps are controlled by 19 design parameters (13 controlling amplitudes, 3 positions of tops and 3 the tightness).

Once computed, the surface mesh deformation is extended into the volume mesh of the block surrounding the pylon thanks to an integral method. Both surface and volume mesh deformation softwares are able to deliver analytical sensitivities of the mesh with respect to each design variable.

5.3. Flow field analysis

The aircraft is examined at cruise conditions with an upstream Mach number of 0.85 and a Reynolds number of 20 million. The flow in each of the four engines is computed

thanks to proper boundary conditions on entry and exit planes. To save computational time and memory, local refinement was used, on the wing and the pylon to optimise only, in order to capture the boundary layer. The mesh comprises around 1.5 million nodes, which is much coarser than usual for such configuration.

The flow computations are achieved using the *elsA* software. RANS equations with Spalart-Allmaras model are solved thanks to a Roe upwind scheme extended to 2nd order accuracy using MUSCL method with Van Albada limiter. Previous proper smoothing of the grid and the use of a 2-level multigrid scheme allows to converge in only 500 iterations, representing 4,000 seconds on one NEC SX-8 processor.

Aerodynamic coefficients are then computed by surface integration to build the objective function. The sensitivities of these coefficients with respect to the mesh and to the computed flow-field are also analytically computed.

5.4. Adjoint state resolution

When requested by the optimiser, the gradient of the objective function is computed using the *'Opt'* module of *elsA*. It solves the discrete adjoint equations with frozen eddy viscosity. The large non-linear system is solved thanks to a LU-SSOR implicit scheme iterated according to a Newton method. Due to the lack of experience and validation of this method on complex viscous turbulent cases, extensive work was performed to assess the accuracy of the computed gradient.

First, numerical dissipation and a pseudo-time term previously introduced in the adjoint equations were shown to be sufficient to prevent the appearance of a diverging trend. The optimal amount of necessary dissipation and CFL-like number was identified, and is highlighted in Table 3.

	k_2	0	0,005	0,02	0,018
CFL	k_4	0	0,001	0,004	0,016
∞			DVI	DVI	DVI
100				DVI	
30			DVI	DVI	
25			CV 4	DVI	
20		DV	**CV 4**	CV	
10		DV	CV 3	CV 3	DVI

Table 3 : Convergence behaviour of adjoint computations. DV: divergence after initial trend to converge; DVI: immediate divergence; CV: converge, the figure indicates the rate of decrease of Log10 (L2 r-Residual) per 1,000 iterations.

Secondly, it was shown that the adjoint state equations do not need to be fully converged to give a good estimate of the searched gradient. 500 Newton iterations are sufficient to reach a converged value, although the residual of the adjoint equations have lost only one order of magnitude.

Finally, based on the above results, the accuracy of the computed gradient was assessed by comparison with the value obtained by second order finite difference. Deriving finite difference data with few or no *a priori* idea of the order of magnitude of the effect of each design parameter requires to carefully choose an appropriate step for each variable. Too small a step causes a small absolute change of the objective function, with associated numerical noise, whereas too large a step may trigger non-linear behaviour of the objective function. It is recommended to investigate several step sizes, and to rely rather on second order finite difference.

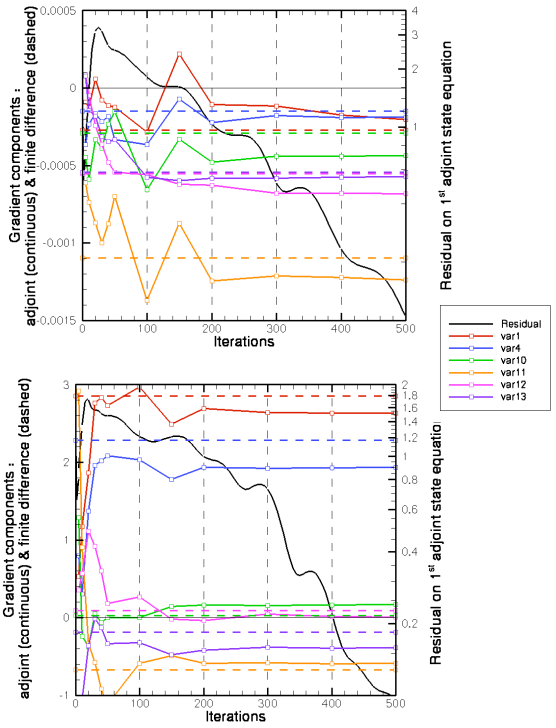

Figure 18: Convergence of adjoint state for two objective functions: C_{Dp} (top) and C_{lp} (bottom).

- In colour, continuous lines: components 1, 4, 10, 11, 12 & 13 of gradient (d.c.m^{-1} and m^{-1}) computed at different stages of the convergence history of the adjoint state.
- In colour, dashed lines: components values according to 2nd order finite difference.
- In black: L_2 residual on 1st adjoint variable.

Six components of pressure lift and drag gradients were assessed by second order finite difference, and compared to adjoint state prediction. The agreement is reasonably accurate for the dominating components of lift and drag gradient, with relative accuracy in the range *5 – 30%*. Smaller components of the gradient exhibit higher relative errors.

5.5. Optimisation results

Based on the above described results, an optimisation was performed on a large civil aircraft configuration, relying on the gradient computed by adjoint method with 500 Newton iterations.

The convergence of this unconstrained problem was satisfactory with most of the drag gain achieved within two optimisation iterations. After 8 iterations, the objective function levels off, with little further evolution during the next 11 iterations. In more details, the first 2 iterations bring 8 design variables to their bounds. During the rest of the optimisation process, 3 more variables reach their bounds, and an optimum is found on the remaining 8 parameters. The objective was decreased by 0.56 drag counts with a total CPU time dedicated to flow and adjoint computations equal to 116 hours.

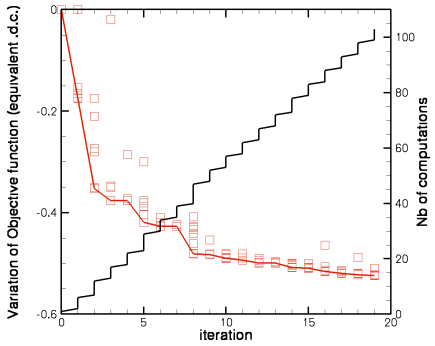

Figure 19: Convergence history of the objective function (in equivalent drag counts) and number of *elsA* calls with respect to the number of gradient iteration of the optimisation algorithm.

For this case, the most influential areas of the pylon are:

- the leading edge fairing that was deflated;
- the whole inboard intersection that was reshaped;
- the front part of outboard intersection which radius of curvature was increased.

Flow field analysis

The surface pressure field was noticeably modified only at local scale. On the pylon inboard side, the velocity peak downstream of the *'béret basque'* is higher and followed by a steeper recompression. On the outboard side, the velocity field is smoother, close to a constant pressure between 5% and 50% chord.

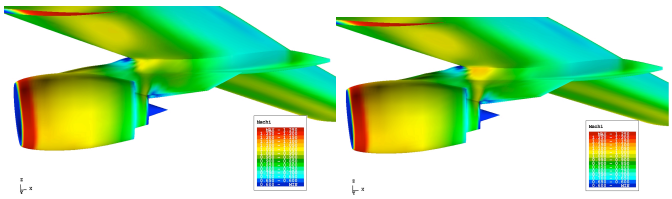

Figure 20: View from inboard of the baseline (left) and optimised (right) pylon, coloured by isentropic Mach number

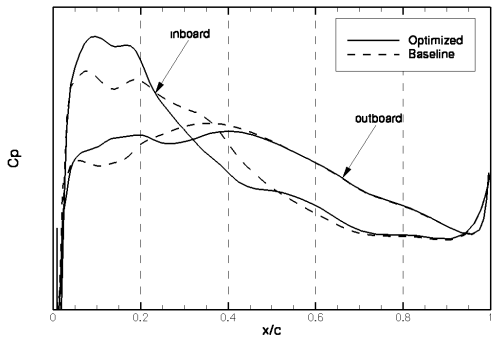

Figure 21: Pressure distribution along the wing-pylon intersection line

Drag production

A volume analysis of drag production was performed, allowing to breakdown between several physical and non-physical contributions, and to identify the area of production [11].

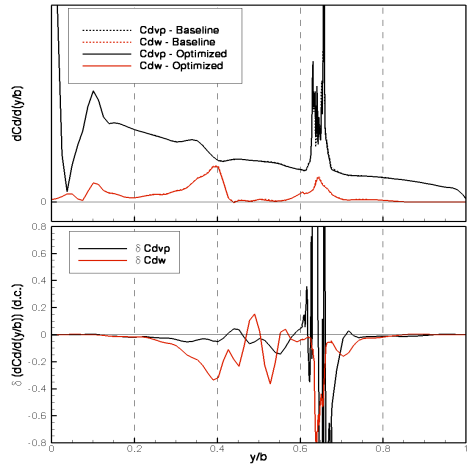

Figure 22: Far-field analysis of irreversible drag production along wing span (top) and difference between optimised and baseline shapes (bottom)

Figure 22 displays the production of irreversible drag (wave and viscous pressure) along the wing span for baseline and optimised shapes and the difference between them. Production of wave drag peaks around engines position because of the shock waves on the nacelles and because of the interaction with the wing. Wave drag was brought down on 60% of the wing span, including on the inboard part. The effect of the optimisation on viscous pressure drag is more local than for wave drag and is quite difficult to analyse because of its rapid evolution while the scanning plane travels through the boundary layers of the outboard pylon.

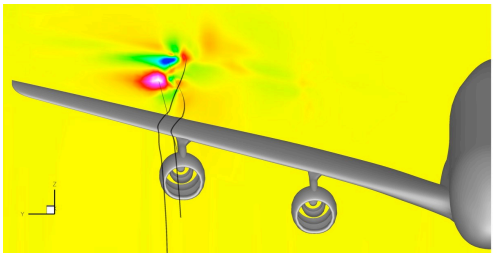

Figure 23: Difference in transverse kinetic energy between optimised and baseline configuration in a cross-plane behind of the aircraft by 10% of its length

The pylon shape locally impacts the distribution of transverse kinetic energy in the aircraft wake (Figure 23) and therefore the induced drag.

Performances

The Table 4 below presents the differences in aerodynamic coefficients between optimised and baseline configurations. Definition of far-field drag components are detailed in [11]. It can be observed that the drag gains are not obtained on the pylon itself, which produces actually more drag, but rather on the wing and nacelle contribution. This may be interpreted as wave drag and engine through flow gains. In the same time, induced drag has risen.

Lift	**-0.8 10^{-4}**
Pitching moment	**-1.10^{-4}**
Near-field drag (in d.c.)	**-0.60**
wing	-0.70
nacelle	-0.32
pylon	+0.13
inner nacelle	0.00
engine cowl	-0.07
hot nozzle and spinner	+0.34
Far-field drag (in d.c.)	**-0.18**
wave	-0.12
viscous pressure	+0.01
induced	+0.08
engine through flow	-0.10
friction	-0.04
spurious	-0.42

Table 4 : performance variations

Unfortunately, due to the coarse mesh used, a large level of spurious drag is generated and varies slightly when the shape is deformed. Regarding the small order of magnitude of the searched performance improvement, this spurious drag variation strongly hinders the accuracy of the drag assessment during the optimisation process.

5.6. Discussion

This work was one of the first attempts to perform an optimisation based on gradient provided by adjoint method on a large realistic complex case, modelled by RANS equations. It confirmed that this method is especially well suited to cases with a large number of parameters: beside the savings in CPU time, the touchy choice of a finite difference step for each variable is avoided.

The reduction of drag by 0.5 drag counts was in line with expected order of magnitude, according to designers' experience. This is however a very small amount with respect to the mesh accuracy, that was severely constrained. This explains the high variation level of spurious drag, that accounts for two third of drag gains. The interpretation of these results must therefore be made precociously as they may not be fully meaningful from a physical point of view.

To go around this problem it is recommended for future work to use far-field drag as an objective function, which requires deriving associated sensitivities in the post-processing. Another identified way of progress concerns the shape parameterisation. The concept of local surface mesh deformation thanks to bumps is well suited for 2D cases but its interest in 3D remains limited, especially when small radius of curvature are concerned.

4. Conclusion

Three applications of aerodynamic optimisation performed with the *elsA* adjoint software have been presented in this paper, covering a wide range of complexity (two- and three-dimensional), and physics (Euler and Navier-Stokes equations, transonic and supersonic flows).

The expected benefits from the adjoint approach, compared to the finite-differences approach could be confirmed, consistent evaluations of the gradient being obtained at a cost which does not increase with the number of design variables. The present results indicate that the adjoint approach in *elsA* is mature enough to be applicable to complex industrial aircraft design problem. Complementary validations of this new technique are nonetheless needed including careful verifications of adjoint-based gradients accuracy and validations for low-speed flows, among others. Efforts to further improve the performance and robustness of the adjoint calculations will also be beneficial to accompany the expectable generalisation of its use in the near future.

Finally, to fully benefit from the adjoint technique in optimisation applications, each module of the optimisation loop needs to be completely differentiated. This calls for a rational optimisation software architecture and will require additional efforts for the pre- and post-processing modules, especially to allow the use of complex optimisation functions.

Acknowledgements

The present work has been supported by the French DPAC and SPAe agencies. The authors would like to thank the European Commission for funding this research work. The helpful technical support of AIRBUS-France is also acknowledged.

References

[1]L. Cambier and M. Gazaix
elsA: An Efficient Object-Oriented Solution to CFD Complexity
40th AIAA Aerospace Science Meeting and Exhibit, Reno, Jan. 2002.

[2]G. Carrier
Multi-disciplinary Optimisation of a Supersonic Transport Aircraft Wing Planform
ECCOMAS 2004, European Congress on Computational Methods in Applied Sciences and Engineering, Jyväskylä, 24 - 28 July 2004.

[3]P.A. Henne, R.D. Gregg
A New Airfoil Design Concept
Journal of Aircraft, Vol.28, No 8, pp.333-345, May 1991.

[4]R.M. Hicks, P.A. Henne
Wing Design by Numerical Optimisation
J. Aircraft 15, pp. 407-412, 1978.

[5]A. Jameson
Aerodynamic Design Via Control Theory
Journal of Scientific Computing, Vol 3. , p 233-260, 1988.

[6]D. A. Lovell
Reduction of Wave and Lift-Dependant Drag for Supersonic Transport Aircraft – EUROSUP
DERA Report DERA/MSS5/CR990055/1.0, 14/11/2000.

[7]J. Peter
Discrete Adjoint Techniques in elsA (part 1): Method/Theory
ONERA-DLR Aerospace Symposium (ODAS), Toulouse, October 2006.

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

[9]G. van Rossum
Python Reference Manual.
Centrum voor Wiskunde en Informatica (CWI), Report number CS-R9525, 1995.

[10]G.N. Vanderplaats
CONMIN – A Fortran Program for Constrained Function Minimization – User's Manual.
NASA TM X 62.282, August 1973.

[11]J. van der Vooren, D. Destarac
Drag / Thrust Analysis of Jet-Propelled Transonic Transport Aircraft: Definition of Physical Drag Components
Aerospace Science and Technology, Vol. 8, No. 7, October 2004.

[12]http://endo.sandia.gov/DAKOTA/software.html