Received May 2000
Revised October 2000
Accepted October 2000

Influence from numerical noise in the objective function for flow design optimisation

Jörgen Burman and B. Rikard Gebart
Division of Fluid Mechanics, Luleå University of Technology, Luleå, Sweden

Keywords *Numerical noise, Design optimization, Grids*

Abstract *The overall pressure drop in an axisymmetric contraction is minimised using two different grid sizes. The transition region was parameterised with only two design variables to make it possible to create surface plots of the objective function in the design space, which were based on 121 CFD calculations for each grid. The coarse grid showed to have significant numerical noise in the objective function while the finer grid had less numerical noise. The optimisation was performed with two methods, a Response Surface Model (RSM) and a gradient-based method (the Method of Feasible Directions) to study the influence from numerical noise. Both optimisation methods were able to find the global optimum with the two different grid sizes (the search path for the gradient-based method on the coarse grid was able to avoid the region in the design space containing local minima). However, the RSM needed fewer iterations in reaching the optimum. From a grid convergence study at two points in the design space the level of noise appeared to be sufficiently low, when the relative step size is* 10^{-4} *for the finite difference calculations, to not influence the convergence if the errors are below 5 per cent for this contraction geometry.*

Introduction

Automatic design optimisation has become increasingly used over the past years, mainly because of the recent advances in efficient and robust computational algorithms for fluid applications and the fast development in computer resources.

The optimisation methods can be divided into several categories; e.g. stochastic methods and gradient-based methods. Stochastic methods, e.g. genetic algorithms, are global optimisation methods that search the entire design space. These methods start with an initial family of solutions and use a random element in the evolution of the design. Poorer and unfeasible new designs are allowed in the short term to avoid local minima.

Because each individual CFD analysis is expensive to perform, gradient-based optimisation algorithms are often used, which require relatively few analyses compared to stochastic methods assuming the number of design parameters and the number of active constraints is not too large. These methods start from an initial design and use the gradient to determine the search direction. A one-dimensional optimisation in this direction gives a new

The financial support of the Swedish National Board for Industrial and Technical Development (NUTEK) and the Polhem Laboratory at the Luleå University of Technology is gratefully acknowledged. We would also like to acknowledge the importance of many fruitful discussions with Mr Hans Mårtensson, Volvo Aero Corporation.

International Journal of Numerical Methods for Heat & Fluid Flow, Vol. 11 No. 1, 2001, pp. 6-19.

point in the design space from where a new search direction may be calculated. Some more efficient methods use an approximation of the Hessian matrix of second derivatives. The gradient can be derived in a number of different ways, where the most straightforward is to use finite differencing. If one-sided differencing is used N additional CFD calculations are required for every iteration where N is the number of design variables. Another way to compute the gradient is to solve the adjoint equation which can be derived from a Lagrange formulation of the optimisation problem (Jameson *et al.*, 1998). With this approach one only has to solve one finite difference equation to get the gradient independently of the number of design variables. The computational cost for solving this adjoint equation is comparable with one flow calculation.

A problem with gradient-based methods compared to stochastic methods is that these methods are not global optimisation methods and are therefore sensitive to numerical noise in the objective function. The occurrence of numerical noise depends on the discretisation of the partial differential equations and causes two problems (Giles, 1997). The first is that as the design changes the grid usually does not vary continuously, particularly if unstructured grids are used. This will produce small discontinuities in the objective function. The second problem is inadequately resolved flow features caused by a too coarse grid or discontinuities in the solution (such as shocks) (Narducci *et al.*, 1994), which will give rise to a small amplitude "wave-pattern" in the design space superimposed on the objective function hypersurface.

For most shape optimisation problems it is not possible to be close to grid independent solutions since each flow calculation is very time consuming, therefore one has to consider the numerical noise in the objective function. One way to reduce the influence from this numerical noise is to construct a response surface of the objective function. A common response surface is obtained from a least-square fit of a quadratic polynomial to discrete values of the goal function.

The present paper considers a shape optimisation of an axisymmetric contraction where the contraction is designed to minimise the overall pressure drop. The optimisation is performed with two different grid sizes, one coarse grid with significant numerical noise in the objective function and one finer grid with less numerical noise. Two optimisation techniques, a response surface model and a gradient-based method (the method of feasible directions), are used to study how the solution is affected by the numerical noise.

Computational details

The geometric dimensions of the axisymmetric contraction were taken from Wang *et al.* (1996) and appear in Figure 1. A commercial code (AEA-CFX) was used for solving the flow problem (AEA Technology, 1999). It is a finite-volume based code using a structured non-staggered multi-block grid. All terms in all equations were discretisised using second-order centred differencing apart from the convective terms. The convective terms in all equations were discretised using hybrid differencing, which is a modification of upwind

differencing where central differencing is used if the mesh Peclet number is less than two, and upwind differencing is used if the mesh Peclet number is greater than two. This method is the default method in CFX due to its positive impact on convergence. The expected overall behaviour of the numerical scheme is somewhere between first and second order accuracy. The turbulence was modelled using the standard k-ε model with wall function boundary condition. The inlet conditions required for this model are the turbulence intensity and the turbulence length scale. The turbulent quantities k and ε at the inlet are then defined by (AEA Technology, 1999):

$$k_{inl} = 1.5(iu_{inl})^2, \tag{1}$$

$$\varepsilon_{inl} = \frac{k_{inl}^{1.5}}{0.3D}, \tag{2}$$

where u_{inl} is the mean inlet velocity, i is the turbulence intensity and D is the dissipation length scale. The turbulence intensity was set to 5 per cent and the dissipation length scale was set to equal the inlet diameter of the contraction. At the inlet a constant axial velocity profile (plug profile) was also specified. For the outlet the pressure was set to a constant value and for all other variables a zero normal gradient was adopted which means that the flow is assumed to be fully developed. The distance from the wall to the centre of the first cell was monitored, and the grid for the initial design in the optimisation was adjusted to have a non-dimensional wall distance for the wall cells of $y^+ \approx 70$, which is in the recommended region $30 < y^+ < 100$ (Hallbäck *et al.*, 1996). The Reynolds number based on the inlet diameter is taken to be 20,000. For the coarse grid optimisation we used a mesh consisting of 84 cells and for the less coarse grid optimisation 1,440 cells were used. These grid sizes are very coarse by today's standard but it serves the purpose of this study since it brings out interesting features of the objective function. Due to convergence problems the outlet was extended from L3 to 3·L3 for the finer grid. Of these 1,440 cells 880 are located in the part of the geometry shown in Figure 1.

For the optimisation we use a commercial design exploration system called iSIGHT (Engineous Software, 1999). iSIGHT contains a number of different optimisation algorithms, e.g. gradient methods, a genetic algorithm, a simulated annealing algorithm, and is easy to couple to different simulation tools through a visual programming language.

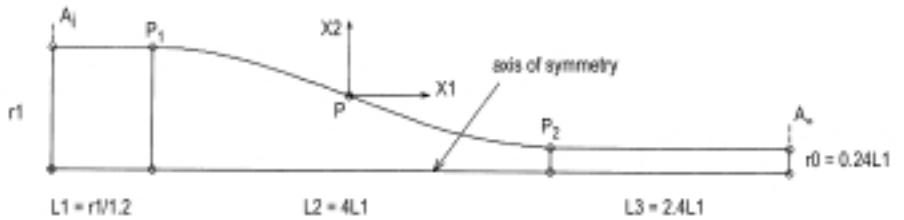

Figure 1.
Initial contraction design with design variables X1 and X2

Geometry parameterisation

The transition region is described by a spline curve where the tangent vectors of the curve at the start point (P_1) and end point (P_2) are specified to be in the axial direction (see Figure 1). The design parameters are the axial and the radial position of a control point (P) on the spline curve. Both the spline curve and the mesh are created by the pre-processor CFX-Build (AEA Technology, 1999). This simple geometry parameterisation will limit the possible shapes of the transition region, but this is of less interest since the intention is not to find the optimal contraction design but to study the influence of numerical noise in the objective function during the optimisation. The parameterisation with only two design variables was chosen because it is then possible to make a surface plot of the objective function in the design space. However, it is believed that an extension of the conclusions from this study will still be valid in a multidimensional design space.

Optimisation methods

A response surface is used to approximate the objective function and constraints. The optimisation is then performed with this approximation and the response surface is updated with a CFD analysis in the proposed optimum. The procedure is repeated until convergence to the optimum of the objective function is reached. The most common approximation is to use a Taylor series expansion (Vanderplaats, 1999):

$$F(\mathbf{X}) = F(\mathbf{X}^0) + \nabla\mathbf{F}(\mathbf{X}^0) \cdot \delta\mathbf{X} + \frac{1}{2}\delta\mathbf{X}^T\mathbf{H}(\mathbf{X}^0)\delta\mathbf{X} + ..., \tag{3}$$

$$\delta\mathbf{X} = \mathbf{X} - \mathbf{X}^0, \tag{4}$$

where $\mathbf{X}$ is a vector with design variables, $\mathbf{X}^0$ is the best design point calculated so far, $\nabla\mathbf{F}(\mathbf{X}^0)$ is the gradient and $\mathbf{H}(\mathbf{X}^0)$ is the Hessian matrix. The gradient and the Hessian are normally only approximated by a fit to randomly spaced design points around $\mathbf{X}^0$. This means that equation (3) will be a polynomial approximation of the objective function. The Response Surface Models (RSM) in iSIGHT (Engineous Software, 1999) use polynomials of first order or second order depending on the amount of data available. The first step in the process is to initialise the RSM. For a first order approximation N+1 designs need to be specified and for a full quadratic approximation (N+1)(N+2)/2 designs are needed, where N is the number of design variables. The initialisation is made by selecting a start design and specifying a fraction of the design space in which initial random designs are generated. There are several methods available based on classical design of experiments and optimal design of experiments (Giunta *et al.*, 1994) that will produce a better quality of fit than random point selection, but when only two design parameters are used is the initialisation technique of less importance.

For the optimisation of the contraction we start with a linear approximation. When the RSM has been initialised $\nabla\mathbf{F}$ can be determined by solving a linear system of N+1 equations. The Sequential Quadratic Programming technique SQP-DONLP is then used to find the minimum of this linear approximation, with move limits on the design variables. A CFD analysis is performed with this optimal design and the approximation is then updated to include one of the terms in the Hessian matrix $\mathbf{H}$. After (N+1)(N+2)/2 CFD analysis all terms in the Hessian matrix $\mathbf{H}$ are included in the approximation. For additional design points least squares fit will be used because the system of linear equations for solving $\nabla\mathbf{F}$ and $\mathbf{H}$ will then be overdetermined.

The performance of the approximation model is estimated using the so-called Trust Region Ratio (Engineous Software, 1999):

$$TRR = \frac{F_1 - F_2}{A_1 - A_2}, \tag{5}$$

where F_1 and F_2 are the values of the objective function (from the CFD analyses) between two iteration cycles, A_1 and A_2 are the approximate values. If TRR = 1 there is a perfect performance of the approximation model, if TRR = 0 there is no improvement between two iteration cycles and if TRR < 0 the new design is worse than the initial one. The value of TRR is used to decide if the move limits should be increased, reduced or unchanged. If the value of TRR is between 0.25 and 0.75, the move limits are unchanged. If the value is larger than 0.75, the move limits are increased by 50 per cent and if the value is below 0.25, the move limits are reduced by 50 per cent.

The second method used was the Method of Feasible Directions (Engineous Software, 1999) which is a direct numerical optimisation technique. With this method, we first find a search direction $\mathbf{S}$ and then move in this direction to find a new design point according to:

$$\mathbf{X}^q = \mathbf{X}^{q-1} + \alpha \cdot \mathbf{S}^q, \tag{6}$$

where q is the iteration number and the scalar α is determined by a one-dimensional search. The method to define the search direction $\mathbf{S}$ depends on whether any constraints are violated. If not, a conjugate direction is used as the search direction:

$$\mathbf{S}^q = -\nabla\mathbf{F}(\mathbf{X}^{q-1}) + \beta_q \mathbf{S}^{q-1}, \tag{7}$$

where the scalar β_q is defined as:

$$\beta_q = \frac{\left|\nabla\mathbf{F}(\mathbf{X}^{q-1})\right|}{\left|\nabla\mathbf{F}(\mathbf{X}^{q-2})\right|}. \tag{8}$$

If one or more constraints are violated the search direction $\mathbf{S}$ is found by minimising:

$$\nabla \mathbf{F}(\mathbf{X}^{q-1}) \cdot \mathbf{S}^q - \Phi\beta, \tag{9}$$

subject to:

$$\begin{aligned} &\nabla \mathbf{g}_j(\mathbf{X}^{q-1}) \cdot \mathbf{S}^q + \theta_j\beta \leq 0 \quad j \in J, \\ &\mathbf{S}^q \cdot \mathbf{S}^q \leq 1, \end{aligned} \tag{10}$$

where Φ is a large positive number, J is the set of active and violated constraints and θ_j is a push-off factor for constraints. If some constraints are active but none are violated, θ_j in equation (10) and β in equation (9) are equal to zero. $\nabla \mathbf{F}$ in equations (7-9) and $\nabla \mathbf{g}$ in equation (10) are calculated using forward finite differencing:

$$\frac{\partial \mathbf{F}}{\partial X_i} = \frac{\mathbf{F}(\mathbf{X} + \varepsilon X_i \hat{\mathbf{e}}_i) - \mathbf{F}(\mathbf{X})}{\varepsilon X_i} \tag{11}$$

where $\hat{\mathbf{e}}_i$ is a vector whose elements are zero except for the ith element which is unity and ε is the relative finite difference step size. There should be no summation over indices i in equation (11). Calculating two points in the search direction, besides the start point of the iteration, and finding the minimum of the quadratic polynomial passing through these three points then performs the one-dimensional search.

The objective function for the optimisation of the contraction is the non-dimensional pressure coefficient (C_p) defined by:

$$F(p(\mathbf{X})) = \frac{1}{\frac{1}{2}\rho U^2} \left(\frac{\int_{A_i} p dA}{A_i} - \frac{\int_{A_o} p dA}{A_o} \right), \tag{12}$$

where U is the axial inlet velocity, A_i and A_o are the cross sections located at the inlet and the outlet of the contraction respectively (see Figure 1). In order to avoid problems with mesh generation the following side constraints on the design variables are used.

$$\begin{aligned} &2.0 \leq X1 \leq 3.0 \\ &0.4 \leq X2 \leq 0.8. \end{aligned} \tag{13}$$

Figure 2 shows the rectangular box formed by the side constraints and the contraction shape at the corner points in the design space.

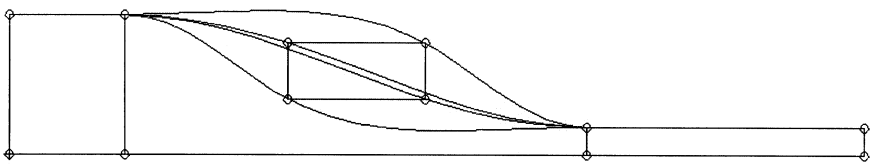

Figure 2. Contraction shapes at the corner points of the rectangular box formed by the side constraints

Results

Before the optimisation of the contraction shape, the pressure drop for the whole design space was calculated from 121 CFD calculations for each grid size. The design points were distributed in a 11 × 11 rectangular grid. Thus, it is possible to create a surface plot of the pressure drop. Figure 3 shows the objective function surface in the design space for the two different grid sizes. As can be seen from this figure, there is a large number of local minima in the objective function for the coarse grid. When the grid size is refined it is not possible to detect any numerical noise, which indicates either that the grid size is adequate to avoid the numerical noise or that the wave length of the "wave pattern" is to small to be resolved by the 11 × 11 design points. Notice the "valley" along a line approximately between the points [X1, X2] = [2.0, 0.6] and [X1, X2] = [2.7, 0.4]. The slope of this line is small as the contraction shapes of the design point along this line, and along other parallel lines, are very similar. The optimum for the coarse grid is estimated by inspection of the 121 CFD calculations to be at [X1, X2] = [2.0, 0.60] and the optimum for the finer grid is estimated to be at [X1, X2] = [2.0, 0.64], where the pressure coefficient is 771.960 and 766.577 respectively.

Coarse mesh results

As mentioned earlier the optimisation starts with a linear approximation of the objective function and as more design points become available the approximation model is updated to a full quadratic polynomial approximation. Table I and Figure 4a show the iteration history where the initial design is [X1, X2] = [2.5, 0.6]. Here, only the outer iterations are shown, i.e. the design points that involve CFD calculations, as the function evaluation time for a point on the response surface is negligible compared to a CFD calculation. The large dots are the design points needed to initialise the RSM. The termination criterion on successive iteration changes in the objective function was set to 10^{-4}. The best design is iteration six with [X1, X2] = [2.0, 0.6138] and

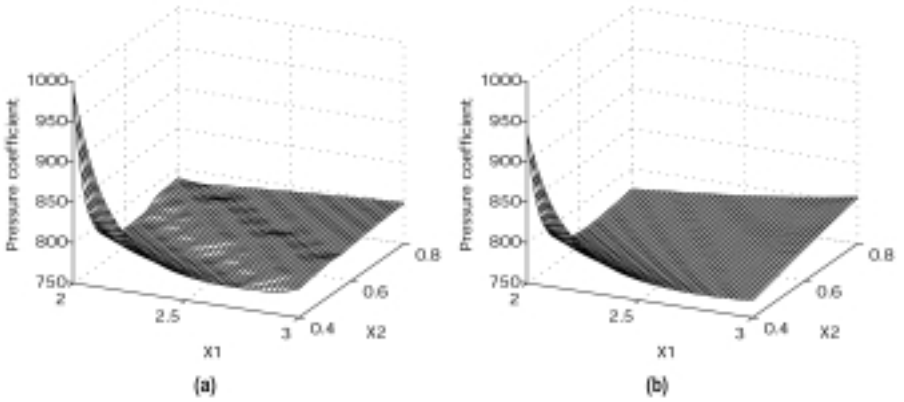

Figure 3.
Numerical noise in the objective function due to discretisation effects for two different grid sizes

C_p = 772.067 to be compared with the estimated optimum at [X1, X2] = [2.0, 0.60] with C_p = 771.960.

To improve the result of the optimisation a restart was made with [X1, X2] = [2.0, 0.6138] as start design and the five closest design points were used to initialise the RSM. Convergence closer to the true optimum was now reached within two iterations, where [X1, X2] = [2.0, 0.5997] and C_p = 771.961. This indicates that the first response surface was too much influenced by the design points far away from the optimum. A total of 11 CFD calculations were needed to reach the final optimum (nine calculations before restart of the algorithm).

The optimisation on the coarse grid was also performed with the Method of Feasible Directions. The initial design was [X1, X2] = [2.5, 0.6] and the relative perturbation parameter (for the gradient calculation was set to 10^{-4}. Figure 4b shows the iteration history where the points connected with broken lines are the outer iterations, the remaining points belong to the gradient calculations and the points needed for the one-dimensional searches. The global optimum [X1, X2] = [2.0, 0.6019] with pressure coefficient equal to 771.955 was reached in 12 outer iterations, which includes totally 63 CFD calculations. The local minima are avoided by the gradient search since most of them are located at the upper right part of the design space and the search path avoids that region. Further optimisations with the gradient method using different start designs showed no problems in reaching the global optimum, since the optimiser always found the "valley" containing the global optimum after the first outer iteration. The convergence criterion on successive iterations was set tighter for the gradient method, 10^{-7}, to avoid false convergence in the flat "valley". The RSM was not able to come as close to the true optimum as the gradient method. However, the difference between the optimum pressure coefficient for the two optimisation techniques was insignificant compared to the improvement of the pressure coefficient from the initial design.

Fine mesh results

Also for the fine grid the contraction shape was optimised with both RSM and the Method of Feasible Directions. The same start design and convergence criteria were used. Table II and Figure 5a show the iteration history for the

Iteration	Design variable 1 (X1)	Design variable 2 (X2)	Pressure coefficient
1	2.5000	0.6000	787.393
2	2.3739	0.6502	787.638
3	2.0000	0.5134	784.573
4	2.0000	0.4500	839.606
5	2.9375	0.7050	797.200
6	2.0000	0.6138	772.067
7	2.0000	0.5772	772.557
8	2.0000	0.5795	772.444
9	2.0000	0.5800	772.411

Table I. Iteration history for coarse grid optimisation using RSM

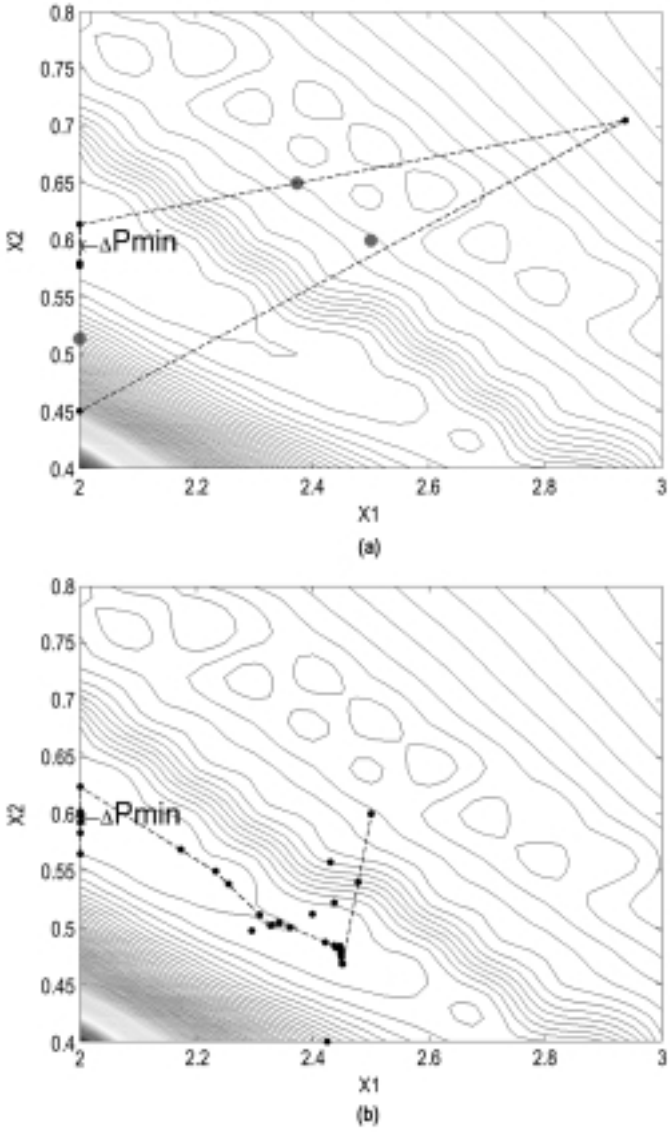

Figure 4.
Iteration history for (a) coarse grid optimisation using RSM; and (b) method of feasiable directions

optimisation using RSM. The three large dots in Figure 5a are the design points needed to initialise the RSM. Convergence to the true optimum was not reached without a restart in this case either. The best design was iteration six with [X1, X2] = [2.000, 0.6361] and C_p = 766.514. After a restart using this point and the five closest points, three more iterations gave [X1, X2] = [2.000, 0.6225] with C_p = 766.441.

Since the numerical noise is less for the fine mesh, the optimisation using the method of feasible directions showed, as expected, no problems in finding the

Iteration	Design variable 1 (X1)	Design variable 2 (X2)	Pressure coefficient
1	2.5000	0.6000	773.331
2	2.3739	0.7003	777.068
3	2.0000	0.4268	860.693
4	3.0000	0.7500	803.879
5	2.5865	0.7050	783.608
6	2.0000	0.6361	766.519
7	2.0000	0.7141	769.512
8	2.0000	0.6639	767.195
9	2.0000	0.6658	767.262
10	2.0000	0.6629	767.162
11	2.0000	0.6548	766.906
12	2.0000	0.6492	766.761
13	2.0000	0.6489	766.755

Table II.
Iteration history for fine grid optimisation using RSM

global optimum (see Figure 5b). The design point [X1, X2] = [2.000, 0.6236] with C_p = 766.440 was reached in nine outer iterations, including 50 CFD calculations.

Also with this grid size, considerably more iterations were needed to reach the optimum design using the gradient-based method. It can also be noticed that the difference in minimum pressure coefficient obtained by the two optimisation strategies is very small.

Grid convergence error

Numerical noise in the objective function will always be present to some extent. The results above for coarse and fine mesh indicate that some of the noise may be associated with poor mesh resolution. It is therefore of interest to try to estimate the magnitude of the numerical errors. Hopefully, with more experience, it may be possible to determine acceptable levels of errors for which it may be assumed that numerical noise is negligible in an optimisation procedure. The first step in a design optimisation effort would then be to estimate the required mesh size for errors below the "threshold" value that should be used for the remainder of the simulations.

Richardson extrapolation has been used to estimate the grid convergence error according to the method presented by Celik and Zhang (1995) where the order of the numerical scheme and the extrapolated value can be derived from:

$$\frac{\phi_{\alpha_2 h} - \phi_{\alpha_3 h}}{\phi_{\alpha_1 h} - \phi_{\alpha_2 h}} = \frac{\alpha_3^p - \alpha_2^p}{\alpha_2^p - \alpha_1^p}, \tag{14}$$

$$\phi_{extrapolated} = \frac{\alpha_2^p \phi_h - \phi_{\alpha_2 h}}{\alpha_2^p - 1}, \tag{15}$$

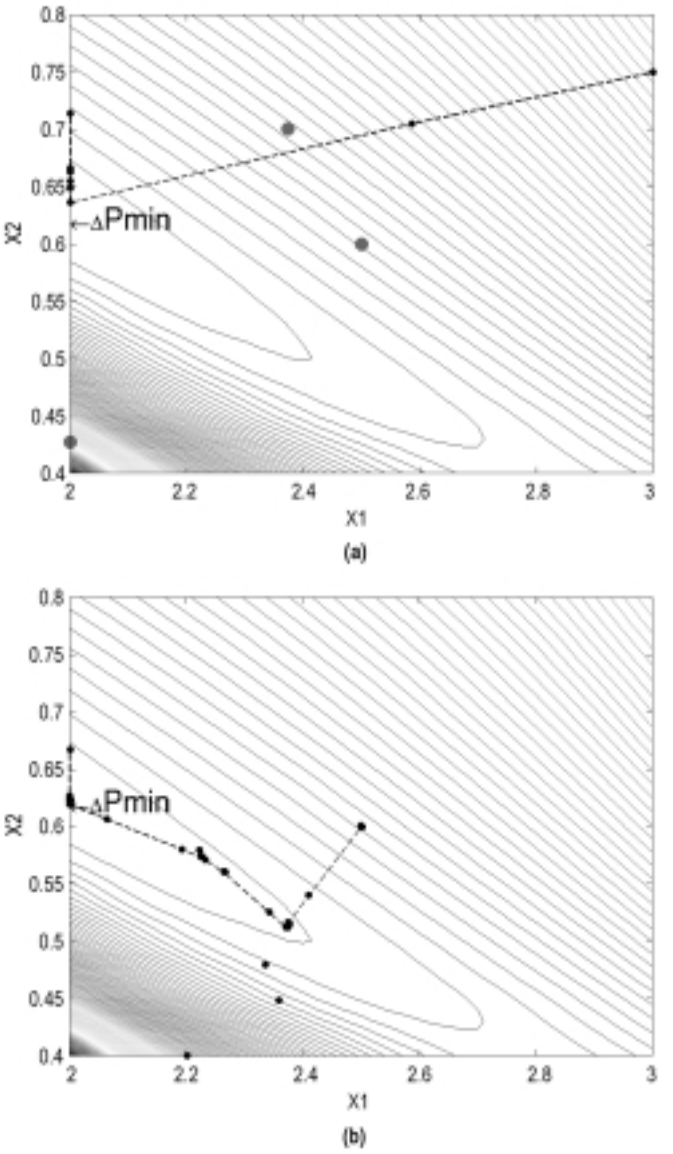

Figure 5.
Iteration history for (a) fine grid optimisation using RSM; and (b) Method of Feasible Directions

where h is the grid cell size, p is the order of the scheme and α is the grid refinement factor (the square root of the grid cell ratio between the finest grid and the present grid). The subscript on α (1, 2 or 3) indicates the grid level. In equation (15) α_1 has been assumed to be 1 while α_2 and α_3 are larger (fewer grid points). Equation (14) is first used to make sure that $1 \leq p \leq 2$, which should be the case for a hybrid differencing method when the solutions are in the asymptotic range (Bergström and Gebart, 1999). When the actual order of the method is known equation (15) can be used to get an approximation of the

exact value. Using this extrapolated value an approximate relative grid convergence error is then defined as (Celik and Zhang, 1995):

$$e_{r,approx} = \left| \frac{\phi_{extrapolated} - \phi_h}{\phi_{extrapolated}} \right|, \quad (16)$$

where ϕ_h is the value from a grid having grid cell size h.

The pressure coefficient was used to check the grid convergence at two points in the design space, one point in the middle of the design space (design point 1), i.e. [X1, X2] = [2.5, 0.6], and the other point close to the optimum (design point 2), at [X1, X2] = [2.0, 0.6]. The pressure coefficient and the maximum y^+ value for the two design points appear in Table III for three different grids with grid refinement factor one, two and four. The grid points in the table include the extended outlet with length 3·L3. Equations (14) and (15) now give $p = 1.40$ and $\phi_{extrapolated} = 736.477$ for design point 1 and $p = 1.51$ and $\phi_{extrapolated} = 737.997$ for design point 2.

By using equation (16), the estimated relative errors in the pressure coefficient for the coarse grid were calculated to 6.9 per cent and 4.7 per cent, for design points 1 and 2, respectively. The corresponding relative errors for the finer grid were 5.0 per cent and 3.9 per cent.

Thus, it appears in the present case with a contraction geometry to be sufficient if the errors are below about 5 per cent, when the relative step size is 10^{-4} for the finite difference calculations, to reduce the level of noise sufficiently to not influence the convergence, i.e. get caught in a local optimum.

Conclusions

The shape of an axisymmetric contraction has been optimised with two different grid sizes to study how the solution is affected by the numerical noise in the objective function. A surface plot of the objective function based on 121 CFD calculations in the design space showed a large number of local minima for the coarse grid. The global optima for the two grid sizes were close to each other in the design space which indicates that the optimisation could be started with the coarse grid and to refine the grid as the optimiser moves closer to the optimum.

The RSM and the Method of Feasible Directions were able to find the global optimum with both the coarse grid and the fine grid. However, the RSM needed much fewer iterations in reaching the optimum. There is also a risk of getting

No. of grid points	Design point 1 [X1, X2] = [2.5, 0.6] Pressure coefficient	y^+	Design point 2 [X1, X2] = [2.0, 0.6] Pressure coefficient	y^+
2,304 ($\alpha_3 = 4$)	751.383	77.3	749.903	75.9
9,216 ($\alpha_2 = 2$)	742.126	70.8	742.181	77.0
36,864 ($\alpha_1 = 1$)	738.618	69.3	739.467	70.0

Table III. Pressure coefficient and maximum non-dimensional wall distance (y^+) for three different grids at two design points

caught in a local optimum with the gradient-based method, although this was avoided here because the "valley" containing the global optimum, which was without local minima, was reached immediately, independently of the start design. The restarts made with the RSM did not seem to be necessary, although the design parameters changed, since the objective function showed little improvement.

According to Giles the problem with numerical noise associated with the discretisation of the partial differential equations may not become significant until one is very close to the optimum solution, in which case it is not important since the solution is sufficiently close to optimal (Giles, 1997). However, for this contraction optimisation problem the numerical noise is distributed in the whole design space and the solution at a local optimum can therefore differ significantly from the solution at the global optimum.

From the grid convergence study it is difficult to suggest a limit on the error that will cause the optimisation to fail since the optimum was found both on the coarse and the fine grids. However, it appears for this contraction geometry to be sufficient if the errors are below 5 per cent with the relative step size of 10^{-4} for the finite difference calculations. The noise level is then reduced enough to avoid getting caught in a local optimum. This error level is probably dependent on the choice of objective function. The pressure drop used here is related partly to the shear stress acting tangent to the contraction wall and partly to changes in dynamic pressure. The region with high shear stress is confined to the boundary layer near the wall and the pressure drop will therefore be sensitive to variations in the non-dimensional wall distance as a consequence of geometry modifications. Calculations with the same geometry but different values on maximum y^+ (between 70 and 90) showed significant changes in pressure drop. The variations in the non-dimensional wall distance as the geometry changes during the optimisation will contribute to the numerical noise. Hence, there is a strong need for grid independent wall functions or turbulence models that resolve the wall layer in shape optimisation. However, the low-Reynolds models, where wall functions are not used, require in most cases too fine grid resolution close to walls for practical use in an optimisation procedure.

Seeing that already a simple problem like this one exhibits a significant influence from numerical noise makes it unreasonable to believe that the objective function of a problem with more complicated flow will be more "well-behaved". Using a response surface method to approximate the objective function seems to be a good way to minimise the influence from numerical noise.

References

AEA Technology (1999), *CFX 4.3 Flow Solver Guide User Guide*, Computational Fluid Dynamics Services, Building 8.19, Harwell Laboratory, Harwell.

Bergström, J. and Gebart, B.R. (1999), "Estimation of numerical accuracy for the flow field in a draft tube", *International Journal of Numerical Methods for Heat & Fluid Flow*, Vol. 9 No. 4, pp. 472-86.

Celik, I. and Zhang, W.M. (1995), "Calculation of numerical uncertainty using Richardson extrapolation: application to some simple turbulent flow calculations", *Journal of Fluids Engineering*, Vol. 17, pp. 439-45.

Engineous Software, Inc. (1999), *iSIGHT Designer's Guide 5.0*, North Carolina Research, Triangle Park, NC.

Giles, M.B. (1997), "Aerospace design: a complex task", *VKI Lecture Series, Inverse Design and Optimisation Methods*, von Karman Institute for Fluid Dynamics, Rhode-Saint-Genèse.

Giunta, A.A., Dudley, J.M., Narducci, R., Grossman, B., Haftka, R.T., Mason, W.H. and Watson, L.T. (1994), "Noisy aerodynamic response and smooth approximation in HSCT design", AIAA Paper 94-4376, *5th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, Panama City, FL.

Hallbäck, M., Henningson, D.S., Johansson, A.V. and Alfredsson, P.H. (1996), *Turbulence and Transition Modelling*, Kluwer Academic Publishers, Dordrecht.

Jameson, A., Martinelli, L. and Pierce, N.A. (1998), "Optimum aerodynamic design using the Navier-Stokes equations", *Theoretical and Computational Fluid Dynamics*, Vol. 10, Nos. 1-4, pp. 213-37.

Narducci, R., Grossman, B. and Haftka, R.T. (1994), "Sensitivity algorithms for an inverse design problem involving a shock wave", AIAA Paper 94-0096, *32nd AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV.

Vanderplaats, G.N. (1999), *Numerical Optimization Techniques for Engineering Design: With Applications*, Vanderplaats Research & Development, Inc.

Wang, Z.X, Tortorelli, D.A. and Dantzig, J.A. (1996), "Sensitivity analysis and optimization of coupled thermal and flow problems with applications to contraction design", *International Journal for Numerical Methods in Fluids*, Vol. 23, pp. 991-1020.