Superconvergent lift estimates through adjoint error analysis

M.B. Giles and N.A. Pierce

Oxford University Computing Laboratory
Oxford OX1 3QD, United Kingdom
email: *giles@comlab.ox.ac.uk*

Key Words: Error analysis, functionals, adjoint equations

Abstract. This paper introduces a new idea, using adjoint error analysis to obtain approximate values for integral quantities, such as lift and drag, which are twice the order of accuracy of the flow solution. The theory is presented for both linear and nonlinear applications and numerical results confirm the effectiveness of the technique for the one-dimensional Poisson equation and the quasi-1D Euler equations.

1 Introduction

In engineering applications of CFD, there are usually a few integral quantities of primary concern, such as lift and drag on an aircraft, total mass flux through a turbomachine, or total heat flux into a turbine blade. The rest of the flow solution is often needed only for qualitative purposes, for example to see if there is a bad flow separation. In this paper we show how the order of accuracy of an important integral quantity can be greatly improved, usually doubled, compared to the accuracy of the flow solution on which the estimate is based. This is accomplished through an error analysis using an approximate solution to the adjoint flow equations. These are the same adjoint equations that are solved to efficiently obtain the linear sensitivity of an objective function in design optimisation [Jameson (95), Jameson (97), Anderson (97), Elliott (97)], but in the present context, the adjoint variables reveal the contributions of flow solution approximation errors to the error in the computed integral. Correcting the leading order error produces a corrected value for the integral which is much more accurate.

This idea is closely related to the *a priori* and *a posteriori* analysis of the *superconvergence* of integral functionals arising from finite element computations in a variety of applications [Babuška (84), Barrett (87), Becker (96), Paraschivoiu (97), Giles (97b), Süli (97), Monk (98)]. However, with these methods the superconvergence arises naturally from Galerkin orthogonality without the addition of a correction term. Previous work by the present authors on doubling the order of accuracy of quasi-1D lift estimates obtained from a first order upwind method [Giles (98)] was based on a discrete truncation error viewpoint [Giles (97c)]. The new approach uses an analytic viewpoint which leads to a much simpler implementation when using more accurate discretisations. We are not aware of other work on the use of adjoint solutions to improve the accuracy of integral quantities through the evaluation of a correction term.

The paper begins by presenting the linear theory and numerical results for the one-dimensional Poisson equation. The nonlinear theory is then presented and applied to the quasi-1D Euler equations. Results are given for subsonic flow and transonic flow, with and without shocks. These demonstrate the effectiveness of the approach, and the paper concludes with a discussion of the challenges to be overcome in extending the technique to multi-dimensional applications.

2 Linear theory

Let u be the solution of the linear differential equation

$$Lu = f,$$

on the domain Ω, subject to homogeneous boundary conditions for which the problem is well-posed. The adjoint differential operator L^* and associated homogeneous boundary conditions are defined by the identity

$$(v, Lu) = (L^* v, u),$$

for all u, v satisfying the respective boundary conditions. Here the notation $(.,.)$ denotes an integral inner product over the domain Ω.

If we are concerned with the value of the functional $J = (g, u)$, where g is a given function defined on Ω, an equivalent dual formulation of the problem is to evaluate the functional $J = (v, f)$, where v satisfies the adjoint equation

$$L^* v = g,$$

subject to the homogeneous adjoint boundary conditions. The equivalence of the two forms of the problem follows immediately from the definition of the adjoint operator.

$$(v, f) = (v, Lu) = (L^* v, u) = (g, u).$$

Suppose that u_h and v_h are approximations to u and v, respectively, and satisfy the homogeneous boundary conditions. The subscript h denotes that the approximate solutions are derived by interpolating the results of a numerical computation using a grid with average spacing h. The functions f_h and g_h are defined by

$$Lu_h = f_h, \quad L^* v_h = g_h.$$

It is assumed that u_h and v_h are sufficiently smooth that f_h and g_h lie in $L_2(\Omega)$. If u_h and v_h were equal to u and v, then f_h and g_h would be equal to

f and g. Thus, the *residual errors* $f_h - f$ and $g_h - g$ are a computable indication of the extent to which u_h and v_h are not the true solutions.

Now, using the definitions and identities given above, we obtain the following expression for the functional:

$$\begin{aligned}(g,u) &= (g,u_h) - (g_h, u_h - u) + (g_h - g, u_h - u)\\ &= (g,u_h) - (L^* v_h, u_h - u) + (g_h - g, u_h - u)\\ &= (g,u_h) - (v_h, L(u_h - u)) + (g_h - g, u_h - u)\\ &= (g,u_h) - (v_h, f_h - f) + (g_h - g, u_h - u).\end{aligned}$$

The first term in the final expression is the value of the functional obtained from the approximate solution u_h. The second term is an inner product of the residual error $f_h - f$ and the approximate adjoint solution v_h. The adjoint solution gives the weighting of the contribution of the local residual error to the overall error in the computed functional. Therefore, by evaluating and subtracting this adjoint error term we obtain a more accurate value for the functional.

The third term is the remaining error after making the adjoint correction. If $g_h - g$ is of the same order of magnitude as $v_h - v$ then the remaining error has a bound which is proportional to the product $\|u_h - u\|\,\|v_h - v\|$ (using L_2 norms), and thus the corrected functional value is superconvergent. If the solution errors $u_h - u$ and $v_h - v$ are both $O(h^p)$ so that halving the grid spacing leads to a 2^p reduction in the errors, then the error in the functional is $O(h^{2p})$.

For simplicity of presentation, we have assumed above that the primal problem has homogeneous boundary conditions, and that the functional is simply an inner product over the whole domain and does not have a boundary integral term. More generally, inhomogeneous boundary conditions and boundary integrals in the functional are both permissible. Inhomogeneous boundary conditions for the primal problem lead to a boundary integral term for the adjoint formulation, and similarly a boundary integral in the primal form of the functional leads to inhomogeneous adjoint boundary conditions. Although the analysis is slightly more complicated, the final form of the adjoint error correction is exactly the same as before, provided the approximate solutions u_h and v_h still exactly satisfy the inhomogeneous boundary conditions. If they do not, then there is an additional correction term to take account of this error.

3 Linear example

The example is the one-dimensional Poisson equation,

$$\frac{d^2u}{dx^2} = f,$$

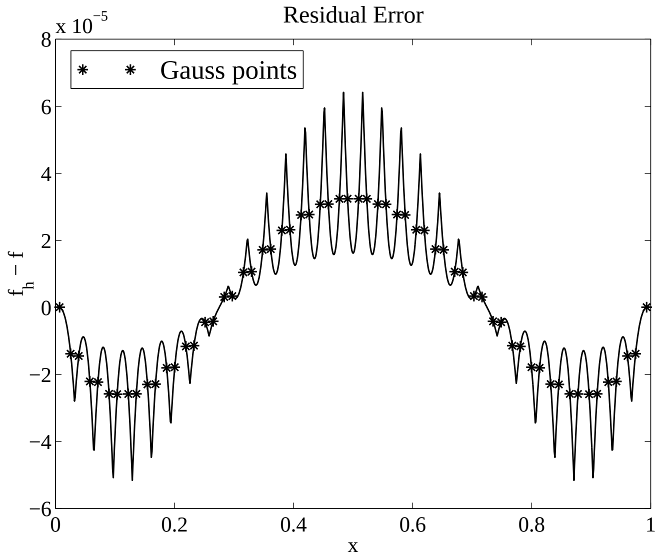

Fig. 1. Residual error for 1D Poisson equation.

on the unit interval $[0,1]$ subject to homogeneous boundary conditions $u(0)=u(1)=0$. This is approximated numerically on a uniform grid, with spacing h, using a simple second order finite difference discretisation,

$$h^{-2}\delta_x^2 u_j = f(x_j).$$

The approximate solution $u_h(x)$ is then defined by interpolation with a cubic spline through the nodal values u_j.

The dual problem is also a Poisson equation,

$$\frac{d^2 v}{dx^2} = g,$$

subject to the same homogeneous boundary conditions, and the approximate adjoint solution v_h is obtained in exactly the same manner.

Numerical results have been obtained for the case

$$f = x^3(1-x)^3, \qquad g = \sin(\pi x).$$

Figure 1 shows the residual error $f_h - f$ when $h = \frac{1}{32}$, as well as the values at the two Gaussian quadrature points on each sub-interval which are used in the numerical integration of the inner product $(v_h, f_h - f)$. Since u_h is a cubic spline, $f_h \equiv \frac{d^2 u_h}{dx^2}$ is continuous and piecewise linear. The best piecewise linear approximation to f has an approximation error whose dominant term is quadratic on each sub-interval; this explains the scalloped shape of the residual error. Figure 2 shows the approximate adjoint solution v_h, illustrating that the residual error in the center of the domain contributes most significantly to the overall error in the functional.

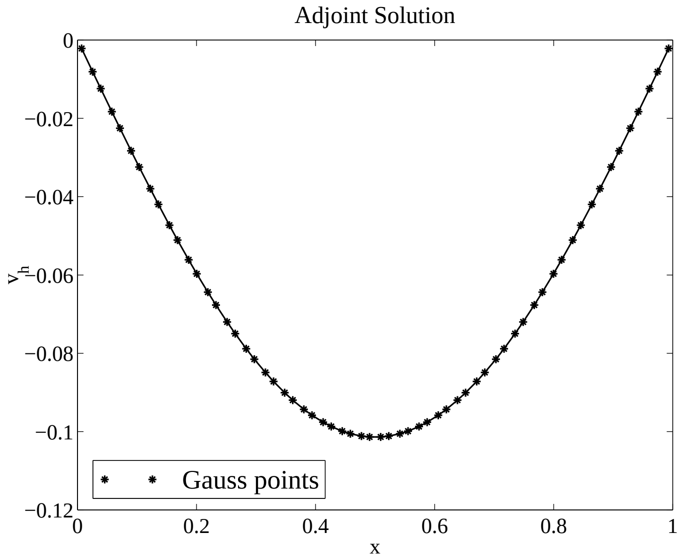

Fig. 2. Adjoint solution for 1D Poisson equation.

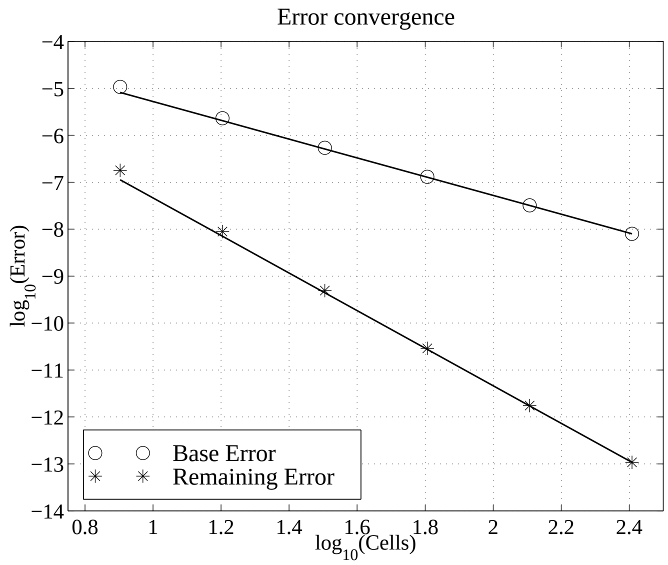

Fig. 3. Error convergence for 1D Poisson equation.

Figure 3 is a log-log plot of two quantities versus the number of cells: the error in the base value of the functional (g, u_h) and the error remaining after subtracting the adjoint error correction term $(v_h, f_h - f)$. The superimposed lines have slopes of -2 and -4, confirming that the base solution is second order accurate while the corrected functional is fourth order accurate. It is also worth noting that on a grid with 16 cells, which might be a reasonable choice for practical computations, the error in the corrected value is over 200 times smaller than in the uncorrected value.

4 Nonlinear theory

Let u be the solution of the nonlinear differential equation

$$N(u) = f,$$

on the domaon Ω subject to certain boundary conditions, and let the functional of interest, $J(u)$, be an integral over the domain of a nonlinear algebraic function of u. The linear differential operator L_u is defined to be the Fréchet derivative [Collatz (66)] of N,

$$L_u\, \tilde{u} \equiv \lim_{\epsilon \to 0} \frac{N(u+\epsilon \tilde{u}) - N(u)}{\epsilon},$$

and, similarly, the function $g(u)$ is defined by

$$(g(u), \tilde{u}) \equiv \lim_{\epsilon \to 0} \frac{J(u+\epsilon \tilde{u}) - J(u)}{\epsilon}.$$

The linear adjoint problem is

$$L_u^* v = g,$$

subject to the appropriate homogeneous adjoint boundary conditions [Giles (97)]. Now consider approximate solutions u_h, v_h which have again been obtained by interpolating the results of a finite volume calculation. The quantities f_h, g_h are defined by

$$N(u_h) = f_h, \qquad L_{u_h}^* v_h = g_h.$$

Note the use of $L_{u_h}^*$, the Fréchet derivative based on u_h which is known, instead of L_u^* based on u which is not known. In addition, the analysis requires averaged Fréchet derivatives $\overline{L}_{(u,u_h)}$ and $\overline{g}(u, u_h)$ defined by

$$\overline{L}_{(u,u_h)} = \int_0^1 L|_{u+\theta(u_h-u)}\, d\theta,$$
$$\overline{g}(u, u_h) = \int_0^1 g(u + \theta(u_h - u)) d\theta,$$

so that

$$N(u_h)-N(u) = \int_0^1 \frac{\partial}{\partial\theta} N(u+\theta(u_h-u))\,d\theta = \overline{L}_{(u,u_h)}\,(u_h-u),$$

and similarly

$$J(u_h)-J(u) = (\overline{g}(u,u_h), u_h-u).$$

Using the above definitions, we obtain the following result:

$$\begin{aligned}
J(u) &= J(u_h) - (\overline{g}(u,u_h), u_h-u) \\
&= J(u_h) - (g_h, u_h-u) + (g_h-\overline{g}(u,u_h), u_h-u) \\
&= J(u_h) - (L^*_{u_h} v_h, u_h-u) + (g_h-\overline{g}(u,u_h), u_h-u) \\
&= J(u_h) - (v_h, L_{u_h}(u_h-u)) + (g_h-\overline{g}(u,u_h), u_h-u) \\
&= J(u_h) - (v_h, \overline{L}_{(u,u_h)}(u_h-u)) + (g_h-\overline{g}(u,u_h), u_h-u) \\
&\qquad\qquad - (v_h, (L_{u_h}-\overline{L}_{(u,u_h)})(u_h-u)) \\
&= J(u_h) - (v_h, N(u_h)-N(u)) + (g_h-\overline{g}(u,u_h), u_h-u) \\
&\qquad\qquad - (v_h, (L_{u_h}-\overline{L}_{(u,u_h)})(u_h-u)) \\
&= J(u_h) - (v_h, f_h-f) + (g_h-\overline{g}(u,u_h), u_h-u) \\
&\qquad\qquad - (v_h, (L_{u_h}-\overline{L}_{(u,u_h)})(u_h-u))
\end{aligned}$$

The first term in the final result is the functional evaluated using the approximate solution u_h. The second term is the adjoint error correction term which is again an inner product of the residual error and the approximate adjoint solution. Since both of these are known, this second term can be computed and subtracted from the first to form a corrected value for the functional.

The last two terms, which cannot be computed since the analytic solution u is not known, form the remaining error in the corrected functional. If the solution error for the nonlinear primal problem and the linear adjoint problem are of the same order, and they are both sufficiently smooth that the corresponding residual errors are also of the same order, then the order of accuracy of the functional approximation after making the adjoint correction is twice the order of accuracy of the the primal and adjoint solutions on which it is based.

5 Quasi-1D Euler equations

The steady quasi-1D Euler equations in conservative form are

$$\frac{d}{dx}(AF) - \frac{dA}{dx}\,P = 0,$$

where $A(x)$ is the cross-sectional area of the duct and U, F and P are defined as

$$U = \begin{pmatrix} \rho \\ \rho q \\ \rho E \end{pmatrix}, \qquad F = \begin{pmatrix} \rho q \\ \rho q^2 + p \\ \rho q H \end{pmatrix}, \qquad P = \begin{pmatrix} 0 \\ p \\ 0 \end{pmatrix}.$$

Here ρ is the density, q is the velocity, p is the pressure, E is the total internal energy and H is the stagnation enthalpy. The system is closed by the equation of state for an ideal gas,

$$H = E + \frac{p}{\rho} = \frac{\gamma p}{(\gamma-1)\rho} + \tfrac{1}{2}q^2,$$

where γ is the ratio of specific heats.

Numerical results have been obtained using a standard second order finite volume method with characteristic smoothing on a uniform computational grid. Except when there is a shock, the approximate solution $u_h(x)$ is constructed from the discrete nodal values u_j by cubic spline interpolation of the three components of U. All other variables are then calculated from these. Evaluation of the residual error $f_h - f$ requires first derivatives of flow quantities; these are obtained by differentiating the cubic spline representation.

The linear adjoint problem is approximated by the 'continuous' method, which involves linearising the nonlinear flow equations, constructing the analytic adjoint equations, and then forming a discrete approximation to these on the same uniform grid as the flow solution [Jameson (95), Jameson (97), Anderson (97)]. An alternative approach which could have been used is the 'discrete' method in which one takes the discretised nonlinear flow equations, linearises them and then uses the transpose of the linear matrix as the discrete adjoint operator [Elliott (97)]. Previous research has shown that both approaches produce consistent approximations to the analytic adjoint solution which has been determined in closed form for the quasi-1D Euler equations [Giles (98)].

Results have been obtained for three test cases: a subsonic flow, a shock-free transonic flow with subsonic inflow and supersonic outflow, and a shocked flow with supersonic inflow and subsonic outflow. The Mach number distributions for these three cases are shown in Figure 4. In each case the functional of interest is the integral of pressure along the duct; this serves as a prototype for the lift in airfoil and aircraft calculations.

5.1 Subsonic flow

Figure 5 shows the error convergence for a subsonic flow in a converging-diverging duct. The base error, which is the error before applying the adjoint correction, is second order, as indicated by the superimposed line of slope -2. This is as expected given the second order truncation error in approximating the nonlinear flow equations. The other superimposed line of slope -4 shows that the error remaining after the adjoint correction is fourth order.

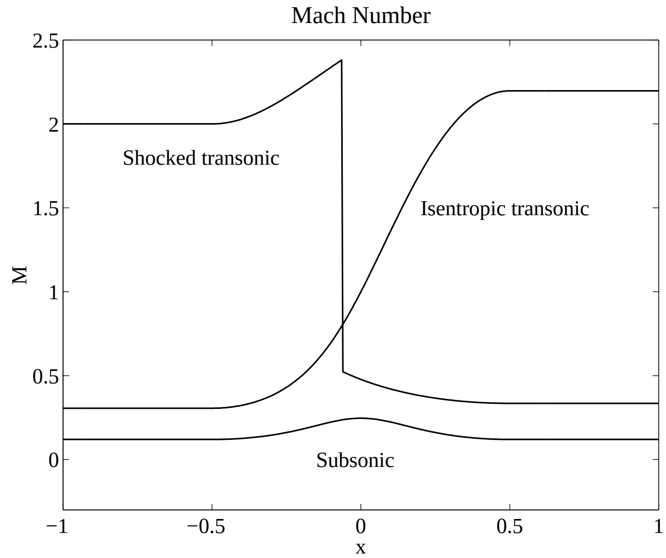

Fig. 4. Mach number distributions for quasi-1D test cases.

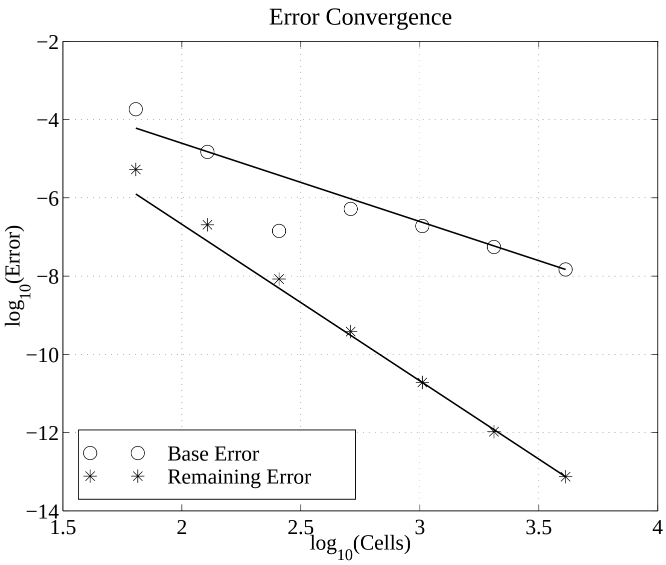

Fig. 5. Error convergence for quasi-1D subsonic flow.

5.2 Isentropic transonic flow

Figure 6 shows the error convergence for a transonic flow in a converging-diverging duct with the throat located at $x = 0$. The flow is subsonic upstream of the throat and supersonic downstream of the throat. Again the results show that the base error is second order while the remaining error after the adjoint correction is fourth order.

The accuracy of the corrected functional in this case is a little puzzling because the adjoint solution has a logarithmic singularity at the throat [Giles (98)], as shown in Figure 7. Therefore, $v_h - v$ is $O(1)$ in a small region of size $O(h)$ on either side of the throat. Based on this, one would expect that the remaining error might be $O(h^3)$ since the numerical results confirm that the residual error for the nonlinear equations is $O(h^2)$. The explanation for the fourth order convergence must lie in a leading order cancellation within the two remaining error integrals, but we do not yet have a complete understanding of this phenomenon.

5.3 Shocked transonic flow

The final example is for flow in a diverging duct, where a shock separates supersonic upstream and subsonic downstream regions. Previous research has proved that the analytic adjoint solution is continuous and has zero gradient at the shock, so the adjoint variables pose no special difficulty in this case [Giles (98)]. The challenge is the reconstruction of the approximate solution $u_h(x)$ from the nodal quantities u_j coming from the finite volume calculation.

The analytic solution is discontinuous at the shock, and satisfies the Rankine-Hugoniot shock jump relations which require that there is no discontinuity in the nonlinear flux F. The discrete solution has a slightly smeared shock, and so if one interpolates the conservative variables U it is clear that locally in a neighborhood of size $O(h)$ the error in the reconstructed solution $u_h(x)$ will be $O(1)$.

To recover a discontinuous approximate solution $u_h(x)$ we instead use the fact that F is known to be continuous at the shock and therefore choose to interpolate the nodal values of F. From these one can deduce the conservation variables U by solving a quadratic equation, one branch of which gives a subsonic flow solution, the other being supersonic. Therefore, given a shock position, one can reconstruct a supersonic solution on the upstream side, a subsonic solution on the downstream side, and automatically satisfy the Rankine-Hugoniot shock jump conditions at the shock itself. To determine the shock position, we rely on prior research [Giles (96)] which shows that the integrated pressure along the duct is correct to second order when using a finite volume method which is conservative and second order accurate in smooth flow regions. Therefore, we iteratively adjust the position of the shock until the reconstructed solution has the same base functional value

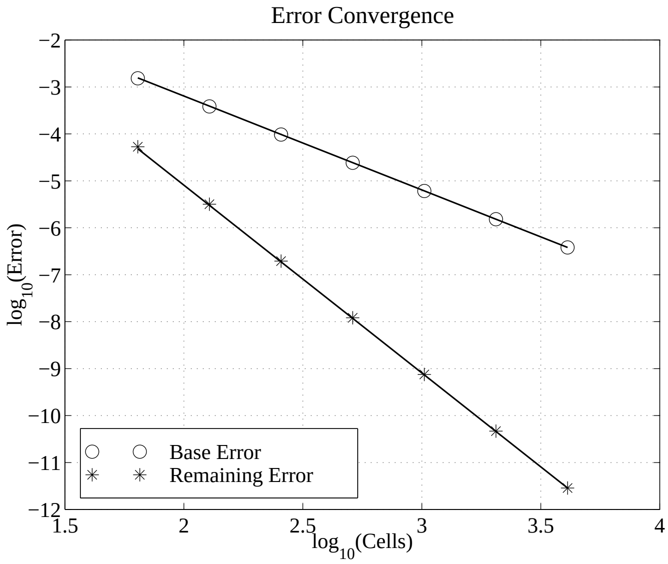

Fig. 6. Error convergence for quasi-1D shock-free transonic flow.

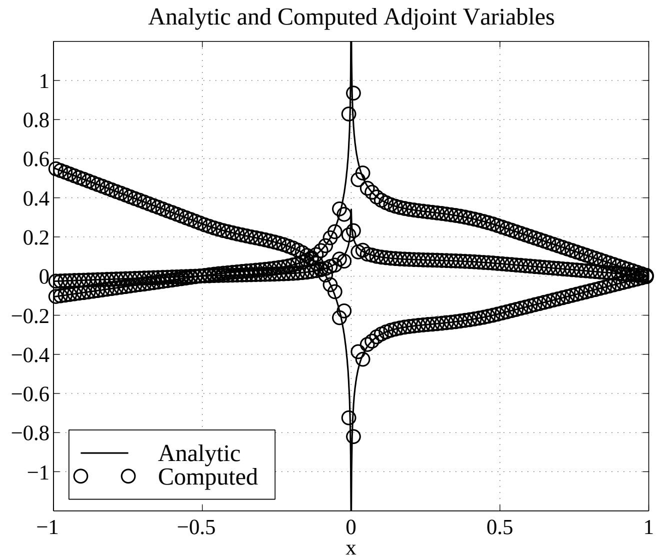

Fig. 7. Adjoint solution for quasi-1D shock-free transonic flow.

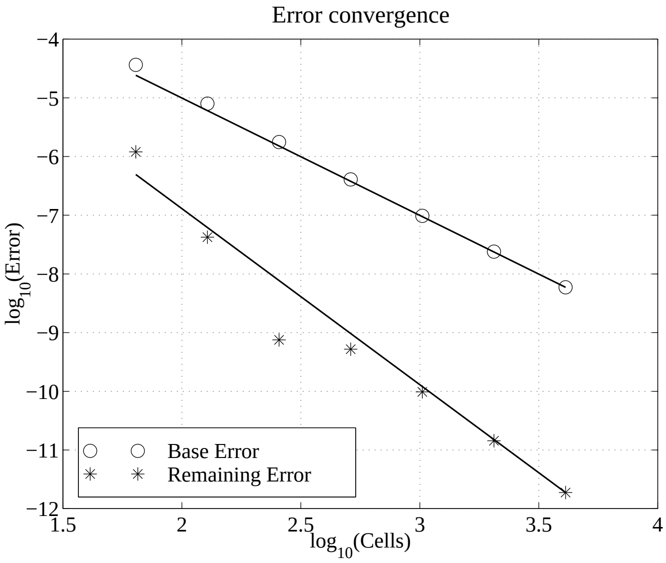

Fig. 8. Error convergence for quasi-1D shocked flow.

(i.e. without the adjoint correction) as the original numerical approximation, thereby obtaining the correct shock position to second order.

Figure 8 shows the error convergence. As expected, the base error is again second order. Because there is still an $O(h)$ error in the approximate solution $u_h(x)$ in the neighbourhood of the shock, the corrected error is now third order, not fourth. However, in future work we hope to recover overall fourth order accuracy, based on the average cell size, by using local grid adaptation at the shock.

6 Concluding remarks

In this paper we have outlined a means of calculating improved estimates of integral quantities such as lift and drag from CFD calculations, by evaluating an adjoint correction term which is an inner product of the residual error in approximating the flow equations and an approximate solution to the corresponding adjoint equations. The numerical results demonstrate the effectiveness of the technique applied to a second order finite volume approximation of the quasi-1D Euler equations. When the flow is smooth, the error in the integrated pressure is fourth order; when there is a shock, it is third order.

The theory is equally applicable to the Euler and Navier-Stokes equations in multiple dimensions. However, there are three important issues to be addressed before similar results can be obtained for airfoil and aircraft applications of engineering interest. The first is the treatment of curved surfaces; to achieve fourth order accuracy for corrected functional such as lfit and drag, it is likely that smooth curved surfaces will need to be approximated in a way which ensures continuity in the surface normal, as opposed to the use of simple linear (or bi-linear) facets. The second issue is the resolution of singularities; the adjoint flow solution in two dimensional airfoil applications has an inverse square root singularity along the incoming stagnation streamline [Giles (97)] and this will need to be well resolved. The final issue concerns unstructured grid calculations which are needed for complex applications. The approximate solution u_h needs to be sufficiently smooth that the error in ∇u_h is of the same order as the error in u_h itself. To achieve this on unstructured grids where the solution error has a significant high-frequency content may require the use of multi-dimensional smoothed cubic splines.

Another interesting direction for future research is *a posteriori* estimation of the error remaining after making the adjoint correction. The goal of such research would be to develop a mathematical framework on which one could base efficent grid refinement indicators, and thereby obtain the value of a functional to the desired level of accuracy and at a minimum computational cost.

Acknowledgments

This research was supported by EPSRC under grant GR/K91149.

References

W.K. Anderson and V. Venkatakrishnan. Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation. AIAA Paper 97-0643, 1997.

J.W. Barrett and C.M. Elliott. Total flux estimates for a finite-element approximation of elliptic equations. *IMA J. Numer. Anal.*, 7:129–148, 1987.

I. Babuška and A. Miller. The post-processing approach in the finite element method – Part 1: calculation of displacements, stresses and other higher derivatives of the displacements, *Intern. J. Numer. Methods Engrg.*, 20:1085–1109, 1984.

R. Becker and R. Rannacher. Weighted *a posteriori* error control in finite element methods. Technical report No. 96-1, Universitat Heidelberg, 1996.

L. Collatz. *Functional analysis and numerical mathematics.* Academic Press, 1966.

J. Elliott and J. Peraire. Practical 3D aerodynamic design and optimization using unstructured meshes. *AIAA J.*, 35(9):1479–1485, 1997.

M.B. Giles. Analysis of the accuracy of shock-capturing in the steady quasi-1D Euler equations. *Comput. Fluid Dynamics J.*, 5(2):247–258, 1996.

M.B. Giles and N.A. Pierce. Adjoint equations in CFD: duality, boundary conditions and solution behaviour. AIAA Paper 97-1850, 1997.

M.B. Giles, M.G. Larson, J.M. Levenstam, and E. Süli. Adaptive error control for finite element approximations of the lift and drag in viscous flow. Technical Report NA97/06, Oxford University Computing Laboratory, 1997.

M.B. Giles On adjoint equations for error analysis and optimal grid adaptation in CFD, In *Computing the Future II: Advances and Prospects in Computational Aerodynamics*, M. Hafez, editor, World Scientific, 1998.

M.B. Giles and N.A. Pierce. On the properties of solutions of the adjoint Euler equations. In *Numerical Methods for Fluid Dynamics VI*, M.J. Baines, editor, ICFD, 1998.

A. Jameson. Optimum aerodynamic design using control theory. *Comput. Fluid Dynam. Rev.*, pages 495–528, 1995.

A. Jameson, N. Pierce, and L. Martinelli, Optimum aerodynamic design using the Navier–Stokes equations. AIAA Paper 97-0101, 1997.

P. Monk and E. Süli. The adaptive computation of far field patterns by *a posteriori* error estimation of linear functionals. Technical Report NA98/02, Oxford University Computing Laboratory, 1998.

M. Paraschivoiu, J. Peraire, and A. Patera. A posteriori finite element bounds for linear-functional outputs of elliptic partial differential equations. *Comput. Methods Appl. Mech. Engrg.*, 150(1-4):289–312, 1997.

E. Süli. *A posteriori* error analysis and adaptivity for finite element approximations of hyperbolic problems. Technical Report NA97/21, Oxford University Computing Laboratory, 1997.