



AIAA 2001-2550

Solution Adaptive Mesh Refinement Using Adjoint Error Analysis

Jens-Dominik Müller, Michael B. Giles

Oxford University Computing Laboratory
Oxford, United Kingdom OX1 3QD

15th Computational Fluid Dynamics Conference
11 June – 14 June, 2001 / Anaheim, California

Solution Adaptive Mesh Refinement Using Adjoint Error Analysis

Jens-Dominik Müller*,
Michael B. Giles†

Oxford University Computing Laboratory
Oxford, United Kingdom OX1 3QD

A solution adaptive mesh refinement method for unstructured meshes is presented. The sensor is based on a smoothed estimate of the residual error, weighted by the the adjoint variables corresponding to the integral functional of most engineering interest. The difficulties in obtaining a smooth representation of the residual are discussed, and preliminary results are presented for two-dimensional inviscid subsonic and transonic test cases.

Solution adaptive mesh refinement methods refine or derefine the mesh locally according to some error criterion. Many practical algorithms have been developed for the Euler and Navier-Stokes equations on a variety of grid types.^{3,4,11-13,15,24} However, adaptation sensors are most often still based on first or second derivatives of one or more flow variables. These simple sensors have various problems. Firstly, the selected sensor may not be a suitable choice for the flowfield or the integral functional one is interested in. E.g. a velocity difference sensor may exhibit large values due to the gradients of a boundary layer, even though it is sufficiently resolved, while the velocity differences in a more interesting vortical separation structure are less pronounced. Secondly, the fact that these sensors are not weighted, or are only crudely limited geometrically, leads to wasted refinement in areas where errors are present but not relevant. E.g. a shear layer shed from an airfoil will be convected downstream of its trailing edge. At one or two chord lengths from the trailing edge a poor resolution of the shear layer does not influence the flow around the airfoil significantly and need not be resolved. Thirdly, these sensors sometimes do not converge for discontinuities. E.g. a shock will become steeper and steeper with refinement and require all the refinement resources available, while there may be more important features to resolve. This can actually lead to seemingly grid-converged solutions that are incorrect.²³

Recently, great progress has been made in using adjoint methods for *a posteriori* error estimation of hyperbolic, convection-diffusion, incompressible Navier-Stokes and compressible Euler equations.^{2,5,6,10,17-22} The use of adjoint methods is based on the fact that in many calculations it is the error in integral out-

put quantities such as the lift and drag which is of most engineering concern. The adjoint solution at each node can be viewed as the linearised effect of a unit source term on the chosen functional. In error analysis, this source term corresponds to the residual error, which measures the extent to which the approximate numerical solution does not satisfy the original differential equation. Thus, the adjoint solution gives the weighted effect of this residual error on the functional of interest. Where the adjoint variables are very small, large residual errors will have little effect on the functional and so the grid need not be adapted. On the other hand, where the adjoint variables are very large, even small residual errors may have a significant effect on the functional and therefore local grid refinement is desirable.

Given this representation of the error in the functional, the adaptive strategy is to refine those parts of the grid which contribute most significantly to the error. Such an approach was initially developed for finite element methods,^{2,6,10,20} and more recently has been used for finite volume methods.^{21,22} In the present paper we apply the approach to the finite volume solution of the 2D Euler equations on unstructured grids. Particular attention has to be paid to the estimation of the residual. We will show that residuals evaluated piecewise linearly over the elements exhibit a high frequency noise to which the standard finite-volume or stabilised finite-element methods are transparent. These high frequencies have to be eliminated by smoothing before a useful adaptation criterion is obtained.

We begin with a review of the adjoint error analysis theory of Giles & Pierce.^{8,19} We then describe in detail the steps required to apply the theory to numerical solutions of the Euler equations to obtain a useful grid adaptation criterion. Numerical results show that this leads to a sequence of adapted grids on which the desired functional converges much more rapidly than with a more uniform grid refinement sequence.

*Res. Officer, email: jdm@comlab.ox.ac.uk

†Professor, email: giles@comlab.ox.ac.uk

Copyright © 2001 by J.-D. Müller, M.B. Giles. Published by the American Institute of Aeronautics and Astronautics, Inc. with permission.

Adjoint error analysis

Let u be the solution of the nonlinear differential equation

$$N(u) = 0,$$

in the domain Ω , subject to the nonlinear boundary conditions

$$D(u) = 0,$$

on the boundary $\partial\Omega$.

The linear differential operators L_u and B_u are defined to be the Fréchet derivatives of N and D ,

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

$$B_u \tilde{u} \equiv \lim_{\epsilon \rightarrow 0} \frac{D(u + \epsilon \tilde{u}) - D(u)}{\epsilon}.$$

and it is also assumed that the nonlinear functional of interest, $J(u)$, has a Fréchet derivative of the following form,

$$\lim_{\epsilon \rightarrow 0} \frac{J(u + \epsilon \tilde{u}) - J(u)}{\epsilon} = (g(u), \tilde{u}) + (h, C_u \tilde{u})_{\partial\Omega}.$$

Here the inner product $(\cdot, \cdot)_{\Omega}$ is an integral over the domain Ω , whereas the inner product $(\cdot, \cdot)_{\partial\Omega}$ is an integral over the boundary. The dimension of the operator C_u (which may be differential) is required to equal the dimension of the adjoint boundary operator B_u^* , to be defined shortly.

The corresponding linear adjoint problem is

$$L_u^* v = g(u)$$

in Ω , subject to the boundary conditions

$$B_u^* v = h$$

on the boundary $\partial\Omega$. The adjoint identity defining L_u^* , B_u^* and the boundary operator C_u^* is

$$(L_u^* v, \tilde{u}) + (B_u^* v, C_u \tilde{u})_{\partial\Omega} = (v, L_u \tilde{u}) + (C_u^* v, B_u \tilde{u})_{\partial\Omega},$$

for all \tilde{u}, v ; see the paper by Giles & Pierce⁷ for how this adjoint identity is derived for the Euler and Navier-Stokes equations.

Now let u_h, v_h be approximate solutions which have been obtained by solving some numerical discretisation of the partial differential equation and its boundary conditions. It is assumed that they are sufficiently differentiable that we may define g_h, h_h by

$$L_{u_h}^* v_h = g_h, \quad B_{u_h}^* v_h = h_h$$

Note the use of the Fréchet derivatives based on u_h which is known, instead of those based on u which is not known.

By defining averaged Fréchet derivatives

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

$$\bar{B}_{(u, u_h)} = \int_0^1 B|_{u+\theta(u_h-u)} d\theta,$$

$$\bar{C}_{(u, u_h)} = \int_0^1 C|_{u+\theta(u_h-u)} d\theta,$$

$$\bar{g}(u, u_h) = \int_0^1 g(u + \theta(u_h - u)) d\theta,$$

so that

$$\begin{aligned} N(u_h) - N(u) &= \int_0^1 \frac{\partial}{\partial \theta} N(u + \theta(u_h - u)) d\theta \\ &= \bar{L}_{(u, u_h)} (u_h - u), \end{aligned}$$

and similarly

$$D(u_h) - D(u) = \bar{B}_{(u, u_h)} (u_h - u),$$

we then obtain

$$\begin{aligned} J(u_h) - J(u) &= (\bar{g}(u, u_h), u_h - u) + (h, \bar{C}_{(u, u_h)}(u_h - u))_{\partial\Omega} \\ &\approx (g_h, u_h - u) + (h_h, C_{u_h}(u_h - u))_{\partial\Omega} \\ &= (L_{u_h}^* v_h, u_h - u) + (B_{u_h}^* v_h, C_{u_h}(u_h - u))_{\partial\Omega} \\ &= (v_h, L_{u_h}(u_h - u)) + (C_{u_h}^* v_h, B_{u_h}(u_h - u))_{\partial\Omega} \\ &\approx (v_h, \bar{L}_{(u, u_h)}(u_h - u)) + (C_{u_h}^* v_h, \bar{B}_{(u, u_h)}(u_h - u))_{\partial\Omega} \\ &= (v_h, N(u_h)) + (C_{u_h}^* v_h, D(u_h))_{\partial\Omega} \end{aligned}$$

In lines 3 and 6 of the above derivation, the errors which are introduced are extremely small, proportional to either $\|u_h - u\|^2$ or $\|u_h - u\| \|v_h - v\|$. Hence, this result shows that the difference between the functional based on the approximate solution, $J(u_h)$, and that based on the true solution, $J(u)$, can be expressed to leading order as the sum of two terms, related to the residual errors in satisfying the nonlinear differential equation and its boundary conditions. In practice, the dominant error of the two is usually the first, $(v_h, N(u_h))$, so we see that the significance of the adjoint solution is that it gives the weighted effect of the residual error $N(u_h)$ on the functional of interest.

Giles & Pierce have used this analysis to correct the computed value for the functional and thereby obtain a higher order of accuracy.^{8,19} Here though we choose to use it to motivate an adaptive strategy. Writing the error in the functional as a sum of terms from each cell in the grid,

$$(v_h, N(u_h)) = \sum_{\alpha} (v_h, N(u_h))_{\alpha} \quad (1)$$

the idea is that we will adapt any cell α for which $(v_h, N(u_h))_{\alpha}$ is greater than some threshold.

Approximation of the Residual Error

The steady two-dimensional Euler equations are

$$N(u) \equiv \frac{\partial F(u)}{\partial x} + \frac{\partial G(u)}{\partial y} = 0. \quad (2)$$

F and G are the usual conservative fluxes in the x and y directions.

Given an approximate solution u_h , the residual flux divergence error on a particular triangular cell E_i is defined as

$$f_i^{elem} = \iint_{E_i} \nabla \cdot (F, G) dA = \int_{\partial E_i} (F, G) \cdot \mathbf{n} ds \quad (3)$$

Given values of u_h at the corner nodes of E_i , the values of F and G on the boundary are obtained by linear interpolation. For faces on a solid boundary, such as an airfoil surface, this is modified to require that there is no mass flux through the face

The resulting plot of the piecewise constant residuals for each element is shown in figure 1. One can clearly see the presence of a high-frequency checkerboard mode. Although not expected, this result is not too surprising, and it is also seen in the 1D results of Venditti and Darmofal.²¹ The underlying nonlinear solution was produced with a node-centered finite-volume code.¹⁴ The discrete flux balance is driven to zero at the nodes where contributions from elements with opposite signs cancel. To first order, centered finite-volume schemes, central cell-vertex schemes and finite-element Galerkin schemes are identical on triangles.¹ Thus, the methodology presented here applies in general, although the problem of noisy element residuals is much more pronounced in finite volume




Fig. 1 Mass flux divergence per element, NACA 0012, Ma .4, 2° angle of attack.

discretisations where upwinding leads to a highly nonlinear artificial viscosity.

A sensor based on these noisy residuals will flag elements in a checkerboard mode, selecting every other element. This in turn destroys the regularity of the mesh and is likely to increase the numerical errors. Since the flow solution itself does not have this checkerboard error, it is not appropriate to be concerned with this residual mode when adapting the grid.

One remedy is to also refine any elements with 2 or more neighbors which are flagged for refinement. Any hanging nodes are removed by tessellation.¹⁶ In this case the smoothing of the grid adds a few extra elements. This strategy is adopted for all of the examples in this work.

Smoothing can also be applied to the elemental flux divergence. The solver for a node-centered discretisation drives nodal residuals to zero, thus the residual flux divergence at the nodes is a smoother quantity. A Galerkin scheme or node-based finite volume scheme distributes an equal share of the residual of element i to each of its three nodes N_i . The flux divergence at node j is then the sum of the contributions from the set of triangles T_j of which it is a node, corresponding to a dual control volume whose area is $A_j = \sum_{i \in C_j} \frac{1}{3} A_i$.

The gathering of the elemental residual to the nodes can be combined with an area-weighted scattering back to the elements, to form a smoothing operation which can be applied repeatedly.

$$f_j^{node} = \sum_{i \in C_j} \frac{1}{3} f_i^{elem} \quad (4)$$

$$f_i^{elem} = \sum_{j \in N_i} \frac{A_j}{A_i} f_j^{node} \quad (5)$$

Note that the total residual error is conserved in this process. Figure 2 shows $v^T f$ after one scatter-gather step. The resulting residual errors are still very noisy but possibly suitable for adaptation. After 30 steps we obtain a smooth approximation to the residual error (Figure 3). To further increase robustness the sensor is evaluated as $|v||f|$ rather than $|v^T f|$. The results show little difference between the two forms.

The smoothing results in compact regions of refinement where the area of the interfaces between adapted and non-adapted elements is small and thus the mesh irregularity introduced by the refinement is minimal. However, as the results in the next section will show, the smoothing also smears out some fainter features and best results are achieved with a small number of smoothing sweeps.

If estimated with sufficient accuracy the product $v^T f$ can also be used to correct the functional to obtain superconvergence on the finest adapted grid.⁸ Since $v^T f$ is an inner product over the entire domain, oscillations will average out and no smoothing is necessary in this case.




Fig. 2 $v^T f$ per element, after one scatter-gather step, log of absolute values.




Fig. 3 $v^T f$ per element, after 30 scatter-gather steps, log of absolute values.

Examples

Subsonic airfoil

The following examples show the subsonic flow around a NACA 0012 profile at a Mach number of 0.4 and an angle of attack of 2° . Figure 10 compares four different refinement strategies. The left column shows refinement based on pressure differences in each element. The pressure sensor is very appropriate for this flow; the flow is isentropic and pressure driven, moreover the pressure is a very smooth field. The second

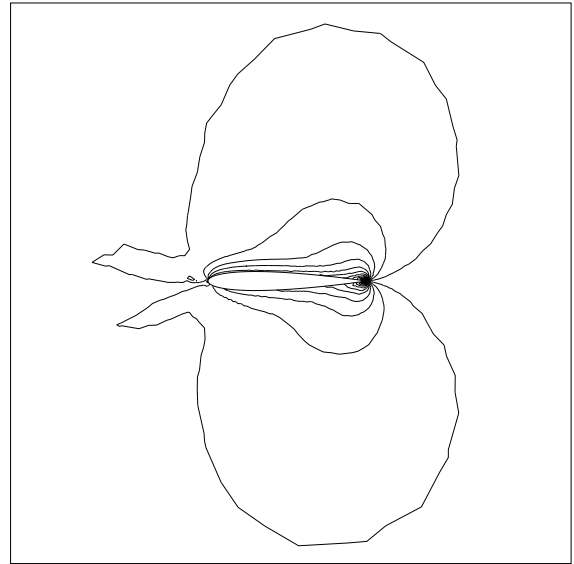


Fig. 4 Subsonic NACA airfoil, adjoint to the energy equation.

column shows refinement using an unweighted sum of residuals, the third column the residuals weighted with the adjoint, the fourth column the adjoint weighting with 30 smoothing iterations. Every adaptation step refined the elements with a sensor value one standard deviation above the average value - and all elements with two or more refined neighbors. In each column, the figure shows the first adaptation and the Mach number contours of the solution, the second adaptation and the fifth adaptation and Mach contours, respectively.

The pressure sensor is ideal for this case and refines in a smooth way where it is needed: at the leading and trailing edges. It can be seen from the large values in the adjoint solution in Figure 4 that the lift coefficient is most sensitive to perturbations in these areas. As opposed to the unweighted residuals, the adjoint weighted residuals do center the adaptation around the profile. Of the total 11% refined elements on the first level, more are situated near the profile. The $|v| |f|$ smoothed 30 times centers the adaptation only around the leading edge. The failure to refine the trailing edge leads to the overshoot in lift prediction shown in Figure 5.

The Mach contours after the first adaptation show a pronounced layer of numerical entropy generated at the leading edge, leading to an unphysical wake. The weighted unsmoothed adjoint sensor does best here. After the fifth adaptation the unweighted residual sensor has refined the domain more or less evenly. Visually the Mach contours show no differences between the sensors except for the pressure and smoothed adjoint sensors that apply no refinement further away from the wing and exhibit slightly less smooth contour lines in those areas.

The main interest in adaptive methods is to improve

level:	1		2		3		4		5	
sensor:	size	err	size	err	size	err	size	err	size	err
p :	1.37	-0.0173	1.80	-0.0021	<i>2.22</i>	<i>0.0008</i>	2.58	0.0013	2.58	0.0013
$ v f $:	1.40	-0.0180	1.77	-0.0020	<i>2.25</i>	<i>-0.0001</i>	2.94	0.0002	4.14	0.0013
$ v f $, 10 sm:	1.32	-0.0176	1.92	-0.0036	<i>2.80</i>	<i>-0.0007</i>	4.13	0.0003	5.81	0.0005
$ v f $, 30 sm:	1.46	-0.0082	2.08	0.0171	3.38	0.0134	5.08	0.0091	7.21	0.0064
f :	1.50	-0.0172	2.49	-0.0028	<i>4.14</i>	<i>-0.0007</i>	6.93	-0.0005	12.19	-0.0001
p , 30%:	1.94	-0.0192	3.79	-0.0032	<i>7.54</i>	<i>0.0005</i>	15.14	0.0010	30.26	0.0003
$ v f $, 30%:	2.14	-0.0195	4.78	-0.0040	<i>10.45</i>	<i>-0.0002</i>	23.24	0.0000	52.24	0.0001
full:	4,00	-0.0200	16.00	-0.0042	<i>64.00</i>	<i>-0.0007</i>	256.00	-0.0002	1024.00	0

Table 1 Convergence of the lift coefficient with refinement for the subsonic NACA 0012. Shown are the grid size relative to the initial grid of 1764 elements and the error in c_L relative to the fully refined solution. Bold numbers indicate a drop below 1% error, italics below 0.1%.




Fig. 5 Grid convergence of the lift for the subsonic test case. The horizontal axis is the inverse of the number of elements in the grids.

the efficiency of a numerical method. To assess the performance of the sensors, Figure 5 plots the lift coefficient c_L over the inverse of the number of elements in each grid. The solid line represents full refinement and the 5th level fully refined grid with 1.8 million elements is taken to be the reference solution. The initial grid with 1764 triangles is at the bottom right, the most efficient methods reach the reference value with the lowest number of elements.

Four sensors exhibit similar grid convergence and converge most rapidly: the pressure sensor, and the adjoint sensors with no smoothing or 10 iterations. Their convergence is also robust, as opposed to the adjoint sensor with 30 smoothing iterations which overshoots the reference value. All of these sensors were used to flag elements one mean deviation above the average. The convergence to the reference solution is not as rapid if refinement with a fixed fraction of 30% is done.

Quantitative results are shown in Table 1. The so-




Fig. 6 Grid convergence for the transonic test case.

lutions are converged to a drop in the L_2 norm of the residuals of about 7 orders of magnitude. All of the solutions except the highly smoothed adjoint drop below 1% error on the second, below 0.1% on the third grid. However, the cost is very different. While the pressure sensor and the adjoint methods need a grid just over twice as fine as the initial one to reach 0.1%, the unweighted residual sensor and the sensors used with fixed fraction refinement require between two and five times as much as the optimal methods.

Transonic airfoil

The test case presented in Figure 11 is the transonic flow over a NACA 0012 profile with Mach number of 0.8 and an angle of attack of 1.25° . It exhibits a strong shock on the upper and a weaker one on the lower surface, resulting in a difference of total pressure across the wake.

In the previous section the pressure sensor was a clear winner: at minimal cost it provided the fastest grid convergence. This will not be the case in general since, for example, a pressure sensor is insensitive to

level:	1		2		3		4		5	
sensor:	size	err	size	err	size	err	size	err	size	err
$ v f $:	1.37	0.0025	1.76	-0.0036	2.35	0.0049	3.40	-0.0018	5.25	-0.0001
$ v f $, 10 sm:	1.43	0.0540	1.97	0.0680	2.88	0.0316	4.51	0.0148	7.28	0.0108
$ v f $, 30%:	2.15	-0.0005	4.83	0.0093	10.79	0.0101	23.56	0.0041	<i>52.10</i>	<i>0.0007</i>
p :	1.47	0.0254	1.80	0.0390	2.14	0.0438	2.63	0.0440	3.47	0.0442
p , 30%:	1.98	0.0030	3.91	0.0108	8.06	0.0059	16.35	0.0021	33.00	-0.0007
full:	4.00	0.0049	16.00	0.0113	64.00	0.0075	256.00	0.0034	1024.00	0

Table 2 Convergence of the lift coefficient with refinement for the transonic NACA 0012. Shown are the grid size relative to the initial grid of 1764 elements and the error in c_L relative to the fully refined solution. Bold numbers indicate a drop below 1% error, italics below 0.1%.

level:	5		6		7		8		9	
sensor:	size	err	size	err	size	err	size	err	size	err
$ v f $:	5.25	-0.0001	8.64	-0.0066	15.23	-0.0055	27.65	-0.0055	51.67	-0.0031
$ v f $, max 4:	5.80	0.0018	9.39	0.0033	15.22	0.0052	23.48	0.0041	39.49	0.0041
p :	3.47	0.0442	5.00	0.0443	7.72	0.0444	12.50	0.0444	20.88	0.0444
p , max 4:	5.13	0.0241	8.19	0.0114	13.02	0.0046	19.58	0.0022	30.25	0.0022

Table 3 Convergence of the lift coefficient for the transonic NACA 0012. Shown are grid size and error in lift coefficient relative to the fully refined solution of level 5, for both the pressure and the adjoint sensor, with and without limiting refinement to four levels.




Fig. 7 Grid convergence for the transonic test case; enlargement around the fully refined solutions.

shear layers. Moreover, the magnitude of the sensor remains $O(1)$ around shocks, independent of refinement. As can be seen in Figure 11, the shock is over-resolved in the final mesh, that is, it is refined beyond what is needed for the accurate prediction of a functional. It absorbs refinement resources needed in other parts of the domain to resolve features which are less visible to the pressure sensor such as the shock on the lower surface or the leading and trailing edges.

This deficiency manifests itself in the grid convergence shown in Figure 6. The pressure sensor fails to converge toward the solution of the fully refined grid.

If the depth of refinement is limited to, say, 4 levels, refinement resources are made available to weaker features after the 4th level and the pressure sensor does converge toward the fully refined solution. Table 3 shows the lift convergence for levels 5 to 9 with and without limiting the maximum depth to 4 levels. Clearly this not a recommended procedure; it merely illustrates where the problem with the pressure sensor lies. Note that even with limited refinement levels the error stagnates after 8 levels of refinement at around 0.2%.

All other sensors including the adjoint ones exhibit a “hump” in the convergence of the lift. The lift increases in the first adaptation stages when the upper shock is refined and the shock moves aft. It decreases when the refinement starts to detect the weaker lower shock. The same phenomenon appears with the fully refined solution; the lift decreases after the second adaptation. Considering the “erratic” convergence even on the fully refined grid, it is not too surprising that the adjoint sensors exhibit a similar behaviour. However, what can be seen is that they do converge toward the fully refined solution, albeit at a slow rate. This can be attributed to the fact that the adjoint sensor overly refines the shock on the upper surface.

Figure 8 shows cuts of the nodal residual across the shock on the upper surface. It can be seen that the magnitude of the oscillations decreases with mesh refinement, but for the first adaptation levels they are larger than the residuals near the leading edge. The adjoint variables are continuous and have a zero gradient in a 1-D shock^{7,9} and consequently the oscillations in the residual error cancel to a large extent on either side of the shock. Smoothing does reduce the




Fig. 8 Unweighted sum of residuals at levels 3, 4, 5, cut across the shock at $y=.25$. Solid line with symbols is level 3, dashed line level 4, dotted line level 5.




Fig. 9 Unweighted sum of residuals at level 3, cut across the shock at $y=.25$. Solid line with symbols are 0 smoothings, dashed line 3 smoothings, dotted line 10 smoothings.

magnitude of these oscillations (see Figure 9) but a low frequency perturbation remains that triggers the adaptation. The plot in Figure 11 corresponding to the adjoint sensor smoothed 10 times shows the broadening of the refinement around the shock due to the smoothing. On the other hand smoothing does smear out the weaker shock such that is not recognised by the adaptation.

It can be concluded that the adjoint sensors do produce a converging sensor for shocked flows, however at a rate that is not optimal.

Conclusions

Mesh refinement based on adjoint sensors has been presented. The advantages of adjoint sensors over conventional finite difference sensors have been described.

Difficulties with the smoothing of the estimated residuals are discussed and a solution was presented. Results show that a smoothed adjoint sensor has better grid convergence for an integral functional than sensors without adjoint weighting or smoothing.

As presented, the method makes no assumptions

about the numerical method underlying the primal and adjoint solutions. It applies generally to three-dimensional hybrid meshes, and is capable of being extended to the Navier-Stokes equations.

Acknowledgements

This research has been supported by funding from the Engineering and Physical Sciences Research Council under grants GR/K91149 and GR/L95700.

We also acknowledge the contributions of M.C. Duta and N.A. Pierce to the development of the adjoint code.




Fig. 10 Grid refinement based on pressure (left), unweighted residuals (middle left), unsmoothed $|v||f|$ (middle right), $|v||f|$ smoothed 30 times (right). The rows show first adaption and their Mach contours, second adaption and fifth adaption with Mach contours.




Fig. 11 Refinements for the transonic flow over a NACA 0012. Shown are grids and Mach contours of the fifth refinement level. Top row: pressure, pressure after 9 levels of refinement, pressure with 30% refinement, initial grid and solution after 5 full refinements. Bottom row: $|v||f|$, $|v||f|$ after 9 levels, $|v||f|$ 30% and $|v||f|$ with 10 residual smoothing iterations.

References

- ¹T.J. Barth. Aspects of unstructured grids and finite-volume solvers for the Euler and Navier-Stokes equations. In *Special Course on Unstructured Grid Methods for Advection Dominated Flows*, Report 787, pages 6.1–6.61. AGARD, 1992.
- ²R. Becker and R. Rannacher. Weighted *a posteriori* error control in finite element methods. In *et al* H.G. Block, editor, *Proceedings of ENUMATH-97*, pages 621–637. World Scientific Publishing, 1998.
- ³M. Berger and P. Colella. Local adaptive mesh refinements for shock hydrodynamics. *J. Comput. Phys.*, 82:64–84, 1989.
- ⁴R.L. Davis and J.F. Dannenhoffer. Adaptive grid embedding Navier-Stokes technique for cascade flows. *AIAA J. Propulsion and Power*, 7(5):792–799, 1991.
- ⁵M.B. Giles. Defect and adjoint error correction. Proceedings of the ICCFD conference, Kyoto, 2000.
- ⁶M.B. Giles, M.G. Larson, M. Levenstam, and E. Süli. Adaptive error control for finite element approximations of the lift and drag in a viscous flow. Technical Report NA 97/06, Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford, OX1 3QD., 1997.
- ⁷M.B. Giles and N.A. Pierce. Adjoint equations in CFD: duality, boundary conditions and solution behaviour. AIAA Paper 97-1850, 1997.
- ⁸M.B. Giles and N.A. Pierce. Improved lift and drag estimates using adjoint Euler equations. AIAA Paper 99-3293, 1999.
- ⁹M.B. Giles and N.A. Pierce. Analytic adjoint solutions for the quasi-one-dimensional Euler equations. *J. Fluid Mech.*, 426:327–345, 2001.
- ¹⁰C. Johnson, R. Rannacher, and M. Boman. Numerics and hydrodynamic stability – toward error control in computational fluid dynamics. *SIAM J. Numer. Anal.*, 32(4):1058–1079, 1995.
- ¹¹R. Löhner, K. Morgan, and O. Zienkiewicz. *Adaptive grid refinement for the compressible Euler equations*, pages 281–297. John Wiley and Sons Ltd, 1986.
- ¹²D. Mavriplis and A. Jameson. Multigrid solution of the Euler equations on unstructured and adaptive meshes. In S. McCormick, editor, *Proceedings of the Third Copper Mountain Conference on Multigrid Methods: Lecture Notes in Pure and Applied Mathematics*. Marcel Dekker Inc, 1987.
- ¹³S. McCormick. Fast adaptive composite (FAC). *Defect Correction Methods: Theory and Applications, Computing Supplementum*, 5:115–121, 1984.
- ¹⁴P. Moinier, J.-D. Müller, and M.B. Giles. Edge-based multigrid and preconditioning for hybrid grids. AIAA Paper 99-3339, 1999.
- ¹⁵J.-D. Müller, T. Schönfeld, and M. Rudgyard. A comparison of the treatment of hanging nodes for hexahedral grid refinement. AIAA Paper 97-1859, 1997.
- ¹⁶J.-D. Müller, T. Schönfeld, and M. Rudgyard. A comparison of the treatment of hanging nodes for hybrid grid refinement. *AIAA-CP-97-1859*, 1997.
- ¹⁷M. 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.
- ¹⁸J. Peraire and A.T. Patera. Bounds for linear-functional outputs of coercive partial differential equations: local indicators and adaptive refinement. In P. Ladeveze and J.T. Oden, editors, *New Advances in Adaptive Computational Methods in Mechanics*. Elsevier, 1997.
- ¹⁹N.A. Pierce and M.B. Giles. Adjoint recovery of super-convergent functionals from PDE approximations. *SIAM Rev.*, 42(2):247–264, 2000.
- ²⁰E. Süli. *A posteriori* error analysis and adaptivity for finite element approximations of hyperbolic problems. In D. Kröner, M. Ohlberger, and C. Rohde, editors, *An Introduction to Recent Developments in Theory and Numerics for Conservation Laws*, volume 5 of *Lecture Notes in Computational Science and Engineering*, pages 123–194. Springer-Verlag, 1998.
- ²¹D. Venditti and D. Darmofal. A multilevel error estimation and grid adaptive strategy for improving the accuracy of integral outputs. AIAA Paper 99-3292, 1999.
- ²²D. Venditti and D. Darmofal. Adjoint error estimation and grid adaptation for functional outputs: application to quasi-one-dimensional flow. *J. Comput. Phys.*, 164:204–227, 2000.
- ²³G.P. Warren, W.K. Anderson, J.L. Thomas, and S.L. Krist. Grid convergence for adaptive methods. AIAA Paper 91-1592, 1991.
- ²⁴D. De Zeeuw and K.G. Powell. An adaptively refined Cartesian mesh solver for the Euler equations. *J. Comput. Phys.*, 104(1):56–68, 1993.