Goal-oriented Mesh Adaptation in an Industrial Stabilized Finite Element Navier-Stokes Code

Frédéric Chalot

Abstract This chapter describes Dassault Aviation's contribution to Workpackage 5 of the ADIGMA Project. The adjoint operator developed in the framework of optimum design is used to estimate the error in the solution with respect to a given target quantity. Local values of this error estimation are used as a criterion to refine the mesh. This yields significant improvement over traditional criteria based on the residual or on gradients of physical quantities. The method is carefully tested using inviscid, transonic, laminar, and high Reynolds number turbulent flows.

1 Stabilized finite element schemes for the RANS equations

Dassault Aviation's Navier-Stokes code, called AETHER, uses a finite element approach, based on a symmetric form of the equations written in terms of entropy variables. The advantages of this change of variables are numerous: in addition to the strong mathematical and numerical coherence they provide (dimensionally correct dot product, symmetric operators with positivity properties, efficient preconditioning), entropy variables yield further improvements over the usual conservation variables, in particular in the context of chemically reacting flows (see [1, 2]).

The code can handle the unstructured mixture of numerous types of elements (triangles and quadrilaterals in 2-D; tetrahedra, bricks, and prisms in 3-D). In practice mostly linear triangular and tetrahedron meshes are used.

The code has been successfully ported on many computer architectures. It is fully vectorized and parallelized for shared or distributed memory machines using the MPI message passing library (IBM SP2 Series, IBM BlueGene, Itanium II- and Xeon-based Bull NovaScale) or native parallelization directives (NEC SX-4) (see [3]).

Frédéric Chalot
Dassault Aviation, 78, quai Marcel Dassault, CEDEX 300, 92252 Saint-Cloud CEDEX, France
e-mail: frederic.chalot@dassault-aviation.com

For details about the numerical method, the reader is referred to Chapter **??**. We just recall the semi-discrete Galerkin/least-squares variational problem which can be stated as:

Find $\boldsymbol{V}^h \in \mathscr{S}^h$ (trial function space), such that for all $\boldsymbol{W}^h \in \mathscr{V}^h$ (weighting function space), the following equation holds:

$$
\begin{aligned}
\int_\Omega \Big(W^h \cdot \boldsymbol{U}_{,t}(\boldsymbol{V}^h) - \boldsymbol{W}^h_{,i} \cdot \boldsymbol{F}^{\mathrm{adv}}_i(\boldsymbol{V}^h) + \boldsymbol{W}^h_{,i} \cdot \widetilde{\boldsymbol{K}}_{ij} \boldsymbol{V}^h_{,j} \Big)\, d\Omega & \\
+ \sum_{e=1}^{n_{\mathrm{el}}} \int_{\Omega^e} \Big(\mathscr{L} \boldsymbol{W}^h \Big) \cdot \boldsymbol{\tau} \Big(\mathscr{L} \boldsymbol{V}^h \Big)\, d\Omega & \\
+ \sum_{e=1}^{n_{\mathrm{el}}} \int_{\Omega^e} \nu^h g^{ij} \boldsymbol{W}^h_{,i} \cdot \widetilde{\boldsymbol{A}}_0 \boldsymbol{V}^h_{,j}\, d\Omega & \\
= \int_\Gamma \boldsymbol{W}^h \cdot \Big(- \boldsymbol{F}^{\mathrm{adv}}_i(\boldsymbol{V}^h) + \boldsymbol{F}^{\mathrm{diff}}_i(\boldsymbol{V}^h) \Big) n_i\, d\Gamma. & \qquad (1)
\end{aligned}
$$

2 Adjoint-based indicators for stabilized finite element methods

In the following sections we describe how the adjoint operator developed in the framework of optimum design can be used to estimate the error in the solution with respect to a given target quantity. Local values of this error estimation are used as a criterion to refine the mesh. This yields significant improvement over traditional criteria based on the residual or on gradients of physical quantities.

2.1 Adjoint-based extrapolation

We assume that we have a solution of the Navier-Stokes equations V^H solution of (1): $R^H(V^H) = 0$ on a given mesh characterized by a mesh size parameter H, where R^H is the discrete residual associated with (1). We can compute some aerodynamic function f of the solution (for instance drag or lift): $f^H(V^H)$. Unfortunately, mesh M^H is too coarse to compute f^H accurately. A finer mesh h whose characteristic mesh size parameter h is smaller than H would yield a better estimate of f, viz., $f^h(V^h)$. The question is "how can we estimate f^h without explicitly computing V^h, which would require the solution of $R^h(V^h) = 0$?" We will seek an estimate of f^h in the form

$$f^h(V^h) \approx f^h(V^h_H) + \cdots$$

where V^h_H is the projection of V^H onto mesh M^h.

The first order Taylor expansion of $f^h(V^h)$ about V^h_H reads:

$$f^h(V^h) = f^h(V^h_H) + \left.\frac{\partial f^h}{\partial V^h}\right|_{V^h_H} (V^h - V^h_H) + O\left((V^h - V^h_H)^2\right) \tag{2}$$

We can expand $R^h(V^h)$ in the same fashion

$$R^h(V^h) = R^h(V^h_H) + \left.\frac{\partial R^h}{\partial V^h}\right|_{V^h_H} (V^h - V^h_H) + O\left((V^h - V^h_H)^2\right) \tag{3}$$

Since V^h is the solution of (1) on M^h, we have

$$R^h(V^h) = 0 \tag{4}$$

Combining (2), (3), and (4), it results

$$f^h(V^h) \approx f^h(V^h_H) - \left.\frac{\partial f^h}{\partial V^h}\right|_{V^h_H} \left.\frac{\partial R^h}{\partial V^h}\right|_{V^h_H}^{-1} R^h(V^h_H) \tag{5}$$

Following the ideas of Giles [5], we introduce the adjoint problem

$$\left[\frac{\partial R^h}{\partial V}(V^h)\right]^T \Psi^h = \left[\frac{\partial f^h}{\partial V}(V^h)\right]^T \tag{6}$$

The adjoint of R^h with respect to the entropy variables V^h was obtained from the original FORTRAN code by automatic differentiation using TAPENADE [4] in reverse mode.

Eq. (5) can be rewritten as

$$f^h(V^h) \approx f^h(V^h_H) - \Psi^h \cdot R^h(V^h_H)$$

In practice, we do not solve the adjoint problem on mesh M^h. Instead, we replace Ψ^h with Ψ^h_H, the projection onto mesh M^h of Ψ^H, solution of the adjoint problem (6) on mesh M^H. This adjoint problem is solved with a preconditioned GMRES algorithm.

Finally we can write the approximate of $f^h(V^h)$ as

$$f^h(V^h) \approx f^h(V^h_H) - \Psi^h_H \cdot R^h(V^h_H) \tag{7}$$

which can be computed cheaply with the projected solution of the adjoint problem on the coarse mesh Ψ^h_H and a mere residual evaluation on the fine mesh using the projection of the coarse mesh solution V^h_H.

2.2 Error estimation and refinement criterion

Similar ideas can be used to place error bounds on the estimated value of the aerodynamic function $f^h(V_H^h)$ with respect to $f^h(V^h)$:

$$f^h(V^h) - f^h(V_H^h) \approx -\Psi_H^h \cdot R^h(V_H^h) + (\Psi_H^h - \Psi^h) \cdot R^h(V_H^h) \tag{8}$$

Let T_H^h denote the correction term in (7):

$$T_H^h = \Psi_H^h \cdot R^h(V_H^h)$$

Then (8) becomes

$$f^h(V^h) - f^h(V_H^h) \approx -T_H^h + (\Psi_H^h - \Psi^h) \cdot R^h(V_H^h)$$

One can show that

$$\begin{aligned} \left|(\Psi_H^h - \Psi^h) \cdot R^h(V_H^h)\right| &\to 0, \quad \text{as } h \to 0 \\ &\leq \left|\Psi_H^h \cdot R^h(V_H^h)\right| \end{aligned}$$

Consequently,

$$\left|f^h(V^h) - f^h(V_H^h)\right| \leq C \left|T_H^h\right|$$

with

$$\begin{aligned} C &\to 0, \quad \text{as } h \to 0 \\ &\leq 2 \end{aligned}$$

Finally

$$f^h(V_H^h) - C\left|T_H^h\right| \leq f^h(V^h) \leq f^h(V_H^h) + C\left|T_H^h\right| \tag{9}$$

If M^h is a very fine mesh on which asymptotic convergence is reached, (9) represents the error bound on $f^h(V_H^h)$ with respect to the exact solution of (1).

As we will see shortly, T_H^h can also be used as a goal-oriented mesh adaptation criterion, which reveals area in the mesh where locals errors have the biggest impact on the value of the target function f^h. If we go back to equation (7), we can see that if the “coarse mesh” M^H (in fact the current mesh) produces a solution V^H which is accurate enough to give a satisfactory value of the target aerodynamic function f^h,

$$f^h(V^h) \approx f^h(V_H^h)$$

and

$$T_H^h \approx 0$$

The idea consists in refining the mesh where local values of $|T_H^h|$ are greater than some specified limit ε. The refinement algorithm goes as follows:

1. on mesh M^H, solve (1) for V^H and compute the solution of the adjoint problem Ψ^H with respect to some target function f^H;
2. generate a finer finer mesh M^h, typically obtained by a uniform (iso-P2) refinement of M^H;
3. project V^H and Ψ^H onto M^h and evaluate $R^h(V_H^h)$. Remark: (1) is not solved on M^h;
4. compute a local mesh adaptation parameter P_i^h at each node of mesh M^h

$$P_i^h = (\Psi_H^h)_i \cdot R_i^h(V_H^h)$$

   where the dot product is extended to the sole number of degrees of freedom at node i;
5. project P^h onto mesh M^H;
6. if $(P_h^H)_i < \varepsilon$, refine the mesh locally and go back to step 1; otherwise the mesh is fine enough and $f^H(V^H)$ is computed with an adequate accuracy.

3 Numerical examples of goal-oriented refinement for 2-D flows

Dassault Aviation computed the same four Mandatory Test Cases defined in Workpackage 2 of the ADIGMA Project and already presented in Chapter **??**. They cover a wide range of applications: from inviscid subsonic and transonic flows (MTC's 1 and 2), to laminar Navier-Stokes (MTC 3), and finally a profile in transonic turbulent conditions (MTC 5). For each test case we compare the baseline results obtained using Dassault Aviation's industrial Navier-Stokes code AETHER on a set of uniformly refined meshes with successive goal-oriented adaptation based on the same initial coarse grid.

Local isotropic mesh enrichment is used: triangles tagged for adaptation are split into four. Nodes added on the boundary are placed on the actual surface; if needed mesh deformation techniques are used to make all elements positive. On the border of a locally refined zone, hanging nodes are connected to the opposite vertex. In order to control the aspect ratios in the adapted meshes, we allow only subdivision of original triangles into four or two. A triangle with hanging nodes on two faces will be split into four, propagating a new hanging node further away. Memory of triangles split into two is kept to avoid later division. This technique ensures the quality of the adapted grids.

3.1 *MTC 1: NACA0012,* $M = 0.50$*,* $\alpha = 2^\circ$*, inviscid*

For this inviscid subsonic test case, the drag coefficient C_D was used as the target quantity. The initial unadapted mesh around the NACA0012 airfoil contains 1106 nodes. Figure 1 presents the original mesh and five successive levels of goal-oriented adaptive refinement together with the matching Mach-number contours.

Kinks in Mach number contours disappear after only two levels of adaptation although refinement mostly occurs in the stagnation, suction, and trailing edge regions. This is an indication that the entropy layer observed in coarse grid solutions is not due to lower order boundary conditions. Instead spurious entropy is generated at the leading edge and is convected along the profile.

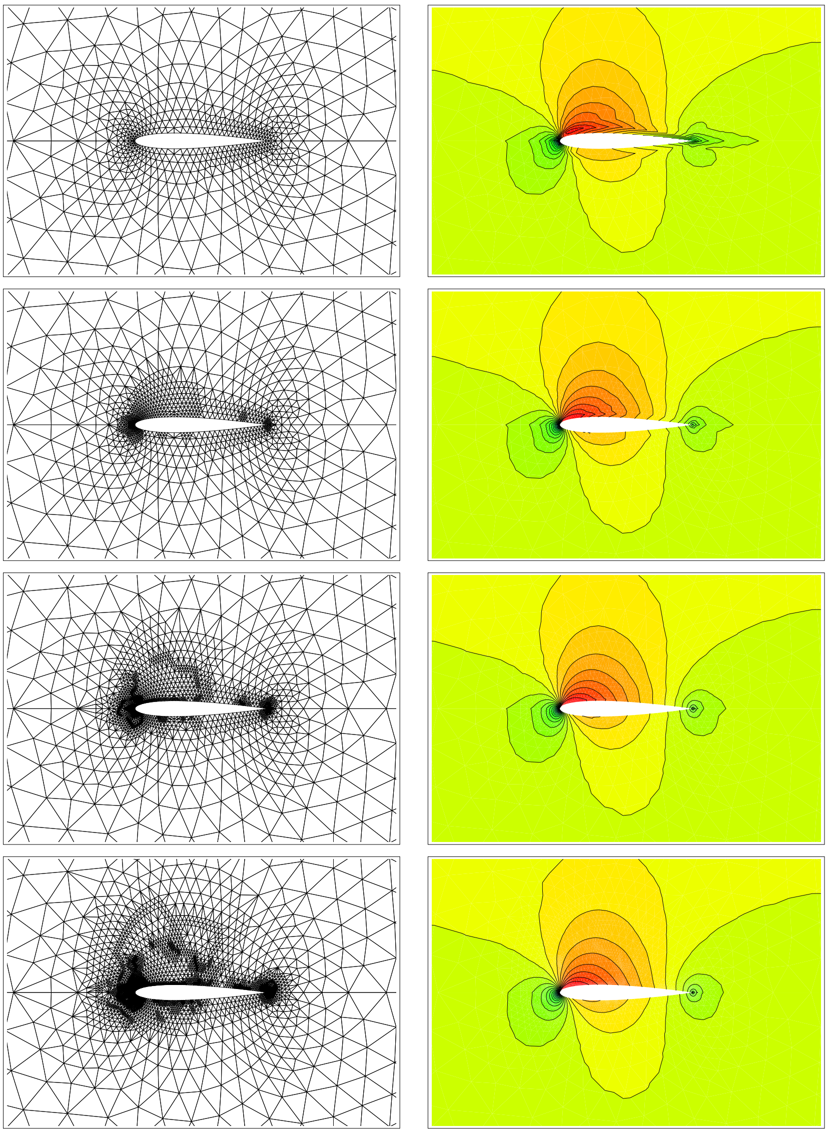

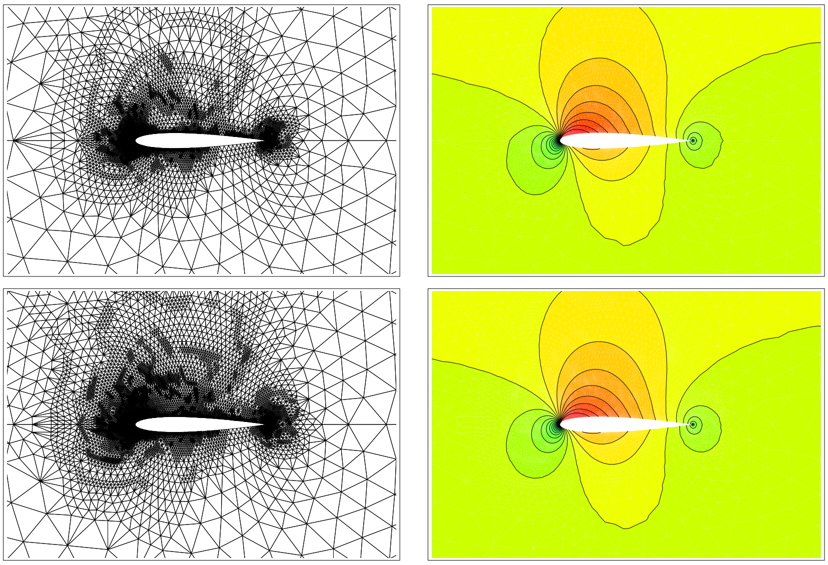

Fig. 1 MTC 1: NACA0012, $M = 0.50$, $\alpha = 2°$, inviscid. Original 1106-node mesh and five successive levels of goal-oriented refinement based on drag with corresponding Mach number contours on the right.

Figure 2 shows the convergence of force coefficients obtained with goal-oriented mesh adaptation (mixed line) compared with those computed with global mesh refinement (solid line).

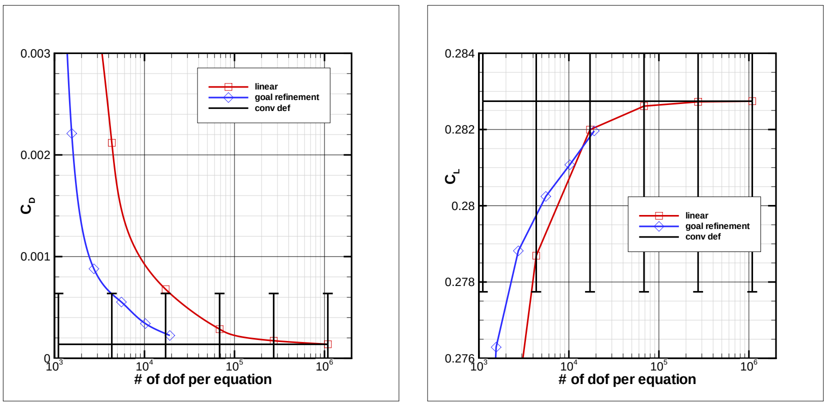

Fig. 2 MTC 1: NACA0012, $M = 0.50$, $\alpha = 2°$, inviscid. Convergence of force coefficients compared with global mesh refinement.

It is striking to see that C_D, which is the target, converges better than C_L as the sizes of adapted meshes increase. The gain in the required number of degrees of freedom for convergence is roughly a factor of 5. CPU gain is of the same order, with an additional advantage for adapted meshes: they have fewer points in the freestream, thus they require fewer time steps to reach convergence at a given CFL number.

It must be noted that the threshold for refinement was fixed at 50% of the mean criterion value. This yields a series of adapted meshes which nearly double in size at each level of adaptation. This figure can certainly be reduced with a stricter refinement criterion limit.

3.2 *MTC 2: NACA0012, $M = 0.80$, $\alpha = 1.25°$, inviscid*

The pressure drag coefficient C_D was also used as the target quantity to produce the adapted meshes for the transonic case MTC 2 shown in Figure 3. Mach number contours corresponding to the initial grid and to the six subsequent levels of adaptation are displayed on the right. The initial mesh is the same as the one for MTC 1.

In the beginning, refinement occurs more or less uniformly in the dependency region. Only when the overall mesh size reaches a reasonable level, the criterion hits more selectively at the shocks. Both leeward and windward side shocks are accurately captured which would have been particularly challenging for a gradient-based adaptation. As in the subsonic case, points are added at the surface of the airfoil, in particular in the acceleration/expansion regions. The slip line at the trailing edge is also detected in the final meshes.

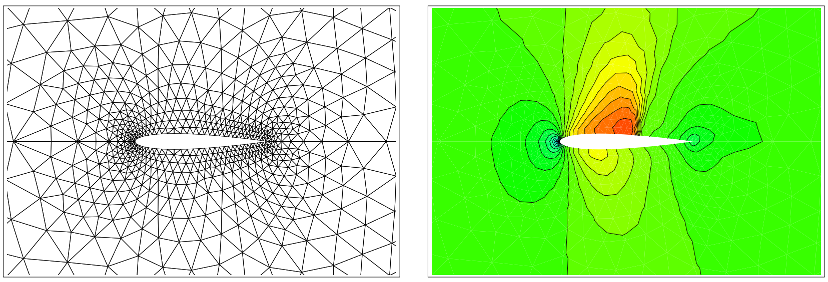

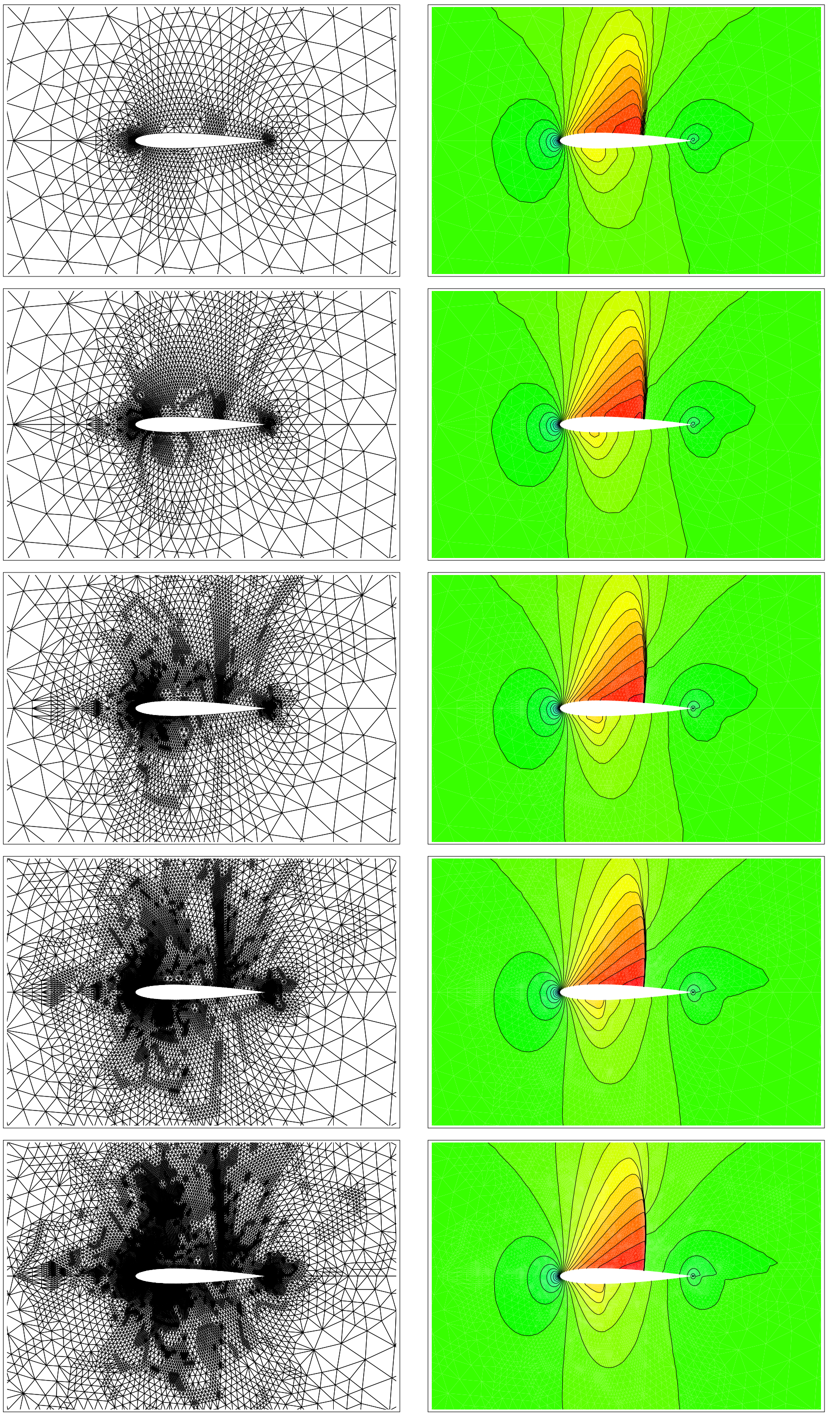

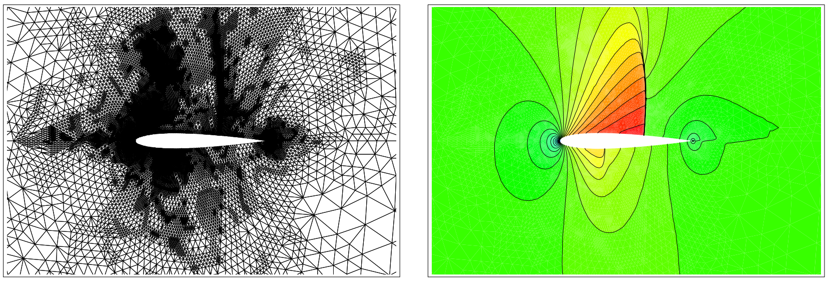

Fig. 3 MTC 2: NACA0012, $M = 0.80$, $\alpha = 1.25°$, inviscid. Original 1106-node mesh and six successive levels of goal-oriented refinement based on drag with corresponding Mach number contours on the right.

Figure 4 presents the convergence of force coefficients. The same criterion threshold as MTC 1 was used. Again it produces a series of adapted meshes which grow too fast in terms of degrees of freedom. Nonetheless the required number of nodes to converge the drag coefficient is reduced by a factor of 2.3. The behavior of lift coefficient is very disappointing. Even on the finest adapted grid (54,220 nodes), it does not meet the convergence criterion. Goal-oriented mesh refinement based on C_L should be tested.

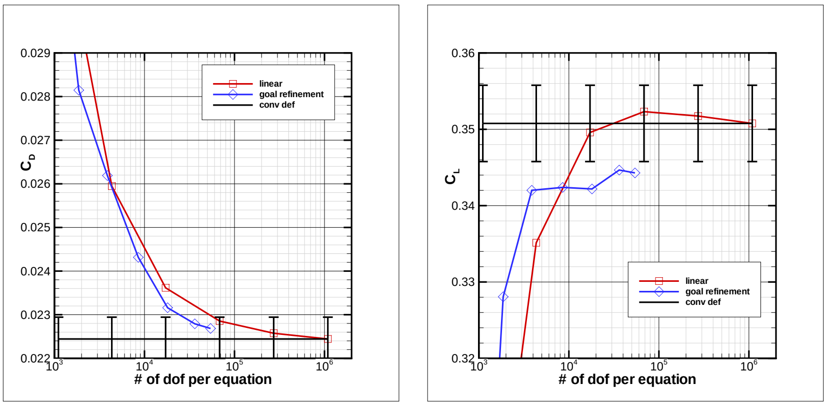

Fig. 4 MTC 2: NACA0012, $M = 0.80$, $\alpha = 1.25°$, inviscid. Convergence of force coefficients compared with global mesh refinement.

3.3 MTC 3: NACA0012, $M = 0.50$, $\alpha = 2^\circ$, $Re = 5,000$

Even though MTC 3 is a viscous test case, we have again used the pressure drag coefficient C_D as the target for goal-oriented mesh adaptation with the same criterion threshold chosen previously.

Figure 5 shows the original 1533-node mesh and four successive levels of goal-oriented refinement based on pressure drag together with the corresponding pressure contours on the right. Refinement occurs at the leading edge, along the profile, and in the wake. Mesh deformation was used at each level of adaptation to place the new nodes along the actual profile.

Only four levels of refinement could be applied. A fifth adapted mesh was generated, on which the computation revealed unsteady. This behavior was observed by other partners in the ADIGMA Project, although all our previous computations for this $Re = 5,000$ test cases always converged to a steady state.

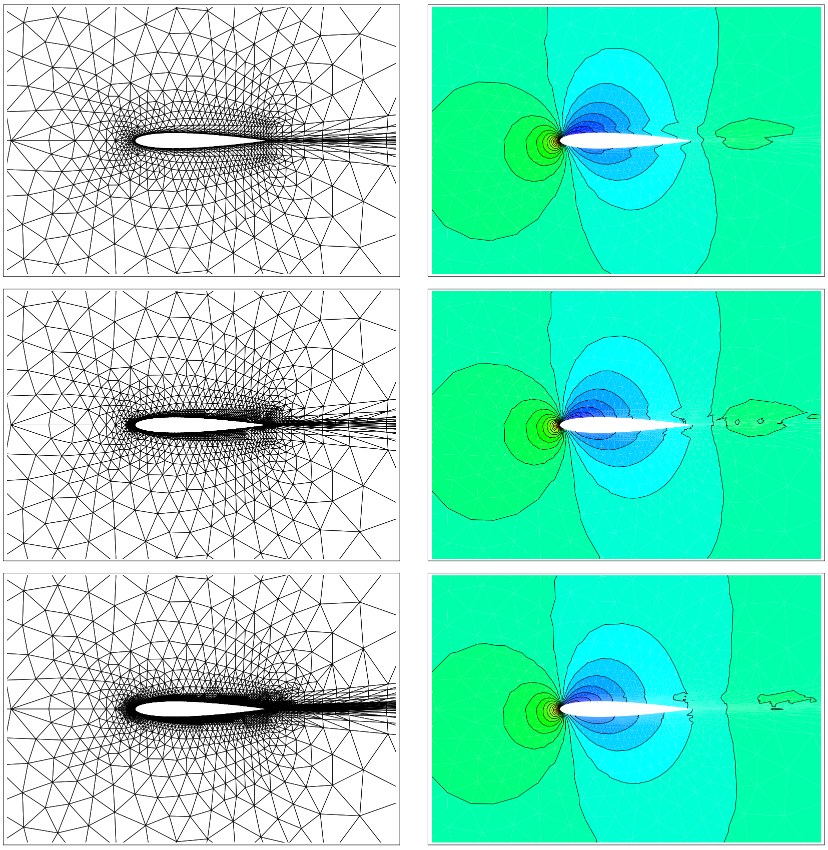

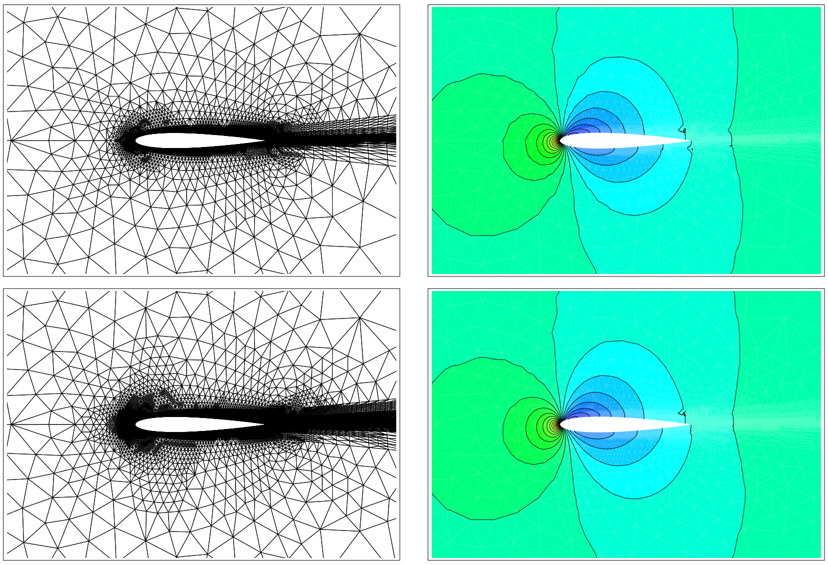

Fig. 5 MTC 3: NACA0012, $M = 0.50$, $\alpha = 2^\circ$, $Re = 5,000$. Original 1533-node mesh and four successive levels of goal-oriented refinement based on pressure drag with corresponding pressure contours on the right.

Figure 6 displays the convergence plots of the force and heat flux coefficients. Once more the criterion threshold seems too lenient. The mesh size increases too rapidly at each refinement step (by a factor of nearly 3). The required number of degrees of freedom to converge pressure drag is barely reduced by 10%.

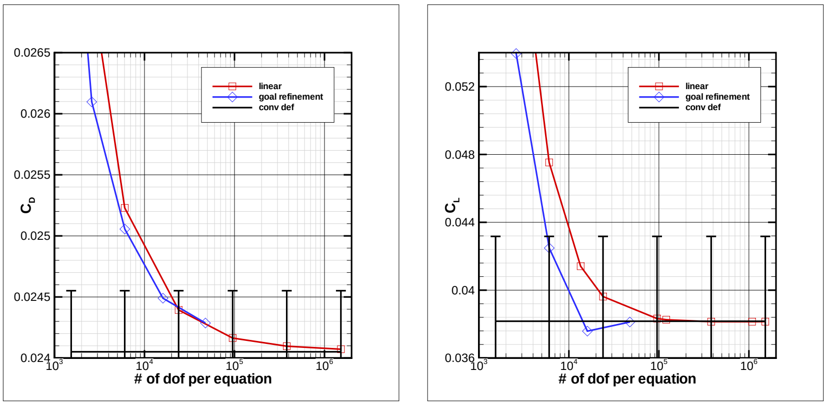

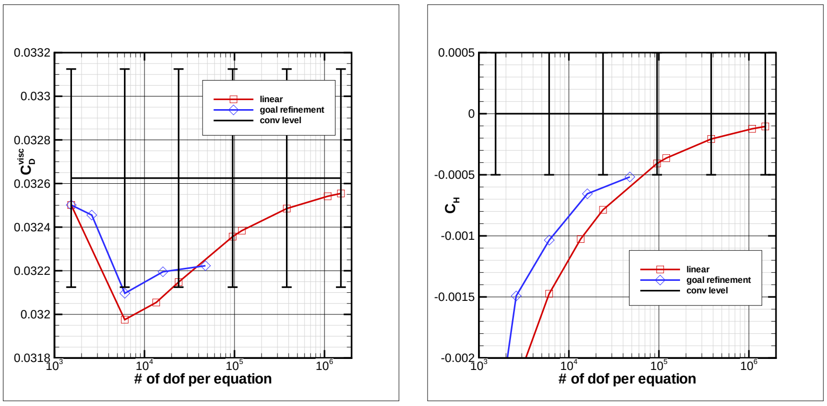

Fig. 6 MTC 3: NACA0012, $M = 0.50$, $\alpha = 2°$, $Re = 5,000$. Convergence of force and heat flux coefficients compared with global mesh refinement.

Although not targeted, lift convergence requirements drop by about 40%, those of friction drag by about 20%. The convergence curves of friction drag and of heat flux tend to flatten out with the finest adapted meshes. This might be an indication of the forthcoming onset of unsteadiness.

3.4 MTC 5: RAE2822, $M = 0.734$, $\alpha = 2.79°$, $Re = 6,500,000$

The final test case is a transonic high Reynolds number RANS calculation past an RAE2822 profile. We have still used the pressure drag coefficient C_D as the target for goal-oriented mesh adaptation. We have altered the criterion threshold from the previous MTC's though. It appeared that too heavy a refinement was applied at each level, which somehow reduced the potential benefit of mesh adaptation. For MTC 5 we have tried to limit the refinement to the top 20% of the elements where the criterion was the largest.

In doing so problems were encountered when refining under-resolved highly-curved boundaries near the leading edge. Too little local refinement would prevent mesh deformation from completing successfully: the locally refined region must be thick enough to allow the deformation of the thinnest elements into the volume. The refinement zone had to be extended respectively to the top 40 and 60% to permit deformation of the first two adapted meshes. For the next two refined grids, no criterion threshold value would yield no negative elements after deformation. We chose to stick to our 20% rule and to skip mesh deformation all together. The final adapted mesh was obtained with the top 20% criterion and with a successful mesh deformation. The five adapted meshes are presented in Figure 7 with the initial 2668-node mesh; Mach number contours are shown on the right next to each corresponding grid.

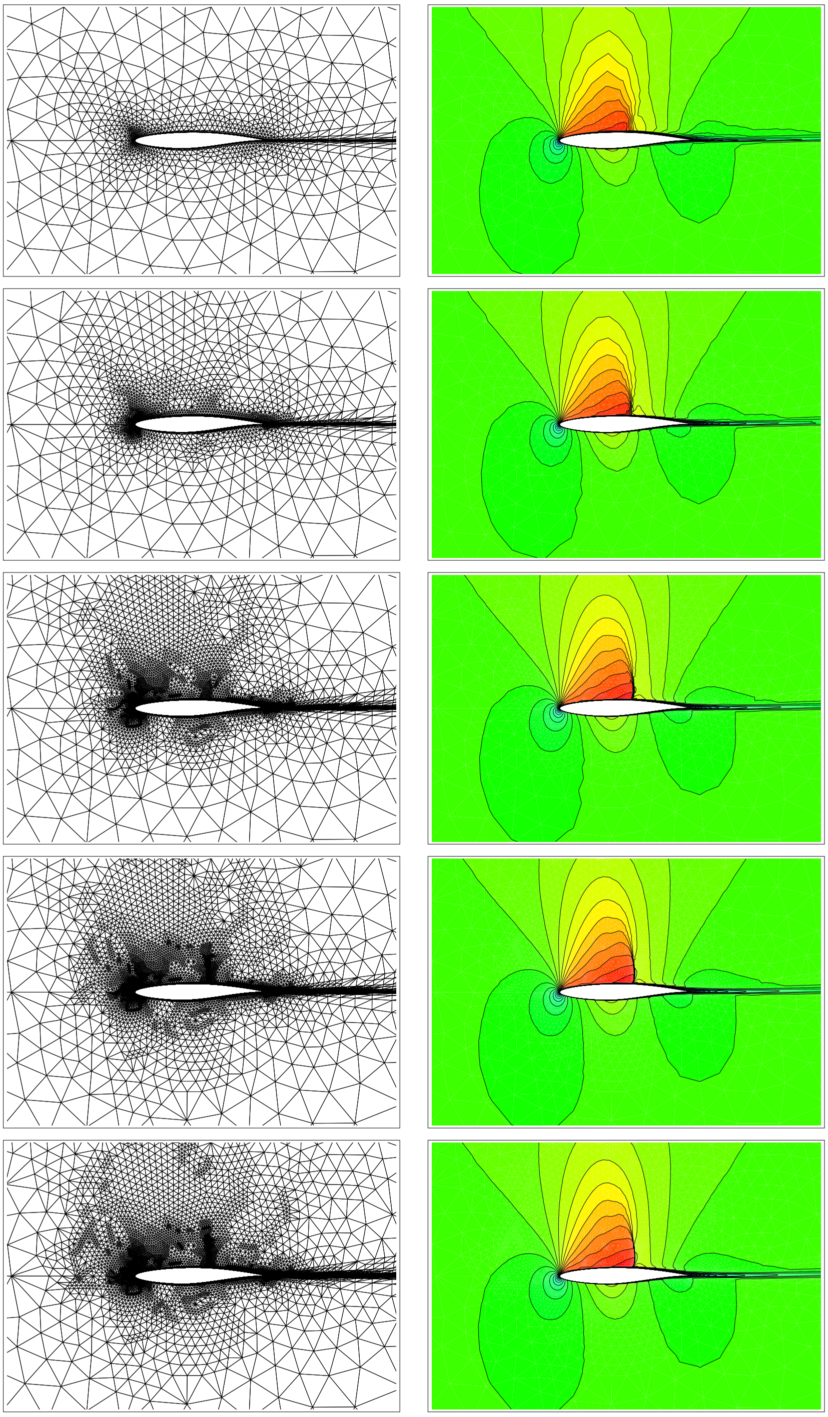

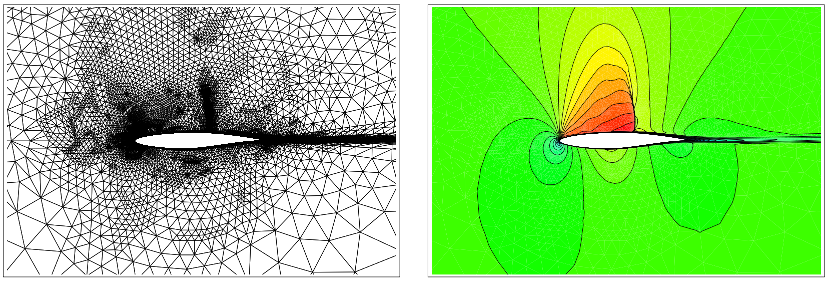

Fig. 7 MTC 5: RAE2822, $M = 0.734$, $\alpha = 2.79°$, $Re = 6,500,000$. Original 2668-node mesh and five successive levels of goal-oriented refinement based on pressure drag with corresponding Mach number contours on the right.

At first goal-oriented adaptation driven by pressure drag refined the leading edge, the suction region and the wake. Then it hit more specifically at the shock area. Although intrinsically inviscid the chosen target seems to have tackled this high-Reynolds number case rather well. Let's have a look at the convergence of force and heat flux coefficients presented in Figure 8 for a more quantitative analysis.

The number of degrees of freedom to converge pressure drag is reduced by a factor of 13 from $130,000$ to about $10,000$. This exemplifies the power of goal-oriented mesh adaptation when used with a controlled refinement criterion level. The requirement for lift is divided by a factor of nearly 3 (from $60,000$ to about $23,000$). Viscous coefficients are less successful. Friction drag converges only slightly faster than with a uniform grid refinement. Heat flux seems to reach an asymptotic value of 10^{-4} and not converge any further. Again it would be worth testing more Navier-Stokes specific target functions, such as friction drag, total drag, or heat flux.

In spite of the aforementioned mesh deformation difficulties, goal-oriented mesh adaptation, even based on a pressure target, can handle high Reynolds numbers. Feature based adaptation would probably miss some of the features of the flow.

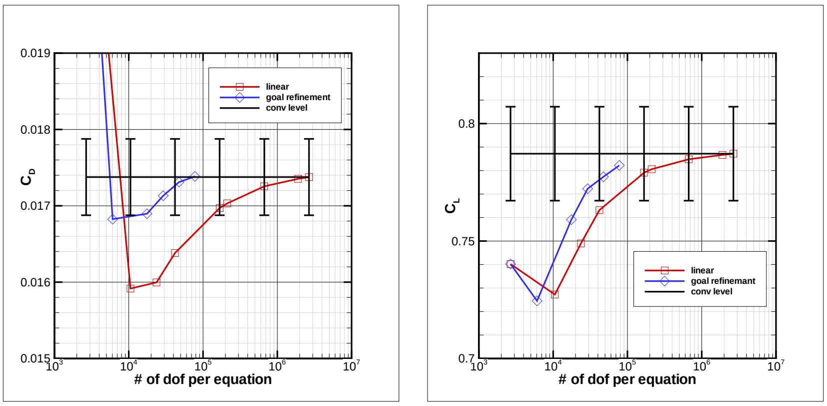

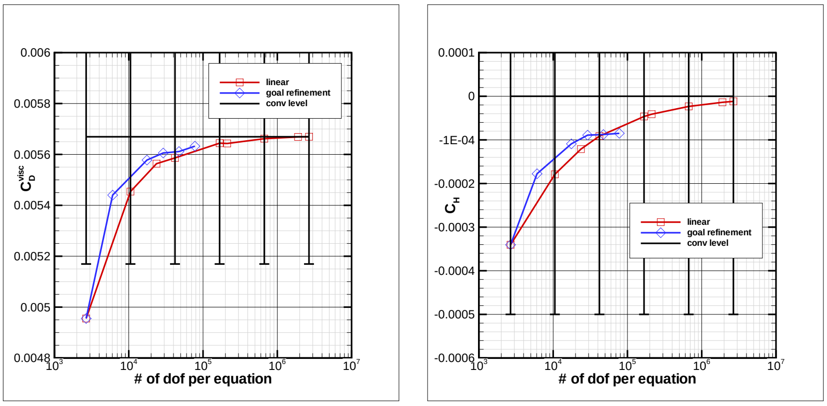

Fig. 8 MTC 5: RAE2822, $M = 0.734$, $\alpha = 2.79°$, $Re = 6,500,000$. Convergence of force and heat flux coefficients compared with global mesh refinement.

4 Conclusion

Substantiated by various test cases, we have showed that goal-oriented adaptation coupled with local isotropic mesh refinement, works for diverse situations: inviscid, transonic, laminar, and high Reynolds number turbulent flows.

Adaptation based on local isotropic refinement requires a decent mesh to start with, that is for instance stretched elements along the wall for Navier-Stokes calculations. On that condition, it can handle high aspect ratios. One must pay attention though to underresolved curved boundary layer regions where mesh deformation to match the surface definition can be an issue.

In order to maximize the gain over uniform grid refinement, the criterion threshold needs to be tuned. Too much refinement at each adaptation level will slow down the whole process. Nevertheless there is always a slight CPU advantage for adapted meshes with an equivalent number of nodes: they have less points in the freestream and thus require fewer time steps to converge at a given CFL number.

The different force coefficients converge at different rates. This is even more true with goal-oriented adaptation. Only the targeted quantity tends to see the benefit of the refinement; the adjoint does not seem to improve the solution globally. Other cost functions should be tested, especially specific Navier-Stokes ones. Multiple targeted adaptation might be the solution. In any case, goal oriented adaptation is still more versatile than feature or residual based refinement: it works for shocks, boundary layers, and wakes, even with a simple pressure cost function.

Local isotropic refinement not viable in 3-D; the number of nodes grows too fast. Anisotropic refinement/derefinement or a remeshing capability are needed for complex 3-D applications. For industrial use, an automated procedure is needed.

Running several meshes for convergence is a turn down for the technique. Human cost must not overwhelm the CPU advantage.

The method should possibly be coupled with higher-order elements for augmented performance. In principle the adjoint obtained through automatic differentiation should work as is with higher order elements.

References

1. F. Chalot and T.J.R. Hughes. A consistent equilibrium chemistry algorithm for hypersonic flows. *Computer Methods in Applied Mechanics and Engineering.* vol. 112, pp. 25–40, 1994.
2. F. Chalot, M. Mallet, and M. Ravachol. A comprehensive finite element Navier-Stokes solver for low- and high-speed aircraft design. Paper #94-0814. *AIAA 32nd Aerospace Sciences Meeting.* Reno, NV, January 10–13, 1994.
3. F. Chalot, Q.V. Dinh, M. Mallet, A. Naïm, and M. Ravachol. A multi-platform shared- or distributed-memory Navier-Stokes code. *Parallel CFD '97.* Manchester, UK, May 19–21, 1997.
4. L. Hascoët and V. Pascual. *TAPENADE 2.1 user's guide.* Rapport Technique N° 300. INRIA. September 2004.
5. N. Pierce and M. Giles. Adjoint and defect error bounding and correction for functional estimates. *Journal of Computational Physics.* Vol. 200, pp 769–794, 2004.