Journal of Fluids and Structures (2002) **16**(8), 1127–1144
doi:10.1006/jfls.2002.0464. available online at http://www.idealibrary.com on IDEAL®

SIMULATION OF AEROELASTIC MESOFLAPS FOR SHOCK/BOUNDARY-LAYER INTERACTION

B. Wood, E. Loth and P. Geubelle

Department of Aeronautical and Astronautical Engineering, 104 South Wright Street, 306 Talbot Lab, University of Illinois at Urbana-Champaign, Urbana, IL 61801, U.S.A.

(Received 12 April 2000; and in final form 11 November 2001)

A novel concept involving an array of mesoflaps that allow for aeroelastic recirculating transpiration has the capability to control shock/boundary-layer interactions. The concept consists of a matrix of small flaps (rigidly fixed at their upstream end and covering an enclosed cavity) which are designed to undergo aeroelastic deflection to achieve proper mass bleed or injection when subjected to gas dynamic shock loads. To investigate the static behavior of the mesoflap system, a loosely coupled aeroelastic finite element scheme was developed. The technique uses an unstructured grid for both the fluid and solid domains to allow for potentially complex geometries. Issues of optimum fluid cycles per aeroelastic iteration, under-relaxation, and adaptive mesh re-gridding versus motion were considered in the context of the flap deflection. The aeroelastic convergence was accelerated and improved by employing such techniques.

1. INTRODUCTION

1.1. Conventional Inlet Bleed Systems

Future aerospace vehicle development will depend on the development of advanced flow control systems for both internal and external aerodynamics. In particular, the interaction between shock waves and turbulent boundary layers is a primary determinant of the performance of high-speed aircraft. Most engine inlets on military aircraft operating at speeds above Mach 2 employ active bleed, which requires ducting of bleed flow to an external surface where it is discharged (Gridley & Walker 1999; Wood *et al.* 1999). The amount of bleed required increases significantly with Mach number and is on the order of 10–15% of the engine mass flow for Mach 3. Associated penalties including drag, weight, and cost of the overall vehicle are directly related to this bleed percentage. Boundary-layer bleed can also be critical with respect to external aerodynamics. For example, bleed on a transonic airfoil can help prevent shock-induced boundary-layer separation which can otherwise substantially increase drag and/or reduce effectiveness of downstream control or lifting surfaces.

However, there are significant performance advantages if one can combine both suction downstream of the shock with flow injection upstream of the shock, while eliminating any net bleed from the inlet. Passive flow control of the shock/boundary-layer interaction (SBLI) region has been suggested as a promising method to reduce the detrimental effects of strong shock waves, especially for supersonic engine inlets. Such a recirculating transpiration system uses a porous surface covering a cavity and the high-pressure

0889-9746/02/081127 + 18 $35.00/0

downstream of the shock allows for bleed, while the low-pressure upstream yields blowing. Since there is no net mass removal, bleed ducting and dumping systems are eliminated. However, success for conventional recirculating transpiration system has been limited in part due to insufficient transpiration rates because of limits on hole geometry and surface porosity (Bur *et al*. 1998), and increased drag penalties for no-shock flow conditions due to the effective roughness of the plate (Laurendeau 1995).

1.2. Flow Control System

The goal of the mesoflap system is to efficiently reduce and prevent flow separation caused by SBLI with no net mass removal, thereby improving the resulting downstream boundary-layer characteristics. The concept consists of a matrix of mesoflaps covering an enclosed cavity as shown in Figures 1 and 2. Each flap is rigidly fixed along a small portion of its upstream end, but can bend aeroelastically at its downstream end based on the pressure difference between the supersonic flow above and the subsonic cavity flow below. Under no-shock conditions, i.e., subsonic flow, the streamwise pressure distribution on the flaps is nearly constant such that no transpiration is induced [Figure 1(a)]. Therefore, the pressure within the cavity is nearly equal to that of the flow above, and the flaps do not deflect. The advantage of this approach is that the roughening of the surface caused by conventional transpiration holes or slots is avoided.

Figure 1(b) shows the system for the conditions of an oblique shock impingement, which induces a strong streamwise pressure variation. In general, the nearly uniform cavity pressure lies roughly between the low pressure of the pre-shocked flow and the high pressure of the post-shocked flow. As such, the flaps downstream deflect downward to allow angled bleed from the boundary layer into the cavity. Similarly, the flaps upstream of the reflection deflect upward allowing flow injection angled into the boundary layer, thus recirculating the flow. Therefore, the cantilever flap deflection due to the aerodynamic pressure distribution achieves the locally desired bleed or injection (both at low aerodynamically efficient angles). By constructing these flaps in a matrix in both the spanwise and streamwise directions (Figure 2), such deflection can accommodate variations in both streamwise position and sweep angle of the impinging shock.

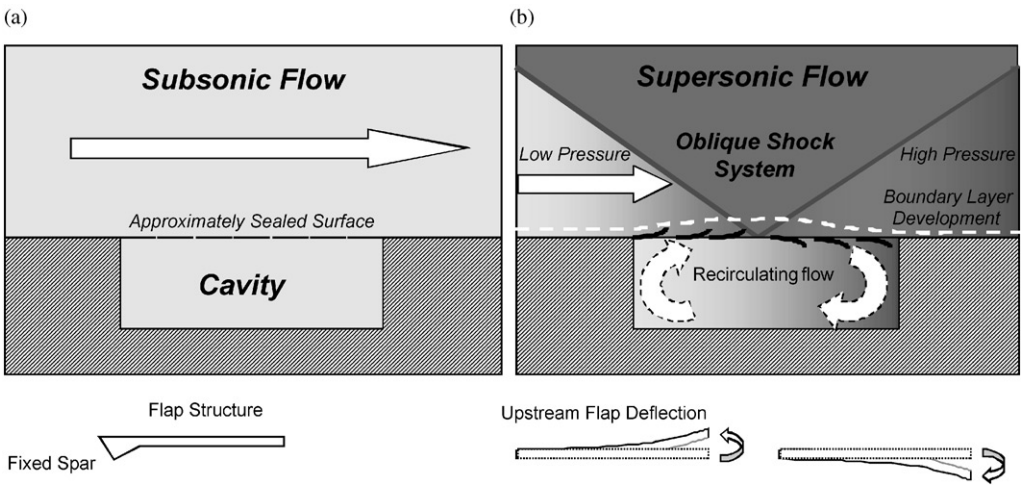

Figure 1. Aeroelastic mesoflaps over a cavity for (a) subsonic flow and (b) supersonic flow with impinging oblique shock.

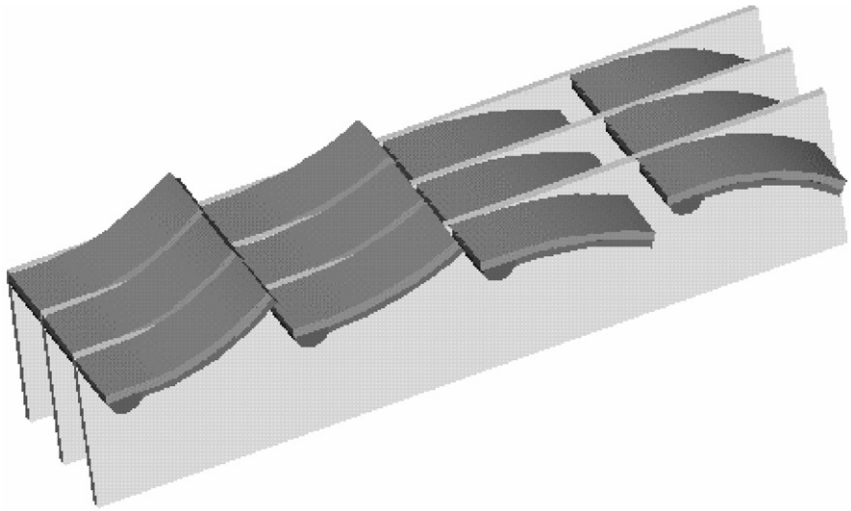

Figure 2. Three-dimensional schematic of mesoflap system showing a portion of mesoflap matrix.

The deflections of these mesoflaps are on the order of a few millimeters. This physical scale for upward deflection is bounded by the limitation that they should be less than the sonic thickness of the oncoming boundary layer to minimize drag generation. However, the deflections should be large enough to ensure sufficient mass flow capturing of the boundary layer and take advantage of the high convection speeds within the sonic region to drive the transpiration. The resulting mesoflap geometry is submillimetric in thickness and subcentimetric in length. Therefore, a large number of flaps would be typically needed over a bleed section for a full-scale inlet application. Experiments have been recently conducted for the oblique and normal shock configurations (Gefroh *et al*. 2000; Hafenrichter *et al*. 2001). These studies have shown that improvement in the boundary-layer health and stagnation pressure recovery are possible with the mesoflap system. In addition, it was noted that, in most cases, the flap deflections are static: only excessively thin flaps that experience large deflection angles (e.g., in excess of 15°) have been found to experience flutter.

As such, the mesoflap concept retains the self-contained simplicity of a conventional passive-transpiration system (eliminating bleed airflow ducting and dumping) while also featuring three additional advantages based on the system's unique structural geometry. (i) The geometry improves bleed and injection aerodynamics (i.e., increased sonic mass coefficient) by employing angled transpiration. (ii) It allows variable streamwise position and sweep angle of the shock while retaining angled transpiration. Therefore, the exact shock impingement locations are not required to be known a *priori* as a function of flight/inlet conditions. (iii) It increases high aerodynamic efficiency under subsonic flow (i.e., it leads to skin friction consistent with that of a solid wall) since the system simply reverts to a nearly smooth flat plate when no shocks are present.

2. COMPUTATIONAL METHODOLOGY

The basic objective of this work is to construct and validate a detailed computational static model of the aeroelastic system to investigate the flow physics and the system

performance under steady state conditions. The analysis performed herein focuses on the complex fluid–structure interaction between the supersonic viscous flow and the cantilever plates. Such a complex coupled description is employed in this case, since the aeroelastic deformations of the mesoflaps are expected to influence strongly the turbulent boundary layer, which in turn describes the gas dynamic loads applied to the mesoflaps as well as the flow transpiration, which governs the system performance. Therefore, the present numerical study combines the fully compressible turbulent flow equations with those of the theory of elasticity. This combination of computational fluid dynamics (CFD) and computational solid mechanics (CSM) has been integrated into a single code. In the following sections, we briefly review the basic features of the CFD and CSM components and then describe the coupling technique used in the aeroelastic code.

2.1. CFD Formulation

The CFD component of the aeroelastic code is based on the finite element code NSU2D (Mavriplis 1991). This code solves the Reynolds-averaged Navier–Stokes equations using a Galerkin finite element approach. The numerical scheme employs a conservative Taylor–Galerkin explicit time stepping approach to march to the steady state solution of the nonlinear system. The code can apply a variety of techniques for convergence acceleration, including local time stepping, residual smoothing, and multigrid algorithms. The present flows are at such high speeds that the flow may be considered turbulent throughout, such that a Reynolds-average Navier–Stokes solution is appropriate. It should be noted that, while turbulence models are typically necessary for solution of such complex flows, they are also empirical and therefore should be validated for similar flows if robust performance is to be expected. The compressible turbulence model employed herein is the Spalart–Allmaras one-equation model which has been shown to robustly simulate boundary layers, free shear layers, and resulting confluent boundary layers with flow compressibility (Dunn & Loth 1999). In the following, the model is also shown to provide reasonable aerodynamic results for oblique shock–wave interactions with a turbulent boundary layer under both attached and separated flow conditions, which are conditions of direct relevance to the mesoflap problem.

The fluid solution domain is composed of two types of meshes: structured (viscous) meshes and unstructured (inviscid) meshes. The latter elements are fully adaptive mesh and are refined locally based on relevant flow gradients. The structured meshes are wrapped around the region immediately surrounding the individual mesoflaps and viscous surfaces, which are each defined with spline curves. The viscous surfaces include the lower wall and the cavity surfaces as shown in Figure 3. The structured viscous grids are created with an advancing layer method with a geometrically increasing wall normal size. Once the

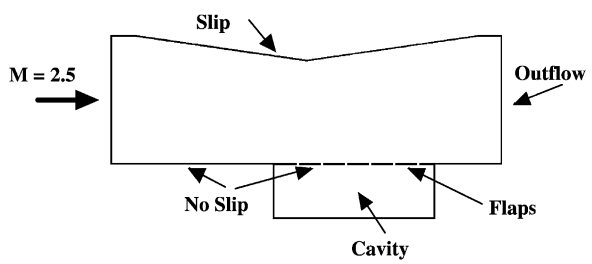

Figure 3. Diagram of boundary conditions for the flap system.

aspect ratio of the viscous grid elements becomes unity, the unstructured grid is advanced to the exterior boundaries using an advancing-front method near the solid surfaces and a Delaunay triangulation method in the rest of the domain. This change from structured to unstructured grids occurs typically at twice the upstream boundary-layer thickness. A typical mesh used in the aeroelastic simulations conducted in the present study is presented in Figure 4. The fluid mesh is composed of approximately 100 000 nodes, with as many as 70 000 of those in the viscous structured regions to provide sufficient resolved (i.e., grid-independent) description of the boundary-layer flow.

Unstructured meshes are used to adapt to compressible flow features. The supersonic nature of the flow necessitates special consideration in regions of large discontinuities, namely shock waves. First-order artificial viscosity is included only in the vicinity of shocks to prevent numerical oscillations, but carefully balanced to avoid excessive dissipation (Mavriplis 1991). Additionally, the adaptive meshing capabilities of the code are exploited to refine the grid near shocks. The unstructured grids are adapted based on density gradients using the technique developed by Dunn & Loth (1999), in which any unstructured element with a gradient exceeding a specified tolerance is divided in two. Each mesh refinement increases the total number of nodes by about 15%.

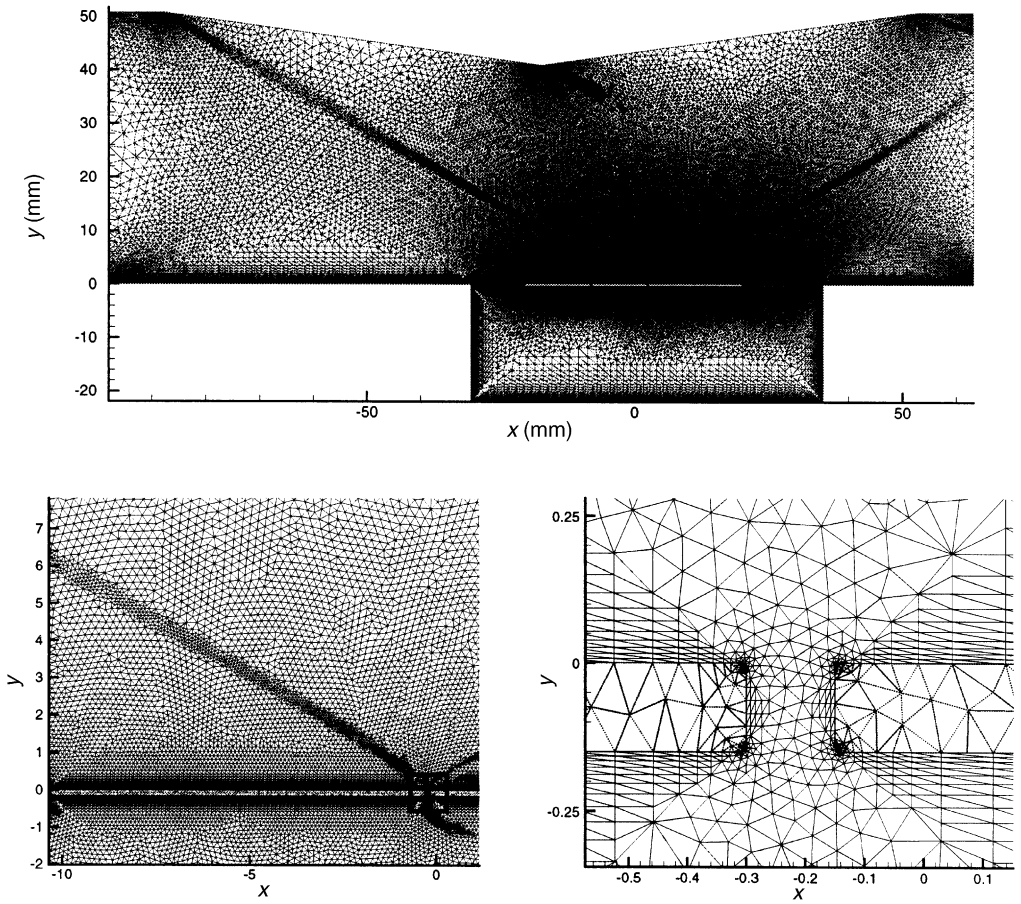

Figure 4. Finite element mesh used to discretize the fluid and solid domains.

The fluid mesh and solver have been validated through detailed convergence studies and comparison with relevant experimental data. The viscous meshes employ an initial grid point placement from the flap surface at about a y^+ of unity and use a 15% geometric increase in normal-wall mesh spacing thereafter. The vertical grid resolution has been validated with turbulent skin friction distributions for subsonic and supersonic flow over a flat plate with a zero pressure-gradient condition (Wood 1999; Wood *et al.* 1999). For an oblique shock interacting with a boundary layer exhibiting flow separation, the minimum horizontal grid spacing needed to achieve grid independence has been found to be one-quarter of the upstream boundary-layer thickness, i.e., $\Delta x < 0{\cdot}25\delta$ (Wood 1999). To validate the pressure distribution for a shock–boundary-layer interaction problem (critical for obtaining the correct aeroelastic behavior), computations were compared with experiments performed at NASA (Porro & Hingst 1993). A diagram of the experimental set-up is shown in Figure 5, where one can see the undesired deficit in the velocity profile near the reflecting oblique shock wave. The shock generator angle was adjusted to create shock waves of various strengths in a Mach 2·5 flow, which then interacted with the 37·3 mm thick boundary layer, with $\mathrm{Re} = 1{\cdot}74 \times 10^7\,\mathrm{m}^{-1}$. The experiments and the simulations examined shock generator angles of 2, 4, 6, and 8°. Pressure profiles along the wall are shown in Figure 6 for these cases, indicating reasonable predictive performance. In addition, flow separation was found to occur only at the two larger angles (6 and 8°) in both the experiment and computations.

2.2. CSM Formulation

To simulate the bending response of the mesoflaps, a conventional plane-strain linearly elastic finite element code has been used. The present results indicate that strains and rotations remain relatively small, justifying the linearization of the problem. Due to the relatively simple rectangular shape of the flaps considered in the present study, the major part of the structural domain can be discretized with a structured mesh composed of triangular or quadrilateral elements. The triangular element option has been favored here due to its greater flexibility in the discretization of the flap rounded corners (Figure 4) and of the more complex geometries (such as the inclusion of spars) envisioned for future designs of the flap array. The triangular element mesh generator used in this study is Triangle (Shewchuk 1996), which relies on a boundary conforming Delaunay triangulation method.

To assess the ability of the structural component of the aeroelastic code to capture accurately the bending response of the mesoflaps, a detailed convergence study has been

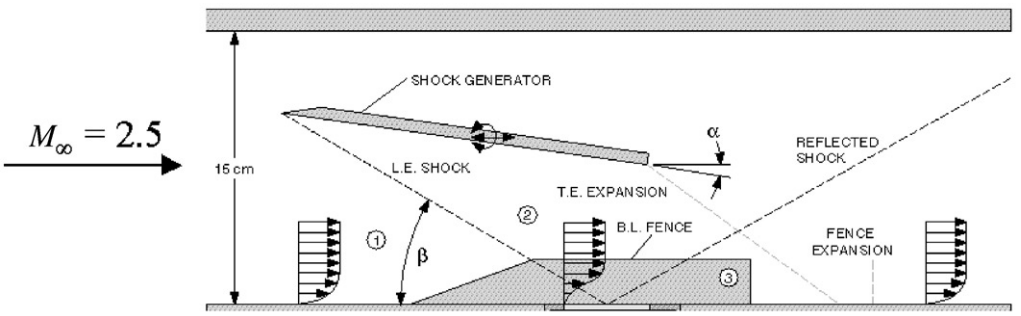

Figure 5. Diagram of SBLI experimental setup.

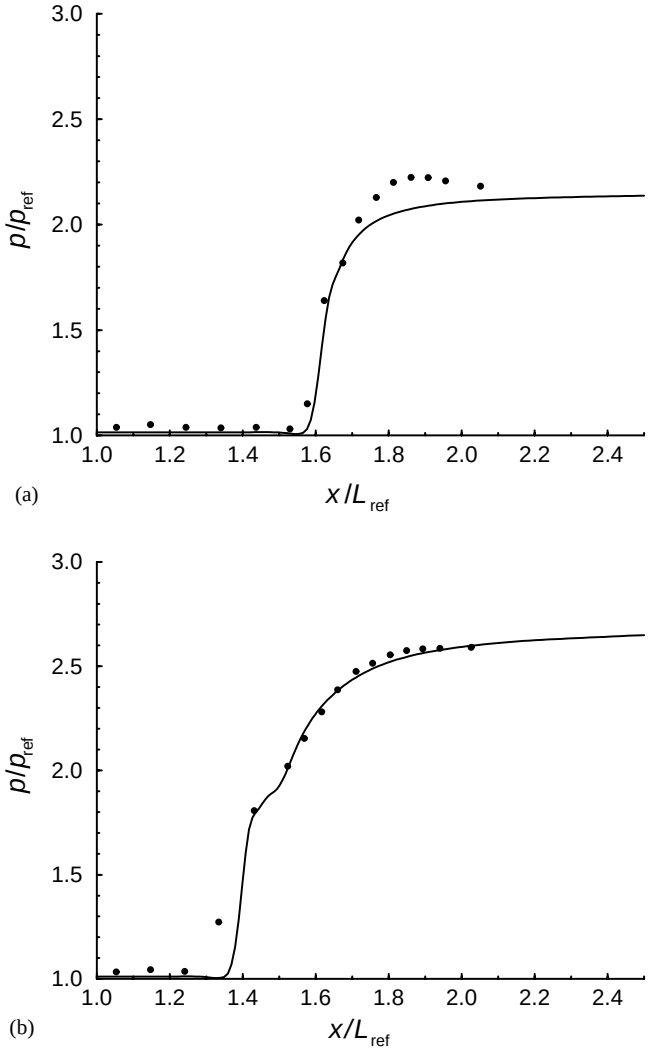

Figure 6. Pressure profile of experiments (symbols) and NSU2D predictions (lines) for shock generator angles of: (a) 6° and (b) 8°; ●, experiment; —, NSU2D Code.

conducted on the classical problem of a cantilever beam subjected to a uniformly distributed transverse load. The study compared the performance of structured and unstructured triangular meshes composed of constant (three-node) or linear (six-node) strain elements, with various aspect ratios. The detailed results are given in Wood (1999). As expected, constant strain elements behaved very poorly in bending and required a very fine discretization of the flaps, while fairly coarse meshes with a few as two or three linear strain elements through the flap thickness (Figure 4) captured the bending response of the flaps very accurately.

Due to the relatively small number of solid domain nodes (each mesoflap interior is typically discretized with approximately 4 000 nodes), a direct symmetric banded solver was used at each aeroelastic iteration. The Gibbs–Poole–Stockmeyer method for automatic bandwidth reduction was employed, which typically reduced the bandwidth from around 1 000 to about 50.

2.3. Aeroelastic Coupling Approach

A schematic representation of the algorithm employed to model the static aeroelastic design problem is shown in Figure 7. As indicated by the aeroelastic iteration loop, a loosely coupled approach (Löhner *et al*. 1995) is employed to simulate the fluid/structure

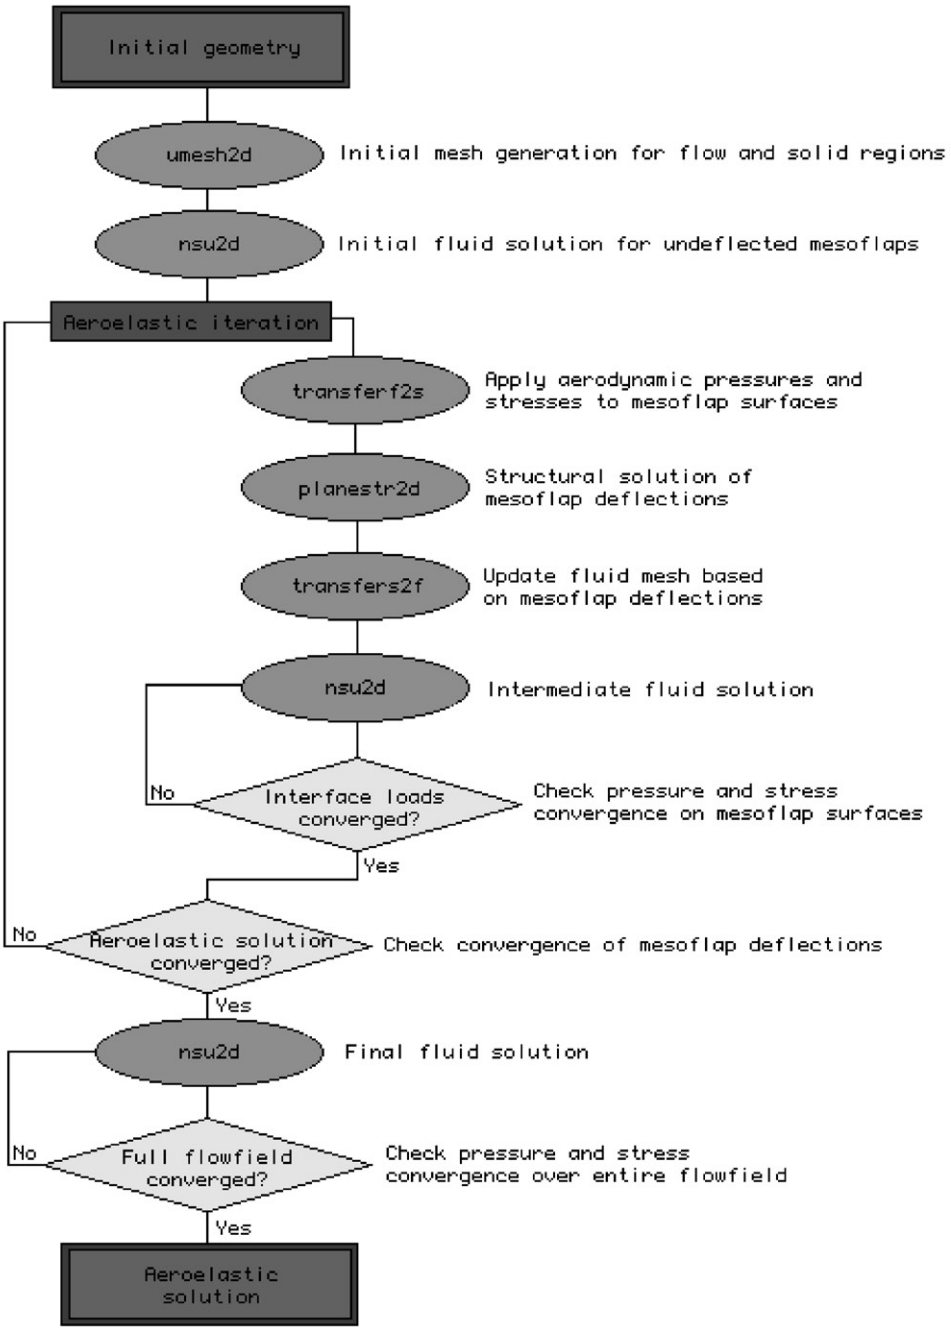

Figure 7. Flowchart of aeroelastic coupling algorithm.

interaction, whereby two different computational modules are employed in turn to solve the fluid and structure fields. There are several reasons why loose coupling is used in the present work instead of strong coupling, in which the solid and fluid nodal values are solved simultaneously. Firstly, the solid field equations are linear and elliptic while the flow equations are primarily hyperbolic and nonlinear. As such, the optimum solution schemes are substantially different. Secondly, the physical time scales for fluid response and structural response are much different such that a simultaneous time-marching (including "pseudo" time-marching to steady state) scheme with direct coupling would result in a much slower convergence rate for this problem. This coupling approach allows for an easier integration of existing fluid and structure codes, as well as easier upgrades of the solvers (for example, to incorporate more complex material behavior in the structure model). It should also be noted that a closely coupled scheme (in time) is not necessary for static analysis.

In the loosely coupled approach, the information relative to the interface values is passed back and forth between the CFD and CSM solvers at chosen intervals (aeroelastic iteration) until convergence is achieved. The aeroelastic iterations are chosen such that the fluid stresses acting on the mesoflap surfaces (not necessarily the entire flow domain) are reasonably converged, after which the structural problem is solved directly to within reasonable precision. The description of convergence criteria for the fluid and the solid fields for each iteration is discussed in the next section. After a converged aeroelastic deflection is obtained, a high-precision computation is completed in the final fluid solution (last ellipse in Figure 7) to ensure that the entire flow field is computationally converged to compute boundary-layer thickness, shape factor, etc.

The interface information (transferred from the fluid solver to the solid solver) consists of the boundary loads associated with the pressure and viscous stresses values computed by the fluid solver at each node. For every edge on the structural boundary, the nearest fluid nodes on either side of the edge center are found (Figure 8). The normal and tangential loads applied to the structural model are then computed through a linear interpolation of the pressure and shear stress, respectively. This interpolation method allows for arbitrary alignment of fluid and structure nodes along the interface. This method is desired since, although the fluid boundary points match up with structural boundary points (see previous section), there are generally additional structural boundary points which do not coincide with fluid nodes.

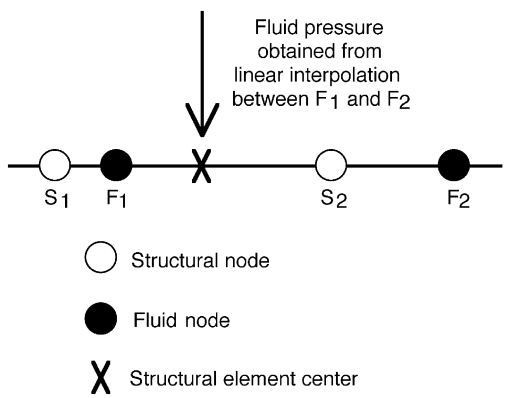

Figure 8. Illustration of load interpolation method.

2.4. Fluid Mesh Adaptation to Aeroelastic Behavior

Once the boundary displacements are computed by the structure solver, they are transferred to the fluid solver. Transferring displacements to the fluid domain involves updating the fluid mesh with a new geometric boundary description. Two different methods have been investigated for this procedure: full remeshing and mesh motion.

Full remeshing simply starts from scratch with the boundary description and generates a new mesh. This method has the advantage of simplicity, and minimal computer time. Additionally, the quality of the mesh (if no grid refinement was used) is essentially as good as the original mesh, since the same parameters, aside from some changes to the boundary description, are used. However, this method requires that the flow solution be interpolated from the old mesh onto the new one, and local errors may be introduced in this process requiring additional convergence cost. For example, if the tip of a flap were to deflect downward by more than its own thickness, a flow value which was previously below the flap would then be above it. Additionally, regenerating the mesh from scratch means that any mesh refinement which was previously performed is lost.

The second technique, mesh motion, involves computing new locations of the internal nodes of the fluid mesh based on the displacements of the boundary nodes. In particular, a linear elastic analogy is used, in which the motion $\boldsymbol{u}$ of the fluid nodes is obtained by treating the fluid domain as a simplified solid satisfying

$$\nabla\cdot([1+\tau]\nabla\mu) = 0$$

and subjected to imposed displacements corresponding to the motion of the fluid–structure interface. In order to limit the distortion of the small elements located in the immediate vicinity of the fluid–structure interface, the "stiffness" variable τ is chosen as a function of element size, similar to the method proposed by Masud & Hughes (1997). The linear formulation of Masud & Hughes (1997) was employed for inviscid grids where the ratio of minimum to maximum mesh element areas was limited to around $\Delta_{\min}/\Delta_{\max} \approx 10^{-2}$. However, the present viscous mesh has area ratios of $\Delta_{\min}/\Delta_{\max} \approx 10^{-7}$ in order to properly resolve the turbulent boundary layer. Such small element size ratios would lead to an excessively ill-conditioned problem for the linear formulation of Masud & Hughes (1997). Therefore, in the present study, the stiffness as a function of the element size (Δ_e) is given as

$$\tau(\Delta_e) = \tau_{\text{ref}}\left(1+(q-1)\exp\left[\frac{-(\Delta_e/\Delta_{\max})}{(\Delta_{\min}/\Delta_{\max})}\right]\right),$$

where τ_{ref} is a chosen reference stiffness, and q and α are user-defined parameters which prescribe a maximum stiffness ratio for $\Delta_e \to 0$. Use of $\alpha = 0{\cdot}1$ and $q = 100$ worked well for the disparity in element sizes appearing in the present simulations. The mesh motion solution is obtained with an iterative conjugate gradient method.

While the mesh motion technique avoids the need to interpolate the solution to the new mesh, as well as retaining any grid refinement previously done, it has been found to be computationally less efficient than complete remeshing. The process of solving the mesh motion differential equation for the large number of nodes in the viscous simulations (e.g., 100 000) requires significantly more computer time than full remeshing. On a single-processor Origin2000, mesh generation typically takes about 30 s, while mesh motion calculations may take 10–15 min. Additionally, regardless of the variable stiffness scheme described above, some mesh distortion of the elements located between the closely spaced flaps limit is unavoidable for large deflections. Figure 9 shows the region between two flaps after a mesh update, including both the full remeshing and mesh motion techniques and

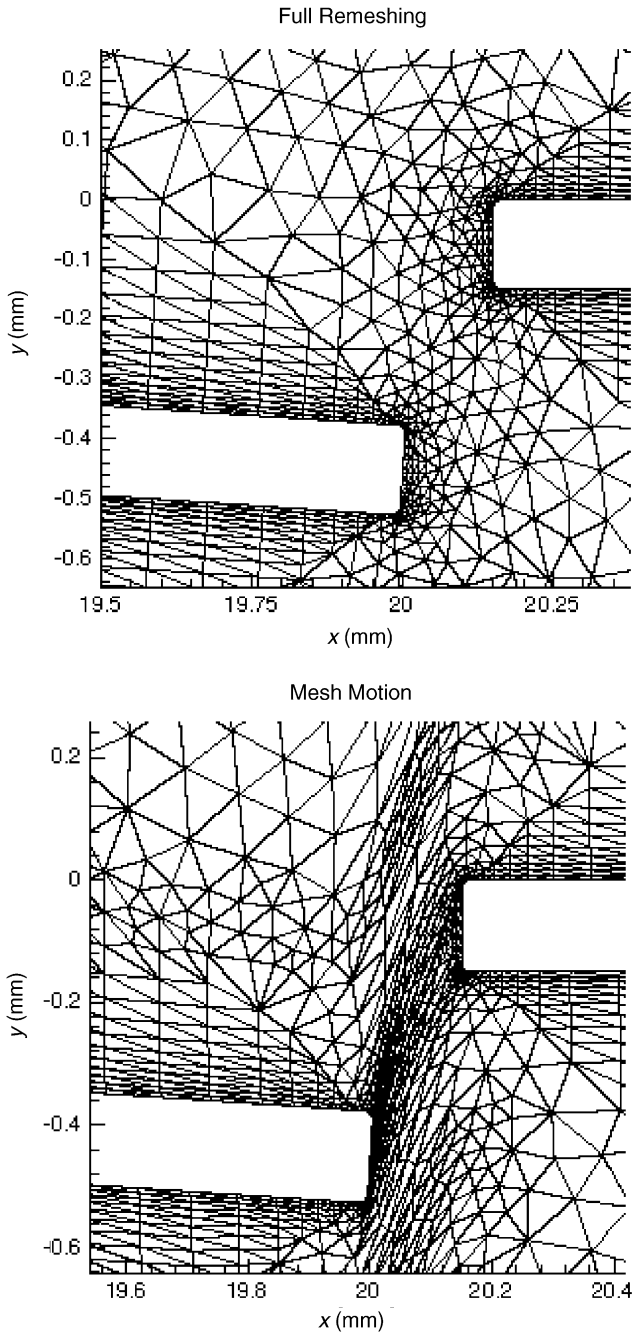

Figure 9. Mesh close-up between two flaps after a mesh update with significant flap deflection, showing the full remeshing (top) and mesh movement techniques (bottom).

illustrating the element stretching clearly. Even with these disadvantages, mesh motion is useful when the solution is close to converging, at which point boundary displacements are relatively small (thus limiting the amount of element stretching) and errors introduced due to full remeshing are significant.

Therefore, both mesh update techniques are used in this study, depending on the amount of displacement that occurs. Full remeshing is used for $\delta u/u_{\text{tip}} > 0{\cdot}2$, since mesh motion would introduce excessive element stretching; here δu refers to the deflection changes (i.e., maximum displacement change between two iterations) and u_{tip} refers to the average flap deflection (at a given iteration). For smaller deflection changes, which should occur closer to convergence, the mesh motion technique is used to preserve the solution between iterations.

2.5. Aeroelastic Convergence Acceleration

Convergence of the entire aeroelastic solution includes convergence of both the structural displacements and the flow solution in the entire flow field. This is achieved by first converging on the interface displacements through the aeroelastic iteration loop (Figure 7) then performing a final fluid solution to converge on the flow field values. For the equations appearing in this section, we define i and j as counters for aeroelastic iteration and fluid cycle, respectively, with N as the total number of aeroelastic iterations and M_i as the total number of fluid cycles for the ith iteration (see Figure 10 for a diagram illustrating these values). In the following, the term "total fluid cycles" refers to $\sum_{i=1}^{N} M_i$, which should generally be minimized (for a given convergence criterion) for the optimum scheme.

It is useful to examine first the fluid convergence behavior of a fixed deflection solution (i.e., CFD only) to provide a baseline for what can be expected in the aeroelastic solution. The convergence of pressure values along the flap surfaces is important during the aeroelastic iterations, while the wall skin friction downstream of the cavity is of primary interest during the final fluid solution. Simulations by Wood (1999) show that the pressure clearly converges faster than the skin friction during the initial part of the solution, although the convergence behavior for both variables is similar afterwards. For example, the pressure error remains below about 6% after around 1 000 cycles, while it takes nearly 2 500 cycles for the skin friction to reach the same level of precision.

Flap boundary displacement convergence is evaluated by examining the change in maximum displacements between the current iteration and the previous one, averaged over all the flaps and normalized by the average tip deflection. We thus define the displacement residual ε_d^i for the ith aeroelastic iteration as $\varepsilon_d^i = \delta u^i / u_{\text{tip}}^i$, with

$$\delta u^i = \frac{1}{n_{\text{flap}}} \sum_{f=1}^{n_{\text{flap}}} |\max\|\mathbf{u}_f^i\| - \max\|\mathbf{u}_f^{i-1}\||$$

and

$$u_{\text{tip}}^i = \frac{1}{n_{\text{flap}}} \sum_{f=1}^{n_{\text{flap}}} \max\|\mathbf{u}_f^i\|,$$

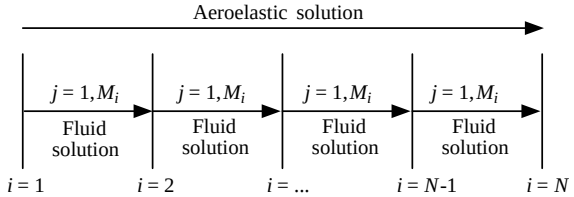

Figure 10. Diagram of variables describing aeroelastic iterations and fluid cycles.

where $\mathbf{u}_f^i$ are the displacements of the interface nodes on flap f for iteration i, and n_{flap} is the number of flaps. The aeroelastic iteration is stopped (considered converged) when δu becomes smaller than a given tolerance (a typical value is 10% of the flap thickness, i.e., $\delta u/t_{\text{flap}} < 0{\cdot}1$), at which point the final fluid solution is run to achieve final fluid dynamic convergence.

Initial simulations with the aeroelastic code produced significant oscillations in the flap deflections throughout the solution. It should be noted that while aeroelastic flutter can generally occur (Kornecki *et al.* 1976), recent experiments (Gefroh *et al.* 2000; Hafenrichter *et al.* 2001) have indicated that flap deflections are typically static. As such, the numerical oscillations are attributed to poor convergence qualities of the loosely coupled approach. This problem was remedied with a combination of the following techniques.

First, an under-relaxation scheme was adopted to transfer only a fraction of the displacement values computed by the structure solver to the fluid model. The new displacements are computed as $\boldsymbol{u}_{f_i} = \omega \boldsymbol{u}_{s_i} + (1-\omega)\boldsymbol{u}_{f_{i-1}}$, where $\boldsymbol{u}_{f_i}$ are the nodal displacement values actually transferred to the fluid model, $\boldsymbol{u}_{s_i}$ are the displacement values computed by the structure solver, and ω is a user-defined under-relaxation parameter (where $\omega = 1$ correspond to the absence of under-relaxation). Figure 11 shows the flap tip deflections versus aeroelastic iteration for a typical solution. A large amount of oscillation in the solution is seen for the case of no under-relaxation, and these oscillations are decreased substantially for smaller values of ω. This is more clearly seen in Figure 12 which shows the evolution of the boundary displacement error normalized by flap thickness, $\delta u/t_{\text{flap}}$, for the same solution. While the case with $\omega = 0{\cdot}25$ converges more smoothly, it also converges too slowly, having the highest error during the latter half of the solution. A value of $\omega = 0{\cdot}5$ provides a good balance between the other two cases, and gives the best convergence for this solution.

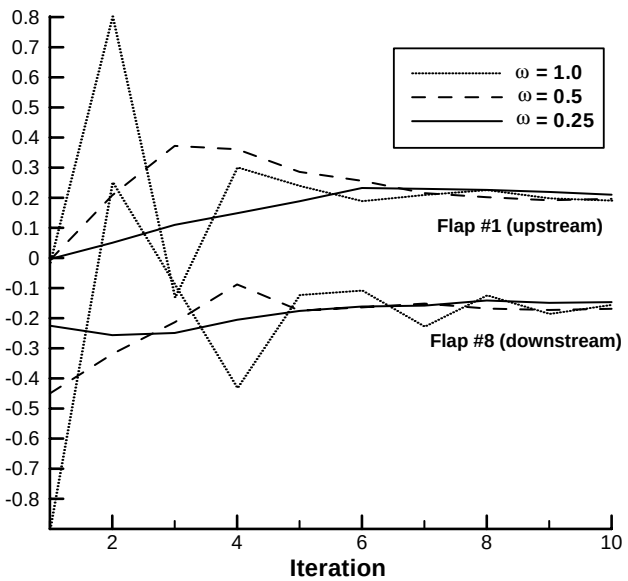

Figure 11. Normalized tip deflection of the flap aeroelastic iteration of the furthest upstream and downstream flaps for different levels of under-relaxation (ω).

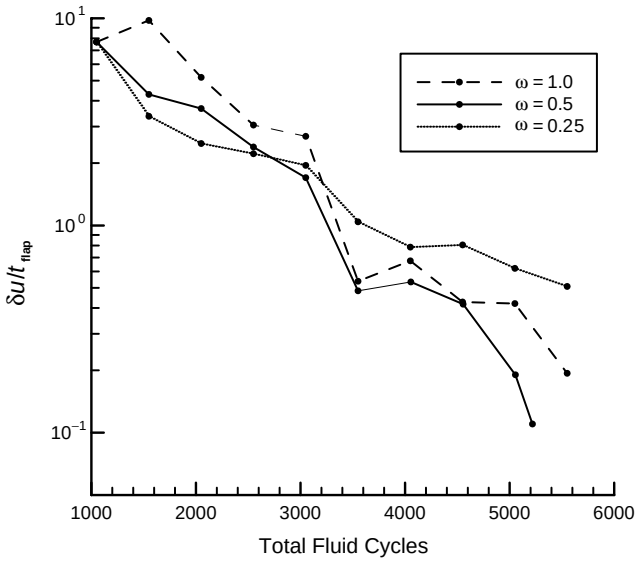

Figure 12. Displacement error versus total fluid cycles for different amounts of under-relaxation.

Since the fluid solution is an iterative process, one major issue which arises is when to stop the intermediate fluid solution during any given aeroelastic iteration. A fully resolved CFD solution of the current problem requires on the order of 5 000 fluid cycles, which uses approximately 12 cpu hours, and is clearly not an option for every iteration. Therefore, some less restrictive criteria for stopping the intermediate fluid solution must be used. First of all, one may note that the only data used from each intermediate fluid solution are the pressure values along the fluid–structure interface. If these values are not changing appreciably, there is no point in continuing with the fluid solution. Therefore, we define the pressure residual for the ith iteration, ε_p^i, as, $\varepsilon_p^i = \delta p_{M_i}^i / \Delta p_{\text{shock}}$ with $\delta p_j^i = \|\mathbf{p}_j^i - \mathbf{p}_1^i\|$, where $\mathbf{p}_j^i$ are the interface pressure values for fluid cycle j and aeroelastic iteration i and Δp_{shock} is the pressure difference across the impinging shock (which will be approximately equal to the pressure difference across the flap). In order to limit the precision of the fluid solution to that of the current aeroelastic solution, the fluid solution is stopped if $\varepsilon_p^i > \varepsilon_d^i$. This yielded a minor but noticeable improvement on the overall convergence acceleration (Wood 1999).

Therefore, in addition to the limit on δp, a limit on the number of limit cycles is also imposed. This limit must be high enough that the fluid solution has enough time to react to the changes in displacement for any given iteration, especially in cases when full remeshing is used (since full remeshing can introduce additional errors along the interface). On the other hand, the limit should be low enough to keep the cpu time required for each iteration to a reasonable amount. For the present problem, a parameteric study (Wood 1999) yielded an optimum cycle cap of 500 for minimizing the *total* cycles required for convergence, i.e., $M_i = \min(500, j\varepsilon_p^i \varepsilon_d^i)$ for all iterations.

In summary, several techniques were examined for improving the aeroelastic convergence including: loose-coupling, under-relaxation, variable convergence criteria for fluid cycles in an aeroelastic iteration, capping of fluid cycles in an aeroelastic iteration, and mesh motion. The combination yielded faster overall convergence. For example, the simulations initially required as many as 30 aeroelastic iterations, which when combined with a full-fluid convergence approach for each aeroelastic iteration would require

approximately 150 000 total fluid cycles. With the aforementioned techniques in place, fully converged aeroelastic solutions are obtained in about 6 000 total fluid cycles.

3. SAMPLE AEROELASTIC SIMULATIONS

The sample flap system examined herein with the above methodology includes oblique shocks interacting with a fully turbulent boundary layer to provide direct relevance to a mixed internal supersonic inlet application. For all of the results shown here, the following flow conditions are used: free-stream Mach number $M_\infty = 2{\cdot}5$, free-stream total pressure $P_o = 5{\cdot}2 \times 10^5\,\mathrm{N/m^2}$ (75 psia), Reynolds number $\mathrm{Re}_\delta = 1{\cdot}16 \times 10^5$ and boundary layer thickness (just before the cavity) $\delta \approx 3{\cdot}3$ mm. The shock wave is generated using an 8° bend in the upper wall, and is positioned such that the inviscid shock impingement location on the lower wall is in the center of the cavity. The boundary conditions and flow geometry are shown in Figure 3. Additionally, the flap material is assumed to be aluminum with elastic modulus $E = 6{\cdot}9 \times 10^{10}\,\mathrm{N/m^2}$ and Poisson's ratio $\nu = 0{\cdot}3$.

The aeroelastic solution presented here involves eight flaps, each 5 mm long, with 75% of the length allowed to deflect while the first 25% is kept rigid. The flap thickness is 60 μm

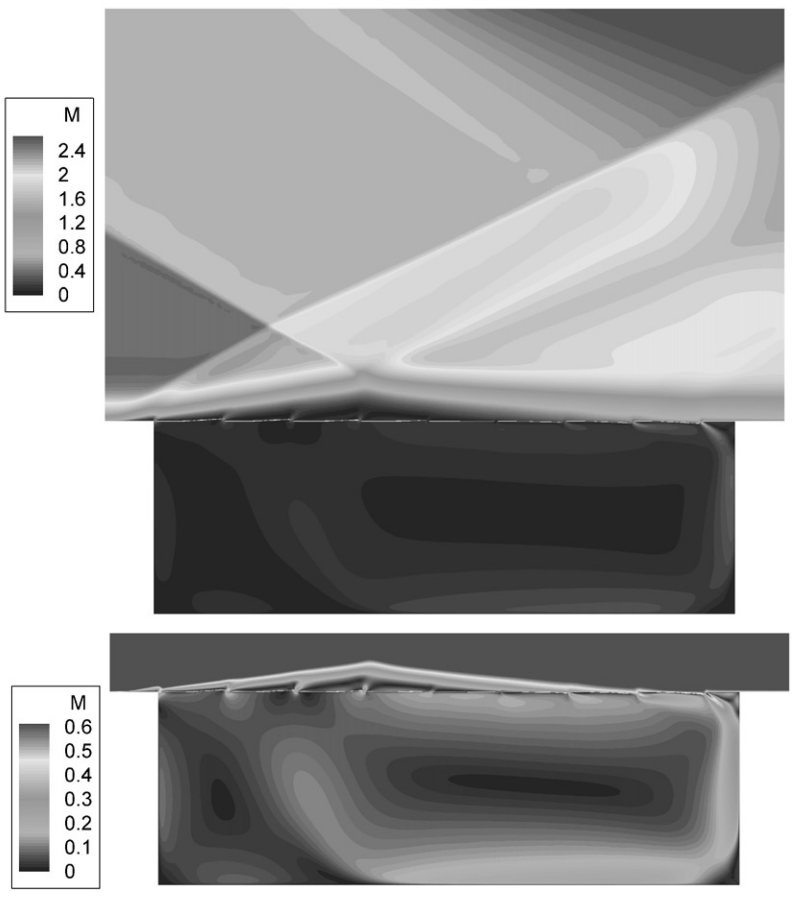

Figure 13. Aeroelastic solution: Mach contours overall (top) and resolved in the cavity region (bottom).

in order to produce significant deflections, and flap spars were neglected in order to simplify the geometry (Wood 1999). This solution is obtained with under-relaxation parameter $\omega = 0{\cdot}5$. After 10 aeroelastic iterations, the final fluid solution used an additional 2 000 fluid cycles, with mesh refinement occurring at the beginning and after 1 000 cycles to capture the impinging and reflected shock waves.

Mach contours are displayed in Figure 13(a), illustrating the impinging shockwave and its interaction with the boundary layer, while Figure 13(b) shows Mach contours resolved inside the cavity. Bleed flow is clearly seen over the last three flaps, and there are actually two recirculation regions inside the cavity, with a smaller reverse recirculation occurring underneath the first two flaps. Pressure contours are plotted in Figure 14(a), showing the increase in pressure after both the impinging and reflected shocks which bends the downstream flaps downward. Note that due to the large amount of separated flow, the pressure over the downstream flaps is significantly lower than the pressure directly behind

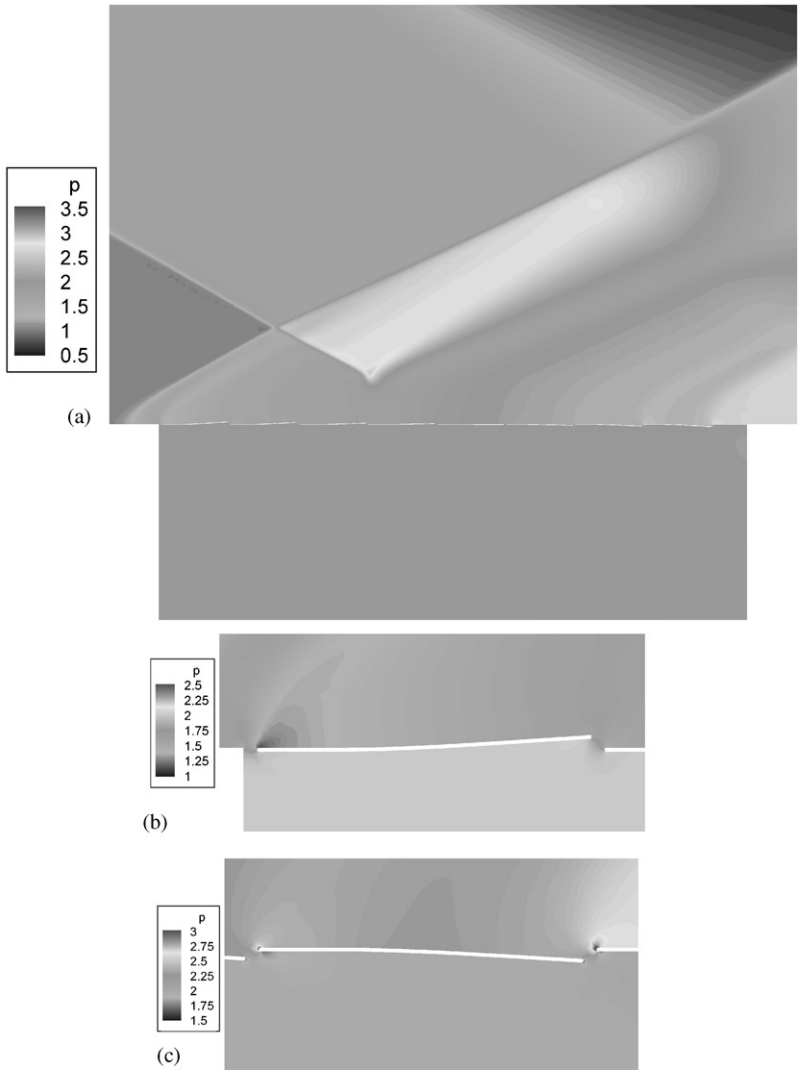

Figure 14. Aeroelastic solution: pressure contours (a) overall and resolved near the (b) first and (c) last flaps.

the reflected shock wave. Figures 14(b) and 14(c) zoom in on the first and last flaps, respectively, showing the details of the pressure distribution around them as well as a good view of the flap deflections. It should also be noted that the average gas temperature on the top of the flaps is essentially equal to that on the bottom, as expected for no-slip conditions since adiabatic conditions were assumed along with a constant turbulent Prandlt number. Since individual flaps are surrounded by a nearly constant temperature, the flap interiors will also be at this temperature. The resulting flow field demonstrates the critical aeroelastic features of the mesoflap system. Future studies will employ the above methodology to gain further understanding of the fluid–structure physics and to conduct parametric studies to optimize the mesoflap system performance, e.g., vary geometry, stiffness, and number of the flaps and vary the geometry of the cavity.

4. CONCLUSIONS

In order to investigate computationally the mesoflap concept under steady state conditions, a static aeroelastic simulation methodology has been developed. The code includes a robust fluid solver for compressible turbulent flows on unstructured grids and a conventional linear elastic finite element structure solver. In addition, procedures were developed for generating high quality matching fluid and structure meshes and for imposing a linear interpolation scheme to transfer fluid loads to the structure. Aeroelastic coupling is achieved with a loose-coupling approach, in which the solution iterates between the fluid and structure solvers until the interface displacements have converged. This loose-coupling approach allows for efficient use of pre-existing code and avoids issues of ill-conditioned problems. Convergence of the aeroelastic solution is significantly improved by using under-relaxation on the interface displacement values, and by choosing an appropriate limit on number of cycles in each intermediate fluid solution. Displacement transfer to the fluid domain is performed with full remeshing for large displacements in order to avoid excessive element stretching, and mesh motion for small displacements to avoid introducing errors by interpolating the solution to a new mesh. The code has then been used to compute an aeroelastic solution of the mesoflap system and the relevant fluid–structure features.

ACKNOWLEDGEMENTS

This work was supported by the Air Force Office of Scientific Research (AFOSR) under contract F49620-98-1-0381, with computational time furnished by the National Center for Supercomputing Applications. The discretization of the structural components has been generated using the code Triangle, created by Jonathan Shewchuk at Carnegie Mellon University (available at http://www.cs.cmu.edu/~quake/triangle.html).

REFERENCES

BUR, R., CORBEL, B. & DELERY, J. 1998 Study of passive control in a transonic shock-wave/boundary-layer interaction. *AIAA Journal* **36,** 394–400.

DUNN, T., & LOTH, E. 1999 Effects of simulated-spanwise-ice shapes on airfoils: computational investigation. AIAA Paper 99-0093.

GEFROH, D., HAFENRICHETER, E., DUTTON, J.C., MCILWAIN, S. & LOTH, E. 2000 Experiments of aeroelastic mesoflaps. AIAA Paper 2000-0355.

Gridley, M.C., & Walker, S.H. 1996 Advanced aero-engine concepts and controls. *AGARD Conference Proceedings*, *86th Symposium*, Seattle, USA, published in Neuilly sur-Seine, France.

Hafenricheter, E., Lee, Y., McIlwain, S., Dutton, J.C. & Loth E. 2001, Experiments on normal shock/boundary layer interaction using aeroelastic mesoflaps. AIAA Paper 2001-0156.

Kornecki, A., Dowell, E.H. & O'brien, J. 1976 On the aeroelastic instability of two-dimensional panels in uniform incompressible flow. *Journal of Sound Vibration* **47,** 163–178.

Laurendeau, E. 1995 Boundary-layer bleed roughness. Ph.D. dissertation, Department of Aeronautics and Astronautics, University of Washington, Seattle, WA, USA.

Lohner, R., Yang, C., Cebral, J.R., Baum, J.D., Luo, H.Pelessone, D. & Charman, C. 1995 Fluid–structure interaction using a loose coupling algorithm and adaptive unstructured grids. In: M. Hafez, K. Ushima (eds) *Computational Fluid Dynamics Review*, pp. 755–776. New York: John Wiley.

Masud, A. & Hughes, T. J. R. 1997 A space–time Galerkin/least-square finite element formulation of the Navier–Stokes equations for moving domain problems. *Computational Methods Applied Mechanics Engineering* **146,** 91–126.

Mavriplis, D., 1991 Multigrid solution of compressible turbulent flow on unstructured meshes using a two-equation model. NASA CR 187513, NASA Langley Research Center.

Porro, A.R. & Hingst, W. R. 1993 Use of surface heat transfer measurements as a flow separation diagnostic in a two-dimensional reflected oblique shock/turbulent boundary layer interaction. AIAA Paper 93-0775.

Schewchuk, J. 1996 Triangle: Engineering a 2-D quality mesh generator and delaunay triangulator. *ACM Workshop on Applied Computational Geometry*, Philadelphia, PA, USA, pp. 124–133.

Wood, B., Loth, E. & Geubelle, P.H. 1999 Mesoflaps for aeroelastic transpiration for SBLI control. AIAA Paper 99-0614.

Wood, B. 1999 Aeroelastic simulation of a novel transpiration system. M.S. thesis, Department of Aeronautics & Astronautice Engineering, University of Illinois at Urbana-Champaign.