

THREE-DIMENSIONAL MESH EMBEDDING FOR THE NAVIER–STOKES EQUATIONS USING UPWIND CONTROL VOLUMES

B. L. LAPWORTH

ROLLS-ROYCE plc, PO Box 31, Derby DE2 8BJ, U.K.

SUMMARY

A numerical model for the compressible Navier–Stokes equations using local mesh embedding is presented. The model solves for three-dimensional turbulent flow using an algebraic mixing length model of turbulence. The technique of control volume upwinding is used to produce a novel treatment, whereby the hanging nodes on the mesh interfaces are left with null control volumes. This yields an efficient discretization scheme which ensures second-order accuracy, flux conservation and stability at the mesh interfaces, whilst retaining a simple interpolative treatment for the hanging nodes. The discrete flow equations are solved using the semi-implicit pressure correction method. The accuracy of the embedded mesh solver is demonstrated by modelling the three-dimensional flow through a cascade of turbine vanes at design and off-design conditions. Mesh embedding gives a saving of 48% in the number of nodes. The embedded mesh solutions compare well with fine structured mesh solutions and experimental measurements. The capability of the embedded mesh solver to perform solution adaptive calculations is demonstrated using a two-dimensional mid-height section of the cascade at the off-design flow conditions.

KEY WORDS Mesh embedding Control volume upwinding Three-dimensionality Turbulent flow Turbine cascade

1. INTRODUCTION

Accurate solutions of the Navier–Stokes equations require fine grid resolution in regions of high flow gradients such as in shear layers or at shocks. In many applications, these regions occupy only a small proportion of the flow domain with the flow varying smoothly elsewhere. This situation is poorly suited to structured meshes where the mesh clustering needed to resolve high-gradient regions frequently results in unnecessary resolution elsewhere in the domain. In addition, the clustering produces high aspect ratio cells in the ‘freestream’ regions of the mesh which are often a cause of instability. The resolution of high gradient regions is also beset by uncertainties in their location, particularly at the outset of a calculation. The location of certain features, such as attached boundary layers, are easily determined *a priori* and may be meshed accordingly, whereas the adequate resolution of features such as shocks and flow separations can only be guaranteed by solution adaptive meshing. Flow adaption using structured meshes has been reported by, among others, Shyy¹ and Hawkins and Kightley.² Here the mesh lines of an initial mesh are redistributed according to flow-dependent weighting functions.³ This approach has the advantage that it may easily be implemented within existing structured solvers. However, the problems of grid clustering are not alleviated and the redistribution may not be able to accommodate multiple adaption criteria where conflicting redistribution requirements are to be resolved.

A more flexible approach is to locally embed fine mesh regions within an initial coarse mesh, creating a mesh with 'semi-structured' characteristics. Adaptive two-dimensional (2D) Navier–Stokes calculations on locally embedded meshes have been reported by Kallinderis and Baron⁴ and Davis and Dannenhoffer.⁵ Mesh embedding is characterized by internal interfaces in the mesh on which hanging nodes are used to facilitate the transition between the coarse and fine regions. Ideally, the discretization at the embedding interfaces should be stable, conserve fluxes and have the same order of accuracy as the discretization away from the interfaces. Additionally, the update procedure for the hanging nodes should be easy to compute. However, Almaras and Baron⁶ have reported that these criteria lead to conflicting requirements which are difficult to satisfy simultaneously. The difficulty presented by the hanging nodes (Figure 1) is that the control volume integration on the fine side of the interface generally provides an update to the hanging node, whereas on the coarse side it is preferable to treat the hanging nodes by interpolation to prevent them entering the coarse side integrations. Both these requirements cannot, except in rare cases, be satisfied by a single value at the hanging node. This leads to either a loss of accuracy or flux conservation when one method of setting the values at the hanging nodes is used in preference to the other. Kallinderis and Baron⁴ used interpolation for the hanging nodes and found that good quality solutions could be obtained even though the interfaces produced small imbalances in local flux conservation. One method of removing flux imbalances is to include the hanging nodes in the coarse side integrations. However, this produces integration rules⁷ which depend on the precise number and location of hanging nodes within the cell and is cumbersome to implement. Kallinderis⁸ also found that conservative integration rules can lead to a degradation in the numerical order of accuracy at the interface. In the present study, a novel treatment for the hanging nodes is developed which utilizes the technique of control volume upwinding.⁹ The hanging nodes are assigned null control volumes and hence have null fine side contributions to their values. The values at the hanging nodes may then be set by linear interpolation without compromising flux conservation or accuracy. Stability is ensured by preserving the upwind bias of




Figure 1. Interface node updates: (a) coarse side integrations prefer linear variation between 1 and 3; (b) fine side integrations provide update to node 2 which does not vary linearly between nodes 1 and 3

the control volume configuration. The present study also extends the mesh embedding approach into three dimensions by using meshes that are stacked in the third direction.

The study is presented as follows: firstly, brief details of the upwind control volume technique of Moore⁹ are given. This is followed by details of the control volume upwinding at the embedding interfaces. Secondly, the implementation of an algebraic turbulence model is presented. Thirdly, details of the numerical solution scheme are given. The numerical model is then used to compute the three-dimensional (3D) flow through a cascade of turbine vanes at design and off-design conditions. The embedded mesh results are compared with structured mesh results obtained using the method of Moore and Moore¹⁰ and with experimental measurements. Finally, the capability to perform adaptive calculations is demonstrated using a 2D mid-height section of the cascade.

2. THE NUMERICAL MODEL

With the assumption of an isotropic effective viscosity (μ_e), the steady Reynolds-averaged continuity and Navier–Stokes equations are

$$\nabla \cdot (\rho \mathbf{u}) = 0, \quad (1)$$

$$\rho \mathbf{u} \cdot \nabla \mathbf{u} - \nabla \cdot \mu_e \nabla \mathbf{u} = \nabla \cdot \mu_e \nabla \mathbf{u}^T - \nabla p, \quad (2)$$

where viscous compressibility effects are taken to be negligible. A constant enthalpy is assumed and the perfect gas law is used to relate pressure, temperature and density. The effective viscosity is the sum of laminar and turbulent parts. Sutherland's law is used for the laminar viscosities and a Prandtl mixing length model is used for the turbulent viscosities. Details of this model and its implementation on embedded meshes are presented in Section 3.

The flow equations are discretized using the upwind control volume scheme developed by Moore⁹ and Moore and Moore.¹⁰ Brief details of this scheme are included for completeness.

2.1. Control volume upwinding

The control volumes allocated to each node are not fixed in advance but are allowed to adjust as the calculation proceeds so that they are aligned as closely as possible with the local flow direction. The alignment is computed by subdividing each cell vertex control volume into eight 1/8th sub-cells. The 1/8th cells are shown in Figure 2(a), which illustrates the case when the cell Peclet number (Pe) is less than 2 and each 1/8th is assigned to the vertex it touches. If there is a convection-dominated flow across the cell ($Pe > 2$), each 1/8th cell is assigned to that vertex which is closest to being in its downstream direction. The cases of flow parallel to a cell edge, flow diagonally across a cell face and flow diagonally across the cell are illustrated in Figures 2(d)–2(b) respectively. Once all the 1/8th cells have been allocated, the control volume for each node is simply the agglomeration of all the 1/8th cells assigned to it.

Over each control volume, the flow equations are discretized using second-order-accurate central differencing without the need to introduce upwinding into the differencing scheme. Where the flow is diffusion-dominated ($Pe < 2$), each node lies at the centre of its associated control volume and central differencing is well known to be stable.¹¹ To demonstrate stability in the case of convection-dominated flow, consider the convection of a scalar variable diagonally across a 2D cell by a constant velocity field (Figure 3):

$$\int \rho \phi \mathbf{u} \cdot d\mathbf{n} \sim \frac{1}{2} \rho (u \delta y + v \delta x) \phi_1 + \frac{1}{2} \rho (u \delta y - v \delta x) \phi_2 + \frac{1}{2} \rho (-u \delta y - v \delta x) \phi_3 + \frac{1}{2} \rho (-u \delta y + v \delta x) \phi_4. \quad (3)$$

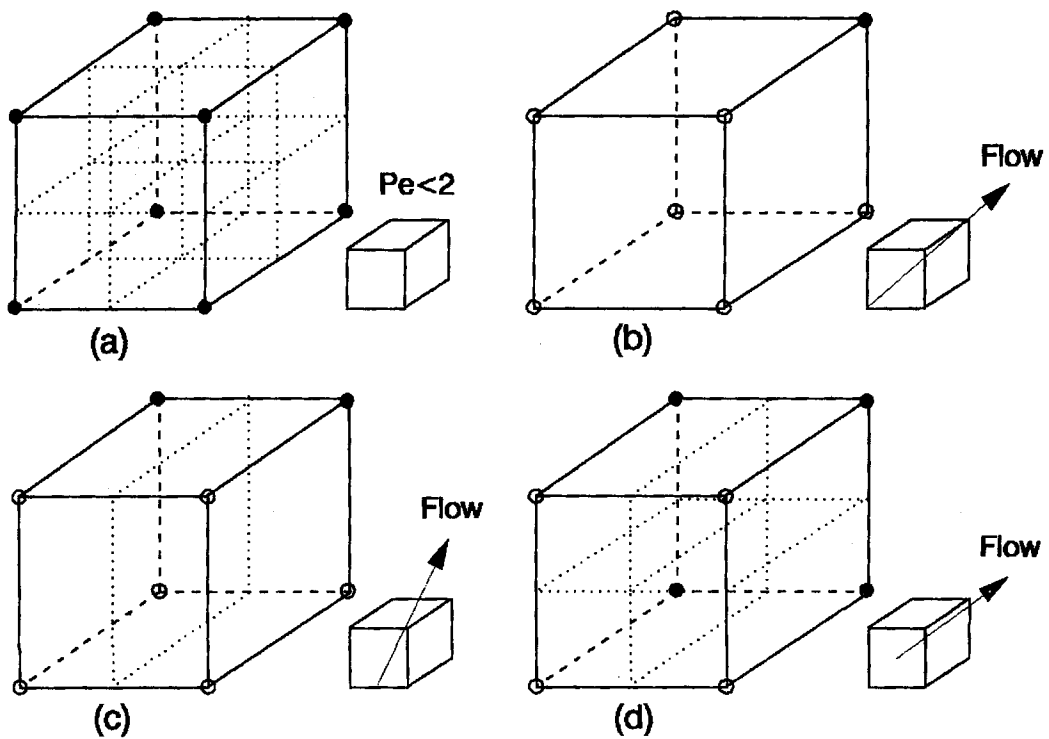


Figure 2. Upwind assignment of 1/8th control volumes: (a) cell $Pe < 2$; (b) flow diagonally across the cell; (c) flow diagonally from edge to edge; (d) flow from face to face




Figure 3. Upwind control volume for flow diagonally across a 2D cell

Since the flow is diagonally across the cell, $u/\delta x \sim v/\delta y > 0$ and the coefficient of ϕ_1 is the dominant term. Hence, the assignment of the control volume to node #1 ensures that the discretization is stable (and second-order-accurate). Diagonal dominance would not be achieved if the control volume were assigned to any of the other vertices of the cell. The stability of other configurations follows by a similar analysis. Fuller details of control volume upwinding are provided by Moore.⁹

2.2. Control volume upwinding at embedding interfaces

The rationale of the interface treatment is to ensure that none of the 1/8th sub-cells are assigned to the hanging nodes. If this is achieved, there is no conflict in setting values at the hanging nodes by linear interpolation, and flux conservation is maintained. Hence, any 1/8th sub-cells that would 'naturally' be assigned to a hanging node must be reassigned to one of its neighbouring 'regular' nodes. For stability the reassignment must preserve the upwind bias of the momentum control volumes. There are three generic upwinding configurations to be considered.

2.2.1. Flow into an interface. In this case, the hanging nodes do not receive any 1/8th contributions since all the fine side 1/8th sub-cells are assigned to nodes in the interior of the fine region and the coarse side 1/8ths are assigned to regular nodes on the interface (Figure 4(a)). Hence, no control volume reassignment is necessary. Flux conservation across the interface is guaranteed by the linear interpolation at the hanging nodes, which ensures that integrating between two adjacent regular nodes (as would be performed on the coarse side) produces the same value as integrating between the hanging node and each of its regular neighbours (as would be performed on the fine side).

2.2.2. Flow out of an interface. In this case, the fine side 1/8ths naturally assigned to the hanging nodes cannot be reassigned to the next node in the strict downstream direction since this lies within the interior of the coarse region. The 1/8ths are, therefore, reassigned to the adjacent regular node on the interface since this is in the next closest downstream direction (Figure 4(b)). Although the reassignment is not in the strict downstream sense, the control volumes for the interface nodes are identical to those that would arise if the interface were not present and the mesh were simply expanding/contracting at this location. This ensures that stability is preserved. Mesh stretching errors do not arise because each 1/8th cell is discretized separately. This is a feature of the basic control volume upwinding scheme.⁹ Hence, second-order accuracy is preserved across the interface. As in the previous configuration, flux conservation is guaranteed by the linear interpolation at the hanging nodes.

2.2.3. Flow along an interface. In this case, the 1/8th sub-cells assigned to the hanging nodes are reassigned to the next downstream (regular) node on the interface (Figure 4(c)). This reassignment is stable since it preserves the upwind bias of the momentum control volumes. The control volumes for the interface nodes are again identical to those that would arise if the interface were not present. For this configuration, the interface lies entirely within the interior of a control volume and hence fluxes are inherently conserved.

Flows diagonally across the interfaces are treated in an exactly analogous manner to the previous cases.

2.3. IAHR test case #1

The accuracy, stability and conservation properties of the interface treatment are demonstrated using the IAHR (International Association for Hydraulics Research) test case #1 defined for the

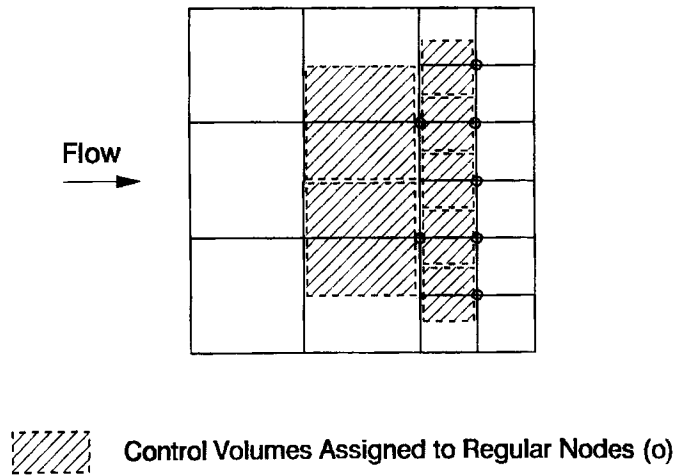


Figure 4(a). Control volume assignment for flow into an embedding interface

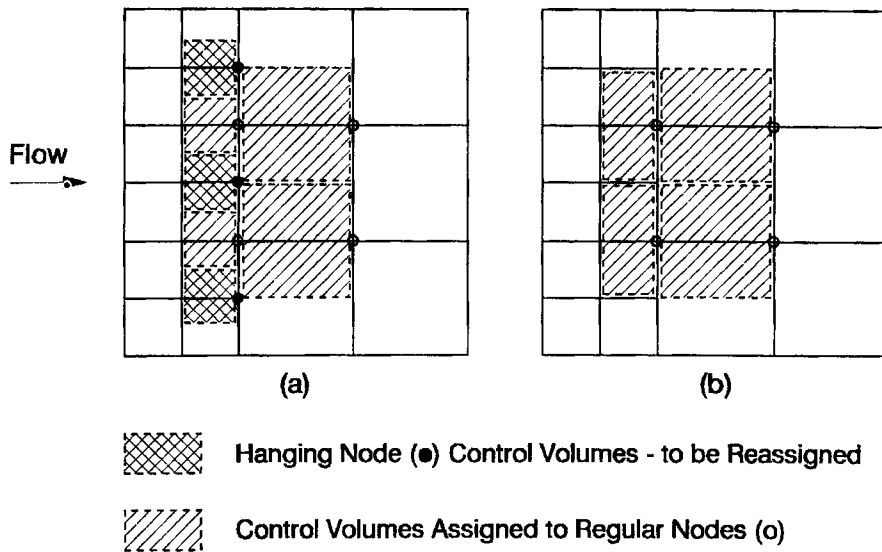


Figure 4(b). Reassignment of hanging-node control volumes for flow out of an embedding interface: (a) before; (b) after reassignment

Working Group on Refined Modelling of Flows, 14–15 May 1981, Berkeley Nuclear Lab., England. In the test case, a passive scalar is transported around a rectangular cavity by the prescribed velocity field:

$$u = 2y(1 - x^2), \quad -1 \leq x \leq 1, \quad 0 \leq y \leq 1,$$

$$v = -2x(1 - y^2), \quad -1 \leq x \leq 1, \quad 0 \leq y \leq 1.$$

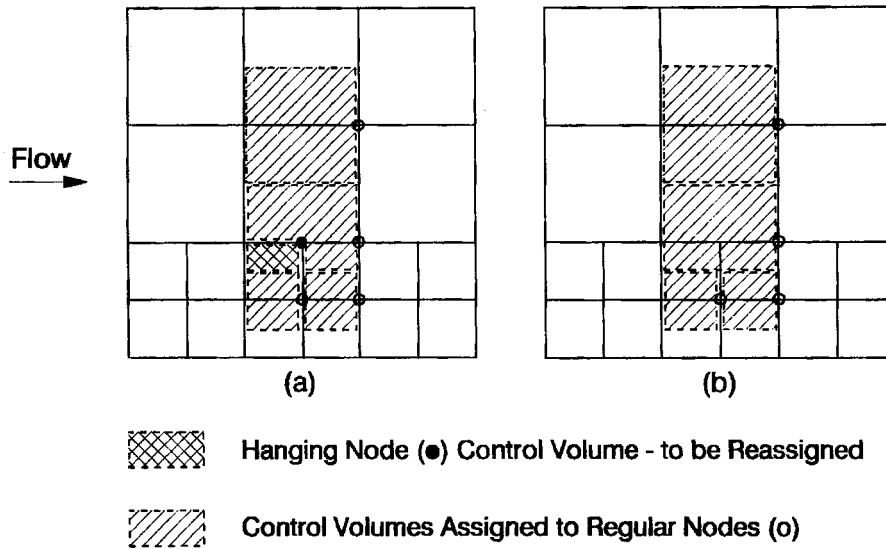


Figure 4(c). Reassignment of hanging-node control volumes for flow along an embedding interface: (a) before; (b) after reassignment

The inlet profile for the scalar φ is given by

$$\varphi = 1 + \tanh(10(2x + 1)), \quad -1 \leq x \leq 0, \quad y = 0.$$

The mesh used for the calculations is shown in Figure 5(a). The embedded fine mesh region is positioned so that interfaces exist where significant gradients of φ are transported into, out of and along the embedded region. Reynolds numbers (based on the cavity height) of $Re = \infty$ and $Re = 200$ are considered in order to test both the inviscid and viscous interface treatment. In both cases, the calculated contours of φ (Figures 5(b) and 5(c)) pass smoothly through the interfaces with no evidence of numerical diffusion or dispersion of the φ profile. Conservation is evidenced by the fact that there is no spurious creation or destruction of φ at the interfaces which would be manifested as a sudden convergence or divergence of the contours. In Figure 6, profiles of φ at the cavity exit are compared with results on a uniform fine structured mesh obtained using the method of Moore and Moore.¹⁰ There is near-identical agreement between the results in both the inviscid and viscous case; further demonstrating the second-order accuracy and conservation of the interface treatment. Exactly identical solutions are not achieved because of differences in the size (but not order) of the truncation errors where coarse cells in the embedded mesh correspond to a cluster of fine cells in the structured mesh.

2.4. Three-dimensional modelling

Three-dimensional embedding is achieved by stacking a series of 2D meshes on which the interface locations are topologically identical. The mesh co-ordinates need not be cloned identically in the stacking direction, allowing 3D variations in geometry to be modelled. The control volume upwinding remains fully 3D in nature, even at the embedding interfaces. The




Figure 5. IAHR test case #1: (a) embedded mesh; (b) $Re = \infty$, contours of transported scalar; (c) $Re = 200$, contours of transported scalar

reassignment of the interface control volumes in three dimensions is exactly analogous to the 2D treatment described previously.

Mesh stacking in the third direction represents the current status of the numerical model. Extension of the model to fully 3D mesh embedding is envisaged as a future development.




Figure 6. IAHR test case #1. Exit profiles of transported scalar (x, embedded mesh; +, structured mesh): (a) $Re = \infty$; (b) $Re = 200$

3. THE TURBULENCE MODEL

The Prandtl mixing length model developed by Moore and Moore¹² is used to compute the turbulent viscosities:

$$\mu_e = \mu_1 + \mu_t \tag{4}$$

$$\mu_t = \rho L^2 \left[\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \frac{\partial u_i}{\partial x_j} \right]^{1/2}, \tag{5}$$

where the mixing length L is taken to be the smaller of 0.08 times the boundary (or shear) layer width (δ) and 0.41 times the distance to the nearest wall (YWAL). Van Driest damping is applied in the near-wall regions.

To determine the edges, and hence thickness, of a boundary or shear layer a search function is defined as

$$VTES = (P_{max}^* - P^*) / (P_{max}^* - p_r), \tag{6}$$

where P^* is the rotary stagnation pressure and p_r is the reduced static pressure. The subscript max denotes the maximum value of P^* at the inlet to the calculation domain. The search function varies from a maximum value of 1 at a viscous wall to a minimum value of 0 in the freestream. Boundary and shear layers exist where the gradient of the search function is greater than the inverse of a characteristic passage width.¹² In the present cascade calculations, the blade pitch is used for the characteristic width.

The quantities required to evaluate μ_t fall into two categories:

3.1.1. Cellwise properties. These are properties such as the gradients of VTES or the du_i/dx_j terms which depend only on the values at the vertex of the cell under consideration. Since the values at the hanging nodes are set by linear interpolation, the cellwise properties can be efficiently evaluated independently of the presence of any hanging nodes.

3.1.2. Global properties. These are properties such as YWAL and δ which depend on integrated values across the boundary or shear layer. The global properties depend only on mesh distances and are built up by a marching procedure in which the current value of YWAL, say, in a given cell is compared with the sum of the distance across the cell and the values of YWAL in the adjacent cells. If any of the latter values are smaller than the current value, YWAL is updated. YWAL is easily initialized in all the near-wall cells and the marching is repeated until converged YWAL values are obtained in all cells. YWAL is stored for all mesh cells and needs only to be calculated at the outset of a calculation. The values of δ are set in a similar fashion during each iteration: here the marching is only in the direction of maximum shear and only continues until the edge of a boundary or shear layer is reached and is very efficient to implement.

At the embedding interfaces the marching must accommodate the transition in mesh size. When the marching crosses from the coarse side to a fine side, the coarse cell does not have a unique cell into which the marching may continue. Rather than considering a number of alternative marching routes, the marching proceeds solely into the fine side cell with the smallest value of YWAL: there is some arbitrariness in this choice but it has the advantage of removing bifurcations in the marching process. When marching from the fine side to the coarse side, no special action is required since the fine cell has a unique coarse side neighbour into which to march.

The turbulence model allows the computed values of μ_t to be scaled by a factor between 0 and 1 in order for predetermined transition locations to be specified.

4. SOLUTION ALGORITHM

The momentum equations are integrated over the upwind control volumes by taking a linear variation in the flow properties between the mesh nodes. The discrete momentum equations take the form

$$\sum_{\alpha=1}^{27} AU_{\alpha} * \delta u_{\alpha} = \mathbf{BU}, \quad (7)$$

where \mathbf{BU} is the momentum equation residual and AU is the 27-point coefficient matrix arising from the convective and diffusive transport terms. The momentum equations are solved using the Gauss-Seidel algorithm, which is adequate given the hyperbolic nature of these equations.

In structured regions of the mesh, the difference stencil relates each mesh node to its 26 immediate neighbours. In embedded regions, the difference stencil relates each interface node solely to its 26 neighbouring regular nodes (Figure 7): the hanging nodes are implicitly replaced by their linear interpolation rules and do not form part of the stencil. A special treatment is also needed where the reassignment of the hanging-node control volumes introduces coefficients which




Figure 7. Two-dimensional difference stencil in the vicinity of an embedding interface

lie outside the 27-point stencil. Prior to solving the momentum equations such coefficients temporarily reside in difference stencils about the hanging nodes. For each sweep through the calculation domain in the Gauss–Seidel procedure, the additional coefficient terms are updated and allocated to the appropriate regular node. This approach allows the difference stencil for AU to remain of fixed size.

The non-linear system of equations is solved using the SIMPLER algorithm¹¹ in which the discrete momentum equations are substituted into the continuity equation in order to form a pressure correction equation which is used to update the pressure and promote continuity satisfaction. The continuity equation is integrated over a control volume with nodes at its vertices. The resulting pressure correction equation is solved using the defect correction approach:

$$\sum_{\alpha=1}^7 AP_{\alpha} * \Delta p_{\alpha} = BP - \sum_{\alpha=1}^8 CRUC_{\alpha} * \rho * \delta \mathbf{u}_{\alpha}, \quad (8a)$$

$$\delta p = \delta p + \Delta p, \quad (8b)$$

$$\delta \mathbf{u} = \sum_{\alpha=1}^8 CGP_{\alpha} * \delta p_{\alpha} / CAU, \quad (8c)$$

where BP is the continuity equation residual, AP is the coefficient matrix for the defect equation, CRUC is the discrete continuity operator, CGP is the discrete pressure gradient operator and CAU is the centre point of the momentum equations.

In the rationale of the defect correction scheme, the matrix AP is an approximation of the full pressure correction matrix which is better posed and hence more assured of convergence. In the present scheme, a seven-point approximation is used in order to make the interface treatment less cumbersome. The above formulation avoids the need to form the full pressure correction matrix which would require more than 27 points in the difference stencil at the embedding interfaces. For cells on the coarse side of an interface, the stencil is limited to seven points by using dummy cells on the fine side which are an agglomeration of the fine side continuity cells. Values in the dummy cells are set as the average of their associated fine cells. On the fine side of an interface there is no ambiguity in the difference stencil. The pressure correction equation is solved using the conjugate gradient algorithm with a Jacobi preconditioner.

5. RESULTS

The embedded mesh model is validated in three dimensions using the cascade of turbine vanes tested experimentally by Hodson and Dominy.^{13,14} The cascade is rectilinear and has a symmetric annulus flare of 3.1° between the leading and trailing edges. The calculations are performed over the lower half of the cascade with a symmetry boundary condition at mid-height. Results are presented for the design condition and the negative (-20.3°) and positive (8.6°) incidence conditions. The present results are compared with the measurements and with results obtained using the structured mesh method of Moore and Moore.¹⁰

The embedded mesh used for the calculations is illustrated in Figure 8. The embedding is based on geometrical considerations in order to provide resolution of the blade boundary layers and wake region. The structured mesh calculations were performed on the equivalent fine mesh which is generated by extending all the mesh lines of the embedded mesh to the domain boundaries. The extended mesh lines were flared to avoid instabilities in the freestream regions. The embedded mesh contains 26 640 nodes compared to 51 504 for the structured mesh—a saving of 48%. A previous numerical simulation of the Hodson and Dominy cascade by Smith¹⁵ found that the numerical modelling was improved by the prescription of laminar and turbulent regions which aimed to mimic the experimentally observed transition locations. The prescription favoured by Smith is used in the current investigations. Briefly, on the suction surface abrupt transitions are specified at 82, 0 and 60% of axial chord for the design, positive and negative incidence cases, respectively. On the pressure surface, abrupt transition at 95% axial chord is specified for the




Figure 8. Embedded mesh for the Hodson and Dominy cascade

design and positive incidence cases, and abrupt transition at 25% axial chord is specified for the negative incidence case.

5.1. Blade surface isentropic Mach numbers

The mid-height surface Mach numbers for the three flow conditions are compared in Figure 9. The embedded and structured mesh solutions show excellent agreement with each other. Good agreement with the measurements is obtained in the design and positive incidence cases. At negative incidence, Smith¹⁵ has shown that the length of the pressure side leading edge separation bubble, and hence the surface Mach number distribution, is influenced by the specified transition locations. The present calculations achieve a similar level of agreement with the test data to that achieved by Smith¹⁵ using the same turbulence model prescription. Further tuning of the turbulence prescription was not attempted since the aim of the calculations is to demonstrate the validity of the embedded mesh model. This validity is exhibited by the excellent agreement between the embedded and structured mesh results.

5.2. Exit whirl angle distributions

Radial profiles of whirl angle downstream of the cascade trailing edge for the three flow conditions are compared in Figure 10. The embedded and structured mesh calculations show good agreement with each other. There is good agreement with the measured mid-height whirl angles and also with the levels of peak underturning. The measured inward migration of the peak underturning as the incidence reduces from positive to design to negative is clearly modelled. However, a second underturning peak at approximately 10% height, which is not apparent in the measurements, is exhibited by both calculations in the positive incidence case, and by the embedded calculation at design incidence. The second peak is as a result of a small loss core that is calculated at 10% height (see Figures 12(a) and 12(b) and the discussion in Section 5.4). There is some evidence of a loss core in the experiment but this is much closer to the endwall in both flow conditions. It is not clear whether the calculated loss cores, and hence the second underturning peaks, are a result of poor modelling of the rate of mixing in the wake region or poor modelling of the migration of the loss cores under the influence of the secondary flows. (Cleak¹⁶ has shown that the state of the endwall boundary layer can have an effect on loss core migration.) In either case, shortcomings in the prescription of laminar and turbulent regions are felt to account for the discrepancies with the measurements. Brief experimentation with the turbulence model prescription verified that these effects can be significant. Other factors such as shortcomings in the algebraic turbulence model and the effects of truncation errors due to the reasonably large radial spacing that has been used may also contribute to the discrepancies with the measurements.

Examination of the differences between the embedded and structured mesh results at design incidence determined that at 10% chord upstream of the trailing edge the whirl angle profiles are in excellent agreement; however, passing through the trailing edge region the second peak develops in the embedded solution but not in the structured solution. Further examination revealed that adjacent to the trailing edge, the structured solution has spuriously high viscosities which are stimulated by the high aspect ratio cells in the region where the boundary layer grid lines extend across the passage into the freestream. Similar levels of viscosity do not arise in the embedded solution because the mesh derefines away from the trailing edge producing cells with much smaller aspect ratios. The higher viscosities in the structured solution clearly produce greater mixing and hence restrain the development of the second underturning peak. This effect does not produce similar differences in the positive incidence case because the deviation angle is




Figure 9. Three-dimensional turbine cascade. Mid-height isentropic surface Mach numbers (\times , embedded mesh; $+$, structured mesh; Δ , measurements): (a) design incidence; (b) positive incidence; (c) negative incidence

different and the loss core responsible for the second underturning peak is more remote from the high-viscosity region.

5.3. Exit loss profiles

Radial profiles of loss downstream of the cascade trailing edge for the three flow conditions are compared in Figure 11. Excellent agreement is achieved between the embedded and structured mesh solutions in all cases. The agreement with the measurements is also good: some discrepancies (particularly in the positive incidence case) are clearly due to the relatively large radial spacing that has been used. Again, the excellent agreement between the two calculations is assumed to validate the embedded model, and improvements in the radial resolution, which would require prohibitive increases in CPU time and storage, were not attempted.

5.4. Exit loss contours

Contours of loss in a downstream axial plane for the three flow conditions are compared in Figures 12(a)–12(c). For the design case, the embedded solution shows two distinct loss cores at 20 and 25% blade height, with similar levels of loss in each core. This is in good agreement with the measurements, even though there is more loss measured in the outer core. In contrast, the structured solution shows only a single loss core in this region. It is probable that the higher levels of viscosity near the trailing edge in the structured calculation have caused the loss cores to mix and coalesce. Both calculations show the additional loss core at about 10% height as discussed in Section 5.2: again the loss core in the structured mesh is more mixed as a result of the higher viscosities. The evidence of a loss core much closer to the endwall in the measurements is also apparent.

At the positive incidence condition, both calculations show two distinct loss cores and also capture the outward migration of the loss cores relative to the design condition. The loss cores from the structured solution are slightly more mixed out but have not coalesced as in the design case. At the negative incidence condition, both calculations show only a single loss core, as in the measurements, and model the inward migration relative to the design condition.

6. SOLUTION ADAPTION

To demonstrate solution adaption using the present model, a 2D mid-height section of the Hodson and Dominy^{13, 14} cascade is modelled at the off-design conditions. Solution adaption in 3D is currently precluded by the requirement that meshes are stacked in the third direction. The adaption is based on first-order differences of velocity: cells where du is greater than the r.m.s. value of du are subdivided in the corresponding direction provided the cell size is greater than a prespecified threshold. To improve mesh 'regularity' and quality the 'halo' cells adjacent to the adapted cells are also refined.

The calculations are initialized using a 2D section of the mesh used previously (Figure 8) and two levels of adaption are used. The sequence of meshes for the negative incidence condition is shown in Figure 13. The adaption produces improved resolution of the leading edge flow and of the pressure side separation region.

6.1. Blade surface isentropic Mach numbers

Isentropic surface Mach numbers for both off-design conditions are shown in Figure 14. In the positive incidence case, the mesh adaption has only a minor influence on the solution—indicating

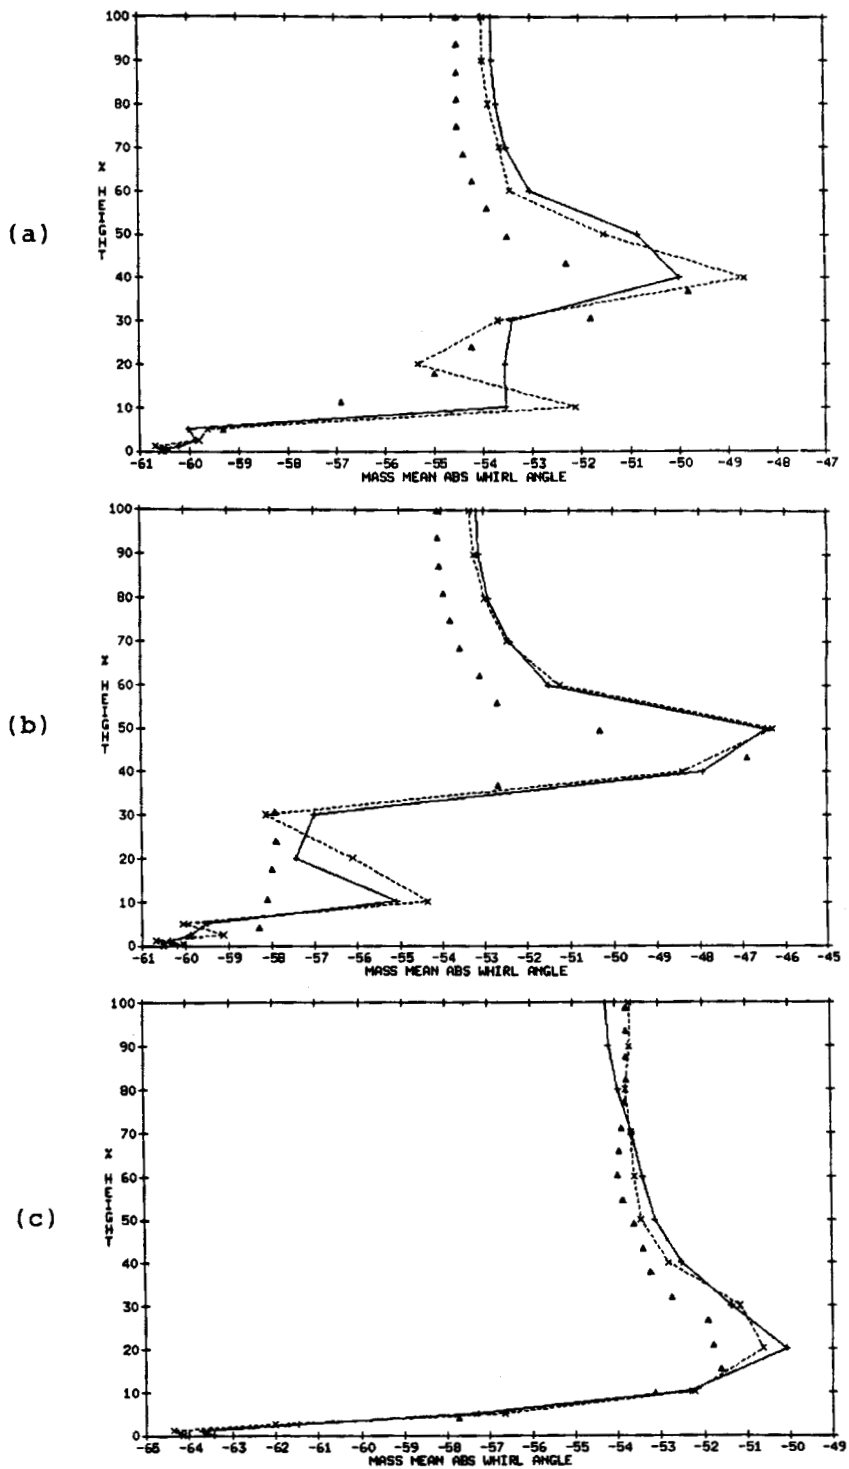


Figure 10. Three-dimensional turbine cascade. Exit whirl angle profile (x, embedded mesh; +, structured mesh; Δ, measurements): (a) design incidence, 143% chord; (b) positive incidence, 136% chord; (c) negative incidence, 136% chord

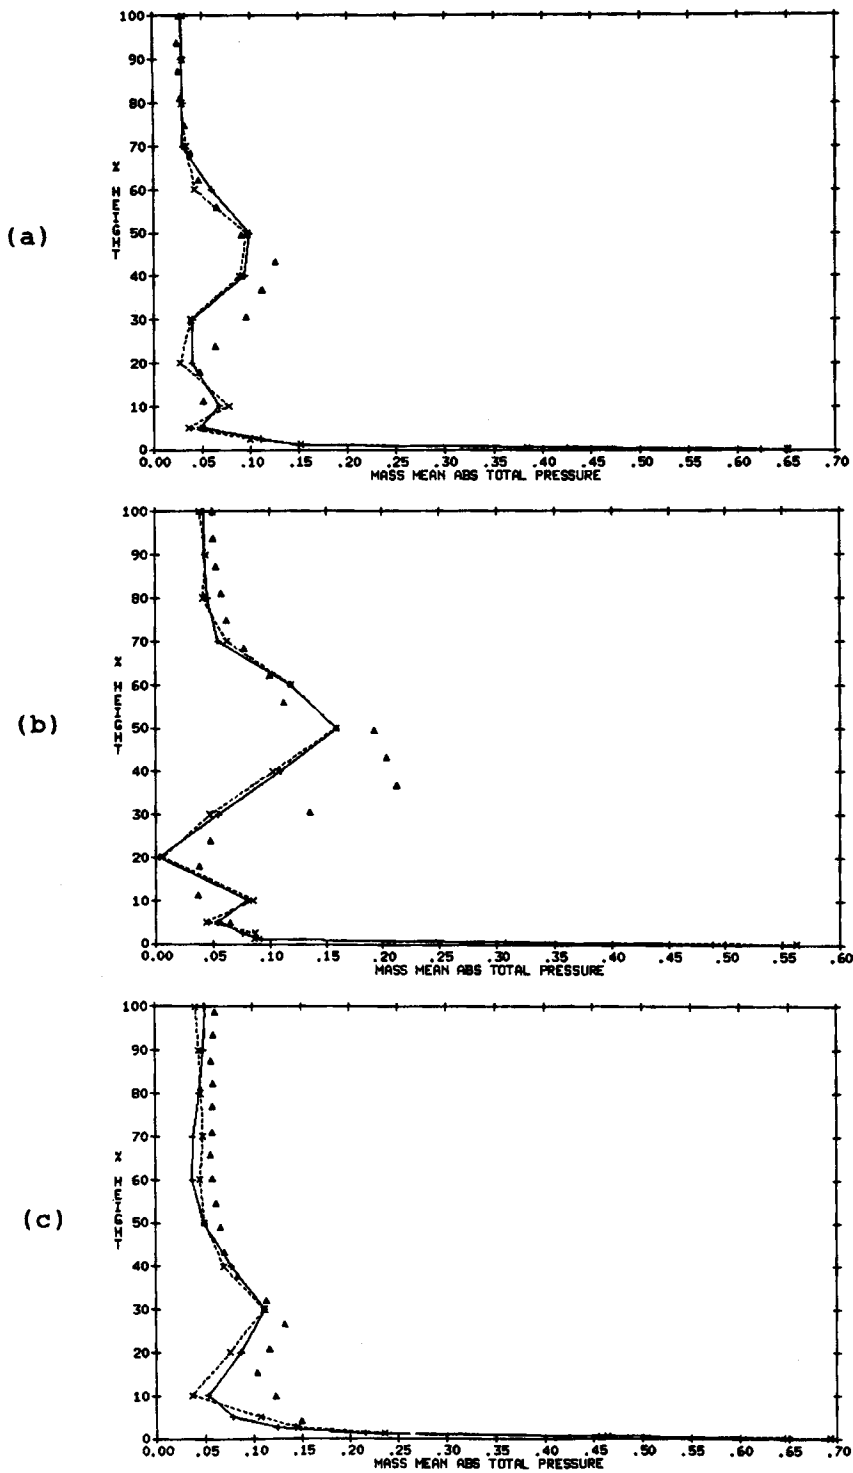


Figure 11. Three-dimensional turbine cascade. Exit loss profile (x, embedded mesh; +, structured mesh; Δ, measurements): (a) design incidence, 142% chord; (b) positive incidence, 136% chord; (c) negative incidence, 136% chord




Figure 12(a). Three-dimensional turbine cascade. Design incidence, contours of loss at 136% axial chord; (a) measured; (b) embedded mesh; (c) structured mesh





Figure 12(b). Three-dimensional turbine cascade. Positive incidence, contours of loss at 136% axial chord: (a) measured; (b) embedded mesh; (c) structured mesh



(a)

(b)

(c)

Figure 12(c). Three-dimensional turbine cascade. Negative incidence, contours of loss at 136% axial chord: (a) measured; (b) embedded mesh; (c) structured mesh




Figure 13. Two-dimensional turbine cascade. Negative incidence, mesh adaption: (a) non-adapted mesh; (b) first adapted mesh; (c) second adapted mesh

that the original mesh was adequate for resolving the surface Mach numbers. Differences in the loading between this calculation and the mid-height results from the 3D calculation (Figure 9(b)) are due to the influence of the annulus flare on the streamtube height which is neglected in the 2D calculations. In the negative incidence case, there is a marked improvement in the agreement with the experiment as the mesh is adapted. The adaption not only improves the Mach numbers on the




Figure 14. Two-dimensional turbine cascade. Surface isentropic Mach numbers (+, non-adapted; X, first adaption; Y, second adaption): (a) positive incidence; (b) negative incidence

pressure surface, but the suction surface Mach numbers also improve due to the extra blockage produced by the increased size of the separation bubble.

6.2. Velocity vectors

Velocity vectors on the sequence of adapted meshes for the negative incidence condition (Figure 15) show a significant increase in the extent of the recirculation zone as the mesh is adapted. Experimentation with the minimum cell-size threshold indicated that the extent of the recirculation is strongly influenced by the resolution of the pressure side leading edge spike region.




Figure 15. Two-dimensional turbine cascade. Negative incidence, velocity vectors: (a) non-adapted; (b) first adaption; (c) second adaption

On the initial mesh, there is a very strong spike (Figure 14(b)) leading to large streamline curvatures which delay the onset of the flow separation, giving only a small recirculation bubble. As the mesh is adapted, the resolution and near-wall spacing in the spike region improves and a better model of the boundary layer growth is achieved. This influences the calculated displacement thicknesses which effectively modify the leading edge shape—reducing the magnitude of the spike

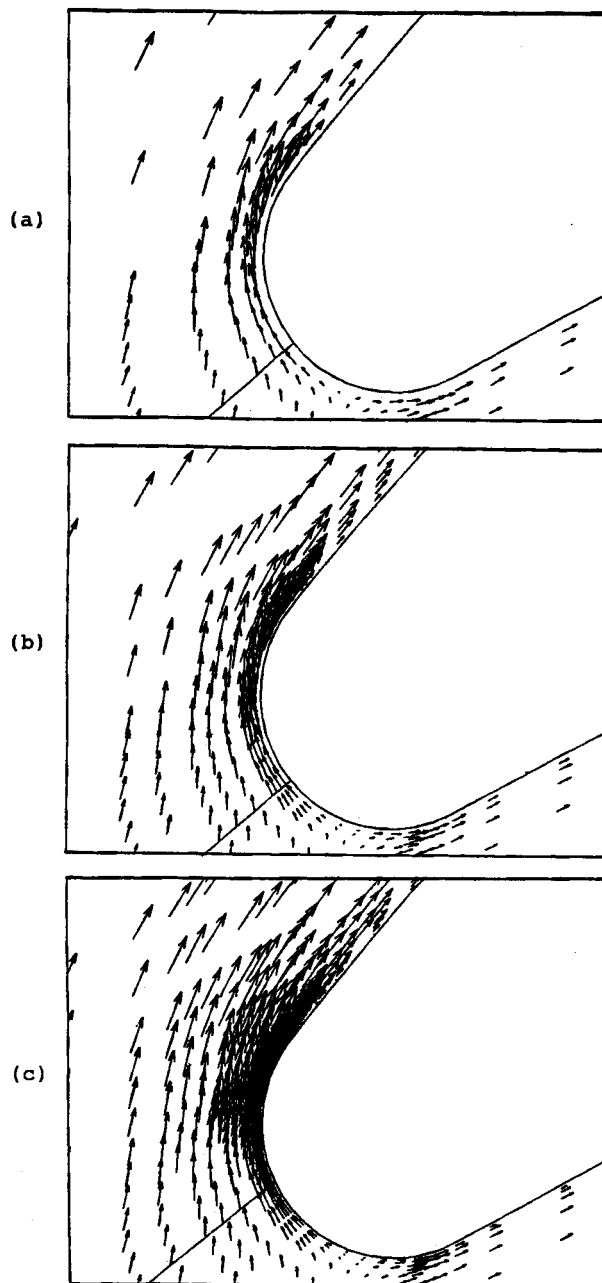


Figure 16. Two-dimensional turbine cascade. Positive incidence, velocity vectors: (a) non-adapted; (b) first adaption; (c) second adaption

and the corresponding streamline curvatures. Hence, the separation point moves forward and the size of the separation bubble increases. The reattachment points computed on the initial, first and second adapted meshes are at 38, 52 and 60% of the axial chord, respectively—compared to a measured value of 45%. However, the value of 60% on the final adapted mesh agrees very well

with the 58% reattachment point computed by Liu and Jameson,¹⁷ who used a 209×65 mesh. The discrepancies with the experiment may be influenced by streamtube height effects, but the use of prespecified transition locations is also likely to have a strong effect.

Velocity vectors on the sequence of adapted meshes for the positive incidence case are shown in Figure 16. There is an improved resolution of the boundary layer regions as the mesh is adapted. The second adaption captures a tiny suction side separation immediately downstream of the leading edge which is not apparent in either of the previous meshes. This separation was clearly observed in the experiments.¹⁴

7. CONCLUSIONS

A mesh embedding algorithm for the 3D Navier–Stokes equations using an algebraic model of turbulence has been reported. A stable, second-order-accurate and conservative treatment for the embedding interfaces has been developed using upwind control volumes. The interface treatment retains a simple interpolative scheme for the hanging nodes and avoids complex integration rules. Three-dimensional modelling is achieved by stacking topologically similar meshes in the third direction.

The embedded mesh solver has been used to compute the 3D flow through a cascade of turbine vanes at design and off-design conditions. Comparisons with results on an ‘equivalent’ structured mesh, obtained using the method of Moore and Moore,¹⁰ demonstrate that the embedded solutions achieve the same solution quality as the structured solutions with a saving of 48% in the number of mesh nodes. Minor differences between the embedded and structured mesh solutions were found to be mostly attributable to the viscosity modelling in the trailing edge region. High aspect ratio cells near the trailing edge which stimulate spuriously high viscosities are eliminated by the mesh embedding.

The solution adaption demonstrates the clear improvements in solution quality and physical modelling that can be achieved. The improvements in the negative incidence case are particularly marked, with the adapted solution showing significantly better agreement with the measurements than the non-adapted solution. The positive incidence case demonstrates the ability of adaption to capture, in a single calculation, flow phenomena (in this case a leading edge separation bubble) that would otherwise require manual remeshing and recalculation.

REFERENCES

1. W. Shyy, ‘Computation of complex fluid flows using an adaptive grid method’, *Int. j. numer. methods fluids*, **8**, 475–489 (1988).
2. I. R. Hawkins and J. R. Kightly, ‘Grid adaption for turbulent flow problems’, *Proc. 6th Int. Conf. on Numerical Methods in Laminar and Turbulent Flow*, Pineridge Press, Swansea, 1989, pp. 1651–1662.
3. J. F. Thompson, ‘A survey of dynamically adaptive grids in the numerical solution of partial differential equations’, *Appl. Numer. Math.*, **1**, 3–28 (1985).
4. Y. G. Kallinderis and J. R. Baron, ‘Adaption methods for a new Navier–Stokes algorithm’, *AIAA J.*, **27**, 37–43 (1989).
5. R. L. Davis and J. F. Dannenhoffer, ‘Adaptive grid embedding Navier–Stokes technique for cascade flows’, *AIAA-89-0204, AIAA 27th Aerospace Sciences Meeting*, Reno, Nevada, 1989.
6. S. R. Allmaras and J. R. Baron, ‘Embedded mesh solutions of the 2-D Euler equations: evaluation of interface formulations’, *AIAA-86-0509, AIAA 24th Aerospace Sciences Meeting*, Reno, Nevada, 1986.
7. J. M. McDill, J. A. Goldak, A. S. Oddy and M. J. Bibby, ‘Isoparametric quadrilaterals and hexahedrons for mesh-grading algorithms’, *Commun. Appl. Numer. Methods*, **3**, 155–163 (1987).
8. Y. Kallinderis, ‘Numerical treatment of grid interfaces for viscous flows’, *J. Comput. Phys.*, **92**, 129–144 (1992).
9. J. G. Moore, ‘Calculation of 3-D flow without numerical mixing’, *AGARD Lecture Series No. 140 on ‘3-D Computational Techniques Applied to Internal Flows in Propulsion Systems’*, 1985.
10. J. G. Moore and J. Moore, ‘Calculation of viscous flow using the Moore elliptic flow program’, *College of Engineering Report JM/85-1*, Virginia Polytechnic Institute and State University, 1985.

11. S. V. Patankar, *Numerical Heat Transfer and Fluid Flow*, McGraw-Hill, New York, 1980.
12. J. Moore and J. G. Moore, 'Calculation of three-dimensional, viscous flow and wake development in a centrifugal impeller', *ASME CP Performance Prediction of Centrifugal Pumps and Compressors*, 1980, pp. 61–67.
13. H. P. Hodson and R. G. Dominy, 'Three-dimensional flow in a low pressure turbine cascade at its design condition', *ASME Paper 86-GT-106*, 1986.
14. H. P. Hodson and R. G. Dominy, 'The off-design performance of a low pressure turbine cascade', *ASME Paper 86-GT-188*, 1986.
15. A. D. Smith, 'A comparison between the experimental investigation and the Moore elliptic flow program predictions of the 3D flow within a rectilinear cascade of 'LA' profile LP turbine blades', *Rolls-Royce Report*, 1989.
16. J. G. E. Cleak, 'Validation of viscous, three-dimensional flow calculations in an axial turbine cascade', *Ph.D. Thesis*, Durham University, 1989.
17. F. Liu and A. Jameson, 'Multigrid Navier–Stokes calculations for three-dimensional cascades', *AIAA-92-0190*, *AIAA 30th Aerospace Sciences Meeting and Exhibit*, 1992.