

AIAA-99-0917
A High-Accuracy Solution-Adaptive
Unstructured Macro-Cell Algorithm for CFD

Robert E. Childs,* E. Burke Treidler,*
John A. Ekaterinaris,* and Patrick H. Reisenthel*
Nielsen Engineering & Research
526 Clyde Ave., Mountain View, CA 94043-2212
childs@nearinc.com

Introduction

This paper describes an unstructured macro-cell (UMC) algorithm for solving partial differential equations, and its application to computational fluid dynamics. A macro-cell is an intermediate-sized domain, consisting of roughly 10 or more grid points in each spatial direction in a structured array. The governing equations are solved on this structured grid within a macro-cell. A complete computational domain is composed of perhaps 10s to 1000s of macro-cells which are arranged in an unstructured manner. Local grid refinement is achieved by splitting a macro-cell, recursively as needed, in a manner similar to that of adaptive mesh refinement (AMR) schemes. The UMC algorithm has features of structured and unstructured grid methods, and it is intended to retain the benefits and avoid the drawbacks of purely structured and unstructured methods. Because the equations are solved on a structured array of grid points, high-accuracy finite-volume and implicit methods can be used. The use of unstructured macro-cells enables the use of several geometry- and solution-adaptive strategies. Thus, a UMC method offers many significant benefits for CFD. The principal difficulty which must be overcome is the need for methods of achieving high accuracy at macro-cell boundaries. This paper addresses the general concepts of the macro-cell solver with emphasis on boundary issues. Companion papers by Treidler et al (1999) and Reisenthel and Childs (1999) describe aspects of the computational efficiency and the solution-adaptation algorithm, respectively.

* Member AIAA

Copyright © 1999 by Nielsen Engineering & Research,
Published by the American Institute of Aeronautics and
Astronautics, Inc., with permission.

Motivation and Background

The goal of this work is to solve the governing equations of a physical process with arbitrarily high accuracy, despite the lack of knowledge of details of that process prior to the calculation. The specific problem of interest is dynamic stall on rotorcraft, in which a leading-edge separation bubble forms, forms a vortex, which then detaches and convects over the aft portion of the rotor blade, producing damaging aerodynamic loads. For incipient or "light" dynamic stall, the problem of greatest concern, the location, phase, and even the formation of a stall vortex cannot be anticipated. Efforts to demonstrate grid-independent calculations of dynamic stall in two dimensions have been limited to about 1/10th of flight Reynolds number. Solution-adaptive methods are needed for these calculations to resolve the unknown, transient location of the stall vortex at reasonable computational cost. High-accuracy methods are also expected to provide better global accuracy at lower computational cost in flows with (Carpenter and Casper,1997) and without shocks.

Most commonly used CFD methods are block-structured-grid methods or unstructured-grid methods. The former are generally well-suited to high-accuracy solvers but not to solution-adaptivity. The latter enable solution-adaptivity but are generally limited to first- or second-order accuracy. Exceptions to these generalizations exist. High-accuracy unstructured methods have been developed by Barth (1993) and others. While block-structured methods are nominally well-suited to high-accuracy methods, the common practice of using overset grids causes the effective computational domain inside individual blocks to be irregular, which complicates the implementation

of high-accuracy boundary methods.

The UMC method is, in some respects, closely related to the spectral-element method pioneered by Patera and further developed by Hendersen and Karniadakis (1995) and Sherwin and Karniadakis (1996). Spectral element methods employ some high-order basis functions, such as Jacobi polynomials, within each element, and there are $O(10)$ interior nodes in each spatial direction. The so-called *hp*-adaptive finite element methods are also related; they permit variable order of accuracy and on-demand cell refinement. Within these methods, there are two general approaches to solution-adaptive refinement. Some methods use triangular and tetrahedral elements, and local refinement is achieved in a manner similar to conventional unstructured methods. In particular, each face of a spectral element is mated to exactly one neighbor. The other approach is to use the so-called “mortar” methods to join dissimilar numbers of elements (Bernardi and Maday, 1998). A common implementation is to permit a single spectral-element face to adjoin two neighboring elements, in each spatial direction, which enables 2-to-1 grid size refinement across the boundary between two elements. These mortar methods are often used in conjunction with hexahedral (brick) elements.

Macro-Cell Concepts

Macro-cells and Micro-cells - The unstructured macro-cell method involves unstructured agglomerations of macro-cells, each containing a structured inner grid. In principal, macro-cells can be of four basic shapes, hexahedra (bricks), prisms, pyramids, or tetrahedra, and hybrid grids combining different types of macro-cells can be used. However, the present work concerns only brick macro-cells. Local grid refinement is achieved by permitting 2-to-1 refinement in any or all of the grid coordinate directions of any macro-cell. There is also a limit of not more than 2-to-1 refinement between adjacent macro-cells, which is arbitrary but reasonable. This benefits accuracy and limits code complexity, at a possible small penalty in computational efficiency. This approach is similar to that used in AMR methods. For example,




Figure 1. Example of AMR-like grid refinement for macro-cell grid.




Figure 2. Macro-cell boundaries for airfoil grid.

Figure 1 shows a grid of several macro-cells with local grid refinement. This approach to grid refinement produces “sympathetic” refinement in neighboring macro-cells, as the transition from a single small macro-cell to larger macro-cells must occur in 2-to-1 jumps. Achieving good accuracy and conservation across a macro-cell junction with 2-to-1 refinement is one of the critical challenges of this work. Although a uniform Cartesian grid is used in Fig. 1, arbitrary curvilinear macro-cell grids may be used. Figure 2 gives the boundaries of macro-cells for an airfoil grid that is refined in a region suitable to resolve dynamic stall. Within each macro-cell is a somewhat conventional finite-volume grid. For linguistic aesthetics, it is referred to as the *micro*-cell grid.

The macro-cell domain decomposition can be used to accommodate difficulties associated with complex configurations; thus, the grid within each macro-cell can be reasonably smooth which improves accuracy. Most significantly, metric discontinuities across the junction between

neighboring macro-cells have no direct effect on accuracy. Sharp features of a configuration are accurately resolved by placing a macro-cell boundary on that discontinuity, rather than wrapping a macro-cell around it.

Points, Lines, Planes, and Volumes - In developing a high-accuracy method, one needs a clear understanding of what the solution “at a grid point” represents. This discussion is presented in three-dimensional terms, and the reduction to two dimensions is obvious. The macro-cell interior points represent volume integrals of the solution, which are identical to finite volume notation. At grid points on the boundaries of a macro-cell, the solution represents planar integrals. Thus, a grid point which is common to neighboring macro-cells represents the same planar integral (excluding 2-to-1 refinement). This is similar to finite difference concepts, but not, strictly, collocation methods. All fluxes are planar integrals. Spatial representation of the solution is factored, that is the variations in the ξ , η , and ζ coordinate directions are separate and independent. Each one-dimensional factor is represented in terms of line integrals and points. All spatial operators (integration, interpolation, etc.) are developed in one dimension, and they deal with line integrals and points. The solution within the volumes is obtained by temporal integration of the integral form of the governing equations, while planar boundary values of the solution are computed with boundary conditions.

Plotting Anomalies - Plotting of results has been performed with a commercial package which interprets the solution “at a grid point” in a finite-difference, rather than integral, sense. It also uses linear interpolation for contour lines, etc. Both assumptions are in conflict with the underlying representation of the computed solution. This can give the appearance of gaps between macro-cells and discontinuities between the solutions in adjacent macro-cells, for example.


Interior Methodologies

Away from MC boundaries, the governing equations are solved with relatively conventional, albeit high-order methods. The grid metrics are

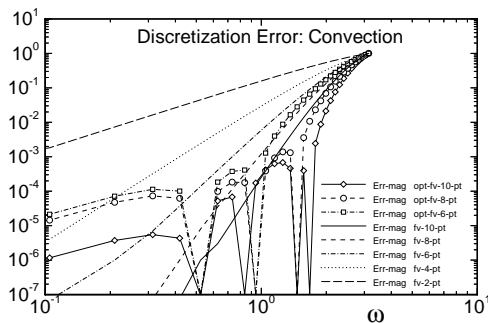
computed with fourth-order finite differences, and interpreted in terms of cell volumes and cell-face areas in the usual manner. Spatial operations are factored, so that independent and similar operations are performed in the ξ , η , and ζ coordinate directions. Inviscid cell face fluxes are computed in two steps: first the cell-centered contravariant fluxes are computed based on the solution within individual cells; then central interpolation is used to estimate the flux averaged over cell face, based on cell-volume fluxes at several points on either side of that face. The UMC solver currently includes operators for fourth- and eighth-order interpolation; these are centered operators which use data from four and eight grid points, respectively. Viscous fluxes are computed with sixth-order discretization, which have roughly the same bandwidth as the eighth-order Euler terms. A nonlinear artificial dissipation related to the methods of Jameson is used. It consists of second- and eighth-order dissipation terms (actually, first- and seventh-order dissipative fluxes). The relative weighting on each term depends on the smoothness of the solution, which is measured via the fourth-derivative of the pressure. It should be noted that the bandwidth of eighth-order interior dissipation is significantly less than the bandwidth of the eighth-order Euler terms (see Figs. 3 and 7). The long-term goal is to implement optimized dissipation schemes which minimize this discrepancy. Details of the underlying high-accuracy finite-volume method were described by Treidler and Childs (1997). Examples of this algorithm used in DNS simulations of supersonic boundary layers are given in Childs and Reisenhel (1995, 1997). Grid motion terms are included in the usual fashion, which adds to the conventional fluxes the contributions due to grid motion.

Algorithm Bandwidth - The accuracy of these schemes can be expressed in terms of various metrics, such as polynomial order. However, order can be misleading near boundaries, and the preferred method here is in terms of bandwidth. Bandwidth is taken to mean the range of Fourier wave numbers, in computational space, resolved to some small relative error. The error threshold used here is 10^{-3} , which is arbitrary but probably reasonable for most CFD work. Wave numbers used to define bandwidth are normalized so that

four grid points per wavelength corresponds to $\omega = \pi/2$. The benefit of measuring accuracy in terms of bandwidth is that it provides a direct correlation to the resolved length scales of the solution. Methods of similar bandwidth resolve similar length scales. The same is not true for the order of the method. The bandwidth of several interior interpolation schemes (for convective processes) is given in Fig. 3. The error shown in Fig. 3(a) is a lagging wave speed at higher wave numbers. Polynomial schemes up to tenth order and optimized schemes are shown. At present, an eighth order polynomial form is used for the interior convective terms. The difference between polynomial and optimized schemes is seen when the error is plotted in logarithmic coordinates. The optimized schemes have a significantly wider bandwidth, to the error tolerance specified, but they have poorer asymptotic convergence.



3(a) Error in convection, linear scale.



3(b) Error in convection, logarithmic scale.

Figure 3. Error versus wave number for several finite-volume (FV) Taylor series (lines alone) and optimized formulae (with symbols).

The principal benefit of viewing accuracy in terms of bandwidth is the direct correlation between bandwidth and resolved length scales. A major step toward achieving seamless accuracy across macro-cell boundaries is to insure that the bandwidths (i.e., resolved length scales) of all near-boundary operators are similar.

Boundary Methods

A variety of techniques are needed to achieve high accuracy up to and across macro-cell boundaries. This section attempts to present these issues in a concise manner. The following are key issues:

- Intermacro-cell boundary conditions are enforced through the planar-averaged solution and the fluxes on the macro-cell boundaries. That is, C_0 continuity is enforced on the solution and fluxes. Solution-gradient data is not directly shared, but it can be implicitly shared via viscous fluxes.
- Spatial discretization is factored, and boundary-normal and boundary-tangent issues are separable.
- Specialized boundary-normal treatment is required at all boundaries.
- Specialized boundary-tangent methods are required where sharing boundary values and fluxes between neighboring macro-cells with 2-to-1 refinement.
- Micro-cell grid refinement in computational space (ξ , η , ζ) is used at the first interior volume adjacent to all boundaries.

The use of only C_0 continuity in the solution at macro-cell boundaries simplifies the programming, but it is now seen as a weak point in the method as presently implemented. In inviscid flow, there is no algorithmic mechanism to suppress slope discontinuities at intermacro-cell boundaries. Although reasonably rare and weak, slope discontinuities can be observed in the solutions.

Near-Boundary Micro-Cell Grid - Figure 4 is a schematic of a pair of adjacent macro-cells. The one-dimensional numerical problem is rendered in two dimensions to give it substance.

Solution data are stored at the micro-cell nodes (“grid points”) associated with the dots, which are planar (boundary) nodes or volumetric (interior) nodes. The two important issues concerning the micro-cell grid are the explicit distinction of volumetric and planar nodes, and the use of a smaller volumetric node adjacent to the boundary. This node is taken to be half the volume of the other interior nodes, in computational space. Note that conventional curvilinear grid stretching, in physical space, can be used in addition to the single half-sized cell in computational space. The stretching in computational space is explicitly reflected in the discretization formulae, while stretching in physical space is handled through the grid metrics as with conventional CFD methods.




Figure 4. Sketch of two adjacent macro-cells.

Boundary-Normal Schemes - Normal to a macro-cell boundary, higher accuracy spatial discretization used in the near-boundary region must be biased, reflecting the availability of data. It may also be optimized to improve the bandwidth and intrinsic stability of the expressions. The class of optimized biased methods used here have been




Figure 5. Error versus wave number for near-boundary polynomial and optimized interpolation schemes.

described previously by Treidler and Childs (1997). Other optimized methods have also been used (e.g., Lele, 1992). In the present work, optimized methods are constructed from a linear combination of polynomial schemes of different stencil size and order. The relative weights of the polynomial schemes are selected to provide a suitable combination of bandwidth, asymptotic convergence rate, and intrinsic stability. Part of this process involves forcing the error to zero at a few intermediate wave numbers chosen to give good overall characteristics. These zeros are visible in plots below. Within the macro-cell grid, there are many cell faces for which optimized biased schemes must be developed. Hundreds of coefficients are required for the various spatial discretization operations, and these are not presented here.

The benefits of using the macro-cell grid scheme with optimized biased discretization are illustrated in a pair of examples. Figure 5 gives the error magnitude versus wave number for interpolation to the second interior face for Taylor series schemes of various orders and an optimized scheme. To the target error tolerance of 10^{-3} used during the optimization, the optimized scheme has a bandwidth of roughly $\pi/2$, which is roughly 50% greater than achieved by the highest order Taylor-series scheme. Figure 6 gives the error magnitude for interpolation to the first three micro-cell faces in a macro cell. Face 1 is on the boundary and thus involves extrapolation, while Faces 2 and 3 are the first and second interior faces. The poorest bandwidth is achieved for pure extrapolation at $\pi/6$. This is the true accuracy only for




Figure 6. Several optimized near-boundary interpolation schemes.

extrapolation to a physical boundary. For extrapolation to an intermacro-cell boundary, the net effect of extrapolation *and averaging* with a neighbor cell actually improves the accuracy to be comparable to the interior schemes. Both interior formulae shown here, and others (not shown), achieve a bandwidth of $\pi/2$. Optimized and, where needed, biased schemes are also used in the 2-to-1 refinement across macro-cell faces, discussed below. Derivative calculations used in viscous terms are treated in a manner similar to that for the inviscid terms. The bandwidth of the viscous terms is generally equal to or better than inviscid terms.

Algorithmic (artificial) dissipation has proven to be one of the more challenging aspects of the macro-cell boundary treatment. The primary concern here is to achieve good bandwidth for the “linear” or background part of the dissipation which is dominant throughout the smooth parts of the flow. The ability to handle shock capturing across macro-cell boundaries is another important issue, but it is not addressed here.

The interior algorithmic dissipation consists of a “seventh-order” flux composed of a seventh derivative and a conventional spectral radius scaling. The critical issue is how to alter the high-order terms near the boundaries where high-order central schemes cannot be used. The possibility of simply stepping down the order of the fluxes, to fifth, and third order, at the boundary is unstable. This concept can be stabilized by using fourth and




Figure 7. Linear algorithmic dissipation at near-boundary and interior points.

sixth order dissipation at near-boundary points, but, of course, that dissipation is split into pairs of third and fifth order fluxes. Thus, the dissipation order at the interior faces moving away from a boundary are: null; 3rd; 3rd+5th; 5th+7th; 7th; etc.. The fluxes are initially scaled so that the sum of the magnitude of the dissipation coefficients is unity. This yields similar but not identical dissipation magnitude at the Nyquist frequency, and similar stability properties, at all near-boundary points. There is also a conventional global dissipation scaling.

Interpreted in terms of the error tolerance demanded for the convective and diffusive terms, the near-boundary dissipation has the worst bandwidth of all elements of the solver. For a dissipation of 10^{-3} , the bandwidth is $\pi/6$; conversely, the dissipation at $\pi/3$ is about 10^{-2} . Thus, algorithmic dissipation near boundaries is expected to be the dominant error in the current form of the macro-cell algorithm. Possible means of increasing the dissipation’s bandwidth have been identified but not implemented.

Boundary-Tangent Conditions - Within the plane of the macro-cell boundary, there are two conditions that must be accommodated: there may, or may not, be relative refinement between neighboring macro-cells, in which case the micro-cell faces do, or do not, match identically across the boundary.

If the micro-cells are exactly mated to a single micro-cell in the neighboring macro-cell, this aspect of the boundary condition is straightforward. The micro-cell-face areas are identical, and the solution and fluxes must be made to be identical. As noted previously, the solution and fluxes must be treated separately, in general, to insure that good accuracy and conservation are achieved. Within each macro-cell separately, the solution in a boundary node is obtained by extrapolation, as noted in the “Boundary-Normal” section. A common value for the face is then computed via averaging. In a calculation without viscous or algorithmic boundary fluxes, the fluxes from neighboring micro-cells are identical, because the solutions are identical, and do not require averaging. In calculations of viscous flows, the

fluxes depend also on the solution gradient. Within each macro-cell separately, the fluxes are computed at the cell face using biased (one-sided) difference operators. A common value for the boundary face flux is then computed via averaging.

The more challenging situation is when there is 2-to-1 refinement between neighboring macro-cells. This grid refinement is performed so that there is a regular pattern between micro-cells, in computational and physical spaces, as shown in Fig. 8. As before, this is a two-dimensional rendering of a one-dimensional problem. Conservation must now be enforced over the sum of all three macro-cell faces sketched in the Figure: the sum of all fluxes in the two smaller macro-cells, MC_1 and MC_2 , must equal the fluxes of the single larger one, MC_3 . As in other interpolation challenges here, accuracy is achieved by obtaining good bandwidth for the transfer of the line-integral solution across the macro-cell boundary. Optimized and, as needed, biased schemes are




Figure 8. Sketch of micro-cell face alignment for 2-to-1 refinement.




Figure 9. Error magnitude for boundary-tangent interpolation at cells with 2-to-1 refinement.

employed. The objectives of conservation and accuracy can be achieved during the optimization process: a linear system used to obtain accuracy is supplemented by an equation which represents exact conservation, and its solution satisfies both objectives. Since optimized schemes are used, it is straightforward to add a single extra condition shared across all interpolation formulae.

Examples of the accuracy achieved with these schemes are given in Fig. 9. This represents the transfer of the solution from the coarser grid to the finer one. Note that the error is measured in terms of wave number on the coarser grid, since the coarser grid will always be the “weakest link” in these conditions. The two curves are for interpolation to cells which are centered within the neighboring cell, and which straddle two neighboring cells. The minimum bandwidth achieved here is about $\pi/3$. This is a limiting factor in the accuracy for *physical* terms, as opposed to the *algorithmic* dissipation, at intermacro-cell boundaries.

Grid-Refinement

In a dynamical solution-adaptive calculation, the level of grid refinement must be adjusted dynamically in response to the adequacy of the grid for the evolving solution. This process involves error estimation and grid de/refinement. Error estimation will be provided by comparison of the energy spectrum of the solution with the error versus wave number characteristics of the spatial discretization. A companion paper by Reisentel and Childs (1999) describes how the maximum entropy method can be used to estimate the energy spectrum of the solution on a macro-cell grid and trigger solution adaption. At present, the error estimation method has not been incorporated in the macro-cell code. For demonstration purposes here, macro-cells requiring refinement are manually identified.

There are many possible strategies for grid refinement, each with strengths and weaknesses. For the present purposes, the de/refinement strategy must preserve high accuracy and conservation for the time-dependent solution. The method currently used involves recursive refinement of the initial

macro-cell grid. Each macro-cell in the initial grid can be split recursively into two, in any or all of the computational coordinates. This approach is related to that used in the AMR methods. To simplify the programming logic, the grid may be redefined only to the initial macro-cell grid, but no coarser. To achieve good efficiency, the initial grid should be as coarse as possible, so that a fine grid may be used only where required.

During the refinement process, the solution must be transferred between coarser and finer micro-cell grids. As in all other spatial operations, multi-dimensional operations are achieved with a sequence of one-dimensional operations. The same one-dimensional interpolation operators used to transfer the solution across boundaries with 2-to-1 are used to interpolate the solution between coarser and finer grids. This operation is conservative to the accuracy with which the metrics are computed; typical errors in the metrics are 10^{-4} for smooth micro-cell grids. The accuracy is identical to that for boundary 2-to-1 interpolation, as shown in Fig. 9.

Other Issues

Time Integration - Time integration is typically performed with an explicit Runge-Kutta algorithm or an implicit/explicit dual time-stepping scheme. Explicit time integration is straightforward and requires no special consideration of the macro-cell boundaries. Some form of implicit scheme is needed to achieve good efficiency for high-Reynolds number flows, but most “conventional” implicit schemes also suffer from errors due to approximate factorization, boundary conditions, and a limited order of temporal accuracy. Also, the use of a high-accuracy implicit step is expensive. The present effort uses a dual time-stepping scheme related to ones used by Venkatakrishnan and Mavriplis (1996) and DeRango and Zingg (1997). A block-implicit approximate-factorization scheme with second-order spatial approximations for all terms is used to compute an initial estimate of the solution at the new time level, Q^{n+1} . Then, the explicit RK scheme is used to iteratively improve the accuracy of Q^{n+1} in the global time-integration scheme

$$\frac{\partial Q}{\partial t}|^{n+1} = \alpha_{+1}Q^{n+1} + \alpha_0Q^n + \alpha_{-1}Q^{n-1}$$

in which the α_i are chosen to give second-order accuracy for an arbitrary time-step history. The error associated with the factored second-order implicit step is generally at moderate to high wave number, and the RK inner scheme gives good convergence rates. It is likely, however, that this approach can be improved upon significantly. Further discussion of implicit methods for high-accuracy schemes is given in a companion paper by Treidler et al (1999).

Turbulence Modeling - The UMC code currently includes the Spalart-Allmaras turbulence model. This model has received preliminary testing and appears to be performing correctly; however, it has not been thoroughly validated. The focus of this paper is on the accuracy of the macro-cell algorithm for nonturbulent flows.

Grid Generation - The macro-cell concept can be used with a range of different grid generation concepts, and may simplify some grid generation issues. At present however, an initial macro-cell grid is generated by subdividing a conventional finite difference grid into separate macro-cells. Generation of imbedded finer micro-cell grid points is handled automatically by the UMC code.

Software Issues - The UMC code is built on “object-based” concepts, is written in Fortran 90, and makes extensive use of the advanced features of that language. The brief description of the code given here touches on important highlights. There are objects which contain descriptions of the specific conditions of the flow being computed and of the surface geometry, and there is an object which is the solution. The “solution-object” is a linked list of macro-cells; thus, macro-cells are added or deleted as dictated by the needs of the solution. There is no limit to the size of this list, nor is memory wasted on “place-holder” macro-cells. Throughout all “higher” levels of the code, the entire solution is treated as a single object. At lower levels of the code, all operations are performed on individual macro-cells. Each macro-cell is a self-contained unit consisting of a

hierarchy of data structures and the specification of all operations to be performed on that data. The combination of object-based programming and macro-cell decomposition make the UMC code ideal for highly parallel machines, although the code has not yet been parallelized.

Most mathematical operators like interpolation and differentiation are written in a modular form, separate from the specific details of the governing equations, and useful in multiple related tasks. As noted previously, a single set of operations is used to perform all 2-to-1 interpolation operations. That interpolation module was developed and validated independent of the flow solver. Thus, the challenging task of writing bug-free high-accuracy spatial operators is performed just once, and then those operators are used as needed.

Results

Results are presented for a few flows to illustrate the accuracy of the flow solver. Test cases are selected to test specific aspects of the macro-cell code.

Acoustic Waves - This paper concentrates on the basic issues of the macro-cell algorithm, principally accuracy, and particularly near macro-cell boundaries and during de/refinement. A good means of examining basic properties is with acoustic problems, wherein the wave-number content of the solution can be specified.

Figure 10 gives the pressure contours as the wave initially passes over a macro-cell boundary containing 2-to-1 macro-cell refinement. The heavy lines are the macro-cell boundaries; the interior grids are not shown. In Fig. 10(a), the primary wavelength of this Gaussian signal corresponds to $\lambda \approx 0.3$, or about 6 grid points per period of a sine wave. This signal should be well within the resolving power of the algorithm and the solution can be seen passing over the macro-cell boundary with no visible disturbance. In Fig. 10(b), the primary wavelength of this Gaussian signal corresponds to $\lambda \approx 0.2$, or about 4 grid points per period of a sine wave. This signal slightly exceeds the theoretical bandwidth of the boundary scheme (predicated on 10^{-3} permissible

error) but the adverse effects are not readily apparent in this solution. Further time integration does reveal minor degradation in accuracy.




Figure 10(a). Contours of pressure for acoustic wave, $\lambda \approx 0.3$.




Figure 10(b). Contours of pressure for acoustic wave, $\lambda \approx 0.2$.

Another test case involved an acoustic pulse initialized inside a square box with inviscid walls. The domain was initially decomposed into four macro-cells, one for each corner of the square. During the calculation, further refinement occurred in one corner, so that there were numerous cases of 2-to-1 refinement across macro-cell boundaries. The peak strength of the pressure perturbation was 20% of the static pressure, which produces peak unsteady Mach numbers of roughly $M' \sim 0.1$. Throughout the duration of the calculation, which

ran 2000 iterations and a nondimensional time of $T = 20$, the total mass in the box was $1 \pm 25 \times 10^{-7}$, with zero mean drift. The total energy was also constant to the same precision. The relative precision for 32-bit arithmetic is $O(10^{-7})$, so the total mass and energy appear to be conserved to machine accuracy.

Grid Motion - The results of a simple test which illustrates the accuracy of convective terms for a moving grid are given in Fig. 11. The problem involves a stationary Gaussian thermal spot in stagnant fluid supported on a moving grid. As the grid rotates, the thermal spot convects around the computational domain, crossing three macro-cell boundaries in the process. The micro-cell grid, shown with dotted lines in one macro-cell, is similar in all macro-cells. The initial solution and the final solution after the grid has completed one full revolution are given by the solid and dashed lines, respectively. Contour levels are identical in both cases. The effects of algorithmic diffusion




Figure 11. Algorithmic diffusion of entropy spot due to grid motion (solid: initial; dashed: after 360° rotation).

can be seen in the azimuthal direction. This type of diffusive error is consistent with the known characteristics of the near-boundary algorithmic dissipation.

Captured Shock - Figure 12 gives the pressure contours from the flow over a NACA0012 airfoil at $M_\infty = 1.4$. The capture shock contains one interior grid point, and there are no significant

oscillations. Shocks have been observed crossing macro-cell boundaries without difficulties. However, the general issue of shocks and boundaries has not been systematically explored.




Figure 12. Example of shock captured with high-accuracy scheme and nonlinear artificial dissipation.

Macro-Cell Refinement Examples - A few examples are given of solutions subject to grid refinement. Figure 13 demonstrates the reduction of error as a grid is refined. The case involves inviscid flow at $M_\infty = 0.5$ and $\alpha = 3^\circ$. In the initial grid, there are 6 grid points around the leading edge, between $\pm 45^\circ$. The entropy, $s = p/\rho^\gamma$, is a good measure of error, as it should be constant ($s = 1/\gamma$) in this flow. Figure 13(a) reveals about a 0.5% error in the entropy on the coarse grid. After one level of refinement, Fig. 13(b), the error is reduced to less than 0.1%. With an additional level of refinement in the wall-normal direction(not shown), the entropy error diminished to less than 0.01%.




Figure 13(a). Entropy contours on initial grid.




Figure 13(b). Entropy contours, after grid refinement.

Periodic vortex shedding from a two-dimensional circular cylinder is used to demonstrate the operation of the UMC code. Macro-cell refinement is manually determined, as the automated methods (Reisenthel and Childs 1999) are not yet implemented. Figure 14(a) shows contours of ρu , the streamwise momentum, in the near-field of fully developed vortex shedding from a cylinder at $Re_D = 400$, $M = 0.25$. This initial grid contains 16 macro-cells, each with 13×13 grid points. The dual time-stepping scheme was used for this calculation and the time-step was fixed at $\Delta t = 0.1$, based on sound speed and radius. The wall normal grid spacing is such that roughly 2 grid points fall within the thinnest region of the boundary layer. Visual examination reveals that the boundary layer and the far downstream wake are under-resolved; however, the near-wake




Figure 14(a). Contours of streamwise momentum on initial grid.

resolution appears adequate. A clear indication of the inadequate resolution is the “overshoot” in momentum, which occurs where (potentially) countergradient algorithmic dissipation dominates viscous transport. Thus, the three macro-cells containing attached boundary layers were refined in the normal direction, while the downstream wake region was refined in both directions, as shown in Fig. 14(b). The thinnest boundary layer regions now contain about 3 grid points, as the boundary layer became thinner with increased resolution. Three boundary layer containing macro-cells are again split in the wall-normal direction. This forces wall-normal refinement in the near-wake macro-cell, to satisfy the 2-to-1 refinement limit between adjacent macro-cells. Figure 14(c) shows these macro-cells and the solution thereon. At this level of refinement, there




Figure 14(b). Streamwise momentum on grid after initial refinement.




Figure 14(c). Streamwise momentum on grid after final refinement.

is no visual evidence of inadequate resolution. The initial and final grids contained 16 and 32 macro-cells, respectively, each with 13×13 grid points.

Three points are made concerning details of the flow. Near the symbol “A” in Fig. 14(b), and at other locations, there is a 2-to-1 jump in the normal grid spacing of the inner micro-cell grid. These boundaries are treated well under most conditions. Here they are essentially “transparent.” At “B” the characteristic overshoot in the momentum, which is due to inadequate resolution, and the algorithmic dissipation, can be seen. At “C” there are discontinuities in the contour lines. This is solely due to the plotting package which employs, in essence, second-order collocation spatial discretization. If this solution is cast onto a uniformly refined macro-cell grid, this anomaly vanishes.

The Strouhal number, St , of this calculation is compared to a correlation of experimental data by Fey et al (1998) and to calculations of Rogers and Kwak (1988) in Fig. 15. The error in St in the present calculation is 2%.

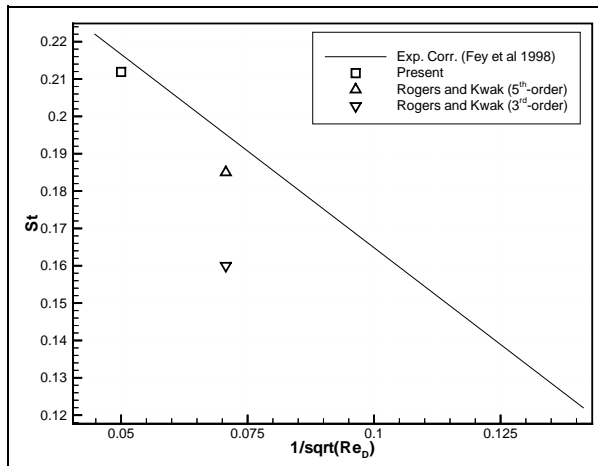


Figure 15. Strouhal number for calculations and experimental correlation of Fey et al (1998).

This example also points out a weakness in the solution-adaptive capabilities of this method. Because of the AMR-style macro-cell refinement and the 2-to-1 limit on refinement between neighbors, an initial grid whose clustering is very

poorly suited to the final flow field will require many levels of macro-cell refinement and, frequently, wasted sympathetic refinement. To achieve good resolution in the far-downstream wake with the initial grid above would be very inefficient.

Conclusions and Future Work

The high-accuracy unstructured macro-cell algorithm developed here is conceptually related to spectral element methods, but it uses finite-volume methods within each macro-cell, rather than spectral basis functions. The spatial accuracy of the macro-cell algorithm is regulated by the minimum bandwidth of the many spatial discretization operators. One benefit of using finite-volume methods is the availability of good time integration and shock capturing algorithms, which are needed for many aerospace applications. The macro-cell algorithm has performed well on a variety of model problems, including shock capturing and the propagation of acoustic and entropy modes across macro-cell boundaries. The algorithm’s performance observed in these tests is consistent with the theoretical accuracy of the method. The dominant error term in the present version of the code is algorithmic dissipation at macro-cell boundaries.

The near-term future work will focus on the needs of dynamic stall calculations at high Reynolds numbers. Two key tasks are (1) to include the turbulence modeling in the dual time integration method, as it currently runs only with the explicit Runge-Kutta scheme, and (2) to incorporate the automated grid-refinement triggering scheme described in Reisinger and Childs (1999). Longer-term goals include the inclusion of C_1 continuity at boundaries, parallelization of the code, and the development a temporally adaptive integration scheme.

Acknowledgements

The authors gratefully acknowledge support for this work provided by the Army Research Office under Contract DAAG55-97-0029.

References

- Barth, T. J. (1993) "Recent Developments in High-Order k-exact Reconstruction on Unstructured Meshes," AIAA-93-0668.
- Bernardi, C. and Maday, Y. (1998) "Spectral, Spectral Element, and Mortar Element Methods, presented at NASA/RTO workshop on High-Order Discretization Methods in CFD, NASA Ames Research Center, 21-25 Sept. 1998.
- Carpenter, M. H. and Casper, J. (1997) "The Accuracy of Shock-Capturing in Two Spatial Dimensions," AIAA-97-2107.
- Childs, R. E. (1993) "Prediction and Control of Turbulent Aero-Optical Distortion using Large Eddy Simulation," AIAA-93-2670.
- Childs, R. E. and Reisenthel, P. H. (1995) "Simulation Study of Compressible Turbulent Boundary Layers," AIAA-95-0582.
- Childs, R. E. and Reisenthel, P. H. (1997) "Dynamical Compressibility Effects in Simulated Turbulent Boundary Layers," AIAA-95-0428.
- Derango, S. and Zingg, D. W. (1997) "Improvements to a Dual-Time-Stepping Method for Computing Unsteady Flows," AIAA J., Vol. 35, No. 8, p. 1548.
- Fey, U., Konig, M., and Eckelmann, H. (1998) "A New Strouhal-Reynolds-Number Relationship for the Circular Cylinder in the Range $47 < Re < 2 \times 10^5$," Physics of Fluids, Vol. 10, No. 7, pp.1547-1549.
- Henderson, R. D. and Karniadakis, G. E. (1995) "Unstructured Spectral Element Methods for Simulation of Turbulent Flows," J. Comp. Phys., Vol. 122, pp. 191-217.
- Henderson, R. D. (1998) "Adaptive Spectral Element Methods for Turbulence and Transition," presented at NASA/RTO workshop on High-Order Discretization Methods in CFD, NASA Ames Research Center, 21-25 Sept. 1998.
- Lele, S. K. (1992) "Compact Finite Difference Schemes with Spectral-Like Resolution," J. Comp. Phys., No. 103, pp. 16-42.
- Reisenthel, P. H. and Childs, R. E. (1999) "Wave Number-based Criterion for Dynamic Mesh Refinement in CFD," AIAA-99-0300
- Rogers, S. E. and Kwak, D. (1988) "An Upwind Differencing Scheme for the Time-Accurate Incompressible Navier-Stokes Equations," AIAA-88-2583.
- Sherwin, S. J. and Karniadakis, G. E. (1996) "Tetrahedral hp Finite Elements: Algorithms and Flow Simulations," J. Comp. Phys., Vol. 124, pp. 14-45.
- Treidler, E. B. and Childs, R. E. (1997) "High-Accuracy Spatial Discretization Schemes for CFD," AIAA-97-0541.
- Treidler, E. B., Ekaterinaris, J. A., and Childs, R. E. (1999) "Efficient Solution Algorithms for High-Accuracy Central Difference CFD Schemes," AIAA-99-0302.
- Venkatakrisnan V. and Mavriplis, D. J. (1996) "Implicit Method for the Computation of Unsteady Flows on Unstructured Grids," J. Comp. Phys., Vol. 127, pp. 380-397.