

Three-dimensional unsteady incompressible flow computations using multigrid

Andrey Belov
Princeton Univ., NJ

Luigi Martinelli
Princeton Univ., NJ

Antony Jameson
Princeton Univ., NJ

AIAA, Aerospace Sciences Meeting & Exhibit, 35th, Reno, NV, Jan. 6-9, 1997

We apply a robust and computationally efficient multigrid-driven algorithm for the simulation of time-dependent three-dimensional incompressible bluff body wakes at low Reynolds numbers (Re less than or equal to 350). The computational algorithm combines a generalized time-accurate artificial compressibility approach, a finite-volume discretization in space, and an implicit backward discretization in time. The solution is advanced in time by performing iterative 'pseudo-transient' steady-state calculations at each time step. The key to the algorithm's efficiency is a powerful multigrid scheme that is employed to accelerate the rate of convergence of the pseudo-transient iteration. The computational efficiency is improved even further by the application of residual smoothing and local pseudo time-stepping techniques, and by using a point-implicit discretization of the unsteady terms. The solver is implemented on a multiprocessor IBM SP2 computer by using the MPI Standard, and a high parallel scalability is demonstrated. The low Reynolds number regime (Re less than or equal to 500) encompasses flow transitions to unsteadiness and to three-dimensionality and attracts considerable attention as an important step on the road to turbulence. In this regime, the slow asymptotics of the wake provide a challenging test for numerical methods since long integration times are necessary to resolve the flow evolution toward a limiting cycle. Our method is extended to three dimensions and applied for low Reynolds number flows over a circular cylinder (Re less than or equal to 250) and a circular semi-cylinder ($Re = 350$). The computational results are found to be in close agreement with the available experimental and computational data. (Author)

Three-Dimensional Unsteady Incompressible Flow Computations Using Multigrid.

Andrey Belov,* Luigi Martinelli,† and Antony Jameson‡
 Department of Mechanical and Aerospace Engineering
 Princeton University
 Princeton, New Jersey 08544 U.S.A.

1 Abstract

In recent years, a number of fast and accurate computational algorithms have been developed for solving steady-state Euler and Navier-Stokes equations. However, several fluid phenomena of great industrial and scientific interest are inherently time-dependent. These include, for example, wake flows, turbulent flows, and free-surface flows. In the present work we apply a robust and computationally efficient multigrid-driven algorithm for the simulation of time-dependent three-dimensional incompressible bluff body wakes at low Reynolds numbers ($Re \leq 350$).

The computational algorithm combines a generalized time-accurate artificial compressibility approach, a finite-volume discretization in space, and an implicit backward discretization in time. The solution is advanced in time by performing iterative, "pseudo-transient" steady-state calculations at each time step. The key to the algorithm's efficiency is a powerful multigrid scheme that is employed to accelerate the rate of convergence of the pseudo-transient iteration. The computational efficiency is improved even further by the application of residual smoothing and local pseudo-time stepping techniques, and by using a point-implicit discretization of the unsteady terms. The solver is implemented on a multiprocessor IBM SP2 computer by using the MPI Standard, and a high parallel scalability is demonstrated.

Bluff cylindrical body wakes are being actively investigated by a number of authors with the focus on the mechanism of momentum and vorticity transport [3], [26], [40] and on the active control of the wake structure [28]. The low Reynolds number regime ($Re \leq 500$) encompasses flow transitions to unsteadiness and to three-dimensionality and at-

tracts considerable attention as an important step on the road to turbulence. In this regime, the slow asymptotics of the wake provide a challenging test for numerical methods since long integration times are necessary to resolve the flow evolution toward a limiting cycle. Our method is extended to three-dimensions and applied for low Reynolds number flows over a circular cylinder ($Re \leq 250$) and a circular semi-cylinder ($Re = 350$). The computational results are found to be in a close agreement with the available experimental and computational data.

2 Nomenclature

R	cylinder radius
C_d	drag coefficient of a cylinder $C_d = \frac{D}{\rho U_\infty^2 R}$
C_l	coefficient of lift
C_{pb}	base pressure coefficient $C_{pb} = \frac{2(p-p_\infty)}{\rho U_\infty^2}$ at 180° from the front stagnation point
\vec{U}_b	body velocity
St	Strouhal number
\mathbf{w}	vector of flow variables
\mathbf{f}_i	Euler flux vectors
\mathbf{s}_i	viscous flux vectors
τ_{ij}	viscous stress tensor
$\mathcal{R}(\mathbf{w}_{ijk})$	flux residual for cell i,j,k
\mathcal{R}^*	modified residual
t	real time
t^*	pseudotime
Δt	implicit real time step
p	static pressure
p_∞	static pressure at the far field
\vec{u}	Cartesian velocity vector
u_i	Cartesian velocity components
$u_{mesh,i}$	mesh Cartesian velocity components
U_∞	free stream velocity, velocity scale
U_n	contravariant velocity
a_n	artificial speed of sound
δ_{ij}	Kronecker delta
λ_i	characteristic wave speeds
X_{ni}	metric coefficients in i,j and k direction
\vec{n}	surface normal vector
V_{ijk}	volume of i,j,k cell
$\Omega, \partial\Omega$	cell element and its boundary

Copyright ©1997 by the Authors. Published by the AIAA Inc. with permission.

*Ph.D. Student

†Professor of Aerospace Engineering, Member of AIAA

‡James S. McDonnell Distinguished University Professor of Aerospace Engineering, AIAA Fellow

3 Introduction

Recent advances in numerical methods and in parallel computing technology make it possible at the present time to simulate stationary viscous flows around complex aerospace geometries in a matter of hours. In particular, the finite-volume method of A. Jameson and L. Martinelli, originally proposed for compressible flow calculations, allows the computation of steady-state solutions of the Reynolds-Averaged Navier-Stokes equations on meshes with $\sim 1.8 \cdot 10^6$ nodes in less than 1.5 hours. However, a wide range of flow problems can not be adequately modeled with the time-independent governing equations. In particular, flow around a maneuvering submarine, sea keeping problems, various wake and turbulent flows fall into this category. Time-dependent computations in three dimensions require both a very fast CPU and a very large memory and often present a challenge even for the most powerful modern computers. The difficulty of unsteady flow simulations is fundamentally more complex than just an increase of the problem dimensionality by one. Unlike the spatial dimensions, the temporal span can not be decomposed in subdomains and parallelized, because the time stepping can only be performed sequentially. This work is motivated by a great practical interest in time-accurate and computationally efficient algorithms for unsteady viscous flow simulations. In particular, we extend an implicit multigrid-driven approach, originally proposed in references [5], [6], to three dimensions and validate it for low Reynolds number flows ($Re \leq 350$) around a circular cylinder and a circular half-cylinder.

In calculations of incompressible flow, the continuity equation imposes a time-independent solenoidal constraint on the velocity field. The governing system of equations does not include a temporal evolution term for pressure thus precluding application of straightforward time marching techniques. This difficulty can be circumvented by casting Chorin's artificial compressibility concept [8] into an implicit time-accurate technique. In particular, the artificial compressibility method is applied at each time step, i.e. an unsteady problem is converted into a number of iterative, "pseudo-transient" steady state computations. The dual time stepping idea dates back to the pioneering work of Peyret [27], and it received considerable attention in literature. In recent years, this strategy has been applied by Taylor and Whitfield [35] by coupling a Newton iteration scheme with an upwind flux-difference splitting scheme, and by Rogers and Kwak[33] using an implicit line relaxation scheme in pseudo-time. Miyake et al. [25] used an explicit, up to second order accurate discretization in time, and a rational Runge-Kutta scheme for the subiterations. As was proposed by Jameson for compressible flow calculations [21], a multigrid

scheme can be employed at each time step to expedite convergence of the subiterations. A similar approach can be applied for incompressible flow computations since, as it has been shown by Rizzi and Eriksson [31], the artificial compressibility preconditioning renders the pseudo-transient system hyperbolic (in the inviscid case), thus allowing for the application of standard compressible flow algorithms to march toward a steady state in pseudo-time.

Our method, originally presented in references [5], [6], for two-dimensional problems, combines the advantages of the artificial compressibility preconditioning, backward implicit time discretization and the very efficient finite-volume multigrid acceleration technique. The resulting scheme has favorable numerical stability properties in physical time, achieves fast convergence in pseudo-time and, due to the local spatial support and explicit nature of the pseudo-transient subiterations, is highly amenable to parallel implementation. Additional algorithmic improvements include a point-implicit treatment of the unsteady terms, approximately nonreflective characteristic far-field boundary conditions, and parallel implementation using the Message Passing Interface (MPI) Standard. The main features of the present algorithm are outlined in the following sections and are explained in more detail in reference [7].

4 Governing Equations

The governing equations of the flow considered are the nonlinear time-dependent incompressible Navier-Stokes equations of a laminar, constant viscosity flow without body forces.

Let p , u_i denote the pressure and Cartesian velocity components respectively, where $i=1,2,3$. Consider a control volume Ω with boundary $\partial\Omega$ which moves with Cartesian velocity components $u_{mesh,i}$. If the total volume enclosed by Ω does not change with time, the equations of motion of the fluid can be compactly written in the following integral nondimensional form

$$\mathbf{I}^m \cdot \frac{d}{dt} \iiint_{\Omega} \mathbf{w} \, d\Omega + \oint_{\partial\Omega} (\mathbf{f}_1 - \mathbf{s}_1) \, \partial\Omega_i = 0 \quad (1)$$

where \mathbf{w} is the vector of flow variables

$$\mathbf{w} = \begin{Bmatrix} p \\ u_1 \\ u_2 \\ u_3 \end{Bmatrix},$$

the \mathbf{f}_1 , \mathbf{s}_1 are the Euler and the viscous flux vectors

correspondingly

$$\mathbf{f}_i = \begin{Bmatrix} (u_i - u_{mesh_i}) \\ u_1(u_i - u_{mesh_i}) + p\delta_{1i} \\ u_2(u_i - u_{mesh_i}) + p\delta_{2i} \\ u_3(u_i - u_{mesh_i}) + p\delta_{3i} \end{Bmatrix}, \quad \mathbf{s}_i = \begin{Bmatrix} 0 \\ \tau_{i1} \\ \tau_{i2} \\ \tau_{i3} \end{Bmatrix}.$$

The dimensionless viscous stress tensor is

$$\tau_{ij} = \frac{1}{Re} \left[\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right]$$

where $Re = \frac{U_\infty L}{\nu}$ is the Reynolds number, ν is the kinematic viscosity, U_∞, L are the velocity and the length scales of the problem, and \mathbf{I}^m is the modified identity matrix annihilating the temporal derivative of pressure from the continuity equation

$$\mathbf{I}^m = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

A cell-centered formulation is employed in the present method so that, when applied independently to each cell in the mesh domain, Equation (1) results in a set of ordinary differential equations of the form

$$\mathbf{I}^m \cdot \frac{d}{dt} \iiint_{V_{ijk}} \mathbf{w} dV + \mathcal{R}(\mathbf{w}_{ijk}) = \mathbf{0}, \quad (2)$$

where V_{ijk} is the volume of the i, j, k cell and the residual $\mathcal{R}(\mathbf{w}_{ijk})$ is obtained by evaluating the sum of the inviscid and viscous flux integrals. The details of the discretization of the viscous fluxes are described in [23] and will not be presented here. The convective fluxes are approximated with central differences in space, and a third order artificial dissipation term is added in conservation form in order to prevent the odd-even decoupling stemming from the central differencing of convective terms.

Equation (2) can be simplified as follows

$$\mathbf{I}^m \cdot \frac{d}{dt} (\mathbf{w}_{ijk} V_{ijk}) + \mathcal{R}(\mathbf{w}_{ijk}) = \mathbf{0}, \quad (3)$$

where \mathbf{w}_{ijk} are the integral averages of flow variables \mathbf{w} taken in each computational cell i, j, k . The cell volumes can be factored out since they do not change in computation. Equation (3) can be approximated implicitly as follows (subscripts i, j, k are dropped for clarity)

$$\mathbf{I}^m \cdot \frac{d}{dt} [\mathbf{w}^{n+1} V^{n+1}] + \mathcal{R}(\mathbf{w}^{n+1}) = \mathbf{0}, \quad (4)$$

where $n+1$ denotes to the current time level $\mathbf{w} = \mathbf{w}^{n+1}$. The d/dt operator is approximated by an implicit backwards difference scheme [12] of the following form

$$\frac{d}{dt} = \frac{1}{\Delta t} \sum_{q=1}^k \frac{1}{q} [\Delta^-]^q, \quad \Delta^- = \mathbf{w}^{n+1} - \mathbf{w}^n \quad (5)$$

In particular, for a third order discretization in time Equation (4) becomes

$$\frac{1}{\Delta t} \mathbf{I}^m \cdot (p_1 [\mathbf{w} V] + p_2 [\mathbf{w}^n V^n] + p_3 [\mathbf{w}^{n-1} V^{n-1}] + p_4 [\mathbf{w}^{n-2} V^{n-2}]) + \mathcal{R}(\mathbf{w}) = \mathbf{0}, \quad (6)$$

where $\mathbf{p} = (11/6, -3, 3/2, -1/3)$. The third order discretization is stiffly stable, while a second order scheme, given by the coefficients $\mathbf{p} = (3/2, -2, 1/2, 0)$, is A-Stable[9].

It is convenient at this point to define the modified residual $\mathcal{R}^*(\mathbf{w})$ as

$$\mathcal{R}^*(\mathbf{w}) = \mathbf{I}^m \cdot \frac{V}{\Delta t} \sum_{q=1}^4 p_q \mathbf{w}^{n-q+2} + \mathcal{R}(\mathbf{w}), \quad (7)$$

which is driven to zero at each mesh point i, j, k , at each time step by marching toward a steady state in the fictitious time (pseudotime) t^* .

4.1 Local Preconditioning

In the limit of incompressible flow sound waves travel infinitely fast in all directions and dominate the flow. The disparity in the sound and convective wave speeds makes the governing system of equations ill-conditioned. In order to render the system of equations more suitable for numerical computation, a local preconditioning matrix \mathbf{Pr} is introduced yielding

$$\frac{d\mathbf{w}}{dt^*} + \mathbf{Pr} \cdot \mathcal{R}^*(\mathbf{w}) = \mathbf{0}, \quad (8)$$

where

$$\mathbf{Pr} = \begin{bmatrix} \Gamma^2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

The parameter Γ is selected to improve the rate of convergence in pseudotime. Since the unsteady terms enter the modified residual as source terms, and do not affect the range of the wave speeds present in the system (8), the preconditioning suggested by Rizzi and Eriksson [30], and later by Dreyer [10] can be used. It is adopted in the following form

$$\Gamma^2 = \max(C_1, C_2(u_1^2 + u_2^2 + u_3^2)),$$

where $C_1 = 0.25$ and $C_2 = 1$. Let us define the scaled contravariant mesh and flow velocities as follows

$$\begin{aligned} U_n &= X_{ni} \cdot u_i \\ u_{mesh_n} &= X_{ni} \cdot u_{mesh_i}, \end{aligned}$$

where n stands for i, j or k and the X_{ni} are the metric coefficients of the corresponding coordinate direction. The corresponding spectrum of eigenvalues

is then given by

$$\begin{aligned}\lambda_1 &= U_n - u_{mesh} \\ \lambda_2 &= U_n - u_{mesh} \\ \lambda_3 &= (U_n - \frac{1}{2}u_{mesh}) + a_n \\ \lambda_4 &= (U_n - \frac{1}{2}u_{mesh}) - a_n\end{aligned}\quad (9)$$

The artificial speed of sound a_n is

$$a_n = \sqrt{(U_n - \frac{1}{2} \cdot u_{mesh})^2 + \Gamma^2 \cdot (X_{n1}^2 + X_{n2}^2 + X_{n3}^2)}.$$

Since Γ can never be zero, the four eigenvalues can not have the same sign. The pseudo-transient flow is therefore always "subsonic", i.e. it corresponds to the proper domain of dependence in incompressible flow computations.

The Jacobian matrices and the left and right eigenvectors, corresponding to the pseudocompressibility approach formulated in three dimensions, can be found in [32].

Equation (8) represents an explicit set of ordinary differential equations to be solved at each time step. It is easy to see that, once the steady state in pseudotime t^* is achieved, the modified residuals $\mathcal{R}^*(\mathbf{w})$ are driven to zero at each mesh point, yielding the solution of Equation (6) at the new, $n+1$ time level. Then, the time levels in Equation (6) are shifted by minus one, and the new steady state problem is set up. Thus, the original unsteady problem (3) is converted into several steady state calculations in t^* . A five-stage scheme is used to advance the solution in pseudotime. Since the details of the pseudotransient evolution are immaterial, a number of powerful convergence acceleration techniques are employed. In particular, a full approximation multigrid, point-implicit local pseudotime stepping and residual averaging are applied as explained in the following sections.

4.2 Artificial dissipation

In order to prevent odd-even decoupling, stemming from the central differencing of convective terms, a third order artificial dissipation term is added to spatial residuals $\mathcal{R}(\mathbf{w})$. This term is constructed in conservative form and premultiplied by the inverted preconditioning matrix as follows

$$\begin{aligned}\mathcal{R}(\mathbf{w}) &= \mathbf{C}(\mathbf{w}) - \mathbf{D}(\mathbf{w}) \\ &+ \mathbf{Pr}^{-1} \cdot (\mathbf{d}_i^+ - \mathbf{d}_i^- + \mathbf{d}_j^+ - \mathbf{d}_j^- + \mathbf{d}_k^+ - \mathbf{d}_k^-)\end{aligned}\quad (10)$$

where $\mathbf{C}(\mathbf{w})$, $\mathbf{D}(\mathbf{w})$ are the convective and viscous terms and \mathbf{d}_n^\pm are the cell boundary fluxes (the n below stands for i, j or k)

$$\mathbf{d}_n^\pm = \epsilon \cdot \delta_n^2(\mathbf{w}_n^\pm), \quad (11)$$

where δ_n^2 denotes the second difference operator and

$$\begin{aligned}\mathbf{w}_n^+ &= (\mathbf{w}_{n+1} - \mathbf{w}_n) \cdot 2 \cdot \min(R_n, R_{n+1}) \\ \mathbf{w}_n^- &= (\mathbf{w}_n - \mathbf{w}_{n-1}) \cdot 2 \cdot \min(R_n, R_{n-1})\end{aligned}$$

where ϵ is a scaling factor, $\epsilon = 1/64$, and R_n is a spectral radius in i, j or k direction, derived from $\lambda_{1,2,3,4}$ by using equations (9)

$$R_n = \max(|U_n - u_{mesh}|, |U_n - \frac{1}{2}u_{mesh}| + a_n) \quad (12)$$

One can see that, should the spectral radii be constant across the neighboring cells, the artificial dissipation term reduces to the sum of scaled undivided fourth differences of the flow variables. A simpler artificial dissipation model is employed on the coarser meshes of the multigrid cycle, which is obtained from the one described above by omitting the δ_n^2 operators in Equation (11).

4.3 Point-implicit pseudotime stepping

A five-stage generalized Runge-Kutta type scheme is used to advance the solution in pseudo-time. The scheme originates from the base scheme optimized for faster rate of convergence in steady state viscous computations [23]. A straightforward application of this scheme to unsteady flow calculations amounts simply to replacing the spatial residual $\mathcal{R}(\mathbf{w})$ with the modified residual $\mathcal{R}^*(\mathbf{w})$.

$$\begin{aligned}\mathbf{w}^{(n+1,0)} &= \mathbf{w}^n \\ \mathbf{w}^{(n+1,1)} &= \mathbf{w}^{(n+1,0)} - \frac{\alpha_1 \Delta t^*}{V} \mathcal{R}^*(\mathbf{w})^{(n+1,0)} \\ &\dots \\ \mathbf{w}^{(n+1,q)} &= \mathbf{w}^{(n+1,0)} - \frac{\alpha_q \Delta t^*}{V} \mathcal{R}^*(\mathbf{w})^{(n+1,q-1)} \\ \mathbf{w}^{n+1} &= \mathbf{w}^{(q)},\end{aligned}\quad (13)$$

The $\mathcal{R}^*(\mathbf{w})^{(n+1,q)}$ denotes the $q+1$ st stage update of the modified residual. It is not simply computed based on the current value of the last update $\mathbf{w}^{(n+1,q)}$ but rather is comprised of a linear combination of convective and dissipative contributions, evaluated on the current q -th and previous $r = 0, 1, \dots, q-1$ stages [23], plus the unsteady terms. It is convenient to separate the unsteady source term at the current time level from the rest of the modified residual $\mathcal{R}^*(\mathbf{w})$ given by the Equation (7)

$$\mathcal{R}^*(\mathbf{w}) \equiv \mathcal{R}^{**}(\mathbf{w}) + \mathbf{I}^{\mathbf{m}} \cdot \frac{p_1 V \mathbf{w}}{\Delta t}, \quad (14)$$

where $\mathcal{R}^{**}(\mathbf{w})$ includes the spatial residual $\mathcal{R}(\mathbf{w})$ and the unsteady source terms coming from the previous time levels only (which remain constant in t^*). It is easy to see from Equations (13) and (14) that

the unknown source term shifts the Fourier symbol of the augmented residual operator to the left along the negative real axis. The shift is proportional to $1/\Delta t$ and leads to a stability bound on the pseudotime step Δt^* relative to the physical time step Δt , if an explicit scheme is used. This stability constraint may have a negative effect on the computational efficiency of the algorithm when fine temporal resolution is required. In the present algorithm, this restriction is relaxed by employing a point-implicit Runge-Kutta method, first proposed in reference [24] for compressible flow computations. The source terms are regarded on the left-hand side of the pseudotime stepping scheme, i.e. treated implicitly. The point-implicit multistage scheme used in the present work can therefore be written as (omitting the $n + 1$ index for clarity)

$$\begin{aligned}
 \mathbf{w}^{(0)} &= \mathbf{w}^n \\
 \mathbf{w}^{(1)} &= \left(\mathbf{w}^{(0)} - \frac{\alpha_1 \Delta t^*}{V} \cdot \mathcal{R}^{**}(\mathbf{w})^{(0)} \right) \\
 &\quad \times \left(\mathbf{I} + \mathbf{I}^{\mathbf{m}} \cdot \frac{\alpha_1 p_1 \Delta t^*}{\Delta t} \right)^{-1} \\
 &\dots \\
 \mathbf{w}^{(q)} &= \left(\mathbf{w}^{(0)} - \frac{\alpha_q \Delta t^*}{V} \cdot \mathcal{R}^{**}(\mathbf{w})^{(q-1)} \right) \\
 &\quad \times \left(\mathbf{I} + \mathbf{I}^{\mathbf{m}} \cdot \frac{\alpha_q p_1 \Delta t^*}{\Delta t} \right)^{-1} \\
 \mathbf{w}^{n+1} &= \mathbf{w}^{(q)}
 \end{aligned} \tag{15}$$

where \mathbf{I} is the identity matrix and, since the continuity equation has no unsteady terms, the point-implicit construction only affects the momentum equations.¹

The point-implicit construction rescales the stability region of the multistage scheme such that the pseudotime step Δt^* is decoupled from the physical time step Δt , reference [24], and can be determined solely as required by the spatial operator $\mathcal{R}(\mathbf{w})$. For Euler calculations a conservative estimate is given by

$$\Delta t^* \leq \frac{\text{CFL} \cdot V}{R_i + R_j + R_k}$$

where spectral radii R_i, R_j and R_k are computed by using equations (12). In Navier-Stokes computations the denominator is augmented by the viscous term R_v , computed as in reference [23].

4.4 Multigrid

The multigrid scheme is a full approximation scheme defined as in reference [20]. This convergence acceleration technique takes advantage of corrections,

¹Note also that the multigrid and the residual smoothing must still operate on the full residual $\mathcal{R}^*(\mathbf{w})$, not on the $\mathcal{R}^{**}(\mathbf{w})$, since it is the $\mathcal{R}^*(\mathbf{w})$ which must reach zero upon convergence.

computed on a series of successively coarser meshes, driven by the residuals accumulated on the finest mesh. These corrections are then bilinearly interpolated on the finer mesh. The time stepping on the coarser meshes in a multigrid cycle effectively speeds up information transfer within the computational domain and provides fast smoothing of the long wavelength error modes on the finest mesh.

The implementation is as follows. Denote the grids by a subscript k . Start with a time step on the finest grid $k = 1$. Transfer the solution from a given grid to a coarser grid by a transfer operator $P_{k,k-1}$, so that the initial state on grid k is

$$w_k^{(0)} = P_{k,k-1} w_{k-1}.$$

Then on grid k the multistage time stepping scheme is reformulated as

$$w_k^{(q+1)} = w_k^{(0)} - \alpha_n \Delta t^* \left(R_k^{*(q)} + G_k \right),$$

where the residual $R_k^{*(q)}$ is evaluated from current and previous values as above. The forcing function G_k is defined as the difference between the aggregated residuals transferred from grid $k - 1$ and the residual recalculated on grid k . Thus

$$G_k = Q_{k,k-1} R^*(w_{k-1}) - R^*(w_k^{(0)}),$$

where $Q_{k,k-1}$ is another transfer operator. On the first stage the forcing term G_k simply replaces the coarse grid residual by the aggregated fine grid residuals. The accumulated correction on a coarser grid is transferred to the next higher grid by an interpolation operator $I_{k-1,k}$ so that the solution on grid $k - 1$ is updated by the formula

$$w_{k-1}^{new} = w_{k-1} + I_{k-1,k} \left(w_k - w_k^{(0)} \right).$$

The whole set of grids is traversed in a W -cycle in which time steps are only performed when moving down the cycle.

4.5 Boundary Conditions

In the present method, the boundary conditions are imposed as follows.

1. A single-level halo of fictitious cells is added to the mesh.
2. The analytical boundary conditions are discretized on the boundary separating the interior of the computational domain and the halo.
3. Based on this discretization, variables at fictitious cells w_e are derived.
4. In the computation, the interior cells adjacent to the boundary are treated identically to the cells in the bulk of the flow.

5. The external variables w_e are updated each time when the interior variables are updated, i.e. on every physical time step, pseudo-time step and stage of the multistage pseudo-time scheme.

4.5.1 Solid Boundary

For viscous calculations, a no-slip condition is imposed on the solid boundary by setting the flow velocity equal to that of the body, $\mathbf{u} = \mathbf{u}_{mesh}$. The velocity in a fictitious cell is then computed as $\mathbf{u}_1 = 2\mathbf{u}_{mesh} - \mathbf{u}_2$, where \mathbf{u}_2 is a velocity vector in the interior cell adjacent to the solid surface. In a cell-centered approach, the pressure at the wall must be computed from the interior values of the flow variables. For high Reynolds number flows and a non-accelerating body the approximation of a zero pressure gradient normal to the surface can be used to compute the pressure. However, since the present work addresses the low Reynolds number regime, $Re \leq 350$, a balance between the normal component of the viscous stress, body acceleration and the pressure gradient is used to compute the pressure gradient

$$\frac{dP}{dn} = n_i \frac{d\tau_{ij}}{dx_j} - \mathbf{n} \cdot \frac{d\mathbf{u}_{surface}}{dt},$$

where summation in i and j is assumed. The second order accurate central difference discretization of the $\frac{dP}{dn}$ at the solid boundary then allows one to compute an approximation for the pressure. Note that in the parallel implementation (discussed in the following section) this boundary condition has a higher computational cost as compared to the zero pressure gradient approximation - an extra message with the velocity component updates has to be exchanged by the processors in order to compute the viscous stress tensor.

4.5.2 Periodic Approximation

The three-dimensional calculations performed in the present work simulate a body of infinite span by employing periodic boundary conditions along the spanwise direction. This assumption imposes two equality constraints which specify the variables in fictitious cells, as illustrated in Figure (1). The limitations of a spanwise periodicity assumption and its effects on the computed flow solution still remain a topic of active research and are discussed in more detail in reference [4]. Note that the present method allows for arbitrary boundary conditions, consistent with the governing flow equations, to be imposed along the spanwise direction which includes a side wall or zero spanwise gradient possibilities. This illustrates a flexibility advantage of the present approach versus the quasi three-dimensional discretization using Fourier modes.

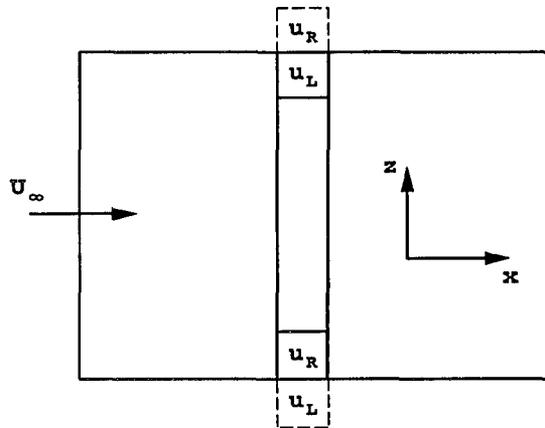


Figure 1: Spanwise Periodic Boundary Conditions (Top View of the Computational Domain).

4.5.3 Far-Field Boundary

On the outer boundaries, approximate non-reflecting far-field boundary conditions are constructed based on the linearized characteristics approach to improve accuracy and the rate of convergence in pseudotime. This formulation does not separate the boundary into inflow and outflow segments, but rather treats the entire domain boundary in a uniform way. This is accomplished by recombining the appropriate information both from the flow interior w^{int} and from the free-stream values w^∞ at each boundary segment² so that incoming and outgoing waves are distinguished and properly treated. This approximation aims at propagating the waves out of the computational domain with minimal impediment from the boundary conditions, thus helping the flow to evolve toward the converged state.

The general procedure described in the beginning of this Section is followed in formulating the far-field boundary conditions. A single level halo of fictitious cells is added to the grid, and the corresponding far field values are initially taken to be equal to the free stream values. Each time the flow variables are updated, the linearized characteristic problem is solved along the direction normal to the boundary of the computational domain.³ In the current implementation, the spanwise velocity component u_3 is treated as a passive scalar, and it is either extrapolated from the interior of the computational domain, or is taken to be equal to its free stream value, depending on the local flow direction. In order to determine p , u_1 , and

²The interior variables are taken at the cell-centers of the outermost radial coordinate layer of the computational domain.

³The formulation assumes that the pseudo-transient waves impinge normally upon the boundary, which is particularly justified for the topology of the O-meshes used throughout this work.

u_2 in the fictitious cells, a two-dimensional similarity transform of the Jacobian matrix [33] is computed on the outer boundary of the domain at each stage of the Runge-Kutta scheme. The characteristic variables are defined as follows

$$\begin{aligned} \mathbf{w}^{char} &= \mathbf{T}^{-1} \cdot \mathbf{w} \\ \mathbf{w} &= \mathbf{T} \cdot \mathbf{w}^{char}, \end{aligned} \quad (16)$$

where $\mathbf{T}, \mathbf{T}^{-1}$ are the similarity matrices defined in reference [33]. Depending on the sign of the corresponding wave speed in Equation (9), the characteristic variables in the dummy cells are either interpolated from the interior of the domain, or computed from the free stream values. In particular,

$$w_e^{char} = \frac{\text{sign}(\lambda_i - \epsilon) + 1}{2} \cdot w_i^{int} - \frac{\text{sign}(\lambda_i - \epsilon) - 1}{2} \cdot w_i^\infty$$

where $i = 1, 2, 3$ and sign is a Heaviside step function. Here ϵ is a small constant ($\epsilon = 10^{-13}$) which serves to prevent oscillations due to roundoff errors when $\lambda_1 \sim 0$ (i.e. when the local flow velocity is tangential to the domain boundary). Finally, the fluxes on the far field boundary are calculated using the velocity and pressure values in the fictitious cells, recovered using the inverse similarity transform (16).

Better approximations to the characteristic far-field boundary conditions can be constructed using the full three-dimensional similarity transform proposed in reference [32]. Implementation of these conditions and the comparative analysis with the present formulation identifies an important area for future research.

5 Parallel implementation

Fast processors and large internal memory are typically required to obtain time-resolved Navier-Stokes solutions in practical geometries. One of the modern approaches to this problem is to employ parallel architectures. An efficient parallel implementation of the computational algorithm and assessment of its scalability are, therefore, integral parts of the flow solver development. These issues are discussed in the following sections.

5.1 Domain Decomposition

The flow solver is parallelized by using the Message Passing Interface (MPI) Standard on an IBM SP2 multiprocessor computer. The single program, multiple data (SPMD) approach is employed. An identical set of instructions is loaded in each processor. To provide each node with its subset of global data, the computational domain is partitioned into $n_p = n_i \times n_j \times n_k$ subdomains, where n_p is the total number of processors used. Since there is no mesh adaptation, a static domain decomposition is

only performed at the beginning of each calculation, and it is sufficient to provide for a nearly perfect load balancing. One master processor collects global information such as the maximum and root mean squared residuals, aerodynamic forces on the body, etc. This minor overhead is the only difference in computational load between the processors.

The computation on each processor can not proceed independently of the other nodes since the updates of the variables next to the subdomain boundary require information stored outside of its partition. Communication between the subdomains is achieved by passing a double or a single-level halo of the boundary variables every time they are updated. The stencil of the artificial dissipation model, Equation (10), extends as far as two cells out of the computational subdomain and thus calls for a double-level halo to be passed on the main mesh. A simpler artificial dissipation model is employed on the coarser meshes of the multigrid cycle which only requires passing of a single level halo. The halo's corner variables must be exchanged due to the stencil of the viscous flux discretization employed in the present method (reference [23]). The stencil is illustrated in Figure (2) for the two-dimensional case. This can be accomplished without direct di-

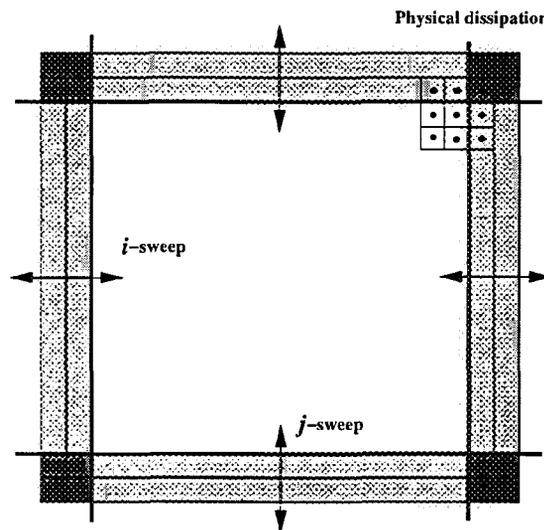


Figure 2: Double Level Halo for the Navier-Stokes Calculations.

agonal communication between the processors by applying single sweeps along each coordinate direction and including the overlapping end values, as shown schematically for 2-D case in Figure 2. On the outer boundaries of the computational domain, void IDs are specified for the missing neighbors and the message passing along these boundaries is disabled. Along the spanwise direction, the computa-

tional domain is considered to be periodically closed on itself which defines the corresponding neighbor ID assignments on the outer spanwise boundaries.

5.2 Nonblocking Communication

The major obstacle in achieving the ideal parallel performance stems from the interprocessor communication costs due to both bandwidth and latency limitations of the available hardware. One of the possible ways of performance improvement is to overlap communication and computation by combining *nonblocking* communication with an appropriate structure of the message passing calls in the program.

In nonblocking communication, two calls are necessary to send a message buffer. The **send start** call returns control to the program immediately after the send has been posted. The program can then proceed with computation, autonomously of the communication call handled by the communication controller. To verify that the outgoing buffer was safely copied out by the controller, an additional **send complete** call is necessary. The nonblocking receive operation consists of the two similar stages, accomplished by correspondingly **receive start** and **receive complete** calls.

Whenever possible, nonblocking communication has been used throughout the flow solver. The following schematic arrangement of the calls has been adopted in the communication subroutine (in two-dimensions), see Figure 2

1. Post the nonblocking receive start calls for all four i, j neighbors.
2. Package the outgoing buffer for the j^- (lower) neighbor.
3. Immediately post a nonblocking send start for the j^- neighbor.
4. Package the outgoing buffer for the j^+ (upper) neighbor.
5. Immediately post a nonblocking send start for the j^+ neighbor.
6. Wait for all pending j -sweep send and receive calls to complete.
7. Retrieve data from the incoming buffers received from j^+, j^- neighbors.
8. Package the outgoing buffer for the i^- (left) neighbor.
9. Immediately post a nonblocking send start for the i^- neighbor.
10. Package the outgoing buffer for the i^+ (right) neighbor.
11. Immediately post a nonblocking send start for the i^+ neighbor.
12. Wait for all pending i -sweep send and receive calls to complete.
13. Retrieve data from the incoming buffers received from i^+, i^- neighbors.
14. Return to the main program.

The very first step of posting the nonblocking receive start calls helps to avoid system buffering and memory-to-memory copying by allowing the information to be provided early on the location of the receive buffer. Note also that the j -sweep completion is enforced before the outgoing i -sweep information can be sent. This ensures the proper transfer of the corner variables, see Figure 2. Finally, the use of the nonblocking send start calls allows to perform the data packaging and retrieving functions concurrently with the communication calls.

5.3 Parallel multigrid

The acceleration techniques, such as multigrid and residual smoothing, which are built in our method rely on global information transfer within the computational domain. Thus, care must be taken to implement these inherently communication intensive techniques on a distributed memory computer. The calculations on the coarser meshes of a multigrid cycle negatively affect parallel performance because of the high task granularity (the ratio of the number of bytes received by a processor to the number of floating point operations it performs). For example, the coarsest three-dimensional mesh may contain as few as 8 cell-centered variables, yet the communication routines have to pass its single level halo with a total of 56 cells. The communication overhead could be reduced by updating flow variables on selected stages of the multistage scheme only, but, as suggested in [2], this leads to the prohibitive decrease in the convergence rate. In the present work, a full multigrid scheme is used, which reproduces exactly the output of the original serial code in the absence of residual smoothing (discussed below). Thus, the boundary variables are exchanged on every mesh in the multigrid cycle and on every stage of the pseudotime stepping scheme. The double-level halo is passed on the finest mesh to provide all necessary values for the computation of the fourth order undivided differences entering the artificial dissipation term. The single-level halo is exchanged on the coarser meshes due to the more compact artificial dissipation model employed at the lower levels of the multigrid cycle.

If the number of processors is increased while the mesh size is kept constant, the progressively finer

subdomain partitioning may reduce the number of multigrid coarsening levels which can be accommodated within each processor's subdomain. The domain decomposition and multigrid scheme may therefore conflict with each other in this limiting case. One possible solution to this problem is to send the coarsest mesh data from each processor to a single "master" node which then will compute the corrections in a serial fashion and redistribute them back to the processors [2]. In the present work, however, the typical subdomain sizes are sufficiently large to allow five multigrid coarsening levels to be accommodated within each subdomain. Consequently, the "master" processor technique has not been implemented in this study.

5.4 Implicit residual averaging

The residual smoothing technique [23] allows a larger pseudo-time step to be used in the pseudo-transient calculation. This technique requires the solution of a tridiagonal system of equations along each coordinate direction of the following form

$$(1 - \epsilon_x \delta_x^2)(1 - \epsilon_y \delta_y^2)(1 - \epsilon_z \delta_z^2) \overline{R^*} = R^*$$

The δ_x^2, δ_y^2 and δ_z^2 represent central second difference operators; the $\epsilon_{x,y,z}$ are constants. Clearly, the resulting augmented residuals $\overline{R^*}$ depend upon R^* on the entire mesh (*global smoothing*). In the shared memory parallel environment, the algorithm can be easily implemented by repartitioning the computational domain parallel to the corresponding sweep direction every time the direction of the inversion sweep is changed. In the distributed memory environment, however, the repartitioning would incur prohibitive communication costs and thus can not be employed. A straightforward coordinate line inversion using the mesh decomposition scheme described in Section 5.1 would cause unacceptable idling times of $n_p \cdot (1 - 1/n_m)$ processors, where m is the direction of a sweep. Special algorithms are available for global inversion of banded matrices on the distributed memory architectures, as summarized in reference [11]. The drawback of these algorithms is that they require approximately 3 times more floating point iterations than the original Thomas algorithm [11] plus the additional communication overhead.

Note that, unlike the tridiagonal inversions used in the time-accurate implicit line-relaxation algorithms, the residual inversion does not have to be performed exactly. It can be safely truncated to comply with the domain decomposition scheme, i.e. confined within each subdomain (*subdomain smoothing*). The truncation only affects the rate of convergence of the pseudo-transient calculations, but has no effect on the resulting converged solution.

The interprocessor boundaries remain fully transparent for the flow and no artificial boundaries are being created. The subdomain smoothing has optimum parallel efficiency and it proved to match the convergence speedup of the global smoothing in the parallel test calculations discussed below. Similar findings have been also reported in reference [1].

6 Results and Discussion

A preliminary validation of the basic algorithm with multigrid has been presented for two-dimensional time-dependent Euler and Navier-Stokes equations in references [5], [6]. Also, time-accurate Euler calculations on unstructured grids have been discussed in [22]. In this work, we extend our algorithm to three dimensions and further validate it versus experimental data available for low Reynolds number wake flows ($Re \leq 350$) around cylindrical bodies. The range of low Reynolds numbers considered in the present work allows to focus on validation of the basic time marching numerical algorithm free of turbulence modelling issues. The slow asymptotics in this flow regime provide a challenging test for numerical methods, since long integration times are necessary to resolve the wake evolution toward a limiting cycle. The computational complexity of time resolved three-dimensional flow simulations is conveyed by the following considerations. Although the implicit formulation used in the present algorithm essentially removes the stability restriction on the time step employed, around 40 time steps are still necessary to accurately resolve each shedding cycle. This may require as many as 4000 time steps to let the flow fully evolve and reveal its asymptotic behavior. Even on a relatively coarse mesh with $96 \times 64 \times 32$ cells, this leads to a problem with as many as $4000 \times 96 \times 64 \times 32 \simeq 800,000,000$ points in the space-time computational domain. The present study demonstrates feasibility of computations of this size using only 6 processors of the IBM SP2 (RS6000/590 wide nodes) and includes an assessment of accuracy, parallel scalability, and overall efficiency of the present computational algorithm.

6.1 Vortex Shedding from a Circular Cylinder

Flow past a circular cylinder is one of the classical problems in fluid dynamics encompassing, at different Reynolds numbers, a wide variety of fascinating fluid phenomena – from parallel vortex shedding and famous von Kármán vortex street ($Re \geq 50$) to an aerodynamic drag crisis ($Re \sim 5 \times 10^5$). The comprehensive experimental and numerical evidence accumulated for this problem makes it one of the most valuable benchmarks for testing and compar-

ing numerical methods. At low Reynolds numbers, a number of vortex shedding phenomena have been accurately captured using our two-dimensional flow solver, as explained in references [5], [7]. These include transition of the flow to the unsteady regime at $47 < Re_c < 48$ and the spanwise independent, so called “parallel” shedding regime [37] at $50 \leq Re \leq 180$. According to reference [37], at Reynolds numbers higher than $Re \sim 180$ the flow becomes unstable to three-dimensional disturbances which makes the two-dimensional approximation inadequate.

Here we present three dimensional results and extend the range of Reynolds numbers considered to $Re \leq 250$. In particular, results of three-dimensional computations, performed at Reynolds numbers 150, 225, and 250, are presented to both validate the multigrid-driven implicit algorithm for flow transition to three-dimensionality and to assess its parallel computational efficiency in three-dimensions.

6.1.1 Relevant Experimental and Computational Data

The unsteady laminar vortex shedding from a circular cylinder at low Reynolds numbers has been studied both experimentally and numerically by a number of authors, and it still remains a topic of active research. A comprehensive experimental data set for these flows can be found in works by Williamson and collaborators [36]–[41]. On the computational side, a detailed numerical study of the vortex shedding at $25 \leq Re \leq 300$ has been performed by Henderson and Barkley by using an unstructured spectral-element method, references [16], [15], [17] and [4].

Experiments conducted by Williamson and Roshko, reference [37], indicate that the flow around a circular cylinder is stationary up to approximately $Re_c = 49$. The majority of experimental results discussed in literature (tabulated in reference [19]) suggest that the transition to vortex shedding occurs in the interval $43 \leq Re_c \leq 49$. The experimental studies reported in references [13], [37] indicate that a spanwise independent *parallel* shedding can be observed in the range of approximately $49 < Re < 180$, provided that special care is taken to control the end effects. Above $Re = 180$, three-dimensional shedding modes are observed. A continuous relation $St(Re)$ of the form

$$St(Re) = \frac{A}{Re} + B + C \cdot Re, \quad (17)$$

where A , B and C are constants, was proposed by Williamson [36] to describe the parallel vortex shedding, which is believed to be accurate within 1%. The relationship (17) is often referred to as the “universal curve”, since, as it was originally proposed by G. Brown and shown in [36], experimental data for

Source	Quantity	A	B	$C \cdot 10^4$
[36]	St	-3.3265	.1816	1.600
[34]	St	-4.50	.212	—
[15]	St	-3.3658	.1831	1.627
[37]	Cpb	-14.3500	.6950	16.920
[15]	Cpb	-10.1090	.6380	20.570

Table 1: The coefficients of the “universal curve” (Equation 17).

the three-dimensional *oblique* shedding also collapse on this curve, provided that the Strouhal number is scaled with the inverse cosine of the shedding angle. Furthermore, a relationship of the same form (17), used with a different set of values for A , B and C , closely approximates the variation of the base pressure coefficient $Cpb(Re)$ with Reynolds number.

The experimental $St(Re)$ data of Williamson [36] are shown in Figure 3 for the range $49 < Re < 250$ by a solid line with the ticks indicating an estimated accuracy of $\pm 1\%$. Several other experimen-

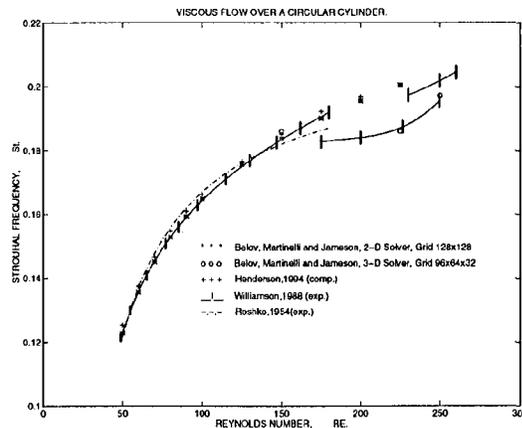


Figure 3: Variation of the Strouhal Frequency with Reynolds Number.

tal and computational results [14], [15], [36], [37] are also presented in Figure 3 with the curve fitting coefficients corresponding to Equation (17) summarized in Table 1. The corresponding variation of the base pressure coefficient with Reynolds number $Cpb(Re)$ is shown in Figure 4 with the experimental errors quoted as being between 1% and 5%, reference [34]. Results of two-dimensional computations of the “universal curve” using our method, reference [7], are shown in Figures 3 and 4 by circles and are in very good agreement with the other experimental and computational data.

According to reference [37], above $Re \sim 180$ the flow becomes three-dimensional and exhibits spanwise vortex loops with a spanwise length scale of

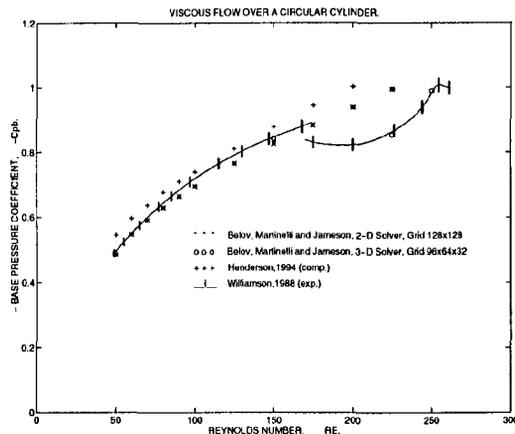


Figure 4: Variation of the Base pressure Coefficient with Reynolds Number.

about 3 diameters (shedding mode “A”). The transition to three-dimensionality is subcritical [17] and is marked by a discontinuous drop in the Strouhal frequency, as shown in Figure 3. The exact reason for this reduction in St is still not fully understood and may be attributed to either primary vortex deformations or to the spot-like “vortex dislocations”, described in references [38], [39]. The second main discontinuity in the $St(Re)$ curve occurs at $Re = 230 - 260$, when the aforementioned vortex loops give way to a finer scale vortex structure (shedding mode “B”), described in detail in references [37], [38]. It is important to note that the double value region of the $St(Re)$ curve corresponding to this transition, Figure 3, is *not hysteretic* but reflects temporal intermittency between the two modes A and B [37].

6.1.2 Three-Dimensional Parallel Computations.

As demonstrated by the results presented in Figure 3 and 4, the two-dimensional computations at Reynolds numbers above 180 fail to capture the transition of the flow to a three-dimensional mode and consequently deviate from the experimental data. In the present work, time-resolved three-dimensional computations of vortex shedding from a circular cylinder at Reynolds numbers 150, 225 and 250 are performed to assess the accuracy and parallel efficiency of the present algorithm in three dimensions.

The computations are performed on an O-H-mesh using $96 \times 64 \times 32$ cells with the far field boundary extended 25 diameters from the center of the cylinder. The aspect ratio of the cylinder is $L/D = 6.4$.

	Grid 128x128	Grid 96x64x32	
	2-D Shedding	Initial 2-D	3-D
Cl	± 0.703	± 0.701	Modulated
Cd_{max}	1.26	1.27	1.20
Cpb	-1.00	-1.02	-0.855
St	0.201	0.205	0.186

Table 2: Comparison of the 2-D and 3-D results, $Re=225$.

When periodic boundary conditions are used to simulate a cylinder of infinite span, the aspect ratio of a cylinder plays an important role by selecting the discrete spectrum of allowed spanwise modes. This can be compared to the nearly continuous spectrum realized in the very high aspect ratio experimental settings. Extra care must be taken in the case of Reynolds numbers close to the transition between 2-D and 3-D flow regimes (in the neighborhood of $Re = 180$) since, according to the spectral Floquet stability analysis [15], [4], the flow is only unstable to a narrow band of spanwise modes. This suggests that certain values of aspect ratio L/D may potentially suppress the modes responsible for the onset of the three-dimensionality and consequently lead to overprediction of Re_c .⁴

The recent study by Henderson [4], addresses this issue in great detail and identifies the range of unstable modes as a function of the Reynolds number. For the computations discussed in this section ($Re = 225, 250$) the modes with wavelengths between approximately 2.8 and 6.5 diameters are expected to fall into the range which is unstable to three dimensional disturbances.

An initial disturbance is introduced in the flow to expedite the onset of the shedding. The disturbance imitates the distributed torsional deformation of a cylinder: spanwise sections of a cylinder are rotated in their planes with a common frequency but with a random number generated amplitude between 0° and 5° degrees and then stopped. The spanwise variation of the base pressure coefficient $d(Cpb) = Cpb_{max} - Cpb_{min}$ is considered as one of the indicators of the flow three-dimensionality. Its time history is shown in Figure 5 for $Re = 150$ and $Re = 225$. The results for $Re = 150$ exhibit a decaying response to the initial disturbance and reach, asymptotically, a two-dimensional shedding, which is in agreement with the experimental predictions. At $Re = 225$ the flow is unstable to the same disturbance: initially the flow develops essentially two-dimensional shedding, and then, the $d(Cpb)$ and the maximum spanwise velocity component gradually increase by two orders in magni-

⁴It is interesting to mention that the transition to the unsteady regime occurs at Reynolds numbers strongly depending on the aspect ratio of a cylinder [29], with the minimum at $L/D \rightarrow \infty$, Section 6.1.1.

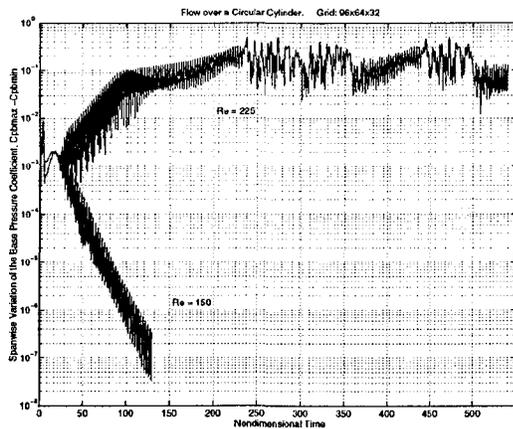


Figure 5: Disturbance Evolution as Measured by the Spanwise Variation of Base Pressure Coefficient $dC_{pb}(t)$. Three-Dimensional Flow around Circular Cylinder, $Re=150$ and $Re=225$.

tude until the fully evolved 3-D shedding results. The time-averaged flow quantities corresponding to the initial parallel shedding regime agree reasonably well with those computed by the two-dimensional solver, as compared in Table 2, given the difference in spatial resolution of the two meshes. The time evolution of the lift coefficient for the $Re = 225$ case is shown in Figure 6. The onset of three-

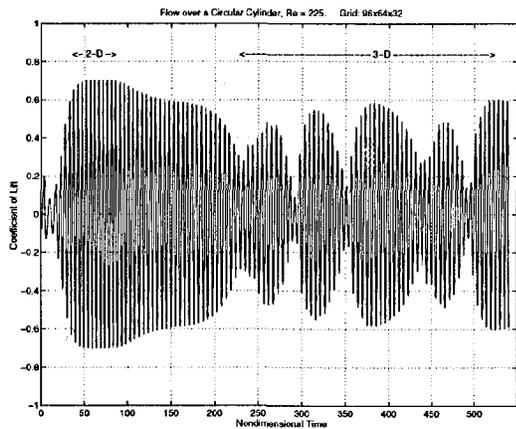


Figure 6: Three-Dimensional Flow over a Circular Cylinder. Evolution of the Lift Coefficient.

dimensionality is marked by a sharp reduction in the lift coefficient, the Strouhal frequency and the base suction ($-C_{pb}$). This is similar to the results of a computation with $Re = 525$, described in reference [3]. A pronounced modulation of the lift coeffi-

cient is observed in Figure 6, presumably indicating the formation of vortices with varying degrees of coherence along the spanwise direction. Also, not only the Cl amplitude, but also the shedding period itself were found to vary with time. This suggests that a single St frequency description is inadequate for this flow regime and possibly explains the double value region in the experimental $St(Re)$ curve [37]. These findings are in line with experimental observations [39] and can be potentially explained by the vortex dislocation mechanism [38]. Nevertheless, higher aspect ratio computations may be necessary to investigate this phenomenon. One can observe two distinctive patterns in time histories of $d(C_{pb})$ and the Cl , superimposed in Figure 7. The corre-

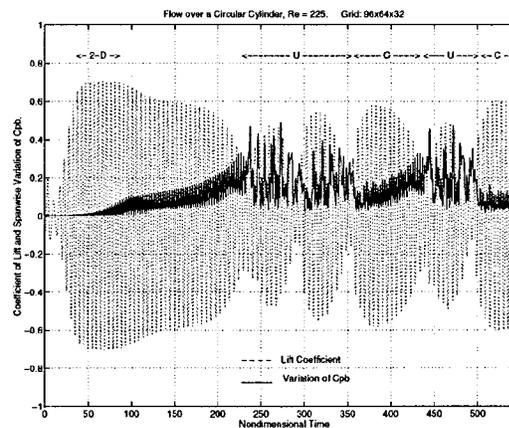


Figure 7: Superimposed Time Evolution of Cl and $d(C_{pb})$, $Re=225$.

lation of $d(C_{pb})$ with Cl seems to suggest an intermittency between two distinctive shedding modes. A “coherent” C-mode, defined by high Cl and regular small-amplitude $d(C_{pb})$ variations, corresponds to essentially two-dimensional flow behavior next to the cylinder. By contrast, low Cl and sporadic large-amplitude excursions of $d(C_{pb})$ (“uncorrelated” U-mode) reveal a highly three-dimensional vortex formation region next to the cylinder’s surface. Thus, the region where transition to three-dimensionality occurs migrates continuously up- and downstream between the cylinder and the near wake. This finding may explain the Cl modulation, the accompanying temporal variation of the shedding period and the double value region of the $St(Re)$ curve observed in the experiments.

Time-averaged results for the three-dimensional shedding mode for Reynolds numbers 225 and 250 are shown in Figures 3 and 4 with circles. Excellent agreement is observed with the lower branches of the experimental data curves corresponding to the shedding mode “A” as defined in references [37], [38].

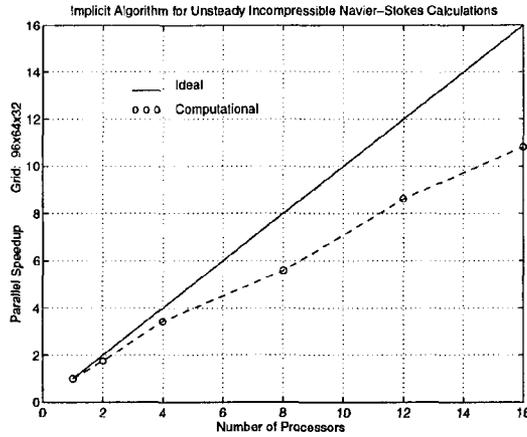


Figure 8: Parallel Speedup in Three Dimensional Calculations.

According to this source, the “A” mode of vortex shedding is characterized by the presence of the streamwise vortex loops with a spanwise wavelength of about 3 diameters. Figures 9 and 10 show that similar structures are observed in our computation with the wavelength equal to 3.2 diameters at Reynolds number 250. Note that this wavelength is indeed within the unstable spectrum reported in reference [4]. The bending and eventual rupture of the spanwise vortices and their interconnecting *ribs*, observed in the computation, appear to be similar to the vortical structures described in literature [4], [3]. Further analysis of the flow regime using a finer mesh and higher aspect ratio cylinder is necessary to establish the grid independence and to investigate the vorticity transport mechanism in more detail, which is one of the directions for future work.

The fully developed three dimensional flow regime was resolved using 40 time steps per shedding cycle with 25 five-level W-multigrid cycles per time step. This required less than 3 min. per time step using 6 wide nodes (RS6000/590) of IBM SP2. This corresponds to $\sim 37 \mu\text{sec}$ per time step per mesh point or ~ 24 min. per convective time scale. The parallel speedup achieved by our algorithm on the relatively coarse mesh used in the present work is shown in Figure 8. Calculations on more spatially resolved meshes (1 to 5 million points) result in lower task granularity and can be performed efficiently on parallel clusters with large numbers of processors (64 and over). Provided that access to the adequate computational resources is available, the present algorithm can enable time resolved three-dimensional computations on practical engineering meshes in less than an hour per convective time scale.

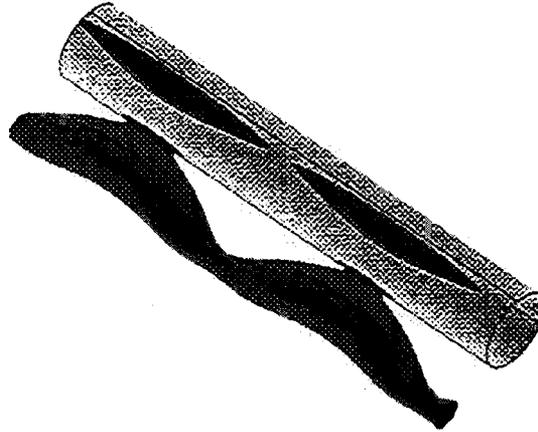


Figure 9: Isosurface of Pressure for the Three-Dimensional Flow around a Circular Cylinder, $Re=250$.

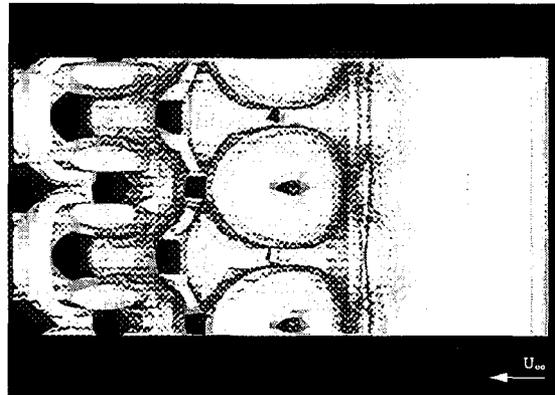


Figure 10: Isosurface of Vorticity Magnitude (Top View), $Re=250$.

6.2 Flow over a Circular Half-Cylinder

The discussion above was mainly focused on the flow characteristics which refer to the body, such as the base pressure and its variation along the span, the lift coefficient evolution, and the Strouhal frequency. On the other hand, it is also important to determine how accurately the wake structure and temporal behavior are captured by the present algorithm. In order to address these issues in more detail, we take advantage of the experimental data collected by I. Huang and G. Brown [18] for the flow over a circular half-cylinder at $Re = 350$. This section presents validation of the present method with an emphasis on the flow parameters characterizing the vortex wake region.

The computations were performed on an O-H-mesh using $96 \times 144 \times 32$ cells with the far field boundary extended 30 diameters from the center of the rear facing part of the semi cylinder. The aspect ratio of a semi-cylinder is $L/D = 6.4$ and spanwise periodicity is assumed. The computational geometry is illustrated in Figure 11. An initial distur-

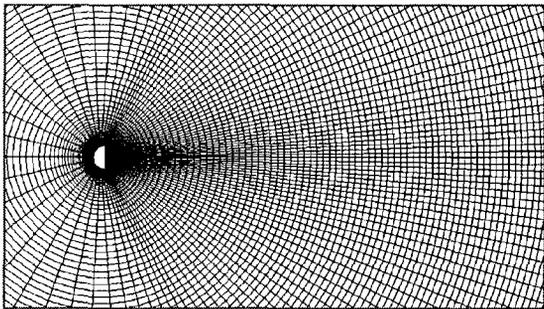


Figure 11: Detail of the O-H Mesh (Side View).

bance is introduced to expedite the onset of the vortex shedding in the same way as explained in the previous section. To further speed up the computation and save CPU time, the solution is advanced in time with time step $\Delta t = 5/24$ (approximately 24 time steps per shedding cycle) until nondimensional time $t = 150$ is reached. At this time level, the flow is considered to be fully developed, statistic data accumulation and time averaging are initialized, and a halved time step is used $\Delta t = 5/48$.

In Figure 12, the computed time averaged centerline streamwise velocity \bar{U} is compared with the time averaged hot wire measurement data [18] available at locations from 2 to 20 diameters downstream of the semi-cylinder. It can be seen from Figure 12, that the computational results are in excellent agreement with the experimental data at all locations except for the first (2 diameters downstream) and the last one (20 diameters downstream). The former

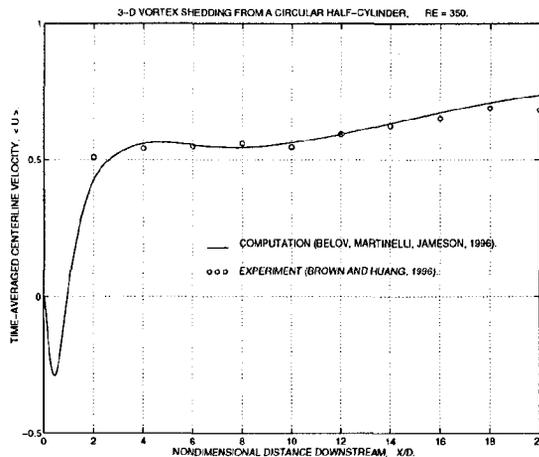


Figure 12: Time Averaged Centerline Streamwise Velocity, \bar{U} .

deviation can be explained by the fact that the hot wire measurements register the magnitude of the velocity and thus refer to the time averaged *absolute value* rather than the true average. The difference between the two averages may therefore be caused by periodic flow reversals observed in the computation in the near field of the half-cylinder ($x/D < 4$) including at the location $x/D = 2$. In the far field, the momentum defect profiles can be expected to flatten out due to diffusion and mixing effects as the downstream distance increases, thus causing the centerline value \bar{U} to increase monotonically with x . The reduction in the experimentally observed value of \bar{U} at $x/D = 20$ is thus counterintuitive and may simply indicate the experimental error bounds.

Important information regarding the accuracy of the wake flow simulations can be discerned from the computed velocity defect profiles, which are superimposed with the experimental ones in Figures 14. The terracing of the computational data is due to the piecewise-constant interpolation used to project the data from the curvilinear mesh to the probe locations. As it can be seen from Figures 14, the spread and magnitude of the profiles are in close agreement with the experimental data for the most part of the flow, except at the locations already identified above $x/D = 2$ and $x/D = 20$. A minor overshoot (less than 2%) consistently observed in the computed \bar{U} profiles may be attributed to blockage effects. These can be estimated by the ratio of the radius of the half-cylinder $r = 0.5$ to the distance to the far field $R = 30$ which is indeed close to the two percent. The accurate capture of the rate of spread of the velocity defect profiles, exhibited in our computations, is particularly informative and suggests that the wake is dominated by physical rather than by numerical

diffusion. The time averaged profiles of the streamwise velocity fluctuation are plotted in Figures 15. The computational results are derived from the time averages of the streamwise velocity and the velocity squared, taking advantage of the following relationship

$$\overline{U^2} \equiv \overline{(\overline{U} + u')^2} = \overline{U}^2 + \overline{u'u'}, \quad (18)$$

where of course $\overline{2\overline{U}u'} \equiv 0$. The computational results exhibit a higher rate of decay of the fluctuation energy (as a function of the downstream distance) than the experimental data in the far field region $x/D > 12$. This suggests that, as the mesh becomes coarser, the streamwise velocity fluctuations are affected by the numerical diffusion.

A time series of the x -component of the velocity has been computed at the location $(x/D, y/D) = (1.0, 0.5)$. A power spectral density plot of this data is compared in Figure 13 with the experimental spectrum, taken at the same location. One can see from

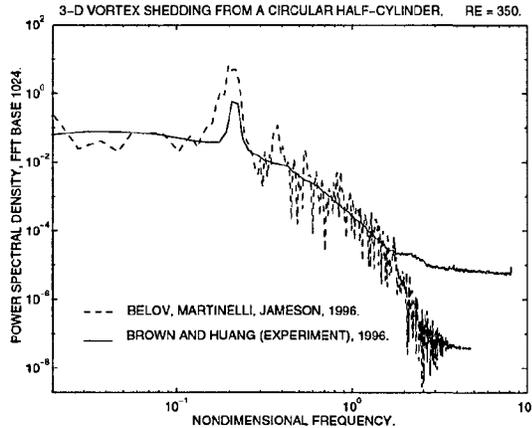


Figure 13: Power Spectrum of Fluctuating Streamwise Velocity at $(x/D, y/D) = (1.0, 0.5)$.

the Figure 13 that the Strouhal frequency is accurately captured in the computation and both low and high frequency regions are in good agreement with the experimental data, up to a nondimensional frequency of approximately $10 \times St$. Beyond this threshold, the experimental curve flattens out, perhaps due to high frequency noise. Another possible explanation of the deviation is that the high temporal frequency modes are effectively damped out in the computation thus indicating the cutoff frequency of the numerical dissipation model employed in this work.⁵ In any case, it should be noted that the high frequency modes have an energy content three to five orders of magnitude smaller than the

⁵The small time scales primarily correspond to the high spatial frequency modes which are filtered out by the artificial dissipation term 11.

energy bearing modes in the lower part of the spectrum and thus can be expected to have a negligible effect on the overall solution. This observation leads to the conclusion that a wide range of the important scales is adequately resolved in our computation. However, the computed spectrum exhibits stronger correlation at the Strouhal frequency as compared to the experiment. This is consistent with a slight overprediction in the computed streamwise velocity fluctuations $\overline{u'u'}$ in the near field of the half-cylinder and suggests a reduced rate of the three-dimensional energy transfer to the spanwise $\overline{w'w'}$ fluctuations because of the relatively low aspect ratio computational settings ($L/D = 6.4$). To address this issue, a study of the aspect ratio and mesh resolution effects on the computed solution is planned and will be presented in a future work.

7 Conclusions

The results of the test calculations performed so far indicate that our solver is both accurate and very efficient for three dimensional unsteady incompressible flows. The key to the algorithm's efficiency is the implicit backward time discretization coupled with the multigrid acceleration technique, which enables fast convergence of the subiteration to be achieved at each time step. A point-implicit formulation in pseudo-time allows improvement of efficiency even further by decoupling the pseudo-time step from the physical time step. The multigrid-driven algorithm is highly suitable for parallel architectures as demonstrated by the speedup obtained on up to 16 processors of an IBM SP2 computer, presented in Figure 8. Preliminary computations performed on a mesh with over one million cells required about $6 \times 10^{-6} \text{ sec}$ per mesh point per subiteration on 32 wide nodes (RS6000/590) of IBM SP2. This suggests that time-resolved computations over geometries of engineering complexity can be obtained in less than one hour per convective time-scale.

Based on the analysis of the computational results presented here, we conclude that the proposed scheme provides an accurate and efficient way of computing time-dependent incompressible flows.

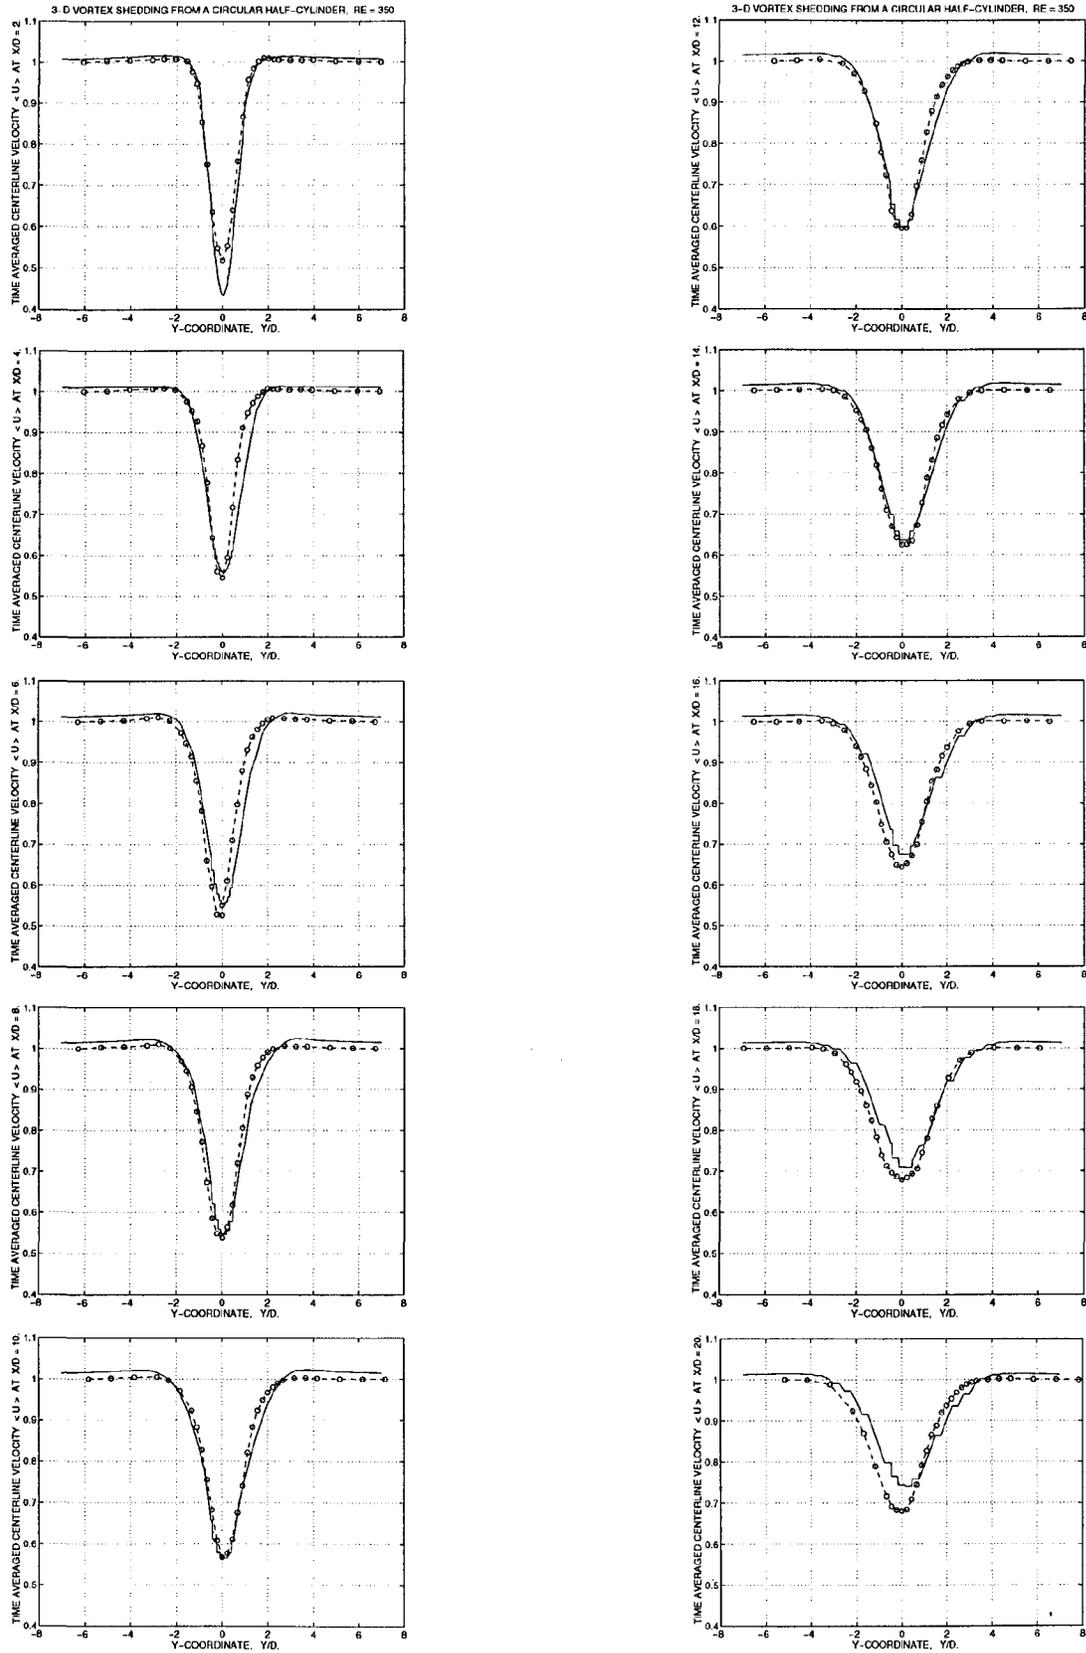


Figure 14: Time Averaged Streamwise Velocity Profiles, \bar{U} , at $x/D = 2, 4, \dots, 18, 20$. 'ooo' - Experiment (Brown and Huang, 1996), '—' Computation (Belov, Martinelli, and Jameson, 1996).

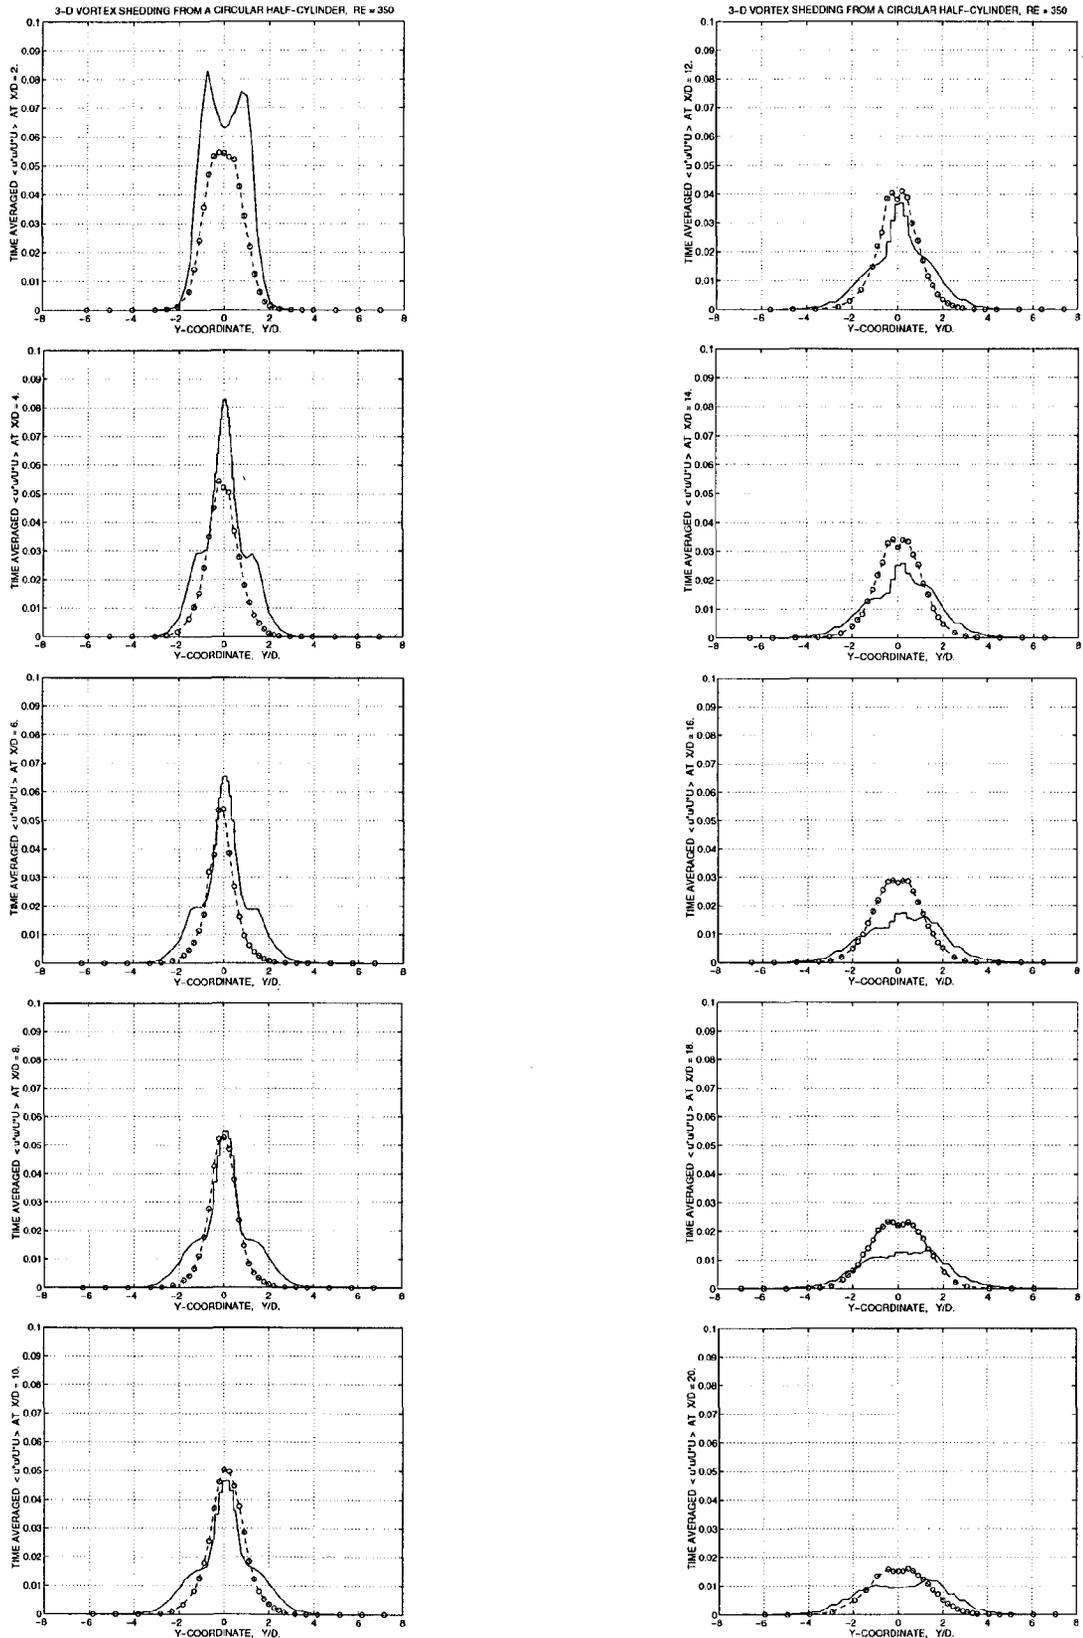


Figure 15: Time Averaged Streamwise Velocity Fluctuations, $\overline{u'u'}$, at $x/D = 2, 4, \dots, 20$. 'ooo' - Experiment (Brown and Huang, 1996), '—' Computation (Belov, Martinelli, and Jameson, 1996).

Acknowledgement

We are grateful to I. Huang and G. Brown for providing us with their experimental data.

Support for this work was provided through by the Office of Naval Research through Grant N00014-93-I-0079, under the supervision of Dr. E.P. Rood.

References

- [1] J. J. Alonso. *Parallel Computation of Unsteady and Aeroelastic Flows Using an Implicit Multigrid-Driven Algorithm*. PhD thesis, Princeton University, 1997.
- [2] J. J. Alonso, T. J. Mitty, L. Martinelli, and A. Jameson. A two-dimensional multigrid Navier-Stokes solver for multiprocessor architectures. In Satofuka, Periaux, and Ecer, editors, *Parallel Computational Fluid Dynamics: New Algorithms and Applications*, pages 435–442, Kyoto, Japan, May 1994. Elsevier Science B. V. Parallel CFD '94 Conference.
- [3] S. Balachandar and R. Mittal. Momentum and vorticity transport by 3-d vortical structures in bluff body wakes. In *The 10th Symposium on Turbulent Shear Flows*, volume 1 of 4, pages pp. 4–7. The Pennsylvania State University, 1995.
- [4] D. Barkley and R. D. Henderson. Three-dimensional floquet instability analysis of the wake of a circular cylinder. *J. Fluid Mech.*, 322:pp. 215–241, 1996.
- [5] A. Belov, L. Martinelli, and A. Jameson. A new implicit algorithm with multigrid for unsteady incompressible flow calculations. *AIAA paper 95-0049*, AIAA 33rd Aerospace Sciences Meeting and Exhibit, Reno, NV, January 1995.
- [6] A. Belov, L. Martinelli, and A. Jameson. A novel fully implicit multigrid driven algorithm for unsteady incompressible flow calculations. In *Computational Fluid Dynamics '94*, ECCOMAS, pages pp. 663–670, Stuttgart, Germany, September 1994. Proc. of the Second European C.F.D. Conf., John Wiley & Sons.
- [7] A. A. Belov. *A New Implicit Multigrid-Driven Algorithm for Unsteady Incompressible Flow Calculations on Parallel Computers*. PhD thesis, Princeton University, 1997.
- [8] A. J. Chorin. A numerical method for solving incompressible viscous flow problems. *J. Comp. Phys.*, 2:pp. 12–26, 1967.
- [9] G. Dahlquist. A special stability problem for linear multistep methods. *BIT*, 3:pp. 27–43, 1963.
- [10] J. Dreyer. Finite volume solutions to the unsteady incompressible euler equations on unstructured triangular meshes. Master's thesis, Princeton University, 1990.
- [11] B. Galanti and M. Wolfshtein. A multi-domain tri-diagonal solver. TAE Report 753, Technion-Israel Institute of Technology, 1995.
- [12] C. W. Gear. *Numerical Initial Value Problems in Ordinary Differential Equations*. Prentice Hall, Englewood Cliffs, 1971.
- [13] M. Hammache and M. Gharib. An experimental study of the parallel and oblique vortex shedding from circular cylinders. *J. Fluid Mech.*, 232:pp. 567–590, 1991.
- [14] R. D. Henderson. Private communication.
- [15] R. D. Henderson. *Unstructured Spectral Element Methods: Parallel Algorithms and Simulations*. PhD thesis, Princeton University, 1994.
- [16] R. D. Henderson. Details of the drag curve near the onset of vortex shedding. *Phys. Fluids*, 7(9):pp. 2102–2104, September 1995.
- [17] R. D. Henderson and D. Barkley. Secondary instability in the wake of a circular cylinder. *Phys. Fluids*, 8(6):pp. 1683–1685, 1996.
- [18] I. Huang and G. Brown. Princeton University, 1996. Private communication.
- [19] C. P. Jackson. A finite-element study of the onset of vortex shedding in flow past variously shaped bodies. *J. Fluid Mech.*, 182:pp. 23–45, 1987.
- [20] A. Jameson. Multigrid algorithms for compressible flow calculations. In *Second European Conference on Multigrid Methods*, Cologne, October 1985. Princeton University Report MAE 1743.
- [21] A. Jameson. Time dependent calculations using multigrid, with applications to unsteady flows past airfoils and wings. *AIAA paper 91-1596*, AIAA 10th Computational Fluid Dynamics Conference, Honolulu, HI, June 1991.
- [22] P. T. Lin. Implicit time dependent calculations for compressible and incompressible flows on unstructured meshes. Master's thesis, Princeton University, Princeton, NJ, November 1994.
- [23] L. Martinelli. *Calculations of Viscous Flow with a Multigrid Method*. PhD thesis, Princeton University, Princeton, NJ, October 1987.

- [24] N. D. Melson, M. D. Sanetrik, and H. L. Atkins. Time-accurate Navier-Stokes calculations with multigrid acceleration. Copper Mountain, CO, July 1993. Proceedings of the Sixth Copper Mountain Conference on Multigrid Methods.
- [25] T. Miyake, Y. Sakamoto, H. Tokunaga, and N. Satofuka. Numerical solution of incompressible flow using two-step, one-stage runge-kutta time integration scheme. Brussels, Belgium, September 1992. 1st European Comput. Fluid Dynamics Conf.
- [26] H. Persillon, M. Braza, H. HaMinh, and C.H.K. Williamson. Prediction of the three-dimensional transition in the flow past a circular cylinder at low reynolds numbers. In *The 10th Symposium on Turbulent Shear Flows*, volume 1 of 1, pages pp. 37-42. The Pennsylvania State University, 1995.
- [27] R. J. Peyret. Unsteady evolution of a horizontal jet in a stratified fluid. *J. Fluid Mech.*, 78:49-63, 1976.
- [28] A. Prasad and C.H.K. Williamson. Three-dimensional effects in turbulent bluff body wakes. In *The 10th Symposium on Turbulent Shear Flows*, volume 1 of 6, pages pp. 1-6. The Pennsylvania State University, 1995.
- [29] M. Provansal, C. Mathis, and L. Boyer. Bénard-von kármán instability: Transient and forced regimes. *J. Fluid Mech.*, 182:pp. 1-22, 1987.
- [30] A. Rizzi. Numerical implementations of solid-body boundary conditions for the Euler equations. *ZAMM*, 58:301-304, 1978.
- [31] A. Rizzi and Lars-Erik Eriksson. Computation of inviscid incompressible flow with rotation. *J. Fluid Mech.*, 153(pp. 275-312), 1985.
- [32] S. E. Rogers, J. L. C. Chang, and D. Kwak. A diagonal algorithm for the method of pseudocompressibility. *J. Comp. Phys.*, 73(2):pp. 364-379, December 1987.
- [33] S. E. Rogers and D. Kwak. Upwind differencing scheme for the time-accurate incompressible navier-stokes equations. *AIAA Journal*, 28(2):pp. 253-262, February 1990.
- [34] A. Roshko. On the development of turbulent wakes from vortex streets. NACA Report 1191, National Advisory Committee for Aeronautics, 1954.
- [35] L. K. Taylor and D. L. Whitfield. Unsteady three-dimensional incompressible navier-stokes solver for stationary and dynamic grids. AIAA Paper 91-1650, Reno, NV, 1991.
- [36] C. H. K. Williamson. Defining a universal and contiguous Strouhal-Reynolds number relationship of the laminar vortex shedding of a circular cylinder. *Physics of Fluids*, 31:2747-2747, October 1988.
- [37] C. H. K. Williamson. Measurements of base pressure in the wake of a cylinder at low reynolds numbers. *Z.Flugwiss. Weltraumforsch*, 14:pp.38-46, 1990.
- [38] C. H. K. Williamson. The natural and forced formation of spot-like "vortex dislocations" in the transition of a wake. *J. Fluid Mech.*, 243:pp. 393-441, 1992.
- [39] C. H. K. Williamson. American Physical Society Meeting, November 1996. Private communication.
- [40] C. H. K. Williamson. Mode 'a' secondary instability in wake transition. *Phys. Fluids*, 8:pp. 1680-1682, 1996.
- [41] C. H. K. Williamson. Vortex dynamics in the cylinder wake. *Ann. Rev. Fluid Mech.*, 28:pp. 477-539, 1996.

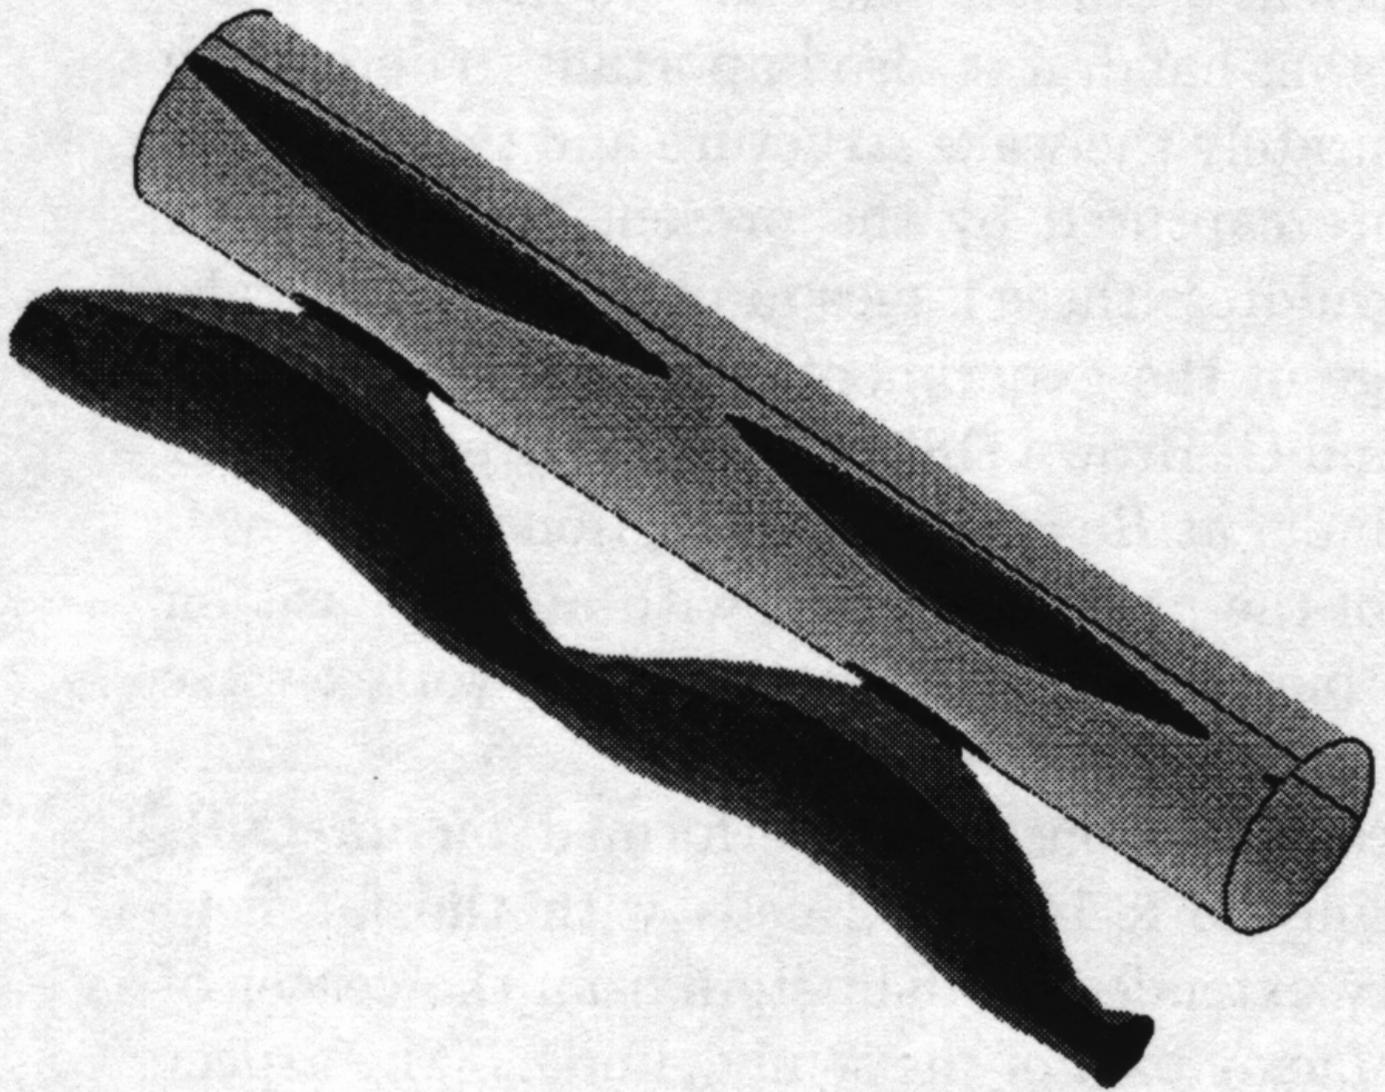


Figure 9: Isosurface of Pressure for the Three-Dimensional Flow around a Circular Cylinder, $Re=250$.

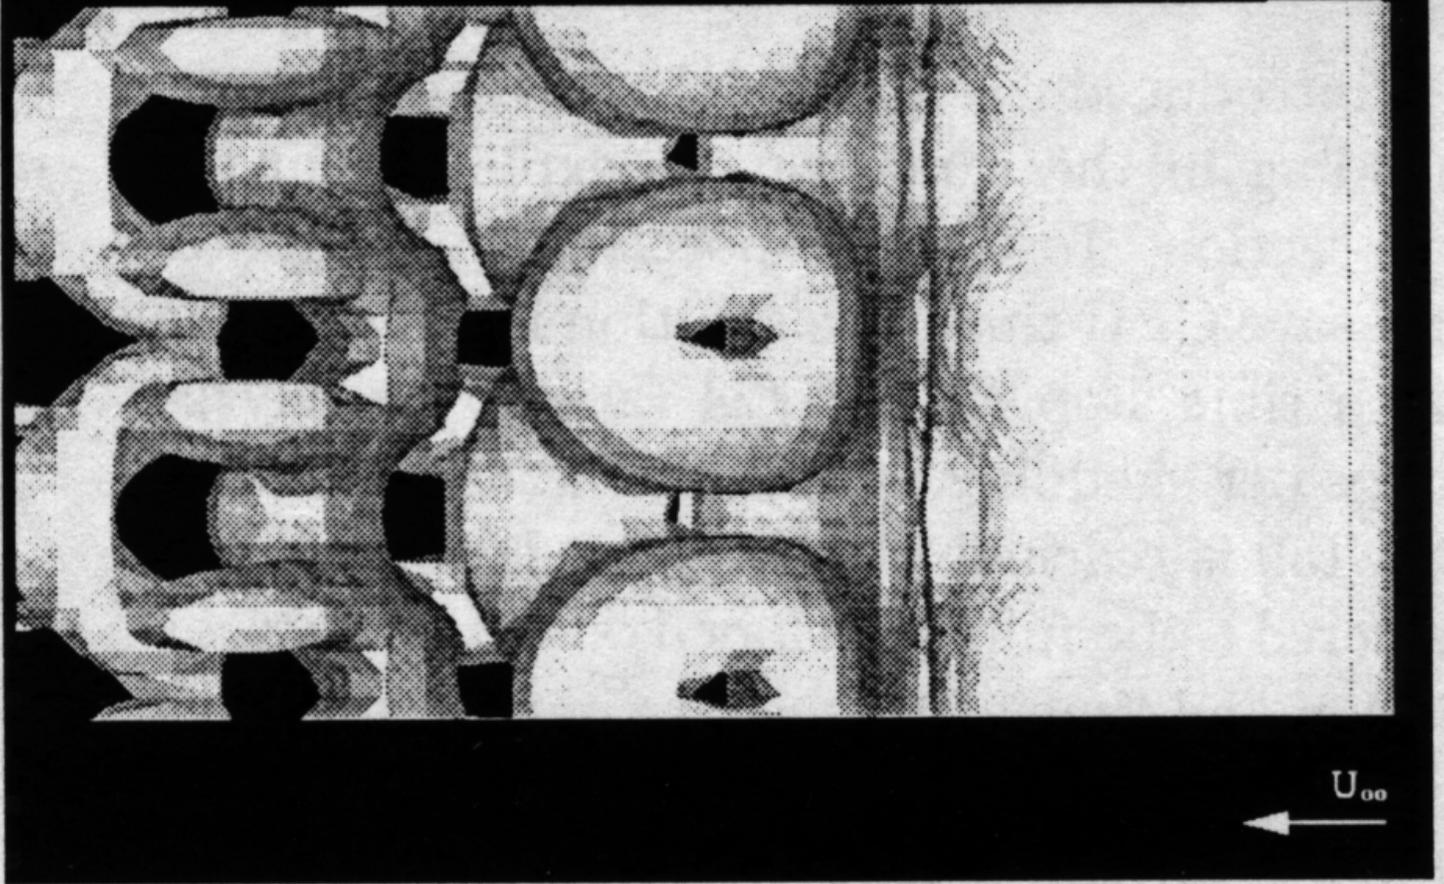


Figure 10: Isosurface of Vorticity Magnitude (Top View), $Re=250$.