NoSync Implementation

The two preceding parallel implementations of the FI technique required a synchronization of processors between consecutive time integration steps, (n) and $(n+1)$. Elimination of this synchronization should increase the speedup, but requires very careful restructuring of the different computational phases and appropriate means of accessing them. An algorithm, hereafter referred to as NoSync, was successfully designed with no synchronization. As shown in Fig. 2b, the flow of execution for this four-phase implementation begins with the calculation of the right-hand side D_{ij} and the recursive coefficients α_{ij} and β_{ij}, from bottom left (en1) to a location (ex1) set at an intermediate i value, $I_{\text{int}} < I_{\text{limit}} = NX - P$. Then, available processors start from $i = I_{\text{int}} + 1$ and $j = 2$ (en2), calculate D_{ij}, α_{ij}, and β_{ij} in the remaining forward x sweep, and perform the backward x sweep on the same row. Phase 2 terminates at the top left corner (ex2). Free processors begin the forward y sweep from bottom left (en3) up to an intermediate j value, $J_{\text{Jint}} < J_{\text{limit}} = NY - P$, and complete the third phase at the right end of J_{int} (ex3). In the fourth and last phase, unassigned processors start at the left end of $J_{\text{int}} + 1$ with the remainder of the forward y sweep, as well as the entire y back substitution, ending at the bottom right corner of the computational domain (ex4).

First, the load distribution in the NoSync implementation increased the work load in phase 1 to more equitably balance the loads between phases 4 and 1. Of course, care must be taken not to alter the load balance between phases 1 and 2 to the point that the race condition is only migrated from one phase interface to another rather than diminishing the possibility of its occurrence. Second, keeping $I_{\text{int}} < I_{\text{limit}} = (NX - P)$ ensures that processors in phase 1 do not interfere with any processor working in phase 4. Furthermore, choosing J_{limit} such that $P - 1$ rows are inaccessible to processors in phase 3 eliminates the possibility of a race condition between phases 2 and 3. Finally, it becomes possible to move the starting point of phase 3 to the bottom left corner and to allow phase 4 to exit on the right side of the computational domain as opposed to the left side in the four corner implementation. As a result, at the next time level the forward x sweep can begin more safely at the left bottom corner without the need for synchronization. Since the NoSync implementation involves four phases, processor interference is possible at any one of the four phase interfaces. By setting phase boundaries at $I_{\text{int}} = NX - P$ and $J_{\text{int}} = NY - P$, the occurrence of processor interference is avoided at the interfaces between phases 4 and 1 and between phases 2 and 3. Proper load balancing and efficient processor routing removed the possibility of processor interference along the other two interfaces.

Parallelization of the Mapping Equations

One may ask how will these implementations perform with more complex equations such as those solved in elliptic grid generation[3]? In this work, the time-dependent version of the two familiar, quasilinear, two-dimensional, curvilinear mapping equations[4] was solved. The solution at every time level represented a larger computational task than that of the heat equation and, as expected, yielded higher speedups. A comparison of resulting speedups from *C$DO* implementations of the heat and mapping equations can be made by examining Fig. 1. The NoSync implementation looks promising for large size problems running with a large number of processors. A continued advantage for the NoSync implementation on computers equipped with more processors is expected, therefore, when solving more complex problems on higher resolution grids (as is typically the case in CFD).

Conclusions

This work was aimed at investigating the suitability of parallel processing for the class of alternating direction implicit techniques on a tightly coupled, shared-memory, parallel architecture. One such technique, that of Beam and Warming, was parallelized on a Sequent Symmetry S81 parallel computer. A significant impediment in achieving near-theoretical speedups with implicit techniques is the requirement for processor synchronization at one or more stages during a time integration. The need for synchronization was eliminated by the careful use of load balancing, phase splitting, and processor routing. This in turn enabled the factored implicit technique to achieve higher speedups than possible with standard parallel compiler implementations. The highest speedups resulted from more complex equations, larger size grids, and more stringent convergence criteria. The principles presented in this work are general and should be applicable to other implicit techniques with minor technique-dependent alterations.

Acknowledgment

Funding for this work by the Department of Defense under DARPA Project 5-25085 of Contract MDA-093-86-C-0182 is gratefully acknowledged.

References

[1]Beam, R. M., and Warming, R. F., "An Implicit Factored Scheme for the Compressible Navier-Stokes Equations," *AIAA Journal*, Vol. 16, No. 4, 1978, pp. 393–402.

[2]Varghese, A. N., "Surface Texture Effects on Ultra-Thin Finite-Width Gas Bearings," Ph.D. Thesis, Civil and Mechanical Engineering Dept., Southern Methodist Univ., Dallas, TX, May 1991.

[3]Thompson, J. F., "A General Three-Dimensional Elliptic Grid Generation System on a Composite Block Structure," *Computer Methods in Applied Mechanics and Engineering*, Vol. 64, Nos. 1-3, 1987, pp. 377–411.

[4]Raad, P. E., and White, J. W., "A 'Transient' Automated Mapping Procedure for Complex Geometries," *Numerical Grid Generation in Computational Fluid Mechanics*, edited by S. Sengupta, J. Häuser, P. R. Eiseman, and J. F. Thompson, Pineridge Press Limited, Swansea, United Kingdom, 1988, pp. 237–245.

Effects of a Rear Stagnation Jet on the Wake Behind a Cylinder

R. Duke,* B. Shrader,* and J. Mo†
Memphis State University, Memphis, Tennessee 38152

Introduction

THE control of the wake flow behind bluff bodies such as cylinders and the subsequent improvement of the aerodynamic performance of the cylinder are interesting subjects for both fundamental and applied research because they have practical applications and illustrate the basics of modern fluid mechanics. It is known that when the flow Reynolds number based on the cylinder diameter exceeds some critical value, the cylinder will experience an oscillating lateral force. This oscillating force is caused by the asymmetric flow pattern in the downstream wake, which is known as the von Kármán vortex street. Because of the massively separated wake, the cylinder also experiences a considerable drag. Experimental evidence shows that if the flow Reynolds number is over 100, the averaged drag coefficient on the cylinder is about 1.0 (Ref. 1). Of the two parts of drag, pressure or form drag is directly related to the wake behavior, whereas the friction drag is very small compared to the form drag. In fact, for a cylinder, the

Received Oct. 17, 1992; revision received Feb. 24, 1993; accepted for publication March 9, 1993.

*Graduate Assistant, Department of Mechanical Engineering. Student Member AIAA.

†Assistant Professor, Department of Mechanical Engineering. Member AIAA.

form drag is about two orders of magnitude higher than the friction drag. Therefore, fundamental research on the behavior of the wake is of tremendous significance.

Flow Visualization Facilities and Setup

The experimental visualization of the flow past a circular cylinder was performed in a recirculating open water channel located at Memphis State University. The flows were visualized at a variety of flow Reynolds numbers based on the cylinder diameter. Figure 1 is a schematic drawing of the open channel facility. The channel is equipped with two pumps, a sluice gate, a flow straightener, and a tail gate; it also has the capability of variable slope. The water flow rate was controlled by using a single pump and adjusting the orifice on the pump. To generate laminar flow, an additional flow straightener was located far upstream that consisted of a 2-mm-thick aluminum sheet with 3-mm-diam holes equidistantly spaced at 3-mm intervals over the entire sheet.[2]

Figure 2 is a sketch of the apparatus used for the experiment. The test cylinder was made from ¾ schedule 40 PVC. As shown here, the cylinder was modified by machining a uniform narrow slit in the cylinder wall. To ensure two-dimensional flow, the slit was 4-cylinder-diameters long. The two ends of the pipe were sealed with permanent plugs and a hole in the top plug was internally threaded to accommodate a pressure line from the pressurized water reservoir. The pressure of the reservoir was continuously controlled by using an air pressure regulator which was connected to a high-pressure line with a maximum pressure of 90 psi. A rotameter was installed and employed to measure the jet volumetric flow rate.

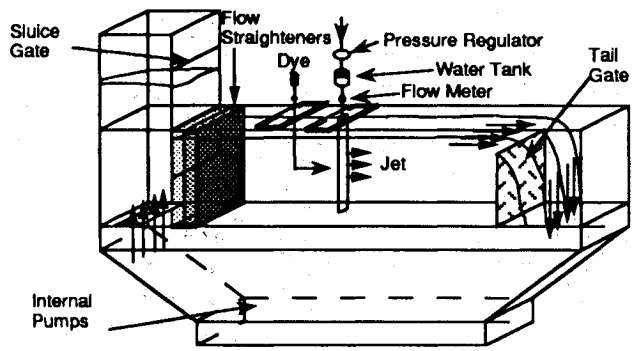

Fig. 1 Schematic of open channel apparatus.

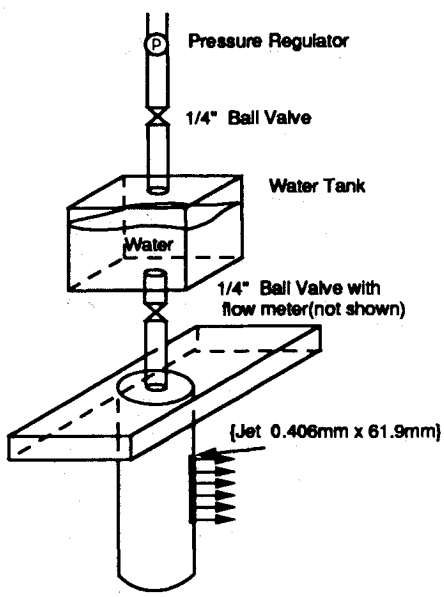

Fig. 2 Experimental setup.

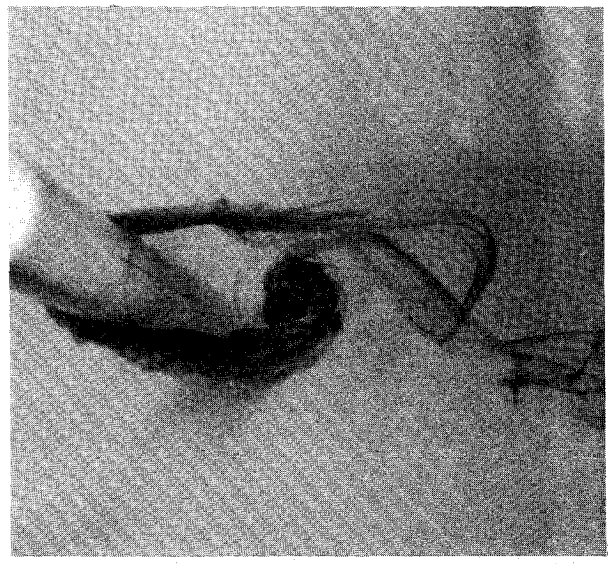

Fig. 3 Oscillating wake flow past a cylinder with $Re = 700$.

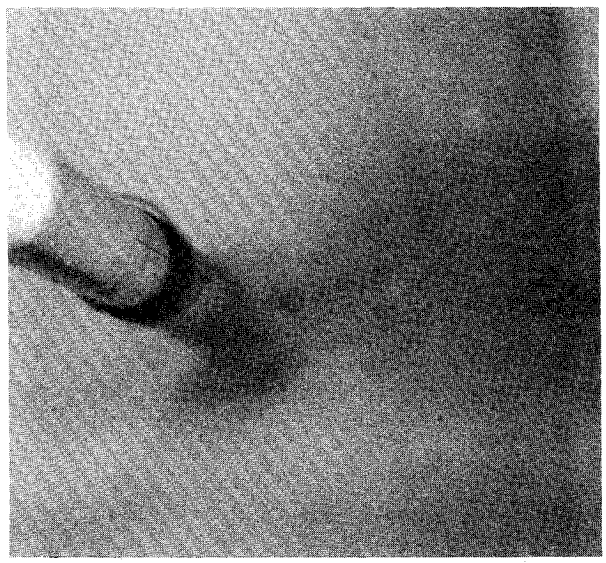

Fig. 4 Fully attached flow over a cylinder with $V^* = 25$ and $Re = 700$.

Table 1 Summary of experimental data

U_∞, cm/s	V_j, cm/s	Re	V^*_{min}
3.18	100.32	909	31.55
5.29	125.40	1511	23.70
6.35	163.00	1815	25.67
8.36	188.08	2389	22.50

Results and Discussion

In the present study, the flow Reynolds number based on cylinder diameter was varied from around 600 to 2500. Over this range of Reynolds numbers, laminar flow was visually confirmed. Figure 3 shows a conventional von Kármán vortex street behind a stationary circular cylinder and provides visual evidence of a laminar flow pattern with a Reynolds number of 700. Figure 4 shows a flowfield with the same upstream flow conditions as shown in Fig. 3, but with the rear stagnation jet operating at a velocity 25 times the freestream velocity. It can be seen that the wake flow behind the cylinder was significantly changed, compared to Fig. 3, and appears fully attached. It was also observed that once the jet velocity exceeded 25 times the freestream velocity, there were no noticeable changes in the wake flow pattern. In Fig. 4, the dye downstream of the cylinder cannot be seen because the high-speed

jet entrains the dye and lowers its concentration. To obtain comparable data for different upstream flow conditions, a new parameter V^* was introduced and defined as the ratio of the jet velocity to the freestream velocity. Table 1 shows a summary of the experimental data collected, in which V^*_{min} is the minimum value of V^* that produced attached flow at the given flow Reynolds numbers.

Conclusions

We have conducted a flow visualization of the flow past a cylinder with and without a rear stagnation jet. It has been found that the massively separated wake behind the cylinder can be completely eliminated and the flow around the cylinder made to behave as inviscid flow by a rear stagnation jet with a minimum dimensionless parameter V^* of 25. Increasing V^* over this critical value showed no further noticeable changes on the flow pattern. The details of the flowfield are under a comprehensive numerical investigation by the authors.

Acknowledgment

The support of the Mechanical Engineering Department at Memphis State University for this research work is greatly appreciated.

References

[1]Hoerner, S. F., *Fluid-Dynamic Drag*, published by the Author, 1965, pp. 3-15.

[2]Shrader, B., and Duke, R., "An Experimental Investigation on the Control of Separated Flow over Cylinders by Means of a Rear Stagnation Jet," AIAA Paper 92-4039, July 1992.

Multigrid Techniques for Hypersonic Viscous Flows

F. Grasso* and M. Marini†
Università di Roma "La Sapienza," Rome 00184, Italy

Introduction

VISCOUS hypersonic flows are dominated by strong shock-wave/boundary-layer interactions that affect the features of the flowfield and may cause boundary-layer separation and laminar-turbulent transition. For an accurate prediction of these complex phenomena, the numerical solution of such flows requires very fine grids, with a loss in the computational efficiency.

In this context, the multigrid technique is a viable (and necessary) tool to accelerate the convergence of the calculation, by allowing the use of large time steps on coarse grids so that disturbances are rapidly expelled from the computational domain. However, up to now only a few applications of the technique to hypersonic flows have been reported.

Turkel et al.[1] have applied the standard full approximation storage (FAS) full multigrid (FMG) method[2,3] to compute viscous hypersonic flows over airfoils and blunt biconic bodies. They have concluded that the smoothing of coarse grid corrections is essential to ensure fast convergence. Leclercq and Stoufflet[4] have developed a characteristic multigrid method to solve inviscid hypersonic flows on unstructured and un-nested grids. Applications of the method to compute flows around airfoils and blunt bodies have shown that the use of an upwind prolongation operator and a geometrical restriction operator is the most efficient strategy for hypersonic flows. Koren and Hemker[5] have developed a damped direction-dependent multigrid method to solve steady Euler equations by means of point relaxation sweeps. They have shown that an upwind prolongation based on a MUSCL reconstruction of the fine grid solution, coupled with a locally damped classical residual restriction operator, is the most appropriate strategy for blunt body flows. Recently, Radespiel and Swanson[6] have investigated various multigrid schemes with semicoarsening to overcome stiffness problems due to the high cell aspect ratios, needed for high Reynolds number flows. They have used standard grid transfer operators coupled with local damping of the restriction operator and an upwind total variation diminishing (TVD) discretization, concluding that sequential semicoarsening gives the best results in terms of convergence rate, even though the computational work per multigrid cycle is about three times that of a full coarsening multigrid approach. Blazek et al.[7] have compared several implicit residual smoothing operators in combination with multigrid, showing that, for hypersonic flows, direction-dependent operators have better damping and convergence properties with respect to central ones.

In the present work, we have developed a multigrid technique for viscous hypersonic flows. The technique belongs to the class of the FAS-FMG methods, and it uses a "V-cycle" multigrid strategy and direction-dependent grid transfer operators. The full Navier-Stokes equations are solved by a finite volume approach with cell centered formulation, based on an adaptive dissipation scheme.[8] Time integration is performed by a multistage explicit Runge-Kutta algorithm coupled with a direction-dependent, implicit residual smoothing, to extend the stability region of time-stepping schemes.[3,8,9] The multigrid algorithm is based on a conservative residual restriction, geometrical solution restriction, and direction-dependent coarse grid correction prolongation.

Numerical Solution

The Navier-Stokes equations (in conservation form) are solved by means of a cell-centered finite volume formulation, whereby the computational domain is decomposed into arbitrary quadrilateral cells (i, j). Space and time discretization are separated by using the method of lines, thus reducing the governing equations to a system of ordinary differential equations for each computational cell. The basic numerical algorithm employs symmetric discretization of both the inviscid and viscous terms.[8]

Adaptive dissipation is added to prevent oscillations and even/odd point decoupling. On the finest grid the adaptive dissipation is a blending of first- and third-order derivatives,[3,9] with the shock sensor definition of Turkel et al.,[1] that ensures a TVD character of the scheme across shocks and improves the accuracy and the rate of convergence. A dissipation model based on first-order derivatives (computationally efficient) has been employed on coarse grids.[3,6] Time integration is performed by using a five-stage Runge-Kutta algorithm.

Implicit Residual Smoothing

To extend the stability bounds of the algorithm, we have analyzed the use of two direction-dependent implicit residual smoothing operators. Let $\boldsymbol{R}^{(k)}_{i,j}$ be the residual of the (discretized) equations at the kth stage of the Runge-Kutta algorithm.

In the first approach, acoustic residual smoothing (ARS), the residual is smoothed in the two directions according to the following algorithm

$$-\epsilon_i b_i \tilde{\boldsymbol{R}}^{(k)}_{i-1,j} + [1 + \epsilon_i(a_i + b_i)]\tilde{\boldsymbol{R}}^{(k)}_{i,j} - \epsilon_i a_i \tilde{\boldsymbol{R}}^{(k)}_{i+1,j} = \boldsymbol{R}^{(k)}_{i,j} \quad (1)$$

$$-\epsilon_j b_j \tilde{\tilde{\boldsymbol{R}}}^{(k)}_{i,j} + [1 + \epsilon_i(a_i + b_i)]\tilde{\tilde{\boldsymbol{R}}}^{(k)}_{i,j-1} - \epsilon_j a_j \tilde{\tilde{\boldsymbol{R}}}^{(k)}_{i,j+1} = \tilde{\boldsymbol{R}}^{(k)}_{i,j} \quad (2)$$

Received April 13, 1992; presented as Paper 93-0771 at the AIAA 31st Aerospace Sciences Meeting, Reno, NV, Jan. 11-14, 1993; revision received Feb. 18, 1993; accepted for publication March 12, 1993.

*Associate Professor, Dipartimento di Meccanica e Aeronautica. Member AIAA.
†Ph.D. Candidate, Dipartimento di Meccanica e Aeronautica.