

FAST MULTIGRID SOLUTION OF THE ADVECTION PROBLEM WITH CLOSED CHARACTERISTICS*

IRAD YAVNEH[†], CORNELIS H. VENNER[‡], AND ACHI BRANDT[§]

Abstract. The numerical solution of the advection-diffusion problem in the inviscid limit with closed characteristics is studied as a prelude to an efficient high Reynolds-number flow solver. It is demonstrated by a heuristic analysis and numerical calculations that using upstream discretization with downstream relaxation ordering in a multigrid cycle with appropriate residual weighting leads to an efficient solution process. Upstream finite-difference approximations to the advection operator are derived whose truncation terms approximate “physical” (Laplacian) viscosity, thus avoiding spurious solutions to the homogeneous problem when the artificial diffusivity dominates the physical viscosity [A. Brandt and I. Yavneh, *J. Comput. Phys.*, 93 (1991), pp. 128–143].

Key words. multigrid, advection-diffusion, recirculating flow, upstream discretization

AMS subject classifications. 65N12, 65N55, 76D05

PII. S1064827596302989

1. Introduction. Efficient multigrid algorithms for the numerical solution of partial differential problems normally require good ellipticity measures on all scales of the problem, which implies that nonsmooth solution components can be resolved by local processing [2]. But problems with small ellipticity measures are marked either by indefiniteness or by anisotropies. In the latter case, there exist so-called characteristic directions of strong dependence. Some nonsmooth components of the solution are then advected along these characteristics, and hence they cannot be resolved locally [1]. A typical example is steady flow at high Reynolds numbers (small viscosity).

When applied to such problems of small ellipticity the usual multigrid algorithms often exhibit a severe degradation of performance compared to that seen in elliptic problems. Indeed, most multigrid codes in use today for solving steady flows at high Reynolds numbers, although yielding a great improvement over previous single-grid solvers, fall far short of attaining the so-called textbook multigrid efficiency for general (even smooth) flows. To regain this efficiency the multigrid algorithm requires modifications that take into account the anisotropic properties of the operator. For example, it was shown in [4] and [9] that using upstream discretization and downstream relaxation ordering yields a fully efficient multigrid solver for flows whose characteristics (streamlines) start at some part of the boundary and end at another without recirculating (entering flows). To obtain efficient multigrid solvers for flows with closed characteristics, however, different modifications were proposed, such as defect-correction cycles and residual overweighting [5]. The main drawbacks of the latter approaches are: (a) they are not likely to generalize efficiently to orders of accuracy higher than one; (b) they require W cycles, which may be substantially more expensive than simple V cycles in parallel computation; (c) they suggest different treatment for different types of flow, viz., recirculating versus entering flows.

*Received by the editors May 3, 1996; accepted for publication September 3, 1996. This work was supported by The Royal Netherlands Academy of Arts and Sciences and The Feinberg Graduate School, by United States–Israel Binational Science Foundation grant 94-00250, by United States Air Force grant F49620-92-J-0439, and by the Carl F. Gauss Minerva Center for Scientific Computation.

<http://www.siam.org/journals/sisc/19-1/30298.html>

[†]Technion—Israel Institute of Technology, Haifa, Israel (irad@cs.technion.ac.il).

[‡]University of Twente, Enschede, the Netherlands (chvenner@wb.utwente.nl).

[§]The Weizmann Institute of Science, Rehovot, Israel (mabrandt@weizmann.weizmann.ac.il).

The upshot of the present work is to obtain a unified approach for both types of flow by employing upstream discretization and downstream relaxation ordering for recirculating flows as well. In section 2 we formulate the simple model problem of advection-diffusion and present the first differential approximation (FDA) to its discretized form. In section 3 we present the two-level cycle and use the approximation of section 2 in a heuristic analysis for a priori prediction of the performance of this algorithm. In section 4 new first-order upstream discretizations for the advection operator are presented, whose first truncation terms approximate isotropic diffusivity. These schemes are shown to eliminate spurious solutions to the homogeneous (i.e., unforced) small-viscosity advection-diffusion equation, such as those reported in [3]. Section 5 presents numerical calculations testing the accuracy of the discretization and the efficiency of the multigrid algorithm and how it compares to the predictions of section 3. Section 6 summarizes the main conclusions and further research plans.

2. The scalar advection-diffusion equation. We study the scalar advection-diffusion equation with closed characteristics as a prelude to the study of flow problems. This equation serves well as a preliminary problem, since the advection part (i.e., momentum equations) is responsible for the degraded performance observed in the solution of the incompressible-flow equations by the usual multigrid algorithm [5]. Also, as is shown for entering flows in [4], the solution process of the advection part of the system can be effectively decoupled from that of the elliptic part that is due to the continuity equation. Hence, efficient solution of the advection problem is a necessary stage in the development of a fully efficient flow-equations solver, and [4] suggests that the resulting advection-problem solver can indeed be used in designing the sought-after flow solver.

The advection-diffusion equation in two dimensions is

$$(1) \quad \begin{aligned} L_\epsilon u &= -\epsilon \Delta u + au_x + bu_y = f, & (x, y) \in \Omega, \\ u &= g, & (x, y) \in \partial\Omega, \end{aligned}$$

where ϵ is a positive constant and a , b , f , and g are given functions of x and y . Equation (1) is discretized on a uniform grid of meshsize h , whose gridlines lie parallel to the x and y coordinates. The characteristic direction of the advection operator in (1) is given (locally) by $dy/dx = b/a = \tan(\phi)$, where ϕ is the (local) angle of nonalignment between the x coordinate and the characteristic direction. We will focus our attention on the particular case where the characteristics defined by a and b form closed loops (as in vortices), one of which may coincide with $\partial\Omega$ (as in internal flows).

Suppose that (1) is discretized by some stable finite-difference discretization of first-order accuracy. The main aspects of the problem can be analyzed by substituting for the discrete operator its FDA—see [8] and also [1, 2, 9]. For the advection-diffusion equation with positive but vanishingly small ϵ we need only consider the advection operator, since the tiny diffusion will be dominated by the artificial diffusivity represented by the truncation terms (except at stagnation points). Let L^h denote a first-order accurate discrete approximation to the advection operator. Then, by a Taylor series expansion, we generally have

$$(2) \quad L^h u^h = a u_x^h + b u_y^h - h(\tilde{T}_1^h u_{xx}^h + \tilde{T}_2^h u_{xy}^h + \tilde{T}_3^h u_{yy}^h) + O(h^2),$$

where u^h denotes the discretized function, h being the meshsize of a uniform grid. Here, \tilde{T}_j^h are functions of x and y , the specific details of which are determined by a

and b and the discretization. The FDA is the approximation of L^h by the differential operator that remains in (2) after the $O(h^2)$ terms are neglected. (Hence, it applies only to sufficiently smooth u^h , since the neglected terms are higher derivatives.)

We now assume for simplicity of the discussion that the equation is normalized such that $a^2 + b^2 = 1$, and introduce a (conformal) local coordinate system, (ξ, η) , where ξ denotes the local “streamwise” coordinate parallel to the characteristic direction, while η denotes the “cross-stream” coordinate that is perpendicular to the characteristic. Thus,

$$(3) \quad \begin{aligned} \partial_\xi &= a\partial_x + b\partial_y, \\ \partial_\eta &= a\partial_y - b\partial_x, \end{aligned}$$

and

$$(4) \quad \begin{aligned} \partial_{xx} &= a^2 \partial_{\xi\xi} + b^2 \partial_{\eta\eta} - 2ab \partial_{\xi\eta}, \\ \partial_{yy} &= b^2 \partial_{\xi\xi} + a^2 \partial_{\eta\eta} + 2ab \partial_{\xi\eta}, \\ \partial_{xy} &= ab \partial_{\xi\xi} - ab \partial_{\eta\eta} + (a^2 - b^2) \partial_{\xi\eta}. \end{aligned}$$

The FDA of L^h in the local coordinate system is therefore

$$(5) \quad \tilde{L}^h = -h (T_1^h \partial_{\eta\eta} + T_2^h \partial_{\xi\eta} + T_3^h \partial_{\xi\xi}) + \partial_\xi,$$

where, by (2) and (4),

$$(6) \quad \begin{aligned} T_1^h &= b^2 \tilde{T}_1^h - ab \tilde{T}_2^h + a^2 \tilde{T}_3^h, \\ T_2^h &= -2ab \tilde{T}_1^h + (a^2 - b^2) \tilde{T}_2^h + 2ab \tilde{T}_3^h, \\ T_3^h &= a^2 \tilde{T}_1^h + ab \tilde{T}_2^h + b^2 \tilde{T}_3^h. \end{aligned}$$

We assume a consistent and stable discretization, which requires that the artificial viscosity operator represented by the first truncated term be elliptic, implying that the $O(h)$ part of the operator in (5) is of positive type. Under special circumstances, such as consistent alignment of the characteristics with the grid, T_1^h (and T_2^h) vanish, and this property is marginally violated. In this case the physical diffusion term becomes important, no matter how small ϵ may be, and the analysis below does not apply. Accordingly, we will assume below that T_1^h is large compared to h , which is the usual case.

3. Two-level error-reduction analysis. We now analyze the error reduction attainable with a two-level cycle using upstream discretization and downstream relaxation ordering.

3.1. Two-level cycle. The proposed two-level cycle for a given discrete problem $L^h u^h = f^h$ is defined as follows.

- Starting with some approximation to u^h , perform ν_1 (small integer) relaxation sweeps.
- Calculate the residuals, $r^h = f^h - L^h \tilde{u}^h$, where \tilde{u}^h is the current approximation to the solution, and transfer them to a twice-coarser grid $2h$, multiplied by a globally uniform weight W .
- Solve the coarse-grid problem, $L^{2h} v^{2h} = W r^{2h}$, for the correction. Here, r^{2h} is the restriction of r^h to the coarse grid.
- Interpolate and add the correction v^{2h} to the fine-grid approximation.
- Perform ν_2 (small integer) fine-grid relaxation sweeps.

In studying the asymptotic performance of the two-level cycle, the number of pre-relaxation sweeps ν_1 need not be distinguished from the number of postrelaxation sweeps ν_2 . (Recall that we associate asymptotic performance with the spectral radius ρ of the iteration matrix, and that $\rho(AB) = \rho(BA)$ for any pair of square matrices A, B of the same dimension.) We denote the total number of sweeps by $\nu = \nu_1 + \nu_2$.

In analyzing the two-level cycle we shall make many simplifying assumptions. The degree to which these assumptions are justified needs to be judged by the degree to which numerical results match the predictions of the analyses.

3.2. The model problem and analysis. We analyze the two-level convergence for the discrete approximation to (1) in the limit of vanishing ϵ for problems with closed characteristics by considering the following model problem on grid h :

$$(7) \quad L^h u^h(\xi, \eta) = f^h(\xi, \eta),$$

where L^h is, as above, the discretization of the advection operator. For the domain of solution and boundary conditions we require periodicity in ξ in order to simulate closed characteristics, and we choose for simplicity of the discussion $\xi \in [0, 1]$. The boundary conditions in the cross-stream direction are not germane in the present context. For simplicity, we let $-\infty \leq \eta \leq \infty$.

The main point of our approach is to use discretization that is purely upstream, and to relax the equation in downstream ordering, starting at $\xi = 0$. Downstream ordering means that we relax a variable only after relaxing all other unknowns which participate in the equation that corresponds to this variable (except, perforce, at $\xi = 0$). Thus, a full relaxation sweep results in the elimination of all the residuals except at a narrow band (of $O(h)$ width) that stretches from $\eta = -\infty$ to ∞ near $\xi = 1$ (which coincides with $\xi = 0$ due to the periodicity). Neglecting the width of this band, we find that the residual function, r^h , which remains after at least one full relaxation sweep has been carried out, can be modeled by

$$r^h(\xi, \eta) = \begin{cases} R^h(\eta), & \xi = 0 \pmod{1}, \\ 0, & \text{otherwise.} \end{cases}$$

The fine-grid error v^h satisfies the residual equation. Wherever the residual vanishes we now revert to FDA, obtaining

$$(8) \quad \tilde{L}^h v^h = 0,$$

with \tilde{L}^h defined in (5). We now add a further simplifying assumption that T_j^h in (5) are independent of η . Hence, we may expand (8) in a Fourier series in η . For an error component $\hat{v}_\omega^h(\xi) \exp(i\omega\eta)$ of frequency ω , (5) and (8) then yield

$$(9) \quad \left(h\omega^2 T_1^h(\xi) + O(h) + \frac{d}{d\xi} \right) \hat{v}_\omega^h(\xi) = 0,$$

where all the $O(h)$ terms but the first can be neglected in the homogeneous equation, since they multiply derivatives and are therefore small compared to $d\hat{v}_\omega^h/d\xi$. The solution to (9) in the interval $(0, 1)$ is therefore given by

$$(10) \quad \hat{v}_\omega^h(\xi) = A_\omega^h \exp \left(-h\omega^2 \int_0^\xi T_1^h(s) ds \right),$$

where A_ω^h is the amplitude of \hat{v}_ω^h at $\xi = 0^+$, which we shall determine shortly. (Superscripts $+$ ($-$) will denote an infinitesimal positive (negative) increment.) In particular, at $\xi = 1^-$, we have $\hat{v}_\omega^h = A_\omega^h D(h, \omega)$, with

$$(11) \quad D(h, \omega) = \exp(-h\omega^2 \bar{T}_1^h),$$

where \bar{T}_1^h is the average value of T_1^h over the entire domain, under the assumption that T_1^h is η -independent (and recalling that the domain length in the ξ direction is 1).

It is important to note that $D(h, \omega)$ is approximately the factor by which a single relaxation sweep amplifies (reduces) an error component that oscillates at frequency ω in the η direction. This is due to the fact that, given upstream differencing, the downstream relaxation ordering yields numerical integration; and $D(h, \omega)$ is the factor by which this integration over the domain reduces the error. Equation (11) implies that relaxation reduces error components with large ω very efficiently, but components that are smooth along η need to be corrected on the coarse grid.

At $\xi = 0 \pmod{1}$, the FDA is no longer useful. Instead, we have a jump in \hat{v}_ω^h that is proportional to the Fourier coefficient of $R^h(\eta)$ corresponding to frequency ω . We denote this jump by δ_ω^h . The periodicity in ξ now implies by (10) and (11) that

$$(12) \quad A_\omega^h = \delta_\omega^h / (1 - D(h, \omega)).$$

Now, following the two-level algorithm, we attempt to approximate the (weighted) residual equation on the coarse grid $2h$. We assume that the same discretization stencil is used on the coarse grid as on the fine. (*Note this important assumption on which the entire method hinges.*) We also assume that the restriction operator is such that the jump condition at $\xi = 0$ is approximated correctly on the coarse grid. In practice this holds, provided that a proper averaging, such as full-weighted residual transfers, is used. Analogously to (10), (11), and (12), respectively, we obtain

$$(13) \quad \hat{v}_\omega^{2h}(\xi) = A_\omega^{2h} \exp\left(-2h\omega^2 \int_0^\xi T_1^{2h}(s) ds\right),$$

$$(14) \quad D(2h, \omega) = \exp(-2h\omega^2 \bar{T}_1^{2h}),$$

and

$$(15) \quad A_\omega^{2h} = W\delta_\omega^{2h} / (1 - D(2h, \omega)),$$

where W is a constant weight to be chosen. Since the stencils of L^h and L^{2h} are the same, we also assume $\bar{T}_1^{2h} = \bar{T}_1^h$, and hence, by (14),

$$(16) \quad D(2h, \omega) = D(h, \omega)^2.$$

Also, since we assume that the restriction operator transfers the jump condition correctly, we have $\delta_\omega^{2h} = \delta_\omega^h$. Equations (12), (15), and (16) now yield

$$(17) \quad \frac{A_\omega^{2h}}{A_\omega^h} = W \frac{1 - D(h, \omega)}{1 - D(h, \omega)^2} = \frac{W}{1 + D(h, \omega)}.$$

Neglecting again the effects of intergrid transfers and aliasing, we assume that the remaining error after adding the coarse-grid correction is $v^h - v^{2h}$. Hence, the fine-grid error is amplified by the factor $(\hat{v}_\omega^h - \hat{v}_\omega^{2h})/\hat{v}_\omega^h$. In addition, ν relaxation sweeps

performed on the fine grid amplify the error by $D(h, \omega)^\nu$, as noted above. The two-level error-amplification factor μ_{tl} is then given by the absolute value of the product of these terms. Note, however, that the determining values of \hat{v}_ω^h and \hat{v}_ω^{2h} in the coarse-grid correction term depend on where one begins relaxing on the fine grid immediately following the coarse-grid correction (since all other fine-grid values at the end of one or more sweeps are determined solely by the values where relaxation begins, due to the upstream differencing and downstream relaxation ordering). If the fine-grid relaxation begins at 0^+ , that is, at or shortly after the point where the residual was nonzero, then we have

$$(18) \quad \mu_{tl} = \left| \left(1 - \frac{A_\omega^{2h}}{A_\omega^h} \right) D^\nu \right| = \left| \left(1 - \frac{W}{1+D} \right) D^\nu \right|,$$

where $D = D(h, \omega)$. However, if the fine-grid relaxation begins at 0^- (since there is some small overlap in the region being relaxed), then

$$(19) \quad \mu_{tl} = \left| \left(1 - \frac{A_\omega^{2h} D^2}{A_\omega^h D} \right) D^\nu \right| = \left| \left(1 - \frac{WD}{1+D} \right) D^\nu \right|.$$

3.3. Optimal residual weighting. By (11) we have $0 < D < 1$, and in order to obtain an h -independent analysis we assume now that D can take on any value in the interval $(0, 1)$. The optimal value of W is that which minimizes the *supremum* of μ_{tl} over $D \in (0, 1)$. For any fixed W , the *supremum* is evidently obtained either for $D \rightarrow 1$ or for

$$D_m \stackrel{\text{def}}{=} D \in (0, 1) : \frac{\partial \mu_{tl}}{\partial D} = 0.$$

From (18) and (19) we can thus obtain D_m as a function of W and ν , from which we can then calculate W_{opt} (the value of W which yields the fastest convergence) for either of these cases. For $\nu = 1$ we obtain $W_{opt} = 16/9$ for (18), yielding $\mu_{tl} = 1/9$, and $W_{opt} = 16/7$ for (19), yielding $\mu_{tl} = 1/7$. For larger ν the tedious calculations need to be carried out numerically. But W_{opt} tends to 2 for both cases rather quickly. This is expected, since this value is the ratio of the Green's functions on the coarse and fine grids for components that are very smooth in the cross-characteristic direction; other components are reduced by relaxation. With $W = 2$ we obtain for both (18) and (19)

$$(20) \quad \mu_{tl} = \left| \frac{1-D}{1+D} D^\nu \right|.$$

Equating the derivative of (20) with respect to D with zero, we get as the only relevant root,

$$(21) \quad D_m = \sqrt{1 + \nu^{-2}} - \nu^{-1}.$$

The two-level error-amplification factor is now obtained by substituting (21) into (20). For sufficiently small ν^{-2} we may neglect this term, obtaining $D_m \approx 1 - \nu^{-1}$ and

$$(22) \quad \bar{\mu}_{tl} \stackrel{\text{def}}{=} \sup_{D \in (0, 1)} \mu_{tl} \approx \left| \frac{\nu^{-1}}{2 - \nu^{-1}} (1 - \nu^{-1})^\nu \right| \approx (2\nu e)^{-1}.$$

In fact, (22) gives an excellent approximation of the maximal μ_{tl} in (20) for any ν , erring by less than 2% for $\nu \geq 2$. By curious coincidence, the same asymptotic

TABLE 1

The number of cycles necessary to reduce the L_2 residual norm by a factor of 10^8 in Example 1. The problem is taken from [10], and MGD9V is the automatic method of de Zeeuw which is used there.

Levels	Grid	$W=1$	$W=2$	MGD9V
4	33×33	16	6	15
5	65×65	26	7	17
6	129×129	41	7	22
7	257×257	65	8	N.A.
8	513×513	108	8	N.A.

two-level error-amplification factor for large ν , $(2\nu\epsilon)^{-1}$, is obtained for the Poisson equation on a rectangle using Gauss–Seidel relaxation in red-black order [7]! Thus, this analysis leads us to expect efficiency that is comparable to that obtained for the Poisson problem.

Example 1. We apply our algorithm to the advection-diffusion problem with closed characteristics used in [10] (originally in [6]):

$$-\epsilon\Delta u + a(x, y)u_x + b(x, y)u_y = 0,$$

on $\Omega = (0, 1) \times (0, 1)$, with $u(x, y) = \sin(\pi x) + \sin(13\pi x) + \sin(\pi y) + \sin(13\pi y)$ on $\partial\Omega$, and

$$a(x, y) = 4x(x-1)(1-2y), \quad b(x, y) = -4y(y-1)(1-2x).$$

For the advection term we use the same discretization as is used on the finest grid in [6] and [10]: standard upstream (SU), defined by

$$(23) \quad \text{SU} = \frac{1}{h} [(|a_{i,j}| + |b_{i,j}|)u_{i,j} - |a_{i,j}|u_{i1,j} - |b_{i,j}|u_{i,j1}],$$

where

$$i1 = i - \text{sign}(a_{i,j}), \quad j1 = j - \text{sign}(b_{i,j}).$$

In [10], $\epsilon = 10^{-5}$. We use this value only at the stagnation point, adding no viscosity elsewhere to maintain upstream discretization. Since the “physical” viscosity is dominated by the artificial viscosity elsewhere anyway, the difference is small. (Alternatively, we could use a much smaller ϵ everywhere.) As in [10], the initial guess for u in Ω is zero, and we cycle until the L_2 norm of the residuals is reduced to at most 10^{-8} times its initial value. We performed this test with 4, 5, 6, 7, and 8 levels, with the coarsest grid always 5×5 , including boundary points (as in [10]). We used V(1,1) cycles throughout, with the usual full-weighted residual transfers and bilinear interpolation. (See section 5 for details on implementation of the relaxation.) In Table 1 we compare our results with $W=1$ and 2 to those reported for MGD9V, the automatic method of de Zeeuw (available up to six levels), using a so-called “sawtooth” cycle with one ILLU (incomplete line LU) relaxation sweep per level. It must be stressed that the efficiency and robustness of this method is convincingly demonstrated in [10], and all the results achieved on ten other problems (including nonrecirculating advection-diffusion) were far better than these.

Evidently, the present method performs very well, with efficiency comparable to that of elliptic equations in this simple test problem. Clearly, the downstream relaxation ordering itself is not sufficient for the recirculation problem (nor is ILLU). Both

MGD9V and the present method with no residual overweighting show clear deterioration as the grid is refined, while with proper overweighting the convergence rate remains excellent even for very fine grids. We reiterate that our residual overweighting approach may not apply to MGD9V, since there different stencils are used on the different grids, requiring different overweighting.

3.4. Several bands of residuals. It may not always be easy to obtain a single band of residuals per vortex. This happens when the relaxation is carried out in piecewise downstream ordering, as would be the case in a domain-decomposition setting, for example. Our analysis can be extended to the case where several bands of nonzero residuals remain after relaxation. It is found that one then requires two coarse-grid corrections, with optimal weights 1 and 2 approximately. This seems to imply that a W cycle is required in this situation. Also, one must then use upstream intergrid transfers, so as to avoid averaging over interfaces between subdomains, which may actually cause divergence.

4. Discretization. Flows in which the streamlines do not start and end at boundaries, but constitute closed curves, require special considerations in the discretization. In such cases, even a very small viscosity plays an important role in determining the main flow throughout the domain. The solution in the limit of vanishing viscosity depends very strongly on how these coefficients tend to zero. In effect, the advection terms determine the behavior of the solution along streamlines, whereas the viscous terms determine its cross-stream form. And since the boundary is often a streamline itself, the propagation of information from the boundary into the domain is governed by the viscous terms no matter how small they may be. This effect is discussed in detail in [3, 9], where it is shown for both the advection-diffusion problem and the incompressible Navier–Stokes equations that solutions with schemes in which the numerical viscosity is anisotropic (having different viscosity coefficients for the cross-stream and streamwise directions), such as standard upstream-difference schemes, may be spurious.

In the most general case it can be shown that even isotropic viscosity is not sufficient for convergence of the solution, and one must actually specify a *uniform* viscosity. We do not know how to do this while retaining the purely upstream structure (but see remarks in section 6). However, for the homogeneous advection-diffusion problem, there are several indications (though no proof) that isotropy suffices. This is shown below and also in [3, 9], where it is also shown (in a numerical example) to suffice for the incompressible Navier–Stokes equations. This is consistent with the fact that the vorticity in the Navier–Stokes equations satisfies a *homogeneous* advection-diffusion equation.

To obtain a discretization scheme that exhibits the appropriate physical-like behavior for vanishing viscosity we must thus either add sufficient explicit isotropic viscosity that will dominate the artificial viscosity of the discrete advection operator, or else derive a discretization of the advection operator that satisfies the condition of isotropy in its lowest-order truncated terms. Since we want our scheme to remain purely upstream, we follow the latter approach.

Consider the standard upstream scheme of (23) and assume for simplicity of discussion $a \geq b \geq 0$. From (2) we have by a Taylor expansion

$$(24) \quad \tilde{T}_1^h = 0.5a, \quad \tilde{T}_2^h = 0, \quad \tilde{T}_3^h = 0.5b.$$

Hence, in order to obtain isotropic artificial viscosity we may either add some approximation of $0.5h(a-b)u_{xx}$, or else subtract some approximation of $0.5h(a-b)u_{yy}$.

In order to retain an upstream scheme we define this additional viscosity at the point $(i-1, j)$.

For *general* a and b we obtain in the first case the isotropic-viscosity upstream scheme IVU1, defined by

$$(25) \quad \text{IVU1} = \text{SU} + 0.5 |a_{i,j}| - |b_{i,j}| |(u_{i,j} - 2u_{i1,j} + u_{i2,j})|,$$

if $|a_{i,j}| > |b_{i,j}|$, and

$$\text{IVU1} = \text{SU} + 0.5 |a_{i,j}| - |b_{i,j}| |(u_{i,j} - 2u_{i,j1} + u_{i,j2})|$$

otherwise. In the second case we obtain scheme IVU2, defined by

$$(26) \quad \text{IVU2} = \text{SU} - 0.5 |a_{i,j}| - |b_{i,j}| |(u_{i1,j+1} - 2u_{i1,j} + u_{i1,j-1})|,$$

if $|a_{i,j}| > |b_{i,j}|$, and

$$\text{IVU2} = \text{SU} + 0.5 |a_{i,j}| - |b_{i,j}| |(u_{i+1,j1} - 2u_{i,j1} + u_{i-1,j1})|$$

otherwise. Here $i1, j1$ are defined as in (23), and similarly

$$i2 = i - 2 \text{sign}(a_{i,j}), \quad j2 = j - 2 \text{sign}(b_{i,j}).$$

The first truncated term in scheme IVU1 is thus $-0.5 \min(|a_{i,j}|, |b_{i,j}|) \Delta u$, while that of IVU2 is $-0.5 \max(|a_{i,j}|, |b_{i,j}|) \Delta u$. Both schemes have isotropic artificial viscosity, but that of IVU1 is smaller, and in fact it vanishes upon alignment of the characteristic directions with the grid.

Both discretizations are stable in downstream-ordered Gauss–Seidel relaxation. The former is a nonnegative-weighted average of the standard first-order and second-order upstream schemes, both of which are stable in this relaxation. The latter produces an M matrix. As expected, there were no stability problems in any of our many numerical calculations.

5. Numerical experiments. We first test numerically the discretizations derived in Section 4 on a model problem for which the standard upstream scheme has been shown to yield spurious solutions [3, 9]. Then, the asymptotic error reduction of two-level and multilevel cycles are investigated for several problems.

5.1. Accuracy test. The accuracy of the different discretizations is tested on the model problem

$$(27) \quad -\epsilon \Delta u + a u_x + b u_y = 0,$$

with a and b given by

$$(28) \quad a = \sin(\pi y) \cos(\pi x), \quad b = -\cos(\pi y) \sin(\pi x).$$

(These coefficients are the same as those of Example 2 below, and a picture of the characteristics appears in Figure 1(a).) The domain of solution is the unit square, centered at the origin, with a square of diagonal 0.5, whose sides form a 45 degree angle with the axes, removed from its center. On the outer boundary $u = 1$ is prescribed, and on the inner boundary, $u = 0$. We solve this problem with the three upstream schemes, SU, IVU1, IVU2, and also the (nonupstream) isotropic-viscosity scheme used in [3], denoted ISO. No “physical” viscosity is added: $\epsilon = 0$.

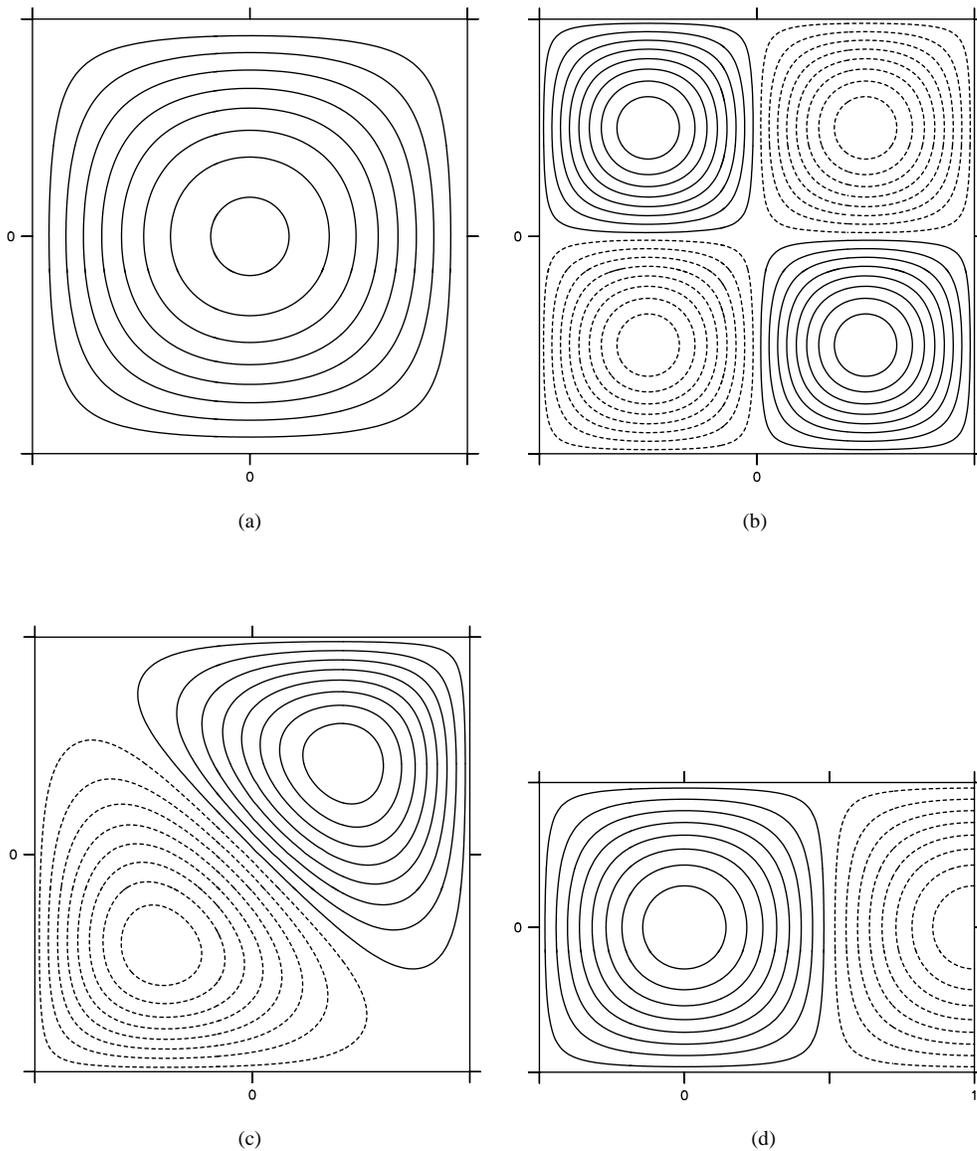


FIG. 1. Characteristics of Examples (a), 2, (b) 3, (c) 4, (d) 5.

These solutions are compared to that obtained with a standard second-order upstream scheme with physical viscosity coefficient 0.001. The latter solution is obtained on a 257×257 grid. In Table 2 we present L_1 norms of the differences between the test solutions at various resolutions and the second-order accurate solution (restricted to the corresponding grid by injection). Since the solution is smooth and the 0.001 physical viscosity dominates the second-order truncation terms at this high resolution, the latter solution is assumed to be very accurate.

Evidently, the three schemes with isotropic artificial viscosity produce convergent solutions, but the SU scheme does not, despite the fact that its average viscosity is

TABLE 2

L_1 difference norms between the solutions obtained with several schemes at different resolutions, and a high-accuracy solution obtained on grid 257×257 (see text for details). The three isotropic-viscosity schemes are seen to yield convergent solutions, but the SU scheme does not.

Grid	SU	IVU1	IVU2	ISO
33×33	0.0993	0.0153	0.0430	0.0265
65×65	0.0865	0.0094	0.0295	0.0187
129×129	0.0750	0.0057	0.0163	0.0098

smaller than that of IVU2 and ISO. IVU1, which has the least artificial viscosity of the four schemes tested, produces the smallest error.

5.2. Efficiency tests. The remainder of our numerical calculations are aimed at testing the performance of the algorithm in various configurations and comparing to the analytical predictions of (19) and (20). We test the SU scheme (as this is the most widely used first-order scheme) and the IVU1 scheme (which is more accurate than IVU2 and also employs just a four-point stencil). In all these tests we use first-order upstream residual restriction and bilinear interpolation of the corrections. The restriction is performed as follows: for all *even* i and j on the fine grid we define $i1, j1$ as in (23), and restrict to the corresponding coarse-grid right-hand side at coarse-grid point $(i/2, j/2)$ the average of the fine-grid residuals at points (i, j) , $(i1, j)$, $(i, j1)$, and $(i1, j1)$. This restriction gives slightly better results than standard full weighting in multivortex problems, since residuals are less likely to be transferred from one vortex to another.

The finest grid in all the tests is 129 by 129 , and six levels are employed except in the two-level tests. We include no “physical” viscosity except at stagnation points ($a = b = 0$), where it is required for well-posedness. We calculate convergence factors as follows: the boundary conditions and right-hand sides are chosen to be zero; (the choice is immaterial for linear problems, but this allows us to normalize the solution by a constant factor every few cycles in order to avoid roundoff errors). The initial solution field is pseudorandom, and 100 cycles are performed. We calculate the convergence factor as the (geometric) average error-convergence per cycle over the last 80 cycles. The averaging is used because in some cases the convergence history is not smooth, and the value corresponding to any particular cycle may not carry much meaning. However, it should be noted that in all cases the convergence factors in the vicinity of the optimal W were not sensitive to the exact choice of W .

Example 2. The first test for efficiency is the problem used above to test the discretization, but without the inner “island” (so as to have pure recirculation everywhere). The characteristics of this problem, which form a single clockwise-rotating vortex, are plotted in Figure 1(a). A relaxation sweep is implemented by sweeping four times over the domain, and in each such sweep relaxing roughly one-quarter of the variables as follows: in the first sweep only variables at locations where both $a(x, y)$ and $b(x, y)$ are nonnegative are relaxed (designated first quadrant); in the second sweep only variables corresponding to locations where $a(x, y)$ is nonnegative and $b(x, y)$ is nonpositive are relaxed (second quadrant); in the third sweep only variables at locations where both $a(x, y)$ and $b(x, y)$ are nonpositive are relaxed (third quadrant); in the fourth sweep only variables corresponding to locations where $a(x, y)$ is nonpositive and $b(x, y)$ is nonnegative are relaxed (fourth quadrant). Finally, all stagnation points (in this case just one) are relaxed last. This entire process comprises a single clockwise sweep. Of course, each quarter-sweep is performed in downstream

TABLE 3

Comparison of error-amplification factors obtained by analytical prediction and two-level numerical calculations of Example 2 for $W = 2$ and W_{opt} . The optimal values of W were obtained from (19).

ν	W	Two-Level Analysis	SU	IVU1
1	2.00	0.172	0.207	0.239
1	2.29	0.143	0.182	0.220
2	2.00	0.090	0.090	0.112
2	2.14	0.072	0.077	0.084
3	2.00	0.061	0.056	0.082
3	2.10	0.048	0.044	0.064

order, i.e., x and y increasing in the first quarter-sweep, x increasing and y decreasing in the second, etc. It is efficient to store the order of relaxation of the entire sweep during a setup-sweep (which costs very little), so that from then on each full clockwise sweep costs nearly the same as an ordinary lexicographic Gauss–Seidel sweep.

Note that we specify *nonnegative* and *nonpositive* in the above description. That is, it should be ensured that the boundaries of the quadrants are included in the quadrant. This was important in some of the tests.

Such a clockwise sweep indeed eliminates all the residuals except along a narrow band that extends from the center of the vortex to the boundary. But if the vortex rotates counterclockwise, several such bands would remain. Hence, in Examples 3–5 below, where both clockwise and counterclockwise vortices exist (as would be the general case), we perform an analogous counterclockwise sweep following each clockwise one. The quadrant where this counterclockwise sweep starts has some bearing on performance. In our experiments we use exactly the reverse order. This allows us to save some work by performing the counterclockwise sweep over just three quadrants, since the fourth quadrant ($a(x, y)$ nonpositive and $b(x, y)$ nonnegative) has just been relaxed in the clockwise sweep. Thus, we begin with the third quadrant (a and b nonpositive), then the second, and finally the first. (Note, by the way, the difference from the usual symmetric Gauss–Seidel: here the sweeps are all performed in downstream ordering within each quadrant, but the quadrants are scanned symmetrically).

In Examples 3–5 (especially 4) there was also some sensitivity to in which quadrant one chooses to start the relaxation sweep. That is, a somewhat different convergence rate and optimal overweighting were obtained if one performed the clockwise sweep as above, than if one relaxed, say, the second quadrant first, then the third, then the fourth, and finally the first. Hence, in these tests we shifted the relaxation starting-point by one quadrant after every cycle so as to obtain results that represent some average or “typical” case.

We first performed two-level tests in order to compare the numerical results with the analysis. We used $W = 2$ and the theoretically optimal W , calculated from (19). This is the relevant value for this example, since there is a one-line overlap between the quadrants. The results are given in Table 3. Evidently, the analysis captures the main features of this problem very well, despite the numerous simplifications, as seen especially in the SU results.

Calculations were also carried out with V(1,0), V(1,1), and V(2,1) cycles, with $W = 1, 2$ and the optimal value, which was determined experimentally. The numerical results are summarized in Table 4. We find that the experimental V-cycle results also match the two-level predictions fairly well in the vicinity of the optimal residual-weighting factors. But the optimal W ’s are somewhat higher than predicted, although

TABLE 4

Error-amplification factors obtained in numerical calculations of Example 2. The optimal values of W were obtained experimentally (see text).

Cycle	W	SU	IVU1
V(1,0)	1	0.795	0.898
V(1,1)	1	0.676	0.831
V(2,1)	1	0.536	0.724
V(1,0)	2	0.280	0.440
V(1,1)	2	0.143	0.302
V(2,1)	2	0.069	0.133
V(1,0)	<i>opt</i>	0.168 ($W = 2.44$)	0.232 ($W = 2.50$)
V(1,1)	<i>opt</i>	0.077 ($W = 2.27$)	0.066 ($W = 2.41$)
V(2,1)	<i>opt</i>	0.039 ($W = 2.15$)	0.041 ($W = 2.25$)

they do show the expected dependence on the number of relaxation sweeps performed. In other problems, reported below, the optimal value varied, but it was always fairly close to 2.

The overall performance when the optimal W is used is very satisfactory. As noted above, this performance is not sensitive to moderate changes in W (see also below).

Example 3. This example features flow with four vortices rather than just one. Here, a and b are given by

$$(29) \quad a = \cos(2\pi y) \sin(2\pi x), \quad b = -\sin(2\pi y) \cos(2\pi x),$$

and the characteristics are plotted in Figure 1(b). The domain, as in all the examples except Example 5, is the unit square centered at the origin. The numerical results using V(1,1) and V(2,1) cycles appear in Tables 5 and 6, respectively. The convergence performance remains excellent, even though some nonvanishing residuals remain after the relaxation at parts of the borderlines between vortices (where the flow leads away from the borderline). Recall, however, that here and below each full relaxation sweep consists of one clockwise sweep followed by three-quarters of a counterclockwise sweep (see description of implementation above), in order to allow for the opposite-sign vortices.

Example 4. In order to test the effect of grid-alignment of the borderlines between vortices we solve a problem in which this borderline is not aligned with the grid. In this problem a and b are given by $a_1 + a_2$ and $b_1 + b_2$, respectively, where

$$(30) \quad a_i = \sin(\pi y_i) \cos(\pi x_i), \quad b_i = -\cos(\pi y_i) \sin(\pi x_i), \quad i = 1, 2,$$

with

$$\begin{aligned} x_1 &= (x + 0.5)^2 + 0.5, & y_1 &= (y + 0.5)^2 + 0.5, \\ x_2 &= (x - 0.5)^2 + 0.5, & y_2 &= (y - 0.5)^2 + 0.5. \end{aligned}$$

This represents a superposition of two opposite-sign vortices. The characteristics are depicted in Figure 1(c), and the numerical performance is shown in Tables 5 and 6. Here we see some loss of efficiency, but the performance is still satisfactory and far better than is usually exhibited in such problems.

Example 5. Here we test a mixed problem, where the flow enters and leaves through the boundary, but there also exists a large recirculation zone. This is obtained

TABLE 5

Error-amplification factors obtained with $V(1, 1)$ cycles for $W = 2$ and W_{opt} . The latter were found experimentally.

Example	SU	IVU1
3	0.072 ($W = 2.00$)	0.179 ($W = 2.00$)
3	0.052 ($W = 2.07$)	0.101 ($W = 2.29$)
4	0.182 ($W = 2.00$)	0.230 ($W = 2.00$)
4	0.130 ($W = 1.88$)	0.206 ($W = 2.05$)
5	0.061 ($W = 2.00$)	0.185 ($W = 2.00$)
5	0.047 ($W = 2.05$)	0.109 ($W = 2.23$)

TABLE 6

Error-amplification factors obtained with $V(2, 1)$ cycles for $W = 2$ and W_{opt} . The latter were found experimentally.

Example	SU	IVU1
3	0.035 ($W = 2.00$)	0.090 ($W = 2.00$)
3	0.023 ($W = 2.06$)	0.042 ($W = 2.24$)
4	0.113 ($W = 2.00$)	0.129 ($W = 2.00$)
4	0.063 ($W = 1.84$)	0.128 ($W = 1.97$)
5	0.030 ($W = 2.00$)	0.091 ($W = 2.00$)
5	0.020 ($W = 2.07$)	0.048 ($W = 2.19$)

by redoing Example 2, but in an extended domain:

$$\Omega = (-0.5, 1) \times (-0.5, 0.5).$$

The mesh is still uniform, and the finest grid is now 193 by 129. The characteristics are shown in Figure 1(d), and the numerical performance is given in Tables 5 and 6. As in the other examples, there was virtually no sensitivity to the order in which the relaxation sweeps were performed (that is, clockwise first and then counterclockwise or vice versa).

Summary. Often, one may not wish to search for optimal residual weighting factors for every problem. Instead, one can simply use the nominal value $W = 2$. In the “realistic” Examples 3–5, one saves at most 25% of the time spent in relaxation by using the optimal value rather than 2, and usually much less. This would also be the case in Example 1 if we were to employ the “symmetric” relaxation. Even with this nominal value the convergence rates are comparable to those of elliptic problems.

6. Conclusions and further research. An experimental approach that is hoped to eventually lead to a fully efficient solver for general high-Reynolds flows has been introduced, analyzed, and tested on the advection-diffusion problem in the inviscid limit. The numerical tests mostly match the predictions very well, indicating that the main cause for slow convergence of the usual multigrid algorithms for recirculating flows has indeed been understood, and a way to eliminate it has been found.

The multigrid V-cycle, using downstream relaxation and upstream discretization, was shown to yield an efficient solver for the tested problem in several simple situations of closed characteristics and in a mixed entering/recirculation problem. The tests were performed with the classical standard first-order upstream discretization scheme and also with a novel first-order upstream discretization, that was shown to preclude the spurious solutions reported in [3].

The present approach is cheaper to implement than that developed in [5], and can straightforwardly be applied to mixed entering/recirculating flows. More important, there is potential for success with high-order discretization, for which the approach of [5] yields an inadequate compromise. However, the results obtained are still preliminary. The effect of the intergrid transfers on the small band of residuals and its consequences in terms of error-reduction efficiency should be investigated over a wide variety of cases, along with a study of how to deal with (or avoid) situations where there remain several bands of nonvanishing residuals per vortex. Then, further research should be directed towards higher-order discretization. For this case too, an effectively upstream discretization needs to be developed—one whose truncation error represents isotropic artificial diffusivity. One approach is to use a predictor-corrector-type discretization, employing an upstream scheme as a (local) driver and a possibly higher-order (not necessarily upstream) scheme as a (local) corrector. Finally, the present approach has only been tested on the advection problem. Experiments with the incompressible Navier–Stokes equations for flows with closed streamlines need to be performed, employing distributive Gauss–Seidel relaxation, as shown in [4]. These will no doubt raise further questions.

It is anticipated that the techniques investigated here will carry over to three dimensions, although the implementation will be considerably more complicated. This is supported by the fact that simple experiments performed with the overweighting methods of [5] in three dimensions exhibited the expected performance.

An obvious drawback of the entire approach is that it is inherently sequential, and efficient parallel implementations are hard to envisage. Some parallelization might be achievable by performing a downstream line Gauss–Seidel relaxation. Also conceivable is a domain decomposition approach which leaves several lines of residuals per vortex.

Another drawback of the present approach is that it is not directly applicable for flows with significant additional viscosity, since this entails using discretizations that are not purely upstream. Methods that deal with such flows as well are presently being investigated.

REFERENCES

- [1] A. BRANDT, *Multigrid Solvers for Non-Elliptic and Singular-Perturbation Steady State Problems*, The Weizmann Institute of Science, Rehovot, Israel, 1981, unpublished.
- [2] A. BRANDT, 1984 *Multigrid Guide with Applications to Fluid Dynamics*, Monograph, GMD Studien No. 85, GMD-F1T, Postfach 1240, D-5205, St. Augustin 1, Germany, 1985, unpublished.
- [3] A. BRANDT AND I. YAVNEH, *Inadequacy of first-order upwind difference schemes for some recirculating flows*, *J. Comput. Phys.*, 93 (1991), pp. 128–143.
- [4] A. BRANDT AND I. YAVNEH, *On multigrid solution of high-Reynolds incompressible entering flows*, *J. Comput. Phys.*, 101 (1992), pp. 151–164.
- [5] A. BRANDT AND I. YAVNEH, *Accelerated multigrid convergence and high-Reynolds recirculating flows*, *SIAM J. Sci. Comput.*, 14 (1993) pp. 607–626.
- [6] J. RUGE AND K. STÜBEN, *Efficient solution of finite difference and finite element equations by algebraic multigrid*, in *Multigrid Methods for Integral and Differential Equations*, D. J. Paddon and H. Holstein eds., *Inst. Math. Appl. Conf. Ser. New Ser.*, Oxford Univ. Press, New York, 1985, pp. 169–212.
- [7] K. STÜBEN AND U. TROTTEMBERG, *Multigrid methods: Fundamental algorithms, model problem analysis and applications*, in *Multigrid Methods*, W. Hackbusch and U. Trottenberg eds., Springer-Verlag, Berlin, 1982, pp. 1–176.
- [8] N. N. YANENKO AND Y. I. SHOKIN, *On the Correctness of First Differential Approximation of Difference Schemes*, *Dokl. Akad. Nauk. SSSR*, 182 (1968), pp. 776–778.
- [9] I. YAVNEH, *Multigrid Techniques for Incompressible Flows*, Ph.D. thesis, The Weizmann Institute of Science, Rehovot, Israel, 1991.
- [10] P. M. DE ZEEUW, *Matrix-dependent prolongations and restrictions in a blackbox multigrid solver*, *J. Comput. Appl. Math.*, 33 (1990), pp. 1–27.