Submitted to Ocean Modelling (August 2001); revised Feb/March 2002

Overcoming the CFL time-step limitation : a simple, effective and stable iterative implicit numerical scheme for slowly evolving advection-diffusion systems.

J.G. Shepherd

School of Ocean & Earth Science
Southampton Oceanography Centre
University of Southampton
European Way, Southampton, SO14 3ZH, UK

Abstract

Numerical calculations of advection-diffusion problems (such as the long-term evolution of deep ocean temperature, salinity and other property distributions) are often severely constrained by the Courant-Friedrichs-Levy (CFL) time-step limit and/or the equivalent diffusive stability limits. Implicit numerical schemes are an effective way of overcoming these limits, but both direct and iterative methods for solving the implicit equations can be difficult and time-consuming to implement especially for two and three-dimensional flows with irregular boundaries. A simple iterative method (similar to the Jacobi relaxation method) for obtaining an approximate solution to the implicit equations at each time step is proposed, which is particularly suitable for long-term integrations of systems which become slowly varying, and approach an actual or approximate (i.e secular or transient) steady-state as they respond to slowly varying forcing. This is unconditionally stable and relatively straightforward to implement, but not exactly conservative. It enables much longer timesteps (up to at least 30 times those constrained by the limits on an explicit scheme) but involves a few iterations, so that overall reductions of computation time of about one order of magnitude are achievable in practice. The method is not restricted to quasi-steady conditions but offers no advantage for transient conditions. It should be particularly useful for long-term geochemical and palaeoclimate calculations in non eddy-resolving models of moderate or high spatial resolution. Its potential advantages and disadvantages are discussed, and some preliminary results are presented.

1) Introduction

Maximum time-steps for conventional explicit numerical methods for solving the advection-diffusion equations are generally constrained by the Courant-Friedrichs-Levy (CFL) limit and/or the equivalent diffusive stability limits, which can be very restrictive and thus make long integrations very costly. This is a particularly frustrating problem for systems which are subject to slowly varying external forcing, and whose solutions thus reach approximate (secular or transient) equilibrium states. Such situations are common for models of natural geophysical and bio-geochemical systems, including the ocean-climate system, for which very long integrations (relative to the system response time) are often required. A particular and well-known problem is the need to spin-up the large-scale deep ocean circulation and the associated temperature, salinity and other property distributions, especially under different (past and future) climate conditions. Whilst most ocean models still use fairly conventional explicit methods (see Kantha and Clayson 2000), a number of special and somewhat *ad hoc* methods of acceleration have been developed, generally using split time-scales and/or distorted physics, particularly to address the deep circulation spin-up problem (Bryan 1984; Killworth, Smith et al. 1984; Danabasoglu, McWilliams et al. 1996; Klinger 2000), and also to overcome difficulties encountered with free-surfaces

(Griffies and Pacanowski 2001). Similar problems arise in dealing with convection and stability-dependent diapycnal mixing, and are also usually solved by special schemes (e.g. Marotzke 1991; Hallberg 2000). The computational penalty of a stability-limited explicit scheme is however particularly severe (and fundamentally inappropriate) for systems such as non eddy-resolving climate models responding to slowly varying forcing, which is the principal focus here.

In addition to the special schemes mentioned above, a number of more general numerical methods are available which allow for longer timesteps, including semi-Lagrangian methods which address the advection (CFL-limited) problem (see e.g. Staniforth and Cote 1991; see also (Leonard, Lock et al. 1996; Lin and Rood 1996)). Implicit methods (see e.g. Roache 1998), are applicable to both advection and diffusion limited problems, and are therefore well adapted to the present problem. Schemes which are both semi-Lagrangian and implicit are well suited to unsteady problems, but for quasi-steady problems and very large Courant numbers an Eulerian framework is preferable. Implicit methods generate block-tridiagonal systems of simultaneous equations, which need to be solved by either direct or iterative methods at each timestep (Anderson 1995). Direct methods require the manipulation of large sparse matrices, and are thus less suitable for the very large systems which arise with three-dimensional models with even moderate resolution. Very efficient iterative methods for solving the linear systems arising from multi-dimensional PDE's are now available, and among these multi-grid methods (Brandt 1977; Briggs 1987) have proved to be extremely effective (Roache 1998). The implementation of such methods however involves substantial complication and effort for complicated geometry, especially with irregular boundaries (Tong and Tuminaro 1999) and for flow-fields in several dimensions (i.e. for most cases of practical interest in geophysical fluid dynamical systems). Other iterative methods which are more easily applicable in such cases range from the alternating directions implicit (ADI) method (see e.g. (Ames 1969)) to various conjugate gradient algorithms and the GMRES method (Saad and Schultz 1986): see also Barrett, Berry et al. (1994), and Saad and van der Vorst (2000) for an up-to-date review. These methods have been developed to secure very effective final convergence to an accurate solution after a substantial number of iterations, and of course they also involve non-trivial implementation and computational costs. These may be worthwhile if accurate solutions are indeed required, although simpler schemes such as SOR can also be competitive in some circumstances (Woznicki 2001). In the present context there is little point in using an elaborate method to find a very accurate solution at each time-step, given that the system of implicit equations is itself only an approximation, on account of the substantial truncation errors due both to limited spatial resolution, and to the use of long time-steps. Speed and stability are in this situation more important considerations than obtaining a precise solution to an approximate statement of the problem. It is therefore appropriate and necessary only to solve the implicit equations approximately at each time-step, and only a few iterations of a fairly simple scheme are affordable if the computational advantage of a long time-step is to be realised.

A practical alternative method for the approximate solution of the implicit equations is proposed here. This is a very simple iterative method, which is straightforward to implement even for complicated geometry in three dimensions. It is robust and unconditionally stable, allowing timesteps which are limited by truncation errors, but which can be more than one order of magnitude longer than those for conventional explicit methods, once the advective-diffusive fluxes have reached approximate balance with the external sources, sinks, and boundary conditions, at an acceptable and moderate computational cost. It is also quite capable of dealing with transient conditions such as the early stages of spin-up from initial conditions, and/or the response to changes of the external forcing terms, but it has a computational penalty (rather than an advantage) unless and until the solutions become slowly varying, as quasi-equilibrium conditions

are approached. The algorithm is closely related to well-known relaxation methods for steady-state systems, and is therefore almost certainly computationally sub-optimal, and inferior to the more advanced methods cited above. It does nevertheless provide a worthwhile computational advantage for a relatively very modest investment in implementation, and it should be useful for a wide range of non-eddy-resolving models, such as the ocean component of intermediate complexity climate models. The structure of the simple method proposed also makes it potentially well-suited for parallelisation for large problems.

2) Derivation of the iterative implicit scheme

A conventional explicit, forward, finite difference scheme for the advection-diffusion equations in two dimensions, employing weighted upstream/centred differences for the advection terms is (see e.g. Roache 1976, 1998)

$$\begin{aligned} C_o' - C_o \;=\; & d_l\,(C_l - C_o) \;+\; d_r\,(C_r - C_o) \;+ d_u\,(C_u - C_o) \;+\; d_d\,(C_d - C_o) \\ & +\; 0.5\,a_l\,\left[(1+w_l)\,C_l \;+\; (1-w_l)\,C_o\right] - 0.5\,a_r\,\left[(1+w_r)\,C_o \;+\; (1-w_r)\,C_r\right] \\ & +\; 0.5\,a_u\,\left[(1+w_u)\,C_u \;+\; (1-w_u)\,C_o\right] - 0.5\,a_d\,\left[(1+w_d)\,C_o \;+\; (1-w_d)\,C_d\right] \end{aligned} \tag{1}$$

where C_o' is the new value of concentration C at the central point, denoted by suffix "o", suffices l, r, u and d indicate left, right, up and down respectively, all fluxes are evaluated at the cell interfaces, the "a" coefficients (for advection) are the Courant numbers uDt/Dx (etc) and the "d" coefficients (for diffusion) are the diffusion numbers $KDt/(Dx)^2$ (etc). I use this somewhat unconventional mnemonic notation in an attempt to avoid the considerable confusion which can otherwise arise from the proliferation of subscripts and superscripts in a multi-dimensional scheme, which involves various time-levels, successive approximations, and interface values as well as estimates at intermediate time levels. The upstream/centred weighting coefficients, denoted by w, which are an important feature of the scheme, are such that sign(w)=sign(a) and $-1 \le w \le +1$. For upstream differences w=sign(a), and for centred differences w=0, but a more general expression is proposed below. The equations are here written for a two-dimensional system, just as a reasonable compromise between complexity and generality, but may easily be modified for more or fewer dimensions.

The maximum time-step using such a scheme is normally limited by the Courant-Friedrichs-Levy (CFL) stability limit, i.e. max{abs(a)}< 1 , or the related diffusive limit max{abs(d)}< 1/2n in n dimensions. These limitations can however be avoided relatively easily by using an implicit scheme, especially for systems which will reach an approximate and/or transient steady-state (and thus become slowly varying in time), such as those with relatively rapid interior diffusive and advective processes, which are expected to reach a secular or transient equilibrium with slowly varying forcing terms (i.e. boundary conditions, sources and sinks). The method proposed is therefore not likely to be useful for eddy-resolving models where the magnitude of the transient terms never becomes small. It is however appropriate for long integrations of geophysical and biogeochemical systems in which turbulent processes are parameterised, especially those with moderate to high spatial resolution, for which the conventional stability restrictions become increasingly restrictive.

To derive an implicit scheme, one replaces all the C values on the RHS by some estimate C* of future values of the concentration, which are appropriate for determining the advective and

diffusive fluxes **during** a time-step, such as an average which may involve the values C' at the next time level, as well as the available (past) values C. This leads to

$$C_o' - C_o = d_l\,(C_l^* - C_o^*) + d_r\,(C_r^* - C_o^*) + d_u\,(C_u^* - C_o^*) + d_d\,(C_d^* - C_o^*)$$

$$+ 0.5\,a_l\,[(1+w_l)\,C_l^* + (1-w_l)\,C_o^*] - 0.5\,a_r\,[(1+w_r)\,C_o^* + (1-w_r)\,C_r^*]$$

$$+ 0.5\,a_u\,[(1+w_u)\,C_u^* + (1-w_u)\,C_o^*] - 0.5\,a_d\,[(1+w_d)\,C_o^* + (1-w_d)\,C_d^*] \qquad (2)$$

A suitable estimate for C^* is a weighted average of past and future values, i.e.

$$C^* = a\,C' + bC \qquad (3)$$

where $b=(1-a)$, and $a=1$ leads to a fully implicit method, $a=0$ leads to a conventional explicit method, and $a=0.5$ yields a second-order two-level centred "semi-implicit" method, analogous to the Crank-Nicholson method for diffusion. The resulting simultaneous equations for the values of C' need to be solved at each timestep. A very simple method for solving these equations can be constructed, which is effective for the relatively "easy" situation of a slowly varying system considered here. This is obtained by collecting all the terms involving future values of the central concentration (C_o') only onto the LHS, and is therefore essentially the conventional Jacobi relaxation method for elliptic and steady-state systems, here applied at each timestep (see e.g. (Anderson 1995), (Fletcher 1991), (Tannehill, Anderson et al. 1997); note that (Roache 1976; Roache 1998) refers to this as Richardson's method). This leads to

$$C_o'\,[\,1 + a\,F] = C_o\,\{1 - bF\} + d_l\,C_l^* + d_r\,C_r^* + d_u\,C_u^* + d_d\,C_d^*$$

$$+ 0.5\,\{(1+w_l)\,a_l\,C_l^* - (1-w_r)\,a_r\,C_r^* + (1+w_u)\,a_u\,C_u^* - (1-w_d)\,a_d\,C_d^*\} \qquad (4)$$

where

$$F = [d_l + d_r + d_u + d_d - 0.5\,\{(1-w_l)\,a_l - (1+w_r)\,a_r + (1-w_u)\,a_u - (1+w_d)\,a_d\}] \qquad (5)$$

For a non-divergent flow-field, we always have $(a_l - a_r + a_u - a_d) = 0$, and thus

$$F = d_l + d_r + d_u + d_d + 0.5\,\{w_l\,a_l + w_r\,a_r + w_u\,a_u + w_d\,a_d\} \qquad (5a)$$

This positive quantity, F, is essentially a **volume replacement factor**, i.e. roughly the number of times the contents of any cell would be replaced, by advective and/or diffusive processes, in a single time-step. This may be considerably larger than one, for long time-steps exceeding the CFL and diffusive stability limits, which are together broadly equivalent to requiring $F<1$.

An explicit expression for C_o' in terms of its values at the previous time-level, and the current estimates for the adjacent points, is therefore

$$C_o' = [\,C_o\,\{1 - bF\} + d_l\,C_l^* + d_r\,C_r^* + d_u\,C_u^* + d_d\,C_d^*$$

$$+ 0.5 \{(1+w_l)\, a_l\, C_l^* - (1-w_r)\, a_r\, C_r^* + (1+w_u)\, a_u\, C_u^* - (1-w_d)\, a_d\, C_d^* \}] \;/\; [\, 1 + a\, F]$$

(6)

Note that in equations (4) and (6) the value of the **central** concentration (C_o) which appears on the RHS is the old value (that at time level t), whilst all the **adjacent** values are averages, which involve both the past values and the current estimates for the future time level ($t+Dt$), i.e. C'_u (etc), via equation (3).

Repeated application of equation (6) constitutes the simple iterative implicit method proposed. For the initial "predictor" step, one can approximate the future values of the concentrations at adjacent points by their values at the previous time-level, as in the DuFort-Frankel method for the diffusion problem ((DuFort and Frankel 1953; Roache 1976)) and simply replace all future (C') values on the RHS, and thus also the average values C^*, with their old values (C).. This is only likely to be a reasonably good approximation in the situation considered here, i.e. where the concentration field (C), is only evolving slowly, as only in this case will the old and new values be similar. A fully implicit initial prediction may also be obtained, if required, by setting $\alpha=1$, to obtain a first approximation which may explicitly be written as

$$C_o' = [\, C_o + d_l\, C_l + d_r\, C_r + d_u\, C_u + d_d\, C_d$$

$$+ 0.5 \{(1+w_l)\, a_l\, C_l - (1-w_r)\, a_r\, C_r + (1+w_u)\, a_u\, C_u - (1-w_d)\, a_d\, C_d \}] \;/\; [\, 1 + F\,]$$

(7)

However, in practice there actually seems to be no advantage in using a different value of α for the first approximation, and equation (6) may be used throughout. It should be noted (a) that all these approximations introduce mixed forward/backward space/time differences on the RHS, and also (b) that in practice the equations for the initial approximation do not need to be coded separately, as the code for equation (6) can be used directly, simply setting $C'=C$ on the RHS.

The initial approximation so obtained can then be refined by one or more further applications of equation (6) as an iterative refinement, using any desired value of α, including $\alpha=1$ (fully implicit), $\alpha=0.5$ ("semi-implicit", for second-order accuracy) or some other value (see below). This constitutes the very simple iterative implicit scheme proposed. Its properties, and in particular the conditions under which the iteration converges are discussed below. The extension to three (or even more) dimensions is obvious and straightforward.

3) Stability, conservation and other properties

For an iterative scheme such as this to be useful in practice, it is necessary that both (a) the iterative sequence converges to a unique solution of the equation for each new time-level which is being solved, and (b) that this solution for the new time-level have satisfactory properties in respect of stability and convergence (with respect to time-stepping) as well as accuracy and consistency. If the iterative scheme converges, equation (6) is solved, and the final new result obtained for each time-level is by definition identical to that which would have been obtained by any other suitable method for implementing the implicit method, as this solves the same equations. Any convergent iterative scheme must therefore inherit the same time-stepping and other properties as conventional direct methods for solving the implicit equations. These properties are well-known (Anderson 1995; Roache 1998) and are therefore not discussed in detail here. Briefly,

for $\alpha>0.5$, implicit methods are unconditionally stable for any time-step (i.e. for arbitrarily large values of the volume replacement factor, F). However, they are not very accurate unless $\alpha\approx0.5$ and the time-step is not too large (i.e. F is not too much greater than 1). They also tend to be very dispersive except under the same conditions. However, for problems of the type considered here, these properties are adequate, and the restrictions are not seriously limiting, as discussed further in section (5), below.

(a) Analysis of the prediction step alone

The iterative convergence of the scheme is discussed in section 3(b), below. However, both the advantages and disadvantages of the method can be understood more easily (and for the worst case) by considering repeated time-stepping using the initial approximate prediction step alone (without any iterative refinement). This helps to establish the context for the method, and is discussed here first. The simplest fully implicit predictor equation (7) may be re-written

$$C_o' = [\ C_o + d_l\, C_l + d_r\, C_r + d_u\, C_u + d_d\, C_d + 0.5\ \{a_l\, C_l - a_r\, C_r + a_u\, C_u - a_d\, C_d\ \}$$

$$+ 0.5\ \{\ w_l\, a_l\, C_l + w_r\, a_r\, C_r + w_u\, a_u\, C_u + w_d\, a_d\, C_d\ \}]\ /\ [\ 1 + F\ ] \tag{8}$$

For a non-divergent flow-field, we always have $(a_l - a_r + a_u - a_d) =0$, and thus for a constant concentration field, $C=const$ everywhere, the first term in curly braces { } on the RHS of equation (8) is identically zero, and we recover $C_o' = C_o$, which is a necessary (but not sufficient) condition for a conservative scheme. In addition, for the special case of a diffusive-only system (where all values of a are zero) we obtain

$$C_o' = [\ C_o + d_l\, C_l + d_r\, C_r + d_u\, C_u + d_d\, C_d\ ] / [\ 1 + d_l + d_r + d_u + d_d\ ] \tag{9}$$

Thus , in this case the predicted future values of C are just weighted averages of the central and adjacent past values, with non-negative weights, which is sufficient for the step to be sign-preserving, and total variation diminishing (TVD), and therefore unconditionally stable for arbitrarily large Dt (and thus large values of d, i.e. $d >> 0.5$). In the limit as $d \rightarrow$ infinity, the process actually becomes a simple relaxation iteration, and for finite d is seen to be a successive under-relaxation scheme. It is therefore to be expected that repeated use of this initial step alone would be rather slowly convergent in time to a steady-state (slower than an SOR method, for example), but this is not important, since it retains the explicit connection with time-stepping, and it therefore allows for coupling with other time-dependent processes (such as those due to biological, geochemical and other physical processes).

When the advective terms are finite, the analysis is more complex. However, for upstream differencing, w=sign(a), and we may write a'=abs(a) for all advection into the cell (i.e. with the convention used here, for a_l and $a_u > 0$, and a_r and $a_d < 0$), and a'=0 otherwise, we obtain

$$C_o' = [\ C_o + d_l\, C_l + d_r\, C_r + d_u\, C_u + d_d\, C_d + \{a_l'\, C_l + a_r'\, C_r + a_u'\, C_u + a_d'\, C_d\ }] \ ...$$

$$/\ [\ 1 + d_l + d_r + d_u + d_d + \{a_l' + a_r' + a_u' + a_d'\ \}\ ] \tag{10}$$

Since the a' terms are non-negative (by definition), this also leads to a weighted average of previous central and adjacent values, with non-negative weights. For the purely advective case

(with all d terms equal to zero), and for large Dt (and thus large a), C_o' now reduces to a weighted average of the upstream values only, as would be expected. The use of upstream differencing would however lead (as usual) to high spurious numerical diffusion, and it is preferable to use a weighted upstream/centred scheme with weights dependent on the Peclet numbers $Pe=(a/d)$ to reduce numerical diffusion (Leith 1965; Roache 1972; Roache 1976; Fiadeiro and Veronis 1977)

The choice of weighting however also affects the stability of the scheme. To carry out a von Neumann analysis of the time-stepping stability it is easier to consider just the one-dimensional case, for which non-divergent flow implies $a_l=a_r$ ($=a$), and for simplicity we also assume $d_l=d_r$ ($=d$). For this one-dimensional case we have

$$C_o' = [\ C_o + d\ (C_l + C_r)\ + 0.5\ a\ (C_l - C_r)\ \} + 0.5\ w\ a\ (C_l + C_r\ )\ ]\ /\ [\ 1 + 2\ d + w\ a\ ] \tag{11}$$

Setting $C_l = C_o \exp(+\mathrm{i}\ f)$ and $C_r = C_o \exp(-\mathrm{i}\ f)$, as usual, and noting that in one dimension $F = (2d + wa)$, the gain G for a single time-step, and for phase shift f is

$$G(f) = C_o'/C_o = [\ 1 + F \cos(f)\ + \mathrm{i}\ a \sin(f)\ ]\ /\ [\ 1 + F\ ] \tag{12}$$

For stability we require $\mathrm{abs}(G) \le 1$, and thus, for large timesteps, and therefore large values of a, d and F, this implies $\mathrm{abs}(a) \le F = (2\ d + w\ a)$.

For the more general case of a single-step, partially implicit scheme (with $0<\alpha<1$), it can be shown in the same way that the gain G' is

$$G'(f) = [\ 1 + F\ \{\cos(f) - b\}\ + \mathrm{i}\ a \sin(f)\ ]\ /\ [\ 1 + aF\ ] \tag{12a}$$

and that the requirement that $\mathrm{abs}(G') \le 1$ leads to the same inequality, irrespective of the value of α. For upstream differences (i.e. $w=\mathrm{sign}(a)$) this inequality is satisfied for all Pe, whilst for central differences ($w=0$) it is only satisfied for $\mathrm{abs}(Pe) < 2$. In fact we require $w>0$ for $Pe<2$, and $w>1-2/\mathrm{abs}(Pe)$ for $Pe>2$. It is therefore easily shown that it is also satisfied for all Pe for various other choices of w, including

$$w = Pe\ /\ (2 + \mathrm{abs}(Pe)) \tag{13}$$

This is a simple approximation to the weighting proposed by (Fiadeiro and Veronis 1977) in order to eliminate numerical diffusion in the steady-state, and it is asymptotically equal to it for large Pe (for which the problem is most serious). This choice of upstream/centred weighting thus leads to an implicit predictor step which has only moderate spurious numerical diffusion, and is stable (although perhaps only marginally so) for any time-step, and for any Peclet number.

Unfortunately, however, these predictor equations are not normally conservative (N. Edwards, pers. comm.) unless and until a precise steady-state has been reached, because the fluxes in opposite directions across each cell boundary are actually computed using different values for the concentration inside and outside the cell (one future and one past value in one direction, and vice versa in the other). The errors so created cancel for the diffusion-only case, when advection is zero (see above), but not otherwise. The issue of conservation is discussed further below.

(b) Iteration stability analysis of the full iterative scheme

The time-stepping application of a simple single-step approximation to the implicit method, as discussed above, would therefore be unsatisfactory despite its attractive stability properties, for the many applications in which conservation of material properties is important. However, if and when it converges, the full iterative solution of equation (6) **is** conservative in the limit, since the fluxes in opposite directions are computed using concentration values which are all evaluated at the same (intermediate) time-level. Repeated applications of equation (6) may thus be used to obtain any desired accuracy and degree of conservation, at the expense of additional computations, provided of course that it converges to the correct solution. Moreover, as discussed above, the iterative scheme is just another method for the solution of the equations generated by a conventional implicit method, and thus if the iteration converges, the usual time-stepping stability and other properties of the implicit method must also apply to the final iterative solution.

We must therefore consider the stability and convergence of the solution of equation (6) under simple iteration. Assume that a solution $\overline{C}$ exists, for the iteration on C' for a new time-level, and write e for the deviation of any current estimate of C from this final solution, i.e $e = C - \overline{C}$. Then, for the intermediate time level denoted by *, and defining $b=(1-a)$ for brevity as before,

$$e^* = C^* - \overline{C} = aC' + bC - (a\overline{C} + bC) = a(C' - \overline{C}) = ae \tag{14}$$

Then, using the superscript n (etc) to indicate the time-level where necessary, since $\overline{C}$ is by definition a solution of equation (6), we may substitute $\overline{C}$ for C' in both sides of that equation, and subtract them from the original to give

$$e_o^{n+1} = \left[d_l e_l^* + d_r e_r^* + 0.5\left\{(1+w_l)a_l e_l^* - (1-w_r)a_r e_r^*\right\}\right] / \left[1+aF\right] \tag{15}$$

since the two terms involving $C_o(1-bF)$ cancel out. Then, using equation (14),

$$e_o^{n+1} = \frac{a}{\left[1+aF\right]}\left[d_l e_l^n + d_r e_r^n + 0.5\left\{(1+w_l)a_l e_l^n - (1-w_r)a_r e_r^n\right\}\right] \tag{16}$$

To evaluate the stability and convergence of this iteration, we use a von Neumann analysis and write as usual $H(f) = e^{n+1}/e^n$, with $e_l = e^{if} e_o$ and $e_r = e^{-if} e_o$. Then, remembering that non-divergent flow in one dimension implies $a_l=a_r$ $(=a)$, and again assuming for simplicity that $d_l=d_r$ $(=d)$, so that we again also have $F=2d+wa$, we find

$$\begin{aligned} H(f) &= \frac{a}{\left[1+aF\right]}\left[2d\cos f + 0.5\left\{2ia\sin f + 2wa\cos f\right\}\right] \\ &= a\,\frac{\left[F\cos f + ia\sin f\right]}{\left[1+aF\right]} \end{aligned} \tag{17}$$

and for stable convergence we require $\text{abs}\{H(f)\} \le 1$ for all f.

It should be noted that this condition for stable convergence of the iteration is remarkably closely related to that derived above, for the time-stepping convergence of the simple single-step approximate formulation. Indeed, for large time steps, and thus large values of *a*, *d* and *F* (which is both the worst case and that of interest to us) we recover just the same condition as before, i.e.

$$\text{abs}(a) \leq F = (2d + wa) \tag{18}$$

for all phase angles f . For moderate *F*, the iteration gain depends on αF , so that a small value of α will allow a larger value of *F* and thus a larger timestep for the same gain, but for very large F values, we find that a cancels, so that ultimately the iteration convergence is not affected by the degree of implicitness.

Using the same methods, it can be shown that the time-stepping gain for the fully converged iterative scheme is

$$G''(f) = \frac{\left[1 - b\left\{F(1-\cos f) + ia\sin f\right\}\right]}{\left[1 + a\left\{F(1-\cos f) - ia\sin f\right\}\right]} \tag{19}$$

which is just a generalisation of the usual result for a partially implicit scheme for the advection-diffusion system. The modulus of this expression approximates to $\pm(1-\alpha)/\alpha$ for large values of *F* and for a wide range of f , so that a value of α in the upper part of the range 0.5 to 1 would be desirable for stability. However, for rapid iterative convergence equation (17) indicates that one should choose α as small as possible, so as to reduce the product αF and thus the value of H. A compromise between these conflicting requirements implies that α should be somewhere in the range 0.6 to 0.8, since a value too close to 0.5, which would achieve second-order accuracy, will provide only marginal time-stepping stability (although this may be adequate).

For a practical choice of α (say, $\alpha \approx 0.6$), and a time-step such that $F \approx 10$, the iteration convergence will be moderately slow ($H \approx 6/7 \approx 0.85$), although this is a respectable rate in the general context of Jacobi relaxation methods, and the time-stepping will also be only moderately stable ($G'' \approx 0.4/0.6 \approx 0.7$). Thus, for very long time-steps, more iterations will be needed to achieve adequate convergence (and thus to preserve adequate conservation of material properties). This could actually destroy any computational advantage of the method, other than its robustness (i.e. unconditional stability with an appropriate choice of weighting *w(Pe)*). Based on practical experience, briefly described in section (4), this problem appears not to be severe, for reasons which are discussed below. Finally, it should be noted that **any** iterative implicit scheme is likely to be non-conservative, and that this potential problem is not special to this simple scheme.

(c) Time-step selecton

With a robust and stable system such as this, it becomes possible to select the time step according to criteria other than numerical stability. One may therefore also implement automatic and variable time-step selection, and allow timesteps to be increased as much as required, as and when a steady state is approached. In the implementation used for the practical tests described below, the time-step has been selected primarily in order to approach and maintain a target level for the maximum change of the concentration of any of the modelled properties across the domain at each time-step. By monitoring these changes, it is possible to increase (e.g. double) and decrease (e.g. halve) the time-step progressively, as required.

If this is done, the method will actually use short time-steps whenever significant transient adjustment is under way, and will use long timesteps only as steady-state (or transient equilibrium) conditions are approached. In this way the scheme is actually able to handle transient situations without difficulty, even though it is not designed for them, and generally has no significant advantage for them, except in those cases where the stability limits are much more restrictive than the limits set by the transient evolution. These may arise where diffusion and advection processes are rapid, but the forcing (including boundary conditions) is such that the tracer fields are rather uniform. It is notable that it is in these ostensibly "easy" cases that the computational penalty imposed by the stability limits is particularly expensive and frustrating.

Limiting the time-step in this way has the further advantage that it automatically also effectively limits the temporal truncation errors involved in the integration, provided that the maximum change per time-step is chosen to be small. A target level of a maximum change (across the domain) of a few percent of the expected range of each property has been found to be satisfactory in practice. This also automatically limits the conservation errors, since these arise only when the concentration field is changing, and are proportional to the magnitude of the change at any point between adjacent time levels, as discussed in the next section.

In order to limit the distortion of the time-scale which is a general feature of implicit methods, as discussed below in section (3e), the time-step is however in practice not allowed to increase without limit as the steady-state is approached. It can capped at a value which limits the volume replacement factor (F) at a level which permits worthwhile acceleration of the calculation, whilst moderating the distortion. Maximum values of F of about 30 have been found to work well in practice. Limiting the value of F somewhat also serves to prevent excessively slow iteration convergence, as discussed above.

The same automatic time-step selection method can of course also conveniently be used for the explicit scheme (obtained by setting $\alpha=0$), and in this case it may require time-steps smaller than the stability limit during significant transient evolution, increasing automatically just to the stability limit as transients decline (because above this limit the instabilities generate unwanted and limiting changes in the property fields).

(d) Conservation of material properties

The method proposed is not absolutely conservative, which could be a serious problem for the modelling of material properties, especially those for which external sources and sinks are small, so that residence times are long, for example salinity in the ocean. However, it is worth noting that whilst strict conservation is a desirable property of a numerical scheme, it is not essential. Indeed, since any numerical calculation can only achieve conservation to machine precision, it is never achievable in practice. For any problem there must be some acceptable level of non-conservation, although this may be quite difficult to determine in practice. Since non-conservation effectively generates spurious sources and sinks, this acceptable level will be greater when the true sources and sinks of the properties are large, compared with the magnitude of the interior diffusive and advective fluxes (whose combined divergence they must balance in the steady-state), and *vice versa.*

However, we may note that the non-conservation arises only (a) when advection is important (since the scheme is conservative for diffusion only), and (b) the property distribution is non-uniform, and evolving (since it is due to the difference between the concentration field at adjacent

points at successive time-levels), and (c) when the iterative sequence is imperfectly converged. It may therefore be reduced by various means, including limiting the change in the concentration fields (as discussed above), choosing parameters which permit rapid iterative convergence, and performing enough iterations per timestep. Constructing a suitable measure of non-conservation, and determining a practical limit for it, and therefore determining the necessary number of iterations, requires further work. For the present I can report only that satisfactory results, as judged by comparison with use of the explicit scheme, seem to be obtained with a few (two to four) iterations per time-step. Given that for large F, as required for long time-steps, convergence may become rather slow, with abs (H) > 0.7, this is at first sight rather surprising. The effects of limiting the time-step (and thus F) during transients, and also choosing α rather close to 0.5, to provide both a more accurate scheme and to improve the convergence rate by reducing the product αF, all serve to ameliorate the problem.

However, for any real problem in more than one dimension, the advective field is non-uniform. The regions of high F are therefore special and likely to be somewhat localised (and this is of course also true if the diffusion coefficients are spatially variable). The problems of slow convergence and non-conservation are therefore likely to be also localised in the same regions, which are actually those where rapid diffusive and advective processes tend to homogenise the spatial distribution. It therefore seems possible that some sort of *de facto* cancellation occurs, such that where the problem is potentially large, adequate convergence suppresses it, whereas where convergence is poor, the problem is not too large anyway. For the moment however this remains mainly speculation.

(e) Time-scale distortion

A further potential disadvantage of this and other implicit schemes is distortion of the time-scale, which is a manifestation of temporal truncation errors, and arises because in a given (and nominally long) time-step, the scheme actually only moves as much material as would an exact scheme with a shorter time-step. Thus, although the integration appears to have been carried out for a very long time, the results actually correspond to reality for some lesser time. The diffusive and advective processes in effect proceed too slowly, and the labels on the time axis are optimistic. This problem is a general feature of implicit methods, and iterative approximations to them, including the DuFort-Frankel method of which this method is an extension (see (Roache 1976), p63 for a good discussion of this phenomenon). The problem is quite fundamental, since it is precisely these reduced fluxes of material which confer the desirable stability properties of such schemes. Conversely, it should be noted that explicit schemes generally move too **much** material (because they do not adjust the fluxes to or from any cell for the changing concentrations in adjacent cells, within a single time-step, whilst implicit schemes over-compensate). It is this which causes explicit methods to over-shoot, producing firstly infeasible negative concentrations, and eventually actual instability. These features can easily be understood by considering the simple case of radioactive decay in a zero-dimensional (box) model. This shows that the fully implicit method generates only hyperbolic rather than exponential approach to an equilibrium in a single time-step. The analysis of this simple system also easily confirms that the semi-implicit scheme is second-order accurate for infinitesimal time-steps, and has good accuracy up to normalised time-steps (lDt) of about 1.5. However for longer finite time-steps, values of α which are larger than 0.5 give superior accuracy. Indeed for larger normalised timesteps (analogous to large values of F in the diffusive/advective system), progressively larger values of α [actually $\alpha \approx lDt/(1+lDt)$] yield better accuracy. One might therefore wish to use (say) $\alpha \approx 0.9$ for F of the order of 10, but

such a large value would degrade the iteration convergence, and a compromise is needed. Since this problem primarily arises in those regions where F is large, so that fluxes (and the imbalance thereof) tend to be large and things change fast, it seems to be preferable overall to use a value of α closer to 0.5, to maintain iteration convergence and conservation, and accept that rapid changes (transients) are not modelled very accurately, but proceed somewhat too slowly. Since the method is not primarily intended for model situations where transients are of the essence, this is a lesser price to pay. Finally, one should note that in those regions where values of F are not so large, the rate at which the processes proceed, and the transients evolve, is modelled more accurately. In fact, where $F \leq 1$, they are modelled at least as accurately, (but with error of the opposite sign), as for an explicit scheme.

4) Results of initial numerical tests

Preliminary tests with a two-dimensional implementation of the scheme, and with steady (fixed) flow-fields have demonstrated extremely good stability, using sequential time-step doubling as a steady-state tracer field is approached, based on a target maximum property change (of a few % of the full natural range of possible values) per time-step. The modelling of dynamically active properties is discussed below. Stability is indeed maintained even with Courant numbers (and/or diffusion numbers) of 100 or more in places, but in practice it is preferable to limit the time-step so that the maximum Volume Replacement Factor (F) is not more than about 30, to avoid excessive distortion of the time-scale, as discussed above.

The results have invariably been found to be indistinguishable from those of the equivalent explicit scheme for the same problems, run close to the stability limit, for both fully and partially implicit calculations. Using 3 or 4 iterations per time-step, the method has still been found to permit cpu time reductions of a factor of 10 or so, even where it is necessary to model a time-consuming "spin-up" transient (for which this method offers no advantage other than robustness, and where it incurs significant costs due to the iterations at each time-step). This arises because with automatic time-step selection the scheme tends to spend roughly the same cpu time computing the transient phase, with short time-steps, as it does computing the later more gradual approach to the steady-state, with much longer time-steps than the explicit method. The advantage would thus be greater when modelling systems for which any such transient phase is absent, or of short duration compared with the total integration time required.

Typical results for the classic "spiral" passive tracer problem with a steady specified flow, using α=0.5 and a bi-quadratic stream-function, are illustrated in Figure 1. These are indistinguishable from those for the fully implicit version, and also from those of the equivalent standard explicit scheme (not shown) The latter required a time-step 20 times smaller, and took 10 times longer to integrate for the same modelled period of time. For this calculation the explicit diffusion coefficient was set to a very small value, so the calculation is for a very large nominal Peclet number. In practice the weighting scheme effectively leads to upstream differencing in this situation, and thus to an effective cell Peclet number of 2 (so the Peclet number based on the model domain is 256).

Typical results for a problem with an active tracer (and thus a non-steady flow) are given in Figures 2 and 3, for the natural convection due to buoyancy caused by variations of temperature in a room with hot and cold walls. These results are also indistinguishable from the equivalent semi-implicit and explicit computations, but did not achieve such a large cpu-time advantage,

because the time-step was in this case limited in practice by an incipient instability. This has been observed in several problems when the Peclet number is large, but only with dynamically active tracers.

Further investigation has shown that this instability is in fact a finite amplitude computational mode (typically of wave-number 2 per unit cell dimension, i.e. checker-boarding, at the Nyquist limit) involving an interaction (over-compensation) between the evolution of the active tracer and the dynamics (i.e. the computed flow). It arises because of the sequential (rather than simultaneous) solution of the equations for the tracer field evolution, and the consequential flow field, which is used in this and indeed in most other methods for geophysical fluid dynamical problems. The mode, which only appears for long time-steps, and thus at high levels of F, which are inaccessible to an explicit method, is a form of process-splitting instability. I have found that it can usually be successfully suppressed by (possibly severe) *ad hoc* under-relaxation (see e.g. Anderson, 1995) of the computed flow at each time-step, at least for the buoyancy/frictional-geostrophic dynamics (Edwards, Willmott et al. 1998) used here, where the fluid has no inertial mass or viscosity (which may be sufficient to suppress it in other types of model).

Further numerical tests (not illustrated) have been carried out using the same 2D implementation, set up to represent the thermohaline circulation in an ocean basin, with similar results. It is again generally possible to achieve about one order of magnitude acceleration of the calculation, compared with the conventional explicit method, yielding results for the steady-state which are indistinguishable. The acceleration achievable is limited in practice either by the time-scale distortion of the transient phase, which exhibits the same qualitative features, whilst evolving somewhat too slowly, or by the process-splitting instability described above.

In addition to these two-dimensional calculations, the method has also been successfully implemented and used in one-dimensional calculations (James Annan, pers comm) and in a three-dimensional model (Neil Edwards, pers comm), allowing substantial (order of magnitude) savings of cpu-time in both cases. In the three-dimensional case under-relaxation was also necessary to suppress a computational mode (very likely of the same type as discussed above).

5) Discussion & Conclusions

The numerical scheme proposed here provides a reasonably effective and practicable solution to one of the most serious and long-standing limitations to the very long integrations of advective-diffusive systems, which are subject to slowly varying forcing, in an efficient and relatively easily implemented way. It is particularly well suited to the long-term modelling of biogeochemical processes in fluid systems (including box models, for which the equations given can easily be adapted). An efficient method for such long integrations using moderately well-resolved models in 2 or 3 dimensions will be extremely helpful for the development of intermediate complexity models of oceanic and atmospheric processes in Earth System Models for climate studies.

However, there is of course "no free lunch", and the method has a number of potential disadvantages, including inaccurate modelling of transients, imperfect conservation of material properties, and distortion of the time-scale of the modelled results. This sounds like a fairly serious list of problems, but as discussed all of them can be ameliorated substantially, and they do not present serious problems to the successful use of the method within its domain of applicability. The inaccuracy of transients is not a problem as long as one only uses long time-steps as and when a steady-state is approached, for example as suggested here by choosing the

time-step so as to limit the maximum change in the concentration fields. If this is done the magnitudes of the inaccuracy and of non-conservation become small, and only a few iterations (similar to corrector steps) are required to reduce the imperfections to an acceptable level. Moreover, the non-conservation errors may not be too serious in practice, since these arise mainly in regions where diffusion and advection are rapid, and tracer concentrations tend to be homogenised. Indeed, for the accurate computation of a quasi-steady tracer field we actually really need the most accurate computations where these processes are **slow**, thus permitting large concentration gradients to evolve, since it is these which create the main features of the concentration field. It should be noted incidentally that some of the other acceleration methods mentioned in the introduction also cause some degree of non-conservation.

A further possible problem is that the number of iterations (N) needed at each time-step to achieve adequate conservation increases at high levels of the volume replacement factor, and eventually becomes proportional to it, thus potentially negating any computational advantage. However, the constant of proportionality depends on the magnitude of abs(H) required, and provided this is not too small, say not less than 0.7, and $\alpha \approx 0.5$, then $N<0.2F$, and no more than 4 iterations are required even for F as high as 20. This conjecture is supported by the simulations conducted so far, but further investigation of this point would be desirable. Even if this does prove to be a significant problem, there seems to be no reason why the iteration number should not be varied across the domain, so that more iterations are only applied where necessary, thus making the punishment fit the crime, and avoiding any unnecessary computational burden.

The final complication is that, in common with other implicit methods, care must be taken to use an implicit treatment of the boundary conditions wherever necessary, lest these generate instabilities which have been suppressed in the interior. This has not so far been found to be a problem for Dirichlet boundary conditions, and should be straightforward for von Neumann and mixed boundary conditions too. Radioactive decay, and other transformations such as those due to chemical reactions, should also be treated implicitly, so far as possible.

The method is also suitable for modelling convective adjustment processes without requiring special coding for them, since it handles very high diffusivities without difficulty, and wherever significant up-welling or down-welling develops it automatically tends towards an upstream replacement algorithm. Whilst presumably it would be somewhat less effective than special algorithms for this process (Marotzke 1991), it should be competitive with a straightforward implicit treatment, and it has the advantage over these methods that no special provision needs to be made for invoking it. It could also presumably be used for the vorticity transport equations in both viscid and inviscid systems (and with a minor modification could even be used as a Poisson equation solver for stream-function calculations), thus conceivably allowing the same acceleration technique to be applied to the solution of the dynamical equations.

It should be noted that the method inherently involves multiple steps (for each time increment), and that while pure upstream differencing is permitted by the stability criteria, its use has not been advocated here because of the high resultant numerical diffusion. The use of weighted upstream differencing serves to reduce numerical diffusion to some extent. However, it would be interesting to determine whether this iterative method could actually be combined with schemes such as the MPDATA method (Smolarkiewicz and Margolini 1998) which also involve upstream differencing, but use a multi-step algorithm to suppress numerical diffusion very effectively, thus allowing accurate calculations even for very high Peclet numbers.

Whilst the method is not intended or designed for seriously unsteady systems (such as those with strong time-dependent forcing), it is actually able to handle transients without particular difficulties, by using the time-step limitation method described. It seems possible that in some circumstances its robustness might also permit it to be used with advantage, even in such transient situations. Although the accuracy is potentially reduced by non-conservation and timescale distortion, these are unlikely to be serious while the timesteps are small, such that F <1. With a fixed number of several iterations per time-step there would be a computational penalty in such circumstances, but if the number of iterations was made variable, so that multiple iterations were only used where and when necessary, this could be overcome. The scheme could then permit Courant (and/or diffusion) numbers to exceed unity in some restricted regions, and thus allow calculations with (say) an average volume replacement factor of the order of one, rather than requiring the maximum volume replacement factor to be always less than one. This would be at the expense of higher truncation and conservation errors in these limited regions. Practical experience, as described above, for simulations which *de facto* involve an unsteady (spin-up) phase, suggests that this is possible. Whether this would be of any practical utility in extreme cases, such as eddy-resolving models, is not clear.

The issue of non-conservation has been addressed above, but the treatment is as yet incomplete. Furthermore, the consistency and dispersion properties of the scheme have not been considered in any detail, as it is assumed that these would be approximately the same as for conventional implicit methods, provided the iteration scheme is adequately converged. In fact, because the iterative scheme can propagate a disturbance only over N grid-points per time-step, rather than over the whole domain as in a normal implicit scheme, it is possible that it could be somewhat less dispersive which may be an incidental advantage. Although suitable choices have been suggested for the various parameters (α, N, w, F_{max}, maximum fractional change per time-step) which control the precise behaviour of the scheme, these are not critically important, and no serious attempt has been made to optimise these. Further work on the optimisation of the scheme, on the time-scale distortion incurred, and also particularly on the evaluation of the levels of non-conservation which occur, and the acceptability of them, would therefore be desirable. The iterative method used is essentially Jacobi relaxation, which is generally slowly convergent, and it is therefore very likely that other more effective iteration methods would further improve the performance. Some form of conjugate gradient method or even the GMRES algorithm (see (Barrett, Berry et al. 1994; Saad and van der Vorst 2000)) should be very suitable, although it is not certain that the trade-off between cost and effectiveness would be favourable. A simpler alternative would of course be just to use Gauss-Seidel relaxation or SOR, although this has not yet been tried. This could exacerbate the non-conservation problem, and would make the method less suitable for parallelisation, although use of a red-black variant (see e.g. (Briggs 1987)) would avert this disadvantage.

The simple iterative implicit method proposed should be a useful addition to the geophysical fluid modelling tool-kit, enabling worthwile acceleration of computations in suitable situations for very moderate implementation effort . An important feature is the use of some upstream weighting of the advection terms which ensures iterative convergence and stability, even for cell Peclet numbers much greater than 2, and also serves to reduce spurious numerical diffusion. The scheme is much faster than an explicit scheme run near its stability limit, and also much more robust, and in some circumstances it might even be worth using it for this reason alone. With appropriate design, it can be coded as a simple alternative to an explicit scheme (involving only the re-setting of a few parameters, and the execution of a few extra lines of alternative code), so that it is

possible to use it, and also to easily check the results against the explicit method, as and when required. Unlike (say) a quasi-Lagrangian, multigrid or ADI method, it involves essentially the same computational structures and sequences as an explicit method, and in some cases (depending on existing program design) it may therefore be possible to implement it as a fairly straightforward modification of an existing explicit implementation. A numerical scheme such as this can of course be used in addition to the special methods for dealing with fast time scales in ocean models discussed in the introduction, if appropriate. It should be useful for a wide range of non-eddy-resolving models, especially those used as components of intermediate complexity climate models, and also for example for the accelerated spin-up of the large-scale deep ocean circulation and property distributions in conventional General Circulation Models (with eddies temporarily suppressed), under different (e.g. past and future) climate conditions.

Acknowledgements

I am grateful to Peter Killworth, Sandy Gourlay and Neil Edwards for several discussions on the merits (and otherwise) of the scheme, and to Neil Edwards and James Annan for implementing the method in other models, as discussed in the text. The critical comments of several referees have assisted me in placing the method much more precisely in both its ocean modelling and its numerical context.

References

Ames, W. F. (1969). Numerical Methods for Partial Differential Equations.

Anderson, J. D. (1995). Computational Fluid Dynamics : The basics with applications. New York, McGraw-Hill.

Barrett, R., M. Berry, et al. (1994). Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods. Philadelphia, PA, SIAM.

Brandt, A. (1977). "Multi-level adaptive solutions to boundary value problems." Mathematics of Computation **31**: 333-390.

Briggs, W. (1987). A Multigrid Tutorial. Philadelphia, PA, SIAM.

Bryan, K. (1984). "Accelerating the convergence to equilibrium of ocean-climate models." Journal of Physical Oceanography **14**: 666-673.

Danabasoglu, G., J. C. McWilliams, et al. (1996). "Approach to equilibrium in acclerated global oceanic models." Journal of Climate **9**: 1092-1110.

DuFort, E. C. and S. P. Frankel (1953). "Stability conditions in the numerical treatment of parabolic differential equations." Mathematical Tables & Other Aids to Computation **7**: 135-152.

Edwards, N. R., A. J. Willmott, et al. (1998). "On the role of topography and wind stress on the stability of the thermohaline circulation." Journal of Physical Oceanography **28**: 756-778.

Fiadeiro, M. and G. Veronis (1977). "On weighted mean schemes for the finite difference approximation to the advection-diffusion equation." Tellus **29**: 512-522.

Fletcher, C. A. J. (1991). Computational Techniques for Fluid Dynamics, Volume 1. Berlin, Springer.

Griffies, S. M. and R. Pacanowski (2001). "Tracer conservation with an explicit free surface method for z-coordinate ocean models." Monthly Weather Review **129**: 1081-1098.

Hallberg, R. (2000). "Time integration of diapycnal diffusion and Richardson number- dependent mixing in isopycnal coordinate ocean models." Monthly Weather Review **128**(5): 1402-1419.

Kantha, L. H. and C. A. Clayson (2000). Numerical Models of Oceans & Oceanic Processes, Academic Press.

Killworth, P., J. M. Smith, et al. (1984). “Speeding up ocean circulation models.” Ocean Modelling(56): 1-4.

Klinger, B. (2000). “Acceleration of general circulation model convergence by exponential extrapolation.” Ocean Modelling **2**: 61-72.

Leith, C. E. (1965). “Numerical simulation of the Earth's atmosphere.” Methods in Computational Physics **4**: 1-28.

Leonard, B. P., A. P. Lock, et al. (1996). “Conservative explicit unrestricted-time-step multidimensional constancy-preserving advection schemes.” Monthly Weather Review **124**: 2588-2606.

Lin, S.-J. and R. B. Rood (1996). “Multidimensional Flux-Form Semi-Lagrangian Transport Schemes.” Monthly Weather Review **124**: 2046-2070.

Marotzke, J. (1991). “Influence of Convective Adjustment on the Stability of the Thermohaline Circulation.” Journal of Physical Oceanography **21**(6): 903-907.

Roache, P. J. (1972). “On Artifical Viscosity.” Journal of Computational Physics **10**: 351-366.

Roache, P. J. (1976). Computational Fluid Dynamics. Albuquerque, Hermosa Publishers.

Roache, P. J. (1998). Fundamentals of Computational Fluid Dynamics. Albuquerque, Hermosa Publishers.

Saad, Y. and M. H. Schultz (1986). “Gmres - a Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear-Systems.” Siam Journal on Scientific and Statistical Computing **7**(3): 856-869.

Saad, Y. and H. A. van der Vorst (2000). “Iterative solution of linear systems in the 20th century.” Journal of Computational and Applied Mathematics **123**(1-2): 1-33.

Smolarkiewicz, P. R. and L. G. Margolini (1998). “MPDATA : a finite-difference solver for geophysical flows.” Journal of Computational Physics **140**: 459-480.

Staniforth, A. and J. Cote (1991). “Semi-Lagrangian integration schemes for atmspheric models: a review.” Monthly Weather Review **119**: 2206-2223.

Tannehill, J. C., D. A. Anderson, et al. (1997). Computairl FluachMdchanics and Heat Transfer. London, Taylor & Francis.

Tong, C. and R. Tuminaro (1999). ML: a multilevel framework for parallel unstructured grid computations. Copper Mountain Conference on Multigrid Methods, Copper Mountain, Colorado.

Woznicki, Z. I. (2001). “On performance of SOR method for solving nonsymmetric linear systems.” Journal of Computairl and Applied Mathematics **137**(1): 145-176.

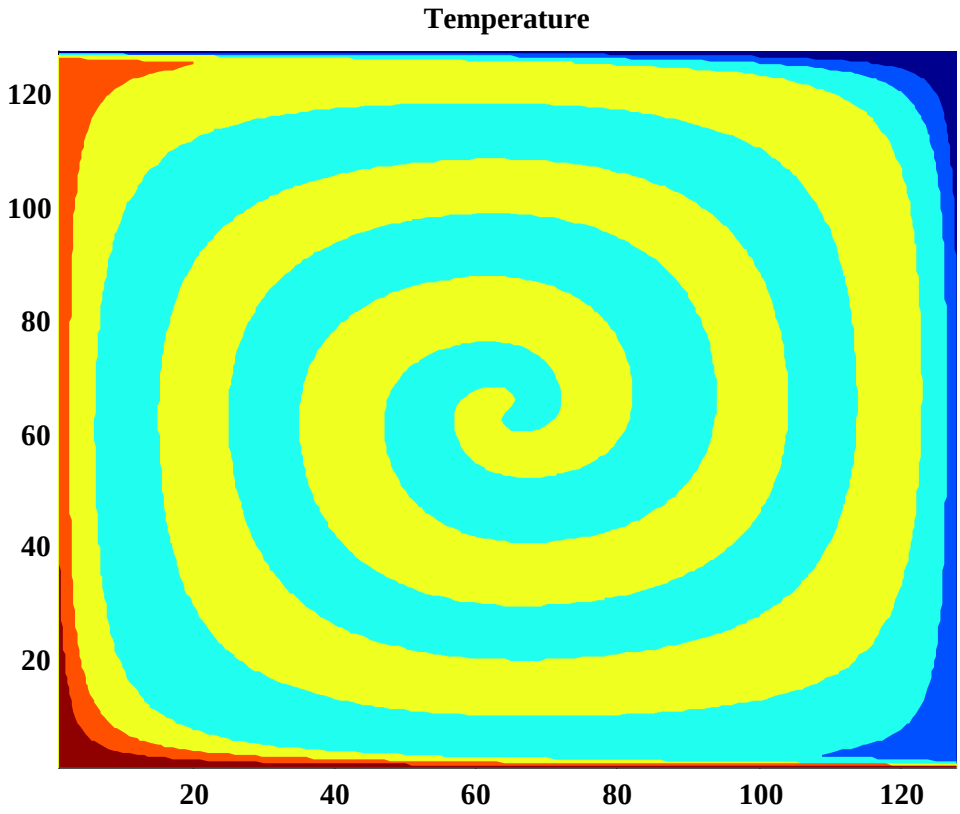

Figure 1 : Results for a classical steady flow diffusive-advective "spiral" problem, using the iterative implicit scheme with α=0.5, a fixed flow-field derived from a bi-quadratic stream-function (of arbitrary magnitude), and minimal explicit diffusion, on a 128 x 128 grid. Dirichlet boundary conditions are imposed at the upper ("cold") and lower ("hot") boundaries, (but there is no convection because the flow-field is fixed). The maximum Courant number attained is about 13. The minimal explicit diffusion implies a very large nominal Peclet number, so the weighting tends to upstream differencing in this limit. This introduces numerical diffusion, so the effective cell Peclet number is reduced to about 2. The results therefore correspond to a solution for a finite domain Peclet number of about 256 (2N for an N x N domain).

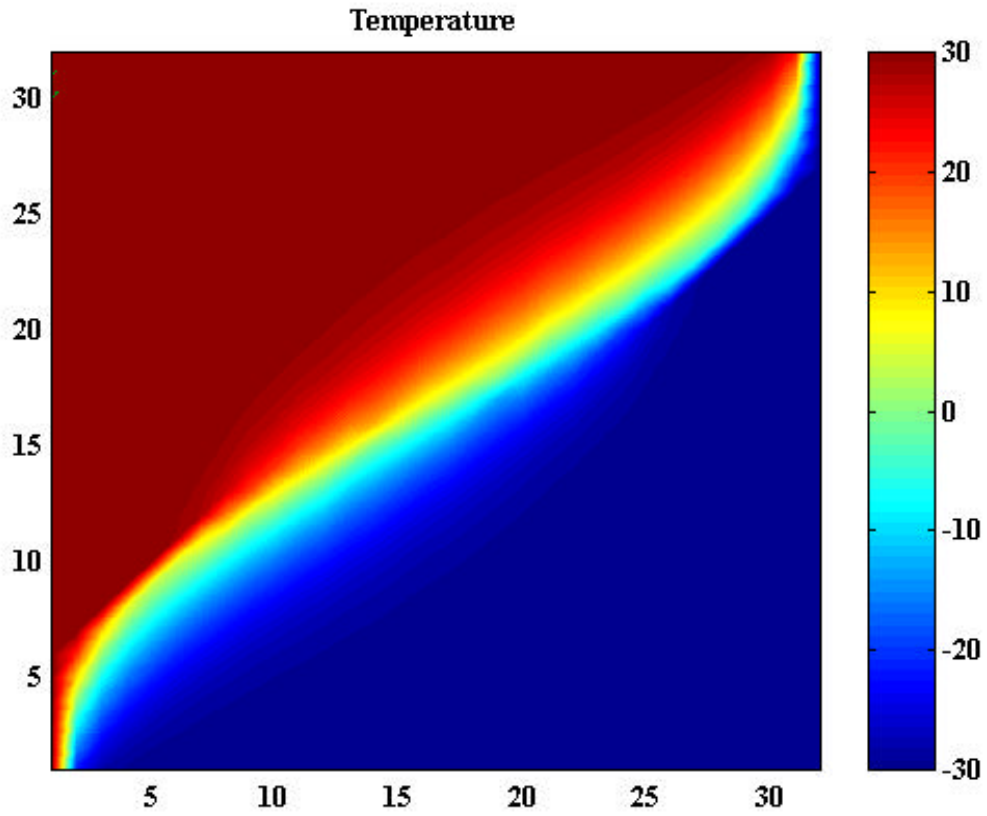

Figure 2 : Results (final temperature distribution) for a dynamical calculation of a natural convection problem, using the iterative implicit scheme and α=1. The flow is determined by buoyancy (with a linear equation of state) and Rayleigh friction only, and the explicit diffusivity is minimal both horizontally and vertically. The boundary conditions are Dirichlet, with a hot left-hand wall (+30C) and a cold right-hand wall (-30C). The grid is 32 x 32, so the effective domain Peclet number (see caption to Figure 1) is about 64.

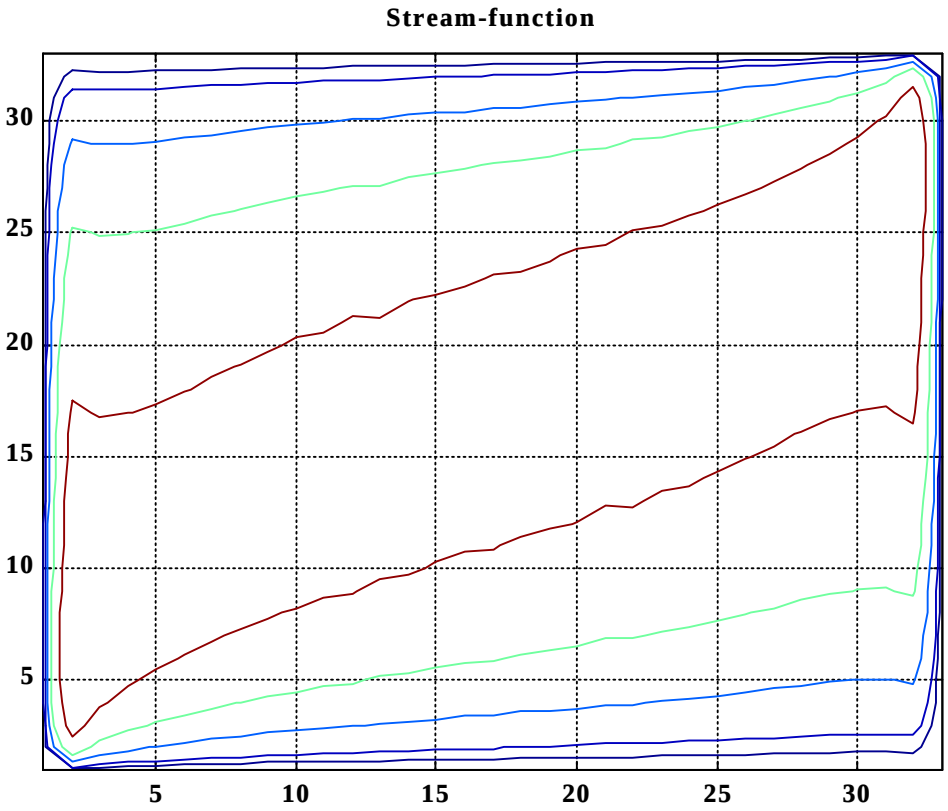

Figure 3 : Contours of the calculated steady-state stream-function (arbitrary units) computed for the natural convection problem of Figure 2, showing the very strong velocities which develop near the vertical boundaries, and weak interior flow. Results which are in all essential respects identical are also obtained using the corresponding explicit scheme, and α=0.5 and α=0.75.