225638

Avoiding BDF Stability Barriers in the MOL Solution of Advection-Dominated Problems

Alan C. Hindmarsh
Lawrence Livermore National Laboratory

Lawrence Livermore National Laboratory

Numerical Mathematics Group
Center for Computational Sciences & Engineering

UCRL-JC-116618, Rev. 1
December 1994

PREPRINT

This is a preprint of a paper submitted to Applied Numerical Mathematics. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

Avoiding BDF Stability Barriers in the MOL Solution of Advection-Dominated Problems*

Alan C. Hindmarsh
Numerical Mathematics Group, L-316
Lawrence Livermore National Laboratory
Livermore, CA 94551
email: alanh@llnl.gov

December 12, 1994

Abstract

Weakly damped oscillatory modes in ODE systems, such as are generated from advection-dominated PDE problems, frequently cause difficulty for BDF integrators, because of their absolute stability regions at high orders. We describe a procedure to detect the stability limit responsible for the difficulty. From it, we form a detection algorithm, which we have added to the order selection logic in an experimental version of the ODE solver VODE. Tests on an advection-diffusion problem demonstrate the effectiveness of this algorithm in lowering order to avoid the stability limit.

1 Introduction

The Method of Lines is a powerful and versatile approach to solving time-dependent partial differential equations, and integrators based on Backward Differentiation Formula (BDF) methods are probably the most common choice for the solution of the resulting ODE systems. But when the Method of Lines is applied to advection-dominated PDE problems (among others), the resulting ODE system is typically characterized by weakly damped but strongly oscillatory modes. When a BDF method of order three or more is used on such a system, the step size may be unduly limited, because the absolute stability region of the method omits a pair of lobes in the left half plane. See Figure 1, which shows the boundary of the absolute stability region for orders 3, 4, and 5 (with horizontal axis exaggerated). Linear stability theory requires $h\lambda$ to lie in the region for each eigenvalue λ of the system and each stepsize h used. Indeed, existing BDF codes very often perform poorly in this situation. They include some means of selecting the method order in a dynamic manner, based primarily on estimated local truncation errors, obtained from approximate scaled

*This work was performed under the auspices of the U.S Department of Energy by the Lawrence Livermore National Laboratory under contract W-7405-Eng-48.

derivatives of the solution of various orders. In some cases, there have been indirect attempts to use the scaled derivatives also to signal a reduction in order when the stepsize is stability-limited. But until now no BDF solver has included a direct attempt to detect the presence of this stability limit.

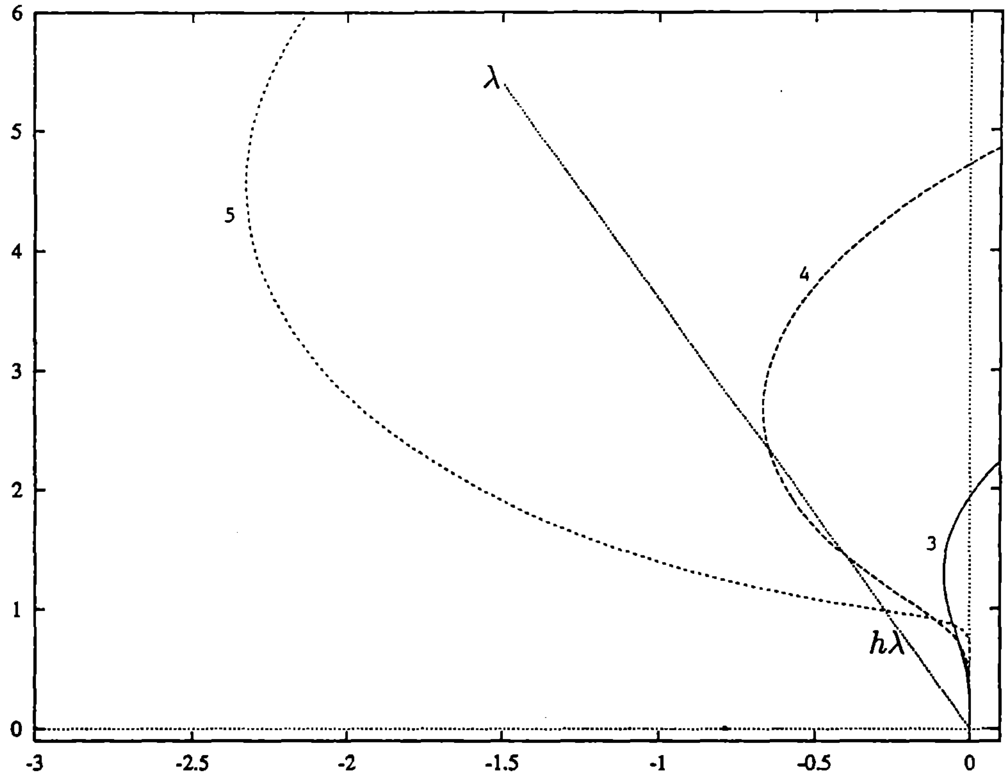

Figure 1: Boundary of the BDF absolute stability region S in the $h\lambda$ plane for orders 3, 4, and 5. For a given nonreal λ, the requirement $\lambda h \in S$ on stepsize h may unduly constrain h.

In an earlier paper [3], we analyze the scaled derivatives arising in the case where the numerical solution is completely dominated by a fixed, mildly damped, oscillatory mode. The analysis was done first for the complex scalar test equation, and then extended to the case of pair of real ODEs with nonreal eigenvalues. Elimination of the unknown parameters leads to a procedure for detecting the stability limit. The procedure is applied to the norms of the scaled derivatives of the solution, of orders $q-1$, q, and $q+1$ if the current order is q, and gives a value for the size of the dominant characteristic root. Stand-alone tests of the procedure on data from a simple ODE of size 2 were very encouraging.

At least three earlier authors have addressed the issue of detecting stability limits. In [4], Krogh gives a sign-change test designed to detect instability for certain PECE methods. It is simple and inexpensive, but relies on method-specific properties of the dominant extraneous characteristic roots. In [5], Skelboe derives approximate relations for the local truncation error in BDF methods, and poses a stability test using the norms of scaled derivatives. It is also simple, but it is based in large part on approximations and heuristics. In [1], Brenan, Campbell, and Petzold, using the order selection strategy in the Adams solver of Shampine and Gordon [6], describe a test for monotonicity of the scaled derivative norms. They use this in the BDF solver DASSL, with the idea that failure of the monotonicity test indicates that a stability boundary has been encountered. But it is also quite imprecise, in that the crossing of the stability boundary may or may not correspond to a violation of monotonicity. Even for the linear scalar model problem, in separate calculations, it has been observed that monoticity may hold outside the stability region, and it may fail to hold inside the region [3].

Here we use the procedure in [3] to form an algorithmic addition to the order selection scheme of the Adams/BDF solver VODE [2], whereby the order is reduced when a stability limitation is detected while using the BDF method. This algorithm has been installed in an experimental version of VODE. We use a simple advection-diffusion PDE in one space dimension to test the algorithm.

2 Stability Limit Detection

For the sake of completeness, we begin with a summary of the procedure for detecting the BDF stability limit. Details are available in [3].

To begin with, consider the BDF solution of the standard scalar test equation, $\dot{u} = \lambda u$, with a complex constant λ, $Re(\lambda) < 0$, at order $q \geq 3$ and fixed stepsize h. In the stability-limiting situation, the discrete solution values u_n are dominated by the n-th powers of a single characteristic root z_1 (a function of $h\lambda$) associated with the BDF method, with $|z_1| \approx 1$. (The absolute stability region is that in which $|z_j| \leq 1$ for all the characteristic roots.) The scaled derivatives $\sigma_n(k) = h^k u_n^{(k)}$, defined by way of the interpolating polynomial of degree q, are also dominated by z_1^n. Working with $\sigma_n(k)$ for $k = q-1, q, q+1$, one can easily eliminate the unknown proportionality constants and solve for $|z_1|$.

In the analogous BDF solution of a real linear system, where one or more of the eigenvalues is capable of causing a stability limit, we can assume that one of them, λ, will be such that $h\lambda$ is the first one to approach or cross the stability region boundary. In that case, the BDF solution is eventually dominated by the powers of a single characteristic root z_1 corresponding to $h\lambda$, so that the solution and its scaled derivatives $h^k y_n^{(k)}$ behave as if they were generated by a 2×2 linear system with eigenvalues λ and $\bar{\lambda}$. This situation was also analyzed in [3]. It is necessary (and reasonable) to assume that a weighted L_2 norm is used for all error-like vectors and that the weights on the two component are equal. It can then be shown that the squared norms $S_n(k) \equiv \|h^k y_n^{(k)}\|^2$ have the form

$$S_n(k) = G_k |z_1|^{2n} [1 + \gamma \cos(2n\theta + \nu_k)], \tag{1}$$

where $\theta = \arg(z_1)$, for constants G_k, ν_k, and γ (the last being independent of k).

In the special case that the system Jacobian is a normal matrix, the constant γ vanishes, and for each k the ratios $S_{n+1}(k)/S_n(k) = |z_1|^2$ give the quantity $|z_1|$ we seek. In the general case, the oscillatory terms in (1) can be eliminated by a procedure applied to the 15 values $S_m(k)$ $(n-4 \leq m \leq n; k = q-1, q, q+1)$. For each k, $R = |z_1|^2$ satisfies a quartic equation with coefficients formed from the five $S_n(k)$. The root can be found by elimination, which yields a linear equation because the quartics have no quadratic term. A set of auxiliary equations yields the G_k. In addition, if the assumption of a single dominant mode is correct, the value of R must also satisfy an equation

$$\frac{G_{q+1}}{G_q}\left(\frac{G_{q-1}}{G_q} - \frac{q^2-1}{4}\right) - 1 = \left(\frac{q-1}{2}\right)(1 - R^{-1}). \tag{2}$$

The function on the left side of (2) is a "barrier function" in the sense that its sign is precisely sensitive to the stability barrier, i.e. to whether $h\lambda$ is inside or outside the absolute stability region.

3 A Modified Order Selection Algorithm

A modified order selection algorithm is now possible for BDF solvers, by adding to the existing order selection rules (based on local errors) a test to detect the presence of a BDF stability limit. The following is a crude algorithm for this, denoted STALD: STAbility Limit Detection. It includes various consistency checks to verify the validity of the dominant mode model.

Algorithm: STALD

1. If both q and h have been constant for at least 5 steps, with $q \geq 3$, collect the 15 values of $S_m(k)$.
2. For each k, look at the variance of the four ratios $S_{m+1}(k)/S_m(k)$. If it is small for each k, get R from these ratios (the corresponding 2×2 Jacobian is normal or nearly normal). If the R values are not consistent, then exit. If a consistent R is found, go to step 4.
3. Form the three quartics $Q_k(R)$, and eliminate to get a tentative R. If the Q_k are dependent (elimination is impossible), exit. Evaluate the $Q_k(R)$ and do Newton corrections to improve R if necessary. If the new values of $Q_k(R)$ are not all small, exit.
4. For the given R, compute the three G_k, and the solution $R = R_B$ of Eqn. (2). If R_B disagrees with R, exit.
5. If $R \approx 1$ or $R > 1$, signal a reduction in order (the stability barrier has been reached or exceeded).

We have modified the VODE solver [2] to include the STALD algorithm. VODE contains both Adams and BDF methods, with the latter going up to order 5. The algorithm involves a number of heuristic (tuning) quantities, and these have been given tentative values, subject to change in light of further testing. The scaled derivative norms are easily available, because VODE uses a Nordsieck history array. For $k = q$ or $q - 1$, we need only multiply the k-th Nordsieck vector by $k!$, and for $k = q + 1$, a convenient multiple of the estimated local error norm (which is already computed) gives the required scaled derivative norm. The procedure for the stability limit determination, represented by steps 2–4 in the above algorithm, is carried out in a separate subroutine, and a flag returned to the central integration step routine. Then, if a positive determination was made, an order reduction is forced. The stepsize is reset to the value that would have been used on an order reduction by the existing criteria, i.e. a value based on the estimated local error at order $q - 1$.

The modified VODE solver was first tested on the 2×2 problem used in the stand-alone tests reported in [3], which has eigenvalues $-10 \pm 100i$. It performed very much as expected. In a typical case, at a point where the integration was proceeding at order 5, the algorithm forced an order reduction from 5 to 4, then shortly thereafter from 4 to 3. Following that there was a dramatic increase in step size, which did not occur in the integration by the unalterd solver.

4 Advection-Diffusion Tests

Here we test the STALD algorithm on a simple advection-diffusion PDE in one dimension with constant coefficients, namely

$$\partial u/\partial t = D\partial^2 u/\partial x^2 - V\partial u/\partial x, \tag{3}$$

on the unit interval in x. We pose a Dirichlet boundary condition $u = .5$ at the left and a Neumann boundary condition $u_x = 0$ at the right. For initial conditions, we pose a polynomial profile, peaked in the center and satisfying the boundary conditions: $u(0,x) = 1 - (2x-1)^2 + (2x-1)^4/2$.

We use a uniform mesh of $M+1$ points (spacing $\Delta x = 1/M$), and apply standard central differencing for the spatial derivatives. Other differencing schemes (such as biased upwind) may be as good or better, but any sufficiently accurate scheme will have eigenvalues near the imaginary axis, and it is only that feature that is relevant here.

We retain discrete values u_j of u at $x_j = j\Delta x$ with $j = 1, \ldots, M$, and represent the boundary conditions by taking $u_0 = .5$ in the ODE for u_1 and $u_{M+1} = u_{M-1}$ in the ODE for u_M. In the resulting ODE system $\dot{y} = Ay + b$ (where b comes only from the left boundary value), the tridiagonal matrix A involves the dimensionless coefficients $d = D/\Delta x^2$ and $a = V/2\Delta x$. The eigenvalues are complex when the grid Peclet number $P = a/d$ exceeds 1, and in that case the spectrum is given by $\lambda_j = 2d[-1 + i\sqrt{P^2-1}\cos\psi_j]$, where the ψ_j are the roots of $P\tan M\psi_j + \psi_j = 0$ $(0 < \psi_j < \pi)$. The two extreme eigenvalues have a slope of $|Im(\lambda)/Re(\lambda)| = \sqrt{P^2-1}\cos\psi_1$. For large M and P, this is approximately P. Thus with reference to the BDF stability regions, a stability limit exists at order 3 or more if $P > 14.4$, at order 4 or more if $P > 3.34$, and at order 5 if $P > 1.27$, roughly.

In our tests with VODE, we fix $V = 20$, and take three values of D and three values of M, as shown in Table 1. We integrate to $t_f = .25$ in all cases, which is well beyond the time $1/V = .05$ required for the profile to exit the domain. Thus we expect that stepsizes will be limited by accuracy requirements until roughly $t = .05$, but then should grow considerably, unless limited either by spurious oscillations or by the BDF stability limit. We use scalar tolerances RTOL = ATOL = 10^{-6} in all cases. For each case, Table 1 gives the total number of steps to completion for the unaltered VODE and the modified VODE + STALD algorithm. The step count is a reasonable cost measure for these problems, as other measures show about the same relative comparison between runs.

The modifed VODE results in a reduced number of steps in most cases, by ratios as large as 3.7. To see in more detail the beneficial effects of the STALD algorithm, the Figures 2 and 3 below show the order q and stepsize h used by the two versions of the solver as a function of time t, for the case $D = .005, M = 200$. In Figure 2(b), the asterisks mark the two order reductions forced by the STALD algorithm. Figure 3 shows the corresponding stepsizes, and their rapid growth following the order reductions with the modified VODE.

In some cases, however, the savings for the modified solver are small or nonexistent. In one case, where $P = 2.5$, this is simply because there is no BDF stability limit. In the others, the solution retains oscillations that limit the stepsize to the cell crossing time $\Delta t_c = \Delta x/V = 1/2a$. This value is approximately equal to the limiting step size when a BDF stability limit exists. The benefits of the STALD algorithm cannot be realized unless and until the stepsizes appropriate for accurate resolution exceed Δt_c.

Another way to avoid the stability limit problem with a solver like VODE is to limit the maximum method order to an appropriate value. So for a further comparison, the next two

V	D	M	P	Steps - VODE	Steps - modVODE
20	.01	100	10	577	399
20	.01	200	5	975	355
20	.01	400	2.5	549	549
20	.005	100	20	647	485
20	.005	200	10	1225	551
20	.005	400	5	1880	505
20	.002	100	50	762	874
20	.002	200	25	1653	1579
20	.002	400	12.5	2487	765

Table 1: Number of steps for VODE and Modified VODE on the advection-diffusion problem

figures show the result of running VODE (unaltered) on the same case, $D = .005, M = 200$, with the use of the optional input MAXORD (maximum order). Figure 4 shows order q vs t, and Figure 5 shows stepsize h vs t, each for the cases (a) MAXORD = 3, and (b) MAXORD = 2. The step counts were 667 in case (a) and 582 in case (b). Thus the use of MAXORD in this case, while giving much lower costs than the default input (1225 steps), still is not quite as efficient as the modified VODE + STALD (551 steps). In some other cases, however, the results using MAXORD were more efficient than those from VODE + STALD. But in any case, the cost is quite sensitive to the value of MAXORD, and in practice one rarely knows the correct optimal value of it.

References

[1] K. Brenan, S. Campbell and L. Petzold, *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*, Elsevier, New York, 1989.

[2] P. N. Brown, G. D. Byrne, and A. C. Hindmarsh, "VODE, a Variable-Coefficient ODE Solver," SIAM J. Sci. Stat. Comput. 10 (1989), pp. 1038-1051.

[3] A. C. Hindmarsh, "Detecting Stability Barriers in BDF Solvers," in *Computational Ordinary Differential Equations*, J. R. Cash and I. Gladwell, eds., Oxford University Press, Oxford, 1992, pp. 87-96.

[4] F. T. Krogh, "A Test for Instabaility in the Numerical Solution of Ordinary Differential Equations," J. A. C. M. 14 (1967), pp. 351-354.

[5] S. Skelboe, "The Control of Order and Steplength for Backward Differentiation Methods," BIT 17 (1977), pp. 91-107.

[6] L. F. Shampine and M. K. Gordon, *Computer Solution of Ordinary Differential Equations*, W. H. Freeman and Co., 1975.

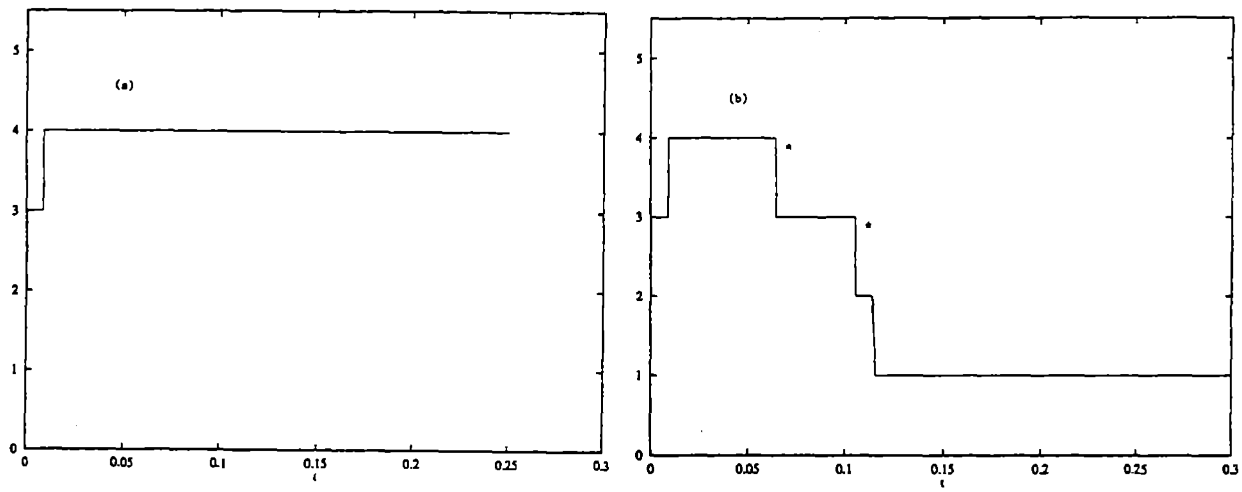

Figure 2: Order q vs t for $D = .005, M = 200$, with (a) unaltered VODE, and (b) modified VODE. The asterisks * mark order reductions forced by the STALD algorithm.

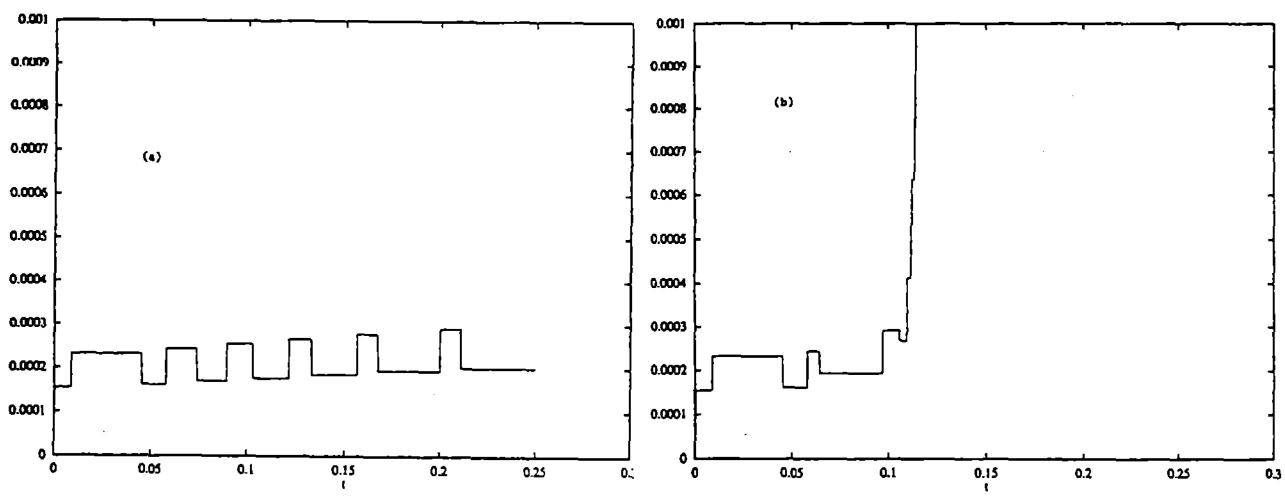

Figure 3: Stepsize h vs t for $D = .005, M = 200$, with (a) unaltered VODE, (b) modified VODE.

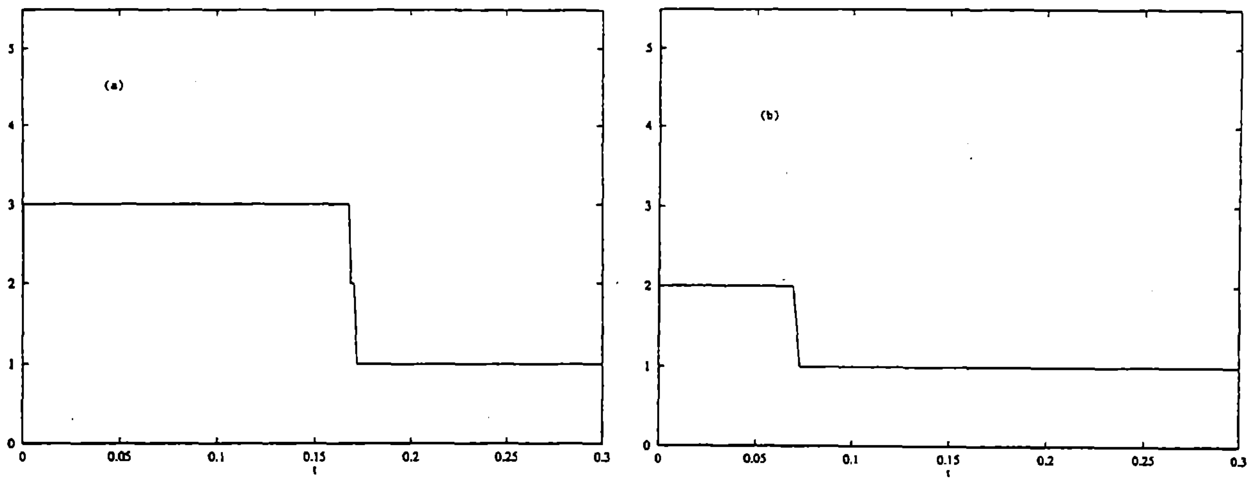

Figure 4: Order q vs t for $D = .005, M = 200$, unaltered VODE with specified maximum order MAXORD. The two cases are: (a) MAXORD = 3, (b) MAXORD = 2.

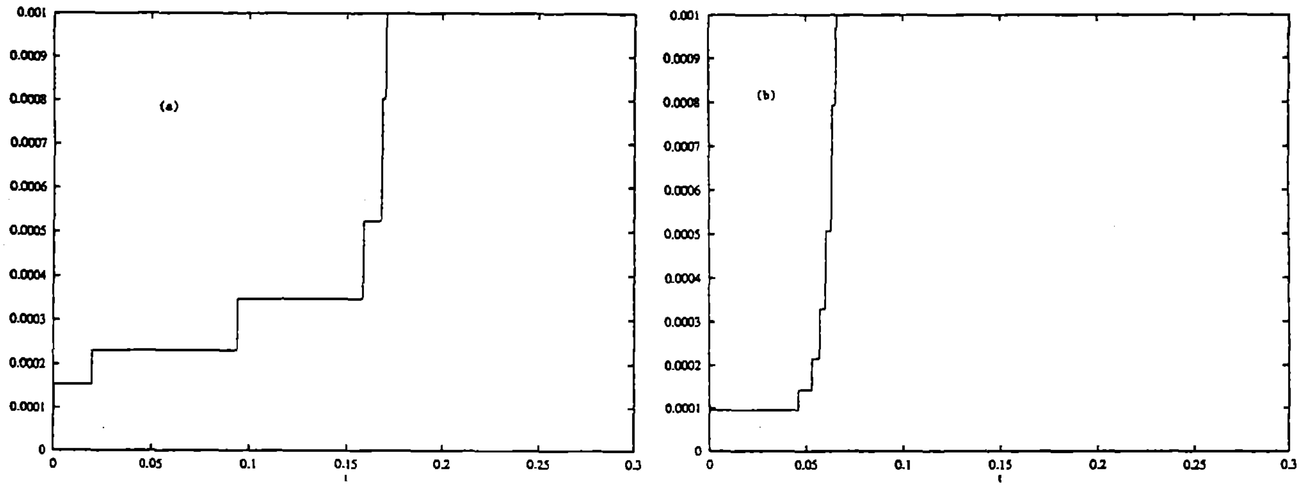

Figure 5: Stepsize h vs t for $D = .005, M = 200$, unaltered VODE with specified maximum order MAXORD. The two cases are: (a) MAXORD = 3, (b) MAXORD = 2.