

Stiff Systems of Kinetic Equations—A Practitioner's View*

ROBERT J. GELINAS

Lawrence Livermore Laboratory, University of California, Livermore, California 94550

Received March 10, 1971

An ever increasing number of current problems in applied science are described by sets of kinetic equations which may suffer from the difficulty known as stiffness when numerical solution of the equations is attempted. In this article the common causes and effects of stiffness are examined from a general user's point of view. Generally available measures for alleviating this type of difficulty are indicated. A specialized routine designed specifically for stiff systems (in that it enacts the measures which are discussed herein) is applied to a kinetic model of photochemical smog. Performance of the specialized routine is compared to a standard Adams predictor-corrector method as well as to the results obtained by another standard kinetics code which was originally applied to this model problem. Finally, the admissibility of quasi-steady state assumptions on selected components in this system model is examined.

INTRODUCTION

Systems of ordinary differential equations having drastically different time constants often impose severe " Δt limitations" upon numerical integration routines in order to insure stability. Such systems are known as time constant limited, or "stiff," systems. To date, the problem of stiffness has been of concern mostly to engineers and to applied mathematicians. There is, in fact, a growing literature of formal methods which are directed at the difficulties of stiff systems. It would be quite impossible to deal properly with even a representative sampling of numerous specific approaches which have been advanced and which may, or may not, be suitable for stiff systems. Therefore that task will not be attempted herein. Instead I would try to call attention to the fact that, even though the difficulties of stiffness have been apparent since the advent of finite difference methods, the problem seems to have been barely recognized yet in physics. Of greater concern is the fact that an ever increasing number of problems in applied physics and chemical kinetics require (or at least would benefit from) the use of stiff solution methods. Among this number are nonequilibrium problems of all sorts; e.g. calculations of the temporal evolution of laser-heated plasmas, diagnostic analysis of discrete line spectra

* This work was performed under the auspices of the U.S. Atomic Energy Commission.

from pinch-discharge plasmas, numerous astrophysical applications, chemical laser kinetics, and kinetic models of photochemical smog. There are, perhaps, a number of applications in transport theory, as well; such as radiation or neutron transport through a thick medium, where numerical methods invariably suffer from increment limitations. In view of this situation, the thrust of the present article is to qualitatively examine, from the point of view of a scientific practitioner, the nature of stiffness, its causes, its effects (or symptoms), and general requirements of numerical methods which are expected to adequately deal with this class of problems. This point of view also places a certain premium on the desirability of general solution routines which can satisfactorily be used by nonexperts in numerical analysis. In this vein, our experience with a routine of C. W. Gear's [1-3] will be cited for the following reasons: the Gear routine is representative of modern stiff methods, it is specifically designed for general users, and we have found that it fulfills the claims made of it. Finally an example of a stiff photochemical smog problem will be presented and discussed.

THE NATURE OF STIFFNESS

A typical set of N ordinary differential equations can be written in the form

$$\dot{y}_i(t) \equiv dy_i(t)/dt = f_i(y_j(t), \dot{y}_j(t), t) \quad i, j = 1, 2, \dots, N. \quad (1)$$

Stiffness is a difficulty which is engendered in the very foundations of finite difference methods and has no bearing whatever on *analytic* solution methods. The emphasis of the present discussion is directed at large sets of ordinary differential equations where the individual equations are generally exceedingly nonlinear. In such cases numerical solution is perhaps the only reasonable alternative.¹ If the system is stiff, it becomes manifest in a readily recognizable way: incremental steps Δt in the independent variable t are constrained to inordinately small values in order to maintain stability of the numerical method. It can also happen that accuracy requirements restrict Δt to very small values, independent of stability [4]; but most often stability is the factor of major concern. At any rate, one cannot tolerate Δt being systematically restricted to such small values that a problem cannot be run to "completion" in a reasonable, economical number of steps. Further, it is difficult to assess in advance the degree of difficulty which is to be encountered with an arbitrary system. With this in mind, it is well to review a bit more closely the problem of error amplification, or stability, in finite difference methods. More complete treatments can be found in numerous places [4-6]; so

¹ When sets become very large, analog or hybrid solution is a delicate matter which merits discussion in its own right. The present article is restricted to only digital solution methods.

present considerations are intended only to be illustrative of the scope of the problem. For this purpose it is appropriate to consider the rather simple case:

$$(dy_i/dt) - \lambda_i y_i(t) = 0, \quad i = 1, \dots, N. \quad (2)$$

The two basic ways of differencing Eq. (2) are *explicitly* (involving only current and past values of y on the right-hand-side of the equation):

$$[y_i(t_{n+1}) - y_i(t_n)]/\Delta t_{n+1/2} = \lambda_i y_i(t_n), \quad i = 1, \dots, N, \quad (3)$$

where

$$\Delta t_{n+1/2} = t_{n+1} - t_n = \Delta t, \quad (4)$$

and *implicitly* (involving future values of y on the right-hand-side of the equation)²:

$$[y_i(t_{n+1}) - y_i(t_n)]/\Delta t_{n+1/2} = \lambda_i y_i(t_{n+1}). \quad (5)$$

The stability of a method is determined by the magnitude of the accumulated error $\epsilon_i(t_n)$ as a function of (increasing) t_n . This error is defined to be the difference between the solution of the finite difference equation and the true solution of the differential equation. There generally exists a difference equation for $\epsilon_i(t_n)$ having one or more characteristic roots [4]³. In linear systems, these roots can be evaluated and the error bounds determined for increasing t_n . For nonlinear systems of equations, generally acceptable, systematic stability theories seem to be unavailable at present. One can also look at the manner in which the error propagates from step to step. For the explicit scheme (Euler's method) above (Eq. (3)) it can be shown that, to first order in Δt ,

$$\epsilon_i(t_{n+1}) = (1 + \lambda_i \Delta t) \epsilon_i(t_n), \quad i = 1, \dots, N. \quad (6)$$

whereas for the implicit scheme (Eq. (5))

$$\epsilon_i(t_{n+1}) = \epsilon_i(t_n)/(1 - \lambda_i \Delta t), \quad i = 1, \dots, N. \quad (7)$$

To maintain a stable numerical solution, $\epsilon_i(t_n)$ must not grow with increasing n . For exponentially growing solutions ($\lambda_i > 0$) neither the explicit nor the implicit approach consistently enjoys a great advantage over the other in regard to the allowable Δt . But for decaying components ($\lambda_i < 0$) the implicit method (Eq. (7)) is stable for *all* Δt (this is known as absolute stability in the negative $\lambda \Delta t$ half-plane) whereas the explicit method (Eq. (6)) is constrained by the smallest (most negative λ_i) time constant $\tau_i = 1/|\lambda_i|$ *throughout* the problem history! Clearly

² There are, of course, numerous variations of these basic approaches, but Eqs. (3) and (5) adequately serve to demonstrate the matter of error amplification.

³ Some of these roots may be spurious, in that they are due only to the particular numerical method and have nothing at all to do with the physical system.

⁴ This is also referred to as A-stability.

when the τ_i span a very large range, explicit methods may exact an extreme toll in the number of operating cycles that are required to obtain asymptotic solutions. This is the basic difficulty posed by stiff systems.

Even so, a sizable portion of current kinetics work incorporates explicit Runge–Kutta methods, often coupled with approximation schemes such as making steady-state assumptions on selected components either initially or during the progress of the problem. This sort of tampering is ill-suited to most nonexperts and hardly seems advisable for uninitiated users. On the other hand, many implicit schemes avoid the stability problem but nevertheless may command great amounts of machine running time for other reasons. This is the case both for implicit methods which require matrix operations in their convergence schemes and those which do not. In the former case, matrix operations are costly. In the latter case, where matrix operations are not required, satisfactory convergence is often obtained only by using very small Δt . Typical of this latter case are numerous implicit corrector methods which are used in conventional predictor–corrector routines. Further in this regard, Dahlquist [7] developed two important theorems: (1) no (k -step) explicit method can be A -stable, and (2) a linear multistep method (this constitutes a class of conventional corrector methods; e.g., Adams method, etc.) can be A -stable only for order-accuracy of two or less. The second theorem suggests the other dilemma in regard to solving stiff systems. That is the problem of maintaining sufficient accuracy throughout the problem evolution. Namely, while a first or second order A -stable method is free from stability limitations, it can happen that, for a reasonable step size Δt , some solution components (those corresponding to the largest λ_i 's) are often approximated rather poorly [4].⁵ The usual recourse is to go to higher order or else cut Δt . An example of this will be seen with a conventional Adams method in a later section. It should be noted that higher-order methods often allow larger Δt 's to be used for a given degree of accuracy per cycle (often referred to as local truncation error), but the range of $\lambda\Delta t$ space over which absolute stability holds is often restricted with higher-order methods (see Refs. [4] and [8]). At this point it may be well to cap-sulize, in some way, general requirements that will alleviate some of the difficulties associated with stiff systems.

REQUIREMENTS OF A SUITABLE METHOD FOR STIFF SYSTEMS

The primary requirement of stiff methods has been defined above. It is that a suitable method must be highly stable and accurate for values of Δt which are large

⁵ For example, Euler's method of Eq. (3) approximates $e^{\lambda_i\Delta t}$ by $(1 + \lambda_i\Delta t)$, which is only first-order accurate in Δt . If the problem is scaled so that $\Delta t < 1$ over most of the evolution, higher-order methods can improve the accuracy for a given Δt .

enough to span the problem history in an economical number of operating cycles. To achieve this, some, or all, of the following considerations must be enacted. The method should be implicit for stability reasons. It should have a variable-order capability. This allows the use of high-order accuracy (greater than third or fourth order) when a low-order method would be too restrictive on Δt . It also allows the method to be maintained at a very low order to insure A -stability⁶ when this is desirable. In most practical cases absolute stability is not required over the *entire* negative $\lambda\Delta t$ half plane. Thus, larger Δt 's can often be realized, along with absolute stability, by operating in a somewhat restricted region of $\lambda\Delta t$ space with a higher-order method. Preferably a general stiff routine would automatically adjust the order-accuracy and Δt optimally in this regard. Implicit routines using matrix methods most often require the Jacobian $\partial\dot{y}_i/\partial y_j$. Numerical evaluation of Jacobians should be performed when needed, so as to relieve the user of tedious programming. Finally, the matrix equation solution method in an implicit routine should take advantage of sparse matrix structure (many zero elements) when it exists.

Having tried a number of less successful approaches on stiff systems (both in plasma physics and in chemical kinetics), I have found the Gear routine to be very satisfactory, to date, for general use in unexplored problem areas. This is not surprising inasmuch as that routine was designed with all of the above-mentioned requirements in mind. Specifically, Gear's method is an efficient implicit corrector routine which uses a standard predictor formula to obtain trial values for the $y(t_{n+1})$'s. It employs a modified Newton-Raphson method to converge the trial solutions. The user specifies local truncation error requirements. The user can specify a maximum allowable order (up to sixth order) and constant Δt , or he can choose to automatically adjust the order and Δt in optimal fashion, as indicated above. Further, Jacobians are evaluated numerically; and a sparseness capability can be employed, if the user so chooses. The Gear package which I have used also reverts to a standard Adams predictor-corrector (to seventh order) with the relevant user's options mentioned above.

It should again be emphasized that I am not so much concerned with a specific method as with the general properties of stiffness and with generally available measures for dealing with this class of problems. It is most important that physicists and/or other practitioners recognize the scope of the problem and understand its elements of cause and effect. As such, the present article is primarily intended to illustrate, by representative example, the factors which are at play in stiff systems.⁷

⁶ A -stability implies absolute stability over the *entire* negative $\lambda\Delta t$ plane.

⁷ I have not only neglected specific mention of stiff methods and routines developed by others [9-11], but one of my own as well, simply because the Gear routine meets our requirements and is well enough documented (in addition to performing up to its claims) so that it serves our discussion herein with emphasis on the rudiments of stiffness from the user's point of view.

A KINETIC MODEL OF PHOTOCHEMICAL SMOG

Theoretical models for any number of chemical kinetic systems demonstrate the features of stiffness. The one which we will presently consider is a theoretical photochemical smog model proposed by Westberg and Cohen [12] to simulate a set of actual smog chamber measurements which were made by Schuck and Doyle [13]. Our concern is directed specifically to the solution of the equations of Westberg and Cohen. We are not particularly concerned at this point with the validity of the model as a good or bad simulation of the experiment, nor with the credibility of the experimental results *per se*.⁸ I simply want to illustrate comparative features of solutions of a given problem by a suitable stiff routine (code name, SMOG, using Gear's method) as opposed to a conventional method, such as an Adams predictor-corrector (code name, ADAM).

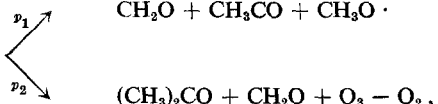
Westberg and Cohen's photochemical smog model is based on the reactions and rate constants given in Table I.

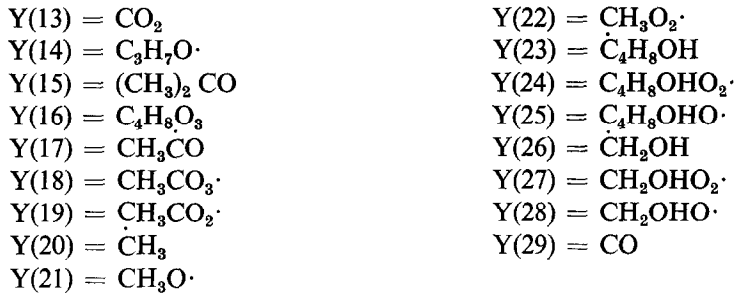
In reactions (35)–(38), R and R' denote the radicals C_3H_7 , C_4H_8OH , CH_2OH , CH_3CO , and CH_3 . Thus, reactions (35)–(37) actually represent five reactions each. Reaction (38) actually represents fifteen reactions. From the above sixty-four reaction processes, twenty-nine rate equations can be constructed. The rates which constitute the equations are unimolecular, bimolecular, or termolecular. For example, the photolysis of NO_2 by solar energy $h\nu$ in Reaction (1) is unimolecular; this rate is expressed as the product of the NO_2 concentration (moles/liter) times the rate constant, 6.7×10^{-3} , given above. Reaction (4) is bimolecular; its rate is expressed as the product of the NO concentration times the O_3 concentration times the rate constant, 9.1×10^6 . Reaction (3) is termolecular; its rate is expressed as the product of the O concentration times the square of the O_2 concentration times the rate constant, 2.4×10^8 . Reactions (18) and (29) involving inert constituents are treated as termolecular. The inert concentration is taken to be 0.04 moles/liter. The dependent variables of these equations are those individual constituents which are both produced and destroyed in the reactions listed. They are:



⁸ Because wall effects were noted in the experiment, one should view the subsequent equations for an infinite, homogeneous system to be little more than a qualitative guide in regard to simulating the experimental results.

TABLE I
 Reactions and Rate Constants (in *l*-mole-sec)
 Underlying Westberg and Cohen's Photochemical Smog Model

Reaction	Rate constant
1. $\text{NO}_2 + h\nu \rightarrow \text{NO} + \text{O}$ (<i>hν</i> represents solar radiant energy)	6.7×10^{-3}
2. $\text{O} + \text{O}_3 + \text{N}_2 \rightarrow \text{O}_3 + \text{N}_2$	2.03×10^8
3. $\text{O} + 2\text{O}_2 \rightarrow \text{O}_3 + \text{O}_2$	2.4×10^8
4. $\text{NO} + \text{O}_3 \rightarrow \text{NO}_2 + \text{O}_2$	9.1×10^6
5. $\text{C}_4\text{H}_8 + \text{O} \rightarrow \dot{\text{C}}_3\text{H}_7 + \dot{\text{H}}\text{CO}$	3.8×10^{10}
6. $\dot{\text{C}}_3\text{H}_7 + \text{O}_2 \rightarrow \text{C}_3\text{H}_7\text{O}_2 \cdot$	1.0×10^6
7. $\text{C}_3\text{H}_7\text{O}_2 \cdot + \text{NO} \rightarrow \text{C}_3\text{H}_7\text{O} \cdot + \text{NO}_2$	3.0×10^8
8. $\text{C}_3\text{H}_7\text{O} \cdot + \text{O}_2 \rightarrow (\text{CH}_3)\text{CO} + \text{HO}_2 \cdot$	1.0×10^4
9. $\dot{\text{H}}\text{CO} + \text{O}_2 \rightarrow \text{CO} + \text{HO}_2 \cdot$	1.0×10^6
10. $\text{HO}_2 \cdot + \text{NO} \rightarrow \text{NO}_2 + \text{HO} \cdot$	3.0×10^8
11. $\text{C}_4\text{H}_8 + \text{O}_3 \rightarrow \text{C}_4\text{H}_8\text{O}_3$	7.0×10^3
12. $\text{C}_4\text{H}_8\text{O}_3$	3.0×10^4
	
(<i>P</i> ₁ and <i>P</i> ₂ are the splitting fractions)	
13. $\text{CH}_3\dot{\text{C}}\text{O} + \text{O}_2 \rightarrow \text{CH}_3\text{CO}_3 \cdot$	1.0×10^6
14. $\text{CH}_3\text{CO}_3 \cdot + \text{NO} \rightarrow \text{CH}_3\text{CO}_2 \cdot + \text{NO}_2$	3.0×10^8
15. $\text{CH}_3\text{CO}_3 \cdot + \text{NO}_2 \rightarrow \text{CH}_3\text{CO}_3\text{NO}_2$	1.0×10^7
16. $\text{CH}_3\text{CO}_3 \cdot + \text{C}_4\text{H}_8 \rightarrow \text{C}_3\text{H}_7\text{CHO} + \text{CH}_3\text{CO}_2$	1.0×10^4
17. $\text{CH}_3\text{CO}_2 \rightarrow \dot{\text{C}}\text{H}_3 + \text{CO}_2$	1.0×10^6
18. $\dot{\text{C}}\text{H}_3 + \text{O}_2 + \text{M} \rightarrow \text{CH}_3\text{O}_2 \cdot + \text{M}$ (<i>M</i> is an inert constituent)	3.0×10^{10}
19. $\text{CH}_3\text{O}_2 \cdot + \text{NO} \rightarrow \text{CH}_3\text{O} \cdot + \text{NO}_2$	3.0×10^8
20. $\text{CH}_3\text{O} \cdot + \text{O}_2 \rightarrow \text{CH}_2\text{O} + \text{HO}_2 \cdot$	2.0×10^4
21. $\text{HO} \cdot + \text{C}_4\text{H}_8 \rightarrow \dot{\text{C}}_4\text{H}_8\text{OH}$	1.0×10^{10}
22. $\dot{\text{C}}_4\text{H}_8\text{OH} + \text{O}_2 \rightarrow \text{C}_4\text{H}_8\text{OHO}_2 \cdot$	2.0×10^6
23. $\text{C}_4\text{H}_8\text{OHO}_2 \cdot + \text{NO} \rightarrow \text{C}_4\text{H}_8\text{OHO} \cdot + \text{NO}_2$	3.0×10^8
24. $\text{C}_4\text{H}_8\text{OHO} \cdot \rightarrow (\text{CH}_3)_2\text{CO} + \dot{\text{C}}\text{H}_2\text{OH}$	3.0×10^3
25. $\dot{\text{C}}\text{H}_2\text{OH} + \text{O}_2 \rightarrow \text{CH}_2\text{OHO}_2 \cdot$	2.0×10^6
26. $\text{CH}_2\text{OHO}_2 \cdot + \text{NO} \rightarrow \text{CH}_2\text{OHO} \cdot + \text{NO}_2$	3.0×10^8
27. $\text{CH}_2\text{OHO} \cdot \rightarrow \text{CH}_2\text{O} + \text{HO} \cdot$	3.0×10^3
28. $\text{HO} \cdot + \text{CO} \rightarrow \text{CO}_2 + \text{H} \cdot$	8.0×10^7
29. $\text{H} \cdot + \text{O}_2 + \text{M} \rightarrow \text{HO}_2 \cdot + \text{M}$	8.0×10^9
30. $\text{HO}_2 \cdot + \text{CO} \rightarrow \text{HO} \cdot + \text{CO}_2$	1.0×10^{-6}
31. $\text{O} + \text{CO} \rightarrow \text{CO}_2$	7.0×10^3
32. $\text{O}_3 + \text{CO} \rightarrow \text{CO}_2 + \text{O}_2$	1.0×10^{-10}
33. $2\text{HO}_2 \cdot \rightarrow \text{H}_2\text{O}_2 + \text{O}_2$	5.0×10^9
34. $\text{HO}_2 \cdot + \text{HO} \cdot \rightarrow \text{H}_2\text{O} + \text{O}_2$	5.0×10^9
35. $\text{RO} \cdot + \text{CO} \rightarrow \text{R} \cdot + \text{CO}_2$	1.0×10^7
36. $\text{RO}_2 \cdot + \text{CO} \rightarrow \text{RO} \cdot + \text{CO}_2$	1.0×10^{-6}
37. $\text{RO}_2 \cdot + \text{HO}_2 \cdot \rightarrow \text{ROOH} + \text{O}_2$	5.0×10^9
38. $\text{RO}_2 \cdot + \text{R}'\text{O}_2 \cdot \rightarrow \text{RO} \cdot + \text{R}'\text{O} \cdot + \text{O}_2$	1.0×10^3



Ordinarily, equations for O_2 , N_2 , and CO_2 are deleted from the set because their concentrations remain constant for all practical purposes. I have retained the O_2 and CO_2 equations simply as an additional checkpoint against untoward solution behavior. I have used splitting fractions with values, $P_1 = 0.70$ and $P_2 = 0.30$. Initially in this problem there are 2.8 parts per million (ppm) by volume




FIG. 1. Pollutant concentration (ppm) vs. time (min).

of isobutene (C_4H_8) and 1 ppm NO_2 as well as the normal air concentrations of O_2 , N_2 , and CO_2 . All other constituents are assumed to be absent initially. Figures 1 and 2 show the evolution of the concentrations of C_4H_8 , NO_2 , CO , and O_3 from $t = 0$ to $t = 80$ min obtained from the program, SMOG.⁹ These curves simply represent the fully time-dependent solutions which were obtained from the set of twenty-nine equations indicated above.

⁹ Westberg and Cohen's results are also shown in these figures and will be discussed in the next section.

Using a CDC 7600 computer, this problem required 0.54 min running time for 707 operating cycles with the stiff routine. In this problem matrix equations were solved by simply using matrix inverses and not exploiting sparse structure. Computing matrix inverses is generally costly and inefficient. For production routines it can be avoided in the interests of economy. However, since a modified Newton-Raphson method is used to converge the trial solutions $y(t_{n+1})$, the matrix inverse was not computed every cycle in the present case. In fact, only 101 calls to the inversion routine were made, which is obviously a tolerable expense in this example.

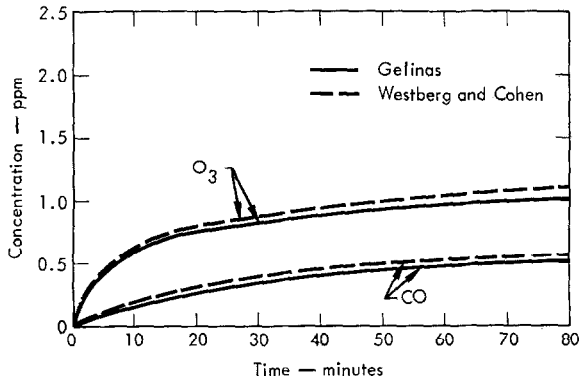


FIG. 2. Pollutant Concentration (ppm) vs. time (min).

A local truncation error constraint of 10^{-4} was sufficient to give reproducible results to three significant figures. The results were reproducible in the sense that the dependent variables *and* effective decay constants agreed to three significant figures with the results of the same problem run with an error constraint of 10^{-6} . Of course, more cycles (1336) and more running time (0.68 min) are required with a constraint of 10^{-6} . The main results quoted herein refer to runs with a local truncation error constraint of 10^{-4} . To decide whether the implicit high-order stiff routine, SMOG, has done anything of real value or not, the same problem should be run with one or more of the conventional methods. A standard Adams predictor-corrector is a reasonable candidate for at least two reasons: it can be made to run more efficiently than, say, Runge-Kutta-Gill in terms of necessary function evaluations at high order accuracies (e.g., fourth order and higher). Second, if worst should come to worst because of very negative decay constants, we have noted that A-stability can be ensured by operating at second-order accuracy, or lower. Thus, if stiffness problems are encountered, at least they can be attacked stably by lowering the order. With this in mind, I ran the above problem using a variable-step size, variable order (as high as seventh) Adams method (this version of our code is named ADAM), which also was available in the "Gear package."

In 4.32 min of running time on the C.D.C. 7600 computer¹⁰, the Adams solution has consumed 97,100 operating cycles in order to advance the solutions to $t = 6.39 \times 10^{-3}$ sec. At this stage Δt was 7.84×10^{-9} sec, and the dependent variables and effective decay constants differed from the solutions from the stiff routine by one to four percent. At about $t = 10^{-5}$ sec the method was reduced and maintained at first-order accuracy in the interest of stability. From this point, a prodigious number of operating cycles (nearly 97 000 cycles) was required to advance the low-order Adams solutions¹¹ from $t = 10^{-5}$ sec to $t = 6.39 \times 10^{-3}$ sec. This is clearly an example of forcing a solution to be A-stable (by lowering the order to two or less); but, in so doing the convergence limitations on Δt (due to the low-order accuracy) were so severe that the problem could not advance to asymptotic solutions in reasonable amounts of machine time. Program SMOG also spent a sizable portion of its run time in this regime; but due to its higher-order capability (in a suitably restricted region of $\lambda\Delta t$ space), it was able to stably advance solutions to late times. Further details will appear in the next section. From the above example alone, it seems plausible that numerous applied problems may be sufficiently stiff that, for all practical purposes, they would have to be considered "off-limits" for conventional methods. Rather than abandon such problems, in the absence of suitable methods, there would be little else to do but take the refuge which is available in the invocation of quasi-steady state assumptions on some of the system components.

QUASI-STEADY STATE ASSUMPTIONS

We noted at least twice previously that conventional numerical methods may suffice, in some sense, if the original set of time-dependent equations is simplified, i.e., simplified by using quasi-steady state values *at each instant* t for some selected constituents rather than retaining in the initial set, and solving, the fully time-dependent equations for those selected constituents. For example, the equation for the oxygen atom concentration $[O(t)]$, in the model we are considering is

$$\begin{aligned} \frac{d[O(t)]}{dt} = & k_1[NO_2(t)] - [O(t)]\{k_2[O_2(t)][N_2(t)] \\ & + k_3[O_2(t)]^2 + k_5[C_4H_8(t)] + k_{30}[CO(t)]\}, \end{aligned} \quad (8)$$

where k_i is the rate constant for the i -th reaction listed previously. By setting the

¹⁰ This is roughly equivalent to 20 min. running time on our C.D.C. 6600 machine.

¹¹ This problem was also run with a Runge-Kutta-Merson routine [14] with even less satisfactory results than the Adams results. Namely, the Runge-Kutta-Merson routine goes unstable at about $t = 2 \times 10^{-3}$ sec., even with a very severe local truncation error constraint of 10^{-8} .

time derivative, $d[\text{O}(t)]/dt$, equal to zero; the instantaneous, or quasi-steady state (QS), oxygen atom concentration at any time instant t is given by

$$[\text{O}(t)]_{\text{os}} = \frac{k_1[\text{NO}_2(t)]}{\{k_2[\text{O}_2(t)][\text{N}_2(t)] + k_3[\text{O}_2(t)]^2 + k_5[\text{C}_4\text{H}_8(t)] + k_{30}[\text{CO}(t)]\}} \quad (9)$$

The usual practice is to simply employ this last expression for $[\text{O}(t)]_{\text{os}}$ at each instant it is required in preference to retaining its time-dependent equation in the set which is to be solved. The question is, "How can one be sure that Eq. (9) is a good representation of the true value which would be obtained at that instant t by actually solving Eq. (8) in conjunction with the entire original set of time-dependent equations?" I have attempted to indicate, by implication, an answer to this question by first solving the fully general set of time-dependent equations for all variables, including $[\text{O}(t)]$. Then, at various times t , I evaluated Eq. (9) independently, based on my instantaneous general solution values for $[\text{NO}_2(t)]$, $[\text{O}_2(t)]$, $[\text{N}_2(t)]$, $[\text{C}_4\text{H}_8(t)]$, and $[\text{CO}(t)]$. At this point I simply compared the latter to the former. From such "head to head" comparison it became obvious, after the fact, when a quasi-steady state assumption on the variable(s) in question was advisable, and when not. But to establish criteria *a priori* as to when a quasi-steady state assumption becomes justified for a given component seems to be fraught with peril; particularly for large systems in which the constituents are coupled in rather complex ways.

When Westberg and Cohen solved their system model, they invoked quasi-steady state assumptions on $[\text{O}]$, $[\text{HO}\cdot]$, $[\text{CH}_3\text{CO}_2\cdot]$, $[\text{CH}_2\text{OHO}\cdot]$, $[\text{C}_3\text{H}_7\text{O}\cdot]$, $[\text{C}_4\text{H}_8\text{OHO}\cdot]$, and on $[\text{CH}_3\text{O}\cdot]$ at each instant t throughout their solution history. Strictly speaking, such a procedure amounts to solving a different problem than the initially posed one. For a first-hand view of this matter I carried out the "head to head" comparison mentioned above by assuming my fully time-dependent stiff solutions (generated by program SMOG) to be true solutions. I then found the earliest times at which the separately calculated quasi-steady state values for the seven components listed above (and in Table I) agreed with the SMOG solutions, if at all. Indeed, if agreement does not exist at some stage, either a component does not physically reach steady state at times spanned by my numerical solution, or the stiff solutions are not close to being true solutions of the differential equations. I found the SMOG solutions and the separately calculated instantaneous-steady state values for the seven selected components to agree to three significant figures¹² at, or before, $t = 19.02$ sec. The SMOG solution had to run 600 cycles to reach this point. In terms of operating cycles, the problem was nearly completed by the time quasi-steady state was a good approximation for all of the seven constituents. Table II summarizes, in chronological order, the approach to quasi-steady state

¹² This is very acceptable in view of our local truncation error constraint of 10^{-4} .

TABLE II
The Approach to Quasi-Steady State for Selected System Constituents

Constituent	Earliest time (in sec) at which general solutions agree with quasi-steady state solution (three significant figures)	No. of operating cycles required to reach this point
0	1.39×10^{-4}	343
CH ₃ CO ₂ ·	0.036	509
HO ·	9.21	587
CH ₃ OHO ·	9.21	587
C ₃ H ₇ O ·	9.21	587
C ₄ H ₈ OHO ·	9.21	587
CH ₃ O ·	19.02	600

for each of these components. It should be mentioned that, prior to $t = 19.02$ sec the disagreement between the general, time-dependent (SMOG) solutions and the corresponding quasi-steady state values for these seven components was in many cases as great as ten orders of magnitude. While this comparison was not carried out for all of the twenty-nine system components, it can reasonably be suspected that the other twenty-two components also experienced departures of their instantaneous steady-state¹³ values from the general, time-dependent solutions over varying portions of the time domain. Neither have we demonstrated that the quasi-steady state values for *all* of the components sooner or later come into agreement with the fully time-dependent solutions. But the above results for the seven components we considered is suggestive that such may be the case.

DISCUSSION

With this background information we can now suggest possible sources of the differences, shown in Figs. 1 and 2, between the present work and that of Westberg and Cohen: First, quasi-steady state assumptions are suspect. They are suspect because the large early errors (ten orders of magnitude) in these seven components can easily introduce errors into the values of some of the *other* twenty-two components. These errors may persist and feed back their influence to other

¹³ I am using the terms "instantaneous-steady state" and "quasi-steady state" interchangeably because both terms are commonly used regarding the same procedure.

components, e.g., C_4H_8 , NO_2 , O_3 , or CO , at varying later times in the problem history. Obviously this effect appears not to have been too serious for Westberg and Cohen in this example. However, the extreme magnitude of the departures seen above suggests the very real possibility of disastrous errors in other examples with other users. This effect would not occur in my own runs with program SMOG because I did not allow any of the (twenty-nine) components to be compromised. That is, I simply calculated the instantaneous steady-state values in an independent "side-calculation," but never allowed these values to be used in the main solution sequence. Secondly, the error amplification properties of Westberg and Cohen's routine are different from those of the Gear routine. I believe they use a modified, explicit Runge-Kutta method; whereas the Gear method is implicit. Thus, even though there are only small differences between our C_4H_8 , NO_2 , O_3 , and CO concentrations at $t = 19.02$ sec, these small differences could propagate quite differently from that point on by the respective methods.¹⁴ We have already indicated the reason for this possibility. Namely, either method only has estimates of local truncation errors and not of the accumulated (or global) numerical errors. Third, and concomitant with the above point is accuracy control. The respective methods are no doubt different in this regard, which also gives rise to different accumulated errors. Finally, we trust that the probability of programming errors is small in the respective routines.

It should be noted here that a somewhat safer way of tampering with a system is to assume quasi-steady state values for a constituent (thereby deleting the time-dependent equation from the initial set) only at times *after* it can be demonstrated to be a valid assumption. But there are two main drawbacks to such a procedure: it is most often more costly in machine time to do this satisfactorily than it is to use a suitable stiff method (at least in the systems I have worked on); and, more seriously, such a procedure leaves one entirely vulnerable to secular behavior of any of the constituents (i.e., nonquasi-steady state behavior can recur and one would never know it).

SUMMARY

This article has been directed primarily to the applied scientist who may be relatively uninitiated in some areas of numerical analysis. The main intent was to alert the general user to the hazards of stiffness and perhaps assist him in recognizing its symptoms; especially since there is rarely a simple advance warning that a system will suffer stiffness difficulties. Through a qualitative description of

¹⁴ This would be true even if all of the other twenty-five components were in agreement at that time ($t = 19.02$ sec), which is, of course, unlikely.

common causes and effects of stiffness and through the force of a current practical example, we have seen somewhat typical stiff behavior. Resort to quasi-steady state assumptions on carefully selected system components was also seen to be somewhat hazardous; particularly in large systems in which the coupling of various components is unfathomable. Means of ameliorating stiffness were indicated; and a routine using Gear's method, which incorporates these means, solved the above problem in a reasonable fashion.

However, along with the good there is also some bad. The severest critic will point out that none of the numerical solutions cited above necessarily represent the true solution of the set of differential equations. This is, of course, accepted and simply emphasizes the need for caution, particularly in this area where extrapolation of general statements from simple particular examples is so very perilous. Perhaps an appropriate attitude in dealing with numerical solutions for such systems at present is that of quantitative analysis in chemistry; namely, a number of samples are tested (in our case, a number of independent solution methods would be applied to a given problem), and carefully considered confidence limits accompany quoted results. In the example cited above, it could only be claimed that the standard (*A*-stable) Adams and the SMOG solutions were within four percent of each other at very early times ($t = 6.39 \times 10^{-3}$ sec), and the Runge-Kutta samples had to be rejected due to observed instabilities. We had no other routine available which yielded asymptotic solutions. Therefore, the only statement of confidence that can be made about the SMOG solutions from $t = 6.39 \times 10^{-3}$ sec to $t = 7000$ sec is that the solutions were stable and reproducible with the application of increasingly restrictive local truncation error constraints. In cases such as this, where analytic solutions are unavailable and manifestly *A*-stable numerical solutions cannot reach late times, little more can be said pending improvements in numerical (or analytic) analysis.

Finally, how is one to know *a priori* that a system may be stiff? At present the practitioner is probably best advised to find out by simply employing a suitable stiff routine¹⁵ on each new class of problems. These solutions can be checked, as far as one's budget allows, with a standard *A*-stable method. If the problem turns out not to be stiff, successive runs should be made by the more economical standard methods, keeping an eye out for the onset of stiffness symptoms as parameter changes are made.

¹⁵ For very large sets of equations current stiff methods can become prohibitively costly because of the great amounts of machine time required for solving large matrix equations. Resolution of this shortcoming will, perhaps, come only with further developments in the numerical analysis of stiff systems.

ACKNOWLEDGMENTS

I am indebted to R. L. Pexton for initially bringing C. W. Gear's work to my attention and to A. C. Hindmarsh for his valuable contributions in incorporating the Gear routine in L. L. L. systems.

REFERENCES

1. C. W. GEAR, The automatic integration of stiff ordinary differential equations, in "Proceedings of the IFIPS Conference," pp. A81-85, 1968.
2. C. W. GEAR, The simultaneous numerical solution of differential-algebraic equations, *IEEE CT-18*, (1971),
3. C. W. GEAR, The automatic integration of ordinary differential equations, *Comm. ACM* **14** (1971), 176-179.
4. J. SEINFELD, L. LAPIDUS, AND M. HWANG, Review of numerical integration techniques for stiff ordinary differential equations, *I and EC Fundamentals* **9** (1970), 266.
5. L. LAPIDUS AND J. SEINFELD, "Numerical Solution of Ordinary Differential Equations," Academic Press, New York, 1971.
6. I. W. SANDBERG AND H. SHICHMAN, "Numerical Integration of Systems of Stiff Nonlinear Differential Equations," *Bell Systems Tech. J.* **47** (1968), 511-527.
7. G. DAHLQUIST, A special stability problem for linear multistep methods *BIT* **3**, 27-43.
8. C. W. GEAR, "The Automatic Integration of Large Systems of Ordinary Differential Equations," Digest Record of 1969 Joint Conference on Mathematical and Computational Aids to Design, Anaheim, CA, 27-58.
9. J. C. BUTCHER, Implicit Runge-Kutta processes, *Math. Comp.* **15** (1964), 50.
10. H. H. ROSENBROCK, "Some General Implicit Processes for the Numerical Solution of Differential Equations," *Computer J.* **5**, 329-330.
11. W. LINIGER AND R. A. WILLOUGHBY, "Efficient Integration Methods for Stiff Systems of Ordinary Differential Equations," *Siam J. Numer. Anal.* **7** (1970), 47-66.
12. K. WESTBERG AND N. COHEN, "The Chemical Kinetics of Photochemical Smog as Analyzed by Computer," Report No. ATR-70 (8107)-1, Aerospace Corporation, El Segundo, CA, 1969.
13. E. A. SCHUCK AND G. J. DOYLE, Air Pollution Foundation Report No. 29, Stanford Research Institute, South Pasadena, CA, 1959.
14. R. L. PEXTON, Lawrence Livermore Laboratory, Livermore, CA, private communication.