Pergamon

PII: 1352–2310(97)00315–4

Atmospheric Environment Vol. 32, No. 4, pp. 791–796, 199
© 1998 Elsevier Science Lt
All rights reserved. Printed in Great Britai
1352–2310/98 $19.00 + 0.0

TECHNICAL NOTE

IMPROVEMENT OF SMVGEAR II ON VECTOR AND SCALAR MACHINES THROUGH ABSOLUTE ERROR TOLERANCE CONTROL

MARK Z. JACOBSON*

Department of Civil and Environmental Engineering, Stanford University, Stanford, CA 94305-4020, U.S.A.

(*First received* 3 *April* 1997 *and in final form* 27 *June* 1997. *Published February* 1998)

Abstract—The computer speed of SMVGEAR II was improved markedly on scalar and vector machines with relatively little loss in accuracy. The improvement was due to a method of frequently recalculating the absolute error tolerance instead of keeping it constant for a given set of chemistry. To test the effects of the modification on speed and accuracy, three Eulerian model simulations were performed on two vector machines, the Cray C-90 and Cray J-916, and a scalar machine, the Silicon Graphics Origin 2000. The first simulation was an urban air pollution case in which gas chemistry, dynamical meteorology, radiation, and species transport were solved together. The second and third simulations were global cases in which chemistry, dynamical meteorology, radiation, and species transport were solved together. For the urban case, the new algorithm reduced SMVGEAR II's computer time by 50–67%, depending on the machine used. For the global cases, chemistry time was reduced by 14–44%. In all simulations, normalized gross errors were less than 1%, and the time required to solve chemistry dropped to near or below 50% of the total model simulation time. Thus, gas chemistry may no longer be the computationally-limiting process in many atmospheric models. The simulations show the usefulness of SMVGEAR II on a scalar workstation with sufficient memory bandwidth. © 1998 Elsevier Science Ltd. All rights reserved.

Key word index: Gear code, photochemistry, global modeling, air pollution, ordinary differential equations.

1. INTRODUCTION

Many numerical schemes have been developed to solve chemical ordinary differential equations in atmospheric models. These include iterative backward Euler schemes (e.g. Curtiss and Hirschfelder, 1952; Shimazaki and Laird, 1970; Hertel *et al.*, 1993), family schemes (e.g. Crutzen, 1971; Turco and Whitten, 1974; Austin, 1991; Elliott *et al.*, 1993; Jacobson, 1994), hybrid predictor–corrector schemes (e.g. Young and Boris, 1977), the quasi-steady-state approximation scheme (QSSA) (Hesstvedt *et al.*, 1978), parameterization schemes (e.g. Jacob *et al.*, 1989), Runge–Kutta–Rosenbrock schemes (Hairer and Wanner, 1991; Sandu *et al.*, 1996), hybrid Newton–Raphson iterative schemes (Gong and Cho, 1993), iterative forward–backward Euler schemes (Jacobson and Turco, 1994), Gauss–Seidel schemes (Verwer, 1994), and others. Codes that have taken advantage of Gear's method (Gear, 1971) include the Livermore Solver for Ordinary Differential Equations (LSODES) (Hindmarsh, 1983), the Variable Coefficient Ordinary Differential Equation Solver (VODE) (Brown *et al.*, 1989), and the Sparse–Matrix, Vectorized Gear Code (SMVGEAR) (Jacobson and Turco, 1994; Jacobson, 1995).

In many urban-scale models that treat photochemistry, transport, and emissions, chemical kinetics has often taken up to 80–90% of total computation time (Odman *et al.*, 1992). When meteorology and radiation are added, the time to solve chemical kinetics has dropped to around 65–75% (Jacobson *et al.*, 1996). Because gas chemistry usually takes more than 50% of computer time in atmospheric pollution models, it has been the limiting computational process in such models. Thus, reducing the time required to solve chemistry to less than 50% of total computer time is a useful goal. Gear's method has frequently been cited as a benchmark for accuracy (e.g. Odman *et al.*, 1992; Gong and Cho, 1993; Hertel *et al.*, 1993; Saylor *et al.*, 1995). Reducing the time required to solve chemistry with a Gear solver to less than 50% of total computer time would be an important step forward.

For this paper, a speed improvement to a Gear-type code, SMVGEAR II, was made with relatively little loss in accuracy. The revised version of SMVGEAR II was implemented into regional and global scale air pollution models that couple chemistry, dynamical meteorology, radiation, transport, and emissions. In both cases, chemistry now takes near or less than 50% of the total computer time. In the sections below, the modification is described, and resulting changes in speed and error are discussed.

2. MODIFICATION TO SMVGEAR II

Gear's method solves the backward differentiation formula (BDF). A detailed description of Gear's solution

*For a copy of SMVGEAR II, please email jacobson@ce.stanford.edu

mechanism as applied to SMVGEAR II is given in Jacobson (1997). Each time step, the change of concentration is determined in an iterative procedure. Two error tests, a local error test and a global error test, are used to control accuracy. The local test checks convergence after each iteration in a time step, and the global test checks convergence after all iterations have been completed for the time step.

In the local error test, a normalized-root-mean-square (NRMS) error is compared to a parameter value that depends on the order of approximation and the iteration number. The NRMS at time step t and iteration m is

$$\mathrm{NRMS}_{t,m} = \sqrt{\frac{1}{K}\sum_{i=1}^{K}\left[\frac{\Delta N_{i,t,m}}{R_{\mathrm{tol}}N_{i,t,1} + A_{\mathrm{tol},t}}\right]^2} \quad (1)$$

where $\Delta N_{i,t,m}$ (molec cm^{-3}) is the change in concentration of species i at time step t during iteration m, K is the number of species, R_{tol} the relative error tolerance, $A_{\mathrm{tol},t}$ the absolute error tolerance for the time step (molec cm^{-3}), and $N_{i,t,1}$ is the concentration (molec cm^{-3}) of species i before the first iteration of the current Gear time step. The relative error tolerance controls errors relative to species concentrations at the beginning of a time step, and the absolute error tolerance controls errors relative to fixed concentrations. For pure relative error tolerance control, $A_{\mathrm{tol},t} = 0$. For pure absolute error tolerance control, $R_{\mathrm{tol}} = 0$.

If the local error test fails but $\mathrm{NRMS}_{t,m}$ is decreasing relative to $\mathrm{NRMS}_{t,m-1}$, iterations continue until the local error test passes. If $\mathrm{NRMS}_{t,m}$ is increasing, convergence is not occurring, and the matrix of partial derivatives is updated before iterations continue. If the convergence test fails again, the time step is reduced until the test passes. Once local convergence has occurred, a global error test is performed to check whether a cumulative NRMS, determined as

$$\mathrm{NRMS}_{t} = \sqrt{\frac{1}{K}\sum_{i=1}^{K}\left[\frac{\sum_m \Delta N_{i,t,m}}{R_{\mathrm{tol}}N_{i,t,1} + A_{\mathrm{tol},t}}\right]^2} \quad (2)$$

exceeds a parameter value that depends on the order of approximation. In this equation, $\sum_m \Delta N_{i,t,m}$ is the net change in species concentration during a time step. If the global error check fails, a new time step is predicted at the same or one order lower approximation. If the error check continues to fail, the time step is reduced. If the global error test passes, the time step was successful. Every few successful time steps, the time step and order of approximation are re-calculated with a time-step estimation scheme.

In SMVGEAR II, the relative error tolerance is variable, but usually set to 10^{-3}. This is the value used for all simulations discussed here. In all previous versions of the code, the absolute error tolerance was set to a constant that depended on the type of chemistry being solved. In Jacobson (1995), $A_{\mathrm{tol},t}$ was fixed to 10^3 molec cm^{-3} for urban chemistry and 10^5 molec cm^{-3} for global chemistry.

A fixed absolute error tolerance has two disadvantages. First, if the concentrations of many species fall below the absolute tolerance, errors can build up rapidly, and the solver can fail. In fact, during several years of testing of SMVGEAR and SMVGEAR II, the only times the solver has failed were when too many concentrations dropped below $A_{\mathrm{tol},t}$. This rarely occurs in the case of gas chemistry but is more likely to occur in the case of aqueous chemistry. The second disadvantage of a fixed A_{tol} is that it unnecessarily slows the speed of obtaining solutions. In the past, $A_{\mathrm{tol},t}$ was set low enough to ensure that chemistry solutions would not fail or build up errors on a three-dimensional grid. Such low values of $A_{\mathrm{tol},t}$ resulted in a slower than optimal speeds for the solver.

For this work, an algorithm was developed to allow $A_{\mathrm{tol},t}$ to vary during a simulation period. $A_{\mathrm{tol},t}$ is now chosen so that 40% of all concentrations are above it at the beginning of a Gear time step. Forty percent is arbitrary, but has been shown to work for urban, regional, and global gas chemistry cases to date. In the new algorithm, $A_{\mathrm{tol},t}$ is bounded for each chemistry case. For urban chemistry, its range is set from 10^3–10^7 molec cm^{-3}. For free tropospheric and stratospheric chemistry, it ranges from 10^5–10^7 molec cm^{-3}. Values of $A_{\mathrm{tol},t}$ larger than 10^7 molec cm^{-3} causes normalized gross errors to exceed 1% in most regions of the atmosphere; thus, such absolute tolerances are not used. Values of $A_{\mathrm{tol},t}$ smaller than 10^3 molec cm^{-3} in urban regions and 10^5 molec cm^{-3} in tropospheric/stratospheric regions increase computer time without much gain in accuracy.

The $A_{\mathrm{tol},t}$ prediction algorithm counts all species with concentration above each of six possible absolute tolerances in each grid cell at the beginning of a Gear time step. The six fixed tolerances are determined as

$$F_{\mathrm{tol},j} = 10\log_{10}F_{\mathrm{tol,min}} + (\log_{10}F_{\mathrm{tol,max}} - \log_{10}F_{\mathrm{tol,min}})(6-j)/5 \quad \text{for } j = 1,\ldots,6 \quad (3)$$

where $F_{\mathrm{tol,min}}$ and $F_{\mathrm{tol,max}}$ are lower and upper bounds, respectively, of the fixed absolute tolerances. For urban gas chemistry, they are set to 10^3 and 10^7 molec cm^{-3}, respectively. For stratospheric chemistry, they are set to 10^5 and 10^7 molec cm^{-3}, respectively. At the beginning of a Gear time step, the number of species concentrations in a grid cell that falls within a fixed error tolerance increment is denoted as I and counted with

$$\begin{cases} N_{i,t,1} > F_{\mathrm{tol},1}, & I_1 = I_1 + 1 \\ F_{\mathrm{tol},2} > N_{i,t,1} \geqslant F_{\mathrm{tol},1}, & I_2 = I_2 + 1, \\ F_{\mathrm{tol},3} > N_{i,t,1} \geqslant F_{\mathrm{tol},2}, & I_3 = I_3 + 1 \\ F_{\mathrm{tol},4} > N_{i,t,1} \geqslant F_{\mathrm{tol},3}, & I_4 = I_4 + 1, \\ F_{\mathrm{tol},5} > N_{i,t,1} \geqslant F_{\mathrm{tol},4}, & I_5 = I_5 + 1 \end{cases} \quad \text{for } i = 1,\ldots,K. \quad (4)$$

The final error tolerance for a grid cell and time step is determined as

$$A_{\mathrm{tol},t} = \begin{cases} F_{\mathrm{tol},1}, & I_1 > fK \\ F_{\mathrm{tol},2}, & I_1 + I_2 > fK \geqslant I_1 \\ F_{\mathrm{tol},3}, & I_1 + I_2 + I_3 > fK \geqslant I_1 + I_2 \\ F_{\mathrm{tol},4}, & I_1 + I_2 + I_3 + I_4 > fK \geqslant I_1 + I_2 + I_3 \\ F_{\mathrm{tol},5}, & I_1 + I_2 + I_3 + I_4 + I_5 > fK \geqslant I_1 \\ & \quad + I_2 + I_3 + I_4 \\ F_{\mathrm{tol},6} & fK \geqslant I_1 + I_2 + I_3 + I_4 + I_5 \end{cases} \quad (5)$$

where f is a fraction of K, chosen as 0.40 here. Thus, the highest tolerance with at least 40% of concentrations above it is chosen as the tolerance for the time step and grid cell. This method is much faster computationally than sorting all species from highest to lowest concentration and setting the tolerance as the concentration of the 40th percentile species. Yet, the counting method, itself, requires a fair amount of computer time if it is done each time step. To minimize this computer time without losing advantage of having a variable tolerance, A_{tol} is recalculated every third Gear time step instead of every step.

3. COMPUTER TIMINGS AND ERROR ANALYSES

Two tests were run to estimate the effects of the new error tolerances on accuracy and speed. The first test was a set of box model experiments applied to judge the accuracy of

Table 1. Normalized gross errors (NGEs) resulting from box-model simulations of urban, free tropospheric, and stratospheric chemistry and time-varying photolysis

	Urban	Tropospheric	Stratospheric
Number of gases	111	91	44
Number of reactions	260	232	158
Time period of simulation	24 h	24 h	10 d
Old absolute tolerance (molec cm^{-3})	10^3	10^5	10^5
NGE w/cutoff 10^6 molec cm^{-3} (%)	0.15	0.25	0.10
NGE w/cutoff 10^3 molec cm^{-3} (%)	0.16	0.28	0.17
New absolute tolerance (molec cm^{-3})	10^3–10^6	10^5–10^6	10^5–10^6
NGE w/cutoff 10^6 molec cm^{-3} (%)	0.25	0.32	0.19
NGE w/cutoff 10^3 molec cm^{-3} (%)	0.35	0.37	0.66
New absolute tolerance (molec cm^{-3})	10^3–10^7	10^5–10^7	10^5–10^7
NGE w/cutoff 10^6 molec cm^{-3} (%)	0.67	0.71	0.37
NGE w/cutoff 10^3 molec cm^{-3} (%)	0.80	0.85	0.92

Note: NGEs are given for three cases. In the first case, old, fixed absolute error tolerances were used. In the latter two cases, new, varying absolute error tolerances were used. In each case, NGEs for two cutoff concentrations are given.

Table 2. Grid statistics, chemical mechanism sizes, and absolute error tolerances for three 3-D test cases

	Urban case	Global case (no Cl)	Global case (w/Cl)
Model used	GATOR/MMTD	GATOR/AGCM	GATOR/AGCM
No. west–east cells	66	72	72
No. south–north cells	42	44	44
No. vertical layers	20	17	17
Total no. grid cells	55,440	53,856	53,856
West–east cell size	0.05°	5.0°	5.0°
South–north cell size	0.045°	4.0°	4.0°
Location of model base	Ground surface	Ground surface	Ground surface
Model top pressure	250 mb	0.425 mb	0.425 mb
Horizontal coordinate	Spherical	Spherical	Spherical
Vertical coordinate	Sigma-pressure	Sigma-pressure	Sigma-pressure
No. chemical species	111	22	39
No. chemical reactions	265	82	137
Simulation time for tests	68 h	24 h	24 h
Old A_{tol} (molec cm^{-3})	10^3	10^5	10^5
New A_{tol} (molec cm^{-3})	10^3–10^7	10^5–10^7	10^5–10^7

using variable tolerances vs a fixed tolerance. In the experiments, urban, free tropospheric, and stratospheric chemistry were coupled with time-dependent photolysis. Sizes of chemistry sets and simulation lengths for each case are given in Table 1. The table also shows normalized gross errors (NGEs) from two sets of simulations in which new absolute tolerances were used and one in which old tolerances were used. The equation for normalized gross error (NGE) is

$$\mathrm{NGE} = \frac{1}{N_{\mathrm{tim}}} \sum_{i=1}^{N_{\mathrm{tim}}} \left[\frac{1}{K_{s,j}} \sum_{i=1}^{K_{s,j}} \frac{|N_{i,j} - E_{i,j}|}{E_{i,j}} \right] \times 100\% \qquad (6)$$

where N_{tim} is the number of one-half hour time intervals, at the end of which, statistics were computed, $N_{i,j}$ and $E_{i,j}$ the predicted and exact concentrations, respectively, of species i at the end of interval j, and $K_{S,j}$ the number of exact concentrations above a minimum cutoff concentration at the end of interval j. "Exact" solutions were calculated with SMVGEAR II using a relative error tolerance of 10^{-9} and an absolute error tolerance of 10^{-9}. All other calculations were carried out with a relative error tolerance of 10^{-3}. Minimum cutoff levels are given in Table 1.

Table 1 shows that the NGE in concentration from the old absolute error tolerances were always less than those from the new tolerances. This was expected since the old tolerances served as lower bounds for the new tolerances. In all new tolerance cases, NGES after the simulations were less than 1%, even when a cutoff concentration of 10^3 molec cm^{-3} was used.

The second test was a set of urban and global Eulerian model simulations in which chemistry was coupled with meteorology, radiation, and transport. These simulations were used to judge the accuracy and speed of SMVGEAR II in realistic simulations. Three sets of simulations, one urban and two global, were performed. The urban simulations were carried out with the GATOR/MMTD air pollution model (Jacobson *et al.*, 1996). The global simulations were performed with the GATOR/AGCM air pollution model, where the AGCM is an updated version of the code of Arakawa and Lamb (1977). Both coupled models solve gas chemistry together with dynamical meteorology, radiative transfer, species transport, emissions, and deposition. Table 2 gives grid and chemical mechanism statistics for each simulation.

Simulations were performed on each of the three computers, a Cray C-90, Cray J-916, and a Silicon Graphics (SGI) Origin 2000. The first two computers use vector processors and the last uses a scalar processor. The three simulations were each performed on a single processor of each of the three computers. The urban simulation was carried out over a 68 h simulation period, and the two global simulations were carried out over a 24 h period. Longer global simulations were also performed on the Cray J-916 and SGI Origin 2000 with little change in the results. Because allocated time on the Cray C-90 was small and used up, longer

Table 3. Comparison of *old* and *new* cpu time to solve three sets of simulations, each on three different computers

	Cray C-90			Cray J-916			SGI Origin 2000		
	Old tol.	New tol.	% dif.	Old tol.	New tol.	% dif.	Old tol.	New tol.	% dif.
Urban case									
Chem. cpu (h/day)	2.24	1.09	− 51.3	14.2	5.1	− 64.1	25.9	8.4	− 67.6
Chem. mflops	364	360		57.4	76.9		31.5	46.7	
Overall cpu (h/day)	3.42	2.28	− 33.3	19.6	10.5	− 46.4	33.3	15.8	− 52.5
Overall mflops	341	330		59.5	71.7		35.0	47.6	
Global case (no Cl)									
Chem. cpu (h/day)	0.24	0.16	− 33.3	1.11	0.68	− 38.7	1.72	0.96	− 44.2
Chem. mflops	347	339		75.0	79.8		48.4	56.5	
Overall cpu (h/day)	0.51	0.43	− 15.7	2.23	1.80	− 19.3	2.87	2.11	− 26.5
Overall mflops	318	311		72.7	74.3		56.5	63.4	
Global case (w/Cl)									
Chem. cpu (h/day)	0.39	0.33	− 15.4	1.81	1.55	− 14.4	2.71	2.22	− 18.1
Chem. mflops	359	351		77.4	74.7		51.7	52.2	
Overall cpu (h/day)	0.70	0.64	− 8.6	3.23	2.97	− 8.1	3.82	3.33	− 12.5
Overall mflops	328	321		71.1	69.2		60.1	61.7	

Note: Timings are given in cpu hours per day of simulation for *chemistry* only and for the *overall* model (chemistry, meteorology, radiation, and transport, etc.). Also shown are speeds of the simulations on the different computers in units of million floating point operations per second (mflops). *Old* refers to simulations before the changes in absolute error tolerance. *New* refers to simulations after the changes in absolute error tolerance. The old and new tolerances are given in Table 2.

Table 4. Percent cpu time required to solve each process in the overall model on the Cray C-90 computer for the three sets of simulations shown in Table 3

	Urban case		Global case (no Cl)		Global case (w/Cl)	
	Old tol. % cpu time	New tol. % cpu time	Old tol. % cpu time	New tol. % cpu time	Old tol. % cpu time	New tol % cpu time
Chemistry	65.4	48.0	46.2	36.5	55.3	51.1
Meteorology	14.8	21.7	6.9	8.1	5.5	5.9
Radiation	10.8	15.6	37.7	44.4	28.2	30.9
Transport	5.7	8.3	7.2	8.6	9.2	10.0
Input/output	2.9	4.4	0.8	1.0	0.7	0.8
Other	0.4	2	1.2	1.4	1.1	1.3
Total	100	100	100	100	100	100

Note: Old refers to simulations before the changes in absolute error tolerance. *New* refers to simulations after the changes in absolute error tolerance.

simulations could not be performed on that machine and comparisons for 24 h are shown. The difference between the two global cases was that, in one case, stratospheric chlorine chemistry was included and in the other, it was not. Both global simulations solved chemistry between the surface and 0.425 mb altitude.

Table 3 compares the cpu time to solve chemistry and all model processes together when the old and new absolute error tolerances were used on each of the three machines. The old and new tolerances used for the simulations are given in Table 2. Table 3 shows that, in the urban case, the new tolerances reduced the time to solve chemistry by 50–67%, depending on the machine used. Improvements in chemistry reduced the corresponding overall computer time for the model simulations by 33–52%. In the global cases, speedups were less significant than in the urban case simply because the old fixed absolute error tolerance in the global cases was already 100 times larger than in the urban case. Nevertheless, the reduction in computer time for chemistry in the global cases were significant, ranging from 14.4 to 44.2%.

SMVGEAR II was originally developed for vector machines. Because it minimizes inner-loop array referencing, it has advantages on scalar machines as well. Array referencing is minimized because the inner loop of all nested loops is the grid-cell loop (with length equal to the number of grid cells in a grid block). Arrays that reference species numbers are outside the inner loop; thus, their use is minimized. Table 3 shows that SMVGEAR II was only 35% slower on an SGI Origin 2000 than on a Cray J-916.

The main reason the solution to chemistry is slower on the Origin 2000 than the Cray J-916 is that the memory bandwidth on the Origin 2000 is smaller than that on the Cray J-916. SMVGEAR II requires storage and retrieval of information in relatively large arrays. The speed at which this information is retrieved is controlled by the memory bandwidth. On the Cray J-916, the memory bandwidth is 1.6 GB/s/processor while on the SGI Origin 2000, it is 0.4 GB/s/processor. The bandwidth bottleneck on the Origin is evidenced in Table 3, which shows that the average mflop speed of all model processes is greater than that of chemistry alone. Most other model processes do not require so much memory retrieval as does chemistry. Table 3 also shows that the mflop speed of urban chemistry, which requires more memory referencing, is slower than that of global chemistry. Despite the bandwidth limitations,

Table 5. Absolute value percent difference in the normalized gross error before and after the absolute error tolerance was modified for the urban case, calculated as $D_{NGE} = 100\%|NGE_a - NGE_b|/NGE_b$, where NGE_b is the NGE with the old absolute tolerance, and NGE_a is the NGE with the new tolerances. NGE_a and NGE_b were determined by comparing predictions to observations from SCAQS surface sites. For the simulations, 111 gases were solved for, but SCAQS data were available for only a few gases or groups

Parameter	No. of comparisons	D_{NGE} (%)
O_3	1246	0.017
NO	583	0.036
NO_2	1576	0.026
NO_x	1691	0.030
CO	1540	0.014
CH_4	469	0.082
SO_2	455	0.032
PAN	232	0.070
PAR	64	0.019
NMOC	61	0.028
Temp.	1096	0.008
Surf. pres.	648	0.009
Rel. hum.	649	0.017
Wind drct.	2448	0.037
Wind spd.	911	0.008
Solar rad.	185	0.002

chemistry and overall model solutions on an Origin 2000 are sufficiently fast to enable its consistent use of SMVGEAR II on this scalar machine.

Table 4 shows the relative percent of computer time spent for each process in each of the three simulations on the Cray C-90 computer. The table shows that the new absolute error tolerances reduced the time to solve chemistry to below 50% of all computer time in the urban case. In one global case, the time required to solve chemistry was already below 50%. In the other, the new tolerances reduced the time to solve chemistry to almost 50%.

Table 5 shows absolute value differences between normalized gross errors in several parameter values before and after the new absolute error tolerances were implemented in the urban case. Normalized gross errors (NGEs) were determined by comparing model predictions with the observations from the Southern California Air Quality Study (SCAQS) of August 26–28, 1987, which is the time the urban simulations were performed for. Once the NGEs were calculated for both simulations, their absolute value percent differences were determined and presented in Table 5. The table shows that the maximum difference arising from use of new absolute error tolerances was less than 0.1%.

Table 6 shows the absolute-value percent difference in final concentrations between the new and old absolute tolerance simulations, averaged over the entire grid and normalized over all species. The table shows that the difference was less than 0.5% in all cases, which is less than the difference in error between the new and old absolute tolerance simulations, shown in Table 1.

Use of the absolute error tolerance prediction scheme gives lower NGEs in a three-dimensional model than in a box model, as roughly demonstrated by comparing Table 6 to the last row in Table 1. The reason is that, in a three-dimensional model, the grid domain is divided into blocks of 500 grid cells. The absolute tolerance prediction scheme decreases the number of iterations and time steps required by the cell in each block with the stiffest equations. Since all cells in a block use the same time step and number of iterations, reducing the time and iterations required by the stiffest cell reduces the time and iterations required by all cells. But, since all cells in a block, except the stiffest, still take more iterations and time steps than they need to converge, their accuracy is affected less than that of the stiffest cell. As a result, the loss in accuracy, averaged over the entire block, is less than the loss in accuracy of the stiffest grid cell.

Table 6. Absolute-value percent difference in final concentrations between the new and old absolute tolerance simulations, averaged over the entire grid and normalized over all species, for each of the test cases described in Table 2

Urban case	Global case (no Cl)	Global case (w/Cl)
0.48%	0.32%	0.36%

4. CONCLUSIONS

SMVGEAR II was modified to allow for the recalculation of the absolute error tolerance in each grid cell of a model after every third Gear time step. Such recalculation is important for speeding solutions to chemistry with relatively little loss in accuracy. The modified version of SMVGEAR II was implemented into a regional model (GATOR/MMTD) and a global model (GATOR/AGCM), and simulations were run for three test cases, each on three machines. The machines compared were the Cray C-90 and Cray J-916, and the Silicon Graphics Origin 2000. The first two machines use vector processors and the last uses a scalar processor. For urban simulations, the new absolute error tolerance reduced computer time for SMVGEAR II by 50–67%, depending on the machine used. For the global simulations, time to solve chemistry was reduced by 14–44%. In all cases, the time required to solve chemistry dropped to near or below 50% of the total model simulation time. Normalized gross errors in all simulations were less than 1%. The success of the code on the SGI Origin 2000 indicates that SMVGEAR II is useful on a scalar machine that has sufficient memory bandwidth.

Acknowledgements—Some of this work was performed on a Cray J-916 and SGI Origin 2000, provided partially by Cray Research, Inc. and Silicon Graphics, Inc., respectively. Cray 90 computer support was also given by the NAS computer facilities in Mountain View, California. This work was funded by grants from the Environmental Protection Agency under assistance agreement 823186-01-0, the Charles Lee Powell Foundation, and the National Science Foundation under agreement ATM-9504481.

REFERENCES

Arakawa, A. and Lamb, V. R. (1977) Computational design of the basic dynamical processes of the UCLA general circulation model. *Methods in Computational Physics* **17**, 173–265.

Austin, J. (1991) On the explicit versus family solution of the fully diurnal photochemical equations of the stratosphere. *Journal of Geophysical Research* **96** (12), 941–12,974.

Brown, P. N., Byrne, G. D. and Hindmarsh, A. C. (1989) VODE: a variable coefficient ODE Solver. *SIAM Journal of Science Statistical Computing* **10**, 1038–1051.

Crutzen, P. J. (1971) Ozone production rates in an oxygen–hydrogen–nitrogen oxide atmosphere. *Journal of Geophysical Research* **76**, 7311–7327.

Curtiss, C. F. and Hirschfelder, J. O. (1952) Integration of stiff equations. *Proceedings of National Academy of Sciences.* U.S.A. **38**, 235–243.

Elliott, S., Turco, R. P. and Jacobson, M. Z. (1993) Tests on combined projection/forward differencing integration for stiff photochemical family systems at long time step. *Computers Chemistry* **17**, 91–102.

Gear, C. W. (1971) *Numerical Initial Value Problems in Ordinary Differential Equations.* Prentice-Hall. Englewood Cliffs, New Jersey.

Gong, W. and Cho, H.-R. (1993) A numerical scheme for the integration of the gas-phase chemical rate equations in three-dimensional atmospheric models. *Atmosmorphic Environment* **27A**, 2147–2160.

Hairer, E. and Wanner, G. (1991) *Solving Ordinary Differential Equations II. Stiff and Differential-Algebraic Problems.* Springer, Berlin.

Hertel, O., Berkowicz, R. and Christensen, J. (1993) Test of two numerical schemes for use in atmospheric transport-chemistry models. *Atmospheric Environment* **27A**, 2591–2611.

Hesstvedt, E., Hov, O. and Isaksen, I. S. A. (1978) Quasi-steady-state approximations in air pollution modeling: comparison of two numerical schemes for oxidant prediction. *International Journal of Chemical Kinetics* **10**, 971–994.

Hindmarsh, A. C. (1983) ODEPACK, a systematized collection of ODE solvers. In *Scientific Computing*, R. S. Stepleman *et al.*, ed. North-Holland, Amsterdam, pp. 55–74.

Jacob, D. J., Sillman, S., Logan, J. A. and Wofsy, S. C. (1989) Least independent variables method for simulation of tropospheric ozone. *Journal of Geophysical Research* **94**, 8497–8509.

Jacobson, M. Z. (1994) Developing, coupling, and applying a gas, aerosol, transport, and radiation model to study urban and regional air pollution. Ph. D. thesis, University of California, Los Angeles.

Jacobson, M. Z. (1995) Computation of global photochemistry with SMVGEAR II. *Atmospheric Environment* **29A**, 2541–2546.

Jacobson, M. Z. and Turco, R. P. (1994) SMVGEAR: a sparse-matrix, vectorized Gear code for atmospheric models. *Atmospheric Environment* **28A**, 273–284.

Jacobson, M. Z. (1997) *Fundamentals of Atmospheric Modeling.* Cambridge University Press, Cambridge (in press).

Jacobson, M. Z., Lu, R., Turco, R. P. and Toon, O. B. (1996) Development and application of a new air pollution modeling system: Part 1. Gas-phase simulations. *Atmospheric Environment* **30**, 1939–1963.

Odman, M. T., Kumar, N. and Russell, A. G. (1992) A comparison of fast chemical kinetic solvers for air quality modeling. *Atmospheric Environment* **26A**, 1783–1789.

Sandu, A., Verwer, J. G., van Loon, M., Carmichael, G. R., Potra, F. A. and Seinfeld, J. H. (1996) Benchmarking stiff ODE solvers for atmospheric chemistry problems I: implicit versus explicit. CWI Report NM-R9603, Netherlands.

Saylor, R. D. and Ford, G. D. (1995) On the comparison of numerical methods for the integration of kinetic equations in atmospheric chemistry and transport models. *Atmospheric Environment* **29**, 2585–2593.

Shimazaki, T. and Laird, A. R. (1970) A model calculation of the diurnal variation in minor neutral constituents in the mesosphere and lower thermosphere including transport effects. *Journal of Geophysical Research* **75**, 3221–3235.

Turco, R. P. and Whitten, R. C. (1974) A comparison of several computational techniques for solving some common aeronomic problems. *Journal of Geophysical Research* **79**, 3179–3185.

Verwer, J. G. (1994) Gauss–Seidel iteration for stiff ODEs from chemical kinetics. *SIAM Journal of Scientific Computing* **15**, 1243–1250.

Young, T. R. and Boris, J. P. (1977) A numerical technique for solving stiff ordinary differential equations associated with the chemical kinetics of reactive-flow problems. *Journal of Physical Chemistry* **81**, 2424–2427.