NASA
Reference Publication
1327
Lawrence
Livermore National
Laboratory Report UCRL-ID-113855

1993

## Description and Use of LSODE, the Livermore Solver for Ordinary Differential Equations

Krishnan Radhakrishnan
Sverdrup Technology, Inc.
Leewis Research Center Group
Alan C. Hindmarsh
Lawrence Livermore National Laboratory
Livermore, CA

National Aeronautics and Space Administration Otfice of Management
Scientitic and Technical Information Program

## Preface

This document provides a comprehensive description of LSODE, a solver for initial value problems in ordinary differential equation systems. It is intended to bring together numerous materials documenting various aspects of LSODE, including technical reports on the methods used, published papers on LSODE, usage documentation contained within the LSODE source, and unpublished notes on algorithmic details.

The three central chapters-on methods, code description, and code usage-are largely independent. Thus, for example, we intend that readers who are familiar with the solution methods and interested in how they are implemented in LSODE can read the Introduction and then chapter 3, Description of Code, without reading chapter 2, Description and Implementation of Methods. Similarly, those interested solely in how to use the code need read only the Introduction and then chapter 4, Description of Code Usage. In this case chapter 5, Example Problem, which illustrates code usage by means of a simple, stiff chemical kinetics problem, supplements chapter 4 and may be of further assistance.

Although this document is intended mainly for users of LSODE, it can be used as supplementary reading material for graduate and advanced undergraduate courses on numerical methods. Engineers and scientists who use numerical solution methods for ordinary differential equations may also benefit from this document.

## Contents

List of Figures ..... ix
List of Tables ..... xi
Chapter 1 Introduction ..... 1
Chapter 2 Description and Implementation of Methods ..... 7
2.1 Linear Multistep Methods ..... 7
2.2 Corrector Iteration Methods ..... 9
2.2.1 Functional Iteration ..... 11
2.2.2 Newton-Raphson Iteration ..... 13
2.2.3 Jacobi-Newton Iteration ..... 15
2.2.4 Unified Formulation ..... 15
2.3 Matrix Formulation ..... 17
2.4 Nordsieck's History Matrix ..... 20
2.5 Local Truncation Error Estimate and Control ..... 28
2.6 Corrector Convergence Test and Control ..... 32
2.7 Step Size and Method Order Selection and Change ..... 33
2.8 Interpolation at Output Stations ..... 36
2.9 Starting Procedure ..... 39
Chapter 3 Description of Code ..... 41
3.1 Integration and Corrector Iteration Methods ..... 41
3.2 Code Structure ..... 42
3.3 Internal Communication ..... 46
3.4 Special Features ..... 46
3.4.1 Initial Step Size Calculation ..... 57
3.4.2 Switching Methods ..... 64
3.4.3 Excessive Accuracy Specification Test ..... 64
3.4.4 Calculation of Method Coefficients ..... 65
3.4.5 Numerical Jacobians ..... 66
3.4.6 Solution of Linear System of Equations ..... 69
3.4.7 Jacobian Matrix Update ..... 70
3.4.8 Corrector Iteration Convergence and Corrective Actions ..... 70
3.4.9 Local Truncation Error Test and Corrective Actions ..... 71
3.4.10 Step Size and Method Order Selection ..... 72
3.5 Error Messages ..... 74
Chapter 4 Description of Code Usage ..... 75
4.1 Code Installation ..... 75
4.1.1 BLOCK DATA Variables ..... 75
4.1.2 Modifying Subroutine XERRWV ..... 76
4.2 Call Sequence ..... 76
4.3 User-Supplied Subroutine for Derivatives (F) ..... 83
4.4 User-Supplied Subroutine for Analytical Jacobian (JAC) ..... 84
4.5 Detailed Usage Notes ..... 85
4.5.1 Normal Usage Mode ..... 86
4.5.2 Use of Other Options ..... 86
4.5.3 Dimensioning Variables ..... 86
4.5.4 Decreasing the Number of Differential Equations (NEQ) ..... 87
4.5.5 Specification of Output Station (TOUT) ..... 87
4.5.6 Specification of Critical Stopping Point (TCRIT) ..... 88
4.5.7 Selection of Local Error Control Parameters (ITOL, RTOL, and ATOL) ..... 88
4.5.8 Selection of Integration and Corrector Iteration Methods (MF) ..... 89
4.5.9 Switching Integration and Corrector Iteration Methods ..... 91
4.6 Optional Input ..... 91
4.6.1 Initial Step Size (H0) ..... 92
4.6.2 Maximum Step Size (HMAX) ..... 92
4.6.3 Maximum Method Order (MAXORD) ..... 92
4.7 Optional Output ..... 92
4.8 Other Routines ..... 93
4.8.1 Interpolation Routine (Subroutine INTDY) ..... 93
4.8.2 Using Restart Capability (Subroutine SRCOM) ..... 94
4.8.3 Error Message Control (Subroutines XSETF and XSETUN) ..... 95
4.9 Optionally Replaceable Routines ..... 95
4.9.1 Setting Error Weights (Subroutine EWSET) ..... 95
4.9.2 Vector-Norm Computation (Function VNORM) ..... 97
4.10 Overlay Situation ..... 98
4.11 Troubleshooting ..... 98
Chapter 5 Example Problem ..... 101
5.1 Description of Problem ..... 101
5.2 Coding Required To Use LSODE ..... 102
5.2.1 General ..... 102
5.2.2 Selection of Parameters ..... 102
5.3 Computed Results ..... 104
Chapter 6 Code Availability ..... 105
References ..... 107

## List of Figures

Figure 3.1.-Structure of LSODE package ..... 45
Figure 3.2.-Flowchart of subroutine LSODE ..... 59
Figure 3.3.-Flowchart of subroutine STODE ..... 61
Figure 5.1.-Listing of MAIN program for example problem ..... 103
Figure 5.2.-Listing of subroutine (FEX) that computes derivatives for example problem ..... 103
Figure 5.3.-Listing of subroutine (JEX) that computes analytical Jacobian matrix for example problem ..... 104
Figure 5.4.-Output from program for example problem ..... 104

## List of Tables

Table 2.1.-Method coefficients for Adams-Moulton method in normal form of orders 1 to 12 ..... 25
Table 2.2.-Method coefficients for backward differentiation formula method in normal form of orders 1 to 6 ..... 26
Table 3.1.-Summary of integration methods included in LSODE and corresponding values of METH, the first decimal digit of MF ..... 42
Table 3.2.- Corrector iteration techniques available in LSODE and corresponding values of MITER, the second decimal digit of MF ..... 42
Table 3.3.-Description of subprograms used in LSODE ..... 44
Table 3.4.-Routines with common blocks, subprograms, and calling subprograms in double-precision version of LSODE ..... 47
Table 3.5.-Routines with common blocks, subprograms, and calling subprograms in single-precision version of LSODE ..... 48
Table 3.6.-Common blocks with variables and subprograms where used ..... 49
Table 3.7.-Description of variables in common block EH0001, their current values, and subprograms where they are set ..... 49
Table 3.8.-Description of variables in common block LS0001, their current values, if any, and subprograms where they are set or computed ..... 50
Table 3.9.-Length LENWM of array WM in table 3.8 for iteration techniques included in code ..... 55
Table 4.1.-Values of ITASK used in LSODE and their meanings ..... 77
Table 4.2.-Values of ITOL used in LSODE and their meanings ..... 78
Table 4.3.-Values of ISTATE that can be used on input to LSODE and their meanings ..... 79
Table 4.4.-Values of ISTATE returned by LSODE and their meanings ..... 79
List of Tables
Table 4.5.-Values of IOPT that can be used on input to LSODE and their meanings ..... 80
Table 4.6.-Optional input parameters that can be set by user and their locations, meanings, and default values ..... 80
Table 4.7.-Optional output parameters returned by LSODE and their locations and meanings ..... 82
Table 4.8.-Useful informational quantities regarding integration that can be obtained from array RWORK and their names and locations ..... 82
Table 4.9.--Minimum length required by real work array RWORK (i.e., minimum LRW) for each MF ..... 83
Table 4.10.-Minimum length required by integer work array IWORK (i.e., minimum LIW) for each MITER ..... 83

## Chapter 1

## Introduction

This report describes a FORTRAN subroutine package, LSODE, the Livermore Solver for Ordinary Differential Equations, written by Hindmarsh (refs. 1 and 2), and the methods included therein for the numerical solution of the initial value problem for a system of first-order ordinary differential equations (ODE's). Such a problem can be written as

$$
\left.\begin{array}{l}
\underline{\dot{y}} \equiv \frac{d \underline{y}}{d \xi}=\underline{\mathrm{f}}(\underline{\mathrm{y}}(\xi), \xi)  \tag{1.1}\\
\underline{\mathrm{y}}\left(\xi_{0}\right)=\underline{\mathrm{y}}_{0}=\text { Given, }
\end{array}\right\}
$$

where $\underline{\mathrm{y}}, \underline{\mathrm{y}}_{0}, \underline{\mathrm{y}}$, and $\underline{\mathrm{f}}$ are column vectors with $N(\geq 1)$ components and $\xi$ is the independent variable, for example, time or distance. In component form equation (1.1) may be written as

$$
\left.\begin{array}{c}
\frac{d y_{i}(\xi)}{d \xi}=f_{i}\left(y_{1}(\xi), \ldots, y_{N}(\xi), \xi\right)  \tag{1.2}\\
y_{i}\left(\xi_{0}\right)=y_{i, 0}=\text { Given }
\end{array}\right\} \quad i=1, \ldots, N
$$

The initial value problem is to find the solution function $y$ at one or more values of $\xi$ in a prescribed integration interval $\left[\xi_{0}, \xi_{\text {end }}\right]$, where the initial value of $\underline{y}, \underline{y}_{0}$, at $\xi=\xi_{0}$ is given. The endpoint, $\xi_{\text {end }}$, may not be known in advance as, for example, when asymptotic values of $\underline{y}$ as $\xi \rightarrow \infty$ are required.

Initial value, first-order ODE's arise in many fields, such as chemical kinetics, biology, electric network analysis, and control theory. It is assumed that the

## 1. Introduction

problem is well posed and possesses a solution that is unique in the interval of interest. Solution existence and uniqueness are guaranteed if, in the region of interest, $\underline{f}$ is defined and continuous and for any two vectors $\underline{y}$ and $\underline{y}^{*}$ in that region there exists a positive constant $\mathscr{L}$ such that (refs. 3 and 4) ${ }^{-}$

$$
\begin{equation*}
\left\|\underline{f}(\underline{\mathrm{y}}, \xi)-\underline{\mathrm{f}}\left(\underline{\mathrm{y}}^{*}, \xi\right)\right\| \leqq \mathscr{L}\left\|\underline{\mathrm{y}}-\underline{\mathrm{y}}^{*}\right\| \tag{1.3}
\end{equation*}
$$

which is known as a Lipschitz condition. Here $\|\cdot\|$ denotes a vector norm (e.g., ref. 5), and the constant $\mathscr{L}$ is known as a Lipschitz constant of $\underline{f}$ with respect to $y$.

The right-hand side $f$ of the ODE system must be a function of $y$ and $\xi$ only. It cannot therefore involve $\underline{y}$ at previous $\xi$ values, as in delay or retärded ODE's or integrodifferential equations. It cannot also involve random variables, as in stochastic differential equations. A second- or higher-order ODE system must be reduced to a first-order ODE system.

The solution methods included in LSODE replace the ODE's with difference equations and then solve them step by step. Starting with the initial conditions at $\xi_{0}$, approximations $\mathrm{Y}_{n}\left(=Y_{i, n}, i=1, \ldots, N\right)$ to the exact solution $\mathrm{y}\left(\xi_{n}\right)\left[=y_{i}\left(\xi_{n}\right)\right.$, $i=1, \ldots, N]$ of the ODE's are generated at the discrete mesh points $\xi_{n}(n=1,2, \ldots)$, which are themselves determined by the package. The spacing between any two mesh points is called the step size or step length and is denoted by $h_{n}$, where

$$
\begin{equation*}
h_{n}=\xi_{n}-\xi_{n-1} . \tag{1.4}
\end{equation*}
$$

An important feature of LSODE is its capability of solving "stiff" ODE problems. For reasons discussed by Shampine (ref. 6) stiffness does not have a simple definition involving only the mathematical problem, equation (1.1). However, Shampine and Gear (ref. 7) discuss some fundamental issues related to stiffness and how it arises. An approximate description of a stiff ODE system is that it contains both very rapidly and very slowly decaying terms. Also, a characteristic of such a system is that the $N \times N$ Jacobian matrix $\mathbf{J}\left(=\partial \underline{f} / \partial \underline{\mathrm{y}}\right.$ ), with element $J_{I J}$ defined as

$$
\begin{equation*}
J_{i j}=\partial f_{i} / \partial y_{j}, \quad i, j=1, \ldots, N, \tag{1.5}
\end{equation*}
$$

has eigenvalues $\left\{\lambda_{i}\right\}$ with real parts that are predominantly negative and also vary widely in magnitude. Now the solution varies locally as a linear combination of the exponentials $\left\{e^{\xi \operatorname{Re}\left(\lambda_{i}\right)}\right\}$, which all decay if all $\operatorname{Re}\left(\lambda_{i}\right)<0$, where $\operatorname{Re}\left(\lambda_{i}\right)$ is the real part of $\lambda_{l}$. Hence for sufficiently large $\xi\left(>1 / \max \mid \operatorname{Re}\left(\lambda_{i}\right)\right.$, where the bars $|\cdot|$ denote absolute value), the terms with the largest $\operatorname{Re}\left(\lambda_{l}\right)$ will have decayed to insignificantly small levels while others are still active, and the problem would be classified as stiff. If, on the other hand, the integration interval is limited to $1 / \max \left|\operatorname{Re}\left(\lambda_{i}\right)\right|$, the problem would not be considered stiff.

In this discussion we have assumed that all eigenvalues have negative real parts. Some of the $\operatorname{Re}\left(\lambda_{i}\right)$ may be nonnegative, so that some solution components are nondecaying. However, the problem is still considered stiff if no eigenvalue has a real part that is both positive and large in magnitude and at least one eigenvalue has a real part that is both negative and large in magnitude (ref. 7). Because the $\left\{\lambda_{i}\right\}$ are, in general, not constant, the property of stiffness is local in that a problem may be stiff in some intervals and not in others. It is also relative in the sense that one problem may be more stiff than another. A quantitative measure of stiffness is usually given by the stiffness ratio $\max \left[-\operatorname{Re}\left(\lambda_{i}\right)\right] / \mathrm{min}$ $\left[-\operatorname{Re}\left(\lambda_{i}\right)\right]$. This measure is also local for the reason given previously. Another standard measure for stiffness is the quantity $\max \left[-\operatorname{Re}\left(\lambda_{i}\right)\right] \xi_{\text {end }}-\xi_{0} \mid$. This measure is more relevant than the previous one when $\left|\xi_{\text {end }}-\xi_{0}\right|$ is a better indicator of the average "resolution scale" for the problem than $1 / \min \left[-\operatorname{Re}\left(\lambda_{i}\right)\right]$. (In some cases $\min \left[-\operatorname{Re}\left(\lambda_{i}\right)\right]=0$.)

The difficulty with stiff problems is the prohibitive amounts of computer time required for their solution by classical ODE solution methods, such as the popular explicit Runge-Kutta and Adams methods. The reason is the excessively small step sizes that these methods must use to satisfy stability requirements. Because of the approximate nature of the solutions generated by numerical integration methods, errors are inevitably introduced at every step. For a numerical method to be stable, errors introduced at any one step should not grow unbounded as the calculation proceeds. To maintain numerical stability, classical ODE solution methods must use small step sizes of order $1 / \max \left[-\operatorname{Re}\left(\lambda_{i}\right)\right]$ even after the rapidly decaying components have decreased to negligible levels. Examples of the step size pattern used by an explicit Runge-Kutta method in solving stiff ODE problems arising in combustion chemistry are given in references 8 and 9 . Now, the size of the integration interval for the evolution of the slowly varying components is of order $1 / \min \left[-\operatorname{Re}\left(\lambda_{i}\right)\right]$. Consequently, the number of steps required by classical methods to solve the problem is of order $\max \left[-\operatorname{Re}\left(\lambda_{i}\right)\right] / \min \left[-\operatorname{Re}\left(\lambda_{i}\right)\right]$, which is very large for stiff ODE's.

For stiff problems the LSODE package uses the backward differentiation formula (BDF) method (e.g., ref. 10), which is among the most popular currently used for such problems (ref. 11). The BDF method possesses the property of stiff stability (ref. 10) and therefore does not suffer from the stability step size constraint once the rapid components have decayed to negligible levels. Throughout the integration the step size is limited only by accuracy requirements imposed on the numerical solution. Accuracy of a numerical method refers to the magnitude of the error introduced in a single step or, more precisely, the local truncation or discretization error. The local truncation error $\underline{d}_{n}$ at $\xi_{n}$ is the difference between the computed approximation and the exact solution, with both starting the integration at the previous mesh point $\xi_{n-1}$ and using the exact solution $y\left(\xi_{n-1}\right)$ as the initial value. The local truncation error on any step is therefore the error incurred on that step under the assumption of no past errors (e.g., ref. 12).

The accuracy of a numerical method is usually measured by its order. A method is said to be of order $q$ if the local truncation error varies as $h_{n}^{q+1}$. More

## 1. Introduction

precisely, a numerical method is of order $q$ if there are quantities $\underline{C}$ and $h_{o}(>0)$ such that (refs. 3 and 13)

$$
\begin{equation*}
\left|\underline{d}_{n}\right| \leq \underline{\mathrm{C}} h_{n}^{q+1} \quad \text { for all } \quad 0<h_{n} \leq h_{o} \tag{1.6}
\end{equation*}
$$

where $\left|\underline{d}_{n}\right|$ is an $N$-dimensional column vector containing the absolute values of the $d_{i, n}(i=1, \ldots, N)$. The coefficient vector $\underline{C}$ may depend on the function defining the ODE and the total integration interval, but it should be independent of the step size $h_{n}$ (ref. 13). Accuracy of a numerical method refers to the smallness of the error introduced in a single step; stability refers to whether or not this error grows in subsequent steps (ref. 7).

To satisfy accuracy requirements, the BDF method may have to use small step sizes of order $1 / \max \left(\operatorname{Re}\left|\lambda_{t}\right|\right)$ in regions where the most rapid exponentials are active. However, outside these regions, which are usually small relative to the total integration interval, larger step sizes may be used.

The LSODE package also includes the implicit Adams method (e.g., refs. 4 and 10), which is well suited for nonstiff problems. Both integration methods belong to the family of linear multistep methods. As implemented in LSODE these methods allow both the step size and the method order to vary (from 1 to 12 for the Adams method and from 1 to 5 for the BDF method) throughout the problem. The capability of dynamically varying the step size and the method order is very important to the efficient use of linear multistep methods (ref. 14).

The LSODE package consists of 21 subprograms and a BLOCK DATA module. The package has been designed to be used as a single unit, and in normal circumstances the user needs to communicate with only a single subprogram, also called LSODE for convenience. LSODE is based on, and in many ways resembles, the package GEAR (ref. 15), which, in turn, is based on the code DIFSUB, written by Gear (refs. 10 and 16). All three codes use integration methods that are based on a constant step size but are implemented in a manner that allows for the step size to be dynamically varied throughout the problem. There are, however, many differences between GEAR and LSODE, with the following important improvements in LSODE over GEAR: (1) its user interface is much more flexible; (2) it is more extensively modularized; and (3) it uses dynamic storage allocation, different linear algebra modules, and a wider range of error types (ref. 17). Most significantly, LSODE has been designed to virtually eliminate the need for user adjustments or modifications to the package before it can be used effectively. For example, the use of dynamic storage allocation means that the required total storage is specified once in the user-supplied subprogram that communicates with LSODE; there is no need to adjust any dimension declarations in the package. This feature, besides making the code easy to use, minimizes the total storage requirements; only the storage required for the user's problem needs to be allocated and not that called for by a code using default values for parameters, such as the total number of ODE's, for example. The many different capabilities of the code can be exploited quite simply by setting values for appropriate
parameters in the user's subprogram. Not requiring any adjustments to the code eliminates the user's need to become familiar with the inner workings of the code, which can therefore be used as a "black box," and, more importantly, eliminates the possibility of errors being introduced into the modified version.

The remainder of this report is organized as follows: In chapter 2 we describe the numerical integration methods used in LSODE and how they are implemented in practice. The material presented in this chapter is based on, and closely follows, the developments by Gear (refs. 10 and 18 to 20) and Hindmarsh (refs. 1, $2,15,21$, and 22). Chapter 3 describes the features and layout of the LSODE package. In chapter 4 we provide a detailed guide to its usage, including possible user modifications. The use of the code is illustrated by means of a simple test problem in chapter 5 . We conclude this report with a brief discussion on code availability in chapter 6.

## Chapter 2

## Description and Implementation of Methods

### 2.1 Linear Multistep Methods

The numerical methods included in the packaged code LSODE generate approximate solutions $Y_{n}$ to the ordinary differential equations (ODE's) at discrete points $\xi_{n}(n=1,2, \ldots)$. Assuming that the approximate solutions $Y_{n-j}$ have been computed at the mesh points $\xi_{n-j}(j=1,2, \ldots)$, these methods advance the solution to the current value $\xi_{n}$ of the independent variable by using linear multistep formulas of the type

$$
\begin{equation*}
\underline{\mathrm{Y}}_{n}=\sum_{j=1}^{K_{1}} \alpha_{j} \underline{\mathrm{Y}}_{n-j}+h_{n} \sum_{j=0}^{K_{2}} \beta_{j} \underline{\mathrm{f}}_{n-j} \tag{2.1}
\end{equation*}
$$

where the current approximate solution vector $\underline{Y}_{n}$ consists of $N$ components,

$$
\begin{equation*}
\underline{Y}_{n}=\left(Y_{1, n}, \ldots, Y_{N, n}\right)^{T}, \tag{2.2}
\end{equation*}
$$

and the superscript $T$ indicates transpose. In equation (2.1), $\mathrm{f}_{n-j}\left[=\mathrm{f}\left(\underline{Y}_{n-j}\right)\right]$ is the approximation to the exact derivative vector at $\xi_{n-j}, \underline{\dot{y}}\left(\xi_{n-j}\right)\left[=\underline{\mathrm{f}}\left(\underline{y}\left(\xi_{n-j}\right)\right)\right]$, where for notational convenience the $\xi$ argument of $f$ has been dropped; the coefficients $\left\{\alpha_{j}\right\}$ and $\left\{\beta_{j}\right\}$ and the integers $K_{1}$ and $K_{2}$ are associated with a particular method; and $h_{n}\left(=\xi_{n}-\xi_{n-1}\right)$ is the step size to be attempted on the current step $\left[\xi_{n-1}, \xi_{n}\right]$. The method is called linear because the $\left\{\mathrm{Y}_{j}\right\}$ and $\left\{\mathrm{f}_{j}\right\}$ occur linearly. It is called multistep because it uses information from several previous mesh points. The number $\max \left(K_{1}, K_{2}\right)$ gives the number of previous values involved.

The values $K_{1}=1$ and $K_{2}=q-1$ produce the popular implicit Adams, or Adams-Moulton (AM), method of order $q$ :

## 2. Description and Implementation of Methods

$$
\begin{equation*}
\underline{\mathrm{Y}}_{n}=\underline{\mathrm{Y}}_{n-1}+h_{n} \sum_{j=0}^{q-1} \beta_{j} \underline{\mathrm{f}}_{n-j} . \tag{2.3}
\end{equation*}
$$

The method is called implicit because it uses the as yet unknown $\underline{f}_{n}$ to compute $\underline{Y}_{n}$. The method order $q$ means that if equation (2.3) is solved with all past values being exact, the resulting $\underline{Y}_{n}$ will differ from the exact solution $\underline{y}\left(\xi_{n}\right)$ to the ODE system by a local truncation error that is of order $\mathrm{O}\left(H^{q+1}\right)$ for small values of $H=$ $\max \left|h_{k}\right|$.

The choice $K_{1}=q, K_{2}=0$ results in the backward differentiation formula (BDF) method of order $q$ :

$$
\begin{equation*}
\underline{Y}_{n}=\sum_{j=1}^{q} \alpha_{j} \underline{Y}_{n-j}+h_{n} \beta_{0} \underline{\mathrm{f}}_{n} . \tag{2.4}
\end{equation*}
$$

The term "backward differentiation formula" is used to describe the method because equation (2.4), upon division by $h_{n} \beta_{0}$ and rearrangement of terms, can be regarded as an approximation for $\underline{y}\left(\xi_{n}\right)$ in terms of $\underline{Y}_{n}, \underline{Y}_{n-1}, \ldots, \underline{Y}_{n-q}$ (refs. 15 and 17).

The two methods can be written in the general form

$$
\begin{equation*}
\underline{\mathrm{Y}}_{n}=\underline{\Psi}_{n}+h_{n} \beta_{0} \underline{\mathrm{f}}_{n}=\underline{\Psi}_{n}+h_{n} \beta_{0} \underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}\right) . \tag{2.5}
\end{equation*}
$$

where $\underline{\Psi}_{n}$ contains previously computed information and is given by

$$
\begin{equation*}
\underline{\Psi}_{n}=\underline{\mathrm{Y}}_{n-1}+h_{n} \sum_{j=1}^{q-1} \beta_{j} \underline{\mathrm{f}}_{n-j} \tag{2.6a}
\end{equation*}
$$

for the AM method of order $q$, and

$$
\begin{equation*}
\underline{\Psi}_{n}=\sum_{j=1}^{q} \alpha_{j} \underline{Y}_{n-j} \tag{2.6b}
\end{equation*}
$$

for the BDF method of order $q$.
The coefficients $\left\{\alpha_{j}\right\}$ and $\left\{\beta_{j}\right\}$ are determined such that equations (2.3) and (2.4) will be exact if the solution to equation (1.1) is a polynomial of degree $q$ or less. Stability characteristics limit $q$ in equation (2.4) to 6 (ref. 10). In LSODE, however, BDF's of order up to only 5 are used because of additional stability considerations (refs. 7 and 23). The coefficients $\left\{\alpha_{j}\right\}$ and $\left\{\beta_{j}\right\}$ for the two
methods are given by Gear (ref. 10) for $q \leq 6$. In equation (2.5), although the subscript $n$ has been attached to the step size $h$, indicating that $h_{n}$ is the step size to be attempted on the current step, the methods used in LSODE are based on a constant $h$. When the step size is changed, the data at the new spacing required to continue the integration are obtained by interpolating from the data at the original spacing. Solution methods and codes that are based on variable step size have also been developed (refs. 17, 23, and 24) but are not considered in the present work.

### 2.2 Corrector Iteration Methods

If $\beta_{0}=0$, the methods are called explicit because they involve only the known values $\left\{\mathrm{Y}_{n-j}\right\}$ and $\left\{\mathrm{f}_{n-j}\right\}$, and equation (2.1) is easy to solve. If, however, $\beta_{0} \neq 0$, the methods are called implicit and, in general, solution of equation (2.1) is expensive. For both methods, equations (2.3) and (2.4), $\beta_{0}$ is positive for each $q$ and because $f$ is, in general, nonlinear, some type of iterative procedure is needed to solve equation (2.5). Nevertheless, implicit methods are preferred because they are more stable, and hence can use much larger step sizes, than explicit methods and are also more accurate for the same order and step size (refs. 4, 10, and 12). Explicit methods are used as predictors, which generate an initial guess for $Y_{n}$. The implicit method corrects the initial guess iteratively and provides a reasonable approximation to the solution of equation (2.5).

The predictor-corrector process for advancing the numerical solution to $\xi_{n}$ therefore consists of first generating a predicted value, denoted by $\mathrm{Y}_{n}^{[0]}$, and then correcting this initial estimate by iterating equation (2.5) to convergence. That is, starting with the initial guess $\underline{Y}_{n}^{[0]}$, approximations $\mathrm{Y}_{n}^{[m]}(m=1,2, \ldots, M)$ are generated (by using one of the techniques discussed below) until the magnitude of the difference in two successive approximations approaches zero within a specified accuracy. The quantity $\mathbf{Y}_{n}^{[m]}$ is the approximation obtained on the $m$ th iteration, the integer $M$ is the number of iterations required for convergence, and we accept $Y_{n}^{[M]}$ as an approximation to the exact solution y at $\xi_{n}$ and therefore denote it by $Y_{n}$ although, in general, it does not satisfy equation (2.5) exactly.
At each iteration $m$ the quantity $h_{n} \dot{\underline{Y}}_{n}^{[m]}$, which is defined here, is computed from $Y_{n}^{[m]}$ by the relation

$$
\begin{equation*}
\underline{\mathrm{Y}}_{n}^{[m]}=\underline{\Psi}_{n}+\beta_{0} h_{n} \dot{\mathbf{Y}}_{n}^{[m]} . \tag{2.7}
\end{equation*}
$$

Now, as discussed by Hindmarsh (ref. 21) and shown later in this section, if $\mathbb{Y}_{n}^{[m]}$ converges as $m \rightarrow \infty$, the limit, that is, $\lim _{m \rightarrow \infty} \underline{Y}_{n}^{[m]}$, must be a solution of equation (2.5) and $\dot{\mathrm{Y}}_{n}^{[m]}$ converges to $\left.\mathrm{f}_{n}\left[=\stackrel{m \rightarrow \infty}{\mathrm{f}} \mathbf{Y}_{n}\right)\right]$, the approximation to $\dot{\mathrm{y}}\left(\xi_{n}\right)$.

## 2. Description and Implementation of Methods

Hence $h_{n} \dot{\underline{Y}}_{n}^{[m]}$ is the $m$ th estimate for $h_{n} \underline{\mathrm{f}}_{n}$ and $\lim _{m \rightarrow \infty} h_{n} \dot{\underline{\dot{Y}}}_{n}^{[m]}=h_{n} \underline{\mathrm{f}}_{n}$. The predicted value of $h_{n} \underline{\mathrm{f}}_{n}$, denoted by $h_{n} \dot{\mathrm{Y}}_{n}^{[0]}$, is also obtained from equation (2.7) (by setting $m=0$ ). In practice, we terminate the calculation sequence at a finite number $M$ of iterations and accept as an approximation to $h_{n} \underline{\mathrm{f}}_{n}$ the quantity $h_{n} \dot{\mathrm{Y}}_{n} \equiv h_{n} \dot{\mathrm{Y}}_{n}^{[M]}$, which is obtained from $\underline{Y}_{n}^{[M]}$ by using equation (2.7). Note that $\underline{Y}_{n}$ is only an approximation to $\underline{\mathrm{f}}_{n}$ because $\underline{Y}_{n}^{[M]}$ does not, in general, satisfy equation (2.5) exactly (see eqs. (2.5) and (2.7)). Moreover, because $\dot{\mathrm{Y}}_{n}^{[M]}$ is defined to satisfy the solution method, in the sense of equation (2.7), it is not necessarily equal to $\underline{f}^{\mathrm{f}}\left(\underline{Y}_{n}^{[M]}\right)$. Therefore $\underline{Y}_{n}^{[M]}$ and $\underline{\mathrm{Y}}_{n}^{[M]}$ do not necessarily satisfy the ODE, equation (1.1). Thus, in practice, to advance the solution, the methods use the $\left\{\dot{\mathrm{Y}}_{j}\right\}$ (e.g., see eqs. (2.8a) and (2.8b)), rather than the $\left\{\mathrm{f}_{j}\right\}$ as written in equation (2.1).

After convergence of the estimates $\underline{Y}_{n}^{[m]}$, we could define $\dot{Y}_{n}^{[M]}$ to be equal to $\underline{\mathrm{f}}_{( }\left(\underline{Y}_{n}^{[M]}\right)$, so that $\underline{\mathrm{Y}}_{n}^{[M]}$ and $\underline{\mathrm{Y}}_{n}^{[M]}$ satisfy the ODE exactly. However, besides being more expensive because it will require one derivative evaluation, performing this operation is actually less stable for stiff equations than using equation (2.7) (ref. 25).

The predicted value at $\xi_{n}, \underline{Y}_{n}^{[0]}$, is generated by a $q$ th-order explicit formula similar to equations (2.3) and (2.4) (refs. 18 and 20):

$$
\begin{equation*}
\underline{Y}_{n}^{[0]}=\underline{Y}_{n-1}+h_{n} \sum_{j=1}^{q} \beta_{j}^{*} \underline{Y}_{n-j} \tag{2.8a}
\end{equation*}
$$

for the AM method of order $q$ and

$$
\begin{equation*}
\underline{Y}_{n}^{[0]}=\sum_{j=1}^{q} \alpha_{j}^{*} \underline{Y}_{n-j}+h_{n} \beta_{1}^{*} \underline{\dot{Y}}_{n-1} \tag{2.8b}
\end{equation*}
$$

for the BDF method of order $q$. In these two equations $\dot{\underline{Y}}_{n-j}$ is the approximation to $\underline{f}_{n-j}$ computed on the step $\left[\xi_{n-j-1}, \xi_{n-j}\right]$. The coefficients $\left\{\alpha_{j}^{*}\right\}$ and $\left\{\beta_{j}^{*}\right\}$ are selected such that equation (2.8a) or ( 2.8 b ) will be exact if the solution to equation (1.1) is a polynomial of degree $q$ or less.

The predictor step for the two methods can be generalized trivially as

$$
\begin{equation*}
\underline{Y}_{n}^{[0]}=\underline{\Psi}_{n}^{*}, \tag{2.9}
\end{equation*}
$$

where $\Psi^{*}$ is given by the right-hand sides of equations (2.8a) and (2.8b), respectively, for the AM and BDF methods.

To correct the initial estimate given by equation (2.9), that is, to solve equation (2.5), LSODE includes a variety of iteration techniques-functional, Newton-Raphson, and a variant of Jacobi-Newton.

### 2.2.1 Functional Iteration

To derive the functional iteration technique, also called simple iteration (refs. 11 and 26) and successive substitution (ref. 27), we rewrite equation (2.5) as follows:

$$
\begin{equation*}
\underline{\mathrm{Y}}_{n}=\Phi\left(\underline{\mathrm{Y}}_{n}\right), \tag{2.10}
\end{equation*}
$$

where

$$
\begin{equation*}
\underline{\Phi}\left(\underline{Y}_{n}\right)=\underline{\Psi}_{n}+h_{n} \beta_{0} \underline{f}\left(\underline{Y}_{n}\right) . \tag{2.11}
\end{equation*}
$$

The $(m+1)$ th estimate, $\mathbf{Y}_{n}^{[m+1]}(m=0,1, \ldots, M-1)$, is then obtained from equation (2.10) by (e.g., ref. 27)

$$
\begin{equation*}
\underline{\mathbf{Y}}_{n}^{[m+1]}=\underline{\Phi}\left(\underline{\mathbf{Y}}_{n}^{[m]}\right)=\underline{\Psi}_{n}+h_{n} \beta_{0} \underline{\mathrm{f}}\left(\underline{\mathbf{Y}}_{n}^{[m]}\right) . \tag{2.12}
\end{equation*}
$$

Now equation (2.7) gives the following expression for $h_{n} \dot{\underline{\dot{Y}}}_{n}^{[m+1]}$ :

$$
\begin{equation*}
\underline{\mathrm{Y}}_{n}^{[m+1]}=\underline{\Psi}_{n}+\beta_{0} h_{n} \dot{\underline{\dot{X}}}_{n}^{[m+1]} . \tag{2.13}
\end{equation*}
$$

Comparing equations (2.12) and (2.13) gives

$$
\begin{equation*}
h_{n} \dot{\hat{Y}}_{n}^{[m+1]}=h_{n} \underline{\underline{£}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \tag{2.14}
\end{equation*}
$$

for functional iteration.
We now define the vector function $g(\underline{y})$ by

$$
\begin{equation*}
\underline{\mathrm{g}}(\underline{\mathrm{y}})=h_{n} \underline{\mathrm{f}}(\underline{\mathrm{y}})+\frac{\underline{\psi}_{n}-\underline{\mathrm{y}}}{\beta_{0}}, \tag{2.15}
\end{equation*}
$$

which, upon using equation (2.7), gives

$$
\begin{equation*}
\mathrm{g}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)=h_{n} \underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)-h_{n} \dot{\mathrm{Y}}_{n}^{[m]} . \tag{2.16}
\end{equation*}
$$

By using equation (2.15) we can rewrite the functional iteration equation (2.12) as follows:

$$
\begin{equation*}
\underline{\mathrm{Y}}_{n}^{[m+1]}=\underline{\mathrm{Y}}_{n}^{[m]}+\beta_{0} \mathrm{~g}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \tag{2.17}
\end{equation*}
$$

Finally the combination of equations (2.14) and (2.16) produces the following functional iteration procedure for $h_{n} \underline{\dot{Y}}_{n}$ :

$$
\begin{equation*}
h_{n} \dot{\mathrm{Y}}_{n}^{[m+1]}=h_{n} \dot{\mathrm{Y}}_{n}^{[m]}+\underline{\mathrm{g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) . \tag{2.18}
\end{equation*}
$$

Equation (2.17) is simple to use, but it converges only linearly (ref. 27). In addition, for successful convergence the step size may be restricted to very small values for stiff problems (refs. $4,10,12,26$, and 28 ), as shown here. By using equation (2.14) we can rewrite equation (2.16) as

$$
\begin{equation*}
\mathrm{g}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)=h_{n} \underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)-h_{n} \underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}^{[m-1]}\right), \tag{2.19}
\end{equation*}
$$

for $m \geq 1$. Hence, equation (2.17) can be rewritten as

$$
\begin{equation*}
\underline{\mathrm{Y}}_{n}^{[m+1]}=\underline{\mathrm{Y}}_{n}^{[m]}+h_{n} \beta_{0}\left[\underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)-\underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}^{[m-1]}\right)\right] . \tag{2.20}
\end{equation*}
$$

By using the Lipschitz condition, equation (1.3), we get the following relation from equation (2.20):

$$
\begin{equation*}
\left\|\underline{\mathrm{Y}}_{n}^{[m+1]}-\underline{\mathrm{Y}}_{n}^{[m]}\right\| \leq\left|h_{n}\right| \beta_{0} g\left\|_{n}^{[m]}-\underline{\mathrm{Y}}_{n}^{[m-1]}\right\|, \tag{2.21}
\end{equation*}
$$

which shows that the iteration converges, that is, the successive differences

$$
\left\|\underline{\mathrm{Y}}_{n}^{[m+1]}-\underline{\mathrm{Y}}_{n}^{[m]}\right\|
$$

decrease, only if

$$
\begin{equation*}
\left|h_{n}\right| \beta_{0} \mathscr{L}<1 \tag{2.22}
\end{equation*}
$$

Now stiff problems are characterized by, and often referred to as systems with, large Lipschitz constants (e.g., refs. 4, 12, and 26), and so equation (2.22) restricts the step size to very small values. Indeed, the restriction imposed by this inequality on $h_{n}$ is exactly of the same form as that imposed by stability requirements on classical methods, such as the explicit Runge-Kutta method (refs. 4 and 26). For this reason, when functional iteration is used, the integration method is usually said to be explicit even though it is implicit (ref. 17).

### 2.2.2 Newton-Raphson Iteration

Newton-Raphson (NR) iteration, on the other hand, converges quadratically and can use much larger step sizes than functional iteration (refs. 27, 29, and 30). Rapid improvement in the accuracy of the estimates is especially important because the corrector is iterated to convergence. The reason for iterating to convergence is to preserve the stability characteristics of the corrector. If the correction process is terminated after a fixed number of iterations, the stability characteristics of the corrector are lost (refs. 4 and 12), with disastrous consequences for stiff problems.

To derive the NR iteration procedure, we rewrite equation (2.5) as

$$
\begin{equation*}
\underline{\mathrm{R}}\left(\underline{\mathrm{Y}}_{n}\right)=\underline{\mathrm{Y}}_{n}-\underline{\Psi}_{n}-h_{n} \beta_{0} \underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}\right)=0 \tag{2.23}
\end{equation*}
$$

so that solving equation (2.5) is equivalent to finding the zero of $\underline{R}$. The quantity $\mathrm{R}\left(\mathrm{Y}_{n}^{[m]}\right)$ is the residual vector on the $m$ th iteration; that is, it is the amount by which $Y_{n}^{[m]}$ fails to satisfy equation (2.5). To obtain the ( $m+1$ )th estimate, we expand equation (2.23) in a Taylor series about the $m$ th estimate, neglect the second and higher derivatives, and set $\mathbb{R}\left(Y_{n}^{[m+1]}\right)=0$ because we seek a $Y_{n}^{[m+1]}$ that produces this result (e.g., ref. 27). Performing these operations and then rearranging terms give the following relation for the NR iteration technique:

$$
\begin{equation*}
\mathbf{P}\left(\underline{\mathrm{Y}}_{n}^{[m+1]}-\underline{\mathrm{Y}}_{n}^{[m]}\right)=-\underline{\mathrm{R}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)=\underline{\Psi}_{n}+h_{n} \beta_{0} \underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)-\underline{\mathrm{Y}}_{n}^{[m]} \tag{2.24}
\end{equation*}
$$

where the $N \times N$ matrix $\mathbf{P}$ is given by

$$
\begin{equation*}
\mathbf{P}=\partial \underline{\mathrm{R}} / \partial \underline{\mathrm{Y}}=\mathbf{I}-h_{n} \beta_{0} \mathbf{J} \tag{2.25}
\end{equation*}
$$

## 2. Description and Implementation of Methods

In equation (2.25), $\mathbf{I}$ is the $N \times N$ identity matrix and $\mathbf{J}$ is the Jacobian matrix, equation (1.5). Comparing equations (2.15) and (2.23) shows that

$$
\begin{equation*}
\underline{\mathrm{R}}(\underline{\mathrm{Y}})=-\beta_{0} \underline{\mathrm{~g}}(\underline{\mathrm{Y}}), \tag{2.26}
\end{equation*}
$$

so that equation (2.24) can be rewritten as follows:

$$
\begin{equation*}
\underline{Y}_{n}^{[m+1]}=\underline{Y}_{n}^{[m]}+\beta_{0} \mathbf{P}^{-1} \underline{g}\left(\underline{Y}_{n}^{[m]}\right) \tag{2.27}
\end{equation*}
$$

The NR iteration procedure for $h_{n} \dot{\mathrm{Y}}_{n}$ is derived by subtracting equation (2.7) from equation (2.13) and then using equation (2.27). The result is

$$
\begin{equation*}
h_{n} \dot{\mathrm{Y}}_{n}^{[m+1]}=h_{n} \dot{\mathrm{Y}}_{n}^{[m]}+\mathbf{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \tag{2.28}
\end{equation*}
$$

This iteration will converge provided that the predicted value is sufficiently accurate (refs. 4 and 12). The prediction method, equation (2.9), provides a sufficiently accurate initial estimate that the average number of iterations per step is less than 1.5 (ref. 7). In fact, the predictor is generally as accurate as the corrector, which is nonetheless needed for numerical stability. However, much computational work is required to form the Jacobian matrix and to perform the linear algebra necessary to solve equation (2.27). Now, because the Jacobian does not appear explicitly in the ODE's, equation (1.1), or in the solution method, equation (2.5), $\mathbf{J}$ need not be very accurate. Therefore, for problems in which the analytical Jacobian matrix is difficult or impossible to evaluate, a fairly crude approximation such as the finite-difference quotient

$$
\begin{equation*}
J_{i j}=\frac{f_{i}\left(Y_{j}+\Delta Y_{j}\right)-f_{i}\left(Y_{j}\right)}{\Delta Y_{j}}, \quad i, j=1, \ldots, N \tag{2.29}
\end{equation*}
$$

is adequate. In equation (2.29), $\Delta Y_{j}$ is a suitable increment for the $j$ th component of $Y$.

Inaccuracies in the iteration matrix may affect the rate of convergence of the solution but not the solution if it converges (refs. 4 and 21). Hence this matrix need only be accurate enough for the iteration to converge. This beneficial fact can be used to reduce the computational work associated with linear algebra, as described in chapter 3.

### 2.2.3 Jacobi-Newton Iteration

Jacobi-Newton (JN) iteration (ref. 31), also called Jacobi iteration (ref. 32), is obtained from Newton-Raphson iteration by neglecting all off-diagonal elements of the Jacobian matrix. Hence for JN iteration

$$
J_{i j}=\left\{\begin{array}{l}
0, \quad i \neq j  \tag{2.30}\\
\partial f_{i} / \partial y_{j}, \quad i=j
\end{array}\right.
$$

This technique is as simple to use as functional iteration because it does not require any matrix algebra. Also, it converges faster than functional iteration but, in general, not as fast as NR iteration.

A method closely resembling JN iteration is implemented as a separate method option in LSODE. It is like JN iteration in that it uses a diagonal approximation $\mathbf{D}$ to the Jacobian matrix. However, the diagonal elements $D_{i i}$ are, in general, different from $J_{i i}$ and are given by the difference quotient

$$
\begin{equation*}
D_{i i}=\frac{f_{i}(\underline{\mathrm{Y}}+\Delta \underline{\mathrm{Y}})-f_{i}(\underline{\mathrm{Y}})}{\Delta Y_{i}}, \quad i=1, \ldots, N \tag{2.31}
\end{equation*}
$$

where the increment vector $\Delta Y=0.1 \beta_{0} \underline{g}\left(Y_{n}^{[0]}\right)$. If J is actually a diagonal matrix, $D_{i i}=J_{i i}+\mathrm{O}\left(\Delta Y_{i}^{2}\right)$, but, in general, $D_{i i}$ effectively lumps together the various elements $\left\{J_{i j}\right\}$ in row i of $J$.

### 2.2.4 Unified Formulation

The different iteration methods can be generalized by the recursive relations

$$
\begin{equation*}
\underline{Y}_{n}^{[m+1]}=\underline{Y}_{n}^{[m]}+\beta_{0} \mathbf{P}^{-1} \underline{g}\left(\underline{Y}_{n}^{[m]}\right) \tag{2.32}
\end{equation*}
$$

and

$$
\begin{equation*}
h_{n} \dot{\mathrm{Y}}_{n}^{[m+1]}=h_{n} \dot{\mathrm{Y}}_{n}^{[m]}+\mathbf{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \tag{2.33}
\end{equation*}
$$

where $\mathbf{P}$ depends on the iteration method. For functional iteration $\mathbf{P}=\mathbf{I}$, and for NR and JN iterations $\mathbf{P}$ is given by equation (2.25), where $\mathbf{J}$ is the appropriate Jacobian matrix, equation (1.5), (2.30), or (2.31).

## 2. Description and Implementation of Methods

The combination of equations (2.32) and (2.33) gives

$$
\begin{equation*}
\underline{Y}_{n}^{[m+1]}-\underline{Y}_{n}^{[m]}=h_{n} \beta_{0}\left(\dot{\underline{Y}}_{n}^{[m+1]}-\underline{\mathrm{Y}}_{n}^{[m]}\right) \tag{2.34}
\end{equation*}
$$

which shows that if $\underline{Y}_{n}^{[m]}$ converges as $m \rightarrow \infty$, so does $\dot{Y}_{n}^{[m]}$. Equation (2.32) shows that if $\underline{Y}_{n}^{[m]}$ converges (to $\underline{Y}_{n}$ ) as $m \rightarrow \infty, \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \rightarrow 0$, and therefore we see from equations (2.15) and (2.16), respectively, (1) that the converged solution satisfies equation (2.5) and (2) that $\dot{\underline{Y}}_{n}^{[m]} \rightarrow \underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}\right)=\underline{\mathrm{f}}_{n}$.

The predictor-corrector methods can be summarized as follows:

Predictor:

$$
\left.\begin{array}{c}
\underline{\mathrm{Y}}_{n}^{[0]}=\underline{\Psi}_{n}^{*}  \tag{2.35}\\
h_{n} \dot{\mathrm{Y}}_{n}^{[0]}=\frac{\underline{\mathrm{Y}}_{n}^{[0]}-\underline{\Psi}_{n}}{\beta_{0}} .
\end{array}\right\}
$$

## Corrector:

$$
\left.\begin{array}{c}
\underline{\mathrm{g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)=h_{n} \underline{f}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)-h_{n} \underline{\mathrm{Y}}_{n}^{[m]} \\
\underline{\mathrm{Y}}_{n}^{[m+1]}=\underline{\mathrm{Y}}_{n}^{[m]}+\beta_{0} \mathbf{P}^{-1} \mathrm{~g}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \\
h_{n} \dot{\underline{\mathrm{Y}}}_{n}^{[m+1]}=h_{n} \underline{\mathrm{Y}}_{n}^{[m]}+\mathbf{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \tag{2.36}
\end{array}\right\} m=0,1, \ldots, M-1 .
$$

### 2.3 Matrix Formulation

The implementation of linear multistep methods is aided by a matrix formulation (ref. 21). This formulation, constructed by Gear (ref. 18), is summarized here.

To solve for $Y_{n}$ and $h_{n} \underline{Y}_{n}$ by using equations (2.35) to (2.37), we need, and therefore must have saved, the $L=q+1$ column vectors $Y_{n-1}, h_{n} \dot{X}_{n-1}, h_{n} \dot{X}_{n-2}, \ldots$, and $h_{n} \dot{\mathbf{Y}}_{n-q}$ for the AM method of order $q$, or $Y_{n-1}, Y_{n-2}, \ldots, Y_{n-q}$, and $h_{n} \dot{\underline{Y}}_{n-1}$ for the BDF method of order $q$. Hence for the AM method of order $q$ we define the $N \times L$ history matrix $w_{n-1}$ at $\xi_{n-1}$ by

$$
\begin{equation*}
\mathbf{w}_{n-1}=\left(\underline{\mathrm{Y}}_{n-1}, h_{n} \underline{\dot{Y}}_{n-1}, h_{n} \underline{\underline{Y}}_{n-2}, \ldots, h_{n} \underline{\dot{Y}}_{n-q}\right) \tag{2.38a}
\end{equation*}
$$

that is,

$$
\mathbf{w}_{n-1}=\left(\begin{array}{ccccc}
Y_{1, n-1} & h_{n} \dot{Y}_{1, n-1} & h_{n} \dot{Y}_{1, n-2} & \ldots & h_{n} \dot{Y}_{1, n-q}  \tag{2.39}\\
Y_{2, n-1} & h_{n} \dot{Y}_{2, n-1} & h_{n} \dot{Y}_{2, n-2} & \ldots & h_{n} \dot{Y}_{2, n-q} \\
\cdot & \cdot & \cdot & & \cdot \\
\cdot & \cdot & \cdot & & \cdot \\
\cdot & \cdot & \cdot & & \cdot \\
\cdot & \cdot & \cdot & & \cdot \\
\cdot & \cdot & \cdot & & \cdot \\
Y_{N, n-1} & h_{n} \dot{Y}_{N, n-1} & h_{n} \dot{Y}_{N, n-2} & \ldots & h_{n} \dot{Y}_{N, n-q}
\end{array}\right)
$$

The updated matrix

$$
\begin{equation*}
\mathbf{w}_{n}=\left(\underline{Y}_{n}, h_{n} \dot{\mathbf{Y}}_{n}, h_{n} \dot{\underline{Y}}_{n-1}, \ldots, h_{n} \dot{\underline{Y}}_{n-q+1}\right) \tag{2.40a}
\end{equation*}
$$

is then constructed at each step $\xi_{n}$. The predicted matrix $w_{n}^{[0]}$ at $\xi_{n}$ is given by

$$
\begin{equation*}
\mathbf{w}_{n}^{[0]}=\left(\underline{\mathrm{Y}}_{n}^{[0]}, h_{n} \dot{\mathrm{Y}}_{n}^{[0]}, h_{n} \dot{\mathrm{Y}}_{n-1}, \ldots, h_{n} \dot{\mathbf{Y}}_{n-q+1}\right) . \tag{2.41a}
\end{equation*}
$$

For the BDF method of order $q$ these matrices take the form

$$
\begin{align*}
\mathbf{W}_{n-1} & =\left(\underline{Y}_{n-1}, h_{n} \underline{\dot{Y}}_{n-1}, \underline{Y}_{n-2}, \ldots, \underline{Y}_{n-q}\right)  \tag{2.38b}\\
\mathbf{W}_{n} & =\left(\underline{Y}_{n}, h_{n} \dot{\underline{Y}}_{n}, \underline{Y}_{n-1}, \ldots, \underline{Y}_{n-q+1}\right) \tag{2.40b}
\end{align*}
$$

## 2. Description and Implementation of Methods

and

$$
\begin{equation*}
\mathbf{w}_{n}^{[0]}=\left(\underline{\mathrm{Y}}_{n}^{[0]}, h_{n} \dot{\mathrm{Y}}_{n}^{[0]}, \underline{\mathrm{Y}}_{n-1}, \ldots, \underline{\mathrm{Y}}_{n-q+1}\right) . \tag{2.41b}
\end{equation*}
$$

The matrix formulations for $\mathbf{w}_{n}^{[0]}$ and $\mathbf{w}_{n}$ are derived as follows: Substituting the expression for $\underline{Y}_{n}^{[0]}$, equation (2.8a) or (2.8b), into that for $h_{n} \dot{\mathrm{Y}}_{n}^{[0]}$, equation (2.35), and then using equation (2.6a) or (2.6b) give

$$
\begin{equation*}
h_{n} \dot{\mathrm{Y}}_{n}^{[0]}=\sum_{j=1}^{q-1}\left(\frac{\beta_{j}^{*}-\beta_{j}}{\beta_{0}}\right) h_{n} \dot{\underline{Y}}_{n-j}+\frac{\beta_{q}^{*}}{\beta_{0}} h_{n} \dot{\underline{Y}}_{n-q} \tag{2.42a}
\end{equation*}
$$

for the AM method of order $q$ and

$$
\begin{equation*}
h_{n} \dot{\mathrm{Y}}_{n}^{[0]}=\sum_{j=1}^{q}\left(\frac{\alpha_{j}^{*}-\alpha_{j}}{\beta_{0}}\right) \underline{Y}_{n-j}+\frac{\beta_{1}^{*}}{\beta_{0}} h_{n} \dot{\underline{Y}}_{n-1} \tag{2.42b}
\end{equation*}
$$

for the BDF method of order $q$. Equations (2.8a) and (2.42a), or (2.8b) and ( 2.42 b ), that is, the prediction process, can be rewritten as the matrix equation

$$
\begin{equation*}
\mathbf{w}_{n}^{[0]}=\mathbf{w}_{n-1} \mathbf{B}, \tag{2.43}
\end{equation*}
$$

where the $L \times L$ matrix $\mathbf{B}$ depends on the solution method. For the AM method of order $q$, it is given by
and for the BDF method of order $q$,

The corrector equation, equation (2.36), can be expressed in matrix form as

$$
\begin{equation*}
\mathbf{w}_{n}^{[m+1]}=\mathbf{w}_{n}^{[m]}+\mathbf{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \underline{\mathbf{k}}, \tag{2.45}
\end{equation*}
$$

where $\mathbf{w}_{n}^{[m]}$, the history matrix on the $m$ th iteration, is given by

$$
\begin{equation*}
\mathbf{W}_{n}^{[m]}=\left(\underline{Y}_{n}^{[m]}, h_{n} \dot{\underline{Y}}_{n}^{[m]}, h_{n} \dot{\underline{Y}}_{n-1}, \ldots, h_{n} \dot{\mathbf{Y}}_{n-q+1}\right) \tag{2.46a}
\end{equation*}
$$

for the AM method and by
2. Description and Implementation of Methods

$$
\begin{equation*}
\mathbf{w}_{n}^{[m]}=\left(\underline{\mathrm{Y}}_{n}^{[m]}, h_{n} \dot{\underline{\mathrm{Y}}}_{n}^{[m]}, \underline{\mathrm{Y}}_{n-1}, \ldots, \underline{\mathrm{Y}}_{n-q+1}\right) \tag{2.46b}
\end{equation*}
$$

for the BDF method, $\underline{\mathfrak{k}}$ is the $L$-dimensional vector

$$
\begin{equation*}
\underline{k}=\left(\beta_{0}, 1,0, \ldots, 0\right), \tag{2.47}
\end{equation*}
$$

and $\mathbf{P}$ depends on the iteration technique, as described in section 2.2.4.
The matrix formulation of the methods can be summarized as follows:
Predictor:

$$
\begin{equation*}
\mathbf{w}_{n}^{[0]}=\mathbf{w}_{n-1} \mathbf{B} . \tag{2.48}
\end{equation*}
$$

Corrector:

$$
\begin{gather*}
\underline{\mathrm{g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)=h_{n} \underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)-h_{n} \dot{\underline{\mathrm{Y}}}_{n}^{[m]}  \tag{2.49}\\
\left.\mathbf{w}_{n}^{[m+1]}=\mathbf{w}_{n}^{[m]}+\mathrm{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \underline{\mathbf{k}}\right]  \tag{2.50}\\
\mathbf{w}_{n}=\mathbf{w}_{n}^{[M]} .
\end{gather*}
$$

### 2.4 Nordsieck's History Matrix

Instead of saving information in the form $\mathbf{w}_{n-1}$, equation (2.38a) or (2.38b), Gear (ref. 18) suggested making a linear transformation and storing the matrix $\mathbf{z}_{n-1}$ given by

$$
\begin{equation*}
\mathbf{z}_{n-1}=\mathbf{w}_{n-1} \mathbf{Q}, \tag{2.51}
\end{equation*}
$$

where the $L \times L$ transformation matrix $\mathbf{Q}$ is nonsingular. In particular, $\mathbf{Q}$ is chosen such that the matrix representation suggested by Nordsieck (ref. 33) is obtained:

$$
\begin{equation*}
\mathrm{x}_{n-1}=\left(\underline{\mathrm{Y}}_{n-1}, h_{n} \dot{\underline{Y}}_{n-1}, \frac{h_{n}^{2}}{2!} \ddot{\underline{Y}}_{n-1}, \ldots, \frac{h_{n}^{q}}{q!} \underline{\mathrm{Y}}_{n-1}^{(q)}\right) \tag{2.52}
\end{equation*}
$$

that is, the $N \times L$ matrix $z_{n-1}$ is given by

In equation (2.53), $Y_{i, n-1}^{(j)}$ is the $j$ th derivative of the approximating polynomial for $Y_{i, n-1}$. Because scaled derivatives $h_{n}^{j} Y^{(j)}{ }_{n-1} / j$ ! are used, $\mathbf{Q}$ is independent of the step size. However, $\mathbf{Q}$ depends on the solution method. The $N$ rows of $z_{n-1}$ are numbered from 1 to $N$, so that the $i$ th row ( $i=1, \ldots, N$ ) contains the $q+1$ scaled derivatives of the $i$ th component, $Y_{i, n-1}$, of $Y_{n-1}$. The $q+1$ columns are, however, numbered from 0 to $q$, so that the column number corresponds to the order of the scaled derivative stored in that column. Thus the $j$ th column $(j=0,1, \ldots, q)$, which we denote by the vector $\mathrm{z}_{n-1}(j)$, contains the vector $h_{n}^{j} \mathrm{Y}_{n-1}^{(j)} j$ !. The Nordsieck matrix formulation of the method is referred to as the "normal form of the method" (ref. 10).

Applying the appropriate transformation matrix $\mathbf{Q}$ to the predictor equation, equation (2.48), gives

$$
\begin{equation*}
\mathbf{z}_{n}^{[0]}=\mathbf{w}_{n}^{[0]} \mathbf{Q}=\mathbf{w}_{n-1} \mathbf{B} \mathbf{Q}=\mathbf{z}_{n-1} \mathbf{Q}^{-1} \mathbf{B} \mathbf{Q}=\mathbf{z}_{n-1} \mathbf{A} \tag{2.54}
\end{equation*}
$$

where

$$
\begin{equation*}
\mathbf{z}_{n}^{[0]}=\left(\underline{\mathrm{Y}}_{n}^{[0]}, h_{n} \dot{\mathrm{Y}}_{n}^{[0]}, \frac{h_{n}^{2}}{2!} \ddot{\underline{\mathrm{Y}}}_{n}^{[0]}, \ldots, \frac{h_{n}^{q}}{q!} \underline{\mathrm{Y}}_{n}^{[0]}(q)\right) \tag{2.55}
\end{equation*}
$$

is the predicted $N \times L$ Nordsieck history matrix at $\xi_{n}$ and

$$
\begin{equation*}
\mathbf{A}=\mathbf{Q}^{-1} \mathbf{B} \mathbf{Q} \tag{2.56}
\end{equation*}
$$

The $L \times L$ prediction matrix A provides a $q$ th-order approximation to $\mathbf{z}_{n}^{[0]}$ in terms of $\mathbf{z}_{n-1}$ and is therefore the lower-triangular Pascal triangle matrix (ref. 10), with element $A_{i j}$ given by

$$
A_{i j}=\left\{\begin{array}{l}
0, \quad i<j  \tag{2.57}\\
\binom{i}{j}, \quad i \geq j
\end{array} \quad i, j=0,1, \ldots, q\right.
$$

where $\binom{i}{j}$ is the binomial coefficient, defined as

$$
\begin{equation*}
\binom{i}{j}=\frac{i!}{j!(i-j)!} . \tag{2.58}
\end{equation*}
$$

Hence


The principal advantage of using the Nordsieck history matrix is that the matrix multiplication implied by equation (2.54) can be carried out solely by repeated additions, as shown by Gear (ref. 10). Hence computer multiplications are
avoided, resulting in considerable savings of computational effort for large problems. Also $\mathbf{A}$ need not be stored and $\mathbf{z}_{n}^{[0]}$ overwrites $\mathbf{z}_{n-1}$, thereby reducing memory requirements.

Because

$$
\begin{equation*}
\binom{i+1}{j+1}=\binom{i}{j+1}+\binom{i}{j} \tag{2.60}
\end{equation*}
$$

and $A_{i i}=A_{i 0}=1$ for all $i$, the product $\mathbf{z A}$ is computed as follows (refs. 10 and 15 ):
For $k=0,1, \ldots, q-1$, do:

$$
\left\{\begin{align*}
\text { For } & j=q, q-1, \ldots, k+1, \text { do: }  \tag{2.61}\\
& z_{i, j-1} \leftarrow z_{i, j}+z_{i, j-1}, \quad i=1, \ldots, N .
\end{align*}\right.
$$

In this equation the subscripts $n$ and $n-1$ have been dropped because the $\mathbf{z}$ values do not indicate any one value of $\xi$ but represent a continuous replacement process. At the start of the calculation procedure given by equation (2.61), $\mathrm{z}=\mathrm{z}_{n-1}$; and at the end $\mathbf{z}=\mathbf{z}_{n}^{[0]}$. The arrow " $\leftarrow$ " denotes the replacement operator, which means overwriting the contents of a computer storage location. For example,

$$
z_{i, 3} \leftarrow z_{i, 4}+z_{i, 3}
$$

means that $z_{i, 4}$ is added to $z_{i, 3}$ and the result replaces the contents of the location $z_{i, 3}$. The total number of additions required in equation (2.61) is $N q(q+1) / 2$. The predictor step is a Taylor series expansion about the previous point $\xi_{n-1}$ and is independent of both the integration method and the ODE.

Another important advantage of using Nordsieck's formulation is that it makes changing step size easy. For example, if at $\xi_{n}$ the step size is changed from $h_{n}$ to $r h_{n}$, the new history matrix is obtained from

$$
\begin{equation*}
\mathbf{z}_{n} \leftarrow \mathbf{z}_{n} \mathbf{C}, \tag{2.62}
\end{equation*}
$$

where the $L \times L$ diagonal matrix $\mathbf{C}$ is given by

$$
\mathbf{C}=\left(\begin{array}{llllll}
1 & & & & & 0  \tag{2.63}\\
& r & & & & \\
& & r^{2} & & & \\
& & & \cdot & & \\
& & & & \cdot & \\
& 0 & & & & \\
& & & & & r^{q}
\end{array}\right) .
$$

## 2. Description and Implementation of Methods

The rescaling can be done by multiplications alone, as follows:

$$
\begin{align*}
& R=1 \\
& \text { For } j=1, \ldots, q, \text { do : } \\
& \qquad\left\{\begin{array}{l}
R \leftarrow r R \\
z_{i, j} \leftarrow z_{i, j} R, i=1, \ldots, N .
\end{array}\right. \tag{2.64}
\end{align*}
$$

The corrector equation corresponding to equation (2.49) is given by

$$
\begin{equation*}
\mathbf{z}_{n}^{[m+1]}=\mathbf{w}_{n}^{[m+1]} \mathbf{Q}=\mathbf{w}_{n}^{[m]} \mathbf{Q}+\mathbf{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \underline{k} \mathbf{Q}=\mathbf{z}_{n}^{[m]}+\mathbf{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \underline{\ell}, \tag{2.65}
\end{equation*}
$$

where $\mathrm{z}_{n}^{[m]}$, the Nordsieck history matrix on the $m$ th iteration, is given by

$$
\begin{equation*}
\mathbf{z}_{n}^{[m]}=\left(\underline{\mathrm{Y}}_{n}^{[m]}, h_{n} \underline{\mathrm{Y}}_{n}^{[m]}, \frac{h_{n}^{2}}{2!} \ddot{\mathrm{Y}}_{n}^{[m]}, \ldots, \frac{h_{n}^{q}}{q!} \underline{\mathrm{Y}}_{n}^{[m]}{ }^{(q)}\right) \tag{2.66}
\end{equation*}
$$

and

$$
\begin{equation*}
\underline{\ell}=k \mathbf{Q} \tag{2.67}
\end{equation*}
$$

is an $L$-dimensional vector

$$
\begin{equation*}
\ell=\left(\ell_{0}, \ell_{1}, \ldots, \ell_{q}\right) \tag{2.68}
\end{equation*}
$$

For the two solution methods used in LSODE the values of $\underline{\ell}$ are derived in references 21 and 22 and reproduced in tables 2.1 and 2.2. Methods expressed in the form of equations (2.54) and (2.65) are better described as multivalue or $L$ value methods than multistep methods (ref. 10) because it is the number $L$ of values saved from step to step that is significant and not the number of steps involved.

The two matrix formulations described here are related by the transformation equations (2.51), (2.54), and (2.65) and are therefore said to be equivalent (ref. 10). The equivalence means that if the step $\left[\xi_{n-1}, \xi_{n}\right]$ is taken by the two methods with equivalent past values $\mathbf{w}_{n-1}$ and $\mathbf{z}_{n-1}$, that is, related by equation (2.51) through $\mathbf{Q}$, then the resulting solutions $\mathbf{w}_{n}$ and $\mathbf{z}_{n}$ will also be related by equation (2.51) through $\mathbf{Q}$, apart from roundoff errors (ref. 21). The transformation does not affect the stability properties or the accuracy of the

TABLE 2.1.-METHOD COEFFICIENTS FOR ADAMS-MOULTON METHOD IN NORMAL FORM OF ORDERS 1 TO 12

| 9 | $\ell_{0}$ | $\ell_{1}$ | $\ell_{2}$ | ${ }^{2}$ | ${ }^{2} 4$ | ${ }_{5}$ | $\ell_{6}$ | ${ }^{7}$ | ${ }_{8}$ | ${ }^{\prime} 9$ | ${ }^{2} 10$ | ${ }^{\ell} 11$ | $\ell_{12}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 1 | 1 |  |  |  |  |  |  |  |  |  |  |  |
| 2 | $\frac{1}{2}$ | 1 | $\frac{1}{2}$ |  |  |  |  |  |  |  |  |  |  |
| 3 | $\frac{5}{12}$ | 1 | $\frac{3}{4}$ | $\frac{1}{6}$ | 1 |  |  |  |  |  |  |  |  |
| 4 | $\frac{3}{8}$ | 1 | $\frac{11}{12}$ | $\frac{1}{3}$ | $\frac{1}{24}$ |  |  |  |  |  |  |  |  |
| 5 | $\frac{251}{720}$ | 1 | $\frac{25}{24}$ | $\frac{35}{72}$ | $\frac{5}{4 B}$ | $\frac{1}{120}$ |  |  |  |  |  |  |  |
| 5 | $\frac{95}{288}$ | 1 | $\frac{137}{120}$ | $\frac{5}{8}$ | $\frac{17}{96}$ | $\frac{1}{40}$ | $\frac{1}{720}$ |  |  |  |  |  |  |
| 7 | $\frac{19087}{60480}$ | 1 | $\frac{49}{40}$ | $\frac{203}{270}$ | $\frac{49}{192}$ | $\frac{7}{144}$ | $\frac{7}{1440}$ | $\frac{1}{5040}$ |  |  |  |  |  |
| 8 | $\frac{5257}{17280}$ | 1 | $\frac{363}{280}$ | $\frac{469}{540}$ | $\frac{967}{2880}$ | $\frac{7}{90}$ | $\frac{23}{2160}$ | $\frac{1}{1260}$ | $\frac{1}{40320}$ |  |  |  |  |
| 9 | $\frac{1070017}{3628800}$ | 1 | $\frac{761}{560}$ | $\frac{29531}{30240}$ | $\frac{267}{640}$ | $\frac{1069}{9600}$ | $\frac{3}{160}$ | $\frac{13}{6720}$ | $\frac{1}{8960}$ | $\frac{1}{362880}$ |  |  |  |
| 10 | $\frac{25713}{89600}$ | 1 | $\frac{7129}{5040}$ | $\frac{6515}{6048}$ | $\frac{4523}{9072}$ | $\frac{19}{128}$ | $\frac{3013}{103680}$ | $\frac{5}{1344}$ | $\frac{29}{96768}$ | $\frac{1}{72576}$ | $\frac{1}{3628800}$ |  |  |
| 11 | $\frac{26842253}{95800320}$ | 1 | $\frac{7381}{5040}$ | $\frac{177133}{151200}$ | $\frac{84095}{145152}$ | $\frac{341693}{1814400}$ | $\frac{8591}{207360}$ | $\frac{7513}{1209600}$ | $\frac{121}{193536}$ | $\frac{11}{272160}$ | $\frac{11}{7257600}$ | $\frac{1}{39916800}$ |  |
| 12 | $\frac{4777223}{17418240}$ | 1 | $\frac{83711}{55440}$ | $\frac{190553}{151200}$ | $\frac{341747}{518400}$ | $\frac{139381}{604800}$ | $\frac{242537}{4354560}$ | $\frac{1903}{201600}$ | $\frac{10831}{9676800}$ | $\frac{11}{120960}$ | $\frac{1}{207360}$ | $\frac{1}{6652800}$ | $\frac{1}{479001600}$ |

TABLE 2.2.-METHOD COEFFICIENTS FOR BACKWARD DIFFERENTIATION FORMULA METHOD IN NORMAL FORM OF ORDERS 1 TO 6

| $q$ | $\ell_{0}$ | $\ell_{1}$ | $\ell_{2}$ | $\ell_{3}$ | $\ell_{4}$ | $\ell_{5}$ | $\ell_{6}$ |
| :---: | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 1 | 1 | 1 |  |  |  |  |  |
| 2 | $\frac{2}{3}$ | $\frac{3}{3}$ | $\frac{1}{3}$ |  |  |  |  |
| 3 | $\frac{6}{11}$ | $\frac{11}{11}$ | $\frac{6}{11}$ | $\frac{1}{11}$ |  |  |  |
| 4 | $\frac{24}{50}$ | $\frac{50}{50}$ | $\frac{35}{50}$ | $\frac{10}{50}$ | $\frac{1}{50}$ |  |  |
| 5 | $\frac{120}{274}$ | $\frac{274}{274}$ | $\frac{225}{274}$ | $\frac{85}{274}$ | $\frac{15}{274}$ | $\frac{1}{274}$ |  |
| 6 | $\frac{720}{1764}$ | $\frac{1764}{1764}$ | $\frac{1624}{1764}$ | $\frac{735}{1764}$ | $\frac{175}{1764}$ | $\frac{21}{1764}$ | $\frac{1}{1764}$ |

method, but roundoff properties and computational effort depend on the representation used, as discussed by Gear (ref. 10).

The first two columns of $\mathbf{z}_{n}$ and $\mathbf{w}_{n}$ are identical (see eqs. (2.38a), (2.38b), and (2.52)), and so $\ell_{0}=\beta_{0}$ and $\ell_{1}=1$. For the same reason the corrector iteration procedures for $\underline{Y}_{n}$ and $h_{n} \dot{\mathrm{Y}}_{n}$ remain unchanged (see eqs. (2.45), (2.47), and (2.65)). However, to facilitate estimation of the local truncation error, a different iteration procedure than that given by equation (2.65) is used. To derive the new formulation, $\mathbf{z}_{n}^{[m+1]}$ is written as

$$
\mathbf{z}_{n}^{[m+1]}=\mathbf{z}_{n}^{[m+1]}-\mathbf{z}_{n}^{[m]}+\mathbf{z}_{n}^{[m]}-\mathbf{z}_{n}^{[m-1]}+\ldots+\mathbf{z}_{n}^{[1]}-\mathbf{z}_{n}^{[0]}+\mathbf{z}_{n}^{[0]}
$$

or

$$
\begin{equation*}
\mathbf{z}_{n}^{[m+1]}=\mathbf{z}_{n}^{[0]}+\sum_{j=0}^{m}\left(\mathbf{z}_{n}^{[j+1]}-\mathbf{z}_{n}^{[j]}\right) \tag{2.69}
\end{equation*}
$$

Substituting the difference $\mathbf{z}_{n}^{[j+1]}-\mathbf{z}_{n}^{[J]}$ obtained from equation (2.65) into equation (2.69) produces

$$
\begin{equation*}
\mathbf{z}_{n}^{[m+1]}=\mathbf{z}_{n}^{[0]}+\sum_{j=0}^{m} \mathbf{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[j]}\right) \underline{\ell}=\mathbf{z}_{n}^{[0]}+\underline{\mathrm{e}}_{n}^{[m+1]} \underline{\ell} \tag{2.70}
\end{equation*}
$$

where $\underline{\mathrm{e}}_{n}^{[m+1]}$ is defined as

$$
\begin{equation*}
\underline{\mathrm{e}}_{n}^{[m+1]}=\sum_{j=0}^{m} \mathbf{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[j]}\right) \tag{2.71}
\end{equation*}
$$

It is clear from this equation that

$$
\begin{equation*}
\underline{\mathrm{e}}_{n}^{[m+1]}=\underline{\mathrm{e}}_{n}^{[m]}+\mathrm{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right) \tag{2.72}
\end{equation*}
$$

Equation (2.70) can be used to rewrite $\underline{g}\left(\underline{Y}_{n}^{[m]}\right)$, equation (2.16), as follows:

$$
\begin{equation*}
\underline{\mathrm{g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)=h_{n} \mathrm{f}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)-h_{n} \underline{\mathrm{Y}}_{n}^{[0]}-\underline{\mathrm{e}}_{n}^{[m]} \tag{2.73}
\end{equation*}
$$

because $\ell_{1}=1$.
Finally, because only the first two columns of $z_{n}$ enter into the solution of equation (2.5), the successive corrections can be accumulated and applied to the remaining columns of $z_{n}$ after convergence. Clearly, not updating all columns of the Nordsieck history matrix after each iteration results in savings of computational effort, especially when a high-order method is used and/or the number of ODE's is large. For additional savings of computer time the history matrix is updated only if both (1) the iteration converges and (2) the converged solution satisfies accuracy requirements.

The predictor-corrector formulation utilized in LSODE can be summarized as follows:

Predictor:

$$
\left.\begin{array}{c}
\mathbf{z}_{n}^{[0]}=\mathbf{z}_{n-1} \mathbf{A}  \tag{2.74}\\
\underline{e}_{n}^{[0]}=0
\end{array}\right\}
$$

Corrector:

$$
\left.\begin{array}{c}
\mathrm{g}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)=h_{n} \underline{\mathrm{f}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)-h_{n} \dot{\mathrm{Y}}_{n}^{[0]}-\underline{\mathrm{e}}_{n}^{[m]} \\
\underline{\mathrm{e}}_{n}^{[m+1]}=\underline{\mathrm{e}}_{n}^{[m]}+\mathbf{P}^{-1} \underline{\mathrm{~g}}\left(\underline{\mathrm{Y}}_{n}^{[m]}\right)  \tag{2.75}\\
\underline{\mathrm{Y}}_{n}^{[m+1]}=\underline{\mathrm{Y}}_{n}^{[0]}+\ell_{0} \underline{\mathrm{e}}_{n}^{[m+1]}
\end{array}\right\} m=0,1, \ldots, M-1 .
$$

$$
\left.\begin{array}{c}
\underline{\mathrm{e}}_{n}=\underline{\mathrm{e}}_{n}^{[M]}  \tag{2.76}\\
\mathrm{z}_{n}=\mathrm{z}_{n}^{[0]}+\underline{\mathrm{e}}_{n} \ell
\end{array}\right\}
$$

### 2.5 Local Truncation Error Estimate and Control ${ }^{1}$

The local truncation error is defined to be the amount by which the exact solution $\underline{y}(\xi)$ to the ODE system fails to satisfy the difference equation of the numerical method (refs. $4,10,12$, and 26). That is, for the linear multistep methods, equation (2.1), the local truncation error vector $d_{n}$ at $\xi_{n}$ is the residual in the difference formula when the approximations $\left\{\underline{Y}_{j}\right\}$ and $\left\{\underline{\mathrm{f}}_{j}\right\}$ are replaced by the exact solution and its derivative. ${ }^{2}$ In LSODE, however, the basic multistep formula is normalized by dividing it by

$$
\sum_{j=0}^{K_{2}} \beta_{j}
$$

[^0]for reasons given by Henrici (ref. 29) and Gear (ref. 10); however, see Lambert (ref. 4). For example, the BDF method of order $q$, equation (2.4), can be expressed in this form as
\[

$$
\begin{equation*}
0=\sum_{j=0}^{q}\left(\frac{\alpha_{j}}{\beta_{0}}\right) \underline{\mathbf{Y}}_{n-j}+h_{n} \underline{\mathrm{f}}_{n}, \tag{2.79}
\end{equation*}
$$

\]

where $\alpha_{0}=-1$. The local truncation error for this method is then given by

$$
\begin{equation*}
\underline{\mathrm{d}}_{n}=\sum_{j=0}^{q}\left(\frac{\alpha_{j}}{\beta_{0}}\right) \underline{\mathrm{y}}\left(\xi_{n-j}\right)+h_{n} \underline{\dot{\mathrm{y}}}\left(\xi_{n}\right) \tag{2.80}
\end{equation*}
$$

where $\mathrm{d}_{n}$ consists of $N$ components

$$
\begin{equation*}
\underline{\mathrm{d}}_{n}=\left(d_{1, n}, \ldots, d_{N, n}\right)^{T} \tag{2.81}
\end{equation*}
$$

If we assume that each $y_{i}(i=1, \ldots, N)$ possesses derivatives of arbitrarily high order, each $y_{i}\left(\xi_{n-j}\right)(i=1, \ldots, N ; j=1, \ldots, q)$ in equation (2.80) can be expanded in a Taylor series about $\xi_{n}$. Upon collecting terms the resulting expression for $\mathrm{d}_{n}$ can be stated compactly as

$$
\begin{equation*}
\underline{\mathrm{d}}_{n}=\sum_{k=0}^{\infty} C_{k} h_{n}^{k} \underline{\mathrm{y}}^{(k)}\left(\xi_{n}\right) \tag{2.82}
\end{equation*}
$$

where the $\left\{C_{k}\right\}$ are constants (e.g., ref. 10). A method is said to be of order $q$ if $C_{0}=C_{1}=\ldots=C_{q}=0$, and $C_{q+1} \neq 0$. The local truncation error is then given by

$$
\begin{equation*}
\underline{\mathrm{d}}_{n}=C_{q+1} h_{n}^{q+1} \underline{\mathrm{y}}^{(q+1)}\left(\xi_{n}\right)+\mathrm{O}\left(h_{n}^{q+2}\right), \tag{2.83}
\end{equation*}
$$

where the terms $C_{q+1}$ and $C_{q+1} h_{n}^{q+1} \underline{y}^{(q+1)}\left(\xi_{n}\right)$ are, respectively, called the error constant and the principal local truncation error (ref. 4). In particular, for the BDF method of order $q$ in the normalized form given by equation (2.79) (refs. 22 and 29)

$$
\begin{equation*}
C_{q+1}=\frac{1}{q+1} \tag{2.84a}
\end{equation*}
$$

For the implicit Adams method of order $q$ in normalized form (ref. 22)

$$
\begin{equation*}
C_{q+1}=\left|l_{0}(q+1)-l_{0}(q)\right| \tag{2.84b}
\end{equation*}
$$

where $\ell_{0}(q)$ and $\ell_{0}(q+1)$ are, respectively, the zeroth component of the coefficient vectors for the AM method in normalized form of orders $q$ and $(q+1)$.

The $(q+1)$ th derivative at $\xi_{n}, \underline{\mathrm{y}}^{(q+1)}\left(\xi_{n}\right)$, is estimated as follows: As discussed in section 2.4, at each step the solution method updates the Nordsieck history matrix $z_{n}$ :

$$
\begin{equation*}
\mathrm{z}_{n}=\left(\underline{\mathrm{Y}}_{n}, h_{n} \dot{\underline{Y}}_{n}, \frac{h_{n}^{2}}{2!} \ddot{\mathrm{Y}}_{n}, \ldots, \frac{h_{n}^{q}}{q!} \underline{\mathrm{Y}}_{n}^{(q)}\right) . \tag{2.85}
\end{equation*}
$$

For either method of order $q$ the last column of $\mathbf{z}_{n}, \underline{Z}_{n}(q)$, contains the vector $h_{n}^{q} \underline{Y}_{n}^{(q)} / q$ !, which is the approximation to $h_{n}^{q} \underline{\mathrm{y}}^{(q)}\left(\xi_{n}\right) / q$ !. Now the prediction step being a Taylor series method of order $q$ does not alter the last column of $\mathbf{z}_{n-1}$, namely the vector $h_{n}^{q} Y_{n-1}^{(q)} / q$ !. Hence the last column of $\mathbf{z}_{n}^{[0]}, \underline{z}_{n}^{[0]}(q)$, contains the vector $h_{n}^{q} Y_{n-1}^{(q)} / q$ !. The difference, $\underline{z}_{n}^{(q)}-\underline{Z}_{n}^{[0]}(q)$, is given by

$$
\begin{equation*}
\underline{\mathrm{Z}}_{n}(q)-\underline{\mathrm{z}}_{n}^{[0]}(q)=\frac{h_{n}^{q}}{q!} \underline{\mathrm{Y}}_{n}^{(q)}-\frac{h_{n}^{q}}{q!} \underline{\mathrm{Y}}_{n-1}^{(q)}=\frac{h_{n}^{q+1}}{q!} \underline{\mathrm{Y}}_{n}^{(q+1)}+\mathrm{O}\left(h_{n}^{q+2}\right) \tag{2.86}
\end{equation*}
$$

by using the mean value theorem for derivatives. However, equation (2.76) gives the following expression for $\underline{\underline{Z}}_{n}(q)-\underline{\underline{Z}}_{n}^{[0]}(q)$ :

$$
\begin{equation*}
\underline{z}_{n}(q)-\underline{z}_{n}^{[0]}(q)=\ell_{q} \underline{\mathrm{e}}_{n} . \tag{2.87}
\end{equation*}
$$

Equating equations (2.86) and (2.87) gives the following approximation for $h_{n}^{q+1} Y_{n}^{(q+1)}$ if higher-order terms are neglected:

$$
\begin{equation*}
h_{n}^{q+1} \underline{\mathrm{Y}}_{n}^{(q+1)} \cong q!\ell_{q} \mathrm{e}_{n} . \tag{2.88}
\end{equation*}
$$

Substituting this equation into equation (2.83) and neglecting higher-order terms give the following estimate for $\underline{\mathrm{d}}_{n}$ :

$$
\begin{equation*}
\underline{\mathrm{d}}_{n}=C_{q+1} q!\mathrm{l}_{q} \mathrm{e}_{n} . \tag{2.89}
\end{equation*}
$$

In order to provide for user control of the local truncation error, it is normalized by the error weight vector $\mathrm{EWT}_{n}$, with element $\mathrm{EWT}_{l, n}$ defined by

$$
\begin{equation*}
\mathrm{EWT}_{i, n}=\mathrm{RTOL}_{i}\left|\mathrm{Y}_{i, n-1}\right|+\mathrm{ATOL}_{i} \tag{2.90}
\end{equation*}
$$

where the user-supplied local relative ( $\mathrm{RTOL}_{i}$ ) and absolute ( $\mathrm{ATOL}_{i}$ ) error tolerances for the $i$ th solution component are discussed in chapter 4. The solution $Y_{n}$ is accepted as sufficiently accurate if the following inequality is satisfied:

$$
\begin{equation*}
\left\|\mathrm{d}_{n}\right\| \equiv \sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{d_{i, n}}{\mathrm{EWT}_{i, n}}\right)^{2}} \stackrel{?}{\leq} 1, \tag{2.91}
\end{equation*}
$$

where $\|\cdot\|$ denotes the weighted root-mean-square (rms) norm, which is used for reasons discussed by Hindmarsh (ref. 15). Equation (2.91) can be rewritten as

$$
\begin{equation*}
\left\|\underline{e}_{n}\right\| \equiv \sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{e_{i, n}}{\mathrm{EWT}_{i, n}}\right)^{2}} \stackrel{?}{\leq} \frac{1}{C_{q+1} q!\ell_{q}}, \tag{2.92}
\end{equation*}
$$

by using equation (2.89). If we define the test coefficient $\tau(q, q)$ as

$$
\begin{equation*}
\tau(q, q)=\frac{1}{C_{q+1} q!\ell_{q}} \tag{2.93}
\end{equation*}
$$

the accuracy test, equation (2.92), becomes

$$
\begin{equation*}
\left\|\underline{e}_{n}\right\|^{?} \leq \tau(q, q) \tag{2.94}
\end{equation*}
$$

If we further define the quantity $D_{q}$ by

$$
\begin{equation*}
D_{q}=\frac{\left\|\mathrm{e}_{n}\right\|}{\tau(q, q)}, \tag{2.95}
\end{equation*}
$$

the accuracy test reduces to

$$
\begin{equation*}
D_{q} \stackrel{?}{\leq} 1 \tag{2.96}
\end{equation*}
$$

The reason for using two variables in the definition for $\tau$ will become apparent when we discuss step size and method order selection in section 2.7.

## 2. Description and Implementation of Methods

### 2.6 Corrector Convergence Test and Control

The test for corrector convergence is independent of both the integration method and the iteration technique and is determined by the magnitude of the successive differences $h_{n} \dot{\underline{Y}}_{n}^{[m]}-h_{n} \dot{\mathrm{Y}}_{n}^{[m-1]}$. To provide for user control of the convergence process, the difference $h_{n} \dot{\dot{Y}}_{n}^{[m]}-h_{n} \dot{\underline{Y}}_{n}^{[m-1]}$ is normalized by the error weight vector $\underline{\mathrm{EWT}}_{n}$, equation (2.90). Now, equation (2.33) provides the following expression for $h_{n} \dot{\mathrm{Y}}_{n}^{[m]}-h_{n} \dot{\mathrm{Y}}_{n}^{[m-1]}$ :

$$
\begin{equation*}
h_{n} \underline{\dot{Y}}_{n}^{[m]}-h_{n} \dot{\dot{Y}}_{n}^{[m-1]}=\underline{\delta}_{n}^{[m]}, \tag{2.97}
\end{equation*}
$$

where we have replaced $\mathbf{P}^{-1} \underline{g}\left(\underline{Y}_{n}^{[m-1]}\right)$ by $\underline{\delta}_{n}^{[m]}$. Now, because

$$
\underline{\mathrm{e}}_{n}^{[m+1]}=\sum_{j=0}^{m} \underline{\delta}_{n}^{[j]}
$$

(see eq. (2.71)) and the test on $\left|\mathrm{e}_{n}\right|$ is $\left|\mathrm{e}_{n}\right| \stackrel{?}{\leq} \tau(q, q)$, equation (2.94), the following test for convergence

$$
\begin{equation*}
\varepsilon_{m} \equiv \sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{\delta_{i, n}^{[m]}}{\mathrm{EWT}_{i, n}}\right)^{2}} \stackrel{?}{\leq} \frac{\tau(q, q)}{2(q+2)} \tag{2.98}
\end{equation*}
$$

is consistent with the local truncation error test. The empirical factor $2(q+2)$ in equation (2.98) guarantees that the implicit equation (2.5) is solved to greater accuracy than that required of the numerical solution (refs. 22 and 25).

To increase computational efficiency, especially when the iteration is clearly not converging, LSODE uses the following convergence test instead of equation (2.98):

$$
\begin{equation*}
\varepsilon_{m}^{\prime} \stackrel{?}{\leq} \frac{\tau(q, q)}{2(q+2)} \tag{2.99}
\end{equation*}
$$

The quantity $\varepsilon_{m}^{\prime}$ is related to $\varepsilon_{m}$ by

$$
\begin{equation*}
\varepsilon_{m}^{\prime}=\varepsilon_{m} \min \left(1,1.5 c_{m}^{\prime}\right) \tag{2.100}
\end{equation*}
$$

where

$$
\begin{equation*}
c_{m}^{\prime}=\max \left(0.2 c_{m-1}, c_{m}\right) \tag{2.101}
\end{equation*}
$$

and

$$
\begin{equation*}
c_{m}=\varepsilon_{m} / \varepsilon_{m-1} \tag{2.102}
\end{equation*}
$$

is the estimated convergence rate (refs. 22 and 25). Clearly at least two iterations are required before $c_{m}$ can be computed. For the first iteration $c_{m}^{\prime}$ is set equal to the last value of $c_{m}$ from the previous step. For the first iteration of the very first step and, in the case of NR or JN iteration, after every update of the Jacobian matrix, $c_{m}^{\prime}$ is set equal to 0.7 . Equation (2.100) assumes that the iteration converges linearly, that is, $\lim \left(\varepsilon_{m+1} / \varepsilon_{m}\right)=$ finite constant $c$, and essentially anticipates the magnitude of $\varepsilon_{m}^{m \rightarrow \infty}$ one iteration in advance (ref. 15). Equation (2.101) shows that the convergence rate of the latest iteration is given much more weight than that of the previous iteration. The rationale for this decision is discussed by Shampine (ref. 25), who examined various practical aspects of implementing implicit methods.

### 2.7 Step Size and Method Order Selection and Change

Periodically the code attempts to change the step size and/or the method order to minimize computational work while maintaining prescribed accuracy. To minimize complications associated with method order and step size selection, the new order $q^{\prime}$ is restricted to the values $q-1, q$, and $q+1$, where $q$ is the current order. For each $q^{\prime}$ the step size $h^{\prime}(q)$ that will satisfy exactly the local error bound is obtained by assuming that the highest derivative remains constant. The method order that produces the largest $h^{\prime}$ is used on the next step, along with the corresponding $h^{\prime}$, provided that the $h^{\prime}$ satisfies certain restrictions described in chapter 3.

For the case $q^{\prime}=q, h^{\prime}(q)$ is computed by setting $D_{q}\left(h^{\prime}\right)\left(=\right.$ value of $D_{q}$ for step size $h$ ) $=1$ (see eq. (2.96)), so that the local accuracy requirement is satisfied exactly. Then because $\mathrm{d}_{n}$ varies as $h_{n}^{q+1}$ (see eq. (2.83)), we get

$$
\frac{D_{q}}{1}=\left(\frac{h_{n}}{h^{\prime}(q)}\right)^{q+1}
$$

or

$$
\begin{equation*}
r_{\text {same }} \equiv \frac{h^{\prime}(q)}{h_{n}}=\left(\frac{1}{D_{q}}\right)^{\frac{1}{q+1}}, \tag{2.103}
\end{equation*}
$$

## 2. Description and Implementation of Methods

where $r$ is the ratio of the step size to be attempted on the next step to its current value. The subscript "same" indicates that the same order used on the current step is to be attempted on the next step.

For the case $q^{\prime}=q-1, \underline{\mathrm{~d}}_{n}(q-1)$ is of order $q$, where the variable $q-1$ indicates the method order for which the local truncation error is to be estimated, and

$$
\begin{equation*}
\underline{\mathrm{d}}_{n}(q-1)=C_{q} h_{n}^{q} \underline{\mathrm{y}}^{(q)}\left(\xi_{n}\right) \tag{2.104}
\end{equation*}
$$

where $C_{q}=\left|\ell_{0}(q)-\ell_{0}(q-1)\right|$ for the AM method and $1 / q$ for the BDF method (refs. 22 and 29). Now, the last column of $\mathbf{z}_{n}, \underline{z}_{n}(q)$, contains the vector $h_{n}^{q} \underline{Y}_{n}^{(q)} / q$ ! (see eq. (2.85)), and so $\mathrm{d}_{n}(q-1)$ is easily calculated. On using the rms norm, equation (2.91), the error test for $q^{\prime}=q-1$ becomes

$$
\begin{equation*}
\sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{C_{q} h_{n}^{q} Y_{i, n}^{(q)}}{\mathrm{EWT}_{i, n}}\right)^{2}} \stackrel{?}{\leq} 1 \tag{2.105}
\end{equation*}
$$

If we define the test coefficient $\tau(q, q-1)$ as $1 / C_{q} q$ !, equation (2.105) can be written as

$$
\begin{equation*}
D_{q-1} \equiv \frac{\sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{\frac{h_{n}^{q}}{q!} Y_{i, n}^{(q)}}{\mathrm{EWT}_{i, n}}\right)^{2}}}{\tau(q, q-1)}=\frac{\sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{z_{i, n}(q)}{\mathrm{EWT}_{l, n}}\right)^{2}}}{\tau(q, q-1)} \leq 1, \tag{2.106}
\end{equation*}
$$

where $z_{i, n}(q)$ is the $i$ th element of $\underline{z}_{n}(q)$. The first variable in the definition for $\tau$ gives the method order used on the current step. The second variable indicates the method order for which the local truncation error is to be estimated.

The step size $h^{\prime}(q-1)$ to be attempted on the next step, if the order is reduced to $q-1$, is obtained by using exactly the same procedure that was utilized for the case $q^{\prime}=q$, that is, by setting $D_{q-1}\left(h^{\prime}\right)=1$. Because $\mathrm{d}_{n}(q-1)$ varies as $h_{n}^{q}$, the resulting step size ratio $r_{\text {down }}$ is given by

$$
\begin{equation*}
r_{\mathrm{down}} \equiv \frac{h^{\prime}(q-1)}{h_{n}}=\left(\frac{1}{D_{q-1}}\right)^{\frac{1}{q}} \tag{2.107}
\end{equation*}
$$

The subscript "down" indicates that the order is to be reduced by 1.

For the case $q^{\prime}=q+1$ the local truncation error $\underline{d}_{n}(q+1)$ is of order $q+2$ and is given by

$$
\begin{equation*}
\underline{\mathrm{d}}_{n}(q+1)=C_{q+2} h_{n}^{q+2} \underline{\mathrm{y}}^{(q+2)}\left(\xi_{n}\right) \tag{2.108}
\end{equation*}
$$

where $C_{q+2}=\left|\ell_{0}(q+2)-\ell_{0}(q+1)\right|$ for the AM method and $1 /(q+2)$ for the BDF method (refs. 22 and 29). This case is more difficult than the previous two cases because equation (2.108) involves the derivative of order $q+2$. The derivative $\underline{y}^{(q+2)}\left(\xi_{n}\right)$ is estimated as follows. Equation (2.88) shows that the vector $\ell_{q} \underline{e}_{n}$ is approximately proportional to $h_{n}^{q+1} Y_{n}^{(q+1)} / q!$. We difference the quantity $\ell_{q} \mathbf{e}$ over the last two steps and use the mean value theorem for derivatives to get

$$
\begin{align*}
\ell_{q} \nabla \underline{\mathrm{e}}_{n} \equiv \ell_{q} \underline{\mathrm{e}}_{n}-\ell_{q} \underline{\mathrm{e}}_{n-1} & \cong \frac{h_{n}^{q+1}}{q!} \underline{\mathrm{Y}}_{n}^{(q+1)}-\frac{h_{n}^{q+1}}{q!} \underline{\mathrm{Y}}_{n-1}^{(q+1)} \\
& \cong \frac{h_{n}^{q+2}}{q!} \underline{\mathrm{Y}}_{n}^{(q+2)}+\mathrm{O}\left(h_{n}^{q+3}\right) \tag{2.109}
\end{align*}
$$

Hence the error test for $q^{\prime}=q+1$ becomes

$$
\begin{equation*}
\sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{C_{q+2} q!\ell_{q} \nabla e_{i, n}}{\mathrm{EWT}_{i, n}}\right)^{2}} \stackrel{?}{\leq} 1 \tag{2.110}
\end{equation*}
$$

where we have again used the rms norm and $\nabla e_{i, n}$ is the $i$ th component of $\nabla \mathrm{e}_{n}$. If we define the test coefficient $\tau(q, q+1)$ as $1 /\left(C_{q+2} q!\ell_{q}\right)$, the error test, equation (2.110), can be rewritten as

$$
\begin{equation*}
D_{q+1} \equiv \frac{\sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{\nabla e_{i, n}}{\mathrm{EWT}_{i, n}}\right)^{2}}}{\tau(q, q+1)} \leq 1 . \tag{2.111}
\end{equation*}
$$

To solve for $h^{\prime}(q+1)$, we use the same procedure as for $h^{\prime}(q)$ and $h^{\prime}(q-1)$. The resulting ratio $r_{\text {up }}$ is given by

$$
\begin{equation*}
r_{\mathrm{up}} \equiv \frac{h^{\prime}(q+1)}{h_{n}}=\left(\frac{1}{D_{q+1}}\right)^{\frac{1}{q+2}} \tag{2.112}
\end{equation*}
$$

The subscript "up" indicates that the order is to be increased by 1.
After a suitable value for the step size ratio $r$ has been computed, the step size $h^{\prime}$ to be attempted next is calculated:

$$
\begin{equation*}
h^{\prime}=r h_{n} \tag{2.113}
\end{equation*}
$$

If the step size and/or the method order is changed, the Nordsieck history matrix has to be modified. For the case $q^{\prime}=q$ and $h^{\prime} \neq h_{n}$, the $q+1$ columns of $\mathbf{z}_{n}$ are scaled, as described in section 2.4 (see eqs. (2.62) to (2.64)). For the case $q^{\prime}=q-1$ and $h^{\prime} \neq h_{n}$, the same scaling is performed on the first $q$ columns; the last column of the old $\mathbf{z}_{n}$ is ignored because it is not needed on subsequent steps.

If $q^{\prime}=q+1, \mathbf{z}_{n}$ must be augmented by a column containing the vector $\left(h^{\prime}\right)^{q+1} \underline{Y}_{n}^{(q+1)} /(q+1)!$. The column addition is done in two stages. First, by using equation (2.88) we derive the following expression for $h_{n}^{q+1} \underline{Y}_{n}^{(q+1)} /(q+1)$ !:

$$
\begin{equation*}
\frac{h_{n}^{q+1}}{(q+1)!} \underline{\mathrm{Y}}_{n}^{(q+1)} \cong \frac{q!\ell_{q} \underline{\mathrm{e}}_{n}}{(q+1)!}=\frac{\ell_{q} \underline{\mathrm{e}}_{n}}{q+1} \tag{2.114}
\end{equation*}
$$

and the new column, $\underline{z}_{n}(q+1)$, is given by

$$
\begin{equation*}
\underline{\mathrm{z}}_{n}(q+1)=\frac{\ell_{q} \mathrm{e}_{n}}{q+1} \tag{2.115}
\end{equation*}
$$

Second, in order to account for any change in the step size, all $q+2$ columns of $\mathbf{z}_{n}$ are rescaled as before.

Another factor that must be considered if the step size and/or method order is changed is that the iteration matrix $P$, equation (2.25), may be altered even if the Jacobian matrix is current. To minimize convergence failures caused by an inaccurate $\mathbf{P}$, it must be updated if the coefficient $h \beta_{0}$ has changed significantly since the last evaluation of $\mathbf{P}$.

### 2.8 Interpolation at Output Stations

It frequently happens that the user requires the solution at values $\xi_{o u t, 1}$,

therefore important that in implementing the solution method provision be made for the efficient computation of the solution at the required output stations. Moreover, the procedure used for these computations should not adversely affect the efficiency of the integration beyond the output station. Such a situation arises, for example, if the method has to adjust the step size to "hit" the output station exactly. Because the Nordsieck history array is used to store past history information, the solution can be generated at the output stations quite easily, as described next.

For each $\xi_{\text {out }}$ the integration is continued until the first mesh point $n$ for which $\xi_{n} \geq \xi_{\text {out }}$, and then the solution at $\xi_{\text {out }}$ is obtained by interpolation. Now the solution and its scaled derivatives up to order $q_{n+1}^{\prime}$ are available at $\xi_{n}$. Here $q_{n+1}^{\prime}$ is the order to be attempted on the next step, that is, $\left[\xi_{n}, \xi_{n+1}\right]$. Hence the solution at $\xi_{\text {out }} \underline{Y}\left(\xi_{\text {out }}\right)$, is computed by using a ( $q_{n+1}^{\prime}$ )th-order Taylor series expansion about $\xi_{n}$ and is given by

$$
\begin{align*}
\underline{Y}\left(\xi_{\text {out }}\right) & =\underline{Y}_{n}+\left(\xi_{\text {out }}-\xi_{n}\right) \dot{Y}_{n}+\frac{\left(\xi_{\text {out }}-\xi_{n}\right)^{2}}{2!} \underline{\ddot{Y}}_{n} \\
& +\ldots+\frac{\left(\xi_{\text {out }}-\xi_{n}\right)^{q_{n+1}^{\prime}}}{\left(q_{n+1}^{\prime}\right)!} \underline{Y}_{n}^{\left(q_{n+1}^{\prime}\right)}=\sum_{k=0}^{q_{n+1}^{\prime}\left(\xi_{\text {out }}-\xi_{n}\right)^{k}} \frac{k!}{Y_{n}^{(k)}} \tag{2.116}
\end{align*}
$$

If we define the quantity $r$ by

$$
\begin{equation*}
r=\frac{\xi_{\text {out }}-\xi_{n}}{h_{n+1}^{\prime}} \tag{2.117}
\end{equation*}
$$

where $h_{n+1}^{\prime}$ is the step size to be attempted on the next step, equation (2.116) can be rewritten as

$$
\begin{equation*}
\underline{Y}\left(\xi_{\text {out }}\right)=\sum_{k=0}^{q_{n+1}^{\prime}} r^{k} \frac{\left(h_{n+1}^{\prime}\right)^{k}}{k!} \underline{Y}_{n}^{(k)} \tag{2.118}
\end{equation*}
$$

Now

$$
\frac{\left(h_{n+1}^{\prime}\right)^{k}}{k!} \underline{Y}_{n}^{(k)}
$$

is the $k$ th column $\underline{z}_{n}(k)$ of $\mathbf{z}_{n}$, and so equation (2.118) can be expressed compactly as

$$
\begin{equation*}
\underline{\mathrm{Y}}\left(\xi_{\text {out }}\right)=\sum_{k=0}^{q_{n+1}^{\prime}} r^{k} \underline{\mathrm{Z}}_{n}(k) \tag{2.119}
\end{equation*}
$$

Because the solution is accurate to order $q_{n+1}^{\prime}$ at $\xi_{n}$ and a ( $q_{n+1}^{\prime}$ )th-order Taylor series expansion is used to compute $\underline{Y}\left(\xi_{\text {out }}\right)$, the latter is also accurate to order $q_{n+1}^{\prime}$.

The solution at $\xi_{\text {out }}$, equation (2.119), can be evaluated by additions and multiplications alone by using Horner's rule (ref. 13):

$$
\begin{align*}
& Y_{l}\left(\xi_{\text {out }}\right)=z_{l, n}\left(q_{n+1}^{\prime}\right), \quad i=1, \ldots, N \\
& r=\frac{\xi_{\text {out }}-\xi_{n}}{h_{n+1}^{\prime}} \\
& \text { For } k=1, \ldots, q_{n+1}^{\prime}, \text { do : }  \tag{2.120}\\
& \qquad\left\{\begin{array}{l}
k^{\prime}=q_{n+1}^{\prime}-k \\
Y_{l}\left(\xi_{\text {out }}\right) \leftarrow
\end{array}\right) z_{l, n}\left(k^{\prime}\right)+r Y_{l}\left(\xi_{\text {out }}\right), \quad i=1, \ldots, N
\end{align*}
$$

The Taylor series expansion method can be used to compute the solution derivative of any order (up to $q_{n+1}^{\prime}$ ) at $\xi_{\text {out }}$. For example, the $\mu$ th-derivative at $\xi_{\text {out }}, \underline{Y}^{(\mu)}\left(\xi_{\text {out }}\right)$, is given by

$$
\begin{align*}
\underline{Y}^{(\mu)}\left(\xi_{\text {out }}\right) & =\underline{Y}_{n}^{(\mu)}+\left(\xi_{\text {out }}-\xi_{n}\right) \underline{Y}_{n}^{(\mu+1)}+\ldots+\frac{\left(\xi_{\text {out }}-\xi_{n}\right)^{q_{n+1}^{\prime}-\mu}}{\left(q_{n+1}^{\prime}-\mu\right)!} \underline{Y}_{n}^{\left(q_{n+1}^{\prime}\right)} \\
& =\sum_{k=\mu}^{q_{n+1}^{\prime}} \frac{r^{k-\mu}}{(k-\mu)!}\left(h_{n+1}^{\prime}\right)^{k-\mu} \underline{Y}_{n}^{(k)}, \tag{2.121}
\end{align*}
$$

upon using equation (2.117). Substituting $\underline{Y}_{n}^{(k)}=k!\underline{Z}_{n}(k) /\left(h_{n+1}^{\prime}\right)^{k}$ into equation (2.121) produces

$$
\begin{equation*}
\underline{Y}^{(\mu)}\left(\xi_{\text {out }}\right)=\frac{1}{\left(h_{n+1}^{\prime}\right)^{\mu}} \sum_{k=\mu}^{q_{n+1}^{\prime}} r^{k-\mu} \frac{k!}{(k-\mu)!} \underline{\mathrm{Z}}_{n}(k) \tag{2.122}
\end{equation*}
$$

### 2.9 Starting Procedure

At the outset of the integration, information is available at only the initial point $\xi_{0}$. Hence multistep methods cannot be used on the first step. The difficulty at the initial point is resolved easily by starting the integration with a single-step, first-order method. The Nordsieck history matrix $\mathrm{z}_{0}$ at $\xi_{0}$ is constructed from the initial conditions $\underline{y}_{0}$ and the ODE's as follows:

$$
\begin{equation*}
\underline{z}_{0}(0) \equiv \underline{Y}_{0}=\underline{y}_{0} \tag{2.123}
\end{equation*}
$$

and

$$
\begin{equation*}
\left.\underline{z}_{0}(1) \equiv h_{0} \underline{\underline{Y}}_{0}=h_{0} \underline{\mathrm{f}}^{\left(\underline{y}_{0}\right.}, \xi_{0}\right) \tag{2.124}
\end{equation*}
$$

where $h_{0}$ is the step size to be attempted on the first step.
As the integration proceeds, the numerical solutions generated at the points $\xi_{1}$, $\xi_{2}, \ldots$ provide the necessary values for using multistep methods. Hence, as the numerical solution evolves, the method order and step size can be adjusted to their optimal values by using the procedures described in section 2.7.

## Chapter 3

## Description of Code

### 3.1 Integration and Corrector Iteration Methods

The packaged code LSODE has been designed for the numerical solution of a system of first-order ordinary differential equations (ODE's) given the initial values. It includes a variable-step, variable-order Adams-Moulton (AM) method (suitable for nonstiff problems) of orders 1 to 12 and a variable-step, variableorder backward differentiation formula (BDF) method (suitable for stiff problems) of orders 1 to 5 . However, the code contains an option whereby for either method a smaller maximum method order than the default value can be specified.

Irrespective of the solution method the code starts the integration with a firstorder method and, as the integration proceeds, automatically adjusts the method order (and the step size) for optimal efficiency while satisfying prescribed accuracy requirements. Both integration methods are step-by-step methods. That is, starting with the known initial condition $\underline{y}\left(\xi_{0}\right)$ at $\xi_{0}$, where $\underline{y}$ is the vector of dependent variables, $\xi$ is the independent variable, and $\xi_{0}$ is its initial value, the methods generate numerical approximations $Y_{n}$ to the exact solution $\underline{y}\left(\xi_{n}\right)$ at the discrete points $\xi_{n}(n=1,2, \ldots)$ until the end of the integration interval is reached. At each step $\left[\xi_{n-1}, \xi_{n}\right]$ both methods employ a predictor-corrector scheme, wherein an initial guess for the solution is first obtained and then the guess is improved upon by iteration. That is, starting with an initial guess, denoted by $\mathrm{Y}_{n}^{[0]}$, successively improved estimates $Y_{n}^{[m]}(m=1, \ldots, M)$ are generated until the iteration converges, that is, further iteration produces little or no change in the solution. Here $\mathrm{Y}_{n}^{[m]}$ is the approximation computed on the $m$ th iteration, and $M$ is the number of iterations required for convergence.

A standard explicit predictor formula-a Taylor series expansion method devised by Nordsieck (ref. 33)-is used to generate the initial estimate for the solution. A range of iteration techniques for correcting this estimate is included in LSODE. Both the basic integration method and the corrector iteration procedure are identified by means of the method flag MF. By definition, MF has the two decimal digits METH and MTTER, and
3. Description of Code

TABLE 3.1.-SUMMARY OF INTEGRATION METHODS INCLUDED IN LSODE AND CORRESPONDING VALUES OF METH, THE FIRST DECIMAL DIGIT OF MF

| METH | Integration method |
| :---: | :---: |
| 1 | Variable-step, variable-order, implicit Adams method of orders 1 to 12 <br> 2 |
| Variable-step, variable-order, implicit backward differentiation formula <br> method of orders 1 to 5 |  |

TABLE 3.2.-CORRECTOR ITERATION TECHNIQUES AVALLABLE IN LSODE AND CORRESPONDING VALUES OF MITER, THE SECOND DECIMAL DIGIT OF MF

| MITER | Corrector iteration technique |
| :---: | :--- |
| 0 | Functional iteration |
| 1 | Modified Newton iteration with user-supplied analytical Jacobian |
| 2 | Modified Newton iteration with internally generated numerical Jacobian |
| 3 | Modified Jacobi-Newton teration with internally generated numerical |
| $\mathrm{b}_{4}$ | Jacobian |
| $\mathrm{b}_{5}$ | Modified Newton iteration with user-supplied banded Jacobian |
|  | Modified Newton iteration with intemally generated banded Jacobian |

${ }^{\text {a }}$ Modified Jacobi-Newton iteration with user-supplied analytical Jacobian can be performed by specifying MITER $=4$ and $M L=M U=0^{b}$ (i.e., a banded Jacobian with bandwidth of 1 ).
${ }^{6}$ The user must specify the lower (ML) and upper (MU) half-bandwidths of the Jacobian matrix.

$$
\begin{equation*}
\mathrm{MF}=10 \times \mathrm{METH}+\mathrm{MITER}, \tag{3.1}
\end{equation*}
$$

where the integers METH and MITER indicate, respectively, the integration method and the corrector iteration technique to be used on the problem. Table 3.1 summarizes the integration methods included in LSODE and the appropriate values for METH. The legal values for MITER and their meanings are given in table 3.2. The iteration procedures corresponding to MITER $=1$ to 5 are described as modified Newton iteration techniques because the Jacobian matrix is not updated at every iteration.

### 3.2 Code Structure

The double-precision version of the LSODE package consists of the main core integration routine, LSODE, the 20 subprograms CFODE, DAXPY, DDOT, DGBFA, DGBSL, DGEFA, DGESL, DSCAL, D1MACH, EWSET, IDAMAX, INTDY, PREPJ, SOLSY, SRCOM, STODE, VNORM, XERRWV, XSETF, and

XSETUN, and a BLOCK DATA module for loading some variables. The singleprecision version contains the main routine, LSODE, and the 20 subprograms CFODE, EWSET, INTDY, ISAMAX, PREPJ, R1MACH, SAXPY, SDOT, SGBFA, SGBSL, SGEFA, SGESL, SOLSY, SRCOM, SSCAL, STODE, VNORM, XERRWV, XSETF, and XSETUN. The subprograms DDOT, D1MACH, IDAMAX, ISAMAX, R1MACH, SDOT, and VNORM are function routines-all the others are subroutines. The subroutine XERRWV is machine dependent. In addition to these routines the following intrinsic and external routines are used: DABS, DFLOAT, DMAX1, DMIN1, DSIGN, and DSQRT by the double-precision version; ABS, AMAX1, AMIN1, FLOAT, SIGN, and SQRT by the single-precision version; and MAXO, MINO, MOD, and WRITE by both versions.

Table 3.3 lists the subprograms in the order that they appear in the code and briefly describes each subprogram. Among these, the routines DAXPY, DDOT, DGBFA, DGBSL, DGEFA, DGESL, DSCAL, IDAMAX, ISAMAX, SAXPY, SDOT, SGBFA, SGBSL, SGEFA, SGESL, and SSCAL were taken from the LINPACK collection (ref. 34). The subroutines XERRWV, XSETF, and XSETUN, as used in LSODE, constitute a simplified version of the SLATEC error-handling package (ref. 35).

The structure of the LSODE package is illustrated in figure 3.1, wherein a line connecting two routines indicates that the lower routine is called by the upper one. For subprograms that have different names in the different versions of the code, both names are given, with the double-precision version name listed first. Also, the names in brackets are dummy procedure names, which are used internally and passed in call sequences. The routine $F$ is a user-supplied subroutine that computes the derivatives $d y_{i} / d \xi(i=1, \ldots, N)$, where $y_{i}$ is the $i$ th component of y and $N$ is the number of ODE's. Finally, the user-supplied subroutine JAC computes the analytical Jacobian matrix $\mathbf{J}(=\partial \underline{f} / \partial \underline{y})$, where $\mathrm{f}=d \underline{y} / d \xi$.

The code has been arranged as much as possible in a "modular" fashion, with different subprograms performing different tasks. Hence the number of subprograms is fairly large. However, this feature aids in both understanding and, if necessary, modifying the code. To enhance the user's understanding of the code, it contains many comment statements, which are grouped together in blocks and describe both the task to be performed next and the procedure to be used. In addition, each subprogram includes detailed explanatory notes, which describe the function of the subprogram, the means of communication (i.e., call sequence and/or common blocks), and the input and output variables.

Each subprogram contains data type declarations for all variables in the routine. Such declarations are useful for debugging and provide a list of all variables that occur in a routine. This list is useful in overlay situations. For each data type the variables are usually listed in the following order: variables that are passed in the call sequence, variables appearing in common blocks, and local variables, in either alphabetical order or the order in which they appear in the call sequence and the common blocks.

TABLE 3.3.-DESCRIPTION OF SUBPROGRAMS USED IN LSODE

| Subprogram | Description |  |
| :--- | :--- | :--- |
| Double- <br> precision <br> version | Single- <br> precision <br> version |  |
| LSODE | LSODE | Main core integration routine. Checks legality of input, <br> sets work array pointers, initializes work arrays, com- <br> putes initial integration step size, manages solutions <br> of ODE's, and returns to calling routine with solution <br> and errors. |
| Computes interpolated values of the specified derivative |  |  |
| of the dependent variables. |  |  |



Figure 3.1.-Structure of LSODE package.

### 3.3 Internal Communication

Communication between different subprograms is accomplished by means of both call sequences and the two common blocks EH0001 and LSO001. The reason for using common blocks is to avoid lengthy call sequences, which can significantly deteriorate the efficiency of the program. However, common blocks are not used for variables whose dimensions are not known at compilation time. Instead, to both eliminate user adjustments to the code and minimize total storage requirements, dynamic dimensioning is used for such variables.

The common blocks, if any, used by each subprogram are given in tables 3.4 and 3.5 for the double- and single-precision versions, respectively. These tables also list all routines called and referenced (e.g., an external function) by each subprogram. Also, to facilitate use of LSODE in overlay situations, all routines that call and reference each subprogram are listed. Finally, for each subprogram the two tables give dummy procedure names (which are passed in call sequences and therefore have to be declared external in each calling and called subprogram) in brackets.

The variables included in the two common blocks and their dimensions, if different from unity, are listed in table 3.6. The common blocks contain variables that are (1) local to any routine but whose values must be preserved between calls to that routine and (2) communicated between routines. The structure of the block LS0001 is as follows: All real variables are listed first, then all integer variables. Within each group the variables are arranged in the following order: (1) those local to subroutine LSODE, (2) those local to subroutine STODE, and (3) those used for communication between routines. It must be pointed out that not all variables listed for a given common block are needed by each routine that uses it. For this reason some subprograms may use dummy names, which are not listed in table 3.6.

To further assist in user understanding and modification of the code, we have included in table 3.6 the names of all subprograms that use each common block. For the same reason we provide in tables 3.7 and 3.8 complete descriptions of the variables in EH0001 and LS0001, respectively. Also given for each variable are the default or current value, if any, and the subprogram (or subprograms) where it is set or computed. The length LENWM of the array WM in table 3.8 depends on the iteration technique and is given in table 3.9 for each legal value of MITER.

### 3.4 Special Features

The remainder of this chapter deals with the special features of the code and its built-in options. We also describe the procedure used to advance the solution by one step, the corrective actions taken in case of any difficulty, and step size and method order selection. In addition, we provide detailed flowcharts to explain the computational procedures. We conclude this chapter with a brief discussion of the error messages included in the code.

TABLE 3.4.-ROUTINES WTTH COMMON BLOCKS, SUBPROGRAMS, AND CALLING SUBPROGRAMS IN DOUBLE-PRECISION

VERSION OF LSODE

| Subprogram [Dummy procedure name] | Common blocks used | Subprograms called and referenced | Calling subprograms |
| :---: | :---: | :---: | :---: |
| LSODE | LSOOOI | DIMACH EWSET <br> F INTDY JAC PREPJ SOLSY STODE VNORM XERRWV |  |
| CFODE |  |  | STODE |
| DAXPY |  |  | DGBFA DGBSL DGEFA DGESL |
| DDOT |  |  | DGBSL DGESL |
| DGBFA |  | DAXPY DSCAL IDAMAX | PREPJ |
| DGBSL |  | DAXPY DDOT | SOLSY |
| DGEFA |  | DAXPY DSCAL IDAMAX | PREPJ |
| DGESL |  | DAXPY DDOT | SOLSY |
| DSCAL |  |  | DGBFA DGEFA |
| DIMACH |  |  | LSODE |
| EWSET |  |  | LSODE |
| IDAMAX |  |  | DGBFA DGEFA |
| INTDY | LSOOOI | XERRWV | LSODE |
| PREPJ <br> [PJAC] | LS0001 | DGBFA DGEFA <br> F JAC VNORM | STODE |
| $\begin{aligned} & \text { SOLSY } \\ & \text { [SLVS] } \end{aligned}$ | LS0001 | DGBSL DGESL | STODE |
| SRCOM | EH0001 LS0001 |  |  |
| STODE | LS0001 | CFODE F JAC PREPJ SOLSY VNORM | LSODE |
| VNORM |  |  | LSODE PREPI STODE |
| XERRWV | EH0001 |  | LSODE INTDY |
| XSETF | EH0001 |  |  |
| XSETUN | EH0001 |  |  |
| BLOCK DATA | EH0001 LSOOO1 |  |  |

TABLE 3.5.-ROUTINES WITH COMMON BLOCKS, SUBPROGRAMS, AND CALLING SUBPROGRAMS IN SINGLE-PRECISION

VERSION OF LSODE

| Subprogram [Dummy procedure name] | Common blocks used | Subprograms called and referenced | Calling <br> subprograms |
| :---: | :---: | :---: | :---: |
| LSODE | LS0001 | EWSET F INTDY JAC PREPJ R1MACH SOLSY STODE VNORM XERRWV |  |
| CFODE |  |  | STODE |
| EWSET |  |  | LSODE |
| INTDY | LS0001 | XERRWV | LSODE |
| ISAMAX |  |  | SGBFA SGEFA |
| PREPJ $[\mathrm{PJAC}]$ | LS0001 | F JAC SGBFA SGEFA VNORM | STODE |
| RIMACH |  |  | LSODE |
| SAXPY |  |  | SGBFA SGBSL SGEFA SGESL |
| SDOT |  |  | SGBSL SGESL |
| SGBFA |  | ISAMAX SAXPY SSCAL | PREPJ |
| SGBSL |  | SAXPY SDOT | SOLSY |
| SGEFA |  | ISAMAX SAXPY SSCAL | PREPJ |
| SGESL |  | SAXPY SDOT | SOLSY |
| $\begin{aligned} & \text { SOLSY } \\ & \text { [SLVS] } \end{aligned}$ | LS0001 | SGBSL SGESL | STODE |
| SRCOM | EH0001 LS0001 |  |  |
| SSCAL |  |  | SGBFA SGEFA |
| STODE | LS0001 | CFODE F JAC PREPJ SOLSY VNORM | LSODE |
| VNORM |  |  | LSODE PREPJ STODE |
| XERRWV | EH0001 |  | LSODE INTDY |
| XSETF XSETUN | EHOOO1 <br> EHOOOI |  |  |

TABLE 3.6.-COMMON BLOCKS WITH VARIABLES AND SUBPROGRAMS WHERE USED

| Common block | Variables (dimension) | Subprograms where used |
| :---: | :---: | :---: |
| EH0001 | MESFLG LUNIT | SRCOM XERRWV XSETF XSETUN BLOCK DATA ${ }^{2}$ |
| LS0001 | CONIT CRATE EL(13) ELCO 13,12 ) HOLD RMAX TESCO(3, 12) CCMAX ELO H HMIN HMXI HU RC TN UROUND ILLIN INIT LYH LEWT LACOR LSAVF LWM LIWM MXSTEP MXXHIL NHNIL NTREP NSLAST NYH IALTH IPUP LMAX MEO NQNYH NSLP ICF IERPJ IERSL JCUR JSTART KFLAG L METH MITER MAXORD MAXCOR MSBP MXNCF $N$ NQ NST NFE NJE NQU | LSODE INTDY <br> PREPJ SOLSY <br> SRCOM STODE <br> BLOCK DATA ${ }^{\text {a }}$ |

"Double-precision version only.

TABLE 3.7.-DESCRIPTION OF VARIABLES IN COMMON BLOCK EHOOOI, THEIR CURRENT VALUES, AND SUBPROGRAMS WHERE THEY ARE SET

| Variable | Description | Current <br> value | Subprogram where <br> variable is set |
| :---: | :---: | :---: | :---: |
| MESFLG | Integer flag, which controls <br> printing of error messages from <br> code and has following values <br> and meanings: <br> 0 No error message is printed. <br> 1 All eror messages are printed. <br> Logical unit number for messages <br> from code | 1 | BLOCK DATA in <br> double-precision version <br> and XERRWV in single- <br> precision version |

TABLE 3.8.-DESCRIPTION OF VARIABLES IN COMMON BLOCK LSOOO1, THEIR CURRENT VALUES, IF ANY, AND SUBPROGRAMS WHERE THEY ARE SET OR COMPUTED²

| Variable | Description | Current value, if any | Subprograms where variable is set or computed |
| :---: | :---: | :---: | :---: |
| CONIT | Empirical factor, $0.5 /(\mathrm{NQ}+1)$ used in convergence test (see eq. (2.99)) | $\cdots$ | STODE |
| CRATE | Estimated convergence rate of iteration | - | STODE |
| EL | Method coefficients in normal form $\left\{l_{1}\right\}$ (see eq. (2.68)), for current method order | -------------------- | STODE |
| ELCO | Method coefficients in normal form for current method of orders 1 to MAXORD | -------------------- | CFODE |
| HOLD | Step size used on last successful step or attempted on last unsuccessful step | ------------------10 | STODE |
| RMAX | Maximum factor by which step size will be increased when step size change is next considered | Normally $10 ; 10^{4}$ for very first step size increase for problem if no difficulty encountered; 2 after a failed convergence or local entor test | STODE |
| TESCO | Test coefficients for current method of orders 1 to MAXORD; used for testing convergence and local accuracy and selecting new step size and method order | ----------------------3 | CFODE |
| CCMAX | Maximum relative change allowed in $\mathrm{H} \times \mathrm{EL} 0$ before Jacobian matrix is updated | 0.3 | LSODE |
| ELO | Method coefficient $\ell_{0}$ (see eq. (2.68)) for current method and current order | ------------------ | STODE |
| H | Step size cither being used on this step or to be attempted on next step | 0.0 | LSODE STODE |
| $\mathrm{HMIN}^{\text {b }}$ | Minimum absolute value of step size to be used on any step | 0.0 0 |  |
| HMEXI ${ }^{\text {b }}$ | Inverse of maximum absolute value of step size to be used on any step | 0.0 | LSODE |
| HU | Step size used on last successful step | - | STODE |

[^1]TABLE 3.8.-Continued.

| Variable | Description | Current value. if any | Subprograms where variable is set or computed |
| :---: | :---: | :---: | :---: |
| RC | Relative change in $\mathrm{H} \times \mathrm{ELO}$ since last update of Jacobian marrix | $\cdots$ | STODE |
| TN | Value of independent variable to which integrator either has successfully advanced solution or will do so after next step | --------------------- | STODE |
| UROUND | Unit roundoff of computer |  | DIMACH in double-precision version and R1MACH in single-precision version |
| ILLIN | Number of consecutive times LSODE has been called with illegal input for current problem | ------3 | Initialized in BLOCK DATA (double-precision version) and LSODE (singleprecision version). Updated in LSODE in both versions. |
| INTT | Integer flag ( $=0$ or 1 ) that denotes if initialization of LSODE has been performed (INIT $=1$ ) or not (INIT $=0$ ) | $\cdots$ | LSODE |
| LYH | Base address for Nordsieck history array YH of length NYHx(MAXORD + 1) | 21 | LSODE |
| LEWT | Base address for error weight vector EWT of length N | LWM + LENWM ${ }^{\text {c }}$ | LSODE |
| LACOR | Base address for array ACOR (of length N ) containing local crtors on last successful step | LEWT + 2N | LSODE |
| LSAVF | Base address for an array SAVF (of length N ), used for temporary storage | LEWT + N | LSODE |
| LWM | Base address for array WM (of length LENWM ${ }^{\circ}$ ), required for linear algebra associated with Jacobian and iteration matrices | $\underset{\text { LYH + }}{\text { NYHx(MAXORD }+1 \text { ) }}$ | LSODE |

[^2]TABLE 3.8.-Continued.

| Variable | Description | Current value, if any | Subprograms where variable is set or computed |
| :---: | :---: | :---: | :---: |
| LIWM | Base address for integer work array IWM | 1 | LSODE |
| MXSTEP ${ }^{\text {b }}$ | Maximum number of steps allowed on any one call to LSODE | 500 | LSODE |
| MXHNL ${ }^{\text {b }}$ | Maximum number of times that warning message that step size is so small that $\mathrm{TN}+$ $\mathrm{H}=\mathrm{TN}$ for next step is printed | 10 | LSODE |
| NHNIL | Number of times that this difficulty with small step size has been encountered so far for problem | - | LSODE |
| NTREP | Number of consecutive times an initialization or "first" call (see table 4.3) has been made to LSODE with same initual and final values for integration interval | - | Initialized in BLOCK DATA (double-precision version) and LSODE (singleprecision version). Updated in LSODE in both versions. |
| NSLAST | Number of steps used for problem prior to current call to LSODE; used to check that the limit of MXSTEP steps is not exceeded | --------------> | LSODE |
| NYH | Maximum number of ODE's to be solved for current problem (This number is equal to the number of ODE's specified on first call to LSODE.) | $\cdots$ | LSODE |
| IALTH | Integer counter, related to step size and method order changes, with following values and meanings: <br> 0 Select optimal step size and method order. <br> 1 If NQU < MAXORD, save vector $e$ (sec eqs. (2.76) and (2.111)) so that an order increase can be considered on the next step. <br> $>1$ Neither of these two operations is to be performed. | - | STODE |

[^3]TABLE 3.8.-Continued.


TABLE 3.8.-Continued.

| Variable | Description | Current value, If any | Subprograms where variable is set or computed |
| :---: | :---: | :---: | :---: |
| JCUR | Integer flag, related to state of Jacobian matrix, with following values and meanings: 0 Jacobian matrix is not current and may need to be updated later. <br> 1 Matrix is current. | ------------------- | PREPJ <br> STODE |
| JSTART | Integer flag, used to communicate state of calculation to STODE, with following values and meanings: <br> 0 This is the first step for the problem. <br> I Continue normal calculation of problem. (This is the value returned by STODE to facilitate continuation.) <br> -1 Take the next step with new values for H, MAXORD, N, METH (see table 3.1), MTIER (see table 3.2), and/or matrix parameters. | ------------------ | LSODE <br> STODE |
| KFLAG | A completion code from STODE with following values and meanings: <br> 0 Step was successful. <br> -1 Requested local accuracy in solution could not be achieved. <br> -2 Repeated convergence test failures occurred. | -- | STODE |
| L | Number of columns of Nordsieck array | $N Q+1$ | STODE |
| METH | Integration method to be used on next step | ------------------- | LSODE |
| MITER | Iteration technique to be used on next step | $\qquad$ | LSODE |
| MAXORD ${ }^{\text {b }}$ | Maximum method order to be used for problem | 12 for Adams-Moulton method and 5 for backward differentiation formula method | LSODE |
| MAXCOR | Maximum number of corrector iterations to be attempted on any one step | 3 | LSODE |
| MSBP | Maximum number of steps for which same Jacobian matrix is used | 20 | LSODE |

[^4]TABLE 3.8.-Concluded.

| Variable | Description | Current value, if any | Subprograms where variable is set or computed |
| :---: | :---: | :---: | :---: |
| MXNCF | Maximum number of corrector convergence failures allowed on any one step | 10 | LSODE |
| N | Number of ODE's to be solved on next step | --- | - |
| NQ | Method order either being tried on this step or to be attempted on next step | ----- | STODE |
| NST | Total number of incegration steps used so far for problem | -- | LSODE STODE |
| NFE | Total number of derivative evaluations required so far for problem | -- | LSODE STODE |
| NJE | Total number of Jacobin matrix evaluations (and iteration matrix LUdecompositions or inversions) required so far for problem | - | LSODE PREPJ |
| NQU | Method order used on last successful step. | --..-- | STODE |

TABLE 3.9.-LENGTH LENWM
OF ARRAY WM IN TABLE 3.8
FOR ITERATION TECHNIQUES
INCLUDED IN CODE

| MTTER $^{\mathrm{a}}$ | LENWM $^{\mathrm{b}}$ |
| :--- | :--- |
| 0 | 0 |
| 1,2 | $\mathrm{~N}^{2}+1$ |
| 3 | $\mathrm{~N}+2$ |
| 4,5 | $(2 \mathrm{ML}+\mathrm{MU}+1) \mathrm{N}+2$ |

[^5]
## 3. Description of Code

The main routine, LSODE, controls the integration and serves as an interface between the calling subprogram and the rest of the package. A flowchart of this subroutine is given in figure 3.2. In this figure ITASK and ISTATE are userspecified integers that specify, respectively, the task to be performed and the state of the calculation, that is, if the call to LSODE is the first one for the problem or a continuation; if the latter, ISTATE further indicates if the continuation is a normal one or if the user has changed one or more parameters since the last call to LSODE (see chapter 4 for details). On return from LSODE the value of ISTATE indicates if the integration was performed successfully, and if not, the reason for failure. The integer JSTART is an internally defined variable used for communicating the state of the calculation with the routine STODE. The variables $\mathrm{T}(=\xi), \mathrm{H}$, and $\underline{Y}$ are, respectively, the independent variable, the step size to be attempted on the next step, and the numerical solution vector. TOUT is the $\xi$ value at which the solution is next required. Finally, TCRIT is the $\xi$ value that the integrator must not overshoot. This option is useful if a singularity exists at or beyond TCRIT and is discussed further in chapter 4.

The subroutine STODE advances the numerical solution to the ODE's by a single integration step $\left[\xi_{n-1}, \xi_{n}\right]$. It also computes the method order and step size to be attempted on the next step. The efficiency of the integration procedure is increased by saving the solution history, which is required by the multistep methods used in the code, in the form suggested by Nordsieck (ref. 33). The $N \times(q+1)$ Nordsieck history matrix $\mathbf{z}_{n-1}$ at $\xi_{n-1}$ contains the numerical solution $\underline{Y}_{n-1}$ and the $q$ scaled derivatives $h_{n}^{j} \underline{Y}_{n-1}^{(j)} / j!(j=1, \ldots, q)$, where $h_{n}\left(=\xi_{n}-\xi_{n-1}\right)$ and $q$ are, respectively, the current step size and method order and $\underline{Y}^{(j)}=d^{j} \underline{Y} / d \xi^{j}$.

The flowchart of STODE is presented in figure 3.3. In this figure NCF is the number of corrector convergence failures on the current step, KFLAG is an internally defined integer used for communication with LSODE, NQ $(=q)$ is the method order to be attempted on the current step, and the integer counter IALTH indicates how many more steps are to be taken with the current step size and method order. The (NQ + 1)-dimensional vector $\ell$ contains the method coefficients and depends on both the integration method and the method order; $\ell_{0}$ is the zeroth component of $\underline{\ell}$ (see eq. (2.68)). The matrix $\mathbf{z}_{n}^{[0]}$ is the predicted Nordsieck history matrix at $\xi_{n}$, and the $N \times N$ iteration matrix $\mathbf{P}$ is given by equation (2.25). The variable $R$ is the ratio of the step size to be attempted next to its current value, RMAX is the maximum $R$ allowed when a step size change is next considered, and HMIN and HMAX are user-supplied minimum and maximum absolute values for the step size to be tried on any step. The ratios RHDN, RHSM, and RHUP are factors by which the step size can be increased if the new method order is NQ -1 , NQ (the current value), and NQ +1 , respectively. Finally, NQMAX is the maximum method order that may be attempted on any step, and the vector $\underline{\mathrm{e}}_{n}\left(=h_{n} \dot{\mathrm{Y}}_{n}-h_{n} \dot{\mathrm{Y}}_{n}^{[0]}\right)$ is proportional to the local truncation error vector at $\xi_{n}$ (see eqs. (2.87) and (2.89)).

### 3.4.1 Initial Step Size Calculation

An important feature of LSODE is that it will compute the step size $h_{0}$ to be attempted on the first step if the user does not provide a value for it. The calculation procedure attempts to produce an $h_{0}$ such that the numerical solution $Y_{1}$ generated at the first internal mesh point $\xi_{1}$ will satisfy the local error test. Now with either solution technique the code starts the integration with a firstorder method. Hence the asymptotic local truncation error $d_{i, 1}$ in the $i$ th solution component at $\xi_{1}$ will be equal to $(1 / 2) h_{1}^{2} \ddot{y}_{i}\left(\xi_{1}\right)$ for both the AM and BDF methods of order 1. Here $h_{1}$ is the step size successfully used on the first step, and $\ddot{y}_{i}\left(\xi_{1}\right)$ is the second derivative of the $i$ th component of $y$ at $\xi_{1}$. To pass the local error test, equation (2.91), the weighted local error vector, that is, $\left\{d_{i, 1} / \mathrm{EWT}_{i, 1}\right\}$, must satisfy the inequality

$$
\begin{equation*}
\sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{\frac{1}{2} h_{1}^{2} \ddot{y}_{i}\left(\xi_{1}\right)}{\mathrm{EWT}_{i, 1}}\right)^{2}} \leq 1 \tag{3.2}
\end{equation*}
$$

where $\mathrm{EWT}_{i, 1}$ is the $i$ th component of the error weight vector for the first step (see eq. (2.90)):

$$
\begin{equation*}
\mathrm{EWT}_{i, 1}=\mathrm{RTOL}_{i}\left|Y_{i, 0}\right|+\mathrm{ATOL}_{i} \tag{3.3}
\end{equation*}
$$

In this equation $\mathrm{RTOL}_{i}$ and $\mathrm{ATOL}_{i}$ are, respectively, the user-supplied local relative and absolute error tolerances for the $i$ th solution component, $Y_{i, 0}$ is the $i$ th solution component at $\xi_{0}$, and the vertical bars $|\cdot|$ denote absolute value.

The test given by equation (3.2) cannot be applied at the start of the step [ $\left.\xi_{0}, \xi_{1}\right]$ because $\ddot{\mathrm{y}}\left(\xi_{1}\right)$ is not known. We therefore modify this test by using $\ddot{\mathrm{y}}\left(\xi_{0}\right)$ as follows: We first define a weighted principal error function at order $1, \bar{\phi}$, with element $\phi_{i}$ given by

$$
\begin{equation*}
\phi_{i}=\frac{1}{2} \frac{\ddot{y}_{i}\left(\xi_{0}\right)}{W_{i}}, \tag{3.4}
\end{equation*}
$$

where

$$
\begin{equation*}
W_{i}=\mathrm{EWT}_{i, 1} / \mathrm{TOL} \tag{3.5}
\end{equation*}
$$




Figure 3.2.-Fiowchart of subroutine LSODE.



Figure 3.3.-Flowchart of subroutine STODE.

## 3. Description of Code

and the scalar tolerance quantity TOL, which is to be determined, is such that $W_{i}$ is a suitable weight for $Y_{i}$, the $i$ th component of $\underline{Y}$. The step size and the local error are then together required to satisfy the inequality

$$
\begin{equation*}
h_{0}^{2}\|\phi\| \leq \mathrm{TOL} \tag{3.6}
\end{equation*}
$$

where $\|\|\|$ represents a suitable norm. We have used a different symbol for the initial step size than in equation (3.2) to indicate that this quantity is not known and must be computed. Because a first-order method will be used on this step, for a sufficiently small step size the numerical approximation $\ddot{Y}_{1}$ at $\xi_{1}$ will not be significantly different from $\ddot{\mathrm{y}}\left(\xi_{0}\right)$, and use of the latter quantity is therefore reasonable. The rationale for introducing TOL will become apparent shortly.

The second derivative $\underline{\ddot{y}}\left(\xi_{0}\right)$ is not generally available, and so the following empirical procedure is used to estimate it. We consider the dominant eigenvalue $(=\lambda)$ of the ODE system and model this component with the simple scalar ODE

$$
\begin{equation*}
\dot{y} \equiv \frac{d y}{d \xi}=\lambda y, \tag{3.7}
\end{equation*}
$$

where $|\lambda| \gg 1$. For this problem, $\phi=(1 / 2) \ddot{y} / W=(1 / 2) \lambda^{2} y / W$. Now, if TOL is chosen such that $y / W$ is of order unity, $\phi$ can be approximated by $(\dot{y} / W)^{2}$ $\left[=(\lambda y / W)^{2}\right]$, which is known. For the scalar ODE this condition is obtained by setting TOL $=$ RTOL and ATOL $=0$ (see eqs. (3.3) and (3.5)). The quantity $\dot{y} / W$ may be regarded as the weighted principal error function for a "zeroth order" method. We use this empirical rule to replace each $\phi_{i}$ by $\left(\dot{y}_{i} / W_{i}\right)^{2}$ so that equation (3.6) can be written as

$$
\begin{equation*}
h_{0}^{2}\left\|\left[\left(f_{i, 0} / W_{i}\right)^{2}\right]_{1}^{N}\right\| \leq \mathrm{TOL} \tag{3.8}
\end{equation*}
$$

where $f_{i, 0}\left[=f_{I}\left(\underline{Y}_{0}, \xi_{0}\right)\right]$ is the first derivative of the ith component at $\xi_{0}$. Because the weighted root-mean-square ( rms ) norm is used in the local error test, equation (3.2), for convenience, we use the following criterion for initial step size control:

$$
\begin{equation*}
h_{0}^{2}\left[\frac{1}{N} \sum_{i=1}^{N}\left(\frac{f_{i, 0}}{W_{t}}\right)^{2}\right] \leq \mathrm{TOL} \tag{3.9}
\end{equation*}
$$

Equations (3.5) and (3.9) together show that $h_{0}(\propto 1 / \sqrt{\mathrm{TOL}})$ is a decreasing
function of TOL. To produce a reliable estimate for $h_{0}$, we therefore select a TOL erring on the high side. A suitable value is given by

$$
\begin{equation*}
\mathrm{TOL}=\max _{i}\left(\mathrm{RTOL}_{i}\right) \tag{3.10}
\end{equation*}
$$

This expression cannot be used if all $\mathrm{RTOL}_{i}=0$. In this case an appropriate value for TOL is given by

$$
\begin{equation*}
\mathrm{TOL}=\max _{i}\left(\frac{\mathrm{ATOL}_{i}}{\left|Y_{i, 0}\right|}\right) \quad \text { for } Y_{i, 0} \neq 0 \tag{3.11}
\end{equation*}
$$

In any case the value of TOL is constrained to be within reasonable bounds as follows:

$$
\begin{equation*}
100 u \leq \text { TOL } \leq 10^{-3} \tag{3.12}
\end{equation*}
$$

where $u$ is the unit roundoff of the computer or the machine epsilon (ref. 13). It is the smallest positive number such that $1+u>1$.

Equation (3.9) cannot be used to compute $h_{0}$ if either each $f_{i, 0}$ is equal to zero or the norm is very small. To produce a reasonable $h_{0}$ in such an event, we include the independent variable $\xi$ as the zeroth component $y_{0}$ of $\underline{y}$ and modify equation (3.9) as follows:

$$
\begin{equation*}
h_{0}^{2}\left[\frac{1}{W_{0}^{2}}+\frac{1}{N} \sum_{i=1}^{N}\left(\frac{f_{i, 0}}{W_{i}}\right)^{2}\right] \leq \mathrm{TOL} \tag{3.13}
\end{equation*}
$$

where we have used the fact that $\dot{y}_{0}=1$. To be consistent with the other $W_{i}$, which are of order $Y_{i, 0}$, the weight $W_{0}$ should be of order $\xi_{0}$; however, we use

$$
\begin{equation*}
W_{0}=\max \left(\left|\xi_{0}\right|,\left|\xi_{\text {out }, 1}\right|\right) \tag{3.14}
\end{equation*}
$$

to ensure that it is not equal to zero. In equation (3.14), $\xi_{\text {out }, 1}$ is either the first (or only) value of the independent variable at which the solution is required or, as discussed in chapter 4, a value that gives both the direction of integration (i.e., increasing or decreasing $\xi$ ) and an approximate scale of the problem. If the quantity $\xi_{\text {out, } 1}-\xi_{0}$ is not significantly different from zero, an error exit occurs. Equation (3.13) gives a reasonable value for $h_{0}\left(=W_{0} \sqrt{\text { TOL }}\right)$ if $\mathrm{f}_{0}=0$.

## 3. Description of Code

The calculation procedure used for $h_{0}$ is therefore given by

$$
\begin{equation*}
h_{0}=\sqrt{\frac{\mathrm{TOL}}{W_{0}^{-2}+\frac{1}{N} \sum_{i=1}^{N}\left(\frac{f_{i, 0}}{W_{i}}\right)^{2}}} \tag{3.15}
\end{equation*}
$$

Several restrictions apply to the step size given by equation (3.15). It is not allowed to be greater than the difference $\left|\xi_{o u t, 1}-\xi_{0}\right|$. Hence

$$
\begin{equation*}
h_{0} \leftarrow \min \left(h_{0},\left|\xi_{\text {out }, 1}-\xi_{0}\right|\right) . \tag{3.16}
\end{equation*}
$$

In addition, if the user has supplied a value for $h_{\text {max }}$, the maximum step size to be used on any step, $h_{0}$ is restricted to

$$
\begin{equation*}
h_{0} \leftarrow \min \left(h_{0}, h_{\max }\right) . \tag{3.17}
\end{equation*}
$$

However, no comparison of $h_{0}$ is made with $h_{\text {min }}$, the user-supplied minimum step size to be used on any step, so that $h_{0}$ is allowed to be less than $h_{\text {min }}$. Finally the sign of $h_{0}$ is adjusted to reflect the direction of integration.

### 3.4.2 Switching Methods

Another useful feature of LSODE is that different integration methods and/or different iteration techniques can be used in different subintervals of the problem. This option is useful when the problem changes character and is stiff in some regimes and nonstiff in others as, for example, in combustion chemistry. Indeed, because stiff problems are usually characterized by a nonstiff initial "transient" region, the ability to switch integration methods is a desirable feature of any ODE package. During the course of solving a problem the method flag MF may be changed both whenever and as many times as desired. As described in chapter 4 changing methods is quite straightforward.

### 3.4.3 Excessive Accuracy Specification Test

At each integration step $\left[\xi_{n-1}, \xi_{n}\right]$ LSODE checks that the user has not requested too much accuracy for the precision of the machine. This condition is said to occur if the criterion

$$
\begin{equation*}
d_{i, n} \stackrel{?}{<} u Y_{i, n} \tag{3.18}
\end{equation*}
$$

is true for all $N$ solution components. In equation (3.18), $d_{i, n}$ is the estimated
local truncation error in $Y_{i, n}$, the $i$ th solution component at $\xi_{n}$. Now the numerical solution $Y_{n}$ at $\xi_{n}$ is judged to be sufficiently accurate if the following inequality is satisfied (see chapter 2):

$$
\begin{equation*}
\left\|\mathrm{d}_{n}\right\| \equiv \sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{d_{i, n}}{E W T_{i, n}}\right)^{2}} \stackrel{?}{\leq} 1 . \tag{2.91}
\end{equation*}
$$

The quantity $\mathrm{EWT}_{i, n}$ is the $i$ th component of the error weight vector, equation (2.90), for this step. Equations (3.18) and (2.91) together imply that if the quantity TOLSF (tolerance scale factor) defined as

$$
\begin{equation*}
\mathrm{TOLSF}=u \sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{Y_{i, n}}{\mathrm{EWT}_{i, n}}\right)^{2}} \tag{3.19}
\end{equation*}
$$

is greater than 1 , the test for excessive accuracy requirement is passed. This test is quite inexpensive, but it can be applied only after the solution at $\xi_{n}$ is produced. It is, however, wasteful to generate a solution only to discover that excessive accuracy has been required, either because TOLSF is greater than 1 or because repeated convergence failures or error test failures occur. The computational cost can be significant if any difficulty is encountered because of the corrective actions-described later in this section-performed by the code. Even if the step is successful, the solution is not meaningful because of roundoff errors.

To avoid these difficulties, the calculation procedure for TOLSF uses $Y_{n-1}$, which is known, so that the test can be applied at the start of each step, including the first. Thus the code ascertains inexpensively if excessive accuracy has been requested before attempting to advance the solution by the next integration step. The value of TOLSF may be used to adjust the local error tolerances so that this condition does not recur. For example, scaling up the $\left\{\mathrm{RTOL}_{i}\right\}$ and $\left\{\mathrm{ATOL}_{i}\right\}$ values by a minimum factor of TOLSF should produce satisfactory values for the local error tolerances if the same type of error control is to be performed (see chapter 4 for details).

### 3.4.4 Calculation of Method Coefficients

The integration method coefficients and test constants used to check corrector convergence and local accuracy, as well as to select method order and step size, are computed in subroutine CFODE. The calculation procedure uses the generating polynomials discussed by Hindmarsh (refs. 21 and 22) to increase portability of the code. The coefficients corresponding to all method orders are computed and stored both at the start of the problem and whenever the user changes the integration method. This feature avoids the computational cost associated with recomputing these quantities whenever the method order is changed.

## 3. Description of Code

### 3.4.5 Numerical Jacobians

If Newton-Raphson (NR) or Jacobi-Newton (JN) iteration is selected, the code will generate elements of the Jacobian matrix by finite-difference approximations if the user chooses not to provide an analytical Jacobian. For the iteration procedures corresponding to MITER $=2$ (full Jacobian matrix) and 5 ("banded" Jacobian matrix, i.e., a matrix with many zero entries and all nonzero elements concentrated near the main diagonal), the element $J_{i j}\left(=\partial f_{i} / \partial y_{j}\right)$ at $\xi_{n}$ is estimated by using the approximation

$$
\begin{equation*}
J_{i j} \cong \frac{f_{i}\left(\left\{Y_{k, n}^{[0]}+\delta_{k j} \Delta Y_{j}\right\}, \xi_{n}\right)-f_{i}\left(\left\{Y_{k, n}^{[0]}\right\}, \xi_{n}\right)}{\Delta Y_{j}}, \quad i=1, \ldots, N \tag{3.20}
\end{equation*}
$$

where $Y_{k, n}^{[0]}$ is the $k$ th component of $Y_{n}^{[0]}, \delta_{k j}$ is the Kronecker symbol,

$$
\delta_{k j}= \begin{cases}0, & k \neq j  \tag{3.21}\\ 1, & k=j\end{cases}
$$

and the increment $\Delta Y_{j}$ in the $j$ th solution component is selected as follows: The standard choice for $\Delta Y_{j}$ is

$$
\begin{equation*}
\left.\Delta Y_{j}=\sqrt{u} \mid Y_{j, n}^{[0]}\right\} \tag{3.22}
\end{equation*}
$$

This equation cannot be used if $Y_{j, n}^{[0]}$ is either equal to zero or very small. Therefore an alternative value, based on noise level, is deduced as follows: Now the error in each $f_{i}$ due to roundoff is of order $u f_{i}$. Hence in replacing $\partial f_{i} / \partial y_{j}$ by the difference quotient, equation (3.20), the resulting element $J_{i j}$ has an error of order $u \mid f_{l} l / r_{j}$, where for clarity in presentation we have replaced $\Delta Y_{j}$ by $r_{j}$. Finally because the method coefficient $\beta_{0}\left(=\ell_{0}\right)$ is of order unity (see tables 2.1 and 2.2), the error $\delta P_{i j}$ in the element $P_{i j}$ of the iteration matrix $\mathbf{P}$, equation (2.25), is approximately

$$
\begin{equation*}
\delta P_{i j} \simeq|h| u\left|f_{i}\right| / r_{j} \tag{3.23}
\end{equation*}
$$

If we introduce the $N$-dimensional column vector $\underline{s}$, with element $s_{j}$ defined as

$$
\begin{equation*}
s_{j}=1 / r_{j}, \quad j=1, \ldots, N \tag{3.24}
\end{equation*}
$$

the matrix $\delta \mathbf{P}$ containing the errors $\left\{\delta P_{i j}\right\}$ is given by

$$
\begin{equation*}
\delta \mathbf{P}=|h| u \mid f \underline{s}^{T}, \tag{3.25}
\end{equation*}
$$

where $|\mathrm{f}|$ is an $N$-dimensional column vector containing the absolute values of the $f_{i}(i=1, \ldots, N)$ and the superscript $T$ indicates transpose. A suitable increment $r_{j}$ is obtained by bounding $|\mathbf{\delta P}|$, as discussed next.

To be consistent with the corrector convergence test, equation (2.98), and the local error test, equation (2.91), we use the weighted rms norm, which for an arbitrary $N$-dimensional column vector $\underline{x}$ is given by

$$
\begin{equation*}
\|\underline{x}\|=\sqrt{\frac{1}{N} \sum_{i=1}^{N}\left(\frac{x_{i}}{\mathrm{EWT}_{i}}\right)^{2}} \tag{3.26}
\end{equation*}
$$

If we introduce the diagonal matrix $\mathbf{D}$ of order $N$, with element $D_{i i}$ given by

$$
\begin{equation*}
D_{i i}=1 / \mathrm{EWT}_{i}, \quad i=1, \ldots, N, \tag{3.27}
\end{equation*}
$$

it is easily verified that

$$
\begin{equation*}
\|\underline{x}\|=\|\mathbf{D} \underline{x}\|_{E} / \sqrt{N}, \tag{3.28}
\end{equation*}
$$

where $\left\|\|_{E}\right.$ is the Euclidean norm, defined for $\underline{x}$ as

$$
\begin{equation*}
\|\underline{x}\|_{E}=\sqrt{\sum_{i=1}^{N}\left|x_{i}\right|^{2}} \tag{3.29}
\end{equation*}
$$

Now the norm of $\delta \mathbf{P}$ is given by

$$
\begin{equation*}
\|\delta \mathbf{P}\|=\max _{\underline{\underline{x}}} \frac{\|\delta \mathbf{P} \underline{\underline{x}}\|}{\|\underline{\underline{x}}\|} \tag{3.30}
\end{equation*}
$$

where

$$
\begin{equation*}
\left\|\delta \mathbf{P}_{\underline{x}}\right\|=|h| u\|\underline{\underline{1}}\| \underline{s}^{T} \underline{\underline{x}} \mid \tag{3.31}
\end{equation*}
$$

because $\delta \mathbf{P}$ is of rank 1. Hence

$$
\begin{aligned}
& \|\delta \mathbf{P}\|=|h| u\|f\| \max _{\underline{x}} \frac{\left|\underline{s}^{T} \underline{x}\right|}{\|\underline{x}\|}=|h| u\|\underline{f}\| \max _{\underline{\mathrm{x}}} \frac{\left|\left(\mathbf{D}^{-1} \underline{\mathrm{~s}}\right)^{T}(\mathbf{D} \underline{\mathrm{x}})\right|}{\|\underline{x}\|} \\
& =|h| u\|\underline{f}\| \max _{\underline{x}} \frac{\left\|\mathbf{D}^{-1} \underline{s}\right\|_{E}\left\|\mathbf{D}_{\underline{x}}\right\|_{E}}{\|\mathbf{D}\|_{E} / \sqrt{N}}=|h| u\|\underline{f}\| \sqrt{N}\left\|\mathbf{D}^{-1} \underline{s}\right\|_{E} \\
& \leq|h| u N\|\underline{f}\| \max _{i}\left(\frac{\mathrm{EWT}_{i}}{r_{i}}\right),
\end{aligned}
$$

which can be rewritten as

$$
\begin{equation*}
\|\delta \mathbf{P}\| \leq \frac{|h| u N\|\mathrm{f}\|}{\min _{i}\left(r_{i} / \mathrm{EWT}_{i}\right)} \tag{3.32}
\end{equation*}
$$

To establish the maximum allowable error in $\mathbf{P}$, we consider the linear system $\mathbf{P} \underline{x}=\underline{b}$, which is the form of the equation to be solved at each Newton iteration, equation (2.24). To first order, the error $\delta \underline{x}$ in $\underline{x}$ due to the error $\delta \mathbf{P}$ in $\mathbf{P}$ is given by (e.g., ref. 13)

$$
\begin{equation*}
\frac{\|\delta \mathbf{x}\|}{\|\underline{x}\|} \leq\left\|\mathbf{P}^{-1}\right\|\|\mathbf{P}\| . \tag{3.33}
\end{equation*}
$$

The norm $\left|\mathbf{P}^{-1}\right|$ is not known but is expected to be of order unity because $\mathbf{P} \rightarrow \mathbf{I}$, the identity matrix of order $N$, when $h \rightarrow 0$ and $\mathbf{P} \sim-h \beta_{0} \mathbf{J}$ when $h \rightarrow \infty$ (see eq. (2.25)). Therefore, a reasonable strategy is to bound $|\delta \mathbf{P}|$ alone by selecting a suitably small value for the relative error that can be tolerated in the Newton correction vector. By using a value of 0.1 percent for this error, we obtain from equations (3.32) and (3.33)

$$
\begin{equation*}
\min \left(\frac{r_{i}}{\mathrm{EWT}_{i}}\right) \geq 10^{3}|h| u N\|\underline{f}\| \equiv r_{0} . \tag{3.34}
\end{equation*}
$$

For additional safety $r_{0}$ is reset to 1 if it is equal to zero. Finally the increment $\Delta Y_{J}$ in the $j$ th variable used to estimate the $\left\{J_{i j}\right\}$ is given by

$$
\begin{equation*}
\Delta Y_{j}=\max \left(\sqrt{u}\left|Y_{j, n}^{[0]}\right|, r_{0} \mathrm{EWT}_{j, n}\right) \tag{3.35}
\end{equation*}
$$

For a full Jacobian matrix the above procedure will require $(N+1)$ derivative evaluations and can therefore become much more expensive than the use of an analytical Jacobian, especially for large $N$. Now $\underline{f}\left(\underline{Y}_{n}^{[0]}, \xi_{n}\right)$ is required by the corrector (see eq. (2.36)), irrespective of the iteration technique. Hence the use of MITER $=2$ requires the evaluation of only $N$ additional derivatives.

In generating the finite-difference banded Jacobian matrix (MTER =5) the code exploits the bandedness of the matrix for efficiency. The number of additional derivative evaluations required to form the Jacobian matrix is only $M L+M U+1$, where $M L$ and $M U$ are, respectively, the lower and upper half-bandwidths of the Jacobian matrix.

If JN iteration with MITER $=3$ is used, the $N$ diagonal elements $J_{i i}(i=1, \ldots, N)$ are estimated by using the approximation

$$
\begin{equation*}
J_{i i} \cong \frac{f_{i}\left(\underline{\mathrm{Y}}_{n}^{[0]}+\Delta \underline{\mathrm{Y}}, \xi_{n}\right)-f_{i}\left(\underline{\mathrm{Y}}_{n}^{[0]}, \xi_{n}\right)}{\Delta Y_{i}}, \quad i=1, \ldots, N \tag{3.36}
\end{equation*}
$$

which requires only one additional derivative evaluation. The increment $\Delta Y_{i}$ is selected as follows: Now equation (2.17) shows that if functional iteration were used, the correction $Y_{n}^{[1]}-Y_{n}^{[0]}$ that would be obtained on the first iteration is equal to the quantity $\beta_{0} g\left(Y_{n}^{[0]}\right)$, where the vector function $\underline{g}$ is given by equation (2.16). The increment vector $\Delta \underline{Y}$ is taken to be 10 percent of this correction:

$$
\begin{equation*}
\Delta Y_{i}=0.1 \beta_{0} g_{i}\left(\underline{\mathrm{Y}}_{n}^{[0]}\right), \quad i=1, \ldots, N . \tag{3.37}
\end{equation*}
$$

Hence the diagonal matrix approximation, equation (3.36), resembles a directional derivative of $\underline{f}$ taken in the same direction as the correction vector above. Also, this approximation gives the correct Jacobian if it is a constant diagonal matrix. If the magnitude of $\Delta Y_{i}$ is less than $0.1 u \beta_{0} \mathrm{EWT}_{i}$, that is, if $\left|g_{i}\left(\underline{Y}_{n}^{[0]}\right)\right|<u \mathrm{EWT}_{i}, J_{i i}$ is set equal to zero.

### 3.4.6 Solution of Linear System of Equations

If NR iteration is used for the problem, a linear system of the form $\mathbf{P} \underline{x}=\underline{b}$ must be solved for the correction vector $x$ at each iteration (see eq. (2.24)). The linear algebra necessary to solve this equation is performed by the LU method (e.g.,

## 3. Description of Code

refs. 5 and 36 ), rather than by explicitly inverting the iteration matrix, which will require prohibitive amounts of computer time (ref. 13). In the LU method the iteration matrix is factored into the product of two triangular matrices $\mathbf{L}$ and $\mathbf{U}$. Solving equation (2.24) then requires the fairly simple solution of two triangular linear systems in succession.

LSODE also includes special procedures for the LU-decomposition of the iteration matrix and the solution of equation (2.24) when the matrix is known to be banded. Compared to a full matrix, it is significantly less expensive to form a banded matrix, perform its LU-decomposition, and solve the linear system of equations (refs. $5,25,26$, and 36 ). An important advantage of LU -decomposing a banded matrix over inverting it is that, besides being faster, the triangular factors $\mathbf{L}$ and $\mathbf{U}$ lie within nearly the same bands as the original matrix, whereas the inverse is a full matrix (ref. 36). This feature makes the computation of the correction vector significantly faster with the LU method than by premultiplying the right-hand side of equation (2.24) with the inverse of the matrix.

If MITER $=3$ is used for the problem, the resulting iteration matrix is diagonal (see eq. (3.36)). Its inverse can therefore be obtained trivially and is used to compute the corrections.

### 3.4.7 Jacobian Matrix Update

The difficulty with Newton-Raphson iteration is the computational cost associated with forming the Jacobian matrix and the linear algebra required to solve for the correction vector at each iteration. However, as discussed in chapter 2 , the iteration matrix need not be very accurate. This fact is exploited to reduce the computational work associated with linear algebra by not updating $\mathbf{P}$ at every iteration. For additional savings it is updated only when the iteration does not converge. Hence the iteration matrix is only accurate enough for the solution to converge, and the same matrix may be used over several steps. It is also updated if three or more error test failures occur on any step. Now $\mathbf{P}$ may be altered if the coefficient $h \beta_{0}$ is changed (see eq. (2.25)) because a new step size and/or method order is selected. In order to minimize convergence failures caused by an inaccurate $\mathbf{P}$, the code updates $\mathbf{P}$ and performs its LU-decomposition (or inversion if MITER $=3$ ) if $h \beta_{0}$ has changed by more than 30 percent since the last update of $\mathbf{P}$. In addition, for MITER $=3$, because $\mathbf{P}^{-1}$ can be generated inexpensively, it is first modified to account for any change in $h \beta_{0}$ since its last update, before the corrections are computed. The reevaluation and LUdecomposition or inversion are also done whenever the user changes any input parameter required by the code. Finally the same $\mathbf{P}$ is used for a maximum number of 20 steps, after which it is reevaluated and LU-decomposed or inverted.

### 3.4.8 Corrector Iteration Convergence and Corrective Actions

Irrespective of the solution method and the corrector iteration technique, the maximum number of corrector iterations attempted on any step is set equal to 3 ,
based on experience that a larger number increases the computational cost without a corresponding increase in the probability of successful convergence (refs. 19, 21,22 , and 25). In addition to performing the convergence test, equation (2.99), at each iteration, STODE examines the value of the convergence rate $c_{m}$, equation (2.102). If $c_{m}$ is greater than 1 , the iteration is clearly not converging. STODE exploits this fact by abandoning the iteration if $c_{m}$ is greater than 2 after the second iteration.

If convergence is not obtained because either (1) equation (2.99) is not satisfied after three iterations or (2) $c_{m}>2$ after the second iteration, the following corrective actions are taken: For NR and JN iterations, if $\mathbf{P}$ is not current, it is updated at $\underline{y}=\underline{Y}_{n}^{[0]}$ and LU-decomposed or inverted, and the step is retried with the same step size. However, if either $\mathbf{P}$ is current or functional iteration is used, a counter of convergence failures on the current step is increased by 1 , the step size is reduced by a factor of 4 , and the solution is attempted with the new step size. The same corrective actions are taken in the event of a singular iteration matrix.

This procedure is repeated until either convergence is obtained or the integration is abandoned because either (1) 10 convergence failures have occurred or (2) the step size has been reduced below a user-supplied minimum value $h_{\min }$. In the event of an error exit the index of the component with largest magnitude in the weighted local error vector is returned to the subprogram calling LSODE.

### 3.4.9 Local Truncation Error Test and Corrective Actions

After successful convergence STODE performs the local truncation error test, equation (2.96). If the error test fails, the step size is reduced and/or the method order is reduced by 1 by using the procedures outlined in section 3.4.10, and the step is retried. After two consecutive failures the step size is reduced by at least a factor of 5, and the step is retried with either the same or a reduced order. After three or more failures it is assumed that the derivatives that have accumulated in the Nordsieck history matrix have errors of the wrong order. Therefore the first derivative is recomputed and the method order is set equal to 1 if it is greater than 1. Then the step size is reduced by a factor of 10 , the iteration matrix is formed and either $L U$-decomposed or inverted, and the step is retried with a new $\mathbf{z}_{n-1}$ that is constructed from $\underline{Y}_{n-1}$ and $\dot{\underline{Y}}_{n-1}=\underline{\mathrm{f}}_{\left(Y_{n-1}\right)}$ ).

This procedure is repeated until either the error test is passed or an error exit is taken because either (1) 10 error test failures have occurred or (2) the step size has been reduced below $h_{\min }$. In the event of an error exit LSODE returns the index of the component with the largest magnitude in the weighted local error vector to the calling subprogram.

If the accuracy test is passed, the step is accepted as successful, and the Nordsieck history matrix $z_{n}$ and the estimated local truncation error vector $\underline{d}_{n}$ at $\xi_{n}$ are computed by using equations (2.76) and (2.89), respectively. Irrespective of whether the step was successful or not, STODE saves the value of the most recent step size attempted on the step so that the user may, if desired, change it.

## 3. Description of Code

### 3.4.10 Step Size and Method Order Selection

In addition to advancing the solution STODE periodically computes the method order and step size that together maximize efficiency while maintaining prescribed accuracy. As discussed in chapter 2, this result is accomplished by selecting the method order that maximizes step size. To simplify the algorithm, the code considers only the three method orders $q-1, q$, and $q+1$, where $q$ is the current method order. For each method order the step size that will satisfy exactly the local error bound is computed by assuming that the highest derivative remains constant. The resulting step size ratios (defined as the ratio of the step size to be attempted on the next step to the current value $h_{n}$ ) are given by equations (2.107), (2.103), and (2.112), respectively, for method orders $q-1, q$, and $q+1$. These equations are, however, modified by using certain safety factors (1) to produce a smaller step size than the value that satisfies the error bound exactly, because the error estimates are not exact and the highest derivative is not usually constant, and (2) to bias the order-changing decision in favor of not changing the order at all, because any change in order requires additional work, and then in favor of decreasing the order, because an order reduction results in less work per subsequent step than an order increase. The formulas used in STODE to calculate the step size ratios are as follows:

$$
\begin{align*}
& r_{\text {down }}=\frac{1}{1.3\left[\left(D_{q-1}\right)^{\frac{1}{q}}+10^{-6}\right]},  \tag{3.38}\\
& r_{\text {same }}=\frac{1}{1.2\left[\left(D_{q}\right)^{\frac{1}{q+1}}+10^{-6}\right]}  \tag{3.39}\\
& r_{\text {up }}=\frac{1}{1.4\left[\left(D_{q+1}\right)^{\frac{1}{q+2}}+10^{-6}\right]}
\end{align*}
$$

In equations (3.38) to (3.40) the factors $1.2,1.3,1.4$, and $10^{-6}$ are strictly empirical. The subscripts "down," "same," and "up" indicate, respectively, that the method order is to be reduced by 1 , left unchanged, and increased by 1 .

To prevent an order increase either after a failed step or when $q=q_{\max }$, the
maximum order allowed for the solution method, $r_{\text {up }}$ is set equal to zero in such cases. Similarly, if $q=1, r_{\text {down }}$ is set equal to zero to avoid an order reduction.

The maximum step size ratio $r=\max \left(r_{\text {down }}, r_{\text {same }}, r_{\text {up }}\right)$ and the corresponding method order are selected to be attempted on the next step if $r \geq 1.1$ after a successful step. Changes in both step size and method order are rejected if the step size increase is less than 10 percent because it is not considered large enough to justify the computational cost required by either change (refs. 10 and 22). After a failed step the method order is decreased if $r_{\text {down }}>r_{\text {same }}$; however, $r=\max$ ( $r_{\text {down }}, r_{\text {same }}$ ) is reset to 1 if it is greater than 1 . Several additional tests, given next, are performed on $r$, if $r \geq 1.1$ after a successful step, but irrespective of the value of $r$ after a failed step, before the step size $h^{\prime}\left(=r h_{n}\right)$ to be attempted next is selected.

If the maximum step size $h_{\max }$ to be attempted on any step has been specified by the user, $r$ is restricted to

$$
\begin{equation*}
r \leftarrow \min \left(r, \frac{h_{\max }}{h_{n}}\right) \tag{3.41}
\end{equation*}
$$

Similarly if the user has specified a minimum step size $h_{\text {min }}$ that may be attempted on any step, $r$ is restricted to

$$
\begin{equation*}
r \leftarrow \max \left(r, \frac{h_{\min }}{h_{n}}\right) \tag{3.42}
\end{equation*}
$$

Finally $r$ must satisfy the inequality

$$
\begin{equation*}
r \leq r_{\max } \tag{3.43}
\end{equation*}
$$

where the variable $r_{\text {max }}$ is normally set equal to 10 . However, for the very first step size increase for the problem, if no convergence or error test failure has occurred, $r_{\max }$ is set equal to $10^{4}$ to compensate for the small step size attempted on the first step. For the first step size increase following either a corrector convergence failure or a truncation error test failure, $r_{\max }$ is set equal to 2 to inhibit a recurrence of the failure.

To avoid numerical instability caused by frequent changes in the step size, method order and step size changes are attempted only after $S$ successful steps with the same method order and step size, where $S$ is normally set equal to $q+1$. However, if an unsuccessful step occurs, this rule is disregarded and the step size and/or the method order may be reduced. Following a failed error test or a failed convergence test with either functional iteration or $\operatorname{NR}$ and $J N$ iterations if $\mathbf{P}$ is current, $S$ is set equal to $q+1$. If three or more error test failures occur on any one

## 3. Description of Code

step, $S$ is set equal to 5 even though the method order is reduced to 1 . Finally following a step for which step size and method order changes are rejected because $r<1.1, S$ is set equal to 3 .

After every $S-1$ successful steps STODE saves the vector e, if $q<q_{\max }$, in order to estimate $\nabla \underline{\mathrm{e}}$, which is required to compute $r_{\mathrm{up}}$ (see eqs. (2.109) to (2.112)). To minimize storage requirements, $\underline{e}_{n}$ is saved as the $q_{\max }$ th, that is, the last, column of $z_{n}$.

### 3.5 Error Messages

The code contains many error messages-too numerous to list here. Every input parameter is tested for legality and consistency with the other input variables. If an illegal input parameter is discovered, a detailed message is printed. Each error message is self-explanatory and complete. It not only describes the mistake but in some instances tells the user how to fix the problem. Any difficulty encountered during execution will result in an error exit. A message giving the reason for termination will also be printed. If the computation stops prematurely, the user should look for the error message near the end of the output file corresponding to the logical unit number LUNIT (see chapter 4).

## Chapter 4

## Description of Code Usage

To use the LSODE package, the following subprograms must be provided: (1) a routine that manages the calls to subroutine LSODE, (2) a routine that computes the derivatives $\left\{f_{i}=d y_{i} / d \xi\right\}$ for given values of the independent variable $\xi$ and the solution vector $\underline{y}$, and (3) if an analytical Jacobian matrix $\mathbf{J}(=\partial \underline{f} / \partial \underline{y})$ is required by the corrector iteration technique selected by the user, a routine that computes the elements of this matrix. In addition, some modifications, discussed below, to the LSODE source itself may be necessary.

### 4.1 Code Installation

### 4.1.1 BLOCK DATA Variables

The user may wish to reset the values for the integer variables MESFLG (currently 1) and LUNIT (currently 6), which are both set either in the BLOCK DATA module (double-precision version) or in subroutine XERRWV (single-precision version). The variable MESFLG controls the printing of error messages from the code, and LUNIT is the logical unit number for such output (see table 3.7). Setting MESFLG $=0$ will switch off all output from the code and therefore is not recommended.

The single-precision version of the code loads initial values for the common block LS0001 variables ILLIN and NTREP (see table 3.8) through a DATA statement in subroutine LSODE. The same procedure is used in subroutine XERRWV for the common block EH0001 variables MESFLG and LUNIT (see table 3.7). However, on some computer systems initial values for common block elements cannot be defined by means of DATA statements outside a BLOCK DATA subprogram. In this case the user must provide a separate BLOCK DATA subprogram, to which the two DATA statements from subroutines LSODE and XERRWV must be moved. The BLOCK DATA subprogram must also contain the two common blocks EH0001 and LSO001 (see table 3.6).

## 4. Description of Code Usage

### 4.1.2 Modifying Subroutine XERRWV

The subroutine XERRWV, which prints error messages from the code, is machine and language dependent. Therefore the data type declaration for the argument MSG, which is a Hollerith literal or integer array containing the message to be printed, may have to be changed. The number of Hollerith characters stored per word is assumed to be 4, and the value of NMES, which is the length of, that is, number of characters in, MSG is assumed to be a multiple of 4 , and at most 60 . However, the routine describes the necessary modifications for several machine environments. In particular, the user must change a DATA statement and the format of statement number 10. The routine assumes that all errors are either (1) recoverable, in which case control returns to the calling subprogram, or (2) fatal, in which case the run is aborted by passing control to the statement STOP, which may be machine dependent. If a different run-abort command is needed, the line following statement number 100 , which is located near the end of the routine, must be changed.

### 4.2 Call Sequence

The call sequence to subroutine LSODE is as follows:

CALL LSODE (F, NEQ, Y, T, TOUT, ITOL, RTOL, ATOL, ITASK, ISTATE, IOPT, RWORK, LRW, IWORK, LIW, JAC, MF)

All arguments in the call sequence are used on input, but only Y, T, ISTATE, RWORK, and IWORK are used on output. Also, Y and T are set only on the first call to LSODE; the other arguments may, however, have to be reset on subsequent calls. The arguments to LSODE are defined as follows:

F The name of the user-supplied subroutine that computes the derivatives of the dependent variables with respect to the independent variable. This name must be declared EXTERNAL in the subprogram calling LSODE. The requirements of subroutine $F$ are described in section 4.3.

NEQ The number of first-order ordinary differential equations (ODE's) to be solved. (The code allows the user to decrease the value of NEQ during the course of solving the problem. This option is useful if some variables can be discarded as the solution evolves as, for example, in chemical kinetics problems for which the reaction mechanism is reduced dynamically.) As discussed later, NEQ can be specified as an array. In this case NEQ(1) must give the number of ODE's to be solved, and the subprogram calling LSODE must contain a dimension statement for NEQ.

A vector of length NEQ (or more) containing the dependent variables. The subprogram calling LSODE must include a dimension statement for $Y$ if it contains more than one component. On the first call to LSODE this vector must be set equal to the vector of initial values of the dependent variables. Upon every return from LSODE, Y is the solution vector either at the desired value (TOUT or TCRIT, see below) of the independent variable or that generated at the end of the previous integration step. In case of an error exit Y contains the solution at the last step successfully completed by the integrator.

T The independent variable. On the first call to LSODE, T must give the initial value of this variable. On every return from LSODE, $T$ is either the independent variable value (TOUT or TCRIT, see below) at which the solution is desired or the independent variable value to which the numerical solution was advanced on the previous integration step. If an error exit occurs, $T$ gives the value of the farthest point (in the direction of integration) reached by the integrator.

TOUT The next value of the independent variable at which the solution is required, if ITASK $=1,3$, or 4 (see table 4.1). For ITASK $=2$ or 5 , LSODE uses TOUT on the first call to determine the direction of integration and, if necessary, to compute the step size to be attempted on the first step; on subsequent calls TOUT is ignored. LSODE permits integration in either direction of the independent variable.

ITOL A flag that indicates the type of local error control to be performed. The legal values that can be assigned for ITOL and their meanings are

## TABLE 4.1.-VALUES OF ITASK USED IN LSODE AND THEIR MEANINGS

| ITASK | Description |
| :---: | :---: |
| ${ }^{2} 1$ | Compute output values of $Y(\xi)$ at $\xi=\xi_{\text {out }}$ by overshooting and <br> interpolation. |
| 2 | Advance the solution to the ODE's by one step and return to <br> calling subprogram. <br> Stop at the first intemal mesh point at or beyond $\xi=\xi_{\text {our }}$ and <br> return to calling subprogram. <br> Compute output values of $Y(\xi)$ at $\xi=\xi_{\text {out }}$ but without over- <br> shooting $\xi=\xi_{\text {crit }}$ <br> $\mathrm{b}_{5}$ |
| Advance the solution to the ODE's by one step without passing <br> $\xi=\xi_{\text {crit }}$ and return to calling subprogram. |  |

${ }^{2}$ User must supply value for $\xi_{\text {out }}$ ( $=$ TOUT).
${ }^{6}$ User must supply value for $\xi_{\text {crit }}(=$ TCRIT). This option is useful if the problem has a singularity at or beyond $\xi=\xi_{\text {crit }}$.

TABLE 4.2.-VALUES OF ITOL USED IN LSODE AND THEIR MEANINGS

| ITOL | Description |
| :---: | :---: |
| 1 | Scalar RTOL and scalar ATOL |
| 2 | Scalar RTOL and array ATOL |
| 3 | Array RTOL and scalar ATOL |
| 4 | Array RTOL and array ATOL |

given in table 4.2. The variables RTOL and ATOL are described next.
RTOL The local relative error tolerance parameter for the solution. This parameter can be specified either as a scalar, so that the same tolerance is used for all dependent variables, or as any array of length NEQ, so that different tolerances are used for different variables. In the latter case the subprogram calling LSODE must contain a dimension statement for RTOL.

ATOL The local absolute error tolerance parameter for the solution. This parameter can also be specified either as a scalar, so that the same tolerance is used for all dependent variables, or as an array of length NEQ, so that different tolerances are used for different variables. In the latter case the subprogram calling LSODE must contain a dimension statement for ATOL.

ITASK An index that specifies the task to be performed. This flag controls when LSODE stops the integration and returns the solution to the calling subprogram. The legal values for ITASK and their meanings are given in table 4.1. If ITASK $=4$ or 5 , the input variable TCRIT (= independent variable value that the integrator must not overshoot, see table 4.1) must be passed to LSODE as the first element of the array RWORK (defined below).

ISTATE An index that specifies the state of the calculation, that is, if the call to LSODE is the first one for the problem or if it is a continuation. The legal values for ISTATE that can be used on input and their meanings are given in table 4.3. The option ISTATE $=3$ allows changes in the input parameters NEQ, ITOL, RTOL, ATOL, IOPT, MF, ML, and MU and any optional input parameter, except HO, discussed in the descriptions of RWORK and IWORK. The integer variables IOPT, MF, ML, and MU are defined below. The parameters ITOL, RTOL, and ATOL may also be changed with ISTATE $=2$, but LSODE does not then check the legality of the new values. On return from LSODE, ISTATE has the values and meanings given in table 4.4.

TABLE 4.3.-VALUES OF ISTATE THAT CAN BE USED ON INPUT TO LSODE AND THEIR MEANINGS

| ISTATE | Description |
| :---: | :--- |
| 1 | This is the first call for the problem. <br> 2 <br> 3 |
| This is not the first call for the problem, and the calculation is to <br> be continued normally with no change in any input parameters <br> except possibly $\xi_{\text {our }}$ and ITASK. <br> This is not the first call for the problem, and the calculation is to <br> be continued normally, but with a change in input parameters <br> other than $\xi_{\text {out }}$ and ITASK. |  |

${ }^{3}$ See table 4.1 for description of ITASK.

TABLE 4.4.-VALUES OF ISTATE RETURNED BY LSODE AND THEIR MEANINGS

| ISTATE | Meaning |
| :---: | :---: |
| 1 | Nothing was done because TOUT $=T$ on first call to LSODE. (However, an internal counter was set to detect and prevent repeated calls of this type.) |
| 2 | The integration was performed successfully. |
| -1 | Excessive amount of work was done on this call (i.e., number of steps exceeded MXSTEP ${ }^{1}$ on this call), but the integration was successful as far as the value returned in $T$. |
| -2 | Too much accuracy was requested for the computer being used, but the integration was successful as far as the value returned in $T$. (If this error is detected on the first call to LSODE (i.e., before any integration is done), an illegal input error (ISTATE $=-3$, see below) occurs instead.) |
| $-3$ | Illegal input was specified. The error message is detailed and selfexplanatory. |
| -4 | Repeated error test failures occurred on one step, but the integration was successful as far as the value retumed in $T$. |
| -5 | Repeated convergence test failures occurred on one step, but the integration was successful as far as the value returned in $T$. |
| $-6$ | Some component, EWT ${ }_{i}$, of the error weight vector EWI ${ }^{\text {b }}$ vanished, so that the local error test cannot be applied, but the integration was successful as far as the value returned in T. (This condition arises when pure relative error control (i.e., ATOL ${ }_{i}$ $=0^{b}$ ) was specified for a variable whose magnitude is now zero.) |

[^6]
## 4. Description of Code Usage

IOPT An integer flag that specifies if any optional input is being used on this call. The legal values for IOPT together with their meanings are given in table 4.5. The optional input parameters that may be set by the user are given in table 4.6. For each such input variable this table lists its location in the call sequence, its meaning, and its default value. The quantities RWORK and IWORK are work arrays described below.

TABLE 4.5.-VALUES OF IOPT THAT CAN BE USED ON INPUT TO LSODE AND THEIR MEANINGS

| IOPT | Meaning |
| :---: | :---: |
| 0 | The user has not set a value for any optional input parameter. ${ }^{2}$ <br> (Default values will be used for all these parameters.) <br> Values have been specified for one or more optional input <br> parameters. |

${ }^{\text {a }}$ See table 4.6 for a list of these parameters.

TABLE 4.6.-OPTIONAL INPUT PARAMETERS THAT CAN BE SET BY USER AND THEIR LOCATIONS, MEANINGS, AND DEFAULT VALUES

| Optional input parameter | Location | Meaning | Default value |
| :---: | :---: | :---: | :---: |
| H0 | RWORK(5) | Step size to be attempted on the first step | Computed by LSODE |
| HMAX | RWORK(6) | Absolute value of largest step size (in magnitude) to be used on any step | $\infty$ |
| HMIN | RWORK(7) | Absolute value of smallest step size (in magnitude) to be used on any step ${ }^{2}$ | 0 |
| MAXORD | IWORK(5) | Maximum method order to be used on any step | 12 for Adams-Moulton method and 5 for backward differentiation formula method |
| MXSTEP | IWORK(6) | Maximum number of integration steps allowed on any one call to LSODE | 500 |
| MXHNIL | IWORK(7) | Maximum number of times that warning message that step size is getting too small is printed | 10 |

[^7]RWORK A real work array used by the integrator. The subprogram calling LSODE must include a dimension statement for RWORK. If ITASK = 4 or 5 , the user must set RWORK(1) = TCRIT (see table 4.1) to transmit this variable to LSODE. If any optional real input parameters are used, their values are also passed in this array to LSODE; the address for each of these parameters is given in table 4.6. Upon return from LSODE, RWORK contains several optional real output parameters. For each such output variable table 4.7 lists its location in RWORK and its meaning. In addition, the Nordsieck history array at the current value of the independent variable (TCUR in table 4.7) and the estimated local error vector in the solution incurred on the last successful step can be obtained from RWORK. Table 4.8 lists the names used for these two quantities and their locations in RWORK. In this table NYH is the value of NEQ on the first call to LSODE, and NQCUR and LENRW are both defined in table 4.7, which also gives their locations in the array IWORK (see below).

LRW Length of the real work array RWORK. Its minimum value depends on the method flag MF (see below) and is given in table 4.9 for each legal value of MF. In this table the integer MAXORD is the maximum method order (default values $=12$ and 5 for the AM and BDF methods, respectively) to be used. The integers ML and MU are the lower and upper half-bandwidths, respectively, of the Jacobian matrix if it is declared to be banded (see table 3.2).

IWORK An integer work array used by the integrator. The subprogram calling LSODE must include a dimension statement for IWORK. If MITER ( $=$ second decimal digit of MF, defined below) $=4$ or 5 (table 3.2), the user must set IWORK(1) $=\mathrm{ML}$ and IWORK (2) $=\mathrm{MU}$ (see descriptions above) to transmit these variables to LSODE. If any optional integer input parameters are used, their values are also passed in this array to LSODE; the address for each of these parameters is given in table 4.6. Upon return from LSODE, IWORK contains several optional integer output parameters. For each such output variable table 4.7 lists its location in IWORK and its meaning.

LIW Length of the integer work array IWORK. Its minimum value depends on MITER (table 3.2) and is given in table 4.10 for each legal value of MITER.

JAC The name of the user-supplied subroutine that computes the elements of the Jacobian matrix. This name must be declared EXTERNAL in the subprogram calling LSODE. The form and description of subroutine JAC are given in section 4.4.

TABLE 4.7.-OPTIONAL OUTPUT PARAMETERS RETURNED BY LSODE AND THEIR LOCATIONS AND MEANINGS

| $\begin{array}{c}\text { Optional } \\ \text { output } \\ \text { parameter }\end{array}$ | Location | Meaning |
| :--- | :--- | :--- |
| HU | RWORK(11) | $\begin{array}{l}\text { Step size used on last successful step } \\ \text { HCUR } \\ \text { TCUR } \\ \text { RTep size to be attempted on next step } \\ \text { Curent value of independent variable. The } \\ \text { integrator has successfully advanced the } \\ \text { solution to this point. } \\ \text { A tolerance scale factor, greater than 1.0, that } \\ \text { is computed when too much accuracy is } \\ \text { requested (ISTATE = -2 or -3, see table 4.4). } \\ \text { To continue integration with the same ITOL, } \\ \text { the local error tolerance parameters RTOL } \\ \text { and ATOL must both be increased by at }\end{array}$ |
| least a factor of TOLSF. |  |  |$\}$ RWORK(13) $\quad$| NST |
| :--- |
| Number of integration steps used so far for |
| problem |
| NJE |

TABLE 4.8.-USEFUL INFORMATIONAL QUANTITIES REGARDING INTEGRATION THAT CAN BE OBTAINED FROM ARRAY RWORK AND THEIR NAMES AND LOCATIONS

| Quantity | Name | Location |
| :---: | :--- | :--- |
| Nordsieck history array for problem | YH | RWORK(21) to <br> RWORK(20 + NYH(NQCUR + 1)) <br> Estimated local error in solution on <br> last successful step | ACOR | RWORK(LENRW - NEQ + 1) to |
| :--- |
| RWORK(LENRW) |

TABLE 4.9.-MINIMUM LENGTH REQUIRED BY REAL WORK ARRAY RWORK (i.e., MINIMUM LRW) FOR EACH MF

| MF | Minimum LRW |
| :--- | :--- |
| 10,20 | $20+\mathrm{NYH}(\mathrm{MAXORD}+1)+3 \mathrm{NEQ}$ |
| $11,12,21,22$ | $22+\mathrm{NYH}(\mathrm{MAXORD}+1)+3 \mathrm{NEQ}+(\mathrm{NEQ})^{2}$ |
| 13,23 | $22+\mathrm{NYH}(\mathrm{MAXORD}+1)+4 \mathrm{NEQ}$ |
| $14,15,24,25$ | $22+\mathrm{NYH}(\mathrm{MAXORD}+1)+(2 \mathrm{ML}+\mathrm{MU}+4) \mathrm{NEQ}$ |

${ }^{2}$ NYH is the number of ODE's specified on first call to LSODE, MAXORD is the maximum method order to be used for problem, NEQ is the number of ODE's specified on current call to LSODE, and ML and MU are, respectively, the lower and upper halfbandwidths of the banded Jacobian matrix.

| TABLE 4.10.-MINIMUM LENGTH REQUIRED BY INTEGER WORK ARRAY IWORK (i.e., MINIMUM LIW) FOR EACH MITER |  |
| :---: | :---: |
| MITER ${ }^{\text {a }}$ | Minimum LIW ${ }^{\text {b }}$ |
| $\begin{aligned} & 0 \\ & 1,2 \\ & 3 \\ & 4,5 \end{aligned}$ | $\begin{gathered} 20 \\ 20+\text { NEQ } \\ 20 \\ 20 \text { + NEQ } \end{gathered}$ |
| ${ }^{2}$ See table 3.2 for description of MTER. <br> ${ }^{\mathrm{b}} \mathrm{NEQ}$ is the number of ODE's specified on current call to LSODE. |  |

MF Method flag that indicates both the integration method and corrector iteration technique to be used. MF consists of the two decimal digits METH, which specifies the integration method, and MITER, which specifies the iteration technique (eq. (3.1)). Equation (3.1) and tables 3.1 and 3.2 show that MF has the following 12 legal values$10,11,12,13,14,15,20,21,22,23,24$, and 25 . If $\mathrm{MF}=14,15,24$, or 25, the values of ML and MU must be passed to LSODE as the first and second elements, respectively, of the array IWORK (see above).

### 4.3 User-Supplied Subroutine for Derivatives (F)

Irrespective of the solution method or corrector iteration technique selected to solve the problem, the user must provide a subroutine that computes the derivatives $\left\{f_{i}\right\}$ for given values of the independent variable and the solution vector. The name ( F ) of this subroutine is an argument in the call vector to LSODE and must

## 4. Description of Code Usage

therefore be declared EXTERNAL in the subprogram calling LSODE. The derivative subroutine F should have the form

> SUBROUTINE F (NEQ, T, Y, YDOT) DIMENSION Y(1), YDOT(1) in FORTRAN 66 or DIMENSION Y(*), YDOT(*) in FORTRAN 77

In addition, if NEQ is an array, the subroutine $F$ should include a DIMENSION statement for it. The routine $F$ should not alter the values in T, NEQ (or NEQ(1), if NEQ is an array), or the first N elements in Y , where N is the current number of ODE's to be solved. The derivative vector should be returned in the array YDOT, with $\operatorname{YDOT}(\mathrm{I})=d y_{i} / d \xi(i=\mathrm{I})$, evaluated at $\xi=\mathrm{T}, \mathrm{y}=\mathrm{Y}$.

If the calculation of $\left\{f_{i}\right\}$ involves intermediate quantities whose current values, that is, at $\xi=\xi_{n}$ (or $\xi_{\text {out }}$ ), are required externally to LSODE, a special calculation, such as a call to the routine F , must be made. The results of the last call from the package to the routine $F$ should not be used because they correspond to a $Y$ value that is different from $\underline{Y}_{n}\left[\right.$ or $\left.\underline{Y}\left(\xi_{\text {out }}\right)\right]$ and a $\xi$ value that may be different from $\xi_{n}$ (or $\xi_{\text {out }}$ ). Here $\xi_{n}$ is the independent variable value to which the numerical solution was advanced on the previous integration step and $\xi_{\text {out }}=$ TOUT. If a special call to subroutine $F$ is made, to reduce the storage requirement, the YDOT argument may be replaced with RWORK(LSAVF), the base address of an N-dimensional array, SAVF (see table 3.8), used for temporary storage by LSODE; LSAVF is the 224th word (6th integer word after 218 real words) in the common block LS0001 (table 3.6). If the derivative $\underline{\mathrm{X}}_{n}$ is required, it can be obtained by calling subroutine INTDY, as explained in section 4.8 .

### 4.4 User-Supplied Subroutine for Analytical Jacobian (JAC)

If the corrector iteration technique selected by the user requires a Jacobian matrix, we recommend that a routine that computes an analytical Jacobian be provided. The name (JAC) of this routine is an argument in the call vector to LSODE and must therefore be declared EXTERNAL in the subprogram calling LSODE. The Jacobian subroutine JAC should have the form

SUBROUTINE JAC (NEQ, T, Y, ML, MU, PD, NROWPD)
DIMENSION Y(1), PD (NROWPD, 1) in FORTRAN 66
DIMENSION Y(*), PD (NROWPD, *) in FORTRAN 77

Here ML and MU are, respectively, the (user-supplied) lower and upper halfbandwidths of the Jacobian matrix if it is banded; and NROWPD, which is set by
the code, is the number of rows of the Jacobian matrix PD. For a banded matrix NROWPD is equal to the extended bandwidth ( $=2 M L+M U+1$ ), and for a full matrix it is equal to the current number N of ODE's. If NEQ is an array, the subprogram JAC must include a DIMENSION statement for it.

This routine should not alter the values in NEQ (or NEQ(1), if NEQ is an array), T, ML, MU, or NROWPD. However, the Y array may, if necessary, be altered. For a full Jacobian matrix $(M T E R=1)$ the element $\operatorname{PD}(\mathrm{I}, \mathrm{J})(\mathrm{I}=1, \ldots, \mathrm{~N}$; $\mathrm{J}=1, \ldots, \mathrm{~N})$ must be loaded with $\left.\partial f_{i} \partial y_{j}\right|_{\xi=\mathrm{T} ; \underline{\mathrm{y}}=\mathrm{Y}}(i=\mathrm{I} ; j=\mathrm{J})$. In this case the arguments ML and MU are not needed. If the Jacobian matrix is banded (MITER $=4$ ), the element $\partial f_{i} / \partial y_{j}(i=1, \ldots, \mathrm{~N} ; i-\mathrm{ML} \leq j \leq i+\mathrm{MU})$ must be loaded into PD $(\mathrm{I}-\mathrm{J}+\mathrm{MU}+1, \mathrm{~J})(\mathrm{I}=i ; \mathrm{J}=j)$. Thus each band of the Jacobian matrix must be loaded in column-wise manner, with diagonal lines of J , from the top down, loaded into the rows of PD. For a diagonal matrix $\mathrm{ML}=\mathrm{MU}=0$, and the diagonal elements must be loaded into a single row of length N . In any case the solver sets all elements of PD equal to zero before calling JAC, so that only the nonzero elements need to be loaded. Also each call to subroutine JAC is preceded by a call to subroutine F with the same arguments NEQ, T, and Y . To improve computational efficiency, intermediate quantities needed by both routines may be saved by routine F in a common block, thereby avoiding recomputation by routine JAC. If necessary, even the derivatives at T can be accessed by JAC by means of this method.

If functional iteration (MTER $=0$ ) or an internally generated Jacobian matrix (MITER $=2,3$, or 5 ) is used, a dummy version of JAC may nonetheless be required to satisfy the loader. This version may be given simply as follows:

SUBROUTINE JAC (NEQ, T, Y, ML, MU, PD, NROWPD) RETURN END

### 4.5 Detailed Usage Notes

It is apparent from the description of the call sequence to LSODE that the code has many capabilities and therefore requires the user to set values for several parameters. To further clarify code usage and assist in selecting values for userset parameters, we provide here a somewhat detailed guide. We first summarize how we expect the code to be normally used and then give detailed usage notes. Additional insight into code usage can be obtained from the discussions by Byrne and Hindmarsh (ref. 17), who examined in some detail the solution of 10 example problems representing a variety of problem types, and by Radhakrishnan (ref. 37), who studied the effects of various user-set parameters on the solution of stiff ODE's arising in combustion chemistry.

### 4.5.1 Normal Usage Mode

The normal mode of communication with LSODE may be summarized as follows:
(1) Set initial values in $Y$.
(2) Set NEQ, T, ITOL, RTOL, ATOL, LRW, LIW, and MF.
(3) Set TOUT $=$ first output station, ITASK $=1$, ISTATE $=1$, and IOPT $=0$.
(4) Call LSODE.
(5) Exit if ISTATE $<0$.
(6) Do desired output of Y.
(7) Exit if problem is finished.
(8) Reset TOUT to next print station and return to step (4).

This procedure will result in LSODE (a) computing the step size to be attempted on the first step, (b) continuing the integration with step sizes generated internally until the first internal mesh point at or, more usually, just beyond TOUT, and (c) computing the solution at TOUT by interpolation. The returned value T will be set equal to TOUT exactly, and $Y$ will contain the solution at TOUT. Because the normal output value of ISTATE is 2 , it does not have to be reset for normal continuation.

### 4.5.2 Use of Other Options

The calling subprogram may also make use of other options included in the package. For example, in step (8) ISTATE could be reset to 3 to indicate that at TOUT some parameters, such as NEQ or MF, have been changed. The task to be performed, indicated by the value of ITASK, can, however, be changed without resetting ISTATE. In the event of integration difficulties parameter values may also be changed in step (5), followed by a return to step (4), if the new values will prevent a recurrence of the indicated trouble.

### 4.5.3 Dimensioning Variables

Irrespective of the options selected, the subprogram calling LSODE must include DIMENSION statements for all call sequence variables that are arrays. Such variables include Y, RTOL, ATOL, RWORK, IWORK, and, as discussed below, possibly NEQ. The solution vector Y may be declared to be of length NEQ or greater. The first NEQ elements of the Y array must be the variables whose ODE's are to be solved. The remaining locations, if any, may be used to store other real data to be passed to the routines F and/or JAC. The LSODE package accesses only the first NEQ elements of $Y$; the remaining elements are unchanged by the code.

The parameter NEQ is usually a scalar quantity. However, an array NEQ may
be used to store and pass integer data to the routines $F$ and/or JAC. In this case the first element of NEQ must be set equal to the number of ODE's. The LSODE package accesses only NEQ(1). However, NEQ is used as an argument in the calls to the routines $F$ and JAC, so that these routines, and the MAIN program, must include NEQ in a DIMENSION statement.

### 4.5.4 Decreasing the Number of Differential Equations (NEQ)

In the course of solving a problem the user may decrease (but not increase) the number of ODE's. This option is useful if some variables reach steady-state values while others are still varying. Dropping these constant quantities from the ODE list decreases the size of the system and hence increases computational efficiency. To use this option, upon return from LSODE at the appropriate time, the calling subprogram must reset the value of NEQ (or NEQ(1)); set ISTATE = 3; reset the values of all other parameters that are either required to continue the integration, such as TOUT if ITASK $=1,3$, or 4 (table 4.1), or are changed at the user's option; and then call LSODE again. If the Jacobian matrix is declared to be banded (MITER $=4$ or 5 , table 3.2) and reductions can be made to the halfbandwidths ML and MU, they will also produce efficiency increases. The option of decreasing the number of ODE's may be exercised as often as the user wishes. Of course, each time the size of the ODE system is decreased the changes discussed above should be made and the resulting number of ODE's can never be less than 1. However, the LRW and LIW values need not be reset.

If, at any time, the number of ODE's is decreased from $N$ to $\mathrm{N}^{\prime}$, LSODE will drop the last $N-N^{\prime}$ ODE's from the system and integrate the first $\mathrm{N}^{\prime}$ equations. It is therefore important in formulating the problem to order the variables carefully and make sure that it is indeed the last $N-N^{\prime}$ variables that attain steady-state values. In continuing the integration LSODE will access only the first $N^{\prime}$ elements of Y. However, the remaining $\mathrm{N}-\mathbf{N}^{\prime}$, or more, elements can be accessed by the user, and so no special programming is needed in either routine F or JAC.

### 4.5.5 Specification of Output Station (TOUT)

The argument TOUT must be reset every time LSODE is called if the option given by ITASK $=1,3$, or 4 is selected. For the other two values of ITASK (i.e., 2 and 5), TOUT need be set only on the first call to LSODE. Irrespective of the value of ITASK, the TOUT value provided on the first call to LSODE is used to determine the direction of integration and, if the user has not supplied a value for it, to compute the step size to be attempted on the first step. Therefore unless the user specifies the value for the initial step size, it is recommended that some thought be given to the value used for TOUT on the first call to LSODE.

On the first call to LSODE, that is, with ISTATE $=1$, TOUT may be set equal to the initial value of the independent variable. In this case LSODE will do nothing, and so the value ISTATE = 1 will be returned to the calling subprogram; however,

## 4. Description of Code Usage

an internal counter will be updated to prevent repeated calls of this nature. If such a "first" call is made more than four times in a row, an error message will be issued and the execution terminated.

On the second and subsequent calls to LSODE there is no requirement that the TOUT values be monotonic. However, a value for TOUT that "backs up" is limited to the current internal interval [(TCUR - HU),TCUR], where TCUR is the current value of the independent variable and HU is the step size used on the previous step.

### 4.5.6 Specification of Critical Stopping Point (TCRIT)

In addition to TOUT a value must be specified for TCRIT if the option ITASK $=4$ is selected. TCRIT may be equal to TOUT or beyond it, but not behind it, in the direction of integration. The integration is not permitted to overshoot TCRIT, so that the option is useful if, for example, a singularity exists at or beyond TCRIT. This variable is also required with the option ITASK =5. In either case the first element of the array RWORK (i.e., RWORK(1)) must be set equal to TCRIT. If the solver reaches TCRIT within roundoff, it will return $\mathrm{T}=\mathrm{TCRIT}$ exactly and the solution at TCRIT is returned in Y. To continue integrating beyond TCRIT, the user must reset either ITASK or TCRIT. In either case the value of ISTATE need not be reset. However, whenever TCRIT is changed, the new value must be loaded into RWORK(1).

### 4.5.7 Selection of Local Error Control Parameters (ITOL, RTOL, and ATOL)

Careful thought should be given to the choice of ITOL, which together with RTOL and ATOL determines the nature of the error control performed by LSODE. The value of ITOL dictates the value of the local error weight vector EWT, with element $\mathrm{EWT}_{i}$ defined as

$$
\begin{equation*}
\mathrm{EWT}_{i}=\mathrm{RTOL}_{t}\left|Y_{\imath}\right|+\mathrm{ATOL}_{i}, \tag{4.1}
\end{equation*}
$$

where $\mathrm{RTOL}_{i}$ and $\mathrm{ATOL}_{i}$ are, respectively, the local relative and absolute error tolerances for the $i$ th solution component $Y_{i}$ and the bars $|\cdot|$ denote absolute value. The solver controls the estimated local errors $\left\{d_{i}\right\}$ in $\left\{Y_{i}\right\}$ by requiring the root-mean-square (rms) norm of $d_{i} / E W T T_{i}$ to be 1 or less.

Pure relative error control for the $i$ th solution component is obtained by setting $\mathrm{ATOL}_{i}=0 ; \mathrm{RTOL}_{i}$ is then a measure of the number of accurate significant figures in the numerical solution. This error control is generally appropriate when widely varying orders of magnitude in $Y_{i}$ are expected. However, it cannot be used if the solution vanishes because relative error is then undefined. Pure absolute error control for the $i$ th solution component is obtained by setting
$\mathrm{RTOL}_{i}=0 ; \mathrm{ATOL}_{i}$ is then a measure of the largest number that may be neglected.
Both RTOL and ATOL can be specified (1) as scalars, so that the same error tolerances are used for all variables, or (2) as arrays, so that different tolerances are used for different variables. The value of the user-supplied parameter ITOL indicates whether RTOL and ATOL are scalars or arrays. The legal values that can be assigned to ITOL and the corresponding types of RTOL and ATOL are given in table 4.2. If RTOL and/or ATOL are arrays, the calling subprogram must include an appropriate DIMENSION statement. A scalar RTOL is generally appropriate if the same number of significant figures is acceptable for all components of Y. A scalar ATOL is generally appropriate when all components of Y, or at least their peak values, are expected to be of the same magnitude.

In addition to ITOL, RTOL and ATOL should be selected with care. Now the code controls an estimate of only the local error, that is, an estimate of the error committed on taking a single step, starting with data regarded as exact. However, what is of interest to the user is the global truncation error or the actual deviation of the numerical solution from the exact solution. This error accumulates in a nontrivial manner from the local errors and is neither measured nor controlled by the code. It is therefore recommended that the user be conservative in choosing values for the local error tolerance parameters. However, requesting too much accuracy for the precision of the machine will result in an error exit (table 4.4). In such an event the minimum factor TOLSF by which RTOL and ATOL should both be scaled up is returned by LSODE (see table 4.7). Some experimentation may be necessary to optimize the tolerance parameters, that is, to determine values that produce sufficiently accurate solutions while minimizing the execution time. The global errors in solutions generated with particular values for the local error tolerance parameters can be estimated by comparing them with results produced with smaller tolerances. In reducing the tolerances all components of RTOL and ATOL, and hence of EWT, should be scaled down uniformly.

There is no requirement that the same values for ITOL, RTOL, and ATOL be used throughout the problem. If during the course of the problem any of these parameters is changed, the user should reset ISTATE $=3$ before calling LSODE again. (ISTATE need not be reset; however, LSODE will not then check the legality of the new values.) This option is useful, for example, if the solution displays rapid changes in a small subinterval but is relatively smooth elsewhere. To accurately track the solution in the rapidly varying region, small values of RTOL and ATOL may be required. However, in the smooth regions these tolerances could be increased to minimize execution time.

### 4.5.8 Selection of Integration and Corrector Iteration Methods (MF)

The choice of the method flag MF may also require some experimentation. The user should consider the nature of the problem and storage requirements. The primary consideration regarding MF is stiffness. If the problem is not stiff, the best choice is probably $\mathrm{MF}=10$ (Adams-Moulton (AM) method with functional

## 4. Description of Code Usage

iteration.) If the problem is stiff to a significant degree, METH should be set equal to 2 (table 3.1), and MITER (table 3.2) depends on the structure of the Jacobian matrix. If the Jacobian is banded, MITER $=4$ (user-supplied analytical Jacobian) or 5 (internally generated Jacobian by finite-difference approximations) should be used. For either of these two MITER values the user must set values for the lower (ML) and upper (MU) half-bandwidths of the Jacobian matrix. The first and second elements of the integer work array IWORK must be set equal to ML and MU , respectively; that is, $\operatorname{IWORK}(1)=\mathrm{ML}$ and $\operatorname{IWORK}(2)=\mathrm{MU}$. For a full matrix MITER should be set equal to 1 (analytical Jacobian) or 2 (internally generated Jacobian). If the matrix is significantly diagonally dominant, the choice MITER $=3$, that is, Jacobi-Newton (JN) iteration using an internally generated diagonal approximation for the Jacobian matrix, can be made. To use this iteration technique with an analytical Jacobian, set MITER $=4$ and $M L=M U=0$.

If the problem is only mildly stiff, the choice METH $=1$ (i.e., the AM method) may be more efficient than METH $=2$ (i.e., the backward differentiation formula (BDF) method). For this case experimentation would be necessary to identify the optimal METH. If the user has no a priori knowledge regarding the stiffness of the problem, one way to determine its nature is to try $\mathrm{MF}=10$ and examine the behavior of both the solution and step size pattern. (It is recommended that some upper limit be set for the total number of steps or derivative evaluations to avoid excessive run times.) If the typical values of the step size are much smaller than the solution behavior would appear to require, for example, more than 100 steps are taken over an interval in which the solution changes by less than 1 percent, the problem is probably stiff. The degree of stiffness can be estimated from the step sizes used and the smoothness of the solution.

Irrespective of the integration method selected, the least effective iteration technique is functional iteration, given by MITER $=0$, and the most effective is Newton-Raphson (NR), given by MITER $=1$ or 2 ( 4 or 5 for a banded Jacobian matrix). Generally JN iteration is somewhere in between. However, storage requirements increase in the same order as the effectiveness of the iteration technique (see table 4.9), and so trade-off considerations are necessary. For reasons of computational efficiency the user is encouraged to provide a routine for computing the analytical Jacobian, unless the system is fairly complicated and analytical expressions cannot be derived for the matrix elements. The accuracy of the Jacobian calculation can be checked by comparison with the $\mathbf{J}$ internally generated with MTTER $=2$ or 5 . Jacobi-Newton iteration requires considerably less storage and execution time per iteration but will be effective only if the Jacobian matrix is significantly diagonally dominant.

The importance of supplying an analytical Jacobian matrix, especially for large problems, is illustrated by Radhakrishnan (ref. 37), who studied 12 test problems from combustion kinetics. The problems covered a wide range of reaction conditions and reaction mechanism size. The effects on solution efficiency of (1) METH, (2) the first output station, and (3) optimizing the local error tolerances were also examined.

### 4.5.9 Switching Integration and Corrector Iteration Methods

The user may specify different values for MF in different subintervals of the problem. This option is useful if the problem changes character and is nonstiff in some regions and stiff elsewhere. Because stiff problems are usually characterized by a nonstiff initial "transient" region, one could use $\mathrm{MF}=10$ in the initial region and then switch to $M F=21$ (the BDF method with NR iteration using an analytical Jacobian matrix) in the later stiff regime. It is very straightforiward to change integration methods and corrector iteration techniques. Upon return from LSODE the user simply resets MF to the desired new value. The other action required is to reset ISTATE $=3$ before calling LSODE again. The lengths LRW and LIW, respectively, of the arrays RWORK and IWORK depend on MF (see tables 4.9 and 4.10). If different methods are to be used in the course of solving a problem, storage corresponding to at least the maximum values of LRW and LIW must be allocated. That is, the dimensions of RWORK and IWORK must be set equal to at least the largest of the LRW and LIW values, respectively, required by the different methods to be used.

### 4.6 Optional Input

In addition to the input parameters whose values are required by the code, the user can set values for several other parameters to control both the integration and the output from the code. These optional input parameters are given in table 4.6, together with their locations and default values. If any of these parameters are used, the user must set IOPT = 1 to relay this information to the solver, which will examine all optional input parameters and select only those for which nonzero values are specified. A value of zero for any parameter will cause its default value to be used. Thus to use a subset of the optional inputs, set RWORK(I) $=0.0$ and IWORK $(\mathrm{I})=0(\mathrm{I}=5$ to 7$)$, and then set parameters of interest to the desired (nonzero) values. The variable H0, the step size to be attempted on the first step, must indicate the direction of integration. That is, H 0 must be a positive quantity for integration in the forward direction (increasing values of the independent variable) and negative otherwise. All other input parameters must be positive numbers; otherwise, an error exit will occur.

To reset any optional input parameter on a subsequent call to LSODE, ISTATE must be set equal to 3. IOPT is not altered by LSODE and therefore need not be reset. Also because the code does not alter the values in RWORK (5) to RWORK (7) and IWORK(5) to IWORK(7), only parameters for which new values are required need to be reset. To specify a default value for any parameter for which a nondefault value had previously been used, simply load the appropriate location in RWORK or IWORK with a zero. Of course, if all variables are to have default values, simply reset IOPT $=0$.

## 4. Description of Code Usage

### 4.6.1 Initial Step Size (H0)

The sign of the step size H 0 must agree with the direction of integration; otherwise, an error exit will occur. Also, its magnitude should be considerably smaller than the average value expected for the problem because the code starts the integration with a first-order method. Of course, the integrator tests that the given step size does produce a solution that satisfies the local error test and, if necessary, decreases it (in magnitude). The only test made on the magnitude of H0 prior to taking the first step is that it does not exceed the user-supplied value for HMAX, the maximum absolute step size allowed for the problem.

### 4.6.2 Maximum Step Size (HMAX)

The user may have to specify a finite value for HMAX (default value, $\infty$ ) if the solution is characterized by rapidly varying transients between long smooth regions. If the step size is too large, the solver may skip over the fine detail that the user may be (primarily) interested in. An example of this behavior is the buildup of ozone and oxygen atom concentrations in the presence of sunlight (ref. 17).

### 4.6.3 Maximum Method Order (MAXORD)

The optional input parameter MAXORD, the maximum method order to be attempted on any step, should not exceed the default value-- 12 for the AM method and 5 for the BDF method. If it does, it will be reduced to the default value. Also, in the course of solving the problem, if MAXORD is decreased to a value less than the current method order, the latter quantity will be reduced to the new MAXORD.

The maximum method order has to be restricted to a value less than the default value for stiff problems when the eigenvalues of the Jacobian matrix are close to the imaginary axis; that is, the solution is highly oscillatory. In such a situation the BDF method of high order ( $\geq 3$ ) has poor stability characteristics and, as the stability plots in Gear (ref. 10) show, the unstable region grows as the order is increased. For this reason MAXORD should be set equal to 3 unless the eigenvalues are imaginary; that is, $\operatorname{Re}\left(\lambda_{l}\right)=0$ and $\operatorname{Im}\left(\lambda_{i}\right) \neq 0$, where $\operatorname{Re}\left(\lambda_{i}\right)$ and $\operatorname{Im}\left(\lambda_{i}\right)$ are the real and imaginary parts of $\lambda_{i}$, the $i$ th eigenvalue. In this case the value MAXORD $=2$ should be used.

### 4.7 Optional Output

The user is usually primarily interested in the numerical solution and the corresponding value of the independent variable. These quantities are always returned in the call variables Y and T . In addition, several optional output
quantities that contain information about the integration are returned by LSODE. These quantitites are given in tables 4.7 and 4.8, together with their locations. Some of these quantities give a measure of the computational work required and may, for example, help the user decide if the problem is stiff or if the right method is being used. Other output quantities will, in the event of an error exit, help the user either set legal values for some parameters or identify the reason for repeated convergence failures or local error test failures.

### 4.8 Other Routines

To gain additional capabilities, the user can access the following subroutines included in the LSODE package: INTDY, SRCOM, XSETF, and XSETUN. Among these, only INTDY is used by LSODE.

### 4.8.1 Interpolation Routine (Subroutine INTDY)

The subroutine INTDY provides derivatives of Y, up to the current order, at a specified point $T$ and may be called only after a successful return from LSODE. The call to this routine takes the form

CALL INTDY (T, K, RWORK(21), NYH, DKY, IFLAG)

where T, K, RWORK(21), and NYH are input parameters and DKY and IFLAG are output parameters. The arguments to INTDY are defined as follows:

T Value of independent variable at which the results are required. For the results to be valid $T$ must lie in the interval [(TCUR $\mathrm{HU}), \mathrm{TCUR}$ ], where TCUR and HU are defined in table 4.7.
$\mathrm{K} \quad$ Integer that specifies the desired derivative order and must satisfy $0 \leq K \leq$ current method order NQCUR (see table 4.7 for location of this quantity). Now, because the method order is never less than 1 , the first derivative $d \underline{Y} / d \xi$ can always be obtained by calling INTDY.

RWORK(21) Base address of the Nordsieck history array (see table 4.8).
NYH Number of ODE's used on the first call to LSODE. If the number of ODE's is decreased during the course of the problem, NYH should be saved. An alternative way of obtaining NYH is to include the common block LS0001 in the subprogram calling INTDY. LSODE saves NYH in LS0001 as the 232nd word-the 14th integer word after 218 real words (see table 3.6).

## 4. Description of Code Usage

DKY Array of length $N$ that contains the Kth derivative of $Y$ at $T$. The subprogram calling INTDY must include a DIMENSION statement for DKY if NYH > 1. Alternatively, to save storage, DKY can be replaced with RWORK(LSAVF)-see section 4.3.

IFLAG An error flag with following values and meanings:
0 Both T and K were legal.
-1 Illegal value was specified for K.
-2 Illegal value was specified for T .

### 4.8.2 Using Restart Capability (Subroutine SRCOM)

The subroutine SRCOM is useful if one is either alternating between two or more problems being solved by LSODE or interested in interrupting a run and restarting it later. The latter situation may arise, for example, if one is interested in steady-state values with no a priori knowledge of the required integration interval. The run may be stopped periodically, the results examined and, if necessary, the integration continued. This procedure is clearly more economical than making repeated runs on the same problem with, say, increasing values of TOUT. To exploit the capability of stopping and then continuing the integration, the user must save and then restore the contents of the common blocks LS0001 and EH0001. This information can be stored and restored by calling SRCOM. The call to this routine takes the form

## CALL SRCOM (RSAV, ISAV, JOB)

where RSAV must be declared as a real array of length 218 or more in the calling subprogram and ISAV as an integer array of length 41 or more and JOB is an integer flag whose value ( $=1$ or 2 ) indicates the action to be performed by SRCOM as follows: JOB = 1 means "save the contents of the two common blocks," and JOB = 2 means "restore this information."

Thus to store the contents of EH0001 and LSOOO1, SRCOM should be called as follows:

## CALL SRCOM (RSAV, ISAV, 1)

Upon return from SRCOM, RSAV and ISAV will contain, respectively, the 218 real and 39 integer words that together make up the common block LS0001. The 40th and 41 st elements of ISAV will contain the two integer words MESFLG and LUNIT in the common block EH0001 (table 3.6). The lengths and contents of the arrays RWORK and IWORK must also be saved. The lengths LENRW and LENIW required for the arrays RWORK and IWORK are saved by LSODE as the 17th and 18th elements, respectively, of the array IWORK (see table 4.7).

To continue the integration, the arrays RWORK and IWORK and the contents of the common blocks LS0001 and EHOOO1 must be restored. The common block
contents are restored by using the previously saved arrays RSAV and ISAV and calling the routine SRCOM as follows:

## CALL SRCOM (RSAV, ISAV, 2)

The user should then set values for the input parameters required by LSODE, and the integration can be continued by calling this routine. Note, in particular, that ISTATE must be set equal to 2 or 3 to inform LSODE that the present call is a continuation one for the problem (see table 4.3).

### 4.8.3 Error Message Control (Subroutines XSETF and XSETUN)

To reset the value of the logical unit number LUNIT for output of messages from the code, the routine XSETUN should be called as follows:

## CALL XSETUN (LUN)

where LUN is the new value for LUNIT. Action is taken only if the specified value is greater than zero.

The value of the flag MESFLG, which controls whether messages from the code are printed or not, may be reset by calling subroutine XSETF as follows:

## CALL XSETF (MFLAG)

where MFLAG is the new value for MESFLG. The legal values for MFLAG are 0 and 1. Specifying any other value will result in no change to the current value of MESFLG. Setting MFLAG $=0$ does carry the risk of losing valuable information through error messages from the integrator.

### 4.9 Optionally Replaceable Routines

If none of the error control options included in the code are suitable, more general error controls can be obtained by substituting user-supplied versions of the routines EWSET and/or VNORM (table 3.3). Both routines are concerned with measuring the local error. Hence any replacement may have a major impact on the performance of the code. We therefore recommend that modifications be made only if absolutely necessary, and that too with great caution. Also the effect of the changes and the accuracy of the programming should be studied on some simple problems.

### 4.9.1 Setting Error Weights (Subroutine EWSET)

The subroutine EWSET sets the array of error weights EWT, equation (4.1). This routine takes the form

## 4. Description of Code Usage

## SUBROUTINE EWSET (N, ITOL, RTOL, ATOL, YH, EWT)

where N is the current value of the number of ODE's; ITOL, RTOL, ATOL, and EWT have been defined previously; and YH contains the current Nordsieck history array, that is, the current solution vector YCUR and its NQ scaled derivatives, where NQ is the current method order. On the first call to EWSET from the routine LSODE, YCUR is the same as the Y array (which then contains the initial values supplied by the user); thereafter the two arrays may be different.

The error weights $\left\{\mathrm{EWT}_{i}\right\}$ are used in the local truncation error test, which requires that the rms norm of $d_{i} / \mathrm{EWT} \mathrm{E}_{t}$ be 1 or less. Here, $d_{i}$ is the estimated local error in $Y_{i}$. The above norm is computed in the routine VNORM (discussed in section 4.9.2) to which the EWT array is passed.

If the user replaces the current version of EWSET, the new version must return in each $\mathrm{EWT}_{i}(i=1, \ldots, \mathrm{~N})$ a positive quantity for comparison with $d_{i}$. This routine is called by the routine LSODE only (tables 3.4 and 3.5 ). However, in addition to its use in the local truncation error test (which is performed in the routine STODE), EWT is used (1) by the routine LSODE in computing the initial step size H0 and the optional output integer IMXER (table 4.7) and (2) by the routine PREPJ in computing the increments in solution vector for the difference quotient Jacobian matrix (MITER $=2$ or 5, table 3.2) and for the diagonal approximation to the Jacobian matrix (MITER = 3). The base address for EWT in the array RWORK is LEWT, which is the 222nd word (the 4th integer word after 218 real words) in the common block LS0001.

If the user's version of EWSET uses current values of the derivatives of $\underline{Y}$, they can be obtained from YH, as described later. Indeed, derivatives of any order, up to NQ, can be found from YH, whose base address in RWORK is LYH ( $=21$ ), the 221 st word (the 3 rd integer word past 218 real words) in LS0001. The array YH is of length $\mathrm{NYH}(\mathrm{NQ}+1)$, where NYH is the value of N on the first call to LSODE. The first $N$ elements correspond exactly to the YCUR array. The remaining terms contain scaled derivatives of YCUR. For example, the $N$ elements J•NYH +1 to $J \cdot N Y H+N(J=0,1, \ldots, N Q)$ contain the $J$ th scaled derivative $H^{J} \underline{Y}^{(J)} / J$ !, where $H$ is the current value of the step size. On the first call to EWSET, before any integration is done, H is (temporarily) set equal to 1.0. Thereafter its value may be determined from LS0001, where it is the 212th real word. This common block also contains NYH as the 232 nd word (the 14th integer word past 218 real words) and NQ as the 253 rd word (the 35 th integer word past 218 real words). Thus if the user wishes to use the Jth derivative in EWSET, it may be obtained by including the following statements:

SUBROUTINE EWSET (N, ..., YH, ..., EWT)
REAL (or DOUBLE PRECISION) YH, EWT, RLS, H, ...
INTEGER N, ILS, NQ, NYH, ...
DIMENSION YH(1), EWT(1), ... in FORTRAN 66
DIMENSION YH(*), EWT(*), ... in FORTRAN 77

COMMON/LS0001/RLS(218), ILS(39)
$\mathrm{NQ}=\mathrm{ILS}(35)$
$\mathrm{NYH}=\mathrm{ILS}(14)$
$\mathrm{H}=\mathrm{RLS}(212)$
The Jth derivative ( $0 \leq \mathrm{J} \leq \mathrm{NQ}$ ) is then given by

$$
\begin{equation*}
\mathrm{Y}_{\mathrm{I}}^{(\mathrm{J})}=\frac{\mathrm{J}!\mathrm{YH}(\mathrm{~J} * \mathrm{NYH}+\mathrm{I})}{\mathrm{H}^{\mathrm{J}}}, \quad \mathrm{I}=1, \ldots, \mathrm{~N} \tag{4.2}
\end{equation*}
$$

where $Y_{I}^{(J)}$ is the Jth derivative of $Y_{I}$. The routine must include a data type declaration and a DIMENSION statement for $\underline{Y}^{(J)}$. To save on storage, these values may be stored temporarily in the vector EWT.

### 4.9.2 Vector-Norm Computation (Function VNORM)

The real (or double precision) function routine VNORM computes the weighted root-mean-square (rms) norm of a vector. It is used as follows:

$$
\mathrm{D}=\mathrm{VNORM}(\mathrm{~N}, \mathrm{~V}, \mathrm{~W})
$$

where N is the length of the real arrays V , which contains the vector, and W , which contains the weights. Upon return from VNORM, D contains the weighted rmsnorm

$$
\sqrt{\frac{1}{\mathrm{~N}} \sum_{i=1}^{\mathrm{N}}\left(\frac{\mathrm{~V}_{i}}{\mathrm{~W}_{i}}\right)^{2}}
$$

This routine is used by STODE to compute the weighted rms norm of the estimated local error. STODE also uses information returned by VNORM to perform the corrector convergence test and to compute factors that determine if the method order should be changed. Other routines that access VNORM are LSODE, to compute the initial step size H0, and PREPJ, to compute the increments in the solution vector for generating difference quotient Jacobians (MITER $=2$ or 5, table 3.2).

If the user replaces the routine VNORM, the new version must return a positive quantity in VNORM, suitable for use in local error and convergence testing. The weight array $W$ can be used as needed, but it must not be altered in VNORM. For example, the max-norm, that is, max $V_{i} / W_{i} l$, satisfies this requirement, as does a norm that ignores some components of V . The latter procedure has the effect of suppressing error control on the corresponding components of $Y$.

## 4. Description of Code Usage

### 4.10 Overlay Situation

If LSODE is to be used in an overlay situation, the user must declare the variables in the call sequence to LSODE and in the two internal common blocks LS0001 and EH0001 in the MAIN program to ensure that their contents are preserved. The common block LS0001 is of length 257 ( 218 real or doubleprecision words followed by 39 integer words), and EH0001 contains two integer words (see table 3.6).

### 4.11 Troubleshooting

In this section we present a brief discussion of the corrective actions that may be taken in case of difficulty with the code. If the execution is terminated prematurely, the user should examine the error message and the value of ISTATE returned by LSODE (table 4.4). We therefore recommend that the current value of MESFLG not be changed, at least until the user has gained some experience with the code. The legality of every input parameter, both required and optional, is checked. If illegal input is detected by the code, it returns to the calling subprogram with ISTATE $=-3$. The error message will be detailed and will make clear what corrective actions to take. If the illegal input is caused by a request for too much accuracy, the user should examine the value of TOLSF returned in RWORK(13) (table 4.7) and make necessary adjustments to RTOL and ATOL, as described in section 4.5.7. If an excessive accuracy requirement is detected during the course of solving the problem, the value ISTATE $=-2$ is returned. To continue the integration, make the adjustments mentioned above, set $\operatorname{ISTATE}=3$, and call LSODE again.

Another difficulty related to accuracy control may be encountered if pure relative error control for, say, the $i$ th variable is specified (i.e., $\mathrm{ATOL}_{i}=0$ ). If this solution component vanishes, the error test cannot be applied. In this situation the value ISTATE $=-6$ is returned to the calling subprogram. The error message identifies the component causing the difficulty. To continue integrating, reset ATOL for this component to a nonzero value, set ISTATE = 3, and call LSODE again.

If more than MXSTEP (default value, 500) integration steps are taken on a single call to LSODE without completing the task, the error return ISTATE $=-1$ is made. The problem might be the use of an inappropriate integration method or iteration technique. The use of $\mathrm{MF}=10$ (or 20) on a stiff problem is one example. The user should, as described previously under the selection of MF (section 4.5.8), verify that the value of MF is right for the problem. Very stringent accuracy requirements may also cause this difficulty. Another possibility is that pure relative error control has been specified but most, or all, of the $\left|Y_{l}\right|$ are very small but nonzero. Finally, the solution may be varying very rapidly, forcing the integrator to select very small step sizes, or the integration interval may be very
long relative to the average step size. To continue the integration, simply reset ISTATE $=2$ and call LSODE again-the excess step counter will be reset to zero. To prevent a recurrence of the error, the value of MXSTEP can be increased, as described in section 4.6. If this action is taken between calls to LSODE, ISTATE must be set equal to 3 before LSODE is called again. Irrespective of when MXSTEP is increased, IOPT should be set equal to 1 before the next call to LSODE.

If the integrator encounters either repeated local error test failures or any local error test failure with a step size equal to the user-supplied minimum value HMIN (table 4.6), LSODE returns with ISTATE $=-4$. The difficulty could be caused by a singularity in the problem or by inappropriate input. The user should check subroutines $F$ and JAC for errors. If none is found, it may be necessary to monitor intermediate quantities. The component IMXER causing the error test failure is returned as IWORK(16) (table 4.7). The values Y(IMXER), RTOL(IMXER), ATOL(IMXER), and ACOR(IMXER) (see table 4.8) should be examined. If pure relative error control had been specified for this component, very small but nonzero values of Y(IMXER) may cause the difficulty.

These checks should also be made if the integration fails because of either repeated corrector convergence test failures or any such failure with a step size equal to HMIN. In this case LSODE returns the value ISTATE $=-5$ along with a value for IMXER defined above. If an analytical Jacobian is being used, it should be checked for errors. The accuracy of the calculation can also be checked by comparing $\mathbf{J}$ with that generated internally. Another reason for this failure may be the use of an inappropriate MITER, for example, MITER $=3$ for a problem that does not have a diagonally dominant Jacobian. It may be helpful to try different values for MITER and monitor the successive corrector estimates stored as the $Y$ array in subroutine STODE.

In addition to the error messages just discussed, a warning message is printed if the step size H becomes so small that $\mathrm{T}+\mathrm{H}=\mathrm{T}$ on the computer, where T is the current value of the independent variable. This error is not considered fatal, and so the execution is not terminated nor is a return made to the calling subprogram. No action is required by the user. The warning message is printed a maximum number of MXHNIL (default value, 10) times per problem. The user can change the number of times the message is printed by resetting MXHNIL, as discussed in section 4.6. To indicate the change to LSODE, the parameter IOPT must be set equal to 1 before LSODE is called.

## Chapter 5

## Example Problem

### 5.1 Description of Problem

In this chapter we demonstrate the use of the code by means of a simple stiff problem taken from chemical kinetics. The test case, described elsewhere (refs. 17,28 , and 38), consists of three chemical species participating in three irreversible chemical reactions at constant density and constant temperature:

$$
\begin{gather*}
S_{1} \xrightarrow{k_{1}} S_{2},  \tag{5.1}\\
S_{2}+S_{3} \xrightarrow{k_{2}} S_{1}+S_{3},  \tag{5.2}\\
S_{2}+S_{2} \xrightarrow{k_{3}} S_{3}+S_{3}, \tag{5.3}
\end{gather*}
$$

with $k_{1}=4 \times 10^{-2}, k_{2}=10^{4}$, and $k_{3}=1.5 \times 10^{7}$. In reactions (5.1) to (5.3), $s_{i}$ is the chemical symbol for the $i$ th species, the arrows denote the directions of the reactions (the single arrow for each reaction means that it takes place in the indicated direction only), and the $\left\{k_{j}\right\}$ are the specific rate coefficients for the reactions. The units of $k_{j}$ depend on reaction type (e.g., ref. 39). If $y_{i}$ denotes the molar concentration of species $i$, that is, moles of species i per unit volume of mixture, the governing ODE's are given by

$$
\begin{gather*}
\frac{d y_{1}}{d t}=-0.04 y_{1}+10^{4} y_{2} y_{3},  \tag{5.4}\\
\frac{d y_{2}}{d t}=0.04 y_{1}-10^{4} y_{2} y_{3}-3 \times 10^{7} y_{2} y_{2}, \tag{5.5}
\end{gather*}
$$

## 5. Example Problem

$$
\begin{equation*}
\frac{d y_{3}}{d t}=3 \times 10^{7} y_{2} y_{2}, \tag{5.6}
\end{equation*}
$$

where $t$ is time in seconds. The initial conditions are

$$
\begin{equation*}
y_{1}(t=0)=1 ; \quad y_{2}(t=0)=y_{3}(t=0)=0 . \tag{5.7}
\end{equation*}
$$

The example problem is interesting because the reaction rate coefficients vary over nine orders of magnitude. Also it can be quite easily verified that at steady state, that is, as $t \rightarrow \infty, y_{1} \rightarrow 0, y_{2} \rightarrow 0$, and $y_{3} \rightarrow 1$. To study the evolution of the chemical system, including the approach to the final state, we integrate the ODE's up to $t=4 \times 10^{10} \mathrm{~s}$, generating output at $t=0.4 \times 10^{n} \mathrm{~s}(n=0,1, \ldots, 11)$.

### 5.2 Coding Required To Use LSODE

### 5.2.1 General

All of the coding required to solve the example problem with LSODE is included (in the form of comment statements) in the package supplied to the user. The MAIN program that calls LSODE and manages output is given in figure 5.1. Figure 5.2 lists the subroutine that computes the derivatives. Because a value of MITER $=1$ is used (fig. 5.1), a routine that computes the analytical Jacobian matrix is required. This routine is given in figure 5.3. The names used for the derivative and Jacobian matrix subroutines are, respectively, FEX and JEX. Therefore these names are used as arguments in the call to LSODE and declared EXTERNAL in the MAIN program (fig. 5.1).

### 5.2.2 Selection of Parameters

Because the problem is stiff, the choice METH $=2$ is made. For the same reason functional iteration, that is, MITER $=0$, is rejected. It is straightforward to compute the analytical Jacobian matrix, which should be used for reasons of efficiency. In any case, the choice MITER $=3$, that is, Jacobi-Newton iteration, must not be made because the Jacobian matrix is not diagonally dominant. The choice MITER $=4$ with $M L=1$ and $M U=2$ could be made but will require more storage than MITER $=1$ (see table 4.9). More importantly the computational overhead for the LU-decomposition of the iteration matrix is more for MITER $=4$ than for $\mathrm{MITER}=1$. Hence the value $\mathrm{MF}=21$ is used.

The number NEQ of ODE's is equal to the number $(=3)$ of chemical species. To minimize storage, the lengths LRW and LIW of the work arrays RWORK and IWORK are set equal to their minimum required values. According to the formulas given in tables 4.9 and 4.10 for $\mathrm{MF}=21$, these lengths are as follows:

```
    EXTERNAL FEX, JEX
    DOUBLE PRECISION ATOL, RWORK, RTOL, T, TOUT, Y
    DIMENSION Y(3), ATOL(3), RWORK(58), IWORK(23)
    NEQ \(=3\)
    \(Y(1)=1.00\)
    \(Y(2)=0.00\)
    \(Y(3)=0.00\)
    \(T=0.00\)
    TOUT \(=.400\)
    ITOL \(=2\)
    RTOL \(=1 . \mathrm{D}-4\)
    ATOL(1) \(=1.0-6\)
    ATOL \((2)=1 . D-10\)
    ATOL(3) \(=1 . D-6\)
    ITASK \(=1\)
    ISTATE = 1
    IOPT \(=0\)
    \(L R M=58\)
    LIH \(=23\)
    \(\mathrm{MF}=21\)
    DO 40 IOUT = 1,12
        CALL LSODE (FEX, NEQ, Y, T, TOUT, ITOL, RTOL, ATOL, ITASK, ISTATE,
    1 IOPT, RWORK, LRW, IHORK, LIM, JEX, MF)
        \(\operatorname{WRITE}(3,20) T, Y(1), Y(2), Y(3)\)
        FORAMT ( 7 H AT \(T=, E 12.4,6 \mathrm{H} \quad Y=, 3 E 15.7\) )
        IF (ISTATE .LT. 0) 60 TO 80
        TOUT \(=\) TOUT \(* 10 . D 0\)
    WRITE (3, 60) IWORK (11), IMORK (12) , IMORX (13)
60 FORNAT (/12H NO. STEPS \(=14,11 H\) NO. F-S \(=14,11 H\) NO. J-S \(=, 14\) )
STOP
\(80 \operatorname{WRITE}(3,90)\) ISTATE
90 FORNAT ( \(/ / / 22 H\) ERROR HALT. ISTATE \(=13\) )
STOP
END
```

Figure 5.1.-Listing of MAIN program for example problem.

```
SUBROUTINE FEX (NEQ, T, Y, YDOT)
DOUBLE PRECISION T, Y, YDOT
DIMEHSION Y(3), YDOT(3)
YDOT (1) = -.04DO*Y(1) + 1.D4*Y (2)*Y (3)
YDOT (3) = 3.D7*Y(2)*Y(2)
YDOT(2) = -YDOT(1) - YDOT(3)
RETURN
END
```

Figure 5.2.-Listing of subroutine (FEX) that computes derivatives for example problem.

$$
\mathrm{LRW}=22+3(5+1)+3(3)+3^{2}=58
$$

and

$$
\operatorname{LIW}=20+3=23 .
$$

Selection of the error tolerances requires some explanation. A scalar RTOL is used because the same number of significant figures is acceptable for all components. However, because $y_{2}$ is expected to be much smaller than both $y_{1}$

```
SUBROUTINE JEX (NEQ, T, Y, ML, MU, PD, NRPD)
DOUBLE PRECISION PD, \(T, Y\)
DIMENSION Y(3), PD(NRPD,3)
\(\operatorname{PD}(1,1)=-.04 \mathrm{DO}\)
\(P D(1,2)=1 . D 4 \star Y(3)\)
\(\mathrm{PD}(1,3)=1.04 * Y(2)\)
\(P D(2,1)=.0400\)
\(\mathrm{PD}(2,3)=-\mathrm{PD}(1,3)\)
\(\mathrm{PD}(3,2)=6 . \mathrm{DF}^{7 * Y}(2)\)
\(P D(2,2)=-P D(1,2)-P D(3,2)\)
RETURN
END
```

Figure 5.3.-Listing of subroutine (JEX) that computes analytical Jacobian matrix for example problem.
and $y_{3}$, an array ATOL, with ATOL(2) much smaller than both ATOL(1) and ATOL(3), is used. For these choices of the RTOL and ATOL types, table 4.2 gives $\mathrm{ITOL}=2$. Pure relative error control cannot be used because the initial values of both $y_{2}$ and $y_{3}$ are zero and, as $t \rightarrow \infty, y_{1} \rightarrow 0$ and $y_{2} \rightarrow 0$. Pure absolute error control should not be used because of the widely varying orders of magnitude of the $\left\{y_{l}\right\}$. Note that because a scalar RTOL is used, the MAIN program does not require a DIMENSION statement for this variable.

The remainder of the program calling LSODE is straightforward and selfexplanatory. Because the output value for ISTATE is equal to 2 for a normal return from LSODE and no parameter (except TOUT) is reset between calls to LSODE, ISTATE does not have to be reset.

### 5.3 Computed Results

The output from the program, obtained on the Lawrence Livermore Laboratory's CDC-7600 computer using single-precision arithmetic, is given in figure 5.4. In addition to the results at the specified times, values for the following parameters, which give a measure of the computational work required to solve the problem, are printed at the end: total number of integration steps (STEPS), total number of derivative evaluations ( $\mathrm{F}-\mathrm{S}$ ), and total number of Jacobian matrix evaluations and LU-decompositions of the iteration matrix ( $\mathrm{J}-\mathrm{S}$ ).

| AT | 4.0000E-01 | $=$ | $9.851726 \mathrm{E}-01$ | 3.386406E-05 | 1.49357E-02 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| AT T | $4.0000 \mathrm{E}+00$ | $Y=$ | $9.055142 \mathrm{E}-01$ | 2.240418E-05 | 9.446344E-02 |
| AT | $4.0000 \mathrm{E}+01$ |  | $7.158050 \mathrm{E}-01$ | 9.184616E-06 | 2.841858E-01 |
| AT $T=$ | $4.0000 \mathrm{E}+02$ |  | $4.504846 \mathrm{E}-01$ | 3.222434E-06 | 5.495122E-01 |
| AT $T=$ | $4.0000 \mathrm{E}+03$ |  | $1.831701 \mathrm{E}-01$ | 8.940379E-07 | 8.168290E-01 |
| AT $T=$ | $4.0000 \mathrm{E}+04$ | Y | 3.897016E-02 | 1.621193E-07 | 9.610297E-01 |
| AT | $4.0000 \mathrm{E}+05$ |  | $4.935213 \mathrm{E}-03$ | 1.983756E-08 | 9.950648E-01 |
| AT T $=$ | $4.0000 \mathrm{E}+06$ |  | 5.159269E-04 | 2.064759E-09 | 9.994841E-01 |
|  | $4.0000 \mathrm{E}+07$ |  | $5.306413 \mathrm{E}-05$ | $2.122677 \mathrm{E}-10$ | 9.999469E-01 |
| AT T $=$ | $4.0000 \mathrm{E}+08$ | $Y=$ | 5.494529E-06 | $2.197824 \mathrm{E}-11$ | 9.999945E-01 |
| AT T | $4.0000 \mathrm{E}+09$ |  | $5.129458 \mathrm{E}-07$ | $2.051784 \mathrm{E}-12$ | 9.999995E-01 |
| AT | $4.0000 \mathrm{E}+10$ |  | -7.170592E-08 | -2.868236E-13 | 1.000000E+00 |

Figure 5.4.-Output from program for example problem.

## Chapter 6 Code Availability

The present version of LSODE, dated March 30, 1987, is available in single or double precision. The code has been successfully executed on the following computer systems: Lawrence Livermore Laboratory's CDC-7600, Cray-1, and Cray-X/MP; NASA Lewis Research Center's IBM 370/3033 using the TSS operating sytem (OS), Amdahl 5870 using the VM/CMS OS and the UTS OS, Cray-X/MP/2/4 using the COS and UNICOS operating sytems and the CFT and CFT77 compilers, Cray-Y/MP/8/6128 using UNICOS 6.0 and CFT77, Alliant FX/S, Convex C220 minicomputer using the Convex 8.0 OS, and VAX 11/750, 11/780, 11/785, 6320, 6520, 8650, 8800, and 9410; NASAAmes Research Center's Cray- 2 and Cray-Y/MP using the UNICOS operating system and the CFT77 compiler; the Sun SPARCstation 1 using the Sun 4.0 OS; the IBM RISC System/ 6000 using the AIX 3.1 OS and the XLF and F77 compilers; several IRIS workstations using the IRIX 4.0.1 OS and F77 compiler; and various personal computers under various systems.

The LSODE package is one of five solvers included in the ODEPACK collection of software for ordinary differential equations (ref. 2). The official distribution center for ODEPACK is the Energy Science and Technology Software Center at Oak Ridge, Tennessee. (ESTSC supersedes NESC, the National Energy Software Center at Argonne National Laboratory, in this activity.) Both single- and doubleprecision versions of the collection are available. Additional details regarding code availability and procurement can be obtained from

> Energy Science and Technology Software Center
> P.O. Box 1020

> Oak Ridge, TN 37831-1020
> Telephone: (615) 576-2606

The ODEPACK solvers can also be obtained through electronic mail by accessing the NETLIB collection of mathematical software (ref. 40). Both single- and double-precision versions of ODEPACK are contained in NETLIB. Detailed instructions on how to access and use NETLIB are given by Dongarra and Grosse (ref. 40).

## References

1. Hindmarsh, A.C.: LSODE and LSODI, Two New Initial Value Ordinary Differential Equation Solvers. ACM SIGNUM Newsletter, vol. 15, no. 4, 1980, pp. 10-11.
2. Hindmarsh, A.C.: ODEPACK: A Systematized Collection of ODE Solvers. Scientific Computing, R.S. Stepleman, et al., eds., North Holland Publishing Co., New York, 1983, pp. 55-64.
3. Shampine, L.F.; and Gordon, M.K.: Computer Solution of Ordinary Differential Equations: The Initial Value Problem. W.H. Freeman and Co., San Francisco, CA, 1975.
4. Lambert, J.D.: Computational Methods in Ordinary Differential Equations. Wiley \& Sons, New York, 1973.
5. Forsythe, G.E.; and Moler, C.B.: Computer Solution of Linear Algebraic Systems. Prentice-Hall, Englewood Cliffs, NJ, 1967.
6. Shampine, L.F.: What is Stiffness? Stiff Computation, R.C. Aiken, ed., Oxford University Press, New York, 1985, pp. 1-16.
7. Shampine, L.F.; and Gear, C.W.: A User's View of Solving Stiff Ordinary Differential Equations. SIAM Rev., vol. 21, 1979, pp. 1-17.
8. Radhakrishnan, K.: A Comparison of the Efficiency of Numerical Methods for Integrating Chemical Kinetic Rate Equations. Computational Methods, K.L. Strange, ed., Chemical Propulsion Information Agency Publication 401, Johns Hopkins Applied Physics Laboratory, Laurel, MD, 1984, pp. 69-82. (Also NASA TM-83590, 1984.)
9. Radhakrishnan, K.: New Integration Techniques for Chemical Kinetic Rate Equations. I. Efficiency Comparison. Combust. Sci. Technol., vol. 46, 1986, pp. 59-81.
10. Gear, C.W.: Numerical Initial Value Problems in Ordinary Differential Equations. Prentice-Hall, Englewood Cliffs, NJ, 1971.
11. Shampine, L.F.: Stiffness and the Automatic Selection of ODE Codes. J. Comput. Phys., vol. 54, 1984, pp. 74-86.
12. May, R.; and Noye, J.: The Numerical Solution of Ordinary Differential Equations: Initial Value Problems. Computational Techniques for Differential Equations, J. Noye, ed., North-Holland, New York, 1984, pp. 1-94.
13. Forsythe, G.E.; Malcolm, M.A.; and Moler, C.B.: Computer Methods for Mathematical Computations, Prentice-Hall, Englewood Cliffs, NJ, 1977.
14. Hull, T.E., et al.: Comparing Numerical Methods for Ordinary Differential Equations. SIAM J. Numer. Anal., vol. 9, no. 4, 1972, pp. 603-637.
15. Hindmarsh, A.C.: GEAR: Ordinary Differential Equation System Solver. Report UCID-30001, Rev. 3, Lawrence Livermore Laboratory, Livermore, CA, 1974.
16. Gear, C.W.: Algorithm 407, DIFSUB for Solution of Ordinary Differential Equations. Comm. ACM, vol. 14, no. 3, 1971, pp. 185-190.
17. Byrne, G.D.; and Hindmarsh, A.C.: Stiff ODE Solvers: A Review of Current and Coming Attractions. J. Comput. Phys., vol. 70, no. 1, 1987, pp. 1-62.
18. Gear, C.W.: The Numerical Integration of Ordinary Differential Equations. Math. Comput., vol. 21, 1967, pp. 146-156.
19. Gear, C.W.: The Automatic Integration of Stiff Ordinary Differential Equations. Information Processing, A.J.H. Morrell, ed., North-Holland Publishing Co., New York, 1969, pp. 187-193.
20. Gear, C.W.: The Automatic Integration of Ordinary Differential Equations. Comm. ACM, vol. 14, no. 3, 1971, pp. 176-179.
21. Hindmarsh, A.C.: Linear Multistep Methods for Ordinary Differential Equations: Method Formulations, Stability, and the Methods of Nordsieck and Gear. Report UCRL-51186, Rev. 1, Lawrence Livermore Laboratory, Livermore, CA., 1972.
22. Hindmarsh, A.C.: Construction of Mathematical Software. Part III: The Control of Error in the GEAR Package for Ordinary Differential Equations. Report UCID-30050, Pt. 3, Lawrence Livermore Laboratory, Livermore, CA, 1972.

## References

23. Byrne, G.D.; and Hindmarsh, A.C.: A Polyalgorithm for the Numencal Solution of Ordinary Differential Equations ACM Trans. Math. Software, vol. 1, no. 1, 1975, pp. 71-96.
24. Brown, P.N.; Byme, G.D.; and Hindmarsh, A.C: VODE, A Variable-Coefficient ODE Solver. SIAM J. Sci. Stat. Comput., vol. 10, 1989, pp. 1038-1051.
25. Shampine, L.F.: Implementation of Implicit Formulas for the Solution of ODEs. SIAM J. Sci Stat. Comput., vol. 1, no. 1, 1980, pp. 103-118.
26. Hall, G.; and Watt, J.M , eds.: Modern Numerical Methods for Ordinary Differentıal Equations. Clarendon Press, Oxford, U.K., 1976.
27. Finlayson, B.A.: Nonlinear Analysis in Chemical Engineering. McGraw Hill, New York, 1980.
28. Lapidus, L.; and Seinfeld, J.H.: Numerical Solution of Ordinary Differential Equations. Academic Press, New York, 1971
29. Henrici, P.: Discrete Variable Methods in Ordinary Differentıal Equations. Wiley, New York, 1962.
30. Shampine, L.F.: Type-Insensitive ODE Codes Based on Implicit A-Stable Formulas. Math. Comput., vol. 36, no. 154, 1981, pp. 499-510.
31. Ortega, J.; and Rheinbolt, W.C : Iterative Solution of Nonlinear Equations in Several Variables. Academic Press, New York, 1970
32. Shampine, L.F.: Type-Insensitive ODE Codes Based on Implicit A( $\alpha$ )-Stable Formulas Math. Comput., vol. 39, no. 159, 1982, pp. 109-123.
33. Nordsieck, A.: On Numerical Integration of Ordinary Differential Equations. Math. Comput., vol 16, 1962, pp. 22-49.
34. Dongarra, J.J , et al.: LINPACK User's Guide. SIAM, Philadelphia, 1979
35. Jones, R.E.: SLATEC Common Mathematical Library Error Handling Package. Report SAND781189, Sandia National Laboratories, Albuquerque, NM, 1978.
36. Strang, G.: Linear Algebra and Its Applications. Second ed., Academic Press, New York, 1980.

37 Radhakrishnan, K.: LSENS-A General Chemical Kinetics and Sensitivity Analysis Code for Homogeneous Gas-Phase Reactions. I Theory and Numerical Solution Procedures NASA RP-1328, 1994.
38. Robertson, H.H.: The Solution of a Set of Reaction Rate Equations Numerical Analysis, An Introduction, J. Walsh, ed., Academic Press, New York, 1966, pp. 178-182.
39. Benson, S.W.: The Foundations of Chemical Kinetics. Robert E Krieger Publishing Co., Malabar, FL, 1982.
40. Dongarra, J J.; and Grosse, E.: Distribution of Mathematıcal Software via Electronic Mail Comm. ACM, vol. 30, no 5, 1987, pp. 403-407.

| REPORT DOCUMENTATION PAGE |  |  | Form Approved OMB No. 0704-0188 |
| :---: | :---: | :---: | :---: |
|  <br>  <br>  |  |  |  |
|  | AGENCY USE ONLY (Leave blank) | 2. REPORT DATE <br> December 1993$\|$3. R | D DATES COVERED Reference Publication |
| 4 title and subtite <br> Descripton and Use of LSODE, the Livermore Solver for Ordinary Differential Equations |  |  | 5. FUNDING NUMBERS |
| 6. AUTHOA(S) <br> Knshnan Radhakrıshnan and Alan C. Hindmarsh |  |  | WU-505-62-52 |
|  | PERFORMing organization nak <br> Natonal Aeronautics and Spa Lewis Research Center Cleveland, Ohio 44135-319 | S) AND ADDAESS(ES) <br> Administration | 8. PERFORMING ORGANIZATION REPORT NUMBER <br> E-5843 |
|  | SPONSORINGMONTORING AGENCY <br> National Aeronautics and Spa <br> Washington, D.C. 20546-00 | NAME(S) AND ADDRESS(ES) <br> Administration | 10. SPONSORING/MONITORING AGENCY REPORT NUMBER <br> NASA RP-1327 <br> UCRL-ID-1 13855 |
| 11. SUPPLEMENTARY NOTES <br> Krishnan Radhaknshnan, Sverdrup Technology, Inc., Lewis Research Center Group, 2001 Aerospace Parkway, Brook Park, Ohio 44142 (work funded by NASA Contract NAS3-25266), and Alan C Hindmarsh, Lawrence Livermore National Laboratory. Livermore, CA 94551 (work funded by DOE Contract W-7405-ENG-48) Responsible person, Edward J Mularz, organization code 2650, (216) 433-5850. Repnnted in October 1994 incorporating correcuons noted on the errata |  |  |  |
| 12a DISTRIBUTIONAVAILABILITY STATEMENT <br> Unclassified-Unlımıted <br> Subject Categories 61 and 64 |  |  | 12b. DISTRIBUTION CODE |
| 13. ABSTRACT (Maximum 200 words) <br> LSODE, the Livermore Solver for Ordinary Differential Equations, is a package of FORTRAN subroutnes designed for the numerical solution of the inttual value problem for a system of ordinary differential equations $I t$ is particularly well suited for "suff" differential systems, for which the backward differentiation formula method of orders 1 to 5 is provided. The code includes the Adams-Moulton method of orders 1 to 12; so it can be used for nonsuff problems as well. In addition, the user can easily switch methods to increase computational efficiency for problems that change character For both methods a vanety of corrector iteration techniques is included in the code. Also, to minimize computational work, both the step size and method order are vaned dynamically This report presents complete descriptons of the code and integration methods, including their implementation It also provides a detailed gude to the use of the code, as well as an illustrative example problem. |  |  |  |


| First-order ordinary differential equations; Stiff ODE's, Linear multistep methods, Adams-Moulton method, Backward differentiation formula method; Smple iteration, Newton-Raphson iteration, Numerical Jacobians: Accuracy; Error control; Mehhod order selection; Step size selection |  |  | 15. NUMBER OF PAGES 122 |
| :---: | :---: | :---: | :---: |
|  |  |  | 16. PRICE CODE A06 |
| 17. SECURITY CLASSIFICATKN OF REPORT Unclassified | 18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified | 19. SECURITY CLASSIFICATION of Abstract Unclassified | 20. LIMITATION OF ABSTRACT |

# National Aeronautics and Space Administration <br> Lewis Research Center 21000 Brookpark Rd. Cleveland, OH 44135-3191 

Official Business
Penalty for Private Use $\$ 300$
POSTMASTER: If Undeliverable - Do Not Return


[^0]:    ${ }^{1}$ Although the corrector convergence test is performed before the local truncation error test (which is done only if the iteration converges), we discuss the accuracy test first because the convergence test is based on it.
    ${ }^{2}$ As discussed in chapter 1, another commonly used definition for the local truncation error is that it is the error incurred by the numerical method in advancing the approximate solution by a single step assuming exact past values and no roundoff errors (refs. 12, 13, and 21 ). That is, $\underline{d}_{n}$ is the difference between the numerical approximation $\underline{Y}_{n}^{*}$ obtained by using exact past values (i.e., $\left\{\underline{y}\left(\xi_{n-j}\right)\right\}$ and $\left.\left\{\underline{\dot{y}}\left(\xi_{n-j}\right)\right\}\right)$ and the exact solution $\underline{y}\left(\xi_{n}\right)$ :

    $$
    \begin{equation*}
    \underline{\mathrm{d}}_{n}=\underline{\mathrm{Y}}_{n}^{*}-\underline{\mathrm{y}}\left(\xi_{n}\right), \tag{2.77}
    \end{equation*}
    $$

    where, for example,

    $$
    \begin{equation*}
    \underline{\mathrm{Y}}_{n}^{*}=\sum_{j=1}^{q} \alpha_{J} \underline{\mathrm{y}}\left(\xi_{n-J}\right)+h_{n} \beta_{0} \mathrm{f}\left(\underline{\mathrm{Y}}_{n}^{*}\right) \tag{2.78}
    \end{equation*}
    $$

    for the BDF method of order $q$. For an explicit method the local truncation error given by equation (2.77) and that obtained by using the definition given in the text above (i.e., the residual of eq. (2.1)) have the same magnitude. However, for an implicit method the two quantities are only approximately proportional to one another (ref. 4), although they agree asymptotically in the limit of small step size.

[^1]:    ${ }^{\text {a }}$ Note that some variables appear in the table before they are defined.
    ${ }^{6}$ Default value for this variable can be changed by the user, as described in table 4.6.

[^2]:    ${ }^{\circ}$ The length LENWM of the array WM depends on the iteration technique and is given in table 3.9.

[^3]:    ${ }^{\mathrm{b}}$ Default value for this variable can be changed by the user, as described in table 4.6.

[^4]:    ${ }^{\text {b }}$ Default value for this variable can be changed by the user, as described in table 4.6.

[^5]:    ${ }^{2}$ See table 3.2 for description of MITER.
    ${ }^{\mathrm{b}} \mathrm{N}$ is the number of ODE's and ML and $M U$ are defined in table 3.2.

[^6]:    ${ }^{2}$ See table 4.6.
    ${ }^{\mathrm{b}}$ See chapter 2.

[^7]:    ${ }^{\text {a }}$ This value is ignored on the first step and on the final step to reach TCRIT when ITASK $=4$ or 5 (see table 4.1).

