

The PORT Mathematical Subroutine Library

P. A. FOX, A. D. HALL, and N. L. SCHRYER

Bell Laboratories

The development at Bell Laboratories of PORT, a library of portable Fortran programs for numerical computation, is discussed. Portability is achieved by careful language specification, together with the key technique of specifying computer classes by means of predefined machine constants. The library is built around an automatic error-handling facility and a dynamic storage allocation scheme, both of which are implemented portably. These, together with the modular structure of the library, lead to simplified calling sequences and ease of use.

Key Words and Phrases: libraries, portability, error handling, dynamic storage allocation, mathematical software, numerical analysis

CR Categories: 5.1

The Algorithm Framework for a Portable Library. *ACM Trans. Math. Software* 4, 2 (June 1978), 177-188

1. INTRODUCTION

We have celebrated the 25th anniversary of computer program libraries by producing another. The library is called PORT. Our resolve to create a portable library can best be explained by sketching a brief history of mathematical subroutine libraries. A look at the tremendous effort that went into the early machine-dependent libraries, each one ultimately thrown on the scrap heap along with its passé computer, and a glance at more recent developments involving either maintenance or generation of several machine-dependent versions may indicate our drive to construct a single portable library. But first a word on the historical setting.

It was in 1951 that Maurice V. Wilkes, David J. Wheeler, and Stanley Gill, all of the University of Cambridge, published their book, *The Preparation of Programs for an Electronic Digital Computer*, subtitled *With special reference to the EDSAC and the use of a library of subroutines* [56]. Their library was in machine language, but in that era the need for moving the library to another computer, and the trauma involved, had not yet been experienced. The very thought of having a library was new.

Since that time many libraries have been developed. Rice [46, ch. 1] gives some historical notes including the remark that the *Communications of the Association for Computing Machinery* published 73 algorithms during 1960-1961. Also in

General permission to make fair use in teaching or research of all or part of this material is granted to individual readers and to nonprofit libraries acting for them provided that ACM's copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Association for Computing Machinery. To otherwise reprint a figure, table, other substantial excerpt, or the entire work requires specific permission as does republication, or systematic or multiple reproduction.

Authors' address: Bell Laboratories, 600 Mountain Ave., Murray Hill, NJ 07974.

© 1978 ACM 0098-3500/78/0600-0104 \$00 75

1961 IBM made the decision to enter the software field [6, p. 121], and the SSP library was launched. Initially IBM provided SSP free with the hardware, but by 1971 IBM was charging (around \$100 per month) for the new SL-Math library [29] developed in Germany for the 360/370/1130/1800 IBM computers.

Other libraries dating from the 1960's, and we mention only a few, include the Monsanto Company's subroutine library started in 1967 [15] and the Boeing library [44] developed about the same time. Also in 1967 the Harwell Atomic Energy Research Establishment, in England, was converting their subroutine library, developed in 1963 for the IBM 7030 (STRETCH) computer, to an IBM 360 version [28]. In December 1967 the Sandia Mathematical Library Project [36] was initiated, and doubtless many other library projects were underway, at that time, across the computer world.

By the early 1970's an appreciation of the magnitude of the effort required to establish libraries was beginning to be felt. CDC, rather than developing a mathematical library from scratch, purchased the high-quality Boeing library. Also in 1970 a commercial library became available when the International Mathematical and Statistical Libraries (IMSL) was incorporated [35] and produced Library 1 in Fortran for the IBM/360-370 series. Another approach to avoiding duplication of effort was taken by the NATS (National Activity to Test Software) group [9], which was established to produce quality special-purpose software packages. The first of these, the eigenvalue-eigenvector package, EISPACK [52], is already in a second release, the package for special functions, FUNPACK, has appeared [14], and MINPACK (minimization package) and LINPACK (linear equations package) are being developed.

These libraries are all written in Fortran; in other countries the language problem (Algol versus Fortran) adds to the difficulties. In England, NAG (Numerical Algorithms Group) [17] has been furnishing for some time a large library in either Algol or Fortran to a variety of computers across the universities of that country.

Throughout the development of these software packages, the effort involved in terms of manpower, time, and money, not only to develop the packages, but to adapt them to particular computers, has been a source of wonder, and the desire to have the software portable has been growing. Even the early (1967) Monsanto library emphasized machine independence [15, p. 143], and Rice [47] has emphasized the importance of making quality software transportable. In fact, Rice makes the rather striking observation that, "It has been estimated that 60 percent to 90 percent of all research and development work is a duplication of previous work, and it is easy to believe that this applies to mathematical software with perhaps an even higher percentage than 90." The development of mathematical subroutine libraries is certainly a case in point.

At Bell Laboratories, Gentleman and Traub [21], in 1967, proposed the development of a machine-independent library using a Fortran-Algol interface to alleviate the language problem. Three sets of programs were implemented in that effort: DESUB (differential equation solution), NSEVB (eigenvalues and eigenvectors of nonsymmetric matrices), and MIDAS (linear equation solution). The work described in the present paper is an independent effort, initiated in 1973. PORT does not include any of the earlier programs, but the experience gained in that work has been useful.

In Section 2 we discuss portability: the definition of the term, the several factors inhibiting portability, the solutions taken by others in some of the libraries we have discussed, and the procedures we have applied toward portability. In Section 3 we describe in more detail the structure of the PORT library—in particular the simplified calling sequences, the error monitoring and error handling, and the dynamic storage allocation—relating each topic to the corresponding approaches used in other libraries. At the end of Section 3 we give an overview of the contents of PORT, that is, the subprograms included in the current edition. To keep the size of the paper down somewhat, we defer discussion of other aspects of the library to another time. Relevant topics not covered here include algorithm selection, implementation, testing, refereeing, documentation, certification, and distribution, as well as installation procedures and maintenance.

The basic library support utilities for PORT, which are discussed in general terms in this paper, are given as ACM Algorithm 528 [19].

2. PORTABILITY

Definitions of “portability” are rampant; the view is so clouded that we might do well to adopt an entirely new term, possibly that used by Waite [54] in another connection: “mobility.” The problem is that there is a continuum of degrees of portability, from “completely” to “not at all.” Clearly an unportable Fortran program could be made even less so by recoding it in assembler language. Brown [12], recognizing this matter of degree, offered the definition, “A program or programming system is called ‘portable’ if the effort required to move it into a new environment is much less than the effort that would be required to reprogram it for the new environment.” The environment, of course, includes the computer, the compiler, the operating system, and the particular computer hardware-software configuration.

Aird et al. [2] define four distinct concepts to span the high end of the spectrum: (1) portable, (2) converter portable, (3) processor portable, and (4) transportable. They call a program “portable, across a set of computer-compiler environments if, without any modification, it can be compiled and executed, according to defined performance criteria, for every member of the set.” The second and third definitions cover the cases where a “master version” of the library is established and a processor is provided that can particularize the program to a particular member of the set of computer-compiler environments. The term “converter portable” is distinguished from “processor portable” by indicating that the unprocessed version runs, without alteration, in at least one environment. Finally they term a program “transportable” if enough information is available to guide a required particularization, perhaps even by hand. (Note that this is an imprecise paraphrase of their very careful definitions.)

The IFIP Working Group (on Numerical Software) (WG 2.5) has proposed, in a working paper draft, the following definitions [18].

Portable. A program is portable over a given range of machines and compilers if, *without any alteration*, it can compile and run to satisfy specified performance criteria on that range.

Transportable. In transferring a program between members of a given range

of machines and compilers, some changes may be necessary to the base version before it satisfies specified performance criteria on each of the machines and compilers. The program is transportable if (1) the changes lend themselves to mechanical implementation by a processor, and (2) the changes, ideally, are limited in number, extent, and complexity.

In PORT we have so far taken the view that the library subprograms must be portable in the IFIP sense, with the exception that the three machine-constant defining functions must be particularized to the host computer once, at installation. Section 2.2 discusses the details of the approach.

What gets in the way of portability? A routine that runs well on one computer may run badly or not at all on another. Its failure may be due to language and compiler differences, to differing word structures, to variations in arithmetic (both static number representation and dynamic performance), or to differences in operating systems. If a library of mathematical subroutines aspires to portability, each of these problems must be solved.

2.1 Overview of Various Libraries and Their Approach to Portability

There is general agreement among the developers of mathematical libraries on one major point: only one source should be maintained. But there is a wide divergence of opinion as to what the source should look like, and particularly how much coded information it should contain. Some would urge that the master source should work in at least one environment without change, that is, be converter portable in terms of the definition cited above. The developers of the IMSL library take this view: the library is based on a source called the basis deck, which exists as an executable deck in one environment, but which contains control information, in the form of keyed comment cards, permitting a converter program to generate a deck for another computer-compiler environment, or to generate a double-precision version of the program [2]. The NATS project in its initial work with EISPACK used a similar approach [10].

The NAG library is based on a master library file system consisting of card-image files and three main utility programs which operate on the files [25]. The programs include editors, selectors, and extractor-comparison programs, all written in Fortran with some assembly code routines. The programming took about 18 man-months to complete. A second version to be written in Algol 68 is under study.

Krogh [37] and Krogh and Singletary [38] have developed a method, called the specializer language, for maintaining a composite source. The approach, which is rather like the IMSL approach, allows code to be specified for different machine environments and for different base precisions.

A few library projects are beginning to take the view that the master tape or composite source should be a more abstract vehicle. The NATS II system, in fact, is based on the concept of the "abstract form" of a program combined with control programs [10]. One of these, the "recognizer", can map Fortran programs written for a variety of computer systems into the abstract form. To generate particularized programs from the abstract form a "formatter" program is used.

Most of the composite or master sources are based on a system of flagged comment cards, control cards, and other types of record-based selection clues.

Some use has been made of macro capabilities to generate particularized versions, and Boyle and Dritz [10] have proposed actually parsing the Fortran input and storing the symbol table and parse tree as the master or composite tape.

In the various schemes for maintaining a master source proposed to date, it has generally been true that either the source, or the programs which generate particularized programs from the coded source, assume the existence of a finite set of specific machines; each program in the master source contains information on the attributes of the given computer environments. If this approach is used, the introduction of a new computer-compiler configuration into the scheme presents a major problem—in essence, portability has been defined only with respect to the initial set of computers.

A better way of approaching the problem, in our view, is to formulate an idealized, but robust, model of a computer from the standpoint of numerical computation. The model should be simple and yet apply to most existing (and many future) computers, and should be consonant with an ANSI Fortran computing environment. Then, given the parameters defining the model, portable software can be written from a truly machine-independent, but model-dependent, orientation. To particularize the model for any given target site, appropriate values for the parameters can be set in a simple way. The model exemplified in PORT is described in some detail in the following sections; further details on floating-point arithmetic in the model are given by Brown [13].

2.2 PORT and Portability

The techniques used in the PORT library to make it easily portable are: (1) programs are written in a subset of ANSI Fortran, and (2) the target environment is specified in terms of machine-dependent parameters. We have evidence of success to the extent that the PORT tape has been compiled and is in use on three IBM 360/370 computers, a Univac 1100, two Honeywell 6000 series, a CDC Cyber '72 system, a Harris S220, and a PDP 11.

2.2.1 Language

The programming language used for PORT subprograms is restricted to the particular, portable subset of ANSI Fortran known as PFORT, described by Ryder [49]. Programs submitted to PORT are always sent through the PFORT verifier program, described in [49], to guarantee their adherence to this language requirement.

There are two non-ANSI Fortran usages made in PORT; both are valid for the usual production system. First, it is assumed that there is no runtime subscript range checking. Second, for some of the subprograms implementing the error handling and stack allocation it is assumed that a variable (local to a subprogram) that is initialized by a DATA statement and then changed within the subprogram retains its most recently assigned value. If overlays are used by a programmer, care must be taken to avoid overlaying these few routines. These issues and others related to portability are discussed in more detail in Appendix A.

Although the ANSI Fortran standard [3] makes the assumption that LOGICAL, INTEGER, and REAL data are allocated one "storage unit" and that DOUBLE PRECISION and COMPLEX data are allocated two "storage units,"

the assumption is frequently violated in minicomputer Fortran systems. To allow for use of the library on these smaller computer systems, we have been careful, in formulating the dynamic storage scheme, to make it independent of the amounts of storage allocated to the different data types.

2.2.2 Specification of Machine-Dependent Quantities

Very early in the development of mathematical subroutine libraries, the importance of isolating the machine-dependent parameters and constants was recognized. Newbery [44, p. 155] notes that the Boeing library provided a single program whose function was to store these values in one place. EISPACK uses only two machine-dependent constants: RADIX, the base of the machine floating-point representation, and MACHEP (machine epsilon), the relative precision of the floating-point arithmetic. However, for more general libraries, it is convenient to have a number of important machine and operating system dependent constants available. Redish and Ward [45], Aird et al. [1], Krogh and Singletary [38], and others have presented lists of machine-dependent constants and parameters. The IFIP Working Group 2.5 has studied the matter in some detail and has proposed a standard set [16].

Once a set of values is decided upon, there still remains the question of getting them into the running programs. Four principal mechanisms come to mind:

- (1) Dynamically sample the host computer using subroutines to discover the base of the arithmetic, the number of digits, and so on, an approach discussed by Malcolm [40], Gentleman and Marovich [20], and George [22].

- (2) Flag or mark the machine-dependent quantities in the master source tape, enabling the correct values to be generated when a particular machine-dependent version of the library is created. This approach is used, for example, by IMSL, [2].

- (3) Use a language (possibly a Fortran preprocessor) that allows global definitions of variables and build a portable scheme based on that language's capabilities.

- (4) Develop library subprograms that can be particularized for each target computer and then called, during runtime, to obtain desired machine-dependent values. Redish and Ward [45] have proposed using this method. Their discussion is excellent. Also Ford and Sayers [17] note that the NAG II system uses a similar approach, permitting machine-dependent numbers to be evaluated by procedure calls to a sublibrary, called the constants and utilities library.

Each of these mechanisms has its problems. The original version of (1) failed on computers, such as the Honeywell 6000 or the ICL 4130, for which the floating-point registers contain more digits than a word in storage. Later versions have had troubles stemming from the fact that assumptions have to be made on computer hardware or compiler design. The techniques used in (2) are discussed above. They have the advantage that the generated version can be tailored to be especially efficient for a given computer-compiler environment (including alters to get around known bugs), but a change of compiler, or a new computer require extensive changes in the master source and the programs that control it. These requirements, together with the requirement that updates and corrections be generated in machine-dependent form before being sent out, make a sizable staff

of maintenance people necessary. Finally, the use of method (3), which may be the way of the future, means giving up (or extending) Fortran, and that is hard to do right now.

In PORT we use the fourth approach: three Fortran function subprograms are provided which can be invoked to determine basic machine or operating system dependent constants. When the library is moved to a new environment, only the DATA statements in these three subprograms need to be changed. Values are provided, in the library, for the Burroughs 5700/6700/7700, the CDC 6000/7000 series, the Data General Eclipse S/200, the DEC PDP 10 (KA and KI processors), the DEC PDP II, the Harris S220, the Honeywell 6000 series, the IBM 360/370 series, the SEL Systems 85/86, the Univac 1100 series, and the Xerox SIGMA 5/7/9; others can be added easily.

We have found this approach advantageous since the source code for the library, except for the three subprograms, remains unchanged from one environment to another. A potential drawback to the approach is the difficulty of writing portable programs for some areas of numerical computation. The field of special functions is the most trying: recursive algorithms are quite portable, but not always adequate, and rational function approximations are machine dependent. We are looking into several techniques, including program generators and other tactics, and we have detected some promising directions.

To return to our mechanism for specifying machine-dependent quantities, the three functions are: I1MACH, which delivers integer constants, R1MACH, which delivers single-precision floating-point (REAL) constants, and D1MACH, which delivers double-precision floating-point constants. The function names stem from the PORT convention that subprograms which will not be called by the casual user are given names with a digit as the second character to help avoid name conflicts. The functions have a single integer argument indicating the particular constant desired. For example, I1MACH(2) is the logical unit number of the standard output unit, so the statements

```
IWUNIT = I1MACH(2)
WRITE (IWUNIT, 9003) . . .
```

will write output (using FORMAT statement 9003) on the standard output unit. As another example, R1MACH(2) is the largest positive single-precision number, so if a program wishes to test, a priori, whether the product $x \times y$ will overflow (where $x, y > 1$), it can include the test

```
IF (Y .GE. R1MACH(2)/X) GO TO overflow
```

(The ultraprecise reader may note that the subsequent multiplication might still overflow by as much as two roundoff units, so the test should be shaded to be safe.)

If the integer argument to R1MACH or D1MACH is out of range, the error-handling facility used in PORT, which is discussed in Section 3, is called to deliver an appropriate message and terminate the run. In I1MACH, the message is output directly to avoid the possibility of a recursive call from the error handler.

The constants provided in the function subprograms cover logical unit numbers and certain properties of integer, floating-point, and character-string quantities. Care has been taken to distinguish the space configuration used by integers from

that used by single-precision (REAL) quantities; for some computers this distinction is required since the concept of a computer word for both types is not valid. (For instance, on the CDC 6000 series, both integers and reals are stored in 60-bit words, yet integers have 48 bits of magnitude and 1 sign bit.) The model of a computer we represent is based entirely on Fortran types, and the values we provide completely specify the model. In fact, some redundancy has been included for purposes discussed a little later.

The following are specified:

Logical Unit Numbers: Standard input unit, standard output unit, standard punch unit, standard error message unit.

Integer and Character Storage: Number of bits per INTEGER storage unit, number of characters per INTEGER storage unit.

Integer Variables: Let the values for integer variables be written in the s -digit, base- a form

$$\pm(x_{s-1}a^{s-1} + x_{s-2}a^{s-2} + \dots + x_1a + x_0),$$

where $0 \leq x_i < a$ for $i = 0, \dots, s-1$. Then we specify the base a , the maximum number of digits s , and the largest integer a^s-1 . Although the quantity a^s-1 can easily be computed from s and the base a , it is provided because a naive evaluation of the formula would cause overflow on most machines. (Storage of integers as magnitude and sign or in a complement notation is not specified since PORT subprograms must be independent of the storage mode.)

Floating-Point Variables: If floating-point numbers are written in the t -digit, base- b form

$$\pm b^e (x_1/b + x_2/b^2 + \dots + x_t/b^t),$$

where $0 \leq x_i < b$ for $i = 1, \dots, t$, $0 < x_1$, and $e_{\min} \leq e \leq e_{\max}$, then for a particular machine we choose values for the parameters t , e_{\min} , and e_{\max} such that all numbers expressible in this form are representable by the hardware and usable from Fortran. Note that the formula is symmetrical under negation but not reciprocation. On some machines a small portion of the range of permissible numbers may be excluded. Also, for 2's complement machines care must be taken in assigning the values; see Section 2.2.5.

Then we specify the base b for both single and double precision, and the number t of base- b digits. In order to accommodate machines (such as the CDC 6000 series) with the b -point on the right we must concede the possibility that the magnitude of e_{\min} may be substantially smaller than e_{\max} . Thus for single-precision floating-point we specify the minimum exponent e_{\min} and the maximum exponent e_{\max} .

For double precision, b remains the same, but t , e_{\min} , and e_{\max} are replaced by T , E_{\min} , and E_{\max} . Normally, we have $E_{\min} \leq e_{\min}$ and $E_{\max} \geq e_{\max}$, and $T > t$. However, in machines such as the CDC 6000 series or the PDP-10 KA processor, where double precision is implemented by software simulation, small double-precision floating-point numbers carry only t base- b significant digits. In such cases, we take E_{\min} to be the exponent of the smallest number with T base- b significant digits, and it may be that $E_{\min} > e_{\min}$.

The 16 values given above are all integers and are obtained by invoking the

function IIMACH with the appropriate argument. The floating-point single-precision and double-precision quantities provided by the functions RIMACH and DIMACH can be derived from the given integer quantities, but are provided for efficiency and convenience.

The single-precision floating-point quantities provided in RIMACH are the smallest positive magnitude $b^{e_{\min}-1}$, the largest magnitude $b^{e_{\max}}(1 - b^{-t})$, the smallest relative spacing between values b^{-t} , the largest relative spacing between values b^{1-t} , and the logarithm of the base b , $\log_{10}b$. The relative spacing is $|(y - x)/x|$, when x and y are successive floating-point numbers. Equivalent values for the double-precision floating-point quantities are provided by DIMACH, with e_{\min} , e_{\max} , and t replaced by E_{\min} , E_{\max} , and T .

2.2.3 A Note on Decimal Input-Output

In some applications, particularly input-output, it is often useful to know the basic relationships between the internal representation of numbers and an external decimal representation. Some of the simpler relationships are summarized below. More detail can be found in Matula [42].

For output, one usually wants to know how much space to allow for the decimal representation of an internal number. In the case of integers, the number s' of decimal places that are needed is given by

$$s' = \lceil s \log_{10}a \rceil,$$

where a and s are defined above, and where $\lceil x \rceil$ denotes the smallest integer not less than x .

For floating-point numbers, the situation is slightly more complex. If the external representation is of the form $m' 10^{e'}$ with $10^{-1} \leq m' < 10$, then (in single precision) the minimum and maximum values of e' are

$$\begin{aligned} e'_{\min} &= \lfloor (e_{\min}-1)\log_{10}b \rfloor + 1 \\ e'_{\max} &= \lfloor e_{\max} \log_{10}b \rfloor. \end{aligned}$$

Here, $\lfloor x \rfloor$ denotes the largest integer not exceeding x . The number of decimal places required for the decimal exponent is therefore $\lceil \log_{10}(\max(e'_{\max}, |e'_{\min}|)) \rceil$. To determine the number of decimal places to allow for m , we observe that integers in the range 0 to $b^t - 1$ can be represented exactly in single-precision floating-point numbers. If these are to be represented exactly on output, then the number t' of decimal places required is $t' = \lceil t \log_{10}b \rceil$. Relations similar to these hold for double-precision numbers. It should be noted that a decimal floating-point system carrying t' significant digits has a smallest relative spacing which is less than or equal to the smallest relative spacing of our assumed internal representation.

For input, one usually wants to know the approximate ranges of decimal numbers which can be represented in the machine. For instance, all integers of s'' decimal digits, where $s'' = \lfloor s \log_{10}a \rfloor$ can be represented internally. Of course, the actual range may be larger, but a more complicated test would be needed.

All single-precision floating-point numbers of the form $m'' 10^{e''}$, where $10^{-1} \leq m'' < 10$ and

$$\lfloor (e_{\min}-1)\log_{10}b \rfloor + 1 \leq e'' \leq \lfloor e_{\max}\log_{10}b \rfloor,$$

can be approximated in the machine. Similar relations hold for double-precision numbers.

2.2.4 Programming Using the Machine-Constant Functions

In most cases it is desirable to avoid repeated calls in a single subprogram to the functions described above. The obvious technique is to retrieve the needed values at the outset, but there are cases where substantial overhead may be incurred, even by this technique. One way to eliminate multiple calls is to use a carefully constructed “first-time” switch. For example, to retrieve I1MACH(9), the largest integer, on first entry to a subprogram, the following coding can be used:

```
DATA IMAX /0/
```

```
IF (IMAX EQ. 0) IMAX = I1MACH(9)
```

To ensure portability it is essential that all values obtained in this way be initialized in a DATA statement. If not, some operating systems (notably Burroughs) will not preserve the values from one subroutine call to the next.

2.2.5 Use of Definition Redundancy to Check Installation

Aside from convenience in programming, the redundancy in the definitions of the machine constants allows a particular installation of the PORT library to be checked for consistency. After the library has been compiled, a special checking subprogram is called to verify that the integer and floating-point constants satisfy the following conditions: (1) $t \leq T$, (2) $E_{\max} \geq e_{\max}$, (3) $E_{\min} \leq e_{\min}$, (4) the largest integer agrees with $a^s - 1$, (5) the floating-point constants agree with those computed from the integer constants, (6) the largest and smallest floating-point values are closed under negation, i.e. $-(-x) = x$.

If a discrepancy is found, a warning message and the values of the quantities involved are printed. Condition (3) may fail on some machines, even though the specifications are correct; in this case the warning message should be ignored. For 2's complement machines, if e_{\min} has been set too small, condition (6) will fail, since the negative of the smallest number will underflow. Note that some of the integer definitions must be used by the checking routine to determine the output unit for the error messages and the correct formats for printing.

3. THE PORT LIBRARY

The proof of the porting, one might say, is in the using. Libraries, to be effective, must take into account the motivations of the users. A user is grateful if the programs are easy to use and well documented; protection against errors is particularly appreciated.

To gain user acceptance, PORT provides (1) simplified calling sequences, (2) careful error handling, (3) dynamic storage allocation, and (4) brief but complete documentation, all implemented in a portable way. The first three of these four topics are discussed below.

3.1 Calling Sequences and Modular Structure

PORT is structured like an onion. The programs most visible, on the outer layer

of the library, are the simplest. The calls to these top-level routines need few parameters and are documented in brief (typically one-page) reference sheets. The top-level routines, in turn, can set default values and call lower level routines containing more parameters. A routine at the second level often is documented and available to the more sophisticated user, who may wish, for example, to influence the details of the step-size monitoring in differential equation solution. Then a second level subprogram may call on a third, perhaps undocumented, 20-parameter, subprogram.

At the innermost level, the picture simplifies again to a more primitive state. PORT includes small subprograms for complex double-precision arithmetic and for the trigonometric functions that are not ANSI Fortran; also the library probably will incorporate some version of the basic linear algebra modules (inner products, norms, etc.) proposed by Hanson et al. [27] (see Lawson [39]). Further routines include those for initializing a vector, for moving arrays and/or changing their type, for determining whether a vector is (strictly) monotone, for finding the ceiling or floor of a floating-point quantity, and for doing internal sorting. All the subprograms are of course implemented portably and can be called directly by the user or by other routines in the library.

Calling sequences in PORT are simplified in another sense, by not including parameters for error indication or for scratch storage. The centralized error-handling procedure of the library eliminates the need for error flags, while its dynamic storage allocation capability eliminates scratch storage arrays in the calls to outer level subprograms.

3.2 Centralized Error Handling

In most libraries, a program which can reach an error state includes in its calling sequence a parameter to indicate, on return from the subprogram, whether an error has occurred. The user is responsible for testing the error flag and taking the appropriate action, but, as we all know, the user frequently assumes that the program "will work this time." A safer, but more extreme approach, is to eliminate the error flag, and, if an error occurs, simply print an error message from within the subprogram and terminate the run.

The picture is really more complicated, because the proper treatment of an error may depend on the degree of severity of the error, the sophistication of the user, and other matters not known to the subprogram in which the error is detected. A recent paper by Goodenough [23] discusses the general matter of "exception handling" in a very thorough manner. Various approaches to error handling have been taken in currently available mathematical libraries: The severity of the error is taken into account in the EISPACK package which uses the sign of the error flag to distinguish "path-terminating" (fatal) errors, from those which indicate that some of the computation can be salvaged, even though trouble has occurred [52]. Some libraries (IBM [29], IMSL [31]) define in greater detail the degree of severity of an error from "warning," through "an error for which the subroutine has taken a default action," to "dangerous but nonterminating error," and finally "terminating error."

For the unsophisticated user, the safest action in all cases is to print a message and stop. The experienced user, on the other hand, usually wants to control the

error handling to some extent, and various techniques for doing so are provided in the different libraries: The NAG library [43], in the calling sequence to the error routine, provides a parameter, IFAIL, which can be set by the calling program to control the action: if IFAIL is input as 0 (hard fail), an error message is printed and execution terminated; if the input value is set to 1 (soft fail), the error routine assigns the current error number to IFAIL, now used as the output parameter, and returns to continue execution. FUNPACK [14] provides the experienced user with a mechanism for accessing error patterns in great detail. Within the library, tables are kept of the frequency of occurrence of each error; the user can monitor a particular error and allow or suppress the printing of error messages. Continuation or termination of the run, at any point, can also be controlled. The SANDIA library [36] uses a technique which is, to some extent, similar to the one used in PORT and described in the following subsection. In their library a set of error routines is provided that allow the user to override the default message printing and termination. One routine is used to set the flags controlling the printing and/or termination, and another routine can be used to access the current setting of the flags. The principle error routine, ERRCHK, does the actual error handling as specified by the current flag settings. A fourth routine is provided for one-time-only printing of an error message. The four routines communicate with each other via a COMMON block and are implemented in a machine-independent version.

3.2.1 Error Handling in PORT

In PORT, only two types of error can occur: "fatal" and "recoverable," and two types of users are catered to. For the unwary user either type of error causes an error message to be printed and the run to be terminated; in the case of a fatal error, a call is made to a dump routine. (The dump routine itself is a local option: at Bell Laboratories, Murray Hill, a symbolic dump is provided that lists the names of the variables and their values when the dump was called and prints out the list of active subprograms [26]). For the user who wishes to recover from an error and to gain control over the error-handling process, a "recovery mode" is provided. At any point in a run the user can enter the recovery mode and, while in this mode, can do any of the following: (1) determine whether an error has occurred, and, if so, obtain the error number, (2) print any current error message, (3) turn off the error state, and (4) leave the recovery mode. Only recoverable errors can be controlled by the user; the fatal errors represent unrecoverable situations or user blunders such as setting an input parameter to an impossible value.

When an error is detected in a PORT subprogram, a call is made to the error-handling routine SETERR. (Of course, a user may also call SETERR in a main program, or user-written subprogram.) The calling sequence is

```
CALL SETERR (MESSG, NMESSG, NERR, IOPT)
```

where MESSG is a Hollerith message; NMESSG is the number of characters in MESSG; NERR is the error number; and IOPT is set to 1 to specify that the error is recoverable, or to 2 to specify that the error is fatal. Unless recovery mode is in effect, SETERR prints an error message and terminates execution.

For the casual user of PORT, the possibility of regaining control after an error will probably not be of interest. The message printed out if an error occurs will usually indicate where a change or correction need be made. For the user who does wish to recover from certain errors and continue the computation, there are subprograms permitting this flexibility. Algorithm 528 [19] provides further discussion of the programs in the error-handling package.

3.3 Dynamic Storage Allocation Using a Stack

Dynamic storage allocation in programming systems (as opposed to memory management in operating systems) is most often found in list processing applications such as LISP. Fortran, which has no mechanism for recursion, is not a natural setting for dynamic storage management, although some work has been done. Barron [5] reports the work of Ayers [4] on implementing a set of stack-handling routines in Fortran. Jensen [34] describes some routines for dynamically providing portions of a storage pool to an application program. He has developed a "modular storage management system" which uses data structures in a hierarchy of records, groups of records, and groups of groups, so that the scheme is more structured than a simple stack. Other libraries of mathematical subprograms do contain sets of programs to implement dynamic storage handling, but the only library we happen to know of, besides PORT, which is completely built around a dynamic storage scheme is STATLIB [11].

The PORT library has integrated a dynamic storage allocator into the basic library structure. We consider this method for providing scratch space greatly superior to other methods; the historical approach of compiling workspace directly into individual subprograms is clearly inefficient, and the other general method of passing names of scratch arrays puts a considerable naming and dimensioning burden on the user. We have found that use of dynamic storage allocation in PORT leads to more clearly structured programs, cleaner calling sequences, improved memory utilization, and better error detection. The allocator is implemented as a package of simple portable Fortran subprograms which manipulate a dynamic storage stack.

In general, the casual PORT user need not be concerned about the operation, or even the existence of the dynamic storage stack; the fact that the PORT subprograms are using the stack is invisible. However, for strict conformance with the ANSI standard, and particularly when overlays are being used, a declaration of the stack in the main program should be included, (cf. the discussion of nonstandard usages in Appendix A.)

In the following subsection we discuss the capabilities included in PORT's storage-allocation package, and give examples of its use. Appendix B discusses the implementation of the storage stack, and Algorithm 528 [19] contains the subprograms for it.

3.3.1 The Stack: Allocation and Deallocation

Allocation and deallocation of space on the stack is carried out through the use of explicit subprogram calls in the subprograms of the PORT library. By the nature of a stack, allocations and deallocations are carried out on a last-in first-out basis. In order to make the stack invisible to most users of library programs,

the package is self-initializing and contains a default stack size which will hold approximately 500 DOUBLE PRECISION data items. If desired, larger amounts of stack space can be reserved for a particular run.

The stack resides in the labeled COMMON region CSTAK. Any subroutine that uses space allocated in the stack must include the following declarations:

```
COMMON/CSTAK/DSTAK(500)
DOUBLE PRECISION DSTAK
```

These ensure that the length and type of the stack are properly and consistently declared in all subprograms, including those which use the allocator and are loaded from libraries. Failure to use these declarations could lead to unexpected difficulties during loading (or link-editing). If needed, most Fortran environments permit a larger stack to be declared in the main program, without adjusting these other declarations to match.

To provide LOGICAL, INTEGER, REAL, and COMPLEX aliases for the stack, the following declarations appear in many PORT subprograms:

```
LOGICAL LSTAK(1000)
INTEGER ISTAK(1000)
REAL RSTAK(1000)
COMPLEX CMSTAK(500)
```

C

```
EQUIVALENCE (DSTAK(1), LSTAK(1))
EQUIVALENCE (DSTAK(1), ISTAK(1))
EQUIVALENCE (DSTAK(1), RSTAK(1))
EQUIVALENCE (DSTAK(1), CMSTAK(1))
```

The dimensions are chosen for the ANSI standard situation with LOGICAL, REAL, and INTEGER variables taking half the space of DOUBLE PRECISION or COMPLEX. If the relative lengths are nonstandard (see Section 2.2), there is one stack-management subprogram that must be modified. (See Algorithm 528 [19].)

PORT contains two basic subprograms, ISTKGT and ISTKRL, for getting and releasing stack space, respectively. The function for getting stack space is

```
INTEGER FUNCTION ISTKGT(NITEMS, ITYPE)
```

where NITEMS is the number of items of type ITYPE to be allocated. The values of ITYPE are as follows:

<i>ITYPE</i>	<i>Item Type</i>
1	LOGICAL
2	INTEGER
3	REAL
4	DOUBLE PRECISION
5	COMPLEX

For example, the statement

```
I = ISTKGT(N,2)
```

returns an index I so that the locations

```
ISTAK(I), . . . , ISTAK(I + N - 1)
```

form the space allocated for N INTEGER items. Similarly, the statement

`I = ISTKGT(N,3)`

returns an index `I` so that the locations

`RSTAK(1), . . . , RSTAK(I + N - 1)`

form the space allocated for `N` REAL items. Further, the statement

`I = ISTKGT(N,4)`

returns an index `I` so that the locations

`DSTAK(I), . . . , DSTAK(I + N - 1)`

form the space allocated for `N` DOUBLE PRECISION items. Space may be obtained for LOGICAL or COMPLEX items in a similar fashion. Note that the space allocated is not initialized to any particular value.

Since no assumption is made about the relative amounts of storage allocated by the Fortran system to the various data types, it is important that allocations not be divided into suballocations for data of different types. Instead, ISTKGT should be invoked separately to obtain space for each of the different types being used.

The subroutine for releasing space is

SUBROUTINE ISTKRL(K)

which simply releases the space obtained by the last `K` ISTKGT invocations.

As a simple example of the use of these two subprograms, consider a "little black box" subroutine LBB (A, N) which returns something in a REAL vector `A` of length `N` and requires two scratch arrays to do so: an INTEGER array of length `2N` and a REAL array of length `N`. LBB would look roughly as follows:

```

SUBROUTINE LBB (A, N)
C
COMMON/CSTAK/DSTAK(500)
C
DOUBLE PRECISION DSTAK
INTEGER ISTAK(1000)
REAL A(1)
REAL RSTAK(1000)
C
EQUIVALENCE (DSTAK(1), ISTAK(1))
EQUIVALENCE (DSTAK(1), RSTAK(1))
C
  II = ISTKGT (2*N, 2)
  IR = ISTKGT (N, 3)
  ⋮
  {code referring to RSTAK (IR + n) and ISTAK (II + m) probably ending with code
  to store the stuff from the real scratch storage into array A}
  ⋮
CALL ISTKRL(2)
C
RETURN
END

```

To avoid messy (and possibly nonstandard) subscript calculations, it is often more convenient to pass the arguments and the allocated scratch space down one more level to a subprogram which does the real work. This not only makes programs more readable and easier to code, but often more efficient too. Thus

the preceding subroutine can be coded as a “shell,” LBB, calling on a “workhorse” subprogram, LIBB, as follows:

```

SUBROUTINE LBB (A, N)
C
COMMON/CSTAK/DSTAK(500)
C
DOUBLE PRECISION DSTAK
INTEGER ISTACK(1000)
REAL A(1)
REAL RSTAK(1000)
C
EQUIVALENCE (DSTAK(1), ISTACK(1))
EQUIVALENCE (DSTAK(1), RSTAK(1))
C
  II = ISTKGT (2*N, 2)
  IR = ISTKGT (N, 3)
C
CALL LIBB(A, ISTACK(II), RSTAK(IR), N)
C
CALL ISTKRL(2)
RETURN
END

```

3.3.2 Initializing the Stack Size

As previously mentioned, the subprograms in the allocation package are all self-initializing so that a user with small requirements need not even know of their existence. However, there will be applications which require a larger stack than that provided by default. In this case, declarations for the stack and an explicit call to an initialization subprogram must be made in the main program. The initialization subprogram is

```
SUBROUTINE ISTKIN (NITEMS, ITYPE)
```

where NITEMS is the number of items of type ITYPE set aside for the stack. For example, to set up a larger stack with 1000 DOUBLE PRECISION items, the following declarations and subroutine call would be put in the main program:

```

COMMON/CSTAK/DSTAK (1000)
DOUBLE PRECISION DSTAK
  ⋮
CALL ISTKIN (1000, 4)

```

(Since the library programs are compiled with a stack of 500 DOUBLE PRECISION items, this is a nonstandard usage but one supported by most Fortran environments—see Appendix A.)

3.3.3 Stack Status: Query and Modification

By design, it is considered a fatal error to attempt to allocate more space than is actually available. The error could have been made recoverable, but it was felt that this would unnecessarily complicate both implementation and use. For those situations when it is desirable to query how much stack remains, the function

```
INTEGER FUNCTION ISTKQU(ITYPE)
```

can be used. ISTKQU returns the number of items of type ITYPE remaining to

be allocated in a single invocation of `ISTKGT`. (As noted in Appendix B, there is a small amount of space overhead associated with each allocation. If the stack is effectively full, `ISTKQU` will return 0). The statements

```
NLEFT = ISTKQU(3)
INDEX = ISTKGT(NLEFT, 3)
```

allocate all remaining space as a single block of `REAL` items.

In some applications it may be necessary to modify the size of the most recent allocation. This can be accomplished with the subprogram

```
INTEGER FUNCTION ISTKMD(NITEMS)
```

which will modify the length of the last allocation to `NITEMS` items and, in a manner similar to `ISTKGT`, return the index of the first item of that allocation. If the last allocation is truncated, only the first `NITEMS` items are preserved. If the last allocation is extended, existing information is preserved but the added space is not initialized.

As an example of the use of `ISTKQU` and `ISTKMD`, the following program fragment reads an indeterminate number of positive `REALs` into the stack. For convenience we assume that a negative data item marks the end of the data.

```

      :
C
C   FIND OUT HOW MUCH STACK SPACE IS LEFT
C   AND ALLOCATE IT ALL.
C
      NLEFT = ISTKQU(3)
      I = ISTKGT(NLEFT,3)
C
C   INITIALIZE COUNT OF ITEMS READ SO FAR.
C
      NITEMS = 0
C
C   READ AN ITEM INTO THE STACK AND TEST FOR END-OF-DATA.
C
10   IF (NITEMS .EQ. NLEFT) GO TO error
      READ (IRUNIT, 100) RSTAK(I)
100  FORMAT (F10.6)
C
      IF (RSTAK(I) .LT. 0) GO TO 20
C
      NITEMS = NITEMS + 1
      I = I + 1
      GO TO 10
C
C   HERE WHEN ALL DATA READ
C   CHECK THAT AT LEAST ONE ITEM WAS READ,
C   AND, IF SO, TRUNCATE THE ALLOCATION.
C
20   IF (NITEMS .EQ. 0) GO TO elsewhere
      I = ISTKMD (NITEMS)
C
C   NOW THE ITEMS ARE IN LOCATIONS
C   RSTAK(I), ..., RSTAK (I + NITEMS - 1)
C
      :

```

The function

INTEGER FUNCTION ISTKST(N)

allows one to obtain certain statistics on the storage allocator. All quantities are measured in terms of INTEGER items. Because there is no fixed relation assumed about the relative sizes of the various data types, the values returned should only be used for observing the status of the stack. The values returned by ISTKST are determined by the argument N as follows:

<i>N</i>	<i>Statistic Returned</i>
1	number of outstanding allocations
2	current active length
3	maximum active length achieved
4	maximum active length permitted

To determine the exact number of INTEGER items required for the stack, one might include the following statements at the end of the main program:

```
IUSED = ISTKST(3)
WRITE(IWUNIT, 100) IUSED
100 FORMAT(1X,13HSTACK USED = ,I6)
```

More detail on the Fortran implementation of the stack is given in Appendix B.

3.4 Mathematical Programs in PORT

The programs in the PORT library are grouped into 12 chapter areas: approximation, computer arithmetic, differential equations, linear algebra and eigensystems, mathematical programming, optimization, probability and statistics, quadrature, roots, special functions, transforms, and utility. Some of the routines, such as the Jenkins-Traub polynomial root finder [32, 33], Singleton's mixed radix fast Fourier transform [51], and three programs from EISPACK have been adapted from versions appearing in the open literature. All of these have been revised to fit PORT's portability, error-handling, and dynamic storage requirements. The bulk of the routines have been developed at Bell Laboratories. These include Blue's quadrature routine [7], QUAD, which can integrate "noisy" integrands or integrands with a singularity, and Schryer's differential equation solver [50], ODES, which is built around an efficient and robust extrapolation algorithm. Warner and Eldredge [55] have provided a program, BURAM, for finding the best uniform rational approximation on a mesh using the differential correction algorithm. This last group of programs is included in the second edition of PORT which is now out. Also included now are programs for finding the roots of a set of nonlinear equations [8] and a random number generator, implemented to be portable and to provide the same random deviates on any computer with at least 16 bits [24]. (The latter uses Marsaglia's mixed congruential-Tausworthe shift method [41].) An extensive spline approximation, interpolation, and integration package has been added to the library, and there are many new utility routines, including sorting and various vector operations.

PORT now contains 546 subprograms. The library is not as large as that sounds; separate single-precision and double-precision versions of the subprograms (when appropriate) are each counted. There are 125 documented pro-

grams—one piece of documentation applies to a given single-precision/double-precision pair. Many of the lower level routines which are not documented in the current edition will be written up for future editions. The tape contains some 40,000 lines of Fortran, including comment cards.

4. SUMMARY

The PORT library project has been under way for three years; the library is now installed on various classes of computers and users have found it to provide a solid foundation for program development. The emphasis, throughout the development of the library, has been put on portability, structure, and ease of use. In the paper we have surveyed some of the techniques used to achieve these ends. In summary, we make the library portable by using a subset of ANSI Fortran, and by using function calls to obtain machine-dependent quantities. We have structured the library around a centralized error-handling procedure, and we have included dynamic storage allocation, implemented in a portable manner. These approaches, together with simple, brief, user reference sheets, have helped us achieve user acceptance at the various installation sites.

APPENDIX A. NONSTANDARD FORTRAN IN PORT

PORT makes two assumptions on its host Fortran environment that are outside the ANSI standard. These are that there is no runtime subscript range checking, and that variables initialized in DATA statements retain their most recently assigned values. These transgressions and their effect on the library are the subject of this appendix.

The assumption on subscript ranges allows dummy arrays in subroutines to be given a last subscript dimension of 1 under the assumption that a larger value can actually be used. The extension of the assumption to cover arrays in COMMON means that the stack can be initialized in the main program to a size larger than its default size of 500 double-precision locations, even though the library has been compiled using the default size. Similarly, if the relative lengths of different data types are nonstandard (see Section 2.2), the library need not be recompiled to reflect the actual ratios.

The second assumption, that a variable which is initialized by a DATA statement in a subprogram, and then changed within the subprogram, keeps the latest value from one invocation of the subprogram to the next, is used in PORT's error-handling and stack allocation packages. In the error handler, a DATA variable of this type is used to hold the error number of the last error that has occurred, another is used for the current mode (recovery or nonrecovery), and another to store the error message pertaining to the last error, if any. In the implementation of the storage stack, such a DATA variable is used as a flag to specify whether the stack has been initialized, either by the user or, in the usual fashion, by the stack subprograms.

These various slippages from standard Fortran usage could cause trouble if overlays are being used. When a user is running a large program in a smaller space using overlays, care should be taken to keep the error-handling and stack allocation subprograms in the main memory link in order that the values of the

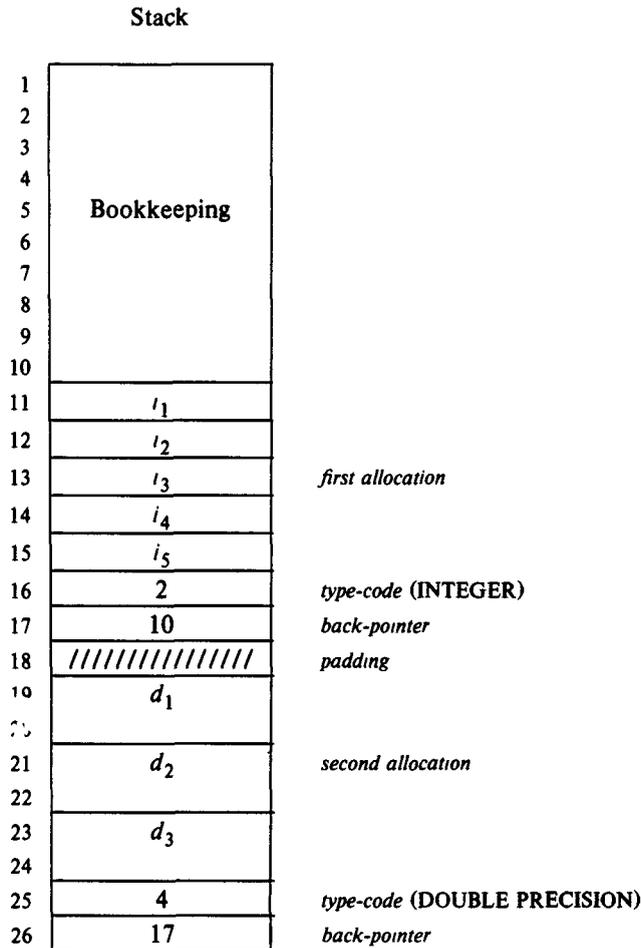


Fig 1 Dynamic storage stack

DATA variables be retained. This is probably a good idea anyway from the standpoint of efficiency, and for the same reason, the three small machine-constant functions should also be kept resident.

Finally, if a user has had to enlarge the storage stack from its default value, or if overlays are being used, the stack should be declared within the main program and the call to the stack initializing routine made in the main program. This will serve to keep the stack in the main link.

APPENDIX B. PORT STORAGE STACK—IMPLEMENTATION NOTES

Each allocation on PORT's dynamic storage stack consists of four parts: initial padding, allocated space, final padding, and control information. The amount of space allocated to the initial padding is less than the space occupied by one item of the type being allocated. The final padding is less than the space occupied by an integer. The padding simply accounts for the differences in the relative

positions of items of different type in the COMMON block CSTACK. The control information takes two integers the first of which contains ITYPE, the type of the allocation. The second word contains the index (in ISTACK) of the second word of the control information associated with the previous allocation. If there is no previous allocation, this contains the space reserved for internal bookkeeping, currently 10 integer locations.

In these 10 locations in the stack, information is stored about the number of allocations currently outstanding, the current active length of the stack, the maximum length achieved so far during the run, etc. Also stored here are 5 integers giving the amount of space allocated to each of the different data types. Since these numbers are used solely for computing subscripts, the unit of measurement is arbitrary and may be words, bytes, bits, or whatever is convenient. By default, the ANSI standard "storage unit" is used. For minicomputer Fortran systems which do not allocate storage as prescribed by the Fortran standard, the subprogram (IOTK00) that initializes these 5 locations should be modified as appropriate. This subprogram and the others in the stack allocation package are contained in Algorithm 528 [19].

Figure 1 shows a schematic of the dynamic storage stack with two allocations outstanding, the first for 5 integers, and the second for 3 double-precision values. The cross-hatching after the first allocation represents the padding needed to align the double-precision quantities on the proper boundary.

At each call to the allocator, the consistency of the stack and the control information stored in it is checked. If an inconsistency is found, SETERR is called to deliver an appropriate message and terminate the run.

Nonstandard Fortran usages in the implementation of the stack are discussed in Appendix A.

ACKNOWLEDGMENTS

Members of both the Computing Science Research Center and the Computing Technology Center at Bell Laboratories have contributed in substantial ways to the development of the PORT library. The authors wish to express their appreciation for this support and to thank, in particular, W. Stanley Brown for his many valuable suggestions on the direction the library should take.

REFERENCES

1. AIRD, T.J., BATTISTE, E.L., BOSTEN, N.E., DARILEK, H.L., AND GREGORY, W.C. Name standardization and value specification for machine dependent constants. *SIGNUM Newsletter (ACM)* 9, 4 (1974), 11-14.
2. AIRD, T.J., BATTISTE, E.L., AND GREGORY, W.C. Portability of mathematical software coded in FORTRAN. Private communication, 1975.
3. American National Standard X3.9-1966 (ISO 1539-1972), FORTRAN, American National Standards Institute (ANSI), New York, 1966
4. AYERS, J.A. Recursive programming in Fortran II *Comm. ACM* 6, 11 (1963), 667-668.
5. BARRON, D.W. *Recursive Techniques in Programming*. American Elsevier, New York, 2nd ed., 1975, pp. 42-45
6. BATTISTE, E.L. The production of mathematical software for a mass audience. In *Mathematical Software*, Academic Press, New York, 1971, pp. 121-130.
7. BLUE, J.L. Automatic numerical quadrature. *Bell Syst. Tech. J* 56, 9 (1977), 1651-1678.
8. BLUE, J.L. Solving systems of nonlinear equations. *Comptng. Sci. Tech. Rep. No. 50*, Bell Labs., Murray Hill, N.J., 1976.

9. BOYLE, J.M., CODY, W J., COWELL, W R., GARBOW, B.S., IKEBE, Y., MOLER, C.B., AND SMITH, B.T. NATS, a collaborative effort to certify and disseminate mathematical software. Proc. 1972 ACM Annual Conf., Vol II, pp. 630-635.
10. BOYLE, J.M., AND DRITZ, K D An automated programming system to facilitate the development of quality mathematical software. Information Processing 74, North-Holland, Amsterdam, 1974, 542-546.
11. BRESFORD, W.M., RELLES, D.A , et al. STATLIB Bell Labs., Holmdel, N.J., 1969.
12. BROWN, W.S. Software portability. In Report of the 1969 NATO Conference on Software Engineering Techniques. J N Buxton, and B. Randell, Eds., *NATO Sci. Comm.*, 1970, pp. 80-84.
13. BROWN, W S A realistic model of floating-point computation. In *Mathematical Software III*, J R. Rice, Ed , Academic Press, New York, 1977.
14. CODY, W.J The FUNPACK package of special function subroutines *ACM Trans. Math. Software* 1, 1 (1975), 13-25.
15. DICKINSON, A W., HERBERT, V P , PAULS, A.C., AND ROSEN, E.M The development and maintenance of a technical subprogram library In *Mathematical Software*, Academic Press, New York, 1971, pp 141-151.
16. FORD, B. Machine characteristics and their parameterisation in numerical software. Private communication, 1976; also FORD, B. Parameterization of the environment for transportable numerical software *ACM Trans. Math Software* 4, 2 (June 1978), 100-103.
17. FORD, B , AND SAYERS, D. Developing a single numerical algorithms library for different machine ranges. *ACM Trans Math Software* 2, 2 (June 1976), 115-131
18. FORD, B , AND SMITH, B.T. Transportable mathematical software: A substitute for portable mathematical software Position paper, IFIP Working Group 2.5 (on Numerical Software), 1975.
19. FOX, P.A., HALL, A.D., AND SCHRYER, N.L. Algorithm 528. Framework for a portable library. *ACM Trans. Math. Software* 4, 2 (June 1978), 177-188
20. GENTLEMAN, W.M., AND MAROVICH, S B More on algorithms that reveal properties of floating-point arithmetic units. *Comm. ACM* 17, 5 (1974), 276-277.
21. GENTLEMAN, W M., AND TRAUB, J.F. The Bell Laboratories numerical mathematics program library project Proc. 1968 ACM 23rd Nat Conf., pp 485-490.
22. GEORGE, J E. Algorithms to reveal the representation of characters, integers, and floating-point numbers *ACM Trans. Math Software* 1, (1975), 210-216.
23. GOODENOUGH, J.B Exception handling: Issues and a proposed notation, *Comm. ACM* 18, 12 (1975), 683-696.
24. GROSS, A. Portable random number generation. Internal Bell Laboratories report, 1976.
25. HAGUE, S.J , AND FORD, B. Portability—prediction, and correction. *Software—Practice and Experience* 6, (1976), 61-69.
26. HALL, A.D. FDS: a Fortran debugging system. Comptng. Sci Tech. Rep. No 29, Bell Labs., Murray Hill, N J., 1975.
27. HANSON, R.J , KROGH, F T , AND LAWSON, C L. A proposal for standard linear algebra subprograms. Tech Memo 33-660, Jet Propulsion Lab , California Inst. of Technology, Pasadena, Calif., 1973
28. HOPPER, M.J. Harwell Subroutine Library, A Catalogue of Subroutines. Theoret. Physics Div., U.K.A.E.A Res Group, Atomic Energy Research Establishment, Harwell, England, 1973
29. IBM System/360 and System/370 Subroutine Library—Mathematics User's Guide. IBM, Germany, 1st ed., Nov 1971
30. IFIP Working Group on Numerical Program Libraries. SIGNUM Newsletter (ACM) 7, 3 (Oct. 1972), 10-11, 9, 3 (July 1974), 3-4, 9, 4 (Oct 1974), 3-4; 10, 2/3 (Nov. 1975), 16, 25.
31. International Mathematical and Statistical Libraries, Library 2, Ed. 5. IMSL, Houston, Tex
32. JENKINS, M.A., AND TRAUB, J.F. Algorithm 419 Zeros of a complex polynomial. *Comm. ACM* 15, 2 (1972), 97-99.
33. JENKINS, M A , AND TRAUB, J F. Algorithm 493. Zeros of a real polynomial. *ACM Trans. Math Software* 1, 2 (1975), 178-189.
34. JENSEN, P S Storage management of numerical processes, Proc Software II Conf , Purdue U , W. Lafayette, Ind , May 1974, p. 265
35. JOHNSON, O.G. IMSL's ideas on subroutine library problems, SIGNUM Newsletter (ACM) 6, 3 (1971), 10-12.
36. JONES, R.E., AND BAILEY, C.B. Brief instructions for using MATHLIB, Vers. 6.0. Sandia Labs., Albuquerque, N. Mex., 1976.

37. KROGH, F.T. A language to simplify maintenance of software which has many versions. *Comptng Memo. No. 360, Jet Propulsion Lab., California Inst. of Technology, Pasadena, Calif., April 1974.*
38. KROGH, F.T., AND SINGLETARY, S.A. Specializer User's Guide. Draft, private communication.
39. LAWSON, C.L. Current status of the SIGNUM Basic Linear Algebra project. *Comptng. Memo. No. 378, Jet Propulsion Lab., California Inst. of Technology, Pasadena, Calif., 1975*
40. MALCOLM, M.A. Algorithms to reveal properties of floating-point arithmetic, *Comm. ACM 15*, 11 (Nov. 1972), 949-951.
41. MARSAGLIA, G., AND ANANTHANARAYANAN, K. Random number generator package—'Super-Duper.' School of Comptr. Sci., McGill U., Montreal, Canada, 1973.
42. MATULA, D.H. In-and-out conversions. *Comm. ACM 11*, (Jan 1968), 47-50.
43. NAG Reference Manual, Mark 2 NAG Central Office, Comptng. Lab., Oxford U., Oxford, England
44. NEWBERY, A C.R. The Boeing library and handbook of mathematical routines In *Mathematical Software*, Academic Press, N.Y., 1971.
45. REDISH, K.A., AND WARD, W. Environment enquiries for numerical analysis SIGNUM Newsletter (ACM) 6, 1 (1971), 10-15.
46. RICE, J R. *Mathematical Software*. Academic Press, N.Y., 1971.
47. RICE, J.R. The distribution and sources of mathematical software. In *Mathematical Software*, Academic Press, N Y, 1971, pp. 13-25.
48. RICE, J R., ED. *Mathematical Software III*. Academic Press, New York, 1977
49. RYDER, B.G The PFORT verifier *Software—Practice and Experience 4* (1974), 359-377
50. SCHRYER, N.L. A user's guide to DODES, a double-precision ordinary differential equation solver *Comptng. Sci Tech Rep No. 33, Bell Labs, Murray Hill, N.J., 1975.*
51. SINGLETON, R.C. An algorithm for computing the mixed radix Fast Fourier Transform. *IEEE Trans Audio and Electroacoustics AU-17* (1969), 93-103.
52. SMITH, B T, ET AL. *Matrix Eigensystem Routines—EISPACK Guide* Springer-Verlag, New York, 2nd ed., 1976.
53. Software II. Informal proceedings of a conference held at Purdue U., W Lafayette, Ind, May 1974.
54. WAITE, W M Building a mobile programming system, *Comptr J 13* (1970), 28-31
55. WARNER, D.D., AND ELDRIDGE, B D An implementation of the differential correction algorithm. *Comptng. Sci. Tech Rep. No 48, Bell Labs., Murray Hill, N.J., 1976.*
56. WILKES, M.V., WHEELER, D.J., AND GILL, S. *The Preparation of Programs for an Electronic Digital Computer* Addison-Wesley, Reading, Mass, 1951
57. Clarification of FORTRAN standards—initial progress. *Comm ACM 12*, 5 (May 1969), 289-294
58. Clarification of Fortran standards—second report *Comm. ACM 14*, 10 (Oct 1971), 628-642.

Received July 1976, revised May 1977