Developing a Numerical Algorithms Library

B. FORD, AFIMA
Cripps Computing Centre, University of Nottingham

IMA Bulletin, No 8, 1972

1. Introduction

ONE characteristic activity of an applied mathematician is the formulation and development of mathematical models. Whilst such models can have many forms, the building process is increasingly oriented towards, and hence aided by, digital computers. The research worker, faced with the problem of reducing his model to computable form, often turns to the known techniques in numerical mathematics. He finds a solution of his mathematical model by using known algorithms, possibly with slight extensions. On other occasions, when the research problem is not amenable to any known method, either new ground is broken or the model is suitably reformulated.

The interest of applied mathematicians in numerical algorithms is therefore considerable and gives rise to a constant demand for reliable numerical library software. In a recent article Fox[1] discussed some of the problems inherent in the choice of algorithms. The present article describes a project in which several computing centres are cooperating in an attempt to create a balanced, general purpose numerical library. Although the project is largely university-based the difficulties faced and the solutions attempted relate equally to government and industrial research establishments.

As the library is being developed in the interests of the people who will use it, I shall start by describing the user community, their aims and how the library can help to achieve them. This is followed by a discussion of some of the major issues faced in the provision of any numerical library. The development of the NAG (Nottingham Algorithms Group) project is described and a concluding comment describes the present state of the project.

2. The user community

The spectrum of users is very wide in their subject backgrounds, their knowledge of computing and their knowledge of numerical analysis. They range from Professors to first term undergraduates. Some will have written major application packages while others will barely be able to write a program. Many will be involved in major research projects. The jobs they wish to compute will vary in running time from a few seconds to the hours required for survey analysis, molecular quantum chemistry or X-ray crystallography. Some users will have a thorough understanding of their program and the algorithms included within it, while others will have been "given the program by a friend" and understand little about it. All have a claim on the computing service and any facility has to be developed with their range of ability, interest and demand in mind.

In particular, since a general purpose numerical library will be of interest to almost all users, every decision relating to proposed contents, documentation, distribution and publicity must be made with their viewpoint and interests in mind. Material must be referred to them frequently for comment and discussion. A careful education programme is also essential.

The majority of the users, regardless of their ability and background, share the common aims of wishing to feed their numbers in and get their answers out, as quickly and effortlessly as possible. Most of them, hopefully, want meaningful answers!

3. How can a numerical library project further these aims?

A numerical library project can help in a number of ways. Here some relevant matters are collected under certain headings rather than in order of importance.

A numerical library will be of assistance to users if:

(*i*) *Algorithms*

(*a*) It contains, in a useful form, an algorithm required for the solution of a particular problem.

(*b*) The algorithms chosen for the library are the best available (as defined below).

(*c*) New, improved or needed algorithms are added to the library as they become available.

(*ii*) *Routines*

(*d*) The numerical software provided is reliable.

(*e*) The library is held in a stable form so that programmers can gain confidence from its use and rely on its availability.

(*f*) Each routine performs efficiently the algorithm that it claims to implement.

(*g*) The source text of the library is available.

(*iii*) *Documentation*

(*h*) Library routines are documented in sufficient detail to make their use, from the information given, straightforward.

(*i*) The documentation is readily available.

(*j*) Advice is available that assists the user to choose a suitable algorithm to solve his particular problem.

(*k*) Criteria are suggested for assessing the significance and accuracy of the calculated results.

(*iv*) *Maintenance*

(*l*) When an error is found in the library software or documentation it is corrected as soon as possible

and users are notified of its existence in the interim.

(v) *Publicity*

(m) The library and its documentation are known to exist.

4. Provision of a numerical library

Although all these features are desirable, in practice some of them are mutually incompatible and some conflict with the declared aims of the users.

4.1 Conflicts in library aims

The collection and comprehension of printed documentation is far from effortless and invariably takes time. It is nevertheless essential if a numerical library is to be used in an intelligent manner.

The individual programmer wants the library to contain routines in the precise form required for his particular problem. Since a routine is included in the library to fulfil a general need, the user will often find that it satisfies his own requirements very nearly, but not completely. So does he adopt it, which takes time and effort, or continue to use his old program? In many instances new algorithms will give not only more accurate answers, or faster convergence, but in addition offer new insights into his problem through improved estimates of stability, measures of the significance of results and of the rounding errors involved.

For a library to be widely used the programmer must have confidence not only in the quality of the software and documentation but also in the continued presence in the library of the routines he uses. The best routine of one year is often the outdated routine of the next. If a routine, once included in the library, is retained for reasons of stability of contents the achievement of other objectives may be impaired. The library may become unwieldy, making excessive demands on storage and search time at program consolidation. Its integrity will be impaired by the continued inclusion of inadequate material and its educational impact may be weakened by apparently conflicting advice.

Another conflict arises from the need to be able to correct errors easily in the text of the routine in the library and the fact that people often want to keep card copies of useful library routines.

There is also the problem of developing and maintaining the confidence of users in a library for which they constantly receive error notices.

4.2 Library criteria that cannot be met

Only the best algorithms are to be chosen for a library, but best for what? By what mathematical and computational criteria can an algorithm be judged? What numerical areas should a library attempt to cover and in what detail? In many areas of numerical analysis the choice of algorithm is problem specific and documented advice can at best be vague and at worst positively dangerous, necessitating expert personal advice. In several areas the programmer is only too grateful for any useful algorithm, yet in other areas the correct choice of algorithm can save hours of computing time and make the difference between correct and meaningless results. The criteria for assessing the significance and accuracy of results are available only in a few routines, often because techniques have yet to be developed which calculate or even estimate such properties.

There are the vital questions of the availability of documentation and public awareness of the existence of the library. The library manual is essential for programmers and users would generally prefer to have their own copy. The cost of providing these manuals is small compared with the annual running costs of a computing centre, but how many people use all of the numerical areas in a library and hence require the whole library manual? Would they not appreciate the manual more if they, personally, had to pay for it? However the manual is distributed, an informed computing community will undoubtedly make more efficient use of the resources made available through the computing service.

The problem of publicity of the library is common to most institutions. So often the information only reaches the people who already know of the existence of the library rather than the intended broader audience. Fox mentioned this point in his article and it is worth emphasising.

4.3 Conditions on library implementation

The common aim of the computer user and the centre providing the service is optimal use of the machine. From a numerical standpoint one looks for optimal use of the machine in terms of the correct choice of algorithm for each particular problem and for careful assessment of the significance and accuracy of the results obtained. Although simply stated these conditions are difficult to fulfil.

There are three areas of practical concern which require attention if the computing service is to be most beneficial.

(a) *Availability of library source text*

For the purposes of error correction and development of the library the fewer people who have the source text of the library the better. However this restriction is obviously unacceptable to the user who often wants to know how a particular routine works. Hence he will require access to the source listing of the library, but generally he will not need a personal copy of the source text. Two groups of users who will need copies of source text are those who need to recode parts of a routine to meet particular needs and those who wish to distribute copies of application programs. In each case the routine is then obviously the responsibility of the person using the program. Programmers will obtain immediate benefits from library improvements and better support from advisory services in numerical methods and in programming, if they use the library version of a particular routine rather than their own version.

(b) *Numerical methods advisory service*

The documentation of the library can only be the first line of advice on algorithm choice and needs to be supported whenever possible by expert numerical advice. This service can also advise users on the reduction of problems to a computable form. Experience suggests that the result of such a service is a considerable saving in machine time and an increased understanding on the part of users of the problems of computation.

(c) *User participation*

If a programmer requires an algorithm that is not available in the library he should be encouraged to contact his local numerical analyst and point out the deficiency. Only if there is a flow of information in both

directions between users and computing services will the library flourish to mutual advantage. A library project must be ready to receive and act on user criticism and advice.

4.4 Library development—the distinction between algorithm and routine

The ability to develop algorithms requires an understanding of the mathematics which underlies a given numerical area and knowledge of techniques for determining required answers given sufficient initial information. In addition to an understanding of the algorithm, its coding into an efficient routine demands a fluent command of the computer language employed, with particular attention to the comparative time penalties required for the various arithmetic and logical operations, the maintenance of numerical accuracy and of the storage used.

An algorithm is independent of language and machine. A routine is dependent on the language, the machine and sometimes even the compiler used in its implementation. Criteria to judge the quality and effectiveness of algorithms are badly needed. At present, much of the testing attempts to judge the algorithm and the routine simultaneously. Several groups in Britain, on the Continent and in North America are seeking to discover the answers to these problems and the NAG project hopes to make a contribution to this discussion. Obviously the simple separation between algorithm and routine suggested above is not always possible or desirable. In some numerical areas the desired routine is provided by a subtle combination of algorithms dependent on a particular language. However the selection of contents on the basis of algorithms appears the best way in which to develop a balanced general purpose library with a useful range of machines.

5. Nottingham algorithms group (NAG)

5.1 History and origins

In the spring of 1970 the Computing Centres of the Universities of Birmingham, Manchester, Leeds, Nottingham and Oxford were notified by the Computer Board for Universities and Research Councils that they would be receiving ICL 1906A machines to replace their present computers. An initial study of the numerical software provided with the machine showed it to be inadequate in the context of an university computing service. It was clear that the development of a numerical library, and its documentation, was needed and for this purpose an exploratory meeting was convened at Nottingham and attended by representatives from the other four interested universities and from the Atlas Computer Laboratory to discuss the creation of a joint algorithms project.

At this first meeting of the group it was agreed that a cooperative effort to create a suitable numerical library was not only desirable but essential if a worthwhile library was to be achieved in the required time scale. The Oxford machine was expected to be fully operational by the spring of 1971, leaving less than a year for the initial work to be completed.

The discussion given in sections 1–4 reflects the consensus from which group decisions were made. In particular from the beginning the library being developed was intended to be part of a widely available computing service.

5.2 NAG library policy

It was decided that a numerical library, even on a large computer, should be a selection of the best material available in the various areas of numerical analysis rather than a collection of all known routines. Firstly, programmers should be encouraged to use software intelligently and secondly attention has to be paid to the storage of such libraries and the time needed to search them. For the reasons already given the choice of material for the library is on the basis of algorithms. Hence agreement had to be reached on what criteria should currently be used for the choice of algorithms in a particular area. Since no rigorous standards are available the selection is of necessity subjective to some degree.

5.3 Criteria for algorithm selection

The criteria, in order of priority, that the group decided upon for selecting algorithms are their usefulness, robustness, numerical stability, accuracy and speed. Where many algorithms are available only the most satisfactory one (or two) are included. Robustness is the ability of an algorithm to solve problems from a wide area of applications. The ordering of accuracy and speed is always a difficult choice. Although generally accuracy is a preferred quality for a library routine it was felt that some algorithms might appear twice in the library, the first version with the emphasis on accuracy and the second with the emphasis on speed.

5.4 Subject contents of NAG library

For the first library the group chose those areas it felt to be essential to any generally useful numerical library. Within these areas it was agreed that descriptive statistics and simple linear algebra (solution of sets of real linear simultaneous equations and the eigenvalue problem for real symmetric matrices) were the most important. The areas chosen were linear algebra, statistics, non-linear optimisation, special functions, sorting, random number generation, ordinary differential equations, roots, approximation theory, quadrature, operational research, interpolation and non-linear simultaneous equations. It was agreed that, at present, no sufficiently general purpose routines were available in the area of partial differential equations to warrant their inclusion.

5.5 NAG library development

Having agreed the general contents of the first library, each of the NAG centres accepted responsibility for two or three subject areas within it. This responsibility involved the following.

(*a*) Determination of the common and likely problem areas in the subject.
(*b*) Determination of the preferred algorithms from the literature and, whenever possible, from authorities within the subject.
(*c*) Suggestion to NAG committee of proposed library contents.
(*d*) Acquisition from the literature, from other libraries and from individuals of routines of the chosen algorithms.
(*e*) Development of useful routines not discovered in existing literature, etc.
(*f*) Translation of routines into Algol or Fortran (when necessary).
(*g*) Testing in detail of the routines.
(*h*) Development of test programs for the routines.

(*i*) Writing of the documentation for the routines.

All numerical software is made available in Algol 60 and ANSI Fortran, since certain incompatibilities between the Trace mechanism in the Algol and Fortran compilers for the ICL 1900 series machines make mixed language programming difficult. This decision also satisfies the different high level language emphases of the six centres.

Test programs are written for all software in the library. The first is a simple test program that appears in the documentation of each routine and simply demonstrates the calling sequence, input and output. The second is a rigorous examination, mathematically and computationally, of a routine or routines. These programs are considered an integral part of the library software. All software is tested and documentation checked in a centre other than that in which it is developed.

There is no input or output within any numerical routine of the library, to aid the transferability of the software between different operating systems and different machines. However an integer parameter is included in the calling sequence of most routines, which may be set to signify that a given mathematical condition has or has not been met. Where intermediate output is of vital importance, for example in non-linear optimisation theory, a procedure for such output, to be written by the user, is included in the calling sequence.

5.6 Routine names

It is important that the Fortran and Algol codings of an algorithm can be immediately related. As the name of a Fortran subroutine or function is permitted only six characters, it was decided for consistency that the name of an Algol procedure should be similarly limited.

It was found to be impossible to give mnemonics in six characters for all routines in the NAG library and an ordered naming policy was therefore developed. The six character name of a routine consists of three parts. The first three characters are the Modified Share Classification Index Entry for the numerical area of the particular routines. The next two characters, which are letters (excluding I and O), are used to order the routines within each area and the final character signifies the language of the program from which the routine should be called (A for Algol, F for Fortran). Hence a typical name would be G05AEA: G05—random number generation; AE—the routine after AD; A—to be called from an Algol 60 program.

5.7 Documentation

The NAG library manual is composed of documents. A document is the smallest unit in the manual and is replaced in its entirety whenever necessary. There are five distinct kinds of document in the manual. The introduction gives advice on the purpose and use of the library manual and the library. The contents document lists the areas of numerical analysis covered in the library. Each entry of the Share index defines a possible chapter in the manual. The chapter introduction consists of an index of the routines in that chapter, background to the contents of the chapter with references and recommendations on algorithm choice. Each routine is described in a document with 14 separate sections, with the Fortran and Algol versions described separately. Finally an index lists in alphabetic order the keywords from section 14 of the routine documents. A routine document is named after the routine it describes, so that the two are readily related.

The manual was designed to have an educational rôle. The beginner would be given the mathematical background and receive advice, when this is possible, on the preferred algorithm for his particular problem in the chapter introduction. The expert would be able to bypass this by use of the keyword index.

The NAG library manual or any part of it can be obtained by users in the NAG centres.

5.8 Error handling

Errors inevitably occur in both the software and documentation and it is essential that any mistake is clearly understood and rectified quickly. It is also important that the integrity of the library source and documentation is maintained by avoiding the creation of a number of slightly different versions.

If a fault is found, users are encouraged to notify their local Advisory Service of the mistake. The numerical analyst (or applications manager) from the Centre then contacts the NAG library co-ordinator who records the error in the library master file, sends out an error alert to each of the relevant centres and dispatches the mistake for correction to the centre who developed the software or documentation. A corrected or extended version of the routine or documentation is then prepared for inclusion in the next release of the library.

5.9 Maintenance and distribution

NAG is responsible for the provision, on magnetic tape, of the source text and semi-compiled version of the NAG library, plus the accompanying NAG library manual, at each of the 6A centres. Within each centre the responsibility for the library and documentation rests with the numerical analyst. The "Co-ordination Manual," which includes error alert documents, defines the precise state of the NAG library and library manual available within the group. A copy of the manual is held by the numerical analyst in each centre. It is hoped to develop and release a new mark of the NAG library at 6-monthly intervals for the next few years.

5.10 Standards

Approximately 35 people have assisted directly in the preparation of software and documentation for the NAG mark 1 library. It has been necessary therefore to provide clear specifications for the writing of library and test material, involving its structure, form and content. Many decisions have been made regarding the internal functioning of the group and its relationship with outside bodies and these are described in the documents that make up the NAG reference manual.

5.11 Support for NAG

A series of seminars is being held to examine the present contents of individual numerical areas within the library and to consider suitable algorithms for inclusion in subsequent marks of the library. The project has received help and support from many numerical analysts who have attended these seminars and furnished routines for the library.

Cooperation between NAG and the American National Science Foundation Linear Algebra project at the Argonne National Laboratory has resulted in many improvements of mutual benefit and a similar activity

has now started with the National Bureau of Standards in Special Functions.

Several university computing centres with IBM, CDC and other ICL machines have shown interest in the library and discussions are taking place for its implementation on these computers.

6. Conclusions

The NAG mark 1 library, which contains some 300 routines, has been in regular use in the 1906A centres since October 1971. Work is in progress on the contents and documentation of the second mark which it is hoped to release in May 1972. The project is a continuing one, depending totally on the cooperation of the people who work in it. The committee consists only of working members who all still share a common purpose. It is expected that many improvements and extensions will be made to the library in the next year or two. In particular numerical areas that are not included at present will be added as suitable algorithms and manpower become available. It is intended to collate the experience gained in the selection of algorithms and the testing of routines in the hope of eventually developing suitable mathematical criteria for algorithm comparison and linguistic insight for the writing of optimal routines. The continued aim of the group is to develop a well balanced general purpose numerical library which will be of assistance to the users in the 1906A and other university computing centres.

Reference

1. Fox, L., "How to get Meaningless Answers in Scientific Computation (and what to do about it)," *IMA Bulletin*, 1971, **7**, 296-302.

This paper was written at the beginning of 1972. The work is now in a more advanced state.

Contents of the Institute's *Journal*

The following papers will be published in future issues of the *Journal*.

1. "The relative sizes of the terms in Chebyshev and other ultraspherical expansions," by D. C. Handscomb.
2. "Method of evaluation of zeros of Bessel functions," by J. Grad and E. Zakrajsek.
3. "The solution of the heat equation in two space variables using Chebyshev series," by P. M. Dew and R. E. Scraton.
4. "A note on an algebra for the K best routes in a network," by E. Minieka and D. R. Shier.
5. "Rank-one and rank-two corrections to positive definite matrices expressed in product form," by K. W. Brodlie, A. R. Gourlay and J. Greenstadt.

Transfer to Associate Fellowship

Graduate members are reminded of the regulations relating to Associate Fellowship. All Graduate members are deemed to comply with the academic requirements for Associate Fellowship and their eligibility for transfer is, therefore, dependent wholly on their professional experience. The minimum period of responsible experience required is three years. It is not possible to lay down simple rules to determine the adequacy of experience, especially for members employed in industry or Government service. For members engaged in academic work, as general guidance, a teacher in a school would be expected to be a head of department or to hold a post which carried a salary on, say, scale 4. For teachers in technical education a post as lecturer Grade 2 would normally be the minimum appointment likely to meet the requirements. In all fields of employment, however, individual circumstances may be such that candidates who do not meet these requirements could be accepted. All applicants are individually considered and the Secretary will always be glad to advise. Forms of application can be obtained from the Secretary.