J.K. Reid 89

T.P. 529

SETTING UP A NUMERICAL ADVISORY SERVICE

by

R. Fletcher and M.D. Hebden

Presented at an IMA Conference on
'Software for Numerical Mathematics and its Applications'
Loughborough University of Technology

16th - 18th April 1973

Theoretical Physics Division,
U.K.A.E.A. Research Group,
Atomic Energy Research Establishment,
Harwell.

April 1973

HL.73/1517

ABSTRACT

This paper is mainly concerned with the experience obtained in setting up and operating a Numerical Analysis advisory service at A.E.R.E. Harwell. Problems in establishing communications between numerical analysts and users are discussed, and in particular we note that the conditions under which a user seeks advice are not those which might have been expected. Examples are presented to show some of the reasons why users do not choose the most appropriate algorithms. Finally we draw some conclusions that we hope will benefit those who are considering setting up an advisory service, and those who take responsibility for a subroutine library.

1. Introduction

Numerical analysts can be divided into two groups: those who regard the subject as a branch of pure mathematics, and those who consider the prime motivation to be towards the solution of practical problems. Most of the latter group would agree that interaction between numerical analysts and people with problems is to the benefit of both parties. We believe it to be essential to keep injecting new problems to prevent numerical analysis from becoming sterile and introspective. We shall describe our experiences at A.E.R.E. Harwell in trying to achieve such interaction, and some of the lessons which we have learned.

Harwell has a professional staff of around one thousand, and they are organized into divisions, representing major scientific and engineering disciplines which are associated with the design of nuclear reactors. The computing service for these divisions is centralized and each has a divisional representative to maintain their interest in the service. Provision of numerical methods has always been considered as a necessary part of this service. However because of the limited manpower, priority in the past has been given to the development and administration of a subroutine library containing a wide ranging set of high quality numerical routines.

Recent increases in staff have resulted in the formation of a Numerical Analysis Group, both to carry out research and to further assist users in numerical problems. The group consists of five people with a strong research interest in numerical analysis, and two with expertise in both numerical analysis and software who carry out the housekeeping associated with maintaining a subroutine library. The setting up of a numerical analysis advisory service was thought to be a good way of helping to disseminate the advantages being obtained though the continuing research in the subject. This has been done in another important way by giving several courses annually in various aspects of numerical methods, each course consisting of about twelve lectures

and practical classes, spread over three successive days.

This paper is concerned mainly with the problems involved in setting up a numerical advisory service. It was felt that the service should be manned by a 'coordinator' or front man, through whom all (or most) queries should be channelled. Details of the way in which the service operates are given in Section 3. However much more important problems must be tackled on the best way of approaching users so that the service becomes well used, and that the advice given is appropriate. Such matters are discussed in Section 2. In Section 4 we present a number of examples, drawn from our practical experience in operating the service, which illustrate these points and the benefits which ensue to the user. Finally in Section 5 we attempt to extract some conclusions that may be valid for other organizations who are considering setting up such a service.

2. Communicating with Users

Initially it was difficult to decide how to set about the task of establishing contacts with users. It seemed clear however that users who were consuming a lot of computer time might be given help to make their programs more efficient, and those getting incorrect answers helped to get correct ones. It also seemed likely that any deficiencies in the subroutine library would be exposed and could be rectified. In an attempt to gain information about users' attitudes, we discussed the projected service with about ten of the representatives of divisions involved in substantial computing, and with a number of major users. The topics which were discussed included:

(i) The nature of the problems solved.

(ii) Which routines in the library were used; and with what failure rate.

(iii) What routines might be added, or improved.

(iv) Whether any problems had been shelved for lack of a known numerical method.

(v) Whether they anticipated our being able to help, either now, or at some future date.

The results of this survey were quite encouraging in that they suggested that to a large extent users were satisfied with the scope of the library. The most useful effect was to inform users of the projected service and to start in the process of creating links between the service and users. However, although a number of instances were given of past occasions on which help would have been appreciated, there was little immediate demand for assistance.

We discussed the outcome of these interviews amongst ourselves, and have come to two conclusions. One is that in the short term we have not realised our original expectation of helping users to speed up their programs. In fact users are not unreasonably averse to spending such effort because of the consequent loss in their time to pursue other research. The user will not spend such effort unless absolutely necessary. More surprisingly, users do not seem to be interested in solving problems which they have shelved in the past through lack of a numerical method. It seems that, in one way or another, users have come to terms with the situation and are now involved in other things. It may well be however that this observation will not apply in a University, where the pressure to solve each problem as it arises and then pass on to the next, is less acute.

The implication is that users will only seek advice under the following conditions:

(a) When their program is so long or so large that they can no longer finance its running.

(b) When their program gives what are known to be incorrect answers.

Bearing in mind the users' reluctance to waste undue effort, it seems most worthwhile to encourage users to seek advice not only in (a) and (b) above, but also when the problem is in the formulative stage. One step in doing this was to send out a circular to all users informing them of

the existence and aims of the service, and in particular stressing the value of a discussion at the formulative stage. However although this met with some response it proved by no means sufficient. The difficulty was and is to establish an atmosphere in which users are prepared to spend time in doing this. It is therefore important to consider why such an atmosphere does not exist a-priori, and how it can be fostered.

One major cause is that numerical analysts often do not encourage users to approach them with problems, largely because to do so makes demands on their time, and prevents them from doing other apparently more interesting or fruitful research. Some numerical analysts take the unfortunate view that although such contacts may be intrinsically valuable, they themselves do not want any part of it. A user may appreciate this and be loth to impose his problem, unless absolutely necessary. The advantage of a formal arrangement such as is described here, is that users no longer feel diffident on this score in making an approach. Furthermore the demands on a numerical analyst's time can be more readily quantified and budgeted for.

However even if the two groups can be brought together, all may not be well. There still remains the real problem of communicating with someone from a totally different scientific background. A language barrier exists between mathematicians and other scientists and engineers and this causes difficulties in understanding each others concepts. It is very important that the scientist is not made to feel stupid merely because he cannot understand some unfamiliar mathematical point. The scientist is mainly interested in solving his problem, to which end he sets up a model which hopefully can be solved. The numerical analyst on the other hand is principally interested in the solution of the model, and its origin is of small concern. One might say that the scientist works in 'problem space' and the numerical analyst work in 'model space'. This also gives rise to a number of difficulties. For instance the scientist may have performed what to the analyst are

undesirable transformations on the model before presenting it for solution. Alternatively the method proposed by the analyst might lead to a solution violating certain physical principles not made explicit in the model. Yet again the scientist may be suspicious of the method of solution of the model because it does not have an obvious interpretation in terms of the problem.

As an alternative to consulting a numerical analyst, a user may seek advice from a member of his own department who has experience in solving similar problems, and with whom he can converse readily. Although this takes some of the load off the numerical analyst, it has its undesirable features, one of which is that the user might not get the most up-to-date advice. Another undesirable situation may occur when a knowledgeable user has developed a program to carry out some task, and the program is used by less experienced users. These latter users are not in a position to incorporate new subroutines into the program, and if an error develops they cannot rectify it. Furthermore there may be restrictions on the scope of the program which the originator has not made explicit. Often the situation is not serious unless the original author is no longer concerned with the running of the program in which case there is nobody who knows that is going on.

As may have become apparent, our way of solving such problems has been to attempt to build up a good atmosphere between numerical analysts and users, so that after receiving advice when stuck ((a) and (b) above), users return at the formulative stage of a later problem if the choice of method is not clear. This effect has been very noticeable. Furthermore we find that users encourage their colleagues to use the service if they themselves have had a good experience with it.

Finally in this section we would point out another useful means by which contact has been established with users. This has been through the numerical methods courses mentioned in the introduction. Users often attend such

lectures in the hope of learning about techniques which may be applicable to their problems. Discussion of such a problem often takes place during the course, and indicates that the choice of a suitable numerical method may not be obvious, and that there are advantages in consulting a numerical analyst. Furthermore students who have discussed problems with the lecturer on such a course, often return to him or to the advisory service at a later date.

3. Operating the Service

Although a formal advisory service has been set up at Harwell, in fact it operates with a fair degree of flexibility and the details to be given here are considered as guidelines rather than hard and fast rules. In the first instance, users are expected to contact the coordinator to establish the best procedure. This may involve the coordinator recommending a method of solution, possibly after studying the problem to a greater or lesser extent. On the other hand the coordinator may suggest that the user sees one of the specialists in the Numerical Analysis Group who back up the service. Alternatively the coordinator may discuss the problem with one or more specialists and then report back to the user. It may be however that the user has received advice from a specialist on a previous occasion in which case he is likely to contact the specialist in the first instance on a later project. In the interests of maintaining good relations, such an arrangement is not ruled out. For most problems a satisfactory line of attack can be decided on fairly readily, and although a number of visits may be necessary to iron out snags or misunderstandings, the total demand on the coordinator's time is no more than a day or two. Of course the user is expected to do the majority of the programming and testing himself, although possibly the advisor might suggest the form that such a program should take.

Less trivial problems involving the development and testing of new numerical methods require more organization. The coordinator's salary is buried in the costs of running the computing service and it is not acceptable

that a large amount of his time should be made available to individual users on this basis. At Harwell there is the provision for contract work to be carried out, so that the user can hire a specialist (subject to availability) to develop a suitable algorithm, should there be funds available. However such arrangements are best regarded as being beyond the scope of an advisory service.

Another question arises as to who should deal with queries about programming (what does this error message mean? - why didn't my job compile? - etc.). It is felt that answering such queries is not the function of the numerical analysis advisory service. In fact at Harwell such questions are handled in the main by one of the two people responsible for the day to day running of the subroutine library. Many university computer installations seem to provide a query answering service of this nature: this is satisfactory and leaves the numerical analyst free to give the advice for which he is most qualified. Nevertheless such queries do arise from time to time, and the advisor has to strike a balance between becoming increasingly bogged down by answering too many such queries on the one hand, or by risking his good relationship with users by answering none on the other.

It had been expected that the operation of the advisory service would lead to feedback on the deficiencies of routines in the subroutine library and the need for new routines. This has in fact happened, and the library has largely met these needs, through the efforts of members of the Numerical Analysis Group. However the prime function of this group is to pursue research into numerical analysis, with the consequence that it has occasionally not been possible to introduce a subroutine for which there has been a request. Perhaps in these circumstances there is a case for contracting out work to research workers in Universities, given that money is made available for developing the subroutine library, or for closer liason with other libraries such as the NAG library.

4. Some Typical Problems

In this section we illustrate a number of points that have arisen in the operation of the service. They are chosen to show some of the reasons why users do not use the best available algorithm. Some of them indicate the value that a short discussion with a numerical analyst can have even when the problem appears straightforward, and indeed a few users have been saved a great deal of fruitless work when defects in the model have been noticed.

The first example arises because the library does not contain a subroutine for solving the problem in its most natural formulation. It is necessary to transform the problem to one for which a suitable subroutine does exist and the user, ignorant of a standard numerical technique, applies a transformation that is inappropriate, for instance one which might introduce instability into the problem. Consider the eigenvalue problem

$$A\underset{\sim}{x} = \lambda B\underset{\sim}{x}$$

with A and B symmetric and B positive definite. The most obvious approach which the user might choose is to solve the equivalent problem

$$B^{-1}A\underset{\sim}{x} = \lambda\underset{\sim}{x} ,$$

but $B^{-1}A$ is no longer symmetric and its eigenvalues may be ill conditioned. If instead we form the factorization

$$B = LL^{T}$$

the problem can be rewritten with a symmetric matrix as

$$(L^{-1}AL^{-T})\underset{\sim}{y} = \lambda\underset{\sim}{y} \quad , \quad \underset{\sim}{x} = L^{-T}\underset{\sim}{y}.$$

This transformation is well conditioned and reduces the problem to one for which a satisfactory subroutine exists. A note is made of the deficiency in the library so that, resources permitting, a more suitable routine can be added to the library.

The next example is chosen to illustrate the way in which users work with familiar concepts, and lack the perception necessary to distinguish the numerically stable approach. The problem is that of finding the only positive value of t such that

$$a = \sum_{i=1}^{n} \frac{x_i}{y_i + t}$$

for given a, n, x_i, y_i. The user was in little doubt that t could be found satisfactorily by rewriting the expression as

$$P_n(t) = 0$$

where $P_n(t)$ is the appropriate polynomial of degree n in t, followed by using a routine for finding the roots of a polynomial. This method is, however, liable to be unstable, and it should be better to regard the problem as a single non-linear equation in t. A short discussion revealed that x_i and y_i were all positive, and that $\Sigma(x_i/y_i) > a$ so that bisection could be used with the initial interval $(0, a/\Sigma x_i)$.

Occasionally, however, the user has made an unsatisfactory transformation to the problem which cannot easily be overcome. An example of this occurs when wishing to constrain a variable x_i to be positive in the minimization of a function whose derivatives are not available. One possible transformation is to minimize with respect to a variable $y_i = \sqrt{x_i}$, which is satisfactory except when the constraint is binding at the minimum. As the numerical analysts have provided no better method the user has no choice but to use a method like the above. This problem illustrates the advantage of the advisory service to the numerical analyst, in pointing out directions in which new research ought to be directed.

Another example shows a longer term project. The problem was: given a histogram with bars of unequal width, and with the height accurate to between five and ten percent, to find a new histogram based on different

unequal widths. The user, again working with a familiar concept, had fitted the cumulative distribution and then had derived the new histogram by taking differences of the fit. When this failed he came to the advisory service for help. After some thought and experiment we devised a method based on determining the function whose definite integrals generate the histogram. This was a case where the correct line of attack was not clear and the collaboration was more in the sense of a small joint research project with the user.

Even when the user has formulated his problem correctly, and is conversant with the methods available, he may have good reason to use what the numerical analyst would consider an inferior algorithm. Two readily accepted reasons for doing so are to reduce the time taken or the storage required. There are, however, a number of other possibilities and two examples will be given. One instance occurred when a routine for non-linear least squares calculations was superceded by a new method. The old subroutine had the facility that the variance-covariance matrix would be passed to the calling program, whereas the new routine lacked such a feature. As a result, some users did not use the superior algorithm. On another occasion lack of portability caused a group to use an inferior algorithm. They were developing programs for use on a wide variety of machines and had contracted to write the programs in standard FORTRAN. As a consequence, any of the Harwell Subroutine Library routines which they required had to be converted to standard FORTRAN, and so instead of using the best possible routine for a problem, they tried to make do with routines that they had already converted.

5. Conclusions

By and large the service at Harwell works well except perhaps that its scope could be extended by making a statistician available to consult with users. Much of the smooth running can be attributed to the provision

of a subroutine library that is well documented* and in which high priority is given to the rectification of errors. Indeed without a good library it is unlikely that such a service would be effective (it is undesirable that the coordinator should have to write or modify software to any great extent). As a result the user who is familiar with FORTRAN and is reasonably numerate is well catered for. In addition we provide courses that cover the spectrum of Numerical Methods so that other users can acquire some expertise.

Less sophisticated users do not fare so well. Because of their lack of knowledge of numerical analysis they are often using subroutines for methods which they do not understand, so they are liable to use a routine when it should not be applied. This would not be too bad if they utilised all the error checks that are available. However their lack of numeracy causes them to fail to spot indications of trouble, and they often take the attitude that results are correct unless they are known to be wrong. To counteract this we feel that it is important that library subroutines print out an error message when an unreliable answer has been produced. Another way by which numerical analysts can help such users is to alleviate the need for them to develop programs. Clearly it is not practicable for a computer centre to provide all users with the program that they want but there are a number of areas where program packages could be provided. A prime area is in the field of non - linear data fitting and other areas are differential equations and linear programming. An example of what can be achieved is the package FATAL** developed by two users in the Health Physics division at Harwell. This package allows the user to fit data, supplied in any form that is convenient, by any approximating function. The only requirement is that the user can write the necessary FORTRAN statements to evaluate his approximating function. An additional refinement is that by using the remote terminal system the user

---

*M.J. Hopper (1971) 'HARWELL SUBROUTINE LIBRARY, A catalogue of subroutines' UKAEA Research Group Report A.E.R.E. 6912; and 'Supplement 1' and 'Supplement 2' (1972).

**L. Salmon and D.V. Booker, UKAEA Research Group Report AERE R7129 (1972).

avoids the need for a knowledge of Job Control Language. This user package is exceptional in that it is well documented and the authors are still concerned with its maintenance. Indeed they incorporated a new non-linear least squares routine into the package shortly after its introduction into the library. However the standard for other such packages is often much lower, and unfortunately users do not have the discrimination to decide whether a package will prove to be reliable. Packages may prove to be poor not only from an inferior numerical method, but from inadequate facilities, poor maintenance, possibly caused by inadequate or non existent documentation, or poor error indications. The only sure way to avoid this misuse of the package is to have it supervised along with the subroutine library to the same high degree of attention.

Given a good library we have found that one person, spending about half his time in an advisory capacity, can cope with almost all of the problems that arise. Occasionally a discussion with a colleague or colleagues is necessary, but the fact that this is possible within a group such as ours at Harwell, we regard as a bonus rather than a necessity. More difficult to handle are jobs that require a large amount of a numerical analyst's time (a few months for example), and exactly how these are handled will depend on the installation. Again these jobs are relatively few. Almost every university has a set-up similar to that at Harwell, in as much as there are specialists who teach Numerical Analysis and users from many disciplines on the campus. It should therefore be possible for them to provide a similar service. There also is no reason why our experience cannot be repeated in any environment with a single numerical analyst who has occasional consultation with experts.

Finally it should be said that co-ordinating a service such as we have described brings the numerical analyst into contact with a wide variety of problems. It is therefore an excellent way for him to apply his training as

a numerical analyst, and is particularly stimulating if his background is of a more theoretical nature.

Distribution of this report is limited to U.K.A.E.A. Libraries only.