

SCIENCE RESEARCH COUNCIL

ATLAS COMPUTER LABORATORY

A ROLE FOR THE ATLAS COMPUTER LABORATORY

A personal view

by

J E Hailstone

26 April 1973

- CONTENTS :
1. The present role
 2. Meeting future demands
 3. Software packages and the "Meeting House" concept
 4. The Atlas Laboratory as the SRC Centre for computing
 5. The ICL "P4" offer
 6. Other roles and activities
 7. An overview and recommendations

INTRODUCTION AND SUMMARY

The paper examines and comments on a number of topics being examined to determine and clarify a future role for the Atlas Computer Laboratory. A set of recommendations is given which are derived mainly from the experiences of computers and computing since the inception of the Laboratory.

The various sections (above) can be summarised as follows:

1. The present role

Outlines the features of the present support for guaranteed users and indicates its relevance to other bodies including the Computer Board.

2. Meeting future demands

Two main arguments are distinguished. First, the support of high quality and specialised services for supported science using production equipment of the highest available power. Second, meeting a similar demand, but extending the role to include development of new techniques and systems for the advanced "pipeline" machine, CDC STAR.

Short term/long term function & done by Computer Board.

- 1 -

*Board: Distributed computing power.
Demand conflicts one way budget.*

3. Software packages and the "Meeting House" concept

The need for "a home from home" for scientific groups with a common interest in a sub-discipline and with the need to use large computing resources is discussed, and the appropriateness of ACL as such a centre indicated. The interest of other Research Councils in such a concept is known and this supports the whole concept.

4. The Atlas Laboratory as the SRC Centre for computing

Some relevant features of this idea are discussed and the problems of bringing the separate interests together are explored. The dangers of dominance by particular groups is exposed.

5. The ICL "P4" offer

A personal view of the involvement with ICL and its implications. In terms of its support for SRC research there is little to argue in favour of this development. A new role as an advanced production or research laboratory with the British Computer effort in mind is considered.

6. Other roles and activities

Subsidiary activities in graphics, artificial intelligence, operations management, etc. do not in themselves provide a mainstream activity but are essential concomitants for an active and high quality Laboratory. The importance of the development of these subsidiary services is stressed.

7. An overview and recommendations

Brief comments on the role of the computer in SRC affairs and concluding with four main recommendations for the Atlas Laboratory to function as an SRC Laboratory for computing. In response to the National need, the main implications of the installation of an ICL P4 are given.

...

THE PRESENT ROLE

The continued pressure on computer resources and in particular the fact that the ACL has distributed some 16.0 (195) hours of computer time guaranteed in the past year shows that the special support for projects of interest is of great value to research workers - reasons for this are that

- (a) the certainty of a useful amount of time and services is important in planning long-term research projects,
- (b) the research worker for whom the computer is the main research tool does not have to compete daily with the very varied university loads; the presence of critical times for student work means that research work suffers,
- (c) the service offered to the large user can be trimmed so that for example an accompanied large core demand can be met without upheaval and resentment by the small user,
- (d) the large user demands additional facilities and the Laboratory has pioneered developments of value. For example, the production of cine films for the display of complex physical processes is one area of development of the SD4020 microfilm plotter which the Laboratory was able to install in 1968 and for the first time to make available this form of graphical output to university research workers.

The continued support of this work depends on the SRC wishing to ensure that its research policies are carried out without a heavy dependence on the services and facilities generally available within the Universities. The policy, particularly in the Regional Centres, developed for the provision of large facilities by the Computer Board is highly relevant. The major provision of two 7600 computers for the immediate future is not seen to be followed by other plans for future provision and with the large numbers of Universities sharing such resources, there is little prospect that the present Atlas load can be absorbed. In the long term the Regional centres could be planned to take on some or all of the Atlas load to the satisfaction of SRC and its supported users, but for the short term, say to 1980, the Atlas role would seem to be essential.

It should be noted that this guaranteed support of users, although a new departure for scientific research workers, brings the class of computer service offered much more in line with the commercial and industrial use of computers which are operated to carry out dedicated tasks and are controlled to ensure that these tasks are carried out with a high probability of success.

This role is consistent with the view taken by nuclear physicists that their work demands unusually powerful facilities and provision is made for university groups, both theoretical and experimental, to have access, as of right, to the 370/195 at the Rutherford Laboratory.

Thus the development over the past eighteen months has been meeting a requirement which can be supported by cogent arguments.

MEETING THE FUTURE DEMANDS

This part of the picture must depend heavily on imaginative ideas about the nature of computing in science and the broad directions that science will take in its demand for computing facilities. It is clear, however, that the broad scientific base of SRC and the wide applicability of the computer will continue to require computer facilities of considerable power. The nature of research precludes a precise specification for demand and the support for installing equipment rests on a good deal of indirect evidence. Two rather different forms of the argument can be distinguished:

- (a) the SRC may recognise that support for its work requires the provision of special facilities;
- (b) the development and exploitation of advanced computing equipment may contribute to an early acquaintance with new scientific areas of activity.

The practice of providing computer services for a wide class of users assumes that there is a good deal of common ground between them and that no one user can influence the service provided to any great extent. The SRC in supporting research in Universities has to rely on this principle for the majority of its grant holders. The SRC may and does wish to support areas of particular promise, but it does not have the necessary power to direct the allocation of resources within the universities. SRC supported work is in direct competition with other activities within each university and there is no guarantee that such work will necessarily get the level of facilities required for a successful implementation. There are many examples of this kind of difficulty; one of the most common is the inability to obtain short test runs requiring large allocations of core store during the day shift in preparation for a large production run overnight. Others arise from the relative isolation of large users in universities with only small and inadequate facilities and the now well known preference by university computer services for small and standard jobs so that the "many" are always preferred to the "few". The experience with the Atlas service showed that those workers who were enthusiastic enough to put up with the inconvenience of visiting the Laboratory, or waiting on a postal service supported by limited operational help from Atlas staff, preferred to use the Atlas facilities. This was because, for a large part of the lifetime of the Atlas service, the power of the Atlas I was such that large users became relatively small users and that the Atlas service was flexible and comprehensive enough to deal with special cases. It is to be expected that the Regional Centres will now be able to deal with much of this work. A dialogue between SRC and the Computer Board is very necessary at this level to ensure that the characteristics of this class of work are catered for; the use of regional centres as an overload facility may take preference over the demands of the specialised user and the low power remote workstations for access to large amounts of computing power may create a "bottleneck" for this class of computer user.

The requirements of the large user are specific and increasingly specialised. The large user tends to be very experienced and can handle quite complex programs, he has applied often sophisticated and new numerical techniques to his problem and the interaction with the program and output needs and directs the provision of new and elaborate equipment. The provision of graphical output via storage tube displays and microfilm plotters is a well known example, but the recent work of Professor J M Stewart and the X-Ray crystallography program shows that at this level of computer programming special problems arise and it is not possible to treat these demands with the labels of average characteristics.

There is a good deal of evidence to suggest that exceptional services at critical stages of a programme of work are vital to the development of some projects. The research proceeds with some intensity and the result is often recognised as deriving directly from the ability to demand high quality services for a finite period. Recent examples of this effect can be seen in the work of Dr J R A Pearson (Cambridge and lately Imperial College, London) on enzyme chemistry and the collaboration of ACL with the Institute of Geological Sciences (Mrs E Gill) in the preparation of a geological data bank management and retrieval system.

The ability to meet special demands of this kind requires sufficient power not only in dealing with the total computer time required, but in handling, rapidly and easily, the demands for services. In addition the availability of sufficiently powerful resources itself stimulates and generates new areas of activity; the work of the Meteorological Office in the early days of the Chilton Atlas computer led eventually to the development of powerful programs needing an IBM 370/195 for day to day forecasts.

If SRC is to support work of this kind, the computing power must be rated as in excess of 25 x ATLAS units. The take up of some 7 ATLAS units in the past year and given the "lead" time to effective installation, together with a reasonable useful life of more than 5 years, shows that enough power must be planned effectively for circa 1980. The prospect of very substantial reductions in the cost per unit of computer time for the most advanced systems is advantageous where large computer time limited jobs are common. In situations where value for money cannot be measured by the profitability of the task, the cost per unit of computing becomes a critically important measure.

To provide specialised services leads to the conclusion that computer systems of the order of IBM 370/195, CDC 7600, or CDC STAR are required and can be the most economical for SRC both in financial terms and in the application of its human resources.

The previous paragraphs have concentrated on the application of computers to specialised projects and has not been concerned with specific equipment. A good service and reliable guarantees of availability can be obtained with the standard production machines available from

IBM and CDC, but I wish to explore the case for the Atlas Laboratory installing the most advanced of the computing systems, the CDC STAR. No other manufacturer has offered equipment of this power, so that a detailed study of the case is necessary. It should be noted that IBM, the only other real contender in this particular range of computing equipment, has not declared its hand and one can only but observe that IBM has to be very careful at this stage to avoid disturbing a well established, but outdated, software system for its computers; change is inevitable during the next five years and it will be a feature of all manufacturers' activities. This remark is pertinent to the argument that the STAR system is different and that it requires new methods of working to be developed. The system is, however, designed to overcome many problem areas of conventional computers:

- (a) operations can be speeded up considerably using the concept of a continuously fed "pipeline" - the vector or string feature;
- (b) the architecture is based on the concept of the dispersal of functions to specialised stations. Simple and possibly repeated functional units carry out tasks independently, but controlled by a simple central system;
- (c) the operating system, reflecting the machine architecture, can be designed and modified in a modular form;
- (d) the machine has good channel capacities for handling the traffic likely to be experienced with a powerful system;
- (e) communications and the connection of terminal equipment has received special attention and the system is expected to handle exceptional numbers of both connections and terminals;
- (f) a "monitor" station for both maintenance and measurement of performance is an integral part of the design.

The concept of a "pipeline" for many repetitive operations leads to the most important of the changes that the system programmers and the users have to face. The concept of the vector process is in many ways a natural one for both the scientific and business user and could lead to some simplification and conciseness of expression in program writing. The limited evidence so far obtained shows that substantial gains can be expected with small changes to existing programs, but that with enhanced basic software the effect could be dramatic. The involvement of ACL staff in the software for this project could be most effective and would contribute significantly to the development of techniques not readily available to any research workers in the U K. In many ways a natural successor to the Atlas I, the machine offers not only an imaginative and forward-looking way of meeting present demands, but a prospect of real scientific contributions in the later years of the 70's. The role of the Laboratory would be broadly based on two main streams of activity:

- (i) the support of research demanding exceptional and possibly advanced facilities;
- (ii) the development of research into computers and the computing process.

Unless the present interest in numerical evaluation, simulation and modelling of complex physical processes is only a passing phase, the preparation and early participation of staff in this kind of activity is necessary. The Laboratory with its already broadly based contact with relevant research disciplines is very suitable as a centre for such an installation. With a reputation of good work in the past it is known that such a role would be applauded by workers in a number of disciplines.

It is known that CDC expect to produce about 20 such machines, and the design and specification indicate that further enhancement of the power of the system is likely. From the SRC point of view there is no difficulty in seeing such a computer as a central source of power, and the timing of an installation for 1975 would not run counter to present plans for computer procurement within SRC; there is a limited involvement with ICL 1900 series and the two main computers, at RHEL a 360/195, at Daresbury a 370/165, are both now obsolete lines and both Laboratories are running down some of their load. A build up of a STAR machine from 1975 onwards is therefore not inconsistent with present developments and has the merit that advanced machinery would be available in 1980.

SOFTWARE PACKAGES AND THE MEETING HOUSE CONCEPT

The emergence of large and complex software packages in the 60's has shown that the immense power of the computer can be made available to many scientists and that useful calculations can be handled by the scientist who does not wish to become involved with the intricacies of the computer hardware and software. The generation, development, testing, and maintenance of these large systems has become a professional job demanding the very best staff and a continuing requirement to keep in touch with the developing science. Experience with the packages developed for XRAY Crystallography, quantum chemistry, time series analysis and statistics shows that a continual dialogue between the users and the writers and support staff is very necessary. The computing process needs to be looked at from a number of points of view and the intention of the researchers kept constantly in mind. These packages of wide applicability emerged as the result of the efforts of particular individuals or small teams and have been recognised as tools of some significance. The continuing availability of these programs through changes of computers, responsibility and expertise is a matter of some importance. This together with the benefits of scientists, within a sub discipline, meeting to discuss and work out ideas and projects in an environment devoted to computing, encourages us to believe that there are real advantages to be gained from the development of teams of research workers assisted by

professional services within the Laboratory. Only a small change in emphasis is necessary to bring this about. The Laboratory is equipped already to handle such a concept and the staff are experienced in dealing with demands of this kind. This concept ties the Laboratory work even more closely with SRC supported science. If the idea of specialised support is accepted as a suitable role for the Laboratory, the generation of suitable topics for the "meeting house" approach could be readily recognised and developed.

The use of the Laboratory facilities by computer departments of the other Research Councils, NERC, SSRC, MRC, is already in being, or discussions are being held with them. Apart from the function of a central computer facility, the indications are that the setting up of a small team at ACL to look after and develop programs and techniques of interest to a particular discipline has appealed very much. Councils' representatives and NERC are seeking to establish such a team. A permanent home for some of their large systems is seen as of great value. It is necessary for a central organisation to be well equipped and, in addition to the availability of computer power and services, a wide range of ancillary equipment is seen as a particularly useful feature for these Councils who individually may never be in a position to encourage and justify use of new and advanced equipment.

Given the provision of adequate accommodation and computer services, this development is consistent with the view that applications are now as important as the provision of hardware. The structure of the Laboratory would need to be adjusted to accommodate such teams and to give them some influence on the direction that the Laboratory's services should take.

THE ATLAS LABORATORY AS THE SRC CENTRE FOR COMPUTING

The case for a centralised source of computing to be sited at the Atlas Laboratory rests on the idea that economies of scale and scarce resources can be obtained. It assumes that there is some common element in the nature of the work and that a simultaneous support of several large computer centres at RHEL, Daresbury and Atlas, is not reasonable. There are several important features which should be considered:

- (a) the level of computing experience between SRC stations is very different - the Nuclear Physics community is most advanced in its application and scale of demand for a small dedicated class;
- (b) the Atlas experience is with a wide class of user disciplines and the support and care of these users requires a sympathetic and flexible attitude;
- (c) the communications equipment available to make a centre really effective is not yet available, although expensive and/or experimental equipment exists;

- (d) an organisation controlled to provide facilities for all four SRC operating divisions must take into account the dominant role played by RHEL and Daresbury in the present provision of computer power;
- (e) the continuing demand for on-line instrumentation facilities at nearly all SRC Laboratories is unlikely to lead to a reduction in local facilities being acceptable even in the long term.

In principle such a centre is feasible and could provide a wide base for expertise across the whole computing process. Its implementation however presents formidable problems because of the present distribution of power and a very strong control would be essential. On balance, the gradual and well planned absorption of different activities from the SRC Laboratories could be achieved over a period of years, say by 1977. The Atlas Laboratory acting as a centre, but with its main computing power residing in nuclear physics establishments is not to be recommended.

I would favour such an approach if a means could be found to build in checks and balances so that all sciences could be treated on equal terms.

THE P4 OFFER

This invitation to purchase at full price the first of the ICL P4 computers, is not an attractive proposition given any of the possible roles for ACL within the SRC areas of responsibility. The Laboratory has examined this proposed machine and finds that it is inadequate to meet the likely demands of ACL. A good deal of special pleading that the Laboratory can offer good and worthwhile services to benefit the British computer industry is not substantiated and it is doubtful that the Atlas requirements are at present sufficiently close to those of ICL's potential customers to make this a sensible proposition even for ICL. As it stands, the "offer", and I am concerned that this word conveys some special and worthwhile benefits, is too vague and the implications for ACL too open-ended.

However, given that a completely new function is defined for ACL which covers the support through an advanced production development laboratory of the "new range" machines, then it is possible to define a series of projects which could fill a gap in the British Computer effort. The absence of an ICL computing research laboratory well equipped and devoted to computer research, particularly on user and operational interfaces, is very noticeable and has accounted for a good deal of the pain and difficulty of producing a high quality final product. To achieve such a goal would be difficult and would commit the Laboratory unequivocally to the ICL thinking and implementation of their ideas for computers to 1980. The Computer Board have carried out some limited studies of this range and may be convinced of the viability of the concept, but a good deal of support is likely to be required to bring the product to an acceptable condition. To assist the British computer effort in this way may be critical and is a clear

indication that ICL resources are barely adequate to launch this "new range". The choice lies between using the Atlas resources for the promotion of scientific research or to defer to a possibly greater national need. I do not think that it is practical to support both half-heartedly, using P4 as a slightly uncertain and inadequate source of power for the guaranteed SRC user.

Under a computer support role the question must be raised concerning the appropriateness of SRC as the controlling organisation. Given limited objectives for the support of P4, a short-term diversion of SRC resources could possibly be justified but in the long term some other controlling body would be necessary.

OTHER ROLES AND ACTIVITIES

There are a number of subsidiary roles and activities which could be associated with these main objectives. Experience has shown that an active working atmosphere attracts hard working and original people and that in the shadow of a large project new and imaginative ideas are generated. The participation in the computing process also offers opportunities to develop good operational management standards. The Atlas Laboratory has contributed substantially to the standards of operating large scientific services and there would appear to be further work in developing and "setting the pace" now that some of the problems are better understood.

Particular attention has been directed to the support of graphical output techniques; the PDP15 installation supporting and complementing the SD4020 microfilm plotter is a good example of this development of a subsidiary but important related activity. The Laboratory has pioneered a lot of work here, which would not have been possible in a less central and smaller organisation. It is the large users who push the frontiers of computer technology and practice into new areas and our experience suggests that there is a real demand for this type of work to continue.

The suggested alternatives for the work in Artificial Intelligence, including a National computing facility, could be handled as subsidiary roles in an Atlas central facility, but on its own Artificial Intelligence does not appear as a main stream activity. Perhaps it can develop into one, but others must deliberate on the probabilities and possibilities of this happening. The argument used is essentially that of the Laboratory's existing role, specialised and high quality services, for particular sciences and disciplines. The content of the subject is closely linked to computer facilities and could benefit from association with large and professional central services.

A strong support for these subsidiary services is necessary for it is here that the real power of the large resources can be developed. The simple dichotomy of "short jobs" and "number crunchers" is inadequate to describe the user profile and these services should be

regarded as an essential ingredient to the successful exploitation of the powerful computers needed by the scientists. The complex arithmetic processes carried out are at present inadequately supported by devices for the user to monitor his progress and it is time that subsidiary services are raised to be on an equal footing with the provision of computing power.

There is great scope here for the SRC to define really effective projects for the eventual benefit of its supported work. The aim should be to ensure that scientific endeavour is not hampered by the unnecessary concern with the details of computer technology.

AN OVERVIEW AND RECOMMENDATIONS

The choice of a suitable role for ACL in the work of SRC is not dictated by a scarcity of ideas. The support of a Computing Laboratory is necessary by the nature of the subject. It is unlikely that computing will cease to be a major tool for all scientific disciplines and, I think, highly probable that it will become an even more critical piece of equipment. The scientific community is pausing to 'catch its breath' after the working up of computer facilities over the past ten years. The immense scope of the computer has probably led to an overrunning of the possibilities open to research-workers - promises of great things have not been fulfilled and there is a realisation that it is harder to extract the benefits from this tool than was at first realised.

To deal with the SRC commitment to science, I would recommend that

- (a) The SRC should define a class of work to be supported - the existing role embracing the "meeting house" concept can be developed into a suitable primary objective. It is necessary however to be specific and an open-ended and vague objective will be no good at all.
- (b) The SRC should move towards setting up a central Laboratory with a planned absorption of SRC projects. This is necessary so that the present structure does not have a heavy influence on the final organisation of the centre or Institute.
- (c) Reorganising the need for central facilities and with the provision of computing equipment to be used in 1978-80, the SRC should seriously consider installing a CDC STAR computer as the main source of computer power for this period.
- (d) Develop the subsidiary services needed to support large computer power and establish here again quite specific objectives.

If, however, in response to the National need, a P4 is to be worked up to a final production level, SRC must abandon a serious attempt to offer services of the kind outlined above. Some more limited objectives to provide a work-load would be necessary and a definite research programme established with ICL. In this event a serious and searching dialogue with the Computer Board would be necessary to ensure the continued and effective take-up of SRC responsibilities for those grant holders heavily dependent on computing.

Finally, it would be a mistake to mix the two activities above too closely, but of course the development of two mainstreams virtually independently is not ruled out.

OTHER GOALS AND ACTIVITIES

RECOMMENDATIONS

The choice of a suitable role for SRC in the work of the National Committee is a matter of some importance. The support of a research programme is necessary in order to ensure that the work is carried out in a manner which will be of maximum benefit to the community. It is suggested that the work should be carried out in a manner which will be of maximum benefit to the community. It is suggested that the work should be carried out in a manner which will be of maximum benefit to the community.

The SRC should move towards setting up a central laboratory with a planned programme of research. This programme should be carried out in a manner which will be of maximum benefit to the community. It is suggested that the work should be carried out in a manner which will be of maximum benefit to the community.

The SRC should move towards setting up a central laboratory with a planned programme of research. This programme should be carried out in a manner which will be of maximum benefit to the community. It is suggested that the work should be carried out in a manner which will be of maximum benefit to the community.

The SRC should move towards setting up a central laboratory with a planned programme of research. This programme should be carried out in a manner which will be of maximum benefit to the community. It is suggested that the work should be carried out in a manner which will be of maximum benefit to the community.