# THE SCIENCE RESEARCH COUNCIL'S CO-ORDINATED RESEARCH PROGRAMME IN DISTRIBUTED COMPUTING



G. Ringland

Rapporteur : Dr. P.C. Treleaven

The SRC has recently initiated a co-ordinated research programme in distributed computing systems. The reasons for initiating such a programme are discussed, as are the differences between this programme and more traditional support mechanisms. The aim of the programme is to improve the efficiency of use of scarce resources, both human and systems. The mechanisms proposed are two-fold: via the provision of channels of communication between groups, and the adoption of common techniques, tools and research vehicles where this will reduce unproductive time. It is also the intention that the programme should encourage collaboration with groups or individuals engaged in similar research outside the programme.

## 1. Background

The traditional role of the Science Research Council (SRC) in funding research in British Universities and Polytechnics is similar to that of many other funding agencies in the academic sphere, i.e. it responds with acceptance or rejection to applications for support of particular projects or programmes from members of its constituency. In many areas of research, however, it has become necessary to provide additional support, namely support of research vehicles, over and above the support for individual projects. Examples of such research vehicles include computer facilities such as Atlas, in the early days of computing, and the contribution of the U.K. through the SRC to the international high energy physics accelerators at CERN.

The arguments for the provision of research vehicles for use by University groups were initially those of the expense of the accelerator or computer, and the need for specialist support staff for its maintenance and operation. Advantages which have been observed to follow from the existence and use of the research vehicles include

- approval mechanisms for access to the vehicle which allow for extensive discussion of the proposed project within the community;
- creation of cross-university collaborations which allow small groups to participate in large projects;
- provision of an infra-structure so that research programmes can take a starting point of known techniques and tools.

These points were very much in the mind of the Computing Science Committee (CSC) of the SRC when it received, two years ago, three applications for research grants to support the building and operation of distributed systems. The amounts of money involved were large, by the standards of academic computer science research,

and the timescale proposed for each project was long enough to allow for building the configuration, developing and testing software, and running experiments. The questions naturally raised were:

- could a research vehicle be provided which met the needs of all three groups?
- alternatively, could one group act as hardware designers and builders, another design and build the software, and another design the applications?

The aim would be in either case to shorten the lead time before results might be expected from the research.

Analysis of the three applications involved showed that neither of these approaches was appropriate: the emphasis of one group was on building blocks for a hierarchical, statically loaded system with tight linkages, another proposed to investigate loosely linked systems in the office environment, the third intended to study dynamic migration of work load in a configuration of a few computers. In other words, the design of the research vehicle itself, and its organisation, was one of the central points to be resolved by the research. How, in this situation, could the CSC afford to support this area of research? Although the cost of hardware is decreasing, useful configurations still require substantial sums, especially if software is provided. And if the configuration does not support any existing software, the effort needed to provide this will be a major expense.

## 2. The Distributed Computing Programme (DCP)

The approach adopted by the CSC was to set up a coordinated programme of research in distributed computing. The aim of the programme was to

- promote relevant Computing Science research of high quality in a positive manner by co-ordinating the efforts of individual academic research teams;
- ensure the best use of funds at a time of financial stringency:
- achieve results of practical value to UK industry by directing research to a key area for the future.

The definition of distributed computing proved more difficult to formulate, but for the purposes of the programme, a distributed computing system was defined as one in which there are a number of autonomous but interacting computers co-operating on a common problem, i.e. including geographical networks of mainframe computers; arrays, hierarchies and rings of micro-processors; and novel forms of computer architecture with highly parallel algorithms for processing and storage allocation. The common thread to these topics is provided by the revolutionary decrease in the cost of processing power, so that the cost of processing is no longer the dominant factor in the provision of computing power. Many consequences follow from this change in emphasis. For instance, software engineering is revolutionised when some problems can be solved by more power in the processor. Similarly, transportability of software rather than its raw processing efficiency becomes an important criterion for many types of application.

The CSC was however more concerned with the implications for the design and building of systems, when programmed rather than programmable devices are the building blocks, and when architectures based on parallel processing are economic. It was concerned therefore to answer questions like:.

- what factors in the organisation of distributed systems affect the operational parameters such as reliability, throughput and response time?
- what language and operating system features allow for information processing on these parallel architectures?

The programme was set up in the academic year 1977-8, and is directed by a panel (given at Appendix 1). The role of the Industrial Co-ordinator is to concentrate on communication between the programme and the outside world, that of the SRC co-ordinator (Professor F.R.A. Hopgood of the Rutherford Laboratory) to concentrate on matters internal to the programme. How does the programme in practice differ from the extremes of responsive funding on the one hand, and mission-oriented research on the other?

The explicit conditions under which awards are made within the programme include stipulations that the investigators agree to talk about work in progress to other researchers in the programme at workshops or meetings of special interest groups. The existence of these communication channels was seen by researchers as a positive benefit of the programme, and there is evidence from the draft schedule outlined for next year (Appendix 2) that this will be a major source of stimulus.

Other advantages which have become apparent since the programme has been in operation include

- early and informal advice from co-ordinators and panel members on other relevant work and possibilities of collaboration/overlap, and on SRC procedures and timescales;
- the existence of an explicit interface mechanism to the programme via the co-ordinators has enabled a range of contacts to be made which would have less likely under traditional arrangements;
- the availability of information, to researchers, on the experience and kit of other groups in and outside the programme;
- the ability to attract extra funding into this research area by responding in an informed way to requests from funding sources for extra information on the context and nature of the proposed research.

The DCP has in fact attracted funding to enable us to establish a pool of hardware items, and some common software, to ease communication problems. Since we do not have the resource, or desire, to parallel existing network developments, we plan to use existing networks, such as the ICF network, in an ad-hoc manner until the availability of PSS (the Post Office's packet switched service), currently planned for the summer of 1979. The hardware items therefore include magnetic tapes and controllers, Diablo printers, and a pool of LSI11 configurations for loan to groups with needs for specific items. There is also a plan to mount and

support one or more parallel processing languages, for experiments in application and operating system construction.

In the next section, the research which the infrastructure is supporting is indicated by describing briefly the scope of projects included in the DCP, and in the final section four of the research vehicles are discussed. The aim of these descriptions is to make explicit the range and nature of the work in the programme, and to solicit aid and assistance in meeting the programme objectives.

## 3. Research Projects in the DCP

The research topics within the scope of the DCP have, for convenience, been categorised under five headings. These headings are used as a guide to the prime thrust of a team's work, rather than an attempt to limit its scope; and many teams find that problems in one area have implications for solutions to problems in others. The headings are:

- theory and languages
- resource management
- architecture
- operational attributes
- design and implementation methodology

### Theory and Language

Distributed computing, in the sense used above, is forcing a re-evaluation of the constructs used to date in computing science theory. Asynchronous and non-deterministic systems require new semantic schemes for their analysis and definition. While in an ideal world such schemes should precede the development of systems and languages, in practice these developments all proceed in parallel. Hence work on the expression of concurrency within programming languages and operating systems is in progress as well as studies of the foundation assumptions.

### Resource Management

Operating systems are concerned with different issues in distributed and centralised systems, but are needed in both. The problems which are more difficult to solve in the distributed environment are primarily those of access - to logical and physical resources such as data files, peripherals and programs controlled by another processor. Work in this area is using communications techniques ranging from wide band loop networks to direct memory access interfaces between components; the aim in each case is to study design trade-offs in strategies for controlling the overall system.

### Architecture

Two main issues are under discussion: the organisation of the computing elements, for instance into an array or hierarchy, and the nature of the interfaces. Work in the programme includes the investigation of an array of processors with both shared and private memory as a tool for data and signal processing, development of a

ring structured data flow computer (see below), and study of X25 based protocols in local and wide-area networks.

Operational attributes

The performance of distributed systems depends on different parameters from that of centralised systems - for instance, while functions may be duplicated on extra computers to increase reliability once an error has been detected, the problems of detecting errors and assessing the damage are more complex in distributed systems. Performance models developed for centralised systems may, perhaps surprisingly, be applicable without major modifications to distributed systems - this is the area tackled by researchers in two universities working together.

Design and Implementation Methodology

By developing software for 'real' applications, it is expected that techniques for specification of the system, for its implementation via a programming language, and for system testing will be investigated. The applications currently being used as test beds in this manner include data processing and numerical analysis. The other aspect of implementation methodology, which may in the long run be of dominant importance, namely that of the user interface to a distributed system, is being studied for the office environment explicitly.

Collaborations

The descriptions of research areas above was intended to provide a guide to the list of research projects listed in Appendix 3. Both the list and the preceding paragraphs may however give a wrong impression, which an outline of just one collaboration may correct.

The research activity into the implications of the architectural concept called dataflow - in which any operation may be executed as soon as all of the items of data required as input operands are available - is spread over four Universities. One group is primarily interested in the language implications of the architecture, another group is building a ring structured data flow machine for use by the collaboration, and the researchers from the other two Universities are working together on the constraints of, and problems in, the approach. The collaboration meets at frequent intervals for seminars and workshops, and believes it is ahead of 'state of the art' in solving many of the technical problems.

## 4. Research Vehicles

The purpose of including in this paper a short description of some of the research vehicles in the programme is to seek out potential users of such systems. In the descriptions below some indication is given of the timescale of hardware availability and the availability of a software environment. Contact may be made with the developers directly, or via the Co-ordinators, all of whose addresses are given in Appendix 4.

CYBA-M

This computer was developed by Aspinall and Dalgless, and is based on Intel 8080's with 16K of private memory and access to shared memory via buses for system files (image memory) and data (global memory). A "CYBA-80" computer controls the configuration, in that it can access the memory of any other processor, and it has perhipherals for program development attached to it. The CYBA-M has 16 processors, and is intended for applications in which there are fewer processes than processors. The applications under development include signal processing, in a BCPL-like language, and there are plans for data processing applications probably in a Pascal-like language.

Dataflow Ring

This computer will be built by Gurd and Watson over the next year. It will be intended to explore dataflow concepts rather than attain high speeds, being implemented in TTL logic. It will be driven by a PDP11/34 to handle I/O, and may be programmed in a language called Lapse currently running on a simulator. The ring will be used as one of the test machines for the work on evaluation of architecture for numerical analysis.

X-25 Protocols

The research effort under Kirstein has been associated for the last few years with the provision of Arpa-net and EPSS access facilities. The thrust of the work is now towards X-25 based protocols: mapping of X-25 protocols onto other networks, high level and terminal protocols for X-25 networks, network testing. The protocol software will be implemented in RTL/2 for a PDP11 and will be available, as is also access to Euronet or SATNET via the facility.

C-Net Ring

A ring network based on Ethernet is being used by Colouris to support his work on the man-machine interfaces at personal workstations for the office environment. The work stations will use PDP11's and software development is in C, the UNIX language. Experiments on the use of colour displays in information processing are in hand, but using non-standard hardware. Software support for the workstation system will be available in 1979-80.

## 5. Concluding Remarks

The author hopes that this talk will fulfil two functions:

i) A report to the academic community.

ii) A mechanism for establishing links with groups outside the programme.

Discussion

Mr. Shelness began by saying that while he enjoyed the talk, he did feel that things were not as rosy as they had been painted and certainly the effect of the programme on the three original research projects was that the delay put the research back considerably. So the effect of the co-ordination was not to improve the speed with which things happened, rather initially it slowed it down. Mrs. Ringland replied that everyone concerned accepted these criticisms. Mr. Shelness added that it was not clear to him whether things would have gone forward any differently had they not been co-ordinated. Although he welcomed the additional SRC funding for people to get together at workshops and conferences, he was not sure whether the massive super-structure of the programme was needed for this to go on.

Mrs. Ringland said she accepted a lot of what Mr. Shelness had said, but added there were advantages to a co-ordinated research programme. Researchers wanting to do work in an area can talk with people who know what other research is going on. The normal thing is that you have to write it all down in 95 copies and send it off and you either hear a yes or no. The DCS programme provides an informal iteration method, that is something not to be taken lightly.

Dr. Treleaven expressed his disagreement with Mr. Shelness, saying that as far as the people working in the area of data flow were concerned, the DCS programme had been extremely useful in helping the people to meet regularly and understand each others ideas. He quoted a recent example of a visit to Britain of Professor Arvind, a data flow researcher from the USA, where the DCS Panel Secretary, Mr. Witty had arranged for the group to spend the day with Professor Arvind at the Rutherford Laboratory.

Professor Page questioned the selection of a topic for a co-ordinated research programme and expressed a view that selection of such a topic normally resulted in a large proportion of the resources being channelled into it. Mrs. Ringland interjected to say that this can happen as a consequence of the area. Continuing, Professor Page said that the virtue of the previous SRC mechanism was that good ideas were not judged by the same group of people.

Professor Neuhold enquired why the SRC was standardising on American PDP/11 computers rather than British equipment adding that in Germany it was very difficult to justify getting foreign computers. In answering, Mrs. Ringland said that in 5 years 90% of the cost of any project will be the software. Any systems house in this country who sell PDP/11 based systems for, say, $50,000 will only pay DEC $1400 for the processor.

Professor Heath returning to the topic of the data flow group, said his group had considerable knowledge of this area and directed graphs in particular, but not had any interaction with the DCS apart from a visit by the Newcastle group.

In replying Mrs. Ringland said that support by committee does not always work, and that what was needed was a lot of financially independent young geniuses. In offering a solution,

Dr. Treleaven enquired whether the SRC might not give some initial funding to virtually anyone who asked, making it clear that the results would be critically examined and those found wanting would be exempt from future grants. Professor Page noted the case of the good shepherd and the bad shepherd. Dr. Treleaven said he knew the reference and offered to lend Professor Page a copy. Remarking on the suggestion, Professor Pyle expressed his opinion that the DCS already contained the flavour of such a mechanism.

APPENDIX I

DISTRIBUTED COMPUTING SYSTEMS

PANEL MEMBERS 1977-1978

I.M. Barron (Chairman)
Professor D. Aspinall
A.C. Conway
Dr. F.K. Hanna
D. Habditch
Professor C.A.R. Hoare
C.J. Hughes
J. McNeil
Professor I.C. Pyle
Professor B. Randell

PANEL MEMBERS 1978-1979

Professor I.C. Pyle (Chairman)
Professor D. Aspinall
A.C. Conway
Dr. F.K. Hanna
D. Habditch
Professor C.A.R. Hoare
C.J. Hughes
Dr. I.R. Taylor

APPENDIX 2

PROGRAMME OF MEETINGS UNDER DCP 1978-1979

There are three forms of meeting which are proposed for the year 1978-1979. These have been labelled special interest groups (SIG's), workshops and conferences.

A special interest group is an ad hoc grouping of researchers, probably mostly funded within the programme. The group meets as needed to discuss a common research programme, e.g. the provision of research vehicles for language work. SIG meetings are intended to be informal and to have less than 20 participants.

A workshop is a larger meeting, at which researchers in one area, e.g. languages, communicate the "state of the art" to workers in other areas. The intention initially was that workshops would be largely attended by researchers within the DCP, but experience of the first few have shown that the subject is such that to be effective such workshops need to include a wider range of speakers and participants. Workshops are hosted by University departments, who also undertake to produce proceedings.

Conferences are more formal gatherings than SIG's or workshops, presenting the public face of the DCP. The aim of conferences is to communicate the nature and scope of results, and their importance, to an audience of non-specialists in distributed computing. The intention is to organise conferences in conjunction with any friendly organisation - the first two will be held in conjunction with the IEE.

The draft schedule of meetings for 1978-9 is as follows:-

| 1978 | SIG | Workshop | Conference |
|---|---|---|---|
| September | | Languages (Warwick) | |
| October | Personal computers | | |
| November | Distributed architecture | | |
| December | | | |
| 1979 | | | |
| January | | | IEE/SRC |
| February | Operating systems | | |
| March | | | IEE/SRC |
| April | | Local networks (PORC) | |
| May | | Theory (?) | |
| June | | Distributed architecture | |
| July | | | Steelman (York) |

APPENDIX 3

## RESEARCH GRANTS AWARDED UNDER THE DISTRIBUTED COMPUTING SYSTEMS COORDINATED RESEARCH PROGRAMME

### Theory and Languages

Dr. P.E. Lauer, Computing Laboratory, University of Newcastle upon Tyne
Design and analysis of highly parallel distributed systems.
$35K over 3 years

Mr. A.J.G. Milner, Department of Computer Science, University of Edinburgh
Application of flow algebras to problems in concurrent computation.
$9K over 1 year

Mr. A.J.G. Milner and Dr. G.D. Plotkin, Department of Computer Science and Articifical Intelligence, University of Edinburgh
Semantics of non-deterministic and concurrent computation.
$37K over 3 years

Dr. P.E. Osmon, Department of Computer Science, Westfield College, London
Implementation of a high level data flow programming language.
$23K over 2 years

### Resource Management

Professor M.V. Wilkes and Dr. R.M. Needham, Computer Laboratory, Cambridge University
Distributed Computing using wide band communications.
$33K over 3 years

Mr. N. Shelness, Department of Computer Science, University of Edinburgh
Architecture for a multiple computer system.
$89K over 3 years

Dr. M.R. Sleep, Department of Computer Science, University of East Anglia
Instruction sets for data flow architectures: a comparative study.
$1.5K over 1 year

Dr. C. Whitby-Strevens and Dr. D. May, Department of Computer Science, University of Warwick
A building block system for Distributed Computing.
$63K over 3 years

### Architecture

Professor D. Aspinall, Department of Electrical and Electronic Engineering, University College of Swansea
Use of microprocessors in information processing systems
$90K over 4 years

Dr. J.R. Gurd and Dr. I. Watson, Department of Computer Science, University of Manchester
A ring structured data flow computer system.
$70K over 3 years

Professor P.T. Kirstein, Department of Statistics and Computer Science, University College, London
Communication protocols in the context of X-25 computer networks.
$208K over 4 years

Dr. B.K. Penney and Dr. M.S. Sloman, Department of Computing and Control, Imperial College of Science and Technology
Communications for distributed process control
$55K over 3 years

Professor B. Randell, Computing Laboratory, University of Newcastle upon Tyne
Investigation of the design of highly concurrent general purpose computing systems
$13K over 2 years

Operational Attributes

Dr. K.H. Bannett, Department of Computer Science, University of Keele
A feasibility study of closely-linked computers.
$6K over 1 year

Dr. I. Mitrani, Computing Laboratory, University of Newcastle upon Tyne
Modelling and performance evaluation for distributed computing systems.
$14K over 2 years

Professor B. Randell, Computing Laboratory, University of Newcastle upon Tyne
Reliability and integrity of distributed computing systems.
$219K over 4 years

Design, Implementation, Applications

Mr. D. Coleman and Ms. J.W. Hughes, Department of Computation, UMIST
Developing a programming methodology for multi-programs.
$12K over 2 years

Mr. G.F. Coulouris, Computing Laboratory, Queen Mary College, London
Distributed systems requirements for effective man machine interaction.
$72K over 3 years

Professor D.J. Evans and Dr. I.A. Newman, Department of Computer Studies, University of Technology, Loughborough
Investigation of the relationship between algorithm structure and parallel architectures
$18K over 3 years

Dr. F.K. Hanna, Electronics Laboratory, University of Kent
Distributed processing systems for interactive knowledge basis.
$24K over 3 years

Professor C.A.R. Hoare, Computer Laboratory, Oxford University
Software engineering
$81K over 4 years

Mr. I. Page, Computer Laboratory, Queen Mary College, London
A high quality display system for effective man machine interaction
$25K over 2 years

Dr. I.C. Wand, Department of Computer Science, University of York
Modula distribution and promulgation
$19K over 3 years

NOTE: The groupings reflect the main thrust of research in each project, but there are obviously instances where one research grant covers more than one field.

## Addresses

COORDINATORS

Professor F.R.A. Hopgood
Atlas Computing Division
Rutherford Laboratory
Chilton
Didcot
Oxon
OX11 OQX
0235-21900 Ext 372

Mrs. G. Ringland
INMOS
c/o 19 Loyd Road
Didcot
Oxon OL1 OX11 8JX
0235-813470

R.W. Witty
Technical Secretary
Atlas Computing Division
Rutherford Laboratory
Chilton
Didcot
Oxon
OX11 OQX
0235-21900 Ext 6218

B. Stephens
Joint Secretary
Computing Science Committee
Science Research Council
P.O. Box 18
Swindon
Wilts
SN1 5BW
0793-26222

GRANT HOLDERS

Professor D. Aspinall
Department of Computation
UMIST
P.O. Box 88
Manchester
M60 1QD
061-236-3311

Dr. K.H. Bennett
Department of Computer Science
University of Keele
Keele
Staffs
ST5 5BG

0782-621111 Ext 359

D. Coleman
Department of Computation
University of Manchester Institute of Science and Technology
P.O. Box 88
Manchester
M60 1QD
061-236-3311 Ext 2248

G.F. Colouris
Department of Computer Science and Statistics
Queen Mary College
Mile End Road
London
E1 4NS
01-980-4811 Ext 751

Professor D.J. Evans
Computer Studies Department
University of Technology
Loughborough
Leicestershire
LE11 3TU
0509-63171 Ext 297

Dr. J.R. Gurd
Department of Computer Science
University of Manchester
Oxford Road
Manchester
M13 9PL
061-273-7121 Ext 5553

Dr. F.K. Hanna
Electronics Laboratory
University of Kent
Canterbury
Kent
CT2 7NT
0227-66822 Ext 396

Professor C.A.R. Hoare
Oxford University Computing Laboratory
Programming Research Group
45 Banbury Road
Oxford
OX2 6PE
0865-58086

Professor P.T. Kirstein
Department of Statistics and Computer Science
University College
London
WC1E 6BT
01-387-7050 Ext 815

Dr. P.E. Lauer
Computing Laboratory
University of Newcastle upon Tyne
Claremont Tower
Claremont Road
Newcastle upon Tyne
NE1 7RU
0632-29233

Dr. I. Mitrani
Computing Laboratory
University of Newcastle upon Tyne
Claremont Tower
Claremont Road
Newcastle upon Tyne
NE1 7RU
0632-29233 Ext 271

Dr. R.M. Needham
Computer Laboratory
University of Cambridge
Corn Exchange Street
Cambridge
CB2 3QG
0223-52435 Ext 256

Dr. P.E. Osmon
Westfield College
University of London
Kidderpore Avenue
London
NW3 7ST
01-435-7141 Ext 521

I. Page
Department of Computer Science and Statistics
Queen Mary College
Mile End Road
London
E1 4NS
01-980-4811 Ext 751

Dr. B.K. Penney
Department of Computing and Control
Imperial College of Science and Technology
180 Queensgate
London
SW7
01-589-5111 Ext 2741

Dr. G.D. Plotkin
Artificial Intelligence Department
Edinburgh University
Hope Park Square
Meadow Lane
Edinburgh
EH8 9NW
031-667-1011 Ext 6443

Professor B. Randell
Computing Laboratory
University of Newcastle upon Tyne
Claremont Tower
Claremont Road
Newcastle upon Tyne
NE1 7RU
0632-29233 Ext 235

N.H. Shelness
The James Clerk Maxwell Building
The King's Building
Mayfield Road
Edinburgh
031-667-10181 Ext 2752

Dr. M.R. Sleep
Department of Computer Science
University of East Anglia
Norwich
Norfolk

Dr. P.C. Treleaven
Computing Laboratory
University of Newcastle upon Tyne
Claremont Tower
Claremont Road
Newcastle upon Tyne
NE1 7RU
0631-29233

Dr. I.C. Wand
Department of Computer Science
University of York
Heslington
York
YO1 5DD
0904-59861 Ext 5570

Dr. C. Whitby-Strevens
Department of Computer Science
University of Warwick
Coventry
CV4 7AL
0203-24011 Ext 2363