technical memorandum Daresbury Laboratory

DL/SCI/TM72T

RELATED SERC AND OTHER INITIATIVES IN NOVEL ARCHITECTURE COMPUTING

by

R.J. BLAKE and M.F. GUEST, SERC Daresbury Laboratory.

JANUARY, 1991

Science and Engineering Research Council
DARESBURY LABORATORY
Daresbury, Warrington WA4 4AD

ISSN 0144-5677

Related SERC and Other Initiatives in Novel Architecture Computing

R.J. Blake and M.F. Guest

Advanced Research Computing Group
SERC Daresbury Laboratory
Daresbury
Warrington WA44AD

Introduction

A growing number of agencies are funding, or proposing to fund, parallel and novel architecture computing initiatives. This paper is an attempt to draw together some of the relevant background information for the Novel Architecture Computing Committee (NACC) which oversees central Science and Engineering Research Council (SERC) Foundation Support to a number of Centres. The role of this support is to provide the infra-structure to encourage the widespread application of parallel and novel architecture computers to problems in science and engineering of interest to the SERC. Therefore, the main emphasis in this review is on programmes which support numerically intensive research applications. Most of the projects have lifetimes of about four years and there is considerable movement in the area from year to year. What we present here is a snapshot of novel architecture computing initiatives as of the summer of 1990. It should be stressed that the nature of some of these programmes may be restructured with the transfer of supercomputing activities from the Computer Board to the Advisory Board of the Research Councils. Contact addresses are provided for those who are interested in pursuing the details, or future progress, of any of the initiatives in more detail.

The SERC has Council-wide and cross-Board activities listed as items 1–2 below. The SERC is involved with other agencies in jointly funding the parallel/novel architecture computing projects listed as items 3–4 below. The Boards of SERC have their own activities listed as items 5–7 below. The Computer Board's initiatives are discussed in item 8. Some American activity is reviewed in section 9.

1. Novel Architecture Computing Initiative
2. Grand Challenge Machine
3. Central Computing Unit, Rutherford Appleton Laboratory
4. SERC/DTI JFIT Programme:
   a. Information Engineering Advanced Technology Programme
   b. Programme on the Application of Parallel Systems
   c. Engineering Applications of Transputers
   d. The Esprit Programme
5. Science Board's Computational Science Initiative
6. Science Board's Advanced Research Computing
7. Engineering Board's CFD Initiative
8. Computer Board's Initiatives
9. Center for Scientific Parallel Programming in the USA

Before presenting details of the initiatives it is perhaps worthwhile outlining some of the organisational details both within the SERC and between the SERC and other agencies.

Organisational Issues

In Figure 1 we show the relationship between the SERC and some of its computing committees. Each of the four SERC Boards has its own theory/computing sub-committees which are represented on the Scientific Computing Advisory Panel (SCAP). The role of the SCAP is to advise the SERC:

1. on future scientific computing developments and the planning of scientific computer provision;
2. on the central provision for scientific computing in the Laboratories: and,
3. on the progress of SERC scientific computing in relation to the the total provision for research computing by SERC, the Computer Board and other bodies.

The Novel Architecture Computing Committee advises the SERC, through the SCAP:

1. on support to chosen centres for the development and provision of techniques to apply novel architecture computers to problems in science and engineering, ensuring that this activity is complementary to and coordinated with the individual Board programmes in novel architecture and other computing; and,
2. on the progress made by centres or groups supported by central funding for novel architectures and to advise on how the scientific and engineering benefits could be enhanced.

The SERC partners the Department of Trade and Industry (DTI) in funding of a number of programmes. The Joint Framework for Information Technology (JFIT) was set up by the DTI and the SERC in the autumn of 1988 to bring together appropriate activities in Information Technology (IT) research, technology transfer and education and training. A key element in this coordination is the advisory committee structure with members drawn in equal numbers from the industrial and academic communities. The Committee structures of the DTI's Information Technology Directorate (ITD) and SERC's ITD, funded through SERC's Engineering Board, are echoed in the JFIT committee structure (Figure 2) which is headed by the Information Technology Advisory Board (ITAB). ITAB's remit is to advise DTI and SERC on all aspects of IT policies and programmes, both UK and European programmes such as ESPRIT, and to act as the executive authority for the award of SERC grants. The main committee concerned with parallel computing is the Parallel and Novel Architecture Sub-Committee of the Systems Architecture Committee.

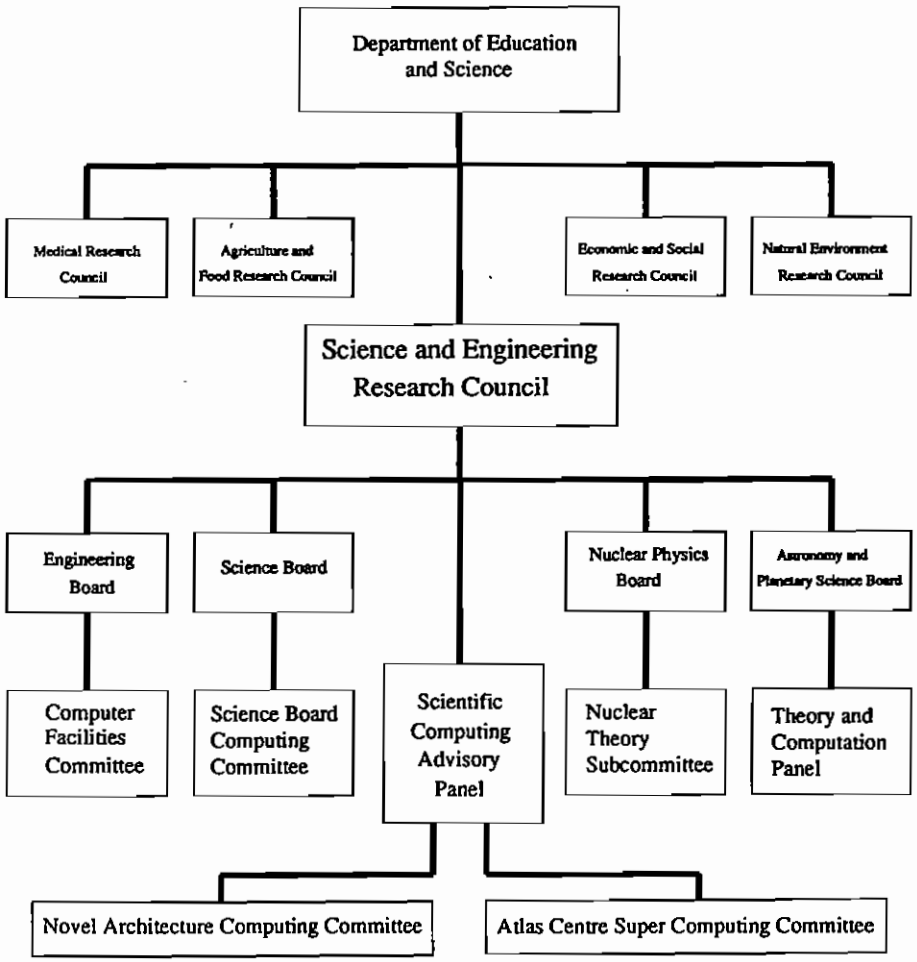

Figure 1: The SERC's Computing Committees.

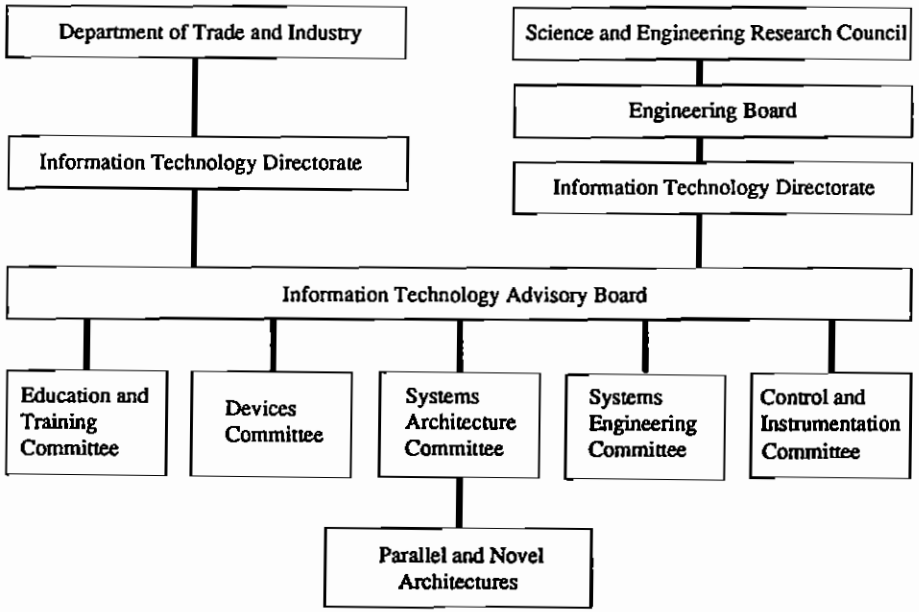

Figure 2: The Joint Framework for Information Technology

The precise role of the other main agency, the Computer Board, in funding computing activities is currently under discussion.

1. Novel Architecture Computing Initiative

In July 1989 Council agreed to provide £2M over 4 years to assist the development of scientific and engineering applications of novel architecture computing at chosen HEI Centres. This central SERC funding is overseen by the Novel Architecture Computing Committee. In the following we briefly review the background to, and purpose of, the Novel Architecture Computing Initiative, and discuss what has been established.

One of the main objectives of the proposal to establish an IRC in parallel computing was to encourage dialogue and coupling between the computer science 'providers' and computational science 'users' of parallel and novel architecture computers. In the event it was decided that a single IRC was not the best way to proceed. A subsequent Town Meeting concluded that there was a need to inject funds to allow for the support of several rolling grants to principal groups in the area. The IRC exercise had confirmed the particular strengths and potential of the groups at Edinburgh, Manchester and Southampton. These three Centres were invited to submit proposals for Foundation Support funding at the level of £600K over 4 years.

The intention of Foundation Support is to provide a level of infra-structure funding to the Centres to provide a facility and associated support for SERC peer-reviewed applications users. Foundation Support was agreed by the NACC to cover (but not necessarily wholly):

1. purchase of hardware and software necessary to provide a facility for users;
2. maintenance costs of hardware and software;
3. personnel to provide an operational base, to maintain and develop hardware and software, and to assist applications users.

Support for research projects carried out by SERC users of the facilities, whether from within the Centre or from outside, is specifically excluded from Foundation Support. Such projects are to be supported by the the appropriate Subject Committee.

The proposals from the three invited Centres focussed on the provision of medium to long term funding of man-power as being the most important factor in their plans to support the development of scientific and engineering applications. Manchester proposed to use the Foundation Support to facilitate a coordinated programme of work in participating universities in the original IRC North-West consortium. A series of application oriented experiments will be conducted with the aim of understanding and classifying novel architecture computing systems and methods for applying them. The Southampton proposal focussed on providing support for projects in areas of Advanced Systems Engineering, generalising the communications abilities of the transputer architecture and attempting to demonstrate theoretical analyses of parallel computation in practical terms. The

Edinburgh proposal focussed on establishing the Edinburgh Parallel Computing Centre using the Foundation Support to investigate the fundamental aspects of parallel computing. This will lead to the development of user support tools aimed at solving the practical problems of parallel computing.

The NACC considered the proposals for Foundation Support at its first meeting in December 1989 and recommended that they be funded more or less in full. Contracts to administer the awards are handled through Daresbury Laboratory, the commencement date was 1 April 1990. Problems with recruitment have caused the start dates for the Southampton and Manchester projects to slip to 1 August 1990 and 1 October 1990 respectively. The Centres have been asked to submit detailed workplans outlining their aims and objectives over the four year period for which funding has been agreed. A set of quantitative performance indicators which will allow the Committee to isolate the added value of Foundation Support are under discussion with the Centres.

One of the main aims of Foundation Support is to encourage the widespread application of parallel and novel architecture computers to research problems in science and engineering of interest to SERC. The Vice-Chancellors, Principals and Directors of Polytechnics have been informed of SERC's actions and invited to express interest from members of their institutions in taking part in the activity, either as new users of the facilities or as groups who might wish to collaborate in the development of the facilities. The NACC has established a mechanism whereby it comments on grant applications which seek to exploit Foundation Support and is actively ensuring that its activities are complementary to those of the individual Board's. The NACC has recognised that the Centres receive funding from a variety of different sources and has invited presentations from the Chairmen of other initiatives in this area to discuss programmes and future plans.

For further details contact:

1. Dr. R.J. Blake, Secretary NACC, Advanced Research Computing Group, SERC Daresbury Laboratory, Daresbury, Warrington WA44AD.

2. Grand Challenge Machine

The possibility of procuring a Grand Challenge machine arose from the unsuccessful Nuclear Physics Board proposal to secure funding, through an additional bid, for the provision of a dedicated concurrent vector super-computer for the solution of problems in Quantum Chromodynamics (QCD). The original proposal was for a machine with 256 vector nodes in a Meiko Computing Surface, each node containing 2 T800 transputers for communications, an Intel i860 chip for vector processing and 8 MBytes of memory. The machine was to be sited at Edinburgh, and the cost to SERC was to be of the

order of £2.85M. An additional capital provision was to be sought from the European Commission for the rest of the financing.

In February 1990, representatives from the Boards and the Scientific Computing Advisory Panel met to discuss the best way forward for possible provision of a QCD machine. Nuclear Physics Board had decided to provide £0.6M for a parallel computer for QCD calculations. A suitable machine would cost about twice as much and the possibility of Engineering Board and Science Board making use of such a machine was therefore being explored. A Working Group, chaired by Prof. P. Burke, was asked to prepare a case for a highly parallel computer to be shared between the Boards but dedicated to a small number of users. Funding from the Science and Engineering Boards would be required at the £0.3M level.

The working group came out in strong support of the provision of a Grand Challenge machine which would be used by Science Board to implement the Car-Parrinello molecular dynamics technique to investigate:

- hydrogen in metals,
- the simulation of MBE growth,
- surface structure and surface chemistry,
- silicates at high pressure and temperature corresponding to the earth's mantle,
- fluids and clusters of metallic elements,

and by the Nuclear Physics Board:

- to investigate the systematics of full QCD calculations,
- to investigate the dynamics of confining gauge theories,
- to explore aspects of the electro-weak theory.

The current proposal is for a machine with 64 twin transputer plus i860 vector nodes in a Meiko Computing Surface with 16 MBytes of memory per node and 4 GBytes of disk space. The machine is to be sited at Edinburgh. The capital cost is of the order of £950K with Nuclear Physics Board contributing some £600K. Science Board, through its Chemistry and Physics Subcommittees and the Materials Commission, is contributing some £250K, and Edinburgh University is contributing £128K. Maintenance costs in the region of £132K have yet to be met in full. The Novel Architecture Computing Committee (NACC) is supporting the proposal by providing £40K towards these maintenance costs and through its Foundation Support programme, from which manpower has been committeed to the main applications projects. European funding is being sought to expand the machine up to 128 nodes.

For further information contact:

1. Prof. D.J. Tildesley, Department of Chemistry, University of Southampton, Southampton SO9 5NH.
2. Dr. R.D. Kenway, Department of Physics, University of Edinburgh, The King's Buildings, Edinburgh EH9 3JZ.

3. Central Computing Unit, Rutherford Appleton Laboratory

The two supercomputers at the Atlas Centre are both shared memory multi-processors. The Cray X-MP/416 has four processors sharing a common fast memory of 16 million 64 bit words, each processor having a peak speed of 235 Mflops. The maximum user memory allocation on the Cray X-MP is 15 million words or 56 MBytes.

The IBM 3090/600E-6VF has six processors each of a peak speed of 116 Mflops and a somewhat more complicated memory structure. Each processor has a cache memory in front of a 256MByte main memory and a 1GByte extended memory that is used for paging. The user has available a maximum job size of 999 MByte under the VM operating system (or 2 GByte under MVS/XA, although this is not currently available at Atlas). The IBM has a virtual memory architecture and problems of data residence and cache coherence are handled by a mixture of hardware and operating system calls so that the user does not need be concerned with memory management.

Both Cray and IBM machines have compiler directives and Fortran extensions to support multi-processing both at the loop level (microtasking) and at the level of user initiated processes (macrotasking). Initial evaluation of both systems on simple multi-tasking constructs suggested that the the IBM routines provide somewhat more flexible and efficient multi-processor support. Both systems provide the necessary semaphores and locks to protect critical data references but needless to say the syntax is not yet standardised.

The IBM 3090/600E installation is supported by a joint study contract with IBM which requires the development of algorithms and methods to exploit both the vector and multi-processor features of the machine. Fourteen 'strategic users' have been selected with problems which are both scientifically exciting and demanding in terms of machine resources and who will therefore be particularly involved in parallel methods.

On both machines the benefits to the user and to the operation of the centre from multi-processing on a single job are greatest when the job occupies an appreciable part of a major system resource such as memory and it is then important to minimise the time that the job is resident in the machine.

A major effort has now gone in to multi-processing on the Fine Resolution Antarctic Model (FRAM), one of the large NERC Community Projects. The FRAM model is about 3.5MWords in size and on the X-MP/48 occupies about half of the available memory. The throughput of the FRAM project is limited by competition for memory from other jobs of a similar memory size, multi- processing allows the FRAM program to use all available CPUs while it is memory resident.

The structure of the numerical model used in the FRAM code is such that the updating of the physical variables is done in two nested loops corresponding to latitude and longitude. The inner loop is a vector triad and chains the add and multiply units while the outer loop is a replication across the grid and is ideally suited to multitasking. Other parts of the code such as the disk IO can be parallelised with a little more effort.

The work of the FRAM group has been very successful in improving turnaround in an empty machine, the wall clock time improves by a factor of more than three on a four processor machine, but the results in a loaded machine with variable job mix were surprising. Multitasking enabled the job to use its allocation of CPU time more quickly but the useful work done did not increase at anything like the same rate! The reasons for this behaviour have been explained by Cray Research and expose a weakness in the way the system scheduling deals with micro-tasked (ie loop level multiprocessing) jobs.

The acquisition of CPUs is done at the start of a time slice and all available CPUs are acquired for the job, the actual number depends on other work in the machine. If the multi-processing parts of the code end before the end of the time slice then the 'spare' CPUs are not released and are accounted to the job. If a job is not totally multi-tasked then it is charged for a variable amount of idle CPU time. The logic in implementing this method was its speed, very little system overhead is involved in setting up a microtasked loop, and its intended use for improving wall clock times in a time critical environment such as weather forecasting.

On a much more optimistic note the UNICOS scheduler supports a newer form of multitasking known as 'auto-tasking' which requires the user only to set a flag for the compiler. Autotasking is claimed to increase throughput on the X-MP/48 by 10-15% by utilisation of otherwise idle CPUs and the charging for idle CPUs is now eliminated.

On the IBM3090 with its capability to run FORTRAN jobs up to 999 MByte the need for multitasking is less obvious. However one of the 'strategic users', Ross Bollens from UCLA working in collaboration with the AMPTE project in the UK has a plasma simulation which runs in 900MByte and multitasking is important in getting the necessary throughput. On the six processor 3090 the elapsed times are about five time less than for a single processor.

From April 3rd- April 6th a course on high performance computing was held at The Coseners House Abingdon, sponsored by IBM as part of our Joint Study Agreement. The course described both the vector and parallel elements of supercomputing and several exercises were provided including a competition. The course was heavily oversubscribed with about 100 applicants for 40 places. The attendees were unanimous in finding the course very valuable and have asked for it to be held regularly. It was amusing that one group of two people on the exercises achieved speed-ups on the matrix multiply of 50% more than the official answer!

For further information contact:

1. Dr. R. Evans, Advanced Research Computing Unit, The Atlas Centre, Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire.

4. SERC/DTI JFIT Programme

At its inception in 1988, the ITAB, the main coordinating committee for the JFIT programme, inherited a number of projects in various states of maturity. Probably the most relevant to the novel architecture computing area is the Information Engineering Advanced Technology Programme (IEATP) for which the first round of proposals were received in autumn 1988. As discussed in section 4.a, priority was given to work which facilitates the implementation of applications on parallel systems. The porting of specific applications was specifically excluded and is the subject of a separate initiative detailed in section 4.b.

The DTI is carrying out a number of technology transfer programmes as part of its Research and Technology Initiative (one of the elements of the Enterprise Initiative). The Advanced IT Technology Transfer programme, which amongst other things aims to ensure rapid uptake by UK industry and commerce of efficient and effective IT systems embodying the benefits of advanced computer architectures. An important part of the programme consists of maintaining a major international centre of excellence at the University of Edinburgh. Other specialist centres to promote particular aspects of parallel and novel architectures have been set up at Bristol, Southampton, Strathclyde, Newcastle, QMW and ULCC. The JFIT programme also supports the Transputer Initiative detailed below in section 4.c.

Inevitably, European programmes now play a much greater role, and the JFIT therefore also covers the support offered by UK Government towards European IT programmes, and particularly ESPRIT which is reviewed in section 4.d.

4.a Information Engineering Advanced Technology Programme

Within this programme the PNA sub-committee has focussed attention on understanding how best to exploit parallel systems by concentrating on:

1. the underlying theory and understanding of emerging system requirements and of related viable solutions;
2. the investigation and evaluation of novel architecture concepts, initially by simulation and analysis; and,
3. the methodologies, techniques and aids required to implement and then exploit ideas in a cost effective, timely and competitive manner.

With some £800k available in 89/90, a dozen projects were funded in the following areas:

- fault tolerant and totally verified systems
- workbench for SIMD machines and a SIMD/MIMD interconnect
- parallel object-oriented languages for transputer based systems and fortran for scalably parallel systems
- declarative programming and knowledge based systems
- neural networks

No major change was made to the themes in the workplan for the second call for projects to be launched in 1990.

For further information contact:

1. Dr. A. Bryden, SERC, Polaris House, North Star Avenue, Swindon SN2 1ET.
2. Dr. S. Reid, Department of Trade and Industry, Kingsgate House, 66–74 Victoria Street, London SW1E 6SW.

4.b Programme on the Application of Parallel Systems

This a a four-year £40M programme, with £15M of Government support (DTI £9M, SERC £6M), to encourage the commercial application of parallel computing. The aim is to facilitate the growth of a strong and broad base of real industrial applications on parallel systems in the UK. This is to be achieved by a four year programme supporting LINK style collaborative industrially based projects covering a broad spectrum of parallel applications. Projects will be sought to cover a broad range of applications areas, such as computational fluid dynamics, financial modelling, remote sensing and image processing.

The research work will be done at number of centres, to be based at universities or polytechnics, which will provide industry with advice and assistance in the application of parallel computing. The role of the centres are:

1. to establish and maintain collaborative arrangements with UK registered industrial user and/or supplier firms,
2. to provide advice and guidance on the matching of particular parallel architectures to specific applications,
3. to provide facilities and assistance for the assimilation and dissemination of the experience gained in application development,
4. to maintain libraries of reusable algorithms, tools and techniques.

The centres will operate commercially and become self-financing after four years. Support for the projects will be available during the first four years. Grant allocations to the centres, based on total centre plus industrial collaborators project costs, will be provide on a tapering basis: up to 50% in years 1 and 2; 35% in year 3 and 20% in year 4. The standard LINK rules will be applied to the centres and industry eligible costs.

Outline proposals were received at the end of July 1990 and during mid-September a short list of centres:

- Bristol Polytechnic and Bristol University consortium
- Edinburgh University
- Liverpool University, Manchester University and Daresbury Laboratory consortium
- London consortium including: Queen Mary, Imperial College, University College London and the University Computer Centre
- Northern Ireland Transputer Centre including Queen's University Belfast and the University of Ulster
- Oxford University in association with Rutherford Appleton Laboratory
- Southampton University
- Strathclyde University
- Thames Polytechnic

were invited to prepare detailed proposals for submission at the end of December 1990. The selected centres will be notified by the end of February 1991 and are expected to start-up from April-October 1991.

For further information contact:

1. Dr. A. Bryden, SERC, Polaris House, North Star Avenue, Swindon SN2 1ET.
2. Dr. S. Reid, Department of Trade and Industry, Kingsgate House, 66–74 Victoria Street, London SW1E 6SW.

4.c Engineering Applications of Transputers

This is a joint initiative between the SERC Engineering Board and the DTI. It is now a 5 year, £2.9M programme which commenced in February 1987. The Engineering Board is contributing 60% of the funding.

The main objectives of the programme are:

1. to promote awareness of the potential of the Transputer, and associated technology within the SERC supported community and UK industry;
2. to enable researchers to acquire the techniques, development tools and systems software for using Transputers in a quick and cost-effective manner;
3. to promote high quality research using Transputers without unnecessary duplication;
4. to transfer as early as possible the benefits of the research to UK industry.

The Programme is coordinated by the Rutherford Appleton Laboratory (RAL). The main components of the Programme are:

a. an Academic Loan Pool of Transputer hardware and software available on a pump-priming basis for periods of up to one year - to date a total of 125 loans to 60 different HEIs have been affected;
b. six Regional Transputer Support Centres (at Belfast, Liverpool, RAL, Sheffield, Southampton and Strathclyde) which are the focal points for UK industry - services offered include basic awareness, advanced training courses, hands-on use of a range of Transputer hardware and software, hot-line support, and consultancy;
c. two Associated Support Centres at Daresbury Laboratory and the University of Ghent, Belgium (the first example of the planned move to a European dimension);
d. funding for specific key development contracts to address the needs for vital base software to facilitate the actual exploitation of Transputers in real engineering applications - a total of 17 such contracts have been let to date;
e. a Software Exchange Library, now operated by the Liverpool Regional Centre, containing software developed under the Loan Pool and the development contracts as well as other software made available to the initiative by suppliers and academics - this Library is available free to all UK academics;
f. a monthly Mailshot with a current circulation of 2800 worldwide - now seen by many suppliers as the 'definitive Transputer Journal';
g. three Transputer Applications Community Clubs in the following areas:
    - Image Processing and Pattern Recognition
    - Molecular Modelling
    - Real Time Control and Simulation

h. an annual Workshop addressing a specific key issue - to date the following subjects have been covered:

   - Software Development Environments
   - Applications
   - Software Standards

i. an annual International Conference on Transputer Applications with an associated Exhibition - the first was held in Liverpool in August 1989 with 350 delegates plus an additional 450 people visiting the Exhibition. The Exhibition attracted 35 suppliers and was the largest of its kind ever held. The Conference and Exhibition this year (Transputer Applications 90) is being held in Southampton on 11-13 July. A total of 115 submitted abstracts were received for the parallel sessions and 75 of these have been accepted for presentation in a total of 5 parallel streams throughout this Conference.

For further information contact:

1. Dr. M.F. Jane, Informatics Department, Rutherford Appleton Laboratory, Chilton, Didcot, Oxon OX11 0QX.

4.d The Esprit Programme

Within the European community many of the research projects investigating the design of parallel computing systems are funded as part of the European Strategic Programme for Research and Development in Information Technology (ESPRIT). ESPRIT I, a five year research programme starting in 1984, with approximately 3000 MECU of funding (1ECU approximately £0.70) and addressed five research areas:

1. Advanced Information Processing,
2. Microelectronics,
3. Software Technology,
4. Office Systems,
5. Computer Integrated Manufacturing.

The second phase of the programme, ESPRIT II, to run for four years from 1989 will continue with the same broad areas of work

The ESPRIT programme funds projects that emphasise research and development in addition to those that emphasise longer term research. All projects are intended as pre-competitive research. The composition of a project varies but typically 6 companies and universities, from different member countries, participate with funding at an average level of 10 MECU over three years. Within the ESPRIT programme there are a number of

large research projects concerned with parallel computing systems. The main emphasis has been on constructing machines for either symbolic or numeric processing. Details of the full ESPRIT II programme are still sketchy; here we concentrate on the numeric processing projects which are probably of most relevance to scientific and engineering applications:

1. GENESIS: aims to develop a highly parallel architecture for very high-performance computing. The GENESIS architecture is a concept for a family of supercomputers with the following goals:

   - high performance obtained through tightly coupled MIMD-SIMD distributed memory architecture,
   - cost effectiveness achieved by exploiting the latest advances in VLSI technology,
   - scalability and extensibility,
   - long product life, high reliability,
   - exploitability, through the provision of a programming environment.

   The aims of the definition phase are to specify: a complete architecture, system requirements, performance goals, the development of applications software, to initiate the development of the next generation of components and to develop a business plan for exploiting the project's results. Academic groups from the universities of Oxford, Southampton and Liverpool are involved in this project.
2. SUPERNODE II: is an extension of the ESPRIT I SUPERNODE project which was concerned with developing a low-cost high performance numeric parallel computer which could be tailored to a user's requirements both in terms of cost and in terms of topology. A SUPERNODE consists of 16 T800 floating point Transputers, with extended local memory, and the machine can be configured in domains varying in size from 1 to 64 SUPERNODES. The ESPRIT II project aims to develop an appropriate operating system and environments for general purpose parallel computers, for a range of tasks. Libraries of numerical routines will be written and image synthesis, simulation of heterogeneous systems, CAD for VLSI, and oil reservoir simulation applications will also be implemented to provide a testbed for the software developments. Academic groups at the universities of Liverpool and Southampton are involved in this project.
3. ENSSEMBLE: is aiming to define, develop, experiment and evaluate cost-effective integrated hardware-software computers composed of up to 100,000 homogeneous processing elements with regular interconnect structures. This project will investigate the development of fine-grain massively parallel systems analogous to the Connection Machine, perhaps making use of VLIW technology but using the AMT-DAP as a baseline architecture.

Besides the main projects, ESPRIT II has a series of Basic Research Actions which concentrate on fundamental research. A Parallel Computing Action inviting universities and

research institutions to make proposals for software projects exploiting parallel computer architectures was announced in April 1989. The total budget was expected to be about 2 MECU with which it was expected that about 20 individual actions would be launched. The purpose of the initiative was to further develop:

1. command of the state of the art in exploitation of parallelism in computing systems;
2. awareness and practice at graduate and undergraduate level, for subsequent technology transfer into industry;
3. maximum synergy between European computer development activities and software development activities;
4. contribution of academic and research organisations to innovation and exploitation of advanced technology developments.

Proposals were invited for basic software (load balancing, performance monitoring), tools (programmers parallelisation aids), environments and applications (CAD/CAM, general process modelling) for parallel computing systems with at least 16 processors.

For further information contact:

1. Dr. D. Worsnip, SERC, Polaris House, North Star Avenue, Swindon SN2 1ET.
2. Dr. A. N. Refenes, Department of Trade and Industry, Kingsgate House, 66–74 Victoria Street, London SW1E 6SW.
3. ESPRIT PROGRAMME — Parallel Computing Action, Commission of the European Communities, DG XIII Telecommunications, Information Industries and Innovation, A25 Room 5/3, rue de la Loi 200, B — 1049 Brussels.

5. Science Board's Computational Science Initiative

The Computational Science Initiative (CSI) of the Science Board is intended to meet the needs of the Board's community for local (distributed) high performance computing as recommended by the Forty Report on Advanced Research Computing. The major aim of the Initiative is to gain new knowledge and understanding in fundamental science through scientific computation, meeting the local advanced research computing needs of individual workers, departmental groups and local consortia in a coordinated initiative across Science Board's interests. The Initiative provides for the development of novel code and algorithms for vector and parallel machines, crucial factors in the effective utilisation of these resources, and fully integrates the activities of the Collaborative Computational Projects (CCPs) into the Initiative through liaison with the Daresbury Laboratory.

The case for the Initiative was prepared by a multidisciplinary panel chaired by Professor D.J. Wallace (Physics, Edinburgh) and reporting to the Science Board Computing

Committee (SBCC). It proposed a 5 year programme, which would require about £21M to provide host and attached processors at some 24 sites, approximately 125 workstations and personnel support. The Initiative was approved by Science Board in October 1986, when it was given £1M pump-priming for the 1986/1987 session. It was announced to the Board's academic community in January 1987 and the first awards were made in July 1987. The Science Board allocated a further £1M to the initiative in 1987/88, £2M in 1988/89, and £2M in the 1989/90 session. The strategy of asking for subject committee contributions (30%) to the funded awards has increased the total sum allocated by the Initiative to some £7.5M. Bids into the Initiative from the SB community have in each funding round greatly exceeded the funds available; in the present round 125 applications totaling some £18M have been received against the £2M available.

The equipment provided to date may be broadly classified into the following categories:

1. Minisupercomputers, typified by the Convex C210, Alliant FX/4, FPS M64/60. Within the numerically intensive modelling work associated with many of the awards, these machines are typically producing 20% single processor X-MP performance.
2. Processor array machines of the DAP architecture, which in applications with a significant integer and logical component can exceed X-MP performance.
3. High performance chips, in particular the T800 transputer and more recently the Intel i860, as building blocks for flexible parallel arrays such as the Meiko Computing Surface and Intel iPSC/860.
4. Superworkstations, such as the Ardent Titan and Silicon Graphics machines for those applications where real-time visualisation is computationally demanding
5. Networks of graphics workstations, typified by the products from SUN and Orion. The lower end offerings are often involved in hosting transputer-based products, or are being used for the pre- and post-analysis of results from central facilities. Several microVAX based networks are also being exploited in similar fashion. The higher end workstations, such as the SPARC-based SUN-4 with graphics accelerators, are associated with more compute intensive activities such as computational fluid dynamics.

The main achievements of the Initiative to date have been summarised in the document "Computational Science Initiative: A progress report, July 1989" (D. Fincham and M.F. Guest, ISBN 1 870669 99), pointing to the scientific productivity generated from the Initiative. The local advanced research computing needs of individual workers, departmental groups and local consortia have been addressed by SBCC through the provision of funds to thirty-four universities and colleges and the award of fifty-six research grants, enabling some twenty major systems to be purchased and about forty workstations. These figures are expected to increase to perhaps thirty systems and sixty workstations following the allocation of awards in the 1989/90 session at the SBCC meeting in June. The

current breakdown of grants to broad subject areas is twenty-nine to Physics, fourteen to Mathematics, nine to chemistry and four to Biology. The initiative continues to play a vital role in supporting the development and scientific exploitation of novel architecture computing. In addition to the timely provision of support for the Edinburgh Concurrent Supercomputer (ECS), some 14 awards covering a range of machine architectures have been provided. These are dominated by transputer-based distributed memory MIMD machines (Bath, Belfast, Durham, Glasgow, ICST, Keele, Kent, Lancaster, Open University, QMC, Sheffield, Southampton) and the massively parallel bit processor SIMD machines typified by the AMT DAP (Bristol, Cambridge). This activity has enabled many sites to benefit from the early availability of leading edge hardware, these sites now being heavily involved in the development of concurrent computer software and in parallel algorithm development, crucial factors in the effective utilisation of these resources.

Support for the novel architecture work is provided by the Initiative's Parallel Computer User Group coordinated by Daresbury Laboratory. The Group's activities are discussed in the next section. Valuable links have been established with both the Computer Board and Engineering Board, and are expected to feature in future co-ordination of the Initiative. Coordinators for the Initiative to date have been Dr. M.F. Guest, Head of the Advanced Research Computing Group (ARCG) within TCS Division at Daresbury, and Dr. D. Fincham (Lecturer in Computational Physics, University of Keele). Increased coordination of the Initiative is in hand, with the establishment of a CSI Panel, reporting directly to SBCC, and the appointment of a 'full-time' coordinator (Dr. M. Allen ,Bristol), with on-going support from M.F. Guest and the ARCG at Daresbury.

For further details contact:

1. Dr. M. Allen, H.H. Wills Physics Laboratory, Royal Fort, Tyndall Avenue, Bristol BS8 1TL.
2. Dr. M.F. Guest, Head of Advanced Research Computing Group, SERC Daresbury Laboratory, Daresbury, Warrington WA44AD.

6. Science Board's Advanced Research Computing

Funded by Science Board, the Advanced Research Computing project at Daresbury Laboratory is focussed on realising the potential of parallel computers and developing methods for visualising the results of large scale scientific and engineering calculations. Visualisation activities are centred around superworkstations like the HP/Apollo DN10020, the Stardent 1520 and the Silicon Graphics 4D/220-GTX Power Series machines. As of March '90, parallel computing activities on distributed memory MIMD machines have been centred around:

1. The Intel iPSC/2 hypercube, configured with 64 scalar 386-based nodes, each with 4 MBytes of memory (32 of which are coupled to 1M Byte vector boards), and 2 GBytes of node-addressable disc space. Remote hosting software allows host programs running on the Convex C220 or superworkstations to treat the cube as an effective co-processor. The Fortran/Unix environment on the hypercube has allowed rapid parallel implementations of a wide range of applications codes summarised below.
2. A somewhat modest T800-based Meiko M10 Computing Surface currently configured with 12 T800 transputers (with a mixture of .25 MBytes to 4 MBytes of memory) and a T414-Mass Store Board with 8 MBytes of memory.

The Intel machine has been jointly funded by Science Board and a collaborative project with ICI plc. Part of the Meiko kit has been loaned by the Engineering Board/DTI Transputer Initiative (in which Daresbury is an Associated Support Centre), with the rest funded by Science Board to support activities on similar kit awarded to university groups through the CSI.

Both the Intel and Meiko installations have recently been significantly upgraded. A second iPSC machine, a 32-node iPSC/860 was installed in July 1990. Based on Intel's i860 64-bit microprocessor, the machine has already proved an order of magnitude faster than the present iPSC/2. The Meiko Computing Surface has been similarly upgraded, through the provision of 4 MK086 boards (4 T800's plus 2 i860's per board) as part of the Engineering Board's CFD Initiative described in section 7 below.

The Advanced Research Computing Group (ARCG) at Daresbury was established to provide a base of knowledge and experience in the application of parallel computing to problems of scientific interest, particularly in those areas covered by the CCPs supported at Daresbury. This involves the evaluation and benchmarking of new computer products, the development of parallel algorithms, the implementation in parallel form of standard CCP programs, the development of systems software to aid scientific programming on parallel machines, and the integration of visualisation techniques into parallel number-crunching.

Many of the CSI awards have been for parallel computers in subject areas relevant to the CCPs, and much of the ARCG work centres on providing direct support for the Initiative's Parallel Computer User Group. This Group has initially been organised into a set of Project teams, each Project covering an area of computational science that looks to parallel processing as a focus for its developments. Nominated representatives from within the Theory and Computational Science Division at Daresbury are responsible for coordinating the activities of the Group, liaising with related university and industrial groups, organising workshops and symposia, and structuring the software development programme of the Group, ensuring that both timescales and deliverables are met. The

Projects currently in place have effected parallel implementations of a wide range of applications codes, for example:

1. Accurate Molecular Modelling: direct SCF, CI and gradient capabilities of the GAMESS electronic structure code and macro-molecular energy refinement (AMBER);
2. Atomic and Molecular Processes: EIKONXS charge transfer package, CPBX Coulomb Projected Born Exchange double and single differential cross-section code;
3. Surfaces and Interfaces: NEWPOOL surface photo-emission code, layer KKR-CPA electronic structure code and the XANES program;
4. Structural and Finite Element Analysis: 2D/3D transient finite element models for elastodynamic problems;
5. Electronic Structure and Materials Science: SCF-ASA-LMTO self-consistent-field bulk electronic structure calculations;
6. Molecular Dynamics and Computer Simulation: systolic loops, replicated data and link-cells algorithms;
7. Computational Fluid Dynamics: Assessing and developing CFD packages;
8. Programming environment and Software Tools: development of the FORTNET harness.

Other Projects in Visualisation, Biological Sciences and Large Molecules and Energy Minimisation will soon be delivering applications codes. In addition, a series of Project meetings has been organised:

1. New Architectures in Condensed Matter Electronic Structure Computation, Daresbury, 30th September 1988;
2. Workshop on Parallel Computing, Daresbury, 24th November 1988;
3. Parallel Computing on the Meiko Computing Surface, Daresbury, 14th June 1989;
4. Workshop on Applications of Atomic and Molecular Physics on Parallel Computers, Durham University, 18th July 1989;
5. Workshop on Parallel Algorithms in Molecular Simulation, Daresbury, 8th November 1989.
6. Workshop on the Intel iPSC/2 Hypercube, Daresbury, 6th March 1990.
7. Workshop on Parallel Codes and Algorithms for the Electronic Structure of Solids, Daresbury, 9th November 1990.
8. Workshop on the Intel iPSC/860 Hypercube, Daresbury, 12th December 1990.

For further details contact:

1. Dr. M.F. Guest, Head of Advanced Research Computing Group, SERC Daresbury Laboratory, Daresbury, Warrington WA44AD.

7. Engineering Board's CFD Initiative

The Theory and Computational Science Division at Daresbury Laboratory has been heavily involved in a number of Engineering Board programmes which aim to encourage the application of novel architecture and parallel computers to problems in engineering with particular emphasis on Computational Fluid Dynamics (CFD). The Advanced Research Computing Group at Daresbury undertook to supervise the acquisition of £1.5M of parallel computing equipment to be placed in academic departments for use in numerical modelling in engineering. Proposals were invited from some 14 vendors covering distributed and shared memory and single and multiple instruction multi-processors. The main conclusion that was drawn from the exercise was that it was too soon to pick a single best parallel machine for engineering (or any other) applications — the currently available architectures and corresponding software environments were too various, and the field was moving too rapidly. Instead, it was recommended that it would be more useful to consider acquiring a set of different machines so that the initiative, viewed as a whole, represented a balanced programme across the different hardware platforms. After consultation with the community, the purchase of the following 5 systems was recommended:

- 2 16 node Meiko i860 distributed memory systems.
- 2 HP/Apollo shared memory systems each comprising DN10040+DN10020.
- 1 shared memory Alliant FX/2808.

These recommendations were accepted and Engineering Board subsequently issued an invitation to heads of departments of engineering and related disciplines for applications from appropriate groups for these systems. A review panel recommended the following allocations to groups actively involved in applying parallel computing methods in a range of different engineering areas:

- Meiko: Dr. Fiddes, Bristol University for research into unsteady flows and turbulence modelling in aerospace applications.
- Meiko: Prof. Owen et. al., Bath University for research into CFD in the built environment, mechanical engineering and process engineering.
- Apollo: Dr. Irving et. al., Durham University for research into new algorithms for the temporal analysis of power transmission and distribution networks.
- Apollo: Prof. Eatock-Taylor, Oxford University for research into methodologies for computational hydrodynamics with application to problems in marine technology.
- Alliant: Prof. M. Leschziner et. al., UMIST for research into thermofluids modelling and particle-bearing froths.

To support and coordinate the Parallel Hardware Initiative, and to encourage the development of appropriate software and algorithms, the Advanced Research Computing Group at Daresbury were asked by Engineering Board Secretariat at Swindon to prepare a case for a Collaborative Computational Project on the Application of Novel Architecture Computers to Computational Fluid Dynamics (CCP12). Science Board has a number of CCPs who's main functions are:

1. to provide for the rapid interchange of information, theory, algorithms, programs, etc,
2. to collect, maintain and develop, relevant items of software,
3. to encourage basic research by providing facilities for the rapid exploration and implementation of new methods and techniques,
4. to assess and advise on associated computational needs,
5. to disseminate information via workshops, newsletters, visitor programmes etc.
6. to foster and coordinate international, especially European, collaborative work.

In the original proposal it was suggested that the CCP should focus on 3 key areas:

1. an investigation of parallel CFD algorithms. For large scale problems direct matrix methods are impracticable, the main emphasis should be on iterative methods: conjugate gradient, multi-grid, etc.;
2. implementing standard CFD codes on parallel systems. Incorporation of the new parallel solvers to provide a rapid demonstration of the effectiveness of parallelism for CFD problems. This project would involve the development of global grid partitioning and synthesis techniques, and static and dynamic load balancing. The problems of portability across shared and distributed memory machines would also need to be addressed;
3. development of real time integrated CFD environments where the adaptive grid generation, solution and graphics are tightly coupled and run in parallel.

Engineering Board Committees considered the proposal and allocated pump-priming funds of 1 man year pa with a sum for T&S and organising meetings and workshops. A bid for 3 man years pa to support the CCP-CFD has been included in Engineering Board's Forward Look.

A Steering Group with representatives from some 22 groups, drawn in the main from the proposals rated highly in the parallel hardware exercise, were invited to attend the first meeting of the CCP-CFD held in September 1990. It is the principle role of the Steering Group to determine the research programme for the project and to promote collaborations. The meeting concluded that in its first year the CCP should focus on a major awareness exercise surveying parallel hardware, languages, operating systems, compilers, program development tools, programming environments and applications libraries within the context

of engineering computations. In terms of research, the CCP was to proceed by collecting and reviewing commercially non-contentious codes which are representative of key applications and begin to implement the codes on a range of different parallel computer architectures. The primary aim is to explore the mapping between different numerical schemes and different architectures with realistic engineering problems and subsequently develop algorithms to realise the potential of powerful computation. With the resolution of property rights issues in the future, the scientific focus will shift towards developing and implementing flagship engineering codes on parallel and novel architecture systems.

For further details contact:

1. Dr. R.J. Blake; Secretary CCP12, Advanced Research Computing Group, SERC Daresbury Laboratory, Daresbury, Warrington WA44AD.

8. Computer Board's Initiatives

The function of Computer Board is not to support particular research activities but to provide a computer service to the academic community as a whole. The parallel computing initiatives detailed below are not intended to test whether the novel architectures are suitable for particular applications, or are worthy of research in their own right, but are to test the response when they are provided as a general user service in university Computer Service environments. It is hoped to gain knowledge of which architectures are suitable for a large range of applications demanding high performance computing resources and to gain experience in providing the necessary support infrastructure for their successful integration into the academic community.

We provide below a review of the two Initiatives in Parallel Processing presently funded by the Computer Board. Responsibility for the initiatives was delegated to the Technical Options Group (TOG) who further delegated to the Parallel Centres Steering Committee (PCSC). The PCSC will, in future, report directly to the Computer Board. This move is in anticipation of the shift of supercomputing to the ABRC who will take over both the Joint Policy Committee and the TOG.

Initiative on High Performance Distributed Facilities (HPDF or HPDS)

The Initiative was originally set up in 1988, with the aim of: "improving the availability of small novel architecture systems of the types which are felt most likely to have significant potential for advanced research computing". In 1989, the Computer Board set aside a total of £250,000 to allow for minor upgrades to the systems already purchased under this initiative. The nine original sites (Edinburgh, Kent, Southampton, Glasgow, Lancaster, Reading, SWURCC, Manchester and Cambridge) were invited to

submit bids for further funding, along with two sites who had not been funded initially (UEA and Newcastle/Durham). Eight bids were received from the original sites (Reading did not re-bid), totaling £451k capital, whilst both UEA and Newcastle/Durham sent bids, totaling £294k capital.A bid was also received from Queens Belfast, which would have been funded by DENI, rather than the Computer Board. The TOG considered the bids in October 1989, and recommended the following levels of funding which were accepted by the Computer Board at the end of October.

SWURCC	£30k
Cambridge	£40k
Edinburgh	£25k
Glasgow	£35k
Kent	£40k
Lancaster/Keele	£20k
Manchester	£30k
Southampton	£30k
	£250k
East Anglia	£75k
Newcastle/Durham	£75k
	£400k

The Board also agreed to recommend to DENI, on TOG's advice, that the bid from Queens Belfast should be funded. Once the awards had been made, the TOG delegated responsibility for the running of the Initiative to the PCSC.

The hardware purchased by the centres consisted of DAPs, and Meiko and Parsys transputer systems. A review of the first year reports indicated there was considerable slippage in the delivery times for some of the hardware and this is unattractive in an area of computing where system lifetimes are short. There had been little significant collaboration between the sites but at least four of the systems had been used for Protein Sequence Analysis.

It is too early to draw conclusions about the user service implications of novel architecture but information that could be gained from the sites under this initiative will include the following:

1. reports on items of software produced by the initiative,
2. reports on software produced,
3. performance information relating to efficiency of ported code and suitability of architectures for particular applications,
4. reports on performance of suppliers,
5. the effect of installing a high performance system on other mainframe and network resources,
6. purchases of Novel Architecture Computers in departments, possibly as a result of use of HPDS facilities.

The HPDS sites are to be assessed in terms of the following set of guideline criteria:

1. speed and efficiency of procurement of hardware,
2. physical availability of the systems, with reference to networking (local, campus, JANET), modes of access (interactive, batch only) and milestones,
3. system service, looking at interactive and batch capabilities, filestore, security, milestones and deliverables,
4. user assistance, focusing on documentation for both advanced and new users, on-line help facilities, training (frequency, duration, content, hands-on, attendance) and milestones and deliverables,
5. promotion of service and efforts to encourage new users,
6. established users, looking at total users, new users, campus network users and remote users, number of departments and different areas of research,
7. relationships with other centres, funded projects, suppliers and industry.

Initiative on Software Environments for Parallel Computers (SEPC)

This initiative was set up in 1989, with the aim of: "defining and, where necessary, developing techniques and software to enable parallel architecture systems to provide an acceptable user service." It was a three year programme and set up the PCSC as the monitoring group for the programme. £600k in capital funding and £500k pa in recurrent funding was made available to support between 3 and 5 sites located in universities. Twenty-four proposals were received in response to the letter announcing the initiative, and the TOG met on the 18 October to sort through them. It shortlisted proposals from the following Universities for further consideration: Edinburgh, Kent, Liverpool, Southampton and QMW. In addition, the TOG shortlisted the proposal from Newcastle/Durham on condition that it was modified to focus on the Newcastle portion of the proposal. A

proposal from Queen's Belfast was shortlisted providing DENI had funds available. The shortlisting was done at the first meeting of the PCSC in December 1989. For this meeting, each one of the shortlisted universities were asked to submit a further more detailed application for consideration. At this meeting it was decided that the following funding should be provided to the universities.

Site	Capital	Recurrent(posts)
Edinburgh	£192k	4
Kent	£128k	4
Liverpool	£75k	3
Southampton	£111k	4
NUMAC	£52k	2
QMW	£42k	2

In their proposals the sites list similar objectives, including:

- support for applications porting,
- documentation, demonstrations and courses
- development of parallel program development tools
- development of general numerical and application specific libraries

It is too early to try and assess any progress the sites may have made against their objectives.

For further information contact:

1. Russell Watts, University of Southampton, Southampton SO9 5NH.
2. Mr. K. Baxter, Computer Board Secretariat, University Research Councils, Grove House, 2–6 Orange Street, Haymarket, London WC2.

9. Center for Scientific Parallel Programming in the USA

This section reviews a proposal to the National Science Foundation, USA to establish a Centre for Scientific Parallel Programming at Rice University. The proposal is spearheaded by a consortium of leading computational scientists, numerical analysts and computer scientists from Rice University, Argonne National Laboratory and California Institute of technology. The research strategy is aimed at developing the algorithms and programming tools which will realise the potential of parallel computers, primarily by addressing the needs of computational scientists and automating the scientific programming process. The proposed programme of work focuses on developing a hierarchy of progressively higher level programming layers:

1. the machine-specific software layer consisting of a Fortran compiler augmented by communications and synchronisation operating system calls,
2. the parallel programming support layer consisting of a machine independent, explicitly parallel, dialect of Fortran with concurrent debuggers and performance monitors,
3. a numerical algorithm layer,
4. a higher level programming environment layer supporting an abstract mathematical language,
5. a scientific application layer for particular research problems.

The rationale underlying the proposal, is strongly reminiscent of that behind the provision of Foundation Support in the UK. The scientific applications range from biology, through physics, chemistry and engineering to general algorithms and numerical analysis. The programme is costed at $23M over 5 years with roughly a third of the funds to be spent on hardware. The proposal was successful and Centre was set up in February 1989.

For further information contact:

1. Dr. B.W. Davies, The Atlas Centre, Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire.

Acknowledgements

Much of the above has been edited from contributions and information received from the authors cited at the ends of the section. We would like to thank all the contributors for their help to date and trust that we can rely on their help in the future to keep the review up to date.