My DIC Rule

SCIENCE AND ENGINEERING RESEARCH COUNCIL
RUTHERFORD APPLETON LABORATORY

COMPUTING DIVISION

DISTRIBUTED INTERACTIVE COMPUTING NOTE 893

issued by
D A Duce

Minutes of the Review Panel Visit to
Drs Gurd and Watson, Manchester University,
30 September 1983(Second Draft)

3 October 1983

DISTRIBUTION:
R W Witty
F B Chambers
Miss G P Jones
Dr D A Duce
Investigators/Gurd and Watson

(see next page)

D R A F T

IN CONFIDENCE

SCIENCE AND ENGINEERING RESEARCH COUNCIL
ENGINEERING BOARD

GR/B/74788

COMPUTING AND COMMUNICATIONS SUB-COMMITTEE

DISTRIBUTED COMPUTING SYSTEMS PANEL

**Minutes of the Review Panel Visit to
Drs Gurd and Watson, Manchester University
30 September 1983.**

1. PRESENT

Panel	Manchester
E C P Portman (ICL - Chairman)	J R Gurd
D A Duce (Academic Coordinator)	I Watson

2. PRIVATE PANEL BRIEFING.

Drs Gurd and Watson hold a 4 year rolling grant under the DCS Programme entitled 'A Multilayered Dataflow Computer System' which commenced on 1 October 1981.

The objective of the Panel visit was threefold:

a. To review progress on the project since the Review Panel visit in September 1982.

b. To consider the necessity for the release of further funds on the present grant.

c. To consider and make recommendations on a submission by Drs Gurd and Watson on 15 September 1983 for a further 4 years support for the project.

3. DISCUSSION WITH DRS GURD AND WATSON

3.1 Progress to Date.

Funds were awarded on the present grant for the construction of a multilayered dataflow system, and for some initial exploration of storage hierarchies. Two RA posts were awarded, one for applications studies and one for hardware engineering. Funds for the construction of the multilayer system were conditional upon the demonstration of the single ring prototype whose construction was funded by a previous ordinary research grant.

The Review Panel visited the project in September 1982 and were indeed satisfied with the performance results then presented for the single ring system. However, it became clear that constructing a multi-ring machine by merely replicating the single ring prototype was not the wisest course of action. The Panel agreed that prior to construction of a multi-ring system three other tasks should be completed:

a. Procurement and installation of a front-end processor system (as requested in the original application).

b. Expansion of the matching store from 16K to 1M words to allow much larger programs to be run.

c. Modelling a 20-ring system in which individual rings are emulated by M68000's and incorporating a 1-bit serial switch between rings of the same topology on that proposed for the multi-ring system.

The Panel also ratified the decision to investigate the theory of programme development for dataflow computers, rather than investigation of applications per se. This work has progressed very well and has lead to one strand of work in the new grant application.

The Panel also recognised the need for a user liaison officer and subsequently an EMR Agreement was awarded to the investigators for this.

3.1.1 Front-end Processor

A Collaborative Research Agreement between DEC and the University of Manchester (approved by the Secretary of the Council in April 1983), enabled the investigators to purchase 2 VAX 11/780 computer systems at a very favourable price. The systems were installed in June and now provide an excellent software development environment, far superior to that previously used. The existing software has now all been moved to the VAX and this was accomplished very speedily.

3.1.2 DEC Research Agreement

The collaborative agreement with DEC has brought users to the machine.

A new single assignment language, SISAL, has been developed as a collaborative exercise between DEC, Lawrence Livermore Laboratory (California) Colorado State University and the University of Manchester. This language will be implemented on a range of parallel architectures and LLL are investing significant effort into rewriting existing codes in SISAL.

LLL staff are visiting Manchester to use the dataflow machine. The joy of having real users of the system is only tarnished by the fact that they are not UK users.

3.1.3 Matching Store Upgrade

The single ring prototype is now operating with the new store boards, though only 2 boards have so far been completed. Already, however, the matching store is 8 times larger than previously. User pressure makes the completion of the remaining boards a priority task and it is hoped the 1M token store will be completed within a month.

It has also been realised that the node store (program store) urgently needs to be expanded. The release of funds for this was requested.

Physical board construction is being carried out by the hardware RA, Bryan Saunders. Board production is done by technicians in the department and very reasonable quality boards are now being produced.

3.1.4 Multi-Ring Model System

The construction of the multi-ring model system in which individual rings are emulated by two M68000's is progressing well. Each ring fits on one (large) board. Four are currently fully operational. A 4x4 switch module has been designed and constructed and is currently being commissioned.

It had been hoped to complete the multi-ring model system within one year, but given the enormity of the task it is understandable that timescales have slipped slightly. This part of the project is being undertaken by John Foley. The Review Panel were very impressed with his abilities and the contribution he has made.

3.2 New Grant Application

The new grant application contains five strongly connected strands of work. Some of these are considered explicitly below.

3.2.1 User Service

Over the last year the project has been responding to the needs of users and this is expected to continue in the future.

Users and a user service will lead to transfer of technology. To support a user service funds have been requested for a system manager and documentation officer. The former was considered to be very essential. The system manager needs to be a senior, responsible person whose role includes that of external interface to the project. A full time system manager would effectively release one of the investigators, who currently has to fulfil this role, for research. The documentation officer will be responsible for software distribution, documentation, preparation of newsletters etc. The Panel strongly supported these requests.

Funds (10K p.a.) have been requested for upgrades to the front end processors to support an expanding user community. Upgrades contemplated are:

Disc space (eg. 500Mb drive)
Terminals
Network System

This request was considered entirely reasonable.

Funds (5K p.a.) have also been requested for further (currently unknown) expansions to the single ring system. This request follows the recent need to upgrade both the token and node stores. Again this was considered entirely reasonable.

3.2.2 Data Structure Store

Data structure storage is considered an essential adjunct to the present system, to support larger applications. The international dataflow community are moving in this research direction. Outline designs for two possible schemes have already been undertaken. The schemes will be simulated on the VAX before hardware is constructed.

3.2.3 Multi-ring Hardware

There was a long discussion of the need for and best approach to the construction of multi-ring hardware and the need to involve industry in this.

a. The present single ring prototype machine has a power of 2-3 mips. Replication of this 4 times will give, say, a 10 MIP machine. This is not impressive by current von Neumann machine standards.

b. The group's research interests and expertise are in computer architecture, not technology.

c. Looking at the prestige angle of the project (UK dataflow expertise), a 10MIP TTL machine would not be particularly impressive. On prestige grounds alone one would reimplement the machine in a more modern technology, CMOS gate arrays for example. However the design effort alone for this would be enormous and could not be provided by Manchester. For example, to implement the machine in 8000 gate array CMOS could require the design of between 10 and 50 chips.

d. The ideal would be for an industrial partner to reimplement the machine in a modern technology, the Manchester team advising on architecture. Few companies have the technology and expertise for this, but there are some.

e. A manufacturer taking on board dataflow now, could have something like a 100-200 MIP production machine with 25-30 parallelism in each of 10-20 rings. A high technology version five years after might have a 5-fold speed increase.

f. To take the dataflow idea to production now would cost £50-100M of which £30-40M is hardware. As a project for short term return of investment it is a non starter. The reason for entering the game now would be a belief that a market will open up and entering now would give a 2-3 year lead. As a comparison of development costs, Mr Portman pointed out that reimplementation of an existing machine in CMOS costs of the order of £5-7.5M.

g. It is instructive to note that in a typical company in the USA 16% of profit is spent on research, whereas in the UK only 4% is so spent. It is also instructive to note that American computer manufacturers are talking seriously to the American dataflow community. Current estimates are that the Manchester research group still have a lead over the rest of the world. A 10MIP machine at Manchester in 2 years time would still have a world lead.

h. Dataflow machines are likely to be interfaced for the foreseeable future to conventional machines.

i. Laboratories such as Lawrence Livermore who have very computationally intensive problems, perceive dataflow as a way out of their present dilemma of having to recode their applications for every new machine in order to extract the last ounce of speed. Dataflow offers a way out because of the incremental expansion capability. Already 10:1 performance improvements have been shown in the Manchester machine by switching in more processors.

3.2.4 Specific Questions

A number of specific points were raised on Mr Newey's behalf by a colleague in a letter to the Review Panel.

a. What sort of debugging facilities will be provided to remote users of the existing machine?

   This phase of the project is covered by a separate application from Drs Gurd and Kirkham, entitled 'A Software Development for Dataflow', and is also the subject of an MSc project. A symbolic debugger for SISAL is being written. This will present hardware events in SISAL language terms.

b. What consideration, if any, will be given to issues of protection in designing the virtual memory system? Will it be possible to prevent a faulty program from destroying evidence leading to identification of a fault? Can the memory be partitioned in such a way as to allow the parallel execution of separate programs?

   The Manchester architecture raises no new protection issues. Existing standard techniques can be applied and are really implementation issues. The prototype machine in fact only has parity on stores. One could, and in a production machine, would, go much further.

   In the present system obscure faults have been traced by recording all tokens circulating in the ring in a circular buffer and tracing back through this history when a fault is detected. The present architecture is sympathetic to this style of monitoring and fault tracing.

   Multiprogramming comes for free. However, it is in a sense a non issue because anyone needing a dataflow machine is liable to do so because they have a single programme that will consume the entire machine!

c. When can a serious example of a heavy computation be put on the machine in order that a more objective comparison with other parallel architectures can be made?

This is a target of the collaborative research with DEC. The Lawrence Livermore studies, which will be in the public domain, should also achieve this. LLL are comparing the dataflow machine with a 4 processor VAX mp, and possibly the HEP at Los Alamos Laboratory. The LLL code is all written in SISAL. The SISAL definition is public domain and the compiler object code may be also. SISAL is compiled to an intermediate form (IF1) which could be translated relatively easily to machine code for other machines.

It was suggested in this context that the optimisation codes developed for the DAP by Dr Dixon at Hatfield Polytechnic might be a useful benchmark. The User Liaison Officer was asked to visit Dr Dixon.

d. Given that the overwhelming majority of CPU intensive scientific and engineering software is written and will continue to be written in FORTRAN, what attention will be given to methods of moving FORTRAN code either directly or via an intermediate language?

None at Manchester. The premise that software will continue to be written in FORTRAN was not accepted by the investigators. LLL for example rewrite codes for new machines because they cannot get the speed from FORTRAN.

Worldwide there is some significant work on the problem of translating FORTRAN to dataflow graphs, which Manchester could pick up if need arose. David Kuck at University of Illinois has probably done most in this field. Arthur Veen at the MC in Amsterdam is also active in the field of compiling conventional languages for the Manchester machine.

e. Will DEC offer the same levels of discount on upgrades to the 11/780 as they provided on the original system?

This is unlikely as the original discount levels were really money in kind.

3.2.5 Travel

The Review Panel considered the request for funds to travel to Japan and supported the request, believing that one has to visit Japan to discover what is happening there. Professor Kowalski's visit, funded by the Royal Society, was cited as an example of the value of a properly planned and coordinated visit.

3.3 Release of Further Funds on Present Grant

Some £80K will be uncommitted under the Equipment and Other Cost headings on the present grant, after the work authorised by the Review Panel has been completed. The release of further funds is urgently requested for:

Node store upgrade
DR780 interface
Switch expansion (multi-ring model)
Small size overflow unit.

4. PRIVATE MEETING

The Panel continued to be impressed with the quality of the project and the excellent progress made. The recommendations and conclusions of the Panel are:

a. The Panel wished the remaining money on the present grant to be released for the construction projects mentioned in 3.3. above.

b. The projects initiated last September (multi-ring model, matching store upgrade) are progressing well and the slight delays give no cause for concern.

c. The new proposal is very clear and well-written. The aims and directions are very sensible and logical. The service side is entirely in line with DCS Panel policy on dataflow. The requests for funds for upgrades to the single ring prototype and front end processor hardware were considered entirely appropriate. The case for a System Manager and Documentation Officer was strongly endorsed.

d. It is clear that an industrial collaborator for the 4-ring project is highly desirable. Without an industrial collaborator it is unfeasible to do other than replicate the single ring system in the existing technology. With industrial collaboration much better and more prestigious solutions would be feasible.

e. The single ring system will soon be running out of speed. Users are likely to be interested in dataflow because of the performance and expansion potential. 2MIPS is too close to what is currently available to be attractive and users would want the promise of something substantially better in the near future. Building a 4-ring system might well be an important factor in attracting meaningful outside interest.

f. A course on Dataflow to encourage new users of the machine was suggested. It was also felt there was a need for someone senior in SERC to approach potential user organisations (eg. AERE, BNOC, etc). Dr Duce was asked to pursue this.

g. The overall recommendation on the new grant was that it should be approved in full. The strands in the project form a synergistic whole. Cutting strands of the project would lead to its ruin. The Panel had no worries about the proposed programme and the investigators' competence to carry it through.

The Panel noted that the Manchester project is the most advanced novel architecture project in the UK. Architecture research is intrinsically expensive. To reject this proposal on financial grounds would be tantamount to an admission that there is inadequate funding to support novel architecture research. New projects in this field ought not to be encouraged. The Panel sincerely hoped that having funded the intitial stages of a project which has given the UK international acclaim in this field, the funds would be found to continue this line of research.

h. The Panel would like to visit the project at Easter 1984 to review progress both on the constructional work and use of the machine.

5. FEEDBACK

The investigators were congratulated on their excellent work and progress and were informed of the Panel's decisions recorded in section 4.

DAD
3.10.83

IC1/14/jg