my DIC file

IN CONFIDENCE

SCIENCE AND ENGINEERING RESEARCH COUNCIL
RUTHERFORD APPLETON LABORATORY

COMPUTING DIVISION

DISTRIBUTED INTERACTIVE COMPUTING NOTE 946

PANEL MEETING

issued by
Dr D A Duce

Minutes of the Gurd and Watson Review Panel Meeting
10 January 1984
(Second Draft)

7 February 1984

DISTRIBUTION:
R W Witty
D A Duce
G P Jones
File Manchester/Gurd & Watson

(see next page)

IN CONFIDENCE

GR/B/74788
GR/B/401966
GR/C/83876
GR/C/83425

Science & Engineering Research Council
Engineering Board - Information Technology Directorate

Information Engineering Committee

COMPUTING SUBCOMMITTEE

Minutes of the Manchester Dataflow Project Review Meeting held on 10 January 1984 at the Institution of Mechanical Engineers, London.

PRESENT

Review Panel:

R Newey (Chairman)
Dr M R Sleep
Prof J Iliffe
Prof D Lewin
Prof R L Grimsdale
E C P Portman

M Hotchkiss (Office)
F B Chambers (DCS Industrial Coordinator)
Dr D A Duce (DCS Academic Coordinator)

Dr J R Gurd (Manchester)
Dr I Watson (Manchester)

FIRST PRIVATE PANEL DISCUSSION

1. The Chairman outlined the history of the project leading to the present grant submissions and the Computing Subcommittee's decision not to support them. The Subcommittee's recommendation appeared in the meeting papers.

The Panel's terms of reference were:

a. To review research in general on dataflow architectures;

b. To discuss with representatives of the Manchester Dataflow team possibilities for future research;

c. To advise the Manchester Dataflow investigators on directions for future research;

d. To report to the Computing Subcommittee.

2. Professor Lewin stressed that the Computing Subcommittee were convinced of the quality and value of the Manchester team and wished to keep them together.

FIRST OPEN DISCUSSION

3. The first open discussion with the Manchester team addressed the issues of dataflow in general:

Where we are now
Where we could go
Where we should go.

4. A distinction was made between the world-wide context, the UK context and the Manchester context.

5. Submissions had been invited from key researchers on the future of dataflow. Very little time had been given for the preparation of responses. It should also be noted that the contributors spoke as individuals, not for their organisations. Submissions were received from Drs Gurd & Watson, Dr Sleep, Dr G Michael (Lawrence Livermore Laboratory, USA), Dr J Dennis (MIT) and M Patel (DEC). The latter submission took the form of a note by Mr Newey on a telephone conversation.

6. The paper by Dr Sleep, surveying novel architectures was taken as the starting point for the discussion. Dr Sleep has an advisory role to the Alvey Directorate on IKBS novel architectures. The following points emerged:

a. There are limits to being able to purchase speed using the von Neumann architecture.

b. Developments of the von Neumann machine make no shift in the underlying model of computation and so are unlikely to find revolutionary solutions to the software crisis.

c. Programming machines based on the parallel composition of von Neumann chips is known to be extremely difficult. This points to the need for architecture research to be hand in hand with research in programming methodology.

d. The systolic approach is useful for problems with certain characteristics, but is not a general approach.

e. There are very few revolutionary approaches. Mago's (USA) machine proposal is probably the most revolutionary.

f. It is worth noting that the dataflow model is applicable at a variety of levels. One could, for example, regard UNIX pipes as very simple dataflow graphs. Occam can be regarded as von Neumann processing plus data driver communication. The Manchester project addresses instruction level dataflow.

g. The language first approach to architecture argues for the sequence:
language -- semantics -- hardware.
The Manchester project followed a more half-way house approach, starting from a computational model. This approach was felt to be the only available methodology for tackling the problem of parallelism.

h. Section B of Dr Sleep's paper contains some guesses of what things will look like in the 1990's. Dr Sleep felt the success of Darlington's ALICE machine lay in CTL (Compiler Target Language) - a workable generalisation of the von Neumann instruction set. This enables software development programmes to proceed independently of research to realise CTL in hardware.

i. Dr Watson expressed concern that Dr Sleep seemed to be suggesting that ALICE CTL should be the basis for all projects. Dr Sleep indicated that CTL was merely an example, but felt that some standard was required at this level, to allow developments to proceed on both software and hardware fronts. Mr Newey pointed out the need for caution when choosing not to keep a field open and inhibiting free ideas. At some point such a decision may be necessary to cause a community to work together and foster symbiotic relationships.

Dr Gurd felt a free reign might be given at lower funding levels, but at higher levels some sort of standard may be required - for example progress in dataflow languages was only made when people wanted to use them.

j. Dataflow research could yield results in:

increasing the store of scientific knowledge
increasing speed at all costs
increasing speed at constant cost

Dr Watson believed dataflow projects around the world had come to the topic from all three directions and believed dataflow might yield results in each.

k. Professor Grimsdale warned against a preoccupation with industry. Professor Lewin felt that if dataflow were purely SERC funded on the basis of addition to scientific knowledge, then value for money was required in the sense that one needed to balance the exploration of the dataflow path against other revolutionary ideas. Mr Newey felt the UK was getting into a scientific knowledge versus industrial exploitation dilemma and that there was a danger that we would go too far into industrial exploitation.

l. Mr Newey drew a distinction between development for research and development for industry. He believed the Manchester construction work to have been development for research.

m. Professor Iliffe wished to know what problems dataflow had solved better than existing approaches and what new problems had been solved.

i. Lawrence Livermore Laboratory look to dataflow as a potential way to satisfy their massive computational requirements (100-1000 times greater than present machines can satisfy). Unsurprisingly dataflow machines do not yet match the power of a Cray. However techniques have emerged.

ii. The general claim for dataflow is that it is good for problems with a lot of irregular parallelism. Results obtained by LLL support this claim.

iii. The solution to the software crisis is not to be found in conventional languages. The dataflow model may be a good way to implement declarative languages.

iv. Millions of man years have been put into a control flow dominated world so there is merit in exploring a heavily dataflow dominated world in the hope that more general theories will result (R Newey).

v. Dr Gurd pointed out that the dataflow architecture had not yet received the kind of investment that builds a Cray. Mr Newey speculated that the UK would never make that kind of investment (guessed at $100M) without firm evidence of the payoff.

n. Dr Sleep pointed out that novel architectures are in a Catch 22 situation - funding for improving the von Neumann architecture is gargantuan compared to the funding given to alternatives. Mr Portman argued that increased performance at constant cost only comes with volume and refinement and hence dataflow was more likely to make an impact in terms of scientific knowledge and speed at all costs. He felt that dataflow should be supported because it offered the promise of higher performance computers without the fearful complexity of super von Neumann machines (such as Cray).

o. Professor Grimsdale felt that the following questions should be addressed:

i. relationship to industry
ii. evidence for cost/performance benefits
iii. role of the machine and class of problems intended to be solved
iv. what dataflow is not
v. effects as a stimlus for computer science research
vi. why build hardware

The Panel decided to apply these questions to dataflow worldwide in the present discussion and to the Manchester project later.

The remaining position papers were then discussed.

MICHAEL'S PAPER

7. a. Dataflow researchers have a responsibility to comment on recycling existing investment in software (eg FORTRAN programs). Dr Gurd believed that Kuck's work provides the most promising evidence that the investment in conventional languages need not be entirely discarded. If a dataflow machine could run FORTRAN at no worse cost performance than a decade ago, say, then it might be worth switching to a dataflow machine.

b. Dr Gurd argued that for a component cost equivalent to a VAX (£40K) the Manchester machine produces VAX performance to within a factor of 2. Dataflow machines could be used as add-on devices giving a cost effective boost for problems with irregular parallelism.

PATEL PAPER

8. Patel's views on multiring systems were discussed. His view appeared to be that the dataflow architecture exchanges the problem of spreading control in parallel von Neumann machines for the problem of spreading data between rings. Dr Gurd argued that data flow did not suffer from an analogue of the problem of splitting a computation into processes. Mr Newey argued that there would still be duals of the containment and load balancing problems. Dr Gurd believed the Manchester machine would have less problems because of the low level of granularity of parallelism. This is a particular property of instruction level parallelism.

DENNIS PAPER

9. The performance comments assume industry will invest!

GURD AND WATSON PAPER

10. a. The language problem was discussed. Dr Gurd argued that the full power of higher order functions might not be required and that implementation of built-in higher order functions might be possible in the Manchester machine. Dr Sleep accepted that the full generality was probably not required but felt there were aspects of functional languages which were not captured by the Manchester machine.

b. Dr Gurd argued that there was merit in exploring the use of single assignment languages. New users find the transition to these languages easier than that to zero assignment (functional) languages. It is also easy to see what code will be generated by a single assignment language.

PROF GRIMSDALE'S QUESTIONS

12. The questions posed by Professor Grimsdale (point 6(o.) above) were then treated in turn, in a worldwide context.

a. Relationship to industry. There are known links between dataflow projects and DEC, IBM (MIT) and NTT (Japanese telephone company). J C Syre is about to start a new project with backing from an unspecified French company. A division of TRW are building a signal processing machine based on the published Manchester papers and the HEP machine contains a lot of dataflow ideas. The prime movers have been researchers assisted by industry. There has been industrial pickup but not exploitation.

b. Cost/Performance benefits. There is no irrefutable evidence that cost/performance ratios will be better for dataflow machines than von Neumann, nor is there any counter evidence. There is little evidence because few machines have been built. Only one paper (Japanese) gives results of a similar degree of comprehensiveness to the Manchester results. Professor Iliffe felt that there is ample a priori and experimental evidence with the von Neumann design.

c. What dataflow is and is not. In the Japanese context dataflow includes reduction. In the Manchester view dataflow is one computational model, reduction is another though they can be tied together in the sense that dataflow can be used to realise reduction. It was agreed that a statement of the meaning of the term 'dataflow' around the world was needed. Dr Sleep was asked to write this.

d. Role of dataflow machines worldwide. Lawrence Livermore Laboratory see applications to numerical problems. DEC see the possibilities for an add on processor for pcb/ula CAD applications. Dr Gurd sees such machines as a general resource for parallel algorithm design.

e. Stimulus for research and ideas. It was believed that dataflow research does and will continue to provide stimuli.

f. Why dataflow? The Manchester team came to dataflow through general considerations. They ruled out SIMD machines as not general enough, and shared store and message passing architectures because of the partitioning problems. Hence they went for instruction level parallelism. Other people have come to dataflow through different routes - Arvind for example came via functional languages. Dataflow seems a good middle position for the moment. It is not obviously wrong.

SECOND PRIVATE MEETING

12. Mr Newey noted that the grant applications under consideration had been formally rejected and hence the next part of the meeting would open with a discussion of research directions. He felt the Panel should resist drifting into a situation in which they felt that the applications should have been approved but were not sure why.

13. Prof Lewin noted the Computing Subcommmittee's concern about what would happen to the project after the next phase of the work. He pointed out that the funds requested for the present phase represented about 50% of the available SERC non-Alvey funds, and it may be necessary to argue for more funds for this area.

14. Mr Portman felt that the multi-ring architecture had not been adequately explored. Simulations have been and are being done, but ultimately it would be necessary to build something. Useful data will only be obtained when the machine is loaded with real problems. Thus he felt that the work should be carried into this phase. Industry needs these results and the work has to be done.

SECOND OPEN DISCUSSION

15. Dr Gurd gave a presentation of the research directions the Manchester team wished to pursue. The description was based on the Programme at the start of the Case for Support for GR/C/83876.

16. The Manchester team regard themselves as having a dual role.

a. The team believe it is sensible to offer other people access to the existing hardware. To this end there is an SERC funded User Liaison Officer, though access to the machine has been provided on a more or less casual basis so far. The proposal requests funds to provide the elements of basic support for external and internal users. This activity is necessarily independent of the other lines. Funds requested cover maintenance of the front-end computers and support staff.

b. The second part of the proposal is also regarded in a service sense. This part is concerned with the construction of a 4-ring dataflow machine and it is felt that there is a National prestige element in having a large dataflow engine for the general community. Manchester would most like to see this built in collaboration with industry. The proposal contains costings for replicating the existing design three times. If SERC wanted Manchester to build a 4-ring system and no industrial support were obtained, then the only feasible avenue would be to replicate the existing design. The costings were included on this basis.

These first two items are regarded as provision of general services. The research context of the proposal lies in the next three sections.

c. Architectural enhancement - data storage hierarchy in the system. Two approaches are possible; firstly to make the matching store into a virtual store (it is not immediately clear how to do this), or secondly to build a specialised structure store - it is clear how to do this and simulations are operational but this approach suffers from the drawback that users need to know a good deal about it.

d. Applications studies. There are two themes here - to maintain the user liaison officer's activities and to explore non-numerical applications.

e. Algorithms research. This activity will explore the semantics of dataflow and program transformations.

Manchester share these research directions with other groups and intend to keep their lead.

18. The second application from Drs Gurd and Kirkham is to continue software support in the languages area. It is intended to appoint a Dutch researcher to the RA post, Wim Bohm. Dr Sleep confirmed Dr Gurd's statement that Wim would be an excellent man for the post.

19. The dataflow project is supported by a rolling grant which presently runs from October 1981 to September 1985 and is valued at £402K. Of the £852K requested in the Gurd and Watson application, about £200K is already committed in an existing grant. Thus the application requests £650K in "new money" and an extension of the project by 2 years. The present grant expires in September 1985.

THIRD PRIVATE MEETING

20. The Panel felt that the previous discussion had clarified the inter-relation of the activities in the research proposals.

21. Dr Sleep, whilst fully agreeing that the work was excellent, felt that the basic hardware design was not yet right.

22. The fact that UK industry is retarded is not a reason for not funding the project. The decision to spend money for the benefit of the UK or other nations was not the Panel's responsibility.

23. The von Neumann architecture will continue to receive investment. All simple enhancements have now been done. With dataflow there is an opportunity to go in a direction in which all such enhancements have not already been made.

24. Low level dataflow may offer some new approaches to traditional problems such as high reliability systems. Such aspects have not yet been tackled.

25. Mr Chambers proposed the following way forward. Support the proposal in the areas of basic support, architectural enhancements, application studies, algorithm research and software development environment, with any possible trimmings. Delete the multi-ring construction activity.

26. Prof Lewin questioned whether the external service should be supported by CS funds or by the users themselves (eg through other SERC subject committees).

27. Mr Portman rehearsed the argument for building the multi-ring system - users will not invest in running real problems unless they can get real answers, ie the machine power is commensurate with existing machines.

28. Mr Newey argued that hierarchic memory management was much more important than a machine 4 times larger. He went on to suggest elements one would like to see in dataflow research proposals. These were refined and appear later in these minutes.

29. Prof Lewin was worried that the principal application areas were numerical and CAD. He would like to have seen non-numerical areas also. He was also concerned at the lack of UK users.

30. The Panel were concerned whether the UK could support more than one novel architecture project.

31. The Panel felt there was insufficient direct customer demand to merit investment also in a 4-ring system in the present technology.

32. The Panel felt that the problem of investment in existing software needed to be addressed and that each dataflow project should have a strategy for benefitting from existing software investment.

33. The Panel discussed the delicate question of the date by which the majority of useful results would have been obtained from the present hardware and the stage at which thoughts should turn to second generation hardware based on the lessons learnt from the present hardware. Professor Iliffe felt that unless there was much stronger evidence in favour of the dataflow approach by October 1985 the lesson to be learnt from the present hardware ought to be to give up. The Panel felt that the useful results could be achieved by October 1985 and that a proposal for second generation hardware could be forthcoming in this timescale.

The Panel then agreed the following recommendations.

Recommendations

34. The Panel considered what they believed should exist in dataflow research. The Panel would expect to find the theme of dataflow research in the UK addressing some substantial part of the following 5 points (in no particular order):

a. Dataflow hardware should be able to handle indefinitely large problems because, for the foreseeable future, the size of problems will be significantly bigger than the low level hardware we can afford to build.

b. Dataflow hardware projects must take advantage of the investment being made in software in closely related fields.

c. Dataflow research should address the problem of the investment already made in conventional software, and offer a strategy for the retention of this.

d. The study of dataflow hardware applied to non-numerical applications (though not to the exclusion of numerical applications).

e. Major hardware projects should provide access to users from computer science and other disciplines.

Points specifically relating to the Manchester project:

35. The panel do not believe that the additional research results obtainable from the construction of a multi-ring system merit the necessary level of investment unless funded from industry.

36. The panel rate point b. above as very important and encourage the Manchester team to address this.

37. Considerable thought should be given now to the next generation of dataflow hardware. The Panel encourage Drs Gurd and Watson, the proven UK team, to consider that subject and will recommend to the sub-committee that the present grant should not be rolled beyond 1985 in its present form. Rather the Panel would expect to see a major new grant application addressing the next generation of dataflow hardware, and in parallel might expect top up grants to maximise the learning from the existing hardware.

FEEDBACK

Drs Gurd and Watson were told of the Panel's recommendations.

38. Mr Newey explained that the Panel wished to see the investigators get as much out of the current hardware as possible in the remaining two years of the present grant. The Panel would not recommend a roll of the grant which would mean adhering to the current hardware beyond 1985, and the investigators should optimise what can be derived from the present hardware. Mr Newey pointed out that the investigators might feel it appropriate to have additional effort during the next 2 years, but that the Panel would not like to see a top up grant for the existing hardware beyond 1985.

39. Drs Gurd and Watson were then invited to discuss any immediate problems/clarifications with the Panel after a brief private meeting.

40. Dr Gurd said that there were some immediate personnel problems for which short term applications had already been submitted. These concern Mr Foley who will do the structure store project and Mr Bohm who will carry on the software environment work. The Panel agreed to recommend immediate award of these grants and that the periods should be extended to September 1985.

41. A second question concerned the release of funds in the present grant. £80K remain unreleased. The proposed usage of these funds includes the construction of a structure store, an overflow processor for the matching store and increasing the instruction store size. The Panel agreed to the release of these funds and added that they should be used with an eye to the future. Mr Hotchkiss was asked to write formally to the investigators.

42. Dr Gurd asked if point d. in the recommendation meant that numerical applications were less favoured than non-numerical. Mr Newey said that the Panel did not intend the point in this sense, rather they wished to point out that they believed non-numerical applications offered an opportunity.

43. Dr Gurd asked if SERC would have funds to cover a new application. Mr Hotchkiss said he believed that the financial position would improve after 84/5 by which stage Alvey policy would also be clear.

44. Dr Gurd said that the Panel's recommendations were seen as a distinct change of emphasis. A lot of effort had been put into getting industrial support for the present project and what assurance was there that a new project would not beat a similar tack. Mr Newey said that one option open to the investigators was that they should not look for a new major grant application, but increase the effort in looking for industrial support for the present hardware. If successful, the realisation of the present hardware beyond 1985 would then be picked up by industry. Alternatively they could take the line that some new generation activity will be undertaken with industrial contacts and support. If funding were given in the Alvey context this would necessarily have industrial support. In an SERC context the research content would be very high.

45. Dr Watson said he was concerned at the lack of forward looking industrial partners. Mr Newey said that Alvey necessarily forces the choice between research and advanced development.

46. Dr Watson explored point b. in the recommendation. Dr Sleep said he felt there should be a tie-in to the Imperial work, though not necessarily through CTL. He felt there were already too many languages and that in the Alvey context standard logic and functional languages were necessary.

47. Dr Watson asked if there was interest in UK industry in declarative languages and if so was it confined to non-numerical applications. Mr Portman believed the answer to both parts was in the affirmative. Mr Newey felt there was interest, but that UK industry did not know how to realise that interest. Mr Portman believed that there are important markets which are not satisfied by present approaches and tools.

48. Dr Watson asked if the view was that UK industry had no interest in speed-at-any-cost numerical machines. Mr Newey said he felt that this was the case. He believed there was no home market big enough to justify the investment. Mr Chambers felt some specialist markets (eg signal processing) did exist.

49. Mr Portman believed that present markets did not justify the investment.