NASA Technical Memorandum 78613

An Assessment of Future Computer System Needs for Large-Scale Computation

Peter Lykos and John White

FEBRUARY 1980

NASA

NASA Technical Memorandum 78613

An Assessment of Future Computer System Needs for Large-Scale Computation

Peter Lykos
Illinois Institute of Technology
Chicago, Illinois

John White
Ames Research Center
Moffett Field, California

National Aeronautics
and Space Administration

**Scientific and Technical
Information Office**

1980

TABLE OF CONTENTS

AN ASSESSMENT OF FUTURE COMPUTER SYSTEM NEEDS FOR LARGE-SCALE COMPUTATIONS

Peter Lykos
Illinois Institute of Technology

and

John White
Ames Research Center

SUMMARY

The report summarizes data, collected from many users, that range from specific computer capability requirements to opinions about the desirability of a national computer facility. Clearly there is a substantial communication gap between the users and the designers of large-scale scientific computers. As a consequence, vendor unwillingness to fully exploit computer technology because of financial risks causes an opportunity for both users (and therefore their sponsors) and manufacturer to be lost. The users' request for "bigger and faster" is too simplistic a statement of need. Users must be more explicit in defining their computational needs and must take into account available technology when designing their formal problem solutions. Considerable attention should be given to improving the user-machine interface; otherwise, increased computer power may not improve the overall effectiveness of the machine user. Significant improvement in throughput will require highly concurrent systems plus the willingness of the user community to develop problem solutions for that kind of architecture (i.e., give up FORTRAN). An unanticipated result was the expression of need for an on-going cross-disciplinary users group/forum in order to share experiences and to more effectively communicate needs to the manufacturers.

INTRODUCTION

A series of aircraft research projects, which involve the building of physical models for wind-tunnel characterization, has been underway at Ames Research Center for some time. The computer has come to be used to a greater and greater extent in the design of the physical models as well as in the design of the experiment and the interpretation of the data. With the concommitant improvements in model design and the increase in available scientific computer power, a point was reached where it seemed reasonable to seriously consider the feasibility of a "numerical wind tunnel," that is, direct numerical solution of a mathematical model.

In the process of analysis there emerged the realization that the commercially available scientific computers are lacking in at least three regards:

1. They are incapable of processing the desired algorithms in tolerable elapsed times.

2. The hardware components presently used are substantially below the current technological limits on speed and size.

3. The machine architectures available are not well suited to the desired algorithms.

While the cognizant Ames research staff pursued these questions with machine designers and computer vendors, a parallel effort was initiated out of Ames' Office of Planning and Analysis to determine to what extent researchers working in other areas were also experiencing constraints on their research due to the limitations of commercially available scientific computers. The latter effort included participation in a series of computer vendors' briefings of NOAA staff on the potential of new machine designs for handling meteorological models, participation in a national symposium on High-Speed Computer Hardware and Algorithms,[1] and informal meetings with computer designers. The principal effort, however, was the preparation and distribution of a comprehensive questionnaire concerning user estimates of their computational needs in the 1985–1990 period. An iterative procedure was used to refine the data and conclusions: Twenty of the respondents, together with representatives from six manufacturers, reviewed a preliminary report and all the responses in a retreat/workshop situation. A preliminary report, with workshop comments added, was then sent to 100 of the original respondents, together with a copy of their original response, for comment. All of these comments were then combined with the preliminary report to make the final report.

METHOD OF PROCEDURE

The first task was to identify users of large-scale scientific computer systems who met the following criteria:

1. Researchers working on socially relevant problems and who were experiencing as an essential constraint the size and speed of the scientific computers commercially available

2. Researchers whose peers would judge that they are approaching significant problems in a meaningful manner

3. Researchers willing to complete a detailed questionnaire starting from their principal area of interest and working toward the machine design implications of the corresponding algorithms

A mailing list of about 5,000 people – academic deans and department chairmen, directors of industrial research laboratories, directors of federal research laboratories, administrators in cognizant federal agencies, and directors of major computation centers – was compiled and a brief questionnaire (appendix A), requesting only names of large-scale users, distributed to them. They responded with about 1,600 names.

The second task, the design of the detailed questionnaire, was started in parallel with the first task. The aid of a few users, computer designers, and directors of major computer centers was sought through interviews. A preliminary questionnaire was pretested on 12 users before the final

---

[1] The proceedings of this symposium were published as: High Speed Computer Algorithm and Organization, edited by D. J. Kuck, D. H. Lawrie, A. H. Sameh, Academic Press, 1977.

questionnaire was written (appendix B). The final version was distributed to the 1,600 people cited earlier plus about 200 others identified in an ad hoc manner.

The design of the questionnaire was such that the first questions asked about a respondent's principal area of interest. Subsequent questions asked increasingly detailed questions about the mathematical model and the algorithm, which led the respondent to consider the system architectural implications. More specifically the questionnaire had the following format.

I. Nature of the problem area of interest: word descriptions, relevance to society, relevance to other disciplines.

II. Model of the problem: mathematical statement using either the laws of physics representing reality or heuristic/statistical analysis of raw data in search of structure; the formal solution to the problem.

III. Algorithms for finding solutions: the process of algorithm selection is an attempt to find the most effective bridge between the problem model and the machine architecture that will be used to work through the formal solution; the process is affected considerably by the self-imposed constraint to use FORTRAN.

IV. System architecture implications of the algorithm: as the cost of computer components continues to decrease, as the relative costs of computer components change, as the various kinds of memories available become larger, and as concurrency of many processors become feasible machine architecture will change and the number of design degrees of freedom will increase correspondingly. There is a corresponding increase in the range of algorithms possible. Conceivably more desirable algorithms can be identified that will drive the computer system designer.

V. Recent experience with current large-scale computers: classification of questionnaire respondents by extent of current large-scale scientific computing.

About 10% of the questionnaires were returned completed. They were sorted by application area and graded using two different schemes:

1. Comprehensiveness of the sections dealing with the substantive problem area

2. Comprehensiveness of the sections dealing with computer system features important to the user

The third task was to summarize the questionnaires and to prepare a preliminary report of the information gleaned from the returned questionnaires.

Following preparation of the preliminary report, the workshop was organized. Invitations were then extended to questionnaire respondents who were selected on the basis of the completeness of their responses and whose problems covered a wide variety of areas. In addition, each of the principal vendors of large-scale scientific computers was invited to send one representative. Prior to the workshop, each participant was sent a copy of the preliminary report together with an agenda (appendix C) for the workshop. A total of 30 persons attended (the attendance list is given in appendix D).

During the first morning of the workshop there were several brief presentations: a report on the purpose of the study, a summary of the questionnaire responses, and two reports on different approaches to new large-scale computations. The participants were then organized into five working groups, each with a complete set of questionnaire responses. Each group contained at least one vendor representative. The groups were then given the following set of discussion questions, with the understanding that the questions were intended to provide a starting point, not to impose limits:

1. Will the needs of the large-scale computations community, as specified by the responses (and the group's own feelings), be met by the normal growth of computational technology?

2. What additional computational capability will be needed to meet these needs?

3. Where it is anticipated that capability will be lacking, what research is needed so that the needed capability will become available?

4. For the large-scale problems, should the machine architecture be related to the algorithm to be solved? If so, how can many different users, with different algorithms, use the same machine?

5. What research is needed in algorithm development to make them faster, more universal (if desired), and to take advantage of highly concurrent devices?

6. Are there "natural partners" who can jointly sponsor the design of a new large computer to meet their common needs? Who are these partners?

Each group elected one of its members to present a report the following morning. The group reports varied widely – some discussed specific questions from the six listed, others discussed more general areas not directly related to any of the questions in particular.

One of the workshop recommendations was that a summary of the workshop reports be appended to the preliminary report and that the preliminary reports be sent to the remaining questionnaire respondents together with a copy of the respondents' original responses, for comment. That was done for about 100 respondents who were selected on the basis of the quality of their original responses. About 25 sent back a substantive comment or suggestion.

About this time, the program organizers of the AFIPS National Computer Conference, scheduled for June, 1979, in New York, invited the project directors to organize a panel as part of the NCC 1979 program. The remarks of the panel were presented in session 5. The immediately following session 16 had essentially the same theme and structure although it was more narrowly focused. The agendas for both sessions are given in appendix E. Session 16 emphasized the great effect on scientific computing of an add-on processor designed to do floating point arithmetic very rapidly. The focus for that session was the Technology Assessment Study of Near-Term Computer Capabilities and Their Effect on Power Flow and Stability Programs (EPRI EL-946, TPS 77-749, Final Report, December 1978). The final report of that study contains an excellent summary of current or near-term computer products that bear on large-scale scientific computing. After the two sessions, which were attended by about 100 persons, there was a discussion of the need for a follow-on activity in order that users of large-scale computers might exchange experiences, across disciplines, and possibly improve their effectiveness in the tasks they were performing.

The final project report was then written, starting with the preliminary report and taking into account all the comments received.

SUMMARY OF FINDINGS

Questionnaire Responses

The basic data received from the questionnaire respondents were collected and are tabulated in the appendixes as follows:

Appendix F– Questionnaire Responses by Applications Area. The breakout by application area is, of course, somewhat arbitrary. The list contains 23 areas, some very broad (e.g., agriculture) and others very specific (e.g., laser fusion). There were 148 responses judged suitable for serious study.

Appendix G– Questionnaire Responses by Affiliation. There were 85 respondents from 55 universities, 20 from 14 private companies, 33 from 18 government laboratories, and 10 anonymous.

Appendix H– Summary of Responses Regarding Discipline-Oriented National Computer Facility. The responses are grouped as for or against that method of resource sharing.

Appendix I– Machine Requirements – from 37 Selected Responses. This set of 37 was selected on the basis of the thoroughness and quality of responses within each application area.

Appendix J– Machine Requirements – Other Responses. Requirements from a second different set of 54 selected responses, this set represented responses that were a cut below the first set of 37 in thoroughness and quality.

Appendix K– Algorithm Components – 37 Selected Responses. Additional information is provided, keyed to specific questionnaire responses, that tells more about the nature of the problem solving methods used and the extent of current machine use.

Appendix L– Special Functions Desired as Hardware Components. The recent rapid growth of use of an attached processor, the floating point AP 120B, has led to renewed interest in other hardware components and some 17 such were called out explicitly by the respondents.

In addition, the basic numerical data have been summarized in figures 1–6. In these figures, each point represents one response. Since not all respondents answered all questions, the number of points in the figures differs from the number of responses summarized. Because it was thought that the essence of most of the machine requirements could be obtained adequately from the data of appendix I, figures 1–5 use only those data. However, the memory requirements were stated adequately in many of the rest of the responses, so figure 6 presents data from both appendixes I

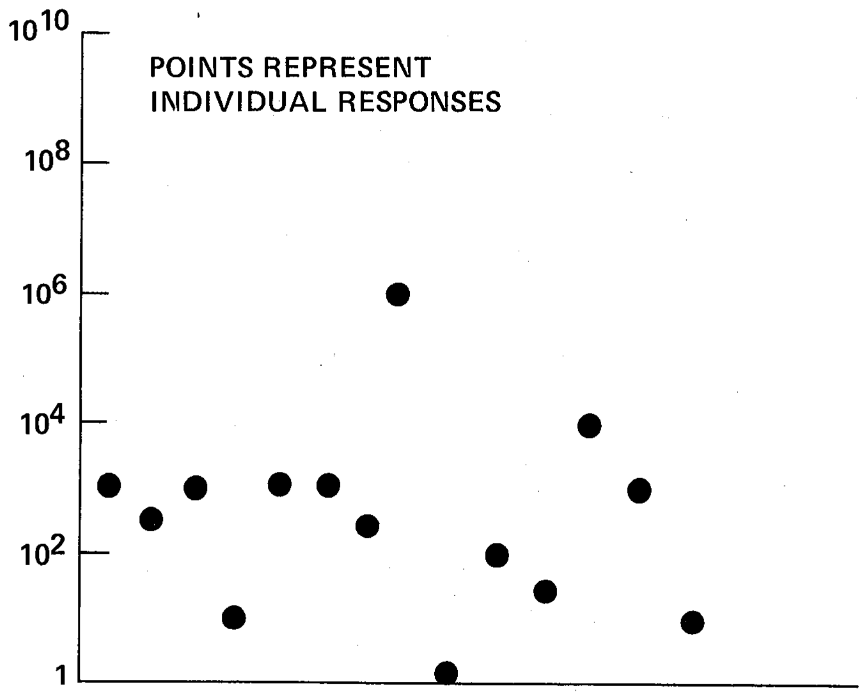

Figure 1.– MFLOPS.

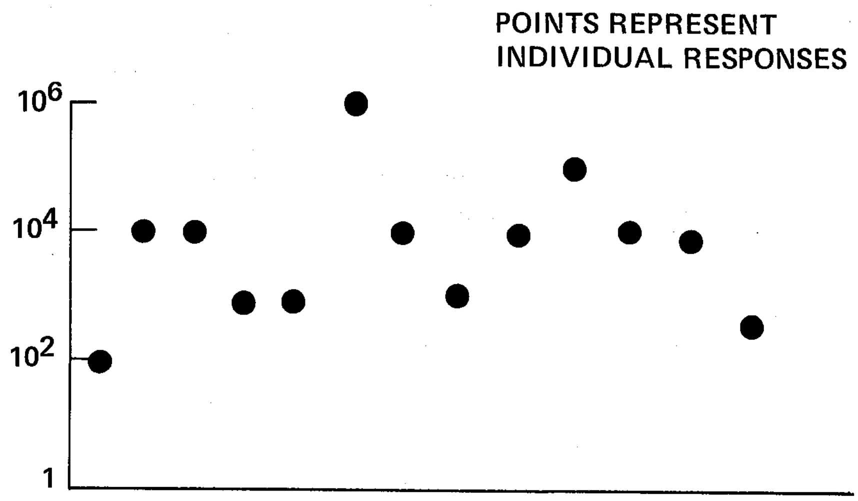

Figure 2.– Vector lengths (words).

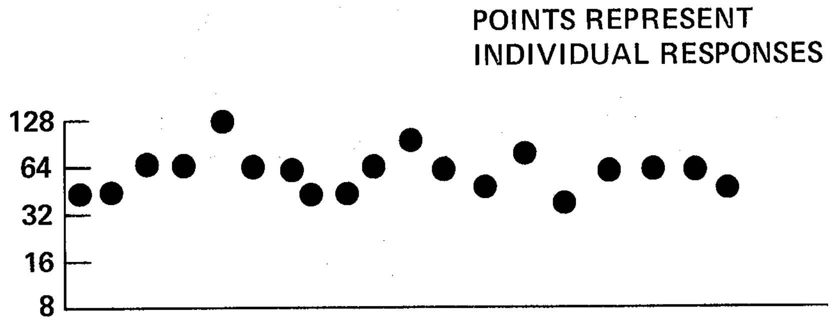

Figure 3.– Word length (bits).

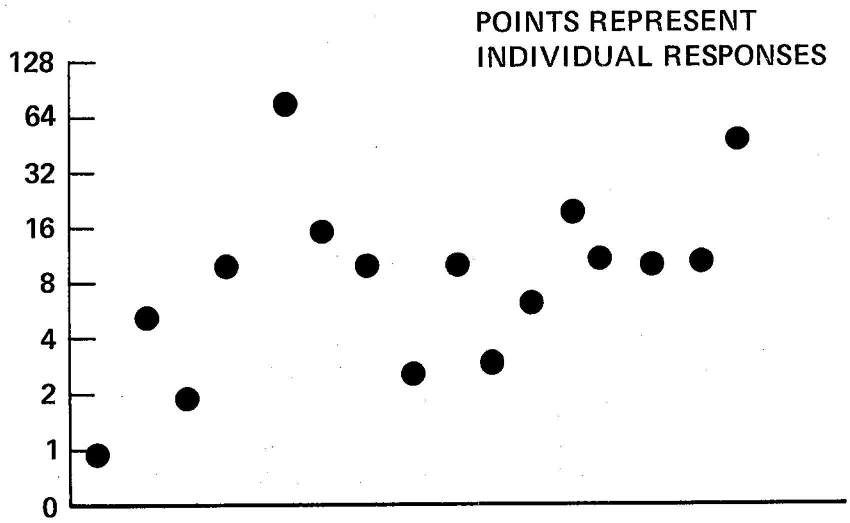

Figure 4.– Operations per memory access.

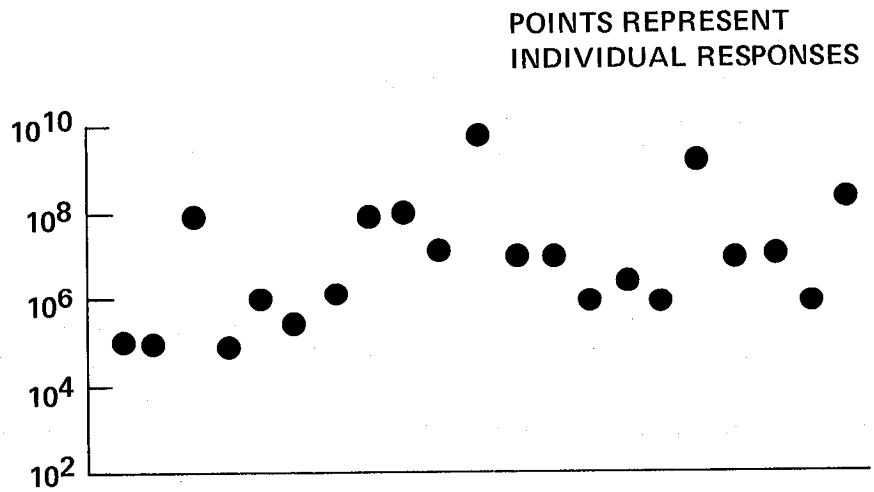

Figure 5.– Active memory size – mesh points times number of variables per point.

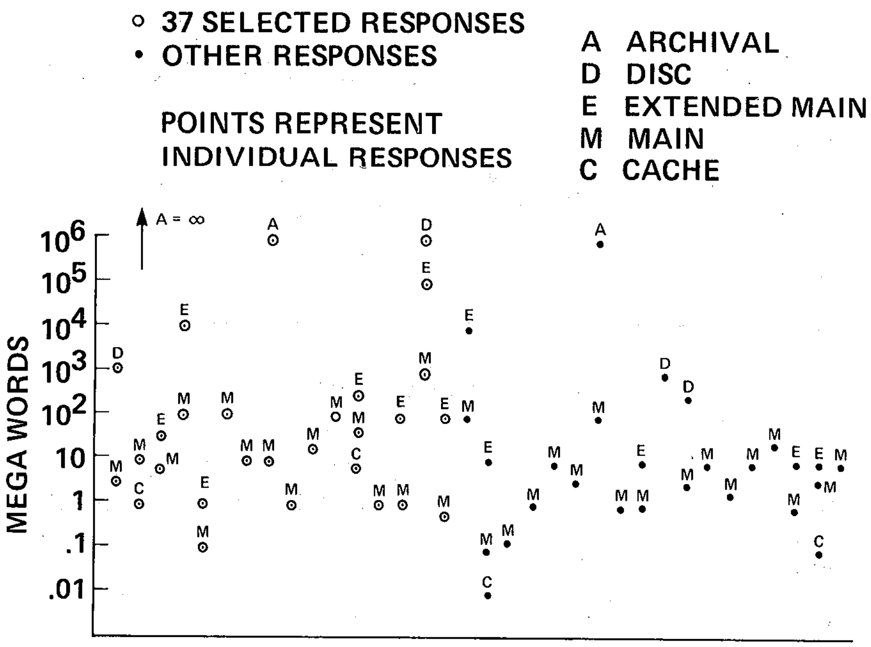

Figure 6.– Memory requirements.

and J. The desired computer speed (see fig. 1) ranged from 1 to 10^6 MFLOPS (millions of floating point operations per second), with most requests in the range of 100 to 1000 MFLOPS. In general, the desire was for faster and faster computational speed. The desired vector length also varied widely (see fig. 2) ranging from 1 to 10^7 words per vector. Most of the lengths were between 10^2 and 10^4.

The desired word length (see fig. 3) was clearly dictated by the word length of existing machines. The requirements were either 32, 48, or 64 bits, with a few requests for multiple precision up to 128 bits. The CDC-popularized word length of 64 bits was greatly preferred over the IBM-popularized length of 32 bits. This question was answered by more respondents than any other. Apparently, they had experience in this area and knew the word-length requirements, even if they could not give other numerical data. (The IBM word length of 32 bits is, of course, generally not recommended for scientific computing that requires high precision, and IBM has integrated extended precision features (64 bit words) into its hardware and software where it is to be so used.) The so-called superminicomputers coming to be used for small-scale scientific problems tend to be 32-bit-word-length devices. The bone of contention seems to be: Is 32 or 36 bits per word usable at all? Is 48 acceptable? Or is 64 (or 60) or 128 (or 120) required?

Figure 4 shows the number of floating point operations done for each word accessed from memory, with the responses varying from 1 to 128; however, most responses were in the range of 2 to 16. It was clear that most were unsure about this parameter, and yet it is important to machine designers: it gives them an indication of how easily they can overlap arithmetic with memory accesses.

Many computer problems can be defined in terms of a computational mesh with a vector of data at each mesh point. Figure 5 shows the product of mesh size times the vector length, which gives an indication of the amount of active memory being used during the major computational periods. This product varies from 10^4 to 10^{11}, with most responses being between 10^4 and 10^7.

The memory requirements are shown in figure 6. Each column on the figure represents the requirements of a single respondent. Thus, the first respondent needs 3 M word of main memory and 10^3 M word of disc memory. A summary of this figure indicates a need for archival memory of 10^6 M words; an extended main memory of 10^3 to 10^4 M words, with a few needs as low as 10 M; a main memory of 1 M to 100 M words; and a cache memory of 0.1 M to 1 M words. Of those who provided data, 15 specifically suggested a memory hierarchy of two or more levels, and 17 specified only a main memory. Most of those 17 will also have a need for tape or disc archives as well, so that even here some level of hierarchy is probably needed.

A brief summary of a few of the nonnumerical responses is indicated below:

1. Language – mostly extend existing language (i.e., FORTRAN)

   - Willing to convert existing programs if the result is an increase in speed by a factor of 2–10.

   - A few willing to tolerate a new language.

   - A few stated that a new language is required, but gave no language specifications.

2. No comments about specific algorithm dependence on machine architecture

3. Not much new in terms of output requested but considerable desire for graphics, moving graphics, and color-3D graphics

4. Special needs – the following were suggested by at least one respondent:

   - Hierarchy of computers; that is, number cruncher and organizer and I/O

   - Broad bandwidth of data transfer between memory and computer – 10^8 words/sec

   - Data sorting

   - Error detection and correction

Ninety-one of the respondents requested special mathematical subroutines, either in hardware or built-in software. In the questionnaire, we suggested the following routines as candidates for special subroutines: complex arithmetic, trigonometric functions, FFT, exponentiation. Most respondents selected from that list, but a few added additional items. All of the special routines mentioned in the responses are listed in appendix L.

Workshop Statements

Before coming to the workshop, each attendee had the preliminary report for study; at the workshop, each group was supplied with a set of all the responses for its deliberations. Each of the working groups presented the results of their deliberations; some of the presentations were structured according to the six questions provided. The responses to those questions (and an additional comment not in response to the questions) are summarized below.

Question 1. Will the needs of the large-scale computations community, as specified by the responses (and the groups' own feelings) be met by normal growth of computational technology?

The feeling was that much of the need would be satisfied by normal growth but that for some of the very largest users no effective alternative exists to the requirement of more speed and memory than will probably be available. The largest users are probably working on nuclear device design, fusion research, weather modeling, and computational fluid dynamics.

Question 2. What additional computational capability will be needed to meet these needs?

The user interface was stressed as an area where new capability is needed. Improved input techniques, software to match the hardware extensions, improved output, and high bandwidth I/O requirements were specific objectives, all aimed at improving the productivity of the researcher, not the machine.

Question 3. Where it is anticipated that capability will be lacking, what research is needed so that the needed capability will become available?

Several different suggestions of needed research were made. A need was cited for hardware designed to handle large ($10^6 \times 10^6$) sparse (about 10^7 nonzero elements) matrices, and to perform a variety of matrix operations including the solution of linear equations. There is also a need for pattern recognition to allow the user to visualize answers in three dimensions. There are further needs for more multiprocessors, including automatic assignment of the multiprocessor capabilities, and better mass storage technology.

Question 4. For the large-scale problems, should the machine architecture be related to the algorithm to be solved? If so, how can many different users, with different algorithms, use the same machine?

There was a general feeling that the machine and the algorithms are definitely related. If one designs a machine too closely around a specific algorithm, however, the machine may have insufficient generality. Microprogramming is one way of restructuring the machine to match various algorithms but not for the time-critical parts of the algorithm. Alternatively, one may need to dynamically reconfigure the machine to match the applications. Also, the use of attached boxes to do specific functions should be considered. If the algorithm can be directly reduced to hardware (and if enough units are to be made), dramatic enhancements in capability can be made. An example of this is the TI "Speak and Spell" toy. Another is the IBM 3838. Research is needed into ways of mapping an algorithm, or problem, directly onto silicon, rather than onto an existing computer system.

Question 5. What research is needed in algorithm development to make them faster, more universal (if desired), and to take advantage of highly concurrent devices?

Further research is needed to develop algorithms that are inherently parallel, rather than the inherently serial ones that are now available. This research should consider the form of the parallelism, and how that form affects the multitude of possible architectures. Development of pattern recognition algorithms, and improved graphic input and output is needed. The new large-scale processors will probably be some form of multiprocessor systems supporting multiple instruction, multiple data streams.[2] To use this type of architecture effectively, a language is needed that will be able to efficiently utilize the parallelism inherent in the algorithms. In addition, there is a need for research into the teaching of parallel algorithms, and for research into means to encourage users to do parallel programming in the present computer environment, which is serial in nature. This implies a way to translate parallel algorithms into serial computations and, probably, vice versa.

Question 6. Are there "natural partners" who can jointly sponsor the design of a new large computer to meet their common needs? Who are these partners?

None of the groups had a good answer to this question, although there was considerable discussion of it. One group divided problems into those that work on a grid, perhaps stretched, in which properties are associated with a point in space, and those in which a grid is neither needed nor wanted and in which properties are associated with a particular particle or entity. Time did not permit further development of this question to identify specific problem areas in which those interested could cooperate in the design of a computer.

There was one additional comment that did not relate directly to any of the questions. It had to do with the opinion that much research is needed into the man/machine interfaces of very large computers. How does one generate and verify the very large amounts of input data that some of these problems require? How does one present the results of a billion-grid-point problem to the user so that he can understand the meaning of the results? That is, how can the user intelligently and easily interact with such a large computer?

Individual Comments

A set of 24 comments selected from those received after the workshop are presented in appendix M. The comments are presented here as they were received.

The comments comprise the following: (1) concerns that too much emphasis on computer power will distract problem solvers from thinking about the problem, and will further perpetuate and institutionalize durrent problem-solving methods; (2) the need to reduce significantly the barriers inhibiting the vendors from exploiting the technology toward larger throughput scientific computers; (3) underscoring the plight of the academic researchers who are severely limited by their own institution's inability to meet the ever-expanding scientific computational needs; (4) the need

[2] In the Sept. 1979 Datamation, pp. 90–94, there is a description of a USN-sponsored project underway at Lawrence Livermore Laboratory. This S-1 project is designed to put super-computers on a chip within 4 to 6 years. The S-1 is a multiple instruction, multiple data stream machine, where each of its 16 uniprocessors (36 bit word length) runs benchmark tests as fast as the Cray-1.

to better interface the user and the system so that increased computer power and algorithm complexity will not decrease the users ability to control and understand what is happening in the system; and (5) comments about specific hardware and software characteristics.

CONCLUSIONS

1. There is a substantial communication gap between the various users of large-scale scientific computers.

2. There is a substantial communication gap between the users and the designers of large-scale scientific computers.

3. It is no longer sufficient for the users of large-scale scientific computers to express their increasing needs with the simplistic statement "bigger and faster."

4. The rapid proliferation of microcomputer- and minicomputer-based systems has as one consequence the rapid disappearance of the computer programmer as an elite class.

5. A new generation large-scale scientific computer will be a highly concurrent device whose power will be realized in direct proportion to the willingness of the user community to develop problem solutions for that kind of architecture, that is, abandon FORTRAN.

6. The computer vendor is not going to develop a product line of supercomputers based on the needs of one laboratory or of one small group of users.

7. Business considerations cause computer manufacturers to build on their sizable present investments in architecture, software, and system structure and thereby detract from investigations into substantially different architectures.

8. This (item 7) is a lost opportunity for the users, for those who depend on the quality and timeliness of their results, and for the computer industry itself.

9. The users must organize in order to determine to what extent there are classes of problems and corresponding algorithms, and to provide some measure of their importance, so that the computer vendors can reduce the uncertainty in estimating the needs of the marketplace and reduce the risks of entering that marketplace to a practical level.

Ames Research Center
National Aeronautics and Space Administration
Moffett Field, California 94035, August 22, 1979

APPENDIX A

REQUEST FOR NAMES OF USERS OF LARGE-SCALE COMPUTERS

Dear Colleague:

The Planning and Analysis Office at NASA-Ames Research Center (home of the Illiac IV) is trying to locate the relatively small number of computer users whose problems challenge or could challenge the very largest and fastest scientific computers, beyond such machines as the CDC 7600 and the IBM 370/195. We hope to learn from them: (1) the algorithmic features of their applications that could help establish design requirements of future very large scientific computer system (VLSCS) architecture, and (2) the extent of their interest in such computers. In addition, we recognize that many problems as presently formulated incorporate restraints on solutions because of limitations of computers currently available or projected extensions of them. We therefore would also like to contact those whose perception of discipline-oriented problems extends beyond currently envisioned extensions of computer technology.

We would appreciate it if you and your associates would note below the name, address and problem-area of any such computer users and return it to us within one week. The users would then be sent a questionnaire about their future needs. Please do not hesitate to list a name, even though you feel it might be forthcoming from some other source.

This request is being sent to a selected subset of the following groups:

(1) Academic departments granting the Ph.D. in engineering, sciences, and business and economics

(2) Computer Manufacturers

(3) Computer centers with large computers

(4) Federal research laboratories and federal agencies

(5) Industrial research laboratories

(6) Professional societies including editors of corresponding journals.

As machine generated lists are being used there may be some duplication of addresses.

Thank you for your help.

Sincerely,

Peter Lykos
Illinois Institute of Technology
Chicago, Illinois

and

John S. White
NASA-Ames Research Center
Moffett Field, California
Co-Principal Investigators

Note: This request is authorized by law (44 U.S.C. 3101 and 42 U.S.C. 2473). While you are not required to respond, your cooperation is needed to make the results of this survey comprehensive, accurate, and timely.

Check here if you want to receive: ☐

☐ A Questionnaire

☐ A Copy of the Final Report, Mid 1977

and include your name and address.

Name:

Address:

Dear Prof. Lykos,

Here are the names and addresses of those who I think should receive your questionnaire.

APPENDIX B

QUESTIONNAIRE

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California
94035

NASA

August 1977

Dear Colleague:

The Ames Research Center of the National Aeronautics and Space Administration and the Illinois Institute of Technology are conducting a survey to determine the projected user needs for very large scientific computer systems (VLSCS) (as opposed to data processing systems) which might become available in the 1985-1990 time period. We hope to learn the algorithmic features of applications which will suggest the design requirements of future VLSCS's, and to determine the extent of national interest in such computer systems. This survey is authorized by law (44 U.S.C. 3101 and 42 U.S.C. 2473). An analysis of the results of the survey will be published.

We are requesting your assistance by asking you to complete and return the attached questionnaire. As our total sample is relatively small and the range of disciplines spanned is relatively large, your cooperation in responding to as much of the questionnaire as possible is needed to make the results of the survey comprehensive, accurate and timely.

Sincerely,

John S White

John S. White
NASA-Ames Research Center

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California
94035

NASA

Form Approved
O.M.B. No. 104-S76002
August 1977

SURVEY OF USERS' NEEDS FOR LARGE-SCALE COMPUTATION

The purpose of this survey is to determine from the users of large-scale scientific computers a clearer picture of their future needs. It appears that the designers and vendors who could create and market very large scientific computers have been inhibited from so doing because of the large uncertainty about design features and the high market risk currently associated with such ventures. This project represents the first formal, sweeping and public attempt to approach this problem from the users' point of view. The users we wish to contact are those whose progress as problem solvers is being limited by the capability of current large-scale scientific computers, such as the widely used CDC-7600 and IBM 370/195, the STAR 100 and the TI-ASC, or near term extensions such as the CRAY 1 and the recently announced Burroughs Scientific Processor.

To start the project, a request for names of large-scale computer users was sent to 5000 individuals with administrative responsibility in industrial research labs, federal labs, and federal agencies as well as to university deans and department chairmen, directors of large computer centers, and editors of appropriate journals. The response to that request is being used as the basis for this mailing which consists of an explanation of the project and a questionnaire on user needs. The responses received to this questionnaire will be reviewed and analyzed. At that time it may prove desirable to convene a small group of users and computer designers to review the preliminary results before the final report is published.

The users being surveyed range from those who only conceptualize the problem and its formal solution to those who, in addition, design the corresponding algorithms and perhaps even generate and "fine-tune" the code in order to take full advantage of the available system architecture. Accordingly, the questionnaire asks separately about the statement of the problem, its formal solution, the solution algorithms, and the architecture, so that more than one person may be involved. The questionnaire has been designed to elicit the following kinds of information:

1. **The importance to society that the problem be solved.** This may be difficult to state, particularly for basic research where the social benefits ultimately to be realized cannot be clearly identified, let alone be made quantitative in terms of social values. Yet the importance of your response hinges on that statement.
2. **The credibility of the problem solver.** The problem and various formal solution methods need to be stated in such a manner that one's peers in the corresponding discipline can make a judgment regarding the importance of the problem within the discipline and of the reasonableness of the formal solution.
3. **The algorithms.** If known, they need to be stated in sufficiently clear terms that the problem-solver's peers can understand them from the perspective of the problem, and the system designer can understand their implications in terms of the system architecture needed to actually implement the algorithms.
4. **The computer architecture implications of the algorithm.** To the extent that you know them, they should be described in terms that system designers can understand. Improvements in system power may come from designing improved software and hardware (e.g., fast Fourier transform hardware) as well as from pure improvements in cycle time.
5. **The responder's computer background in large-scale scientific computation.**

If you are experiencing more than one problem area/computational bottleneck, you should reproduce the form and submit one for each area of interest. Should you prefer to respond in a different manner which you feel would be more informative, please do so. We anticipate that much will be learned from the responses to the questionnaires which was not anticipated by the designers of the questionnaire nor by those who previewed it.

The purpose of the project is to find out what the users need in the way of very large-scale scientific computers. Hopefully, you can and will contribute to this first step toward defining and meeting that need.

NAME (optional):

TITLE:

DEPARTMENT:

ORGANIZATION:

Please return within 30 days to:

Prof. P. Lykos
Illinois Inst. of Tech.
Chicago, IL 60616
(312) 567-3430

If portions of your response are to be quoted in the final report, would you like the opportunity to review them before publication? ☐ Yes ☐ No

If a small workshop were to be held to review the preliminary analysis and provide more definitive conclusions, would you be willing to participate? ☐ Yes ☐ No

Would you like a copy of the final report? ☐ Yes ☐ No

Please read the entire questionnaire first before you generate your response. Remember, we are concerned with the projected needs for a computer which would not be available before the 1985-90 time period. Please consider your long-term requirements, and not necessarily a short-term extension of your present effort. Note that the responses to the questionnaire will be in the public domain, but the individual names and organizations will be deleted prior to any public release. You are free to distribute copies of the questionnaire and its prelude to other professionals whose problems require large scale computers for their solution.

This questionnaire has five sections:

I. Nature of the problem area of interest
II. Model of the problem
III. Algorithm for finding solutions
IV. System architecture implications of the algorithm
V. Recent experience with current large-scale computers

Please feel free to add additional comments in any area, even though they may not be in direct response to the questions. If any response requires more space than allotted, please be sure to identify your extra pages.

I. **Nature of the problem area of interest to you.**

A. State the nature of the specific problem of interest to you and why it is important scientifically or in the national interest that it be solved. What is the value of its solution economically? What are the consequences of it not being solved?

B. In the context of your specific problem of interest, comment on the interface between your discipline/technology and others, and the benefits to be realized in that larger sphere.

C. Give your perception of the long-term (1985-90) problem towards which the specific problem you are addressing is leading, and identify the conceptual barriers (such as mathematics, logic, algorithms, model of physical phenomena, etc.) which restrict you from attacking the total problem now.

D. List what you perceive to be the advantages and disadvantages of using a national center for computation in your area of interest, assuming remote access to be available. Such a center would be along the lines of the ERDA Magnetic Fusion Energy Computer Facility (based at Lawrence Livermore Laboratory) or the evolving National Resource for Computation in Chemistry (J. Chem. Info. and Comp. Sci. **15**, 137 (1975)).

E. What procurement cycles are likely to take place for your laboratory or computation center between now and 1988? Will these procurements be justified primarily on your group's research requirements—if not, please list the other groups or major research laboratories which are likely to be involved.

II. **Model of the specific problem of interest to you in the 1985-1990 period.**

A. Please give a concise, but reasonably complete statement of the specific problem you would like to solve. (If your specific problem does not lend itself to a mathematical statement or representation, but rather you need to adopt a phenomenological approach, please describe the techniques to be applied to the data to extract the structure you believe is inherent in the chosen collection of data.)

B. State the source of the data base for your model and comment on the availability of that information. Will the data be available in computer readable form? How large is the required data base, and what is the precision of its elements?

III. **The algorithms (or procedures) for solution of the problem as defined by the model.** The evolution of computer architecture may tend to be away from single serial machines to multiprocessor and parallel devices. A major barrier to the realization of the potential of multiprocessor machines may be that numerical algorithms currently in use have been designed around serial machines, not highly concurrent machines. Present-day computers may not be able to use algorithms as you state them. The answers below may be used as guides for the design of future computers.

A. Give the method of solution.

1. Describe the algorithms working from categories such as Monte Carlo, coupled partial differential equations, finite element, linear programming, signal processing, list processing, pattern recognition, information retrieval, numeric or symbolic, etc.

2. Describe the structure of the problem data base (matrix, relational, hierarchical, flat file, others) and characterize its parts with respect to use (is access random, ordered by row or column, by row and column and/or other dimension, indirect, other), size, change in composition or state, etc.

3. Outline the specific algorithms in sufficient detail that the impact on the computer architecture can be determined.

4. What simplifying assumptions have been included? Should they be refined or eventually replaced by calculations?

5. Computer designers give considerable attention to the relative amounts of array operations and recurrence operations in the job stream. Accordingly, rather detailed analyses of your algorithms in this regard are desired.

   Do the calculations involve array operations A(I) = f(B(I) . . .), recurrence operations A(I) = f(A(I–1) . . .), or both?

   Do these operations apply over entire arrays or subarrays or do they terminate before an entire subarray has been processed?

   Within the loops, what is the frequency of branching based on the index, or on values computed internal or external to the loop?

   How many constants and variables are used inside the loops?

   How deeply are the loops nested?

   What is the frequency of table look-ups? Are they used for computational speed up, or to access input data?

6. If you can, please give some idea of the potential for concurrency and parallelism such as vector operations. Some of the relevant characteristics include the following:

   Are the matrix or vector operations typically over full or sparse matrices or arrays?

   Are special storage mappings utilized (i.e., for symmetric, triangular or band matrices, etc.)?

   If sparse matrix techniques are utilized, what addressing mechanisms are used (i.e., bit maps, index vectors, threaded lists, etc.)?

   What mix (rough percentages) of vector lengths are anticipated in the problem?

7. Some problems can be defined in terms of a large number of mesh points, and a vector of characteristics at each mesh point. If your problem has these characteristics, how many mesh points are there?

   How many variables are associated with each mesh point and what are the mathematical functions used (e.g., SIN/COS, EXP, LOG, SQRT, etc.)?

8. If you can, give some measure of the ratio of vector results to scalar results. Actual measurements would be extremely valuable.

9. List any other algorithm attributes you believe indicate a basis for extending or modifying current large processor computer architectures.

10. Which current processor facilities fit well with the algorithms but need to operate faster?

B. If you, the problem solver, are to be an integral part of the solution process, describe the frequency with which you need to be presented with an interim result or branch point for a decision, what output would be desired, and what type of modifications to the computationsl process would result?

C. Do you believe that new architectural designs will be required to solve your problem such as vector pipelined (e.g., CRAY I), multiprocessor, parallel array (e.g., ILLIAC IV), hybrid, etc.? Please feel free to give specific architectural ideas that will make your problem run best.

IV. **System architecture implications of your algorithm.** A computer may be thought of as a memory, one or more central processing units consisting of an instruction decode unit and one or more functional execution units and corresponding connectivity requirements. To the extent that you feel able, both answer the following questions and prioritize (vital [1] ; important [2] ; desirable [3] ; unimportant [4]) their importance.

A. Operational characteristics of the system architecture.

☐ 1. What sort of memory hierarchy is needed, and how much memory at each level? What are expected hit ratios in such a hierarchy if automatic data staging is to be included? What type of staging algorithms would be most useful? Is the accessing pattern regular enough that user controlled I/O would be beneficial?

☐ 2. What bandwidth is needed to transmit data to and from CPU, I/O, and the various levels of the hierarchical memory? What memory access time or delay expressed in clock periods is acceptable?

☐ 3. Is concurrent access to the data required from more than one program? At what level would such sharing occur?

☐ 4. What special access requirements to the memory are needed (such as sweeping through a multi-dimensional matrix in several directions)? Can the problem be structured to stream data sequentially or is high speed random access always required?

☐ 5. Would a special structure in the memory simplify your algorithms? What structures? Any new data management techniques? What type of back-up memory is needed?

☐ 6. What word length is required (precision and exponent range of floating point numbers, and range of integers)? Would two or more word lengths, or variable length numeric precision, be required? If so, what are the lengths and how is the extent of their usage distributed? What sort of accuracy and round-off error controls are needed (e.g., symmetric rounding, guard digits, etc.)? What level of priority do you associate with the availability of instructions to support or facilitate multiple precision software packages?

☐ 7. What fixed point precision(s) is (are) required in the computations?

☐ 8. What sort of exception handling procedures and error handling mechanisms are needed (e.g., overflow/underflow exception flags, precise versus imprecise interrupt of pipeline execution units, special formats for zero, $\pm\infty$ and indeterminate results, dynamic arithmetic instruction error detection, instruction retry, etc.)?

☐ 9. How many floating point operations are executed per operand accessed from memory?

☐ 10. What is the throughput requirement in floating point arithmetic operations per second and/or in scalar or control instructions per second?

☐ 11. Do you use bit or character strings? Are they fixed or variable length? What is the average and maximum length? What operations do you perform on these strings?

☐ 12. What are the constraints on solution time, and are these caused by real-time considerations or daily production schedules, etc.? What would happen if such solution or production run took twice the time?

☐ 13. What error rates in CPU, memory and peripherals are tolerable?

☐ 14. Would network access to data bases be required? If so, what bandwidth is required in such a network?

☐ 15. Other specifications?

☐ B. What is the required system reliability, in terms of Mean Time Between Failure, fail safe ability, recoverability, or other? Is check point/restart required, or is re-run acceptable? Is transaction level recovery required for concurrent access to data?

☐ C. Are there any special physical constraints on the computer system, such as size, power weight, etc.?

☐ D. Are man-machine interface features important to the process of solution such as visual, alphanumeric, graphical, audio (two-way), tactile (manual/analog), response time, etc.? State the format desired for your output: Printed summaries, graph(s), moving graphics, color, sound, hard copy, CRT display only, etc. What is the rate at which intermediate results are saved and the ratio of snapshot output to total output? Can you give an estimate of the volume of the total output?

☐ E. Would it be desirable to have special components for such items as complex arithmetic, trigonometric, functions, fast Fourier transform, exponentiation, etc.? What special capabilities are needed?

☐ F. What are the two or three most important system improvements you need? Be as specific as possible.

☐ G. Do you require/believe that the new architecture must be supported transparently (without conversion of existing programs); that the current languages be extended syntactically to support the new machine architecture; or that a new higher level language is needed (what new features should it have)? If either of the latter, what improvement in performance is needed to justify the implied expense and inconvenience involved?

☐ H. What new features should the systems have in its operating system; (collection of utilities, user/system dialog, vector/matrix algebra, debugging tools, etc.)?

☐ I. What about security and privacy? What is your concern?

V. **Current Computer System Experience.**

A. List the specific computer systems you have used recently, give some measure of the extent of such use (hours, or dollars, per year and number of jobs per year), and identify the associated problem areas.

B. List separately large-scale computer use whereby two or more segments of the problem were accommodated sequentially by two or more different machines (whether by networking, tape handling, or other).

C. What programming language do you use, and which version?

D. What type of operating system do you use (batch, interactive, network, etc.)?

E. If any of your code has been analyzed for percentage of floating point operations, branching, integer arithmetic, partitioning relevant to memory requirements, etc., give a synopsis of the result. What tools were used in the analysis?

F. Have the limitations of the computer system available imposed an essential constraint on the solution of your problem? If so, list the simplifying assumptions made to allow a solution. What is the impact of those assumptions on the use or applicability of the results?

G. What interest do you have in exploring use of parallel architecture such as the ILLIAC IV (64 parallel processors)?

APPENDIX C

WORKSHOP

Future Computer System Needs for Large-Scale Computations
Pajaro Dunes, CA
16–18 November 1978

AGENDA

Thurs. 16 Nov.		Arrival at Pajaro Dunes. Informal discussion. Buffet late supper at House 103
Fri. 17 Nov.	7:30 A H103	Breakfast
	8:30 A H133	Introduction. Purpose of Survey. Goals of Workshop
	9:00 A H133	Description of NASA Ames Research Center. "Numerical Wind Tunnel" project. R. F. Bailey, NASA
	9:40 A H133	"A Quantum Chemist and Fifty Coupled LSI 11's (Cm*)" N. Ostlund, Carnegie-Mellon University
	10:00 A H133	Review of questionnaire responses
	11:00 A	Form working groups to discuss questionnaire, responses, and draft report. Individual groups plan their activities.
	12:00 N H103	Lunch
	1:00 P	Working groups meet
	6:00 P H103	Dinner
	7:30 P	Working groups reconvene as desired
Sat. 18 Nov.	7:30 A H103	Breakfast
	8:30 A H133	Reconvene as a group. Receive reports from working groups
	~12:00 N H103	Lunch and end of workshop

APPENDIX D

WORKSHOP ATTENDANCE LIST

Prof. Robert P. Futrelle
Genetics and Development Dept.
University of Illinois
Urbana, IL 61801

Prof. Gilles Cantin
Mech. Eng. Dept.
Naval Postgraduate School
Monterey, CA 93940

Prof. D. E. Ellis
Dept. of Physics and Chemistry
Northwestern University
Evanston, IL 60201

Dr. George D. Purvis, III
Chemical Physics
Battelle Columbus Laboratories
Columbus, OH 43201

Dr. Keith A. Taggart
L-6, MS-531
Los Alamos Scientific Laboratory
Los Alamos, NM 87545

Prof. Don E. Harrison, Jr.
Dept. of Physics and Chemistry
Naval Postgraduate School
Monterey, CA 93940

Dr. C. P. Henderickson
Physicist, A-Division
Lawrence Livermore Laboratory
Livermore, CA 94550

Prof. R. F. Stellingwerf
Physics Dept.
Rutgers University
New Brunswick, NJ 08903

Prof. Robert L. Haney
Meteorology Department
Naval Postgraduate School
Monterey, CA 93940

Prof. J. H. Ferziger
Mechanical Engineering
Stanford University
Stanford, CA 94305

Dr. Kenneth D. Senne
Asst. Group Leader, Group 41
MIT Lincoln Laboratory
Lexington, MA 02173

Dr. F. R. Bailey
Ames Research Center, NASA
Moffett Field, CA 94035

Dr. J. J. Erpenbeck
T-Division, MS 457
Los Alamos Scientific Laboratory
Los Alamos, NM 87545

Dr. Lynn F. Ten Eyck
Inst. of Molecular Biology
University of Oregon
Eugene, OR 97403

Dr. Peter M. Silvaggio
NRC Associate
Astrophysical Experiments
Ames Research Center, NASA
Moffett Field, CA 94035

Dr. C. W. Chapman
Geophysics Division
Chevron Oil Field Res. Co.
3282 Beach Blvd.
La Habra, CA 90631

Dr. B. L. Bivens
Lockheed Calif. Computers
P.O. Box 551
Burbank, CA 91520

Dr. Neil Ostlund
Dept. of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Dr. Morris G. Huck
Dept. of Agronomy and Soils
Auburn University
Auburn, Alabama 36830

Dr. Frank Bogden
Analytical Mechanics Branch
Westinghouse Research Laboratories
Bldg. 401 2X13
Pittsburgh, PA 15235

Mr. Wayne Hathaway
Computation Division
Ames Research Center, NASA
Moffett Field, CA 94035

Mr. John White
Ames Research Center, NASA
Moffett Field, CA 94035

Dr. Peter Lykos
Department of Chemistry
Illinois Institute of Technology
Chicago, IL 60616

Dr. Ram Gupta
Sperry UNIVAC
2121 Wisconsin Ave., N.W.
Washington, D.C. 20007

Dr. Steven Lundstrom
Burroughs Corp.
Great Valley Laboratories
P.O. Box 517
Paoli, PA 19301

Dr. Kornel Spiro
Amdahl Corp.
1250 East Arques
Sunnyvale, CA 94086

Dr. Edward Michehl
Control Data Corp.
4201 Lexington Ave. N
Arden Hills, MN 55112

Dr. Don Gibson
Dept. A 74, Bldg. 707
IBM – P.O. Box 390
Poughkeepsie, NY 12601

Dr. Richard Hendrickson
Cray Research, Inc.
1440 Northland Drive
Mendota Heights, MN 55120

APPENDIX E

1979 NATIONAL COMPUTER CONFERENCE

Partial Agenda

Monday • June 4 2:30-4:00 Georgian Ballroom A, SC • SESSION 5

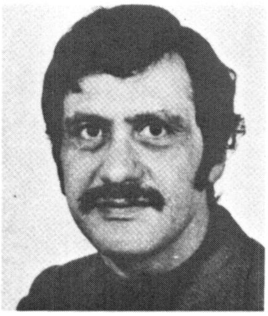

Session Leader:

Peter Lykos
IIT
Chicago, IL

AN ASSESSMENT OF FUTURE COMPUTER SYSTEM NEEDS FOR LARGE-SCALE COMPUTATION

The chairman of the session will describe a joint project with NASA Ames Research Center to assess future needs for large-scale computation. Senne will discuss his attempt to compare the performance of several scientific computers using an application important in signal processing. Ostlund will review his efforts as a theoretical chemist to discover new algorithms for old chemistry problems using the highly concurrent fifty LSI-11 Cm* computer system at Carnegie-Mellon University. Taggart will examine the unfulfilled aspirations of scientific problem solvers at the DOE National Scientific Laboratory at Los Alamos. The goal of this effort is to begin the process of encouraging users of large-scale scientific computers to self organize and to make their needs known to computer designers.

Panelists:

Neal Ostlund
Carnegie-Mellon University, Pittsburgh, PA

Ken Senne
Lincoln Laboratory, Lexington, MA

Keith Taggart
Los Alamos Scientific Laboratory, Albuquerque, NM

SESSION 16 • June 4 4:15-5:45 Georgian Ballroom A, SC • Monday

ARRAY PROCESSING: AN INNOVATIVE APPROACH TO SCIENTIFIC COMPUTING

Array processors are a relatively recent approach to provide cost-effective scientific processing via attached processors. This panel will discuss history, philosophy, and implementation of array processors with particular emphasis on problem-solving environments. The first environment will be computer tomography which has allowed the medical profession to have a greater diagnostic capability. The second environment will be cockpit flight simulators where advances have allowed full aircraft simulation for use in pilot training. This panel also will discuss the basic mathematical requirements of such environments and their adaptation to array processing through the mathematical solutions of sparse matrices.

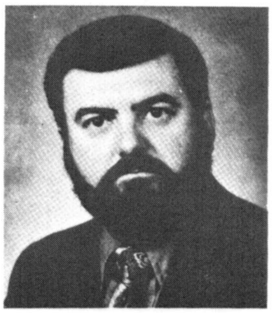

Speaker:

Array Processor Applications to Sparse Matrix Problems
Christopher Pottle
Cornell University, Ithaca, NY

Panelists:

Barry K. Gilbert
Mayo Clinic, Rochester, MN

William Ross
Northrop Aircraft Div., Hawthorne, CA

Ken Thurber
Univac, St. Paul, MN

Session Leader:

Roy D. Gwin
Floating Point Systems, Inc.
Portland, OR

APPENDIX F

QUESTIONNAIRE RESPONSES BY APPLICATION AREA

Topic	Number of responses
Agriculture	2
Astronomy	8
Biology	15
Business/economy	4
Chemical reactions	15
Computational fluid dynamics	11
Computer software	1
Data management/graphics	2
Electric motor/generator design	1
Education	1
Electrical power	1
Laser fusion	1
Mathematical methods	15
Meteorology (weather)	11
Military (warfare models)	3
Nuclear engineering	5
Oil studies	6
Particle physics	10
Plasma physics	6
Quantum chemistry	7
Solid state metallurgy	10
Statistical mechanics/matter models	9
Structures	5
Total	149

APPENDIX G

QUESTIONNAIRE RESPONSES BY AFFILIATION

Note: One response from each institution, company or agency unless indicated otherwise by numbers in parentheses.

Academic

- Arizona, U. of
- Arkansas, U. of
- Auburn U.
- Bryn Mawr C.
- California, Berkeley, U. of
- California, Davis, U. of
- California Inst. of Tech.
- California, Los Angeles, U. of
- California State Polytechnic C., San Luis Obispo
- Case Western Reserve U. (3)
- Clarkson C. of Tech.
- Colorado, U. of
- Connecticut, U. of (2)
- Cornell U. (3)
- Delaware, U. of
- Denver, U. of
- Drexel U.
- East Washington U.
- Eckerd C.
- Harvard U. Business School
- Illinois, U. of (3)
- Iowa State U. of Science and Tech. (3)
- Louisiana State U., Baton Rouge (3)
- Massachusetts Inst. of Tech. (2)
- Michigan, U. of
- Michigan State U. (2)
- Michigan Technological U. (4)
- Nebraska, U. of
- Nevada, U. of
- New York U. (2)
- North Carolina Agricultural and Technical State U.
- Northwestern U. (2)
- Notre Dame U.
- Oregon, U. of
- Pennsylvania State U., Hazleton
- Princeton U.
- Purdue U. (3)
- Rochester, U. of (2)
- Rockefeller U.
- Rutgers, State U. (3)
- Southeastern Massachusetts U.
- Southern Florida, U. of
- Stanford U. (3)
- SUNY, Buffalo
- SUNY, Stony Brook
- Texas, U. of
- Toledo, U. of
- Utah, U. of
- Virginia, U.
- Virginia Polytechnic Inst. and State U.
- Washington, U. of (2)
- West Michigan U.
- William and Mary C.
- Wisconsin, Madison, U. of (2)
- Wisconsin, Stevens Point, U. of
- Yale U.

Industry

- Aerospace Corp. (3)
- Air Products & Chemicals
- Armstrong Corp.
- Battelle, Columbus, Ohio
- Chevron Oil Research (3)
- Data Research Inc.
- Exxon Research
- GM Research Laboratories
- IBM, Palo Alto, California (3)
- Lilly Research (2)
- Lockheed, California
- Pan American
- Science Applications
- Sperry Univac

Government

- AEDC
- ANL (3)

Army Corps of Engineers
Brookhaven
H. Diamond (3)
Kitt Peak
LLL (4)
Los Alamos (3)
NASA-Ames Research Center (4)
NASA-Kennedy Space Center
National Radio Obs.
Naval Underwater Systems
Naval Postgraduate School (3)
NOAA (2)
NRL
Oak Ridge
USA Ballistics Research
West Point

Anonymous (10)

APPENDIX H

SUMMARY OF RESPONSES REGARDING DISCIPLINE-ORIENTED NATIONAL COMPUTER FACILITY

For National Computer Facility

1. Easy communication among all users; forum for exchange of information and ideas; opportunity for collaboration

2. Larger computer (main memory); faster computer

3. Optimized collection of subroutines; standard programs for commonly used algorithms; large complex code created, maintained, and shared

4. Skilled consultants and systems programmers; specialized algorithms

5. Large accurate data base; focus for large data bases (not necessarily processor)

6. Only one machine to learn

7. Users in small organization can access maximum power

8. Leads development of new machines

9. Assemble great deal of computer science expertise; opportunity to explore application of more esoteric C.S. methods

10. Reduce duplication of effort

Against National Computer Facility

The following items summarize the negative comments received regarding a national computer facility. They reflect problems that the respondents anticipate as a result of establishing such a facility. As such, they are items that must be addressed in organizing a national facility, and they can serve as a guide to problem areas that should be avoided.

1. Slower turnaround; not interactive; program update slower and more costly; reliability of remote access difficult to ensure

2. Expensive to transfer large amounts of data

3. Inhibits diversity; inhibits innovation

4. Competes with local facility; university now subsidized on-campus – not so at national center; loss of regional goodwill and support; may give university reason to totally abdicate responsibility to provide computer facilities

5. Inhibits communication with computer-support staff; remote system inherently more difficult to use; more distant scheduling priorities; requires on-site use for orientation

6. For-profit organizations could not use; security problem; security more difficult to maintain

7. Larger bureaucracy, more red tape; another layer of proposal/grant bureaucracy

8. Conservative central system; no non-CDC hardware for special purpose; could become institutionalized, too rigid to respond properly to new developments; system and software inflexibility

9. Overcrowding before a major discipline-oriented conference

10. Users already using large machines would dominate

11. Diverts funds from other researchers to those who already have plenty; another layer of overhead for the user to pay

12. VLSIC's put very powerful computers in anyone's hands; counter to long-term trend to distributed systems

13. Change of programs from local to remote not trivial undertaking

14. Freeze out younger scientists; possible domination by one group

15. Concentration of expertise

APPENDIX I

MACHINE REQUIREMENTS – 37 SELECTED RESPONSES

APPENDIX I

Area	Response No.	Throughput, MFLOPS	Storage, M words	Architecture	"Hardwired" subroutines	New HOL
Quantum chemistry	1	10^3 or more	D=10^9, C=10^3		Sparse mat. op'ns.	Extend existing
Solid state metallurgy	4		C=1		Trans, C, FFT	
	6		A=∞, M=10		Mat. inv., Eigen, T,	
	9	200	E=35, M=6		Exp., diag. sparse mat.	
Particle physics	5				Exp, C	
	9	10^3	E=10^4, M=100		FFT, random, int-fract.	Extend existing
	10	100			Mat, TRID inv.	Existing
Chemical reactions	3	10	E=1, M=0.1		Yes	Extend
	4	Low	M=100		C,4 prec. hdwre	Extend
	13			‖ Processing	T	Existing
Statistical mech./ matter models	5	F	L			
	6	F			FFT, C, T	
	8	10^3	M=10		FFT, Exp.	2:1 for Conv.
Plasma physics	6	10^2	A=10^6, M=10	Assoc. mem.	Yes	4:1 for Conv.
	7	10^2–10^3	L	Assoc. mem.	C, T, FFT, Poisson Sol.	Rewrite crtcl.
Nuclear engineering	3	F				Extend
	4	500				Extend
Astronomy	2	10^6	M=1		C, T, FFT, Mat op'n.	Auto conv.
	3	F	L			
Electrical power	1	1	L	‖, Matrix	C, T	Perhaps new
Meteorology (weather)	3	F			Vector-matrix	3:1 for Conv.
	6				C, T, Random #	Extend (or new)
Oil studies	2	F	M=20		No	10:1 for Conv.
	5	100	M=100	‖	T	Extend or new
Computational fluid dynamics	6	25–50	L		FFT, Exp	New
	11	10^9	E=300, M=40, C=8	Partial MIMD	Multiple arith.	Extend
	8	F	L		FFT, C	New
Military	2				FFT, C, T	Extend
Biology	12	1	C=1		C, T, FFT, Fix, Flt	Extend
	13	F			T, Matrix	Extend
	15	10^4	M=1, E=100	{ Assoc. mem. 10^3 processors	T	Extend
Structures	1	F	L		Yes	Extend
	4	10–10^3	D=10^6, E=10^5, M=10^3		C, T, FFT, vector	Extend
Business/ economics	1		L			New
Agriculture	1				FFT	Extend
Mathematics methods	3	10	E=100, M=0.5		Vector	Extend
	7				Convolution	
Summary		1–10^6		‖ Processing Assoc. memory	FFT, T, C, Matrix Vector, Convolution Eigen, Random Fix, Flt, Poisson Mat. Inv. Sparse Matrix	Mostly extend Some new 2→10=1 for conversion

F=faster

C=Cache
M=Main
E=Extended Main
D=Disc
A=Archival
L=Larger

T=Transcendental
C=Complex

APPENDIX I – CONCLUDED

Area	Response No.	Output	Vector length	Word length bits	Special needs	Operations/ access	Mesh size
Quantum chemistry	1	CRT Graph	100K	48	Sorting $Bw \approx 10^7$ words	1	
Solid state metallurgy	4 6 9	 CRT CRT, G, C		 48 64	H {Graphical language and processors	1–5	10^4P 10V
Particle physics	5 9 10	 G, C G	1–400 $100–10^4$ 100	60 64 40	 $Bw=10^8$ w/sec	Several 2	10^4P 20V 10^7P 15V 10^4P 8V
Chemical reactions	3 4 13	 G	$10^2–10^4$ 10^3 10^4	40–64 128 40	Sort, block mem. access	Few 1–10	 10^4P 10^2V
Statistical mech./ matter models	5 6 8	 CRT G	10^3 $10–10^3$	64 Var 60	 Lower cost	100 2	10^4P 20V 10^6P 10^3P 100V
Plasma physics	6 7	 CRT, G, C		64 32	Bandwidth	1 4–20	10^6P 20V 10^7P 11V
Nuclear engineering	3 4	CRT, G, μ	10^7 10^3	48–60 48	Fast Libraries	10 2	10^7P 20V 10^4P 52V
Astronomy	2 3	G	$10^2–10^6$	48	Matrix speed	3	10^6P 5V $10^6–10^9$P 15V
Electrical power	1	CRT	$10–10^4$	64	Error det and corr	10	10^4P 10^3V
Meteorology (weather)	3 6	G		32 96	Bandwidth Easy rmt access	2–3 1	10^6P 15V 10^6P
Oil studies	2 5	G, CRT CRT		 64	Error det.	2–5 Few	10^4P 100V 10^4P 500V
Computational fluid dynamics	6 11 8	G G G, Movies,	16–64 500 $32–10^6$	48 32–48 64	Bandwidth 3D memory access	1–6 5–50 20	10^6P 4V 10^7P 20V 10^8P 20V
Military	2	G, Tactile, C	$50–10^4$	32–96			
Biology	12 13 15	G G G	$10^3–10^5$	128 (rarely) 36	Integer arith. 64 bit 10^8Bw error tol.	10 several 10	 10^6P 20V
Structures	1 4	G, interact. CRT	 $1–10^4$	120 {64 128 {variable	μ code to represent nonlinear lower cost		10^6P 6V 10^4P 6V
Business/ economics	1			64			
Agriculture	1	G, CRT			Char. string, H		20V
Mathematics methods	3 7	G, Movies	 $32–10^4$	60 and 120 64	 ‖ computers	 2–12	10^5P 6V 10^4P
Summary		G, CRT No specifics	$10 \rightarrow 10^7$	2 @ 32 9 @ 48 13 @ 64 6 @ 96 or 120	Sorting Bandwidth $\simeq 10^8$ Lower cost Error det and corr	1–10 + {1 @ 20 1 @ 100	PxV $10^5 \rightarrow 10^{10}$

G=Graphics
C=Color
μ=Microfiche

H=Hierarchy of computers

P=Pts.
V=Var.

APPENDIX J

MACHINE REQUIREMENTS – OTHER RESPONSES

APPENDIX J

Area	Response No.	Throughput MFLOPS	Storage M words	Architecture	"Hardwired" subroutines	New HOL	Output
Quantum chemistry	5				FFT, Matrix		G, CRT
Solid state metallurgy	1		$D=10^4$, $M=10^2$	∥ or pipeline	Math fns.	Extend	
	2						
	3				FFT, T		
	5						
	11				FFT, T, Matrix		G
Particle physics	2, 3	F	L			Extend	G
	4	F	L		C	Extend	G, CRT
Chemical reactions	1	20			FFT, Mat. Inv.		G
	5, 6				C, T, FFT	Extend	G, CRT
	7		D=10, E=0.1, M=0.01		Exp, Sort, Polyno	Extend	
	10	10^3	M=0.2		T	New	
	11				T	Extend	G, Movie, CRT
Statistical mech./ matter models	1	10^4			T, Matrix	Extend	G, Movie, CRT
	2		L		Yes	Extend	G, Movie
	9		$M=10^6$				
Plasma physics	1		L, $M=10^6$		FFT, T	Extend	G, Movie, CRT, μ
	3		L, $M=10^7$		C, FFT	No	G
	4				Yes	Extend	CRT, μ
	5	F	L		C	No	
Nuclear engineering	1	F	L		T, FFT	Extend	
	2	500				Extend	
Astronomy	5						
	6	50			FFT		G, Color
	8	10^3		∥ CPU			
Miscellaneous (Data management)	2	10	$D=10^6$, $M=10^2$	Assoc. mem.	T	Extend	Moving Color, G
Meteorology (weather)	1		M=1		FFT, C	New, vector	G
	2		$D=10^3$, M=1		FFT, C, TriD	Extend	
	4	10^3–10^6					
	7		$D=10^3$		FFT, Matrix		CRT, G, Color
	8	10^3	$D=10^5$, M=3		T		
Oil studies	1				FFT, C, T	Extend	G, CRT
	3		M=10		FFT, TriD Inv	Yes	Color
	4	10^3	M=2			Extend	
Computational fluid dynamics	2		L	Matrix arith.	No	Extend	G, CRT
	3	10^4	L		C, T, FFT	Yes	CRT
	4	10^2–10^3	M=10			Yes	
	9	100	M=30		Yes	Extend	G
Military	3				APL processor	APL	G, Movie, Color
Biology	1		D=10, M=1				G (Molec. str.), μ
	5				Exp	No	
	6	10	D=10, M=4, C=0.1	Array memory	C, T, FFT, Fix, Flt	Extend	G
	8	F		∥ Processing	Exp, Matrix	Extend	G, CRT
	11	F	L		FFT, C, T	No	3-DG
Structures	5	F	L		No	Extend	
Business/ economics							
Agriculture	2	F	L		Yes	New	G, Movie, Color, μ
Mathematics methods	2, 5			Large scale core address	Yes	Extend	G
	8	F	L			Extend	
	9			∥ Operation	Inner loop	New	
	12		10^7		C, FFT, Matrix	No	
	13				Matrix		

F=Faster

C=Cache
M=Main
E=Extended Main
D=Disc
A=Archival
L=Larger

T=Transcend.
C=Complex

G=Graphics
C=Color
μ=μfiche

APPENDIX J – CONCLUDED

Area	Response No.	Vector length	Word length bits	Special needs	Opinions/access	Mesh size	Responses without data
Quantum chemistry	5			Bandwidth		10^5P	2,3,4,6
Solid state metallurgy	1	10^5	48		2	10^4P 100V	7,8
	2		60			10^7P 5K V	
	3	200–10^4	32–64	Universal JCL		10^3P 20V	
	5						
	11		60		5–25		
Particle physics	2, 3	10–100	64				1–10^6 Array Proc.
	4	50				10^7P 3V	6,7,8
Chemical reactions	1		64, 128			$7(10)^3$P 3V	2,8,9,12,14,15
	5, 6	40–60	32 and 64			10^9P 100V	
	7	200	48		6		
	10	10^3–10^6	20	Error rate	50		
	11	10^2–10^4	64		4	$2(10)^3$P 100V	
Statistical mech./ matter models	1	1, 10^5	60	3D Graphics			
	2	100	64		10	10^4P	3,4,7
	9					10^4P 3V	
Plasma physics	1		60			600P 25V	2
	3		60			10^4P 2V	
	4	300	60	Imp. sci. routines		10^4P 6V	
	5	500	64		1–2		
Nuclear engineering	1		64		100	$2(10)^5$P 30V	
	2	100–200	48	Lower cost		10^6P 10V	5
Astronomy	5	100				10^5P 10^3V	1,4,7
	6	500	32				
	8		64		1–4	10^7P	
Miscellaneous (Data management)	2	10^4–10^6		Bandwidth=10^8		10^6P 20V	Ed1, Sw1, Ec1, Dm-1
Meteorology (weather)	1		64 or 128				
	2	25–100			8	10^5P 10V	5,9,10,11
	4					10^8P 10V	
	7			On line mass store		10^3P 30V	
	8	50–200	60				
Oil studies	1		64				
	3		32		1–5	10^7P 6V	
	4		32			10^4P 20V	
Computational fluid dynamics	2	10^3	64		10–20	10^6P 5V	1,5,7,10
	3	10–100	64			10^6P 9V	
	4	10^6	32			10V	
	9	64	64			10^6P 3V	
Military	3		14	Interactive		$3(10)^5$P	1
Biology	1						3,2,7,10,14,4,9
	5					10^5P 100V	
	6		60		2	10^6P 3V	
	8	1000	64		5–10		
	11		30			10^5P 2V	
Structures	5					$3(10)^3$P 3V	2,3
Business/ economics							2,3,4
Agriculture	2		8		100–1000/bit	10^7 P	
Machematics methods	2, 5	10^4	48			10^3P 6V	1,4,6,10,11,14,15
	8	50	64			10^3P 4V	
	9	100		Bit picking			
	12		36				
	13	10^3	72				

H=Hierarchy of computing

P=No of Pts
V=No of Var

56 responses with no machine requirements

APPENDIX K

ALGORITHM COMPONENTS – 37 SELECTED RESPONSES

APPENDIX K

Area	Response No.	Numerical soln. O.D.E. and P.D.E.	Matrix algebra			Monte Carlo	Extent of machine use
			Dense (full)	Sparse	Spec. struct.		
Quantum chemistry	1	Shroedinger eqn.		Yes			200 hr/yr (370, 1108, CDC 7600)
Solid state metallurgy	4	Schroedinger eqn.		Yes		Yes	1,000 hr/yr 7600
	6	Schroedinger eqn.	Yes	Yes			$100K/yr 195
	9	Schroedinger eqn.		Yes			$8K/yr 195
Particle physics	5	Integro DE	Yes				
	9	Coupled PDE	Yes	Yes	Yes	Yes	1500 hr/yr CDC 7600
	10	Fokker-Planck	Yes		Banded	Yes	200 hr/yr CDC 7600
Chemical reactions	3	2nd order ODE	Yes			Yes	1000 hr/yr 360/67
	4	2nd order ODE	Yes				1000 hr/yr 7600
	13						
Statistical mech./ matter models	5		Yes			Yes	70 hr/yr 7600
	6		Yes	Yes		Yes	
	8		Yes	Hashing		Yes	900 hr/yr 7600
Plasma physics	6	Yes	Large 10^6				200K/yr 360/91 and 7600
	7	Yes	Yes	Yes		Yes	Star 100
Nuclear engineering	3	2 or 3D	Yes				Star 100 100hr/wk
	4	Coupled linear		Yes			$20K/yr 195
Astronomy	2	Coupled PDE			Inversion and eig. anal. of lg. banded matrixes		40 hr/yr 7600
	3	Coupled PDE	Full				$50K/yr 370
Electrical power	1	Linear prog.		Yes	Triangular		$30K/yr 360/65 370/158
Meteorology (weather)	3	Navier-Stokes	Yes				300 hr/yr 360/65
	6	(irreg. domain)	Yes			Srch. algo—Yes	few $K 1108
Oil studies	2	PDE-finite el.	Yes	Yes	Banded		370/158-168
	5	Coupled PDE		Yes—S			500 hr/yr 370/168
Computational fluid dynamics	6	Coupled PDE					Cray 1 100 hr/yr
	11	Coupled PDE		No	Banded		1000 hr/yr CDC 7600
	8	PDE of time march		Banded			7600 100 hr/yr
Military	2	Coupled PDE		S		Inputs	$50K/yr 370/168
Biology	12	PDE					200 hr/yr 1108
	13		Ass. address and freq. tbl. look up	Yes			$70K/yr 168
	15	Coupled ODE					$70K/yr Cyber 175
Structures	1						360/91 450 hr/yr
	4	Finite elements	Yes	Yes			370/168 20 hr/yr
Business/ economics	1	Gauss-Seidel		Yes		Yes	B7700 used daily
Agriculture	1	Coupled PDE					370/158 2000 jobs/
Mathematics methods	3	Num. soln. Bayes Law	Yes			Yes	Star 100 equiv. 15hr/yr
	7	Fin. el. not. ‖		Dynamic			$700K/yr CDC 7600

S=Special mapping

APPENDIX L

SPECIAL FUNCTIONS DESIRED AS HARDWARE COMPONENTS

Transcendental (trig., exp., log., sqrt)	46
FFT	37
Complex arithmetic	31
Matrix operation-vector	14
Sparse matrix operations	2
Eigenvalue-eigenvector	1
Pseudo-random number generator	2
Integer-fractional part	1
Tridiagonal matrix operation	3
Quad precision	1
Poisson solver	1
Fix, float	2
Convolution	1
Polynomial evaluation	1
APL processor	1
Inner loop (hybrid)	1
Combination mult-add instruction	1
None needed	3

APPENDIX M

COMMENTS FROM QUESTIONNAIRE RESPONDENTS AND OTHERS

1. The steps required in solving a complex problem are: (1) problem, (2) formal solution, (3) algorithm, (4) program, and (5) computer architecture. Stability, convergence, and accuracy of the results depends critically on the match between the *formal solution* and the *computer architecture.*

2. It is essential that the problem solver have a fundamental understanding of the *computation processes* and the *physical problems* to begin with.

3. One must be careful that the pressure to design and build and use even bigger and faster computers does not stem from the myth, "Give me a bigger and faster computer and I'll give you the solutions you want." It is not true, even in computational fluid dynamics, that all the difficulties are due to the limitations of size and speed of currently available computer systems. In fact, the introduction of a super complex giant computer may help to perpetuate and multiply the larger number of fudge factors now in use to reproduce known results.

4. More discussion of the relative speeds, capacity, etc., of the several levels of hierarchical memory would be useful.

5. It is agreed that there is a substantial communication gap between users and designers of large-scale scientific computers. The Pajaro Dunes Workshop was a good first step toward initiating some dialogue, and this kind of activity should be continued. NASA could play a constructive role in stimulating dialogue in the development of commercial supercomputers, which could serve an ever greater need in the user community. Strong stimulus in the supercomputer area is needed if our country is to stay at the forefront.

6. Because of our staggering investment in software, compatability is a very important consideration. The second key point is economics. If the cost of computation drops, the number of computations will increase, enabling us to perform analyses presently considered impractical.

7. It could be important to get general agreement in this area (regarding special mathematical subroutines either in hardware or built-in software) because hardware implementation of some elementary functions has already been discussed by IBM.

8. More emphasis should be placed on this possibility (of hierarchies of computers, e.g., number cruncher, organizer and I/O) since its commercial feasibility has already been demonstrated by array processors appended to minicomputers.

9. Some discussion of exponent bit length in floating point number representation and use would be desirable.

10. *Any* characteristics cited (regarding a new language) would be useful.

11. I have used C.S.M.P. (continuous systems modeling program) extensively and find that programmers who are accustomed to straight FORTRAN, PL-1, COBOL, or other computational languages have some difficulty in adjusting to the parallelism of the language. CSMP treats the single processor IBM 370 *as though* it consisted of a number of independent integrators and differentiators, just as an analog computer would operate. The device of centralized integration with very short time steps permits computation with engineering precision.

The basic structure of the language statements, however, is such that it should be readily adaptable to any array of parallel processors. Yet it is also FORTRAN-compatible. I believe a footnote about the relationship between IBM's existing proprietary "parallel processing" languages (CSMP and PDEL for partial differential equation language) and the need expressed at the conference might be appropriate at this point.

12. The focus should be on improving the productivity of the *researcher,* not the machine.

13. Overall this report looks like something that the computer hardware companies should have done (or be doing) themselves as part of a marketing analysis. Why did not they undertake this survey? What was their attitude toward the workshop and this study? Are they interested in the results? (The answers are: we don't know, very cooperative, very interested in the results.)

14. The statistics gathered seem to suggest the need for improved computational machinery/hardware, but the important questions of how, when, and where have not been addressed.

15. Since I replied to the questionnaire, this university acquired and installed an IBM 3033. It has now been in operation for 9 months. It took only 6 months for the user community to saturate this machine. We will probably have to live with it for 10 years. The saturation was accomplished, not by administrative use nor by student's small jobs, but by a modest number of major users. There is evidently a large reservoir of unmet demand. When I reflect on our experiences with the new machines, I first appreciate that there has been a really major improvement in regard to what I would like to accomplish, and at this point, do not see much immediate hope of further improvement.

The decreasing rate of progress in this regard (development of more powerful machines for scientific computation) suggests that . . . the real question is whether anyone is willing to put up the money necessary to enhance the rate of progress. The requirements of business and of scientists have diverged. There is much more money to be made by selling machines to banks, etc., than to scientists.

16. It is reasonable to expect that if someone builds a "numerical wind tunnel," others will find additional uses for such a machine. This is the beauty of mathematics – there are few problems so unique that a powerful tool will not find wide application even if it is rather narrowly conceived at first.

17. Most researchers in these areas (of biology and biological chemistry) find that use of computing is limited largely by practical concerns such as funding for machine time or for personnel, or by inadequacies of parameterization for the models used, or by difficulties of interpretation. In this last case, you will have noticed how frequently responses to your questionnaire included comments about the need for graphics; and this confirms my personal impression that machine

speeds have already exceeded the capacity of many users to effectively digest their results in the form presently supplied.

18. I have read your report with great interest, trying to relate the various requirements with known development efforts in the large processor field. The latter do not seem to match many areas of concern in your project.

19. I agree with all the conclusions. Should your group consolidate the stated specific needs and publish them for the benefit of the manufacturing community, then perhaps the main-frame vendors R & D groups might respond with ideas of a phased building block approach to establish and promulgate continued technological upgrading of some initial state-of-the-art baseline system. To get a number of vendors involved in the same baseline system under an A-109 type of procurement approach, with all competitors interfacing their own best state of the art with one another, would certainly be an interesting project.

20. The current . . . codes use all of the resources of a 7600, STAR-100, or a CRAY-1 for up to 20 hours per run. The only hope of improvement lies in more computing power. A machine with 100 times the power of a 7600 is probably just barely adequate. But more computing power brings more problems. How does one keep track of what the code is doing and how does one analyze the results?

21. After a 5-year period of development we had a good machine to sell. But a 2–3 year window for sales is the best available! It was a monumental task although there was no specific back breaker. The cost was about evenly divided between hardware and software. The next step is to shoot for a machine 50 to 100 times the 7600. It would require a major investment without government funding.

It probably makes more sense to design an extended family of computers from the lowest up to 32 times that. Should accept a normal job stream and an extended FORTRAN should be O.K. The range could start at the VAX-level, which has found a good market, and build from there.

22. Machine development should be driven by a societal need.

23. The users have been bent around particular computers which has influenced the problem statement, even though FORTRAN is used, specific FORTRAN programs are structured around the machines architecture and are hence not transferrable. Need technology independent programming.

24. Users group is needed to couple applications, algorithms, and architecture.

How about a publication. Call it "MEGAFLOPS." Recognize that each contribution on a particular point is likely to render obsolete the previous such. Let people talk about their problems, complaints, suggestions, and wish-list.

Pattern the approach after what the users of small computers are doing.

25. Please be sure to include a statement of need for large-scale string processors.

1. Report No.
NASA TM-78613

2. Government Accession No.

3. Recipient's Catalog No.

4. Title and Subtitle
AN ASSESSMENT OF FUTURE COMPUTER SYSTEM NEEDS FOR LARGE-SCALE COMPUTATION

5. Report Date
February 1980

6. Performing Organization Code

7. Author(s)
Peter Lykos and John White

8. Performing Organization Report No.
A-7929

9. Performing Organization Name and Address
Illinois Institute of Technology
Chicago, Ill. 60616 and
Ames Research Center
Moffett Field, Calif. 94035

10. Work Unit No.
791-40-14

11. Contract or Grant No.

12. Sponsoring Agency Name and Address
National Aeronautics and Space Administration
Washington, D.C. 20546

13. Type of Report and Period Covered
Technical Memorandum

14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract

The report summarizes data, collected from many users, that range from specific computer capability requirements to opinions about the desirability of a national computer facility. Clearly there is a substantial communication gap between the users and the designers of large-scale scientific computers. As a consequence, vendor unwillingness to fully exploit computer technology because of financial risks causes an opportunity for both users (and therefore their sponsors) and manufacturer to be lost. The users' request for "bigger and faster" is too simplistic a statement of need. Users must be more explicit in defining their computational needs and must take into account available technology when designing their formal problem solutions. Considerable attention should be given to improving the user-machine interface; otherwise, increased computer power may not improve the overall effectiveness of the machine user. Significant improvement in throughput will require highly concurrent systems plus the willingness of the user community to develop problem solutions for that kind of architecture (i.e., give up FORTRAN). An unanticipated result was the expression of need for an on-going cross-disciplinary users group/forum in order to share experiences and to more effectively communicate needs to the manufacturers.

17. Key Words (Suggested by Author(s))
Large-scale computers, Scientific computing, Resource sharing, Computer system design, Human engineering of computers, Computer networking

18. Distribution Statement
Unlimited

STAR Category – 59

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price*
Unclassified	Unclassified	57	$5.25

*For sale by the National Technical Information Service, Springfield, Virginia 22161

NASA-Langley, 1980

National Aeronautics and
Space Administration

Washington, D.C.
20546

Official Business
Penalty for Private Use, $300

THIRD-CLASS BULK RATE

Postage and Fees Paid
National Aeronautics and
Space Administration
NASA-451

NASA

POSTMASTER: If Undeliverable (Section 158
Postal Manual) Do Not Return