SIGART NEWSLETTER Number 42 October 1973

SIGART NEWSLETTER Number 42 October 1973 1

CONTENTS 1a

SIGART NEWSLETTER 1b

The SIGART Newsletter is a bimonthly publication of the Special Interest Group on Artificial Intelligence of the Association for Computing Machinery. The Newsletter reports on projects being conducted by the artificial intelligence research community and generally reviews current progress in the

state-of-the-art. Correspondents report news from local SIGART Chapters and other AI Centers.

1b1

SIGART CHAIRMAN: Bob Balzer 1b1a

USC Information Science Institute
4676 Admiralty Way
Marina del Rey, California 90291
Telephone: 213-822-1511 1b1a1

NEWSLETTER EDITOR: Steve Coles ASSOCIATE EDITOR: Rich Fikes 1b1b

Artificial Intelligence Center
Stanford Research Institute
Menlo Park, California 94025
Telephone: 415-326-6200 ext. 4601, 4620 1b1b1

ASSISTANT EDITOR FOR ONLINE OPERATIONS: Kirk Kelley 1b1c

Augmentation Research Center
Stanford Research Institute
Menlo Park, California 94025
Telephone: 415-326-6200 ext. 3630 1b1c1

The Editors encourage contributions from authors, including Letters to the Editor (AI Forum), Technical Contributions (1 to 6 pages), Abstracts (preferably 100-200 words), Book Reviews, Bibliographies of Special Topics in AI, News Items (Conferences, Meetings, Course Announcements, personals, etc.), Advertisements (New Products or Classified Advertising), Puzzles, Poems, Cartoons, etc. Material may be reproduced from the Newsletter for non-commercial purposes with credit to the author and SIGART. 1b2

Anyone interested in acting as editor for a special issue of the Newsletter devoted to a particular topic in AI is invited to contact the Editor. Letters to the Editor will be considered as submitted for publication unless they contain a request to the contrary. Technical papers appearing in this issue are unrefereed working papers, and opinions expressed in contributions are to be construed as those of the individual author rather than the official position of SIGART, the ACM, or any organization with which the writer may be affiliated. 1b3

You are invited to join and participate actively. SIGART membership is open to members of the ACM upon payment of dues of $3.00 per year and to non-ACM members upon payment of dues of $5.00 per year. To indicate a change of address or if you wish to become a member of SIGART, please complete the form on the bottom of the last page of this issue. 1b4

Copy deadline for the December Issue: November 26th. 1b5

CHAIRMAN'S MESSAGE 2

In this issue Steve and I have included a questionnaire (pp. 5-6) designed to supply information needed to provide guidelines in the expansion of our current activities. It focuses on the role of conferences and on the Newsletter. There will be many opportunities for us to sponsor conferences (or sessions at conferences) in the next few years. As the field continues to grow and proliferate, the need for such actvities to keep ourselves up-to-date becomes more important. The information you provide in the questionnaire will help determine how we use these opportunities. Accompanying these opportunities is the responsibility of organization and participation. I hope that our future activities will not be limited by a lack of participation. Included in the questionnaire is a section eliciting such partcipation. 2a

The second major issue in the questionnaire is the role of the Newsletter. It was conceived of and has continued to function, highly successfully, as a quick response centralized source of activities in the field, such as the activities of various centers, summaries of recent conferences, abstracts of current papers, etc. Because of this success, suggestions have been raised to include discussions of controversial topics and the entire texts of refereed papers (e.g., see ACM President Anthony Ralston's recent editorial in the COMMUNICATIONS, p.459, August 1973). We are limited by our budget from materially expanding the size of the Newsletter without increasing our revenue. This might be done either by raising dues or by attaching a charge to an optional supplment. These and related choices are covered in the questionnaire. 2b

The questionnaire represents a major attempt to assemble the information required to plan the activities and functions of SIGART for the next few years. I urge each of you to aid us by taking a few minutes of your time right now to fill it in, tear it out, and return it promptly. 2c

R.M.B. 9/25/73 2c1

QUESTIONNAIRE 3

NAME ---------- ADDRESS ----------

---------- ORGANIZATION ---------- 3a

(Please Circle Response) 3a1

1. CONFERENCES 3b

Would you like to have national SIGART-sponsored conferences? Yes No 3b1

A. Type of sessions desired (indicate preference by percent)-- 3b2

(1) Technical Papers . . . . . %
(2) Tutorials. . . . . %
(3) Panel Discussions . . . . . %
(4) Working Sessions . . . . . %
(Invited participants only) 3b2a

B. Scope (Circle one or order preferences)-- 3b3

(1) Separate conference (e.g., as a supplement to IJCAI-75).
(2) As part of National ACM.
(3) As part of NCC.
(4) Joint Conference with another SIG (e.g., SIGPLAN). 3b3a

C. Time of year (describe best time period as month or season)-- 3b4

D. Frequency (such as alternate years with IJCAI)-- 3b5

E. Location (specify city if desired)-- 3b6

(1) East Coast
(2) Mid-America (e.g., Chicago, St. Louis, Dallas)
(3) West Coast 3b6a

F. Would you be willing to help organize such conferences? Yes No 3b7

If yes, could you serve as 3b7a

(1) General Chairman

(2) Session Chairman
(3) Local Arrangements
(4) Referee
(5) Other (please specify) 3b7a1

2. NEWSLETTER 3c

A. Do you read the Newsletter? 3c1

(1) Just Browse
(2) Certain types of articles in detail (Circle which types)-- 3c1a

Chairman's Message	Progress Reports	Abstracts	
Editors' Entry	Chess		
Advertisements			
News Items	Conferences	Cartoons	3c1a1

(3) Cover-to-cover 3c1b

B. Is the On-line Newsletter important to you? Yes No 3c2

Have you used it? Yes No 3c2a

C. Is the SIGART Newsletter an appropriate forum for controversial issues (or should we restrict ourselves to publishing abstracts, etc.)? 3c3

Appropriate Inappropriate 3c3a

D. Would you like to see a more formal newsletter with refereed articles? (Because of budget restrictions we cannot substantially increase the size of the newsletter; inclusion of long articles [such as the full text of the Lighthill Report], whether refereed or not, is not possible under the present arrangement). Yes No 3c4

(1) If yes, would you accept a dues increase of $5.00 to finance the publication of refereed articles? Yes No 3c4a

(2) Would you accept an optional charge of $5.00 to finance the publication of a supplement containing refereed articles? Yes No 3c4b

(3) Would you be willing to serve as a reviewer? Yes No 3c4c

(4) Would you be willing to allow money to be used to commission a professional technical writer or pay an honorarium to one of our own members to write a definitive

survey article on some aspect of AI?
Yes No 3c4d

(5) As an added incentive, should money be used to offer a prize award for the best paper in any given year? Yes No 3c4e

3. COMMENTS 3d

At the same time you return this questionnaire, you may wish to enclose an additional sheet either to comment on recent developments in your own AI research, or for suggestions for improving the Newsletter, or anything else you feel is constructive and relevant to SIGART. 3d1

Please mail completed questionnaire to 3d2

Steve Coles, Editor
SIGART Newsletter
Artificial Intelligence Center
Stanford Research Institute
Menlo Park, California 94025 3d2a

EDITOR'S ENTRY 4

1. New Conventions for Footnote, Bibliographic, and Page References. 4a

To establish greater consistency between the on-line and hard-copy versions of the Newsletter, we are instituting a new format for footnote, bibliographic, and page references starting with this issue. This will help us avoid the considerable effort needed for page layout in the hard copy version as well as facilitate direct access to references by on-line users using the link command in NLS. 4a1

Instead of appearing at the bottom of their respective pages, footnotes will now be accumulated at the end of the articles to which they pertain. The format of a reference will be an asterisk followed by N, followed by an integer, all contained in angular brackets. For example, <*N3> designates the third note. Bibliographic citations will be similar except that the letter 'R' will be used instead of the letter 'N'. 4a2

References to page numbers in the same issue will now be replaced by references to NLS statement numbers, those funny number/letter combinations you've been seeing in the right hand margins. The format of a statement reference will be 's.(<integer> <letter>)'. Thus, "see s. (5h)" will refer to statement 5h. 4a3

2. IJCAI-73 4b

The recently-held International Joint Conference on Artificial Intelligence at Stanford University was our most successful conference ever. Attendance numbered over 730 registrants <*N1> and a small profit was accumulated. It would be impossible to summarize in a few paragraphs the enormous number of technical activities and events that took place, so I won't even try. I believe it would be fair, however, to say that for those who attended, the spirit of the conference was permeated by an infectious atmosphere of intellectual excitement. Perhaps the greatest frustration at the conference was in deciding at each point which session to attend while interesting papers were being presented simultaneously in three parallel sessions. 4b1

Copies of the formal proceedings (Advanced Papers of the Conference) are available postpaid for $15.00 (a superb value) by writing to... 4b2

Stanford Research Institute
Publications Department
333 Ravenswood Avenue
Menlo Park, California 94025 4b2a

For the benefit of those who could not attend, a listing of papers scheduled for presentation at the Free Sessions with author names and addresses is given at s. (5c) of this issue. Copies of these papers may be obtained by writing to the authors directly. Summaries of the Special Sessions, SS1-SS4, appear starting at s. (5b) of this issue. A number of video tapes were also made during the Conference of special events such as the invited Tutorial Lectures, the Panel Discussions, and the Computers and Thought Lecture delivered by Patrick Winston of MIT. These tapes are now available from Stanford University (See s. (5a)). Depending on demand, these tapes may be made available in the near future in the form of 16mm standard movie films. Please write to Nils Nilsson of SRI if you have interest in the films (it appears that the cost of a one-hour lecture on film might be $125 and rentals would be correspondingly less). 4b3

Through the efforts of Warren Teitelman and Phil Jackson of Xerox Parc, several AI computer programs (obtained over the ARPA Network) were demonstrated live during the Conference near the main registration desk on a TI terminal obtained for that purpose. These programs included PIVOT (Deutch), SCHOLAR (Carbonell and Warnock), CHESS (Greenblatt), DOCTOR

(Weizenbaum), INTERLISP (Teitelman), and ENGROB (Coles and Robinson). 4b4

A special word of thanks must go the International Conference Committee who worked so hard to make IJCAI-73 a successful meeting. Deserving of special praise for a job well done are Nils Nilsson, who worked long hours as Program Chairman, and Les Earnest, who successfully unraveled the labyrinth of local arrangements. 4b5

3. IJCAI-75 4c

As was decided by the Executive Committee, the next IJCAI will be held in Leningrad, U.S.S.R. The General Chairman will be Dr. Erik Sandewall of Uppsula University, Sweden; the Program Chairman will be Prof. Patrick Winston of MIT. 4c1

4. Report on the BBC TV Lighthill Debate 4d

As we speculated in the preceding issue (Item 4 of the Editor's Entry), a copy of the color video tape (transcribed to meet American TV standards) of the controversial debate "The General Purpose Robot is a Mirage" was obtained from BBC television and had five showings at the IJCAI introduced by Professor Donald Michie of the University of Edinburgh. 4d1

To review briefly, the 2 1/2 hour debate was filmed in London at the Royal Institution on July 4th, and an 80-minute edited version was aired in Britain on BBC-2 on August 30th. The debate, moderated by Sir George Porter (a Nobel laureate and Director of the Royal Institution), was a part of the BBC's excellent "Controversy" Series. After intoducing the panelists, consisting of Professors Donald Michie of the University of Edinburgh, John McCarthy of Stanford University, and Richard Gregory of Bristol University, Sir Goerge began with a short review of AI while Shakey, the SRI robot, was flashed on the screen. 4d2

After this, Sir James Lighthill was introduced. In sharp contrast with the other panelists, he emerged from behind large wood-panelled doors to the sound of considerable applause from the audience. I don't really know whether this was an obligatory gesture by the British audience out of respect for Sir James's title, but I can say that it was not well received by the American audience watching the tape at IJCAI. His animated presentation, which lasted some 15 minutes, defined AI as a subject attempting to bridge studies of the brain on the one hand with industrial automation on the other, and generally followed the argument he espoused in his "Lighthill Report."

<*N2> He urged continuing support for the legitimate areas of brain investigation as well as studies of industrial robots, but predictably condemned the "bridge" part of AI as doomed to failure. His main line of attack centered on the "combinatorial explosion," which he claimed would permanently render the goals of AI unattainable. In subsequent discussions he conceded that AI goals might be achievable in principle, but ultimately he dismissed them, with a characteristic wave of his hand, as being beyond the range of our lifetimes. 4d3

Prof. Michie in reply showed a film of the Edinburgh Versatile Assembly Program building a toy car and defended the whole purview of AI as an independent field of meritorious scientific investigation. Both he and John McCarthy politely suggested that Sir James was not very well informed. Prof. McCarthy then went on to list a series of major scientific and engineering accomplishments by the field of AI. In my judgement he was at his best when he said, "I certainly do know why this field is called 'Artificial Intelligence'. It was I who invented that phrase back in the mid '50s..." Later Prof. Michie challenged Sir James to double the stakes in an existing wager <*N3> if he so genuinely believed in the force of his argument, but he graciously declined. Later still Prof. Gregory, whose specialty is the physiology of perception, stressed the substantial benefits of AI research and generally supported the position set forth by Profs. Michie and McCarthy. 4d4

In contrast to earlier programs in the Controvery series, the audience present for this debate was an unusually prestigious group, including such distinguished professors as Needham of Cambridge, Thring of Queen Mary College, London, Strachey of Oxford, and Lederberg of Stanford (a Nobel laureate who just happened to be passing through). Many of the audience contributions which stressed the importance of intelligent robots for industry and the value of AI research in general were edited out of the final tape. 4d5

In conclusion, many observers, including Prof. Michie, felt that Lighthill had lost the battle, and so did I. But I am less certain how the general British TV audience may feel about the arguments presented. They may respond more to sweeping generalities than to the more technical and sometimes academic details brought up during the discussion. Also, when the format requires that three persons are independently charged with responsibility to refute the position of one other person, and each has in principle only one-third the time of the first, and none of them is given explicit authority to summarize and to speak on behalf of the other two, it is difficult for the three to do more than separately chip away at local weaknesses,

thereby diluting the force of their rebuttal. Therefore, the audience never has a single individual with whom they could identify in assessing a global and systematic refutation of the original argument. There may always be the lingering doubt that the first man may really have had something worthwhile to say, and he just didn't do a good job at presenting it. Perhaps a more important question is how the Science Research Council (British AI's principal supporting agency) will view these arguments when they allocate resources in the future. This remains to be seen. 4d6

5. Natural-Language Memory Structures Studied In Seminar Series At Stanford AI Project 4e

Professor Ken Colby of the Stanford University AI Project has organized a series of seminars expected to meet approximately every two weeks during the coming year to examine and contrast different approaches to natural-language memory structures. Regular attendees have been Horace Enea, David Smith, Terry Winograd, Chuck Reiger, Gordon Bower, Yorick Wilks (all from Stanford), Lotfi Zadeh (Berkeley), Steve Coles, Sharon Baranofsky (SRI), Dan Bobrow, Sharon Kaufman-Diamond (Xerox), as well as a half dozen Stanford graduate students interested in computational linguistics. 4e1

Thus far Chuck Reiger has presented his work on MARGIE (a graph-oriented system in collaboration with Roger Shank) and Horace Enea has presented is work on PARRY (a production-oriented system in collaboration with Ken Colby). At the next meeting Steve Coles will discuss the ENGROB System (predicate calculus oriented). Since there are many scientific persuasions represented among the participants (ranging from dependency graphs to procedural representations) as well as motivations (ranging from psychiatry to speech understanding to children's stories), the seminar should prove to be very interesting. 4e2

6. Change in Newsletter Reporters 4f

Rob Kling has recently moved to the University of California at Irvine and will thus become their reporter. To replace him at the University of Wisconsin will be Norman Sondheimer, a graduate student in the Computer Sciences Department. 4f1

7. New Policy for On-Line Newsletters 4g

We have archived the October 72 Newsletter in the on-line SIGART Directory at SRI-ARC to conserve disk space and subsequently will maintain on disk only one year's worth of

Newsletters. Procedures for retrieving older On-Line Newsletters are available in <SIGART, NEWS, ARCHIVED;w>. 4g1

We have recently inaugurated a statistics package to gather data on the frequency of use of the On-Line Newsletter. We expect to publish these figures perodically in future issues. 4g2

8. Two New Programs Available on the Network 4h

A. Stanford University's AI Project has recently made available an Associated Press wire service program <*N4> allowing one to access current news stories on-line, using boolean combinations of key words derived from a dictionary of 1200 words. For example, 4h1

```
@TELNET<CR>
#SU-AI<CR>
 LOGINNET/GUE<CR>
.RAPE<CR>
.(NIXON+TAPES)*WATERGATE-AGNEW
```
4h1a

will itemize stories in reverse chronological order, if there are any, that have to do with either Nixon or his tapes in relation to watergate, but which do not mention the vice president. Typing a number will show you the Nth latest stories on that list, while a <CR> will list all of them. 4h1b

```
.RHOT<CR>
```
4h1c

will list news stories as they are generated by AP. 4h1d

B. "#BBN21<CR>" will provide a random quotation by a famous person or well-known computer scientist. For example, 4h2

"I've been rich and I've been poor. Rich is better."
-- Mae West 4h2a

If anyone has or knows of other programs of general interest, or of particular interest to the AI community, please let us know. 4h3

9. Slagle's Book Published in German 4i

Jim Slagle's book, ARTIFICIAL INTELLIGENCE: THE HEURISTIC PROGRAMMING APPROACH, is being published by Verlag Moderne Industrie, Munich. 4i1

10. Error in Header 4j

We regret the systematic error in the date of the header portion of each page in the last issue. However, the outside cover correctly reads August 1973. 4j1

L.S.C. 10/1/73 4k

NOTES 41

<N1> For reference, about 600 attended the first IJCAI in Washington in 1969, while about 400 participated in the second London IJCAI in 1971. 411

<N2> Public Relations unit; Science Research Council; State House; High Holborn; London WCIR 4TA, England. 412

<N3> For details of the bet, regarding computer chess, see the SIGART Newsletter, No. 36, October 1972, p.26, Item 1. 413

<N4> For more details see "Reading the Assocated Press News" by Martin Frost, Stanford Artificial Intelligence Laboratory, Operating Note 72, July 23, 1973. 414

FOLLOW UP ON IJCAI-73 5

A. VIDEO TAPES OF TUTORIAL LECTURES AVAILABLE THROUGH STANFORD UNIVERSITY 5a

The Stanford Instructional Television group is offering for sale IJCAI-73 black-and-white video tapes in various formats. The tapes are of lectures recorded live at IJCAI-73. Certain of the 1/2" tapes are also available on a rental basis. All tapes are in NTSC (525 scan line) format. A price list is given below (rental prices are given in parentheses for a two-week rental period). 5a1

The attached purchase and rental forms may be used in ordering the tapes. All orders should be accompanied by full payment. Tapes will be sent air mail. 5a2

Besides the regular technical sessions, listed in the last issue, IJCAI-73 had two other types of sessions: Special and Free Sessions. 5a3

B. SUMMARY OF SPECIAL SESSIONS AT IJCAI-73 5b

Four special sessions were organized by individuals who wished to explore certain specialized topics in an informal setting with the help of invited panelists and any interested conference attendees. 5b1

SS1
Social Implications: A Look at Some Immediate Issues
Chairman: Andee Rubin, USC Information Sciences Institute and MIT 5b2

The discussion, based on a carefully prepared three-page list of questions, focused on issues in the area of social implications of AI which are relevant in the immediate future--specifically short range applications, funding agencies, and in particular the relationships between them. In other words: Who is funding current artificial intelligence research, and what are their reasons? 5b2a

Steve Crocker of ARPA, Val Tareski and John Pasta of NSF, and Tom Wachowski of AFOSR were present, and the first part of the session consisted mainly of the audience asking them questins about the amount of money available for such research, the criteria by which it is distributed, the future outlook, and differences among the agencies in such policies. 5b2b

After a while, when the audience felt satisfied that they were aware of the basic issues and facts in that discussion--and were becoming somewhat impatient with mere fact-finding--talk turned to the reasons different agencies proposed for funding AI and what uses they envisioned for specific projects, such as speech-understanding systems, sophisticated vision systems, or automatic programming systems. In particular, participants were interested in what effect they might have on the decision-making process if they weren't satisfied with the directions of certain agencies. 5b2c

The discussion eventually moved to an even more abstract level dealing with the differences between basic and applied research, with issues concerning the possible good and bad applications of pure research and their relevance in deciding whether to follow a certain line of work. Quite a number of people expressed interest in remaining involved in at least talking to other AI researchers about these issues; SICCAS (Coomputers and Society), SESPA (Science for the people) and CPP (Computer People for Peace) were mentioned as possible organizations through which to work. I also collected a list of interested people during the following two days. Anyone who did not sign that list and is interested in being included in a list which will be distributed to everyone on it and which may serve as an informal communication mechanism, please send your name and address to: 5b2d

Ms. Andee Rubin
AI LAB
545 Technology Square
Cambridge, Mass. 02139 5b2d1

or use the ARPA-NET to ANDEE@MIT-AI. At some reasonable interval after this Newsletter has been distributed, I will compile that list and send it out. The attendees' general reaction to the session was favorable, due mainly to the high degree of audience participation and the presence of some of the people who really make the decisions. The discussion was lively and often heated, but its primary virtue was that people felt that they could ask pointed, specific questions and get some kind of answer. In addition, though the discussion finally did move to those more abstract issues of pure/applied research with which many such discussions start out, it did so from the "other direction," so to speak...that is, from the more specific issues of money and priorities. 5b2e

SS2
Formalisms for Artificial Intelligence
Chairman: Carl Hewitt, MIT A.I. Laboratory 5b3

Formalisms for artificial intelligence were discussed by four speakers at this session. Allen Newell began by discussing both the use of production languages for representing problem solving processes and the structure of his MERLIN knowledge net. Richard Weyhracuch followed with a discussion of some work that is based in the predicate calculus and considers the generation of a style of informal proofs that people use and find convincing. Carl Hewitt elaborated on his ACTORS talk of the previous day; his comments stressed the modularity of ACTOR based systems and the elegance obtained from the simplicity of their structure. Finally, Alan Kay gave a brief presentation describing his SMALL TALK language which is similar in spirit to the LOGO language and is implemented using ACTORS. 5b3a

SS3
A.I.: A Discussion of its Impact on Science, Technology, and Society
Chairman: Robert H. Anderson
USC Information Sciences Institute and RAND 5b4

This special session was formed because of the relative paucity of papers being presented at IJCAI-73 dealing with the relationship between A.I. and real-world problems. It was held during the last morning of IJCAI so that the

state-of-the-art in A.I., as defined in the presented papers, could form a background for the discussion. 5b4a

The main question addressed was: Why hasn't more attention been paid, to date, to real world applications of AI? Potential answers discussed were: 5b4b

It would be a diversion from the main research goals and interests of AI. 5b4b1

(The consensus was that this was not a significant factor.) 5b4b1a

We don't have easy access to people who have a deep understanding of relevant problem areas. 5b4b2

(A consensus also thought this was not a significant factor.) 5b4b2a

The power of AI is in the knowledge of the application area which is applied, but we don't yet know how to formalize a significant body of knowledge. 5b4b3

(This was felt to be an important consideration.) 5b4b3a

An important part of the panel discussion was a presentation by Ed Feigenbaum on 'The Nature of an Application of Heuristic Programming Techniques'. A partial list of considerations he raised as being relevant to particular applications in 1973 is as follows: 5b4c

Problem Formulation: Does there exist a generally agreed-upon best way to represent the problem domain, and a set of primitive solution elements out of which solutions will be discovered by combination and search? Can the generation of solutions be conceived as a combinatorial process? 5b4c1

Knowledge Base: Does a model exist which can provide the semantics behind the symbol manipulation? Is there an expert who is highly knowledgable, motivated, computer-oriented, through whom knowledge can be acquired? Is the universe of facts small, and almost uncoupled from other possible universes? 5b4c2

Problem Difficulty: Is the application too difficult or too simple? Can measurements be made to measure the progress of program solutions? Can progress be made

incrementally? Is it ill-structured enough to require heuristics? 5b4c3

Resources (People, Motivation, Money, Computers): Can they be obtained and sustained over long periods of time? 5b4c4

Members of the panel discussing these issues were: 5b4d

Prof. Saul Amarel, Rutgers University
Dr. Dan Bobrow, Xerox Parc
Mr. Stephen Crocker, ARPA-IPT
Prof. Ed Feigenbaum, Stanford University 5b4d1

SS4
Automatic Programming
Chairman: Robert Balzer, USC Information Science Institute 5b5

The session brought together participants from several different ongoing projects. The current activities of each were described and an attempt was made to identify the problems which must be faced in the future. 5b5a

As a way of establishing a common ground for discussion, I used the viewpoint established in "A Global View of Automatic Programming," [3IJCAI, p. 494]. In this view, there are four phases: (1) Problem Acquisition, in which the knowledge of a domain is acquired and structured for use in the problem solution; (2) Process Transformation, in which the domain knowledge is used to transform the problem statement into a detailed procedural form; (3) Model Verification, in which the structure and behavior of this procedural form is examined; and (4) Automatic Coding, in which the procedural form is converted into an efficient program. Using this structure we get the following breakdown (only some of these were represented at the special session and are summarized below): 5b5b

First 3 areas: ISI, Rutgers
Last 3 areas: IBM, MIT
Middle 2 areas: Winograd 5b5b1

Process Transformation:	Program Synthesis, CMU, Waldinger-Green
Model Verification:	Program Verification
Automatic Coding:	Earley, Darlington, Harvard 5b5b2

MIT (Bill Martin): Focus on building high performance knowledgeable application systems in the domain of management information systems. This domain has been

abstracted and described in relational terms. A method has been devised for translating from this relational level into PL/1 programs to perform the desired tasks. As part of this translation process, run-time cost figures for these generated programs are developed and used for either an interactive or heuristic search for an organization that minimizes these costs. A very large relational model (for inventory control) is being constructed which will be used to capture the user's problem definition by instantiation through a dialog jointly directed by the user and the system. 5b5c

IBM/Yorktown (Martin Mikelsons): This work is focused on the domain of business data processing and has two main parts. The first is a document flow language which describes an application in terms of the definition of, the movement of, and processing performed on a set of documents. A translator is being built to convert such descriptions into efficient running programs. The second component is attempting to construct these descriptions from a high level discourse by questionning the user. 5b5d

ISI (Bob Balzer): This project is investigating and attempting to define mechanisms necessary for dynamic acquisition of semantic models of a domain and for effectively using them to transform loose program specifications into a more precise form capable of being evaluated. A broad range of domains is being examined to determine requirements for this domain-independent approach. The focus in acquisition is on techniques for spotting inconsistencies or incompletenesses in a semantic model and directing dialog with the user toward correcting these problems. In transformation, a precise executable model for an arbitrary domain is being defined and techniques developed for transforming loosely specified, ill-defined, problems into precise form. These techniques include: filling in missing relationships between the objects of the domain, using type constraints to disambiguate statements, using information-retrieval techniques to determine the objects to be associated with an operation, determining how to handle override conditions and constraints, and spotting 'extra' information in an invocation statement which is not 'used up' and determining how it should modify the interpretation of invoked actions or subactions. 5b5e

CMU (Jack Buchanan) and Stanford (Dave Luckham): This project is concerned with verification and synthesis via a formal mechanism based on Hoare's logic system for describing programs. The programs to be written are

specified by formal input-output descriptions. Conventional problem-solving techniques are used to produce the program. The problem of representation and acquisition of the necessary programming methods has been a major research focus. Hoare's formalism for program semantics has been found suitable for describing programming methods. A new effort is underway to allow the user to present these methods informally and have the system translate a description into the formalism. The program synthesizer now running can write several-page programs and has some general programming knowledge built in. The programs written automatically have been robot control programs and arithmetic algorithms. Two difficult classes of program structures are handled: contingency programs (necessary for planning with unreliable operators) and complex loop structures. A more interactive problem-acquisition system using essentially "constructive" programming methods is under development. 5b5f

SRI (Waldinger) and Stanford (Green): This effort is concerned with program writing through use of "expert" built-in knowledge in the domain of list-processing, ranging from the fundamentals up to pattern matching, tree search, graph search, graph-matching, etc. Most of the knowledge is relatively "pure" programming knowledge with such domain-dependent knowledge as is necessary. The emphasis is on codification of the considerable body of list-processing programming knowledge. The target system is expected to have a deep understanding of programming as shown by its program-writing ability, by its line of reasoning in creating a program, and by its own discussion of why it made each choice and the factors involved. An interesting feature of this research effort is some emphasis on 'human' methods of program specification, such as example input-output pairs, generic examples, and annotated traces. Prototype systems have been developed that can write short (less than 7 line) programs, including sort, merge, reverse, flatten (a list), etc. The programs have been specified by several alternate methods. The best program written automatically to date is a square-root program, in which the system shows a good line of reasoning in producing the program. The tentative one-year goal of the project is the automatic synthesis of a six-page concept formation program that employs simple list-processing techniques. 5b5g

Stanford (Winograd): This effort is concerned with defining and ultimately building a smart assistant for programming complex problems. It might be considered an advanced bookkeeping system. The help it provides would be based on

built-in knowledge of programming, of functions available in the system, and extra comment-like information provided by the user about his program and data structures. From such knowledge, the system would be expected to perform error checking, be able to answer questions about the static or dynamic structure of the program interactively, and automatically select the appropriate function or method for simple operations. 5b5h

C. LISTING OF FREE SESSIONS AT IJCAI-73 5c

Session 9
PROBLEM SOLVING AND PSYCHOLOGY
Tuesday, August 21, 1973 9:00AM - 12:00 NOON 5c1

Mr. Gregory Gibbons
Department of Mathematics
Naval Postgraduate School
Monterey, California 93940
"The Q-Size Rule: A Method for Eliminating Redundancy in Heuristic Search" 5c1a

Professor Marco Somalvico and A. Vincentelli
Instituto di Elettrotechnica ed Elettronica
20133 Milano
Plazza Leonardo DaVinci, 32, Milan, Italy
"Theoretical Foundations of State-Space Approach to Problem-Solving" 5c1b

Dr. James M. Perry
Computer Science Group
The University of Connecticut
Storrs, Connecticut 06268
"Abstract Problems: A Formal Development for Automatic Problem Generation and Solution" 5c1c

Mr. Larry Rosen
Department of Psychology
University of California at San Diego
La Jolla, California 92037
"The Role of Similarity Relations in the Multi-Alternative Choice Process" 5c1d

Dr. George F. Luger and G. A. Goldin
Graduate School of Education
University of Pennsylvania
Philadelphia, Pennsylvania 19147
"The Use of Artificial Intelligence Techniques for the Study of Problem-Solving Behavior" 5c1e

Mr. S. Romani and A. Newell
Department of Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213
"On Generating Problems" 5c1f

Session 12
FORMALISMS AND AUTOMATIC PROGRAMMING
Tuesday, August 21, 1973 1:30 PM - 5:00 PM 5c2

Mr. Robert Kling
The University of Wisconsin
Computer Science Department
1210 West Dayton Street
Madison, Wisconsin 93706
"Fuzzy Planner" 5c2a

Earl Sacerdoti and Rene Reboh
Artificial Intelligence Center
Stanford Research Institute
Menlo Park, California 94025
"GLISP" 5c2b

R. O. Anderson
Churchill College
Cambridge CB3 ODS, England
"A New Logic of Analogy" 5c2c

Dr. Robert Balzer
Institute for Information Science of USC
4676 Admiralty Way, Suite 522
Marina Del Rey, California 90291
"Language Independent Programmer's Interface" 5c2d

Jack R. Buchanan
Computer Science Department
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213
"An Automatic Programming System" 5c2e

Dr. Richard Waldinger and Karl Levitt
Artificial Intelligence Center
Stanford Research Institute
Menlo Park, California 94025
"Reasoning About Programs" 5c2f

Dr. Peter Deutsch
Xerox Palo Alto Research Center
3810 Porter Drive

Palo Alto, California 94304
"Recent Results in Automatic Program Verificiation" 5c2g

Session 16
VISION
Wednesday, August 22, 1973 9:00 AM - 12 NOON 5c3

Dr. Jerome A. Feldman
Computer Science Department
Stanford University
Stanford, California 94305
"Gunnar Grape's Dazzling Model-Based Line Fitter" 5c3a

Dr. Walton A. Perkins
Computer Science Department
Stanford University
Stanford, California 94305
"A Corner Finder for Visual Feedback" 5c3b

Professor Azriel Rosenfeld
Computer Science Department
University of Maryland
College Park, Maryland 20742
"A Report on the U.S.-Japan Seminar on Picture and Scene Analysis (Kyoto, July 23-27, 1973) 5c3c

L. F. Pau
The Institute of Math, Statistics, and Operations Research
The Technical University of Denmark
Lyngby, Denmark
"Confusion in the Vision of Simple Geometric Objects" 5c3d

Craig Cook and Tom Binford
Artificial Intelligence Laboratory
Stanford University
Stanford, California 94305
"Aspects of Automatic Camera Focusing" 5c3e

Dr. Alan R. Johnston
Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, California 91103
"Proximity Sensing for Manipulation and a Laser Range Finder for a Robot" 5c3f

Session 19
APPLICATIONS
Wednesday, August 22, 1973 1:30 PM - 5:00 PM 5c4

Mr. James Doran
Science Research Council
Atlas Computer Laboratory
Chilton, Didcot
Berkshire 0x11 0QY, England
"Heuristic Search Applied to Problems of Archaeo-Logical Inference" 5c4a

Dr. Tim O'Shea
Department of Computer Science
The University of Texas at Austin
Austin, Texas 78712
"Some Experiments with an Adaptive Self-Improving Teaching System" 5c4b

Professor Billy Claybrook
Virginia Polytechnic Institute
and State University
Blacksburg, Virginia 24061
"Experiments with polyfact: A Learning Program that Factors Multivariable Polynomials" 5c4c

Dr. Saul Amarel
Department of Computer Science
Livingston College
Rutgers University
New Brunswick, New Jersey 08903
"AI Approaches to Diagnosis and Treatment of Glaucoma" 5c4d

Robert S. Engelmore, Research Associate
Computer Science Department
Stanford University
Stanford, California 94305
"Application of AI to Scientific Inference: Determining the Structure of Crystallized Proteins" 5c4e

Professor W. D. Maurer
Department of Electrical Engineering and Computer Sciences
University of California
Berkeley, California 94720
"Symmetric Floating Approximation and Program Correctness" 5c4f

Session 23
ROBOTS AND OTHER
Thursday, August 23, 1973 9:00 AM - 12 NOON 5c5

Dr. N. G. Zagoruiko
Institute for Mathematics

USSR Academy of Sciences
Siberian Division
Novosibirsk 90, U.S.S.R.
"Discovering of Empirical Regularities within the Frame of General Recognition Theory" 5c5a

Mrs. M. V. Aristova and Dr. M. B. Ignatiev
c/o Dr. S. I. Samoylenko
USSR Academy of Sciences
Moscow, U.S.S.R.
"The Concept of the Structure of Highest Levels Control by Robot-Manipulators" 5c5b

Dr. Laurent Siklossy and C. Dawson
The University of Texas
Austin, Texas 78712
"Automatic Generation of Hierarchies of Goals in Robot Worlds" 5c5c

John Birk and Donald Franklin
Department of Electrical Engineering
Kelley Hall
University of Rhode Island
Kingston, Rhode Island 02881
"Minimizing Robot Work Time for Pitching Objects" 5c5d

Professor Teuvo Kohonen
Department of Technical Physics
Helsinki University of Technology
SF-02150
Otaniemi, Finland
"New Analog Associative Memories" 5c5e

Dr. Alois Glanc
Department of Computer Science
Queens College
Flushing, New York 11367
"Design Considerations of the Golem System and Implied Problems in Robotologic" 5c5f

Session 26
NATURAL LANGUAGE
Thursday, August 23, 1973 1:30 PM - 5:00 PM 5c6

Dr. Sharon Kaufman-Diamond
Xerox Palo Alto Research Center
3406 Hillview Avenue
Palo Alto, California 94304
"On Story Understanding as a Task for AI" 5c6a

Dr. Nagib A. Badre
IBM Corporation
Thomas J. Watson Research Center
Yorktown Heights, New York 10598
"CLET--A Computer Program that Learns Arithmetic from an Elementary Textbook" 5c6b

Professor Adele A. Abrahamson
Department of Psychology
Rutgers College
New Brunswick, New Jersey 08903
"Deep Semantic Structures for Natural Language Processing" 5c6c

Dr. Perry L. Miller
20B-208
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
"Locally-Organized Parsing: For Spoken and Text Input" 5c6d

Mss. B. Nash-Webber and M. Bates
Bolt Beranek and Newman
50 Mouton Street
Cambridge, Massachusetts 02138
"Syntactic and Semantic Suport for a Speech Understanding System" 5c6e

NOTES ON THE SOCIAL IMPACTS OF ARTIFICIAL INTELLIGENCE by Rob Kling
Department of Information Sciences
University of California
Irvine, California 6

IJCAI-73 devoted three sessions to the aspects of AI funding, applications, and impact. In addition, one paper (Firschein, et al, 1973) described a Delphi study of plausible AI applications and their possible impacts. I'd like to share some of my reflections on this timely burst of attention devoted to these social concerns. 6a

Impacts of AI based Technologies 6b

The prediction, analysis, and assessment of technological impacts is a critical and frequently misunderstood area. It has been colored by a strong emphasis upon "philosophical" and social speculation (Chartrand, 1972; Martin and Notman, 1969) thought to be irrelevant by pragmatically inclined scientists. Often, simple pieties substitute for careful analysis. In addition, each of us holds some view about the impact of his work. We amateurs often seem no less accurate than people

writing, or speaking, under the rubric of "social impact." Thus we have a hard time believing that especially prescient studies of social impacts are likely or possible. 6b1

Certainly no one will say that social impacts should be ignored. Like affirmations of motherhood and apple pie, periodic attention to "social impacts" serves a ritual function, showing that we are sensitive, concerned, and responsible scientists, rather than crass opportunists, narrowly pursuing our own professional and personal interests. It is in our interest that we speak about "social impacts"; and that the likely impacts of our technologies be, on the whole, labelled benign. Thus, the analysis of social impacts is a necessary, but neither a serious nor dispassionate venture for the AI community. 6b2

A special session, chaired by Lou Fein, met to discuss the "social implications" of AI in a public forum. I won't attempt to summarize their discussion here, except to note that it emphasized AI as a form of theoretical psychology or a set of remote technological possibilities such as an automated courtroom judge. The likely products of AI research, that were identified in the Delphi study reported earlier in the conference (Firschein, et al, 1973), were neglected, and the societal contexts in which computing is now exploited were ignored. 6b3

I'd like to emphasize AI as a technology, since I believe that the mission-oriented AI sponsors such as ARPA, AFOSR, and NIH are as interested in new technologies as in new understanding of human cognitive processes and theories of organized complexity. We are quite self-serving when we emphasize the latter areas that have little short-run payoff and consistently ignore the technologies we develop as a by product of our less applied research.<*N1> (In the long run, the image of man as a cognitive information processor may be the deepest AI impact of all. C.f. Weizenbaum, 1972.) In the short run (10-20 years) the spin-off technologies supported by mission-oriented sponsors may be most critical. 6b4

To give an example, the Delphi group believed that the development of automatic diagnostic equipment or a personal biological model to aid patient monitoring in intensive care units would be of "high-potential significance." Currently we have special problems with medical care in this country. High quality care is easily available to the middle and upper classes, while many of the urban poor and people living in rural areas simply do not have good care available despite their pressing needs. Nevertheless, as a culture we place a

premium value on health. Any health related research carries an angelic halo and an aura of intrinsic good. Despite major attempts to provide more equitable access to medical care based on needs through such programs as community clnics, Medicare, etc., the social structure of medical access has changed little (Alford, 1972). (Alford's acute analysis portrays the political economy of health care as a complex social system in a stable configuration which is highly resistant to fundamental structural changes.) In terms of his analysis which poses a set of "independent" doctors (and the AMA) vying with a research-administrative-educational plus medical-industrial complex, the development of costly technologies may mostly aid the latter in their contest with the former. An increased quality of medical care may simply diffuse over time as a byproduct of such developments and simply trickle down to those most in need over a long period of time. In such a system, expensive technological aids are probably not the most cost-effective means to provide higher quality care to those most in need. 6b5

Whatever the accuracy of this model, it is based on a careful analysis of social change in the relevant societal sectors. An alternative analysis would have to be similarly rooted in the dynamics of change in the medical system in order to be credible. Unfortunately, such deeper analyses are lacking in all too many of our discourses on the social impacts of AI-based products (or X--for all too many X's). 6b6

One assumption implicit in the preceding observations is that the impacts of AI-based technologies are subject to the very same dynamics of social impact as are other computer-based technologies. As we AI researchers separate ourselves from the larger computing community <*N2>, we easily indulge in the fantasy that our impacts will be of a different order than related technologies. This is a very tenuous assumption. For example, to the extent that AI based artifacts are costly and require sophisticated environments for use, they will be used by organizations and groups with substantial power. To the extent that these technologies are "intrinsically" influence-enhancing (Kling, 1973), they will increase the gap bewtween the weak and the powerful, the rich and the poor. Consider federal uses of computing. Everyone exploits computers for routine uses such as budgetary analysis. Everyone automates his payroll. But many of the most advanced computing applicatons go on in DOD. Likewise, the first municipal departments to automate include finance, police, and to a leser extent, welfare. The military and the police are unique in that they hold the only legitimate authority to exercise physical force in this society. Unfortunately, they

have been noted as among the more repressive institutions in our society as well. I am not aware of military or police computing which renders them more humane organizations. I emphasize these concerns with mundane computing, since they help identify those who may be some of the major beneficiaries of any computing technology, especially expensive technologies such as AI. It is not that such organizations will necessarily exploit computing in singularly anti-social ways. Rather, computing aids may provide (differentially) greater power without rendering them more humane. 6b7

In a market economy, expensive sophisticated computing aids are an elite-enhancing technology. For example, the Mead General Corporation is developing an elegant system to retrieve legal precedents for inquiring lawyers. (In many ways it typifies the kind of large--public?--data bases upon which sophisticated question-answering systems are hoped to operate.) Such a technology could help equalize the quality of legal aid available to all by providing easy access to a large body of legal precedent to private lawyers who do not own or are remote from large law libraries. At present this system, which has been given a de facto monopoly by the state Bar Associatons, is sold on an ability-to-pay basis. Unfortunately, the cost is approximately $50K per year and can be afforded only by the larger law firms. In principle, the Mead system could be made generally available.<*N3> That would require a kind of underwriting similar to that which provided electricity to the rural population earlier this century. Without such underwriting, its high cost will prohibit widespread use. Insofar as it is a useful technology, it will simply provide better legal aids to those already well supplied.<*N4> 6b8

In the decades of the 1960's and the 1970's the major computer users are large bureaucratic organizations. Within particular organizations, computing aids automate routine work and provide information and skills that were previously prohibitively expensive or unavailable. Indirectly, computing effects organizational structure, effectiveness, and efficiency. Impacts on the larger society are mediated by whatever missions, policies, and styles of action computer-using organizations employ. While policies are influenced by the available technologies, technological possibilities do not propagate automatically. For example, the FBI does not update its arrest files in the National Criminal Information System to include the disposition of cases, even though it is technically easy to do so. The short-range impacts of AI-based technologies are intimately bound with the policies of the groups that can afford to use them. Any meaningful assessment of AI-based technologies must be situated in a context which

describes likely users and their policies. Such studies require careful, dedicated empirical work, not just part-time avocational efforts. (Unfortunately, only a trickle of support is available for "technology assessment" in contrast with "technology development"). 6b9

Funding Sources 6c

A special session at IJCAI-73 was devoted to such questions as: Who funds AI research, and to what extent? (See s. (5b1)) What do each of the funders expect in return for their support? 6c1

Many people were concerned that ARPA is the major single supporter of AI research and that ARPA would expect (direct?) (long-term) military payoffs from the research it funds. Spokesmen from ARPA, NSF, and AFOSR tried to place their research support in the context of their agency's mission. However, the expectations of the funding agency may be less central than we assumed in that session. For example, even if NSF were to fund all AI research under the rubric of basic science, military scientists and engineers at centers such as (or under contract to) the Rome Air Development Center, Wright-Patterson Air Force Base, and Fort Manmouth would attempt to exploit whatever technology they could to solve military problems (that range from inventory and scheduling to weapons systems, counter-measures, and intelligence). In fact, they would not be doing their jobs very well if they did not attempt to exploit the available technologies, including AI. 6c2

For example, while the projects that AFOSR funds indicate technologies that are of special interest to the Air Force, the concern that many of us share of not aiding the development of ever more devastating and automatic weapons systems is not assuaged by simply working with NSF support rather than Air Force Support. After all, scientists working on Air Force development projects read the project reports and articles in the open literature just as we do. Some deeper analysis of the connection between technical development, military needs, and the control of technology is needed. A central issue is not simply who funds AI, but who are the likely users of the techniques we develop and principles we uncover. (I don't intend to claim that if AI were supported almost exclusively by any one sponsor, regardless of whom, whether DOD, NSF, or a private philanthropist, that such centralized funding could not shape the course of AI development.) 6c3

A Cautionary Note on the Delphi Study 6d

The Delphi Study selected a panel composed of AI experts and

engineers to predict a set of potential AI-based products, predict their likely time of appearance as commercial systems, assess areas of likely application (e.g., library fact and reference retrieval), and assess the desirability of each application. This is the first comprehensive study of possible AI impacts and deserves recognition as a serious venture. Unfortunately, the methods used have two major flaws which limit the conclusions that may validly be drawn from this work. 6d1

1. Some descriptions of possible applications are removed from the human-social and organizational contexts. The expert is left to situate these applications in some meaningful context. For example, "increased utility of data bases, since data is better used and organized," is torn from the context of any particular data and data user. In addition, it triggers our attitudes that, in general, "better information" is preferable to "poorer information." Likewise, "greater opportunity for censorship" is remote from any sense of who may censor whom. Nevertheless, we don't like censorship in any form. 6d1a

2. Assessing impacts, such as those sketched above, requires some special skills in sociological, political, and economic analysis. Some impacts may be subtle, but powerful. For example, it appears that "increased utility of data bases" lends influence and power to the data user (Downs, 1967; Kling, 1973). To the extent that public bureaucracies are major users of large data-base systems, they will gain power over individual citizens and technically unsophisticated citizen's groups. (Such power shifts are typically not at the forefront of our attention when we think of computer impacts, but they may be substantial and difficult to reverse.) 6d1b

3. The Delphi technique is designed to synthesize a concensus of expert opinion. The experts selected for this study were sophisticated technologists. They may, or may not, be equally expert at assessing social impacts, and their social assessments must be viewed with that in mind. I appreciate that the Lockheed-SRI group, which carried out the study, was working with limited funds. Even though a two-panel study which includes both technical and assessment experts might be more useful, it would also be more costly and time consuming. 6d1c

Epilogue 6e

These notes are suggestive and incomplete. I welcome any

elaboration or commentary either via SIGART or through personal correspondence. 6e1

References 6f

1. Alford, Robert, "The Political Economy of Health Care: Dynamics without Change," POLITICS AND SOCIETY, pp. 1-38, Winter 1972. 6f1

2. Chartrand, R. L., COMPUTERS IN THE SERVICE OF SOCIETY, (Pergamon Press, 1972). 6f2

3. Downs, Anthony, "A Realistic Look at the Final Payoff from Urban Information Systems" PUBLIC ADMINISTRATON REVIEW, September 1967 (Also, reprinted in INFORMATION TECHNOLOGY IN A DEMOCRACY, ed. by Alan Westin, 1971.) 6f3

4. Firscheim, Oscar, et. al, "Forecasting and Assessing the Impact of Artificial Intelligence and Society," Proceedings 3rd International Joint Conference on Artificial Intelligence, pp. 105-120, (1973). 6f4

5. Kling, Rob, "Urban Data Systems and Shifts of Power," Internal Memo, Public Policy Research Organization, UC Irvine (July 1973). 6f5

6. Martin, James and Norman, THE COMPUTERIZED SOCIETY, (Prentice Hall, 1969). 6f6

7. Weizenbaum, Joseph, "On the Impact of the Computer on Society", SCIENCE, Vol. 176, pp. 609-614 (May 12, 1972). 6f7

8. Wells, H. G. "The Machine Stops" in OF MEN AND MACHINES, ed. by A.O. Lewis (Dutton, 1968). 6f8

General References on Computer Impacts for AI Buffs 6g

1. Gotlieb And Borodin, SOCIAL ISSUES IN COMPUTING, (Academic Press 1973). 6g1

2. Westin, Alan, ed., INFORMATION TECHNOLOGY IN A DEMOCRACY, (Harvard University Press, 1971). 6g2

Notes 6h

<N1> Speech-recognition and question-answering systems are the kind of artifacts I have in mind. 6h1

<N2> Some members of the computing community, especially those

who have attempted to develop integrated information systems for cities, hospitals, and complex organizations have begun to appreciate the difference between technical possibility and likely impact. 6h2

<N3> For the present argument, I assume that the Mead system is a useful aid. Certainly this is the hope of its designers and users. In its present state of development it is actually of marginal use, but promises more utility in the "near future." 6h3

<N4> In the 1980's powerful computers may become rather inexpensive--on the order of several hundred dollars for a fancy CPU and 65K of memory. While such devices may become easily accessable, they may not have much effect upon the cost of a (legal) informaton system, where the expense of gathering the data is a major fraction of the cost of service. 6h4

PROGRESS REPORT FROM SUNY AT BUFFALO by Teiji Furugori
Department of Computer Science
State University of New York at Buffalo 7

A ROBOT TO LEARN TO DRIVE A CAR 7a

Following Prof. David G. Hays's suggestions concerning a network form for the representation of knowledge in the human mind, we are constructing a robot that will create and use a human memory based on the linguistic experience of learning to drive a car. The main functional components in the model are a Parser, an Integrator, a Planner, and an Effector. The robot is to be simulated in a general-purpose digital computer. 7b

The robot has different functions during teaching time, driving time, and thinking time (not yet designed). During teaching time, the robot learns how to drive a car; it takes input sentences and builds its memory. The Parser and Integrator perform this function. During driving time, the robot receives commands and tries to drive its car on a highway. The Planner and Effector translate analyzed commands into programs of actions. During thinking time, the robot would change its memory structure to store its knowledge more consistently and efficiently. 7c

PRIMITIVES 7d

A few primitive concepts and transformations associated with them enable the robot to connect linguistic meaning directly to perceptions of the highway and the motor control mechanisms of the car it is driving. They are: 7d1

Events Entities Attributes 7d1a

turn	brake	left	right
push	accelerator	front	behind
release	speedometer	near	far
see	steering wheel	high	low
	lane		
	car		

7d1b

The command 'Turn the steering wheel to the left!' requires linguistic analysis, but then can immediately be translated into the control program TURN (STEERING WHEEL, LEFT). Some other commands that can be translated directly are 'Push the brake!', 'Push the accelerator!' and 'Release the accelerator!' The corresponding control programs are PUSH (BRAKE), PUSH (ACCEL) and REL (ACCEL).

7d2

Other primitives are used in conditional commands: 'If you see that the speedometer is high, then...' becomes ON(SPEED(HIGH))..., and so on.

7d3

For the system to work we need to specify the car the robot uses, other cars on the highway, and the highway itself. These are represented in a physical simulation; information about them is stored in the robot's cognitive network.

7d4

TEACHING TIME

7e

The teaching system assumes the primitives and consists mainly of giving the robot principles of driving, causal and sequential relations, and metalingual definitions. For example, the instructor might say, 'Increase speed means speed up. Speed up means that you push the accelerator.' To push the accelerator is primitive; hence the effect of the instruction is to make the internal representation of 'Push the accelerator' serve as a definition for speed up, and also, directly or indirectly, as a definition for the expression increase speed.

7e1

A more complex example is 'Change lane to the left means that you turn left, then when you come to the left lane you turn right.' The latter part of the instruction is to serve as a definition for change lane. The definition has two main clauses, linked sequentially (by then); the second main clause contains a condition (arriving in the left lane) and an action. This complex is stored symbolically in the cognitive network; the whole of it is a definition in which the smallest parts are primitives.

7e2

DRIVING TIME

7f

We now test the robot's ability to drive a car on a highway. Suppose the command is 'Increase speed!' The robot finds a metalingual definition, namely 'speed up', but this is not primitive. Another metalingual link leads to 'push the accelerator', which, being primitive, can be executed. The Planner having done this work passes the command to the Effector, which is in direct communication with the highway simulation. 7f1

However, circumstances may cause the robot to delay or alter execution of a command. 'Change lane to the left!' is impossible if a car is in the left lane and too close, for example. The Integrator leaves an internal representation of each command in touch with a symbolic description of the present situation. If the robot knows the principle that turning sharply while driving fast can cause trouble, it can match this principle against the composite of the current command and the current situation. When the match shows a conflict, the robot either ceases to pursue the purpose fixed by the command (and issues a report) or changes the situation, so that it can carry out the command. 7f2

The Planner develops out of instructions during teaching time a hierarchy of purposes; when conflicts occur, it follows the principle of altering its operation of the vehicle so as to violate only lower-order constraints and only temporarily, if possible. Recursive planning is necessary; each analysis of a command in a situation can reveal conflicts calling for new, internally generated commands that must be analyzed in turn. 7f3

A PROGRESS REPORT ON PROJECT CONSIM by Joe K. Clema

Department of Mathematics and Computer Science
Colorado State University
Fort Collins, Colorado 8

PROJECT CONSIM (Conflict Simulation) is a research effort devoted to emulating the processes of human decision making in complex situations requiring value judgements. The goals are (1) the development of a general computer decision aid and (2) a better understanding of how humans learn to improve in their decision making. The computer decision aid is designed to be prescriptive rather than to include the flaws of some particular human or group of humans, although "CONSIM I" did include the capability for incorporating some human imperfections. The model is of a mathematical and statistical nature and includes the following techniques: (1) linear programming, (2) alpha-beta search methods, (3) bayes's theorem, (4) linear evaluation function, (5) pattern recognition techniques, (6) heuristic procedures,

(7) computer modeling and simulation of human decision processes, (8) game theory, (9) utility theory, and (10) decision theory. 8a

PROJECT CONSIM first undertook research efforts in the area of international affairs. It was obvious that to validate the heuristic methodology employed in this project it must be obvious to the researchers whether the computer decisions are "good" or "bad." Thus, the effort during the last year has moved from international affairs to parlor games where good and bad decisions are quite easily related to wins and losses. 8b

The conflict simulation methodology has been applied in a simplified version of chess and has proven itself successful in a limited fashion. The measure of success was to play the program against a random opponent and count the number of moves in each game. If the "learning" program lost, a penalty of sixty moves was assigned, which was also the upper bound before beginning a new game. A learning procedure as described by Slagle, Samuel, and others was used. The number of moves required to win a first game (no learning) was significantly higher than latter games played with learning. The problem is basically to find a methodology enabling a computer program to re-adjust coefficients and in so doing find better coefficients enabling the program to recognize patterns and distinguish "good" situations from "bad" ones. A rather new methodology was employed utilizing a Bayesian approach and linear programming to achieve a successful re-evaluation of the coefficients. 8c

Present work involves streamlining the learning program and development of two subroutines which will both be equiped with a "learning" capability. These subroutines will play against each other and the results should prove quite interesting. The program has also been used to play against human opponents via an interactive terminal. Even though the program looks only one full move (2 plies) ahead, it has had some success against good human opponents, and nearly won a game! The improvement in play from game to game has been quite apparent at times. This project has only scratched the surface of the various mathematical, statistical, and programming techniques available. Further work involving parallel processing and non-linear evaluation functions should provide a vehicle for a general automated decision aid. 8d

[Ed. Note: Dr. Clema has informed us that this project has resulted in five refereed publications, a Doctoral Thesis, and several other publications.] 8e

AI FORUM 9

During the past several months I have become an interested

follower of the progress in AI, notwithstanding the articles of Sir James Lighthill at Cambridge. 9a

It appears that research on AI is extremely limited in the U.S. as compared with, say, Japan. It also appears that almost no effort has been directed towards using existing heuristic programming techniques in current business applications. The large financial resources of the business community remain untouched by those now involved in AI research, possibly because business programming has up to this time been handled by rather simple, unrelated algorithms. 9b

I am now employed as a banking systems engineer by the Electronic Data Systems Corporation. Heuristic programming is just about the farthest thing from current banking operations that bankers can imagine--but it is clear that the problems that will arise when the "checkless society" starts functioning can not be handled efficiently by the current patchwork programming efforts. 9c

I would greatly appreciate hearing from those who are considering AI techniques in business environments. I am convinced that several industries have quietly developed to the point where AI now offers a reasonable alternative to much larger programming staffs. 9d

Charles L. Bernier
P.O. Box 1027
Whiteville, N.C. 28472 9d1

CHESS 10

1. RESULTS OF THE FOURTH ANNUAL U.S. COMPUTER CHESS TOURNAMENT

ACM-73 Atlanta, Georgia August 26-28, 1973

by
Ben Mittman
Northwestern University
and
Monty Newborn
Columbia University 10a

Northwestern Remains Undefeated 10a1

CHESS 4.0, a completely rewritten version of Northwestern University's computer chess program, won its fourth consecutive title in the Fourth U.S. Computer Chess Championship sponsored by the ACM at ACM-73 in Atlanta, Georgia last August. CHESS 4.0 was written by David Slate and Larry Atkin, systems programmers at Northwestern's Vogelback Computing Center, with the help of Keith Gorlen, now with the U.S. Public Health

Service. It ended the tournament with three wins and one draw (against the Dartmouth chess program) for a total of 3 1/2 points in the four-round Swiss-style tournament. 10a2

Three programs ended in a tie for second place with 3 points each: OSTRICH from Columbia University, CHAOS from Sperry-Rand Univac, and TECH II from MIT. They are to have a play-off for the second-place trophy. 10a3

Other programs were entered by representatives of Carnegie-Mellon, Bell Labs, University of California at Davis, University of California at Berkeley, University of Southern California, Georgia Tech, and the College of William and Mary. Complete results and copies of the game scores appear below. 10a4

The tournament was organized by Prof. Monty Newborn of Columbia University and Prof. Ben Mittman of Northwestern. The tournament director was Mr. David Levy, an international master from London. Messrs. Newborn, Mittman, and Levy are currently trying to arrange for the first world computer chess championship at IFIPS-74 in Stockholm next August. 10a5

Chess programmers from any nation are urged to contact: 10a6

Prof. Ben Mittman
Vogelback Computing Center
Northwestern University
Evanston, Illinois 60201
U.S.A. 10a6a

[Ed. Note: Because of the larger number of games played in this tournament, we will publish only the games from the fourth and final round in this issue. We will be publishing games from the first three rounds, as well as play-off games, in future issues. 10a7

Aside from the first game of the fourth round (CHAOS vs. CHESS 4.0), the games in this round are in my opinion uniformly inferior and hardly worth playing if good chess is one's major interest. Nevertheless, it might be a worthwhile exercise if you're curious about the "machineomorphic" pattern of inferior play. 10a8

My own observation is that these chess programs do not seem to be susceptible to a common failing of mediocre human play--viz., "tunnel vision." I've never made a psychological study, but it seems to me that a poor human player tends to become so caught up in the local tactics of a highly goal-directed sequence that he frequently fails to capitalize on new targets

of opportunity as they present themselves. The programs, on the other hand, always appear to maintain a global perspective and invariably postpone the natural development of a plan in order to execute an irrelevant check or other forcing move, even when it contributes literally nothing to the plan in progress, which is subsequently resumed just as though the distraction never occurred. Occasionally, this non-anthropomorphic feature of the programs is surprising, since it leads to something good and gives the illusion that the programs are better than they really are. In conclusion, it appears that poor human players and poor machines have something to learn from one another.] 10a9

2. FOLLOW UP ON THE MATCH 10b

After the match CHESS 4.0 and CHAOS both played a simultaneous exhibition against Mr. Charles Kalme, a Senior Master, who was a consultant to the USC team. Mr. Kalme played at queen odds, i.e., he removed his queen from the board. He defeated CHAOS (winning a $100 bet) and was defeated by CHESS 4.0. 10b1

3. RECENT ARTICLES ON COMPUTER CHESS 10c

(A) "Some Necessary Conditions for a Master Chess Program" by Hans J. Berliner, Proc. IJCAI-73, pp. 77-85. 10c1

(B) "Can a Computer Beat Bobby Fischer?" by Benjamin Mittman, DATAMATION, pp. 84-87, June 1973. 10c2

4. PRELIMINARY TESTING OF THE EFFECTIVENESS OF THE CICHELLI DEPTH-2 AND REFUTATION HEURISTICS <*N1> 10d

Although the heuristics were designed to improve alpha-beta pruning in middle game play, a test of the program's integrity and heuristics was made on 10 two-move-mate problems. Two runs were made, one with and one without the heuristics enabled. The CPU time and move counts were tabulated for the runs. 10d1

Visiting a node necessarily entails a call to SELECTMOVE, MAKEMOVE, LISTMOVES, and eventually REVERSEMOVE. If the heuristics were applied, a call to PRERATE provided pointers for the DEPTH-2 data and statically-ordered plies with both refutation and DEPTH-2 data. Without the pointers no dynamic ordering was possible. 10d2

The positions were loaded identically for the two runs so that the search trees would be identical if no reordering occurred. Without the heuristics, the program solved the 10 problems by

generating 74,485 nodes at 35 nodes per second. The overall improvement is 212% 10d3

Of particular note is that in general the larger the search is without the heuristics, the greater the improvement with the heuristics. For example, problem 10 had an improvement of 735% 10d4

Problem 1: Bell <*R1>, Figure 3. 10d5

Whites: (11 pieces)
Bf1, Rg1, Kh1, Pe2, Pg2, Pg3, Bh3, Pg4, Pe6, Pg6, <*N2> 10d5a

Blacks: (11 pieces)
Pe3, Pe4, Pe5, Pf5, Pg5, Pf6, Pe7, Pg7, Bf8, Rg8, Kh8. 10d5b

Solution: Pg4*f5 10d5c

Nodes with the heuristics: 6
Nodes without the heuristics: 6 10d5d

Problem 2: Bell, Figure 1b. 10d6

Whites: (11 pieces)
Kb2, Qf2, Nb3, Pf3, Ph4, Ra5, Pg5, Bc7, Ba8, Nf8. 10d6a

Blacks: (7 pieces)
Pd3, Pd4, Pf4, Nd5, Qe5, Kf5, Pf6. 10d6b

Solution: Qf2*d4 10d6c

Nodes with the heuristics: 1388
Nodes without the heuristics: 3994 10d6d

Problem 3: Bull <*R2>, No. 41. 10d7

Whites: (8 pieces)
Ba1, Kh1, Pe2, Pg2, Ph6, Ba8, Nc8, Qg8. 10d7a

Blacks: (9 pieces)
Pa2, Ph2, Pe3, Kf4, Pe5, Pe6, Pb7, Pe7,Ph7.. 10d7b

Solution: Nc8*e7 10d7c

Nodes with the heuristics: 222
Nodes without the heuristics: 832 10d7d

Problem 4: Bull, No. 44. 10d8

Whites: (5 pieces)
Ne1, Pe2, Pe4, Qe7, Ke8. 10d8a

Blacks: (2 pieces)
Ke5, Pe6. 10d8b

Solution: Pe2-e3 10d8c

Nodes with the heuristics: 997
Nodes without the heuristics: 1258 10d8d

Problem 5: CHESS LIFE <*R3>, No. 1. 10d9

Whites: (9 pieces)
Re1, Bd4, Rc5, Ne5, Ph5, Pg6, Ka7, Pd7, Qf8. 10d9a

Blacks: (7 pieces)
Ba1, Pg4, Qh4, Bb5, Pc6, Ke6, Pb7. 10d9b

Solution: Qf8-c8 10d9c

Nodes with the heuristics: 6461
Nodes without the heuristics: 2839 10d9d

Problem 6: CHESS LIFE, No. 2. 10d10

Whites: (11 pieces)
Qd2, Pe2, Nf2, Bh3, Rc4, Nf5, Pc6, Pd6, Pe7, Kf8, Bh8. 10d10a

Blacks: (9 pieces)
Pc3, Pf4, Nb5, Qd5, Pa6, Ke6, Ph6 Rd7, Bf7. 10d10b

Solution: Nf2-e4 10d10c

Nodes with the heuristics: 8303
Nodes without the heuristics: 14706 10d10d

Problem 7: CHESS LIFE, No. 3. 10d11

Whites: (11 pieces)
Rc1, Ka2, Pg2, Bb3, Rc4, Bh4, Pb6, Pf7, Nd8, Qf8.. 10d11a

Blacks: (10 pieces)
Pc2, Pe3, Pe4, Pe5, Rg5, Rh5, Kd7, Ng7, Be8, Ng8. 10d11b

Solution: Rc4*c2 10d11c

Nodes with the heuristics: 3013
Nodes without the heuristics: 10521 10d11d

Problem 8: CHESS LIFE, No. 4. 10d12

Whites: (7 pieces)
Bb1, Ne1, Rc2, Bd4, Pb5, Qc5, Ka6. 10d12a

Blacks: (12 pieces)
Qh2, Pb3, Pf3, Pg3, Ke4, Ph4, Pd5, Pf5, Bb6, Pa7, Rc8, Ng8. 10d12b

Solution: Qc5-d6 10d12c

Nodes with the heuristics: 5177
Nodes without the heuristics: 2849 10d12d

Problem 9: CHESS LIFE, No. 5. 10d13

Whites: (10 pieces)
Bc1, Kf1, Pg2, Ne3, Bf3, Ra4, Ne5, Re6, Qc7, Pf7.. 10d13a

Blacks: (9 pieces)
Ng1, Rb4, Kf4, Rh4, Ph5, Pc6, Pf6, Pe7, Pg7.. 10d13b

Solution: Bf3-g4 10d13c

Nodes with the heuristics: 5144
Nodes without the heuristics: 11279 10d13d

Problem 10: CHESS LIFE, No. 6. 10d14

Whites: (11 pieces)
Rd1, Pe2, Pf2, Nc4, Pg4, Rf5, Kd6, Ne6, Bh6, Bb7, Qe8.. 10d14a

Blacks: (8 pieces)
Nh1, Qd3, Rh3, Ke4, Pd5, Bg5, Pf6, Ph7.. 10d14b

Solution: Rf5*d5 10d14c

Nodes with the heuristics: 3013
Nodes without the heuristics: 26201 10d14d

References 10d15

<R1> Bell, A. G., "How to Program a Computer to Play Legal Chess", (THE COMPUTER JOURNAL, May 1970). 10d15a

<R2> Bull, T. P., CHESS PROBLEMS, (O. A. Brownson, Rockdale, Illinois, [date unknown]). 10d15b

<R3> CHESS LIFE AND REVIEW, "Two-Move Awards," March 1972. 10d15c

Notes 10d16

<N1> See previous issue, SIGART Newsletter, pp. 32-36, June 1973. 10d16a

<N2> [Ed. Note: Using algebraic chess notation, board columns are labeled with the lower case leters from "a" to "h" while rows are numbered from "1" to "8". Thus, the lower left-hand square of the chess board is labeled "a1" while the upper right-hand square is labeled "8h". The coordinate system orientation remains the same for both white and black. 10d16b

CONFERENCES 11

1. COMPUTER SCIENCE CONFERENCE 11a

February 12-14, 1974, Detroit Hilton, Detroit, Michigan 11a1

This conference is sponsored jointly by the ACM and a number of universities and industrial organizatons, in cooperation with the Computer Society of the IEEE and the Computers in Education Division of ASEE. Partial support has also come from NSF. The conference is primarily devoted to short, current research reports. Invited papers will be presented by Juris Hartmanis of Cornell University discussing complexity theory, Thomas Cheatham of Harvard University on extensible languages and automatic programming, Hurbert Dreyfus of the University of California at Berkeley on "The Ever Incomplete Robot," and Frederick Brooks of the University of North Carolina on computer graphics. The banquet presentation will be given by John Opel, Senior Vice President of IBM. 11a2

Participants in the conference need only submit abstracts. The deadline for abstracts is December 1, 1973. The printed program will be prepared directly from the abstracts submited. This form of presentation was well received during the first Computer Science Conference held last February in Columbus, Ohio. 11a3

In addition to the technical program and invited speakers, a number of other activities are planned. Two professional societies will hold meetings at the end of the Computer Science Conference: the Special Interest Group on Computer Science Education of ACM and the Computers in Education Division of ASEE. 11a4

For additional details on this conference see earlier Newsletters (p.45, June 1973 and pp. 41-42, August 1973). 11a5

Information concerning all aspects of the conference can be obtained from: 11a6

Seymour J. Wolfson, Chairman
Computer Science Conference
Computer Science Section
Wayne State University
Detroit, Michigan 48202 11a6a

2. 1974 NATIONAL COMPUTER CONFERENCE AND EXPOSITION 11b

The 1974 National Computer Conference and Exposition will be held May 6-10 in Chicago, Illinois, according to an announcement by the American Federation of Information Processing Societies, Inc. The 74 NCC will be the year's largest gathering of the world-wide computer community and is expected to attract more than 35,000 attendees to its sessions, seminars, and extensive exhibit program. Overall planning for the conference will be handled by a Chicago-based Steering Committee under the direction of Dr. Stephen S. Yau, Conference General Chairman. Dr. Yau is Chairman of the Computer Sciences Department, Northwestern University, Evanston, Illinois. 11b1

3. PROCEEDINGS OF 1973 NATIONAL COMPUTER CONFERENCE 11c

The Proceedings of the 1973 National Computer Conference and Exposition are now available from the American Federation of Information Processing Societies, Inc. The 920-page hard-cover volume contains more than 160 technical papers and abstracts covering a wide range of topics in computer science, technology, methods, and applications featured at the 73 NCC held in New York, June 4-8. The price for the Conference Proceedings, Volume 42, is $40. A reduced rate of $20 is available for prepaid orders from ACM members stating their affiliation and membership number. The Proceedings may be ordered from: AFIPS Press, 210 Summit Avenue, Montvale, New Jersey 07645. 11c1

4. SECOND TEXAS CONFERENCE ON COMPUTING SYSTEMS 11d

November 12-13, 1973, Austin, Texas 11d1

The Texas Conference on Computing Systems is an annual forum for the presentation of state-of-the-art practice in computing systems and research results. This tentative list of sessions and invited papers indicates the scope and depth of the conference. Sessions will feature surveys of state-of-the-art practice as well as reports on recent specific developments by means of invited and contributed papers. 11d2

A partial list of invited participants follows: 11d3

PROGRAMMING LANGUAGES
Chm Harlan Mills (IBM)
S.R. Kosaraju (John-Hopkins) 11d3a

APPLICATIONS OF COMPUTATION THEORY
Chm K.S. Fu (Purdue)
Philip Lewis (General Electric Research Lab)
C.L. Liu (university of Illinois) 11d3b

OPERATING SYSTEMS
Chm James C. Browne (University of Texas)
George H. Mealy (Harvard) 11d3c

SYSTEM EVALUATION AND OPTIMIZATION
Chm C. V. Ramamoorthy (UC Berkeley)
William Lynch (Case-Western Reserve)
John Tarter (University of Alberta, Canada) 11d3d

DATA MANAGEMENT AND INFORMATION RETRIEVAL
Chm Paul deMaine (Penn State)
Robert Simmons (University of Texas) 11d3e

MANAGEMENT OF COMPUTING FACILITIES
Chm Timothy Rueğli (University of Texas)
K. Knight (University of Texas) 11d3f

COMPUTER COMMUNICATIONS AND NETWORKS I
Chm Mani Chandy (University of Texas)
Julius Aronofsky (SMU)
Donald Aufenkamp (National Science Foundation)
Eric Manning (University of Waterloo, Canada)
Jerry Weeg (University of Iowa)
Paul Green (IBM) 11d3g

COMPUTER COMMUNICATIONS AND NETWORKS II
Chm Robert Kuhn (ARPA)
L. Kleinrock (UCLA)
Eric Manning (University of Waterloo, Canada) 11d3h

RELIABILITY AND DIAGNOSIS
Chm Stephen Szygenda (University of Texas)
Herbert Chang (Bell Laboratories)
Francis Mathur (University of Missouri) 11d3i

COMPUTER ARCHITECTURE
Chm Joseph Watson (Texas Instruments)
Michael Flynn (John-Hopkins) 11d3j

MINICOMPUTER SYSTEM
Chm Frank Spiznogle (Texas Instruments)
John Allan (University of Texas) 11d3k

For more information contact Program Chairman: 11d4

Professor Terry Welch
Department of Electrical Engineering
University of Texas
Austin, Texas 78712. 11d4a

5. FIRST INTERNATIONAL JOINT CONFERENCE ON PATTERN RECOGNITION 11e

October 30 - November 1, 1973, Washington, D.C. 11e1

The First International Joint Conference on Pattern Recognition is intended to bring together scientists and engineers to report their latest research and developments, and to discuss the directions and goals for future work in pattern recognition. It is hoped that the conference will assist in coordinating the many activities in pattern recognition which are presently narrowly compartmented along lines of specialization. 11e2

Sixty-five papers covering such diversified subjects as Mathematical Methods, Character Recognition, Biomedical Applications, Picture Processing, Speech, Syntactic Methods, Adaptive Pattern Recognition, Scenes and Structures, and Remote Sensing will be presented by leading scientists and engineers from ten different countries. In addition, there will be two workshops: one entitled, "Gap between Theory and Practice" and the second, "Problems in Pattern Recognition Research." 11e3

The conference is being sponsored by the following societies: ACM, IEEE, IFIPS, OSA, PRS, and SPIE. 11e4

For further information contact: 11e5

Louis S. Rotolo
Pattern Recognition Society
P.O. Box 629
Silver Spring, Maryland 20901
202-625-2121 11e5a

ABSTRACTS 12

CARNEGIE-MELLON WORKING PAPERS IN SPEECH RECOGNITION - II
Department of Computer Science
Carnegie-Mellon University

Pittsburgh, Pennsylvania
August 1973 12a

This report contains three previously published papers and two unpublished ones: 12a1

D. R. Reddy, L. D. Erman, and R. B. Neely, "A Model and a System for Machine Recognition of Speech," IEEE Trans. Audio and Electroacoustics, AU-21 (3), June, 1973. 12a2

D. R. Reddy, L. D. Erman, R. D. Fennell, and R. B. Neely, "The HEARSAY Speech Understanding System: An Example of the Recognition Process," Proc. of the IJCAI-73, Stanford, Calif., August, 1973. 12a3

L. D. Erman, R. D. Fennell, V. R. Lesser and D. R. Reddy, "System Organizations for Speech Understanding: Implications of Network and Multiprocessor Computer Architectures for AI," IJCAI-73, August 1973. 12a4

Janet M. Baker, "A New Time-Domain Analysis of Human Speech," April, 1973. 12a5

James Baker, "Machine-Aided Labeling of Connected Speech," April, 1973. 12a6

ANALYSIS OF THE ALPHA-BETA PRUNING ALGORITHM by S. H. Fuller, J. G. Gaschnig, and J. J. Gillogly
Department of Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 12b

An analytical expression for the expected number of bottom positions examined in a game tree using alpha-beta pruning is derived, subject to the assumptions that the branching factor N and the depth N**D of the tree are arbitrary but fixed, and the bottom positions are a random permutation of N**D unique values. A simple approximation to the growth rate of the expected number of bottom positions examined is suggested, based on a Monte Carlo simulation for large values of N and D. The behavior of the model is compared with the behavior of the alpha-beta algorithm in a chess-playing program and the effects of correlation and non-unique bottom position values in real game trees are examined. 12b1

THE EQUIVALENCE OF REDUCING TRANSITION LANGUAGES AND DETERMINISTIC LANGUAGES by Mario Schkolnick
Department of Computer Science

Carnegie-Mellon University
Pittsburgh, Pennsylvania 12c

The class of reducing transition languages introduced by Eickel, Paul, Bauer, and Samelson was shown by Morris to be a proper superclass of the Simple Precedence Languages. In this paper we extend this result showing that in fact, the first class is equivalent to the class of Deterministic Context-Free Languages. 12c1

LABELLED PRECEDENCE PARSING by Mario Schkolnick
Department of Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 12d

Precedence techniques have been widely used in the past in the construction of parsers. However, they imposed restrictions on their grammars that were hard to meet. Thus, alteration of the rules of a grammar was necessary in order to make them acceptable to the parser. We have shown that, by keeping track of the possible set of rules that could be applied at any one time, one can enlarge the class of grammars considered. The possible set of rules to be considered is obtained directly from the information given by a labeled set of precedence relations. Thus, the parsers are easily obtained. Compared to precedence parsers, this new method gives a considerable increase in the class of parsable grammars, as well as an improvement in error detection. An interesting consequence of this approach is a new decomposition technique for LR parsers. 12d1

MODEL VERIFICATION AND IMPROVEMENT USING DISPROVER by L. Siklossy and J. Roach
Department of Computer Sciences
University of Texas at Austin 12e

Confidence in the adequacy of a model is increased if tasks that are impossible in the world are shown to correspond to disprovable tasks in the model. DISPROVER has been used as a tool to test, in worlds of robots, the impossibility of tasks related to various conservation laws (objects, position, model consistency, etc.) and time constraints. The adequacy and sufficiency of operators can be established. Interacting with DISPROVER, the model designer can improve his axiomatization. The frontier between "acceptable" and "ridiculous" axiomatizations is shown, in many examples, to be a most tenuous one. 12e1

ON THE PREPROCESSING OF RADIOGRAPHIC IMAGERY by Y. P. Chien and K. S. Fu

School of Electrical Engineering
Purdue University
West Lafayette, Indiana 12f

The main idea behind the preprocessing of picture patterns is to effectively reduce the large amount of data, so that it will be easy for us to extract significant features. In other words, in the preprocessing phase, we would like to reduce the so-called irrelevant data and preserve the significant information for a later classification stage. These all seem to be necessary for us to develop a preprocessing technique that has the capability of singling out the specific object of interest from the background. This idea can also be termed "field of vision." In order to remove irrelevant data, we should extract the desired object as a discrete entity and "zoom in" on this particular region. Then the set of features extracted in this region would not be meaningless. 12f1

The concept of "field of vision" plays an important role in most pictorial pattern recognition problems. As in the case of automated computer diagnosis of chest x-rays, we would like to first locate the lung region and heart region, and then proceed to extract the set of features in each region. If we have no knowledge of the location of regions of lung or heart, it would be meaningless to talk about whatever features we extract. Thus, the first step in preprocessing the picture patterns is to locate the region of interest. This is equivalent to picture segmentation or boundary detection. The segmentation of binary pictures does not pose any problem. However, for pictures of multi-grey levels, the boundary for defining the region of interest is not easy to detect. 12f2

RECOGNITION OF X-RAY PICTURE PATTERNS <*N1>by Y. P. Chien and K. S. Fu
School of Electrical Engineering
Purdue University
West Lafayette, Indiana 12g

The "field of vision" is a very important concept in pictorial pattern recognition. In this paper, it is shown that this concept should be used as the preliminary step in preprocessing if automatic picture processing is ever to be of practical use. The preprocessing technique suggested in this paper will enable us to segment the picture into subregions so that it is possible for us to "zoom" into the specific objects. Thus the set of features we extract in that region would at least be meaningful. 12g1

In this paper, x-ray pictures with venus hypertension are used

as an application example. In the preprocessing phase, we first locate the coarse boundary of the lung by registering 5 key points on the lung boundary. Then the detailed lung boundary could be easily located by any existing boundary-finding technique. However, even in the present case, it is shown that the set of texture features extracted from the coarse lung region could be used effectively to screen out abnormal pictures. 12g2

<N1> This paper will be presented at the 1973 Annual conference of the Society of Photographic Scientists and Engineers, May 6-11, 1973, Rochester, New York. 12g3

ON SPEAKER IDENTIFICATION USING COARTICULATION OF NASAL CONSONANTS WITH VOWELS by Lo-Soun Su and K. S. Fu
School of Electrical Engineering
Purdue University
West Lafayette, Indiana 12h

A new approach which used the statistical properties of the nasal spectra was used to quantitatively study the coarticulation of nasal consonants with the vowels in isolated /h 'CVd/ utterances. The mean spectra difference of the nasal followed by front vowels and by back vowels was used as the acoustic measure of coarticulation. The coarticulation of [m] was found to be strongly speaker dependent in particular. This coarticulation, which generally reflects each individual speaker's idiosyncratic characteristics and is not likely to be consciously modified in natural speech, was proved to provide good acoustic clues for speaker identification. Speaker identificaton was performed using the correlation decision criterion and the results indicate that coarticulation clues are much better than the nasal spectrum, which has already been proved to be one of the best acoustic clues for speaker identification. 12h1

APPLICATION OF THE TREE SYSTEM APPROACH TO CLASSIFICATION OF BUBBLE CHAMBER PHOTOGRAPHS, by B. K. Bhargava and K. S. Fu
School of Electrical Engineering
Purdue University
West Lafayette, Indiana 12i

This report concerns the application of the tree system approach to classification of bubble chamber photographs. The report contains details of the classification procedure and results obtained from real photographs. The results were very encouraging, and various programs are being improved at present to get greater efficiency of computation and to reduce the ambiguity in classificaton. 12i1

DYNAMIC SCHEDULING OF LARGE DIGITAL COMPUTER SYSTEMS USING ADAPTIVE CONTROL AND CLUSTERING TECHNIQUES <*N1>by Richard A. Northouse
Electrical Engineering and Computer Science Department
University of Wisconsin
Milwaukee, Wisconsin
and
King-Sun Fu
School of Electrical Engineering
Purdue University
Lafayette, Indiana 12j

This research is directed toward the development of a scheduling algorithm for large digital computer systems. To meet this goal, methods of adaptive control and pattern recognition are applied. As jobs are received by the computer, a pattern recognition scheme is applied to the job in an attempt to classify its characteristics, such as a CPU-bound job, an I/O job, a large memory job, etc. Simultaneously, another subsystem, using a linear programming model, evaluates the overall system performance, and from this information an optimized (or desired) job stream is determined. When the processor requests a new job, it is chosen from the various classifications in an attempt to meet the optimal (or desired) job stream. 12j1

After the jobs are completely processed, their characteristics are compared to the projected classification produced by the pattern discriminant function. The results are then returned to the discriminant function to update the decision mechanism, a minimum-distance discriminant function. From a systems point of view, this results in an adaptive or self-organizing control system. The overall effect is a dynamic scheduling algorithm. 12j2

Simulation studies indicated that the scheduler was able to adapt to changing work loads, and it improved the turnaround times significantly. These simulation studies were based on a multiprocessor-uniprogram environment. 12j3

<N1> IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-3, No. 3, pp. 225-234, May 1973. 12j4

A UNIFICATION ALGORITHM FOR TYPE THEORY by Gerard P. Huet
AIRIA-Laboria
Domaine de Voluceau
78150-Rocquencourt, FRANCE 12k

A semi-decision algorithm to search for unification of formulas in

-order type theory is presented, and its correctness proved. It is shown that the search space is considerably simpler than one for most general unifiers. This allows our algorithm to have good directionality and convergence properties. 12k1

Available as "Rapport de Recherche," No. 23, Laboria, Juillet 1973. 12k2

GADGET: A PROGRAM THAT GENERATES PROGRAMS FOR TESTING SOME PROPERTIES ABOUT GRAPHS AND SETS by Anne Adam
Third Cycle Thesis - University of Paris VI
Laboratoire de Mathematiques Appliquees
Universite de CAEN, FRANCE 121

This program accepts mathematical statements concerning properties of graphs and sets. It generates a program corresponding to the data in FORTRAN IV. In addition to being a compiler for a new high-level language, it has the following capabilities: 1211

(1) simplification rules are applied to the data; 1211a

(2) theorem-proving is used for recognizing valid statements or contradictions; 1211b

(3) some properties of the source statements such as reflexivity and symmetry are used for optimizing loops; and 1211c

(4) dominance relations in the generated program are used for simplifing branches. 1211d

About fifty programs have been generated by GADGET thus far. 1212

SCENE ANALYSIS FOR BREADBOARD MARS ROBOT FUNCTIONING IN AN INDOOR ENVIRONMENT by Martin D. Levine
Jet Propulsion Laboratory
Pasadena, California
September 1, 1973 12m

This report deals with the problem of computer perception in an indoor laboratory environment containing rocks of varios sizes. Such sensory data processing is required for the NASA/JPL breadboard mobile robot that is a test system for an adaptive variably-autonomous vehicle that will someday conduct scientific explorations on the surface of Mars. Scene analysis is discussed in terms of object segmentation followed by feature extraction, which results in a representation of the scene in the robot's world model. 12m1

ON THE SEGMENTATION PROCESS IN SCENE ANALYSIS by Martin D. Levine
Department of Electrical Engineering
McGill University, Canada
Report No. 73-25, August 1973 12n

The problem of scene analysis in artificial intelligence is concerned with obtaining a three-dimensional description of the objects in a digitized representation of the scene. A basic paradigm for scene analysis is presented which includes as sub-models the areas of picture processing, picture analysis, and pattern recognition. Central to this problem is the segmentation process which partitions the picture into subsets of points which constitute atoms. It is argued by reference to the literature and a proposed scene taxonomy that all procedures to date except one result in atoms which can be categorized by humans as recognizable objects. These procedures are classified according to whether they are context-independent or context-dependent. The more general scene analysis paradigm is then discussed in detail and associated with the study of human psychologial behavior. 12n1

AUTOMATED PRESCREENING OF CERVICAL CYTOLOGY SPECIMEN by Ronald S. Poulsen
Department of Electrical Engineering
McGill University, Canada
Ph.D. Thesis, March 1973 12o

In this thesis the cervical cytology prescreening problem is examined in detail and a dual-resolution image processing method is proposed for automating the analysis of the routine cervical smear. The feasibility of this method is established through a comparison of the computer results with those of a manual study involving a large number of cervical specimens from patients with cancerous or precancerous conditions of the uterine cervix. 12o1

In this research an interactive image processing system has been used to develop a model of the abnormal cell class specifically in the context of cervical smears. These studies demonstrate that the image processing system developed here is capable of detecting the vast majority of isolated and slightly overlapping abnormal cells occurring in routine cervical smears and hence, is capable of prescreening these specimens into suspect and normal categories. 12o2

COMPUTER DETERMINATION OF TOTAL LUNG CAPACITY FROM X-RAY IMAGES by J. Lawrence Paul
Department of Electrical Engineering

McGill University, Canada
M. Eng. Thesis, March 1973 12p

This thesis describes an attempt to automatically calculate total lung capacity (TLC) from the posteroanterior and lateral chest X-ray images by computer. The lung model is that of Barnhard, which assumes that the cross-sectional shape of each lung is elliptical. Major and minor axes of each ellipse are determined from the lateral and posteroanterior projections respectively, and the integration of the elliptical cylindroids gives total thoracic volume. From this, values representing heart volume, hemi-diaphragm domes, blood volume, and tissue volume are subtracted, leaving the total lung capacity. 12p1

The major portion of this work describes the feature extraction procedures used to locate such features as the heart, lung outline, and diaphragm in both projections. The feature extration techniques employ feature locaton by means of directional signatures, edge detection by weighting the derivatives of the picture points with a Gaussian function, and binary conversion. Machine calculated results are consistently lower than hand calculated results, but the high correlation coefficient (0.971) suggests that true TLC may be linearly related to machine calculated values. 12p2

CANADIAN A.I. SOCIETY FORMED

Canadian Society for Computational Studies of Intelligence
(Societe Canadienne des Etudes d'Intelligence par Ordinateur)
by
Zenon Pylyshyn
Departments of Psychology and Computer Science
The University of Western Ontario
London, Canada 13

The AI group at the University of Western Ontario, London, Canada recently invited a number of people from several Canadian universities to get together at U.W.O. to discuss the present state and future prospects for AI research in Canada. The response was more than we had hoped for. More than 30 people arrived on May 23, 1973 for a three-day meeting. Tutorial and position talks were given on six topics: Mathematical Studies (T. Pietrykowski, Waterloo; R. Reiter, U.B.C.); Psychological Modeling (Z. Pylyshyn, U.W.O.; G. Baylor, J. Gascon, U. Montreal); Pattern Recognition and Machine Perception (E.W. Elcock, U.W.O.; W.A. Davis, U. Alberta; M. Levine, McGill; R. Paulsen, McGill); Languages for AI (J. Mylopoulos, Toronto; E.W. Elcock, U.W.O.); AI Education (R. Rosenberg, U.B.C.); and Applications in Teleconferencing (B. Bridgewater, P. Allard, C. Billowes,

Communications Research Centre, Department of Communications, Ottawa). 13a

Besides giving Canadian AI researchers an opportunity to find out what others in the country were doing, the most important result of the meeting was the formation of a new society dedicated to the advancement of research in human and machine intelligence. The society is to be called The Canadian Society for Computational Studies of Intelligence [Societe Canadienne des Etudes d'Intelligence par Ordinateur]. Besides promoting development and applications of research the society will maintain links with government and industry, will organize meetings, will set up special study groups to investigate issues relevant to the interests of the Society, and will publish a bulletin as a means of communication among members. A steering committee chaired by E. W. Elcock (Department of Computer Science, U.W.O.) with J. Hart (U.W.O.) as secretary/treasurer will oversee the initial development of the Society. Committee members include, G. Baylor, W. Davis, J. Gascon, M. Levine, J. Mylopoulos, T. Pietrykowski, R. Reiter, R. Rosenberg, and Z. Pylyshyn. 13b

One of the committee's first tasks will be to undertake a study of the computer resource needs of the Canadian AI community. Study groups have also been set up to investigate sources of research funding and to look into the possibilities for exchange of educational materials. 13c

Anyone who is interested in this fledgling society is invited to write the Editor of the bulletin requesting that his name be put on the mailing list: 13d

Professor R. Reiter
Computer Science Department
University of British Columbia
Vancouver, B.C., CANADA 13d1

If you wish to become a member send a $3.00 membership fee to the secretary/treasurer (Professor J. Hart, Computer Science Department, University of Western Ontario, London, Canada). 13e

ARTIFICIAL INTELLIGENCE JOURNAL
[Ed. Note: The following letter was handed out to each registrant at IJCAI-73 and is reprinted here for your further consideration.] 14

On behalf of the Editorial Board, listed below, we invite all registrants at IJCAI-73 to submit, and encourage colleagues to submit, research papers. We also invite IJCAI authors to submit their papers for possible publication in the more permanent form

of our Journal in addition to publication in the Conference Preprints. 14a

ARTIFICIAL INTELLIGENCE is the only international journal centrally devoted to AI. In spite of the great amount of activity in this field, and the considerable volume of research results reported at AI conferences and in university theses, laboratory reports and internal memoranda, the Journal is barely able to obtain enough high-quality material to maintain quarterly production at present. This seems to be largely due to very many workers in the field being entirely content, as far as communication is concerned, with exchange of internal reports and occasional attendance at conferences. Useful as these are, they need to be supplemented by recourse to a well-recognized regular journal for at least three good reasons: 14b

Firstly, workers in the field can be saved a great deal of time and trouble by having a high-quality sieve for selecting, generally, the best core of current work and results. Secondly, a properly supported and recognized journal provides a running overview of the state and development of the subject as a whole for specialists in particular departments of it. Last and by no means least, AI is a subject of such significance and implications for others--e.g., computer scientists, psychologists, linguists, and philosophers--who do not have access to the "inner" grapevine, that it would be irresponsible not to use the obvious means of a central journal to keep them abreast of the subject. 14c

Bernard Meltzer, Editor-in-Chief
Bertram Raphael, Associate Editor 14c1

Editorial Board: Professor M. Aiserman, Professor S. Amarel, Professor W. Bledsoe, Dr. R. Burstall, Dr. M. Clowes, Dr. E. Elcock, Professor E. Feigenbaum, Professor R. Gregory, Professor J. McCarthy, Professor D. Michie, Professor A. Newell, Dr. N. Nilsson, Dr. J. Pitrat, Dr. E. Sandewall, Dr. D. Walker, Professor T. Winograd. 14d

AI ON TV 15

The offerings for the new Fall season are pretty slim. As reported earlier, <*N1> "The Six Million Dollar Man" has been slated to appear on ABC on Saturdays once a month from 8:30-10:30 PM. However, "Genesis II" does not appear to have made it for this Fall. 15a

The only other offerings, "The Starlost," is set in the year 2285 and described by its creator, writer Harlan Ellison, as "the story of three young people discovering their world, and their place in

it." A Canadian production, it stars Keir Dullea, Gay Rovan, and Robin Ward. 15b

<N1> SIGART Newsletter, No. 39, p. 36, April 1973. 15c

ASIMOV'S "I, ROBOT" 16

Long out of print and in great demand, Isaac Asimov's famous science fiction classic has now been reprinted in paperback by Fawcett Crest Publications, Inc., Grenwich, Connecticut; 75 cents. Most of the ten short stories contined in "I, Robot" were originally published during the decade of the 1940s in ASTOUNDING SCIENCE FICTION. However, there is very little about them that seems "dated," as is the case with many other now quaint stories from this time period. Of course, "The Rest of the Robots" (Pyramid Books, New York, 1964, paperback, 75 cents), which picks up where "I, Robot" leaves off with eight more stories exploring the theme of Asimov's now famous "three laws of robotics," is still availale. 16a

ADVERTISEMENTS 17

REF LSC 23-JAN-75 20:45 25201

SIGART NEWSLETTER Number 42 October 1973

(J25201) 23-JAN-75 20:45;;;; Title: Author(s): Richard E. Fikes, L. Stephen Coles/REF LSC; Distribution: /DVN([INFO-ONLY]) GCE([INFO-ONLY]) KIRK([INFO-ONLY]) ; Sub-Collections: NIC; Clerk: DVN; Origin: <SIGART>OCT73.NLS;53, 15-OCT-73 12:45 KIRK ; ####;

GSG 23-JAN-75 23:23 25203

Show of network journal delivery.

This is a test to show Yngvar how the netowrk journal system can delivery the text of a journal message to his message file, as well as his initial file. This should appear shotly in his mailbox at SRI-AI. 1

GSG 23-JAN-75 23:23 25203

Show of network journal delivery.

(J25203) 23-JAN-75 23:23;;;; Title: Author(s): Geoffrey S. Goodfellow/GSG; Distribution: /YL([INFO-ONLY]) ; Sub-Collections: NIC; Clerk: GSG;

YL 24-JAN-75 02:25 25204

From Yngvar Lundh, Norway.

Computer Aided Design etc, 24-JAN-75 02:15 1

Ed: 1a

Jake Feinler mentioned your name to me when I asked her if she knew about anyone interested in Computer Aided Design. I am interested in improving the design process for logic large scale integrated circuits. Do you have any comments? 1b

I can be reached by journal to YL or sndmsg to norsar-tip@sri-arc attn: Yngvar Lundh. (I shall be out of town next week). 1c

Regards from Yngvar (at the Norwegian Defence Research Establishment, 2007 Kjeller, Norway). 1d

YL 24-JAN-75 02:25 25204

From Yngvar Lundh, Norway.

(J25204) 24-JAN-75 02:25; Title: Author(s): Yngvar Lundh/YL;
Distribution: /EPS; Sub-Collections: NIC; Clerk: YL;

JBP 24-JAN-75 10:54 25205

NSW filename syntax

I have learned the following about the syntax of NSW filenames from a phone call with Bob Millstein: 1

<filename> ::= <component-list> / <component-list> "/" <property-list> / "/" <property-list> 1a

<component-list> ::= <component> / <component> "," <component-list> / <component> "..." <component-list> 1b

<property-list> ::= <property> / <property> "," <property-list> 1c

<component> ::= <alphanum> 1d

<property> ::= <alphanum> 1e

Note: The length of an alphanum is not yet set but is likely to be in the range 9<m<32. The total number of components is not yet set but is likely to be approximately 10. 1f

JBP 24-JAN-75 10:54 25205

NSW filename syntax

(J25205) 24-JAN-75 10:54;;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /NPG([INFO-ONLY]) NSW([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG NSW; Clerk: JBP;

JBP 24-JAN-75 11:10 25206

More on printer mess

As long as people are griping about the printer situation i'll add my pet peeve, PEOPLE WHO DO PICK UP THERE OUTPUT BUT LEAVE THE BREAK PAGE! The reason that this is upsetting is it causes the preceeding listing of someone to appear to belong to other than the rightful owner. But since the offenders leave their names behind they can in the future expect to recieve nasty reminders of their thoughtlessness.
--jon.

1

JBP 24-JAN-75 11:10 25206

More on printer mess

(J25206) 24-JAN-75 11:10;;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /SRI-ARC([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: JBP;

JHB 24-JAN-75 12:06 25207

FORMAT FOR USER DEVELOPMENT TRIP AND/OR COURSE REPORTS

Subject to revision.

FORMAT FOR USER DEVELOPMENT TRIP AND/OR COURSE REPORTS 1

As our staff and user population continue to grow, it becomes more imperative that we keep each other informed about who is is getting what service, and the status of what's going on at respective sites that we contact. This form contains general headings that cover the information needed. My last report (HJOURNAL, 25151, 1;w) can serve as a model for both content and level of detail. Of course, headings can be added as well as special notes. 1a

Distribution should include UD (the ident for User Development), JCN, RLL, and DCE. You may send it to all SRI-ARC if you feel it is of general interest, however, the reports are for internal distribution only and not for clients (due to danger of misinterpretation and comprimise of certain client's privacy). 1b

Each field is optional depending on services rendered. Mandays of time spent on the particular client are important, and should be in tenths. Where continuous service is provided, eg. at ARPA, a weekly summary is more appropriate. In that case just the new information, particularly the courses taught, should be included. 1c

USER DEVELOPMENT REPORT: TRIP TO [note general area or site name] 2

1. [name of client site, eg. RADC] (# of mandays [no. of days of attention given, eg. 1.5])
[Note that some clients have more than one site, so both should listed] 2a

2. Persons (users or not) contacted [use uppercase if they have a directory] 2a1

3. COURSE: 2a2

A brief description of course outline used and the areas of NLS covered, eg. "Journal interrogate, substitute for editing". Note how far you were able to get in the time you were there, and exceptions if any. 2a2a

4. ASSISTANCE: 2a3

If a formal course was not given, or other kinds of assistance were offered in addition to a formal course, then briefly describe here. (eg. answering specific questions or helping a user thorugh an area she is having difficulty with.) 2a3a

5. APPLICATION ([note number of slots here for reference]) 2a4

Generally describe the application the particular site is evolving if it hasn't been described before, or if there are new developments, (eg, document production and remote collaboration,..) 2a4a

6. ISSUES : 2a5

Problems, obstacles, etc, that are important to the usefullness of NLS in the site environment. Note any action items here, but coordinate with action persn separately. 2a5a

7. DISCUSSIONS: 2a6

points of interest brought out in conversations with users, (eg, personnel transfers, changes in financial condition, new positions for users, etc,) 2a6a

JHB 24-JAN-75 12:06 25207

FORMAT FOR USER DEVELOPMENT TRIP AND/OR COURSE REPORTS

(J25207) 24-JAN-75 12:06;;;; Title: Author(s): James H. Bair/JHB;
Distribution: /UD([ACTION]) DCE([INFO-ONLY]) JCN([INFO-ONLY]
) RLL([INFO-ONLY]) ; Sub-Collections: SRI-ARC UD; Clerk: JHB;
Origin: < BAIR, REPORTFORM.NLS;2, >, 20-JAN-75 10:56 JHB ;;;;####;

RWW 24-JAN-75 13:30 25208

Getting to the NSW Ontime

If for any reason the NSW system is not to be delivered ontime, I am particularly concerned that we not be the cause. Our reputation for building solid useful systems and delivering as promised is our most important asset. As I mull over our present status I feel we are pushing very close to the wire. The only reasonable thing that I think we can do is to look at each of the three areas we are involved in and ask ourselves whether we are trying to deliver more than is really necessary this first year. 1

Because of our long experience in the business we can see many things that are needed more clearly than others involved and may be setting goals for ourselves with respect to the initial set of features to be delivered that are more than reasonable. Therefore I ask each team to to look at what they are doing, identify those functions that are on the critical path to having an initial system ready for integration testing in April, then ask whether in what is planned there are more bells and whistles than are really needed for an initial system and concentrate on getting those working. 2

For example, in the Frontend there many be features for fancy tool slewing, debugging, window control, L 10 fetures that can be delayed until the first version is running. 3

Similarly in NLS, we might be able to get along without the OSI initially or something else that would free manpower for other top priority tasks. 4

In the Protocol area there may also be some bells and whistles on the initial version of PCP that we could put off until later. What to do about the application level packages is out of our direct control except to push Carlson to get somebody assigned to deal with the issues. 5

Further we need to define what information we need from MCA, ADR or wherever that is critical to our work. 6

Late next week I will set up meetings with each group to review with you what tasks you feel meet the above guidelines and how we should allocate our resources. Remember that some resources are going to be needed to successfully move off our PDP 10, and others will be needed for tasks as yet unknown between now and July. Thanks Dick 7

RWW 24-JAN-75 13:30 25208

Getting to the NSW Ontime

(J25208) 24-JAN-75 13:30;;;; Title: Author(s): Richard W. Watson/RWW; Distribution: /NPG([ACTION]) JBP([ACTION]) DVN([ACTION]) POOH([ACTION]) DCE([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG; Clerk: RWW;

KEV 24-JAN-75 15:16 25209

elf status

dick, is this satisfactory? if so, one of us can distribute it more widely.

KEV 24-JAN-75 15:16 25209

elf status

This document is intended to update (24575,) which was issued 11/22/74. I expect that this document will be re-issued periodically and will use the following convention to indicate when statements are added to the file: 1

statements not preceded by any astericks were contatined in the original version of this document. 1a

statements preceded by one asterick (*) were added in the first update. 1b

statements preceded by two astericks (**) were added in the second update. 1c

etc. 1d

Asterick dates: 2

no astericks - 11/22/74 2a

*) 1/25/75 2b

The following is a list of not yet completed ELF and ELF related tasks required by SRI-ARC for its NSW work, and our understanding of the current status of these tasks. 3

The ELF KERNEL 3a

We need a TEST and a TESTS (test specific) system call so we can check for the occurence of an event without being put to sleep. 3a1

Status: 3a2

Dave Retz has indicated that it would be trivial to implement these two system calls, but has not yet gotten around to doing it. 3a2a

*) These system calls have been implemented. 3a2b

The ELF EXEC 3b

We need the ELF EXEC in a working and reliable state. 3b1

We need to get a better understanding of the relationships that exist between the ELF KERNEL, the ELF EXEC, and user processes running on ELF. Specifically, it appears that from a users point of view, some system calls are part of the KERNEL and some system calls are part of the EXEC. Since it will

eventually be necessary for us to replace the ELF EXEC with an NSW EXEC, we need to know how to separate the ELF EXEC into two parts: 3b2

that part of the EXEC that implements system calls, and 3b2a

that part of the EXEC that serves as the ELF command interpreter. 3b2b

Status: 3b3

The ELF EXEC is supposed to be fully operational by Dec. 1, and documentation on its structure has been promised, but no date set for the documentation. 3b3a

*) The ELF exec is apparently in pretty good shape, although I get the impression there is still some debugging of it going on. We still don't have an understanding of its organization and its relationship to the KERNEL. 3b3b

ELF Network Programs 3c

We need a working NCP in ELF. 3c1

We need a working TELNET in ELF. 3c2

Status: 3c3

The ELF NCP and TELNET programs are supposed to be fully operational by Dec. 1. 3c3a

*) NCP and TELNET are advertised to be in very good shape. 3c3b

ELF Virtual Memory 3d

We need the virtual memory implementation of ELF. Without this capability, only 28K of the memory on an 11 is usable. 3d1

Status: 3d2

The virtual memory features of ELF are not expected to be ready until at least Jan. 1, 1975. 3d2a

* - Virtual memory ELF is still in the debugging stage and I don't have any current promised dates. 3d2b

Loading ELF 3e

We need to be able to "boot load" ELF into an 11 from over the network. 3e1

Status: 3e2

Eric Mader of BBN is currently working on this procedure. However, his boot loading procedures appear to require the use of experimental NCP programs. I am not sure of the current state of his work with regards to completion of this task. 3e2a

*) I haven't spoken to Eric about this recently, but my impression is that BBN is currently boot loading ELF into their 11 over the Net. 3e2b

Loading User Programs 3f

We need to be able to load user processes from over the network. There appear to be several ways to do this: 3f1

1) Have a user FTP that runs on ELF that can get a remote file and store it in core (by using the Inter Process Port capabilites of ELF) rather than on a disk. This seems to be the most desirable approach. 3f1a

2) Have a server FTP that runs on ELF that can receive a remote file and store it in core (by using the Inter Process Port capabilites of ELF) rather than on a disk. In this case we would TELNET to the remote host that holds the file we wish to load and then use FTP on the remote host to send the file to ELF. 3f1b

3) Have a dedicated ELF process (a process that is part of the ELF operating system) that is always listening on a specific socket for files sent to it from a remote host. This process would then store the received file in core. This seems to be the least desirable approach in that it requires initiating action on a remote host and that the functions performed by this process are so similar to those that would be performed by a user FTP that it seems senseless to have a special separate process. 3f1c

All of these methods seem to require the pre-existance of a process that is waiting to load, via an IPP, the remote file. It would be desirable to have a (load) system call that would set up this process with the approriate address space and IPPs. The FTP server or user process could then issue this sytem call at the right time. 3f2

elf status

Status: 3f3

Full server and user FTP processes are planned for ELF, but will probably not be fully operational until Spring, 1975. It appears that we will have to write our own code for the process that will load remote files into core via IPPs. 3f3a

*) No progress has been made in this area that I am aware of. Additionally, it appears that SCRL is not interested in providing much support in this area. 3f3b

ELF Debugging 3g

We need the ELF debugging process. A debugging process, which has the ability to monitor other processes, has been designed for the ELF operating system. Our debugging plans call for the use of this process. 3g1

Status: 3g2

Eric Mader of BBN is writing and implementing the ELF debugging process. He thinks he will be finished around mid December, 1975. 3g2a

*) Eric is almost done with this. The implementation of the debugging process has pointed out some deficienies in ELF and it is my understanding that Dave and Eric have resovled most, if not all, of the problems. 3g2b

Space Allocation 3h

Given the memory limitations of an 11, it might be nice to have system buffer pool calls. 3h1

Status: 3h2

ADR agreed at the recent NSW meeting to investigate this path. 3h2a

*) No progress in this area that I am aware of. As we get further along in our 11 frontend implementation, we will be in a better position to specify exact needs. 3h2b

PCP 3i

We need the PCP routines for the implementation of the NSW. 3i1

Status: 3i2

SRI-ARC has most of the design work done and will be starting implementation soon. 3i2a

*) Same as before. 3i2b

Documentation 3j

There is a need for more documentation about ELF from both a user's point of view, and from a system programmer's point of view. 3j1

Status: 3j2

Dave Retz has plans for eventually getting around to doing all the required documentation, however, it appears that as usual in the programming world, documentation will not be available until after many of the programming tasks are completed. 3j2a

*) We are still hurting here. 3j2b

General Requirements 3k

In general we need an ELF that is reliable and bug free so we can devote ourselves to NSW task without being sidetracked into debugging of ELF. 3k1

*) It is also mandatory that the I/O system be efficient enough to support 8-16 terminals running at 1200 baud or better. 3k2

Status: 3k3

It is hard to make any statement about the reliability of a system that is not yet in full operational use. 3k3a

*) It is also hard to make any statement about the efficiency of a system that is not yet in full operational use. 3k3b

The following is our understanding of which groups have responsibility for the above tasks: 4

SCRL Tasks 4a

The ELF KERNEL 4a1

The ELF EXEC 4a2

The ELF Network Programs 4a3

It appears that the 4 programmers working on ELF are overburdened, and that they are doing the best that is humanly possible. It may be desirable to loan them an ADR person to assist in the current development of ELF. (It's possible that this loaned person could be assigned to assist in getting the needed documentation completed.) 5a

At the recent (NOV. 5-6) NSW meeting ADR indicated that it would like to freeze an NSW version of ELF, possibly as early as next month. By that time, as indicated above, many of the features needed by SRI-ARC for its tasks will not be available and therefore to freeze an NSW version of ELF at this time seems premature. 5b

*) It is still too eraly to freeze ELF. 5c

KEV 24-JAN-75 15:16 25209

elf status

(J25209) 24-JAN-75 15:16;;;; Title: Author(s): Kenneth E. (Ken)
Victor/KEV; Distribution: /RWW([ACTION]) ; Sub-Collections:
SRI-ARC; Clerk: KEV; Origin: < VICTOR, CURRENT-ELF-ISSUES.NLS;1,
>, 24-JAN-75 15:14 KEV ;;;;
####;

DVN 24-JAN-75 17:31 25210

Could You Send Frank Brignoli a Few Pages of the JOVIAL Manual?

Frank G. Brignoli is interested in publishing some exec documentation through COM and is interested in what you have been doing with the JOVIAL Manual. Would it be possible for you to send him a Xerox of a few pages of one of your drafts? He is particularily interested in changing type faces to show meta language. His address is: (FGB)

Naval Ship Research and Development Center
Code 1833
Bethesda, Maryland 20034

1

DVN 24-JAN-75 17:31 25210

Could You Send Frank Brignoli a Few Pages of the JOVIAL Manual?

(J25210) 24-JAN-75 17:31;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /DLS([ACTION]) JOAN([ACTION] dpcs notebook please) FGB([INFO-ONLY]) ; Sub-Collections: SRI-ARC DPCS RADC; Clerk: DVN;

KIRK 25-JAN-75 12:18 25211

How about some collaboration with our co-contractor

> JBP's description of Bob Millstein's syntax for NSW filenames 25205 looks like it could map directly into our current file-name scheme without having to obsolete all of our old filelinks. One change would be necessary: the use of the comma and period must be reversed. The NSW filename can be made up of a list of components separated by periods or a list of properties separated by commas. Unfortunately, the list seperated by commas must always be preceeded by a slash. This may be of such great potential hassle to the prime NSW tool (NLS) that we should find out how amenable the works manager would be to making the switch. It could be that the period and the comma are completely arbitrary and they wouldn't mind changing it considering how much conversion trouble it would save us and considering that they have yet to document the filename syntax design. I'm afraid of the duplication that would be necessary (not just in documentation) if NSW-NLS and NLS-9 are not the same. 1

KIRK 25-JAN-75 12:18 25211

How about some collaboration with our co-contractor

(J25211) 25-JAN-75 12:18;;;; Title: Author(s): Kirk E. Kelley/KIRK;
Distribution: /EKM([ACTION]) JBP([ACTION]) DSM([ACTION])
RWW([ACTION]) NPG([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG;
Clerk: KIRK;

KIRK 25-JAN-75 17:49 25212

Bug with SID's

SID's are supposed to represent statements in the order with which the user entered them into a file he/she created. Instead, they start numbering from 02 instead of 01. The logically consistant SID for statement 0 is 00. Instead it is 01. As a result, there is a very bad bug that causes the user to go into an infinite loop if he/she specifies 00 in Jump to link or Jump to Name any. Using 00 as an address in an editing command or in TNLS can also cause some very bad things to happen. 1

KIRK 25-JAN-75 17:49 25212

Bug with SID's

(J25212) 25-JAN-75 17:49;;;; Title: Author(s): Kirk E. Kelley/KIRK;
Distribution: /FDBK([ACTION]) ; Sub-Collections: SRI-ARC FDBK;
Clerk: KIRK;

KIRK 25-JAN-75 18:03 25213

Bug in Set External command

The CML takes a LINK parameter and it should take a FILELINK parameter. 1

KIRK 25-JAN-75 18:03 25213

Bug in Set External command

(J25213) 25-JAN-75 18:03;;;; Title: Author(s): Kirk E. Kelley/KIRK;
Distribution: /FDBK([ACTION]) ; Sub-Collections: SRI-ARC FDBK;
Clerk: KIRK;

DVN 25-JAN-75 21:52 25214

Phone Log: Dr. Marge Lambie of Bonneville Power Authority

Following < gjournal,25065, > I received a letter from Dr. Marge Lambie of Bonneville Power Authority. It thanks us for our respones, said they were still considering what to do (apparenty they got quite a few responses), suggested I phone, and emphasized their interest in running some system on their CDC 6500. 1

Today I telephoned. 2

Dr. Lambie gave the impression of being intelligent, knowledgeable about computers but not text handling, and of lacking time for this project. 2a

I think it was a good thing I called. She had read our response, but I think it had gotten rather lost in her mind ,and I was able to highlight NLS's special qualities. 2b

Doug had suggested he might be in the Northwest anyway and would be glad to stop and talk with her. She encouraged that and asked him to call her when he had an itinerary. (503-243-3361) 2c

I encouraged her to visit here but she said their travel budget was very tight; in lieu of that she asked to receive the movies. 2d

Again she brought up running NLS on the CDC 6500. I said it was possible but it would take study to know what effort was involved. She asked what effect NLS would have on the 6500's performance. I said we would have to know more about what the 6500 was doing to answer that. 2e

In discussing alternative accesses I mentioned the ARPA net. She had not heard of it, asked if they could get on, I said I thought it would be difficult but gave her Craig Field's name and address. 2f

DVN 25-JAN-75 21:52 25214

Phone Log: Dr. Marge Lambie of Bonneville Power Authority

(J25214) 25-JAN-75 21:52;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /JOAN([ACTION] dpcs notebook please) DPCS([INFO-ONLY]) RLL([INFO-ONLY]) ; Sub-Collections: SRI-ARC DPCS; Clerk: DVN;

KEV 26-JAN-75 15:17 25215

elf and nsw

sent via sndmsg to carlsrtom, <postel>nsw-steering, and
<postel>nsw-pi

elf and nsw

This document is intended to update (24575,) which was issued 11/22/74. I expect that this document will be re-issued periodically and will use the following convention to indicate when statements are added to the file: 1

statements not preceded by any astericks were contatined in the original version of this document. 1a

statements preceded by one asterick (*) were added in the first update. 1b

statements preceded by two astericks (**) were added in the second update. 1c

etc. 1d

Asterick dates: 2

no astericks - 11/22/74 2a

*) 1/25/75 2b

The following is a list of not yet completed ELF and ELF related tasks required by SRI-ARC for its NSW work, and our understanding of the current status of these tasks. 3

The ELF KERNEL 3a

We need a TEST and a TESTS (test specific) system call so we can check for the occurence of an event without being put to sleep. 3a1

Status: 3a2

Dave Retz has indicated that it would be trivial to implement these two system calls, but has not yet gotten around to doing it. 3a2a

*) These system calls have been implemented. 3a2b

The ELF EXEC 3b

We need the ELF EXEC in a working and reliable state. 3b1

We need to get a better understanding of the relationships that exist between the ELF KERNEL, the ELF EXEC, and user processes running on ELF. Specifically, it appears that from a users point of view, some system calls are part of the KERNEL and some system calls are part of the EXEC. Since it will

eventually be necessary for us to replace the ELF EXEC with an NSW EXEC, we need to know how to separate the ELF EXEC into two parts: 3b2

that part of the EXEC that implements system calls, and 3b2a

that part of the EXEC that serves as the ELF command interpreter. 3b2b

Status: 3b3

The ELF EXEC is supposed to be fully operational by Dec. 1, and documentation on its structure has been promised, but no date set for the documentation. 3b3a

*) The ELF exec is apparently in pretty good shape, although I get the impression there is still some debugging of it going on. We still don't have an understanding of its organization and its relationship to the KERNEL. 3b3b

ELF Network Programs 3c

We need a working NCP in ELF. 3c1

We need a working TELNET in ELF. 3c2

Status: 3c3

The ELF NCP and TELNET programs are supposed to be fully operational by Dec. 1. 3c3a

*) NCP and TELNET are advertised to be in very good shape. 3c3b

ELF Virtual Memory 3d

We need the virtual memory implementation of ELF. Without this capability, only 28K of the memory on an 11 is usable. 3d1

Status: 3d2

The virtual memory features of ELF are not expected to be ready until at least Jan. 1, 1975. 3d2a

* - Virtual memory ELF is still in the debugging stage and I don't have any current promised dates. 3d2b

Loading ELF 3e

We need to be able to "boot load" ELF into an 11 from over the network. 3e1

Status: 3e2

Eric Mader of BBN is currently working on this procedure. However, his boot loading procedures appear to require the use of experimental NCP programs. I am not sure of the current state of his work with regards to completion of this task. 3e2a

*) I haven't spoken to Eric about this recently, but my impression is that BBN is currently boot loading ELF into their 11 over the Net. 3e2b

Loading User Programs 3f

We need to be able to load user processes from over the network. There appear to be several ways to do this: 3f1

1) Have a user FTP that runs on ELF that can get a remote file and store it in core (by using the Inter Process Port capabilites of ELF) rather than on a disk. This seems to be the most desirable approach. 3f1a

2) Have a server FTP that runs on ELF that can receive a remote file and store it in core (by using the Inter Process Port capabilites of ELF) rather than on a disk. In this case we would TELNET to the remote host that holds the file we wish to load and then use FTP on the remote host to send the file to ELF. 3f1b

3) Have a dedicated ELF process (a process that is part of the ELF operating system) that is always listening on a specific socket for files sent to it from a remote host. This process would then store the received file in core. This seems to be the least desirable approach in that it requires initiating action on a remote host and that the functions performed by this process are so similar to those that would be performed by a user FTP that it seems senseless to have a special separate process. 3f1c

All of these methods seem to require the pre-existance of a process that is waiting to load, via an IPP, the remote file. It would be desirable to have a (load) system call that would set up this process with the approriate address space and IPPs. The FTP server or user process could then issue this sytem call at the right time. 3f2

Status: 3f3

Full server and user FTP processes are planned for ELF, but will probably not be fully operational until Spring, 1975. It appears that we will have to write our own code for the process that will load remote files into core via IPPs. 3f3a

*) No progress has been made in this area that I am aware of. Additionally, it appears that SCRL is not interested in providing much support in this area. 3f3b

ELF Debugging 3g

We need the ELF debugging process. A debugging process, which has the ability to monitor other processes, has been designed for the ELF operating system. Our debugging plans call for the use of this process. 3g1

Status: 3g2

Eric Mader of BBN is writing and implementing the ELF debugging process. He thinks he will be finished around mid December, 1975. 3g2a

*) Eric is almost done with this. The implementation of the debugging process has pointed out some deficienies in ELF and it is my understanding that Dave and Eric have resovled most, if not all, of the problems. 3g2b

Space Allocation 3h

Given the memory limitations of an 11, it might be nice to have system buffer pool calls. 3h1

Status: 3h2

ADR agreed at the recent NSW meeting to investigate this path. 3h2a

*) No progress in this area that I am aware of. As we get further along in our 11 frontend implementation, we will be in a better position to specify exact needs. 3h2b

PCP 3i

We need the PCP routines for the implementation of the NSW. 3i1

Status: 3i2

SRI-ARC has most of the design work done and will be starting implementation soon. 3i2a

*) Same as before. 3i2b

*) DEC IMP 11A Imp Interface 3j

*) We need the appropriapte ELF software driver fo the DEC IMP 11A Imp interface that is to be part of the NSW PDP11. 3j1

Status 3j2

*) SCRL has indicated that it would prefer NOT to get involved in this task. Perhaps this task should go to ADR. 3j2a

Documentation 3k

There is a need for more documentation about ELF from both a user's point of view, and from a system programmer's point of view. 3k1

Status: 3k2

Dave Retz has plans for eventually getting around to doing all the required documentation, however, it appears that as usual in the programming world, documentation will not be available until after many of the programming tasks are completed. 3k2a

*) We are still hurting here. 3k2b

General Requirements 3l

In general we need an ELF that is reliable and bug free so we can devote ourselves to NSW task without being sidetracked into debugging of ELF. 3l1

*) It is also mandatory that the I/O system be efficient enough to support 8-16 terminals running at 1200 baud or better. 3l2

Status: 3l3

It is hard to make any statement about the reliability of a system that is not yet in full operational use. 3l3a

*) It is also hard to make any statement about the efficiency of a system that is not yet in full operational use. 3l3b

The following is our understanding of which groups have responsibility for the above tasks: 4

SCRL Tasks 4a

The ELF KERNEL 4a1

The ELF EXEC 4a2

The ELF Network Programs 4a3

The ELF Virtual Memory Features 4a4

Documentation 4a5

SRI-ARC Tasks 4b

Loading User Programs Over the Network 4b1

We assume we have responsibiltiy for writing any user code necessary for the loading of user programs; it is not clear who has responsibilty for getting an FTP running or for getting any new system calls needed for the support of loading user programs over the network. 4b1a

PCP 4b2

ADR Tasks 4c

Memory Space Allocation 4c1

Maintainance of ELF after it is developed 4c2

*) ELF driver for DEC IMP 11A Imp Interface 4c3

BBN Tasks 4d

Loading ELF over the Network 4d1

The ELF Debugging Process 4d2

Conclusions 5

It appears that the 4 programmers working on ELF are overburdened, and that they are doing the best that is humanly possible. It may be desirable to loan them an ADR person to assist in the current development of ELF. (It's possible that this loaned person could be assigned to assist in getting the needed documentation completed.) 5a

elf and nsw

At the recent (NOV. 5-6) NSW meeting ADR indicated that it would like to freeze an NSW version of ELF, possibly as early as next month. By that time, as indicated above, many of the features needed by SRI-ARC for its tasks will not be available and therefore to freeze an NSW version of ELF at this time seems premature. 5b

*) It is still too eraly to freeze ELF. 5c

KEV 26-JAN-75 15:17 25215

elf and nsw

(J25215) 26-JAN-75 15:17;;;; Title: Author(s): Kenneth E. (Ken) Victor/KEV; Distribution: /NPG([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG; Obsoletes Document(s): 24575; Clerk: KEV; Origin: < VICTOR, CURRENT-ELF-ISSUES.NLS;1, >, 26-JAN-75 15:06 KEV ;;;; ####;

DVN 25-JAN-75 21:52 25216

Phone Log: Dr. Marge Lambie of Bonneville Power Authority

Following < gjcurnal,25065, > I received a letter from Dr. Marge Lambie of Bonneville Power Authority. It thanks us for our respones, said they were still considering what to do (apparenty they got quite a few responses), suggested I phone, and emphasized their interest in running some system on their CDC 6500. 1

Today I telephoned. 2

Dr. Lambie gave the impression of being intelligent, knowledgeable about computers but not text handling, and of lacking time for this project. 2a

I think it was a good thing I called. She had read our response, but I think it had gotten rather lost in her mind ,and I was able to highlight NLS's special qualities. 2b

Doug had suggested he might be in the Northwest anyway and would be glad to stop and talk with her. She encouraged that and asked him to call her when he had an itinerary. (503-243-3361) 2c

I encouraged her to visit here but she said their travel budget was very tight; in lieu of that she asked to receive the movies. 2d

Again she brought up running NLS on the CDC 6500. I said it was possible but it would take study to know what effort was involved. She asked what effect NLS would have on the 6500's performance. I said we would have to know more about what the 6500 was doing to answer that. 2e

In discussing alternative accesses I mentioned the ARPA net. She had not heard of it, asked if they could get on. I said I thought it would be difficult but gave her Craig Field's name and address. 2f

DVN 25-JAN-75 21:52 25216

Phone Log: Dr. Marge Lambie of Bonneville Power Authority

(J25216) 25-JAN-75 21:52;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /JOAN([ACTION] dpcs notebook please) DPCS([INFO-ONLY]) RLL([INFO-ONLY]) ; Sub-Collections: SRI-ARC DPCS; Clerk: DVN;

KEV 26-JAN-75 15:17 25217

elf and nsw

sent via sndmsg to carlsrtom, <postel>nsw-steering, and
<postel>nsw-pi

elf and nsw

This document is intended to update (24575,) which was issued 11/22/74. I expect that this document will be re-issued periodically and will use the following convention to indicate when statements are added to the file: 1

statements not preceded by any astericks were contatined in the original version of this document. 1a

statements preceded by one asterick (*) were added in the first update. 1b

statements preceded by two astericks (**) were added in the second update. 1c

etc. 1d

Asterick dates: 2

no astericks = 11/22/74 2a

*) 1/25/75 2b

The following is a list of not yet completed ELF and ELF related tasks required by SRI-ARC for its NSW work, and our understanding of the current status of these tasks. 3

The ELF KERNEL 3a

We need a TEST and a TESTS (test specific) system call so we can check for the occurence of an event without being put to sleep. 3a1

Status: 3a2

Dave Retz has indicated that it would be trivial to implement these two system calls, but has not yet gotten around to doing it. 3a2a

*) These system calls have been implemented. 3a2b

The ELF EXEC 3b

We need the ELF EXEC in a working and reliable state. 3b1

We need to get a better understanding of the relationships that exist between the ELF KERNEL, the ELF EXEC, and user processes running on ELF. Specifically, it appears that from a users point of view, some system calls are part of the KERNEL and some system calls are part of the EXEC. Since it will

eventually be necessary for us to replace the ELF EXEC with an NSW EXEC, we need to know how to separate the ELF EXEC into two parts: 3b2

that part of the EXEC that implements system calls, and 3b2a

that part of the EXEC that serves as the ELF command interpreter. 3b2b

Status: 3b3

The ELF EXEC is supposed to be fully operational by Dec. 1, and documentation on its structure has been promised, but no date set for the documentation. 3b3a

*) The ELF exec is apparently in pretty good shape, although I get the impression there is still some debugging of it going on. We still don't have an understanding of its organization and its relationship to the KERNEL. 3b3b

ELF Network Programs 3c

We need a working NCP in ELF. 3c1

We need a working TELNET in ELF. 3c2

Status: 3c3

The ELF NCP and TELNET programs are supposed to be fully operational by Dec. 1. 3c3a

*) NCP and TELNET are advertised to be in very good shape. 3c3b

ELF Virtual Memory 3d

We need the virtual memory implementation of ELF. Without this capability, only 28K of the memory on an 11 is usable. 3d1

Status: 3d2

The virtual memory features of ELF are not expected to be ready until at least Jan. 1, 1975. 3d2a

* - Virtual memory ELF is still in the debugging stage and I don't have any current promised dates. 3d2b

Loading ELF 3e

We need to be able to "boot load" ELF into an i1 from over the network. 3e1

Status: 3e2

Eric Mader of BBN is currently working on this procedure. However, his boot loading procedures appear to require the use of experimental NCP programs. I am not sure of the current state of his work with regards to completion of this task. 3e2a

*) I haven't spoken to Eric about this recently, but my impression is that BBN is currently boot loading ELF into their i1 over the Net. 3e2b

Loading User Programs 3f

We need to be able to load user processes from over the network. There appear to be several ways to do this: 3f1

1) Have a user FTP that runs on ELF that can get a remote file and store it in core (by using the Inter Process Port capabilites of ELF) rather than on a disk. This seems to be the most desirable approach. 3f1a

2) Have a server FTP that runs on ELF that can receive a remote file and store it in core (by using the Inter Process Port capabilites of ELF) rather than on a disk. In this case we would TELNET to the remote host that holds the file we wish to load and then use FTP on the remote host to send the file to ELF. 3f1b

3) Have a dedicated ELF process (a process that is part of the ELF operating system) that is always listening on a specific socket for files sent to it from a remote host. This process would then store the received file in core. This seems to be the least desirable approach in that it requires initiating action on a remote host and that the functions performed by this process are so similar to those that would be performed by a user FTP that it seems senseless to have a special separate process. 3f1c

All of these methods seem to require the pre-existance of a process that is waiting to load, via an IPP, the remote file. It would be desirable to have a (load) system call that would set up this process with the approriate address space and IPPs. The FTP server or user process could then issue this sytem call at the right time. 3f2

Status: 3f3

Full server and user FTP processes are planned for ELF, but will probably not be fully operational until Spring, 1975. It appears that we will have to write our own code for the process that will load remote files into core via IPPs. 3f3a

*) No progress has been made in this area that I am aware of. Additionally, it appears that SCRL is not interested in providing much support in this area. 3f3b

ELF Debugging 3g

We need the ELF debugging process. A debugging process, which has the ability to monitor other processes, has been designed for the ELF operating system. Our debugging plans call for the use of this process. 3g1

Status: 3g2

Eric Mader of BBN is writing and implementing the ELF debugging process. He thinks he will be finished around mid December, 1975. 3g2a

*) Eric is almost done with this. The implementation of the debugging process has pointed out some deficienies in ELF and it is my understanding that Dave and Eric have resovled most, if not all, of the problems. 3g2b

Space Allocation 3h

Given the memory limitations of an 11, it might be nice to have system buffer pool calls. 3h1

Status: 3h2

ADR agreed at the recent NSW meeting to investigate this path. 3h2a

*) No progress in this area that I am aware of. As we get further along in our 11 frontend implementation, we will be in a better position to specify exact needs. 3h2b

PCP 3i

We need the PCP routines for the implementation of the NSW. 3i1

Status: 3i2

SRI-ARC has most of the design work done and will be starting implementation soon. 3i2a

*) Same as before. 3i2b

*) DEC IMP 11A Imp Interface 3j

*) We need the appropriapte ELF software driver fo the DEC IMP 11A Imp interface that is to be part of the NSW PDP11. 3j1

Status 3j2

*) SCRL has indicated that it would prefer NOT to get involved in this task. Perhaps this task should go to ADR. 3j2a

Documentation 3k

There is a need for more documentation about ELF from both a user's point of view, and from a system programmer's point of view. 3k1

Status: 3k2

Dave Retz has plans for eventually getting around to doing all the required documentation, however, it appears that as usual in the programming world, documentation will not be available until after many of the programming tasks are completed. 3k2a

*) We are still hurting here. 3k2b

General Requirements 3l

In general we need an ELF that is reliable and bug free so we can devote ourselves to NSW task without being sidetracked into debugging of ELF. 3l1

*) It is also mandatory that the I/O system be efficient enough to support 8-16 terminals running at 1200 baud or better. 3l2

Status: 3l3

It is hard to make any statement about the reliability of a system that is not yet in full operational use. 3l3a

*) It is also hard to make any statement about the efficiency of a system that is not yet in full operational use. 3l3b

The following is our understanding of which groups have responsibility for the above tasks: 4

SCRL Tasks 4a

The ELF KERNEL 4a1

The ELF EXEC 4a2

The ELF Network Programs 4a3

The ELF Virtual Memory Features 4a4

Documentation 4a5

SRI-ARC Tasks 4b

Loading User Programs Over the Network 4b1

We assume we have responsibiltiy for writing any user code necessary for the loading of user programs; it is not clear who has responsibilty for getting an FTP running or for getting any new system calls needed for the support of loading user programs over the network. 4b1a

PCP 4b2

ADR Tasks 4c

Memory Space Allocation 4c1

Maintainance of ELF after it is developed 4c2

*) ELF driver for DEC IMP 11A Imp Interface 4c3

BBN Tasks 4d

Loading ELF over the Network 4d1

The ELF Debugging Process 4d2

Conclusions 5

It appears that the 4 programmers working on ELF are overburdened, and that they are doing the best that is humanly possible. It may be desirable to loan them an ADR person to assist in the current development of ELF. (It's possible that this loaned person could be assigned to assist in getting the needed documentation completed.) 5a

elf and nsw

At the recent (NOV. 5-6) NSW meeting ADR indicated that it would like to freeze an NSW version of ELF, possibly as early as next month. By that time, as indicated above, many of the features needed by SRI-ARC for its tasks will not be available and therefore to freeze an NSW version of ELF at this time seems premature. 5b

*) It is still too eraly to freeze ELF. 5c

KEV 26-JAN-75 15:17 25217

elf and nsw

(J25217) 26-JAN-75 15:17;;;; Title: Author(s): Kenneth E. (Ken) Victor/KEV; Distribution: /NPG([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG; Obsoletes Document(s): 24575; Clerk: KEV; Origin: < VICTOR, CURRENT-ELF-ISSUES,NLS;1, >, 26-JAN-75 15:06 KEV ;;;; ####;

JCN 28-JAN-75 05:23 25218

KWAC Meeting III in February at ARC

At the last Workshop Architects' Meeting at ARC in September, we discussed the idea of having the next meeting early in 1975. I've talked with several of you about what dates seem best. I get the feeling that some will be attending the IEEE Compcon Conference in Southern California during the week of February 23rd. Thus, we are now planning to hold the KWAC meeting at SRI starting Tuesday, February 18th and running through Friday the 21st. Note that Monday the 17th is a holiday (Washington's BD). 1

Frank Brignoli of NSRDC has offered to coordinate the formation of a tentative agenda. He will be contacting you for your suggestions. 2

Some tentative topics we would suggest are: 2a

Talks of specialized use of NLS by some architects 2a1

Potential system additions desired by KWAC collectively 2a2

Plans for further expansion of the user community 2a3

Please let us know if you will be able to attend and when you plan to arrive here (if you wish) so we can do a little planning of extra-curricular socialization via the Journal to my ident: jcn or to NORTON@SRI-ARC. 3

JCN 28-JAN-75 05:23 25218

KWAC Meeting III in February at ARC

(J25218) 28-JAN-75 05:23;;;; Title: Author(s): James C. Norton/JCN; Distribution: /KWAC([ACTION]) RA3Y([INFO-ONLY]) RLL([INFO-ONLY]) SLJ([INFO-ONLY]) ; Sub-Collections: SRI-ARC KWAC; Clerk: JCN; Origin: < NORTON, MEETING,NLS;1, >, 23-JAN-75 18:45 JCN ;;;;####;

JEW 28-JAN-75 11:40 25219

Answers to ISI PCP Questions

Processor Priority 1

I prefer at the moment, I think, to let your suggestions for dynamically variable processor priority, and call-queue depth assignment via ITDPRCS, sit in my post-implementation queue. Although I can see the possibility of their utility, they tend to make processors much fancier than I intended, and I don't think I'm willing to commit to that yet. I'd rather wait until we have some experience with a running NSW system first. 1a

Inter-Entity Synchronization 2

The event and signal subroutines defined in PCPTNXINT are indeed for intra-process synchronization (between CF-PF and PF-PF). 2a

The lock procedures defined in PMP, however, are provided for INTER-process as well as intra-process synchronization. A data store is locked by a particular processor within a particular process. Thus locking a data store for write prevents read/write attempts by other processors within the locker's process, and other processes within the tree. 2b

We may want to add a BOOLEAN argument to LCKDATA to allow locking by an entire process with free access by ALL processors, not just one. 2b1

The lack of consistency you see between the methods by which SYLOCK and LCKDATA report the successful setting of the lock disappears if you take a larger view. In BOTH cases an event is signalled if you decide to wait. In one case, the event is specified as an argument to SYLOCK; in the other, it's specified as an argument to the CALPRO procedure by which you invoked LCKDATA. 2c

It is indeed the responsibility of PKDSMN to check the lock associated with the data store to be manipulated. The system code can't possibly do it, since it knows absolutely nothing about the data stores within a user package (not even whether one of a particular name exists). I will state that fact explicitly in Version 3. 2d

A user package presumably maintains a control block that contains such things as the name and current value of each data store within the package. This control block must also contain a lock for each data store. 2d1

I also need to define a new USS that LCKDATA can call to lock/unlock a data store or at least hand me the address of the ECB associated with it. 2d2

Answers to ISI PCP Questions

Temporary Data Stores 3

Temporary data stores, since they are contained within PSP, are implemented entirely by PSP. 3a

Temporary data stores were intended for use in conjunction with CALPRO's argument- and result-list masks, providing a place where a caller could temporarily save the results of one procedure and then use them as arguments to a subsequent procedure. Furthermore, their use was only thought practical when the intermediate results were fairly large, in which case efficiences would probably result from not having to ship them back to the caller and later to the (new) callee. None of this is meant to suggest that any other use ISI may see is necessarily unreasonable, but rather just to provide some background. 3b

Consistent with the above, a temporary data store is "known" only to the process containing the procedure that made the call to CRTTMP which created it. That fact allows, for example, two inferiors of some process P to each independently create a temp with name "TEMP" without running into a name complict, which is just what one wants PROVIDED the intended application is as suggested above. 3c

The word "known" refers to who PCP will permit to reference an entity, whether it be a procedure, data store, process, etc.; rather than what portions of the user code happen to be aware that the entity has been created. 3c1

If one section of user code creates a temp whose name is chosen at random at run-time, then of course that temp is not "known" (in a DIFFERENT sense of the word, one which I never use) to other sections of the code until its name is communicated to them. However, from PCP's standpoint, the temp is known (i.e. addressable) from anywhere within the process. 3c2

From your questions, I gather you want to use temps just like one uses builtin data stores, i.e. you want them to be addressable by any process that has a PH for the process that contains the temp. We could, of course, add a BOOLEAN argument to CRTTMP, which specifies the scope of the data store you wish to create. Is there a reason that you can't simply build in the data stores you need at compile time, rather than defining them as temps at run time? 3d

I strongly resist your desire to give user code control of temps by means other than RDDATA/WRDATA. I think you need to use standard data stores, over which you already have complete control. 3e

It turns out, now that I think of it, that since you DO have complete control of the data stores in your own package, you CAN define any additional ones you need at run-time (i.e. you can implement temporary data stores in packages other than PSP). Such data stores would, however, violate the definition given in the PCP document (i.e. data stores exist throughout the life of a process). If you decide you want to do this, I would like to know about it; it may be entirely reasonable. 3f

PSI System 4

The details of system-code PSI usage aren't entirely worked out yet, but I will if possible use just one channel. I expect my usage of the PSI system to be confined mainly to implementation of SYLOCK and SYUNLK. 4a

I may need to provide some SSS's by which user code can manipulate the PSI system. I can't, short of using the JSYS trap mechanism, PREVENT their doing so directly. 4b

Processor Composition 5

Each PF will in fact contain a shared copy of the same SAV file, and thus will contain the same packages (procedures and data stores) as every other PF. Doing otherwise seems a bad idea and would be difficult to implement. 5a

I cringe at the thought of building an entire "system" as a single process. A process is a COMPONENT of a system, and a system of any size will consist of SEVERAL processes. It seems unnecessary (at least at this point) to develop disciplines for solving the address-space-size problem within a single process, when one of the powerful things about PCP is that it permits you to solve that problem by partitioning the system into SEVERAL processes. 5b

New Arguments in RSMPRO 6

The new arguments supplied to a previously-called procedure via RSMPRO bear no necessary relationship, either in number, form, or content, to the original arguments supplied via CALPRO. 6a

The arguments supplied via CALPRO are available to the caller until he makes a permanent return, while those supplied via RSMPRO are available only until he makes his next temporary return (or of course a permanent one). 6b

JEW 28-JAN-75 11:40 25219

Answers to ISI PCP Questions

(J25219) 28-JAN-75 11:40;;;; Title: Author(s): James E. (Jim) White/JEW; Distribution: /RWW([INFO-ONLY]) NPG([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG; Clerk: JEW; Origin: < WHITE, MANMSG.NLS;4, >, 28-JAN-75 11:35 JEW ;;;;####;

CHI 28-JAN-75 17:16 25220

vacation postponed

Barbara and I have decided to postpone our vacation trip to Europe originally scheduled for the month of May until this fall or next spring. We may take a week off sometime before July, but in general I will be here all the time until NSW works in July. 1

CHI 28-JAN-75 17:16 25220

vacation postponed

(J25220) 28-JAN-75 17:16;;;; Title: Author(s): Charles H. Irby/CHI;
Distribution: /SRI-ARC([INFO-ONLY]) ; Sub-Collections: SRI-ARC;
Clerk: CHI;

JBP 28-JAN-75 18:14 25221

Your changes to modify -- insert address

trying insert address after your changes i get the message
"unrecognizable interpreter op code"
--jon.

1

JBP 28-JAN-75 18:14 25221

Your changes to modify -- insert address

(J25221) 28-JAN-75 18:14;;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /KIRK([ACTION]) ; Sub-Collections: SRI-ARC; Clerk: JBP;

JBP 28-JAN-75 18:25 25222

RFCs live -- do you care ?

In the past six months the distribution of network related documents has been different that in the preceed era when the NIC was in full bloom. In particular the mechanisms for distributing that series of network memos called "requests for Comments" has broken down. There still are RFCs but they are not distributed via the NIC (or even journalized), many are available on-line tho. The distribution of RFCs
is the responsibility of the author, and when the list of Technical Liaisons (the one person at each "site" who is the network contact) is about 100 names the authors are not likely to be happy about distributing more copies. So if you really want to see RFCs as they coome out please let me know and i'll form a list to notify of any RFCs i get, and try to work out a distribution method.
--jon.

1

JBP 28-JAN-75 18:25 25222

RFCs live -- do you care ?

(J25222) 28-JAN-75 18:25;;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /SRI-ARC([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: JBP;

GSG 28-JAN-75 18:57 25223

RFC distribution list.

your network type rFc, or whatever distribution Lists so I cAn get the 1

scoop when it comes out. [Geoff] 2

------- 3

GSG 28-JAN-75 18:57 25223

RFC distribution list.

(J25223) 28-JAN-75 18:57; Title: Author(s): Geoffrey S. Goodfellow/GSG ; Distribution: / JBP; Sub-Collections: NIC; Clerk: GSG;

Pickens Thesis Proposal

20-JAN-75 1330-PST PICKENS at USC-ISIB: Warning! This is big!
Distribution: VICTOR AT SRI-ARC
Received at: 20-JAN-75 13:39:20 1

Introduction 1a

The goal of this thesis is to bring together the areas of debugging, distributed process structuring and modular programming in order to outline useful tools for dynamic debugging and monitoring. There is no reason to expect that the final thesis will conform chapter by chapter to the material here outlined, but it is presented in this manner anyway to give the writer a glimmer of hope that the end may someday be achieved. 1a1

In the first chapter the thesis surveys current work on distributed computing, modular programming, and debugging. (The conciseness of this chapter's definition is certainly misleading.) 1a2

In Chapter II a Meta Communications Description Language is defined which has among its features the following attributes: 1a3

1) One or more communications streams (or channels) may be monitored, interpreted, and correlated at a level potentially very high compared to the normally available bit and byte level. 1a3a

2) The descriptive mechanism is essentially that of a human engineered BNF in which the user-defined grammar describes in meta-rules the syntax of single channel communications. 1a3b

3) Local storage in the form of partially enumerated meta-rules as well as integer and character variables is available for user manipulation. 1a3c

4) Multiple channel communications may be correlated through the use of State-Variables and State-Changing constructs. 1a3d

5) Semantics may be associated with the Meta-rules which validates user defined conditions (e.g. value out of range). 1a3e

6) Breakpoints and breakpoint processing are associated with combinations of meta-rules and are used both in semantic validation and causing formatted output to the user or backup storage. 1a3f

7) The communications description language contains macro programming and library storage constructs to relieve the

ultimate burden on the user for creating meta-rules and validation/monitor/debug programs. 1a3g

Chapter III applies the techniques of Chapter II to several example applications and extends the discussion to cases in which more than two channels must be monitored. Three environments which come to mind are 1) various modules programmed in DCS, 2) the communications protocols for processes within the National Software Works, and 3) defining a distributed machine whose control is specified via a Graph Model of Computation (a la Gostelow). The goal of this latter example is not so much to define an efficient GMC machine as it is to demonstrate that the macro-descriptive facilities are powerful enough to enable monitoring and debugging in units higher than individual messages, e.g. token flow and vertex initiation. 1a4

Chapter IV expands the discussion of communications description to modules in general. It is shown that useful descriptive mechanisms can be defined for other than string formatted arguments. The technique of affiliating syntax and semantics checking with module activation is shown to be a practical realization of checking programs dynamically by assertions (James King's paper in Debugging Large Scale Systems). The primary requirement is that traps may be placed at the level of communications between modules. This dynamic binding of traps and monitoring is shown to be far superior to the extant techniques of compiling in monitor instructions. Chapter IV also discusses the use of emulation to check the module's stored behavior. 1a5

Chapter V raises the issues associated with distributed processes. The details are not known yet, but questions as the following might be discussed: 1) Can the monitor be distributed itself? 2) What happens to the notion of time? 3) What extra facilities may be required because of the loose coupling of the processes? 3) At what level might distributed operating systems be monitored? 1a6

Chapter VI discusses how the previous proposed tools might be integrated with more conventional tools. For example, distributed DDT (and its high level language counterparts) should exist to allow evaluation of modules at the sub-module level. Issues exist as to how this function might be distributed and what kind of an interface it should have to the user (ideally the interface should be near the level of the source language the module was written in). The transition between inter-module monitoring and sub-module DDTing should be smooth. It should also be possible,for example, to activate

sub-module breakpoints as a function of inter-module communications conditions. 1a7

Somewhere in the thesis cost will have to be discussed. The traditional tradeoff between monitoring/debugging and free-running execution is speed. The ideal debugging system is one which implies zero overhead on the operating programs. 1a8

Naturally this issue exists for the techniques proposed in this thesis. It should be determined if there are any environments in which the overhead is negligible. In some cases it may be possible to keep up with real time (such as in a hardware meta-message composer and analyzer). In others it will be necessary to impose the constraint that the intercommunications between modules be asynchronous and amenable to slowed message rates. 1a9

Introduction 1b

One very basic tool for debugging processes which communicate via messages is a Communications Description Language. As defined here the language is interactive, programmable, and capable of expressing most interprocess communications. It's form is that of a BNF modified to include dimensional data declarations and state validation/changing information. The fundamental programmable elements of the language are the grammar declarations (GD) which contain dimensional, type, and meta-linguistic declarations, and the partial enumerations (PE's) which affiliate symbolic labels with partial enumerations for meta-rules defined within the GD. 1b1

The language is meant to be used both to synthesize and to analyze communications streams. For synthesis a structured editor coupled with the PE's and GD allows flexible construction of both partial and complete messages. For analysis the PE's and GD allow basic syntax checking and breakpoint synchronization on incoming data. Additional semantics routines allow more sophisticated analysis of the communications as well as primitive emulation mechanisms. 1b2

The Meta Communications Description Language contains several features which won't be described in detail as of yet. Summarized they are: 1b3

1) Local storage is available for character strings and integer variables. 1b3a

2) Macro programming facilities and backup library storage exist. 1b3b

3) Language constructs exist to control the setting of breakpoints, breakpoint processing, and output formatting. 1b3c

4) Semantics processing may be associated with both the GD's and PE's. 1b3d

Grammar Declarations 1c

The root of the entire communication description language is in the declarations. The Grammar Declarations portion is composed of the following elements: 1c1

1) A set of BNF Meta-rules and Meta Symbols to define the structure of the communication strings. No particular enumerations exist within the Grammer Declarations (i.e. no terminals). 1c1a

2) Dimensional declarations. Since most communications by their nature are either fixed length fields or variable length fields (derivable from another fixed field) it is necessary to declare (where appropriate) the dimensions of the non-terminal symbols. The two types of dimensional declarations are fixed (e.g. n Bytes, n Words) and variable (e.g. fixed length "length" field followed by variable length "data" fields). 1c1b

3) Type declarations. In some cases it may be desirable to associate data types with non-terminals. This feature would probably be most useful for formatted output but might also be useful in writing programs to manipulate the messages. 1c1c

4) State Information. The entities which correlate multiple channels are state variables. A channel's state may be required in some cases to correctly interpret the syntax of its messages. The two types of state veriable constructs in the Grammar Declarations are a) current state and b) next state. Current state is used as part of the syntactic recognition of messages. Next state is the new state achieved after all processing associated with the metarule is completed. For now all state variables are global. Eventually it may desirable to define sub-states or states local to individual channels. 1c1d

5) Special recognition metasymbols (i.e. <<metasymbol>>) This construct allows the message parser to differentiate those fields which are required for syntactic recognition from those which aren't. In a particular metarule only a few fields (if more than one) contribute to the recognition process. Generally the most that is required is an opcode

field. The other fields within the metarule identify syntactic units which are not required for recognition and, therefore, need not be qualified by scanning partial enumerations (unless a particular message is being searched for). By introducing special metasymbols (symbols surrounded by an extra pair of brackets, see the example above) the recognizer which processes message streams can differentiate between invalid messages (no matching opcode) and unnumerated messages. (See the enumerations for <MSG> in example 1, Appendix C.) 1c1e

6) Constructs for iteration (see example 3 in Appendix C). In order to describe variable format fields (i.e. keyword positioning as opposed to fixed field positioning) iteration is required. With iteration it should be possible to describe any structure representable by recursion in a form well suited for human consumption. Recursion is ruled out as it tends to be incomprehensible for casual consumption, and one of the goals is that the description language be interactive and well human engineered. 1c1f

The precise syntax associated with the Grammar Declarations is stated in Appendix A. 1c2

Partial Enumerations 1d

All defined terminal symbols associated with the Grammar Declarations are contained in the Partial Enumerations. The PE structure is defined to be general enough that it may be used both for synthesis and analysis and may contain only partially defined rules. Every non-terminal symbol in the GD has potential entries within the PE. 1d1

Each partial enumeration contains as minimal information the non-terminal which it enumerates, a symbolic identifier to identify this enumeration, and a value, i.e. the partially enumerated right side of the referenced meta-rule. [To enumerate is simply to replace non-terminal symbols with specific strings. In a general meta-rule not all of the referenced non-terminals need be enumerated.] Two other items associated with a particular enumeration are: 1d2

1) An indication of the enumeration's environment. Some enumerations may be applicable for synthesis, some for analysis. 1d2a

2) Context qualification. In some cases referral to a meta rule will be ambiguous unless higher levels of the parsed rule are delineated. For example, a meta-rule which

describes a filename may be used in more than one type of communication string, and it may be desired to define non intersecting groups of enumerations based on the category of message. In such cases qualification is required to identify the higher level meta-rules which uniquely specify the desired context (See Appendix B). 1d2b

In general the value of a particular metarule enumeration is a mixture of metasymbols extant in its GD definition, literial entries (e.g. numeric or alphameric), symbolic labels for other partial enumerations, and , significantly, descendant metarules as defined within the GD. As an illustration of the latter consider the following set of rules (see also the examples in Appendix C): 1d3

<A> ::= <B> <C> <C> ::= <D> ::= <E> 1d4

An enumeration for <A> might be; 1d5

Enumeration-A1 = <B> <D> 1d6

Or, another might be; 1d7

Enumeration-A2 = <B> <E> 1d8

Appendix B contains more detailed specifications of partial enumerations. 1d9

Editing 1e

Rule construction and editing is highly structured. In defining partial enumerations, for example, the user is given not a sequential string editor, but rather a structured editor. His options for selection, composition, and replacement are in terms of the defined meta rules. For example, to enumerate a given metarule the user might type something like ENUMERATE <metarule x> RETURN. In response he would see the metarule's unenumerated definition. At this point he would have several editing options: e.g. [SPACE] skips to the next meta symbol within the rule, [?] lists the current enumerations for the rule, [LF] jumps one level lower into the definition of the current metasymbol, a numeric selects one of the displayed enumerations for editing, [R] begins a literal replacement of the current metasymbol, etc. The exact definition of the editor is not important at this stage, but it is important to point out that the structured presentation of the message syntax to the user is one of the more useful features of the communications description language. 1e1

Appendix A -- Syntax of Grammar Declarations 1f

General Comments 1f1

Grammar declarations are BNF-like meta-rules which are composed of meta-symbols, data declarations, and state information. Each declaration is composed of a single Left Side meta-symbol (i.e. context free), followed by the assignment symbol "::=", followed by one or more meta-symbols or, in the case that the meta-rule is at the lowest syntactic level, a group of data declaration keywords. Metasymbols are denoted by surrounding valid identifiers (definition left vague for now) with brackets, three examples are: 1f1a

<MESSAGE>, <OPERAND>, <LENGTH> 1f1b

Multiple alternatives for metarules are indicated by writing each alternative on a new line in the same format as the original metarule, but minus the metarule head symbol. An example follows: 1f1c

<A> ::= <B> ::= <C> 1f1d

Data Declarations 1f2

Metarules which are nearest the message data fields are termed data declaration metarules. The form of a data declaration metarule is as follows: 1f2a

<metasymbol> ::= "type" "number of units" "dimensional unit" 1f2b

"type" is optional, but the next two fields are required. Some examples of possible data types are INTEGER, HEX, CHARACTER, and ANY which denote integer, hexidecimal, character, and "don't care" data respectively. 1f2c

The next field can be either a fixed numeric or a simple expression. In case of expressions the appearance of a metasymbol implies that the symbol's value is to be computed from the current string (it's type must be integer). Addition, subtraction, multiplication and division are supported. An example of a variable length value is as follows: 1f2d

<FILENAME> ::= CHARACTER (<LENGTH>-2) BYTES 1f2e

The final field identifies the units in which the length is

measured. Some possible units are BITS, HEXDIGITS, BYTES, and WORDS. 1f2f

State Information 1f3

State validation and state changing information may also be associated with a metarule. The construct ";STATE=" requires that the current state be identical to the one indicated or else recognition will not occur. Logical conditions on one or more state variable are allowed. The construct ";NEWSTATE=" changes the state to the new state indicated after processing for this metarule has completed. Both constructs are optional and would probably only be used to coordinate the actions of multiple communications channels. A simple example follows: 1f3a

<INQUIRY> ::= <QUERY> <A> ;NEWSTATE = WAIT-A ::= <QUERY> <B> ;NEWSTATE = WAIT-B 1f3b

<RESPONSE-A> ::= <ANSWER-STRING> ;STATE = WAIT-A
<RESPONSE-B> ::= <ANSWER-STRING> ;STATE = WAIT-B 1f3c

Iteration 1f4

Iteration may be used to describe multiple instances of a metarule. An iteration may be defined either as having upper and lower bounds or as occuring an exact number of times. In either case the iteration specification immediately follows the metasymbol. Numeric fields may be any valid numeric (including evaluated metasymbols). The form of the bounded iteration is: 1f4a

<metasymbol> (lower,upper) 1f4b

he form of the fixed iteration is: 1f4c

<metasymbol> ( ) 1f4d

See example 3, Appendix C for an illustration of how iteration might be applied. 1f5

Appendix B -- Syntax of Partial Enumerations 1g

Partial enumerations depend upon metarule declarations for their basic structure, but in fact may appear quite different from their unenumerated definition. They are distinguished from grammar declarations by the detail which they contain. Partial enumerations are really partial parse trees. The head node (which has a symbolic name) and first level descendants

correspond to the left and right sides respectively of the original grammar declaration. But the first level descendants themselves have descendants. The most correct way to display a partial enumeration would be to actually show the partially enumerated parse tree. In most cases, however, it is satisfactory to show only the terminal nodes of the partial enumeration (See the example given above in the original discussion of partial enumerations). 1g1

The left side of an enumeration consists first of one or more metasymbols defining the context of the enumeration and second of a symbolic identifier. Normally only one metasymbol is denoted on the left side and it identifies the metarule which is being enumerated. In cases where more detailed context qualification is required the format is simply to list from left to right with the most distantly related metasymbol on the left all the required ancestors of this metarule. The general form for the left side of a partial enumeration is therefore; 1g2

< >...<g-parent>,<parent>,<this metarule> Symbolic Name = 1g3

One other item which may appear on the left side is an indication that the eumeration is to be used for analysis (:ANALYSIS) or synthesis (:SYNTHESIS , Default is both). This indication is useful in separating enumerations which are required only for synthesis of messages from those used only in analysis of messages. 1g4

As indicated previously, the right side contains enumerations for one or more of the metasymbols in the original metarule. An enumeration for a metasymbol is either symbolic (i.e. referring to the left side symbol of another partial enumeration) or literal (i.e. either the right side from another partial enumeration or actual message characters). Since the right side is really a partially enumerated parse tree it is possible to have an arbitrary mixture of symbolics and literals for each enumerated metasymbol. 1g5

Appendix C -- Examples 1h

Examples 1 and 2 demonstrate two ways to define the same grammar. Example 3 describes a keyword oriented message and in the process, illustrates an application for iteration. In all three examples square brackets, [ ] ,are used to denote partial enumerations. As the GD metarule which the PE applies to is implied by position (i.e. the PE is immediately below the corresponding GD metarule) a further simplification is effected by not listing the GD metasymbol on the left side. Partial

enumerations are listed with the grammar declarations for convenience only. 1h1

Comment on First Two Examples 1h2

The message grammar which is described by the first two examples is quite simple. A message in this grammar consists of a 1 byte length field followed by a one byte opcode field followed by a variable length operand field. The two specific messages defined are OPENFILE (opcode = 1, operand = filename) and READFILE (opcode = 2, operand = disk address). Examples 1 and 2 both describe the same grammar, but example 1 is more consise and easier to read. The first example conforms more closely to the fields of the message with all higher level interpretations appearing in the partial enumerations. The second example attempts to bring the OPENFILE and READFILE groupings into the grammar declarations. The first method appears to be preferable from the point of view of simplicity, though a final judgement cannot be made at this time. 1h3

Example 1 1h4

<MSG> ::= <LEN> <<OpCODE>> <OPERAND> 1h4a

[OPENFILE = <LEN> 01 <FILENAME>] [READFILE = <LEN> 02 <DISKADDRESS>] 1h4b

<LEN> ::= INTEGER 1 BYTE <OPCODE> ::= INTEGER 1 BYTE <OPERAND> ::= <FILENAME> <DISKADDRESS> <FILENAME> ::= CHARACTER (<LEN> - 2) BYTES <DISKADDRESS> ::= INTEGER 2 BYTES 1h4c

Example 2 1h5

<MSG> ::= <OPENFILE> ::= <READFILE> <OPENFILE> ::= <LEN> <<OPCODE>> <FILENAME> <READFILE> ::= <LEN> <<OPCODE>> <DISKADDRESS> <OPCODE> ::= INTEGER 1 BYTE 1h5a

[<OPENFILE>,<OPCODE> OPENOP = 01] [<READFILE>,<OPCODE> READOP = 02] <LEN> ::= INTEGER 1 BYTE <FILENAME> ::= CHARACTER (<LEN> - 2) BYTES <DISKADDRESS> ::= INTEGER 2 BYTES 1h5b

Example 3 1h6

The message grammar described by this example is also fairly simple. A message may consist of from 0 to 3 instances of either of two types of fields. No order is imposed and the fields differ in length and type. Each field type is

distinguished by a single byte key. Iteration is required to represent this grammar. It is assumed that some mechanism exists for determining the beginning and end of a message. The definition follows: 1h6a

<MSG> ::= <UNIT> (0,3) <UNIT> ::= <KEY> <DATA> 1h6b

[UNIT1 = KEY1 <FIELD-TYPE-1>] [UNIT2 = KEY2 <FIELD-TYPE-2>] 1h6c

<KEY> ::= INTEGER 1 BYTE 1h6d

[KEY1 = 00] [KEY2 = 01] 1h6e

<DATA> ::= <FIELD-TYPE-1> ::= <FIELD-TYPE-2> 1h6f

Appendix D -- Metadescription of DCS I/O Handler 1i

This appendix gives a fairly detailed definition of the message interface to the UCI DCS I/O Handler (IOH). A more detailed specification for IOH may be found in "DISTRIBUTED COMPUTER OPERATING SYSTEM, Programming Guide" (Internal document to the Dept. of Information and Computer Science at UC Irvine). However, it is suspected that the following definition is more readable as to the IOH message structure. Due to the requirement of conciseness only three IOH commands are defined, but they are chosen to illustrate a wide range of message formats allowed within the IOH. Following the definition of the IOH message grammar is an example of a completely parsed message showing the derived parse tree. 1i1

<MSG> ::= <<DPN>> <<OPN>> <LEN> <BODY> 1i2

[IOH-REQUEST = IOH <OPN> <LEN> <IO-REQUEST>] [IOH-RESPONSE = <DPN> IOH <LEN> <IO-RESPONSE>] 1i3

<DPN> ::= <PN> <OPN> ::= <PN> <PN> ::= <CLASS> <MACHINE> <SEQUENCE> 1i4

[NUCLEUS = 1 <MACHINE> 01] [IOH = 1 <MACHINE> 02] [CP = 1 <MACHINE> 03] [SEQUENCE = 1 <MACHINE> 04] [CHECKER = 1 <MACHINE> 05] 1i5

<CLASS> ::= INTEGER 1 HEXDIGIT <MACHINE> ::= INTEGER 1 HEXDIGIT <SEQUENCE> ::= INTEGER 1 BYTE 1i6

<LEN> ::= INTEGER 2 BYTES 1i7

<BODY> ::= <IO-REQUEST> ::= <IO-RESPONSE> ::= ANY (<LEN>) BYTES 1i8

```
<IO-REQUEST> ::= <LFN> <<OPCODE>> <ARGUMENT>                                    119

[ READ-SYM = <LFN> OPC-RS <IO-COUNT> ;NEWSTATE=READ] [
OPEN-FILE = <LFN> OPC-OF <OPEN-FLAGS> <FREE-FORM>
;NEWSTATE=OPEN] [ CLOSE-FILE= <LFN> OPC-CF <CLOSE-FLAGS>
;NEWSTATE=CLOSE] . . . ETC.                                                    1110

<LFN> ::= INTEGER 1 BYTE                                                       1111

<OPCODE> ::= INTEGER 1 BYTE                                                    1112

[ OPC-RS = 00 ] [ OPC-OF = 14 ] [ OPC-CF = 25 ] . . . ETC.                     1113

<ARGUMENT> ::= <IO-COUNT> ::= <OPEN-FLAGS> <FREE-FORM> ::=
<CLOSE-FLAGS>                                                                  1114

<IO-COUNT> ::= INTEGER 2 BYTES                                                 1115

<OPEN-FLAGS> ::= HEX 2 BYTES                                                   1116

[ OCORD = 00 ] [ OCOWR = 01 ] [ OCOWN = 02 ] . . . ETC.                        1117

<CLOSE-FLAGS> ::= INTEGER 2                                                    1118

[ CLOSE-EOF = 0] [ CLOSE-NO-EOF = 1]                                           1119

<FREE-FORM> ::= CHARACTER (<LEN> - 4) BYTES                                    1120

<IO-RESPONSE> ::= <LFN> <CC> <VALUE> ;NEWSTATE=IDLE                            1121

[ ;STATE=READ READ-REPLY = <LFN> <CC> (ANY (<LEN>-2) BYTES ] [
;STATE=OPEN OPEN-REPLY = <LFN> <CC> <OPEN-FLAGS> <FIXED-FORM>]
[ ;STATE=CLOSE CLOSE-REPLY= <LFN> <CC> NULL ] . . . ETC.                        1122

<CC> ::= INTEGER 1 BYTE                                                        1123

[ ECNRM =0 ] [ ECDNR =1 ;STATE,NOT="RESERVE,OR,OPEN,OR,TRANSFER
] . . . ETC.                                                                   1124

<VALUE> ::= ANY (<LEN> -2) BYTES ::= <OPEN-FLAGS> <FIXED-FORM>
::= NULL                                                                       1125

<FIXED-FORM> ::= <DEV-NAME> <FILE-NAME> <FILE-EXTENSION>
<DIR-NAME> ANY 2 BYTES <WILD-MASK> <RFN> ANY 5 BYTES                           1126

<DEV-NAME> ::= CHARACTER 6 BYTES                                               1127

[ DCS ::= 'DCS' ] [ DSKAON ::= 'DSKA0' CHARACTER 1 BYTE ] [
LPT0 ::= 'LPT0' ] . . . ETC.                                                   1128
```

```
<FILE-NAME> ::= CHARACTER 6 BYTES <FILE-EXTENSION> ::=
CHARACTER 2 BYTES <DIR-NAME> ::= CHARACTER 6 BYTES <WILD-MASK>
::= CHARACTER 2 BYTES <RFN> ::= INTEGER 1 BYTE                    1129
```

JBP 28-JAN-75 20:01 25224

Pickens Thesis Proposal

(J25224) 28-JAN-75 20:01;;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /KEV([INFO-ONLY]) JEW([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: JBP;

KIRK 28-JAN-75 21:49 25225

VISITLOG: SRI Electronics Explorer Scout Post

Tuesday Jan 28, 1975 8:00-9:00

KIRK 28-JAN-75 21:49 25225

VISITLOG: SRI Electronics Explorer Scout Post

On Tuesday 28 JAN, the SRI Electronics Explorer Post came to visit ARC. They were given a demonstration of the system and Mouse-Keyset code cards. After the demo, we adjurned to see the new line-processor terminal. Considerable interest was shown in using NLS but unfortunately, there were other things on the meeting agenda and the Explorers re-convened their meeting in Conference room A of Building 1. Perhaps in the future, we can have them back with enough ARC man-power to observe novices learing how to use NLS. Potentually valuable feedback particularly in use of the Help system ... I will keep in touch with Don Limuti. The following people attended: 1

Jim Schrempp 2

Jeff Frye 3

Bob Weatherford 4

Mark Bondy 5

Chris Parkinson 6

Larry Abbott 7

Don Limuti 8

Darrell Silkensen 9

Mark Merrill 10

Frank Merrill 11

Mike Campi 12

Bruce Bullard 13

Dan Mendez 14

Henri van Wandelen 15

Clark Martin 16

Tom Anderson 17

Fred Funk 18

Bob Trick 19

John McCammon 20

VISITLOG: SRI Electronics Explorer Scout Post

Sam Thomas	21
Ross Harden	22
Ron Harden	23
Wayne McNinch	24

KIRK 28-JAN-75 21:49 25225

VISITLOG: SRI Electronics Explorer Scout Post

(J25225) 28-JAN-75 21:49;;;; Title: Author(s): Kirk E. Kelley/KIRK;
Sub-Collections: NIC; Clerk: KIRK;

RA3Y 29-JAN-75 09:15 25236

accounts

Dean and I have gone over his program. I am now setting thingss up with Jeff, so that when Tymshare reports come in, I will be able to work them. Per your instructions, I have not contacted Tymshare about the accounts changes yet. We should do that ASAP, however. 1

accounts

(J25236) 29-JAN-75 09:15;;;; Title: Author(s): Raymond R. Panko/RA3Y; Distribution: /RA3Y([ACTION]) JCN([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: RA3Y;

POOH 29-JAN-75 10:07 25237

test again

this is the tird test to see if a sndmsg item can be sent via nls sendmail or whatever it is suppoed to do.

POOH 29-JAN-75 10:07 25237

test again

(J25237) 29-JAN-75 10:07; Title: Author(s): Ann Weinberg/POOH;
Distribution: /POOH; Sub-Collections: SRI-ARC; Clerk: POOH;

POOH 29-JAN-75 10:07 25238

test again

this is the tird test to see if a sndmsg item can be sent via nls
sendmail or whatever it is suppoed to do. 1

------- 2

3

POOH 29-JAN-75 10:07 25238

test again

(J25238) 29-JAN-75 10:07; Title: Author(s): Ann Weinberg/POOH;
Distribution: /KIRK; Sub-Collections: SRI-ARC; Clerk: POOH;

ISI Request for Changes in Text Searches

28-JAN-75 1224-PST ROTHENBERG at USC-ISIB: Modification to LOOKUP ?
Distribution: NLS-HELP:
Received at: 28-JAN-75 17:44:48 1

I've sent a couple of notes asking about modifying LOOKUP for doing case-independent searching. I'm not sure what the disposition of that is, but I thought I'd see how it was going, and also mention that (as may or may not have been apparent) what we really need is a switch which allows selecting EITHER case-independent or (as presently done) case-dependent searching. 1a

Jeff 1b

Modification to lookup 2

Jeff--- 2a

I have modified lookup (as a user program procedure to replace the one which is in the running system) to take an additional parameter to specify case-dependent or independent searches. It works, but is not incorporated in the running system. 2b

I have also specified an addition to the L10 language which would permit case independent searching in FIND statements and also in expressions of the form "IF *string1* != *string2*". (The exclamation before the Boolean operand in exps or before the object of a FIND as in FIND [!"text"] indicates case independent searching.) This is obviously better for your purposes than a modification in lookup using the current L10 as it would permit case independent content analyzer patterns. I have changed the library functions for L10 which do these sorts of searches, again in user programs, but have not changed the Tree Meta compiler, though I have spoken to Don Andrews about it. 2c

I have temporarily put this work aside in order to push out the new File System and expect to get back to it in about two weeks. 2d

Harvey 2e

HGL 29-JAN-75 10:19 25239

ISI Request for Changes in Text Searches

(J25239) 29-JAN-75 10:19;;;; Title: Author(s): Harvey G. Lehtman/HGL; Distribution: /RWW([INFO-ONLY]) NPG([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG; Clerk: HGL;