~~AIP File Copy~~

NASA Technical Memorandum 102837

7N-61
175264
P-39

Investigation of an Automated Cray User Services System (Cray User Services "Expert System")

George P. Callas, Michael Younkin, and
Catherine H. Schulbach

(NASA-TM-102837) INVESTIGATION OF AN AUTOMATED CRAY USER SERVICES SYSTEM (CRAY USER SERVICES EXPERT SYSTEM) (NASA) 39 p

N93-72466

Unclas

Z9/61 0175264

October 1990

National Aeronautics and
Space Administration

NASA Technical Memorandum 102837

Investigation of an Automated Cray User Services System (Cray User Services "Expert System")

George P. Callas, Ames Research Center, Moffett Field, California
Michael Younkin, Control Data Corporation, Minneapolis, Minnesota
Catherine H. Schulbach, Ames Research Center, Moffett Field, California

October 1990

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035-1000

SUMMARY

An experimental expert Cray user help system, which incorporates a natural language interface and a commercially available expert-system shell, was developed at NASA Ames Research Center to investigate the feasibility of augmenting the existing user help system. In this report the system is described and the experimental results are presented. A program describing staffing and funding requirements to develop a production system is outlined, and the results of a benefits analysis are presented to assess the potential of an expert Cray user help system.

INTRODUCTION

Supercomputer growth in the decade following the installation of the first Cray 1 in 1976 was nearly exponential, with new installations rapidly approaching 100 per year (fig. 1). In the first decade, Cray Research, Inc., dominated the market and Control Data Corporation (CDC) provided the only competition. In 1983 three Japanese manufacturers—FUJITSU, HITACHI, and NEC—entered the supercomputer race, and to strengthen its market position, CDC formed a supercomputer division, ETA Systems, Inc. By 1988, the worldwide installed base exceeded 360 supercomputers. Sales of these machines (each of which costs between $5 million and $25 million) have increased more than 25% a year during the past decade and were projected to surpass $1 billion a year by 1988 (ref. 1).

Moreover, the number of supercomputer users has exceeded the growth in supercomputers. This expanding growth may be likened to an iceberg where today's requirements represent the tip and the latent demand represents the base. This demand is driven by the virtual explosion that has occurred in the use of personal computers (PCs) and minicomputers throughout all industry segments during the same time period. Each new application creates a demand for more computer time or power until the existing computer system becomes overloaded. There also is a tendency for each new function a computer provides to quickly become a requirement, which also grows, and eventually a second or "new and bigger" computer is needed. Because the trend in the price/performance ratio favors supercomputers, economics will ultimately shift computer demand toward supercomputers.

More computer power is also fostered by network activity and its growth throughout the industry. Data traffic over telephone lines is growing at an annual rate of 20% (ref. 2). A recent study among members of the Fortune 1000 group projected more than a doubling in the number of T1 systems (1.544 megabits/sec) in use by participating members by 1992 (ref. 3). The top 100 value-added resellers of local area networks (LANs) sustained a growth rate of more than 50% a year since 1985, and in 1988, the group installed over 9500 LANs (ref. 4). A national high-speed (3 billion bits/sec) broadband network is being planned for accessing supercomputers (Senator Gore introduced Senate Bill S1067, The National High Performance Computer Technology Act in 1989). Recently the interagency Federal Research Internet Coordinating Committee became the National Network Council whose charter is to oversee the development of the National Research and Education Network, the high-speed (gigabit) nationwide fiber-optic network for the research community.

As networks make more information available to a greater number of people, even greater processing power will be required to manipulate the massive amounts of data.

To fulfill the demand for processing power, computers in general must become more accessible, or "user friendly," because many of the potential users will not have the experience gained by those who "grew" with computers. Furthermore, as computer science matures, users expect computers to become easier to use. Thus, research and development in areas that can help make computers user friendly is both important and timely.

Recently, the most innovative user friendly developments have emerged for personal computers. Although not surprising, because ease of use is important to the sale of PCs, the cost benefit associated with their efficient use is small compared to that associated with supercomputers because of the relative operating cost (cost per hour). Because computer time is valuable, it is both reasonable and prudent to support the development of user friendly elements for the larger, more expensive computers and, in particular, the supercomputers.

Development of an "expert user help system" is of interest because it has the potential of providing more service at less or equal cost by combining two important functions: an on-line help system, consisting of a library of on-line user manuals; and the user-services system, consisting of one or more computer experts, generally accessible to the users only by telephone and during prime time. An on-line expert user help system would provide around-the-clock service and has the potential of reducing the workload of the existing user-services group.

The evolution and success of expert system shells and natural language parsers as tools for expert system development, and the availability of these tools for a broader range of computers—PCs, VAXes, and mainframes—makes the development and maintenance of an expert user help system conceivable and potentially cost effective.

These facts, coupled with activities and developments at the Ames Research Center (ARC), motivated the research program described in this report. These activities and developments include (1) the expanded use of supercomputers at NASA and particularly at the Numerical Aerodynamics Simulation facility; (2) the user feedback obtained from user-services activities associated with the supercomputers and the complex computer network; (3) the emergence of the supercomputer as the principal tool for a major portion of research at Ames; and (4) the creation of a new Information Sciences Office chartered to explore applications for Artificial Intelligence.

EXPERT CRAY USER HELP SYSTEM PROGRAM

Program Objectives

Objectives of the Expert Cray User Help System (ECUHS) program are to (1) develop a framework for a Cray user help system using the concepts developed for expert systems; (2) implement a preliminary, but limited automated user help system for the Cray to demonstrate the concept and the utility of an automated Cray User Help System; and (3) draft a preliminary program plan that

outlines probable costs, key participants, staffing, and a schedule for developing an expert Cray User Help System that can be installed at any Cray site.

Background

Historically, the high cost of supercomputers has created a unique user scenario where access to the supercomputer is limited through high-speed gateways consisting of either a computer or a high-speed workstation. Large complex networks are not uncommon because a supercomputer can service a large user base.

The ARC network for the Advanced Computer Facility (ACF), used as a model for this research program, is representative of today's complex networks. At the time of this study, the ACF had approximately 700 research users on more than 60 computers that were interconnected through the ARC network to a Cray X-MP/48 supercomputer (which was subsequently upgraded to a Cray Y-MP/832). (A simplified schematic of the ARC network at the time the work was done is shown in fig. 2.) The major link to the Cray was the Network Systems Corporation Hyperchannel, which was accessed by five gateway computers: the four VAXes named IO, EULer, CERes, and EARth and the CDC Cyber 810. Other networked VAXes, MicroVAXes, and PDP-11-series computers were linked to the gateways either through local Ethernets, the CI network, or the CDC LCN network. Although not shown in figure 2, off-site access was also provided, either by high-speed satellite links through the Cyber 810, or by high-speed telephone lines through a host VAX.

It is important to note that any of the remote computers can communicate with the Cray and that each computer can communicate with any other computer on the network. This is especially important in establishing the topology of the expert user help system.

Approach

Skill levels of research users of the ARC ACF vary from expert to novice and, although an expert user help system can be useful to an expert, the system developed for this program is targeted for the intermediate and novice users. ACF user services statistics for 1 year were analyzed and candidate user questions were categorized to establish the framework of the proposed ECUHS. Because an important program objective was to demonstrate the concept by using the ARC ACF, a VAX located on the ARC network was selected for the implementation of the demonstration system.

Other system guidelines were developed: the user-system interface would require an interactive terminal (preferably the one on which the user submits jobs to the supercomputer); communication with the ECUHS would incorporate natural language queries; the protocol required to invoke the ECUHS would be the simple call "Crayexpert"; the use of help menus would be minimized; and the system response to queries would be limited where possible to one terminal screen.

These guidelines led to the specification of VAX-compatible development tools such as an expert system shell, natural language parser, and relational data base. Although the initial intent was to procure these as separate tools and perform the required integration during system development,

an integrated expert system shell, containing all of the required elements, was discovered and procured. This not only resulted in a significant procurement cost savings, but it also reduced the system development time by several months.

Since the expert system was conceived as an evolving entity that was expected to grow with system use, work to develop the knowledge base began when the program was conceived and continued throughout the program's life. This required the development of system management software to capture user queries, expand the knowledge base, and record performance parameters for system upgrades and improvement.

System Description

A block diagram of the expert system developed for this program is shown in figure 3. The expert system shell, called Guru, was procured from Micro Database Systems, Inc., and consists of a suite of *integrated software* that includes a natural language interface, an expert system shell, a relational data base, a manager with a multiple index processor, and application software including a word processor and spreadsheet among others (shown as the blocks at the bottom of fig. 3). Guru also provides "hooks" for incorporating other programs, procedures, or subroutines. The dashed lines and arrows indicate bi-directional flow illustrating the extent of the software integration. For example, either a procedure or a spreadsheet entry can call the expert system. Moreover, the expert system can call a procedure, execute a spreadsheet, add to or extract from a data base, or consult another expert system.

System administration software, developed and interfaced with the Guru software, captures both performance information and the user query and is used to address and upgrade the system through the terminal shown on the upper left of figure 3. Users on the network communicate with the expert help system through their terminals by using natural language queries. User queries are processed through one of three possible paths depending on the nature of the query. For example, questions regarding complex actions (such as locating a lost job) are handled by the expert system. Special requests that require the execution of a program (checking the status of the Cray), a procedure (how to make a Cray job), or the display of a form (how to change your password) are routed through the path shown on the right of figure 3. Simple information requests (such as how to obtain an account) are obtained by default from the relational data base, the middle path of figure 3.

Results

The Cray user services statistics were analyzed over a period of 1 year to establish a measure of the potential benefits from ECUHS and to guide the development of the initial knowledge base for the preliminary system. During the test year, Cray user services received 1784 calls requesting information or assistance. The calls were sorted into the 17 subject areas listed in the first column of table 1. An estimate of the percentage of questions an expert help system could process is shown in the third column, and the right column indicates the product of the other two numerical columns. It was estimated that nearly half (777) of the questions could be processed by an expert help system. This suggests that an effective expert Cray user help system could significantly reduce the user

services workload and increase their effectiveness by answering routine queries, thus allowing user services personnel to concentrate on problems that require their skill or that are beyond the scope of the expert system.

Many of the user questions categorized in table 1 are trivial, which is astounding considering that users were supplied with an excellent user guide containing answers to nearly all of those questions. This fact reinforces the need for an on-line expert user help system and suggests that the ECUHS can simply duplicate the answers provided in the user guide for those questions requiring simple responses. User statistics were monitored continuously throughout program development, and periodic analyses were made to determine if significant changes were detected in the above categories or in the complexity of user questions. To date, no significant change has been noted.

The preliminary knowledge base was limited to questions within the more manageable categories listed in table 1. Development and testing of the preliminary ECUHS required only 10 months. However, the initial system performance was marginal because of the instability of the Guru software (VAX Guru was in Beta testing) and also because of the limited knowledge base. Within the following year, the software vendor implemented several improvements to the VAX version of Guru, which improved both system stability and functionality. The introduction of additional select users provided a larger query base which was used to expand the knowledge base and make the system responsive to a broader class of queries.

During system development, it was observed that both the implementation of Guru software upgrades and the expansion of the knowledge base affected system performance; therefore, a standard query base was developed and used as a test suite for subsequent upgrades of the ECUHS. Appendix A lists the 99 queries in the test suite. The test suite was also used to determine system response in both single- and multi-user modes. Examples showing the ECUHS display response to representative questions in Appendix A are presented in Appendix B.

The system administrator software automatically logs the query, its source and user name, the answer source and record, and the response time (the time elapsed between when the query was received and when the user responded to the options presented after the query was answered). Appendix C provides a description of the system administrator software and the data base it generates. During development, the system administrator data base was analyzed to determine system performance and to collect new queries to guide the growth of the knowledge base. No query was encountered that was considered beyond the scope of the expert system software or of the ECUHS architecture.

At the conclusion of the final test period, the system administrator data base was sorted to eliminate duplicate queries. Throughout the ECUHS development, 164 unique queries were recorded and the system response to each was analyzed. The ECUHS responded correctly to 99 of the unique queries (the test suite); it gave an acceptable response to an additional 42, and it failed on the remaining 23. The 42 queries for which the ECUHS response was considered "merely acceptable," and the 18 queries for which the response was considered unsatisfactory could be answered correctly by upgrading the knowledge base. The remaining five unsatisfactory responses can be corrected by increasing the vocabulary of the natural language interface.

The limited scope of this program did not allow detailed investigation of all problems associated with the growth and maintenance of an expert user help system, such as the ECUHS. For example, the final test suite was evolutionary and consisted primarily of all satisfactorily answered questions logged during the preliminary test. As the system grew, new queries occasionally altered the system response to old queries and were therefore included in the test suite. No attempt was made to omit any of the initial queries or to determine if a smaller test suite would prove satisfactory.

Low system overhead and a robust response, to both single- and multi-user queries, are critical to the success of an on-line user help system. The ECUHS system overhead was determined by monitoring the VAX "Top CPU" utility with as many as four simultaneous ECUHS users. Figure 4(a) depicts the results of a typical test performed on July 7, 1988, with a light load on the EARth VAX. Figure 4(b) depicts the results on July 6, 1988, with the EARth VAX nearly saturated. Four subjects—the second, third, fourth, and eighth (fig. 4(a))—were logged onto the ECUHS from the remote node EARth, on the ARC network. The ECUHS software resided on the node MARs (fig. 2). (Both nodes are Digital Equipment Corporation (DEC) VAX 8650 CPUs). All subjects on the EARth node used queries from the test suite of Appendix A.

For the figure 4(a) tests, a maximum of 20% of the VAX CPU was sufficient to service the ECUHS system in multi-user mode. Throughout the entire test period, the maximum overhead for each of four active processes never exceeded 24% of the VAX 8650 CPU, and the examples in figure 4 are typical of tests run at various times and dates. Note that the first entry in figure 4(a) shows the VAX was idle 46% of the time, indicating ECUHS resource requirements were reasonable even in the multi-user mode. During periods when the VAX CPU was heavily loaded, ECUHS allocation was reduced approximately 50%, which resulted in a degradation of performance to all four users, (fourth, sixth, seventh, and eighth entries of fig. 4(b)).

The ECUHS query response time depends on several parameters including the network delay time, the number of network users, the number of node users, and the complexity of the query. System response time is important to the ultimate success of an ECUHS; therefore, a test was devised to obtain a characteristic measure of this overall delay time as a function of the time of day. The test consisted of four queries with one representing the minimum response time encountered and tested in each of the three loops that the parser can select, and one that invokes the expert system loop and requires the most complex response, thereby, requiring the maximum time. Tests were conducted on different days of the week and in different months to randomize network anomalies and unusual node workload.

Figure 5 presents the delay times obtained on a typical session (April 7, 1988) as a function of the time of day. The test suite was initiated on the hour and repeated three times. The minimum and maximum response times observed during each session are recorded in the scatter chart of figure 5. The ECUHS response was affected by network and user workload and, in the worst case, required roughly 55 sec. With minimum traffic, the ECUHS required between 10 and 25 sec to respond to a query. Although the minimum 10-sec delay, associated with simple responses, is considered marginal, methods exist to reduce this to the acceptable threshold of 3 to 5 sec, which is typical of today's robust on-line help systems.

In these tests, the maximum response time was associated with a complex query that immediately presented the user with system status. For this type of query the delay may be acceptable, based on similar responses seen in today's on-line systems when long menus are displayed or when user interaction is required. The delay caused by system overload, demonstrated by the hourly variations in minimum and maximum delays, may also be acceptable as users learn to accept the slow response time associated with an overloaded node or multi-user system.

Comparison of On-Line Help Systems

Today, nearly all large computer systems include an on-line help system to aid the user in systems communication and operation. There is an increasing trend to also include on-line help systems on small computer systems, engineering workstations, and PCs. Most of these help systems are key word or menu driven and are particularly useful as an alternative to the user manual when the index or table of contents must be scanned to locate the subject material. Recent developments, such as the extended/ programmable keyboard and windowing (multiple windows), simplify the use of on-line help systems and enhance their utility through special key functions and improved graphics displays.

Although the philosophy driving the ECUHS design avoids menus, several menu-driven, on-line help systems were reviewed; desirable features were identified and incorporated into the ECUHS; and awkward or undesirable features were either modified to conform to the ECUHS format or were omitted. In general, the ECUHS is more comparable to key word on-line help systems; however, comparisons of system overhead and implementation are meaningful measures for all on-line help systems, including menu-driven systems.

Two DEC on-line help systems, the VAX VMS and VAX ULTRIX, were analyzed and compared to the demonstration and projected ECUHS. The VAX VMS is a menu-driven system, and the VAX ULTRIX is a key word-driven system. Table 2 summarizes pertinent statistics of four systems analyzed in this study. Two VAX VMS help systems are included in the table, one from the EARth 8650 VAX and the second from the SATurn 785 VAX. These were included to demonstrate the variation in help-system size from one computer system to another. The difference is due to the number of software products installed on each system and not to the power of the CPU. For the VAX VMS environment, the software for each application includes a help file that can be installed as part of the existing help system with the appropriate menu entries. The EARth VAX has 187 primary menu entries, and the SATurn VAX has 168. Nearly all menu items have examples or subtopic menus and each subtopic often includes additional subtopics (tree structure) and examples.

The EARth VAX help system has 1782 second-level entries, excluding the example option which is often the most helpful and is what the user prefers. The EARth VAX VMS help system requires over 12 million bytes of disk space and is considered very good. The key word dictionary entry comparison was made by considering each menu entry as a key word entry. This assumption is reasonable, because help on each menu entry can be invoked directly by prefixing the menu item with "help." Exceptions occur when a submenu entry is keyed to the subject entry and requires compounding key words to access the area directly. Examples include key words such as "debug" or "library" which have first-level menu entries, and second-level menu entries under Fortran, Basic,

Pascal, etc. Menus are most helpful when a user wants to scan the help subject index or cannot recall a key word.

The VAX ULTRIX help system requires over 85 million bytes of disk space and has 754 key word dictionary entries. This is roughly three times more information per entry than for the VAX VMS help system. More information is not always desirable, because the user is often seeking a simple reply which can be obscured or overlooked while scanning a complex or multiple-page response. Key word entries are generally more direct and preferred by most users because the information they seek is associated with the subject or the active program. Although subject menus are sometimes helpful, they are generally not included in key word on-line help systems.

The demonstration (or test) ECUHS requires 827,136 bytes of disk storage and has 738 dictionary entries. The ECUHS entries include both key words and phrases, which make direct comparisons with other menu or key word systems difficult. By design, menus were avoided in the ECUHS; however, a subject menu can be invoked through the expert system by requesting what the system knows. The estimated storage requirements and dictionary entry size for the projected ECUHS were based on the experience gained in developing the demonstration system and the table 1 projections. Although the projected storage requirements are roughly 10 times the demonstration system (10 million bytes), the dictionary entries are projected to increase only by a factor of 5 to roughly 2000 words. Fewer dictionary entries are expected because most of the required redundancy has been included in the demonstration system, and the remaining question categories deal with job-control language and programming, which are specific. On the other hand, the expert system response for the remaining question categories will require more text, particularly where examples are required.

CONCLUSIONS AND RECOMMENDATIONS

Program results are encouraging and no major obstacles were encountered that would prevent the development of a production Expert Cray User Help System (ECUHS). The technology of existing expert system shells and natural language programs is adequate to support the development of a production system and, with minor effort, the ECUHS can be expanded to meet production system requirements.

It should be emphasized that the proposed system is not intended to replace user services with an expert help system; rather, the expert help system is viewed as an extension of user services. It would be available 24 hours a day and would handle most trivial questions, thus off-loading user services and allowing them to work on the complex problems that require their skills. Moreover, the expert help system ultimately must be maintained by the user services staff, who possess the knowledge base required by the system. This maintenance becomes an important consideration in a cost-benefits analysis regarding the implementation of such a system. The rudimentary staffing/cost analysis performed in this study indicates that a one-and-a-half (1 1/2) man-year effort and a software license cost of approximately $7,000 a year, would be adequate to upgrade the ECUHS into a production environment. Because the system has the potential of reducing the staffing requirements of

user services enough to compensate for the ECUHS maintenance requirements, the development of a production ECUHS is considered cost effective.

This effort might have been initiated except that the Ames Research Center Advanced Computer Facility elected to undertake a major conversion to the UNIX operating system. The decision to convert the ACF into a UNIX environment, and the Cray operating system from COS to UNICOS, was made when the ECUHS program was approximately 50% complete, so the ECUHS program direction and remaining program elements were altered to accommodate the change. Care was taken to implement only those features that would be valid with little or no change for conversion of the ECUHS into the UNIX environment.

Although Micro Data Base Systems, Inc., plans to port the Guru software into the UNIX environment, no VAX ULTRIX version was available at the conclusion of our program; therefore, cost and staffing requirements for the development of a UNIX ECUHS were not determined.

The ECUHS development revealed an interesting problem associated with complex networks that has far-reaching implications to developers of software, network hardware, and network protocol. Both hardware-implementation requirements and software-license cost considerations dictated that the ECUHS software reside on a single computer system. To make ECUHS available to all networked computers requires implementing automatic remote log-ins from remote non-host computers. A remote log-in generates many small message packets and tends to overload some network devices. Moreover, Guru software features, such as the hot key concept which is becoming more popular in workstation software, also generate small network packets, that is, a single byte, adding to the network traffic.

Developers of network hardware and network software protocol must be cognizant of the above developments and requirements and design them to ensure a robust network. Conversely, developers of application software should also be cognizant of network limitations and should develop software that can, at the user's option, bridge interactive elements such as hot keys, single character transmissions, etc.

In conclusion, the program objectives were accomplished and the results exceeded all expectations. The use of an existing expert system shell with a natural language interface, relational data base, and integrated utilities reduced development time by several months and resulted in an integrated demonstration system. The demonstration system framework is suitable for the production system, and the ECUHS software can be modified and expanded to fulfill the production system requirements.

Demonstration system tests indicate that a production ECUHS can be as responsive and robust as a menu-driven system. For a multiple-computer networked environment, such as the ACF, an ECUHS residing on a single computer has two advantages over existing on-line help systems: (1) help-system upgrades are both easier and cheaper because only one computer system is affected; and (2) software disk requirements for an ECUHS are reduced by a factor equal to the number of processors on the network over menu-driven help systems, and by an additional factor of 5 to 8 over key word help systems. An ECUHS provides an accurate log of user questions with the corresponding system response. This enhances problem tracking and simplifies statistical analysis on queries, thereby making upgrades more responsive to new problems as well as identifying obsolete help data that can be deleted.

APPENDIX A

ACCUMULATED DESIGN QUERY BASE

Three sources were used to derive the query base from which the knowledge base for this project was developed.

1. The user-services weekly status report logs, which synopsize user questions and the response provided by user services.

2. The ACF User Guide, the manual that is distributed to each user who has an ACF account. It provides an extensive and detailed description of the ACF hardware and software.

3. Selected test subjects who were interested in the project and willing to test or challenge the capability of the ECUHS.

The query data base listed below was used to develop test software that was automatically executed after each expert-system shell software upgrade, and after each major program upgrade to ensure that the ECUHS performed as expected. (Queries are worded exactly as they were received from users.)

1. How do I get Cray User Manuals?

2. How do I get Cray manuals?

3. How do I get Cray manuals

4. Where are Cray manuals?

5. Where can I get reference manuals?

6. Where do I get Cray user manuals

7. Help

8. Subject

9. How do I get started

10. How do I get started on the Cray?

11. How do I contact user services

12. How do I access the CSF

13. Get a new account

14. How can I get an account on the Cray?

15. How do I get a user account

16. How do I get a userid?

17. How do I get an account

18. How do I get an account on the Cray?

19. I am a new user on the Cray

20. How do I get a password?

21. How does one change his password

22. How can I change my password

23. How do I change my password

24. I want to change my password

25. Show me how to change my Cray password

26. Transfer a file

27. Cray xwindows interface

28. Submit a job

29. I want to send in a fortran program

30. Cray jcl?

31. Can you give me a sample job

32. Please give me sample jcl

33. Sample

34. How do I create a Cray job

35. How do I create a Cray job?

36. How do I create a Cray job deck

37. How do I create a job for the Cray?

38. How do I make a job

39. How do I run a Cray job

40. How do I run a Cray job?

41. How do I run a fortran program on the Cray

42. How do I run on the Cray

43. How do I send a job to the Cray

44. How do I submit a job

45. How do I submit a job to the Cray

46. How do I submit the job

47. How do I compile my program

48. How do I write a job

49. How do I use the CSF

50. What languages available

51. What size Cray can I run on here?

52. What type of Cray?

53. What are the logons

54. Where is Cray user services

55. Telephone number for user services

56. User services

57. What are the job classes

58. What is the operator's phone extension

59. What is the Cray schedule

60. What is the schedule for the Cray

61. What is the memory size on the Cray?

62. How much memory is on the Cray

63. Is the Cray available

64. Is the Cray up

65. Is the x-mp available

66. Is the xmp available

67. Job status

68. Cray schedule

69. Can I get more disk space

70. Getting more disk space

71. My job aborted

72. My job bombed

73. My job did not return

74. My job didn't run.

75. My job failed

76. I lost my job

77. My job is lost

78. Lost job

79. My output did not come back

80. Where is my job

81. Whre is my job

82. Where is my output?

83. Where is my job now

84. Where do I get my output

85. Where do I pick up my output

86. Where to I pickup my output

87. Tape information please

88. Tapes

89. Drop a job

90. How do I drop a job

91. How do I kill my job

92. I want to kill a job.

93. I want to kill my job.

94. I want to kill that job

95. I want to stop a job

96. What do you know

97. What do you know about

98. Who wrote you

99. Words

APPENDIX B

ECUHS SCREEN RESPONSES

The ECUHS screen response examples described in this Appendix demonstrate the output obtained from the three major paths shown in the block diagram of figure 3. Where appropriate, the response is contrasted with the current procedure to demonstrate the potential benefit that can be derived from an on-line expert help system.

Any user with access to the Ames Research Center Network can invoke the ECUHS with the command "crayexpert." The ECUHS recognizes first-time users and requests the user's last name. Once provided, a one-screen description of the ECUHS is displayed. This screen is presented only for the first log-in; thereafter, ECUHS recognizes the network log-in ID and proceeds directly to the query screen consisting of the header "Ames Research Center Natural Language Cray Help," followed by the question "Your Request?" The examples that follow include the query screen information, followed immediately by the query in italics and the ECUHS response to the query. In an active session, the ECUHS clears the header screen and query before painting the response on the display screen.

The display screen shown in figure 6 lists the ECUHS response to the question "How do I change my password?" The ECUHS executes a procedure that provides the information (the right path of fig. 3).

The information contained on the display screen was taken from the ACF User Guide and was generated using the Guru Forms Generator. A user requesting the information from Cray User Services is prompted with the pertinent information and referred to the page in the ACF User Guide that contains the information. It is debatable which is more expedient for the current ECUHS implementation of the question, the call to Cray User Services or the invocation of "crayexpert." An alternate implementation, more fitting for an expert system, might invoke an interactive session that guides the user through the password change. The demonstration system contains several examples that respond with information from the ACF User Guide, instead of the more helpful information that would be part of a final on-line expert help system.

The next example, figure 7, demonstrates the ECUHS response to another important but simple question: "How do I get a Cray account?" During 1 year (March 3, 1986 - March 3, 1987), there were 112 inquiries to user services regarding accounts or passwords. This example was selected to demonstrate the Natural-Language Interface regarding phrasing. Several VAX users were canvassed to determine how they would request a Cray account. Table 3 lists the 23 unique requests that were obtained. The natural-language interface responded correctly to all the questions.

In this example, the ECUHS presents account information and phone numbers of the new-accounts administrator, which is the same information issued by user services. This information is retrieved from the data base, the middle path of figure 3. The new-accounts administrator instructs potential users to complete a form that requires management signatures. The ECUHS could just as

easily provide an image of the form on the user's screen or request the required information from the potential user. Once that is supplied, the ECUHS could submit the information through either hard copy or electronic mail to the appropriate manager for the required signatures to initiate the new account.

The next example demonstrates the ECUHS response to the query "Where is my Job?" The ECUHS response is shown in figure 8 and was generated by the expert system, the left path of figure 3. In this example, the user has submitted five Cray jobs from a remote VAX networked to a VAX running the Cray station software. The network interface software assigns a Terminal Identification Number (TID) to each job upon submission. The user must remember his/her TID numbers because the user job name does not appear on networked jobs in the CSTAT (job status) display. To assess the benefit of the ECUHS, one must contrast the figure 8 response with that in figure 9 which represents samples of the normal remote VAX screen output invoked by the CSTAT command. The normal CSTAT output generally fills several screens which the user must scan to locate the TIDs associated with the user's jobs. In this example, the CSTAT command filled five screens. Three of the five TIDs are near the beginning (fig. 9(a)); however, the fourth is on the fourth screen. The fifth job is located on the fifth screen, noted only by its TID number (080138G9) (fig. 9(b)). Most users overlook this information and call user services to report a lost job. In fact, this problem resulted in over 50 calls to user services during the year.

With the expert help system in place, the user would log on to the terminal, invoke Guru; type "Job Status," "Where is my Job?," "CSTAT," or similar questions; and the system would return the information in figure 8. Note that the user did not have to type an ID or other information because the ECUHS uses the log-in ID to determine the TID number. Furthermore, all of the information about the user's jobs is in a compact report that includes the status of the job in the VAX Station TRANSIT Queue. The time required to get the complete Guru report is less than that required to scan the table for the jobs, and it is considerably less than that required to solicit a response from Cray User Services. In this example, an expert help system obviously simplifies the user's job considerably.

A more complex situation exists if the job goes to the BAD queue, shown at the bottom of figure 9(b). An experienced user may understand what caused the problem and know how to correct it. The inexperienced user must call user services because this problem is not covered in the manual. For example, a mistyped VAX password will create this problem, but the cryptic comment does not relate this information or guide the user to a solution. Figure 10 displays a "Bad Queue" response and illustrates how the ECUHS has drawn on the knowledge base derived from what the user services expert did to solve the problem. User services need not have been contacted if the expert help system had been in place.

The last example, figure 11, demonstrates a complex yet helpful feature that can be provided by the ECUHS. Figure 11(a) depicts the screen response to the query "How do I run a Cray Job?" The screen was invoked through a procedure, the right path of figure 3. The information duplicates the ACF User Guide Section 2.2.2; however, an interesting or helpful option, which generates a sample job in the user's VAX directory, is offered and can be invoked by selecting the second option shown at the bottom of the screen. Depressing the letter "c", <Hot key c> on the user's keyboard, commands the ECUHS to create the sample job with the default name, "ACRAYJOB.LIS," and

place it into the user's VAX directory, depicted by the top four lines of figure 11(b). Next, the user directory and the sample job are invoked using the VAX commands "dir" and "ty" which generate the directory listing and the sample job, figures 11(b) and 11(c).

APPENDIX C

SYSTEM ADMINISTRATION

One important user service function that receives little recognition is the personal contact and exchange that takes place as a result of a user's call. This communication, whether by phone or in person, provides valuable feedback from users. The implementation of an on-line expert system could jeopardize the "personal touch" derived from this contact. Users who prefer personal communications will undoubtedly continue using them, unless discouraged. Others will find communication with an inanimate object, such as an expert system, more expedient and less inhibiting. Nevertheless, their questions and concerns must be solicited and the information must be captured to achieve a clear understanding of user concerns and opinions.

These concerns, coupled with the belief that an expert system must be allowed to grow and evolve, led to the development of system administration software. This software is not intended to replace or provide the feedback derived from verbal communication with users, but rather to capture user questions, help expand the knowledge base, and record performance parameters for system upgrades and improvement. It is also intended to illustrate the types of information that can be obtained by an automated system.

The ECUHS system administration software consists of several utilities that includes a procedure to capture user data after each query is entered and send it to the natural language interface for processing. The procedure also stores the data in a data table for later analysis using the utilities. This capture of user phrasing, vocabulary, and subjects will allow the program to grow in the direction of its use. Typical output from the utilities is illustrated in figure 12.

The system administration output, shown in figure 12, includes user data consisting of the date and time of the query, the user name and node that the user is logged on to, and the specific question the user asked. The second line indicates whether the system answered the user query, the source of the answer, the record number for data base answers, and finally, the time between when the query was initiated by the user (user presses the return key on the keyboard), and the time the user initiates any keyboard entry after the query response is displayed on the screen. This timing loop is used to obtain the system response time to various questions by initiating keyboard input immediately after the answer is displayed on the screen. The times shown in the figure 12 example are meaningless because users typically take some time to view or read the screen response before initiating further keyboard input. This feature is only useful to the system administrator when it is used as indicated above.

Utilities exist to retrieve all recorded data, all data since last report, all data collected since a specific date and time, and data collected from a specific node or user. The system administrator can also retrieve the report to satisfy any Boolean expression using the data entries.

20

REFERENCES

1. Leghart, Frank: Supercomputer Technology. Computer Technology Research Corporation, 1989.

2. Lusa, John M: Industry Highlights: Record 16,500 Attend SUPERCOMM; GTE's CEO Says Data Traffic is Growing at 20% Rate. Networking Management, A PennWell Publication, July 1989, pp. 12-14.

3. Lusa, John M: Industry Highlights: Fortune 1000 Telecom Spending to Get Increasing Share of Operating Budget in Next Five Years. Networking Management, Nov. 1988, p. 15.

4. Baker, Gary: Best Sellers: North America's Leading LAN Value-Added Resellers. LAN Magazine, March 1989, pp. 39-56.

Table 1. User Services Calls Applicable to Expert System.

Question category	Total number of calls	Percent applicable	Projected number
Information requests/questions	189	25	47
Program/system help requests	324	30	97
Operator/analyst requests	370	10	37
Programmer software errors	98	5	5
System faults	56	0	0
Hardware	15	0	0
JCL	69	70	48
VAX/CRAY interface (transfers)	57	95	54
User ID/account/password	112	95	106
Lost jobs	50	95	48
Kill program	78	100	78
Quota information/questions	26	90	23
CYBER 835 questions	85	90	80
Files-problems/questions	58	60	34
Tape-problems/questions	128	70	90
Disk questions-(SSD)	33	50	16
Miscellaneous	36	40	14
Totals	1784		777

Table 2. System Level Help System Comparison.

	Storage requirements (bytes)	Menu entries (first-level)	Key word dictionary entries
VAX VMS help system			
EARth 8650 CPU	12,709,888	≈187	>2000*
SATurn 785 CPU	6,604,288	≈168	>2000*
VAX ULTRIX help system			
PIOneer 785 CPU	85,565,000	0	754
VAX VMS test ECUHS			
MARs 8650 CPU	827,136	0	738
VAX VMS projected ECUHS			
EARth 8650 CPU	≈10,240,000	0	2000**

*Each menu entry can be invoked directly with the "help" prefix.
**Projected key word entry dictionary size.

Table 3. Natural Language Phrasing.

Need Cray Account
Get a new account on the Cray
Help access Cray
I want to know how to get on the Cray
Access the Cray
Get a Cray account
Get a new account for the Cray
Need Cray account
Want a user number for the Cray
Account number for the Cray
Username for the Cray
User ID for the Cray
Logon for the Cray Please
I need an account
Can I have an account on the Cray
How do I get on the Cray
Help Cray
How do I get an account number
Help account
Help Cray account
How would you get a Cray user name
I want an account on the Cray
How do I run on the Cray

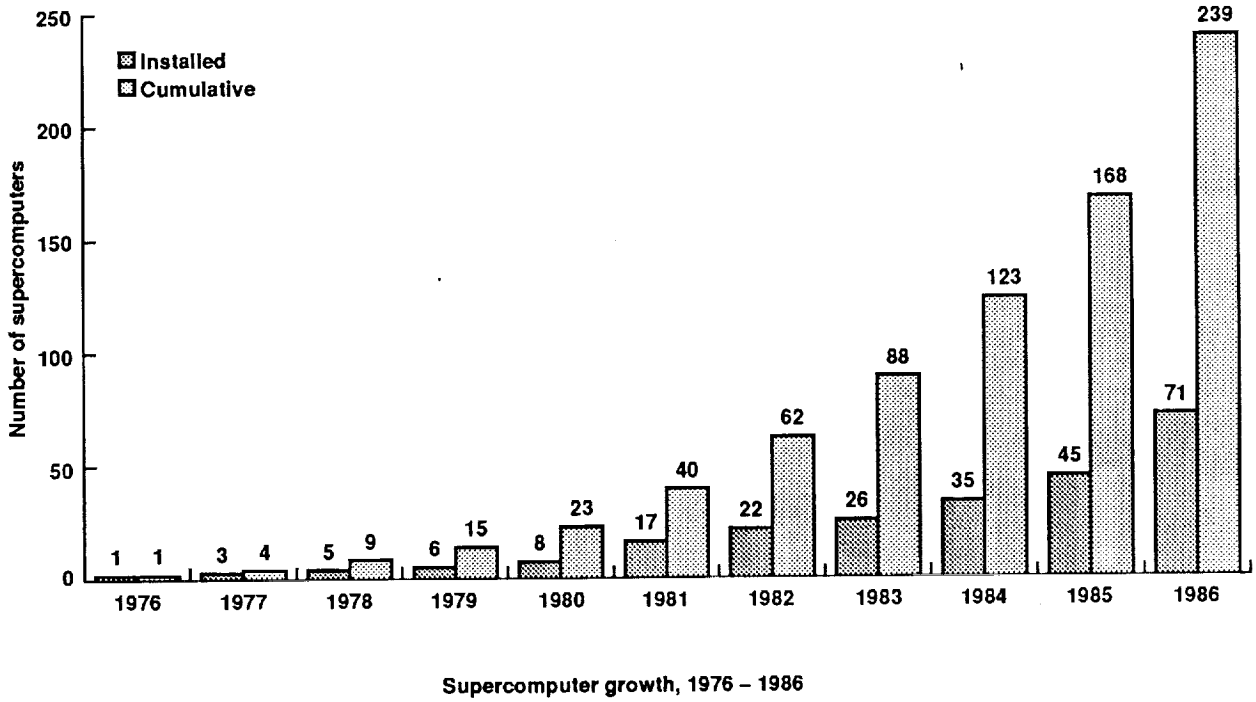

Figure 1. Supercomputer growth, 1976–1986.

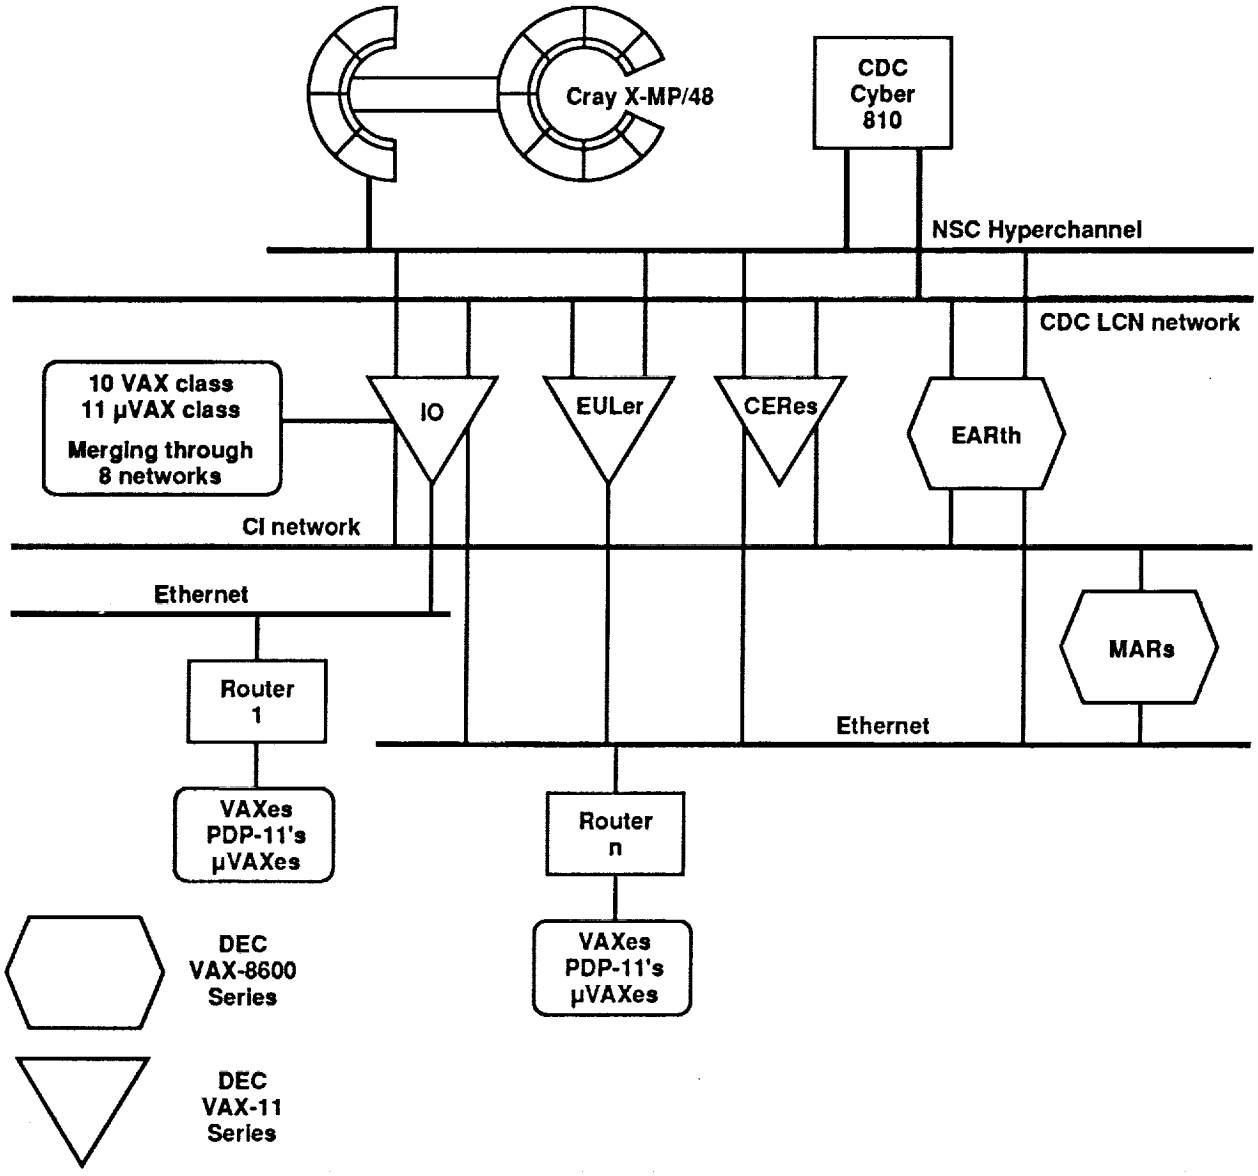

Figure 2. Ames computer network simplified schematic.

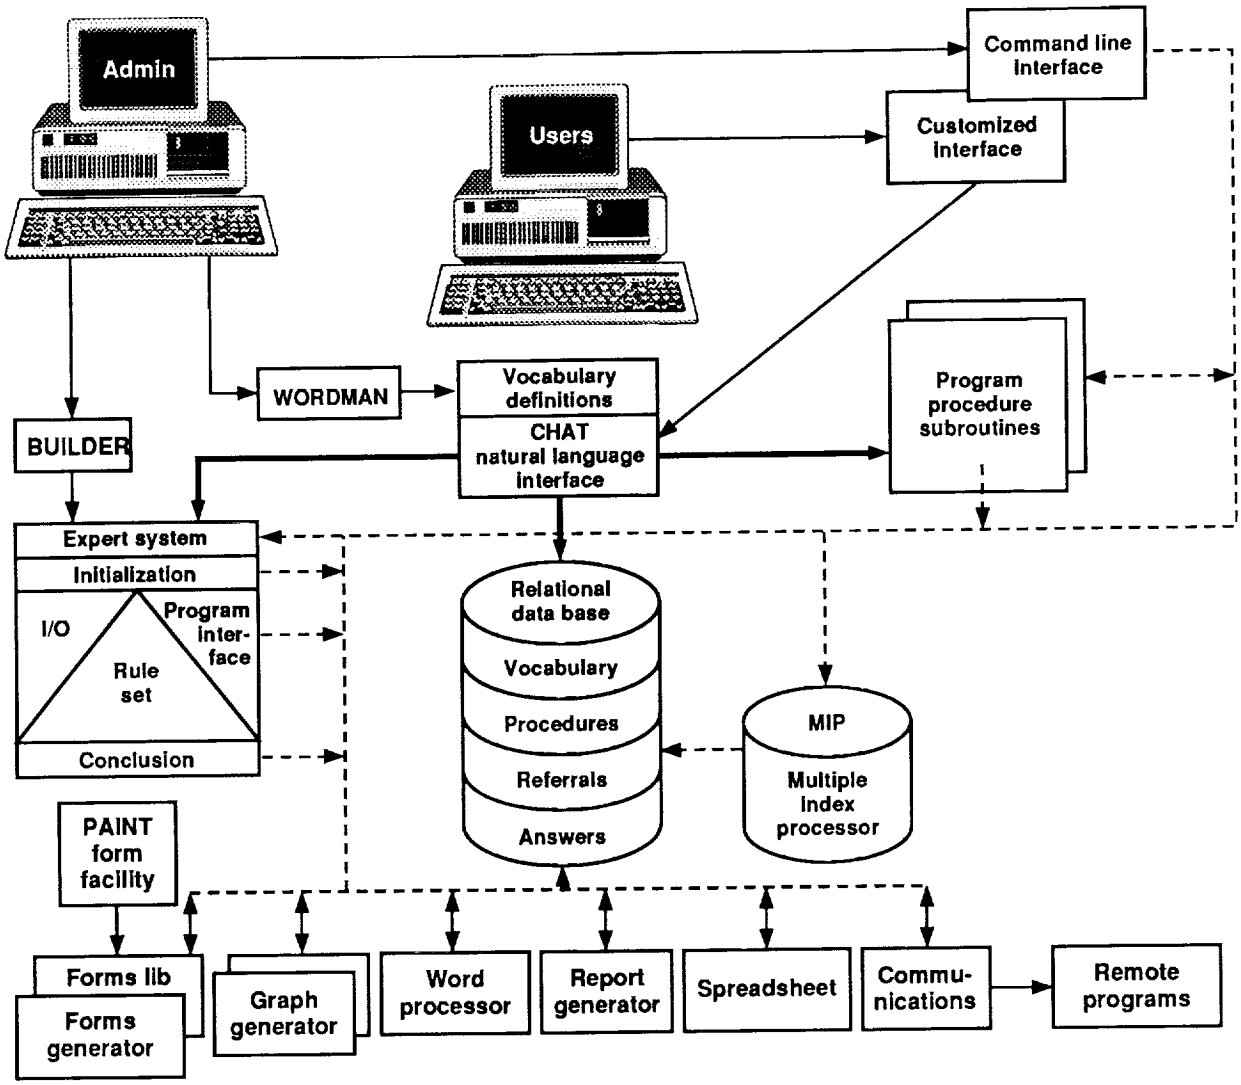

Figure 3. System block diagram.

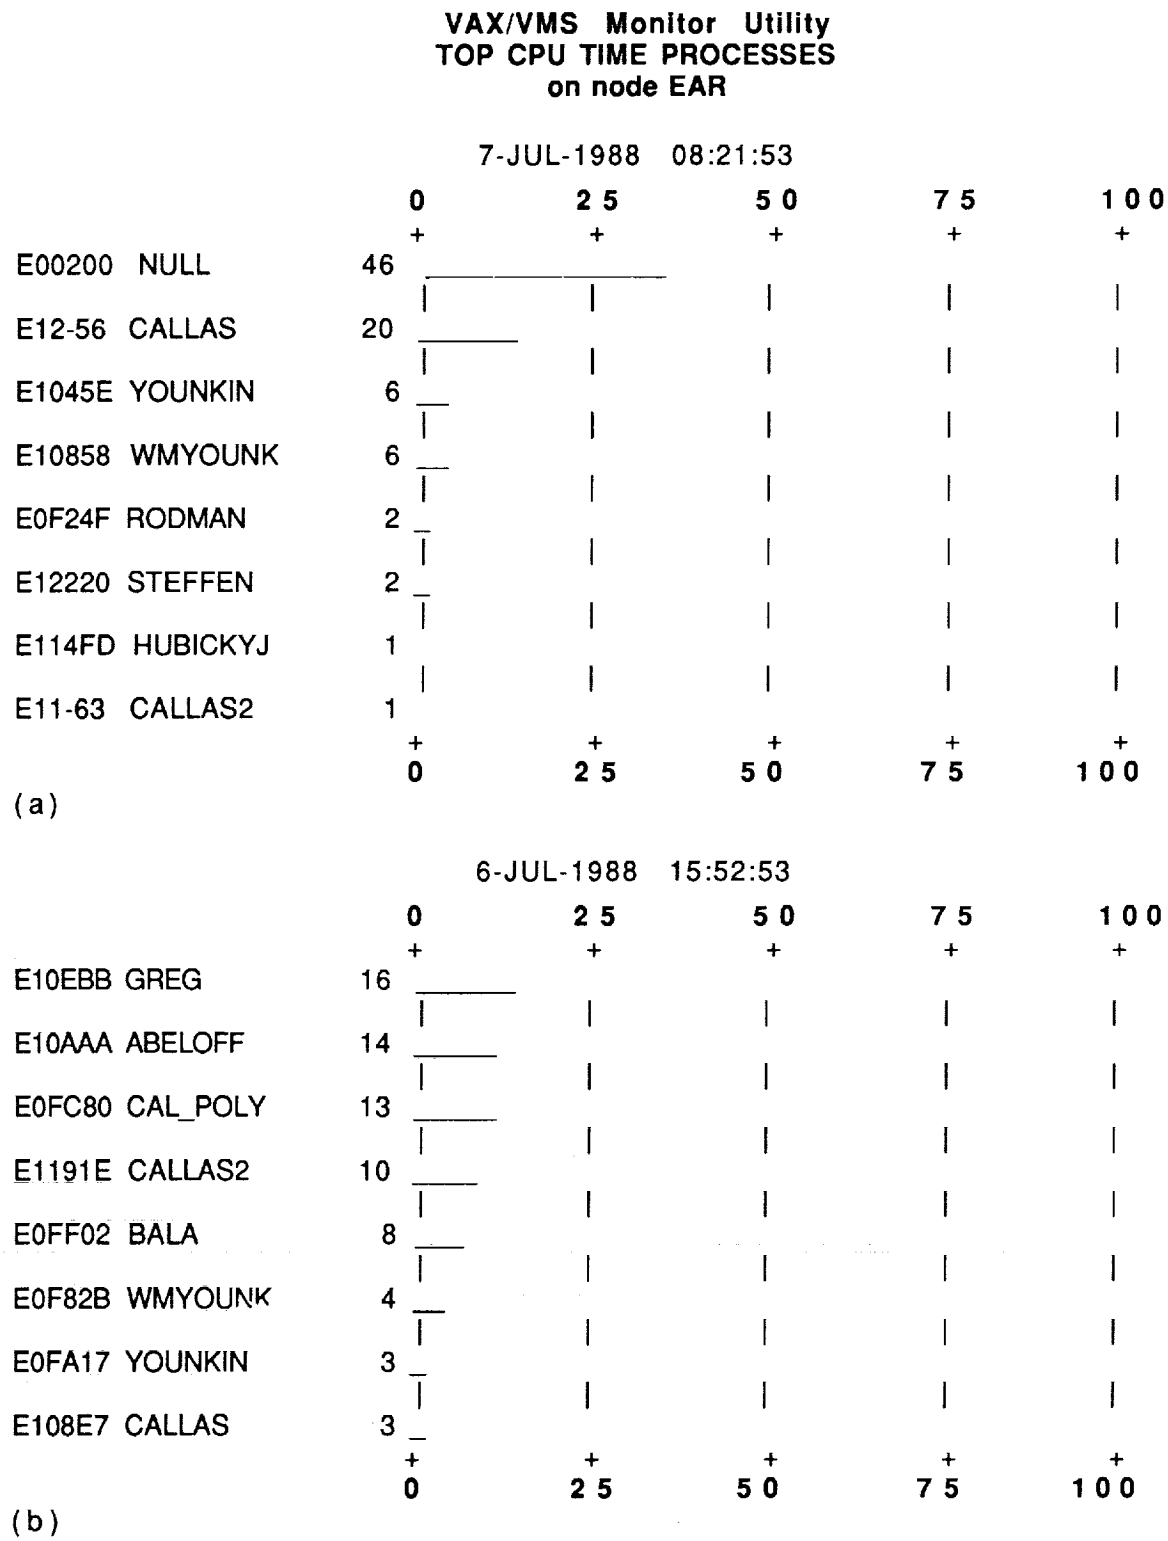

Figure 4. Processor time. (a) For top 8 of 37 EARth VAX 8650 users; (b) for top 8 of 73 EARth VAX 8650 users.

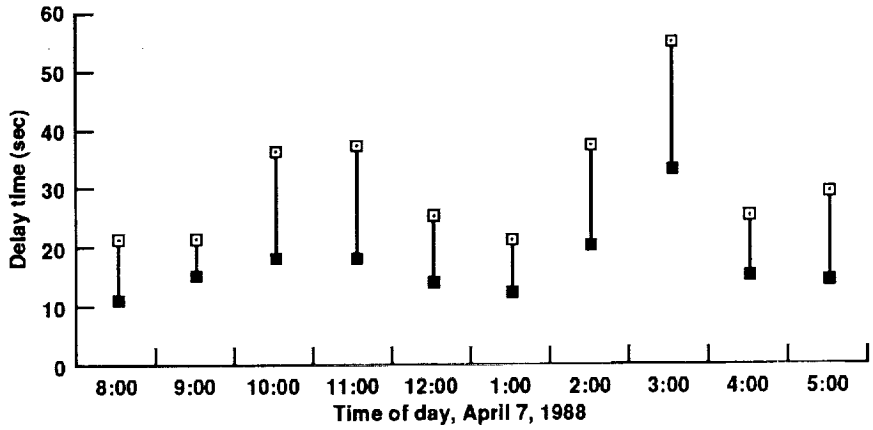

Figure 5. Variation of delay time with time of day.

Ames Research Center Natural Language Cray Help

Your Request? *How do I change my password?*

A Cyber 810 userid is automatically installed when a Cray userid is installed. The Cyber 810 and Cray userids, passwords and account codes will be identical. When installation of the userid, password and account code(s) is complete, the user will be notified. Users are urged to change their passwords during initial access to the Cray and periodically thereafter.

To CHANGE PASSWORD include the following two statements in a Cray job file:

```
AMESLIB.
PASSWORD,NPW=newpassword
```

where "newpassword" is user-specified password of 4 to 7 characters.

The PASSWORD utility changes both the Cray and Cyber front-end passwords. Cray and Cyber front-end passwords must be identical because of Cray utilities which communicate with the Cyber.

wait Logging request, one moment

Enter "Q" to QUIT or RETURN to CONTINUE Your Request?

Figure 6. Expert Cray user help system screen response to query "How do I change my password?"

Ames Research Center Natural Language Cray Help

Your Request? *How do I get an account?*

ACCOUNT CODE GET
Obtain account code from your Ames organization RESOURCE MONITOR

ACCOUNT NEW GET
To get an account on the Cray X-MP, call User Services at x6515.

ACCOUNT NEW
To inquire about an account on the Cray 2, Call NAS User Services: x4444.

ACCOUNT NEW
To open a VAX account, contact your Ames organization RESOURCE MONITOR

Logging request, one moment . . .

Enter "Q" to QUIT or RETURN to CONTINUE _

Figure 7. Expert Cray user help system screen response to query "How do I get an Account?"

Ames Research Center Natural Language Cray Help

Your Request? *Where is my job?*

You have 4 job(s) in the Cray queues:

Jsq	DC	Dataset	Class	Status	Pri	Time Used	Time Limit	Field Length	ID	TID
1884	IN	CRON	DMP	WAIT-TIM	10.0	36	1000	75	CV	080138CA
26867	IN	DIPHPR	BIGDISK	QUEUED	6.0	0	*****	974	CV	08013812
31867	IN	3D	SSD5	ROLLED	6.0	3000	*****	1257	EV	080138AA
37974	IN	BKINTS	SSD3	Q-RSOURC	8.0	0	6000	1949	CV	080138N8

One Job with your TID in VAX Station queues:

CrayVAX Station Transit Queue - TRANSIT

Ord	TID
179	080138G9

wait . . .

Logging request, one moment . . .

Enter "Q" to QUIT or RETURN to CONTINUE

Figure 8. Expert Cray user help system screen response to query "Where is my job?"

Operator message queue - OPMSG

--- SYSTEM STATUS -------

Cray is UP
VAX station is UP

ACF Cray Queues - CRAY

Jsq	DC	Dataset	Class	Status	Pri	Tim Used	Tim Limit	Field Length	ID	TID
1884	IN	CRON	DMP	WAIT-TIM	10.0	36	1000	75	CV	080138CA
18496	PR	MMC2		HOLDING	13.8	0	0	0	EA	AABjvOAA
26867	IN	DIPHPR	BIGDISK	QUEUED	6.0	0	****	974	CV	08013812
31867	IN	3D	SSD5	ROLLED	6.0	3000	****	1257	EV	080138AA
31891	IN	M6XXX	SSD4	ROLLED	6.0	90	599	3186	CV	1A0320SY
31908	IN	ECOBPTC	SSD5	ROLLED	6.0	1268	6500	1752	CV	187320AJ
31910	IN	LU	SSD5	ROLLED	6.0	2690	7200	855	CV	1A0B20E0
31979	IN	SCX02	SSD5	ROLLED	6.0	2501	6000	727	EA	AAWPFaAA
32057	IN	M12D	N5	ROLLED	6.0	3745	****	1560	EA	AANPFaAA
•										
•										
•										
37974	IN	BKINTS	SSD3	Q-RSOURC	8.0	0	6000	1949	CV	080138N8
•										
•										
•										

Figure 9. Normal job status display. (a) Invoked by "CSTAT" command.

```
  •
  •
  •
37974   IN   BKINTS   SSD3      Q-RSOURC    8.0       0      600    1949     CV    08138N8

CrayVAX   Station  Input  Queue  -  STIN
  Ord    TID            Username        Output

CrayVAX   Station  Output  Queue  -  STOUT
  Ord    TID            Username        Output

CrayVAX   Station Acquire  Queue  -  ACQUIRE
  Ord    TID            Username        Input

CrayVAX   Station  Transit  Queue  -  TRANSIT
  Ord    TID
  179    080138G9

CrayVAX   Station  Bad File Specification Queue  -  BAD
  Ord    TID            Username        Output

   44    A87FF3876     RFRFML          JPU$DRA1:[JUPNETWRK.RFRFML]PRINT.OUT
                disk quota exceeded,  size  =  1689  blocks
```

Figure 9. (b) Fifth screen. Concluded.

```
No jobs with your TID number in Cray queues


One job with your TID in VAX Station queues:

CrayVAX    Station  Bad File Specification Queue  -  BAD
  Ord     TID            Username          Output
    1    08013881        YOUNKIN           DRA1:[YOUNKIN.CRAY]TEST.OUT
                 login information invalid at remote node


Possible typo on password when job was submitted.
Please enter your VAX password and I will attempt a RENAME
using the CSTAT facility.

Enter password or Q to Quit==>
```

Figure 10. Expert help system bad queue display.

Ames Research Center Natural Language Cray Help

Your Request? *How do I run a Cray Job?*

ACF 2.2.2 CREATING A JOB x.40

Using an available editor, create a file containing Cray job statements. In the following example, the characters < > delimit user-supplied parameters. There is no "end-of-record" on the Cray, but "end-of-file" and "end-of-data". The /EOF must be in upper case, begin in column 1, with no trailing nulls. Items in [ ] are optional. For default values see ACF Quick Ref Guide Pg 78.

```
JOB,JN=<jobname>,T=<seconds> [MFL=<max_field_length>] [DISK=<megawords>].
ACCOUNT,AC=<account code>,US=<userid>,UPW=<password>.
*    NAME=name, PHONE=phone, EMAIL=email.
*
*          JOB COMMENTS

*  .<Cray control statements>

/EOF                          { /EOF is interpreted as Cray end-of-file }
.
/EOF                          { Cray end-of-file }
  .<file 2 of $IN>
.
/EOF                          { last /EOF is interpreted as Cray end-of-data }
```

Ask another question	<Return or Hot key A>
Copy a sample job to your directory?	<Hot key C>
Return to system?	<Hot key R>

Figure 11. Expert Cray user help system screen response. (a) Response to query, "How do I run a Cray job?.

c

Is file name of ACRAYJOB.LIS ok? *yes* <Hot key y

File ACRAYJOB.LIS has been copied to your directory. Use editor to replace entries in lower case letters.

$ *dir*

Directory RKD0:[CALLAS]

ACRAYJOB.LIS;1 (All but this directory file have been omited from this example.)

$ *ty acrayjob.lis*

```
JOB, JN=jobname, T=seconds. MFL=mfl. DISK=disk.  REQUIRED JOB CARD (1st CARD).
ACCOUNT, AC=account_code, US=userid, UPW=password.  REQUIRED ACCOUNT INFO (2nd)
*   NAME=name,  PHONE=phone,  EMAIL=email.  REQUIRED COMMENT LINE (3rd)

*                              COMMENT LINES

*   THIS IS AN EXAMPLE OF CRAY "COS" JCL
*   COMPILE AND RUN ONLY.. NO UPDATE, NO SSD.
*   lower case words must be replaced with user information
*
CFT,I=$IN,ON=DNX.   COMPILE FORTRAN
ACCESS,DN=DFILE1, PDN=datafilname, ID=your_id.     OBTAIN A DATA FILE
ASSIGN,DN=DFILE1, A=FT01.   ASSIGN FORTRAN ALIAS
ASSIGN,DN=DFILE2, A=FT02.   SETUP FOR OUTPUT TO A FILE
LDR,MAP=OFF,SET=INDEF.   LOAD AND EXECUTE
SAVE,DN=DFILE2, PDN=secondfilname, ID=your_id.     CATALOG FOR FUTURE USE
/EOF
```

Figure 11. (b) Sample job. Continued.

```
      PROGRAM SAMPLE
C THIS SAMPLE PROGRAM IS FILE 1 OF THE $IN FILE
      INTEGER PAR
      PARAMETER (PAR=250), (IPAR=250)
      DIMENSION A(PAR,PAR), B(PAR,PAR), C(PAR,PAR)
      COMMON ABC X(PAR), Y(PAR), Z(PAR)
ASSIGN,DN=DFILE1, A=FT01.
DATA X, Y, Z / PAR*0.0, PAR*25.0, PAR*50.0/

C    THIS LOOP WILL VECTORIZE
      DO 10 I = 1, IPAR
         A(I,I) = X(I)
         B(I,I) = Y(I)
         C(I,I) = Z(I)
   10       CONTINUE
C    THIS LOOP WILL NOT VECTORIZE
      DO 60 I = 1, IPAR
        WRITE (2,1000) X(I), Y(I), Z(I)
   60       CONTINUE
      STOP
  1000  FORMAT (3F12.3)
      END
/EOF
<file 2 of $IN>"
/EOF

$
```

Figure 11. (c) Sample job. Concluded.

DATE	TIME	NODE	USERNAME ANSWERED	QUESTION SOURCE	RECORD	TIME(SEC)
11/02/88	14:58	NEP	PHILLIPS Yes	How do I get more disk space? PR Parsuser		21
03/16/89	09:14	BFM	SYSTEM Yes	SCHEDULE UPTIME DB TABLE6	48	59
04/02/89	23:17	MAX	CALENDAR No	where is the cray located? DB TABLE6		24
04/07/89	12:11	HJA	HJAMGR Yes	cstat ES FINDER	41	
06/14/89	09:15	NEP	AMESFHP No	tss DB TABLE6		7
08/23/89	10:31	RAY	MOEMGR Yes	submit DB TABLE6	25	30

Figure 12. Example system administration output.

NASA
National Aeronautics and Space Administration

Report Documentation Page

1. Report No. NASA TM-102837	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Investigation of an Automated Cray User Services System (Cray User Services "Expert System")		5. Report Date October 1990
		6. Performing Organization Code
7. Author(s) George P. Callas, Michael Younkin*, and Catherine H. Schulbach		8. Performing Organization Report No. A-90193
		10. Work Unit No. 997-04-71-05
9. Performing Organization Name and Address Ames Research Center Moffett Field, CA 94035-1000		11. Contract or Grant No.
		13. Type of Report and Period Covered Technical Memorandum
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546-0001		14. Sponsoring Agency Code

15. Supplementary Notes

Point of Contact: Catherine H. Schulbach, Ames Research Center, MS 200-1,
Moffett Field, CA 94035-1000
(415) 604-3180 or FTS 464-3180

*Control Data Corporation, Minneapolis, Minnesota

16. Abstract

An experimental expert Cray user help system, which incorporates a natural language interface and a commercially available expert-system shell, was developed at NASA Ames Research Center to investigate the feasibility of augmenting the existing user help system. In this report the system is described and the experimental results are presented. A program describing staffing and funding requirements to develop a production system is outlined, and the results of a benefits analysis are presented to assess the potential of an expert Cray user help system.

17. Key Words (Suggested by Author(s)) Expert system Automated help system Natural language interface	18. Distribution Statement Unclassified-Unlimited Subject Category - 61

19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 42	22. Price A03

NASA FORM 1626 OCT 86

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