PROCEEDINGS

FALL 1984

PROCEEDINGS

FOURTEENTH SEMI-ANNUAL
CRAY USER GROUP MEETING

October 1-4, 1984

Hyatt Regency Hotel
San Francisco, California

Host: Technology Development
of California

Ron Deiss, Local Arrangements
Karen Friedman, Proceedings Editor

Prepared for publication and printed at the National Center for Atmospheric Research, Boulder, Colorado (NCAR).*,†

Cover photograph courtesy of Karen Friedman.

Thanks go to Kathy Lucero of NCAR for her assistance in typing various contributions to these Proceedings.

* The National Center for Atmospheric Research is operated by the University Corporation for Atmospheric Research and is sponsored by the National Science Foundation.

† Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

CONTENTS

PRELIMINARY INFORMATION

PRESENTATION ABSTRACTS

WORKSHOP REPORTS

ADDITIONAL REPORTS

Contents (cont.)

ADDITIONAL INFORMATION

CRAY USERS GROUP

BOARD OF DIRECTORS

1984 - 1985

TITLE	NAME	ORGANIZATION
President	Michael Schomberg	AERE-Harwell
Vice President	Helene Kulsrud	IDA
Treasurer	Robert Price	Westinghouse
Secretary	Karen Friedman	NCAR
Member, Board of Directors	Jacqueline Goirand	CISI
Member, Board of Directors	David Lexton	ULCC
Member, Board of Directors	Joe Thompson	LANL

MEMBERS OF THE PROGRAM COMMITTEE FOR FALL 1984, SAN FRANCISCO

Ron Deiss	-	TDC
Mick Dungworth	-	CRI
Jacqueline Goirand	-	CISI
Gary Jensen	-	NCAR
Helene Kulsrud	-	IDA, Chair
Mostyn Lewis	-	Chevron
David Lexton	-	ULCC
Jerry Melendez	-	LANL
David Sadler	-	CRI
Margaret Simmons	-	LANL
Dean Smith	-	ARCO
Mary Zosel	-	LLNL

PROGRAM

CRAY USER GROUP
FALL 1984
SAN FRANCISCO

TUESDAY OCTOBER 2	**WEDNESDAY OCTOBER 3**	**THURSDAY OCTOBER 4**
8:30 Welcome F. Greene (TDC) 8:45 Keynote R. Bailey (NASA-AMES) 9:15 CUG Steering Committee J. Goirand (CISI) 9:30 Report of the Nominating Committee and Elections A. Cowley (NCAR)	8:30 Operations I G. Jensen Performance Evaluation and Optimization J. Goirand	8:30 Graphics H. Kulsrud Operations III G. Jensen
10:30 BREAK	**10:00 BREAK**	**10:00 BREAK**
10:55 Presentation of New Officers 11:00 CRAY Corporate Report D. Whiting (CRI) 11:30 CRAY Software Report M. Loftus (CRI)	10:30 LINCS J. Fletcher (LLNL) 11:00 NETEX HYPERchannel J. Hughes (NSC) 11:30 CRI View of Networking D. Mason (CRI)	10:30 COS D. Lexton CTSS J. Melendez
12:00 LUNCH	**12:00 LUNCH**	**12:00 LUNCH**
1:30 I/O M. Lewis Languages M. Zosel	1:30 Operations II G. Jensen Networking D. Smith Multitasking M. Zosel	1:30 Operations Report G. Jensen (NCAR) 2:00 User Requirements S. Niver (BCS) 2:30 Fortran in the '90s J. Martin (LLNL)
3:00 BREAK	**3:00 BREAK**	**3:00 BREAK**
3:30 Front Ends D. Smith Libraries M. Simmons	3:30 Program Committee H. Kulsrud	3:30 CRAY-2 with CTSS D. Storch (NMFE) 4:00 Digital Scene Simulation L. Yaeger (DP) 4:30 Next Conference
5:00 CLOSE	**4:00 CLOSE**	**4:45 END**

PRESENTATION ABSTRACTS

A LARGE-SCALE COMPUTER FACILITY FOR COMPUTATIONAL AERODYNAMICS

F. R. Bailey

NASA Ames Research Center
Moffett Field, California

NASA has initiated at the Ames Research Center the Numerical Aerodynamic Simulation (NAS) Program. The objective of the Program is to develop a leading-edge, large-scale computer facility, and make it available to NASA, DOD, other Government agencies, industry and universities, as a necessary element in disciplines. The Program will establish an initial operational capability in 1986 and systematically enhance that capability by incorporating evolving improvements in state-of-the-art computer system technologies as required to maintain a leadership role.

The NAS Program had its genesis in the efforts of theoretical aerodynamicists at the Ames Research Center during the mid-1970's. Its purpose was to exploit computational aerodynamic research to improve the basic technology of aircraft and aerospace vehicle design. The combination of improvements in aerodynamic flow simulation modelling and supercomputer performance have permitted the rapid advancement of computational aerodynamics to the stage where it is practical to address non-linear inviscid flow simulations about complex aerodynamic configurations, and detailed turbulent flow simulations about simple configurations. The necessary next step in the advancement of computational aerodynamics is to combine detailed turbulent flow models and complex configurations, to provide the capability to accurately simulate realistic designs over the entire flight envelope.

The goal of the NAS Program is to provide the necessary computer throughput, memory capacity, and system architecture to reach this next step. In particular, the NAS Program will achieve a supercomputer capability of 250 million floating-point operations per second sustained throughput and a 256 million-word memory capacity in 1985/86, and a four-fold increase in these capabilities by 1987. This new supercomputer capability includes the Cray-2, which will be integrated into the NAS Processing System Network (NPSN). The NPSN is designed to support supercomputer-generated, large-scale aerodynamic flow simulations with powerful scientific workstations, mass storage facility, high-resolution graphics facility, and high-bandwidth remote communication links to remote sites nationwide. To aid in its productive use, the system will provide unique software features including a UNIX™ operating system environment across all systems, common graphics services, a library of aerodynamic simulation models, and a rich set of local area network data communication services.

CRAY RESEARCH CORPORATE REPORT

Don Whiting

Cray Research, Inc.
Mendota Heights, MN

1983 HIGHLIGHTS

- $169,690,000 TOTAL REVENUE
- $26,071,000 EARNINGS
- 1551 EMPLOYEES
- 65 SYSTEMS INSTALLED
- 25 SYSTEMS ORDERED
  - 15 DOMESTIC
  - 10 INTERNATIONAL
- 20 SYSTEMS INSTALLED
  - 9 DOMESTIC
  - 11 INTERNATIONAL
- 3 NEW COUNTRY INSTALLATIONS
  - SWEDEN - SAAB
  - CANADA - CMC
  - NETHERLANDS - SHELL

MID-1984 HIGHLIGHTS

- 14 SYSTEMS ORDERED
  - 10 DOMESTIC
  - 4 INTERNATIONAL
- FIRST CRAY-2 ORDER
- 11 INSTALLATIONS

1984 HIGHLIGHTS

- X-MP ANNOUNCEMENT
  - X-MP/11, 12, 14
  - X-MP/22, 24
  - X-MP/48
- SSD-5 128 MW
- DD-49 DISKS

1984 EXPECTATIONS

- 26 CONTRACTS TOTAL
  - 16 DOMESTIC
  - 10 INTERNATIONAL
- 30 INSTALLATIONS TOTAL
  - 21 DOMESTIC
  - 9 INTERNATIONAL
- 84 SYSTEMS INSTALLED
- 59 CUSTOMERS

CRAY SOFTWARE STATUS

Margaret A. Loftus

Cray Research, Inc.
Mendota Heights, Minnesota

The following software has been released since the last Cray Users Meeting.

2.0 VM Station
1.12 MVS Station Bugfix 2 including support for MVS/XA
1.12 COS/1.11 CFT Bugfix 3
2.03 VAX Station
1.13 NOS Station
2.0 PASCAL
1.13 COS/CFT
1.13 Bugfix including support for single processor CRAY X-MP

Feedback on the 1.13 Release has been very good and ease of installation has been excellent. Several sites are now running 1.13 in production and many others are in the process of upgrading.

The major CRAY software effort today is the 1.14 Release. We are beginning the stabilization phase for the 1.14 Release and plan to release the software by the end of 1984. It is another very significant release containing the following major features:

- CRAY X-MP 4 Processor
- 8M words
- New X-MP/4 instructions - compressed index and gather scatter
- Additional CFT performance improvements
- DD49 support
- On-line tape enhancements
  - End of volume user processing
  - Multitape mark read/write
  - 2.5 M Byte tape blocks
  - Partial implementation of IBM multifile (process any single file of a multifile)
  - Installation options for enforcement of ring processing
  - Tape positioning and back space
  - Parallel tape mounting
- Dispose/Acquire from 80 M Byte disk
- Support for new IOS expander peripherals (disk and tape)
- The portion of the Subsystem project that provides networking hooks
- Support of CRAY X-MP hardware performance monitor
- Large SSD support (128 M wds)
- Two very high-speed channels to the SSD
- A multitasking tool - common block cross reference (FTREF)
- Update PL audit utility
- FORTRAN callable SORT package
- Non-ANSI flagger
- Contiguous disk allocation options

The status of other major software efforts is as follows:

<u>NFT</u> - Working on optimization, progressing very well, initial release is still planned for end 1985.

<u>Permanent file archiving</u> - planned for 1.15 release.

<u>Apollo Station</u> - planned for 1Q85 release.

<u>Interactive CYBER NOS Station</u> - planned for 2Q85.

<u>C under COS</u> - 1st half 1985 release.

<u>ISP</u> - available end 1984 for beta testing; release 1Q85.

<u>On-line tape enhancements</u>

- full multifile support
- optional automatic volume recognition
- IBM 3480 support

<u>SUPERLINK</u>

- New MVS Station based on ISP protocol

<u>CRAY-2</u>

- A single processor CRAY-2 is installed in Mendota Heights.
- We are in the process of bringing up UNIX on the actual hardware.

LINCS

John G. Fletcher

Lawrence Livermore National Laboratory
Livermore, California

LINCS -- the Livermore Interactive Network Communications Structure (or Standard or System) -- is a hierarchy of communication protocols designed at the Lawrence Livermore National Laboratory and used by the Laboratory's Octopus computing network and the extension of that network called Labnet. Since meaningful communication requires some degree of commonality between the sender's and the receiver's view of their environment, LINCS defines salient features of that view. The view conforms to an underlying philosophy about the nature of computing and computer-oriented communication, aspects of which are discussed here. Further information about LINCS and a more extensive list of references may be found in an earlier paper [1].

Although LINCS has been designed to be usable in any general purpose computing environment (and many special purpose ones), a brief summary of the Octopus/Labnet environment is appropriate. The network operates continuously (24 hours/day, 365 days/year), serving over 2500 mostly scientific (but also clerical and administrative) users. The network's computers and peripherals are of heterogeneous manufacture and use varied native operating systems; these change continually as technological advance and expanding needs dictate the acquisition of higher-performance facilities and the retirement of obsolescent ones. Current equipment includes a Cray XMP, four Cray-1's, three CDC 7600's, and a multitude of smaller computers. On-line storage exceeds three terabits, mainly because of a Braegen automated tape library and a CDC 38500. Output amounts to about 15 megapage equivalents per month, chiefly from two IBM 18000 line/min printers and five FR80 microfilm recorders. There are between 1600 and 2000 remote terminals, ranging from simple hardcopy terminals and "dumb" softcopy terminals to "smart" workstations and personal computers. Communication speeds extend up to the 50 megabits/sec of the NSC HYPERchannel.

This work performed under the auspices of the U. S. Department of Energy by the Lawrence Livermore National Laboratory under contract number W-7405-ENG-48.

UNIFORMITY

A primary component of the LINCS philosophy, the one emphasized in this paper, is that the environment, both for (human) users and for programs, should be as uniform and homogenous as possible. In particular, above some level there should be no distinction between a program's accessing a resource or a service that is remote (i.e., on another computer across the network) versus one that is local (i.e., provided by another process or the system within the customer's computer). This is achieved by having local services be accessed in the same way as is familiar for remote services, namely by sending and receiving messages containing requests, replies, and data, rather than by making system "calls" specialized to each service. This means that a customer seeking to access, say, a file does not need two sets of procedures, one to be used when the file is local and the other when the file is remote; similarly a server, such as a file server, does not need differing procedures for responding to local and remote customers.

At a higher level there also should be no distinction between an access that is internal (i.e., provided by procedures, possibly library procedures, within the customer's process) versus one that is external (i.e., local or remote as defined above). This does not mean imposing on the external accesses the discipline of the familiar internal subroutine call; namely, invoke the service, and then wait until the service is complete and the result obtained before continuing (although this characterizes many so-called "remote procedure call" mechanisms that have been proposed). Instead it means that internal accesses should have the two generalities often associated with external accesses. First, the services should be available, not only as synchronous routines (where the customer and the server execute in sequence), but also as asynchronous tasks (where the customer and the server can execute concurrently in parallel). Second, it should be possible to interact repeatedly with a server that remembers what the customer is "talking about" (saves its state) between one interaction and the next; that is, coroutines and cotasks, as well as subroutines and subtasks, should be available.

Needless to say, high-level languages that are generally available are not up to this job, and internal tasks and coroutines must be provided in ways that lack the naturalness of subroutines. We are using the language C most extensively because of its expected wide (but, alas, not universal) availability, even though coroutines and tasks must be provided with the aid of assembly language routines that take some effort to develop for each new computer. If another language with significantly better features becomes generally available, we might well change our preferred language. Old-time assembly language buffs among us are somewhat astonished that high-level languages widely touted as being for systems implementation lack valuable features, such as coroutines, that are always readily available in lower-level languages.

COMMUNICATION

With sub- and co-routines and tasks all available in the internal environment of a process, our intent (not yet fully achieved) is that the same constructions and formats be used for external accesses. Of course what is really happening is that the invocation of a routine or task is encoded in some fashion into a request message and sent to the external process; later a reply message is decoded into the results of that invocation. The activity of forming and transmitting messages is carried out by the six lower layers of the seven-layer LINCS protocol hierarchy, the seventh and highest layer, the application layer, being the one that carries out the logic of the customer or server.

LINCS has adopted the nomenclature of the seven-layer OSI Model. However, that should not lead one to suppose that we adhere to the developing standards built on that model. A reader of the defining document of the model [2] will find the part that LINCS has used (Section 6.2) on a few pages in about the middle of the document, following a general discussion of layering that implies an excessive and impractical similarity in structure among layers. These few pages define, in a brief paragraph each, the functionality of the seven layers. Unfortunately the document does not stop there but goes on, twice more (in increasing detail) describing specifications for each layer, specifications which the developing standards tend to follow. These specifications depart from the principles of the earlier parts of the document. In particular, certain functions are repeatedly performed by two or more successive layers, apparently in the belief that one layer cannot do the job well enough. The result is that the standards tend to be overly complex, unnecessarily redundant, unacceptably inefficient, filled with superfluous options, and lacking a coherent point of view. All this no doubt arises from design by committee.

LINCS assigns each communication function being performed to exactly one of the seven layers, with the third (network) layer being a watershed. This layer performs exactly those functions, in particular routing, that must be carried out in a uniform way by every node through which a packet passes as it "hops" from node to node across the network. Lower layers (link and physical) are concerned only with the channels that connect adjacent nodes to provide single hops (including multipoint channels, such as HYPERchannels and Ethernets, which often are given the misnomer of "local area network"); these protocols can differ from channel to channel. Higher layers (transport, session, presentation, and application) concern only the two end processes of a transmission.

MESSAGES

The interface of the application layer to the presentation layer below consists of several primitives. These include open (a message stream to or from another process) and close (such a stream). Sending is accomplished with two primitives: send, which specifies the information to be sent, and obtain, which waits until the application can reuse the storage containing that information. Having two primitives (which in many situations would be executed in immediate sequence as a macro) provides two benefits: queueing of data buffers can occur in the application's memory without blocking the application, and data buffers can be allocated either by the application or by the presentation layer (depending on the order in which send and obtain are issued). Similarly, receiving uses two primitives: give, which specifies where received information is to be stored, and receive, which waits until the information is available. Other primitives include abort (sending or receiving) and status (of an open stream).

When an external routine or task is invoked, the parameters passed in the invocation are encoded by the presentation layer into a byte sequence, placed into a message buffer, and sent; the receiver's presentation layer decodes the sequence and stores the parameters obtained in memory, where the invoked module can access them as variables. LINCS uses an encoding in which each parameter is encoded as a token, a sequence of bytes consisting of four fields: type (e.g., integer, character string, capability), usage (e.g., file identifier, first bit address, count), length (how many additional bytes make up the token), and value (those additional bytes, expressing the value of the parameter in a computer-oriented form, e.g., binary for integers). The routines that parse and generate tokens are fairly simple and efficient for typical machines (although would become more complex on less typical machines, such as decimal machines or ones with a word size not a multiple of 8).

We believe that, because of this simplicity and efficiency, token encoding is far more appropriate for interprocess communication than are

ASCII character strings of a form such as might be transmitted to or from an interactive terminal, although such is a common practice in many places (largely for reasons of habit and history). Most computing processes should not have to be burdened with deciphering forms of expression oriented to the needs of human beings any more than most human beings should be burdened with deciphering computer-oriented forms. The interface,being the binary integer/abstract data structure world of computing processes and the modified natural language world of human computer users,is provided by specialized processes called command language interpreters, which translate between the two forms of expression.

PORTABILITY

Another aspect of the uniformity that LINCS tries to provide is portability: applications, libraries, and systems programmed for one kind of computer should also be able to run on another. Using a high-level language is only part of meeting this need. The other part is providing a uniform environment for the high-level procedures. To this end we have defined SMILE -- the System/Machine-Independent Local Environment. The idea is that this environment can be implemented, perhaps using assembly language, both on "bare" machines and layered on top of existing operating systems; then procedures designed to run within the environment can execute in all such situations without change. The environment is at a relatively low level; message passing facilities, for example, are not part of the environment but are something implemented in high-level language within the environment.

The primary facility supplied by SMILE is memory, either real or virtual, depending on the machine. However, because such events as page faults require process rescheduling, SMILE must also provide multitasking. SMILE consists of about a dozen primitives, such as (for dealing with memory) allocate (a block) and deallocate, and (for dealing with tasks) sleep, wake, create (a new task), and synchronize.

SUMMARY

We of course believe that our design is best for our needs. Although we also believe that it would be suitable for many other needs, we realize the improbability of our approach being widely adopted by others in the near future. In particular, we have no reasonable expectation that a vendor will supply a system such as described here. However, we do believe that it is reasonable to ask that vendors (e.g., Cray) provide systems with the following characteristics:

- They are modular, so that we can easily discard the modules that do not conform to our needs and replace them with our own, while keeping those (particularly low-level hardware-oriented ones) that are usable. It is sad to contemplate, amid all the discussion about structuring programs and organizing programming, that the one idea that stands head and shoulders above all the others, the one idea that most surely is correct, namely modularity, is the one that seems to be ignored most widely. A commercial system all too often is a monolithic, tangled conglomeration, modifiable only after great effort and through the use of many special cases. We ask that this cease to be so.

- The modules provided should have clean, simple, well documented interfaces.

- The design should conform to a clear, parsimonious, easily understood model; it should not be a heterogeneous collection of programs such that knowing how one thing is done gives one no clue as to how other, similar things are done.

- A really good, widely available systems implementation language should be provided.

References

[1] Watson, R. W., "Requirements and overview of the LINCS distributed operating system architecture," Cray User Group Proceedings 13 (April 13, 1984), pp. 24-33.

[2] ISO/TC97/SC16, "Data processing -- open systems interconnection -- basic reference model," Computer Networks 5, (April, 1981), pp. 81-118.

CONTINUOUS PROTOCOL FOR LONG DISTANCE COMMUNICATION LINKS

J. Hughes

Network Systems Corporation
Brooklyn Park, Minnesota

Abstract

This paper describes a continuous protocol for use with long distance, full duplex communications links. Though mainly designed for the transmission of data across a satellite link, this protocol will work very well on any high-speed, full duplex communications link with or without large propagation delays. Discussed are connection, data transfer, record numbering, acknowledgment message format, flow control, equilibrium conditions and error considerations.

Introduction

This paper describes the data transport protocol of NETEX.[1] NSC has decided to use this protocol for all transport connection regardless of the actual media or its propagation delay.

A prototype of this protocol was first used in 1981 to check out NSC capabilities for digital satellite communications on an SBS satellite.[2] This prototype was very effective in proving the feasibility of high-speed (up to 6.3Mb), high efficiency (95%) satellite data transfer.

The objectives of this protocol are to provide a reliable, high speed means of data transmission that can function on varied communication configurations, provide efficient data movement during periods of high errors and be tolerant of temporary line dropout.

A satellite configuration can be considered a long distance, full duplex communication link. The satellite link is sometimes referred to as a "transmit and pray" link, because before one record is acknowledged, the next one is sent. With a six megabit satellite link, as many as four million bits may be in the air between any two ground stations at a time.

Application Interface

Figure 1 shows the primitives for communicating with NETEX.[3]

NETEX is a communication subsystem with complete responsibility for data transfer. This allows the applications that drive NETEX to reuse its buffers when their data is safely in NETEX, even before the data is sent. There are two implications of this type of preacknowledgment:

1. Preacknowledgment means the application does not need to worry about the intermediate steps of data delivery. This allows the transmitting application to fill the communication channel without complex buffering.

---

1 NETEX is a trademark of Network Systems Corp.

2 A good primer on satellite protocols, and their pros and cons can be found in an SBS system brief entitled "Satellite Transmission Protocol For High Speed Data" by James L. Owning.

3 Further information about the NETEX interface can be found in the NETEX general information manual (NSC publication 43999069) and the NETEX implementation manuals (i.e. NETEX for MVS, publication 4399H210).

NETEX Application Interface

Primitives

OFFER	Defines a service
CONNECT	Connects to a service
CONFIRM	Accepts a connect
WRITE	Sends data
READ	Receives sent data
CLOSE	Sends the last data
DISCONNECT	Aborts a connection

Figure 1: NETEX primitives

An example of this would be a file transfer application which has several million words of data to transmit. The application can issue write after write without any concern about the delivery or buffering of the data.

2. Preacknowledgment also requires the receiving application program to acknowledge that it has successfully completed its task (not just that it received the data).

   An example of this is a banking transaction application where a transaction is received, and then processed. The reception of the transaction is an intermediate step in its processing, what we really care about is a peer-to-peer acknowledgment that the transaction is processed. We cannot assume that if the transaction was delivered it will be processed. This requires the receiving application to notify the sending application of the transaction's completion.

Connection

Connection in this protocol consists of sending the connect and waiting for something to come back from the other side. Figure 2 shows the connect sequence. Normally the connect message is sent and a confirm message is received (relating to the connect and confirm primitives).[4] Until the connection is established, a synchronous protocol is used. This causes each transport connection to suffer two propagation delays during establishment.

Error recovery during the connect phase involves repeatedly sending the connect message. If the connect message is lost, resending the connect message accomplishes the desired recovery. A special case involves the failure of the return channel. If the return channel fails, the first connect completes, and the confirm from the other side is lost. The result of this is that the subsequent connect attempts are ignored because the connection is established. The retransmission of the confirm follows the standard error recovery associated with the continuous protocol (this is because the connection is fully established when the connect message is processed). This change over from synchronous to continuous must be accomplished with care.

Data Transfer

Until the connection is complete, the protocol is not a continuous one. Once the connection is complete a different environment is in effect.

NETEX is a full duplex communications vehicle. The protocol is made up of two simplex connections with very little overlap. To understand the protocol I would like to discuss it in terms of its simplex pieces. In a simplex configuration, the data flows from the transmitting application, to the transmitting NETEX (transmitter), to the communication configuration, to the receiving NETEX (receiver), and finally to the receiving application. The terms transmitter and receiver will be used later to describe the two simplex pieces of NETEX.

NETEX represents the access method and is responsible for the end-to-end protocol. That is, once the transmitting application gives the data to NETEX, it can forget about that record and reuse its buffer. Also, the receiving application's only responsibility is to accept the data from NETEX and send a peer-to-peer acknowledgment when the application data has been processed.

Record Numbering

This protocol provides each record with two sequence numbers.

Each record from the transmitting application is given a logical record number (LRN). With this number the receiver can ensure that the receiving application gets the records in the proper order. The logical record number is also used in flow control.

When the record is transmitted. it is given a different number called the physical block number (PBN). This number is unique for every retransmit of a given block. The advantage of this scheme is that if a record is retransmitted, it is given a new physical block number. This allows a record to be ACKed or NAKed more than once, without confusion.

By using 16 bit binary values for PBN and LRN and the ACK scheme discussed later, this protocol can have 2^{16} blocks outstanding at any one time.

---

[4] The actual connection sequence that goes on in NETEX is more complex than this; only the transport service is discussed here.

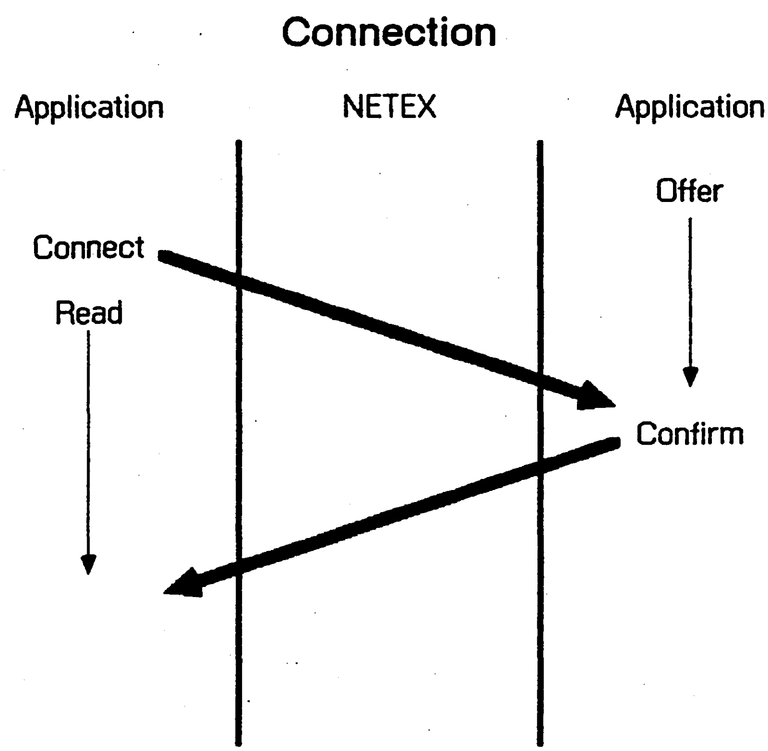

Figure 2: NETEX Transport Connect

Acknowledgment Message Format

In other protocols, lost acknowledgments represent a significant problem. In the synchronous protocols a lost ACK or NAK causes a timeout to occur. In restart on error protocols, lost NAKs cause timeouts. This protocol is designed to be tolerant of lost acknowledgments.

To make the protocol tolerant of lost ACKs, every ACK is stand alone and they can be sent at any time, even multiple times. Every ACK message, by itself, is capable of reestablishing synchronization between both sides of a connection. Any number of ACK messages can be lost, but when one does make it over the link, it contains enough information for resynchronization.

The acknowledgment message contains several fields, among them are:

1. The last physical block (PBN) received.
2. A bit map representing the status of the most recent received block and the 15 blocks with successively lower PBN numbers (0 = ACK, 1 = NAK).
3. The highest logical record number (LRN) that the receiving transport is willing to accept. (This is called the proceed.)
4. The last physical block (PBN) that this side has sent. (This is called the chase PBN.)

These items are kept current in the receiver, and sent whenever one of several conditions are met:

1. Whenever outbound data is sent, it contains the inbound data's ACK information. (The inbound data's ACK information is "piggybacked" with the outbound information.)
2. A separate ACK message is sent if the highest PBN received off the link changes by 2 since the last time an ACK was sent.
3. A separate ACK message is sent if the highest LRN we can accept changes by 2 since the last ACK was sent.
4. A separate ACK message is sent if one second has expired since the last ACK and there is data outstanding. (This is a message chase.)
5. A separate ACK message is sent if one second has expired since the last ACK and a block remains unACKed.
6. A separate ACK message is sent if 30 seconds have expired since the last ACK. (This is an idle.)

Basically, whenever the receiver feels like sending an ACK, one is sent. Multiple ACKs are not a problem (except for performance). The timeout values used above are for a typical line with a propagation delay of less than .4 second. In NETEX, these timer values can be adjusted.

Skipping an ACK transmission (possibly due to an error condition) does not present problems because the last sixteen physical blocks of ACK/NAK information are returned in every ACK message. Succeeding ACK messages will provide the information previously lost.

The transmitter keeps all records waiting for an acknowledgment on a queue called the ACK queue. When ACK information arrives, the bit significant information is processed in such a way that if the bit indicates ACK, the buffer that was transmitted with that PBN is removed from the ACK queue and released. If the bit indicates NAK, the block is resent with a new PBN assigned. (If the block is not on the ACK queue, it is assumed that the block was ACKed or NAKed by some previous ACK information.) All blocks with a PBN older than the oldest bit significant information are considered NAKed and retransmitted. This indiscriminate retransmission can only occur after large line outages.

The last item in the ACK information is the chase data. In the ACK information, the last PBN sent provides an indication of the highest PBN that should have been received. This field is used to update the ACK information and (if necessary) send an ACK.

The typical ACK scheme is graphically represented in figures 3 and 4. These show the timing of when and how the ACK information gets back to the sending transport. It is important to note that timely acknowledgment is not mandatory in this protocol. The only purpose of positive acknowledgment is to free buffers. Applications are not directly affected by the freeing of buffers.

The protocol adapts to the application to reduce the number of standalone ACK messages that are required. Applications that use one-way data transfer see a reduction of 50% because every other block is ACKed. The skipping of blocks is not a problem because multiple blocks are ACKed at the same time. Applications that are synchronous see a 100% reduction of standalone ACK messages because the ACK information is sent along with the outbound data.

ACK Scheme

One-way data transfer (stream)

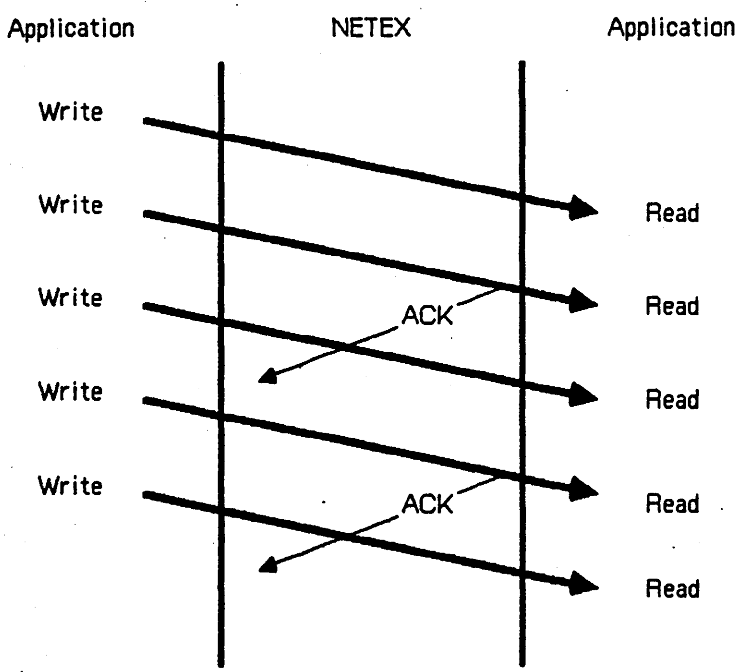

Figure 3: ACK scheme, Data Stream

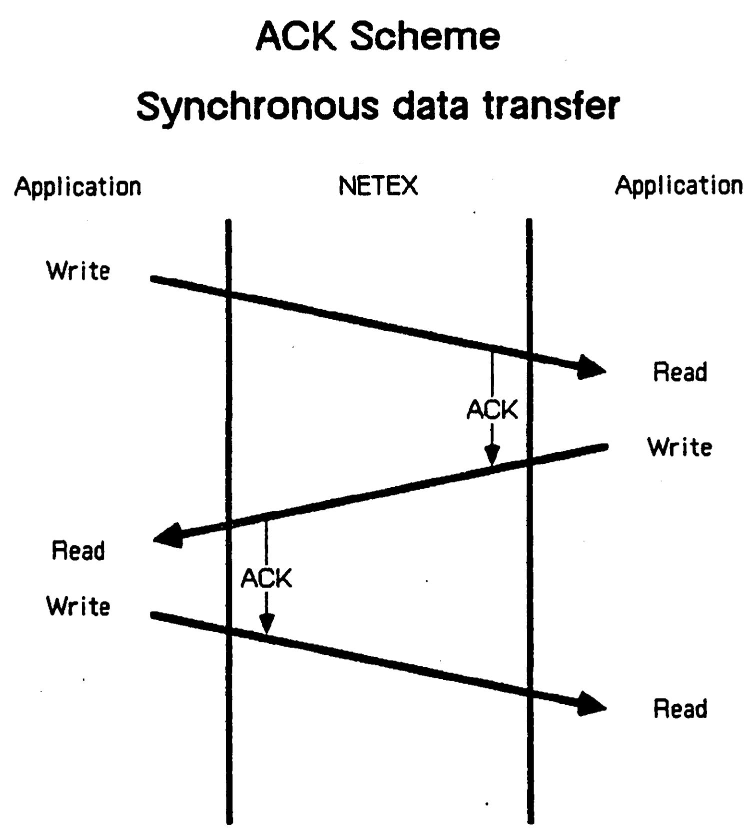

Figure 4: ACK scheme, Synchronous

Flow Control

There are two problems associated with flow control:

1. Overdriving the transmission link, and
2. Overrunning the receiving application.

Overdriving the transmission media is possible in some multiple link configurations. This is due to the differences in transmission speeds of the various links. If the second or third link in the configuration is slower than the first link, the first link can overdrive the slower link. The solution is to "throttle" the transmissions to the rate of the slowest link in the configuration.

Overrunning the receiving NETEX is not possible with this protocol. The receiver queues all of the records received across the link in logical record number order, until the receiving application takes the data. The transmitter is designed so that it will never transmit a record with a logical record number larger than the receiver can accept in its buffers. The transmitter is provided with the receiver's buffer status via the proceed field in the ACK information. As long as the proceed information is increased and never decreased, there will always be buffers available.

As an example, assume the receiver has 20 buffers, and the last record given to the receiving application is numbered 115. Since the receiving transport knows it has 20 buffers, it sends 135 as the proceed in the ACK information. The transmitter will then transmit only up to the last proceed that it received in the ACK information.

NETEX will allocate enough buffers to handle twice the propagation delay. This amount of buffering is enough to keep the link as close to a 100% effective data rate as possible on an error free line. A minimum of three buffers will be allocated for lines with low propagation times. To calculate the number of buffers, the following formula is used:.

ceiling((2*prop*linerate)/(blocksize)+1) = NumBuffers

Twice the propagation delay is required because it takes 1 propagation delay to get the data to its destination, another propagation delay to return the new proceed count. One is added to compensate for the fact that ACKs are sent after every other block is received. The other factors are the time it takes to send a block (assuming a 100% effective line).

For example if the line is T1 (1.544Mb), application blocksize is 32K bits and propagation delay .312 seconds (to and from a geosynchronous satellite).

ceiling((2*.312*1544000)/(32768)+1) = NumBuffers

ceiling(30.4) = 31

Equilibrium Conditions

An application that transmits data one direction (i.e. a file transfer) has three distinct equilibrium conditions which can occur. I am using a single direction data transfer because it is the simplest use of a satellite that can efficiently utilize the line. Totally synchronous applications cannot utilize the link efficiently.[5] Multiple synchronous applications can share the same line and (when there is enough activity) efficiently use the line.

1. A "transmitting application bound" condition occurs when the transmitting NETEX does not get enough work to keep the link busy. This condition is characterized by the receiver being empty and the transmitter having only a small number of records waiting on the ACK queue, without any records ready to transmit.
2. A "receiving application bound" condition is characterized by the transmitter having to suspend the transmission of data because the receiving application is not taking the data fast enough from the receiver. The receiver's buffers will be stacked, and the transmitter will have its buffers ready to transmit.
3. A "link bound" condition occurs if the sending application provides the data faster than the link can take it, and the receiving application accepts the data faster than the link can send it. This is characterized by the receiving NETEX being empty. Whenever a record is received, it is immediately passed to the waiting application. The sender's ACK queue contains two propagation delays worth of data.

Ideally, enough buffers should be available on both sides of the link so that the ACK queue is

[5] Without knowledge of the peer-to-peer communication that is going on, one cannot preacknowledge the far application's response, and therefore cannot eliminate the two propagation delays that the real data takes to go across and get the result back.

never full. The number of buffers that are in the ACK queue at any one time is equal to the number of buffers transmitted during two propagation times: the amount of time it takes to send a buffer across the link and receive an acknowledgment back. The maximum number of buffers that the transmitter can ever have in the ACK queue is equal to the number of buffers in the receiver.

Error Considerations

The two major problems associated with long distance communications links are large propagation times and undetected transmission errors.[6]

There are basically two types of errors which can occur on the link:

1. Random errors

   Random errors may occur in isolated records of data. The method the protocol uses to handle such errors is message chase. Message chase is used to inform the receiver that a data record has been dropped. By using an ever increasing PBN ensures that if a block is dropped, the next block's PBN tells the receiving NETEX (and hence the transmitting NETEX) that there was an error and which block it was.

   If the transmitting NETEX doesn't have additional user data to chase the transmitted data with, and an ACK has not been returned within one second, the sending NETEX sends a chase message to notify the receiving NETEX of the lost block condition.

   During the time that the receiver is waiting for the bad block to rearrive, two propagation delays worth of data will be received. This data cannot be given to the receiving application because its LRN is above the block in error. This data is held pending the retransmission of the bad block. When the bad block is received, it is placed back in the logical sequence, and the application is given the data as fast as possible.[7]

2. Lost communication

   Lost communication is caused when the link stops working for an extended period of time. Satellites are extremely susceptible to these types of outages. Rain and thunderstorms are typical reasons for a satellite line to temporarily stop working.

   When the transmitter notices that it has not received ACKs for a certain period of time (i.e., 30 seconds), it sends out an idle message. This idle message tickles the receiver letting both sides know that they are still functioning. In addition, if nothing is heard in an extended period of time (i.e., 90 seconds), the receiver knows that communication has been lost and alternate paths may be tried.

   If communication is interrupted, the transmitter continues to send an idle message every timer interval to see if the link has been reestablished. When the line is reestablished, (for example turning up the gain on the ground station), the resulting idle message contains enough information for the transmitter to clean house and get started again.

Random errors in the equilibrium conditions discussed before pose certain performance implications.

If the connection is link bound, random errors cause the link to temporarily dry up. During the time that the receiving NETEX is waiting for the block in error, the blocks with higher LRN are queued up. When the block in error arrives, that block and all the blocks that were queued up waiting for the block in error are given to the application as fast as the application can take it.

The net result of this error is an extension of the process by two propagation delays plus one block time. The two propagation times are caused because by the time the transmitter learns that the block is in error, the transmitter has already sent the full proceed. The two propagation times are the time necessary, after the delivery of the block in error, to tell the transmitter of the new proceed.

The restart on error protocol extends the process by the same two propagation delays. The difference between the restart on error protocol and the continuous one is that the time extended on the

---

6 A paper which exhaustively discusses HYPERchannel satellite hardware is "Performance of HYPERchannel Networks, Parameters, Measurements, Models and Analysis, Part II, Satellite Connected HYPERchannel Networks" by J. Heath and W. R. Franta, October 1982.

7 One way to improve high-speed lines is with forward error correction (like SECDET in memory). This topic is discussed in a paper entitled "Forward Error Correction In The SBS Satellite Communication System" by Curtus L. Greer, Satellite Business Systems, McLean, VA.

continuous protocol is line time that is available for other productive use. The time used by the restart on error protocol is for retransmission of blocks that were received out of sequence.

If the number of available buffers were doubled, the effect of a single block in error would be limited to a single block transmit time. This is because the transmitter will not reach its proceed until four propagation times, and by that time the receiver will receive the bad block and will have updated its proceed and notified the transmitter. By lying about the line's propagation time (telling NETEX that it is twice as long as it really is) will cause NETEX to double the number of buffers and provide much better throughput on a dirty line.

Random errors occurring on either transmitting or receiving application-bound connections have no effect on a file transfer time. Retransmission times are overlapped with the time that it takes for the application (the one which is slowest) to handle the data.

Summary

This protocol was designed to efficiently utilize high-speed, long propagation lines which can lose messages.

Continuous protocol provides the ability to keep data moving efficiently over long distance high-speed lines.

Link independence is provided by the protocol's obliviousness to the hardware's characteristics. This allows multiple links of different types to be used concurrently without problem.

Double numbering provides each physical transmission with a unique number so that it can be NAKed multiple times without causing extraneous retransmission.

Standalone acknowledgments provide the transmitter enough information to completely resynchronize on any acknowledgment message.

Adaptive protocol provides a means for reducing or eliminating standalone ACK messages by adapting to the application.

CRI VIEW OF NETWORKING

DON M. MASON

CRAY RESEARCH, INC.
MENDOTA HEIGHTS, MN

The CRI View of Networking presentation summarizes CRI's network strategy over the past two years, discusses recent Cray Operating System changes supporting this strategy and poses several questions which may lead to future change in the strategy.

CRI's current strategy calls for:

- Continuing development and support of stations.
- Following network standardization activity.
- Providing "hooks" in COS for customer sites or third parties to add network support.

CRI continues to support station development by enhancing current stations and developing new stations. Enhancements are being made to the VM, NOS, NOS/BE, VAX, MVS, DG and IOS stations. New developments for Apollo, Superlink/MVS and UNIX stations are underway.

CRI continues to follow network standardization activity. Although our limited resources keep us from being active committee participants, we do keep abreast of new and emerging standards including defacto standards established by industry use.

Beginning with COS 1.14, several networking hooks are being incorporated into the system to support such applications as Network System Corporation's NETEX. Subsystems is the feature name given to these software hooks. Subsystems is intended to provide system services which enable operating system - like functions to be implemented as user codes.

The following are the major subsystem capabilities:

1. Event Recall (in COS 1.14)
2. Inter-Job Communication (in COS 1.14)
3. Channel Driver/Shell Driver (in COS 1.14)
4. Job Operator Communication (future)
5. Enhanced Queue Manipulation (future)

Event recall gives a user job the ability to initiate several events and give up control of the CPU. Upon event completion, the user job is recalled and may easily identify the event causing the recall. Events include channel driver I/O completion, inter-job transfer complete, timer expired or operator message received (future).

Inter-job Communication is the ability for two or more user jobs to pass data amongst themselves. The data transfer will be direct from one job's memory to another job's memory if both are simultaneously resident. If the jobs will not fit together in memory, the data transfer is buffered in system buffers. Jobs may adjust pointers dynamically to control the movement of data to different local buffers in their field length.

A data transfer requires simple steps. First, one job must tell the system it is willing to communicate. Second, another job must attempt to open a path to the willing job. The willing job may accept or reject the open request. If accepted, both jobs may communicate symmetrically.

The Channel Driver/Shell Driver is a set of capabilities which give a user job direct access to a CRAY channel. The Channel Driver is a set of user job interface functions which are used to pass data to and from the site-supplied IOS driver. Functions include read, write, read ahead, write behind and a general function. All but the general function (which is passed directly to the driver) are used to move data. Read ahead causes two reads so that the second data block may be immediately available for the next read or read ahead.

The Shell Driver provides an environment for a CRAY site to write its own IOS driver (currently in the MIOP only). In the past CRI has discouraged sites from writing IOS code which may interfere with the dynamics of other CRI supported data transfers. The Shell Driver manages all data transfer to and from the IOS and fields IOS interrupts. The user site-supplied driver only needs to concern itself with the functioning of the device on the CRAY channel.

Job/Operator Communication is a capability for the user job to exchange messages with the operator terminal. Communication supported includes: information only - a job message sent to a terminal, reply requested - a job message sent to a terminal with a reply solicited and unsolicited message - a message sent to a job without request. Events may be associated with message receipt for the last two types of message.

Enhanced Queue Manipulation is the ability for a privileged user job to manipulate CRAY queues (i.e. input or output queues). Datasets may be added or deleted from these queues. Search criteria are provided to examine the queues. For example, a queue may be searched for all entries matching a particular station ID. The capability may also be used to kill, drop or rerun jobs.

Network Systems Corporation has initiated development of level 2 NETEX for the CRAY. NETEX will be implemented using the CRI subsystem features.

In order to evaluate our networking direction for the future, CRI will be holding a CRAY Technical Council meeting in January of 1985 focusing on the networking topic. Our field organizations will be asking for CRI customer input on their future network needs. Information is needed on:

- The elements customers foresee being in a network which includes CRAY(S). Functionality and data rates of these elements are important parameters.

- The specific hardware customers foresee in a network, including both specific devices and interface hardware.

- The operational environment customers foresee, including user interface, system interface and security.

Fortran in the 90's

Jeanne Martin

Lawrence Livermore National Laboratory
Livermore, California

There are many who expect the use of Fortran to die away by the end of the century, if not sooner. There are others, who perhaps are responsible for hundreds of thousands of lines of Fortran in heavily used production codes, who want it to exist unchanged forever. The Fortran standards committee, ANSI/X3J3, has plans based on neither of these scenarios. Instead they envision an orderly evolution of the language to meet the needs of scientific programmers in the computing world of the future.

* * *

The first Fortran standard appeared in 1966; the second (although it has been called Fortran 77) actually appeared in 1978 and became an international standard in 1980. If X3J3's current milestones are maintained, the proposed draft of the third standard will be completed in late 1985. Ordinary processing of a standard generally requires two years following completion of the first draft. This 2-year period includes time for public comments to be received and processed by X3J3. The processing of comments is likely to result in revisions being made to the draft standard. The revised draft is then reviewed by the public, and if it is deemed acceptable by a general consensus, it will be adopted and published. This means that, if all goes well, the earliest possible date for the next standard is 1987. A more likely date is 1988. Given a couple of years for production quality compilers to be put in place, the effective years of this standard will be the 1990's.

It is clear that a scientific programming language designed in the late 50's is not adequate for the 90's. Newer programming languages have come into existence that are "better." If it were not for the existing code and expertise, it would perhaps be more reasonable to abandon the old and start from scratch in a new language. But there is a large body of useful, existing code that would take many years and much effort to reprogram.

We do not use the same natural language today that was used in the fifties, either. There are many new technical terms and idioms, and older ones have fallen into disuse. But we do not abandon the English language and take up another because it is more regular or more easily extendable. Instead the language evolves to meet modern needs. Can programming languages evolve in a similar fashion?

If not, there is little choice but to exist in a multi-language world. There are a number of efforts underway to permit "interlanguage communication" or "mixed language programming." The Department of Energy Language Working Group is planning a Workshop on this topic in May of 1985. Most successful efforts to date, that allow interlanguage communication between two or more programming languages, have begun with languages that are somewhat similar and proceeded to homogenize them even more. Even so, when this sort of intermingling has been accomplished, the trend seems to be to use it as a bridge to move entirely to one or the other of the languages. For various reasons, including the frustrations of having to deal with the idiosyncrasies of two compilers, and debugging in two languages, mixed language programs do not appear to be stable for very long.

What X3J3 is attempting to do, and it is an experiment, is to provide for this sort of migration in one cycle of the standardization process, while remaining within the confines of a single language.

* * *

It is the users of the language, the scientific programmers, who will be most affected by the evolution of the language. Therefore, it is the users who should have the greatest influence on any new standard for the language. As noted in a recent editorial in Computers and Standards (Vol. 3, No. 1, 1984), "...the importance of formal standards occurs when users participate in the committees, not just the vendors... The value of the formal process is to have the users determine which standards and when they will be changed. True, the users vote for standards with their purchases. And true, the formal process takes too much time. But the best solution calls for more user participation in formal standard committees, users who insist on quick standards and who readily move from one standard to the next in order to gain the greatest economic advantage over time."

At the present moment, X3J3 has 37 members that can be categorized as follows:

hardware vendors	19
software vendors	3
users	9
academics	5
other	1

One aspect of the economic advantage of widespread use of standard programming languages is the ability to acquire and make use of software developed by others. Even when an application area is unique, use can be made of development and maintenance tools that apply to the language, as well as libraries for graphics, mathematics, real-time, etc. The value of acquired software is enhanced if it is written in a language that is reliable, maintainable, and readily optimizable to the hosting hardware. A language standard that has universal application is not apt to be appreciated by a particular vendor to the same extent that it will be by a general user of the language. The vendors tend to be more conservative, because they have an economic interest in preserving the status quo.

* * *

Nevertheless, the committee, as a body, has recognized some of the limitations of Fortran 77 which the next standard will attempt to overcome. They are:

(1) its dependence on storage association

(2) its lack of data structures

(3) its limited means of language extension (only subprograms)

(4) the card orientation of its input form

(5) its limited use of structured control constructs (only IF-THEN-ELSE).

(6) its low-level mechanisms for handling arrays (DO-loops)

Overcoming these limitations might seem to require a revolution rather than an orderly evolution. It might seem that the integrity of untold millions of lines of existing code would be seriously impacted, but that is not the case.

No Fortran 77 features will be removed from Fortran 8X.

Instead, it is proposed that the next standard contain considerable duplicate functionality. A number of Fortran 77 features will be marked as deprecated. Deprecation implies that the feature could be removed from Fortran 9X. This allows for a transition period of approximately 11 years. As soon as new compilers are available, those who choose to, can begin writing in a more modern language without necessarily having to revise existing code. It is X3J3's intention that new Fortran 8X features be compatibly incorporated in standard-conforming Fortran 77 programs. Any exceptions to this general policy will be clearly indicated in the standard document. At the end of the 11-year transition period, any code that has not been updated through ordinary maintenance and revision would have to be converted (possibly with some aid from automatic conversion tools).

* * *

Some of the new features that are being added to the language are:

array processing

dynamic storage allocation and recursion

programmer defined data types

facilities for modular data and procedure definitions

environmental inquiry intrinsics

control of numerical precision

new source form

new control structures

Some of the well known language elements that are deprecated are:

old source form

Computed GO TO

Assigned GO TO

Arithmetic IF

PAUSE

COMMON

EQUIVALENCE

ENTRY

statement functions

DOUBLE PRECISION

BLOCK DATA

A great deal of effort has been put into defining the array features. These provide facilities that no other programming language has yet attempted. They raise the level of the language and allow better optimization on all machine architectures, but particularly those architectures that supply hardware for vectorization.

Although X3J3 has looked hard at event handling, it seems unlikely at this time that event handling facilities will be included in Fortran 8X. One of the reasons is that they may conflict with some, as yet undefined, multi-tasking facilities. Multi-tasking is seen as the major task for Fortran 9X, as array processing has been for Fortran 8X.

X3J3 is attempting to bring into the language, those newer facilities that will most benefit scientific programmers. With the exception of some of the array features, some flavor of all of these newer features have been tried in other languages. The creative challenge for X3J3 has been fitting the pieces together in a Fortran-like framework.

* * *

In a brief presentation, it is not possible to explain the new features or even the motivation for including them in the standard. The Fortran Information Bulletin attempts to do this more fully. The last three pages are a questionnaire. If you have any interest in the future of Fortran, I urge you to complete the questionnaire and return it to Andy Johnson. It has been said that such a questionnaire is valid only if no one pays any attention to the results, its object being to entice people into reading the source material. The Chair of X3J3 does have another objective in asking for the return of the questionnaire. That is, without such a vehicle, X3J3 may get comments only from those who oppose Fortran 8X. Supporters seem not to realize the value of communicating their support. X3J3 is publicizing its plans at this time to forestall delays during the two-year processing period. You may be aware that, for COBOL, that period has stretched to four years and is climbing. The adoption of a standard is a consensus process and negative reaction can cause significant delay or defeat the process.

There are some (including members of X3J3) who feel that Fortran 8X should be limited to the addition of the array features and perhaps a few control constructs and environmental inquiry intrinsics. If this is the general view of all users, then the standards committee should be made aware as soon as possible. When the Fortran Information Bulletin was sent to X3J3's parent committee, X3, for permission to publish, there were two negative votes (IBM and DEC). These two voted NO, not because they were opposed to publication, but because they do not favor the direction X3J3 is taking.

In the past Fortran language standards have only standardized existing extensions. With this revision, there will be an attempt to evolve the language into one that is modern, reliable, and portable. Unless users fully understand the rationale for this change of emphasis, there may be some unwarranted resistance, particularly if the only document that gets publicized is the draft standard. In an attempt to reach users, the Fortran Information Bulletin has appeared in an ACM publication that is a combined issue of Signum Newsletter and FORTEC Forum, and X3J3 has begun holding Fortran Forums to explain its intentions to the public. Forums have been held in Geneva, Switzerland, Ft. Collins, Colorado, and Idaho Falls, Idaho. They are currently planned for MIT, SLAC, an ACM meeting in New York City, EdCompCon in San Jose, California, Ft. Lauderdale, Florida, and Oxford, England. Others will, no doubt, be scheduled.

Users, who have the most to gain by the adoption of Fortran 8X, are inadequately represented on the standards committee. Membership is open to those who have experience with and an interest in Fortran. Members must attend 2 out of every 3 meetings. There are 4 meetings a year. Only one member per organization is permitted. If you have any interest in becoming a member, please contact me. If you do not wish to make that much of a contribution, you can still help by filling out and returning the questionnaire.

"Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore Laboratory under contract number W-7405-ENG-48."

DIGITAL SCENE SIMULATION(S) ON CRAY SUPERCOMPUTERS

Larry S. Yaeger

Digital Productions
Los Angeles, California

As many Cray representatives and some of the user community may be aware, Digital Productions for a little over two years now has been using the computational and I/O speed of Cray super-computers to produce computer graphics at an unprecedented level of sophistication and realism. These examples of Digital Scene Simulation(s) have been seen in a motion picture, "The Last Starfighter," and in various television commercials, including the Cleo award winning spot for the Sony Super Walkman.

In order to provide an understanding of the nature of our use of the Cray X-MP/22300, I would like to briefly describe the work environment at Digital Productions, and the manner in which a typical commercial project flows through the company:

First of all, clients approach us with a product that they wish to advertise, and, usually, some idea as to the basic look and content of the commercial they would like us to produce. This idea can take the form of a fairly detailed storyboard, such as with the Fuji videotape commercial, or of a most basic concept, as was the case with the spot for the Star Frontiers role-playing game. In the latter case, an in-house Art Director (AD) will be assigned the task of designing the look and sequence of actions for the commercial, and will produce a detailed set of storyboards to communicate these ideas to the client and to the rest of the company. Most commonly, a rough storyboard is provided by the client, which is turned into a final detailed storyboard by our AD's in association with the client's artists. As soon as the commercial project is initiated, a Producer is assigned to the project, who will handle client contact, contract negotiations, non-CGI (Computer Graphics Imagery) aspects of the production, and who will track the progress of the work while it is in-house.

From the storyboards, detailed engineering-style drawings will be made of the various objects in the scene by our Designer Encoders (DE's). These detailed drawings are then encoded, or digitized, on Talos/Calcomp tablets with an accuracy of 1/1000 inch, interfaced to a VAX 11/782, using software developed in-house (in FORTRAN). The encoding procedure produces a three-dimensional polygonal object description that will be read and understood by our rendering algorithms. Each object's data file consists of a few counts of informational elements in the file, a set of (x, y, z) coordinate triples describing the location of the points, and a connectivity list describing the sequence of "connect-the-dot" moves that defines the polygonal mesh. DE's are also responsible for uploading these objects from the VAX to the Cray, and setting up "scene" files (back on the VAX) for each object, which position the various parts of the object correctly with respect to each other.

At this point, the objects and corresponding scene files are turned over to Technical Directors (TD's). A TD is responsible for obtaining both the look of the scene (as defined by the AD) and the motion called for in the storyboards. A black and white vector display device, from IMI, is used to position the various objects in a scene, and to define the motion of objects over the length of the scene. The IMI's work under a UNIX shell, and the motion "PREVUE" system is in-house software written partially in VAX FORTRAN and partially in C on the IMI's. The TD uses DP3D - the principle rendering algorithm at DP, written in-house (in FORTRAN and CAL) - to assign the appropriate surface properties and lighting characteristics to the objects in the scene to achieve the desired look, saving these attributes in comprehensive, VAX-based scene files. These images are computed on the Cray, and transferred across the VAX/Cray station to Ramtek frame buffers, and displayed on Ramtek/Ikegami monitors. Our station was, of necessity, written in-house, as it was implemented prior to the existence of Cray UK's VAX/VMS station. It is written primarily in VAX FORTRAN, with a few key modules in MACRO (VAX assembly language). Considerable feedback between the TD and the AD, during this period of look and motion definition, results in the final nature of the images generated. The black and white vector motion footage, color and density wedges of the shaded images, and, sometimes, low resolution shaded film footage, are all seen and signed-off on by the client.

A great degree of flexibility and control over the special capabilities in DP3D - such as fringing, interpolation, transparency, textures, reflections and so on - is provided by a powerful, procedural language called FIFTH, and a macro-expansion language, similar to GPM and TRAC (that we refer to as TRAC), which process all input to DP3D. FIFTH is, again, an in-house development, featuring incremental compilation, ALGOL68-like syntax, stack orientation, infix and reverse polish notation systems, an indirect threaded code compiler, and very compact generated code. FIFTH is used to schedule events during a filming run - such as turning on a light at frame 48, or ramping up the opacity of an object over frames 78 to 100. FIFTH also handles the interface between the PREVUE defined motions and DP3D shaded images, via "movie" files containing object transformation matrices which are applied on a frame-by-frame basis. TRAC provides a general purpose macro-expansion capability, allowing system- or project-wide nouns and verbs to be defined, such as (RED) implying the (r, g, b) triple (1., 0., 0.). All filming jobs develop a corresponding FIFTH file (or set of files), as well as a set of SCENE files.

The final step in Digital Scene Simulation is then to make an entry in our filming queue which points to the appropriate SCENE and FIFTH files for that job. The queue manager, a set of VAX DCL procedures referred to as ROLLEM, will run DP3D, allowing the SCENE and FIFTH files to assign the correct lighting, motion, and special effects on a per-frame basis, and send the images to film. During final filming, two copies of DP3D are run simultaneously, each confined to slightly under 1/2 of the available memory to prevent thrashing with each other and with a frame SPOOLER. One copy of DP3D typically computes odd frames, and the other computes even frames. Interlock with the SPOOLER is provided by a semaphor file (referred to as the SEMAFILe) which contains slots for each cpu. DP3D can only save its current frame to disk, decrement the count for its cpu, and proceed to the next frame, if its count is greater than zero; the SPOOLER, co-residing in main memory, grabs each frame's disk file, spools it to the Digital Film Printer (DFP) Recorder, and then increments the count for the cpu upon which the frame was generated. The DFP Recorder is a special purpose piece of hardware, designed and built at III, which consists of a camera mounted over a high-precision, slow-scan or "flying-spot" CRT, along with various support, and diagnostic electronics. A high-speed interface between Cray's I/O Subsystem and the DFP Recorder was designed and built in-house at Digital Productions, using the Cray's 100 Mbyte/sec channel.

To give some indication of the magnitude of the problem being solved for high-detail, high-resolution Digital Scene Simulation, a few "back of the envelope" type calculations can be quite revealing. Our normal filming resolution (for 35mm Academy aperture) is 2560 x 2048, or 5.2 million pixels. With 24 bits of color (8 each for Red, Green, and Blue), each image comprises 15.6 Mbytes of data. Since a given frame may also frequently consist of a foreground, a mid-ground, and a background pass, and since it is necessary to spool disk files of the images in order to take advantage of both cpus, we can easily be moving as much as 62.4 Mbytes of data per frame. Even though there are many operations in the rendering process that are not "per-pixel" (such as the polygon-to-edge transformation, and elements of the hidden surface algorithm), or which may be low precision integer arithmetic (rather than floating point), it can still be estimated that somewhere between 1 and 10,000 floating point operations (flops) are required to calculate each color of each pixel. This implies a range of 15.6 Mflops to 156 Gflops of computation is required per frame (and remember that normal film is shown at a rate of 24 frames per second). Assuming an average of approximately 2500 flops per pixel-color, and a sustained computational rate of 150 Mflops/sec (this is a reasonable single head rate), a single frame of film then will take about 260 cpu seconds (about 4.3 cpu minutes) to render (single-headed). At this rate, but halving the time to account for double-headed calculation, approximately 50 minutes of wall time is required to produce a second of film.

And 30 minutes of animation, such as was required for The Last Starfighter, would require approximately 62 perfect, 24-hour days. In fact, the bulk of the final filming for The Last Starfighter was carried out over a three-month period, on a nearly dedicated system.

The numbers above are for our typical filming conditions. Just for a moment consider an extreme case of a very high-detail scene averaging 5000 flops per pixel-color, being filmed on a larger film area at high resolution, such as 65mm at 6000 x 4600, with greater color precision, say 36 bits (12 per primary)... These assumptions result in estimates of approximately 125 Mbytes of data per image element (foreground, midground, background, spool file), 500 Mbytes of data per frame, 414 Gflops to render each frame, and a single-headed frame time of 2760 seconds or 46 minutes. A second of film at this rate requires almost 1/2 a day double-headed, and 30 minutes of film would take approximately two years. Clearly, there is a great need in computer graphics imaging for the existing power of supercomputers, and a similarly great need for even more powerful computers in the years to come.

As can be seen in the various numeric figures above, I/O is as significant as cpu utilization in Computer Graphics Imaging. Cray's 100 Mbyte/sec High-Speed Channel is essential to our work. This high data transfer rate allows us to record a frame in approximately 7.5 seconds, to scan a frame in approximately 3.5 seconds, and we will soon be able to transfer a 1280 x 1024 x 20 bit image to a Ramtek Film Recorder frame buffer in about 1/8 sec. Compare these rates to a typical frame time of 3 or 4 minutes on other standard precision color film recorders, and to a period of 20 to 40 seconds (depending on system load) to transmit 4 Mbytes (1280 x 1024 x 24 bits) across the VAX station to a Ramtek frame buffer. Though use of DEC's UNIBUS allows us to purchase off-the-shelf interfaces, its 1.5 Mbyte/sec data rate, DEC's MASSBUS data rate of 2.0 Mbyte/sec, and IBM's block mux data rate of 2.5 Mbytes/sec are all insufficient to meet our heavy I/O usage requirements.

Another somewhat unique feature of the computing environment at Digital Productions is the typical job mix on the Cray. First of all, all of our work is done on an interactive basis, thus requiring relatively short time slices to provide quick response. Secondly, the bulk of our system resources go to large, cpu-intensive jobs (DP3D being exercised by the Technical Directors) with a limited number of small jobs (Software Developers doing some editing and compiling). Since these large (approximately 1 Mword) jobs must roll with some frequency, a high-speed roll device (fast I/O again showing up as a requirement) is an absolute necessity. Just 4 to 6 Megaword jobs rolling to disk with a time-to-roll which is greater than their in-memory time slice will cause COS (or any operating system) to thrash itself to death. We have only achieved a functional interactive environment by rolling job images to the IO Subsystem's buffer memory (which we have maximally configured to 8 Mwords). This roll to buffer memory is accomplished by a local modification to COS that, with some enhancements, should become a supported feature of COS. There are currently some holes in the roll to buffer memory code (which is actually roll to "named device" code) that still cause difficulties with Cray restarts (after a crash), frequently necessitating deadstarts, and thus the loss of all jobs on the system. This problem needs attention from CRI. The possibility for rolling to the new dual 1000 Mbyte/sec (total 2000 Mbyte/sec) SSD is an exciting prospect for Cray interactive computing.

In general, though the 1.12 scheduler is a marked improvement over the 1.11 JSH, we believe that there is still considerable room for improvement in the COS scheduler, especially with regards to its use in an interactive environment. In particular, we feel that implementation of the following items would significantly improve our throughput at Digital Productions: 1) Extrema for in-memory and on-disk times should be a function of job priority; since these thrash-locks end up controlling the system in an interactive environment with large jobs, they must be a function of job priority, in order to properly allocate system resources; 2) Thrash-locks should also be dynamic, should generally have the capability of being unique to a job, and should be a function of the speed of the roll-device; obviously, thrash-lock parameters that work well for a high-speed roll device such as buffer memory, are not appropriate for jobs that roll to disk.; 3) Roll-device should be specifiable on a job by job basis (perhaps priority, job class, and/or Station ID driven); buffer memory is a precious, as yet unschedulable resource, and batch jobs running in the background on an interactive system could be rolled to disk (with their own set of thrash-locks, of course); 4) For both job rolling and resource allocation, a hierarchy of devices should be specifiable; i.e., a request for a resource that cannot be answered should be able to be satisfied by a hierarchy of devices - e.g., a site might wish to keep a list that specified SSD datasets to overflow first to BMR, then to striped disk, and then to regular disk; 5) We would also like to be able to completely discard the roll image of certain classes of jobs (such as by Station ID) while they are in memory; batch jobs run for time sales have a great need to be recoverable across restarts, but our typical interactive jobs do not need to be so recoverable, thus we could free up on the order of 2 million words of buffer memory (associated with the jobs currently in memory) by discarding their roll images; 6) Small memory jobs should be grouped for rolling against large jobs; successively rolling-in multiple small jobs against a few large jobs, instead of bringing in all the small jobs simultaneously, greatly increases the time all jobs spend rolled out; 7) It should be possible to lock a job in memory; one specific application of this at DP is the need to lock the frame SPOOLER into memory for playback-limited jobs (i.e., ones which require more time to record on film than they take to compute); 8) Job aging (probably as a function of priority), such that jobs (coming in at the appropriate priority) would actually be scheduled at a higher priority for a defined period, with the actual scheduling priority dropping with age; this would better support a mixed set of short and long batch jobs running in conjunction with the interactive jobs, by letting the quick test cases move through the system rapidly, while preventing excessive resource utilization by the longer-running jobs; 9) In order to support a wide variety of site-specific requirements, JSH needs a user exit built in, to allow sites to implement their own variations on the scheduling algorithm; 10) The two functions currently assigned to JSH (in-memory versus roll-in/roll-out scheduling) should be broken out into separate code modules, with in-memory scheduling performed often and efficiently, and swapping scheduling being performed less frequently, but utilizing the longer data sample, and performing more intelligent heuristics (including those listed above).

As was mentioned previously, we work with our own VAX/Cray station. Between the current lack of operator privileges in Cray UK's VAX/Cray station, its high cost, and its lack of certain features used on a daily basis in our interactive work, particularly input and output redirection, and interactive graphics, we have found little incentive to change. If these capabilities were added to Cray's station, and the price reduced we might consider switching.

A few other items of special interest and heavy use at Digital Productions are:

Buffer Memory - As mentioned before, buffer memory serves as a high-speed roll device, dramatically improving our interactive response and throughput. It is also a necessary element of our Digital Film Printer process, due to the need for an extremely fast, uninterruptable output streaming rate to the high-speed DAC's driving the display in the recorder, and for the similar input streaming from the scanner. We need the capability to prioritize and schedule this precious resource, a need not currently being addressed by CRI. We also need to be able to force contiguous allocation for some datasets, which need is being addressed in COS 1.14. The ability to upgrade the memory chips in the IOS (as is now possible in the SSD), would be a tremendous boon to DP.

SEGLDR - DP3D itself is heavily segmented; the memory thus gained translates to rendering speed quite effectively, and even permits the set up and rendering of scenes which would be otherwise impossible. In addition, we use SEGLDR to provide exotic memory-intensive capabilities, such as a fractal geometry code (GEOVIEW), a fluid-dynamics driven particle rendering system (VORTEX), and alternative rendering algorithms (e.g., a ray-tracing scheme) in segments which may be invoked at the main command level in DP3D. General improvements in the stability of SEGLDR, in error-trapping and reporting, and in capabilities (such as making the DUP directive actually work) would have a significant impact on production for us. We are especially impacted, in a negative fashion, by the current incompatibility between SID and SEGLDR.

Memory management - We have achieved considerable gains in system throughput by limiting DP3D jobs to the minimum amount of memory required at each stage of use. Accordingly, we are sensitive to changes in and functionality of the CFT-callable MEMORY routine. Also, almost all significant data structures in DP3D are POINTERed.

SHIFT and MASK operations - Also related to memory conservation issues, it should be noted that almost all significant data structures in DP3D are packed up, with multiple items per word. Accordingly, automatic vectorization of these functions by CFT will provide a valuable improvement in rendering speed; we are pleased to hear that this is finally being addressed in CFT 1.13 (if we can only get it to successfully compile our code).

Statement functions - These are heavily used to provide readable code, hiding the bit-picking and multiple-level address indirection with constructs that compile in-line, and do not, of themselves inhibit vectorization.

CAL code - We have recently found ourselves in the traditional situation of a single module consuming most of the cpu cycles for a typical frame. Though this is a large module (with lots of in-line, non-modularized code written purely for computational efficiency), we are working rapidly to move the various elements of this module into CAL. We are able to realize factors of 3 or 4 in typical cases, where the FORTRAN already vectorized to some degree, and factors of 7 or 8 where a technique supports vectorization in CAL that was not achievable in CFT. These enormous gains are a necessity for us, so CAL-coding is a major activity at DP. We would benefit significantly from better support of CAL programming and timing analysis tools by Cray. We have observed significant timing discrepancies between SPY and CFT's FLOWTRACE. The source of these discrepancies is as yet unresolved, as is the level of accuracy that may be ascribed to each. We would very much like to see CYCLES updated to the X-MP, and both SPY and CYCLES supported by CRI.

Interprocess communication - Our applications frequently need to relay their current state to other processes. So far this has been implemented almost independent of COS, by the contents of specially designated datasets, and by the event suspension associated with unique accesses of permanent datasets. Again, we are pleased to hear that this will be addressed in COS 1.14.

In conclusion, I would like to recognize the significance of Cray's supercomputer hardware to our existence as an edge-of-the-art undertaking, and as a viable economic enterprise. Certainly the level of results obtained to date, in terms of both quantity and quality, would not have been possible on a lesser machine. Cray is also to be congratulated on its recognition of the need for enhanced I/O capabilities to match the increased CPU performance of new machines. We look forward to continuing mutual growth.

WORKSHOP REPORTS

I/O WORKSHOP

Mostyn Lewis

Chevron Oil Field Research Company
La Habra, California

As there had been no responses to requests to talk at the I/O workshop, only the chairman spoke. This might be taken as a lack of interest in the area of Cray I/O but from the very large attendance at the workshop, this does not seem to be the case and it is hoped that there will be a more enthusiastic response at the next CUG.

It was the contention of the chairman that I/O is very important on the Cray and ought actively be discussed. Megaflops might be more glamorous than Megabytes and surely Cray have a super 'number cruncher', but often in a real performance aspect it is the humble Megabyte that has Megaflop by the tail. Bandwidth is crucial to performance in data processing, not just the ability to run CPU kernels at high rates.

I/O, that very famous bottleneck, looms larger as technology proceeds. Larger main memories, more numerous processors, and bigger SSDs mean more I/O. Consider swapping 8 megaword job images: at 10 megabytes per second (say on a DD-49) this would take 6.4 seconds and at 3 megabytes per second (DD-29 realistically) this is greater than 21 seconds! Four CPUs obviously means four times the capacity for I/O and as COS is single-threaded, the slow down experienced may be a true bottleneck.

It is evident that I/O must be taken seriously and surely all sites have some illuminating experiences to share.

As an example, it was shown that the most desirable type of random I/O, standard FORTRAN 77 direct access I/O, was the least efficient (slowest) of all.

LANGUAGES WORKSHOP

Mary E. Zosel

Lawrence Livermore National Laboratory
Livermore, California

The languages workshop featured talks on the array extensions to standard Fortran and the CRI Pascal improvements. These are summarized below.

In addition, an informal survey was taken to determine how the various sites are progressing in changing to the new calling sequence conventions. As of the Spring 84 CUG meeting, very little conversion had taken place. The current results were encouraging: twelve sites had completed conversion, seven sites were in progress, and only one site had not yet started. Concern was expressed about conversion of third party software.

CURRENT STATUS OF FORTRAN 8X

Loren Meissner
University of San Francisco

Jeanne Martin
Lawrence Livermore National Laboratory

Among the new facilities being proposed for Fortran 8X, the array processing features are of special interest to scientific users. Arrays and array sections are viewed as atomic computational objects on which operations can be performed directly, rather than on the array elements individually. All scalar operations are extended to conformable arrays. Both user-written and intrinsic functions may return array values. Such operations may be masked by conforming logical arrays. Virtual arrays may be defined by linear functions on top of actual arrays. The ability to manage and control storage of arrays has also been significantly enhanced.

These proposed features and others such as programmer defined data types, numerical precision specification, and facilities for modular data and procedure definitions are described in a recent article, "Status of Work Toward Revision of Programming Language Fortran" by Jerrold Wagener, appearing in a combined issue of two Association for Computing Machinery publications, SIGNUM Newsletter, Volume 19, Number 3, July 1984 and FORTEC Forum, Volume 3, Number 2, June 1984.

PASCAL 2.00 PERFORMANCE

Karen Spackman
CRAY Research, Inc.

The first release of CRAY's supported Pascal, Pascal 1.00, was primarily a functional release. CRAY was interested in making Pascal available for internal use and for CRAY's users. While we were concerned with performance, particularly for scalar code, this was not a primary emphasis for the first release; performance is an area where we expected improvements in subsequent releases.

The Pascal 2.00 release includes some language extensions and improvements to scalar performance. The language extensions include: IMPORTED, EXPORTED, and COMMON data; STATIC data; VALUE statement; VIEWING statement; and SIZE OF function. The performance improvements include: a separate optimization pass which does common subexpression elimination; increased use of B and T registers; partial evaluation of Boolean expressions; dead code elimination; additional strength reduction; avoiding references to the literal pool; and peephole optimizations on generated code.

The basic test used during the development process to measure the effect of the performance improvements was the Pascal compiler. This is a good measure to use because we are interested in improving the performance of the compiler, because the Pascal compiler is a program of substantial size, and because writing a compiler is a typical use for Pascal. Using the Pascal compiler itself as a performance measure has the disadvantage that some of the performance improvements are attributable to algorithmic improvements in the compiler. While it is easy to identify which performance improvements are attributable to generated code changes between subsequent versions of the compiler, it is extremely difficult to determine how much of the performance change between releases is the result of generated code improvements.

Four Pascal programs were run with specially built compilers that turned on some of the performance improvements individually. These programs were also run with all of the performance improvements turned on (as in the released version of Pascal 2.00), and three of these programs were also run with Pascal 1.00. (The fourth program cannot be run with Pascal 1.00 because it uses Pascal 2.00 features.) These timings indicate that the biggest performance improvements resulted from the increased use of B and T registers. This confirmed the results observed with the compiler during the development of Pascal 2.00. Fortran versions of two of these programs were run with CFT 1.13 with vectorization disabled; the Pascal versions were slightly faster indicating the Pascal generated scalar code is roughly comparable to CFT's.

Additional work in improving performance of Pascal generated code is planned for future releases. We intend to develop a set of performance tests for Pascal and would welcome contributions to this. B and T register usage will be further improved by reusing those assigned to temporaries. Invariants will be moved out of loops. Instruction scheduling and vectorization are planned. For non-reentrant code, parameter lists may be built at compile time. Expanding user procedures and functions in-line is also planned.

PASCAL OPTIMIZATION EXPERIMENTS

	Sieve	Primes	Pascsub	Maude
Pascal 1.00	.0301	1.94	5.79	--
2.00 with no Optim.	.0312	1.04	6.30	2.55
2.00 with Peephole	.0306	1.03	5.98	2.52
2.00 with O = O +	.0278	1.05	6.46	2.49
2.00 with B&T Regs	.0144	.92	4.73	2.31
2.00 with All Optim.	.0144	.91	4.87	2.32
CFT 1.13	.0149	.94	--	--

PASCAL TESTING RESULTS

	1.00		2.00	
	%Passed	No.	%Passed	No.
Conformance	80	186	97	212
Deviance	92	159	95	249
Error Handling	43	93	76	67
Level 1	91	56	93	29
Quality	73	57	91	60
Local	86	225	93	425
Overall	80	776	94	1042

FRONT ENDS WORKSHOP

Dean W. Smith

ARCO Oil and Gas Company
Dallas, Texas

This session consisted entirely of presentations by CRAY Research personnel on the current status and future development of stations and station like products. Presentations were provided by John Renwick-CRI Mendota Heights, Brian Walsh - CRI Mendota Heights, John Flemming - CRAY UK, Stuart Ross - CRAY UK. In addition to the status reports on stations in the CDC, VAX and IBM environments, there were presentations on the ISP and SUPERLINK products.

NOS Release 1.13, John Renwick

NOS release 1.13 was released 8/84 and included as major features: BB and BD dataset formats; NOS 2.2 compatibility; improved transfer rates; and an automatic RELOG - station will automatically attempt logon after the link has been broken.

NOS/BE Station Release 1.13, John Renwick

NOS/BE release 1.13 is scheduled for availability October of 1984. It includes the following major features: the same features as the NOS release 1.13; support for DISPOSE,...TID=C; and CSTAT and CJOB output to an alternative file.

NOS/BE Station Release 1.14, John Renwick

NOS/BE release 1.14 is scheduled for release 3rd quarter of 1985. The following features were announced: interactive support through station; multiple stations running on the same NOS or NOS/BE system.

Transfer Rate Improvements

Direction	Transfer 1.12 BF1	Rates 1.13	Transfer Improvement
To CRAY	0.19	0.25	1.3 x
To CYBER	0.15	0.41	2.7 x

Performance characteristics on CRAY XMP and CDC CYBER 73 with 1 X PPS.

VM Station Release 3.0, John Renwick

Scheduled availability is October of 1984. Some of its announced features include interactive graphics support for graphics terminals; user exit dispose - dispose to special applications by passing the passing data segments to other virtual machines via VMCF; submission of files in NETDATA format; and a number of improvements to CRSTAT facility.

SUPERLINK/ISP Status, Brian Walsh

An integrated product line for the MVS environment was announced as part of the ISP talk. The MVS ISP product has been renamed to SUPERLINK/ISP; its successor product is SUPERLINK/MVS and is discussed below. The ISP is currently scheduled for first customer ship in December of 1984.

A host of ISP configurations were discussed - with perhaps the most interesting being a networking configuration of multiple CRAY mainframes and a variable number of MVS frontends. An immediate consequence of such a network could be the interaction or sharing of data by different CRAY mainframes.

Although the ISP product has only been announced in the MVS front end environment, some interest was expressed by attendees of the session for ISP facilities in other environments.

SUPERLINK/ISP Requirements, Brian Walsh

The following requirements were announced for the installation of the ISP:

1. MVS release 1.3.3, JES 1.3.3.
2. No RACF release requirement.
3. COS 1.14.
4. IOS.
5. A station.
6. 2 or more FEI's, check on exact model number.
7. 2 or more IOP channels.
8. The ISP products on CRAY and IBM mainframes.

MVS Station Release 1.12, John Flemming

Release 1.12 is the current release of the MVS station. The base release was available in March 1984; bugfix 2 was available in June 1984 and is the first release with official MVS/XA support.

There was some concern expressed at the session regarding the stability of the CSS 1.12 product, especially with regard to the parameters being passed to user exits. John acknowledged that there had been some problems, especially in the area of user exits. Hopefully, the stability problems with the MVS station product can be better addressed with CRAY UK's new computer center coming online this fall.

MVS Station Release 1.13, John Flemming

Release 1.13 is scheduled for availability in December 1984. Some of its major features will be: support for MVS/XA; support for JES2 Release 1.3.4; new user SMF records for LOGON/LOGOFF, data transfers, and start-up; support for larger segment sizes across the link; IBM 3800 printer JCL support; MVS allocation messages in user's COS job log; support for new COS 1.13 commands JSTAT, RSTAT and FLUSH; and HSM support to restore migrated datasets prior to allocation.

MVS Station Release 1.14, John Flemming

Release 1.14 is scheduled for availability in 2nd quarter 1985. Some of its major features will be: interactive support as a VTAM application; loosely coupled CPU support; RACF support for IOS tapes; integrated RACF Release 1.6 implementation; data set I/O using QSAM; and the JES2 source code modifications installed as user defined exits.

APOLLO Station, John Flemming

A prototype is installed at Mendota and a station product is scheduled for general availability in 1st quarter 1985. Some of its major features will be: interactive support; batch job processing; staging subset of dataset types.

VAX Station, John Flemming

1. Release 2.04, support for VMS 4.0, available November 1984.
2. Release 2.05, bugfix, available February 1985.
3. Release 3.01, full functioning station support, May 1985

Some concern was expressed to me regarding the lack of information and time spent on issues involving the VAX environment. With approximately 20 VAX station installations - nearly the number of MVS installations - there should be increasing interest and time spent on the VAX station, its development and its use.

SUPERLINK/MVS, Stuart Ross

The SUPERLINK product signals a major shift in CRI's station philosophy. It promises to draw the CRAY CPU into the IBM environment as an equal partner in an integrated computing facility, utilizing a gateway approach which will attempt to present the CRAY as a full partner.

SUPERLINK/MVS will be architecturally based on the SUPERLINK/ISP product; but will continue to provide station-oriented facilities. SUPERLINK will be the successor product to both the SUPERLINK/ISP product utilizing the enhanced transfer rates and the external interfaces of the ISP; while maintaining the station operator facilities, NJE support, interactive capabilities, online tape support, and data set transfer facilities.

From the user's point of view it is CRI's intention to provide easy access to the CRAY, route user's jobs to the CRAY using standard IBM NJE facilities, provide simple access to IBM datasets, establish a CRAY exposure throughout a VTAM network, provide an external interface to the CRAY, and support interactive JOB facilities.

Post Script

I would like to take this opportunity to invite all interested parties interested in making presentations at future front end sessions at CUG to come forward. Topics of interest can include, and are not limited to, user experiences with stations, problems encountered with stations, unique modifications made to stations, or beta test experiences.

Disclaimer: This report is taken from my notes of the front end session. Any errors or inconsistencies are probably mine and you should contact the speakers for any clarifications.

LIBRARIES

Margaret L. Simmons

Los Alamos National Laboratory
Los Alamos, New Mexico

The first talk, entitled, "Bridge to the Future: the LTSS Shell," was presented by Barbara Atkinson of the Lawrence Livermore National Laboratory. In it she described how Livermore is using its current LTSS user interface as a shell to bring users into the new operating system environment of NLTSS.

The second paper, entitled "Architectural Interfaces to Mathematical Libraries," was presented by Dr. K. W. Neves of Boeing Computer Services (BCS).

A summary of this talk follows:

Modern vector computers such as the Cray-1 and Cray X-MP offer the engineer and scientist the opportunity to improve their productivity greatly. By offering more memory and greater computational speed, these computers allow for the solution of otherwise infeasible problems. The performance capabilities of these types of machines have not come without some challenges for their users. Speed is obtained through innovative hardware designs that exploit parallelism and concurrency in the very heart of computation. This has created a situation where routine conversion of scientific software from one "brand" of FORTRAN to another does not achieve the "payoff" in performance one expects from looking at quoted MFLOP maximum rates. In order for the engineer or scientist to be prepared to take advantage of this kind of computational power, it is necessary to understand how these machines work. The architecture of today's supercomputers (Cray, Cyber 205, Fujitsu VP-200, Hitachi S-810) can be broken down into the following components:

1. Memory (primary and secondary)
2. Memory access (port, buffer, and bandwidth characteristics to and from vector units)
3. Scalar processing
4. Instruction processing
5. I/O characteristics
6. Vector definition (contiguous, regular, and random storage/access)

If we focus on the Cray series of computers alone, we see remarkable differences in architecture and, hence, performance in just looking at the above characteristics. The Cray-1 series is a "one-path-to-memory" machine. This can degrade optimal vector unit performance by a factor of three in memory-bound computational kernels. The Cray X-MP, in addition to providing a second CPU, provides three vector paths to memory which can greatly improve performance. The new X-MP line provides indirect addressing in vector mode at impressive speeds and low startup times. This redefines the architectural definition of "a vector" through its access/storage mechanism.

These types of changes in architecture can make a substantial change in application software performance. The issues of vectorization are all important. Examples show that a high percentage of vectorization is necessary for high performance, but not sufficient. The "quality of vectorization," (i.e., vector length, algorithm selection, and memory management) are really the key to unleashing the power of modern computers. Time and time again we see benchmark programs improving by factors of 8 or 10 in performance after the first few runs on a vector machine. Quite often changes in less than 5 or 10 percent of the FORTRAN code itself can lead to an order of magnitude improvement.

In an attempt to rationally cope with the vectorization issues that face us in today's and future designs, BCS has developed both software and a methodology for handling the impact of architectural changes. Through a structured software approach we provide a means by which the user community can "corral" architectural effects by use of a set of lower-level software. Boeing application programs themselves liberally use (via the FORTRAN "CALL") mathematical libraries. The standard library at Boeing is BCSLIB, a BCS-developed standard portable mathematical subroutine library. BCSLIB, in turn, rests on BCS/VECTORPACK, a set of lower-level computational kernels tuned to vector architecture. This kernel library is "called" both by BCSLIB routines and by user programs directly. The Cray version of BCS/VECTORPACK is called CRAYPACK and is coded in CAL to provide the "best" implementation of the "best" algorithms for Cray-1 and X-MP architectures. The subroutines are "machine intelligent," which means at run time they can recognize whether they are on Cray-1 architecture or Cray X-MP architecture. All the subroutines in both BCSLIB and CRAYPACK are provided in FORTRAN as well, for use on IBM, CDC, VAX, and a host of other mainframes.

If we look at the evolution of mathematical software, VECTORPACK can be considered the next evolutionary step to providing standard FORTRAN interfaces to special functions like SIN, COS, TAN, and X**Y. It also extends the concept of defining the hardware environment. For years, providers of mathematical software have localized machine-dependent constants such as word length (radix and mantissa), machine epsilon (the smallest number you can add to one and get a number bigger than one), largest floating point number (single and double), and a host of other constants. The localization was done by providing this type of information in subroutine form. When software migrated to new computers, just the localized subroutines needed to be changed. Hence, a measure of portability was created. VECTORPACK provides this capability for "architectural" differences.

The level of computational kernels in CRAYPACK is dictated by performance. For example, very low-level kernels such as sparse and dense vector operations are provided. On the other hand, important gains in performance (a factor of 10) can be achieved by including entire algorithms (such as for 2-D FFTS). It is perhaps of interest that the software discussed (BCSLIB and CRAYPACK) is commercially available for Cray owners from BCS.

Examples are given to show how very subtle differences in computer architecture can greatly impact algorithm performance. The effectiveness of a vector computer on a given scientific application is shown to be a complex mixture of maximum computation rates, vector startup time, scalar degradation, and ability of the algorithm to be vectorized. The structured approach has proven to be effective in maximizing efficiency without sacrifice of accuracy, portability, and ease of use.

Both talks were timely, interesting, and well received. Because of this interest, the planned discussion on multitasking in libraries had very little time. We hope to be able to explore this problem in detail at the Spring 1985 CUG.

A request for wish list input brought only one request: that SAVEd variables under multitasking conform to ANSI78. There is some confusion as to whether they currently do.

OPERATIONS WORKSHOP REPORT

Gary Jensen

National Center for Atmospheric Research
Boulder, Colorado

I began working on the first Operations Workshop in the Summer of '81. That first meeting was held in Dallas at the Spring '82 Cray Users Group meeting. Close to 40 people attended that first meeting, but only 9-10 were operations people. Attendance by operations people grows with each meeting and is now in the high 30's.

The Operations Workshop (now known as The Special Interest Group on Operations) has as its purpose to provide a forum for Operations Managers and staff to present and discuss issues relevant to operating a Cray site. At each of the last three CUG meetings, we have held three sessions of the workshop. With four and one half hours of meetings at each CUG, this has to be the most intense topic discussed.

A regular feature of the workshop is the Cray Report where Cray Representatives are available to discuss hardware-related issues in an open forum. With Stuart Drayton in the saddle for International Field Service now, I think that this will even get better. I have the feeling that Stuart likes a good "scrap" once in a while, and I think that this may be right up his alley. At the next two CUG meetings, it is my plan to devote one session to the issue "To PM, or Not PM". This issue has been discussed at many of our meetings and seems to be getting hotter with time.

Another regular feature are the Site Reports from three of four different Cray sites. Different sites are invited to tell the group about their particular operation. They discuss hardware configurations, staffing, statistics, facilities and special factors that increase the degree of difficulty at their site. At this meeting we heard from Lockheed, Grumman, and Boeing.

All in all, interest in these sessions is increasing.

As important as it is to hear the information presented, these meetings give Operations people the opportunity to meet others with the same interests and problems. Knowing who the people with like interests are, at other Cray sites, makes it easy to call and discuss individual problems. There has been a lot more conversation between sites since these meetings began. Prior to the Workshop, there was almost none.

I know that in my case, I have made many friends around the world, with like responsibilities. I treasure the friendships that I have made at these meetings. I am in contact with many other Operations people throughout the year. I am convinced that these meetings are to our advantage, as well as Cray's.

This workshop also gives Cray Hardware and Maintenance people a forum to discuss their wants, needs and plans for the future. It can also give them a place to discuss things that they may want group opinions on. I hope that these people will look to this workshop as the place to discuss these things.

At this meeting, we tried a new angle at the first session. We devoted an entire session to a panel discussion on the topic of "Networks from the Operations Standpoint". Jerry Esch of Sandia chaired the session, with Jim Hughes from Network Systems Corporation, Fran Pellegrino from Westinghouse, Fred Montoya from Los Alamos and Robert Niffenegger from NCAR on the panel. Each presented their ideas on the problems of operating a network and the need for network control systems. I thought it was great.

At the next two meetings, we intend to have a showdown on the PM issue. Some sites are reducing PM to a very few hours or none at all. These sites feel very strongly about this. Cray representatives feel as strongly against this practice. I think that this workshop is the place to discuss and possibly settle this issue. It may not be to everyone's satisfaction, but the issues will be known as well as the feelings of the users.

The Operations Workshop is only one special interest group of CUG. Our ambitions may be to increase the number of sessions we hold in the future, but that is the limit of our desires. We enjoy meeting at CUG and the relationship this gives us with all you software types.

And let's face it! We need to work on that relationship. It is very important for Operations and Systems staff members from any single site to come to a meeting together and get to know each other. It is also important that we get to know Cray. The bottom line is that these meetings create a feeling of camaraderie among us all.

In the future I will be looking for someone overseas to help me with the European meetings. I am also looking for a volunteer from a US site to help on this side of the ocean. Since we usually meet once on each side of the Atlantic each year, the load on these colleagues would only be for one meeting each year. This would make it easier for the new Board of Directors to give me the boot if they wanted to.

I would like to thank the following people that participated in this series of sessions:

1. Those already mentioned;

2. Fred Montoya of Los Alamos and Lou Saye of Cray for the Reliability Information and presentation;

3. Dee D. Foote of Lockheed Missile and Space for his insights into Lockheed;

4. James Poplawski of Grumman Data Systems for his presentation on activities at Grumman;

5. Robert Cave of IDA for his experiences installing INTERMEM Memory on their CRAY1A;

6. Bernard Pria of SNEA for his explanation of CALLSOFT;

7. Jim Roetter of Boeing Computer Services for his presentation on BCS.

The next meeting will bring us into focus with the activities of several more sites as well as to introduce us to a lot of new information. I don't see how anyone can afford to miss these meetings.

Please contact me at NCAR, Boulder, Colorado 80307-3000, USA if you have anything that you are willing to share with the rest of us. I look forward to the next meeting. See you then.

CUG HARDWARE RELIABILITY SURVEY

TOTAL EVENTS BY REASON

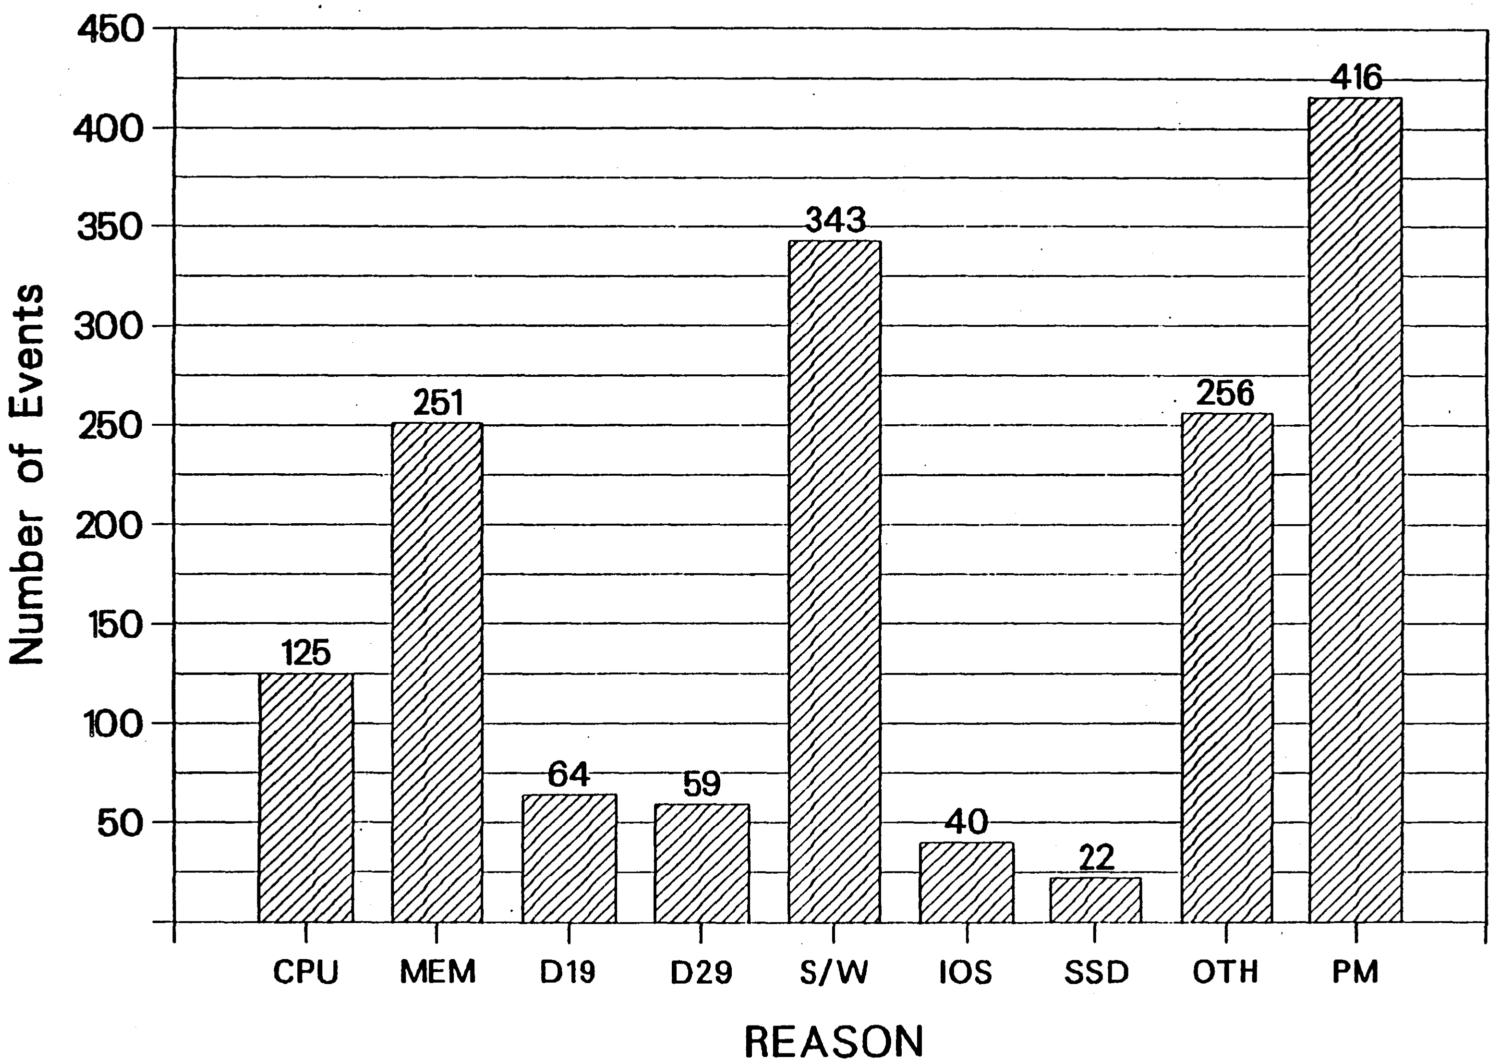

CUG HARDWARE RELIABILITY SURVEY

AVERAGE EVENTS BY REASON

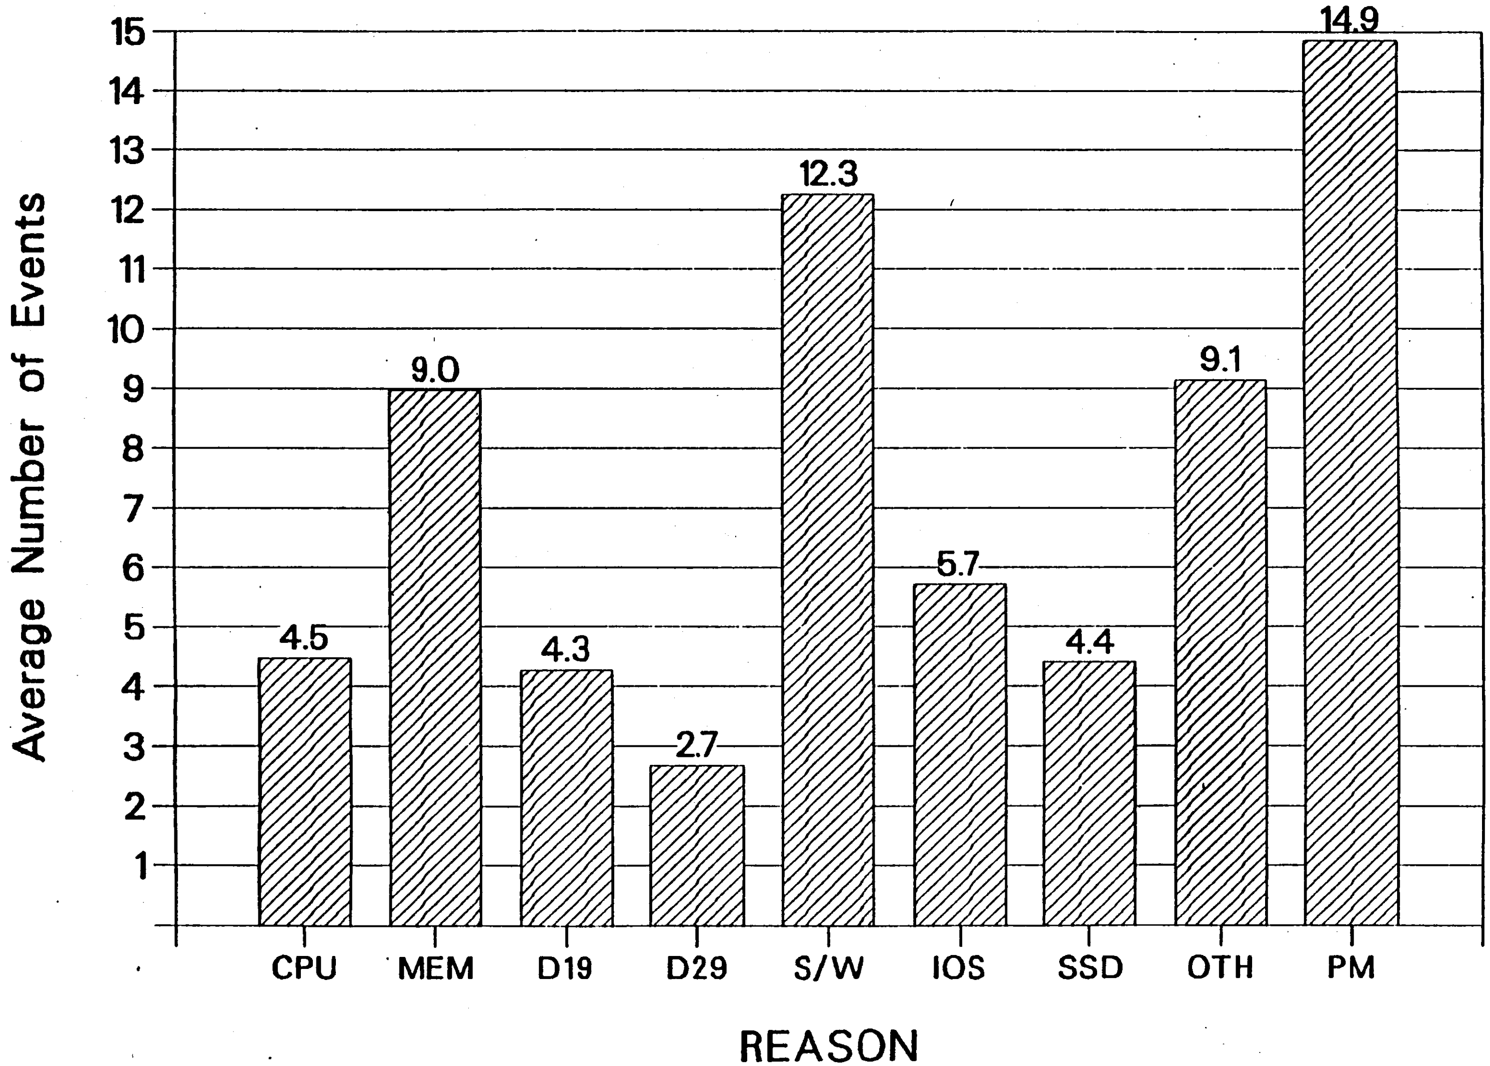

CUG HARDWARE RELIABILITY SURVEY

TOTAL DOWNTIME BY REASON

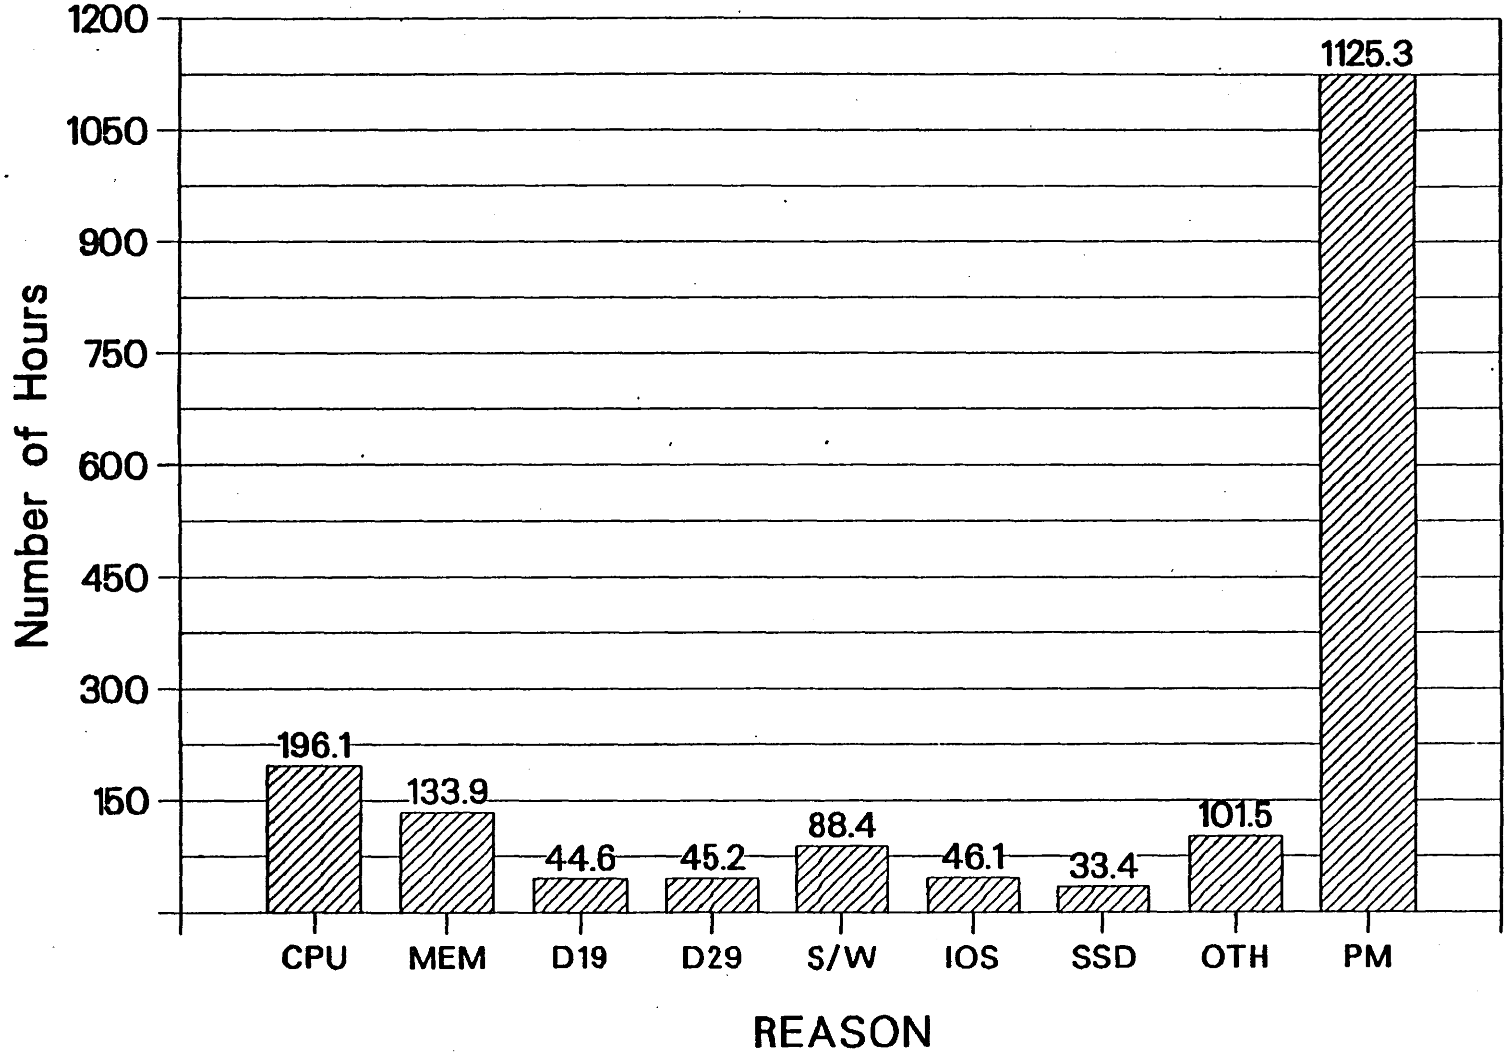

CUG HARDWARE RELIABILITY SURVEY

AVERAGE DOWNTIME BY REASON

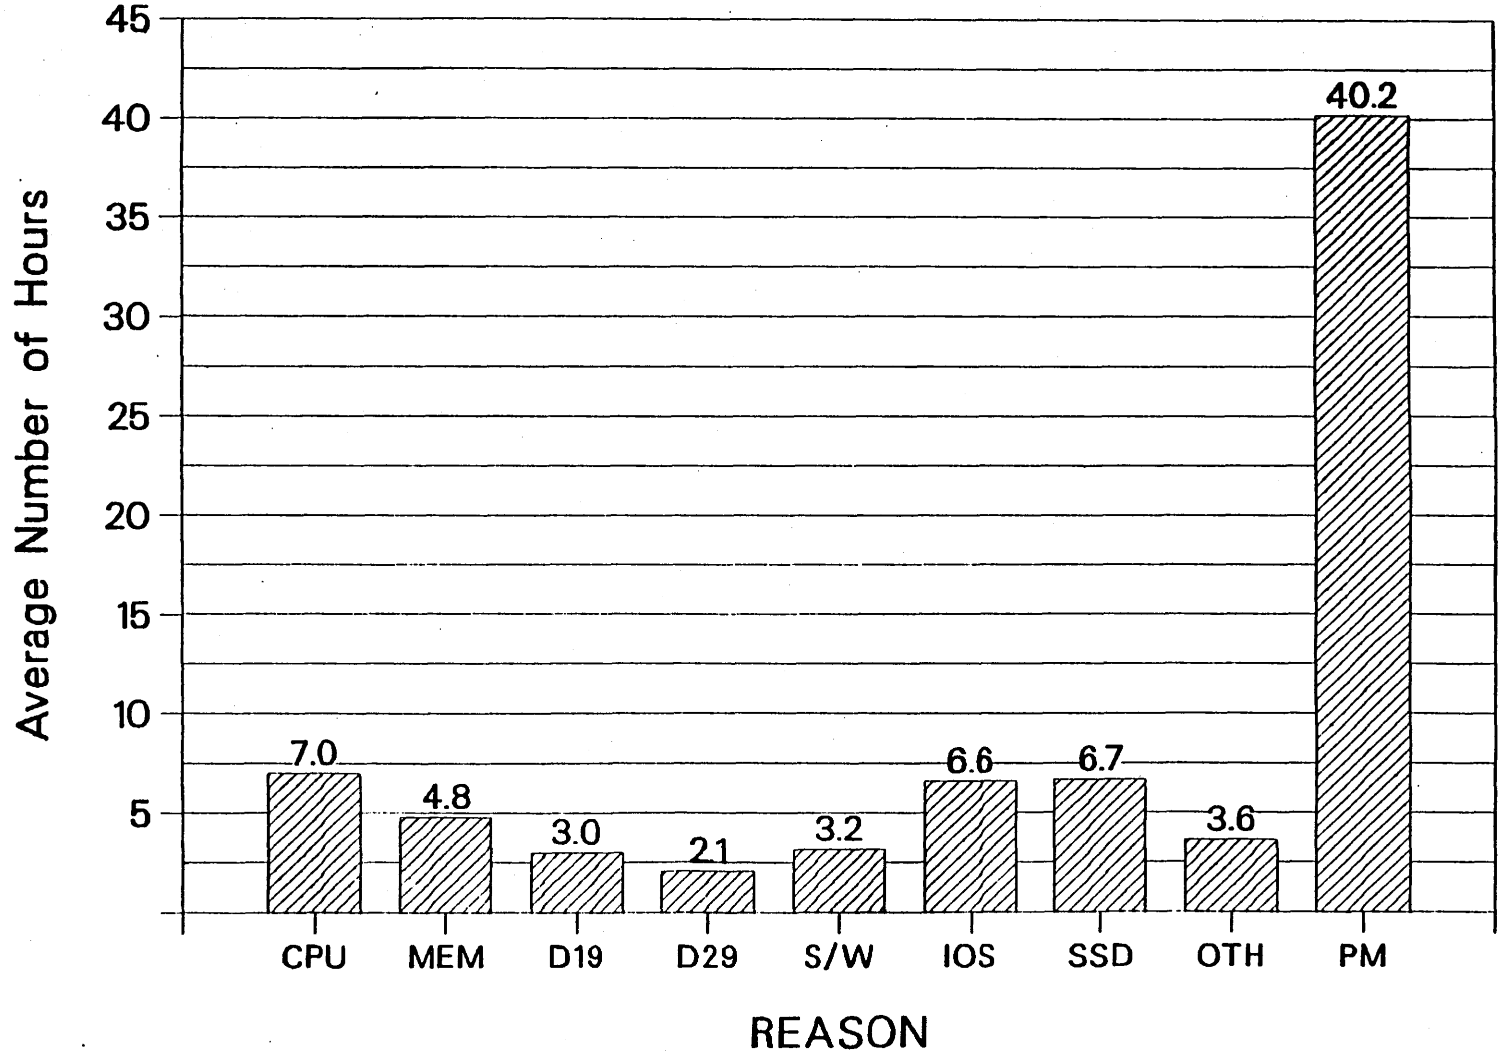

CRAY HARDWARE RELIABILITY SURVEY

Events

S/N	TYPE	CPU	MEM	D19	D29	S/W	IOS	SSD	OTH	PM
003	1S	6	1	22	0	49	0	0	7	26
004	1A	4	15	1	0	8	0	0	0	12
005	1A	0	2	0	0	0	0	0	1	42
006	1A	3	8	3	0	17	0	0	0	0
010	1A	17	15	23	1	18	0	0	94	0
011	1S	2	5	0	0	0	1	0	0	26
014	1S	10	4	9	0	11	0	0	10	24
016	1A	1	3	0	0	7	0	0	0	12
017	1A	1	10	0	0	0	0	0	0	13
018	1S	3	13	0	3	3	0	0	0	12
019	1B	0	2	1	0	0	0	0	0	0
022	1S	0	47	0	0	12	0	0	6	7
023	1S	8	4	0	2	4	0	0	0	13
027	1S	3	8	0	0	0	14	0	1	41
030	1S	3	36	0	1	3	0	0	3	26
032	1S	2	18	2	0	30	0	22	49	28
033	1S	0	2	3	0	9	0	0	0	1
035	1S	16	16	0	3	5	0	0	0	12
038	1S	10	3	0	20	41	0	0	3	6
039	1S	4	2	0	0	2	0	0	63	28
043	1S	6	8	0	1	50	0	0	7	6
049	1S	4	0	0	2	0	0	0	0	0
050	1S	5	6	0	1	0	0	0	0	1
054	1S	3	3	0	3	0	2	0	0	26
107	XMP	1	1	0	4	6	2	0	10	12
110	XMP	7	7	0	13	63	16	0	2	25
113	XMP	5	3	0	3	0	5	0	0	13
114	XMP	1	9	0	2	5	0	0	0	4
TOTALS		125	251	64	59	343	40	22	256	416
AVERAGES		4.5	9.0	4.3	2.7	12.3	5.7	4.4	9.2	14.9

CRAY HARDWARE RELIABILITY SURVEY

Hours

S/N	TYPE	CPU	MEM	D19	D29	S/W	IOS	SSD	OTH	PM
003	1S	8.0	.1	19.9	0	20.2	0	0	4.5	51.0
004	1A	4.5	8.2	1.0	0	1.1	0	0	0	41.9
005	1A	0	.3	0	0	0	0	0	2.4	48.9
006	1A	7.2	2.7	1.1	0	0	2.0	0	29.2	0
010	1A	22.4	6.8	6.8	1.7	3.8	0	0	.7	0
011	1S	5.1	1.5	0	0	0	.2	0	0	55.0
014	1S	14.25	1.2	11.3	0	2.5	0	0	6.2	44.9
016	1A	1.3	.9	0	0	.8	0	0	0	33.8
017	1A	.5	2.0	0	0	0	0	0	2.0	0
018	1S	4.5	10.8	0	2.2	.8	0	0	0	51.4
019	1B	0	2.6	.1	0	0	0	0	0	0
022	1S	0	13.7	0	0	2.3	0	0	4.1	14.7
023	1S	12.5	4.3	0	3.0	3.5	0	0	0	52.0
027	1S	10.5	4.3	0	0	0	28.5	0	2.0	110.0
030	1S	8.8	2.2	0	1.2	.3	0	0	15.0	52.0
032	1S	2.3	24.5	1.7	0	9.9	0	33.4	16.3	82.5
033	1S	0	12.0	2.3	0	1.1	0	0	0	1.9
035	1S	20.8	3.7	0	.9	.7	0	0	0	46.8
038	1S	3.9	2.5	0	14.1	6.2	0	0	.3	19.9
039	1S	8.7	.4	0	0	.3	0	0	12.3	73.3
043	1S	10.6	3.1	0	.9	7.8	0	0	1.8	22.2
049	1S	9.5	0	0	1.1	0	0	0	0	0
050	1S	11.4	4.0	0	2.4	0	0	0	0	3.0
054	1S	5.9	1.9	0	.5	0	1.8	0	0	91.0
107	XMP	4.8	2.3	0	4.3	.8	2.1	0	4.6	64.4
110	XMP	13.6	4.3	0	6.4	23.1	10.7	0	.3	96.9
113	XMP	5.0	3.1	0	6.5	0	2.9	0	0	39.0
114	XMP	.1	10.6	0	.5	1.4	0	0	0	31.7
TOTALS		196.1	133.9	44.6	45.2	88.4	46.1	33.4	101.5	1125.3
AVERAGES		7.0	4.8	3.0	2.1	3.2	6.6	6.7	3.6	40.2

EXPERIENCE WITH THE INTERMEM 8330 STORAGE UNIT FOR THE CRAY-1

BOB CAVE
INSTITUTE FOR DEFENSE ANALYSES
OCTOBER 4, 1984

INTRODUCTION

WHY CHOOSE ANOTHER MANUFACTURER?

- PRICE COMPETITION

PLUG COMPATIBLES FOR SUPERCOMPUTERS

- SMALL MARKET
- ADVANCED TECHNOLOGY
- OCCASIONALLY A SMALL VENDOR WILL TAKE THE RISK

INTERMEM 8330 FEATURES

RAM DISK FOR THE CRAY-1
- 8 TO 32 MILLION WORDS (64K CHIPS)
- 32 TO 128 MILLION WORDS (256K CHIPS)
- VOLATILE MEMORY - DYNAMIC MOS RAM

FOUR ACCESS PORTS
- ABOUT 200 MEGA-BITS TRANSFER RATE ON CRAY-1 SYNCHRONOUS CHANNEL
- ALL FOUR PORTS SIMULTANEOUSLY ACTIVE WITH ABOVE RATE
- CRAY-1 CHANNEL LIMITED

PHYSICAL CHARACTERISTICS
- CABINET 64" x 30" x 72"
- AIR COOLED
- 400 OR 60 CYCLE CONDITIONED POWER

MAINTENANCE FEATURES
- DEGRADABLE IN 1M WORD UNITS, ADDRESS CONTINUITY PRESERVED
- FAILING PORTS CAN BE TURNED OFF
- MAINTENANCE UNIT - OFF-LINE DIAGNOSTICS, ERROR LOGGING.

USAGE AT IDA

- DIRECTORY TABLES
  - DOUBLE WRITE TO PROTECT AGAINST POWER LOSS
- SWAP IMAGES
- UTILITIES, EXECUTIVE AND EDITOR
- APPLICATIONS USAGE
  - ACCESSED LIKE AN ORDINARY DISK
  - USER MUST REQUEST ALLOCATION ON THE RAM DISK.

EFFECT ON PRODUCTIVITY

SYSTEM EXERCISER EMULATING PRIME SHIFT INTERACTIVE USER IS ABLE TO RUN ABOUT TWICE AS MANY SHORT JOBS WITH INTERMEM 8330.

USER BANDWIDTH TEST. MADE DURING DAYTIME PRODUCTION TIME. TEN TRANSFERS AVERAGED IN EACH CATEGORY.

DIRECTION	BLOCK SIZE	DISK MBITS/SEC	IMEM MBITS/SEC	IMEM RATE/DISK RATE
WRITE	512	1.350	21.130	15.649
WRITE	2048	7.758	63.341	8.164
WRITE	9216	18.048	135.956	7.533
WRITE	92160	29.224	187.111	6.403
READ	512	1.010	21.897	21.673
READ	2048	7.632	62.113	8.139
READ	9216	18.044	123.896	6.866
READ	92160	21.349	136.461	6.392

INSTALLATION EFFORT

EXPERIENCE WITH SERIAL # 1's

"PREVIOUS PAINFUL EXPERIENCES DO NOT MAKE THE NEXT ONE LESS PAINFUL."

"EACH PAINFUL EXPERIENCE PRESENTS ITS OWN UNIQUE VARIATIONS."

INTERMEM DID NOT HAVE ACCESS TO A CRAY-1 FOR FULL TESTING.

NEGOTIATED USE OF IDA CRAY-1 FOR FINAL HARDWARE DEBUGGING.

IDA WROTE CRAY-BASED MAINTENANCE DIAGNOSTICS.

HARDWARE DELIVERED DECEMBER 28, 1983, ACCEPTED JUNE 22, 1984.

PROBLEMS -

EXCEPT FOR A FEW EARLY FAILURES, CRAY-BASED DIAGNOSTICS AND INTERMEM MAINTENANCE STATION WERE NOT USEFUL.

MOST FAILURES WERE DEMONSTRABLE ONLY UNDER THE PRODUCTION OPERATING SYSTEM.

OPERATING SYSTEM WAS CHANGED TO PROVIDE ENHANCED DIAGNOSTIC INFORMATION.

ALMOST ALL PROBLEMS WERE WITH THE TRANSFER PORTS.
- SOFT ERRORS
- INTERMITTENT DATA GARBLING

RELIABILITY

PERFECT OPERATION SINCE JUNE 4 EXCEPT FOR THREE EPISODES OF SOFT FAILURES WHICH RESULTED IN NO DOWNTIME.

INTERMEM COMPANY

- SMALL
- COMPETENT
- THEY KEEP THEIR WORD

Performance Evaluation and Optimization Workshop

J. Goirand

Compagnie Internationale de Services en Informatique (CISI)
Paris, France

Dr. Joseph A. Parker, from TDC (Technology Development of California) presented an analysis of supercomputer user load characteristics. Through the analysis of data from a number of supercomputer installations, and the results of discrete simulations, conclusions were drawn concerning differences and common features which characterize user loads in differing environments. The environments studied include both batch and interactive usage, CRAY-1S and CRAY X-MP, networked and stand alone processors.

Then two talks about optimization were presented.

Harry L. Nelson from LLNL (Lawrence Livermore National Laboratory) talked about Tournament Chess and optimization of the CRAY-BLITZ program. Tournament Chess is a real-time problem; competitors are allowed a limited amount of time to decide about which moves to make. He intends to win this year's ACM tournament by throwing more computer power at the problem than the other programs. He has arranged for dedicated use of the XMP/48 at Mendota Heights for the four tournament rounds. Moreover, his code will be multitasking using all four processors. The bottom line is: with FORTRAN, the player loses; with CAL, he draws; with multitasking, he wins!...

Jean-Claude Adam from CCVR (Centre de Calcul Vectoriel pour la Recherche) talked about his experience on optimization of particles codes in plasma physics.

After a short description of the equations to be solved, the problem of optimization of the corresponding code was examined. The optimization is aimed at reducing the computer time, increasing the CPU usage and maximizing the size of the physical problem that can be solved in a given amount of memory. The increase in computing speed is obtained both by maximizing the number of operations in vectorizable loop and organizing data to improve the transfer rate from memory. The CPU usage is increased by using asynchronous I/O and packing two words of data into one CRAY word. This last point also maximizes the size of the physical system that fits into the memory. The sustained rate is obtained above 50 Mflops and CPU usage is 90% to 100%. Future trends were briefly examined.

NETWORKING WORKSHOP

Dean W. Smith

ARCO Oil and Gas Company
Dallas, Texas

The networking session opened up with talks on networking topics by Eugene Goldberg and Dick Watson. The floor was then opened up to a lively discussion on networking where questions were directed from the floor to the speakers and several CRI employees involved in networking.

Introduction to Protocols, Eugene Goldberg, IDA

Eugene Goldberg's talk consisted of an overview of why protocols exist, their definition, and criteria used in making their choice. Protocols exist to support the connection of computers to their peripherals. They imbed in the data information about the data, establishing their grammar, and giving the data meaning. Protocols have in the past been based on the master/slave, front end/back end distinctions.

Eugene's basic conclusion was that no protocol can be best in all cases. The choice is based on trade-offs of cost, performance, hardware and the capabilities required of the network.

1. Generality and portability cost - in both initial costs and in performance characteristics.
2. Isolation of protocol levels is paramount for their utilization.
3. Speed and accuracy of the medium is a design concern.
4. All protocols in a network need not be the same.
5. How many of the computers in the network are under one's control.

Finally, the choice is driven by the application. It determines: the maintenance of queues, the computation of check sums, and the complexity of the recovery schemes.

Eugene concluded his talk by posing the following question: in an environment containing CRAY's, front ends, and personal computers, which is the peripheral?

Dick Watson, Lawrence Livermore

Dick Watson reviewed with us the Livermore Interactive Network Communication System (LINCS) utilized at Lawrence Livermore. LINCS is based on a "software bus" architecture designed to be easily implemented on a variety of systems.

Dick stressed the following points:

1. A network should be implemented as portable as possible.
2. A network must be well layered - LINCS has 7 layers.
3. The external interfaces should be well defined. LINCS utilizes well defined modular interfaces which enables new components to be added to the network with a minimum of modifications.
4. A network should be designed to enable the implementation of the network on an existing system by imbedding the network within existing architecture.

One point brought up as an aside, but which caused considerable concern in the session, was the revelation that the maximum number of I/O's achievable by a CRAY/IOP and an NSC connector is only 520 per second. This is a point of considerable concern when considering the maximum data flow rate possible with the current CRAY hardware. It would be extremely interesting if this value can be independently reproduced.

Networking Review

CRI's networking requirements will be

undergoing a technology review this January and we've received a verbal commitment by Don Mason to review with us the results of their review at the next CUG.

Goals of the Networking Workshop

One purpose of the Networking workshop is to provide input to Cray Research as to our future networking requirements with CRAY computers. I believe this breaks down into 3 categories with which we can influence CRI's software and hardware development:

1. Hardware

   a. What kinds of devices will we need access to by a CRAY?

   b. Will the CRAY's networking capabilities be limited by its current hardware configuration?

   c. What data speeds will we require across the network?

   d. What types of CPU's will exist in the network? What are to be their relationships? Do we need a CRAY to CRAY interface?

2. Applications

   a. What new user applications will be possible in 5 years?

   b. What system applications will we require to provide networking and network resource management?

   c. How can we control the processing of jobs, i.e. load balancing, in a network that may contain several super computers?

   d. Since the level of expertise and hardware is variable, what direction can CUG give CRI for our requirements for software on front ends?

3. Protocol is likely to be the most difficult of the 3 aspects to arrive at consensus within the Networking Workshop. It is unlikely that we will be able to resolve some of the fundamental differences between the bus-based networks and the IBM gateway type network. A bus-based network characteristically has many applications on many machines sharing a single bus; whereas a gateway network utilizes a few centralized stations directly connected to the CRAY.

   a. What protocol models can we expect Cray Research to employ?

   b. Where are the network interfaces to be?

   c. Where can CRI best direct its resources?

CRI is interested in our networking ideas and requirements. It is our ultimate benefit to provide them with as much valid information on our requirements as possible.

Networking Future

I believe that we can begin to see some aspects that will drive networking development in the coming years.

1. The NETEX/Subsystem service facility - this capability will likely be utilized by other 3rd party vendors.

2. The ISP currently announced only for the MVS environment - it is interesting to note that there were several questions regarding CRI's plans to offer ISP products for other front ends.

3. An awareness of the distinctions between control information and data, particularly as it regards paths and capacity.

4. The presentation of super computers as functionally compatible entities in a computer network.

Post Script

I would like to take this opportunity to invite all interested parties interested in making presentations at future networking workshop sessions at CUG to come forward. Topics of interest can include, and are not limited to, tutorials on networking topics, overview of current network, or any unique networking plans that you may have.

With this in mind it might be useful to determine what constitutes a station workshop issue versus a network workshop issue: if it concerns how one uses a CRAY today then it probably is station topic; if it concerns how one would use a CRAY five years from now, then it probably is a networking topic.

Disclaimer: this report is taken from my notes of the front end session. Any errors or inconsistencies are probably mine and you should contact the speakers for any clarifications.

MULTITASKING WORKSHOP

Mary E. Zosel

Lawrence Livermore National Laboratory
Livermore, California

Four talks covering a spectrum of multitasking activities were presented in this session. Abstracts of these talks are included here. CRI again emphasized that they would welcome input on features to aid debugging of multitasking applications.

DEBUGGING PROGRAMS UNDER EIGHTH-EDITION UNIX

Thomas J. Killian
AT&T Bell Laboratories

The Eighth Edition of the Unix(TM) operating system provides powerful debugging techniques both at the kernel and applications level. The kernel supports the /proc file system which makes the address space of every process in the system available as a file. Read and write access are permitted; ioctl's are available for process control, such as stop/go, and selective intercepting of signals. There is also an ioctl which returns an open file descriptor for the process' text file, giving immediate access to new, detailed symbol tables produced by the C compiler; these tables include variable names and types, structure definitions, source line numbers, block levels, and source file names. Thus the full path from process image to source code is available without further information from the user.

The window-based interactive debugger pi, developed by T. A. Cargill, is the first major user of /proc. It can control multiple processes dynamically and asynchronously. Each of its windows is a projection of a process onto a particular view, such as source code, stack frames, global variables, assembler instructions, etc.

These techniques may be extended to multitasked processes if the kernel is designed with the proper primitives. We will explore how this can be done on a Unix system for the Cray.

Reference: T. J. Killian, Processes as Files, USENIX Association, Summer Conference Proceedings, 1984, pp 203-207.

RESULTS OF MULTITASKING EFFORTS AT NASA AMES RESEARCH CENTER

Catherine Schulbach
NASA/Ames

Multitasking can be used effectively on problems of interest to Ames. Speed-ups over 1.9 can be achieved. The problems encountered in modifying codes to make use of multitasking were due to using a development system and to the side effects occurring in old Fortran codes. With the addition of task local common, the constructs provided are adequate for doing multitasking. NASA Ames would like to see the addition of high-level multitasking constructs to Fortran and is pursuing the development of such constructs.

DEBUGGING MULTITASKED APPLICATIONS

Peter A. Rigsbee
CRI

Cray Research has invested resources over the last year in the area of tools and utilities useful in debugging multitasked applications. This presentation summarizes these and presents examples of listing outputs.

In the area of design analysis, FTREF provides global cross-reference information on a CFT application. Reading the CFT module cross-reference listings, FTREF produces listings on (a) common block usage, (b) the static calling tree, and (c) usage of CRI lock variables.

A number of specific test tools and techniques have been developed or are under development. CFT will provide an option to preset local, stack variables to an unusable value, which allows identification of uninitialized variables. TSKLIST, a subroutine, will display the current status of all user tasks. TSKTUNE can be used to limit the number of physical processors to one, thereby eliminating a class of timing problem. An unsupported timer trace package outputs task sequencing information.

DEBUG and DUMP have been enhanced for multitasking. DEBUG will dump symbolic data for all active COS tasks and has a new parameter to dump data for all user tasks. In addition, status and statistics maintained on tasks, stacks, and the heap are output. DUMP will dump registers for all active COS tasks. Finally, if a program aborts due to deadlock among the user tasks, the library will automatically output status of each task.

SID is internally single-threaded, and locks itself whenever entered for a breakpoint. Other tasks continue execution, but will be suspended if they, too, reach a breakpoint.

Future work on debugging has two parts. First, Cray has assigned an analyst to explore multitasking debugging aids. The emphasis will be on tools that help people modify and debug a multitasked program; input from such users is a key component. Second, Cray is beginning an effort for a new debugging system. This new system is planned to support a number of needed features, including Pascal, C, segmented codes, and multitasking.

LOW-LEVEL PRIMITIVES FOR SUPPORT OF MULTITASKING

Bonnie Toy
LLNL

The model of multitasking being adopted at LLNL differentiates between a task and a process. A process is preemptively schedulable by the operating system and is associated with a processor, while a task is a logical expression of the parallelism in a job. The user specifies tasks and synchronizes access to shared data via calls to library routines. The operating system needs no knowledge of the tasking structure or synchronization required by a job. Some unfortunate interactions with the optimizer may occur because tasking and synchronization are implemented as ordinary library calls, and are not recognized as special by the compiler.

Based on this model, a set of tasking primitives has been designed and is being implemented. The set includes the CRI tasking library routines, but also contains primitives which might be used by a compiler doing implicit parallelization, and primitives with an interface appropriate for strongly typed languages such as Pascal. The LLNL tasking library package is being first implemented for the CTSS operating system to run on the Cray-1 machines. This will allow users to experiment with simulated multiprocessing. The CTSS version of the library will be ready by the end of 1984, and the XMP version, which will run on the NLTSS operating system and implement true multiprocessing, will be available in the first quarter of 1983.

GRAPHICS WORKSHOP

Helene E. Kulsrud

Institute for Defense Analyses
Princeton, NJ

The Graphics Workshop has three functions: to inform users of the status of graphics packages currently available on CRAY computers, to present in depth reports on popular packages, and to make requests from CRI for future graphics needs. John Aldag presented an update of his previous status report (given at the Oxford Meeting). In addition, John showed some slides and videotapes which had been made on CRAY computers with various programs. He concluded his talk with a tape made by LUCAS films on the XMP.

A report on NCAR graphics was given by John Humbrecht of NCAR. The graphics project at NCAR has developed a portable device independent graphics system. It has a broad range of utilities including map projection, contouring, 3D prospective representations, topography and others. The system is written in a subset of ANSI/77 Fortran. It is installed on computers ranging from CRAYs to PDPs, driving a variety of graphic displays. Included in the talk were examples of the NCAR utilities and animations using them.

Newt Perdue of NASA/Ames showed the film that the keynote speaker had been unable to present. There were no requests to CRI.

A BRIEFING PAPER ON THE STATUS OF GRAPHICS ON CRAY COMPUTER SYSTEMS

John Aldag

Cray Research, Inc.
Mendota Heights, Minnesota

INTRODUCTION

This paper summarizes the present status of computer graphics on CRAY computer systems. It reviews the present computer graphics support activities within Cray Research, Inc. and describes the known graphics software systems available on CRAY computers for the major application areas. Further information can be obtained by contacting the Applications Department at Cray Research, Inc.

Graphics Support at Cray Research

For many scientific and engineering applications, computer graphics is a natural extension of computer processing. The problems being analyzed on CRAY computers, such as 3-D fluid flow, molecular dynamics, or large deformation dynamic response in solids, are often so complex that full visualization is the only way to understand the results. It is natural, therefore, that CRAY computer systems should provide state-of-the-art graphics capabilities both for scientists and engineers whose graphics are a tool for interpreting analysis results, and for graphics professionals whose graphics are end products in themselves.

In 1981, seeing a rapidly growing need for high quality computer-generated graphics in many areas of application, Cray Research, Inc. began a project to support the graphics requirements of CRAY users. Today, through the Cray Research Applications Department, the company provides support for graphics in two ways. First, we are actively working with vendors and authors of the best and most widely used graphics software to ensure that state-of-the art graphics software is available on CRAY computer systems. We are using this software as it is used by a typical user. Secondly, Cray Research is acquiring a select set of state-of-the-art graphics hardware devices and using them with graphics generated on CRAY computers. Graphics hardware systems presently being used for applications support include RAMTEK 9400 and 9460 high-resolution color display systems, a MATRIX camera system, RAMTEK 6211 and 6221 color terminals, an Evans and Sutherland PS-300, SEIKO GR-2414 and GR-1104 terminals, a Raster Technologies Model 1/20, 2 Apollo workstations, and IBM 3250 and 5081 systems.

Our active participation in the development of graphics on CRAY systems ensures compatibility between CRAY computer systems and the graphics hardware and software which will be used with them. It allows Cray Research to respond to the present graphics requirements of CRAY users, to anticipate the user's future needs and to prepare for them, and finally, to foresee new applications and marketplaces which will demand the power of a CRAY for graphics.

Graphics Software

The graphics software systems available on CRAY computers are described briefly here. A more complete list and description can be found in the Directory of Applications Software, which is published periodically by the Applications Department.

GENERAL GRAPHICS

All the major device-independent line drawing graphics systems are available on CRAY computers. These include CGS, DI-3000, DISSPLA, the NCAR graphics library, and TEMPLATE. These standard graphics systems allow applications analysts to focus their attention on development of functional and effective application programs without being concerned with the operational details of the various graphics output devices. DI-3000 is available for testing and demonstration on the CRAY systems at the Cray software development facility in Mendota Heights, Minnesota (MH). It provides the foundation graphics system for a variety of applications at Cray Research. Preliminary arrangements have been made to use and support DISSPLA from ISSCO in a similar way.

A raster device-independent graphics system called UNIRAS is also available on CRAY systems and can be demonstrated in MH. Techniques for displaying two- and three-dimensional graphics appropriate for a range of applications are available through this system at a subroutine callable level. UNIRAS is used extensively at Cray Research to display the results of benchmarks and other software vendor and customer tests. UNIRAS has recently been used to display the estimated density function of a simulated universe evolved from a "big-bang" concept. A movie of these results enhanced the 3-D nature and made the results much more understandable. It was, in fact, shown on U.S. Public Television as part of a program on astrophysics.

Where appropriate, the graphics devices can be driven interactively from the CRAY with either the device-independent systems or device specific libraries. Alternatively, graphics data in the form of device-specific codes or a device-independent metafile can be transferred to the CRAY front-end for later display, saved on CRAY on-line disks for later use or transferred to CRAY on-line tapes for display on stand alone, high-resolution graphics recording devices. The Los Alamos National Laboratory system is one example of a fully integrated graphics system using many of the techniques noted above.

HIGH-BANDWIDTH GRAPHICS

Relatively high-bandwidth graphics output has been demonstrated at Cray Research via two methods:

1. Direct - DISPOSE. By making small changes to the VAX station, it is possible to treat graphics devices on the VAX as disk or tape drives and to DISPOSE graphics data directly to the device. Using RAMTEK 9400 and 9460 graphics devices, data rates in the range of 100-200KB/sec have been achieved using this technique. The technique has been incorporated as an option into the MOVIE.BYU, PATRAN, UNIRAS, and CSADIE software.

2. On-line graphics. Analagous in some ways to on-line tapes, this method uses a IBM-channel-to-UNIBUS conversion module from AUSCOM, Inc. to drive a RAMTEK 9465 directly from the CRAY/IOS. The data rate is estimated to be 200-300KB/sec. COS additions to support this capability are being tested and may be available in COS 1.14.

Where the application demands it, Cray Research will supply the technical specifications required to connect high-performance graphics hardware to high bandwidth CRAY computer channels. The RAMTEK Corp., Special Systems Division in Napa, CA, developed such an interface as specified by Digital Productions. Called the RAMTEK Film Recorder and interfaced to a 100MB/sec CRAY channel, the system will deliver graphics to a specially built frame buffer at a rate of about 40MB/sec.

GEOSCIENCE GRAPHICS

Major geosciences graphics programs which execute on CRAY computers are:

1. CPS-1, a widely used contour mapping system. CPS-1 can be demonstrated at Mendota Heights.

2. QUIK, a system for modeling seismic data acquisition in three dimensions using a ray-tracing algorithm. QUIK on CRAY computers can be demonstrated at MH.

3. UNIRAS, the raster graphics system noted above. Application-oriented subsets of UNIRAS provide support for the display of seismic, geological and LANDSAT data. It has been used recently to display the results of a system for wave equation modeling of seismic data acquisition.

   The modeling software, developed by Dr. Dan Kosloff at the University of Tel Aviv, uses a 2-D elastic and a 3-D acoustic algorithm and is highly optimized for CRAY systems. Conversion of the 3-D elastic algorithm is planned.

4. CSADIE, an image processing system originally developed at Los Alamos National Laboratory. CSADIE provides vectorized algorithms for image manipulation, compositing, BOOLEAN operations, edge detection filtering and enhancement. It can be used as a package or integrated at a subroutine level to an existing system. It is available from and supported by the Applications Department.

COMPUTER-AIDED MECHANICAL DESIGN GRAPHICS

Several well known programs for the display of solids are functional on CRAY computers:

1. MOVIE.BYU, which allows a designer to build a mechanical model using panels or solid elements and is best known for its animation capability. MOVIE.BYU, executing on a CRAY, was recently used to render 270 views of a Ford prototype automobile for presentation in a holographic display at the 1984 SIGGRAPH conference. The technique allowed a conceptual design to be evaluated in part without the construction of a physical model.

2. PATRAN, a system for pre-and post-processing of solids models for FEM Analysis. In an attempt to integrate mechanical design and analysis, the PATRAN system was converted for

execution on CRAY systems with the cooperation of PDA Engineering. It is supported by CRAY systems, by PDA and Cray Research for a wide range of graphics devices.

AERODYNAMIC FLOW GRAPHICS

A basic system for display of areodynamic flow results from the FLO programs is being developed. It will employ DI-3000 for its device-independent foundation.

GRAPHICS RESEARCH AND DIGITAL IMAGE SYNTHESIS

Several noteworthy and publicized graphics research efforts relied on the processing power of a CRAY computer:

1. Carla's Island, by Nelson Max at Lawrence Livermore National Laboratory. This short computer-generated movie demonstrates a ray tracing algorithm for simulating the reflection of sunlight from the surface of water.

2. MATHSCAPE, the image on the 1982 SIGGRAPH poster was generated using a CRAY by Mel Prueitt at Los Alamos Scientific Laboratory.

3. CHRYSALIS, a movie about the impact of computer graphics on our understanding of complex phenomenon. This movie was made at Los Alamos Scientific Laboratory and contains nearly ten minutes of graphics generated on CRAY systems.

4. Three film sequences in the SIGGRAPH '83 film show were generated on CRAY systems.

   The "EconoMars Earthtours" film from Patrick Weidhaas at Lawrence Livermore National Laboratories demonstrates the capability of digital terrain simulation. Data for the area of interest is drawn from a 7-gigabyte database of elevations for the continental U.S. and displayed from the chosen perspective. The application is interesting -- to help meteorologists understand how local terrain affects near surface wind patterns.

   "When Mandrills Rules the Heavens" from Sandia National Laboratories and "Mandala" from then Seibu Promotion in Tokyo are beautiful for their graphics and production. Both used a ray-tracing algorithm to model the physics of light on simulated but captivatingly beautiful models. Seibu (now Sedic) also produced the introductory film to the nightly news broadcast on Japanese Television NHK.

5. Presentations at SIGGRAPH '84:

   A computer generated hologram of a Ford prototype automobile mobile was displayed at the Cray Research vendors' exhibit.

   Four sequences in the first computer-generated OMNIMAX film were produced on CRAYs. The film premiered at the St. Paul Science Museum during the SIGGRAPH convention.

   A number of computer-generated film pieces generated on CRAYs were shown at the annual film shown. The highlight was "The Adventures of Andre and Wally B" from Lucas Film, generated in part on the CRAY X-MP/48. This piece used a full ray-tracing algorithm and demonstrated motion-blur in fast motion action parts. It was one of the first major applications to use the 'C' compiler on the CRAY.

6. "The Last Starfighter" is a reality, with nearly 25 minutes of film generated at Digital Productions on a CRAY X-MP/22. This feature length movie is now showing around the U.S. In addition, DP has produced numerous short sequences for commercial product promotion. A recent article in the Proceedings of the IEEE (January 1984) describes the unique hardware and software system built around a CRAY computer to design, render, display and film these digitally simulated images.

SUMMARY

Cray Research, Inc. actively seeks to support the graphics applications of CRAY users. CRAY computer systems have been used to generate graphics for a broad range of application areas and for many different graphics hardware devices.

In addition, the processing power of CRAY computers has opened new avenues for the use of computer-generated graphics in scientific and engineering applications and digital scene synthesis.

COS Workshop

David Lexton

ULCC
London, England

1. COS Experience Panel

Jim Sherin (Westinghouse) said that it had taken one year to get 1.12 into production. COS 1.11 had been modified to run 1.12 binaries. Integration of local modifications had been done using dedicated discs until the installation of SN40 in April 1984. During the past 6 months the major local projects were:

(a) connecting the two Crays by sharing directories through a Cray/NOS dual station;

(b) providing the same level of user support on COS 1.12 as on 1.11;

(c) providing the ability to run old binaries under COS 1.12 to cater to users who have to deal with the NRC for certification of codes.

Introduction of the new calling sequence at the same time as COS 1.12 was felt to be too much for users to bear. Westinghouse therefore intend to make the transition over the next 6 - 9 months with the aim of having all users on the new calling sequence by the time COS 1.14 is available. Even so, a continuing requirement to support the old calling sequence was anticipated. Westinghouse plan to bypass COS 1.13.

Lothar Wollschlager (KFA) reported that their X-MP SN104 was running COS 1.12 BF2 with Privacy and Security turned on. There had been no hardware crash for 4 months. It was planned to put 1.13 BF1 into production on October 15th and then investigate the use of multitasking. Local modifications include one to CSP to take the first six characters of the job name as the user number and four new system calls in EXP. Three of these are for the accounting system and get a job's priority, change a job's priority and rerun a job. One is for on-line diagnostics and crashes the system if the same diagnostic fails twice.

Ray Benoit (Environment Canada) said that their machine was the first Cray in Canada and had been installed for one year, their main customer being the Canadian Meteorological Office. The first release of COS 1.12 had proven unusable with the new calling sequence and PDM privacy. The 1.12 NOS station had been even worse. Stability had been reached at BF2 although problems with the station and on-line tapes remained. The new calling sequence was now the default, after much testing. CRI documentation on the new calling sequence was plentiful but unsuitable for the end user. PDM privacy is expected to be a major undertaking and will require much testing. The scheduler has been modified to allow high priority jobs to dominate. The approach is similar to that of ECMWF and provides three priority bands: foreground, midground and background. It is intended to do the same for memory requests. During testing, the 1.13 NOS station has shown some significant bugs.

Kent Koeninger (TDC) said that TDC had made a very smooth transition to COS 1.13 in September. It was more stable than 1.11 and much more stable than 1.12. CFT 1.13 on the other hand had both aborted and compiled wrongly.

Larry Yaeger (Digital Productions) reported that it had been necessary to retrofit tape code from 1.14 to 1.12 BF2 and that lack of priority control had caused job scheduler problems.

Bill Kelly (TDC) referred to a problem with SCP asking the station for information on on-line tapes. Westinghouse had a fix for this.

2. Error Recovery Project

Walt Anderson (CRI) reported on the current status.

3. SPR Processing

Gayle Smith (CRI) described the new SPR procedures.

4. COSSIC Report

Dave Lexton (ULCC) reported on the meeting of COSSIC held on Monday 1st October 1984. The recommendations on the future of COS that had been agreed by the workshop in Paris were being implemented. CRI had made it clear that UNIX was an experiment whose future depended upon its performance. In any case, CRI considered that there should be no implications for COS development. At the same time comments from COSSIC on UNIX would be welcome. The Board of Directors subsequently agreed that the committee should be concerned with Cray-supported operating systems in general. The role of the committee had been discussed and the desirability of involving it in a not too formal design review of new software features was agreed.

On recent experience with COS, the problems in making the move to 1.12 were discussed as was the fact that sites are now getting significantly behind CRI in the level of COS they are running. The desirability of minimising or removing interdependencies between COS and the product was stressed.

Volunteers for membership of the committee were sought so that the members are now:

R. Benoit, Dorval
C. Hilberg, ECMWF
C. Kimble, Boeing
M.R. Lewis, Chevron
D. Lexton, ULCC (Chairman)
D.M. Mason, CRI (Cray contact)
J.J. Sherin, Westinghouse
L. Wollschlager, KFA
L. Yaeger, Digital Productions

ADDITIONAL REPORTS

PRESIDENT'S REPORT

M. G. Schomberg

AERE-Harwell
England

This 14th Meeting of CUG marks a turning point for the Cray User Group. CUG is now operating under new Bylaws and on Tuesday you elected the first Board of Directors. I would like to remind you of what I said regarding the reasons for new Bylaws and the change in the way in which CUG is governed. These reasons are:

- to cope with the growing size of CUG membership;
- to provide better continuity between CUG meetings;
- as a step towards incorporation;

and I would add a fourth

- to have greater influence with Cray Research on the more strategic issues which are of concern to the membership.

Your new Board of Directors consists of:

- Michael Schomberg as President
- Laney Kulsrud as Vice President
- Karen Friedman as Secretary
- Bob Price as Treasurer
- Jacqueline Goirand)
- David Lexton) Directors-at-large
- Joe Thompson)

On behalf of myself and the other members of the Board I would like to thank you most sincerely for the trust you have placed in us. We are indeed honoured.

I can report that your newly elected Board has held two meetings and has made considerable progress in establishing policy positions on a wide range of issues. These policy positions will be distributed, together with the new Bylaws, with the Proceedings of this Conference.

You will recall that under the new Bylaws there is one class of membership: Installation Membership. Associated with this is the concept of the Installation Delegate who is the designated spokesperson for that installation. It is vital that we have the name of the Installation Delegate for each installation - and that person should be in a position to truly represent that Installation. It is the Installation Delegate who is responsible for voting, not only for the Board of Directors, but on major policy issues.

The Programme Committee has been re-constituted under the Chair of the Vice President, Laney Kulsrud. This committee consists of individuals who are responsible for different parts of the CUG meeting programme. There are at present some 12 major areas which have been filled by suitably qualified "volunteers." In some instances, these members will be working with one other person in the same area; sometimes a small sub-committee will be formed. Also, it has been necessary to recognise that there are very close links between areas. Some of the Programme Committee members may also be workshop chairman. The Programme Committee is aiming to plan at least two conferences ahead with the objective of making the program even more valuable to delegates.

The User Requirements Committee has been progressing under the chairmanship of Steve Niver. Its objective is to consider user requirements from the Special Interest Committees, Workshops and Installations and perform some screening and routing with the aim of determining the really important issues which CUG should be taking up with Cray, especially of a more strategic level. We also hope that Cray will find that they will usefully interact with this committee if they wish to seek the consensus view of the membership on proposals or priorities. These issues, whether from the membership or from Cray, will be subject to a postal ballot by Installation Delegates.

Incorporation

This is now becoming an even more important issue - especially with changes in the US tax laws. It is important in the context of limiting personal liability. Our Treasurer, Bob Price, will be pursuing this issue during the next few months.

In Conclusion

Our aim is to make CUG a more effective organisation, to increase the value of conferences and to interact effectively with Cray. We also aim to ensure that it is fun.

CUG STEERING COMMITTEE REPORT

Karen Friedman

National Center for Atmospheric Research
Boulder, CO

The Fall 1984 meeting of the CRAY Users Group (CUG) Board of Directors was brought to order at 2:00 pm on Monday, October 1, 1984, at the Hyatt Regency Hotel in San Francisco, CA, by Jacqueline Goirand, current CUG Chair.

Members present were:

	Name	Organization	Position on Steering Committee
*	Jacqueline Goirand	CISI	Chair
	Mary Amiot	CRI	CRAY Administrative Representative
	Mick Dungworth	CRI	CRAY Technical Representative (current)
	Dave Sadler	CRI	CRAY Technical Representative (future)
*	Glenn Lewis	TDC	Host, Fall 1984
*	Karen Friedman	NCAR	Chair, Publications
*	Gary Jensen	NCAR	Chair, Operations Workshop
*	Helene Kulsrud	IDA	Chair, Program Committee
*	Gene Schumacher	NCAR	Member-at-Large
	Michael Schomberg	AERE-Harwell	By-Laws Collaborator
*	Joe Thompson	LANL	Representing Gerry Melendez, Chair, CTSS Workshop
*	Mostyn Lewis	Chevron	Chair, I/O Workshop
	Ann Cowley	NCAR	Chair, Nominating Committee
*	Bob Price	Westinghouse	Host, Fall 1983
*	Sven Sandin	SAAB-Scania	Host, Spring 1985
*	Stephen Niver	BCS	Chair, Requirements Committee
*	Dean Smith	ARCO	Chair, Networking and Front Ends Workshops

* - Voting Member

The Committee welcomed Sven Sandin as the host for the Spring 1985 CUG meeting.

The minutes from the Spring 1984 CUG meeting were read and approved.

Ann Cowley announced the following slate of persons selected for nomination to office to the Board of Directors under the new Bylaws:

President	- Michael Schomberg
Vice President	- Helene Kulsrud
Treasurer	- Bob Price
Secretary	- Karen Friedman
Directors-at-Large	- Joe Thompson Dave Lexton, ULCC Jacqueline Goirand

It was suggested that the list of nominations be made available to the conference participants so that they may make additional nominations by petition, should they so desire.

The next slate of officers will be elected at the Fall 1985 CUG Meeting in Montreal, Canada. The voting procedure will be one vote per site, with the Installation Delegate casting the vote. The Installation Delegate shall receive a voting form in his/her meeting registration materials.

Helene Kulsrud discussed the role of the Program Committee. She stated that the group had made a great deal of progress in working together as a committee.

Personnel changes within the Committee included the resignation of Andy Anderson (ARCO) and the absence of David Kelvin (AERE-Harwell) at the CUG-San Francisco meeting. Dean Smith volunteered to chair both the Networking and Front Ends workshops previously chaired by David Kelvin. Jacqueline Goirand was still chairing the

Optimization workshop but was seeking someone to chair the Performance Evaluation session. Margaret Simmons (LANL) has returned as chair of the Libraries Workshop and Mary Zosel (LLNL) is chairing the Multitasking session. Dave Sadler is replacing Mick Dungworth as the CRI representative on the Program Committee. Helene thanked Mick for his assistance and welcomed Dave onto the Committee. Sven Sandin was the newest site member on the Program Committee since his company, SAAB-SCANIA, is hosting the next CUG meeting.

The format of the workshops and general sessions for CUG-San Francisco seems to be well planned with good overlap in subject matter between concurrent sessions.

The Call for Papers that was enclosed with the registration materials generated no response. Helene intends to enclose a Call in the next mailing, directed at the Technical Contact. This person would contact those persons at his/her organization who had not yet given a presentation at one of the CUG meetings. It is hoped that this effort will generate more response than in the past.

The CUG-San Francisco meeting is the first where a representative outside the CUG community has been requested to give a presentation (Network Systems Corporation). Helene recommended that if this trend continues, a policy should be developed as to the suggested format and content for these speakers, i.e., no marketing or sales information would be mentioned. She would also like to see a fund developed for keynote or other speakers who would be particularly interesting topically but could not afford to travel to the meetings.

Stephen Niver assumed the responsibility for the Wishlist Committee from Ann Cowley. That Committee, which is still in the organization mode, is now called the Requirements Committee. There is no longer a Wishlist. All work done on Requirements issues will be discussed at the CUG meetings, with site representatives bringing pertinent issues to the sessions. Stephen stated that only six sites responded to his balloting for Committee representatives. He is attempting to attract more representation on the Committee, especially from foreign sites. He hopes that the Installation Delegate will take the major responsibility for discussion of current items and the introduction of new items on the Requirements List, especially beyond the software considerations.

In the area of finance, Jacqueline stated that all expenses incurred for the CUG-Paris meeting were covered by receipts. Bob Price stated that each site should try to have a zero balance after each meeting to simplify the transferring of money, especially between foreign and U.S. sites, and the reporting of this money to the IRS.

Michael Schomberg stated that the new Bylaws should be printed in the next issue of the Proceedings, and congratulated Karen Friedman on a job well done on the Proceedings from CUG-Paris.

The future CUG meetings are as follows:

Date	Location	Host
Spring 1985 (May 7-9, 1985)	Stockholm, Sweden	SAAB-Scania
Fall 1985 (October 1-3, 1985)	Montreal, Canada	Environment Canada
Spring 1986	Seattle, WA	Boeing Computer Services
Fall 1986	West Germany	DFVLR
Spring 1987	New York	Grumman Data Systems
Fall 1987	Bologna, Italy	CINECA
Spring 1988	Sunnyvale, CA	Lockheed Missile and Space Co.

With no further business, the meeting was adjourned at 3:30 pm.

REPORT OF THE PROGRAM COMMITTEE MEETING

Helene E. Kulsrud

Institute for Defense Analyses
Princeton, NJ

The meeting of the program committee was held at 3:30, on October 3. Eleven areas of interest were specified and various members agreed to be responsible for each. These members will organize one or more sessions (previously called workshops) in their areas, or joint sessions when the subject matter overlaps. These areas are:

Cray Operating Systems	- David Lexton
CTSS	- Jerry Melendez
Front Ends	- Dean Smith
Graphics	- Helene Kulsrud
I/O	- Mostyn Lewis
Languages	- Margaret Simmons
Libraries	- (European)
Multitasking	- Mary Zosel
Networks	- Dean Smith
Optimization	- Ann Cowley
Performance Evaluation	- J. Goirand

Michael Schomberg then read the policy procedure on the Program Committee that was enacted by the Board of Directors. The members of the Program Committee are to consist of the heads of the special interest groups, the past, present and next arrangements chairs, the past program committee chair, the technical proceedings editor, and other CUG members. The Vice President of CUG will serve as chair of the program committee. As a result of this policy, six special interest groups, were formed.

Cray Operating Systems - Chaired by David Lexton

CTSS - Chaired by Jerry Melendez

Languages and Libraries - Chaired by Margaret Simmons

Networking and Front Ends - Chaired by Dean Smith

Operations - Chaired by Gary Jensen

Performance, Evaluation and Improvement - Chaired by Ann Cowley

Each SIG would attempt to find a Vice-Chair on the opposite side of the ocean (except for CTSS). This would help with travel problems and obtaining speakers from the countries near the meeting site.

The problem of speakers not being informed of the meeting arrangements was raised again. It was suggested that the members of the program committee receive copies of the hotel registration forms from the arrangements chair to send on to the speakers.

Mary Zosel and Helene Kulsrud agreed to form a committee to select papers for the general sessions. A large number of suggestions for the general sessions were put forward by the committee and the other CUG members. In order to help people determine whether they will attend a CUG meeting, the preliminary schedule will be included in the first mailing. The call for papers will be sent separately to the technical representatives. Sven Sandin took responsibility for finding a keynote speaker. Michael Schomberg agreed to prepare a statement to be sent to non-CUG people who are invited to talk at CUG.

USER REQUIREMENTS COMMITTEE REPORT

Stephen Niver

Boeing Computer Services
Seattle, Washington

The role of User Requirements Committee (URC) is to represent, in the general sense, to Cray Research Incorporated (CRI) the feelings of the Cray User Group (CUG) with regard to specific items. It is our intent that items pursued by the URC not be bug fixes or minor enhancements but larger projects. To do this, the URC mails ballots to the CUG membership. The general procedure is described below.

Ideas for inclusion on the ballots may come directly from sites. Sites are invited to bring descriptions of items for consideration to each CUG or may be forwarded by a CUG committee. CRI may also suggest items for which they would like some feedback.

The URC will meet at CUG to discuss each proposed item. At this time, the results of the previous ballot and the CRI response will be discussed and recommendations as to the disposition of some items will be made. Items may be deleted from the list, for example, because CRI has implemented the request or because there was little interest in the item.

The URC will sponsor a session at each CUG to report to the membership the results of the latest poll.

Shortly after each CUG, a new ballot will be sent to the CUG membership and the cycle will begin anew.

Following are the results of the 1984 voting displayed in several different formats.

Figure 1 shows the items and their point distribution in the order in which they appeared on the ballot.

Figure 2 gives the results sorted by total number of "points" allocated to items. This figure shows the overall interest in each item.

Figure 3 gives the results sorted by items having the most number of sites allocating any points. This distribution gives an indication of how many sites are interested in an item.

Figure 4 gives the results sorted by the average points per item given that the item received any interest at all. This figure gives an indication of how strong the interest is for those sites that have any interest.

1984 CUG USER REQUIREMENTS BALLOT RESULTS
TOTAL RESPONSES

FEATURE TITLE	TOTAL POINTS	PERCENT POINTS	AVERAGE RESP	NUMB RESP
MULTI-MAINFRAME SUPPORT	125.0	10.4	31.2	4
CRAY-TO-CRAY COMMUNICATIONS	125.0	10.4	31.2	4
COS CODING STANDARDS	85.0	7.1	7.7	11
USER EXITS	300.0	25.0	25.0	12
INSTALLATION AREAS	150.0	12.5	13.6	11
TAPE DATASET PRIVACY	40.0	3.3	13.3	3
SYSTEM TUNING	215.0	17.9	19.5	11
SOFTWARE CONFIGURABILITY	148.0	12.3	18.5	8
ENHANCED SCILIB (WRITE-IN)	12.0	1.0	12.0	1

FIGURE 1

1984 CUG USER REQUIREMENTS BALLOT RESULTS
RESPONSES SORTED BY TOTAL POINTS

FEATURE TITLE	TOTAL POINTS	PERCENT POINTS	AVERAGE RESP	NUMB RESP
USER EXITS	300.0	25.0	25.0	12
SYSTEM TUNING	215.0	17.9	19.5	11
INSTALLATION AREAS	150.0	12.5	13.6	11
SOFTWARE CONFIGURABILITY	148.0	12.3	18.5	8
MULTI-MAINFRAME SUPPORT	125.0	10.4	31.2	4
CRAY-TO-CRAY COMMUNICATIONS	125.0	10.4	31.2	4
COS CODING STANDARDS	85.0	7.1	7.7	11
TAPE DATASET PRIVACY	40.0	3.3	13.3	3
ENHANCED SCILIB (WRITE-IN)	12.0	1.0	12.0	1

FIGURE 2

1984 CUG USER REQUIREMENTS BALLOT RESULTS
RESPONSES SORTED BY NUMBER OF RESPONSES

FEATURE TITLE	TOTAL POINTS	PERCENT POINTS	AVERAGE RESP	NUMB RESP
USER EXITS	300.0	25.0	25.0	12
INSTALLATION AREAS	150.0	12.5	13.6	11
COS CODING STANDARDS	85.0	7.1	7.7	11
SYSTEM TUNING	215.0	17.9	19.5	11
SOFTWARE CONFIGURABILITY	148.0	12.3	18.5	8
CRAY-TO-CRAY COMMUNICATIONS	125.0	10.4	31.2	4
MULTI-MAINFRAME SUPPORT	125.0	10.4	31.2	4
TAPE DATASET PRIVACY	40.0	3.3	13.3	3
ENHANCED SCILIB (WRITE-IN)	12.0	1.0	12.0	1

FIGURE 3

1984 CUG USER REQUIREMENTS BALLOT RESULTS
RESPONSES SORTED BY THE AVERAGE OF THOSE RESPONDING

FEATURE TITLE	TOTAL POINTS	PERCENT POINTS	AVERAGE RESP	NUMB RESP
MULTI-MAINFRAME SUPPORT	125.0	10.4	31.2	4
CRAY-TO-CRAY COMMUNICATIONS	125.0	10.4	31.2	4
USER EXITS	300.0	25.0	25.0	12
SYSTEM TUNING	215.0	17.9	19.5	11
SOFTWARE CONFIGURABILITY	148.0	12.3	18.5	8
INSTALLATION AREAS	150.0	12.5	13.6	11
TAPE DATASET PRIVACY	40.0	3.3	13.3	3
ENHANCED SCILIB (WRITE-IN)	12.0	1.0	12.0	1
COS CODING STANDARDS	85.0	7.1	7.7	11

FIGURE 4

Conference Closing Fall 1984 at San Francisco

J. Goirand

Compagnie Internationale de Services en Informatique (CISI)
Paris, France

On behalf of the CUG Membership I would like to formally thank the following:

Glenn LEWIS, Joanie ADAMS and their colleagues from TDC for organizing the most successful conference.

Laney KULSRUD who together with the members of the Programme Committee have worked very hard to prepare the excellent programme for this conference.

Ann COWLEY, Bob PRICE and Andy ANDERSON who have organized the elections of the Officers and the Directors.

All the members of the special interest groups and particularly Karen FRIEDMAN who is responsible for the Proceedings, which is a very difficult task.

CRAY Research, particularly Mary AMIOT and CRAY Western region, for their cooperation with the conference and for hosting the receptions on Monday and Tuesday evenings.

Mick DUNGWORTH for his very enthusiastic participation as representative of CRAY Research to the Steering Committee during last CUG meetings and good luck for his new responsibilities.

Finally I would like to say thank you to you the participants without whom we would not have a conference.

It is now time to say good luck and long life to our organization with the new By-Laws.

ADDITIONAL INFORMATION

CRAY USER GROUP MEETING PARTICIPANTS BY ORGANIZATION

Organization	Representatives
Advanced Computer Tech. Res. P.O. Box 18343 2020 S.W. Freeway Houston, TX 77098	Rob Blain
A.E.R.E. Harwell Oxfordshire OX11 ORA, England	Michael Schomberg A.E. Taylor
Air Force Weapons Laboratory AFWL/ADP Kirtland Air Force Base, NM 87117	Dann Q. Brewer
ARAMCO EXPEC Computer Center P.O. Box 10308 Dhahran, Saudi Arabia	Wayne Schmaedeke
ARCO Oil and Gas Company 2300 West Plano Parkway Plano, TX 75075	Karen Hendricks Dean W. Smith
AT&T Bell Labs 600 Mountain Avenue Murray Hill, NJ 07974	Tom Killian
Boeing Computer Services P.O. Box 24346 Seattle, WA 98124	Ronald D. Hochnadel Conrad Kimball Steven A. Niver James E. Roetter Howard Schmeising
Boeing Computer Services 565 Andover Park West Tukwila, WA 98188	Kenneth W. Neves
CEA-CEC BP 27 94190 Villeneuve St. Georges France	Martine Gigandet
CGG 6 Rue Galvani BP 56 91301 Massy Cedex France	Gou Dedranche

Chevron Oil Field Research Company P.O. Box 446 La Habra, CA 90631	Annabella A. M. Deck Mostyn R. Lewis Mark I. Sherman Robin C. Thornton
CINECA 6/3 Magnanelli 40033 Casalecchio Di Reno Bologna, Italy	Sanzio Bassini Giovanni Erbacci Gianna Fabiani Marco Lanzarini
CISI BP 24 91190 Gif-Sur-Yvette France	Jacqueline Goirand Regis Schoonheere
Cray Canada 4141 Yonge Street, Suite 302 Toronto, Ontario Canada M2P 2A8	Martin Buchanan John Maas
Cray Research - France 7 Rue de Tilsitt 75017 Paris France	Alex Messer
Cray Research GmbH Perhamerstrasse 31 8000 Munich 21 West Germany	Walter Holzmaier
Cray Research, Inc. 11710 Beltsville Drive Suite 500 Beltsville, MD 20705	John Stephens
Cray Research, Inc. Lowater Road Chippewa Falls, WI 54729	Jerry Brost Stuart Drayton Dick Morris Lou Saye Gary Shorrell Don Whiting
Cray Research, Inc. 5858 Westheimer, Suite 500 Houston, TX 77057	Larry Stewart

Cray Research, Inc. 1440 Northland Drive Mendota Heights, MN 55120	Sonya Anderson Sheila Bonamarte Nic Catrambone Mick Dungworth Brian Gaffey Neil Haggard Dick Hendrickson Thea Hodge Kathy Hollander Margaret Loftus Don Mason John Renwick Peter Rigsbee Dave Sadler Gayle Smith Karen Spackman Brian Walsh
Cray Research, Inc. 608 Second Avenue S. Minneapolis, MN 55402	Mary Amiot Walt Anderson Bob Ewald
Cray Research, Inc. 5776 Stoneridge Mall Road The Atrium, Suite 350 Pleasanton, CA 94566	Frank Chism Donna Derby Dave Kimball Brian Mackie Joel Newsom David Sundstrom Howard Watts Mike Wilhelm Bing Young
Cray Research UK Ltd. Cray House, London Road Bracknell Berkshire RG12 2SY England	John Flemming Malcolm Hammerton Stewart Ross
Digital Productions 3416 S. La Cienega Blvd. Los Angeles, CA 90048	Larry Yaeger
Ecole Polytechnique C2VR 91128 Palaiseau Cedex France	Jean Claude Adam Jean Pierre Andre
Electricite de France 1 Avenue Du General De Gaulle 92140 Clamart France	Yves Souffez

ELF-Aquitaine Tour Aquitaine Cedex 4 Paris La Defense 92080 France	Bernard Pria
Environment Canada 2121 Trans-Canada Highway Dorval, Quebec, Canada	Raymond Benoit
European Centre for Medium Range Weather Forecasts Shinfield Park, Reading, RG2 9AX England	Claus Hilberg
Exxon Production Research Company P.O. Box 2189 Houston, TX 77001	Doug Spragg
Ford Motor Company P.O. Box 2053 ECC Bldg. Dearborn, MI 48121-2053	Neil St. Charles Bob Treharne James G. Viculis
Framatome Tour Fiat Place de la Coupole Cedex 16 92084 La Defense France	Truong Thanh Xuan
General Motors Research 267 R.A.N.B. GM Technical Center Warren, MI 48090-9055	Ronald Kerry
Government Communications HDQ Room No. F/1210 Oakley Priors Road Cheltenham GL52 5AJ England	Alan F. Phillips
Grumman Data Systems Corporation 1111 Stewart Avenue Bethpage, NY 11714	Timothy Ertl James Poplawski George Rasor John Riordan
Grumman Data Systems Corporation Grumman Blvd. Bethpage, NY 11714	Lucien Kraner Noreen Wolt

Institute for Defense Analyses
Thanet Road
Princeton, NY 08540

Robert Cave
Eugene Goldberg
Helene E. Kulsrud
Lee P. Neuwirth

KFA Juelich
Postfach 1913
5170 Juelich
West Germany

L. Wollschlaeger

Lawrence Livermore National Laboratory
P.O. Box 808
Livermore, CA 94550

Barbara Atkinson
David Fisher
John E. Raneletti
Roger Skowlund
Robert Strout
Bonnie Toy
George Vranesh
Mary Zosel

Lawrence Livermore National Laboratory
4259 Emory Way
Livermore, CA 94550

Harry L. Nelson

Lockheed Missiles & Space Company
1111 Lockheed Way
Sunnyvale, CA 94086

Lee Coven
D. D. Foote
R. D. Parish
Jack Sherman
T. Doug Telford

Los Alamos National Laboratory
PO Box 1663
Los Alamos, NM 87545

Chris Barnes
Frank W. Bobrowicz
John Dragon
K. Jerry Melendez
Fred J. Montoya
Norman R. Morse
Nicholas J. Nagy
Floyd Segura
Margaret Simmons
Charles A. Slocomb
Joseph L. Thompson

MFE Computer Center
Lawrence Livermore National Laboratory
P.O. Box 5509
Livermore, CA 94550

Pamela Farnstrom
Michael Ganzberger
Cynthia Phillips
David Storch

NASA/Ames Research Center Moffett Field, CA 94035	Eric Barszcz John Barton Arthur Cullati Larry Flynn Frank S. King Hugh LaMaster Eugene Miya James "Newt" Perdue Clifford E. Rhoades, Jr. Karl Rowley Catherine Schulbach Marcie Smith
NASA Lewis Research Center 21000 Brookpark Road Cleveland, OH 44135	Charles W. Putt
National Center for Atmospheric Research P.O. Box 3000 Boulder, CO 80307	Ann D. Cowley Karen Friedman John Humbrecht Gary Jensen Robert Niffenegger Eugene Schumacher
National Security Agency Ft. George G. Meade, MD 20755	Robert Fruit Bruce Steger
Naval Research Laboratory 4555 Overlook Avenue, S.W. Washington, D.C. 20375-5000	Harvey Brock Carolyn Bryant Gary Flenner Judith Flippen-Anderson Rufus "Mac" McCulloch Dale Pfaff Ted Young
Network Systems Corporation 7600 Boone Avenue North Brooklyn Park, MN 55428	Jim Hughes
Phillips Petroleum Company 461 Information Center Bartlesville, OK 74004	David W. Glover Michael Guidry
SAAB-SCANIA AB S-58188 Linkoping Sweden	Stig Logdberg Sven Sandin

Sandia National Laboratories Kirtland Air Force Base Albuquerque, NM 87185	Phillip Campbell Ron Domres Arnold Elsbernd G. L. Esch Frank Mason Jack Tischhauser John Van Dyke
Sandia National Laboratories P.O. Box 969 Livermore, CA 94550	Dona L. Crawford R. E. Huddleston Hilary D. Jones Gordon Miller
Schlumberger Doll Research Old Quarry Road Ridgefield, CT 06877	Bob Snow
Senator fur Wissenschaft und Forschung Projektgruppe Parallelrechnen Heilbronnerstrasse 10 D-1000 Berlin 31 West Germany	Jurgen Gottschewski
SOHIO Petroleum Company 5400 LBJ Fwy., Suite 1200 Dallas, TX 75240	W. C. Simmons
Sverdrup Technology, Inc. ASTF Arnold Air Force Station, TN 37389	John L. Roberson
Technology Development of California Advanced Computational Facility NASA Ames Research Center Moffett Field, CA 94035	Mary Fowler Unfard Gibson Bill Kelly Shaun Kenney R. Kent Koeninger Terry Larson Glenn E. Lewis David Robertson Todd Welch Donald G. Zarlengo
Technology Development of California 2431 Mission College Blvd. Santa Clara, CA 95054	Frank Greene Harry Heard Joe Parker Paul Richards Russ Saunders Dick Wilson
United Information Services, Inc./CDC 2525 Washington Avenue Kansas City, MO 64108	Nate Losapio Leroy Tidwell

University of London Computer Centre 20 Guilford Street London, England	Christopher Lazou David Lexton
University of Minnesota University Computer Center 2520 Broadway Drive Lauderdale, MN 55113	Elizabeth Stadther
University of Minnesota University Computer Center 227 Experimental Engineering Blvd. Minneapolis, MN 55455	Kevin Matthews
Westinghouse Electric Corporation Power Systems Computer Center P.O. Box 355 Pittsburgh, PA 15230-0355	James R. Kasdorf Fran Pellegrino Robert Price James J. Sherin

CUG Site Contact List
January 1985

Air Force Weapons Laboratory
Computation Services Division
AFWL/ADP
Kirtland AFB, NM 87117

<u>Installation Delegate and Technical Contact</u>
Mike Gleicher (505)844-9964
MS 53

<u>Operations Contact</u>
Robert Meisner (505)844-0424
MS 53

Arabian American Oil Company
EXPEC Computer Center
Dhahran, Saudi Arabia

TELEX: 601220 ARAMCO SJ

<u>Installation Delegate</u>
Wayne Schmaedeke (011)966-3-87-65155
X-2660

<u>Technical Contact</u>
Alfred Anderson (011)966-3-87-61188
X-2650

<u>Operations Contact</u>
Gene McHargue (011)966-3-874-1945(or 3830)
Box 10356

Arnold Engineering Development Center
Central Computer Facility
Arnold Air Force Station, TN
37389

<u>Installation Delegate</u>
Randall Thomas (615)455-2611, x. 7263
AEDC MS 100

Atlantic-Richfield Oil & Gas Company
2300 Plano Parkway
Plano, TX 75075

TWX 910 861 4320

TELEX 73-2680

Facsimile Transmission DMS 1000(214) 422-3657

Installation Delegate
Dean Smith (214)422-6415
PRC - C2292

Technical Contact
B.Y. Chin (214)422-6627
PRC - 2211

Operations Contact
Chuck Murphy (214)422-6612
PRC - 5141

Atomic Energy Research Establishment
Harwell, Oxfordshire
OX11 ORA, England

TELEX 83135 ATOM HA G

Installation Delegate
A. E. Taylor 0235-24141, x.3053
H 7.12

Technical Contact
Don Sadler 0235-24141, x.3227
Bldg. 8.12

Operations Contact
Michael Schomberg 0235-24141, x.3263
Bldg. 8.12

Atomic Weapons Research Establishment
Aldermaston
Reading, RG7 4PR
England

TELEX 848104 or 848105

Installation Delegate
L. M. Russell 07356-4111, x.6678

Technical Contact
P. A. Janes 07356-4111, x.4045

Operations Contact
M.D.P Fasey 07356-4111, x.6491

AT&T Bell Laboratories
600 Mountain Avenue
Murray Hill, NJ 07974

TELEX 13-8650
Facsimile (201)582-2608
(201)582-6934

Installation Delegate and Technical Contact
Peter Nelson (201)582-6078

Operations Contact
Randolph Bell (201)582-6368

Boeing Computer Services Company
Post Office Box 24346
Seattle, WA 98124

Installation Delegate
Stephen Niver (206)763-5073
MS 7A-23

Operations Contact
Jim Roetter (206)763-5510
MS 7C-12

Centre de Calcul Vectoriel Pour la Recherche
Ecole Polytechnique
91128 Palaiseau Cedex
France

TELEX: 691596

Installation Delegate — (6) 841 82 00, x.4153, 2506
Tor Bloch

Technical Contact — (6) 908 4061
Maurice Benoit

Operations Contact
Paulette Dreyfus

Centre Informatique de Dorval
(Environment Canada)
2121 Trans-Canada Highway
Dorval, Quebec
Canada H9P1J3

Installation Delegate and Technical Contact
Raymond Benoit — (514)683-9414

Operations Contact
Gary Cross — (514)683-8152

Century Research Center Corporation
3, Nihombashi Honcho 3-chome,Chuo-ku
Tokyo, Japan 103

TELEX 252-4362 CRCNET J

Installation Delegate
Mike(Mitsuru) Maruyama — (03) 665-9901

Technical Contact
Kazuyoshi Fukushima — (03) 665-9901

Chevron Oil Field Research Company
3282 Beach Blvd.
La Habra, CA 90631

TELEX: 176967 via San Francisco

Installation Delegate and Technical Contact
Mostyn Lewis — (213)694-7494

Operations Contact
John Kunselman — (213)694-7029

Commissariat a l´Energie Atomique/CEL-V
BP 27
94190 Villeneuve St. Georges
France

Installation Delegate
Henri Dauty (1)569-96-60, x.6386

Technical Contact
Martine Gigandet (1)569-96-60, x.6184

Operations Contact
Claude Riviere (1)569-96-60, x.6484

Commissariat a L´Energie Atomique/CEV
Centre D´Etudes de Vaujours
Unite de Calcul
BP 7
77181 Courtry
France

Installation Delegate
Bruno Compoint (1) 868-8413

Technical and Operations Contact
Joseph Harrar (1) 868-8688

Compagnie Generale de Geophysique
1, Rue Leon Migaux
BP 56
Massy CEDEX
91301
France

TELEX: CGGEC 692442F

Installation Delegate
Claude Guerin (6) 920.84.08

Compagnie Internationale de Services
en Informatique
BP 24
Gif-sur-Yvette
91190
France

TELEX CISIPSC 691 597 F

Installation Delegate and Technical Contact
Jacqueline Goirand (6) 908.38.41

Operations Contact
Regis Schoonheere (6) 908.38.41

Consorzio Interuniversitario per la Gestione
Del Centro di Calcolo Elettronico dell'Italia
Nord-Orientale
6/3 Magnanelli
Casalecchio di Reno
40033
Bologna, Italy

Installation Delegate
Marco Lanzarini 39-51-576541

Cray Research, Inc.
608 2nd Avenue South
Minneapolis, MN 55402

Administrative Contact
Mary Amiot (612)333-5889

Technical and Operations Contact
Dave Sadler (612)452-6650

DFVLR
WT - DV
D - 8031 Wessling
West Germany

Telephone: (0)8153/281
TELEX: 526401

Installation Delegate
Martin Wacker (0)8153/28-930

Technical Contact
Peter Herchenbach (0)8153/28954

Digital Productions
3416 S. La Cienega Blvd.
Los Angeles, CA 90016

Installation Delegate	
Gary Demos	(213)938-1111
Technical Contact	
Larry Yaeger	(213)938-1111
Operations Contact	
Gordon Garb	(213)938-1111

Electricite de France
1 Avenue du General de Gaulle
A2-004
92140 Clamart
France

TELEX 270 400 F EDFERIM

Installation Delegate	
Yves Souffez	(1) 765 40 18
Technical Contact	
Bertrand Meyer	(1) 765 41 50 or (1) 765 41 05

European Centre for Medium Range
Weather Forecasts
Shinfield Park
Reading RG2 9AX
Berkshire, England

TELEX 847908

Installation Delegate	
Geerd-R. Hoffmann	44-734-876000, x.340
Technical Contact	
Claus Hilberg	44-734-876000, x.323
Operations Contact	
Eric Walton	44-734-876000

Exxon Production Research Company
P. O. Box 2189
Houston, TX 77001

TELEX: 910-881-5579 (Answer back: USEPRTX HOU)

<u>Installation Delegate</u>
T.A. Black (713)965-4203
N-121

<u>Technical Contact</u>
J.E. Chapman (713)965-4689
N-121

<u>Operations Contact</u>
D.N. Turner (713)965-4407
N-180A

Ford Motor Company
Engineering Computer Center
MD-1, Room 208
PO Box 2053
Dearborn, MI 48121

<u>Installation Delegate</u>
Neil St. Charles (313) 845-8493 or
(313) 322-8493

General Motors Research
General Motors Technical Center
12 Mile and Mound Roads
Warren, MI 48090-9055

<u>Installation Delegate and Operations Contact</u>
Ronald Kerry (313) 575-3214

<u>Technical Contact</u>
Dean Hammond (313) 575-3372

Government Communications Headquarters
Priors Road
Cheltenham, Gloucestershire
GL52 5AJ
England

<u>Installation Delegate and Technical Contact</u>
Alan Phillips 0242-521491, x.2301
F/1210, Dept. X34C

Grumman Data Systems
1111 Stewart Avenue
Bethpage, NY 11714

Installation Delegate and Technical Contact
James Poplawski (516)575-2934
MS B34-111

Operations Contact
Steven Hornacek, Jr. (516)575-4273

Institute for Defense Analyses
Thanet Road
Princeton, NJ 08540

Installation Delegate and Operations Contact
Robert Cave (609)924-4600

Technical Contact
Helene Kulsrud (609)924-4600

KFA Julich
Postfach 1913
5170 Julich 1
West Germany

TELEX: 833556 KFA D

Installation Delegate
Friedel Hossfeld 02461-61-6402

Technical and Operations Contact
L. Wollschlaeger 02461-61-6420

Koninklijke/Shell Exploratie & Produktie Laboratorium
Volmerlaan 6
2288 GD Rijswijk (Z.H.)
The Netherlands

TELEX KSEPL NL 31527

Installation Delegate and Technical Contact
A.E. Stormer 070-112741
LS-219

Operations Contact
A.A.H. Kardol 070-112601
LS-208

Konrad Zuse-Zentrum fur Informationstechnik Berlin
Heilbronnerstrasse 10
D 1000 Berlin 31
West Germany

TELEX: 183798

<u>Installation Delegate</u>
Jurgen Gottschewski (030)-3032-233

Lawrence Livermore National Laboratory
PO Box 808
Livermore, CA 94550

TWX 910 386 8339 UCLLL LVMR

<u>Installation Delegate</u>
Richard Zwakenberg (415)422-3750

<u>Technical Contact</u>
Patrick H. Gray (415)422-4049

<u>Operations Contact</u>
George Vranesh (415)422-4008

Lockheed Missile and Space Co.
1111 Lockheed Way
Sunnyvale, CA 94086

TELEX: 346409

<u>Installation Delegate</u>
Jack Sherman (408)742-8993

<u>Technical Contact</u>
Doug Telford (408)742-0948

<u>Operations Contact</u>
Jerry Roninger (408)742-5831

Los Alamos National Laboratory
P. O. Box 1663
Los Alamos, NM 87545

Installation Delegate
Charles Slocomb (505)667-5243
MS B294

Technical Contact
John Dragon (505)667-4812
MS B220

John L. Norton (505)667-5000
MS B257

Operations Contact
Tom Trezona (505)667-4890
MS 252

Max Planck Institute fur Plasmaphysik
8046 Garching
Bei Munchen
West Germany

TELEX 05/215 808

Installation Delegate and Technical Contact
Johann Gassmann 089-3299-340

Merlin Profilers Limited
1 Duke Street
Woking, Surrey
United Kingdom

Installation Delegate
Paul Blundell

Technical Contact
Andy Wright

Mitsubishi Research Institute, Inc.
2-3-6, Otemachi
Chiyoda-ku
Tokyo, Japan 100

TELEX 222-2287 MRI J

Installation Delegate
Nobuhide Hayakawa (03) 270-9211

Technical and Operations Contact
Shuichi Yamagishi (03) 270-9211

Mobil Exploration & Producing Services, Inc.
PO Box 900
Dallas, TX 75221

Installation Delegate
Beverly Jackson (214)658-4409

MOD, RARDE
Fort Halstead
Sevenoaks, Kent, TN14 7BP
England

TELEX: 95267

Installation Delegate and Technical Contact
Bob Youldon 0732-55211, x.3086
Bldg. 511

NASA/Ames Research Center
Moffett Field, CA 94035

Installation Delegate
James "Newt" Perdue (415) 694-5189
MS 233-1

NASA/Lewis Research Center
21000 Brookpark Road
Cleveland, OH 44135

Installation Delegate
William McNally (216) 433-4000, x.6650
MS 142-2

National Center for Atmospheric Research
P. O. Box 3000
Boulder, CO 80307

TELEX 45694

Installation Delegate
Bernie O'Lear (303)497-1268

Technical Contact
Eugene Schumacher (303)497-1264

Operations Contact
Gary Jensen (303)497-1289

National Magnetic Fusion Energy
Computer Center
P. O. Box 5509, L-561
Livermore, CA 94550

TELEX 910-386-8339

Installation Delegate
Hans Bruijnes (415)422-4012

Technical Contact
F. David Storch (415)422-4004

Operations Contact
Marilyn Richards (415)422-4397

National Security Agency
Ft. George G. Meade, MD 20755

Installation Delegate
Bruce Steger (301)688-6275
T335

Technical Contact
C. Thomas Myers (301)688-6275
T335

Operations Contact
Richard W. Ader (301)688-6198
T152

Naval Research Laboratory
4555 Overlook Avenue S.W.
Washington, DC 20375

Installation Delegate
Harvey Brock (202)767-3886

Nippon Telegraph and Telephone Public Corporation
Information Processing and Communication
Services Section, Technology Division
9-11, 3-chome, Midori-cho
Musashino-shi, Tokyo 180
Japan

Installation Delegate
Hiroshi Yamazaki

ONERA
Calculateur Aeronautique
BP 72
Chatillon Sous Bagneux
92322
France

TELEX: ONERA 260 907F

Installation Delegate
Jean-Pierre Peltier (1) 6571160, x.2094

Technical Contact
Daniel Colin (1) 6571160, x.3098

Operations Contact
Jean Erceau (1) 6571160, x.2465

Phillips Petroleum Company
Technical Systems Division
Information Services
Bartlesville, OK 74003

Installation Delegate
Bill Giles (918)661-5488
461 Information Center

Technical and Operations Contact
Arvin Todd (918)661-6426
340 Information Center

Rechenzentrum der Universitat Stuttgart
Pfaffenwaldring 57
7000 Stuttgart 80
West Germany

TELEX: 07255445

Installation Delegate
Walter Wehinger 0711-685-5391

Royal Aircraft Establishment
Bldg. R16
Farnborough, Hants
GU14 6TD
England

TELEX: 858134

Installation Delegate
J.M. Taylor (0252)24461, x.3042

Technical Contact
D. Swan (0252)24461, x.2714

Operations Contact
L. Shepherd (0252)24461, x.2375

SAAB-Scania
Aircraft Division
S-58188 Linkoping
Sweden

TELEX: 50040 SAABLGS

Installation Delegate and Technical Contact
Sven Sandin 4613 182357

Sandia National Laboratories
Albuquerque, NM 87185

Installation Delegate
Jack Tischhauser (505)844-1041
Department 2640

Technical Contact
Frank Mason (505)844-2332
Division 2641

Operations Contact
Kelly Montoya (505)844-1234
Department 2630

Sandia National Laboratories
PO Box 969, East Avenue
Livermore, CA 94550

Installation Delegate and Technical Contact
Dona Crawford (415)422-2192
D8235

Operations Contact
M.H. Pendley (415)422-2965
D8236

Schlumberger-Doll Research
Old Quarry Road
PO Box 307
Ridgefield, CT 06877

TELEX: 643359

Installation Delegate

Bob Snow	(203)431-5527

Technical Contact

Raymond Kocian	(203)431-5522

Operations Contact

Josephine Murray	(203)431-5524

Shell U. K.
Rowlandsway Wythenshawe
Manchester M22 5SB
United Kingdom

TELEX: 668613

Installation Delegate

David Cheater	061-499-4357

SNEA
Rue Jules Ferry
Pau 64000
France

TELEX: Petra 560 804F

Installation Delegate, Technical and Operations Contact

Michel Morin	59-834146

SOHIO
Geophysical Data Center
1 Lincoln Center
5400 LBJ Freeway
Suite 1200-LB 25
Dallas, TX 75240

Installation Delegate

Ron Koehler	(214)386-8000

SVERDRUP Technology, Inc.
Arnold Air Force Station, TN 37389

Installation Delegate and Operations Contact

John L. Roberson ASTF MS900	(615)455-2611, x-5294

Technology Development of California and
NASA Ames Research Center
Ames Research Center, 233-3
Moffett Field, CA 94035

Installation Delegate, Technical and Operations Contact
Glenn Lewis (415)965-6550

United Information Services, Inc.
2525 Washington
Kansas City, MO 64108

Installation Delegate and Operations Contact
Nate Losapio (816)221-9700, x.6535

Technical Contact
John McComb (816)221-9700

University of London
Computer Center
20 Guilford Street
London WC1N 1DZ
England

TELEX: 8953011

Installation Delegate
Richard Field (01)4058400

Technical Contact
John Down

Operations Contact
Lawrie Tweed

University of Minnesota Computer Center
2520 Broadway Drive
Lauderdale, MN 55113

Installation Delegate
John Sell (612)373-7878

Technical Contact
Linda Gray (612)376-5603

Operations Contact
Elizabeth Stadther (612)373-4920

Westinghouse Electric Corporation
Energy Systems Computer Center
P. O. Box 355
Pittsburgh, PA 15146

Installation Delegate

Robert Price (412)374-5826
MNC 206E

Technical Contact

Jerry Kennedy (412)374-4399
Nuclear Center 180

Operations Contact

R.W. Kunko (412)374-4674
Fran Pellegrino

CUG Membership and Update Information

The following two forms can be used for new and update information on CUG Installation Delegates and Technical and Operations Contacts. The first form is for Installation Delegate only. The second form is for Technical and Operations Contacts. "Installation Delegate" replaces what was termed "Administrative Contact." Technical and Operations Contacts remain the same.

Once you have completed either or both forms, please send the information to:

Karen Friedman
NCAR
PO Box 3000
Boulder, CO 80307
USA

Transoceanic sites should send all communications by air mail postage.

MEMBERSHIP APPLICATION / CHANGE

/ / Application / / Change

All requests must bear the signature of the Installation Delegate. The Installation Delegate is the only person qualified to vote for the Member Installation. In the case of a change of Installation Delegate, both the old and new Delegate should sign.

__ __ __ __ __ __ __ __
Installation Code of up to eight alphabetic characters.
(CUG reserves the right to assign or alter codes.)

Installation Delegate (please print)

Installation Name

Street Address

City / Town / State

Postal / Zip Code

Country

Telex

Telephone (Area/City Code - Number - Extension)

(over)

CRAY RESEARCH INCORPORATED COMPUTERS INSTALLED OR ON ORDER:

Mainframe (Type/Configuration) eg. 1-S/2300, X-MP/24	Serial Number	Installation Date eg. 12 NOV 82	Operating System eg. COS, CTSS
________________	________	____________	____________
________________	________	____________	____________
________________	________	____________	____________
________________	________	____________	____________
________________	________	____________	____________
________________	________	____________	____________
________________	________	____________	____________
________________	________	____________	____________

____

Signature, Installation Delegate

Mail to:

Karen Friedman
NCAR
Post Office Box 3000
Boulder, CO 80307
USA

CUG Site Contact Update Form

(Please update our site mailing list entry with the applicable information)

Site: ____

____

____

____

____

TWX: ____

TELEX: ____

Facsimile Transmission: ____

Telephone: ____

FTS (Federal Telecommunications System): ____

TECHNICAL CONTACT

Name: ____

Room: ____

Telephone: ____

FTS (Federal Telecommunications System): ____

OPERATIONS CONTACT

Name: ____

Room: ____

Telephone: ____

FTS (Federal Telecommunications System): ____

Mail to:

Karen Friedman
NCAR
PO Box 3000
Boulder, CO 80307
USA

CRAY USER GROUP

BYLAWS

ARTICLE I - NAME

1.1 Name of Association. The name of this association is CRAY User Group. The abbreviation, CUG, may also be used.

ARTICLE II - PURPOSES

2.1 Principal Purposes. The principal purposes of the Association shall be to provide an open forum to promote the free interchange of information and ideas which are of mutual interest and value to users of CRAY computers, and to provide a formal communications channel between members of the Association and Cray Research, Inc. The Association shall be an international organization for the purposes of administration.

2.2 Achieving Purposes. To achieve these purposes the Association shall:

2.2.1 Organize and conduct meetings, discussion groups, forums, panels, lectures and other similar programs concerned with research and development and the exchange of technical data.

2.2.2 Publish, as appropriate, the results of its research and make such publications available to the public on a non-committal and a non-discriminatory basis.

2.2.3 Establish and continually improve standards for communicating computer science research results and programming information to interested members of the public.

2.3 Conducting Business. To achieve these purposes, the business of the Association shall be conducted as appropriate at Meetings of the Members (as specified in Article V of these Bylaws), by the Board of Directors (as specified in Article VI of these Bylaws) and when permitted, by mail ballot (as specified in Article V Section 5.9 of these Bylaws).

2.4 Not for Profit Association. The Association shall operate as a not for profit organization.

ARTICLE III - DEFINITIONS

3.1 The Association. The CRAY User Group, an unincorporated group.

3.2 CRAY Computer. Any Cray Research, Inc. computer system.

3.3 Installation. Any organization on whose premises one or more CRAY Computers shall be in operation, or for which a CRAY Computer shall have been purchased or leased or is on order.

3.4 Member or Member Installation. An installation that has been accepted for membership pursuant to Article IV.

3.5 Installation Delegate. The employee of a Member Installation designated to serve as that installation's official spokesperson at any function of the Association and to cast that installation's vote on all matters on which the installation may have the right to vote.

3.6 Installation Participant. Any bona fide user of a Member Installation's services. These eligible participants include remote users (within the same company) of a Member Installation contingent upon validation by the Installation Delegate and upon approval of the Membership Committee. An Installation Participant is authorized to participate in Association functions on behalf of that Member Installation.

3.7 Installation Representative. An Installation Delegate or Installation Participant.

3.8 CUG Visitor. Any individual who is not an Installation Representative and who is invited to attend a specific function of the Association by an Installation Representative, Officer or Board of Directors member and whose attendance at such function shall have the prior approval of the Board of Directors.

3.9 Anniversary Meeting. The general meeting of the Members that is held closest to October 15 each year.

ARTICLE IV - Membership

4.1 Membership. There shall be one class of membership: Installation Membership. Any Installation shall be eligible to become a Member.

4.2 Multiple Memberships. More than one Membership may exist within the same organization with the approval of the Board of Directors. In general, however, more than one mainframe per computation headquarters does not automatically qualify an organization for multiple Membership.

4.3 Additional Classes of Members. The Members, by a majority vote, may create one or more additional classes of membership and affiliation and may prescribe the designations, voting rights, if any, powers and privileges for each such class.

4.4 Application for Membership. An Installation desiring to become a Member shall submit a written membership application to the CUG Secretary. The completed application shall provide such information as shall from time to time be prescribed by the Board of Directors.

4.5 Qualification as Member. The Membership Committee shall review each application for membership. If satisfied with the bona fides thereof shall notify the prospective Installation of the said admission. The Membership Committee shall consider all applications received more than four weeks prior to a General Meeting no later than the next General Meeting. The application may be rejected if, in the judgement of the Membership

Committee, the admission would be detrimental to the objectives of the Association.

4.6 Obligations of all Members. Each Member Installation shall instruct and cause its Installation Representatives to abide by the Bylaws and the rules and regulations of the Association as they may from time to time appear.

4.7 Grounds for Loss of Membership. An Installation shall lose its membership within thirty days after receiving written notice from the Secretary to the Installation Delegate that one or more of the following shall have occurred (such notice to state the basis for revocation of membership):

4.7.1 The Member shall cease to be an Installation.

4.7.2 The Board shall have determined that the Installation has failed to abide by the Bylaws or rules and regulations of the Association.

4.8 Appeal. Within ninety (90) days of the receipt of notice sent pursuant to Section 4.7, the recipient Member may appeal in writing (addressed to the President) to the Board of Directors to have the notice set aside. The sole basis upon which such appeal may be made shall be:

4.8.1 Proof satisfactory to the Board that the ground(s) set forth in the notice is (are) not valid; or

4.8.2 A statement of extenuating circumstances.

The Board of Directors shall act upon an appeal within ninety (90) days of its receipt and shall notify the appellant in writing of its decision within thirty (30) days thereafter.

4.9 Withdrawal. A Member may voluntarily withdraw from the Association at any time by giving written notification of the desire to so withdraw, said notification to be signed by the Installation Delegate and to be directed to the Secretary. Such withdrawal shall become effective upon receipt thereof by the Secretary.

4.10 Rights of Members. The right to vote for the election of members of the Board of Directors and officers and to vote on all issues is conferred solely upon the Members. Only an Installation Delegate or Participant shall be eligible to be a member of the Board of Directors or to hold elective or appointive office in the Association.

4.11 Fees. An annual fee payable by Member Installations shall be determined by the Board of Directors. Registration fees shall be payable to cover the cost of meetings, such fees to have been approved by the Board of Directors.

ARTICLE V - MEETING OF MEMBERS

5.1 General Meetings. The General Meeting of Members shall be held twice in each calendar year at approximately six (6) month intervals at such time and place as shall be determined by the Board of Directors and designated in the notice or waiver of notice of the meeting. Dates and locations for General and Special Meetings shall be decided by the Board of Directors and, whenever possible, announced a year in advance. At the Anniversary Meetings the Members entitled to vote shall elect officers and directors as prescribed by Article VIII. At any general meeting the Members may transact such business as may properly come before the meeting.

5.2 Special Meetings. Special Meetings of the Members may be called at any time by the Board of Directors. Upon receipt of a petition (stating the purpose of the proposed meeting) signed by at least one-third of the Installation Delegates of Members entitled to participate, the President shall call a Special Meeting.

5.3 Notice. Written Notice of General and Special Meetings of the Members of the Association shall be given by the President or the Secretary, and sent to each Member Installation entitled to participate thereat, by mail addressed to the Installation Delegate at the address appearing on the records of the Association, not less than thirty (30) days before the time designated for such meeting.

5.4 Waiver of Notice. Any meeting and any action otherwise properly taken thereat shall be valid if notice of the time, place and purposes of such meeting shall be waived in writing, before, at or after such meeting by all Members to whom notices should have been received but were not as provided in these Bylaws.

5.5 Quorum. The presence in person of not less than one-third of the Installation Delegates or their designated Installation Participants entitled to vote shall be necessary and sufficient to constitute a quorum for the transaction of business at any meeting of the Members. When a quorum is once present, it is not broken by the subsequent withdrawal of any Members. A majority of the Members at any meeting, including an adjourned meeting, whether or not a quorum is present, may adjourn such meeting to another time and place.

5.6 Organization. At every General Meeting of Members, the President, and in the absence of the President, the Vice President, shall act as the Chair of the meeting. The Secretary shall act as Secretary of the meeting. In case none of the officers above designated to act as the Chair or the Secretary to the meeting respectively, shall be present, the Chair or Secretary of the meeting, as the case may be, shall be chosen by a majority of the votes cast at such meeting by the Members entitled to vote at the meeting.

5.7 Voting. The Members shall have the exclusive right to vote on all matters pertaining to the general affairs of the Association on which a vote of the Members is required or deemed by the Board of Directors to be desirable. Each Member in good standing and entitled to vote shall be entitled to one vote. Votes for each Installation shall be cast by the Installation Delegate or, if absent, by a designated Participant for such Installation at all meetings of Members, or by mail ballot cast with the Secretary on specific matters, including the removal of officers and directors.

5.8 Action by a Majority Vote. All questions submitted to the Members, except as otherwise provided by law or by the Bylaws, shall be decided by a majority of votes cast by Members entitled to vote who shall have voted thereon.

5.9 Procedures for voting by Mail. Voting by mail shall be permitted for any item of business except election of officers and members of the Board of Directors. All proposals to be acted upon by mail shall be proposed by at least four (4) Members or proposed by action of the Board of Directors. Proposals shall be addressed to the Secretary. The Secretary shall thereupon cause proposals to be sent to all Members of the Association. All Members may, within sixty (60) days, submit comments with respect to said proposals to the Secretary, who shall group or categorize such comments as shall be deemed appropriate, and cause representative commentary to be sent to all Members of the Association. The proposal shall be put to a vote by a mail ballot which shall be enclosed with the said commentary (if any). All mail ballots shall be cast and signed by the then acting Installation Delegate of the Member eligible to vote on the proposal and submitted to the Secretary within thirty (30) days of mailing. If at least two thirds of the Members eligible to vote shall cast their vote, a quorum shall have been achieved. If a quorum is achieved and a simple majority of those eligible Members voting shall vote in favor of the proposal, it shall be approved.

5.10 Restriction. No Installation Delegate, Installation Participant, CUG Visitor or guest of the Association shall engage in employment recruiting or interviewing at any Meeting of Members. Meetings shall not be used for marketing or other commercial purposes.

ARTICLE VI - DIRECTORS

NOTE: ARTICLE VI is subject to the provision of ARTICLE XVI (Consent in Lieu of Meeting).

6.1 Powers. The Board of Directors shall exercise all powers of the Association, except as otherwise expressly provided by law or by these Bylaws. The members of the Board of Directors shall act only as a board and the individual members shall have no power as such. Among such powers are:

6.1.1 Develop and execute Association policy.

6.1.2 Interpret and implement decisions of the Association Members and the Board of Directors.

6.1.3 Approve the Association budget and designate an independent certified public accounting firm to audit the Association financial records.

6.1.4 Establish all fees for the Association.

6.1.5 Approve the use of the Association name, in whole or in part, by individuals or other organizations.

6.2 Number. The number of directors of the Association shall not be less than seven. However, the number of directors may be increased by the Members at any General Meeting. The Association's President, Vice President, Secretary and Treasurer shall automatically become directors when elected to their office. The retiring President shall automatically become a member of the Board of Directors upon the election of his or her successor as President and shall remain a Director-at-Large for the length of the incumbency of the successor. In addition to the aforementioned officers, the Board of Directors shall have two other Directors-at-Large. Any eligible person may be re-elected as a director one or more times.

6.3 Term of Office. The tenure of the President and Vice President will be for one year. The tenure of the Secretary and Treasurer will be for two years and will expire on alternating years. The tenure of the first Secretary elected will be for one year in order to initiate the cycle of alternating terms. The remaining two Directors-at-Large will serve for two years with tenure expiring on alternating years. The tenure of one of the first Directors-at-Large elected will be for one year in order to initiate the cycle of alternating terms. Each director shall continue in office for the terms noted above until the General Meeting at which the successor is elected and qualified. The term of office of any director may be terminated at any time, with or without cause, by an affirmative vote of two thirds of the votes cast by Members entitled to vote and who shall have voted thereon.

6.4 Qualification. To qualify as a director of the Association each individual must be a bona fide employee of a Member and remain so for the entire term of office.

6.5 First Meeting. Each duly constituted Board of Directors may hold its first meeting for the purpose of organization and the transaction of other business, if a quorum be present, without notice of such meeting, on the same day(s) and at the same place the general meeting of Members having elected said Board of Directors is held, and as soon as practicable after such General Meeting. Such first meeting may be held at any other time and place as specified in a notice as hereinafter provided in Section 6.7 of this Article for Special Meetings of the Board of Directors, or in a waiver of notice thereof.

6.6 Regular and Special Meetings. Regular meetings of the Board of Directors may be held at such places and times as may be fixed from time to time by resolution of the Board of Directors to conduct such business that may properly come before it; and unless otherwise required by resolution of the Board of Directors, notice of any such meeting need not be given. The President or the Secretary may call, and upon written request signed by any three (3) directors, the Secretary shall call, Special Meetings of the Board of Directors. Meetings of the Board of Directors shall be held at the place designated in the notice or waiver of notice of such meeting.

6.7 Notice of Special Meetings. Notice of Special Meetings of the Board of Directors shall be in writing, signed by the President or the Secretary, and shall be sent to each director by mail or telex addressed to arrive at his last known address at least twenty (20) days before the time designated for such meeting.

6.8 Waiver of Notice. Any meeting of directors and any action otherwise properly taken thereat shall be valid if notice of the time, place and purposes of such meeting shall be waived in writing (including telex, cable or wireless) before, at, or after such meeting by all directors to whom timely notices were not sent as provided in these Bylaws.

6.9 Quorum. A majority of directors in office, personally present, shall be necessary and sufficient to constitute a quorum for the transaction of business at any meeting of the Board of Directors, but a smaller number may adjourn any such meeting to a later date. At least one day's notice of such adjourned meeting shall be given in the manner provided in Section 6.7 of this Article to each director who was not present at such meeting.

6.10 Action by Majority Vote. Except as otherwise expressly required by law or by these Bylaws, the act of a majority of the directors present at a meeting at which a quorum is present shall be the act of the Board of Directors.

6.11 Filling Vacancies. Any vacancy in the Board of Directors, whether caused by death, resignation, disqualification, removal, increase in the number of Directors or otherwise, may be filled for the unexpired term by a majority vote of the remaining directors, or by the Members at a special meeting called for such purposes. The individual selected must meet the same qualifications as a nominee for a directorship.

6.12 Reports to the Membership. The actions of the directors at any meeting of the Board of Directors shall be reported to the membership within sixty (60) days of that meeting.

6.13 Submission of Matter to Mail Vote of the Members. The Board of Directors may submit any matter to a mail vote of the Members, when required or deemed advisable or desirable by the Board of Directors. Any such mail vote shall be pursuant to Sections 5.7, 5.8 and 5.9.

ARTICLE VII - OFFICERS

7.1 Officers. The officers of the Association shall be the President, Vice President, Secretary and Treasurer, each to have such duties or functions as are provided in these Bylaws or as the Board of Directors may from time to time determine. One person may not hold any two or more of the foregoing offices.

7.2 Nominations and Elections. Nominations and elections shall be in accordance with Article VIII.

7.3 Terms and Qualifications. The term of office of the President and Vice President shall be one year. The term of office for the Secretary and Treasurer shall be two years. The terms commence with the election of each officer and end when the successor is elected and qualifies and may be terminated at any time with or without cause, by an affirmative vote of two thirds of the votes cast by Members entitled to vote and who shall have voted thereon.

7.4 Resignations. Any officer may resign at any time, orally or in writing, by notifying the Board of Directors or the President or the Secretary of the Association. Such resignation shall take effect at the the time therein specified, and, unless otherwise specified, the acceptance of such resignation shall not be necessary to make it effective.

7.5 Vacancies. A vacancy in any office caused by death, resignation, removal, disqualification or other cause may be filled in accordance with Section 6.11 for the unexpired portion of the term of the Board of Directors at any regular or special meeting.

7.6 The President. The President shall be the chief executive officer of the Association and shall have general supervision over the affairs of the Association, subject, however, to the control of the Board of Directors. The President shall, if present, preside at all General Meetings, and at all meetings of the Board of Directors. In general, the President shall perform all the duties incident to the office of the chief executive officer of an association and such other duties as are provided for by these Bylaws and as from time to time may be assigned to the President by the Board of Directors.

7.7 The Vice President. At the request of the President, or in the President's absence the Vice President shall perform all the duties of the President and in so acting shall have all the powers of and be subject to all the restrictions upon the President. The Vice President shall perform such other duties as may from time to time be assigned to the Vice President by the President or by the Board of Directors.

7.8 The Secretary. The Secretary shall act as Secretary of all meetings of the Board of Directors, and of the Members of the Association, and shall keep the minutes thereof in the proper book or books to be provided for that purpose. The Secretary shall cause all notices required to be given by the Association to be duly given and served; shall have charge of the other books, records and papers of the Association; shall cause the report, statements and other documents required by law to be properly kept and filed; shall maintain a current list of Members and be responsible for membership applications; shall act as editor for correspondence received for publication and distribute this information at intervals not greater than six (6) months; and shall, in general, perform all the duties incident to the office of Secretary and such other duties as may from time to time be assigned to the Secretary by the Board of Directors or by the President.

7.9 The Treasurer. The Treasurer shall collect, and keep account of all moneys received and expended for the use of the Association. The Treasurer shall deposit sums received by the Association in the name of the Association in such depositories as shall be approved by the Board of Directors; prepare appropriate financial reports for review by the Board of Directors; and be a member of the Finance Committee.

7.10 Standing Committees. The following standing committees of the Board of Directors are permanently established.

7.10.1 Membership, the Chairperson of which committee shall be the Secretary, and shall consist of not less than two (2) officers or directors appointed by the President.

7.10.2 Finance, the Chairperson of which committee shall be the President, and shall consist of the Treasurer and at least one other Director appointed by the President. The Treasurer serves as Secretary of the committee. The duties of the committee shall consist of preparation of the Association's budget for approval by the Board of Directors, and review and approval of requests for non-budgeted expenditures, supervision of accounting methods and procedures, and the preparation of and delivery to members of an annual report of the Association's financial status.

ARTICLE VIII - ELECTIONS

8.1 Nominations. The immediate Past President, if available, shall be the Chair of the Nominating Committee and select its remaining members. In the event that the immediate Past President is not available to serve in that capacity, the President shall select the Chair. The Nominating Committee shall consist of five (5) members. Any Installation Representative is eligible to serve. No member of the Nominating Committee may be a candidate for a Board of Directors position. The members of the Nominating Committee shall select candidates for each of the positions of officer or director to be filled at the next scheduled election. The Nominating Committee shall determine how many candidates it will nominate for each position. Current Association Officers are not eligible to serve on the Nominating Committee. The Nominating Committee shall cease to exist upon filing its report to the Members.

8.2 Report. On the first day of the Anniversay Meeting, the Nominating Committee shall report the names of candidates for each office scheduled to be filled by the election.

8.3 Nominating by petition. Any eligible individual may be nominated for any office or as a director by a Petition signed on his or her behalf by not less than three (3) Installation Delegates. Nominating petitions and assurances from the candidates (as defined in section 8.4 of this article) must be submitted to the Secretary no later than 6:00 p.m. (local time) on the day preceding the elections for office.

8.4 Qualification and Assurance of Candidates. At the time of nomination, each candidate must be a bona fide employee of a Member. The Chairperson of the Nominating Committee shall require in writing from each candidate for office for the Board of Directors (whether such candidate has been named by the Nominating committee or by Petition) a written statement by which the candidate offers assurances that, if elected, he or she will diligently fulfill the duties of the office or the position on the Board of Directors for which nominated during the term thereof. A candidate by petition must submit an assurance statement with the completed petition on his or her behalf.

8.5 Withdrawal from Candidacy. Any duly nominated candidate may withdraw his or her name from nomination by submitting a written request to such effect to the Secretary at least one hour prior to the first ballot for such position at the general meeting of Members.

8.6 Election Procedure. At the Anniversary Meeting the Chairperson of the Nominating Committee shall announce the names of those persons who have been nominated for each office and for positions as directors, who have given the requisite written assurances of performance in the event of election, and who have not withdrawn. If a nominee for an office shall be unopposed the President shall declare such individual elected. As to those persons who are opposed for office and for candidates for the Board of Directors, an election shall be held by written ballot. The Secretary shall cause ballots to be distributed to the Installation Delegate or the alternate Representative of each of the Members represented at the general meeting at which elections are held.

8.7 Vote Required for Election to Office. When more than one (1) candidate is nominated to each office (except for the Directors-at-Large), the winning candidate must receive a majority of the votes cast (for that office) in order to be elected to that office. In the event that no candidate receives a majority, the two (2) candidates receiving the greatest plurality will remain in nomination for that office, and a runoff election (following the applicable rules of Election Procedures, Section 8.6 of this(Article) will be held for all offices in which majorities were not obtained. Ties will be broken by a runoff election

When vacancies for the Directors-at-Large are to be filled by election and there are more candidates than vacancies, the following procedure will be used. The candidates will be ranked according to the number of votes received, most to least. The first candidate on the list will fill one directorship-at-large, the next candidate, the next directorship-at-large, etc., until all directorships-at-large are filled unless the last candidate to fill a directorship-at-large is tied with the next candidate on the ordered list; in which case, these two candidates will be entered in a runoff election, and the candidate receiving a majority of the votes cast will fill that directorship-at-large.

8.8 Voting formula. Each Member has the right to cast one vote for each position to be filled. Absentee ballots for elections are not permitted.

8.9 Extension of Term of Office. In the event that a successor for an officer or director-at-large whose term has otherwise expired is not elected at an Anniversary Meeting, the then present holder of the office or directorship shall continue in office until a successor is qualified and takes office.

ARTICLE IX - COMMITTEES

NOTE: ARTICLE IX is subject to the provisions of ARTICLE XVI (Consent in Lieu of Meeting).

9.1 Committees. The Board of Directors may from time to time create or terminate standing and ad hoc committees and may determine the names of such committees and the qualification of the members of such committees; and, to the extent permitted by law, may delegate the powers and duties of the Board of Directors to such other committees, and, to such extent, may otherwise determine such powers and duties. The Board of Directors may elect the members of such committees or may authorize the President and/or any other officer or officers to select the members of any such committee.

ARTICLE X - LIMITATIONS ON CIRCULATION OF INFORMATION

10.1 Scholarly and Scientific Endeavor. Persons affiliated with members may refer to, and excerpt material from communications among the Association's members, publishing scholarly articles or giving educational courses or conducting scientific experiments. This section shall be construed liberally for the purposes of advancing scientific research, education and scholarship in the public interest, but shall be construed restrictively to avoid commercialism, journalism, editorializing and notoriety.

10.2 Persons in Computer Field. The President shall, at the request of any individual engaged in the computer field (other than in the news or communication media related thereto) and having a legitimate interest in information disseminated by the Association, make available to any such individual at cost (within reasonable bounds as to quantity of material furnished), matter formally disseminated to Members.

10.3 Information Referrals. It is the policy of the Association to disseminate information and data freely to those having a legitimate interest therein pursuant to Section 10.1 and 10.2 hereof. In furtherance of this policy, all Members shall refer all inquiries or requests with respect to publications and data of the Association to the Association, and such inquiries and requests will be acted upon by the Association in accordance with Sections 10.1 and 10.2.

ARTICLE XI - CONTRACTS, CHECKS, DRAFTS, BANK ACCOUNTS, VOTING OF SECURITIES, ETC.

11.1 Execution of Contracts. The Board of Directors, except as otherwise provided in these Bylaws, may authorize any officer or officers, agent or agents, in the name and on behalf of the Association to enter into any contract or execute and satisfy any instrument, and any such authority may be general or confined to specific instances.

11.2 Checks, Drafts, etc. All checks, drafts and other orders for payment of money out of the funds of the Association shall be signed on behalf of the Association by two officers one of whom to be the Treasurer or, if unavailable, the President.

11.3 Deposits. The funds of the Association not otherwise employed shall be deposited from time to time to the order of the Association in such banks, trust companies or other depositories as the Board of Directors may select.

ARTICLE XII - BOOKS AND RECORDS

12.1 Books and Records. There shall be kept at the principal place of employment of the Treasurer correct books of account of all the business and transactions of the Association.

12.2 Other Books and Records. All books and records not covered by Section 12.1 shall be kept in the custody of the Secretary.

ARTICLE XIII - SEAL

13.1 Seal. The Board of Directors shall provide a seal which shall be in the form of the plan (top) view of three (3) CRAY-1A's, each oriented so as to form the word: "CUG". The three (3) words: "CRAY USER GROUP" will be superimposed on and within the lower half of each symbol, one word per symbol.

ARTICLE XIV - PARLIAMENTARY AUTHORITY

14.1 Parliamentary Authority. "Robert's Rules of Order, Revised" shall prevail, except that where they conflict with these Bylaws, the Bylaws shall govern.

ARTICLE XV - AMENDMENTS OF BYLAWS

15.1 Proposals. Proposed amendments may be directed to or initiated by the Board of Directors. Written notice of the proposed change, the originator, and the recommendation of the Board of Directors, if applicable, will be sent to each Member Installation by mail addressed to the Installation Delegate at the address appearing on the records of the Association, not less than thirty (30) days before the time designated for a General Meeting of the Members.

15.2 **Voting Procedure.** These Bylaws, or any one or more of the provisions thereof, may, at any duly constituted General Meeting of the Members, be amended by changing, altering, suspending, supplementing or repealing the same, by an affirmative vote of two thirds of the votes cast by Members entitled to vote and who shall have voted thereon at such meeting, but only in accordance with a proposed amendment duly published and mailed according to the provisions of Section 15.1 of this Article.

ARTICLE XVI - CONSENT IN LIEU OF MEETING

16.1 Any other provisions of these Bylaws to the contrary notwithstanding, any action required or permitted to be taken at any meeting of the Board of Directors or of any committee may be taken without a meeting, if prior to such action a written consent thereto is signed by all members of the Board or of such committee, as the case may be, and such written consent is filed with the minutes or proceedings of the Board of Directors.

ARTICLE XVII - DISSOLUTION

17.1 **Procedure.** Dissolution of the Association is proposed and approved in the same manner as an Amendment of the Bylaws.

17.2 **Liabilities.** All liabilities and obligations of the Association shall be paid, satisfied and discharged, or adequate provisions shall be made therefore.

17.3 **Remaining Assets.** Remaining assets shall be transferred or conveyed to one or more domestic or foreign corporations, trusts, societies or other organizations engaged in charitable, religious, eleemosynary, benevolent, education or similar activities, but in no case, shall any part of the assets be distributed to members of the Association.

27 January 1984

CRAY USER GROUP

POLICY STATEMENTS

1. Permanent Special Interest Committees

1.1 Permanent Special Interest Committees will be established as appropriate.

1.2 Each Permanent Special Interest Committee will appoint a Chair, approved by the Board of Directors, and at least one other officer who will co-ordinate items for consideration by that Committee.

1.3 Permanent Special Interest Committees will meet during General Meetings.

1.4 Permanent Special Interest Committees will play a key role in bringing forward matters of importance to the Advisory Council.

1.5 Each Permanent Special Interest Committee will play a key role in planning of General Meetings.

2. General Meetings

General Meetings will last a minimum of three working days, preceded by one or two days (as necessary) of Committee and Board of Directors meetings. At least two sessions shall be set aside at each General Meeting for transacting the business of the Association. One of these sessions shall be scheduled for the first half day of the General Meeting and the second for the last half day. Scheduled sessions may include reports by Cray Research, Inc., papers by Installation Representatives on applications experience with hardware and software, and workshops for Special Interest Committees.

3. Future General Meetings

The Vice President has certain delegated responsibilities connected with future general meetings.

3.1 The Vice President shall locate future meeting sites and dates through negotiation with potential hosts for such meetings either in reaction to invitations or through solicitations of invitations. The Vice President shall negotiate appropriate accommodations, dates and arrangements with various potential hosts and shall maintain a contingency concern for the next 2 or 3 meetings. In recommending sites, careful consideration should be given to travel arrangements, hotel accommodation, and physical meeting facilities. The general principle should be to have a balance between meetings held in the United States and those held outside the United States.

3.2 The Vice President shall report to General Meetings regarding recommendations for future meetings.

3.3 The Vice President shall chair the Program Committee.

4. CUG Resolutions

4.1 Any Installation Representative who desires to place a resolution before the membership at a General Meeting will submit a typewritten copy of the proposed resolution to the Secretary at least 24 hours prior to the scheduled plenary session for consideration by the Board.

4.2 The Board will submit the resolution to the membership for discussion at the plenary session together with the Board's recommendations.

4.3 The Board may defer action on the resolution pending further investigation by a CUG Committee or the Board. When such action occurs, the submittor is expected to participate in any committee activities relative to that resolution.

4.4 The Board may submit the resolution to the entire membership for consideration via mail ballot in which case one copy of the resolution and one ballot paper will be mailed to each Installation Delegate pursuant to sections 5.7 and 5.9 of the Bylaws.

4.5 Resolutions considered by the Board will be published in the minutes of the Board Meeting. Action taken with regard to resolutions referred to committees will be published in the reports prepared by those committees.

5. Program Committee

A Program Committee chaired by the Vice President and comprising the host of the previous General Meeting, the host of the next (or current) General Meeting, the host of the following General Meeting, the past Program Chair, the Chairs of the Special Interest Committees, co-opted members and the Proceedings Editor. A non-voting Cray Representative will be invited to attend.

The Program Committee is responsible for the program of General Meetings and for ensuring that all the necessary arrangements are in place.

6. Advisory Council

The Advisory Council serves as the operating committee of the Association. It is composed of the Board of Directors plus the Chairs of the Special Interest Committees, and such other members which the Board of Directors see fit to co-opt.

Its purposes are:

(a) to implement the policies instituted by the Board of Directors;

(b) to co-ordinate the technical and administrative aspects of the

Association;

(c) to formally represent the Association's views to Cray Research, Inc. in the manner defined in section 9.3 of these Policy Statements.

The Advisory Council will meet following the Board meeting on the day prior to a General Meeting. In addition it may meet again after the closing session of a General Meeting if there is business arising from the General Meeting.

7. Board of Directors

It is the policy of the Association that every effort will be made to ensure that the Board of Directors is truly representative. In particular every effort will be made to ensure that not more than two (2) members of the board are from the same organization. In addition every effort will be made to ensure that at least two (2) board members are from organizations within the United States and at least two (2) board members are from organizations outside the United States.

8. Board of Directors Meetings

The Board shall meet at least twice each year for the purpose of receiving reports from committees, approving the agendum for the next General Meeting and transacting any other Board business which may arise. In general these meetings will take place on the day prior to a General Meeting. Additional meetings of the Board of Directors will be arranged as required.

9. Communication with Cray Research, Inc.

9.1 The Program Committee or the Board of Directors may invite employees of Cray Research, Inc. to give presentations on specific topics at General Meetings.

9.2 Cray Research, Inc. will be invited to identify a named employee to represent Cray Research, Inc. at Advisory Council meetings. The President (or his deputy) has the right to ask the Cray Research, Inc. representative to withdraw from an Advisory Council Meeting if this is agreed by two other Directors.

9.3 The Advisory Council shall identify those issues of concern to the membership for which it is desirable to have discussions with Cray Research, Inc. at a senior level. The Advisory Council shall nominate one, or at most two, knowledgeable individuals who are able to present the views of the membership to Cray Research, Inc. The President will then make arrangements with Cray Research, Inc. for the Association President plus the nominated representative(s) to have discussions at a suitable level with Cray Research, Inc.

10. CUG Requests

10.1 All formal committee requests to Cray Research, Inc. will be submitted to the Board of Directors. The Board of Directors will study the requests, approve or reject them, or return them to the appropriate Committee for clarification or elaboration. If the Board feels that a given request is important or controversial enough, it will submit it to the entire membership for a vote.

10.2 Only requests approved in this fashion will constitute official Association requests. The Board of Directors will, after approval, assign a unique number to each such request and will forward it to Cray Research, Inc. The Association will expect Cray Research, Inc. to reply promptly to each such request. The Secretary will keep an official log of all such official requests and will report periodically on the status of each. The status will be considered as open until such a time as the Board and/or the Committee decides that the item has been resolved.

10.3 The above procedures are not meant to impede normal and necessary interaction between committees and Cray Research, Inc. and do not apply to normal requests for information.

11. Proprietary Information

11.1 Committee activities frequently require that proprietary information be made available to committee members. The protection and use of this information must be carefully controlled for this concession to be accepted.

11.2 All information to be protected under this policy statement must be specified as proprietary information. In the case of written material, these requirements can be met by physically attaching a covering letter announcing to the reader that the attached material is proprietary information and must be treated as per this policy statement.

11.3 Proprietary information must not be discussed outside of normal CUG activities and related support activities. Unless explicitly restricted to the members of a particular committee, this information can be discussed at all committee meetings, Board of Directors meetings, and ad hoc meetings of CUG committee members. It may be discussed only with personnel within the CUG organization and authorised personnel of the originator of the material. These restrictions are terminated when the originator presents the information at an open session at a General Meeting.

11.4 This policy statement pertains only to information obtained through CUG activities; it does not pertain to information acquired directly from the manufacturer.

11.5 Violations of the guidelines set

forth in this policy statement may be grounds for loss of membership.

12. Manufacturer's Representation to the Board

12.1 The Board requests that Cray Research, Inc. appoint an official representative to attend such functions of the Board as may be requested by the President.

13. Disclosure of Information

13.1 In accordance with the purposes of the Association, any information not explicitly specified as Proprietary Information (see 11) may be freely discussed among CUG members.

14. Travel Expense Funding

It is the policy of the Association to reimburse Officers, Board Members and other members in any approved enterprise for travel expenses incurred in conjunction with activities on behalf of the Association where these expenses cannot be borne by the individual's parent organization. Travel expense funding is regarded as the exception and not the rule and the cheapest form of travel should be used where possible. Expenses incurred by persons acting in their own name(s) or in the names of their parent organization shall not be reimbursed.

15. Multiple Installation Membership

15.1 Installations having three (3) to five (5) CRAY mainframes shall be eligible for two (2) Installation Memberships and Installations having six (6) or more CRAY mainframes shall be eligible for three (3) Installation Memberships. The annual membership fee is payable for each Installation Membership.

16. CUG Visitors

16.1 The Board delegates authority to the President or Vice President and Secretary, acting concurrently, to review and approve the attendance of CUG Visitors at CUG activities subject to the satisfaction of the Board.

16.2 CUG Visitors are those who, at the request of an Installation Representative, Officer or Board Member, are going to make a contribution to a scheduled session at a CUG activity. Visitors will, therefore, be limited as to which session(s) they may attend and/or participate. It is the responsibility of the inviting Installation Representative, Officer or Board Member to inform the Visitor of his/her rights and responsibilities and to insure his/her proper conduct. In addition, the content of formal presentations by CUG Visitors will require advance approval by the appropriate Committee Chair.

17. Mechanical Recording Devices

17.1 Mechanical recording devices may not be employed at any of the open or closed meetings of the Association.

18. Definition of an Installation

18.1 It is the policy of the Association that there should not be more than one Installation Membership per CRAY mainframe computer. Where more than one organization could qualify for membership on the basis of a single CRAY computer the conflict will be resolved by the Board of Directors.

7 January 1985