

CSR TASK REPORT


by David D. Clark

The attached report covers progress of the Computer Systems Research Division of the Laboratory for Computer Science in the period 1 January to 31 March 1977.

Several events of note occurred this quarter:

- The local net made substantial progress in both the hardware and protocol area.
- Multics is now a tool bearing host in the NSW.
- Four papers were written and submitted to conferences, including part of the final report on the kernel design project.

This note is an informal working paper of the M.I.T. Laboratory for Computer Science, Computer Systems Research Division. It should not be reproduced without the author's permission and it should not be cited in other publications.



Kernel Design Tasks

I. Multi-Level Traffic Control

No progress was made on the implementation of multi-level traffic control, due to lack of time by the relevant people. If practical, we will attempt to continue this project during the spring.

II. Separating the Functions of Page Control and Segment Control

Drew Mason continued his study of the separation of the functions of page control and segment control within the Multics virtual memory manager. Due to unavoidable delays, production of his Master's Thesis is taking longer than expected. It will be submitted within the next few weeks.

III. Study of System Initialization

This project was completed during this quarter, except for the publication of the Technical Report.

IV. Provision of "Breakproof" Environment for User Programming

This project was completed during this quarter. It is documented as LCS Technical Report 175.

V. A Case Study of Intermodule Dependencies in a Virtual Memory Subsystem

This project was completed during this quarter. It is documented as LCS Technical Report 174.

VI. Input/Output in a Security Kernel

His thesis proposal signed, Gene Ciccarelli has been pursuing the major part of his thesis research. Currently he is designing the basic communication-multiplexing mechanism within the kernel and the user-domain programs to use it. This is evolving into an implementation to prove the main ideas in the thesis. Gene has also been developing precise communications security requirements for the kernel, network, and buffer-hosts on the network. He is expecting to complete the design and implementation during the next quarter.

VII. Command Usage Monitoring

During this quarter, Humberto Rodriguez has implemented a command-level monitor to measure the characteristics of a typical Multics user. The monitor creates a file for each terminal session and accumulates frequency counts, resource usage and think times for most of the commands in Multics. Also, a think time histogram is being maintained to provide a clue as to what the distribution is. Currently, the data is being collected and the results will be reported in his Bachelor's thesis in May.

VIII. Responding to Errors in a Computer System

Harry Forsdick completed a draft of a paper summarizing his work in the area of detecting and responding to errors. This paper represents a final report on the work he performed while a member of our group.

IX. Preparation of Final Report

During this quarter, work started on preparation of the final report for the kernel design project. Jerry Saltzer, Mike Schroeder, and Dave Clark have prepared a paper (submitted to the next Symposium on Operating Systems) that summarizes our conclusions and insights about this project. This paper is available as RFC 140. Doug Wells is currently studying the current system to attempt to make some estimate of the total impact on the size of the kernel of the system were all of our projects carried through fully. His results will be combined with the paper and a final task summary to be written next quarter to form the final report.

ARPANET Related Tasks

I. Improvements to Network Mail Facilities

During this quarter, we reached an agreement with MIT Information Processing Services to furnish us with Multics time for the continued development of read_mail and its friends, in return for which we would prepare a version of read_mail for installation in (at least) the author-maintained library. Charlie Davis "installed" a new version of read_mail at the end of this quarter which fixed a number of minor bugs and introduced several small new features. In addition, Charlie brought Steve Swernofsky's Version 4 Message Segment installation up-to-date with several changes that had been made to the Version 3 Message Segment software installed in the system, and the Version 4 software is once again ready for checkout on the development machine.

Charlie's efforts on improved mail facilities were substantially diluted this quarter by his work on software for the new CSR Print remote IO Daemon printer. This work took up most of Charlie's time during the January Independent Activities Period.

Finally, Ken Pogran participated in a discussion of possible enhancement to the installed Multics mail software at a Multics mail software Change Review Board meeting, at the request of several of its members. The topic of discussion was whether incremental improvements should be made to the installed "mail" command, or whether they should be held in abeyance pending installation of a subsystem such as read_mail. Ken indicated that read_mail was of relatively low priority at CSR, and that CSR could not, without assistance from CISL or IPS, expend the effort required to install and maintain it in Honeywell's Standard Service System.

II. Modifications to Multics ARPANET Software

During this period, MIT installed the Network Library software that moves administrative control of the Multics NCP from ring 0 out to ring 1. The changes necessary to remove the old interfaces were submitted to MIT. These changes complete the development necessary for NSW support. No further support of the Multics ARPANET software is planned by CSR, except for training and second level support of IPC and HISI personnel.

Steve Kent and Doug Wells submitted the changes necessary to support "registered dial ids". After MIT installs these changes, incoming ARPANET mail will be received by "dialing" to a special pre-existing process, rather than creating a new process for each message. Because the overhead of creating these new processes is substantially more than the cost of processing the actual mail text, these changes should substantially reduce CSR's mail-related expenses.

III. Development of National Software Works

Multics is now an NSW Tool Bearing Host! During this period, the protocol for interhost transfers was implemented and checked out in the Multics and TENEX File Packages. Also, a special version of the Multics qedx editor was developed and made available as the first Multics tools.

During the next quarter, the MSG communication system will be augmented to include the capability to dynamically create MSG processes. Also, the Multics Foreman will be changed to match the "Interim Reliability Specification". At that point, a complete multi-host system, which includes TENEX and Multics components, will be released for evaluation by potential users. Also, that is the point where CSR will hand over maintenance and further development of Multics NSW software to the NSW group at RADC.

Research in Distributed Systems

I. LCS Local Network

Progress continues on the development of the LCS Local Network. Two undergraduate projects were begun, plans were made to acquire a PDP-11/35, and several meetings were attended by those working on the Local Net.

A. Hardware

Ward Harriman, a senior in Course VI, began an undergraduate thesis project under Dave Clark. Ward is implementing the Data Stream Protocol (DSP) on a Motorola M6800 microprocessor system which will provide direct access to the LCS Network for a small number (1-4) of terminals. During this quarter, Ward designed his DSP implementation, and, in conjunction with Dave Clark and Ken Pogran, determined that the M6800 was the most suitable microprocessor system. The microprocessor hardware was ordered at the end of this quarter, and Ward expects to complete his project by the end of the Spring term.

Tim McMahon, a junior, came on board in mid March as a UROP student supervised by Ken Pogran to implement a new ARPANET IMP interface for the PDP-11. This interface will be used for the ARPANET-LCS Net Gateway machine, described below. It is intended to be a low-cost substitute for the expensive DEC "special product", the IMP11-A. The design for the "IMP11-B", as it is being called, is being adapted from the Multics ABSI designed by Rick Gumpertz in 1973 and from the PDP-11 Direct Memory access interface used in the Local Network Interface being developed by Dave Farber's group at UC-Irvine. The IMP11-B will be compact, fitting onto a single PDP-11 "Hex" Small Peripheral Controller board.

Ken Pogran traveled to UC-Irvine in January to meet with Farber's group and discuss LNI progress. It now appears that the first LNI's will be available to LCS early in the second quarter. By the end of the first quarter, the LNI had been bread-boarded at Irvine and sent out to a wire-wrap shop, with the first wire-wrapped LNI board due back during the first days of the second quarter.

During this quarter, arrangements were made to transfer an ARPA-owned PDP-11/35 from System Development Corporation in California to LCS for use with the LCS Network. This machine, equivalent to a PDP-11/40, will serve as the Gateway machine between the LCS Net and the ARPANET. It will also be used for initial checkout of the Local Network Interface when it arrives, and will be the testbed for the Unix DSP implementation.

Roy Planalp rejoined our group this quarter, after deciding to major in Course VI (EECS) rather than Course V (Chemistry). Roy has tackled the task of learning how to "download" code to a DEC GT-40 terminal, which is really a PDP-11/05 in disguise, from Unix. Our original plans called for use of the 11/05 as the machine to which

the first Local Network Interface (LNI) would be attached. Roy was to write programs to exercise the LNI, and eventually, to allow the GT-40 to communicate with the 11/70 Unix using the first two LNI's and DSP. However, it now appears that the 11/35 will be available for LNI checkout, and Roy will re-direct his efforts accordingly.

B. Protocols

Dave Clark and Dave Reed attended a meeting on the Transmission Control Protocol (TCP) held in Washington, D.C. This meeting was primarily attended by those involved in implementing the recently-developed TCP Version 2 which is to be used in inter-networking experiments. Our group's interest in TCP is twofold: first, we intend to implement TCP on Multics, and, secondly, (and more important) as DSP is a derivative of TCP and must coexist with it, we are interested in its status and evolution. Dave Clark and Dave Reed elicited considerable interest in DSP from those at the TCP meeting. As a result of this meeting, Dave Clark is preparing a document contrasting DSP and TCP, which will be distributed to the TCP community.

Dave Reed finished a document describing his proposal for a high-level service addressing mechanism on the local network. Some of the ideas in this proposal are relevant to his proposed thesis topic dealing with the semantics of names in a distributed system.

Several considerations related to adding security to TCP via encryption were developed by Steve Kent and Dave Reed during a visit by Carl Sunshine. Steve is preparing a document describing these ideas to be presented to the next TCP meeting as a proposal.

Discussions were initiated with members of Steve Ward's group concerning an implementation of DSP on their PDP-11/70 Unix system. Dave Clark produced a "sample" DSP implementation as a talking point, and Dave, Roy Planalp, and Ken Pogran discussed how DSP should fit into the Unix environment with Steve Ward, Chris Terman, John Sieber, and Jim Gula. These discussions resulted in a commitment from Steve Ward to make modifications to the Unix kernel which will ease the implementation of DSP and Unix.

II. Study of Objects in Current Systems

Karen Sollins continued her study of objects and began looking into some of the issues surrounding the naming of objects.

III. Name Resolution in Distributed Systems

This quarter Allen Luniewski began an investigation into the design of a processor's low level addressing mechanism. The goal is to have a mechanism that "naturally" allows the binding together, in a naming sense, of objects that reside on potentially distinct systems. The principal activity has been background reading in the areas of distributed systems and capability based systems.

IV. Eventcounts

Dave Reed, with Raj Kanodia, finished a paper describing the use of eventcounts and sequencers in synchronization, and submitted it to SPSP-6. Raj is working on another paper to contain primarily ideas related to eventcounts in distributed systems, which they hope to submit to IEEE-TOC.

V. Measurements of Sharing in Multics

During this past quarter, Warren Montgomery continued to gather statistics on sharing in Multics. The results of this project were summarized in a paper submitted to the Sixth Symposium on Operating Systems Principles. A copy of this paper was issued as RFC 141. RFC 134 was also issued discussing in greater detail the methods used to gather this information.

VI. Study of Forces Leading to Decentralized Computer Systems

During January Cecilia d'Oliveira commenced an S.B. thesis on the topic of decentralization of the computer facility in various organizations. During IAP a preliminary literature survey was accomplished leading to a development of a list of forces proposed as relevant in user decisions regarding decentralization. Research continued through February and March and goals for further research were defined. A thesis proposal "A Conjecture about Computer Decentralization" was distributed at the end of March as RFC 139. If everything goes according to schedule the thesis will be completed sometime this summer.

VII. Information Security in Distributed Computing Systems

Paul Karger is investigating approaches to security in distributed systems, concentrating on the host to network and host-host interfaces. While much work has been undertaken to provide communications security, primarily through encryption, the host-host security protocols are less well understood. Paul is particularly examining the issues of networks consisting of unsecured host processors in which one-way communication is desired from low security hosts to high security hosts, but not vice versa. Paul expects to complete an S.M. thesis in this area in May.

VIII. User Level Naming in Distributed Systems.

David Reed is working on a thesis in the area of naming at the user level in distributed systems. The basic goal is to be able to name and locate objects which have multiple "equivalent" versions in the distributed system where multiple versions exist because of redundancy for robustness or due to natural evolution (in an "upward-compatible" manner) of an object that is copied to be used. He is also looking at the linkage between (program) modules, since human-level naming maps down into the inter-module naming

mechanism. Especially important in the thesis will be concerns with autonomy and robustness. Dave should finish the thesis proposal before this summer begins.

IX. Communication Security Using Encryption

Steve Kent prepared a paper for submission to the Fifth Data Communication Symposium on the topic of "Encryption Based Protection for Interactive User-Computer Communication". He also spoke before a federal committee on the use of the NBS encryption standard for communication purposes.

X. Seminar on Distributed Systems

This semester Jerry Saltzer is running a graduate seminar that is exploring in depth various issues related to the problems of distributed systems. Several other members of the group are devoting considerable time and effort to this course.

XI. Paper on Name Binding

Jerry Saltzer developed a new draft of his chapter of the 6.033 notes dealing with name binding in computer systems. This material may provide a basis for extension to name binding in distributed systems.

XII. The Concepts of Utility, Level of Coherence, and Distribution of Tasks in a Network of Cooperating Computer Systems.

Liba Svobodova composed a preliminary version of an RFC, outlining several ideas on research in distributed systems that might be worth investigating. First is the idea of the utility as a facility functionally distinguishable from the self-contained cooperating member systems. The utility does not have to be a centrally controlled system. In fact, the parts of the utility could be controlled by individual members. The distinction between utility and member to member cooperation is drawn on the basis of availability. The utility stands for those facilities that are always available to the system members independently of which system members are in a cooperating mode. The notion of "level of coherence" is introduced, and a question is raised of what hardware and software features are necessary to support a specific level of coherence, and what are the permissible variations in the hardware and software of future members of a distributed system.

An important area of research is the distribution of processing tasks over the nodes of a distributed system. The feasibility or preferability of different distribution schemes is strongly influenced by security requirements and performance requirements and limitations.