

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 1976.

During January 1976, the group held its annual internal research review. This three day event, which gives every member of the division the opportunity to present his current activities, seems to be an excellent way to keep our people in touch with what the rest of the group is doing. It also serves to keep outsiders up to date; this year's meeting was attended by various representatives of our research sponsors and other research labs.

In this quarter, the new version of the Network Control Program was finally installed.

Although no theses were completed this quarter, significant progress has been made on several fronts and next quarter should bring the completion of a number of significant projects.

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 referenced in other publications.



Kernel Design Tasks

I. Restructuring of Page Control.

Bob Mabee continues the rewriting of Andy Huber's page control, with the intention of learning as much as possible about the performance issues of the PL1 implementation. Unfortunately, NSS was installed before Bob completed his reworking, so he is now additionally delayed by the necessity of upgrading to NSS.

II. Multi-level Traffic Control.

Several chapters of David Reed's thesis were written during this quarter. It should be completed during the next. His thesis includes a fairly detailed high-level design for a multi-level process implementation. In addition, Bob Frankston has been added to the project, to work out the details of an actual implementation that can be integrated with the current Multics with minimal upheaval. Bob has been investigating both the set of issues to be dealt with in coding the new traffic control for the 68/80 architecture and the set of issues related to interfacing parts of the system such as page control and I/O device control, which must be converted to run as processes in the new regime. The detailed design of the first level of process implementation is complete; the design of the second level and the new page device control interfaces remain to be completed in the next quarter.

III. Restructuring of the Answering Service.

Since the last report, the code for the experimental redesign of the Answering Service has been debugged. Warren Montgomery is now in the process of writing his Master's thesis based on this design. A substantial proportion of the thesis has been written at this time and the thesis is expected to be completed for the May 7th thesis deadline.

IV. Fast Processes in Ring 0.

It now appears that Honeywell will implement and install a version of fast processes which is functionally rather similar to ours. We therefore hope to use this implementation as the basis of our rewrite of the Network Control Program.

V. A Methodology for Designing a Certifiably Correct Computer System.

Rich Feiertag, who is now working for SRI on the west coast, is continuing the writing of his thesis.

VI. Study of System Initialization.

Allen Luniewski continued his investigation of system initialization during the past three months. The principal activity this past quarter was the consolidation and organization of ideas stemming from RFC's 68 and 98. The writing up of the research, which will



become his Master's thesis, was begun.

VII. Support of User Defined Object Types.

Doug Hunt is continuing work on his EE thesis entitled "Building Blocks for an Object-Oriented Virtual Memory System". The thesis deals with modular decomposition of large systems as a technique to make verification possible. In particular, the thesis focuses on the inter-module relationships which may exist. It is claimed that just a few relationships, each with simple semantics, are sufficient to express inter-module dependencies in a large number of cases. To support this claim, a virtual memory subsystem--similar to a Multics virtual memory enhanced to support type extension--is considered as a case study. A modularization for the virtual memory is suggested and justified, and the inter-module relationships are specified.

VIII. Study of System Reliability and Recovery from Errors.

The past quarter Harry Forsdick has been involved in several activities. First, he has finished the work on the first draft of an RFC about a model of the occurrence and handling of errors in a computer system. Further work has been done on revising the presentation of those ideas. Based on the comments received from the draft RFC, he has started to look at the way errors are handled in the current implementation of Multics. He has traced down the seg-fault path in an attempt to develop a characterization of the structure of Multics with respect to errors. Several insights on the placement of error checks have come out of this study and will be documented in a future RFC.

IX. Multics Performance Benchmark.

No progress was made on the project (with one exception) during the quarter, and it appears that the development must be abandoned due to lack of effective manpower. There are not people with the time and interest to manage the project, and the undergraduates working of the project have not performed usefully. The need for the end product is waning, so no resuscitation is proposed. The one area in which progress has been made is the production of a working version of the "aware" software by Roy Planalp.

X. Cryptographic Synchronization in Secure Terminal I/O.

During this quarter, Stephen Kent completed a draft of his thesis on the use of encryption-based protocols for securing interactive user/computer communication. The final draft of the thesis will be completed in early May.

XI. Separating the Functions of Page Control and Segment Control.

During the winter, Drew Mason has been refining his ideas on separating page control and segment control. It now appears that the two modules can be separated, resulting in a virtual memory manager having the following nice properties: (1) segment control,

or some higher module, is entirely responsible for allocating and freeing pages of segments, (2) there is no hidden communication between segment control and page control through the Active Segment Table, (3) segment control can only perform operations on pages through well-defined subroutine calls to page control, and (4) page control knows nothing about the segments to which pages belong. His goal is to implement some of his ideas as his Master's thesis, to be completed by January, 1977. He is currently producing a thesis proposal.

XII. Organization of the Virtual Memory Mechanism of a Computer System.

As of March 31, Phil Janson has completed the writing of the first draft of his Ph.D. thesis on a method for producing modular, structured software to support the virtual memory mechanism of a computer system. The first part of the thesis describes the type extension concept that is recommended as the basis for organizing a virtual memory mechanism. A virtual memory mechanism should be regarded as implementing abstract information containers (e.g. segments) out of physical information containers (e.g. core blocks and disk records). The second part of the thesis explains how that type extension concept can be exploited to organize a virtual memory mechanism. In particular, it shows how one can implement the programs and the address space of the mechanism itself without violating modularity and structure. The last part of the thesis illustrates the use of the method by applying it to the redesign of the virtual memory mechanism of Multics. The modularity and the structure of the resulting system are evaluated.

XIII. Eventcount Process Synchronization Model.

Early in the quarter, Raj Kanodia and David Reed finished a draft of a paper describing the eventcount synchronization mechanism. This was published as a CSR group RFC, and distributed somewhat more widely than the usual RFC distribution. Comments have been received from various people on the document, notably from Dijkstra (lukewarm, but he hasn't finished reading it thoroughly), Parnas (who has added it to the list of synchronization mechanisms being considered for the operating system being designed at Darmstadt), Roger Needham (who liked it), and Dennis Frailey of Southern Methodist University (who has worked on related issues). The paper will eventually be revised and submitted for journal publication, after Raj and David finish other tasks currently occupying them, and after a reasonable number of comments have been received.

XIV. Provision of "Breakproof" Environment for User Programming.

During this quarter Jeff Goldberg has produced a completely new draft of his thesis proposal. He has been modifying the current Multics command processor to call programs either in the ring it is executing in, or in a higher ring. The problem of having multi-ring search rules for the command processor was studied and a workable solution developed. Finally, the problems related to multi-ring signalling were examined. An RFC will appear soon on signalling

with subsystems in mind.

XV. Input/Output in a Security Kernel.

This quarter, Gene Ciccarelli has been studying the Multics Arpanet software, trying to understand how it works, what division of responsibility exists between these programs, and where complexity arises. He is interested in learning how this complexity relates to the choice of network protocols, demands of security and protection, and efficiency. Although the structure of the network software mirrors somewhat the levels of Arpanet protocol, there are instances where this layering is violated, and he has been studying the reasons for these cases. He intends to examine the functions performed by the kernel software, and will try to isolate what functions logically need to be supervisory, which can be relegated to user-ring software, and how this division relates to protocol choice.

He is also looking at some other communication areas in Multics, e.g. the multiplexing of terminal communication from the main processors to the front-end processor. He is looking for structures and functions that are similar to those related to the Arpanet, and considering whether a restructuring of these systems would allow some unification and clarification of responsibility, lessening the overall complexity of input/output software in Multics. He intends to consider different network protocols that might encourage such simplification, while not sacrificing functionality or efficiency.

Technology Transfer and Network Related Tasks.

I. Improvements to the Network Control Program.

A revised Network Control Program (NCP) was finally submitted and installed this quarter. This version of the NCP fixes a number of bugs in making and breaking connections, uses the I/O Assignment Manager so that a user's connections are closed when his/her process terminates, adds a tracing mechanism to aid in detecting problems within the NCP, and changes the handling of Host-Host interrupts so that they can be distinguished from TTYDIM interrupts. Slight changes were made to both the NCP and the IMPDIM to improve their paging behavior during data transfer. Changes were made to the gates into the hardcore NCP to allow access control to be provided by Ring 1 modules.

II. Production of MPM Network User's Supplement.

Little progress was made on this during the quarter.

III. Improvements to Network Mail Facilities.

Work continues on the read_mail program described in the previous task report. Roy Planalp, an undergraduate, is now working on read_mail under the direction of Ken Pogran. Roy has completed several major changes to read_mail, and a new version of the program should be installed early in the second quarter. In addition, read_mail is now available on the Multics system at Rome Air Development Center.

Steve Swernofsky continues his re-writing of the net_mail command and the mail queuing facility. A good deal of progress was made during the January Independent Activities Period; Steve has completed the design and initial coding of the modules, and is presently debugging them.

Ken Pogran continues his membership on ARPA's Message Service Committee, a committee intended to standardize and improve upon the facilities used to send mail between the various hosts on the ARPA Network.

IV. Specialized Inter-Host Tools for the Development Machine.

Doug Wells and Nancy Federman have completed the first design and test phase in the quest for a command to move files between the development and service machine. A special I/O driver has been implemented which will open FTP connections to the development machine. Any dprint request specifying that device class will actually cause the file to be shipped over the ARPANET to the development machine. The current work involves allowing a request to be submitted even if the development machine is down, allowing variables for such arguments as destination host, and modifying the FTP server so it will actually dprint the transferred file.

V. New Login Protocols for Answering Service.

Arthur Gottlieb, who was working on this project, was forced to stop owing to other time commitments. While the project is still appropriate, we have no manpower to resume it.

VI. Exploration of Limits of Bandwidth of ARPANET.

Raj Kanodia continues to pursue this experiment, but there are no results to report as yet. We expect that numbers will be available next quarter.

VII. Development of National Software Works.

During this quarter, the National Software Works participants managed to agree on a set of inter-host protocols. A Multics trial implementation is under way. This trial implementation has brought out a number of weaknesses in Multics' resource control and interprocess communication. Work in the next quarter will include the completion of this implementation and investigation into augmenting Multics' capabilities in these areas.

VIII. Backup Using a Data Network.

Art Benjamin has been continuing work on his thesis about backup using a data network. The major accomplishment during this quarter has been the development of a file system model consistent with the operation of a list driven backup mechanism in a network environment.

IX. Implementation of RSEXEC on Multics.

During this quarter, Tony Smith has continued the design of subroutines to support a user interface to RSEXEC on Multics. He expects to finish by the semester thesis deadline.

Doug Wells has eliminated a few bugs from our RSEXEC server implementation, and with the installation of the new NCP (see task I), which was required for the successful operation of the server, we are now running the server on the development machine mini-service as an experiment.

X. Study of a Local Network for the Laboratory for Computer Science.

As part of our on-going study of a local, high-speed data communication network for the Laboratory, Prof. Dave Farber, of the University of California at Irvine, gave a seminar on the Distributed Computing System and its ring communication network.

Miscellaneous Tasks.

I. Terminal Support.

Bob Frankston and Doug Wells have been investigating terminal support for Multics. The research into improved support for terminals has been motivated by the inadequacies of the current installed Multics support, particularly in the area of supporting terminals using the system via the ARPA network. By implementing the support software in the user ring, we are able to better examine issues in terminal support by being able to easily change it in the light of our experience.

Support for the Terminet 1200 and VT52 terminals has been implemented to increase their usability as Multics terminals. Building upon Doug Wells' work, Bob Frankston has implemented a program that permits Multics to perform all terminal support functions, including character echoing, in a process dedicated to a single terminal, demonstrating the feasibility of such an approach and providing a basis for distributing some of the functionality.

Dave Reed wrote a version of qedx that is much more useful on high speed display terminals such as the HP2640 and VT52, which has gained wide acceptance in the group.

II. Renovation of Conference Room.

Ken Pogran wrote a memo to the Laboratory administration suggesting changes to be made to the Laboratory's Fifth Floor Conference Room as part of its upcoming renovation.

III. Cable TV for Lab.

Ken Pogran met with Charles Patton of the Center for Advanced Engineering Study, engineer for M.I.T.'s cable TV system, to discuss the design of a TV distribution network for the Laboratory.