

PROJECT MAC

November 6, 1973

Computer Systems Research Division

Request for Comments No. 42


CSR TASK REPORT

by David D. Clark

This RFC is the first of what will hopefully be a fairly regular series of task reports for the Computer Systems Research group at Project MAC. The purpose of these reports is to foster communication among the members of our group by bringing to their attention all the ongoing activities of the group.

Any additions or corrections should be brought to my attention, so that the next version of the task list can be made more accurate.

This note is an informal working paper of the Project MAC Computer Systems Research Division. It should not be reproduced without the author's permission, and it should not be referenced in other publications.



A. Certification Tasks.

I. Formal Representation of System Algorithms and Data Bases.

The goal of this project is to create some system representation which can be used as a standard for certification and as an aid in explaining the system to others.

Several projects are currently underway in this area. A representation of system data bases and related algorithms in the Vienna Definition Language is being performed by R. Bratt. As an initial experiment he is making a case study of the KST. A similar description with directory control using English as the descriptive language is being performed by D. Hunt. Finally, B. Greenberg, as part of his thesis, has devised a new language related to PL/I for describing data bases which attempts to avoid implications concerning the implementation of the data base structure. E. Wada is helping guide these projects.

The technique of structured programming provides one way to describe the modules of an operating system. Two case studies are underway to determine the relevance of structured programming to certification: B. Greenberg is describing page control and D. Reed is describing traffic control. These structured descriptions will probably be finished by the first or second week in November, at which point the group can assess the utility of the descriptions as a teaching aid, as a standard for certification, or as the certification language itself.

II. Control of Module Interconnectedness.

R. Feiertag is studying models which will allow the control of the interconnections required between parts of the supervisor. To begin this study, he will choose some part of the system and explore ways to modularize it. This preliminary task is predicted to take perhaps two months.

III. Removal of Linker from Ring 0.

The status of this project is as follows. (See RFC 23 for introduction to nature of task.) P. Janson will soon issue a Multics Technical Bulletin which outlines the implementation technique. He has just distributed his thesis proposal as an RFC. He has also prepared design notes which describe, for each module in the linker, the nature of the modification required. Implementation, which can begin in a few weeks, will be done by P. Janson and R. Kanodia, who will modify ring 0 to support and initialize the new linker, and by E. Thomas, who will modify the linker itself to run in the user ring.

IV. System Initialization.

In order that the system be certifiable, it is necessary to certify the "initial state" of the system. D. Reed is considering how to modify current system initialization to make this "initial state" easier to certify.

V. Recovery from errors.

Currently Multics attempts to recover without crashing from certain but not all system errors. Further, some of the recovery procedures, for example, the directory salvager, are rather ad hoc. It is thus rather difficult to certify that system errors do not compromise security. B. Williams is studying the various error recovery tools on Multics, and is comparing this aspect of Multics to the Burroughs systems, which appear to have effective and orderly error recovery tools.

VI. Study of Multics Security Holes

An attempt is made periodically to catalog all known ways to violate the security of Multics, including ways to crash the system, because study of how bugs arise and what is required to fix them gives insight into the problems of certification. D. Hunt is updating this list, and will eventually publish something similar to RFC 5.

VII. Audit of ALM Segments.

E. Thomas is commencing an audit of all ALM modules in the system with the goal of cataloging the reasons why ALM need be used in the supervisor. This project is just starting.

VIII. Census of Ring 0.

V. Voydock has prepared a summary of the size of all ring 0 modules, sorted by various criteria. This task is essentially completed; the results are published as RFC 37.

IX. System-Wide Segment Numbers.

It is possible that hardware which could support segment numbers on a per-system rather than a per-process basis would result in a simpler system. V. Voydock is currently studying this problem, by means of a simplified system model in which infinite memory is available.

X. Control of Access Required to Meter System.

Currently in order to meter the system, it is necessary to have access to phcs_, which is excessive, in view of the large number of people with legitimate need to meter. To control the access which goes with the ability to meter ring 0, R. Mabee is designing a Ring 1 module which will filter metering requests. This will be described in a Multics Technical Bulletin which should be out in the first week of November; the implementation should take two weeks.

XI. "Quick" Processes

If Multics supported a class of processes which could be scheduled rapidly enough that they could be used to respond to i/o interrupts, then interrupt handlers need no longer be part of the system nucleus. This would clearly simplify certification. A task will be started soon to implement this sort of process on Multics. Initially, R. Mabee will work on this task.

XII. Study of Process Coordination Issues.

R. Feiertag will write up the conclusions he has reached related to process coordination. He does not propose to pursue this task further; it appears at this point that the design of a process coordination mechanism might constitute a Master's thesis.

XIII. Integration of Network and tty Buffers.

Currently, I/O buffers are managed separately for every class of device on the system. The bulk of ring 0 would be reduced if buffering and related I/O functions for several classes of devices could be integrated. As part of the tty dim re-write being undertaken by R. Snyder, the network and tty ring 0 software may be integrated. No one as yet has been assigned this task.

XIV. Paper on Protection Mechanisms.

J. Saltzer and M. Schroeder are writing a tutorial paper on protection mechanisms. It includes a discussion of capability systems. A draft of this should be completed by January, 1974.

B. Performance Tasks

I. Disk Paging Traffic

The thesis of B. Greenberg concerns the measurement of disk paging traffic in order to extend the linear paging model to this region. All experimental work on this thesis has been performed and he is now in the process of writing up this material. His thesis should be completed in the next two or three months. As part of his thesis, he has prepared the structured program representation of page control which is described as one of the certification tasks.

II. Network Performance

As part of the integration of network and typewriter buffering in ring 0, it is necessary to know some of the characteristics of messages sent over the network. There are currently meters in ring 0 which collect this information but there does not exist a user ring program to extract this information. A. Benjamin has taken on the task of writing a program to read and analyze this information.

III. Measurement of Service System Performance

For the past few weeks several members of the CSR group, in particular B. Greenberg, D. Reed, and D. Clark, have been measuring the performance of the service system in an attempt to discover why it is not running as well as was initially expected. The most fruitful area for study has been the paging behavior of the system, and much effort has gone into understanding the effect of the bulk store on the 6180 system. This task is not completed by any means, but little work is being done at the present time. New paging meters which will be added to the system may give further clues as to why the system spends such a large percentage of its time paging.

IV. Multics Working Set Estimator.

The bachelor's thesis by D. Reed which was submitted last June proposed an algorithm for estimation of working set size which could be used by the traffic controller on Multics to set the maximum number of eligible processes. No further research in this area is anticipated at the present time because it is not obvious that a working set estimator is required for the 6180 Multics. However, a report of research to this point will be prepared.

V. Hardware Measurements of the 6180 Processor.

A task is underway to make direct measurements of the performance of the 6180 processor by connection of counters to the hardware while the service system is in operation. These measurements should reveal such things as the raw instruction rate of the machine, the extent to which the machine is memory bound (which will assist in determining whether a cache would be useful), and the extent to which the suppression of overlap by certain instructions is slowing the processor down. D. Gifford should make the hardware measurements sometime in the first week in November.

VI. Model of Disk Traffic

L. Scheffler is writing a Master's thesis which discusses a probabilistic model of moveable-head disk subsystems. His optimizing dim is being used as a test case for this model. The thesis should be done this term.

C. Network Tasks

I. Implementation of new TELNET protocol.

Multics software must be modified to support a revised definition of the network TELNET protocol. (The TELNET protocol is used when logging into a remote computer.) D. Wells is coding this revision, which will be done by the first of the year.

II. Implementation of new File Transfer Protocol.

The definition of the File transfer Protocol has also been revised; the Multics software is in this case being revised by K. Pogran. This task will also be completed by the first of the year.

III. Implementation of Level 0 Graphics Protocol.

K. Pogran started the implementation of this protocol, and has gotten an approximate version of it to execute. M. Gross has taken over the task of finishing the implementation, which will hopefully take less than two months.

IV. Multics Programmers' Manual, Network Users' Supplement.

M. Padlipsky is beginning to prepare a document which will augment the MPM for those users who are involved with the network. Because of the nature of the network, it is anticipated that a large portion of the document will be introductory material rather than description of modules. An outline of this supplement exists presently. It is anticipated that a rough draft might be ready by shortly after the beginning of the year.

V. Attachment of Multics to the New IMP.

There is at the present time a reportedly operational IMP located in Building 39. It is our intention to connect Multics to this IMP. Currently, we are in the process of constructing the cables necessary to connect between the two. K. Pogran is overseeing the manufacture of these cables, which should be completed in the first week of November. As soon as these cables are completed, we can test the ability of Multics to communicate with the new IMP. We expect that the formal cutover will occur three to four weeks after we have successfully communicated.

VI. Attachment of Development machine to ARPA network.

It is intended to attach the Multics development machine to the ARPA network. This will allow testing of new network software without disrupting the return connection to the service machine, and will also allow new means of access to the development machine for system development. In order to make this attachment, a second ABSI is needed; see task B.VII. Also, if the IMP in Building 39 is to be used, rather than the IMP at Project MAC, an additional board is needed for the IMP.

VII. Status of the Asynchronous Bit Serial Interface.

The asynchronous bit serial interface, or ABSI, is the hardware interface between Multics and the network. Currently, we are running with a prototype ABSI which was manufactured by R. Gumpertz as part of his Bachelor's thesis. K. Pogran has taken over the responsibility for bug fixes on this prototype. He has also taken over the task of finishing certain aspects of the design. In particular, a design of a maintenance panel for the ABSI was not completed. Another task is to collect together all the implementation information necessary so that CSR can hire a technician to replicate the ABSI so that we can have a connection the development machine, and so that we can also obtain a physical piece of hardware which is more stable than the prototype which we are now using. It appears that the Information Processing Center is interested in funding the manufacture of a more stable ABSI for the service machine, whenever this information has been collected.

VIII. Specification of the Uniform User Level Protocol.

A user logging into a new system over the network is faced with the difficulty of learning the command language of this new system. In order to avoid this difficulty, M. Padlipsky is drawing up the specifications for a network-wide command language which would give the user a uniform interface to all systems on the network. A preliminary description of this command interface currently exists. The first piece of this uniform user level protocol to be introduced to the network community is a text editor called eds which is described in Chapter 4 of the MPM.

VIII. Implementation of RSEXEC on Multics.

RSEXEC, which is short for resource-sharing executive, is a network protocol which allows the user to exercise an equivalent function on several systems by the issuing of one command. For example, a "who" command issued to a RSEXEC would list all the users logged in every system currently running an RSEXEC monitor. Its principal function is to allow easy sharing of the file systems so that a user can type "list" and have printed what files he has on every system he uses. One view of RSEXEC is that it is an alternative means of achieving the same goals as the unified user-level protocol. In order to explore the strengths and weaknesses of RSEXEC, we propose to implement an RSEXEC monitor to run on Multics. Since RSEXEC so far has only been implemented for PDP-10s, implementation on a different hardware may prove a good test of the RSEXEC concept. M. Gross and W. Montgomery will be working on this implementation.

IX. Maintenance of Network Software.

The Computer Systems Research Division has responsibility for the software which runs the network on the service system. Maintenance tasks such as bug fixing continually occur. All members of the Network group contribute to this effort.

See also B. II and A. XIII.

D. Miscellaneous Tasks

I. Implementation of LISP on Multics.

The project of getting a version of LISP running on Multics has been going on for quite some time. At this point, we are essentially ready to install a version of LISP which is compatible with that which runs on the PDP-10s at Project MAC. D. Moon, D. Reed, and J. Gross will be responsible for the continued support and maintenance of this version of LISP.

II. Implementation of MACSYMA on Multics.

On about the first of November a version of MACSYMA was executed on the Multics system. It still appears to have a few bugs, but in the near future it is hoped that a version will be available which is stable enough to allow comparison of cost and efficiency between MACSYMA on Multics and on the PDP-10. D. Moon and A. Sunguroff are currently working on this project.

III. I/O Switch Enhancement.

In response to MSB-113, several members of the CSR group have been considering ways to increase the efficiency of the I/O switch without sacrificing the functional capability. Over the past several weeks two meetings have been held, involving a fairly large number of people, to discuss modification to the I/O switch.

D. Clark is currently writing a memo describing the results of these various discussions.

IV. Marketing of Services in a Computer Utility.

R. Frankston is writing a Master's thesis discussing those mechanisms which must be added to a computer utility, such as Multics, to create a marketplace where users can buy and sell services. Central to this marketplace is the creation of a financial transaction manager. The thesis should be completed this term.

V. MPM and SPS Document preparation.

Typing and editing of MPM, SPS, and other Multics documentation continues to be performed by S. Grant and O. Carey, under the direction of D. Jones. MPM Part II revision 15 is scheduled to go into production during November, and Part I updated through revision 13, will become a Project MAC Technical Report in November. This task is scheduled to be transferred to HISI by January 1, 1974.

VI. Study of SPS Utilization.

Currently, there are more SPS updates distributed than there are updated SPSs. D. Hunt has performed a survey to locate all valid copies of the SPS; hopefully this survey can result in a consistent policy for the placement of SPSs in our offices.