

PROJECT MAC

October 24, 1975

Computer Systems Research Division

Request for Comments No. 91

CSR TASK REPORT
by David D. Clark

The attached report covers progress of the Computer Systems Research Division of Project MAC in the period 1 July 1975 to 30 September 1975.

With this report, we are instituting the practice of identifying some of the significant events of the reporting period.

During this period, several such events occurred:

Andy Huber has produced a running version of his multi-process page control.


Steve Kent has completed the design of a practical use of encryption as a means for user-computer authentication and communication.

Dave Redell has led an in-depth study of directory control and the AST.

Jerry Rudisin completed his implementation of an RSEXEC server on Multics.

These events are discussed in detail in the report.

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.



kernel Design Tasks

1. Separating the Directory Heirarchy from the Segment Catalog.

As previously reported, this topic was being investigated by Drew Mason as a possible area for Master's thesis research. During the past three months, Drew has been a member of the study group which has been investigating the directory control subsystem and the Active Segment Table. As a result of this investigation, Drew has concluded that an attack on the complex interactions centered in the Active Segment Table can best be launched from a new position. Thus, he is now considering a new topic concerning the relationship between page control and segment control, which is discussed as task XV. Research in the structure of a flat file system is being continued by Loug hunt. See task IX.

- II. Restructuring of Page Control.

Work in this task is nearly completed. Andy Huber has developed a complete design for a page control subsystem structured as parallel cooperating processes, and he has written programs to implement these algorithms. This code has undergone extensive testing, both in a simulation environment and on the development machine, and has successfully performed all the page control functions. It is still too early, however, to evaluate the performance of this code.

- III. Restructuring of Traffic Control.

This project has made little progress during the reporting period, since Dave Reed was not in Cambridge over the summer. He has now returned, stimulated, we presume, by the invigorating change of atmosphere, and intends to work vigorously on the project during the fall semester.

- IV. Restructuring of the Answering Service.

Warren Montgomery's redesign of the answering service has proceeded to the point where considerable code has been written, although certain design decisions still remain. Warren has also written a draft RFC discussing possible implementations of domains. Warren expects that this project will be completed by the end of the fall term.

V. Fast Processes in King O.

The fast processes, which were implemented by Bob Mabee, are currently being used as a basis for the multiple process implementation of page control. This provides a further demonstration of the validity of the design and implementation. When Bob can find the time, he will prepare a revised version of the Multics Technical Bulletin which describes the fast processes, at which point, it appears, they can be installed as part of the standard Multics System.

VI. New I/O Buffer Strategy.

It is necessary to continue to report lack of progress on this project, as Dave Clark and Raj Kanodia continue to be occupied by other more pressing topics. Hope is again renewed, however, that progress will be made in the next quarter towards a paper describing this mechanism, and towards an integration of this mechanism into revised network interface software.

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

Rich Feiertag has published, as RFC 87, his thesis proposal on this topic. He has also produced preliminary drafts of other chapters, which have received limited distribution within the group. Rich is currently completing his plans to move to the west coast, where he will finish his thesis while he is working for SRI.

VIII. Study of Multics System Initialization.

There is little progress to report on this project over the summer, since Allen Luniewski was not in residence in Cambridge working on it. He has, however, prepared a new draft of his thesis proposal. With the thesis proposal out of the way, he intends to work vigorously on the thesis research during the fall term, and hopes to complete it in this quarter.

IX. Support of User Defined Object Types.

Loug hunt continues to develop a system architecture in which directory objects are not implemented as a fundamental part of the system, but rather are a higher level object defined in terms of base level objects within the system. He is attempting to specify the components of a system in which such extension is possible. In particular, he is interested in describing an appropriate interaction among the page manager, the segment manager, and the module which maps

between unique ids and objects. Thus, he is considering a system based on a flat file system.

Phil Janson is currently examining other possible research topics, and is not developing his ideas in this area further.

X. Study of Relationship between Security and Reliability.

Harry Forsdick is among those who were absent for the summer, so that there is no progress to report. He is, however, among those who have returned from their summer's absence with enthusiasm and fresh insight. He is hopeful of results to report in the next period.

XI. Multics Performance Benchmark.

Over the summer, considerable progress was made on the version of the Multics benchmark which uses the ARPANET to provide the experimental load for the system under test. Arthur Gottlieb made operational the "aware" software initially developed at Ford. The software is now in the condition such that it can be run on the experimental machine during the test to provide incremental readings on all normally used Multics meters. He has also, in the process of making these modifications, familiarized himself with "aware" so that he is in a position to make further modifications as necessary.

Barry Grant spent the summer doing two things. First, he completed the coding and checkout of a program to monitor login sessions and gather statistics on command usage. This program may now be placed into operation. Second, he took over the script driver program initially coded by Tony Smith, and familiarized himself with it so that he is now in a position to debug it, and to create test scripts for it. He is currently in the process of creating a set of scripts for this driver which exactly duplicate the effect of the scripts used in the absentee job version of the benchmark test. These scripts will be used to attempt to establish a correlation between the two forms of the performance benchmark.

A new undergraduate research opportunities program (UROF) project is being organized for the fall to continue this project of system metering, and to make use of some of these tools which were developed during the last year. Liba Svobodova will be in charge of this project.

XII. Cryptographic Synchronization for Securing Terminal I/O.

This research has not previously been reported in any quarterly task report, but was described as a new task in the annual report of the division last June. It involves the use of an encryption scheme to provide bi-directional authentication between the user and a computer, as well as to provide security for the user's data itself. The research topic is not in the encryption system itself (the LUCIFER system is being assumed) but rather in the protocols that must surround error recovery, lost messages, and user-requested interruptions. Steve Kent has completed the implementation of a test system in which he can evaluate the protocols which he has devised for these functions, and is now experimenting with this set of programs.

XIII. Group Study of Directory Control and the Active Segment Table.

During the summer, Dave Redell led a group in a detailed investigation of directory control and the Active Segment Table (AST). The purpose of this study was to understand the reasons behind the large bulk and complexity of these systems, and to propose strategies for simplification. One immediate goal is to eliminate the rule that requires that when an AST entry is active, its directory parents must also be active. If this rule is eliminated, the mechanism that enforces the rule can be eliminated, and knowledge of the directory structure will not penetrate so deeply into segment and page control. It is expected that the output of this study will be a variety of proposals for research in various areas related to these topics. Two immediate spinoffs are mentioned in the tasks below.

XIV. Study of the File System "Quota" Mechanism.

The quota mechanism, the mechanism that regulates the number of records that may be placed in any directory or directory subtree, appears to contribute considerable additional complexity to several parts of the Multics system. The reason is that the creation or removal of a single page may require changing quota values all the way up the directory hierarchy at page fault time. Nancy Federman is making a detailed study of the quota implementation on Multics, with the hope that she may be able to propose a simpler mechanism that will provide suitable functionality while at the same time not interacting so strongly with many of the system functions we are attempting to simplify.

XV. Separating of Page Control and Segment Control Functions within the Active Segment Table.

Currently, an entry in the Active Segment Table contains two sorts of information. It contains the page table words, which are needed by page control, and it contains other variables which describe segment attributes rather than attributes of the individual pages. These variables are used both by page control and segment control, and this causes needless interaction between the two systems. Drew Mason believes that it is possible to partition these variables in such a way that in every case they are referenced by either page control or segment control, but not both, and he believes that this separation will lead to a significant simplification in the algorithm associated with page control. In particular, it is possible that a thorough separation of page control and segment control in this way would lead to a simplification in the locking strategies required during the handling of a page fault and a segment fault. Drew is considering this topic as a possible one for Master's Thesis research.

XVI. Design of a Multi-Domain Kernel.

Phil Janson is considering the question of how multiple domains should best be used in the design of a kernel. He is attempting to determine whether the use of several domains within a kernel would make that kernel simpler or easier to certify. As an initial proving ground for some of his ideas, he is exploring the area of directory control, segment control and page control, to see whether multiple domains can be usefully exploited to simplify them. A document describing his proposed system structure exists in draft form.

XVII. Removal of Linker from Ring 0.

Phil Janson has prepared a paper describing the problems and implications of removing the Multics linker from ring 0. The paper discusses the general problems of linking between modules in a multi-domain environment. The paper will be presented at the upcoming Symposium on Operating Systems Principles in Texas.

XVIII. Paper Describing the CSK Kernel Design Project.

Mike Schroeder has prepared a paper which discusses the overall goals of our current Kernel Design Project. It considers the defining nature of a security kernel, and role of system validation. The paper, titled "Engineering a Security Kernel for Multics," will be

presented at the upcoming Symposium on Operating Systems Principles.

XIX. A Lock Free Strategy for Process Synchronization.

Over the summer, Raj Kanodia documented in RFC 85 a strategy for communication and coordination between processes which does not require a global lock on a data base. The facility provided by this mechanism is a queue in which many processes may place messages, and one process may remove these messages. Protection against interference among the various processes is provided by the read-alter-rewrite memory cycle instructions on the processor. This scheme, if generally used, would eliminate the necessity to certify individually the correct operation of various interacting and cooperating processes within the kernel of Multics.

XX. Elimination of the ITT.

Bob Mabee has as time permits been working on the elimination of the Interprocess Transmission Table (ITT). A memo on this redesign should appear in the future.

XXI. New Process Control Mechanism.

Raj Kanodia and Dave Reed are writing a paper describing a new mechanism that they have devised for process co-ordination. This mechanism, which uses monotonically increasing counters to record event occurrences, is now being used to control fast processes. (see task V).

Technology Transfer and Network Related Tasks.

I. Production of MPM Network Users Supplement.

The official first edition of the Network Users Supplement has not yet been issued, and we have decided that we will not attempt to publish it for another several months. Our intention in postponing the official publication date is that we wish to document the network at a point where the interfaces are in a stable and consistent condition, and we expect that we will achieve such a state in three or four months. In particular, at that point the new user interface to hardcore routines will be documented and installed, and the new user interfaces to the network user_telnet and user_ftp commands and the network IOSIMS will be available for use. Documentation for these new modules is currently in preparation.

II. Preparation of Network Program Logic Manual.

Ken Fogran, the editor of the Network PLM, has been prevented from working on this task by several other projects which have occupied all his time, especially the study of the Project MAC Network and the IMP-TIP swap. We fear that the progress of the Program Logic Manual will continue to be slow as long as high priority tasks continue to emerge faster than we can deal with them.

III. Implementation of RSEXEC on Multics.

Jerry Kudisin has completed the design and implementation for an RSEXEC server on Multics. The server appears to implement all its functions correctly when user commands are transmitted to it from Multics; however, there are still a few problems when the Tenex user environment attempts to interact with it. We are currently attempting to eliminate these difficulties, at which point we may install the RSEXEC server as an experimental service. With the server now operational, we are turning our attention to the implementation of the user environment for RSEXEC on Multics, which we hope to design and implement over the next three to six months.

IV. Improvements to the Network Control Program.

Loug Wells developed a new version of the NCP, which uses the I/O Assignment Manager to control the allocation of sockets to processes. It also contains minor modifications which will enable RSEXEC to operate reliably, and fixes several known bugs. This code will be installed shortly.

V. New Programs for User Interfaces to the Network.

Over the past several months, Doug Wells has been performing a coordinated redesign of the user_telnet and the user_ftp commands, and the various network IOSIMS, such as net_ascii_. The result of this redesign will be the versions of these programs which we propose to install and which we will document in the first edition of the Network Users Supplement. These programs are now recoded, and are in the checkout and documentation phase.

VI. Redesign of Network Mail Facilities.

Steve Swernofsky continues his recoding of the network mail programs. He has completed the redesign of the net_mail command, used for sending mail over the network. This command now has several new features, including the ability to send mail to prespecified mailing lists. He is now working on the mail queueing mechanism, changing the interface between it and the mail command. As part of this recoding he will modify the mail queuer so that it will operate in the multi-tasking environment which we intend soon to bring up in the Network_Server process.

VII. Multi-tasking in the User Ring.

Arthur Gottlieb is currently familiarizing himself with the multi-tasking software originally coded by Doug Wells. His first project will be to use the software to get a multi-tasking environment operating inside the Network_Server process. We will then introduce into this environment the RSEXEC server which Jerry Rudisin coded last summer. The other immediate use for this multi-tasking software is as part of the metering software described as Kernel Design Task XI.

VIII. Development of Network Host Database.

As discussed in the task report for the first quarter, Nancy Federman has implemented a facility which maintains an association between the host number, host name, and attributes of every host on the network. The subsystem is currently being installed. Write-arounds have been provided so that the old interface to obtain this information, host_id_, can still be used.

IX. Development of National Software Works on Multics.

This quarter has seen little activity in the area of the National Software Works, largely because the NSW project is undergoing a period of reorganization, which has provided us with little to do. Our effort has been

confined to attending several meetings at which we have discussed the future directions of NSW and tried to define what our role in this project will be. It appears that in the next quarter, a new protocol definition will emerge, which we can evaluate and, as appropriate, implement.

X. Distributed Data Base Management.

Art Benjamin is continuing his investigation of distributed file systems and the use of a data network for improving the reliability of on-line data storage. During the last quarter, his thesis proposal was signed, and he has implemented an experimental program which allows files to be transmitted between Multics and the Datacomputer. This research may also provide a means to explore an idea which came out of the AST study group (also discussed as task XIII of the Kernel project), which is that backup should be driven by messages generated at segment deactivation time this would simplify segment control by eliminating the need to propagate "date-time modified" up the hierarchy, and would eliminate the sweep of the hierarchy now performed by backup.

XI. Bringing the Phoenix Multics Site onto the ARPANET.

We continue to attempt to expedite the attachment to the ARPANET of the Multics site which is maintained in Phoenix by Honeywell. At this point, considerable interest exists within the Honeywell community, and we are hopeful that tangible progress can be made in the near future. As part of our effort in this area, Ken Pogran went to Phoenix to present an overview of the ARPANET and its benefits to a system such as Multics.

XII. Rearrangement of Cambridge Area IMPs and TIPs.

In order to solve a bandwidth problem which was afflicting the Computer Corporation of America, ARPA found it necessary to relocate the various IMPs and TIPs which are near M.I.T. As a result of these moves, the IMP which was located in building 39 has now been moved to CCA at 575 Technology Square, and the TIP which was there has now moved to the ninth floor of Project MAC, next to our original IMP. This move, which was coordinated for M.I.T. by Ken Pogran, was performed with only a minimum of upheaval, given the large number of phone lines which had to be relocated as part of the move. The CSR Division now can take advantage of the proximity of the TIP by directly wiring certain of our terminals to the TIP ports. Ken is experimenting with the feasibility of this means of terminal attachment.

XIII. Study of a Local Network for Project MAC.

During the summer, Ken Pogran has been evaluating the relative merits of several schemes for providing a local network to serve the various present and future computers of Project MAC. The study is a direct outgrowth of the committee which met last spring to consider whether Project MAC should at this point develop such a network and is expected to result in specific recommendations for implementation strategies. The conclusions of this study should be reported in a document fairly soon.