

PROJECT MAC

October 11, 1974

Computer Systems Research Division


Request for Comments No. 62

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 1974 to 30 September 1974.

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. Removal of Linker from Ring 0.

During the summer, Dick Bratt, Phil Jansen, and Elaine Thomas completed the coding and checkout of the user ring linker. A system with this linker has been produced which will successfully execute all commands in the standard system performance measurement test. We are currently attempting to improve the performance of the linker to the point that it can be installed. Initial tests of the first version of this new linker indicate that the performance test runs 7 to 10 percent slower with the new linker; Phil Jansen has produced a slightly reorganized version of the new linker which is expected to run perhaps 2 percent slower than the standard system. Andrew Mason is currently studying the linker to determine how its performance can be made to surpass that of the currently installed linker. A concrete proposal for the installation of this linker will follow from the completion of this study.

II. Removal of Name Space Management from Ring 0.

Dick Bratt has decomposed this project into two sub-tasks. The first sub-task is the redesign to the programs and data bases which implement the actual functions of KST management. Coding of these programs has been largely completed. The second sub-task is the redesign of those modules in Ring 0 which need to convert path names to segment numbers either because they have been handed a path name as an argument, or because they have a pathname imbedded in them internally. Dick is currently attempting to devise a strategy which, by leaving the reference name manager in Ring 0, will allow him to test out his ideas without recompiling all of these programs, about 40 in number. It is possible that within the next four to six weeks a preliminary version of the system will be available which can be used to demonstrate the viability of his overall strategy, and to determine the approximate performance improvement which will be obtained when the new linker uses this mechanism.

III. Removal of the Storage Hierarchy from Ring 0.

Doug Hunt is currently considering the question of whether the concept of hierarchy need be in the kernel of Multics. It is possible that the storage system can be constructed as a single-layer catalog of segments, indexed by unique I.D. The association of names to unique I.D.'s, and the use of segments as directories to contain these associations, can be placed in an outer ring, either ring 1 or the user ring. This study is still in its initial stages, and no forecast can be made

of its timetable.

IV. Restructuring of Page Control

Andy Huber is continuing his research into possible ways of decomposing page control into a number of cooperating processes. It is hoped that within the next two or three weeks a description of his version of page control, coded in the GSPL language, will be available. We are attempting to stay in step with the current interest in page control redesign at Honeywell.

V. Restructuring of Traffic Control

Dave Reed is continuing his research into the restructuring of traffic control in order to speed up the switching from one process to another. A new document should be available in the next few weeks.

VI. Restructuring of the Answering Service

Warren Montgomery is considering the question of what components of the answering service must, in principle, be certified. He has proposed a modification to the process creation strategy which would allow most of the answering service to be considered as uncritical, and which would allow a process to be created by any untrusted process, rather than only by the answering service. It is hoped that a description of his proposal will be available in the near future.

VII. Fast Processes in Ring 0

Bob Mabee is proceeding rapidly on the project of providing an environment within ring 0 in which processes which make very restricted demands on their environment can execute very inexpensively. He has implemented a class of process which has a legitimate stack, which can abandon the processor by means of the wait and notify mechanism, and which can take page faults, but which is restricted from taking segment faults or adding segments to its address space. The wired storage required from these processes has been reduced by two strategies. First, there is one descriptor segment per processor, used by any of these small processes when it runs. Second, the PDS has been split into two components, one of which is not needed by the small processes. Approximately one-fourth of a page of storage is required for these processes, and the rest of that page is used for stack. Thus, if a process requires less than three-quarters of a page of stack, there is one page required in core when that process runs.

As a test of these processes, Bob is currently rewriting the interrupt side of the typewriter dim so that it runs as one of these processes rather than directly as a result of the interrupt. Approximately nine pages can be unwired as a result of this modification, which should cause some system performance improvement; this improvement will be somewhat offset by the necessity of going through the traffic controller for each 355 interrupt, approximately seven times a second. Other benefits of this which can be additionally achieved by further modification of the typewriter code are a simplification of the locking strategy and an increase in the flexibility of the data base structure, since tables no longer need be wired down. An initial version of this typewriter manager process has run for an extended period on the development machine.

VIII. Study of Security Holes

As part of our continued interest in attempting to discover ways to circumvent the security of Multics, Dave Reed and Harry Forsdick have identified one general class of security loop-hole, and have automated a search of the system to determine all modules which might be susceptible to this particular attack. A detailed discussion of the technique used in this analysis will be published as soon as the bugs discovered have been fixed.

IX. Multi-tasking In The User Ring

We are continuing, as time permits, the implementation of a user-ring environment in which the user's process is shared amongst several tasks. Doug Wells has implemented a version of user ring IPC which allows several tasks to abandon the processor by calling block. This will be used as part of an experimental command environment, an initial version of which is already implemented, in which the command processor runs each command in a separate task. One obvious advantage of this structure is that when a quit is signaled in the user's process the result is only that one task is suspended and control is returned to the listener. Any new task started will run on a different stack, thus signal handlers for various tasks never become confused, and it is possible to restart any task in any order with any number of other tasks suspended. Experimentation with this user environment will allow us to specify a simple but effective quit mechanism within the kernel of the system, in place of the complicated one currently existing there. Another use for this multi-tasking facility is to multiplex a server process currently being used as part of the network environment. It is our intention that in the very near future we will add several new functions to

this process, and rather than code these as event call channels, we will use this multi-tasking IPC so that each function can be coded as if it were the only component of the process. The multi-tasking IPC could also be used to simplify the structure of the Answering Service, currently a very highly multiplexed process, which is within the security perimeter of the current system.

X. Development of IO Buffer Strategy

Raj Kanodia and Dave Clark are continuing their development of a buffering strategy which uses the virtual memory as a storage allocation scheme. They have attempted to insure that the implementation being proposed is compatible with the I/O interfacier currently being installed in the system by Honeywell. It is possible that some improvement in the scheme currently proposed will result if the fast processes described in Section VII are integrated into the buffering strategy. Hopefully, a document describing this new version of the buffering strategy will be available in a few weeks, and can be installed in the system as part of the IMP DIM, fairly soon.

XI. Study of Burroughs Operating System.

Ben Williams is studying the Burroughs operating system to see how their error recovery procedures may be applied to Multics. The result may be error recovery procedures which are easier to understand and certify. This project is described in the next section of the report.

XII. Formulation of Criteria for Inclusion of Modules Within the Kernel

Rich Feiertag is continuing his attempt to identify a set of general rules which will specify those modules which must be within the kernel of an operating system. As part of this research he has produced RFC 56, A Security Kernel Model for Page Control, which is a discussion of the criteria for inclusion within the kernel of that particular portion of the system. His intent is to perform a similar study for various other portions of the system, and out of these various studies to evolve a general theory of the structure of a kernel.

B. Technology Transfer and Network Related Tasks.

I. ARPA Net Connection to Development Machine.

A major accomplishment of the summer was a successful connection of the ARPA Net to the Multics Development Machine located at HISI. This project had two major sub-components. (In fact, the project had something in excess of thirty sub-tasks). One task was hauling a cable from the location of the development machine on the third floor of 575 Tech Square to the ninth floor of 545 Tech Square, the location of the IMP. This required locating appropriate ducts for the cable and required a great deal of detective work on the part of Ken Pogran. The other task was the replication of our hardware interface, the Asynchronous Bit Serial Interface, or ABSI. Rick Gumpertz, who returned for the summer, assisted us in producing an up-to-date wire list and set of schematics. We obtained another set of parts from Honeywell, and had ESL wrap and solder another circuit board. The completed board had only one manufacturing defect, a chip soldered in backwards, and the final testing was delightfully without complications. The existence of this second connection to the ARPA Network has allowed us to test for submission a variety of modifications to the Ring 0 software, and to discover and diagnose a number of lingering bugs within the network software. We are currently attempting to make this connection to the network more widely useful for other research purposes.

II. Redesign of ABSI by Electronic Systems Laboratory.

John Ward, of the Electronic Systems Laboratory at M.I.T., is currently overseeing a redesign of our ABSI, for the purpose of using commercially available parts and circuit boards. This version, when completed, will be used on the M.I.T. service machine, and also by Rome Air Development Center. Since ESL will have the knowledge necessary to maintain this new version of the ABSI, we will be relieved of the day-to-day responsibility which we have been forced to hold for the service of the ABSI now in operation. Since this new ABSI is manufactured from commercially available parts, it should in principle be easier to maintain, thus the overall result of this design should be a more stable network connection, and one for which we have no immediate responsibility.

III. Installation of New Ring 0 Network Software.

A new version of the NCP has been installed, which fixes certain bugs, and also attempts to behave better under the circumstances of resetting a foreign host. We suspect that problems in this area have prevented us from

communicating with certain hosts under various conditions. A new version of the IMP DIM is being prepared for installation; this also fixes several known bugs and has several new features. One feature of particular interest is a slight modification to the buffering strategy for messages being written to the network, which increases the size of blocks used in a linked list. This simple modification appears to have increased the speed of large message transmission by more than a factor of 2. The maximum rate of transmission for any one message is now approximately 27 kilobits per second, and the total output rate for the system is about 40 kilobits per second. This compares to a maximum rate of approximately 14 kilobits per second with the old version of the IMP DIM.

IV. Installation of User Interface to TELNET and File Transfer Protocol.

Doug Wells is just installing the telnet command, which allows a user to log into a foreign host on the network via the TELNET protocol. He is currently preparing for installation his version of the user_ftp command, which allows files to be moved from one site to another using the file transfer protocol.

V. Network IOSIMs.

We are currently constructing network interface modules which will allow standard Multics IO streams to be attached to the network. Because of the various ways of using the network, it appears that at least two different such modules are required. Doug Wells has coded and de-bugged a network IOSIM, called net_ascii_, which will support a Multics IO stream over which characters are being transmitted. When the data being transmitted over a stream is not a sequence of characters, but rather a sequence of bits organized into bytes of some size, a different IOSIM is required. The tentative specification of such an IOSIM exists, but we are uncertain how to integrate its interface into the one specified by iox_. It is also unclear the extent to which we wish to support such an IOSIM.

VI. Production of MPM Network Users' Supplement.

Work has again commenced on the NUS, which has sat dormant since the departure of Mike Padlipsky. Dave Clark and Ken Pogran are currently attempting to put together a draft of the entire document, which can be distributed to interested friends for comment. Hopefully, this draft will be available before the end of October.

VII. Preparation of Network Program Logic Manual.

Ken Pogran, the editor of the Network PLM, has been prevented from working on this task by several other projects which have occupied all his time, including fabrication and connection of the ABSI on the development machine and preparation of the NUS. Hopefully, these tasks will be completed in the near future, so that he can once again devote some amount of time to the PLM.

VIII. Installation of Network Software in Network Library.

We are proceeding slowly with the installation of various pieces of network software in the new Network Library. It is our ambition that without a great effort on our part we can have all of the appropriate network software installed in this Library by the time that a version of the network software must be shipped to RADC, for operation of their new network interface.

IX. New Project MAC Mail Service.

For the benefit of those who receive network mail on Multics, and who do not regularly log into the system, we have provided a new service to insure prompt delivery of mail. An absentee job run daily dprints mail for those desiring the service, and has it delivered to Project MAC with a special header which causes the mail room to distribute the listing with the daily mail.

X. Technology Transfer for the ARPA Office.

Planning has begun for an all-day demonstration of Multics and of the ARPANET software for a group of ARPA customers within the defense Department. The demonstration, to be held in ARPA's Washington office, is being coordinated by Ken Pogran, with Jerry Whitmore of HISI.

XI. Study of the Burroughs Operating System.

Ben Williams continues his study of the Burroughs operating system in order to learn whether any of their procedures for recovering from or continuing to run in spite of errors are applicable to the Multics system. His current intent is to produce a document which is a parallel description of various mechanisms in the two systems. This document will probably be produced at the end of the current semester. The Burroughs Corporation is cooperating in this technology transfer, as they expect the report to be of assistance to them as well.

C. Miscellaneous Tasks.

I. Creation of System Load Generator Using the ARPA Network.

All test load generators which have existed in the past for Multics have suffered from one of two drawbacks, either they use absentee jobs to generate their load, which does not exercise the I/O system, or they require a large number of data sets, which is very expensive for the load generator. To circumvent these problems, a UROP project has been organized to develop a load generator which drives Multics through the ARPA Network. This project, being organized by Dave Reed, Harry Forsdick and Ken Pogran, is just commencing now.

II. Use of TERMINET 1200 and ODEC as Multics Printers.

Jeff Goldberg is exploring what must be done to the current Multics printer daemon software in order to make various printers accessible through the standard dprint command. Since these devices connect to Multics in various ways not initially envisioned, such as through the ARPA network, some modification to the I/O driver and I/O coordinator software may be required. Hopefully, this project will be completed in a month.