Saltzer

TO: J. C. R. Licklider

FROM: F. J. Corbató, J. H. Saltzer FJC. JHS.

DATE: November 18, 1969

SUBJECT: The Research Program of the Computer Systems Research Group

This memo outlines the research plans of the Computer System Research Group through the period to July 31, 1970 along with discussion of how these plans develop after that. For several years now the Computer System Research Group has been almost totally engaged in the development of the Multics system. Since this system has been introduced for general use at M.I.T., it is expected to rapidly mature and be consolidated during the next academic year. As a result the research activities of the group will undergo strong reorientation and transformation. This reorientation can be grouped into three main themes.

The first theme is the use of the Multics system as a living experiment in the efforts to discover what are effective means of providing the service of a computer utility to a sophisticated community. Not only are issues like cost/performance, system reliability, and system availability of paramount concern to users, but it is also important to determine what aspects of software functions and facilities are desireable from the user point of view. Efforts will be made to improve the system in these areas as well as to study the interactions with the user community. The impact of sharing and privacy on a community relations is one of the key points of concern. In addition, members of the group will act as consultants to software subsystem designers in other groups which are developing software packages and designs which will enhance the overall system. In this vein also the group will propose to have similar relationships with Bell Telephone Laboratories and the Rome Air Force Development Center because they have different classes of users than those at M.I.T. A significant part of the living experiment is the determination of what degree of maintainability and operational simplicity is necessary in order to transfer system maintenance and operation to another organization such as the Information Processing Center or RADC. The transference of this basic set of responsibilities to IPC is being targeted for July of 1970. (As is discussed later, the Computer System Research group will retain responsibility for future development work.)

The second major theme consists of developing a deeper understanding of the Multics system and all of the issues which it introduces. This effort in turn divides into two categories, namely increasing the understanding within the group and subsequently transmitting this knowledge to others in the field. It is expected that the principle thrust of increased understanding will be done in conjunction with an increasing number of students rather than exclusively with sponsored research staff. (Part of this system understanding, of course, will still require a sponsored research staff since a great deal of it will come from the simplification and improvement of the Multics System, which is in daily operation.) In general, it is expected that much of the insight, understanding, and techniques which have been developed within Multics will be in turn introduced into the undergraduate curriculum (e.g., 6.233) as well as in graduate seminars, etc. In addition, there will be effort made to develop an abstracted conceptual model of the system, with quantitative parameters, such that is possible to predict the impact of major system changes such as faster core memories, increased disk channel capacity, etc. This aspect of the work will lead directly to papers in conferences and research journals as well as publication of the material in the Multics System Programmers' Manual, the Multics Programmers' Manual, and the System listings.

The third major theme of the computer system research group will be in the augmentation and improvement of the Multics system by exploiting its ability to evolve in many directions without starting over. In particular, there are many new features which should be added to significantly enlarge the convenience and range of applications which are possible within Multics. One such feature is the homogeneous embedding in the system of common graphical devices such as ARDS terminals, PDP-8 computers with graphical terminals, etc. A second feature will be insuring that communication with the ARPA Network is straightforward and without any essential difficulty. Finally, there is a large range of research problems uncovered by the implementation on which work will proceed subject to the availability of students and resources. Significant among these research topics are:

1. Ability to reconfigure dynamically. (Currently this operation takes a minimum of several minutes and requires the user to log in again.)

2. Ability to gracefully grow secondary storage by a factor of 50-100. Currently used techniques do not scale up as well as desirable. Observed device reliability is not compatible with current techniques extended to more than 200-300 million words of storage.

3. User interface to storage management subsystem, and to the storage hierarchy. One can at present specify anything desired, but frequently with more trouble than would seem intrinsically necessary. Search mechanisms seem ad hoc. Entry points are handled as a special case.

4. Ability to have many small files cheaply. Requires putting 2 or more files on one physical storage track or whatever. So far, Multics has three special purpose mechanisms which accomplish the function, but no general view of the situation.

5. The complexity of the protection mechanism seems too high. This problem might be sorted out by a better software/hardware interface, although one must be careful not to merely push the complexity into the hardware.

6. Proposed hierarchical organization of processes should be explored to see if it helps simplify Multics organization.

7. Need to find out how to support a process more cheaply, or understand why it seems to be so expensive now. Currently it is not practical to give every user 3 processes (a step which would simplify organization significantly) because it requires 3 address spaces, which in turn require much gearing up to maintain.

8. Need to explore carefully strategies for using the GIOC, to maximize its effectiveness with as little expensive core memory as possible. Also need to learn enough to precisely specify how GIOC-type organization should evolve in order to take advantage of its intrinsic economy of mechanism.

9. In general, the entire system needs to be brought up to "publication qaulity" both for publication and maintenance. More insight is needed into why it is so difficult to automatically maintain this quality during system construction.

10. Should explore and predict the effect of replacement or augmentation of the swapping drum with bulk core memory. Should also explore the potentialities of the "buffer" memory concept when applied to a large time-sharing system.

11. Need to more smoothly integrate the I/O system into the process environment. At present, it is not clear how two parallel, cooperating processes would interface into I/O system; problems of how I/O stream attachments are propagated among processes and handled following intentional interruptions (quits) by the user or long periods of inactivity such as between an automatic logout and subsequent restarting of the process.

12. Need to study addressing mechanisms to evaluate cost of mixed single- and double-word strategies for indirect addressing, and to inquire if 645 has appropriate number of base registers, index registers, associative registers, etc.

13. Need a much cleaner view of fault handling and the way in which it interfaces with the user (e.g., via the PL/I condition/signal features.)

14. Need a more systematic view of the way interrupt signals are propagated through the system. Current hardware techniques for this seem to constrain reconfiguration ability, and hardware fanout techniques do not seem to be useful for software at present. A modest redesign of this area might provide substantial economies both in hardware and software.

15. Need a more systematic view of redundancy, error detection and retry and response to unfixable errors discovered within the supervisor. Software techniques are at a par with typically used hardware techniques, but not with the most sophisticated hardware techniques available.

16. Should explore potential consequences of making segments which are never shared into a special case. This strategy might reduce somewhat the size of the tables which must be resident in core memory.

17. Need to generally review why each item currently resident in core memory is actually there. Some carefully engineered strategies might reduce the minimum requirement for core size.

18. Should workout strategies for file migration across a range of types and speeds of secondary storage.

19. Need to develop a more systematic view of computer system fabrication with respect to modeling tools, higher level languages, and software control of the hardware manufacturing process.

INTERDEPARTMENTAL

MASSACHUSETTS INSTITUTE OF TECHNOLOGY *CAMBRIDGE, MASS. 02139*

from the office of

November 13, 1969

MEMORANDUM

To: Multics Administrative Distribution

From: Jerome H. Saltzer

Subject: Telecommunication Planning for Multics

It has become apparent recently that Multics' current performance is straining the capacity of its telecommunications facilities. When twenty or more users are on the system, the probability of busy signals on dialing 8 is quite high. Several busy signal reports are filed each day with the 645 operator. These busy signals have two origins:

1. The 8-level access trunks are very poorly "graded" at the telephone exchange. As a result, busy signals are encountered by some callers when as few as 15 of the 32 8-level trunks are in use. At least 4 trunks are only accessible from places which have never originated a call to Multics (e.g., Lincoln Lab.).

2. Since Model 37 teletypes are not widely available, at peak usage times as many as 25 or 26 1050 or 2741 terminals are dialed in. This is essentially the maximum that can be expected, no matter how well the telephone exchange is graded.

There are at least three imminent developments which will probably result in attempts to dial up larger numbers of users:

1. The two processor, 384K word memory system seems to be close to ready for production use. The only performance measurement made so far is that this system can create 30 processes (part of the standard benchmark) in about 2/3 the time of the 1 cpu, 250K system, a remarkable result considering that process creation should generate much interference for supervisor data bases. It would appear conservative to estimate that the number of dialed lines could initially increase by 50% for this larger system.

continued

2. As the Limited Service System (BASIC) becomes a reality, the 6.231 restricted class system becomes polished, and the new accounting system permits two restricted users on in place of one regular user, we should expect more dialed-up lines. Assuming a mix of half the load coming from regular users, a 50% increase in the number of dialed-up lines will result. Note that at peak hours (e.g., after midnight on the day before problem sets are due) small users in large numbers may be almost the only customer thereby doubling the need for lines.

3. Several performance improvements are scheduled or planned, including a speeded up page fault handler, a pre-paging, post-purging strategy, a scheduler with shorter time allotments at high priority and longer safety from pre-emption at low priority, DSU-270 Disks, use of PL/I to produce smaller supervisor code, and a new typewriter control module. These modifications together probably represent a 50% improvement in number of lines, even after swallowing some of the improvements in the form of better response.

Taking these three developments together, we should expect a factor of three in the potential number of dialed-in users, all within the next 6 to 12 months. Thus the need for telephone planning.

The Present Configuration

In review, the present service machine line configuration is as follows:

```
32 lines on dataswitch level 8, for 1050/2741 use
14 lines on dataswitch level 7, for M37 use
 8 lines on GE PBX, for M37 use
 2 lines on individual dataswitch numbers for operation use.
56   Total of lines giving access to Multics.
```

In addition there are 32 individual dataswitch numbers used for the development machine and as spares, and six "high-speed" ARDS access lines, currently not connected.

The large gap between the number of users presently able to use Multics and the number of access ports comes from four sources:

1. Breakage between 1050/2741 and M37 ports; one must provide capacity for peaks of usage from either source (cost $\sim$4 ports).

continued

2. Breakage between GE PBX and MIT dataswitch; again, one must provide for reasonable usage from either source. (cost ~8 ports)

3. Overcapacity for M37 Teletypes. There are more Multics M37 ports than there are installed M37 teletypes at GE and MIT together. (excess ~8 ports)

4. Ideal grading on 1050/2741 ports would allow probable usage of no more than 80% of them at any one time. (excess ~6 ports)

Short Term Planning

It would seam wise to presume that the present flood of busy signals is a warning, and that some of the effects mentioned earlier will begin to be felt very soon. The following plan of action is recommended as soon as possible:

1. Recompile the present teletype control module and its tables to accept dial ups from 80 ports, rather than 56.

2. Request that the telephone company move 4 level-7 trunks and 12 individual numbered lines over to level 8 as soon as possible, raising the number of lines on level 8 to 48. This step should include careful review of the grading arrangements to insure that the present busy-signal troubles are not repeated.

3. Order from General Electric Company three additional teletype channel groups (24 channels) for use with 1050/2741 terminals. These will be used for the additional 16 level-8 lines, and for 8 more individually-numbered lines, for users who encounter busy signals when dialing 8.

In summary, this will leave us with:

48 lines on dataswitch level 8 for 1050/2741s
10 lines on dataswitch level 7 for M37s
8 lines on GE PBX, for M37s
2 lines on individual dataswitch numbers for operations (M37)
8 lines on individual dataswitch numbers for users who have trouble dialing to level 8.
76

and the 6 ARDS lines.

continued

We may estimate that up to 45 1050/2741 terminals can obtain access in this configuration, in addition to 12 MIT Model 37's and 8 GE Model 37's. This configuration may hold long enough to do some longer range planning. (It also leaves 12 individual lines for the development machine, an adequate number).

Longer-range Planning

Only the outlines of planning inputs are proposed here:

1. We must get more accurate prediction of arrival dates and performance effect of the three user-increasing effects mentioned above. A presumption of capacity for 100 dialed-up lines (at times) 12 months from now is a likely result.

2. We must assess the future of the Model 37 Teletype terminal at MIT to determine the ratio of M37 to 1050/2741 ports to request from the telephone company. The simplest arrangement would be if Model 37 use remains static, in which case we merely expand level 8 so that at 80% usage it will handle all additional lines.

3. We must assess demand for lines. For several reasons, it would seem best to presume demand equal to capacity.

   a. There seems to be no point in assuming demand greater than capacity.

   b. If one assumes demand below capacity and has guessed wrong, there may be a several month lead time to add more lines, during which time expensive 645 capacity is forced to be idle by the lack of access.

   c. It is of considerable value to be able to talk about very large numbers of (even very light) users as an accomplished fact in convincing potential customers of Multics usefulness. We could not demonstrate today 60 BASIC users, simply because the access ports are not available. Classes such as 6.231 can provide this kind of demonstration, but only if the access ports are there.