

Management Plan for Project MAC
with Special Reference to Multics

> Introduction

At the beginning of 1969, Project MAC is at a critical juncture. During the last two or three years, more and more of the resources of Project MAC have gone into the Multics effort. During 1968, Multics used \$2.1 million of the \$3.0 million of Project MAC's general (Task 1) funds and of the \$3.7 million total.* At the end of 1968, the Multics operating system was beginning to function well enough that some Multics system programming could be carried out within the Multics system, but the bulk of the Multics programming was still being done on CTSS. According to an estimate made by the Multics Group, a convincing operational demonstration of Multics could be made sometime between the middle and the end of 1969. If Multics were opened to general use during 1969, perhaps two years of user-oriented support programming (problem-oriented languages, program packages, etc.) would be required to bring the Multics software base up to a level rivalling that of CTSS. Thus, as seen from one point of view, Multics seems to be late, still in the mid-course of development, and so expensive as to have seriously distorted the research effort of Project MAC.

*The G.E. 645 computer system is costing about \$1.0 per year, rental.

On the other hand, the technical goals of Multics are widely recognized to be extremely important. The recent ARPA Review Committee appreciated the goals and attested to the importance of the Multics software design and the Multics operating-system software. The Review Committee obviously considered the Multics software to be very valuable -- so valuable that it was willing to consider the Multics effort already a success -- so valuable that the Review Committee turned its attention during the second day of its visit from the problems of development of Multics to the problems of "exploitation," i.e., of making Multics available and useful to the world outside M.I.T.

Although the formal report of the Review Committee is not yet available, the review has had a major effect. It has lifted the attention of the Multics Group and the administration of Project MAC up from the problems and difficulties of carrying forward the development of Multics, day by day and week by week, to broader issues -- especially to the question of making Multics come into being in such a form as to ensure to it a wide availability and the advantages of modern technology, high performance, and low cost. In short, the Review Committee stretched the time scale of optimization of Multics from the few weeks before the next test point to the few years before Multics will be used by a hundred organizations. At the same time, the Review Committee

called attention to the imbalance in the over-all research program of Project MAC caused by putting such a large fraction of the over-all funding into the computer system research (Multics) effort.

There may or may not appear to be an inconsistency in the preceding trio of sentences: Ensure a useful long-term future to a very valuable software development -- and economize in the short term. Whether or not those courses appear to be -- or are in fact -- inconsistent depends upon several factors that will be discussed later.

Guidelines and Constraints

Because Multics is at present such a large part of Project MAC, planning for Multics implies planning for MAC. That planning should be guided by a recognition of certain guidelines and constraints and by a wariness of being misguided or inappropriately inhibited by others. What are the valid guidelines and constraints?

1. A main objective^T is to make a "Multics"
 - a. Widely available
 - b. With a high ratio of performance to cost
(which implies advanced technology and/or
low pricing)
 - c. Soon.
2. The essential thing is to achieve a "desirable subset" of the Multics goals.* We should define

*For present purposes, the "desirable subset" will define "Multics."

"desirable subset." It is not essential ^{or} on a priori grounds to implement all 30 of the features in Corbato's list, but it may be (or: and it may not be) easier/quicker to implement them all than to implement only a subset.

3. The GE 645 computer does not satisfy item 1.6. We cannot continue to use the 645, under the present charging schedule, past the middle of 1969.*

4. General Electric might develop a computer (Let us call it the "655.") that would satisfy item 1. It would make sense to continue with the 645 and to order a 655 follow-on if:

- a. G.E. assured/convincing** M.I.T. and ARPA (and BTL if BTL is interested) that G.E. will develop such a machine with advanced technology and that it would make it widely available at an early date (to be specified) and at a price consistent with a low level of G.E.-prepared software support -- i.e., with the fact that much of the software support would have been paid for by ARPA.

*Unless ARPA provides additional money for that purpose.

**We should specify what assurance would convince us.

1/3/69

b. For the interim, G.E. reduced the cost to M.I.T. of the 645 computer -- say, to one-half its present rental cost or to a direct sale cost of a million dollars plus a yearly maintenance contract cost of \$100,000. (Those figures are for discussion.)

*Not
nearly
needed*

5. We are in the process of approaching G.E. to discuss the foregoing points. We should approach other computer manufacturers in the near future -- but not until discussing our management plan with ARPA. We should make a list of manufacturers to approach: DEC, Univac, IBM, Burroughs, SDS,

What we need from each/some manufacturer is an assurance, convincing to M.I.T. and ARPA, of advanced technology and performance, wide availability, low price, good maintenance, and early delivery. We must decide what those attributes mean in terms of specific criterion values.

6. To cover the likely possibility that no manufacturer will meet the criteria, we should investigate the use (with modification if/as necessary) of an existing computer, other than the 645, for further development of Multics. We understand that following this course at M.I.T. (beyond just investigating it) would probably mean losing not only "momentum" but a large part of the

Multics team. We should therefore assess the chances that another contractor could carry through to success along this course. We should also determine whether we would, if Multics terminated, want to undertake a new project aimed at realizing a limited ("Subdesirable?") subset of the Multics objectives on a low-cost computer with assured performance and availability.

7. If we have a consensus, we should develop a statement on the subject of "code compatability."

8. We should develop a good projection of Project MAC's need for computer access during the next two or three years. This projection should take into account the conditionabilities of Multics. It should also take into account the plans of others to use CTSS and other available computers.

9. In connection with item 8, we should examine the availability and usefulness to others in Project MAC of the PDP-6/10 system operated by the Artificial Intelligence/Intelligent Automata group. The PDP-10 should be gotten into a useful state as soon as possible, and full advantage should be taken by Project MAC of the PDP-6/10 resource. A plan should be made for connecting the PDP-10 into the currently operating PDP-6 time-sharing system and/or for connecting it, if the new PDP-/10 materializes, into a new PDP-10 system that we have

proposed to ARPA for development as a special facility for interactive dynamic modeling.*

10. The over-all cost of computer support should not exceed 1/3 of the total budget of Project MAC. In figuring the over-all cost, purchased computers will be amortized over over 60 months. It is fair for the System Research Group to use more than a pro-rate^a share of computer support, but see item 11.

11. Unless ARPA decides to provide special funds for continuation of the Multics effort, the Computer Systems Research Group will have to retrench gradually to a level at which its total annual cost, including computer support, is not greater than \$1 million. That figure is conditional upon ARPA's continuing to fund the main task of Project MAC at \$3 million per year.

12. A plan will be made (and incorporated into the over-all management plan) to intensify research in other areas in order to round out a balanced program of research in Project MAC.

13. A proposal should be prepared to continue Multics at a high rate and to achieve the objectives specified in item 1. One version of this proposal should be

*This facility would be made available to the behavioral-science group that is proposing to use CTSS extensively in data-base work.

prepared for ARPA. Another should be prepared for G.E. We should decide whether other sources are likely enough to be worth approaching.

14. The "M.I.T. community" of computer users, including users outside Project MAC, is vital to our having a success with Multics comparable to the success we had with CTSS. The needs of the community now include "exportability."

Project MAC's plans that involve the community should ensure that ^{its} their exportability needs will be met.

15. Project MAC's computer plans and operations almost necessarily interact with other computer plans and operations at M.I.T. Our management plan therefore has to have the approval of the Provost and, if he so indicates, of the Information Processing Board.

16. The Multics concepts and software must be organized and documented and/or published.

Alternative Courses of Action

Several different courses of action are in keeping with the foregoing guidelines and constraints. Some of them are:

G.E. All the Way with ARPA Support. We approach G.E. now (we are doing so.) re items 1, 3, and 4. G.E. agrees by March 1 to meet the conditions. The rental cost of the 645 comes down to \$500,000 per year. The use of CTSS declines to zero. The other costs of the Computer System Research Group come down to \$500,000. The overly intense pressure of recent months decreases to a level of mere urgency. G.E.'s assurances meet (to a sufficient extent) the needs of members of the M.I.T. community for exportability. Development of Multics continues on the G.E. 645 for the G.E. 655.

G.E. All the Way with G.E. Support. Assume the beginning of the foregoing scenario except that G.E.'s assurances, while satisfactory to M.I.T., are not satisfactory to ARPA -- which may be unlikely. Assume further that G.E. is convinced of the value of Multics and is willing to finance continuation of the Multics effort. ARPA cuts support of Project MAC to about \$2 million, eliminating Multics but permitting Project MAC to develop its program in other areas. G.E. takes on support of Multics. Then there are two alternative courses.

On one, the Computer System Research Group stays at approximately present size and Multics develops in such a way as to involve the M.I.T. community. On the other, most of the Multics effort moves into G.E. while the principles of the Computer System Research Group concern themselves with organization and publication of the Multics concepts.

Switch to Another Computer and Develop Multics for it on the G.E. 645. Assume beginning of first scenario except another manufacturer's computer is selected. Assume ARPA makes additional funds available so that G.E. 645 may be retained during interim. Then there are two or more courses (how many?), depending upon degree of "code compatibility." If code compatibility is essentially complete, the Multics effort continues without much change -- other than to adjust to the cost guideline -- and the software is transferred when the new machine is ready. If the new machine is not code-compatible with the 645, work has to be undertaken to develop an EPL translator (compilers) for the new machine, and someone has to worry about conversion of non-EPL code.

A New Computer But No 645. As I understand it, this is viewed by the Multics Group not as a continuation of the present effort with a sidestep, but as a new project. The 645 leaves -- three months, I think it is, after we give notice to G.E. Some of the members of the Computer Systems Research Group leave,

1/3/69

discouraged or at any rate no longer interested. Presumably G.E. and B.T.L. do not join the new effort, but let us not assume (without finding out) that they do not. Either we obtain a Multics-adapted computer from a manufacturer or we obtain a Multics-adaptable computer computer and adapt it, ourselves or with help from Lincoln or a subcontractor. If possible, we make arrangements for a computer -- same type except not Multics-adapted -- for interim use in Project MAC, and we program the new system on it. If the new machine is not 645-code-compatible, we introduce as much simplification as possible into the design of the Multics software and begin the programming more or less from scratch -- wiser but without any/much transferable code. If the new machine is 645-code-compatible, then it is such a bad thing, not to have the 645 that we insist on keeping the 645, which nullifies the defining assumption of this scenario, or we go to B.T.L. to use a 645 there, or (with nothing to program on) a few key people use the interim to organize, document, and publish and to make any changes in design that seem very important on the basis of experience. Under no interpretation does this latter version, based on a 645-code-compatible follow-on but no interim 645, make much sense. The former version -- for example, modifying a PDP-10 and in the interim programming on an unmodified

PDP-10 -- was, however, regarded as reasonable by some member of the Review Committee, and we should have a clear position vis-a-vis it.