January 6, 1968

D R A F T

MANAGEMENT PLAN FOR PROJECT MAC WITH SPECIAL REFERENCE TO MULTICS

I. INTRODUCTION

During the year 1969, it is anticipated that what has recently been the largest single component of Project MAC research, namely the Multics computer system research project, will reach some significant turning points. These turning points arise from several sources -- the time that it has taken to accomplish the research, budgetary pressures, and most significantly, pending accomplishment of the basic research objectives and the resulting tapering-off of this activity and opportunity to expand and initiate other activities. The plan described here has been constructed in order to be prepared for these turning points, and to meet several possible contingencies.

During the course of the past three years, the investment in computer system research by project MAC has climbed from about 1/3 of the yearly budget to about 2/3 of the 1968 Task 1 budget. This investment is expected to drop drastically over the next two years as Multics becomes a tool used by Project MAC and the rest of M.I.T., rather than a vehicle solely for computer system research. Since these are significant portions and changes in the apportionment of the total MAC budget, any discussion of plans for MAC must include specific mention of plans for Multics and the computer system research group. Thus this plan is divided into two primary sections: The plan for Multics and computer system research, and the plan for the remainder of MAC research.

II. MULTICS AND COMPUTER SYSTEM RESEARCH

In review, Multics is a Multiplexed Information and Computing Service; in short, a "Computer Utility." The intent of the computer system research group has been to explore the consequences of implementing a comprehensive computer utility by building a prototype which is to serve as a cornerstone facility for the Project MAC community to build upon, as that community itself further

DRAFT

explores issues of interactive computer usage, networking, data bank manipulation and experiments in coupling man and machine. The 7094 compatible time-sharing system (CTSS), also developed by the computer system research group, has served exactly this function for the past four years. The limitations of CTSS (primarily price and availability, as well as lack of expandability, generality, and ease of experimentation with new ideas) and the ways in which Multics overcomes these limitations are well documented elsewhere.

By the end of 1968, Multics was beginning to function well enough that some Multics system programming could be carried out within Multics itself; the system began daily operation in September, 1968. According to estimates made by the computer system research group, by June 1969, there will be performance and reliability sufficient to base a summer workshop on Multics, similar to the MAC summer study of 1963 which used CTSS as a vehicle.

Friendly users can be accomodated well before that date, probably in early spring of 1969. Thus among the turning points which are at hand is the ability to allow users from outside the computer system research group access to the Multics system.

In December, 1968, ARPA appointed a panel of computer system specialists from other organizations to review the Multics project in detail, with an eye to evaluation of the worth of the technical objectives and of the methods being used to reach for these objectives. Although the formal report of this panel is not yet available, discussions with the panel indicate that it fully agrees with, appreciates, and supports the goals of the Multics project, and

D R A F T

that further it feels that the specific operating system implementation is technically sound, and valuable. In fact, there seemed to be a consensus of the panel members that they are willing to consider Multics already a success. This view, while academically very interesting, may be considered premature by other workers in the field who, not being privileged (or having time) to explore the inner workings of the system prefer to listen to acclaims of enthusiastic system users in conferring success.

The panel went considerably further in raising several other issues, which are addressed in this plan, and in making a number of suggestions. The specific suggestions of the panel have been accepted as inputs to our thinking, inputs which indicate to us the nature of the problems the panel and ARPA see. Although the plan, as will be seen, does not necessarily follow to the letter all of the panel's specific suggestions, it does address the problems which appeared, in our best judgement, to be raised by the panel. This plan also takes into account interactions with the rest of Project Mac, and the M.I.T. community as a whole, as well as with the General Electric Company and Bell Telephone Laboratories, who have cooperated in the development of Multics.

The immediately following points of this section discuss the basic plan of action for completing Multics and putting it into operation, as well as the assumptions which underly that plan; a set of contingency plans which take into account as best as can be discerned possible changes in the underlying assumptions and describes their implications on timing, cost, and probability of success; and finally a discussion of plans to meet a specific objective mentioned by the panel: that the investment in development of Multics have some fruits beyond usage of Multics at M.I.T.

D R A F T

Basic Plan of Action

In order to comprehend the basic plan of action it is necessary to envision the goals for harvesting the fruit of Multics research. These goals can be described at four levels, with an approximate time scale for each:

1. Starting in 1969, Multics becomes the primary community computer facility at M.I.T. Project MAC.
2. Over a period of three to five years from now a dozen or more installations around the United States are using Multics or its direct successor.
3. In the period from five to ten years, one hundred or more installations are using Multics-like systems including imitations and evolutions by others.
4. After 10 years a majority of large-scale computer installations in the world have facilities which can be traced directly back to the Multics design.

This set of goals is ambitious, although compatible with the potential significance of Multics as expressed by the ARPA Technical panel. An important aspect of the plan of action described here is how to achieve them within realistic constraints of time, budget, and available personnel, making maximum use of the available resources, including resources of the General Electric Company and Bell Telephone Laboratories.

D R A F T

Two fundamental assumptions underly this basic plan. A later section, entitled "contingencies", indicates alternate (and by implication, less desirable) plans which could be followed if either of these two assumptions prove unrealizable.

These two cornerstone assumptions are:

1. The General Electric Company, on a time scale acceptable to all concerned, commits itself to what may be termed "follow-on" support for the 645 computer. That is, GE commits itself to develop and provide a successor hardware system, using modern technology, with competitive prices, and which Multics can be evolved to operate on. The details of what is an acceptable time scale of delivery, performance, cost, and specification of "modern" remain to be proposed and negotiated, but one important constraint is that the hardware system be available to anyone who wishes to purchase it.

2. The ARPA will permit M.I.T. Project MAC the flexibility in its internal budget allocation to provide the necessary support to "get the current system out the door", so to speak, and into productive operation in which it can begin to pay its own way.

As implied by the words "productive operation", one of the immediate points of the plan is to begin to develop a user population as soon as possible. There are at least four important aspects to this part of the plan:

D R A F T

1. It will provide early feedback on the usefulness of certain facilities.
2. It will provide the reassurance that Multics has the necessary performance, reliability, and functional capabilities.
3. It will help to minimize development and underwriting costs since convinced users will be willing to purchase machine resources.
4. A satisfied user population is the most effective technique of convincing the world of success.

There are three immediately apparent sources of users for this facility. Tapping all of these sources does require commitments from the people involved and their own communities, and, of course, their commitments will partly depend on their view of the future availability of Multics, both at M.I.T. and elsewhere. These three sources are:

1. The Project MAC community, replacing current use of CTSS and providing a facility to use as other activities are expanded.
2. M.I.T. as a whole, drawing to start upon current users of the remaining 7094 CTSS, and later upon the exhibited long-range appetite for computation of the M.I.T. community.
3. The Cambridge and Boston academic community as a whole. In particular, there are projects at Harvard University and elsewhere which would make use of an advanced interactive facility. Frequent requests for use of CTSS, which have had to be turned down in the past are an indication of the depth of this potential market.

We now turn to a specific time-table and budget proposal. The chart of figure 1 suggests approximately the budget apportionments implied by this time-table. The following specific plans are proposed:

1. Development of Multics continues at its present pace through 7/1/69, when a tapering-off operation begins, coincident with completion of enough facilate to gracefully absorb paying customers.
2. In early Spring, 1969 " friendly " users (Project MAC or M.I.T. people willing to put up with problems in return for a chance to get an early start) begin to develop their own subsystems using Multics as a tool.
3. The use of CTSS as a development tool by the computer system research group will taper from its present level to zero by 7/1/69, thus freeing up CTSS time and also hardware budget commitments.
4. A summer workshop is proposed, similar in spirit to the Project MAC summer study of 1963. The purpose of this workshop is to bring visitors from around the country to see and try Multics, and to explore possible directions of exploitation of the system.
5. As of 7/1/69, the door is opened to paying users from Project MAC and the M.I.T. community. (There are budget issues of underwriting unused resources and of underwriting "introductions" to the system for skeptical users, which are not addressed here.)
6. By January, 1970, the Computer System Research group's usage of the 645 system has dropped to about 25% of the available resource.

DRAFT

This drop is possible for three reasons:

a. Multics performance is substantially improved over today's version.

b. The pace of change to the system can be slowed down as the pressure of getting a system which can support paying users is removed.

c. As the system research group staff moves onto Multics, from CTSS it can begin using much more finely tuned and simpler system building techniques, abandoning the present brute force approaches.

7. As of January, 1970, the Computer System Research Group staff budget is also reduced, as the staff size is slowly tapered down, consistent with reductions in programming activity. The remaining staff of the group will pursue the following lines:

a. Maintenance and debugging support for the operational system, and user consultation.

b. Building up and fleshing out of Multics features, languages, and primitives.

c. Refinement of the 645 Multics design for improved simplicity, understandability and performance.

d. Writing of papers and books on Multics organization, construction, usage, and specifications. (This is part of a more general publication plan, described below in the section on exportability.)

e. Planning of the hardware changes appropriate in the computer system which is to follow the GE 645; and getting ready any corresponding software changes. Since the future course of Multics is expected to be evolutionary, rather than complete rebuilding, it is not important that moving to the successor of the 645 will require an out-of-line programming investment. We are attempting to play to one of Multics' strong points here, the ability to evolve to meet change.

f. Interfacing with visitors who wish to learn about Multics. Such visits can include both short briefing sessions and long range "internships."

It is important to realize that this timing indicated above is estimated, not absolute, and it is optimistic in the sense that the schedule suggested can be met if no unexpected hazards arise. History suggests that a safety factor should be included in planning. As best as can be foreseen today, the date of 7/1/69 suggested above could slip as much as three months, to 10/1/69; similarily, the date of January 1970 could slip as far as six months, to July 1970. Again, these are felt to be the maximum safety factor needed to guarantee success.

Possible budget projections for Multics/MAC 1/3/69

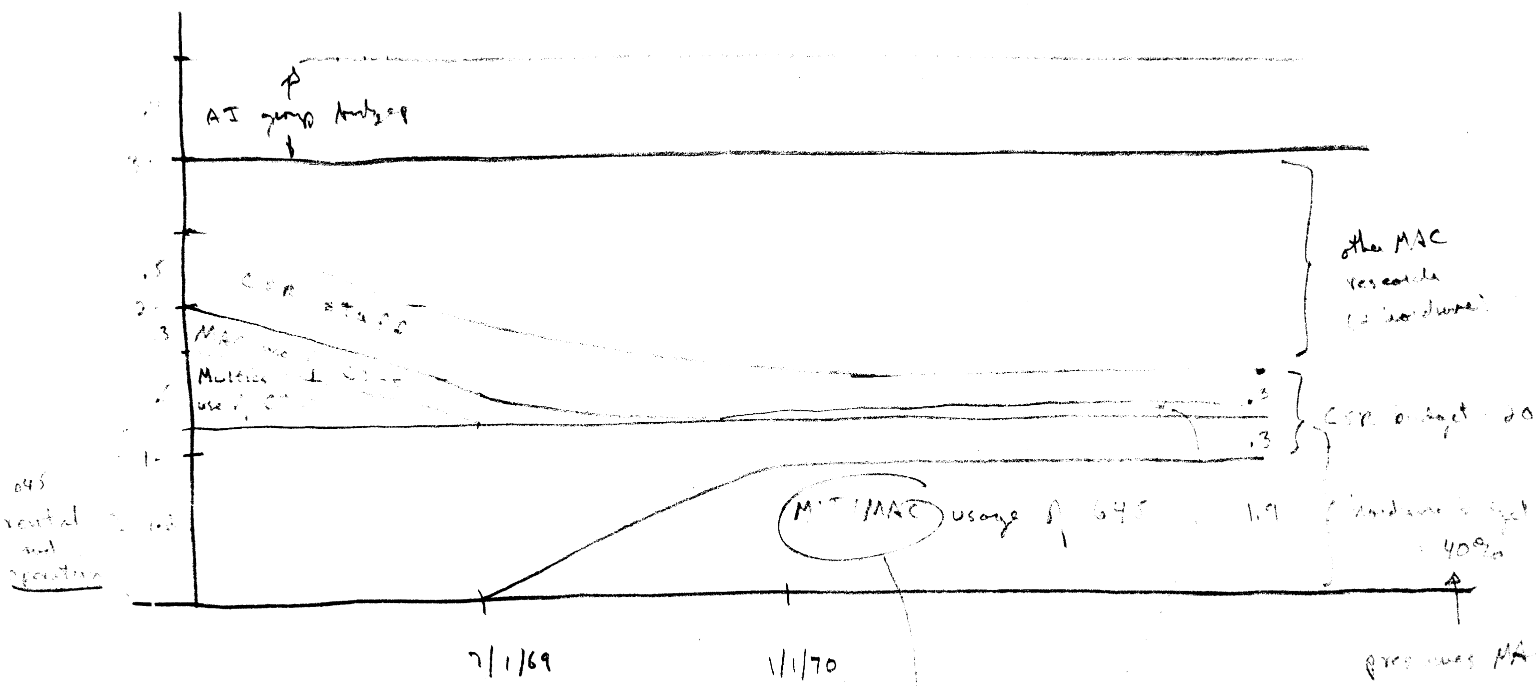

since 645 capacity

Flexibility

CSR budget as function of MAC

67% 20%

50%

* specialist support of Multics

What percentage is non-MAC?

Contingency Plans

The basic plan outlined above is built on two assumptions, and if either of these assumptions ~~could~~ cannot be realized for any reason, the course of Multics research will be strongly affected. To discuss any contingency plan sensibly, it should be viewed as a salvaging operation; an important question is: what is being salvaged? There are four general kinds of things being salvaged:

1. A dollar investment: A back-of-the-envelope ~~given~~ guess is that ARPA, through Project MAC, has invested about $6-8 million, Bell Laboratories has invested about $4-5 million, and General Electric has invested something more than $10 million in hardware and software support of Multics.

2. A man-year investment of perhaps an average of 50 persons working for four years pulling together an ext~~ensively~~remely complicated set of ideas. This intellectual investment has begun to be transferred to paper in a form usable by others, but much more transfer is still possible.

3. A belief in the objectives of Multics and an assumption that a system accomplishing these objectives is potentially a major national resource.

4. The elimination of a wide-spread belief among skeptics that systems with all the objectives of Multics cannot be made to work. Living, working proof of this ~~work~~feasability will encourage others to ~~initiate.~~ imitate. On the other hand, lack of this proof could easily deter other projects from being started.

The final point above emphasize a belief stated earlier that it is extremely difficult to sell ideas in the computer field without a working model in use by a pilot population. This belief, as will be seen, strongly affects the contingency plans proposed here.

Before launching into the various alternatives, one final background issue must be clearly recognized: the psychological climate in which Multics is currently couched strongly affects what can be accomplished, since the success of the project so far has rested strongly on the willingness of capable people to work very hard toward ideas they believe in:

1. The M.I.T. programming staff has for several years been working with a series of peak efforts of extraordinary caliber; this staff wants very much to see Multics in operation, serving users. This team has never viewed Multics as a 9-to-5 project, but as a living, important issue worth personal sacrifice. As a result, hesitation in support transmitted to this staff can have a deadly effect.

2. The M.I.T. Electrical Engineering Department is probably reaching the limit of its patience over the intensity of faculty support of Multics. A very tight shortage of faculty in computer science, coupled with the need for their full-time attention to Multics at the present pace leads to this concern. The project cannot be allowed to stretch out indefinitely.

3. For Project MAC as a whole, there is the problem that as CTSS disapp~~roves~~ears, the project no longer has a unifying vehicle; a clear picture of the future can be difficult to project.

4. Potential M.I.T. users, faced with uncertainty as to the future of Multics are ~~virtually~~ naturally shy to enter into any commitments on that system, thereby contributing to the uncertainty of its future.

5. The General Electric Company, and Bell Telephone Laboratories will notice if M.I.T., ARPA, and MAC ~~are~~ do not see a role for Multics, and will ask why they should see one also.

This list of psychological issues is not intended to frighten, but merely to point out that there are a whole collection of factors that must be considered in contingency plans.

The primary contingency that is foreseen is that the General Electric Company may fail to provide a satisfactory reply to proposals for hardware for the 645 equipment. What constitutes a satisfactory reply has been mentioned, though not worked out in detail above. One issue on this reply, however, is how soon the reply is received. By setting an unrealistically early deadline there is danger of pressing General Electric to the wall and forcing out an unsatisfactory answer because of short notice. On the other hand, a later date makes planning ahead very difficult, and ~~un~~ will deter potential users from jumping aboard. At this time, requesting a reply date (to a question that remains to be asked formally) of March 31, 1968, seems reasonable.

What are the possible responses ~~of~~ to a negative response, or a response too late to be of use, from G.E.? There are many, but seven particular paths are expl~~ained~~ored here to provide some positive guidance as to costs and implication. The first three of these paths, in our judgement, ~~has~~ have a reasonable choice of "success" in the sense of meeting most of the goals outlined before. The fourth path~~s~~ seems to have a high possibility of failure; the fifth and sixth plans are really shut down arrangements, and the seventh is ~~classified~~ considered to demonstrate reasons for its ~~impossibility~~ infeasibility, since it might on the surface look practical or even promising. The plans are placed in order by desirability, with the one felt to be the best placed first.

1. General Electric Company proposes to pick up the hardware cost of the 645 ~~or~~ as a technique of buying time to delay its own decision. This path permits development to proceed at substantially full speed, but a deadline on the length of delay must be imposed in order to allow potential users ability to plan ahead.

2. Another manufacturer agrees to do a follow-on evolutionary version of the 645. The current 645 system then remains at Project MAC as the primary facility for MAC and M.I.T. users who are now confident of the future of the system. The plan is basically the same as before, with another manufacturer stepping in. This plan has the obvious defect that it is not clear what motivation another manufacturer would have to take on such a project, after abandonment by G.E. If the manufacturer instead requires that instead of an evalutionary follow on he can only

he can only provide a modified version of somethin hin his own line, which is not especially compatiyle with the 645, we have just jumped to plan 4, which we see below.

3. Project MAC and M.I.T. coming on with the 645 installation in spite of lack of follow ? support, and produce a "unique" system. This approach is taken on the faith that a unique system with satisfied users will not remain unique for long, as evidenced by wide current use of the GE-265 and SDS-940 time-sharing system, both of which becme widely and through demand these marketing.

4. Another manufacturer provides a "non-compatible" but modern replacement for the 645 computer, requiring extensive software remaking. This plan includes a substantial amount of non-research work, namely the extension software remaking for different hardware; such a project is more suitable for a software know than for M.I.T. On the other hand, if hardware changes are done by others, and software changes are done by others, there is considerable danger that this Multics design will be ill-implemented, since there would not be a widely understood neeed to initiate. Therefore, this plan is given little chance of success.

DRAFT

5. Graceful shutdown. In this plan, the GE645 remains on site until the system is basically complete and a comprehensive set of measurements are taken (about 9-12 months). This period is followed by a period of publication of results. This plan suffers from a high dollar cost (equal to that of continuing the project) and many of the ideas of Multics will probably never be accepted by the computer community at large. Instead, a myth of impossibility of the goals could be engendered.

6. Non-Graceful shutdown. Here, the 645 computer is unplugged immediately the staff disbanded, and writeup of results is attempted by whoever can be convinced to rerun. All of the above psychological issues come home in force on this plan, and it probably is the best way to issue vertually complete loss of the investment to date.

7. Horse-switching plan. Here, the 645 is returned immediately, and a replacement machine is ordered by M.I.T. Project MAC as a software house attempts to reimplement Multics on this new machine affecting with what ever system happents to come with that machine if any. A period of tool-holding must begin and a project, or compiler, and other system building tools are developed, first within the delivered system and later transferred into the new Multics. This plan is described in over to poit out that there is practically a guarantee of failure implicit in the corrupt in that the best way to complete Multics is to start over. Lack of confidence in the ability to deliver would undermine sections with ARPA, MAC and the world at large; even a new project out a new staff would be severely handicapped from the start by the apparent "failure" of the previous effort.

In summary, this spectron of possible plans, including prognosis ranging from wild success through failure into disillusionment represent certain possible avenues which have been explored. There are other plans which are intermediate between some of these, in many cases their inplementation may be partly drawn from the above comments.

DRAFT December 6, 1968

<u>Exportability</u>

Exportability is a word which has recently come into vogue as a succinct way of discussing the ability to communicate to others the fruits of programming efforts. We observe that this communication can, in fact, be done at at least three levels.

1. The first level of exportability is in the realm of ideas and insight. The research results which are embodied in Multics as well as the many sub-system and application program developments which are possible in a Multics environment will need to be discussed and explored in talks, articles, papers and eventually text books. It is anticipated that this process, which is already underway, will contine with increased intensity, and will result in a diversity of publications ranging over a) basic system documentation of the Multics System Programmers Manual (MSPM), b) various manuals educating users on the better ways to exploit the Multics environment, c) manuals explaining the managerial issues involved in administering and operating a computer utility, and d) papers discussing unsolved problems in ~~the~~ computer utility design. However, it is not expected that the process of communicating ideas will be fully successful unless in addition there is an attempt to have face-to-face discussion and contact with visiting technical persons. One way that this could be accomplished would be to have in summer 1969 a several week long workshop in which invited computer specialists

from throughout the country were able to familiarize and acquaint themselves with the full import and techniques involved in developing and using the Multics System.

The above sequence of steps, of course, resemble very closely those which occurred at the initiation of Project MAC when it began to exploit the Compatible Time-Sharing System (CTSS). Starting with the summer study of 1963, the influence of CTSS has been widespread; it has specifically influenced and been imitated in a variety of subsequent time-sharing systems and has consistently been used as a benchmark of comparison for other efforts. A great deal of this success we believe is explicitly due to the ability that persons have had to study the system, explore the implications, and to evaluate work of the users of CTSS. Similarly we expect that an operating Multics system would serve as a continuous laboratory for the export of ideas to the influential persons of the computing world.

2. A second level of exportability is in the ability to exchange programs and subroutines between one computer installation and another. This is straightforwardly accomplished by means of language-based exportability. Thus if it is desired to distribute a new subroutine development to someone at another computer installation and it has been written with appropriate caution using a language such as PL/I, then it is possible to recompile and use the program (with perhaps a minor amount of editing) even though a radically different computer system is in service. Even at this early period of language development, it appears definite that the Multics system will have implementations of at least the following

exportable languages:

a) Fortran IV

b) PL/I

c) AED

d) Snobol

e) BCPL

3. The third level of exportability is at the level of the application program sub-system. Here the pertinent issues regarding exportability are the precise characteristics of the host computer system environment. At one extreme the environment change is nil and exportability is trivial if both installations are using identical brands of equipment, models, configurations and operating systems. Since this is almost never true even among IBM customers there is inevitably some work attached to exporting sub-systems. This work increases as systems become more different in critical ways such as word lengths, character sets, etc. to the point where it is not possible to carry a system from one machine to another if certain features such as a file system in secondary storage or interactive time-sharing terminals are not available. Nevertheless there are a great many sub-systems such as DYNAMO, the BCPL translator and the AED translator which can exist in several computer environments and are exportable. Our expectation with the Multics system is that as the value of the features which it embodies become more widely understood and appreciated, that there will be new computer systems

which will have a Multics-like environment. Of course, for those organizations such as BTL and RADC which are attempting to be pace-setters, there will be straightforward communication, exchange and exportation of all programming work that is useful to the overall community since a principal asset of a Multics installation regardless of its configuration is that the user and the subsystem programs all see the identical environment. We do not consider it unreasonable to assume that as Multics becomes more successful that the number of Multics systems will increase as more and more organizations feel the straight-jacketing constraints of contemporary computer systems.

III. PLANNING FOR THE REST OF PROJECT MAC

(This section is almost totally incomplete. The following questions are some which could be answered in this section:

1. What are the projects

2. What is the hardware support planned for them--what machine and how much.

3. What are the individual budgets and the hardware]development/staff breakdowns of each.

4. To what extent to these projects use j CTSS
   Multics
   IPSC machines (360/67, 360/65)
   other machines

5. Plan for absorbing the budget released by the Computer System Research jgroup.

6. Is there any plan to couple task II equipment with task I work.