file
Mtg/IPC
Communications

From: Information Processing Users Committee
To: IPC Administration and the Information Processing Advisory Committee
Subject: Summary of Meetings and Recommendations
Date: May 18, 1971

The Users Commitee

A small committee of MIT computer users was formed in the fall of 1970. The purpose of this committee is to act as spokesman for serious user problems and to advise the IPC on matters which will have impact on the users of its systems. The current membership of the committee consists of

George Dixon, John Klensin, Marty Jack,
William Mathews, Wren McMains, Douglas Miller,
David Ness, Walter Nissen and Jack Pugh

Membership in the committee is open to any interested IPC user who feels he can made a substantial contribution. A few specific remarks should be made about the current committee:

1. Its members have day-to-day contact with all four major computer systems operated by IPC. That is, they are actual "hands-on" users. Some committee members use all four systems.

2. Most of the committee members are experienced computer professionals with over five years of involvement with computers at MIT. Some have spent more than ten years at MIT.

3. The committee has representatives from the academic, research, and student communities.

Summary of Meetings

Among the meetings held during this first year of the IPC Users' Committee were meetings devoted to:

1. The overall computer picture at MIT
2. MULTICS relaibility
3. The CP/CMS time sharings system on the 360/67
4. The future of batch processing and the 360/65
5. Available graphics software and hardware
6. Economic pressures on users and the center
7. MULTICS metering and performance measurement

Detailed summaries of most of these meetings are available. For copies, please contact W.D. Mathews, Project TIP, Room 14S-310, x 5687. The rest of this document summarizes many of the opinions voiced during the past year.

Computer Systems at MIT

IPC currently maintains four large computer systems at MIT.

1. MULTICS, a time sharing system running on the GE-645.
2. CP/CMS, a time sharing system running on the IBM 360/67.
3. OS/MVT, a batch processing system running on the IBM 360/65-40.
4. CTSS, a time sharing system running on the IBM 7094.

There seems to be a trend over the past years, and still continuing, to acquire still more time-sharing capacity and relegate less and less work to batch processing.

Types of Users

Computer users at MIT are incredibly diverse. Some have only the most limited contact with IPC's systems while others make long term, heavy use. To mention a few of these user types, we have students doing thesis and course work, research scientists doing experimental and applications work and system administrators doing testing and production. Ho study has yet been made of the way these user types depend on IPC and there has been no measurement or quantification of their demands.

MULTICS Anxieties

The MULTICS system has been available to the MIT community for the past 18 months. During that time it has progressed from a shakey and fitful system with a capacity of less than twenty simultaneous users to a current capacity of over forty users. There is considerable enthusiasm among the users with respect to many of the MULTICS design features. There is also high praise for the tuning effort being made and the dilligence of its maintenance staff in documenting crashes and attempting to bring the machine into a stable service situation. We would be seriously remiss, however, if we did not point out that MULTICS is also the source of some deep-seated anxieties in the user community. To name a few of these:

1. Unique Hardware with no follow-on machine announced. The MULTICS hardware configuration exists only at MIT and the Rome air force base. A third MULTICS-like configuration is being installed by Honeywell in Waltham. Given the present hardware, it would be highly unrealistic to predict many more MULTICS systems. Yet, follow-on hardware which would make the software more efficient and attractive, while constantly rumored, has not been officially announced. At this rate, MIT will be the only university with a MULTICS.

2. Unique system software with no institutional commitment beyond the first three years. The software services offered on MULTICS are available nowhere else. Consequently, programs written on MULTICS and relying on these services can be transferred to other environments and machines only with great difficulty. The converse is also true, software services commonly offered elsewhere are not available on MULTICS. This makes importation of subsystems written elsewhere a tricky maneuver. To our knowledge, no one has yet succeeded in a large scale transfer of software either into or out of MULTICS. At the same time, the institute has officially committed itself to MULTICS only for a period of three years, half of which has already passed. This casts some doubt on the advisability of developing subsystems on MULTICS with any pretence that these functions will be an enduring contribution to the community.

3. System Availability. Two components of system availability trouble the users. The first component is an old one: the system is often down due to crashes. There were over a thousand crashes during 1970, each one taking the system out of service for a half hour hour or more. It is further disquieting to note that roughly half of these crashes have been attributed to hardware. If this is true, users of the system will have to expect an exceptional amount of down time even when software development settles down. The second component of availability is overload. When the system comes up it almost immediately fills to its capacity of connected users. This is true not just during the day and during the night but as late as one or two o'clock in the morning. This is not a system for neophytes. No casual walk-up dial-in users here. The devotees know all the angles: forty or so of them are dialed in as soon as the machine comes up and have been battling over it ever since. Some users insist that being among the elect isn't so rosy either. Being connected doesn't mean service will be anything spectacular. These users feel that the present load of forty-plus users is unreasonably high and that this number is maintained only for publicity reasons and not out of a realistic regard for the workable capacity.

4. Limited Number of Compilers. PL/1 (again a unique PL/1 which for all we know may never be matched on other MIT machines) is the only solidly available programming language. This rules out importation of programs already written in fortran, cobol, AED etc. even if those programs demand only those services already provided by MULTICS.

5. Erratic Compiler Support. Even this one available compiler changes not just from day to day but sometimes from minute to minute. Allegedly, the compiler is undergoing improvement, each installation representling a "fix" to some bug. Often enough, however, this flyswatting leads to still other and more glorious bugs being introduced. The user is reminded to "recompile if he compiled anything between the hours of x and y on such and such a date". One wishes for a less crusading spirit (or at least that the recompilation would be done free of charge (or at least that the compiler group was large enough to debug suspected compiler problems as quickly as users complain of them (or at least that they would acknowledge and answer user complaints (or)))). Another rumor has it that all code will have to be recompiled when the version 2 compiler is introduced.

6. Unstable Standard Services. One of the concepts in MULTICS is that there is a set of standard services always available to the user. In theory these services are unchanging. In fact serveral standard service routines have been decommissioned and removed. In a curious rewrite of history, the claim was made that they had been announced as part of the standard service in error. Other standard services have changed in their action, invocations, or outputs. These changes are made unilaterally by the development team and without authorization from IPC or consultation with users. There is a paradox of temporality about the MULTICS system. Things which are "temporarily" wrong seem to endure forever. Things which the users thought were permanent turn out to be less permanent than what is customarily called "temporary".

7. Lack of Provision for detachable strorage. Besides such simple problems as the lack of any tape storage cabinets and the lack of operating procedures for mounting tapes, no 9-track tape drive is available. Although mountable disk packs are part of the configuration, procedures for user mounted disk packs have not yet been established.

8. Absentee User Facility not yet Available. It is not possible to submit jobs to be done except on-line via the console. This is a serious restriction for production-minded users.

9. Process saving still not available. It is not possible to save a process from session to session. This is especially serious in cases of system mishap. All processes on the machine are lost.

10. Multisegment files not well supported. Quirks in the implementation of multisegment files make their use somewhat unsatisfactory. This is of concern to users of large collections such as the census data.

11. Limited on-line storage. Although MULTICS was once advertised as indefinitely expandable, the reality is that its on-line file storage is very limited and does not equal the capacities of the other systems.

12. High and unpredictable costs. Charges for processing on MULTICS are not only high in dollars per hour but very unpredictable. The same process run twice may yield costs differing by more than a factor of two. It is not clear that the high hourly rate is justifiable in terms of machine power. Certainly the fact that on-line storage costs no less than it did on CTSS five years ago is regrettable.

13. Unknown and unadvertised limits and restrictions. Little is known about the practical limits of MULTICS. For example, what happens if a directory has 1000 files? What if a file has 100 different names? What if a link is ten levels deep? Some of these situations may be impossible or highly inefficient but there is no advisory documentation to refer to.

14. Insufficient MULTICS consultation. Nearly every substantive question about MULTICS is passed from IPC to the MULTICS development team or the GE group. It seems that a better in-house knowledge of MULTICS should be established and with it some control over implementation priorities.

CP/CMS Incongruities

The time sharing system available on the 360/67 serves twenty to twenty-five simultaneous users. It is an especially valuable asset to the MIT computing community in that it is the only on-line machine which runs 360 programs. This compatibility is very attractive to users at MIT developing elsewhere. Earlier fears that this system would be removed have largely been dispelled. In some respects, however, the system is treated as a poor relation in the IPC family. The following incongruities have been noted:

1. Although its main selling point is its compatibility with other 360's, different versions of the assembler and compilters exist on the two large 360 machines at IPC. This means there is no guarantee, for example, that a program which compiles on one machine will compile on the other.

2. In spite of the fact that the card punch for the 360/67 is only a few feet away from the card reader for the 360/65, no administractive provision has been made for the possibility of submitting a batch job to the 360/65 from the 360/67.

3. In spite of the fact that there is a direct channel-to-channel connection between the '67 and the '65 this has not been exploited to allow a remote job entry extension of capabilities for '67 users.

4. In spite of the fact that the two machines physically share some secondary storage devices, no regular way of moving data from the '67 to the '65 or back again has been established.

5. Only one staff member is assigned to maintain the system. User enthusiasm for the service does not seem to be matched by proportional commitment by the IPC staff.

CTSS

CTSS, the oldest time-sharing system on campus, is still not dead. For the record, it has some positive things going for it.

1. It provides consistent service with little down time.

2. It is owned outright by MIT and its operating cost is very low.

3. It is the only system to consistently earn revenue for IPC.

It is generally supposed that MIT cannot afford to have as much computer equipment as it now has and that the work being done on CTSS can be absorbed into MULTICS. We feel it is worth stating, however, that from the user point of view:

1. MULTICS (with its incessant crashes and down time, unavailable absentee user facility, non-detachable storage, lack of process saving, and general overload) is emphatically not yet a functional substitute for CTSS.

2. No evidence has been brought forth to show that MULTICS is cost effective. Indeed, it may well be that functions transferred from CTSS to MULTICS will generally cost more.

OS/MVT Apprehensions

As previously mentioned, the trend in computer services at MIT is away from batch processing. It is not clear whether this change reflects a basic alteration of computer consumption patterns or is in response to the ready supply of time sharing services on the campus. Nor has the permanence of this trend been established. Nevertheless, it is clear that the current batch capability is underutilized and an adjustment of the basic hardware configuration may be in order. Four alternatives were outlined to the committee. Briefly these were:

1. Change the current 360/65-40 configuration to a 750 K byte 360/65 with no 360/40. Maximum user core would be about 320K. Cost to IPC of all systems would be $4.3 million,

2. Drop the 360/65-40 and upgrade the 360/67 to a 1-million byte machine. Run a split shift arrangement with CP part of the time and OS/MVT the rest. Cost to IPC of $3.9 million per year.

3. Leave the 360/67 as it is and trade-in the 360/65-40 for a 370/155 to be shared with Harvard. Cost to IPC would be about $3.8 million per year.

4. Drop the 360/67 and upgrade the 360/65 to a 1-million byte machine with terminal access under some system such as CPS. Cost to IPC of $3.6 million per year.

These alternative raised many questions and apprehensions. First of all, why are there not many more alternatives? Is the Draper Labs 360/75 heavily utilized at any time other than during moonshots? Does the accounting scheme for the 360/65 at LNS rule out any load sharing? Cannot a merger with Lincoln Labs' facility be effected? Are the routine record processing needs of MIT (registrar, accounting etc.) truly at variance with the service being provided by IPC? Could their work be done at IPC? Could IPC work be done there?

Another apprehension is that no adequate profile of the batch and CP/CMS users is available. When the users' committee offered to help interview users and develop such a profile it was met initially with enthusiastic encouragement from IPC but was ultimately told that IPC would do it all itself.

Further, the plans presented generally did not mention alterations in salary, yet most of the personnel at IPC are involved in supporting the batch system. If we cut down on batch services will we still need thirty or so operators? Will we still need more IPC programmers assigned to batch services than to all three time sharing services combined?

It is worth noting that the question of IPC's next move on batch processing is met with extraordinary scepticism in some parts of the user community. In spite of a recent exhortation in "the bulletin" to ignore all rumors (imagine that!) and an insistence that everything is still in the planning stage and in due time will be revealed, some users solemnly profess certainty that the decision has already been made. Others shrug with resignation and conviction that whatever the decision is it will be wrong.

Mega Questions

Certain policy questions have also become evident to the committee. These questions are not specifically related to any single machine.

1. Does IPC intend to act as a broker for services at other universities and service bureaus when it is cheaper to go outside for such services than to acquire them here.

2. How does IPC intend to handle rebates for service failures on its on-line systems.

3. Can a systems of royalties be set up to reimburse projects for subsystems developed and introduced onto MIT computers.

4. Can rebates be given to persons who spend their time tracking down system bugs.

5. How should users outside MIT go about joining the MIT computer community.

6. What policy will IPC make concerning the acquisition of software to insure that all users share equally in the services available. The acquisition of APL by the electrical engineering department provides an interesting case. This software is available to EE students: other members of the MIT community are excluded.

Recommendations

We expect that appropriate staff members at IPC will address themselves to the issues raised in this report and will make a written reply to the Users Committee in due time. A few of the obvious recommendations that follow from our meetings are:

1. That the committees which now exist such as the IPAC and the Users Committee be constituted in a more formal manner with more clearly defined responsibilities. The Users Committee often feel that it operates in a vacuum. But an ineffectual Users Committee is surely worse than no Users Committee at all. We suggest that occasional joint meetings of the Users Committee and the IPC Advisory Committee would be appropriate.

2. That the IPC make commitments in writing concerning institute support (or lack of it) for its systems. The current situation erodes the morale of computer professionals both at IPC and throughout the campus. It also makes budgetary projections for computer use impossible. Hopefully, commitments can be made for periods of five years or more.

3. That some uniform measure be taken of all systems and that this measure be reported regularly. Initially this could be a very simple measure such as a report of total hours available, total hours sold, total income and total expenses for each system reported in the bulletin each month.

4. That new goals be set for the performance of the MULTICS system and that these goals reflect a sober attempt to turn the current experimental system into a predictable service.

5. That if a decision has not already been made concerning the future of batch processing at MIT, that a task force be set up to look at the impact each scheme will have on the types of users at MIT.