HONEYWELL INTEROFFICE CORRESPONDENCE

DATE: April 6, 1973

TO: R.F. Montee, Phoenix/MS C-10

FROM: R.E. Hoffman, Waltham/MS-505

DIVISION: DSO

cc: TA Ball
CT Clingen
JC Whitmore

RECEIVED
APR 10 1973
R. E. MONTEE

SUBJECT: Multics Meeting at MIT

As a result of several recent phone conversations between Bob Scott of MIT and Bob Henderson and Jim Renier, a meeting was held in Scott's Conference Room on 4/5/73. In attendance were:

MIT	DSO	PCO
Bob Scott	Jim Renier	John Couleur
"Corby" Corbato	Ted Ball	Bob Montee
Jerry Saltzer	Jerry Whitmore	Charlie Clingen
Bob Daley	Bob Hoffman	

Scott kicked-off the discussion by stating that, in the course of MIT's current activity concerning the management and utilization of 6180/Multics, two very important needs have become apparent:

1. The need for a system with an extremely large memory.

2. The need to reduce MIT's unit cost of providing memory on the 6180/Multics system.

Because of MIT's vested interest in Multics, Scott felt compelled to identify these "two main barriers to Multics success" to Honeywell's top management. He went on to suggest that the resolution of both requirements lies in acquisition of a new Multics memory which is "as large as possible, as cheap as possible, as reliable as possible, with delivery as soon as possible".

Renier's response was that "the HIS people are here today to make certain that we throughly understand MIT's memory problem and that we will respond to MIT's stated needs within a reasonable time frame". Scott then handed-out draft copies of their large-memory RFP to all attendees.

Saltzer cited several specific examples of current projects which now require or will require the rapid random access of an extremely large memory:

1. An Artificial Intelligence project involving the structure of the English language and the ability to relate sentences and words to known (stored) facts.

2. An Automatic Programming project in which computer programs are generated by a computer program -- a program generator for Inventory Control is now operational.

3. "MACSYMA" - the Project-MAC Symbolic Manipulator -- an operational system (on a PDP-10) which reputedly manipulates functions better than the average college graduate. (MACSYMA has been moved onto 6180 from 645.)

4. An example from the medical area in which a half-dozen of the world's leading experts on kidney disease imbed all of the facts and their combined expertise into a (stored) fact base; inquiry into this data base by an "average" physican enables him to reach the same conclusions that a "real expert" would. Corbato cited the conclusion that a typical Doctor's knowledge is not very deep but is very broad -- this situation is ideally suited to the application of a computer's ability to rapidly scan, find "the match" and then display the in-depth facts of a subset.

The basis of all these projects is LISP (which is operational on both the 645 and 6180); a super-fast processor is not required for efficient implementation, just a helluva big memory.

Daley commented on IBM's ability to provide all the "right" hardware (including a 4-megabyte memory) at the "right" price to do these jobs, but he conceded that IBM doesn't yet have Multics. MIT compares an IBM-370/165 at $117K/month with an H6180 at about the same price but only half the performance and less memory. They expect their Mod-168 (with 3-megabyte memory) to have 20% better overall performance than the Mod-165.

Scott declared MIT's preference for acquiring the solution to their memory problem from Honeywell if this is at all possible; he stated that MIT would not be sending-out the RFP to the six vendors (not including HIS) on the bidder's list until they hear what Honeywell will do and will not do in this matter. Corbato then explained that "the current momentum of Multics (acceptance) is marginal at best when the users don't see any enhancements coming " ---MIT feels that memory seems to be the most fruitful area for an encore to our January/1973 public announcement.

Renier's response was that we (HIS) understand and want to be sensitive to MIT's requirements. He pointed-out that, in order to survive and grow, we must be profitable -- new developments deserve premium dollars for premium features. We fully appreciate that this (memory suggestion) is a very "strong indicator" of one of the key areas of enhancement, not only for Multics but also for all of our large-scale futures; we recognize the importance of advanced development and "advanced" users (such as MIT), we will explore the matter among HIS senior management and get back to MIT with our comments and decisions.

Renier then presented the analogy to the 200/2000 marketing success, but Corbato countered with the commentary that the "Multics market" is not being developed, whereas the success of the 200/2000 has been built in the "1401 market" which was already in existence. Daley pointed-out that the 370/145 would be "an ideal system on which to run a small Multics" -- (I feel that this was a veiled attempt by Daley to prompt HIS to answer Corbato by stating our plans to announce 6160, 6140, etc.). The rumors from IBM indicate that they will have Multics within three years, matched feature for feature. Renier reiterated that when we're leading, we expect to get the bucks.

There followed an extended discussion (between Couleur and Corbato primarily) on the philosophy of Multics architecture and the suggested use of cache memory. Montee stated that we have looked at high-density core packaging but that we need extra address lines to be able to have more than 64K-words on each of the two memory ports on an SCU. The current packaging on 6180 will allow two-million words (2048K) of directly-addressable core; redesign of the SCU to accommodate more address lines will take 14 months after funding is approved.

Montee further suggested that the use of a cache store would overcome the problem we are experiencing with certain slow instructions. He cited that we have sold (but not yet shipped) a 2K-word cache for installation on a 6070 at NED. In the testing of this cache, Phoenix has seen reductions of 25-15% in run time, depending on the mix of direct/indirect instruction modification. Additional (and substantial) reductions should occur in a Multics environment because of the heavy indirection.

After being assured by Couleur and Montee of their understanding of the problem, Renier promised to review the matter with Norm Feldman on 4/25/73 and subsequently with Bob Henderson. Scott assured us that they (MIT)"will tolerate the problem (and not make waves) if we see some certainty of a solution". The discussion then shifted from memory futures to current performance. MIT restated that their recent benchmarking and metering has disclosed that the 6180 processors are not as fast as they had originally anticipated. There was some misunderstanding and disagreement over just how much "detune" now exists in the system. Montee stated that the 6180 should run only 5-7% slower than a 6080 on basic ops. Daley reported that he and his people are seeing ratios of 2.00 to 2.02 times faster (than 645) on compilations.

When asked about Honeywell's plans to measure performance, Couleur replied that we are willing to accept MIT's numbers for the moment because we are now concentrating on getting the software to work (however slowly); the measuring and tuning and investigation of instruction anomalies will be deferred until sometime after the availability of an operational 6180 in Phoenix. Corbato's response was "fine--you should publish these intentions" and Scott commented that "we never really expected more than two times (the performance of a 645)".

There was general agreement that the only apparent and productive "quick fix" to improve performance would be to install a cache on the 6180. PCO reported that an 80-nanosecond cache has been designed, can run in a uni-processor system now, is built into the data hubs and is a field-retrofit item. (IBM's cache on the Mod-165 is 120-ns.)

Couleur and Montee agreed to provide all appropriate inputs both to Feldman and to Renier prior to their 4/25 meeting in Phoenix at which time they will discuss the ramifications of all of the various issues raised in this report. After the meeting in Scott's Conference Room, there was an extended discussion at lunch on a variety of topics, including some clever strategies for implementing cache in a multi-processor system.

Bob Montee -- this is one man's opinion of what transpired. Please feel free to add, correct, and delete before you publish the "official" report of the visitation.

R.E. Hoffman
Account Manager, MIT

REH/s

Jerry Saltzer
file - Multics
Administrative

Distribution

C. A. Conover

C. W. Dix

N. N. Feldman

J. Gentell (CISL)

J. Krasula (McLean)

J. Tibbs (McLean)

Attendees

Report on Multics Meeting at MIT
R. E. Hoffman
4/6/73