

file

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

PROJECT MAC

Reply to: Project MAC
545 Technology Square
Cambridge, Mass. 02139

Telephone: (617) 864-6900 x6201

To:	R. Hoffman	} meeting attendees	F. J. Corbató
	R. Montee		C. T. Clingen
	W. Hooper		E. Fredkin
	R. Chevalier		M. D. Schroeder
	R. Scott		W. A. Martin
	R. Daley		
From:	J. Saltzer		

Date: September 26, 1973

Subject: Notes of meeting, 9/25/73, to discuss engineering details of HISI large memory systems proposal.

A meeting was held at 2:30 p.m. in R. Scott's office, to obtain more technical information and background on the proposal from HISI for an 8-million word LSI memory and CPU cache system for the M.I.T. Multics site. These notes record the discussion as I recall it. The primary discussion consisted of Daley and Saltzer asking questions of Montee. The discussion bounced back and forth among several points, so I have taken the liberty of rearranging the order of these notes to discuss one point at a time.

The overall proposal was briefly reviewed. The proposal would have two major phases. In the first phase, the present 384K word 0.5 μ sec. core memory and the 2 million word bulk core memory would be replaced with 2 million words of LSI memory using a technology of 1024 bits per chip. This phase would be essentially finished in about 12 months. In the second phase, to start about 12 months later, the 2 million words of LSI memory would be upgraded to (actually replaced by) a technology using 4096 bits per chip. Expansion above 2 million words could proceed as rapidly as desired after that time, to a limit of 16 million words.

In both phases, the performance specifications of the memory would be identical: the speed of the memory, as seen from inside the CPU (known locally as the time between strobes \$INT and \$DA) is anticipated to be near 1.0 μ sec., with specifications guaranteed of 1.2 μ sec. This is both the access time and the cycle time to a 72-bit double word, and compares with a time of 650 ns. for the present core memory. Apparently access time at the chip is about 350 ns. from initial strobe, but cycle time at the chip is 1.0 μ sec., and that cycle time has been allowed to control the specifications throughout the system, since a CPU cache is assumed.

In both phases, a 36-bit word will be represented by a 40-bit error-detecting-and-correcting code. Details of the code are not specified. (Since 4 check bits are not sufficient to correct all one-bit errors in a 36-bit word, some multiple-word scheme such as grouping 2 words to get 72 information bits and 8 check bits must be in use, but the point was not followed up.) It was reported that a reliability model of an 8 million word configuration indicates that on the average, 1 chip per day will fail, but that with error correction, uncorrectable failures will occur no more often than once per week. Details of the maintenance strategy assumed in the model were not known, but Montee promised to obtain them.

The overall maintenance strategy is that the error detection/correction circuitry will report failing chips; when convenient a 2 million word memory controller can be configured off-line, a 512K word module disconnected, and the remaining memory returned on-line. The offending board is removed, the chip replaced and the board reinserted (about a one-hour job). Possibly a spare-board strategy may be necessary, but it is felt that more experience will be needed to tell. An "exerciser module" will also be supplied, similar in spirit to the exerciser used for the current bulk store.

Supply of the 1024-bit per chip technology seems certain; the memory is identical to that currently being delivered as HISI 6025 main memory, and several suppliers are under contract. Supply of the 4096-bit chips is less certain. One supplier (AMS?) has delivered samples which are apparently satisfactory. If the 4096-bit chips fail to materialize on time, an internal HISI program to develop 2048-bit chips will be used as backup.

A floor plan was exhibited which showed how the present 6180 space could be rearranged to house 8 million words if 4096-bit chips are used. The same floor plan accomodates only 2 million words constructed of 1024-bit chips. There was space in the plan for four processors. Scott pointed out that if 2048-bit chips were used, more space would be required to get to 8 million words but the problem is solvable.


Several hardware changes are required to allow addressing of a large primary memory. The CPU, IOM, and memory controllers are architecturally organized to allow 24 bits of address, but are currently implemented only for 21 bits. These 21 bits consist of 3 bits used to choose a memory controller and 18 used to select a word in that memory controller. The proposal is to widen the 18 bit address paths to 21 bits rather than, for example, expanding the number of memory controllers to 64. The memory controller requires the largest overhaul. It currently consists of 17 boards. Eleven of the 17 would need to be replaced, two completely new board designs are required, and 6 board designs need minor modification.

In the CPU and IOM, things are simpler. Two boards need to be modified, and about 10 backpanel wires added; this change can be done in the field relatively easily. Carry lookahead is used, so no slowdown of address preparation logic is anticipated. (These changes do not include the addition of the cache, discussed below.)

Discussion then turned to the cache proposal. The cache is to be a 4-level set associative design, with 128 columns, each column containing four 4-word blocks and corresponding to a set of 1/128th of the absolute memory range of the system. A simulation of GECOS programs using both an LRU and a round-robin replacement algorithm showed a substantially better hit ratio with round-robin, so this strategy is planned. The strategy proposed by Webber of bypassing the cache for shared, writeable segments (detected by software which notifies the hardware by setting a bit in the segment descriptor) will be used.

Not very much performance information was available. The following numbers will be obtained by Montee: the speed of the CPU if the cache had 100% hits when running Multics, and the speed of the CPU if the cache had 0% hits when running Multics and using a 1.2 μ sec. memory. Considerable discussion was held on the problem of predicting what actual hit ratio can be expected, but there are apparently no tracing tools available which could obtain this information. There are two reasons for not assuming that 6070 experience can be extrapolated to the 6180: 1) Multics is known to run an instruction mix radically different than GECOS -- so different that the CPU instruction rate is only 50 to 70% that of GECOS, and 2) the strategy of omitting shared writeable segments from the cache will probably affect the hit ratio. The proposal to install a cache immediately on one CPU to see what happens is not viable, since the only cache available at present is a 6070 cache, and it will not work in either a 6080 or a 6180, which use a more complex port selection design. At the time I had to leave the meeting, this problem was still unresolved.

One final observation. During the discussion of CPU speed it was pointed out that one of the reasons why the 6180 did not speed up as much as expected when compared with the 645 is that the 645 had a buffer in the path between the pointer registers and memory, so that address preparation overlap could be accomplished during instructions which manipulate pointer registers. Since the 6080 EIS machine did not have this buffer in the corresponding path, it was not implemented in the 6180 either. Thus all pointer register instructions (which are frequent in Multics) operate with address preparation overlap inhibited. This missing register apparently explains why some code sequences (especially the Multics Call/Save/Return Sequence) operate only 1.5 times as fast on the 6180 as they did on the 645. There are apparently no plans to add a buffer register to the 6180.

A handwritten signature, possibly reading "Jew", written in dark ink. The signature is stylized and appears to be a personal mark or initials.