INTERDEPARTMENTAL

MASSACHUSETTS INSTITUTE OF TECHNOLOGY *CAMBRIDGE, MASS. 02139*

from the office of

Price quotations deleted from this version JHS.

To: Ed Fredkin

From: J. H. Saltzer

Date: November 28, 1973

Bob Scott has received a new large memory proposal from HISI. Technically, the proposal appears identical to the earlier proposals, but the entry price has been significantly reduced. I have been puzzling how to describe the proposal succinctly because individual aspects are bundled and hard to sort out, but the following seems like a fair description:

1. Today's 6180, with 2 cpus, one IOM, 384k of 0.5 μsec memory, and 2M words of bulk store, has a purchase price of $4.8M. (There are also disks, tapes, printers, and communication equipment worth about $1.5M).

2. HISI's current proposal would replace the bulk store and 384k memory with a 2M word 1.2 μsec primary (MOS) memory, add a cache to each cpu,

3. After that replacement, HISI offers to provide additional 1.2 μsec MOS memory in 2-million word increments

As for system performance, there are two effects: the 2M word memory would eliminate about 85% of the paging activity, producing a 20-30% performance improvement. The combination of cache and 1.2 μsec memory produces a complex effect, depending on the hit ratio of the cache. As shown by the accompanying memo, any hit ratio higher than .54 would probably not degrade performance; a hit ratio in the range near 0.8% would probably increase performance by about 30%. Thus the proposal described in point two would probably produce a system with a better price/performance ratio, but perhaps with more capacity than sellable at current IPC rates.

This apparently increased price/performance ratio is actually the result of a comparison with a system which is operating well below its expected performance. The 6180 was anticipated to be a 0.95 mips processor, rather than a 0.65 mips one, and the bulk store was expected to be about twice as fast, so that paging delays are swallowing up about 10% more of the system time than expected.

To summarize,

	CPU speed	Useful time after paging	Computation available to users
expected	.95 mips	70%	.66 mips/cpu
available	.65 mips	60%	.39 mips/cpu

Thus, the computation available to the user, on which price/performance should be judged, is about .39/.66 = .59 of that expected. Given these observations, the proposal to add a cache and replace the bulk store with directly addressable 1.2 μsec memory would bring the system performance back up to the range originally anticipated. To add a third line to the above chart, we might assume a 0.8 hit ratio:

	CPU speed	Useful time after paging	Computation available to users
with cache hitting 0.8	.85 mips	90%	.76 mips/cpu

Recommendations

1. Since the actual hit ratio to be expected by a cache is uncertain, and since even the most optimistic assumptions merely bring system performance near to that originally anticipated, I suggest that we recommend to IPC that they insist that HISI perform the first part of the large memory upgrade without additional charge. (That is, install a cache in each cpu and replace the bulk store and 384k core with MOS memory). This should leave IPC prices and utilization unchanged, but return the system to a position where its first-shift load is 60% of its capacity rather than 100%.

xc: F.J. Corbató
W.A. Martin
J. Moses
R.H. Scott
R.C. Daley

INTERDEPARTMENTAL

MASSACHUSETTS INSTITUTE OF TECHNOLOGY *CAMBRIDGE, MASS. 02139*

from the office of

To: Ed Fredkin

From: J. H. Saltzer

Date: November 28, 1973

Bob Scott has received a new large memory proposal from HISI. Technically, the proposal appears identical to the earlier proposals, but the entry price has been significantly reduced. I have been puzzling how to describe the proposal succinctly because individual aspects are bundled and hard to sort out, but the following seems like a fair description:

1. Today's 6180, with 2 cpus, one IOM, 384k of 0.5 μsec memory, and 2M words of bulk store, has a purchase price of $4.8M. (There are also disks, tapes, printers, and communication equipment worth about $1.5M).

2. HISI's current proposal would replace the bulk store and 384k memory with a 2M word 1.2 μsec primary (MOS) memory, add a cache to each cpu, and increase the $4.8M price to $5.6M.

3. After that replacement, HISI offers to provide additional 1.2 μsec MOS memory in 2-million word increments at 1.6¢/bit, or $1.1M per 2M words. Thus, a 4M-word system would cost about $6.7M, and an 8M-word system about $8.9M. (Not counting $1.5M in peripherals.)

I have described all the prices in terms of purchase, for simplicity; HISI is prepared to rent or sell, as desired. IPC has set its current prices and budgets around the price and performance of the $4.8M system. Either an IPC price increase on the order of 20% or a 20% increase in billable usage (about $30k/month) would be needed to pay for such a system.

As for system performance, there are two effects: the 2M word memory would eliminate about 85% of the paging activity, producing a 20-30% performance improvement. The combination of cache and 1.2 μsec memory produces a complex effect, depending on the hit ratio of the cache. As shown by the accompanying memo, any hit ratio higher than .54 would probably not degrade performance; a hit ratio in the range near 0.8% would probably increase performance by about 30%. Thus the proposal described in point two would probably produce a system with a better price/performance ratio, but perhaps with more capacity than sellable at current IPC rates.

This apparently increased price/performance ratio is actually the result of a comparison with a system which is operating well below its expected performance. The 6180 was anticipated to be a0095 mips processor, rather than a 0.65 mips one, and the bulk store was expected to be about twice as fast, so that paging delays are swallowing up about 10% more of the system time than expected.

To summarize,

	CPU speed	Useful time after paging	Computation available to users
expected	.95 mips	70%	.66 mips/cpu
available	.65 mips	60%	.39 mips/cpu

Thus, the computation available to the user, on which price/performance should be judged, is about .39/.66 = .59 of that expected. Given these observations, the proposal to add a cache and replace the bulk store with directly addressable 1.2 μsec memory would bring the system performance back up to the range originally anticipated. To add a third line to the above chart, we might assume a 0.8 hit ratio:

	CPU speed	Useful time after paging	Computation available to users
with cache hitting 0.8	.85 mips	90%	.76 mips/cpu

Recommendations

1. Since the actual hit ratio to be expected by a cache is uncertain, and since even the most optimistic assumptions merely bring system performance near to that originally anticipated, I suggest that we recommend to IPC that they insist that HISI perform the first part of the large memory upgrade without additional charge. (That is, install a cache in each cpu and replace the bulk store and 384k core with MOS memory). This should leave IPC prices and utilization unchanged, but return the system to a position where its first-shift load is 60% of its capacity rather than 100%.

2. For each $25k/month ($300k/yr) in additional revenue which Project MAC can provide or locate, an additional 2M word memory box could be attached. This price is quite reasonable, and some arrangement along these lines should be pursued. This path requires the minimum effort in system engineering to a large memory system.

xc: F.J. Corbató
W.A. Martin
J. Moses
R.H. Scott
R.C. Daley

To: 6180 Performance File

From: J. H. Saltzer

Date: November 28, 1973

Subject: Cache hit ratios required for 6180 performance

I. CPU model

The measured speed of the 6180 running Multics is 0.65 mips, or 1500 ns/inst. That speed is with a 650 ns. memory access time as viewed inside the 6180 cpu. If we assume (from instruction timings with maximum overlap) that the speed with a zero-access time memory is 1.25 mips, or 800 ns/instructions, we conclude that the 650 ns. memory produces about 700 ns. of delay per instruction, or 1.1 memory delays per instruction. We thus crudely estimate that the time per instruction, t_i, is

$$t_i = 800 + 1.1t_m \quad \text{nanoseconds}$$

where t_m is the average memory access time and we have grossly simplified lookahead and overlap details.

II. Cache model

Assume a 100 ns. cache and a 1000 ns. main memory. Assume that access to main memory is serial with cache access, so that the time for access when the cache fails to hold the desired item is (100 + 1200) ns. If the hit ratio (success rate) is r, the average access time is then:

$$t_m = r \cdot 100 + (1-r) \cdot 1300 \text{ ns.}$$

III. Hit ratios required for particular performance levels

We now calculate values for some particular possible performance levels.

for this hit ratio	the average access time would be	and the cpu would run at	giving a performance relative to the 6180 of
0	1300 ns.	0.45 mips	.69
.5	700 ns.	.63 mips	.97
.54	650 ns.	.65 mips	1.0
.6	580 ns.	.70 mips	1.07
.7	460 ns.	.77 mips	1.18
.8	340 ns.	.85 mips	1.30
.9	220 ns.	.96 mips	1.47
1.0	100 ns.	1.1 mips	1.69

minimum hit ratio required not to retrogress

Note that these performance predictions are very sensitive to some untested assumptions, particularly

1. cache speed, assumed to be 100 ns.
2. primary memory speed, assumed to be 1200 ns.
3. cpu speed with zero-access time memory, assumed to be 1.25 mips.

INTERDEPARTMENTAL

MASSACHUSETTS INSTITUTE OF TECHNOLOGY *CAMBRIDGE, MASS. 02139*

from the office of

To: Robert Scott

From: J. H. Saltzer

Date: April 4, 1973

Subject: Second Draft of RFP for large memory for Multics

Enclosed is a revised copy of the large-memory RFP, containing response to comments from Fredkin, Corbató, Daley, Burner, and yourself.

A quick survey of local opinions suggests that the following six companies appear to be technically equipped to make proposals:

Cambridge Memories
Intel
AMS
Toko
Fairchild
IBM

Of them, the I.P.C. experience in purchasing add-on memory from Cambridge Memories indicates that delivery capability of that company is real. It should also be noted that Ford Motor Company is apparently discussing add-on memory for their Honeywell 6080 with Fairchild.

DRAFT II

April 2, 1973
J. H. Saltzer

REQUEST FOR PROPOSALS

The Massachusetts Institute of Technology Information Processing Center, in conjunction with M.I.T. Project MAC, invites specific proposals for engineering, manufacture, installation and maintenance of a large primary memory system for the Honeywell 6180 computer system operated by the M.I.T. Information Processing Center. This system is run as a computer utility using the Multics operating system.

The current system consists of 2 Honeywell 6180 processors, 384k 36-bit words of 0.5 microsecond core memory, a 2048k word bulk core memory (1.5μsec) and associated I/O and disk storage equipment, supplied by Honeywell. This request envisions replacing the 384k word primary memory and the 2048k word bulk store with a one-level primary memory of 8 million 36-bit words.

The remainder of this document provides details which should be considered by a proposal.

Size

The complete system is to consist of 8 million 36-bit words of directly addressable, random access, primary memory. Provision should be made for possible later expansion to 16 million 36-bit words.

Performance

Access and cycle times on the order of 500 nanoseconds at the system controller interfaces are usable by the Honeywell 6180 processors. There is, however, a potential tradeoff between memory price and access time which may be exploited by prospective vendors. For example, a 750 nanosecond memory system with three processors might achieve the same effective performance as a 500 nanosecond memory system with two processors, but with lower total cost.

Technology

It is presumed that to meet the requirements of this request, an electronic memory system based on an integration level of 2048 or more bits per chip will be required. Since the schedule calls for phased delivery over a period of four years, it is acceptable for later delivered modules to use a higher integration level.

Configuration

The planned configuration is based on the standard Honeywell 6000 system organization. Figure one illustrates one possible arrangement. An alternative configuration would consist of only four system controllers, each with two memory modules of 1M words. The planned configuration may be operated with two, three, or four central processors. The designation (H) in figure one indicates that the module will be supplied by Honeywell (or some other previously arranged supplier) and the designation (V) indicates that the module is to be supplied by the vendor responding to this request. The designation (H/V) on the System Controller modules means that the vendor may propose to supply this module, or plan to use the standard Honeywell System Controller. The remainder of this document will assume the latter choice.

Modularity

It is important that the complete 8-million word system consist of at least eight completely independent units, in order that any individual module may be shut down and disconnected from the reaminder of the system while the remainder continue in active operation. This requirement suggests that independent power supplies and attention to switching transients which could affect modules still in use will be required.

Figure One

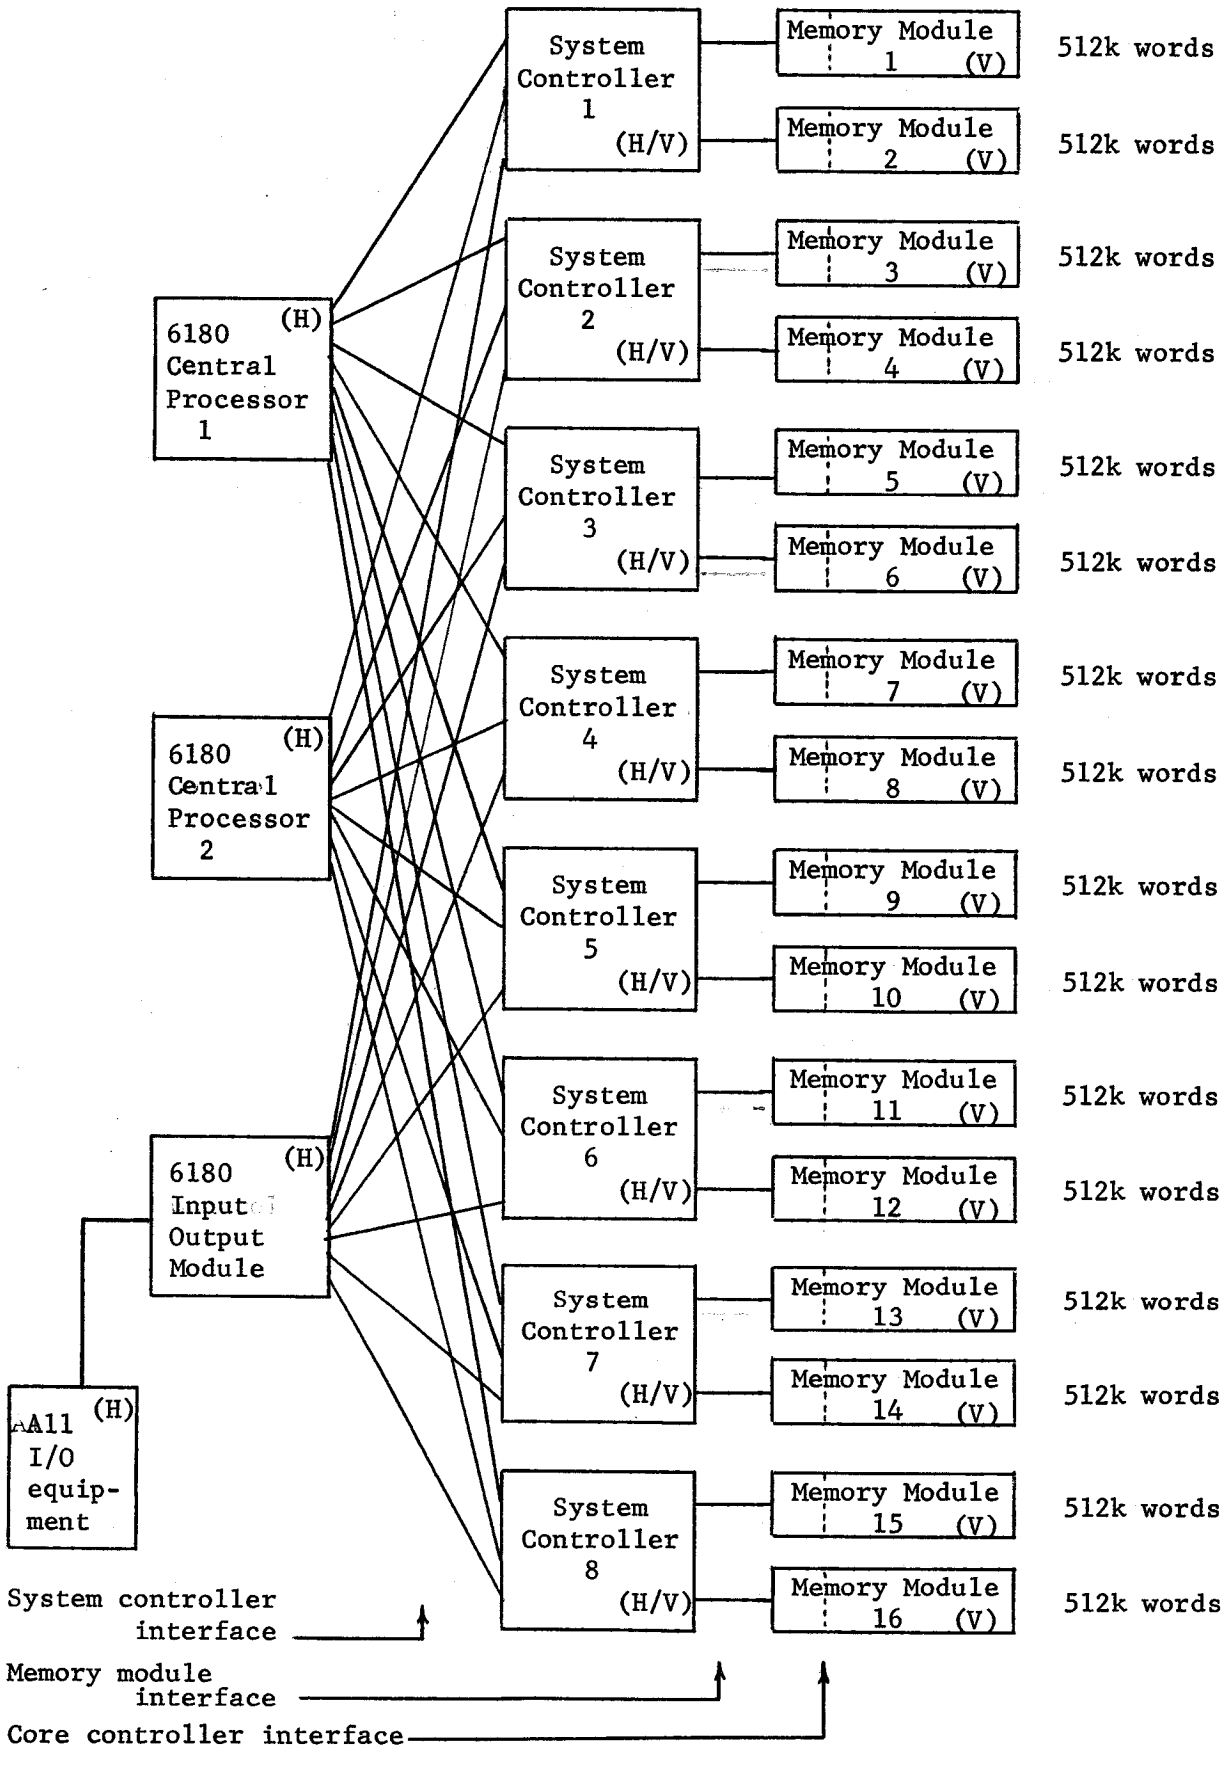
6180 Central Processor 1 (H)
6180 Central Processor 2 (H)
6180 Input Output Module (H)
All I/O equipment (H)
System Controller 1 (H/V)
System Controller 2 (H/V)
System Controller 3 (H/V)
System Controller 4 (H/V)
System Controller 5 (H/V)
System Controller 6 (H/V)
System Controller 7 (H/V)
System Controller 8 (H/V)
Memory Module 1 (V)
Memory Module 2 (V)
Memory Module 3 (V)
Memory Module 4 (V)
Memory Module 5 (V)
Memory Module 6 (V)
Memory Module 7 (V)
Memory Module 8 (V)
Memory Module 9 (V)
Memory Module 10 (V)
Memory Module 11 (V)
Memory Module 12 (V)
Memory Module 13 (V)
Memory Module 14 (V)
Memory Module 15 (V)
Memory Module 16 (V)
512k words
System controller interface
Memory module interface
Core controller interface

Schedule

An initial two-million word batch would be scheduled for installation sometime in 4th quarter, 1973. Delivery of the remaining memory would be in uniform increments over a four-year period, perhaps with a 1-million word increment each 6 months.

Completed proposals should be in the hands of Mr. Robert Scott, Director of Information Processing Services, by June 1, 1973.

Interface

There are at least three interface levels to which attachment to the Honeywell 6180 system may be made: the "core controller" interface, the "memory module" interface, and the "system controller" interface. The "core controller" interface is the most primitive, and closest to the memory devices, but attachment to that interface would require use of the Honeywell core controller logic, which currently is not obtainable as a separate product. The specification of the "memory module" interface may be easier to obtain, but attachment to that interface would require that the vendor supply the equivalent of the Honeywell core controller logic. Attachment to the "system controller" interface requires that the vendor supply the equivalent of the Honeywell System Controller, a substantial piece of logic, including multiple ports for attachment to processors and I/O controllers. Of the three interface possibilities, it is suggested that "memory module" interface be considered first.

Details of the interface specification will have to be obtained by direct negotiation between the vendor and Honeywell. This negotiation may be aided by Honeywell's recent agreement with the General Services Administration to supply technical information required for attachment of memory to the Honeywell 6000-line, which uses the same interfaces as the 6180. In addition,

Project MAC's position as a defense contractor may permit it to take advantage of this agreement directly.

6180 System modifications

Some minor modifications to the Honeywell 6180 Processors, IOM, and System Controllers may be required to provide high-order address lines for which architectural provisions have been made. A completed proposal should include a plan for installation of these modifications, either by Honeywell or by the vendor.

Reliability

The overall reliability of the 8-million word memory system must be at least equivalent to that of common 256k word systems today. An MTBF of 168 hours (1 week of continuous operation) is suggested. This requirement would appear to force an increase in reliability level by a factor of 32. However, certain special properties of the system may be used to make this reliability specification easier to meet: the rate of accesses to the memory system by two parallel processors will be about the same as to a 512k word system, so the average density of access to any one module will be reduced by a factor of 16. Thus, although reliability expressed in mean time between failures is much higher than usual, when expressed in mean number of references between failures, it may actually be quite ordinary. Note however, that we are discussing only averages; any one module must also be prepared for a continuous stream of successive accesses for an indefinite time.

It would appear that error correction techniques will probably be required. It is essential that some method be devised for keeping count of the number of errors successfully corrected, so that the operating system may have an early warning of possible trouble. For example, the highest-order word

in each memory module might maintain a count of the number of errors successfully corrected in that module since module initialization.

In addition to error correction, uncorrectable errors must be detected with high probability, and reported to the operating system via the standard H-6000 "parity failure" interface.

The lifetime of the equipment should be at least 5 years (45,000 hours) when operated continuously, 24 hours per day, 7 days per week.

Maintenance

The proposal should include a maintenance plan for the memory system, including any equipment needed, such as a memory load simulator, and whether or not a spare memory module is planned as part of the system. Maintenance costs and availability of experienced maintenance personnel should be addressed.

The proposal should also address itself to the maintenance of the Honeywell 6180 processors and I/O system, presumably to be carried out by Honeywell on terms acceptable to its Field Engineering Department.

Acceptance Test

Each module will be subject to a separate acceptance test. One week of continuous operation in production service on the 6180 without a hardware failure will be considered as initially passing the reliability specification. In addition, if later experience indicates that a previously accepted module is actually below specified reliability, acceptance of later modules will be delayed until the specification is met.

Tests confirming that access and cycle time are within specification will also be performed.

Engineering Contact

This request for proposals provides only sketchy technical details. It is anticipated that any complete proposal will require considerable ironing out of technical details, and will therefore require direct contact between the responsible engineering personnel of the proposer, and those of M.I.T.

Bidder's Technical Conference

To maximize the rate of information transfer to prospective bidders under this request, M.I.T. will hold a technical conference for bidders on May 1, 1973. At this conference, a brief review of salient features of this RFP will be given, together with any additional considerations which seem significant. Prospective bidders will than have an opportunity to inquire about requirements not described here, and to discuss negotiability and tradeoffs among the explicitly stated requirements.