COST PERFORMANCE ASPECTS OF CCD FAST AUXILIARY MEMORY

Dileep P. Bhandarkar

Texas Instruments Incorporated
Dallas, Texas

ABSTRACT Charge coupled devices (CCD's) have been mentioned as potential fast auxiliary memories in multiprogrammed computer systems with virtual memory. Declining MOS RAM costs will tend to allow computer designers to put more random access memory in their systems. However, due to program locality only a portion of the memory is actively used. If CCD's cost R (>1) times less than RAM's, the size of the random access main memory can be reduced and a CCD paging store can be provided as a back up. Active program segments can now be brought into the RAM as and when needed. Often the question has been raised in industry as to what value of R is needed to justify this application of CCD's.

A queueing network analysis is presented in this paper. The model developed is used to establish a criterion for cost-effectiveness. A relationship is derived between RAM/CCD cost ratio R and the page exception characteristics of the program environment. A rough rule of thumb is suggested. It states that cost-effectiveness is achieved provided the page miss ratio is less than σ, where σ is the ratio of the processor cycle time to CCD block read time. This criterion is applied to known miss ratio characteristics to obtain a critical value of R ranging from 1.2 to 2.0 for $\sigma = 0.005$ and 0.002 respectively. This indicates that a projected cost ratio of two to four can make fast auxiliary CCD memories attractive for block read times ranging from 200 to 500 microseconds.

INTRODUCTION

Several memory technologies with diverse cost-performance characteristics have emerged as potential storage media in digital computers. At the high performance end of the product spectrum, core memory, MOS and bipolar random-access-memory provide a medium for active program storage. Cheaper backing store is or soon will be available with magnetic disks, charge coupled devices (CCD) electron beam addressable memories (EBAM), and magnetic bubble domain devices (MBD). Due to the significant price and performance difference between primary and secondary memory technology, computer designers have used multi-level storage hierarchies where the objective is to keep the current information in the fastest devices and the rest of the information in

Table 1
Cost-Performance of Competing Technologies (Circa 1978)

Technology	Access Time to First Bit	Typical Serial Block Size (Bits)	Maximum Serial Transfer Rate	Parallelism of Data Transfer Path	Cost/Bit (¢)
MBD	0.5 - 1.5 ms	256 - 2048	100 - 500 kHz	4 - 8	0.02 - 0.1
CCD	0.1 - 1.0 ms	256 - 1024	1 - 10 MHz	up to 8 or 16	0.02 - 0.1
MOS RAM	200 ns	1	2 MHz	8 - 64	0.05 - 0.2
EBAM	30 μs	1024 - 8192	10 MHz	1 - 8	0.02 - 0.1
FHD Slow	17 ms	1024 - 8192	2 - 5 MHz	1	0.015 - 0.020
Med	8 ms	1024 - 8192	4 - 6 MHz	1 - 2	0.025 - 0.030
Fast	4 ms	1024 - 8192	4 - 24 MHz	1 - 4	0.100

the slower and less expensive devices. The individual cost and performance characteristics and the user environment (program behavior) will affect the optimal memory system design. This paper presents an overview of the cost-performance characteristics of some of the technologies and identifies the computing environment best suited to them. Moving head disks (MHD) are generally too slow for active program auxiliary storage. They are a secondary memory to store inactive program segments or files. Fixed head disks or head per track drums have been used as fast auxiliary memory between primary and secondary memory.

COMPETING TECHNOLOGIES

The different technologies can be functionally classified into Random Access and Block Access Memories. In RAM's, each bit of storage is individually accessible; in BAM's, the smallest accessible unit is a serial stream of bits. Table 1 summarizes the typical cost-performance characteristics of memory systems using various technologies [cf. 1]. Some of the non-quantitative features of these technologies are discussed below. The traditional BAM market is served by magnetic disks while core memories are most commonly used as RAM's.

Magnetic Bubble Memories

Magnetic Bubbles* provide non-volatile block access storage, competing directly with fixed-head disks. If the projected cost estimates are met, MBD's should virtually capture all the market for small capacity (< 10M bits), low access time FHD's. They also pose a serious threat to the medium per-

*This paper is concerned only with large bubbles (5μ in diameter).

formance FHD's because they offer greatly improved performance at comparable prices. The market for large FHD's (> 100M bits) with relatively slow access times will not be significantly affected as they have a price advantage of a factor of two or three over MBD's. An important feature of MBD's is their modular expansibility. Users will be able to upgrade relatively easily the capacity of their storage to meet increasing requirements. Their non-volatility gives them an edge over CCD's. However, in fast auxiliary memory applications, volatility could be tolerated due to the existence of backup storage in the form of a moving head disk. Note that a three-level hierarchy is not necessary for MBD's and FHD's in small computer systems. The stopability and asynchronous operation of MBD's can be exploited in future computer architectures.

Charge Coupled Devices

CCD memory has the best potential for very high speed operation. Most current computers are not designed to handle fast peripheral devices, and novel architectural ideas are needed to exploit their high performance. Future super computers (a la IBM 360/195, Texas Instruments ASC, and CDC 7600) could take advantage of this technology. The best application area is in small capacity and/or high performance storage. CCD's are easy to interface and can be made modularly expandable. Their cost per bit will be approximately two to four times less than the prevailing MOS RAM cost [2]. In conjunction with a RAM buffer, CCD's can be used to provide cost-effective fast auxiliary memory in computers where a large low-cost virtual address space is desired at an accompanying small degradation in overall computer speed. The pervasiveness of this market will depend strongly on the price ratio of CCD's to RAM's. CCD's can also be used very effectively as an intermediate level storage between fast RAM's and slow, inexpensive moving-head disks.

Fixed-Head Disks

Fixed head disks can be broadly classified into three categories: low capacity/high performance, high capacity/low performance, and medium. The high performance disks usually provide multiple heads per track to reduce the access time. Increase in transfer rate can be achieved by using more than one read/write head in parallel. The average FHD has only one head per track and accesses a single track at any given time by electronically selecting a read/write head. The price increase associated with multiple heads per track and parallel transfer is significant and restricts the capacity to less than 20 megabits. Advances in integrated head assemblies could drive the cost of FHD's down. These disks should eventually be displaced by a less expensive technology like magnetic bubbles or high density MOS RAM's. The high capacity, low performance FHD's will be replaced by MBD's in applications where higher performance at a comparable cost is needed. The modularity of MBD's will also affect the market. If the projected 20 millicent price/bit is achieved for MBD's, the medium FHD market will be seriously threatened. CCD's offer competition at the high performance end if volatility can be tolerated.

Electron Beam Addressable Memories

Electron Beam Addressable Memories (EBAM) are a potential candidate for high capacity fast auxiliary memory. The storage medium is capable of retaining data for over a month in the absence of power. In order to achieve a low cost/bit storage capacity has to be in the vicinity of 100 million bits. This would restrict the applicability of this technology. to applications requiring large capacity, low access time and high data rates.

Random Access Memories

Core is being slowly replaced by semiconductor random access memory. Though MOS has little speed advantage over core, declining costs of LSI memories will make MOS RAM prices comparable to core prices. For this reason core memories will not be considered in this paper. The single transistor per cell MOS RAM has the best potential for high density random access main memory. Declining RAM costs will affect the market for the serial access memories. As RAM costs go down computer designers will tend to put more RAM in their systems, thereby reducing the need for fast, expensive serial access memories. The total system cost-performance may well be better served by an inexpensive medium performance secondary storage.

SYSTEM PERFORMANCE

This section describes a methodology for evaluating the effect of various memory hierarchy alternatives on the overall system performance measured in terms of the instruction execution rate. A simple model of a multiprogrammed computer system with a two-level hierarchy will be presented. Multiprogramming refers to allowing the central processor (CPU) to switch from task to task in order to achieve good utilization of all resources. Typically, a given task is allowed to execute until it requires information from the backup storage or any I/O peripheral. While information is being transferred from secondary storage to primary memory, the CPU performs a task switch and starts execution of another task which references only the primary memory. The degree of multiprogramming is the number of jobs, tasks or processes actively using the main memory resources of the computer system. A discussion of the optimal degree of multiprogramming and its effect on system performance is beyond the scope of this study.

As mentioned earlier, declining RAM costs will tend to increase the size of main memory. However, due to program locality [cf. 3] only a portion of main memory is actively used. The availability of low cost paging devices may offer a cost-effective alternative to the system designer. A smaller main memory can be used in conjunction with a paging store or fast auxiliary memory (FAM) to achieve the same performance as the case in which the entire program resides in main memory. Criteria for cost-effectiveness are explored in this section.

Consider a multiprogrammed computer system C. In this system the entire program is loaded into main memory from a peripheral device like moving-head disk or tape unit. A program is executed until interrupted by the occurrence of an I/O

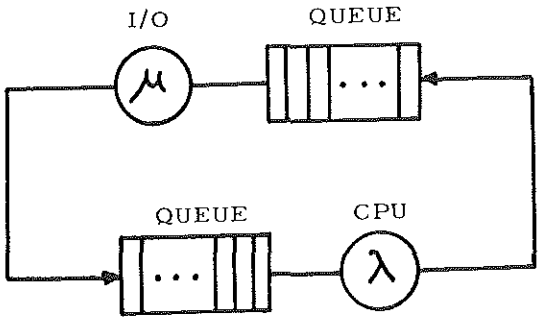

Figure 1 : Queueing model of a multiprogrammed computer system

request to a peripheral device or program termination. Figure 1 shows a queueing model of such a system. This analysis assumes the service times of the CPU and I/O servers to be distributed exponentially with mean $1/\lambda$ and $1/\mu$ respectively.* The CPU utilization, U_c , is given by

$$U_c = \frac{1-\rho^M}{1-\rho^{M+1}} \quad [\text{cf. } 4],$$

where $\rho = \frac{\lambda}{\mu}$, and M is the degree of multiprogramming. M is assumed to be equal to 8 for this analysis. The average execution interval is given by

$$\frac{1}{\lambda} = l \times t_u$$

where t_u is the average time interval between successive memory accesses, and, l is the number of memory references between I/O requests.

*This assumption is widely used in the analysis of the performance of computer systems. It simplifies the mathematics involved while providing useful insight into the dynamics of the queueing phenomena.

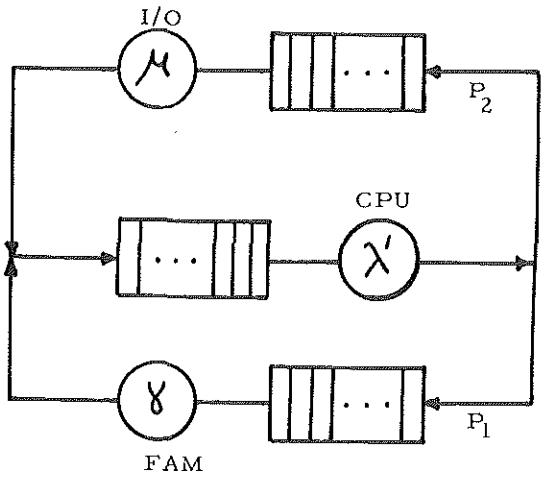

Figure 2 : Queueing model of a paged multiprogrammed computer

An alternate system design, C′ , that uses paging, can be modeled by the queueing network is shown in Figure 2. The cost per bit of main memory is R (>1) times the cost per bit of the paging store or fast auxiliary memory. In this system execution is terminated by an additional event called page fault. The main memory in this system is smaller than the program size. Only a fraction, f , of the program resides in main memory. A page fault occurs when the program references a page not currently in the resident set. Page faults are characterized by the probability of their occurrence -- sometimes also referred to as miss ratio, α . The average number of memory references between page faults is $(1-\alpha)/\alpha$. The rate, δ , at which the CPU accesses the paging store is given by

$$\delta = \frac{\alpha}{1-\alpha} \times \frac{1}{t_u} .$$

Thus, in Figure 2,

$$\lambda' = \lambda + \delta \quad ,$$

P_1 = Probability { Execution terminated due to page fault} = δ/λ' ,

P_2 = Probability { Execution terminated for other I/O } = λ/λ' .

The service rate, γ , of the paging device is equal to $1/T$, where T is the sum of the latency and transfer times. General expressions for the equilibrium distribution of jobs in such networks have been obtained by Gordon and Newell [5] . The use of these expressions gives the utilization, U_c' , of system C' as

$$U_c' = \frac{\sum_{i=1}^{M} (x^i - y^i)}{\sum_{i=1}^{M+1} (x^i - y^i)}$$

where $x = \frac{\lambda}{\mu}$, and $y = \frac{\delta}{\gamma}$.

Note that when a page fault occurs it may be necessary to restore a currently resident page into the paging device before a new page is brought into main memory. In this analysis, such a page fault will be considered equivalent to two page faults. Also, by definition, when α approaches zero the page fault rate δ goes to zero.

Let α^* be the miss ratio for which U_c' is within 5% of U_c . Let the corresponding fractional size of the resident set be f^* (<1). Then, the system C' has the same cost as system C, if the per bit cost ratio of the main memory to paging store is

$$R^* = \frac{1}{1-f^*} \quad .$$

Note that this simplified analysis does not include the cost of implementing a memory management scheme for the paging activity. Thus, the cost advantage has to be greater than that mentioned above in order to justify the added complexity. The value of R^* can be calculated from known miss ratio characteristics of program environments.

The value of α^* is a function of U_c , the utilization of the unpaged system, and the product $\gamma \cdot t_u$, denoted by σ . Figure 3 is a plot of U_c' against α for $\sigma = 0.001$ and U_c varying from 0.5 to 0.9 . Similar curves can be plotted for other values of σ . Based on these curves an approximate rule of thumb can be stated that the critical miss ratio, α^* , is almost equal to the value of σ for less than 10% degradation in performance. The corresponding value of f^* is a function of program behavior.

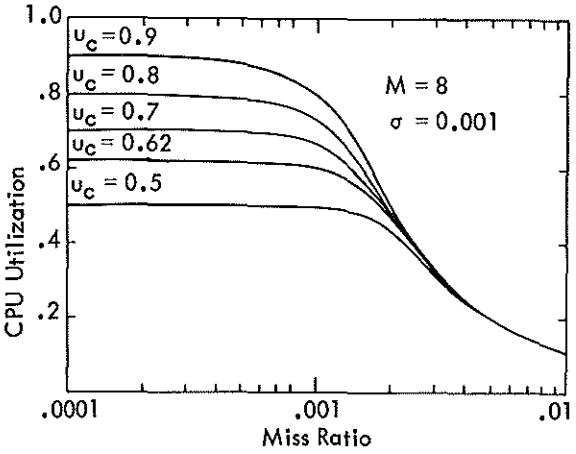

Figure 3. Performance of a Paged System

Table 2.
Critical Miss Ratio

u_c	Critical Miss Ratio α^*		
	σ = .005	σ = .002	σ = .001
0.9	0.003	0.0013	0.0007
0.8	0.004	0.0015	0.0008
0.7	0.005	0.0020	0.0010
0.6	0.006	0.0023	0.0012
0.5	0.007	0.0030	0.0015

CCD memories are expected to be between a factor of two to four times less expensive than equivalent MOS RAM's [2] . Let us assume that the memory management scheme is implemented in hardware at a cost of 10% of the original memory cost. If R is equal to 2 then cost-effectiveness is achieved for environments where $f^* < 0.65$. Table 2 shows the value of α^* that must be achieved at the above mentioned value of f^* .

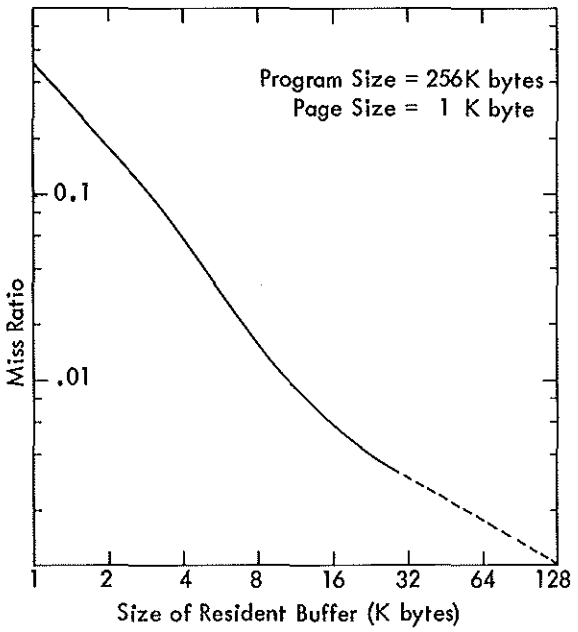

Figure 4. Mattson's Miss Ratio Characteristic

Table 3.
Minimum Cost Ratio for Cost-Effectiveness

u_c	Minimum Cost Ratio R*		
	σ = .005	σ = .002	σ = .001
0.9	1.29	1.92	5.18
0.8	1.23	1.69	3.51
0.7	1.21	1.47	2.5
0.6	1.19	1.38	2.02
0.5	1.18	1.29	1.69

For a given program environment this analysis can be used to predict the minimum cost ratio R* required for cost effectiveness. Figure 4 shows the miss ratio characteristics of Mattson [6] . Table 3 shows the minimum value of R for various system configurations. Once again, a 10% cost is attributed to memory mapping; i.e., $R^* = 1/(0.9 - f^*)$. Since the critical cost ratio ranges from 1.2 to 5.0 , CCD's are likely to be attractive paging devices if the projected cost ratios of two to four are achieved for $\sigma > 0.002$. For typical values of t_u equal to 0.5 and 1 microsecond, this corresponds to a block read time of 250 to 500 microseconds. For CCD register lengths of 1K bit this can be achieved with data rates near 5 MHz.

Further performance improvements can be obtained by maintaining some of the frequently used files such as system tables and directories, compilers, and utility programs in a fast access device such a as CCD's. This reduces the request frequency to the file system which is usually a moving-head disk with access times ranging from 30 to 100 milliseconds. Thus, the CCD now acts as a buffer between the CPU and the file system. Table 4 shows the

Table 4
Performance Improvement due to Buffering of Systems Software

q	CPU Utilization u_c'				
0.0	0.9	0.8	0.7	0.62	0.5
0.1	0.9296	.8516	.7641	.6825	.5533
0.2	.9610	.9058	.8321	.7544	.6194
0.3	.9817	.9795	.8971	.8319	.6995
0.4	.9930	.9782	.9492	.9053	.7925
0.5	.9979	.9929	.9814	.9602	.8875

q = Probability of finding system software in CCD buffer

improvement in utilization due to the CCD buffer when the buffer satisfies a fraction q of all requests to the file system. This improvement is insensitive to the access time of the buffer provided it is at least a factor of 10 better.

The two approaches outlined in this section can be combined to increase the effective size of the main memory. Any performance degradation due to paging can be compensated by the improvement due to permanent residence of certain systems software on the fast paging device.

CONCLUDING REMARKS

The analysis presented in this paper makes many simplifying assumptions in order to derive some general results. There are many page allocation strategies for multiprogramming systems with virtual memory [7]. These strategies are used effectively to obtain the most efficient use of system resources [8]. Some of the issues involved are optimal page size, optimal degree of multiprogramming, time slicing, dynamic memory allocation, page replacement strategies, and scheduling. These problems can be analyzed for more specific system architectures.

A cost ratio of two to four for block read times of 200 to 500 microseconds is likely to result in lower memory system cost using a fast auxiliary memory. The amount saved depends on the size of the memory. The ultimate attractiveness of that cost saving depends on its magnitude relative to the overall system cost.

REFERENCES

1. R.R. Martin and H.D. Frankel: Electronic Disks in the 1980's, Computer, February 1975, pp. 28.

2. J.E. Carnes and W.F. Kosonocky: Charge-Coupled Devices for Computer Memories, Proceedings AFIPS NCC, May 1974, Vol. 43, pp. 833.

3. P.J. Denning: The Working Set Model for Program Behavior, CACM, Vol. 11, No. 5, May 1968, pp. 323-333.

4. F.S. Hillier and G.J. Lieberman: Introduction to Operations Research, Holden-Day, San Francisco, 1967.

5. W.J. Gordon and G.F. Newell: Closed Queueing Systems with Exponential Servers, Operations Research, Vol. 15, No. 2, April 1967, pp. 254-265.

6. R.L. Mattson: Evaluation of Multilevel Memories, IEEE Transactions on Magnetics, Vol. MAG-7, No. 4, December 1971. pp. 817.

7. D.D. Chamberlin et al: An Analysis of Page Allocation for Multiprogrammed Systems with Virtual Memory, IBM Journal of Research & Development, Vol. 17, No. 5, September 1973, pp. 404-412.

8. Y. Bard: Characterization of Program Paging in a Time Shared Environment, IBM Journal of Research & Development, Vol. 17, No. 5, September 1973, pp. 387-393.