New horizons for magnetic bulk storage devices

by FRANK D. RISKO

Bryant Computer Products
Walled Lake, Michigan

INTRODUCTION

There are many types of auxiliary bulk storage devices available on the market now and certainly many more will be available in the future. The second, and in some cases, third generation bulk storage devices are making their appearance on the computer peripheral market now. These devices, much to the surprise of many people, are still rotating magnetic media devices. This paper discusses only state of the art bulk storage devices of the magnetic recording variety. However, since a look at the future is in order, toward the end of the paper, we will attempt to conjecture as to what future memories may look like.

The question of using bulk storage evolves around certain major factors. One factor is cost. This is the primary reason why such items as micrologic or integrated circuit[1] flip-flops have not emerged as bulk storage devices. The cost per bit of these types of devices is extremely high and even with LSI (Large Scale Integration) or woven screen memories, it is still an expensive proposition. Magnetic thin film devices and magnetic cores are both used in most computer systems for the so-called scratch pad, or high speed memory. But again, these devices—in terms of bulk storage—are much too costly.

To put the area of bulk storage into perspective, this paper considers small storage as those of less than five million bits. (The majority of these are head per track devices.) The next memory size range considers a medium range memory to go from 5 million to 50 million bits. (This is the beginning of the positionable-head devices.) Large magnetic memories are then considered from 50 million to 500 million bits. Bulk memories obviously fall above the 500 million bits capacity and have ranges going up to 10 billion bits. Memories above this range which could go up to the trillion* and possibly the quadrillion bit range should be considered large bulk memories. Since a common term makes comparison easier, the bit has been generally used in this paper because character length is manufacture-dependent. When appropriate, character is used and is assumed to have eight bits.

Since the above ranges of memory—particularly in magnetic storage devices—cover a lot of applications, discussion is limited first to state of the art memories in bulk storage category of the contact and non-contact type and then extends these to the near future and beyond. Many papers have been presented in the past[1-7] which have documented many aspects of bulk memories. This paper touches only on some of the more inportant aspects which are affecting present-day thought and equipment design.

Performance factors

There are two major performance factors which must be considered in the discussion of bulk storage devices. These performance factors can be defined as hardware-derived and system-derived.[4] The hardware derived performance factors consider the major areas of:

- — Capacity
- — Access Time
- — Latency
- — Data Transfer Rates
- — Cost per Bit (or bits per dollar)

The system derived performance factors consider such items as:

- — Throughput

*Possibly a new unit of measure such as the Mega-Mega Bit (MMB) for trillion bit capacity should be used since the "mega" is already an accepted term.

— **Indexing Procedures**
— **Memory Allocation**
— **Chaining Provisions**
— **File Activity**
— **Provisions for Queuing**
— **Checking Techniques**
— **Format**

Each of the various types of memory devices such as tape, disc or drum excel in one of the areas listed above more so than in another. Obviously the most important points must be considered in order to optimize the utility of a particular type of system. Because it would be almost impossible to discuss all of the just mentioned factors, this paper is limited to throughput, memory allocations, file activity and format.

Hardware factors

Capacity

Capacity is considered as twofold since this could be taken as total fixed on-line capacity, or total unit capacity, where the unit could be a reel of tape or a replaceable disc, or disc pack. The disc has major advantages when considered as a bulk store. Using any of these three devices it is possible to have virtually rooms full of records up to an almost infinite store. However, the access time of a particular record in that store includes the operators' "fetch" time when comparing it to an on-line fixed store. Most people do not consider "fetch" time when comparing storage devices. Further comparison between these devices is not meaningful because the disc pack outperforms the tape in the area of random access and longevity, and all of these devices have a limited on-line unit capacity.

Consider the fixed capacity on-line store. If enough units (floor space provided) can be put in tandem with the proper controller, it is possible to have in excess of five billion characters available. This means "wall-to-wall" data cells, disc files or drums.

Access time

Access times for very large stores are generally longer than 75 milliseconds. The reason for this is that they are generally moving head devices, or moving strips such as CRAM. The simple fact that an electro-mechanical or electro-hydraulic positioning device is used puts these devices into

TABLE I

MANUFACTURE		UNIT NO.	CAPACITY MILLIONS OF BITS	ACCESS TIME MILLISECONDS MIN.	AVG.	MAX.	END USER COST DOLLARS	BITS/DOLLARS
Bryant	1	4000 Mod 1	750	60	140	180	260K	2,880
	2	4000 Mod 2	1600	60	160	205	305K	5,240
	3	4000 Mod 2a	4000	60	160	205	440K	9,100.
	4	4000 Mod 2b	5000	60	160	205	480K	10,400
	5	Phd 340*	340	30	70	105	125K	2,720
	6	Phd 170*	170	30	70	105	105K	1,820
	7	CD-145*	145	-	17	33	210K	690
CDC	9	818	3750	40	146	226	310K	12,100
	10	814	1200	34	65	110	205K	5,850
	11	6638	1000	25	70	115	325K	3,080
Data Products	12	dp/f5045111	870	50	150	250	227K	3,840
	13	dp/f5085	5000	15	85	150	450K	11,100
Burroughs	8	B475	102	-	20	40	95K	1,080
IBM	14	2302-4	1800	50	115	180	355K	5,070
	15	2311-12	43	30	98	185	26K	1,650
	16	2314	1656	25	88	165	252K	6,580
	17	2321**	3200	50	550	600	140K	22,800
NCR	18	353-5**	360	3	110	125	58K	6,220
RCA	19	70/568-11**	4488	136	500	550	238K	18,800
	20	2488-8**	2400	130	500	640	135K	17,800
Univac	21	Fastrand FRII*	920	5	92	155	168K	5,480

* DRUM
** STRIPS OR CARDS
All other units are discs.

the long access time ranges. Some head per track devices are only limited by rotational latency for access time, but they generally fall into the large storage area rather than the bulk storage area as we mentioned earlier.

If we consider the larger replaceable disc drive type device, such as the IBM model 2314 with eight disc packs and single data channel operation, we find that we have multiple-seek features for the 3.2-billion bit storage. If, for a moment,

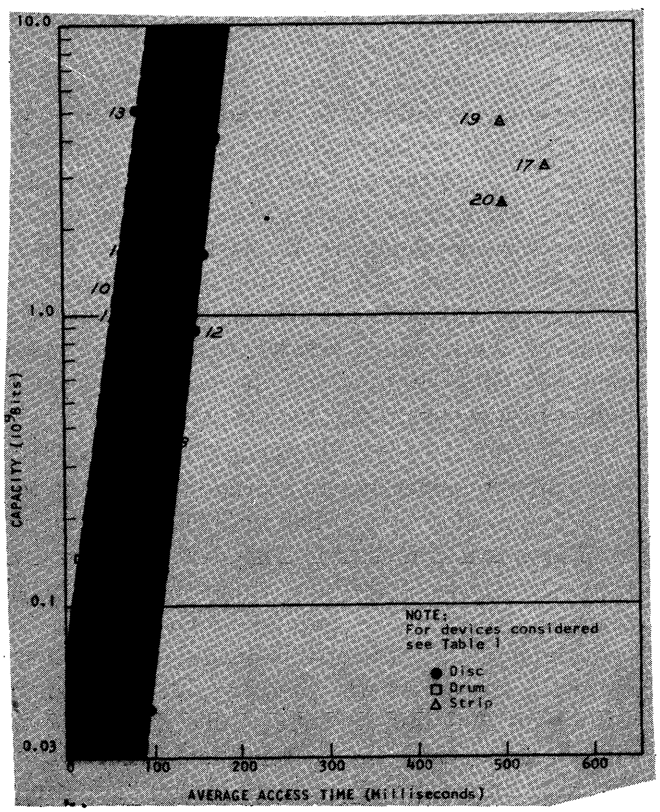

FIGURE 1-Bulk store—capacity vs. access time

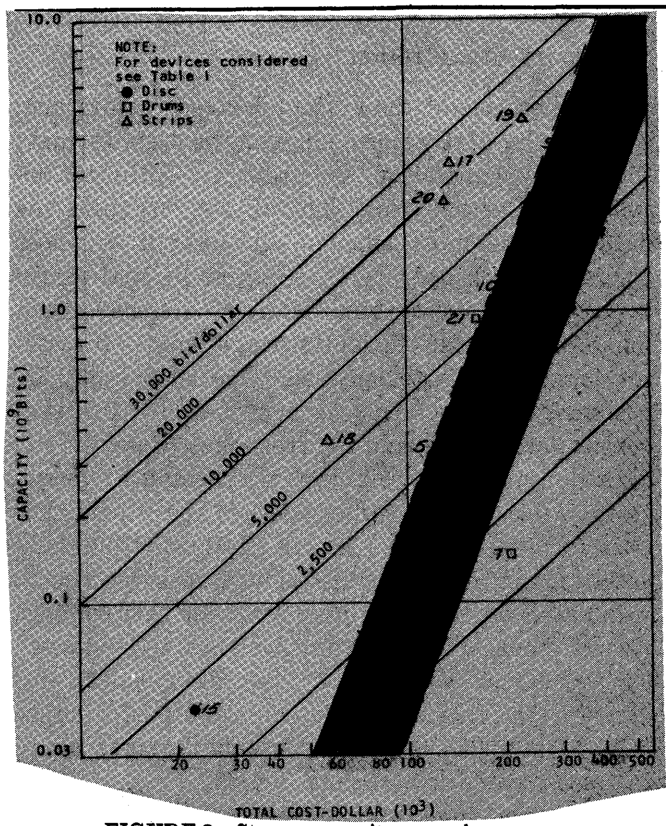

FIGURE 2 Storage capacity vs. end user cost

we want to consider only cost versus access time, we can get the cheapest bulk storage device by buying a unit which has very low cost, but with a long access time. The IBM model 2321 gives 22,800 bits per dollar but at 550 milliseconds access time. A similar price comparison for a Bryant Series 4000 Model 2B Disc File gives us 10,400 bits per dollar and an average access time of only 160 milliseconds. In other words, on one hand you can reduce the cost of a memory system by 50%, but the access time is increased by over 300%. See Table I for a comparison of various manufacturers capacities, access time and costs.

To fully appreciate some of these hardware performance characteristics, we may examine Figure 1 which shows the capacity of the various electro-mechanical types of devices versus access time. This area shows that the higher the capacity of a given store the longer its average access time. Naturally, everyone wants not only a larger store, but wants access to that store in the minimum amount of time, so the area trend line can be expected to move to the left in the future. Figure 2 indicates the capacity versus end-user cost on a bits-per-dollar basis. Again, the trend is end-user cost reduction with increasing capacity. The trend area shows the cost by capacity of the various types of mass memory devices. Note that the greatest economy is gained with strip-type devices at an increase in access time. Also, that all the units outside the trend area to the left in Figure 2 are strip devices. However, the need for increasing speed to obtain data will reduce this trend and cause the industry to move toward the faster access devices into the high bits per dollar range. One of the major problems with the strip device is its poor reliability and limited magnetic coating life.

Latency time

Latency time and data transfer rates are inseparable. Latency time is defined as the time of a single cycle of the magnetic memory device. Data transfer rate is directly related to how many bits-per-inch have been packed onto the given device surface. Since we are always striving for maximum capacity, we are usually recording at maximum packing density. In order to reduce latency, we need as high a surface speed as possible, which increases the transfer rate. In a typical serial device, the data transfer rate could be five Mega-bits per second. This would be typically a rotating device with a mean diameter of 10 inches, turning at 3600 RPM with data written at 2500 bits-per-inch. Lower surface speed units with parallel data transfer can achieve similar rates.

System factors

System-derived performance factors provide a most interesting, as well as practical aproach to the bulk storage areas. With time sharing users becoming more sophisticated, a larger amount of on-line storage is required. This larger on-line storage, however, necessitates a higher throughput. One of the ways throughput (or transactions per second) can be improved upon in a bulk store is to increase the total number of time sharing accesses to that store.

Throughput

Consider for a moment an imaginary time-sharing installation which has a given number of users. To use this computer on a time-share basis it should be obvious that there must be some amount of storage allocated for the users. Storage requirements are a function of how the user is actually achieving his necessary results with his so-called "own computer." A sophisticated user would probably need on the order of 200,000 characters of storage as opposed to a less sophisticated user who would only require on the order of 50,000 characters of storage. Generally this means that somebody has to wait during the access-time period depending on queuing of the various computer user requests. It could be feasible to reduce waiting time for a given number of users by doubling the computer store access capability. (Additionally, security must be provided so users cannot access each other's data.) Since these people are not accessing the same data, it then seems practical for each user to have his own independent access mechanism. This could be taken care of very easily by going to a group of disc drives (let us say 50 drives for 50 users) in our imaginary time-share system. Obviously, the cost of such a system would be prohibitive because the time-share user cannot afford this luxury.

The next best thing would be to have a number of accesses of positioners on the same device. Some average number of users will always be on-line (let us assume this number to be 25). With 25 on line, how many will be needing access to the bulk storage continuously during any given period of time? The number would probably average 10 to 15. Because of the very nature of the binary application, it would be most feasible to have a binary number of accesses. This may be 8 or 16 positioners. Let us now assume that we do have 16 available; this allows the system performance to increase by a factor of 16 (in regard to throughput) over what a normal single positioner machine would give. One of the unfortunate features of a time-shared system is the fact that at a given point in time we usually have more than one user wanting to access the same positioner. No end to the dilemma. . . .

Memory allocation

Having more than one user accessing a device, in a time-share application, memory allocations for each user is limited in order to serve as many people as possible and still not over-extend the available memory. This is particularly important because the computer itself will want some of the available memory for internal executive and swapping programs. Therefore, it is sometimes necessary for time-share systems to redistribute data in order to make a more economical access and thereby maintain reliable throughput time.

Memory organization

Another side of the system-derived factors are: what is the best record length; the best format; the best organization, in order to optimize the system? The user may not always organize his data for the most efficient throughput or most efficient search operations, so an additional strain is put on the total system. Many studies have been made of the most efficient storage length (or sector length) and storage use. One such study by Brenner[8] discusses this as a probability function versus data length, and the frequency of use versus data length. In weighing performance versus storage use, Brenner showed that for a performance of 90%, a record length of 600 to 700 characters is optimum. His decision to use this length was based on the absolute change in performance versus storage use. Figure 3 shows that storage use decreases very quickly as data length increases. In other words, to achieve a performance approaching one hundred percent, it is necessary to have a very long record length.

Of course, one of the major problems with record length is the fact that storage use de-

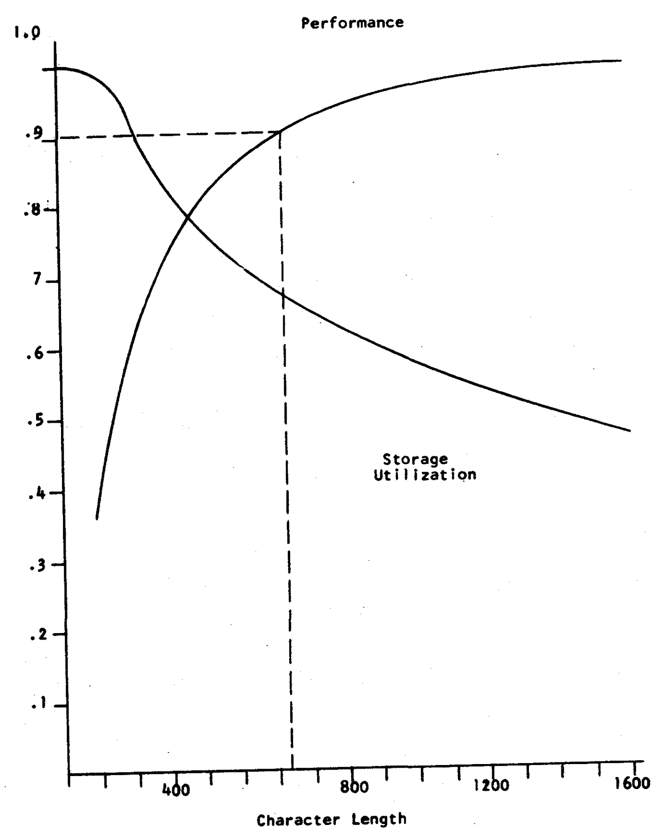

FIGURE 3 Performance and storage utilizati n

creases as length increases because of various over-head functions. This is reflected in an independent (unpublished) study which showed that the most feasible length is approximately nine-thousand bits. In order to make the bit quantity binary, the amount with the nearest factor is actually 256 twenty-four bit words, which gives us approximately 6000 bits as the most useful (or optimum) length. This study is in approximate agreement with the study made by Brenner.

The system and programming aspects must be extended for a given unit in order to allow for the proper indexing, chaining and queuing techniques which are required by the mechanical portions of the storage device. Obviously a head-per-track, parallel system would provide the ultimate in all these areas, but would be very expensive because of the electronics required to handle such a system.

Present and future trends

The present hardware state of the art of the bulk memory field has decreased to 5-mil track widths with 7.5 mil centers between tracks while packing densities have increased upwards to 2500 bits-per-inch. The 5-mil track widths give approximately 128 tracks-per-inch. This gives us a quarter of a million bits per square inch storage capacity as the state of the art magnetic recording of today. However, magnetic recording development is nowhere near being finished in this particular area.

Magnetic recording in the order of 20,000 bits-per-inch have been demonstrated[9,10] and will, no doubt, be the future packing densities by the mid-1970's. Additional hybrid electronic circuits and recording methods will definitely be necessary to achieve these ends. At the present time we're talking of practical playback levels at the head on the order of one to three millivolts. By 1970 to 1975, the playback ranges will be down to the order of 10 to 30 microvolts. This will require amplifiers which must be absolutely immune to noise and spurious signals. These values, I might add, are being achieved experimentally today in the laboratory and this is where we were approximately five years ago. I believe during 1961 and 1962 that Shew from IBM was doing magnetic recording on the order of 1000 to 1500 bits-per-inch under laboratory conditions working with 1-mil wide pole pieces.[11] At that time, state of the art was generally 300 bits-per-inch.

Positioners

A major concern about use of the 1-mil pole piece is that of positioning tolerances necessary so that off-track and peak-shifting problems do not affect playback. The positioning mechanism will have to be at least two orders of magnitude better than present day standards in order to accommodate this type head. Development work at Bryant shows that the next generation bulk stores will make use of a 2.5 to 3 mil track and will be spaced on 3.75-mil centers. This means that we should achieve approximately 256 tracks-per-inch. Considering only a modest increase in packing to 4000 bits-per-inch, we will arrive easily at one-million bits-per-square inch by 1970. The major breakthrough, of course, is the positioning device itself, which must be fantastically more accurate in the area of track selection and positioning repeatability.

To continue our view of future magnetic devices, I would venture to say that a minimum of 16 multiple access moving head-bulk stores, with total capacities of 50 billion bits, should be with us by 1972. The capacity from then on will probably increase logarithmically and the next jump, I would say without any hesitation, would be from 50-billion to 100-bilion bits and then from 100-billion to a 100-trillion bit magnetic memory by 1975. This again will be a magnetic memory in which certain as yet unimagined manufacturing techniques will have been achieved.

Head developments

By 1975 the head will not be a discrete individually-assembled element but will be batch-fabricated heads similar to today's micrologics. The preamplifiers will no doubt be deposited as part of the head assembly. I also believe that many new techniques will have to come from the metallurgical and the chemical industries to give us the needed techniques and materials in order to achieve 16 to 32 heads per pad without interfering crosstalk and frequency limitations. At the present time there are many devices on the market which have from 9 to 20 heads per pads, the Burroughs disc file for one, the Data Disc and Bryant units (which have 9) for another.

The person who wants a tremendous bulk store

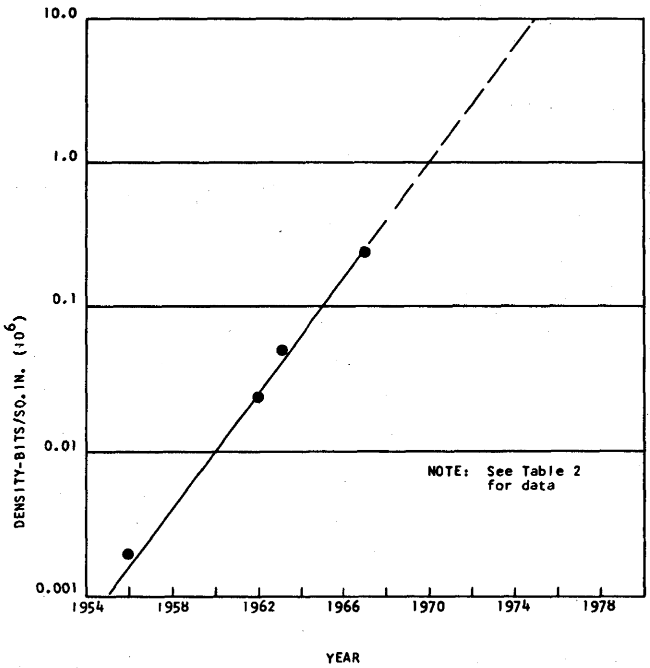

FIGURE 4 Packing density vs. year

is not going to pay any more in 1975 for a trillion (Mega-Mega) bits than he would for a trillion bits right now. The major difference will be that a user will probably only need 1/16 the floor space to hold that device and, no doubt, at a much lower power consumption. The growth of the industry as shown in Table 2, is plotted as a function of packing density (bits/sq. in.) versus the year associated with it.[4-6] This is shown in Figure 4, and is interesting, because the growth has been increasing by an order of magnitude every five years from 1956. At this rate a million bits-per-square inch will be commonplace by 1970.

TABLE 2

YEAR	DENSITY (Bits/Sq. In.)
1956	2,000
1962	25,000
1963	50,000
1967	250,000

The obvious area of improvement along with the batch fabrication techniques envisioned for positioning head devices will be to position the heads over a much shorter distance. The shorter this distance the greater the improvement in access time. Some present day devices move as much as two inches and some as little as a quarter of an inch. I believe that our 1975 devices will use electrical positioners and that the distances traveled will probably be on the order of less than 50-thousandths inch. The reason is that the track width being only 1-mil, or possibly 2-mils, will not allow a positioning accuracy for movements greater than that.

If we consider the present cost of these various devices the future cost ratio will not really change. In other words, the cost per bit is not going to go up appreciably, nor is it going to drop. The manufacturing and material costs will increase, causing the unit cost to increase, but, by the same token, the technology in terms of packing density and bits-per-square inch, as well as bits-per-cubic inch (which, of course, plays an important role in floor space) will also change such that the average cost per bit (or bits per dollar) will not change, but will remain relatively constant.

CONCLUSION

Magnetic type devices are definitely here to stay for another ten years before they give way to the laser/optical type devices. The disc packs provide useful bulk storage and will continue to dominate the small-storage market. The disc packs have a definite place in the bulk memory market although in a different way. The on-line bulk store will be more necessary in the future and will be available with multiple accesses. Some of the problems to be overcome are reliability and improved performance and elimination of the redundant store which is necessary for large systems. Another area of growth will be the medium range (head per track) devices for use in swapping requirements for time-share installations.

Optical type devices are a definite way to go, but they may get competition from LSI devices. The photochromic and laser or optical-type devices, which will far exceed total capacity per square inch (25 million-bits-per-sq. in.) have a long way to go to reach the present state of the art of magnetic recording.

Although they are capable of larger capacities in smaller areas, the technology to develop them is still in the future. Magnetic recording of digital information has a good solid 15 years, maybe 18 years, of technology behind it, whereas attempting to employ lasers in reading and writing (*not just read only*) is a tremendous technological achievement that must be made before the laser technique or beam-memory techniques will become feasible.

At the present time there are many laser read-only memories in operation. I don't want to detract from this achievement, but I feel very strongly that—in order for them to achieve the same state of the art as the magnetic recording media has now (and will have by 1970)—some major technological advances will have to occur. Major research effort is being devoted to this area and the advancement could occur sooner. The holographic displays, although they haven't been publicized much lately, appear also to be feasible and will probably give the laser photochromic devices a good run for their money. However, they still must be nondestruct read/write.

I would like to quote Marvin Camras from the 50th anniversary issue of the IRE Proceedings in 1962[12] extrapolating progress to the year 2012.

> "The first magnetic memory devices resembled tape recorders, just as the first automibiles resembled horse carriages. Eventually tape recorders evolved into the standard memory pack which is presently manufactured in large quantities by specialists. The memory pack of today (2012) is a sealed box about the size of a package of playing cards. It holds upwards of 10^{20} bits of information, and has no mechanically moving parts; the recording, readout, and scanning are all electronic. The storage density is so high that the information of entire libraries is condensed into a few cubic feet. Additional modules are added to extend the capacity to any required level."

I would say we still have a long way to go to achieve that prediction, but the way things can and do happen, a major breakthrough could at any moment occur; I certainly do believe that it can be done.

REFERENCES

1 R A HENLE L O HILL
Integrated computer circuits-past present and future
Proceedings of the IEEE December 1966 pp 1849-60

2 L C HOBBS
Effects of large arrays on machine organization and hardware/software tradeoffs
Proceeding-Fall Joint Computer Conference 1966 pp 89-96

3 L C HOBBS
Present and future state-of-the-art in computer memories
IEEE Transactions on Electronic Computers vol EC-15 August 1966 pp 535-549

4 A S HOAGLAND
Mass storage revistied
AFIPS Proceeding-Fall Joint Computer Conference 1967 pp 235-60

5 L C HOBBS
Review and survey of mass memories
AFIPS Proceeding-Fall Joint Computer Conference 1963 pp 295-310

6 T H BONN
Mass storage A broad review
Proceedings of the IEEE December 1966 pp 1861-70

7 N NISENOFF
Hardware for information processing system: today and in the future
Proceeding of the IEEE vol 54 no 12 December 1966 pp 1820-35

8 F H BRENNER
On designing generalized file records for management information system
AFIPS Proceeding-Fall Joint Computer Conference 1967 pp 291-303

9 K A NORRIS
A technique for high density digital magnetic recording
Presented at the Western Electronic Show and Convention August 24 1967

10 C W STEEL J C MALLINSON
A computer simulation of unbiased digital recording
Abstracts of 1968 Intermag Conference Washington D C p 3 3

11 L F SHEW
High density magnetic head design for non-contact recording
IRE Transactions on Electronic Computers December 1962 pp 764-83

12 M CAMRAS
Magnetic recording and reproduction-2012 AD
Proceeding of the IRE vol 50 no 5 May 1962 p 639

BIBLIOGRAPHY

1 M S FINEBERG O SERLIN
Multiprogramming for hybrid computation
AFIPS Conference Proceedings 1967 vol 31 Washington D C Thompson Book pp 1-13

2 L R GLINKA R M BRUSH A J UNGAR
Design through simulation of a multiple-access information system
AFIPS Conference Proceedings 1967 vol 31 Washington D C Thompson Books pp 437-47

3 A S HOAGLAND
Mass storage
Proceedings IRE vol 50 May 1962 pp 1087-92

4 R L PETRITY
Current status of large scale integration technology
AFIPS Conference Proceedings 1967 vol 31 Washington D C Thompson Books pp 65-85

5 C S POLAND
Advanced concepts of utilization of mass storage
Proceedings of IFIP Congress 65 vol 1 Washington D C Spartan Books 1965 pp 249-54

6 J A RAJCHMAN
Computer memories—possible future developments
RCA Review
vol 23 June 1962 pp 137-51

7 D E SPELIOTIS
Magnetic recording theories; Accomplishments and unresolved problems
IEEE Transaction and Magnetics vol MAG-3 September 1967 pp 195-200