AD-A033 715 WESTINGHOUSE DEFENSE AND ELECTRONIC SYSTEMS CENTER B--ETC F/G 9/2
MASS MEMORY ORGANIZATION STUDY.(U)
SEP 76 D L FARNSWORTH, C P HOFFMAN, J J SHUTT F30602-75-C-0317
UNCLASSIFIED NYD-56237-A002 RADC-TR-76-254 NL

1 OF 3
AD
A033715

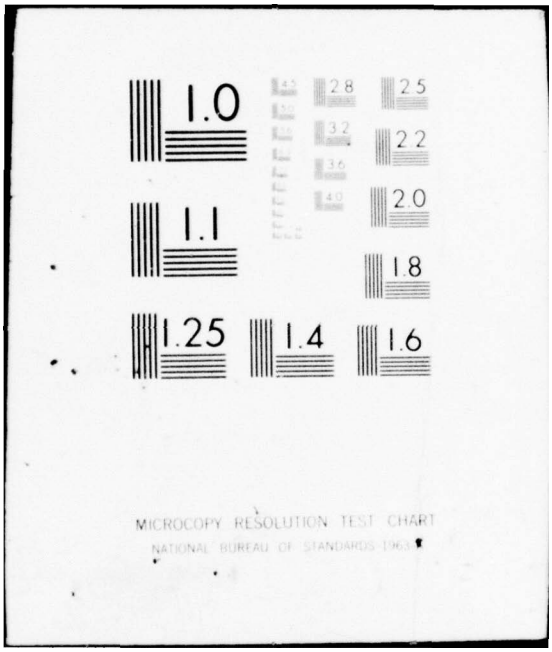
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963

ADA033715

RADC-TR-76-254
Final Technical Report
September 1976

MASS MEMORY ORGANIZATION STUDY

Westinghouse Electric Corporation

Approved for public release;
distribution unlimited.

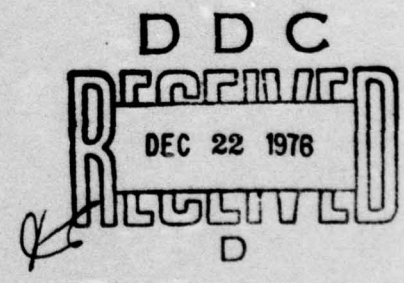

ROME AIR DEVELOPMENT CENTER
AIR FORCE SYSTEMS COMMAND
GRIFFISS AIR FORCE BASE, NEW YORK 13441

This report has been reviewed by the RADC Information Office (OI) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public including foreign nations.

This report has been reviewed and is approved for publication.

APPROVED: William J. Hofstra

WILLIAM J. HOFSTRA
Project Engineer

APPROVED: Robert D. Krutz

ROBERT D. KRUTZ, Col, USAF
Chief, Information Sciences Division

FOR THE COMMANDER: John P. Huss

JOHN P. HUSS
Acting Chief, Plans Office

Do not return this copy. Retain or destroy.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE	READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER: RADC-TR-76-254	2. GOVT ACCESSION NO. / 3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle): MASS MEMORY ORGANIZATION STUDY.	5. TYPE OF REPORT & PERIOD COVERED: Final Technical Report. July 1975 - May 1976.
	6. PERFORMING ORG. REPORT NUMBER: NYD-56237-A002
7. AUTHOR(s): D. L. Farnsworth, C. P. Hoffman, J. J. Shutt	8. CONTRACT OR GRANT NUMBER(s): F30602-75-C-0317
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Westinghouse Electric Corporation, Defense and Electronic Systems Center, Baltimore MD 21203	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 63728F, 55500128, 01
11. CONTROLLING OFFICE NAME AND ADDRESS: Rome Air Development Center (ISCA), Griffiss AFB NY 13441	12. REPORT DATE: September 1976
	13. NUMBER OF PAGES: 284 (277 pp.)
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office): Same	15. SECURITY CLASS. (of this report): UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE: N/A

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

Same

18. SUPPLEMENTARY NOTES

RADC Project Engineer: William J. Hofstra (ISCA)

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Mass Memory, Secondary Memory, Backing Memory, Storage Device, Mass Store, Secondary Store, Backing Store, Associative Processor, Multiprocessor, Parallel Processor, Configurable Processor, Multi-Port Memory, MNOS, STARAN, Block Accessed, Block Oriented and Nonvolatile

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The purpose of this study is to develop a recommended organization for gigabit storage device having speed/bandwidth capabilities compatible with an associating type computer processor. Data acquired under this effort is based on the use of nonvolatile MNOS technology integrated circuits. Full exposition is made of the design principles applicable to production of mass memories having optimum data rates and minimal costs for future associative processor applications. Extensive appendices are included on technology impact factors.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

390 358

bpg

EVALUATION

This contractual effort forms the basis for a more detailed design of a computer backing memory having the capacity of bulk storage devices and the speed/bandwidth capabilities associated with array processors. It will lead to the ultimate realization of such a memory for use with the STARAN Associative Processor in the investigation of techniques capable of solving high data rate problems for the Defense Mapping Agency and the Air Force in conformance with RADC TPO V, C^3 System Availability.

William J. Hofstra

WILLIAM J. HOFSTRA
Project Engineer

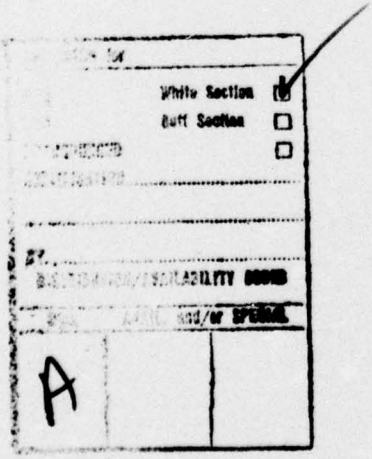

DDC
RECEIVED
DEC 22 1976
D

TABLE OF CONTENTS

	Page
1. INTRODUCTION	1
1.1 Study Format	2
1.2 Program Overview	6
2. DEPLOYMENT DIRECTIVES	13
2.1 STARAN Overview	13
2.2 Interface Configuration	19
2.3 Control Provisions	25
2.4 Timing Considerations	30
3. DETAILED INVESTIGATIONS	37
3.1 Global Hierarchy	37
3.1.1 Baseline Module	40
3.1.2 Intelligent System	44
3.1.3 Full Capability Configuration	51
3.2 Storage Organization	54
3.2.1 Access Modularity	55
3.2.2 Data Routing	62
3.3 Data Management	76
3.3.1 Access Functions	81
3.3.2 Instruction Repertoire	84
4. RELIABILITY EVALUATION	89
4.1 Modeling Approach	92
4.2 Detailed Analyses	104

	Page
5. EXPLORATORY DEVELOPMENTS	113
5.1 Device Test Stand	113
5.2 System Test Bed	116
5.3 Miscellaneous Breadboarding	118
6. CONCLUSIONS AND RECOMMENDATIONS	123
6.1 Technology Prospects	124
6.1.1 Support	124
6.1.2 Status	125
6.1.3 Grooming	126
6.2 System Recommendations	128
6.2.1 Electrical Design	128
6.2.2 Physical Packaging	132
6.2.3 Prototype Development	137
6.3 *Prototype Follow-On Plan*	140

LIST OF ILLUSTRATIONS

Figure		Page
1	Mass Memory Organization Study Plan	4
2	Computer Function Timing Charts	8
3	Idealized SIMD Multiprocessors	8
4	STARAN Computer Interfaced to Recommended MNOS Mass Memory	10
5	STARAN Associative Array Processor Block Diagram	14
6	STARAN System Dual-Channel Communication Network	18
7	STARAN Interface Channel Line Repeater Network Definition Diagram	20
8	Dual-Rail MM Repeater Interface to STARAN	22
9	Line Repeater Functional Equivalent Circuits	22
10	STARAN/Mass Memory Black Box Interface Definition Diagram	26
11	STARAN DMA Channel Timing Relationship	31
12	STARAN PIO Channel Timing Relationships	32
13	PIO Channel Read Transfer Sequence Diagram	34
14	PIO Channel Write Transfer Sequence Diagram	34
15	Direct Access Baseline MM Module Block Diagram	41
16	Indirect Access Intelligent MM System Block Diagram	47
17	Associative Search Access Directory Stack Block Diagram	49
18	Directory Stack Functional Sequence Diagram	50
19	Application Configured MM System Block Diagram	52
20	Sliding-Window Access Concept Diagram	56
21	MM Storage Block-Access Organization	58

Figure		Page
22	Chaptered Text Modularity Visualization of 1 Gbit MNOS Mass Memory	61
23	Port Generalized Cross-Bar-Switch Data Routing Network	64
24	Port Bounded Cross-Bar-Switch Data Routing Network	66
25	Port Coordinated Data Routing Network	68
26	Port Dualed Main Memory Bussing Structure	70
27	Full Permutation Capability Triangular Data Routing Network	71
28	Simple (End-Around) Ordered-Combination Data Translation Network	73
29	Dual-Porting Storage Module Interface	75
30	Indirect Memory Address Directory Record Fields	77
31	MM Indirect, Block-Relative, Post-Indexed Addressing	79
32	Access Time Considerations	80
33	DMA Data Channel Instructions	85
34	Nonmemory Reference Instruction Format	86
35	Comparison of Typical Error Rate Experience	91
36	Augat Drawer Assembly	95
37	Thermal Management Constraints	96
38	Reliability Analysis "Unitized" MM Block Diagram	100
39	Developmental MNOS Device/Board Exerciser (Top) and Sample TV Monitor Display Pattern (Bottom)	114
40	MNOS Device Screening Test Patterns	115
41	Test Bed MM System Organization	117
42	Dual-Port Application Groomed 16 Kbit MNOS Memory Chip	127
43	Mass Memory Organizational Block Diagram	129
44	Port Buffer Unit Organizational Block Diagram	131
45	Mass Storage Unit Organizational Block Diagram	133
46	Two-Port Bussing Organization of Memory Planes	134
47	Full-Capacity, Dual-Ported, Commercially Packaged Mass Memory System Configuration Diagram	135

LIST OF TABLES

Table		Page
1	AP Mass Memory Organization Study Guidelines	12
2	STARAN CIOU to MM Interface Line Listing	24
3	Projected 16 Bit DMA Device Selection Address Assignments	26
4	STARAN PIO Transaction Sequences	30
5	Baseline MM Instruction Primitives	43
6	MM Interconnect Technique Tradeoff Summary	75
7	DMA Address Channel Data Storage Mode Op-Code Table	83
8	DMA Data Channel Status Modify Extended Storage Mode Op-Code Table	83
9	MM Reliability/Maintainability Assessment Factors	90
10	Reliability Modeling Terms	99
11	List of Principal MM Quarter-Word SECDED Parameters	103
12	MM System Reliability Model Relationships	108
13	Nonredundant Unit Reliability Model Relationships	108
14	PBU Reliability Model Relationships	109
15	DMU Reliability Model Relationships	110
16	MSU Reliability Model Relationships	111
17	MM Ground Benign Reliability Analysis Summary	112
18	MM Ground Fixed Reliability Analysis Summary	112
19	Panel Switch Entry System Instructions	119
20	Variation in Number of Memory Planes with Capacity & Device Size	138
21	Mass Memory Prototype Guidelines	139

1. INTRODUCTION

The purpose of this study is to definitively identify the organizational design philosophies most suitable to a reasonably priced, ruggedized, random access, 10^9 bit mass memory which is directly compatible with the speed/bandwidth capabilities of an associative type computer processor. Data acquired through this effort is based on the use of nonvolatile MNOS technology devices to provide the storage of large quantities of rapidly accessible data and instructions necessary to the high operations speeds of an associative array processor (AAP). The study reports on the general characteristics of different approaches to memory implementation. Alternatives are compared in the context principally of proven AAP designs (e.g., STARAN) but also are evaluated in light of their expansion capabilities. Information assimilated from the study is used to establish a prototype mass memory storage device design having optimum data rates and minimal costs for future associative processor applications.

The purpose of this section is to provide a degree of familiarity with the background and orientation of the respective investigations and to briefly summarize the study findings. The items covered include an overview of the study procedures followed, a review of the application considerations for an AAP mass memory, and a general description of the fundamental concepts relevant to large data base, high-speed storage devices. Although the directives of the effort have been guided by the near-term memory needs of STARAN based AAP computers, the results obtained appear to be directly extrapolatable to numerous other parallel and multiprocessor systems which are being actively pursued.

1.1 STUDY FORMAT

The organizational sequence of investigations undertaken during the mass memory study is based on the directives set forth in RADC's original statement of work (SOW, see Appendix A). The main thrust of this effort consists of an evolutionary development of the unique organizational principles which are applicable to production of very high capacity MNOS memories optimally configured to dynamically interact on a 1 to 1 speed basis with associative or other types of multiprocessors. Presently, the performance and throughput capabilities of these high-speed computer facilities are not being realized because of the IO bottleneck which exists at the interface to currently available mass storage devices (e.g., drum, disk, and tape units). Although generalized tradeoffs are made, detailed design treatments are directed solely to satisfying the operational requirements of STARAN* S-1000 type AAP computer facilities (see Appendix B). The reason for this is simply that STARAN has an established performance record, it permits definition of a working set of interface conditions by which alternatives can be evaluated, and it is for facilities at RADC incorporating such equipment that the first mass memory will most likely be procured.

Technology and system alternatives are judged primarily from their impact in six broad memory criteria categories. In order of assessment priority, these are:

- Performance
- Reliability
- Cost
- Power
- Physical (size/weight)
- Environment.

---

*TM of Goodyear Aerospace, Akron, Ohio.

Included under the performance category are all factors affecting capacity, access time, data rate, read/write capability, and storage integrity. Reliability evaluations are based on the established practices of MIL-HDBK-217B. Factors considered affecting reliability include: interdevice data transmission error rates, memory and logic component failure rates, total component counts, power down dormancy provisions, exercisable system redundancies, and physical operating environment. In the technology area, cost considerations encompass relative maturity, device availability, and application grooming. These influence the initial mass memory development investment needed. From an aggregate system perspective, costs have been reviewed from the standpoint of both purchase price and cost of subsequent ownership (as reflected by reliability and maintainability). Interest in the power drain, physical parameters, and environmental constraints of a mass memory centers on the manner in which they influence system performance, reliability, and cost.

Figure 1 denotes the principal elements of the approach taken to the mass memory organizational study. Initial phases of the effort consist of an investigation of the unique requirements imposed on interface between an associative or multiple parallel processors and a high-capacity backing store (mass memory), and the manner in which command and control signals are transferred across the interface to affect memory accesses. The investigation concludes with a definitive set of interface condition specifications based on the STARAN system to which follow-on developments are to be addressed. This is followed by a detailed inquiry into a diverse set of detailed memory design concepts. Activities during this period include attendence at the 1975 Sagamore "Computer Conference On Parallel Processing", to ensure considerations of the most recent developments in this field, and the assimilation and coordination of supplemental information supplied by RADC.

Pursuant to this, a series of interrelated examinations and analyses is conducted to establish the advantages and disadvantages of each technology

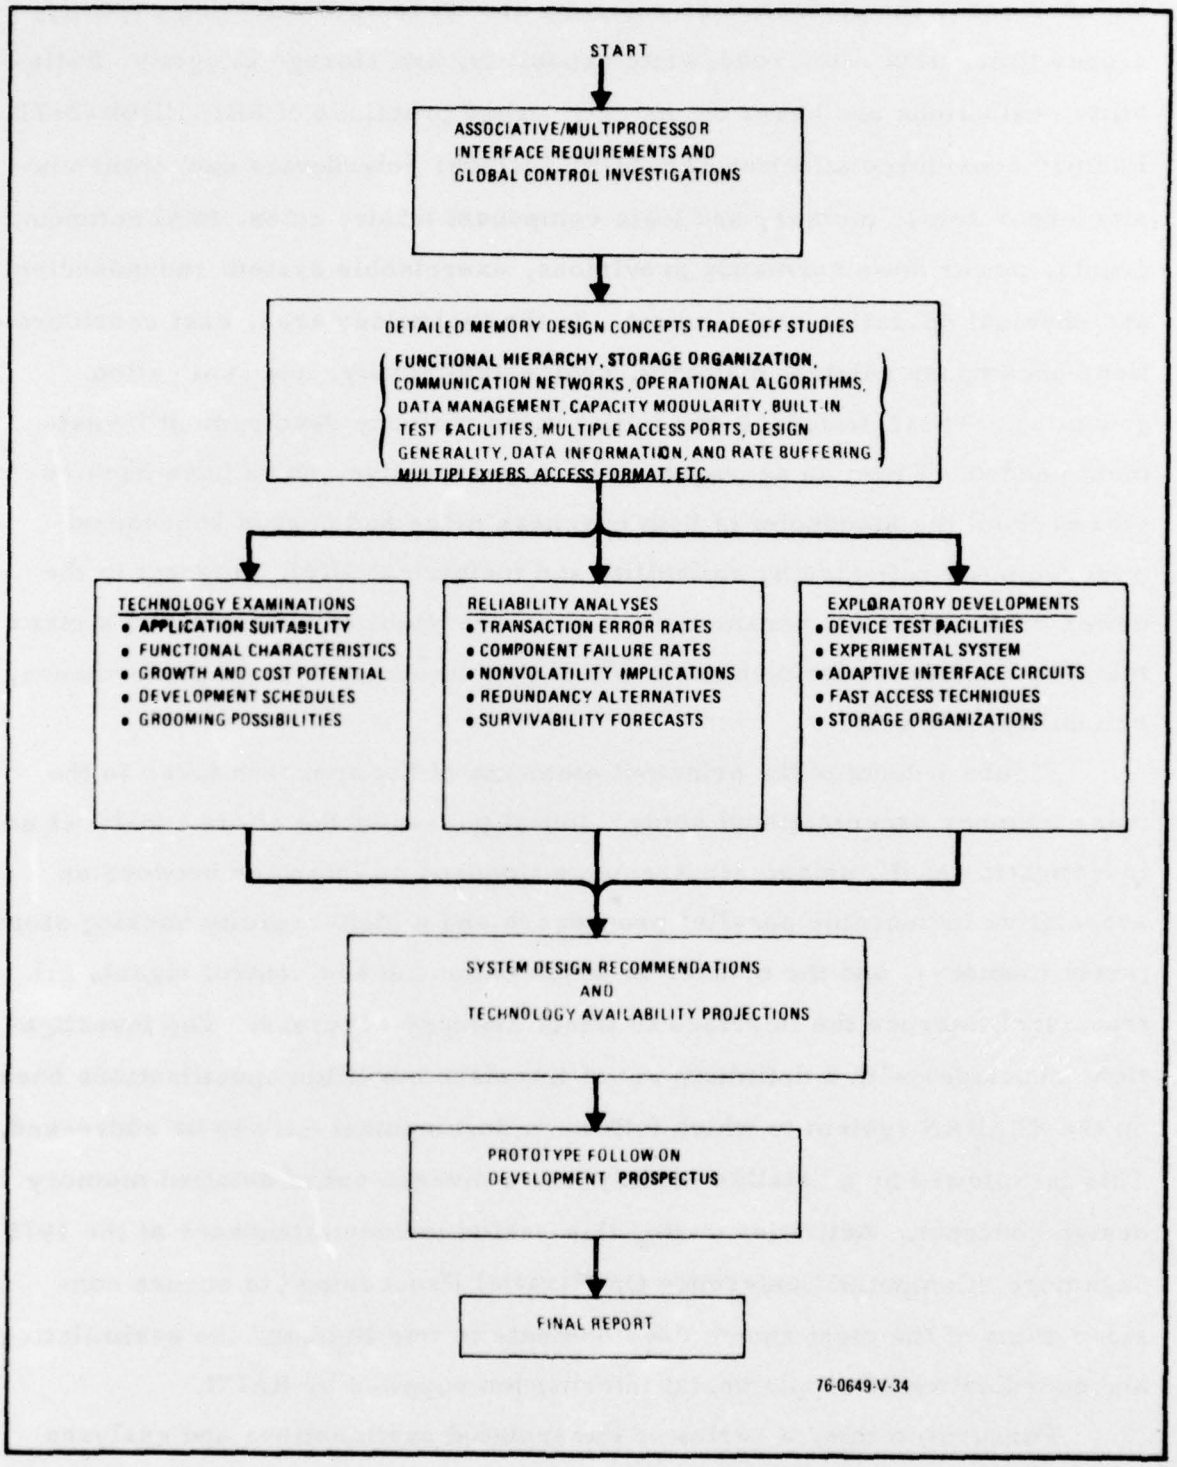

Figure 1. Mass Memory Organization Study Plan

and system option. Care is exercised throughout this phase to consider the impact of technology parameters on system and subsystem design, and to ascertain the true performance, reliability, and cost trades of those designs. The three principal areas addressed are technology, reliability, and exploratory breadboarding. Extended coverage is given in the appendices of this report to the many faceted nature of memory technology selections. Appendix D provides a comprehensive review of technology alternatives. An in-depth description of the prime MNOS technology is presented in Appendix E. Finally, a set of specification charts for a representative 2 Kbit MNOS memory chip presently being manufactured by Westinghouse is given in Appendix F.

Successful implementation of two noteworthy pieces of equipment was realized during this period: a memory chip test and screening unit and an experimental memory system test stand which is interactively exercisable via a CRT control terminal tie-in. Work on the reliability aspects of preliminary system configurations led to major refinements by identifying the areas of weakness. In accomplishing the respective analyses, a useful set of device failure rate tables, reliability analysis equation programs, and miscellaneous other data was amassed. For reference, these items are included in Appendix C. Results from each area of investigation are brought together to arrive at a recommended mass memory organization based on the projected availability of appropriately organized MNOS technology memory devices. This in turn is used to generate a follow-on prospectus for development of a prototype memory.

For purposes of this study, RADC's time frame of interest for undertaking development of a mass memory is mid-1976, with delivery anticipated approximately 3 years thereafter. Considering the extent of system design work involved, there will be approximately 1 to 1-1/2 years lag before the memory components necessary for the system must be available in quantity. Accordingly, the recommended memory approach provides a prototype conceptual design which reflects a 1978 production state of the art,

rather than present day data storage device availabilities. All aspects of the suggested design are derived from study evaluations made in the context of a STARAN AAP application. Technical interchanges with RADC personnel and Goodyear consultants provided the required background information on the deployment parameters.

1.2 PROGRAM OVERVIEW

Through advanced developments of multiprocessing techniques using various forms of distributed logic, major changes have been introduced into computer organizations. These changes result in significant improvements in both computational speed and efficiency of machine programming. Efficiency, here, is considered in the sense of the user program and its performance, not with respect to how busy some internal subcomponent of the system is. This enhanced performance is secured as a direct consequence of the parallel or concurrent operations inherent in a multiprocessor.

Conventional computers accomplish a degree of parallelism that is transparent to the machine language program through the use of Cache type buffer memories and wired-in overlapping concurrency controls. The resulting machine language instruction set is not an ideal vehicle to control the physical resources of the system, since the concurrent control is achieved automatically, without the possibility of any global optimization.

Control of the system resources is made visible to the programmer in a multiprocessor type computer by the microprogramming operations it facilitates. An essential feature of this type of system is the availability of a fast storage medium: fast in this sense being with respect to ordinary combinatorial logic operations. Therefore, the access time of the storage devices must be on the same order as the cycle time of the processor. It is also essential that this fast storage possess the ability to be erased and rewritten with no loss of speed.

Due to the slow access time of current bulk storage devices (discs, drums, etc.), a major bottleneck exists at the I/O interface of the generic

class of multiprocessors, of which STARAN is a member. Consider the timing ing chart of figure 2 (a) for a conventional machine instruction. Due to slow and independent access of the instruction and data blocks, a major portion of the instruction execution time is spent in overhead accessing operations. Contrast this with the situation shown in figure 2(b), which illustrates a multiprocessor type of microprogrammed instruction execution sequence. An implicit assumption is made here of homogeneity of memory; that is, that data and operands are all located in the same fast storage as the instructions. This assumption is valid when the storage is interactively rewritable.

Dominant among the multiprocessor design philosophies that have been explored, and the one on which STARAN is based, is the single-instruction stream/multiple data stream (SIMD) organization. The reason for this is simply that it is more readily optimized for specialized applications than is the more generalized multiple-instruction stream/multiple data stream (MIMD) organization. There are three basic kinds of SIMD machine organizations currently under development: (1) the array processor, (2) the pipeline processor, and (3) the associative processor. Block diagrams of each organization are illustrated in figure 3.

An example of a system employing an array type processor is the ILLIAC IV. A pipelined form of processor is employed in CDC STAR-100 computers. Finally, of primary interest to the present study, an associative type of processor is incorporated in STARAN based computing systems. Although these systems differ basically in their organization of resources and their communication structures, their gross characterizations are quite similar. Since distinct parcels of storage are effectively assigned to individual processing elements in each, a mass memory that has been configured expressly for interfacing properly with a STARAN associative array processor could probably be adapted for use with either of the other types with only interface unit modifications.

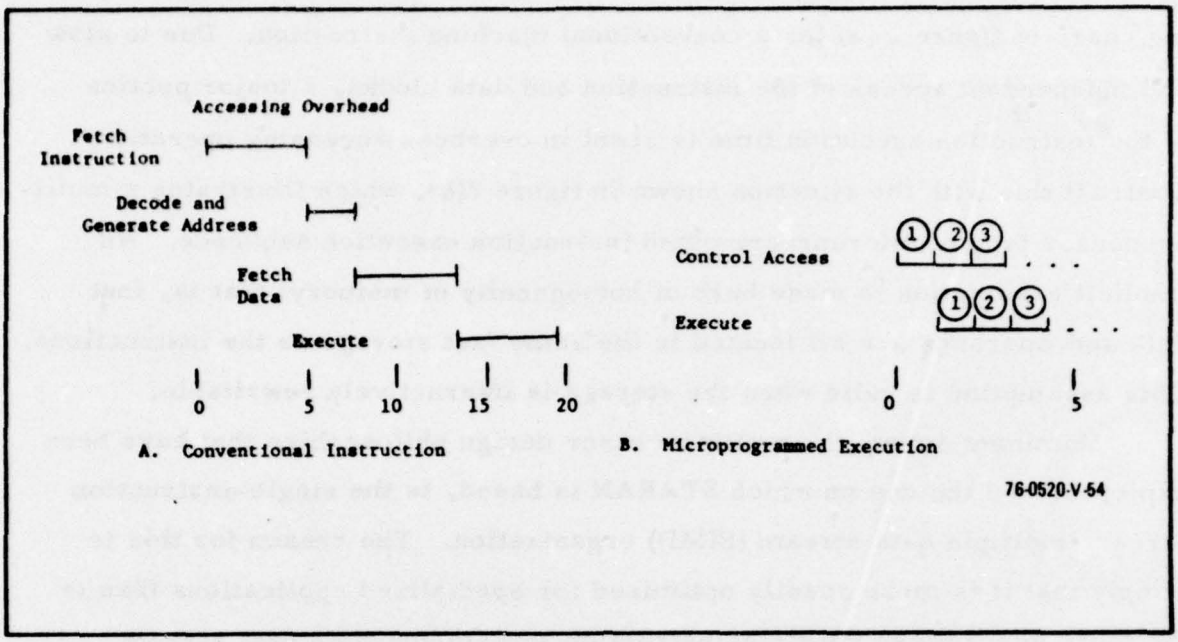

Figure 2. Computer Function Timing Charts

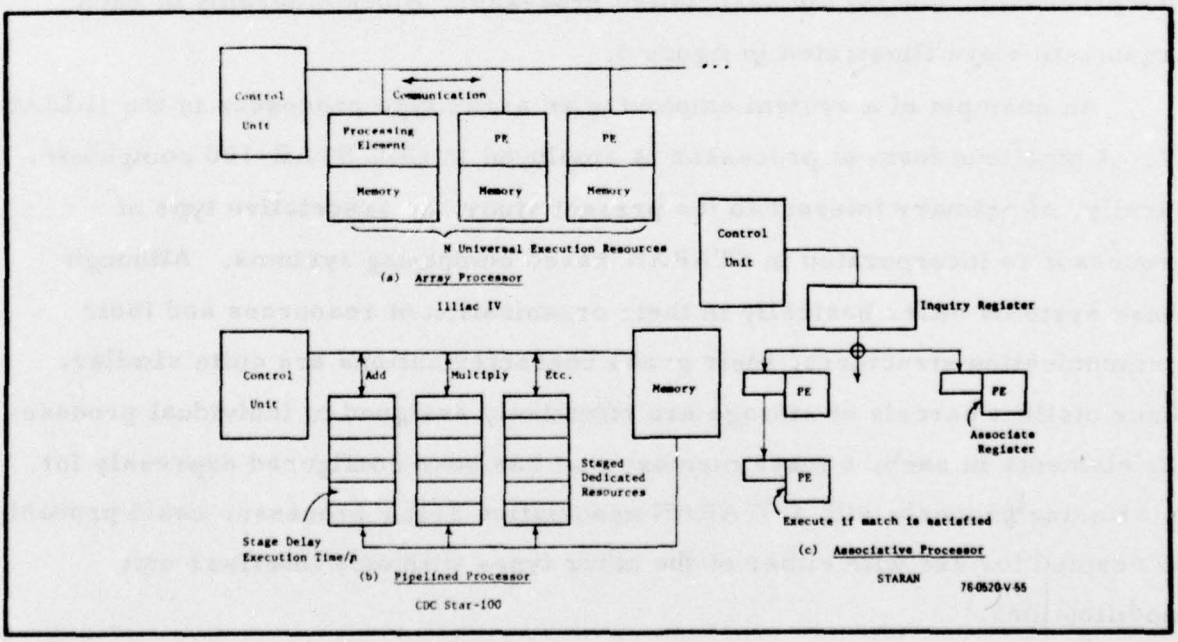

Figure 3. Idealized SIMD Multiprocessors

In the array processor, the execution resources are replicated η times, and operations within successive elements are performed in a parallel fashion. The pipelined processor is similar to the array processor, except that operations within individual elements are serially pipelined. Both array and pipelined processors work with directly addressed data block pairs. Thus, the location of each pair of operands must be known in advanced.

The STARAN type of associative processor is also similar in many ways to the array processor; but, with it, the execution of a control unit instruction is conditional upon a match existing to an inquiry pattern presented to all units. Having been developed for and therefore being immediately available to RADC, the STARAN computer has been used in setting up the guideline requirements desired from a fast mass storage device. Based on these requirements, the design concepts are evolved for a 1 Gbit (10^9 bit) mass memory (MM) having speed/bandwidth capabilities matching those of the STARAN computer. Configuration of the recommended memory system is based on the use of nonvolatile MNOS memories developed by Westinghouse to provide the storage of large quantities of rapidly accessible data and instructions necessary to the high operating speed of STARAN.

Figure 4 presents a global description of the central organization of this machine and the closely mirrored architecture that has been evolved for the mass memory. The STARAN associative array processor (AAP) contains four distinct processing elements (APA's), with each arranged in a square content addressable array pattern of 256 words x 256 bits. Data transfers between the collective APA's of STARAN and the MM will occur in block form, with each block consisting of up to 256 serial bit "lines" flowing along each of 1,024 parallel bit channels. The machine cycle time between each of the 256 successive lines of 4 parallel data words (i.e., 1,024 data bits) will typically be on the order of 300 to 450 nsec. This means that the minimum interface bit rate bandwidth requirement for data transfers between the memory and processor is greater than 2.3 GHz (i.e., $1024 \div 450 \times 10^9 = 2.28 \times 10^9$). The entire capacity of the four processor arrays, consisting

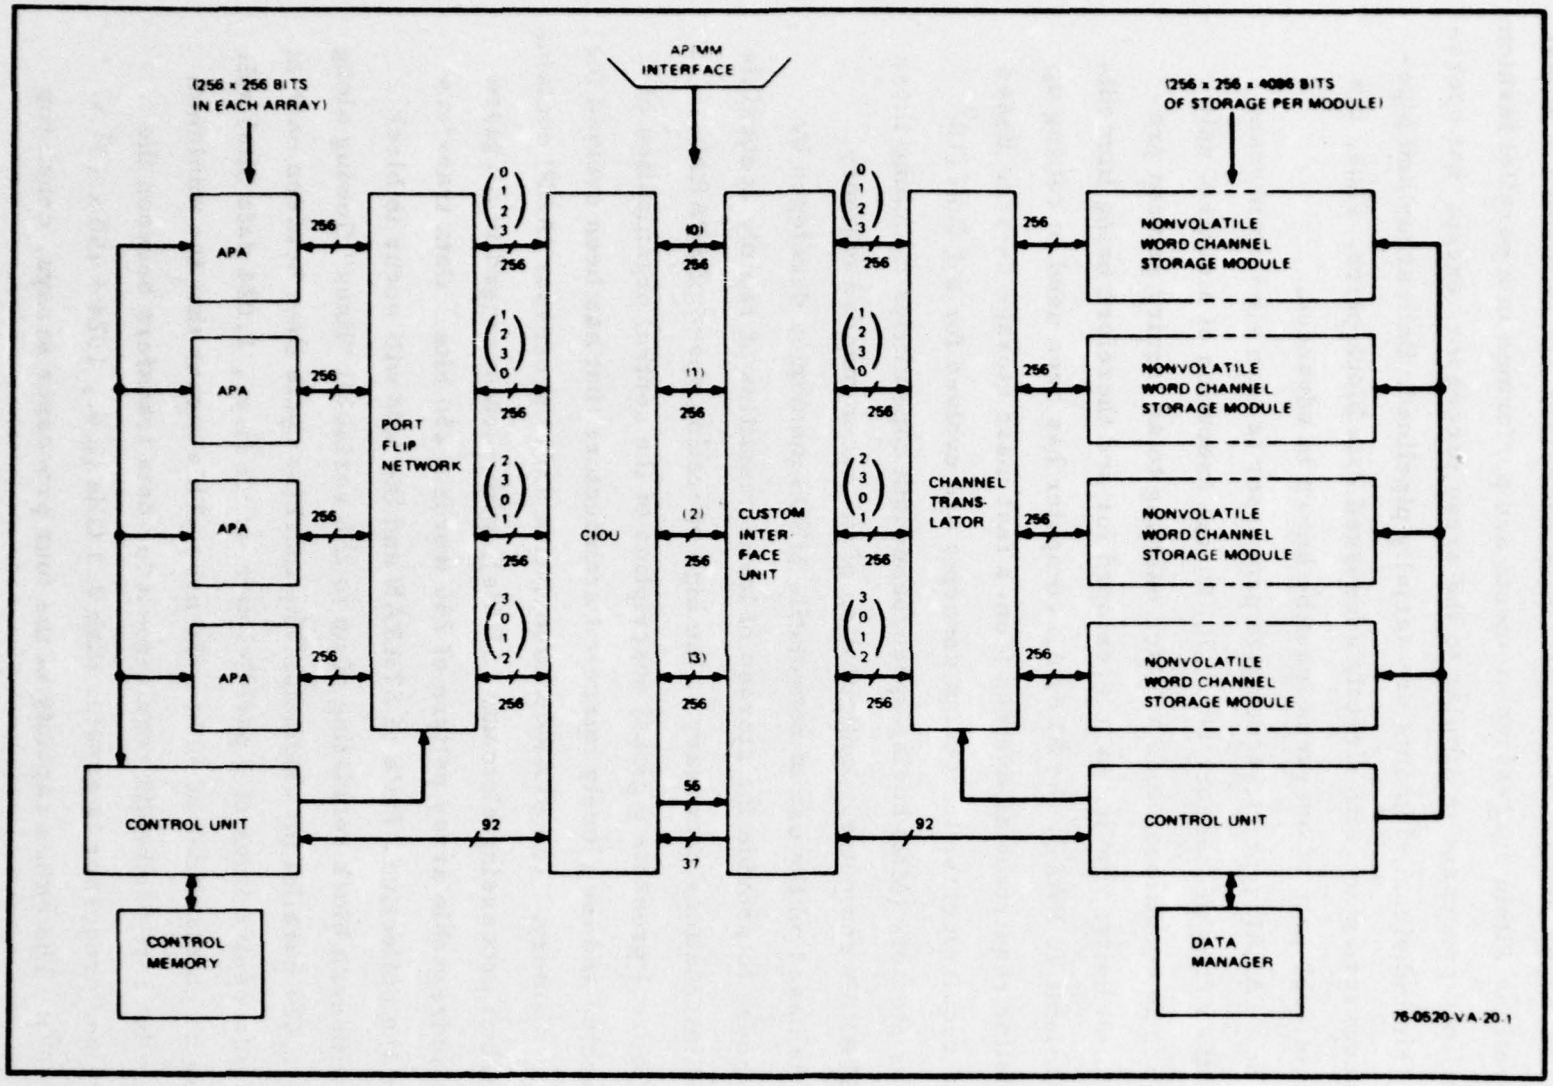

Figure 4. STARAN Computer Interfaced to Recommended MNOS Mass Memory

of a total of 256 Kbits, can thus be transferred between the AAP and the MM in less than 115 μsec.

The degree to which the I/O bottleneck that constricts performance of STARAN is alleviated by the mass memory system depends heavily on the configuration conformity existing between the AAP and the MM. As can be seen, the macroscopic organizational structure of the recommended MM closely reflects that of the AAP. Structuring the MM in this manner was found to provide optimum compatibility, performance, and expansion flexibility.

Storage within the mass memory is organized into word channel modules in a manner compatible with the block transfers taking place with processor arrays across the I/O interface. Storage in each of the word channels is grouped for access as 4096 randomly selectable blocks of 256 x 256 = 64 Kbits, with the combined MM capacity provided by the 4 channels together being 4 x 256M = 1 Gbit. Allowance is made for transfers of fewer than 4 parallel APA words--in either direction and between any ordered set of MM word modules and AAP arrays--by virtue of the port flip network in STARAN and the channel translator in the mass memory. Presentation of the myriad aspects of the design philosophy that this memory structure entails and the tradeoff considerations in adopting this approach are discussed in Sections 2 and 3 of this report.

At the outset of this study it was understood that flexibility would have to be maintained as to the parameters of the desired memory system. The reason for this is simply that the memory design objectives available at that time were generated solely from preliminary global considerations of the backing store needs of associative array processor systems. They necessarily had to be made somewhat abbreviated and imprecise to avoid arbitrarily excluding potentially viable device or design alternatives. These objectives, listed in table 1, are therefore limited in utility to first-order evaluation guidelines.

TABLE 1

AP MASS MEMORY ORGANIZATION STUDY GUIDELINES

Memory Capacity	$>10^9$ Bits
Type of Data Access	Random 65,536 Bit Data Blocks
Form of Data Access	1-64 Slices of 1,024 Bits
Data I/O Format	1,024 Parallel Bit-Serial Data Slice Channels
Data Access Time	<150 Nanoseconds Between Successive Data Slices (The use of precueing and buffer stores is permitted to eliminate latency effects in getting first data slice of block to I/O port)
Data Processing Rate	$>6 \times 10^9$ Bits/Second (This is the approximate rate equivalency resulting from the processing of 1,024-bit data slices at 6×10^6 slices/second)
Operational Modes	Read and Write - Selectable (No stipulation made regarding fast data exchanges requiring concurrent read-write sequences over I/O lines)
Nonvolatile Storage	"Days" of Powerless Data Retention
Reliability	High - (exact target not given) - (Mil qualification is anticipated for follow-on prototype)
Unit Cost	Low
Development Funding	Minimum ("Grooming" only of mature technology is permitted to optimize it to application)
Procurement Timeframe	~1978 - (within 3 years of study)

2. DEPLOYMENT DIRECTIVES

For purposes of this study, all detailed design tradeoffs related to the organization of high-speed/high-capacity mass memories have been made in the context of supplying a 1 Gbit backing store, through the use of MNOS technology devices, which will match the speed-bandwidth requirements of STARAN associative array processor digital computer systems developed by Goodyear under USAF auspices. This section sets forth the functional, operational, and interface requirements toward which design considerations have been directed.

The detailed information on the manner in which specific control signals are or might be interfaced to the STARAN system was assimilated through a series of interactive meetings among Goodyear, RADC, and Westinghouse, and the use of RADC supplied documentation on the STARAN CIOU. Selected reprint sections of this documentation are included for reference in Appendix B. Specifications derived from memory based characteristics are based on operating parameters of MNOS devices which were made available to the study through the BORAM program at Westinghouse.

2.1 STARAN OVERVIEW

The Goodyear developed STARAN associative array processor forms the basis for all design requirements employed in the study of mass memory organizations. A simplified block diagram of the internal organization of this machine is shown in Figure 5. As seen here, it consists of six basic elements:

- Sequential control
- External function logic
- Program pager
- AP control memory

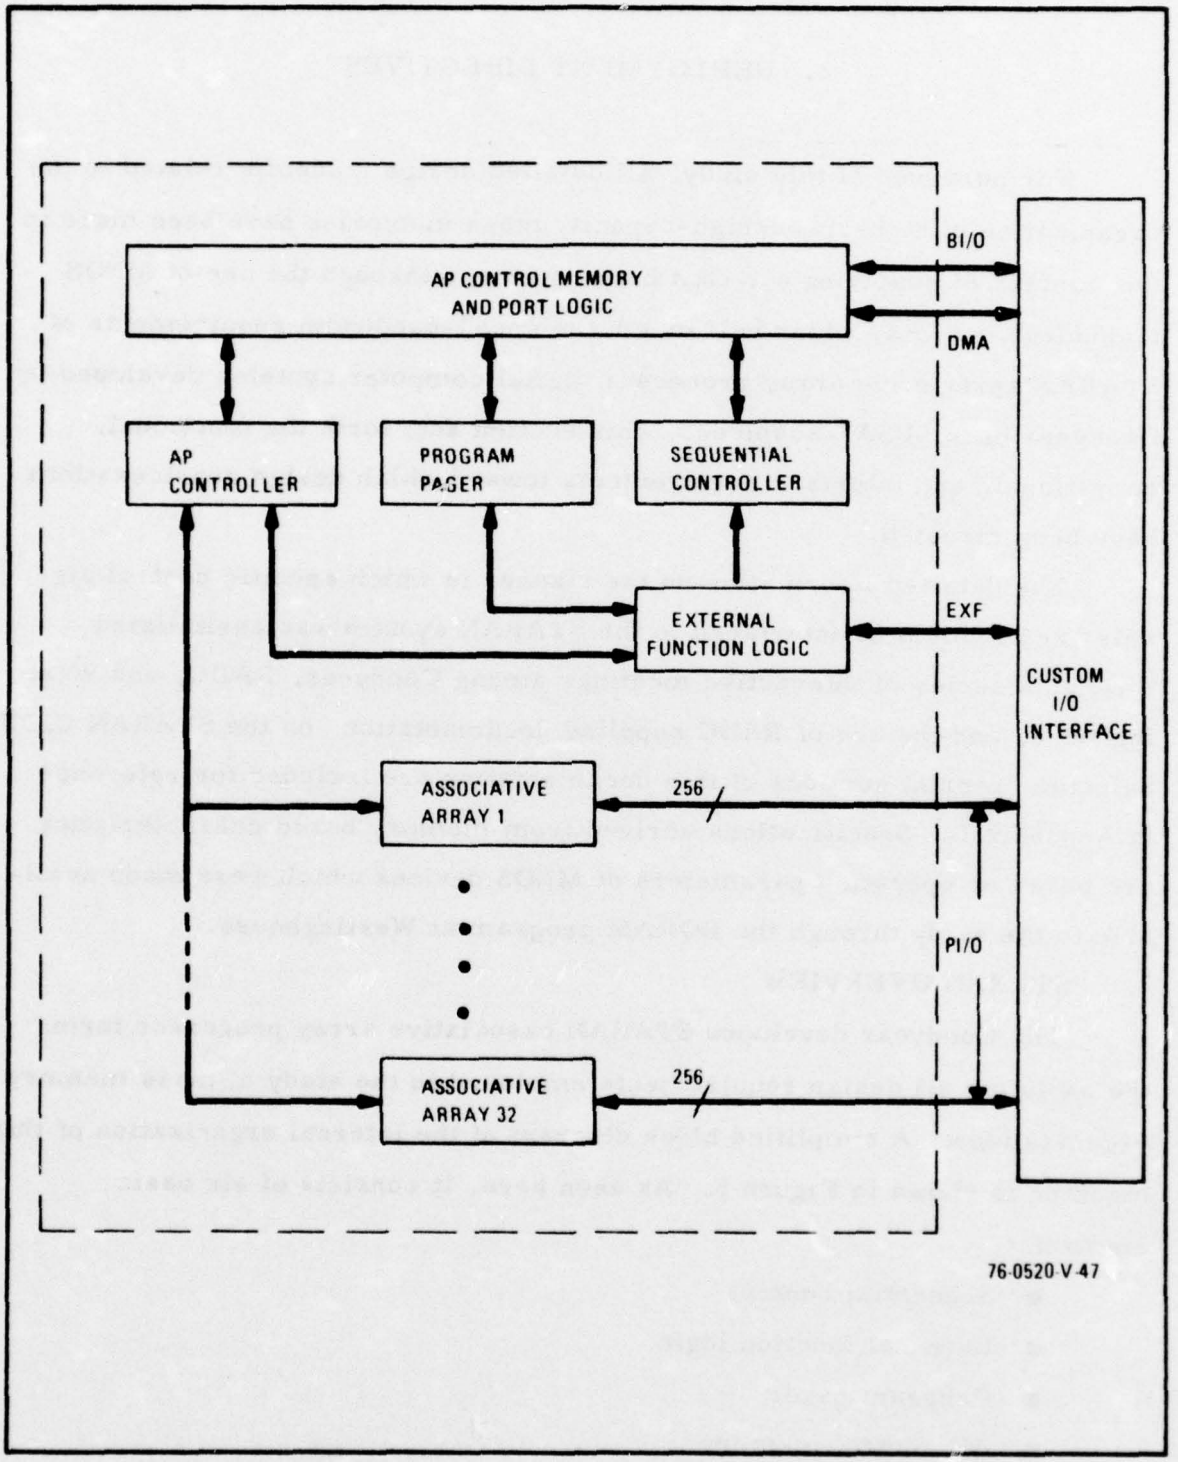

Figure 5. STARAN Associative Array Processor Block Diagram

- Associative processor control
- Associative arrays.

The elements of general interest to the design and development of an interactive mass memory are the associative arrays, the AP controller, and the sequential controller--particularly as to the manner in which they each interface to and interact with the custom I/O unit.

The sequential control unit in STARAN consists of a PDP11 type minicomputer. Its provisions include an associated memory, interface logic to connect it to other elements, and peripheral unit tie-in capabilities facilitated by way of interfaces to its bidirectional UNIBUS. Communication between the sequential controller and the associative processor takes place over various busses in the AP controller as well as the 16 parallel bit minicomputer.

The form of communication between sequential control and the associative processor which has direct bearing on mass memory operations is direct access to AP control memory. By way of this link, words in the AP control memory are given sequential processor addresses to facilitate transfer of data and instructions between AP control and sequential control. Another capability which may prove to be of importance is the sequential controller's central position in the handling of interrupts.

The major function of AP control is to control operation of the associative arrays in STARAN. AP control fetches successive instructions from the AP control memory. It uses a 16 bit program counter to keep track of the instruction address while a 32 bit register contains the instruction itself. Some instructions perform array operations. Others perform AP control functions. Internal control registers affect array operations as well as other elements in the system.

The single most important element in the system is the associative array. It is largely the characteristics of these arrays which provide STARAN with the four major features that distinguish it from a conventional computer:

- Multidimensional access to array memory
- Content addressable (associative) memory
- Simple processing unit at each word of memory
- Unique data permutation network within memory

Each array contains 65,536 bits of temporary solid-state storage organized in a matrix of 256 words by 256 bits. Parallel access to the array can be made in either the bit or word directions. An entire word of 256 bits or a bit "slice" taken from bit position "η" of all 256 words can be accessed in parallel over the PI/O interface lines. A basic configuration of the system will contain only one associative array, but up to 32 can be included. The configuration that is addressed by the present study is with four arrays present.

Four interface options are available for custom tailoring of the processor facilities to specific applications:

- Buffered input/output (BI/O)
- Direct memory access (DMA)
- External function logic (EXF)
- Parallel input/output (PI/O).

The buffered I/O is used principally for tieing different types of peripherals into the processor control memory and interacting with user control programs. Direct memory access allows host computer storage to act as part of the control memory. The external function logic channel facilitates coordination of interface operations between different elements of the processor and external devices. Each associative array has 256 inputs and 256 outputs. These are used for interarray communication and to allow efficient communication with a high-bandwidth external mass memory unit via the parallel I/O channels. The two interface channels that will be involved in communications with and control over a mass memory are DMA and PIO.

In addition to these high-speed custom IO interface channels, the low-speed UNIBUS interface can be used to provide the associative processor

with access to not only the sequential processor but also to any peripherals that are tied in to the UNIBUS. Various peripherals including a DECwriter, paper tape reader/punch, and reader, line printer, alphanumeric CRT display, Disk Operating System (DOS), and magnetic tape unit can be integrated into the STARAN system via the UNIBUS. Thus two communication channels exist in a STARAN computer system: a high-speed parallel processor channel (tied to multiple associative arrays) and a low-speed sequential processor channel. The general configuration of the system is as indicated in figure 6.

At present, instructions as well as original entry of data blocks into the computer (such as from the DOS) are instituted via the low-speed channel. Indicated in this diagram is one possible way in which the mass memory may ultimately be deployed in a STARAN complex: specifically, with separate, concurrently operable access ports to each of the system communication channels. Based on this possibility, the firm requirement that all organizational developments, addressed toward providing an associative array processor compatible mass memory, include provisions for two noninteractive IO ports. Justification for this treatment is found not only from consideration of STARAN but also from review of general development trends in multiprocessor facilities, nearly all of which call for the presence of a multiported mass memory.

Due to plans for a configurable processor facility still being in a state of flux, no final decision was made by RADC regarding disposition of the second mass memory access port. All interface considerations that have been made accordingly apply unequivocably only to one of the MM IO ports. The second port may have a totally different set of interaction capability requirements. Consequently, to prevent application oriented restrictions which are imposed individually on the separate I/O ports from necessitating possibly drastic future redesign to adapt the MM to a different environment, specification was made that a custom tailorable interface (i.e., a CIO unit) be interposed at each MM port to adapt its operations uniquely to the application.

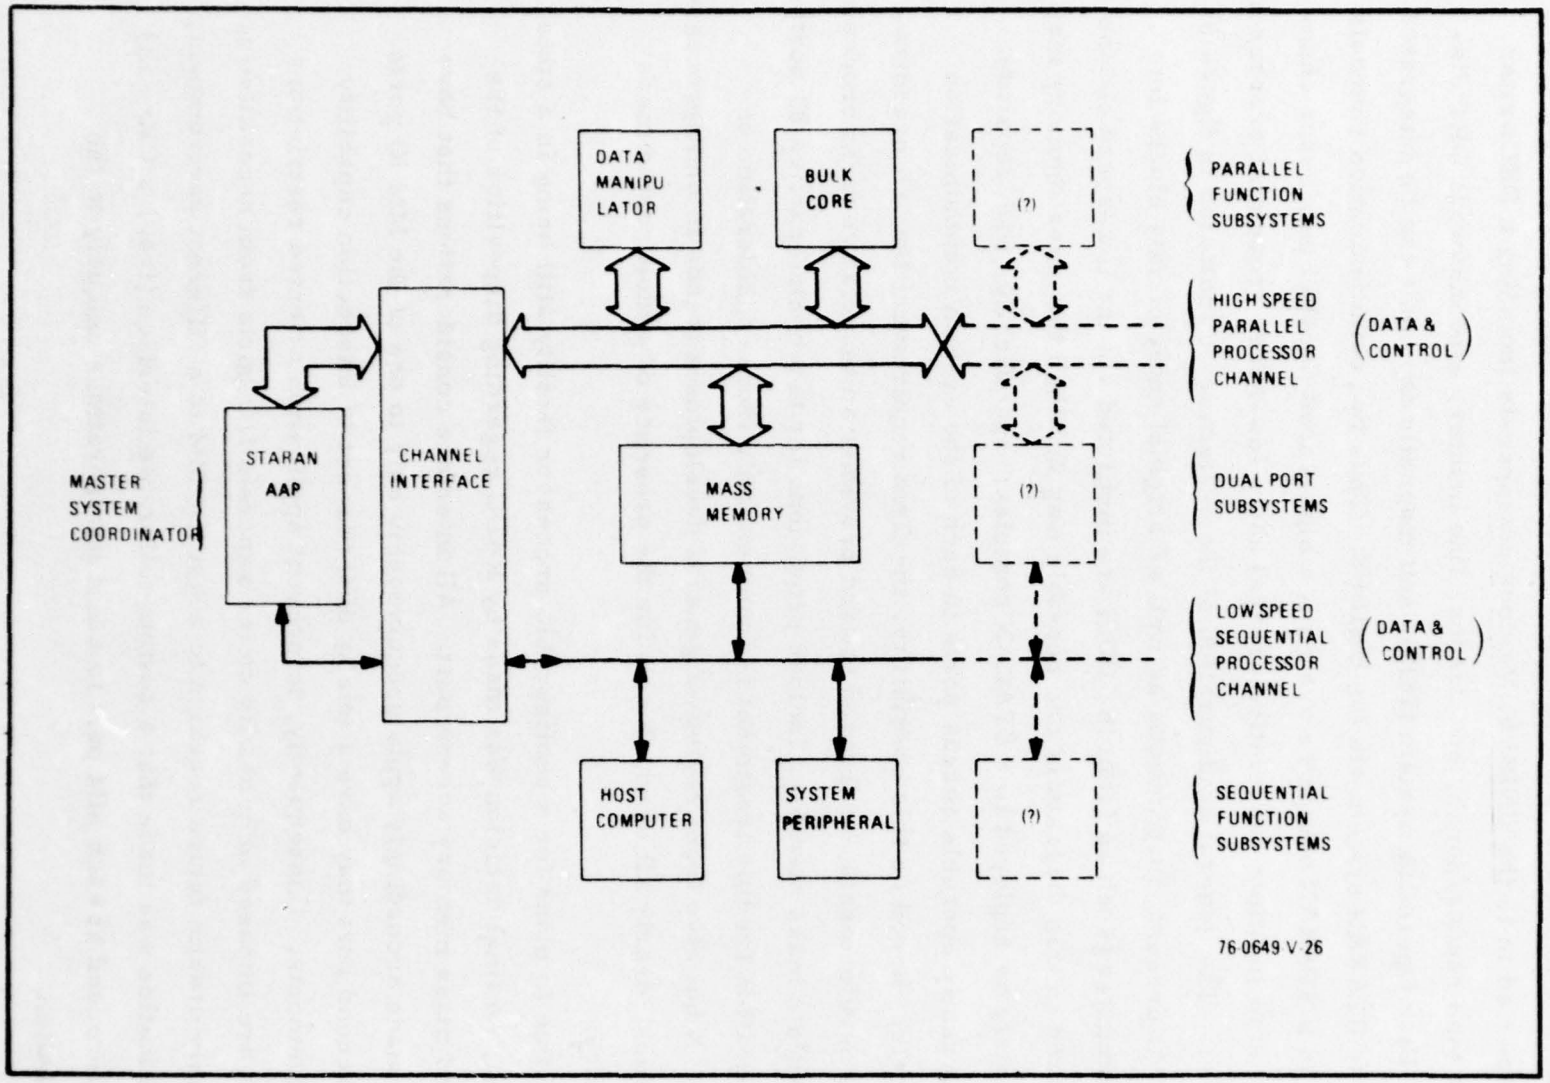

Figure 6. STARAN System Dual-Channel Communication Network

2.2 INTERFACE CONFIGURATION

In a STARAN based associative processing computer system, all subsystems (like the mass memory) are totally subordinate to STARAN. To allow for unrestricted expansion in the number and type of specialized, functionally subordinate systems or peripherals that can be tied into STARAN to enhance the aggregate capabilities of the computer system, special restrictions are imposed on the manner in which devices connect into the CIOU bussing structure. In particular, to prevent undue noise and loading problems from arising as more and more support facilities are added, it is required that each new facility provide for line-receiver termination of all lines affected, and the rebroadcast line-driver transmission of these signals over lines having characteristics identical to those of the original CIOU (specifically, differential twisted-pair ECL levels). Through this "highway" mechanization philosophy, the limitation on the number of peripherals is strictly dependent on the total propagation delay that can be tolerated. A centralized "radial" distribution approach, in contrast, would become limited by the concentration of cabling. A block description of the AAP computer bussing structure that results with chained repeaters is shown in figure 7.

Since active, differential driver devices are involved, the lines interfaced from/to STARAN's CIOU are undirectional in nature. Interconnect cable bundles between STARAN and the mass memory, as well as those between each of the other downstream subordinate systems, are thus composed of two distinctly identifiable line groupings or bus rails. One carries "broadcasts" of commands, controls, and data from STARAN. The other carries "responses" from the designated support mechanism. The bus structure is therefore based on a "dual-rail" organizational principle. Concurrent, noninterfering (at the bus level) operations can thus be instituted, with signals flowing simultaneously in opposite directions over the two rails. Goodyear is understood to be considering possible future use of a bidirectional I/O

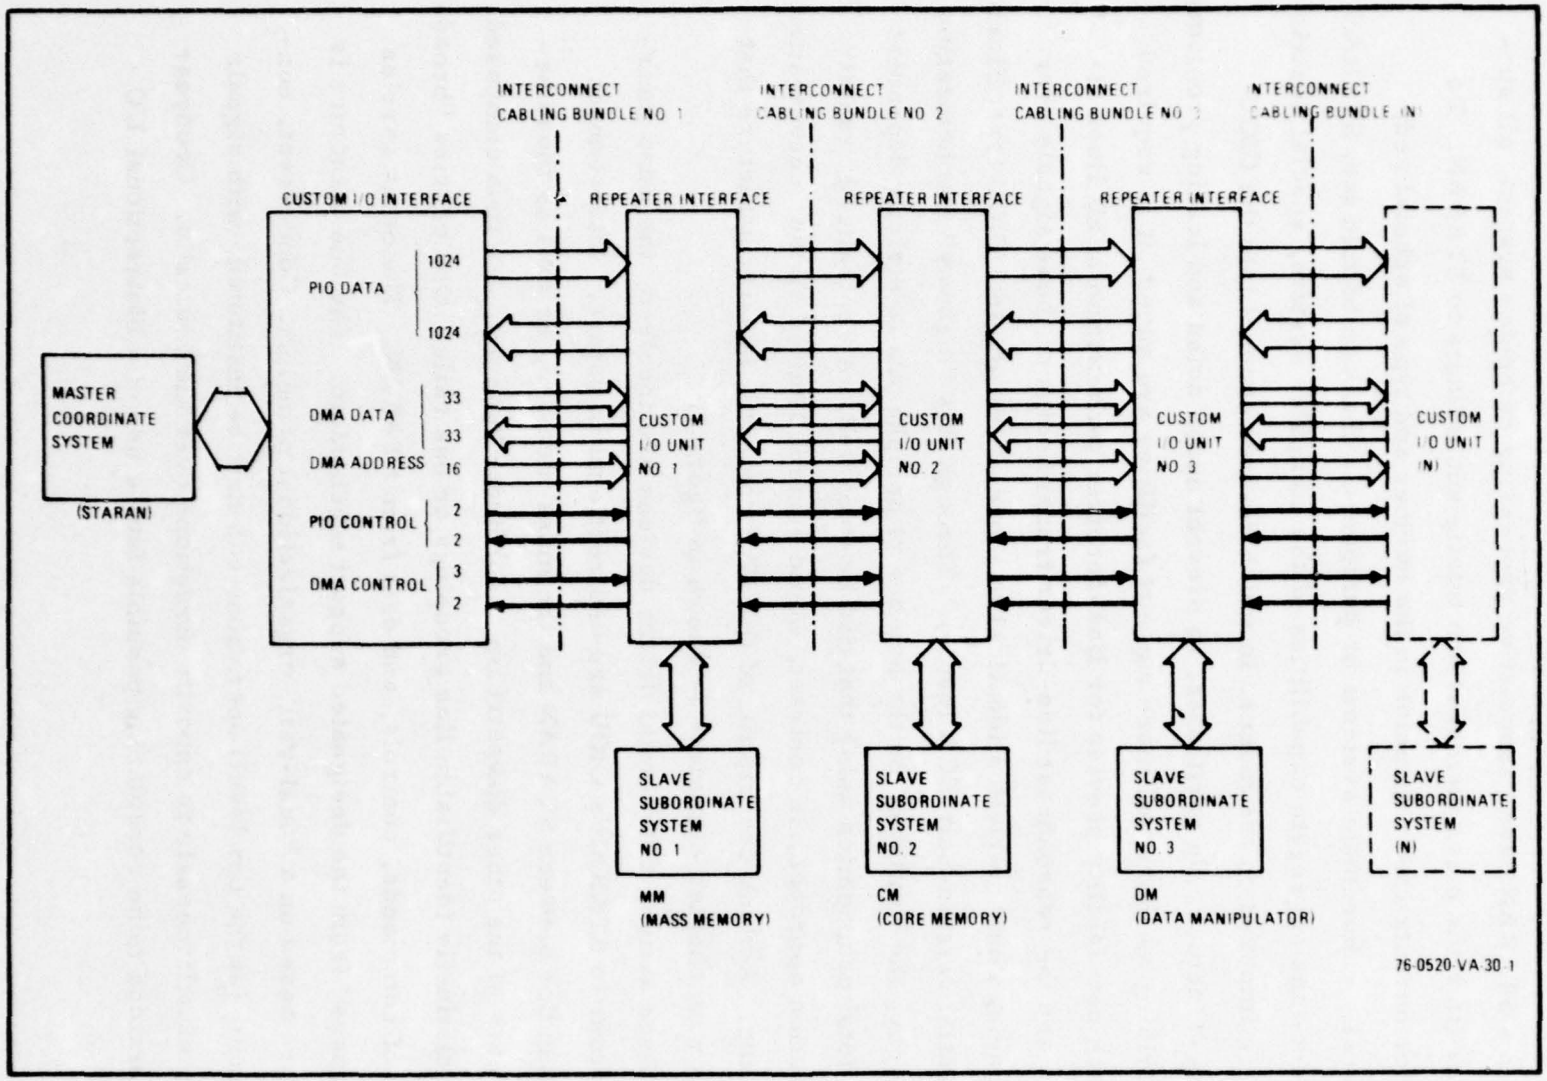

Figure 7. STARAN Interface Channel Line Repeater Network Definition Diagram

line configuration for STARAN interfaces, but for the present, all developments will be based on separate unilateral input and output twisted line pair cabling.

The repeater interface units in the STARAN based AAP computer system will have two distinct repeater modules or sections, as indicated in figure 8 for the MM: one for the broadcast rail and one for the response rail. The manner in which these repeaters are mechanized is markedly different. To facilitate the concurrency of control options over two peripherals that are made possible by the independent sequences that occur over the DMA and PIO channels, the broadcast rail must not be interrupted in any fashion. Signals derived from the broadcast rail repeater are merely sensed and tapped off in a "T" fashion between input and output, with the respective signals passing on through. On the other hand, response rail signal lines must be gated or switched in such a manner that no device further downstream along the communications network can interfere with operations taking place with the mass memory (or other subsystem) when it is selected. Although this does not eliminate all possibility of interference in the communications network (e.g., devices further upstream than the one selected can cause interference), it does significantly reduce it.

Figure 9 depicts the functional equivalent circuits of (a) the tapped throughput broadcast repeater and (b) the gated return response repeater. It is anticipated that hardware mechanization of the repeater interfaces in the manner indicated during an initial mass memory development program will lead to a standardizable repeater module that can be second-source procured in duplicate form as needed for future expansions involving the introduction of additional support devices into the STARAN system.

A final point of note regarding the subordinate system repeater interface modules is that each must obtain its power from supplies which are separate from the main supplies of the given subsystem. This facilitates taking that unit off-line without causing the entire communication channel to

TABLE 2

STARAN CIOU TO MM INTERFACE LINE LISTING

DMA CHANNEL

Outputs (Inputs to MM)		Inputs (Outputs from MM)	
No. Twisted Pair Lines	Line Group Name	No. Twisted Pair Lines	Line Group Name
1	Address Present	1	Access Acknowledge
1	Write/Read Select		
1	Data Acknowledge	1	Data Present
16	Selection Address		
32 + 1	Access Mode Data + Parity	32 + 1	Status Monitor Data + Parity
52	Total	35	Total

PIO CHANNEL

Outputs (Inputs to MM)		Inputs (Outputs from MM)	
No. Twisted Pair Lines	Line Group Name	No. Twisted Pair Lines	Line Group Name
256	Port ϕ Array Read Data	256	Port ϕ Array Write Data
256	Port 1 Array Read Data	256	Port 1 Array Write Data
256	Port 2 Array Read Data	256	Port 2 Array Write Data
256	Port 3 Array Read Data	256	Port 3 Array Write Data
1	BCW Sync Bit	1	Reflected Continue
1	Function Continue	1	Data Strobe
1026	Total	1026	Total

76-0520-T-48-1

become inoperative. A daisy-chained power sense signal line pair (one in and one out) is provided in the STARAN system to determine whether indeed the communications link is fully active at all points along its length. Each intermediate subsystem repeater throughputs this signal in each direction, while the most distant or end unit is strapped to reflect it.

Out of the four interface channels available from the CIOU, only two are projected as being actively tied to the Mass Memory: namely, DMA and PIO. The DMA channel provides STARAN with the means of transferring commands to and monitoring the status of the Mass Memory access control memory. Originally, the EXF channel was considered by RADC to be the appropriate medium for this function, but it was found through consultations with Goodyear that the cycle time of that channel was too slow for the application, and further, that there were too few control lines remaining uncommitted. Data transfers between each of the 256 bit by 256 word associative arrays and the mass memory are effected via the PIO channel. Table 2 lists the data and control function lines that comprise each of these channels.

Interconnection of the mass memory into the PIO channel is accommodated via a circuitry block known as the "port flip network" (PFN). Functionally this "flip ckt" is analogous to an 8-port "cross-bar-switch" network. Each port consists of 512 parallel lines split evenly into two 256 unidirectional input and output line groups. In the RADC installation, four of these ports are presently assigned to STARAN associative arrays. The remaining four are available for use in servicing the MM storage section. Master authority over PFN operating sequences resides with the PIO (parallel input-output) controller which issues sequentially executed program instructions. By tieing the MM into the AAP system in this manner, it can (for ease of programming and control) be viewed as simply an additional set of AP memory arrays which can be manipulated for storage.

TABLE 2

STARAN CIOU TO MM INTERFACE LINE LISTING

DMA CHANNEL

Outputs (Inputs to MM)		Inputs (Outputs from MM)	
No. Twisted Pair Lines	Line Group Name	No. Twisted Pair Lines	Line Group Name
1	Address Present	1	Access Acknowledge
1	Write/Read Select		
1	Data Acknowledge	1	Data Present
16	Selection Address		
32 + 1	Access Mode Data + Parity	32 + 1	Status Monitor Data + Parity
52	Total	35	Total

PIO CHANNEL

Outputs (Inputs to MM)		Inputs (Outputs from MM)	
No. Twisted Pair Lines	Line Group Name	No. Twisted Pair Lines	Line Group Name
256	Port ϕ Array Read Data	256	Port ϕ Array Write Data
256	Port 1 Array Read Data	256	Port 1 Array Write Data
256	Port 2 Array Read Data	256	Port 2 Array Write Data
256	Port 3 Array Read Data	256	Port 3 Array Write Data
1	BCW Sync Bit	1	Reflected Continue
1	Function Continue	1	Data Strobe
1026	Total	1026	Total

76-0520-T-48 -1

2.3 CONTROL PROVISIONS

Collectively, the 32 bit (plus parity) data input and output buses and the 16 bit device selection address output bus of the DMA channel will serve to convey all access request information to the Mass Memory. All transfers will be initiated and carried out under AP control. For purposes of addressing, the MM access control memory is considered an extension of AP control memory and is assigned a specified block of addresses. Table 3 itemizes the complete list of address assignments that have been projected. Hexadecimal addresses 8000 through FFFF are reserved for bulk core, 0800 through 7FFF are set aside for extended AP control, and the lowest ordered hexadecimal address (0000 through 000F) is allocated to MM control. Addresses 0010 through 07FF remain as yet unassigned. All elements of the STARAN system that can access the AP control memory can also access the MM control memory via the DMA interface channel.

Based on the interface line listing of table 2 for the DMA and PIO channels, a "black box" representation of the interconnections from STARAN to the MM appears as shown in figure 10. All accesses to information contained within the mass memory begin with a DMA read or write operation. The "address present" signal, originating in AP control, initiates all data transfers over the DMA channel. The "access acknowledge" signal, originating in the MM control unit, terminates all DMA data transfers.

The DMA "write" operation refers to the transfer of data from STARAN to the MM control memory which signifies the op code and "memory reference" instruction for an access to the main storage section of the mass memory. The DMA "read" operation refers to the transfer of data from the MM control memory to STARAN. This data constitutes readouts from status registers, ledgers, and other control functions associated with the monitoring of activities and conditions within the mass memory.

TABLE 3

PROJECTED 16 BIT DMA DEVICE SELECTION
ADDRESS ASSIGNMENTS

Bulk Core Memory	8000 to FFFF	1111111111111111 to 1000000000000000
Extended AP Control Memory	0800 to 7FFF	0111111111111111 to 0000100000000000
(Unassigned)	0010 to 07FF	0000011111111111 to 0000000000010000
MM Access Control Memory	0000 to 000F	0000000000001111 to 0000000000000000

76-0649-TA-10

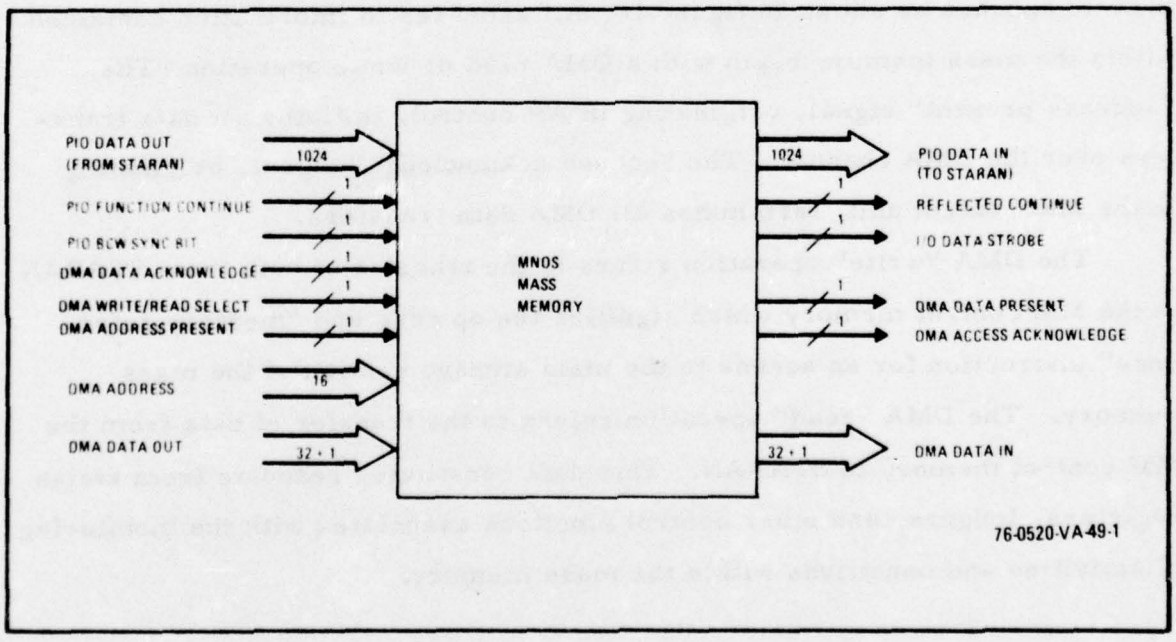

Figure 10. STARAN/Mass Memory Black Box Interface Definition Diagram

Sequences which transpire over the PIO interface are facilitated only after a storage access request has been entered and acknowledged over the DMA channel. Following acknowledgement of the specific access request, PIO data transfer sequencing is allowed to begin. A "read" operation over PIO involves the transfer of data from one or more of the four available 256 bit wide PIO array ports to the mass memory for entry into storage. Alternately, a "write" operation involves the readout of MM stored data and transfer of that data back into the associative arrays via their respective PIO ports.

All arrays in the STARAN system do not always take part in every PIO data transfer operations. Internal to the system there are two control registers which specify the operational states of the respective APA's. The array "assignment" register determines which arrays will take part in external IO operations under PIO control and which ones will remain under AP control. Of those under AP control, the array "select" register further determines which ones will be active and which will remain quiescent during any given program sequence or machine instruction cycle. The information contained in the "assignment" register must be provided to the MM via the DMA access command to facilitate proper selection of only those segments of memory actually needed during a given I/O cycle. Further masking on a bit basis is possible, but this and the multitude of other functions occurring solely within STARAN should be of absolutely no concern to the MM. Only those operations having direct bearing on conditions at their interface are of any importance.

In general, the control mode existing between all sections of an associative processing computer system is that of an asynchronous "handshake" wherein both members must agree on their respective roles in any interaction between them. This results in the transfer of data across the MM-to-AP interface taking place with what might be described as a "gimme-gotcha" type clocking. Initial location and setup of the first data slice preparatory

to transfer following an access request must be accomplished asynchronously by the MM without any further clock or strobe inputs. This necessarily implies that the MM has its own internal timing clock, but that bit sequencing control shift from internal to external control once the first "slice" has been positioned at the MM IO port.

In requesting access to external storage, the associative processor generates a set of commands specifying the starting data-slice address, the number of slices to be transferred, and the type of access desired (e.g., read, write, read/write, clear, etc.). It then sets up the proper function and address line states and "cues" the memory that an access is pending. Assuming the STARAN system is representative, the PIO controller will originate all calls for data transfers to or from the MM and will output the necessary address and control information via the 32 bit wide DMA port. Multiple strobings of the 32 bit DMA "word" can be employed if more than 32 bits are needed to fully specify the access conditions.

The Address Acknowledge line is used to signal that the input address has been entered and recognized and that the requested starting data slice is available for reading at the output port and/or writing at the input port. Once a cued address has been acknowledged by the memory, initiation of data transfer operations is controlled exclusively by the processor via the Function Continue line.

In a generalized multiprocessor or multicomputer system the MM may receive access requests from a number of different devices. In the STARAN system, though, this is not the case. The reason for this is simply that, although I/O data may be steered between the MM and any one of a number of other structures (such as the host computer memory or supporting peripherals), all access requests sent to the MM will originate in the STARAN AP controller and be sent to the MM via the PIO controller and the DMA port. All request priorities will be settled within the STARAN AP controller before

the PIO controller issues successive address commands. Similarly, routing of the MM data to or from specific subsystems will be accomplished via the port flip network, again under direct control of the PIO program.

No provision is made in the DMA or PIO channels whereby the Mass Memory can signify that it is busy (for whatever reason) and therefore cannot grant a pending access request. This condition is handled in STARAN by way of the simple expedient of considering a request invalid after 20 msec, aborting the request, and going ahead to other business. Should it prove of major importance in improving overall system communications, it may be possible to allow the control sections of the MM to tie onto the sequential processor UNIBUS as an additional element in its interrupt system. This would provide a means whereby the mass memory could communicate its unavailability, the reason for it, and request the type of service needed. Consideration of the options will need to be reviewed further during detailed algorithm development work for a prototype Mass Memory.

To ensure proper cycling of the machine functions in the CIOU interface, special provision has been made to tap-out the BCW (buffer control word) sync bit for use in all PIO interface transactions with the Mass Memory. Table 4 sets forth the manner in which the BCW sync bit is related to PIO sequences and how it must be interpreted and used by the MM control structure. In particular, operations are keyed to use of the BCW sync bit by the mass memory as an incrementing enable command which is to be executed upon receipt of the next Function Continue pulse. An additional stipulation made is that Reflected Continue must always be returned to STARAN. Failure to do so is the one source of system lockup present in the PIO interface. Immediate return without action by the MM occurs when BCW sync is low, and delayed return after all requested internal MM cycle time has relapsed occurs when BCW sync is high.

TABLE 4

STARAN PIO TRANSACTION SEQUENCES

AA READ MM WRITE	AA WRITE MM READ	AA READ/WRITE MM WRITE/READ
AA ⟶ MM	MM ⟶ AA	AA ⟷ MM
No Clock	Clock	Clock*
1 Read Array (No Sync) (N-1) Read Array (With Sync) 1 No Op (With Sync)	N Write Array (With Sync)	N Read Array (No Sync) N Write Array (With Sync)
Single Instruction Loop Plus Single Entry Instruction and Single Exit Instruction	Single Instruction Loop	Two Instruction Loop * [Clock Only During Write Instruction]

76-0520-T-52

2.4 TIMING CONSIDERATIONS

In the STARAN computing system, operating sequences are derived from considerations of both devices (sender and receiver) which take part in any communications process. All functions are scheduled asynchronously. No "master clock" exists to keep the internal workings of two autonomous devices in step. Both parties must mutually agree to all interactively common transactions by way of a prescribed handshaking hierarchy of function requests and execution acknowledgements. The internal processes of the two interconnected machines, in this case the AAP and the MM, set the maximum and minimum rates at which information can be exchanged.

Figures 11 and 12 describe the basic timing relationships which exist for transfers over, respectively, the DMA and PIO channels. Internal STARAN machine cycles require nominally 150 nanoseconds. Since five machine cycles minimum are required to set up and complete a DMA "write" sequence, the highest rate at which requests can be generated is about

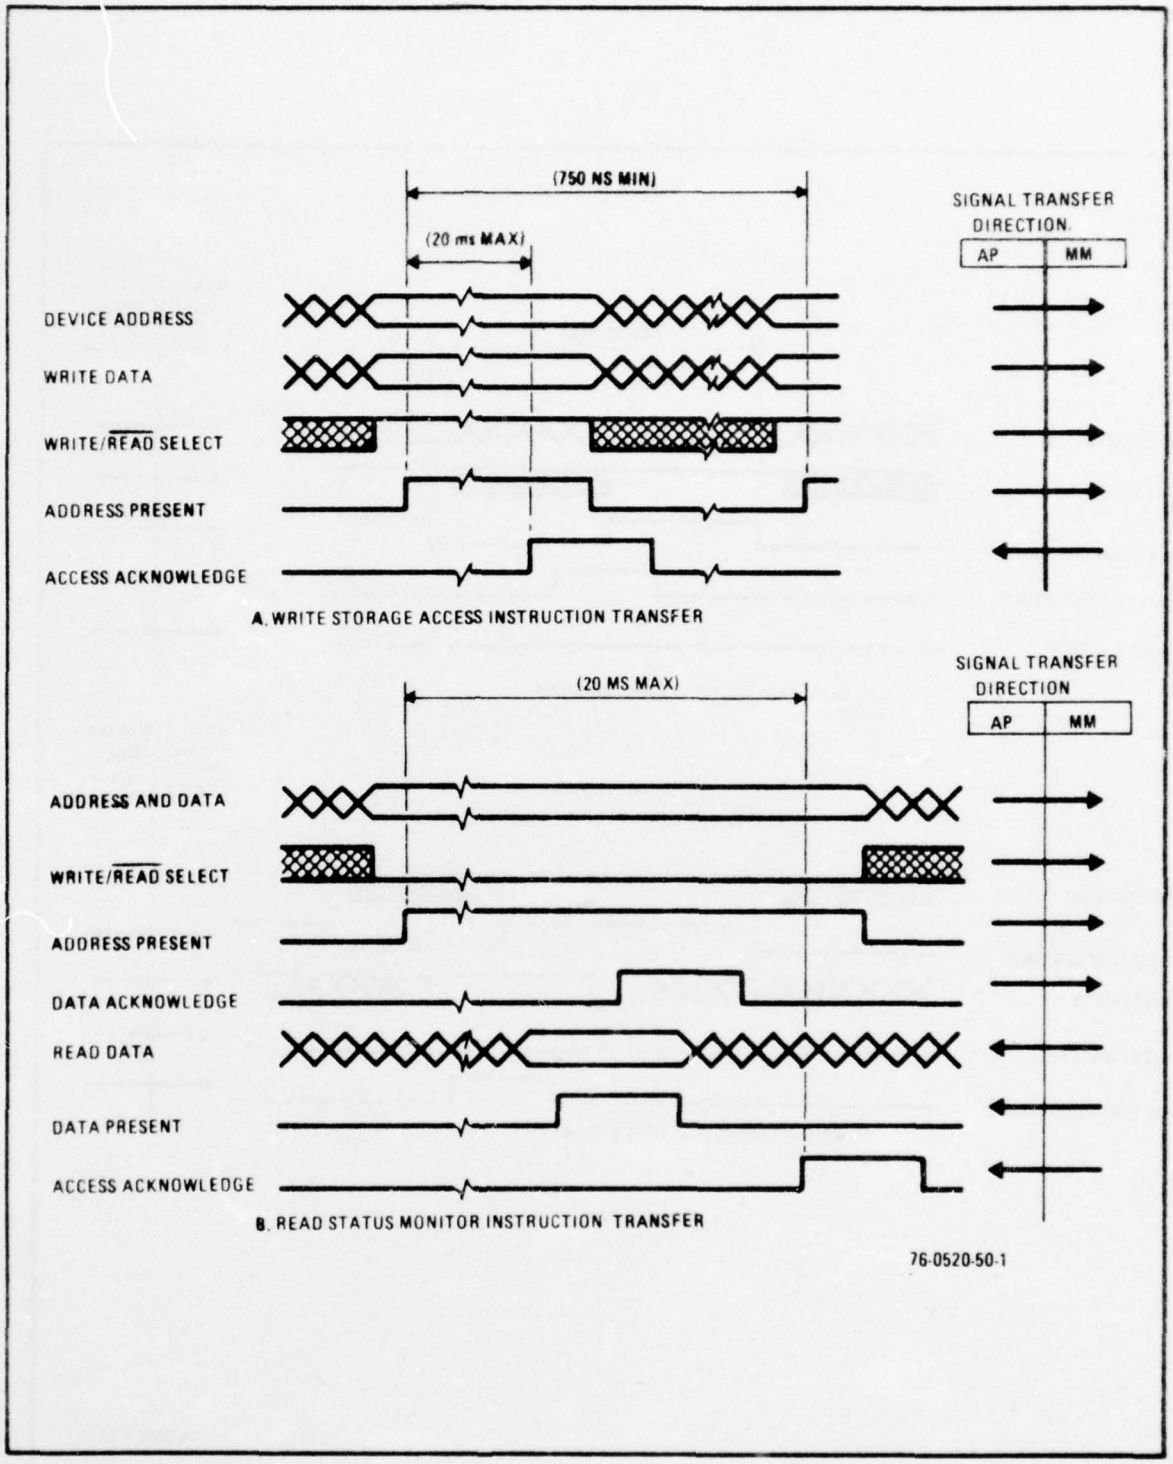

Figure 11. STARAN DMA Channel Timing Relationship

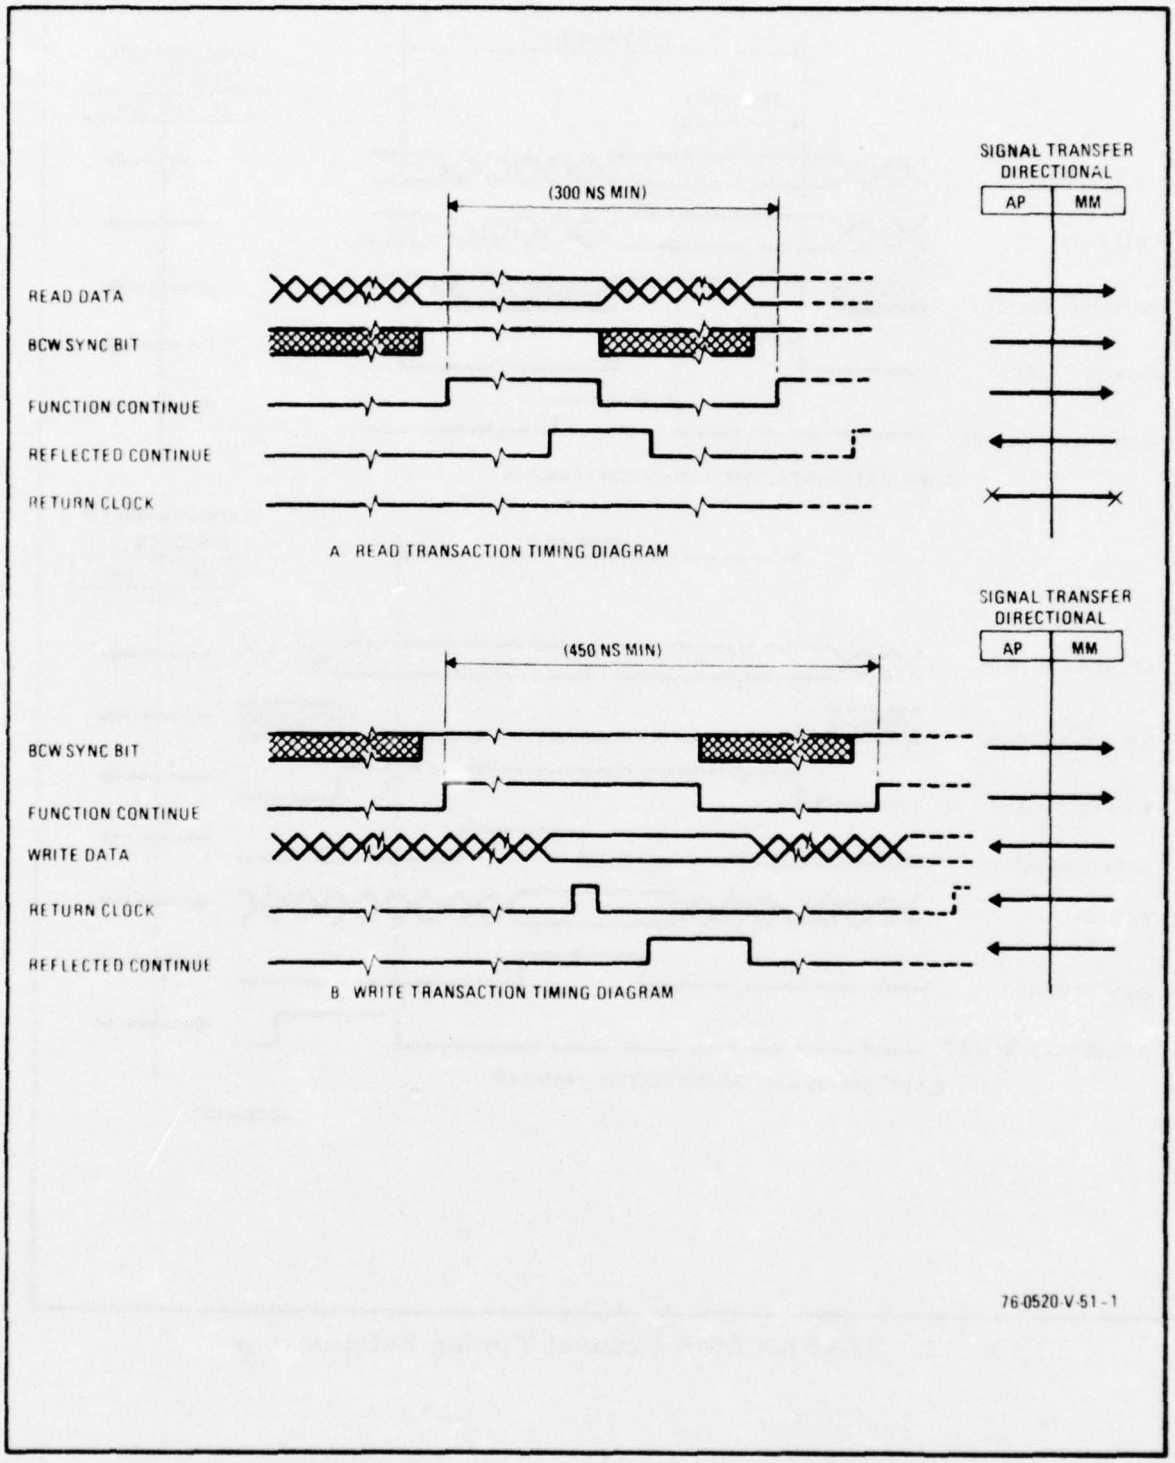

Figure 12. STARAN PIO Channel Timing Relationships

5 x 150 = 700 nsec. A "read" sequence, on the other hand, involves a total of at least nine basic cycles; so, the highest rate for that operation is on the order of 9 x 150 = 1.35 μsec. For all transfer sequences, an upper bounding limit of 20 msec is placed on the permissible time to between issuance of an address present transaction request and receipt of an address acknowledge function done signal. This provision is incorporated into STARAN's control structure in the form of a one-shot timer. It serves to prevent "deadly-embrace" type system lockups.

For data transfer sequences over the PIO channel, the maximum rates are determined by considerations of the shortest permissible instruction cycle time -- in the face of long intercabinet cabling -- in the case of read transactions, and by the shortest reliable write-in time of the STARAN associative arrays during a write transaction. As noted, the instruction cycle time is on the order of 150 nanoseconds. The associative arrays can readout in this length of time, but they require at least 300 nanoseconds (i.e., two machine cycles) to secure reliable write-in. The respective transfer sequences are shown in figures 13 and 14.

The consequence of these timing considerations is that the original guideline objective of accomplishing data transfers between an associative array (AA) and the mass memory (MM) within 150 nsec, or one basic machine cycle, cannot presently be accomplished by STARAN due to physical plant parameters. As to cabling, the same restrictions should apply to the MM. The specification on data flow rate has therefore been cut back to 300 nsec/bit on read (write into MM) and 450 nsec/bit (read out of MM).

Minimum access "queueing" time is dependent on the parameters of the MM storage medium. Using the MNOS technology, the first slice prequeue time will consist of three component intervals during any I/O mode involving a read operation, but only the first of these three for a write-only I/O operation. The three intervals are as follows:

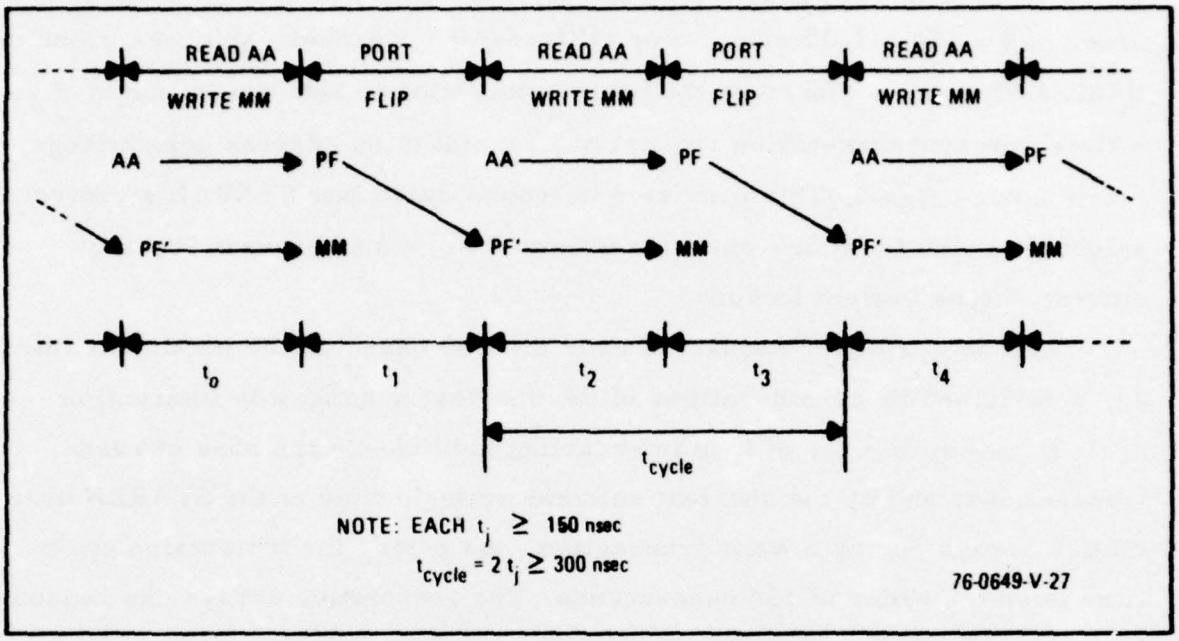

Figure 13. PIO Channel Read Transfer Sequence Diagram

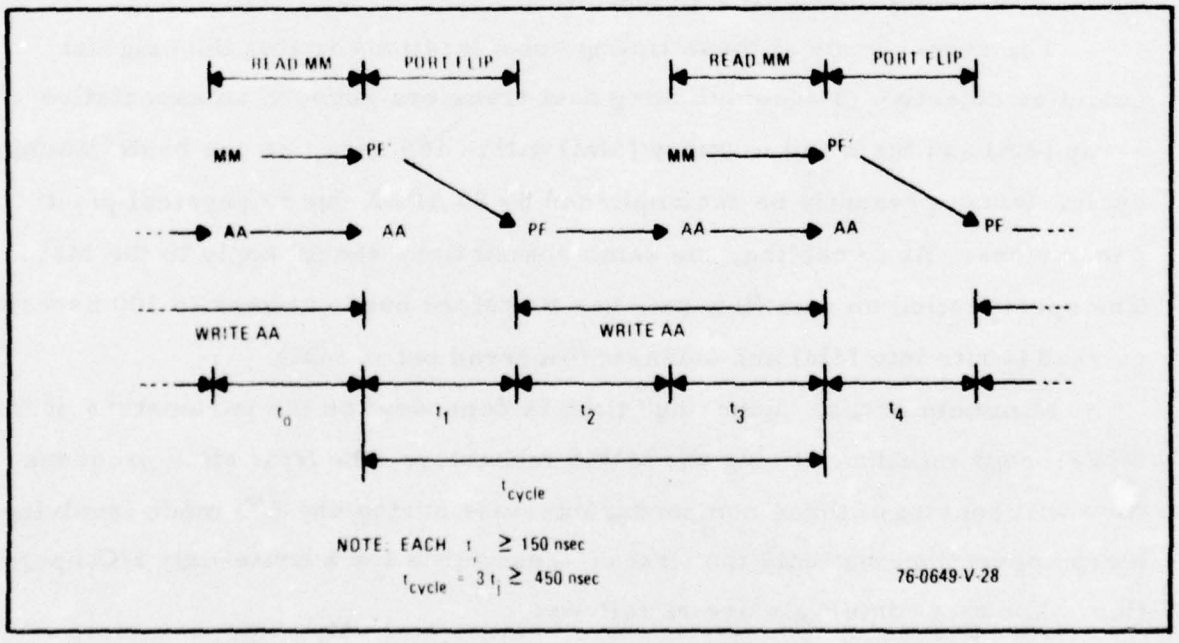

Figure 14. PIO Channel Write Transfer Sequence Diagram

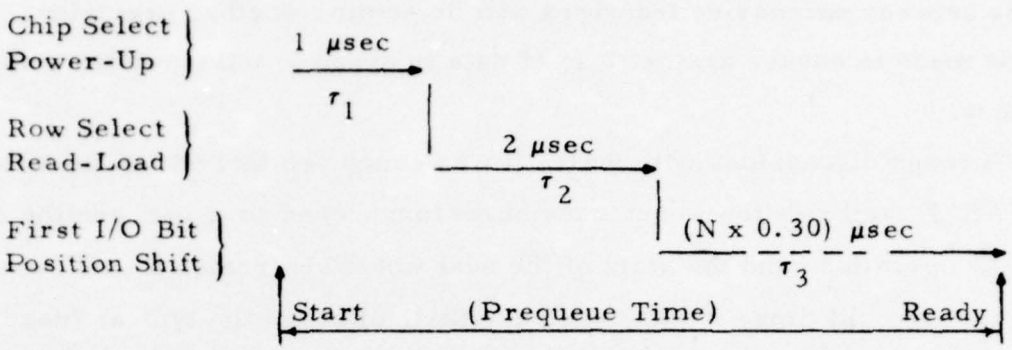

Approximately 1 microsecond is allotted to powering-up the memory chips which will take part in the I/O operation. Another 2 microseconds (typical) is then required to select a specific storage matrix row of data bits and load these bits into the on-chip I/O shift registers.

The final increment of prequeing time consists of shifting the data through the register to position the desired bit at its I/O port. The actual time involved is a function of the shift rate, the register length, and the position of the desired first bit. Presently available registers have proven capable of shifting faster than 200 nsec/bit, so the established 300 nsec/bit is a very safe specification. With the longest registers currently projected for MNOS memory chips up through a capacity of 64 Kbits being 64 stages, the minimum prequeue time will be on the order of 3 microseconds while the maximum first slice setup time will be less than $(3 + (64 \times 0.3)) = (3 + 19.2) = 22.2$ microseconds. This does not, of course, allow for any miscellaneous bookkeeping overhead time that may precede the address being applied to the storage section of the mass memory. Allowance for this tends to indicate that up to 50 μsec total memory queue time may be encountered in end-use installations.

Two further factors relating to the timing of events across the AA/MM interface which will impact on the memory organizations are the average turn-around time between successive PIO transactions and the maximum asynchronous interrupt or idle time between successive word transfers. Turn around time affects the storage device duty factor, which in turn influences average power drain and system reliability. The length of write

periods between successive transfers will determine whether provisions must be made to ensure against loss of data in dynamic intermediate transfer registers.

Through discussions with RADC it was concluded that within the existing STARAN facilities the worst case burst turnaround time between the end of one IO operation and the start of the next would be greater than 20 μsec. In the case of hold time, it was concluded that, although the typical "dead" time in IO sequences would be less than 200 μsec, the outside possibility of it exceeding a few milliseconds does exist. As a consequence, the memory controller must include all necessary provisions to accomplish any periodic "refreshing" of data that may be necessary with extended waite periods.

3. DETAILED INVESTIGATIONS

The need for a large data base storage device has been promulgated by the expanded development of very high speed multiprocessor type computers. These machines have proven so powerful and so fast that they remain idle for extended periods due to the bottleneck that exists at their interfaces to available bulk storage devices like drums, disks, or tape units. Until recently, with the maturing of the MNOS technology, it was not practical to build a solid-state replacement for these memories. It is toward the expressed goal of providing a 1 Gbit storage unit having full capabilities matched to the STARAN type AAP computer that the tradeoff evaluation of alternate organizational and design concepts conducted here is directed.

Six fundamental considerations distinguish the requirements of a mass memory intended to operate in consort with an associative array processor (AAP) like STARAN from more conventional memories which find application with sequential processors:

- Data Base Size
- Dual Access Ports
- Access Port Width
- Block Access Format
- Storage Manipulation
- Function Management.

These factors are addressed in this section. Discussions center around the last three design points, with allowances made where appropriate for the first three.

3.1 GLOBAL HIEHARCHY

In undertaking an investigation of the AAP compatible mass memory design concepts, it is necessary to first identify the global hierarchy of the memory; i.e., to identify and compartmentalize its various functional parts

or subsections so that they can be given individual and coordinated attention. By way of contrast, consider first the configuration of sequential computer storage units. According to established practice, these memories typically contain the following units:

- Random Access Memory (RAM)
- Memory Address Register (MAR)
- Memory Buffer Register (MBR)
- Memory Interface Logic (MIL)

In this class of memory, access to stored data is accomplished on a single bit, or multiple parallel bit, totally random basis. Successive data words (multiple parallel bits) are accessed only by executing a complete access instruction cycle to generate a new address. For handling large quantities of data, the overhead time between successive words makes totally random access very inefficient and useful only in low-speed environments.

To provide greater processor to memory data transfer efficiency, access to high-capacity backing stores (i.e., mass memories) is founded on the principle of block transfers using only a single access instruction cycle. With this approach, a memory address identifies a "block" of sequentially ordered data word storage spaces rather than just a single word. In this class of memory, the equivalent set of functional units is:

- Block Accessed Memory (BAM)
- Memory Address Register (MAR)
- Data Transfer Sequencer (DTS)
- Memory Interface Logic (MIL)

Receipt of an access request causes the first word of the addressed data block to be set up ready for transfer through the memory interface logic. Successive words in the block are shifted across the processor-memory interface at high speed without any further overhead penalty in access instruction cycles in the user device. Note that, at the basic memory module

level, the DTS consists of nothing more than a wire carrying clock or strobe pulses to the storage section.

Using MNOS memories, access cueing times to the first word of a data block on the order 1 μsec and data transfer times between successive words on the order of 200 nsec are easily accommodated. Since this capability exceeds the requirements for the RADC AAP computer, there is no need to incorporate multiplexing (time interleaving) at any point in the memory structure solely for the purpose of boosting the IO speed capabilities of the memory. Consequently, the optimum organization can be more freely established on the basis of minimizing costs and maximizing reliability — two mutually complementary goals to a practical mass memory.

Although block-oriented memory accessing concepts are effective in eliminating the overhead "fetch" times of word accessing approaches, special considerations of the AAP deployment environment must also be accounted for in delineating the MM hierarchy. These include:

- independence of two access ports
- configuration of IO interfaces
- expansion of storage capacity
- reformatting of stored data
- management of data base.

The presence of two concurrently activatable ports impacts heavily on the communication and control networks within the memory. Structuring of the IO interface to provide multiple 256 parallel bit words influences the manner in which storage is modularized to allow for expansion of the storage capacity.

Allowance for reformatting of the stored data comes from consideration of both the types of user peripherals which may be used to initially load the mass memory and from the need to manipulate data among the multiple IO word channels. It bears on the design of data routing and bussing structures and on whether a supporting buffer memory should accompany the man storage unit. Data management provisions constitute an overhead option of the

basic AAP mass memory which converts it from an interactively "dumb" direct accessed memory module to an "intelligent" indirect accessed memory system.

3.1.1 Baseline Module

Taking into account the unique demands placed on its formulation, figure 15 presents the block diagram of a baseline mass memory module capable of being deployed in general multiprocessor environments. As seen here, the global structuring of functions calls for six distinct types of elements:

- Nonvolatile Main Memory (NMM)
- Memory Interface Logic (MIL)
- Memory Function Controller (MFC)
- Access Request Controller (ARC)
- Volatile Buffer Memory (VBM)
- Access Priority Resolver (APR).

Stipulation of nonvolatility in the main storage section provides the requisite data integrity against power loss. Furthermore, it contributes to simplification of support circuitry (through elimination of continuous refreshing), minimization of power drain, and maximization of reliability: all of which improve the memory's cost-of-ownership factor, expressed in terms of ¢/bit/year of operating life.

In single-ported form, no priority resolution is necessary between overlapping access requests entered via a pair of counterpoised ports. The APR unit thus constitutes an application oriented option in the baseline MM module, depending on the need for concurrently operable access ports. Resolution of access priorities among multiple users connected to common port is handled by either the ARC unit or by a higher ordered control structure outside of the memory (somewhere in the computer complex communication network).

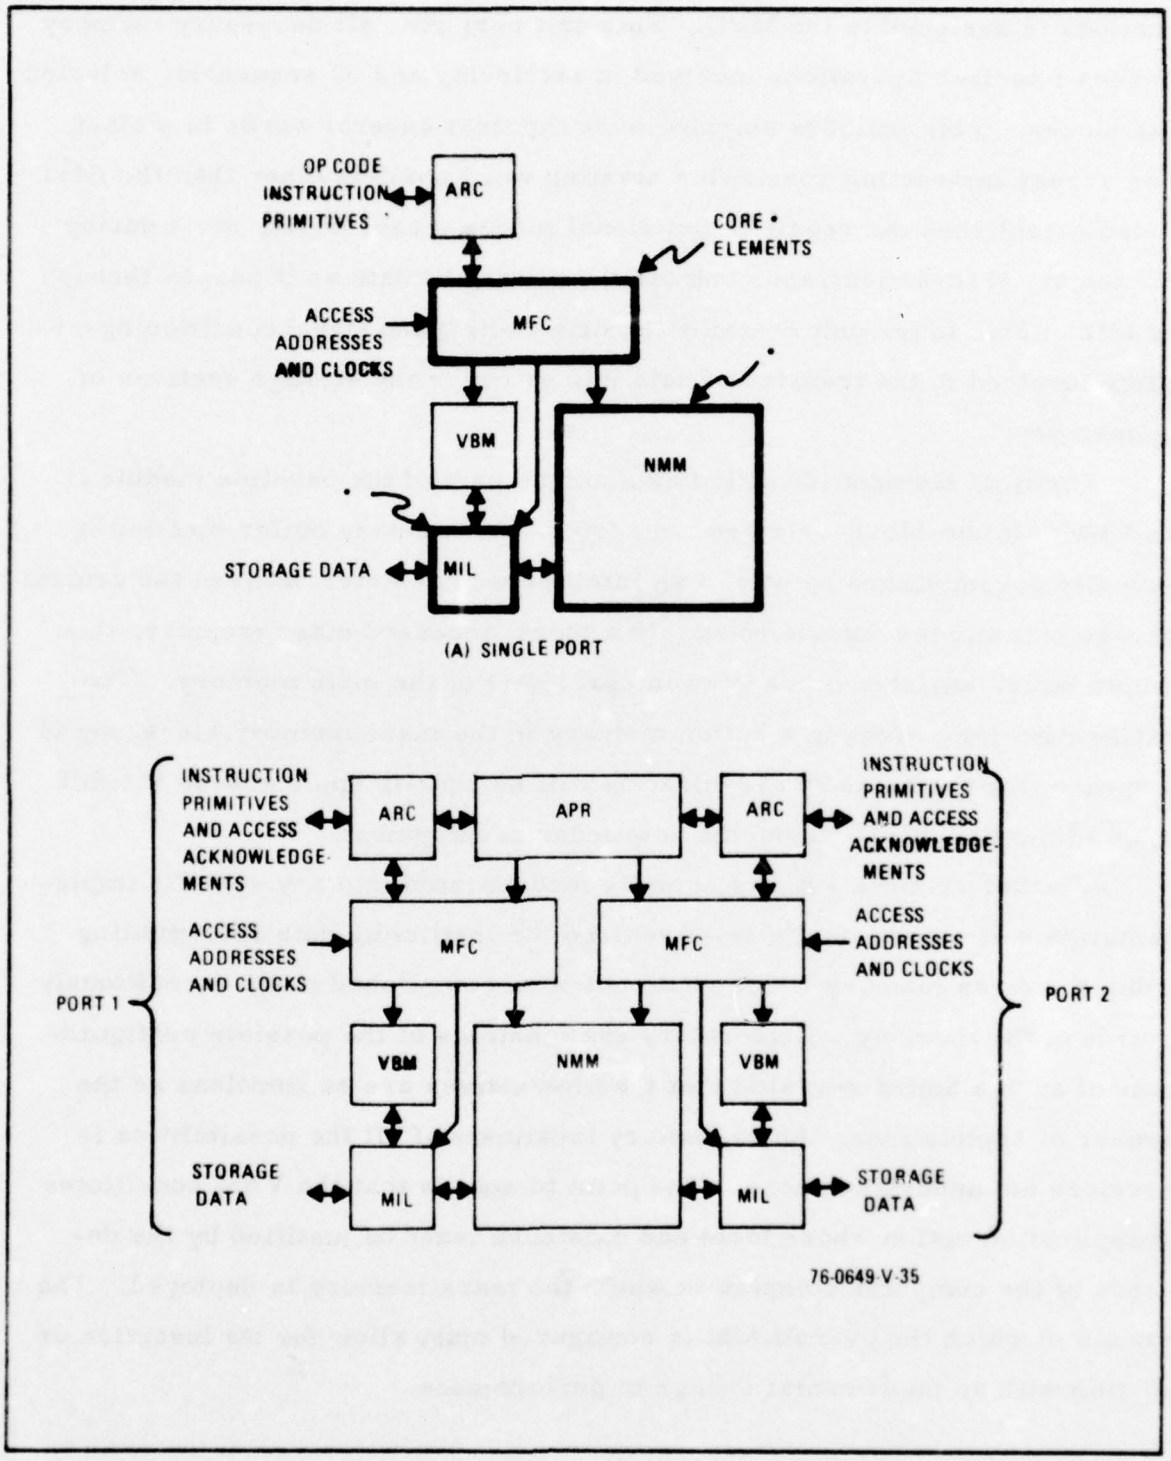

Figure 15. Direct Access Baseline MM Module Block Diagram

Responsibility for direct control of the storage section (i.e., NMM) functions is assigned to the MFC. This unit performs all necessary memory address interface operations involved in retrieving and IO sequencing selected data blocks. This includes skipping over the first several words in a block if an access instruction specifies a starting word position other than the first. It also establishes the requisite functional modes (read, write, etc.) during all memory transactions, and controls the routing of data as it passes through the MIL. This latter unit contains the data routing and signal conditioning circuitry involved in the transfer of data into or out of the storage sections of the memory.

The final element identified as a unique part of the baseline module is the VBM. In non-block-oriented memories, the memory buffer function is normally accomplished by way of an interspersed register between the central storage unit and the outside word. In a block-accessed mass memory, this simple buffer register is made an integral part of the main memory. The justification for providing a buffer memory in the mass memory hierarchy is to ensure that the memory organization will be equally applicable in the full range of multiprocessor systems now under development.

Whether or not a VBM is actually incorporated into any specific implementation will depend on the relative need for instituting data reformatting within the mass memory which could not be accomplished more expeditiously outside of the memory. Exploratory examinations of the possible configurations of such a buffer revealed that the alternatives are as limitless as the number of applications. An expository treatment of all the possibilities is therefore not undertaken here. The point to note is that the VBM constitutes an application option whose form and existence must be justified by the demands of the computer complex in which the mass memory is deployed. The manner in which the overall MM is configured must allow for its insertion or deletion with no fundamental change in performance.

TABLE 5

BASELINE MM INSTRUCTION PRIMITIVES

No	Instruction	Block Address	Storage Operation	Clock
0	READ	XXXX	OUT ← NMM	Ext
1	WRITE	XXXX	NMM ← INP	Ext
2	CLEAR	XXXX	NMM ← "0"	Int
3	DUMP	XXXX	OUT ← NMM ← "0"	Ext
4	MODIFY	XXXX	OUT ← NMM ← INP	Ext
5	DUPLICATE	XXXX,YYYY	NMM ← NMM'	Int
6	RELOCATE	XXXX,YYYY	NMM ← NMM' ← "0"	Int
7	NO-OP	-	(NONE)	(None)

76-0649-T-36

For the baseline MM module, eight instruction primitives have been postulated. These are shown in table 5. It will be observed that only four primitives specify operations which involve transfers between the memory IO and the memory storage section (viz., primitive numbers 0, 1, 3, and 4). For execution of these instructions, the user device must supply a word incrementing clock (function continue strobe). Sequencing during execution of the other four does not require intervention by an external device. It is handled automatically by action of the MFC.

Composition of this list of primitives is based on consultation with RADC. Inclusion of the "duplicate" and "relocate" instructions (No. 5 and 6) was felt to be necessary to facilitate both safety duplication of important data files and restructuring of the data base during the "garbage collection" algorithms inherent in data management. This list is by no means intended to be an irrevocable specification. It has, however, been used in evaluating the capabilities of alternate memory mechanization philosophies.

Observe that provision for VBM primitives has not been made. This is intentional. Designs for the mass memory are based on eliminating the need for an interface buffer store between the AAP and the NMM. Use of a buffer would compromise the effective data rate of the system since continuous, fast data exchanges could not be instituted. Only steering logic will be used to select between data flow paths. Consideration of the uses for a buffer memory are restricted to the means by which data is initially loaded into the system from its supporting peripheral devices. The structure at this point has neither been fully defined nor apparently thoroughly investigated. The various aspects of this peripheral interface problem ultimately will have to be resolved to realize the full computational power of the STARAN AAP system. This depth of system design involvement is, however, beyond the scope of the present study.

3.1.2 Intelligent System

As a stand-alone memory, the baseline MM module will satisfy the guideline performance requirements of a 1 Gbit backing store for an AAP installation where there exists a single, central control unit through which it receives all access commands and which oversees all access requests sent via either port to ensure that storage space is allocated in an orderly and nonoverlapping manner. If this is not handled in an automatic fashion (transparent to the programmer), then conflects will assuredly occur at some point between the data file space assignments made by different users.

The type of information to be stored (whether data or instruction) is, of course, of no direct consequence to the MM organization. Depending on the application, user software routines may make use of microprogramming techniques, reserve whole sections of storage exclusively for reference program instructions, or utilize the memory totally for data storage. Although AAP software assignments are not a direct concern of this study, the functional elements present in the hierarchy of the MM bears heavily on the alternatives visible to the programmer. In this regard, it is to be noted that the

baseline module is based on the use of "direct" data block addressing. It is, therefore, interactively "dumb." It can provide only autonomous responses to processor unit commands.

In the baseline memory, the simple functions described are normally mechanized without frills of any kind. As the size of the available data base and number of potential users expands, the sophistication of the addressing structure must increase. As indicated, a potentially catastrophic consideration arising from indiscriminate use of direct memory addressing in a large data base system is that it is very difficult to incorporate sufficient safeguards among the distributed users to prevent inadvertent destruction of another user's data file as a result of storage space overwrites.

To uniquely access any of the 4 million (256 bit) words in a 10^9 bit data base MM requires the presence of 22 address bits (2^{22} bits = 4 Mbits) for entry and retrieval of successive file words. Furthermore, sufficient overhead memory and interuser communication must be provided as necessary to keep track of still available spaces. Conversion from direct to indirect addressing and inclusion of appropriate service facilities as an integral part of the MM control functions relieves this burden from each user program and greatly enhances the system throughput due to a major reduction in turnaround cycle time.

By centrallizing the "overseer" or "bookkeeping" functions of a data management facility within the mass memory hierarchy in the manner indicated in Figure 16, allocation of space to respective users can be accomplished with continuous checks made to ensure that neither file encroachment nor nonallocated space overwriting occurs. If this responsibility is assigned to a minicomputer tied into the MM control structure, changes in its software routines can be made to affect the manner in which it accomplishes its supervisory tasks. It is accordingly anticipated that this unit will consist of a sequential processor like that in STARAN. This will facilitate rapid modification of the global control functions of the MM.

As seen in figure 16, the global structuring of elements involved in adding an intelligent data management overhead facility to the baseline MM module calls for the use of four distinct types of elements:

- Nonvolatile File Memory (NFM)
- Nonvolatile Resources Memory (NRM)
- Data Management CPU (DMC)
- Cache Access Memory (CAM).

Just as with the main memory section, nonvolatility of storage is necessary in the data file and resource allocation directories of the MM data management unit. Without this proviso, power outages would result in loss of all file location and space assignment information. Recovery from this condition would be extremely difficult, if not impossible, unless extensive "header" information was entered with each block of data in the NMM section so that it could be scanned and used to restore the management directory contents.

Insertion of file headers preceding blocks of data is the standard procedure followed with present rotating magnetic storage units like drums, disks, and tapes. For them, it is virtually mandatory due to their totally sequential access structure. Due to this, the "search" time to locate a requested block of data is frequently very long (extending to several minutes with some tape units). This approach to data base management could also be applied to a solid-state mass memory. It would however, have the obvious drawbacks of (1) reducing the available storage space for data by the amount needed for the headers and (2) increasing the acquisition time due to the need for a serial search of a very large data base. These conditions -- particularly the latter -- are unacceptable from the ground rules of this study.

Minimization of search time is one of the key parameters which must be addressed in developing the control algorithms involved in access request processing and data management. It is for this express purpose that specification is made for the introduction of a small associatively searched "cache" memory to serve as an access ledger for the most recently transacted file

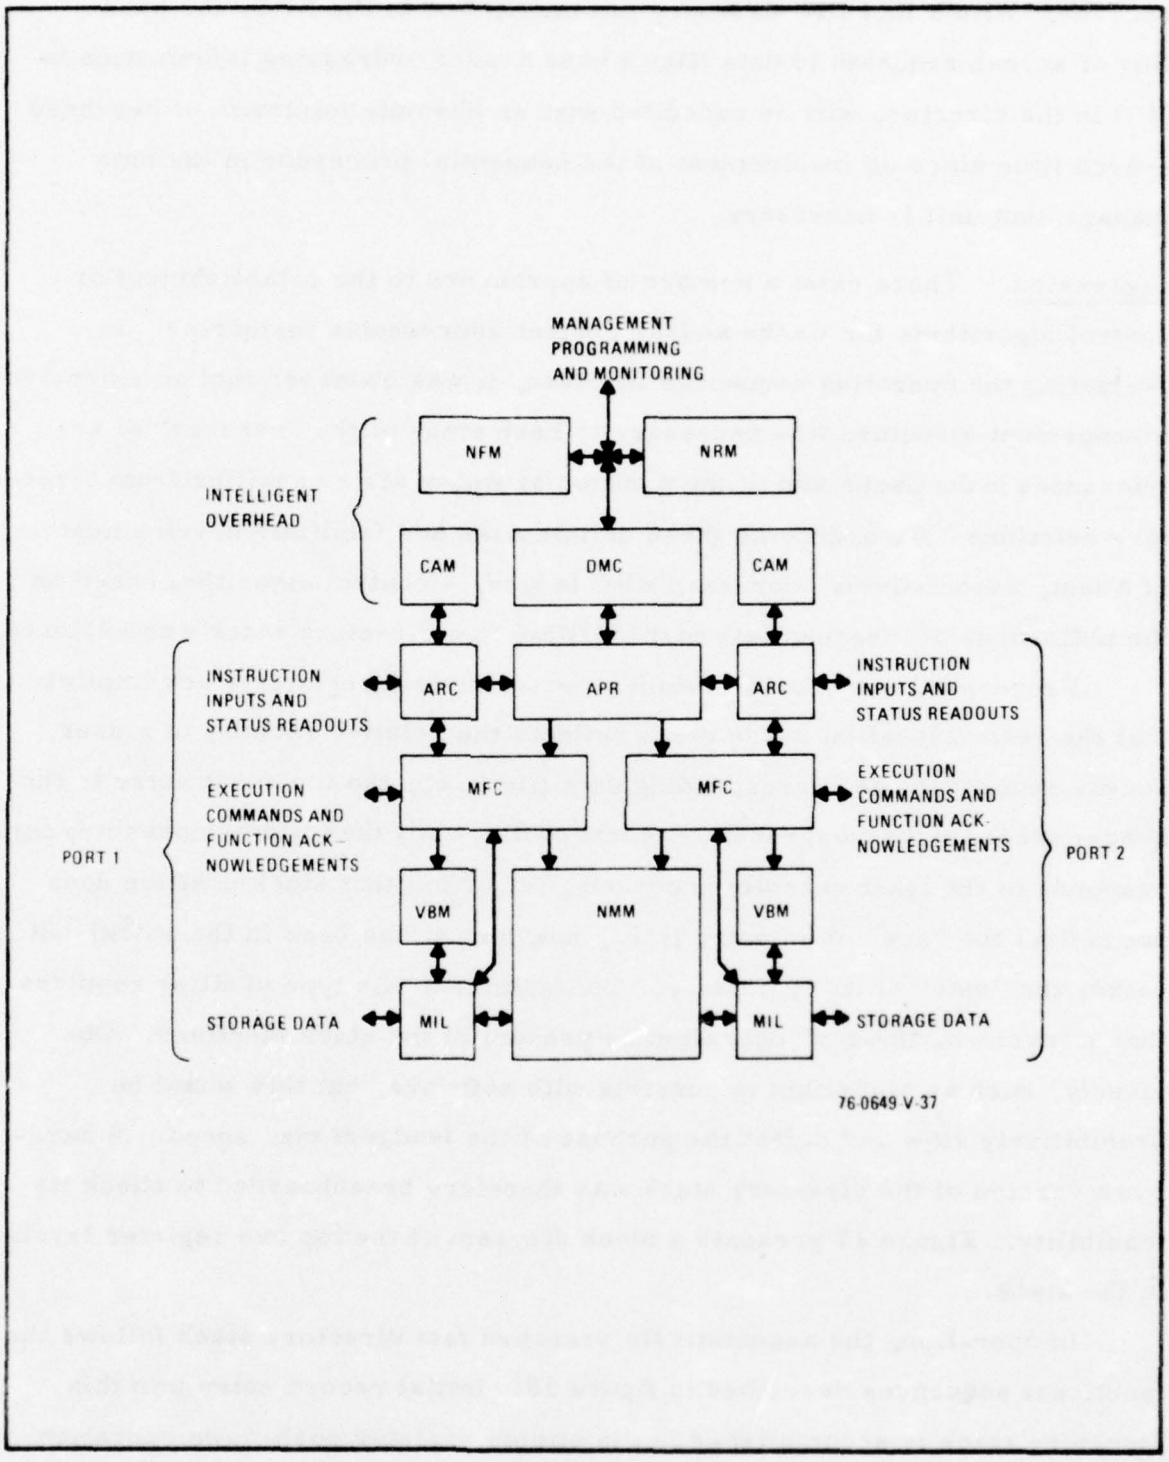

Figure 16. Indirect Access Intelligent MM System Block Diagram

records. With a fast file directory intimately tied to the ARC, the handling of access requests to data files whose header addressing information is still in the directory will be expedited with an absolute minimum of overhead search time since no involvement of the sequential processor in the data management unit is necessary.

Digression. There exist a number of approaches to the establishment of control algorithms for Cache and/or content addressable memories. In evaluating the operating sequences involved, it was observed that an extensive management structure was necessary to keep track of the "recency" of entries made in the Cache and of the position of empty slots resulting from directory deletions. To overcome these deficiencies and facilitate development of a fast, associatively addressed ARC ledger, a control algorithm based on the philosophy of "frequential ordered" filing in a directory stack was explored.

Frequential entry of file header records into a register stack implies that the record location in the stack reflects the relative recency of a user access request to the corresponding data file (i.e., the top-most entry is the header record of the most recently accessed file, while the bottom-most entry corresponds to the least recently accessed file). Note that stack position does not reflect the "age" of an entry (i.e., how long it has been in the stack) but rather the "date" of its last usage. To implement this type of filing requires that a "cycle-to-the-top" operation be present in the stack functions. Obviously, such an algorithm is possible with software, but this would be prohibitively slow and defeat the purpose of the leadger: viz, speed. A hardware version of the directory stack was therefore breadboarded to check its feasibility. Figure 17 presents a block diagram of the top two register levels in the stack.

In operation, the associatively searched fast directory stack follows the functional sequences described in figure 18. Initial record entry into this directory stack is accomplished in the simple register push-down operation of part (A) of this figure. Each new record is loaded into the top register

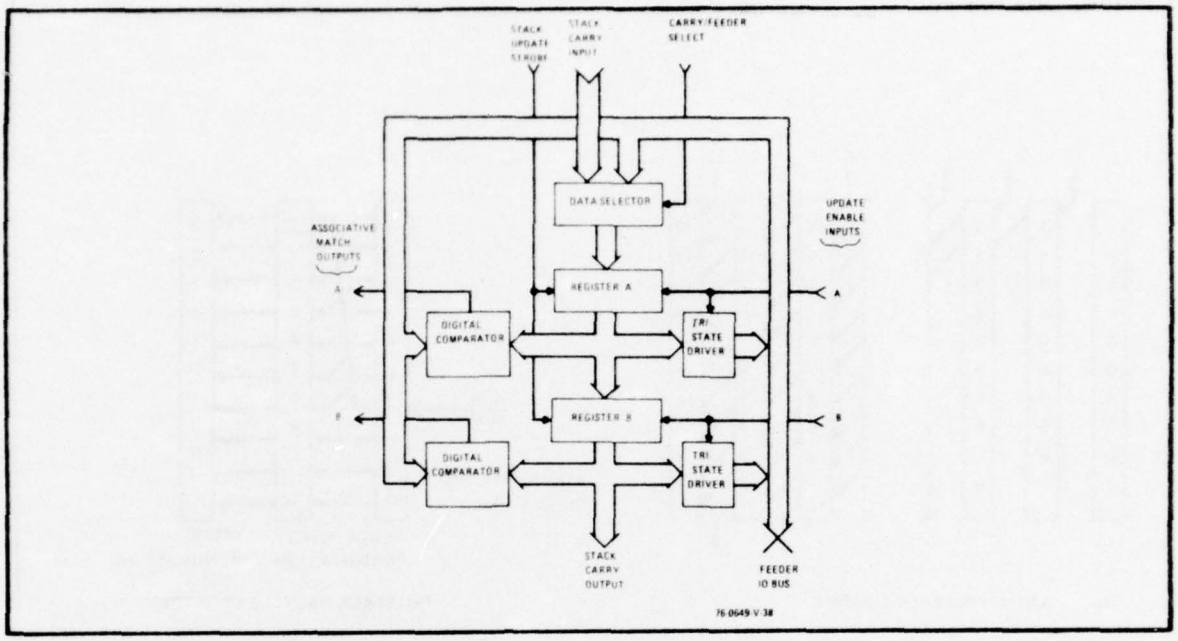

Figure 17. Associative Search Access Directory Stack Block Diagram

with all prior entries being parallel strobed into the next lower ordered register. Ultimately the stack will become filled and, if further entries are not inhibited, the bottom-most record will be lost out the bottom when the stack overflows. This not only presents no problem, it is precisely what is supposed to happen since the stack is intended as a Cached fast search directory for only the "Top Ten" most recently accessed records.

During operation of the system when various records that have been stored are accessed for readout, the fast file directory stack is updated according to the cycle-to-the-top permutation sequence described by (B) of figure 18. The numbering order here is inverted from that of part (A) simply for illustrative convenience. The numbering sequence of (A) represents specific record positions as they move into the stack. The numbering sequence of (B), on the other hand, represents relative record "recency" of being accessed order (i.e., the most recently accessed record is always

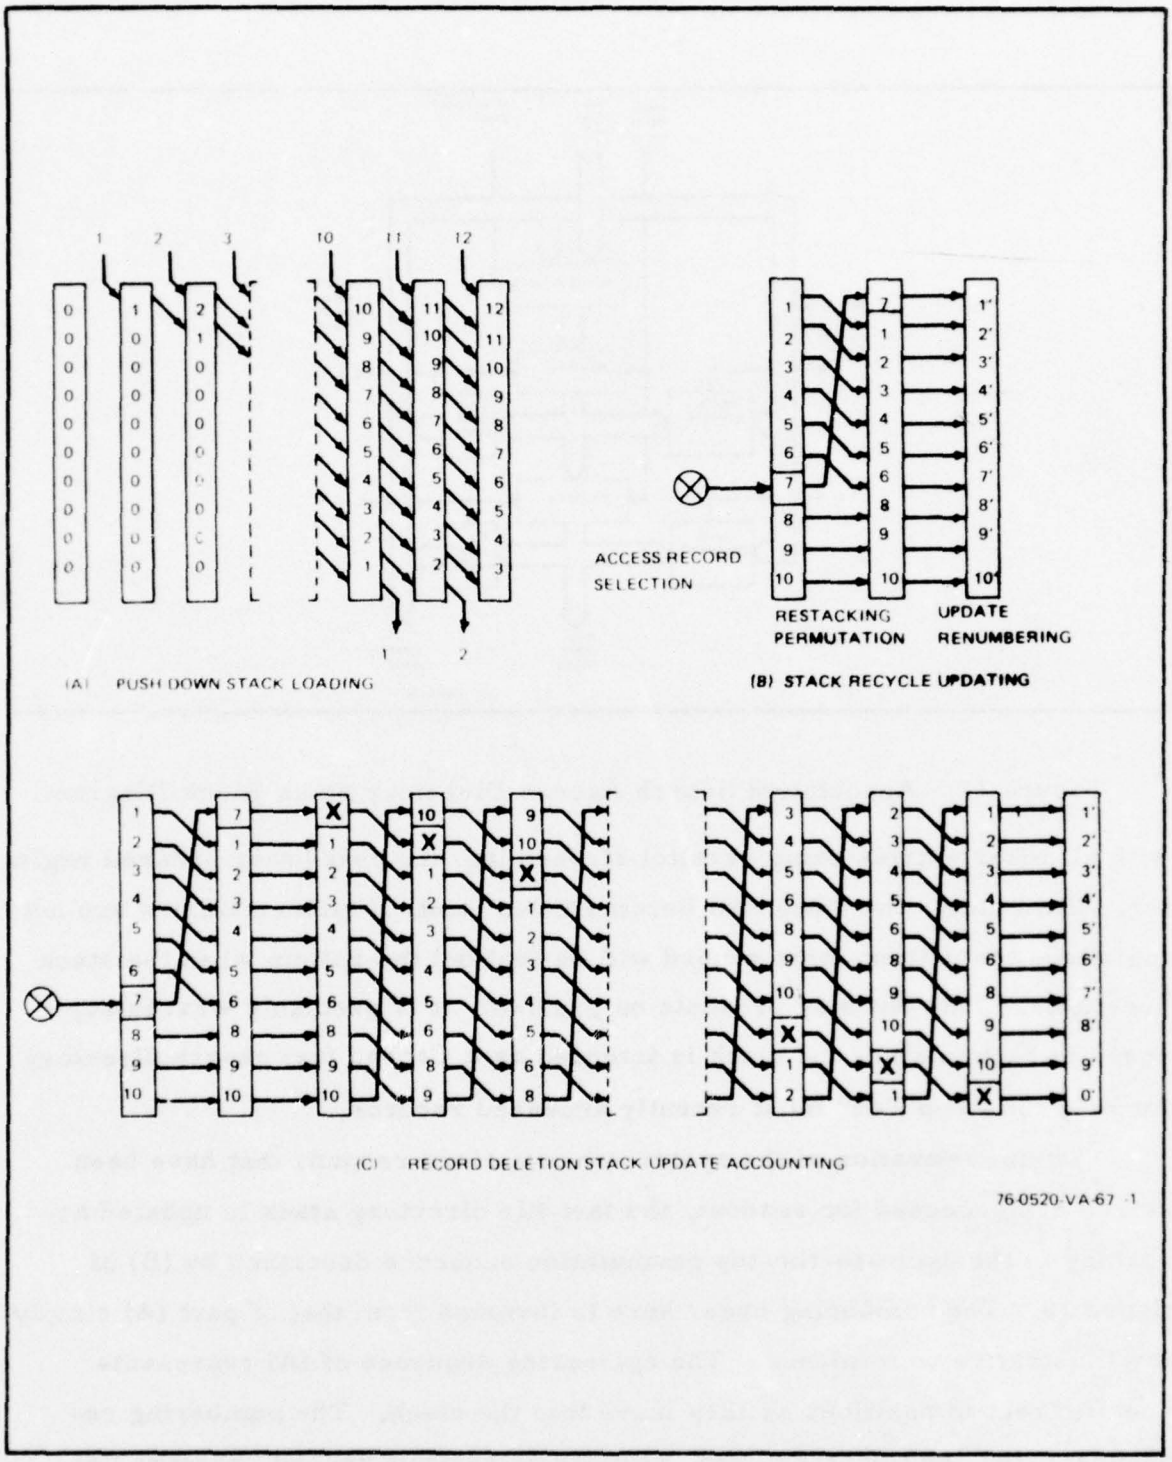

Figure 18. Directory Stack Functional Sequence Diagram

placed at the top and assigned the "frequential" filing number 1.) This updating operation takes place with only a single executive strobe pulse.

Deletion of a record following a clear command causes the frequential style directory to follow the update accounting sequence shown in (C) of the figure. This procedure requires a number of clock pulses identical to the number of registers in the stack and always results in the empty slot in the stack being forced to the bottom-most position (or positions if more than one record is cleared). Consequently, future entry of new records will simply refill the stack in the original push-down manner without losing any prior records until an overflow condition occurs.

The simple and automatic (due to its commitment to hardware) update reordering (or "garbage collection") algorithms possible with the frequential filing technique constitutes an extremely attractive feature. Results obtained with a 4 register stack breadboard built up using standard TTL functions confirmed in every respect the functional algorithm of frequential filing using a cycle-to-the-top updating permutation Operation was observed at a stack update strobing rate of 10 MHz (no attempt was made to measure the maximum achievable rate). For final mechanization of the MM ARC CAM, serious consideration of hardware committed frequential filing is anticipated.

3.1.3 Full Capability Configuration

In establishing the various levels and options in the global hierarchy of a mass memory design possessing universal compatibility with STARAN type AAP computers, two additional facilities options must be included: namely,

- Custom Communication Channel (CCC)
- System Test Stand (STS).

These are shown in Figure 19 combined with a basic two-port MM module and an intelligent overhead management facility to form an application configured, full capability mass memory installation.

All transactions, regardless of their origination or destination points, are coupled through the CCC. This interface unit contains not only all

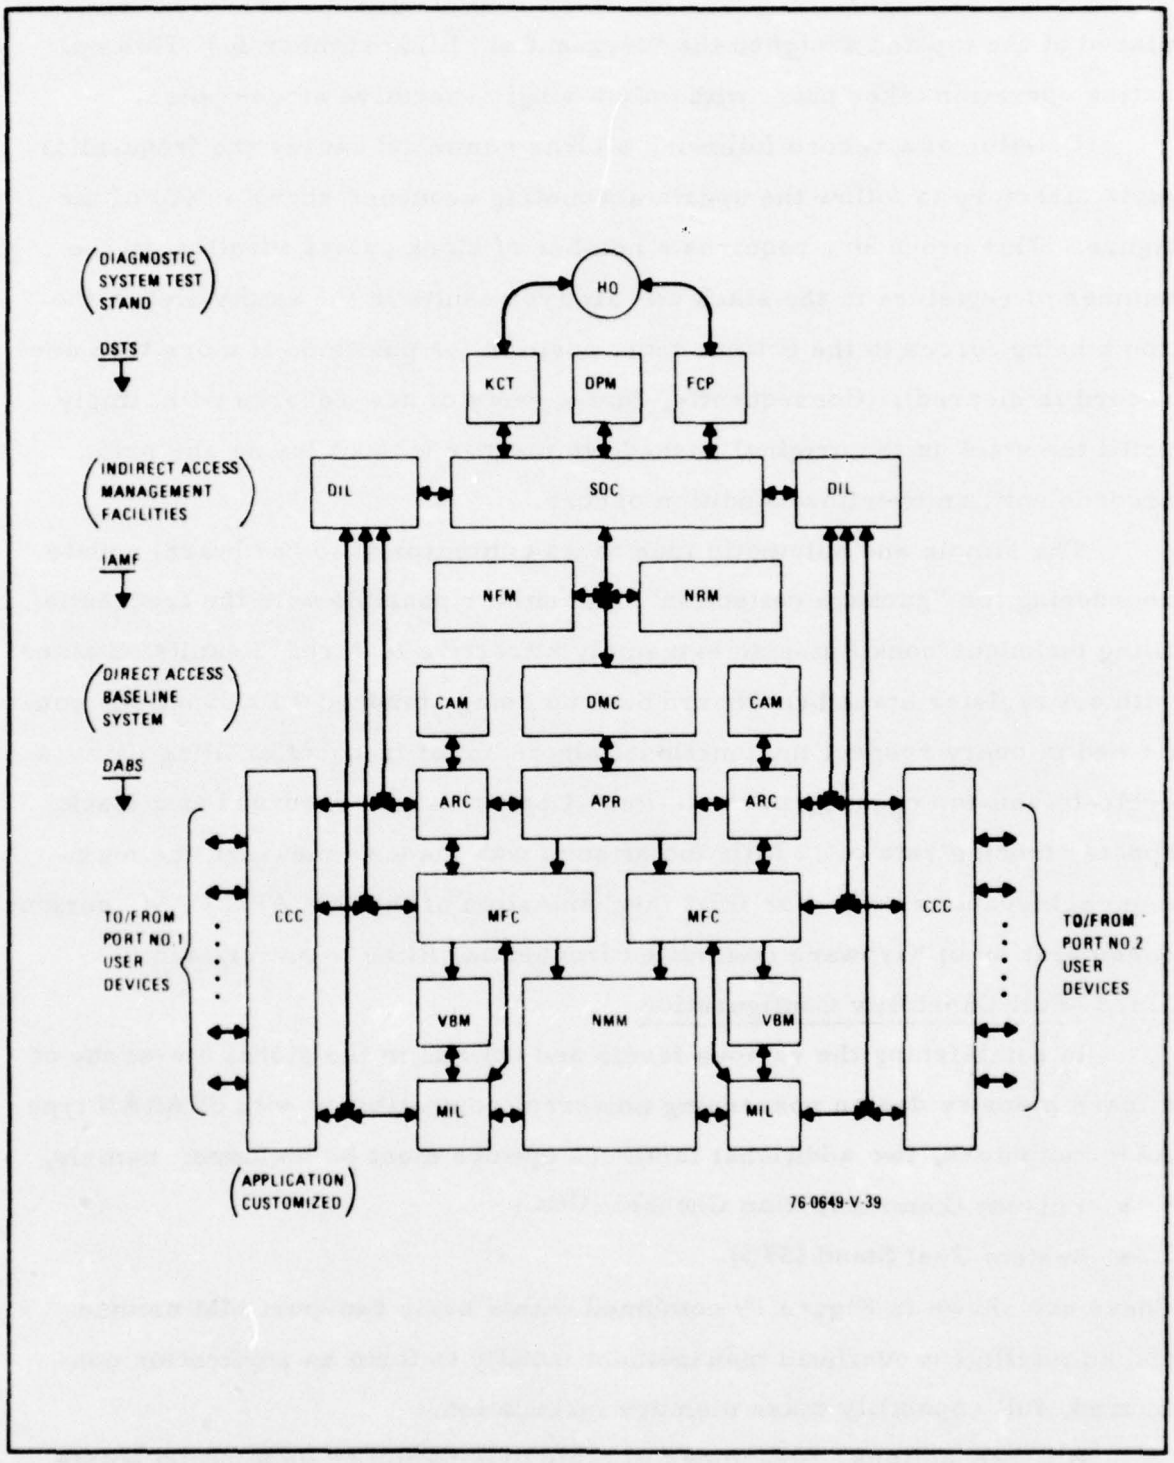

Figure 19. Application Configured MM System Block Diagram

requisite intersystem line driver and receiver circuits but also any command format conversion functions necessary to ensure that all devices tied into the MM receive recognizable signal patterns and sequences. Since the interactive control structures of the mass memory have been established on the basis of being directly compatible with STARAN, the interface unit in a STARAN based AAP installation like that at RADC will contain only the line repeater functions that have been stipulated. For use of the MM in an alternate, generalized multiprocessor environment,other functions to achieve transaction compatibility could also be included.

The final deployment option with which the mass memory can be equipped is built-in-test facilities. In many cases it will be found that these provisions are actually much more important than some of the other basic MM module add-ons. The sole purpose of the STS is to enhance system maintainability by minimizing the time to isolate and correct failures occurring anywhere in the memory. To accomplish this function, the envolvement of six key elements is necessary; specifically,

- Diagnostic Interface Logic (DIL)
- System Diagnostic CPU (SDC)
- Diagnostic Program Memory (DPM)
- Function Control Panel (FCP)
- Keyboard Control Terminal (KCT)
- Human Operator (HO).

The DIL provides the tie-in medium whereby test sequences can be performed on the MM system. The SDC oversees all built-in-test (BIT) operations in accordance with the programmed instructions contained in its CPM. Monitoring and readout of test results by the HO is accommodated via the FCP and KCT. It is anticipated that end use mechanizations of the system test stand will incorporate one of the microprocessors currently being manufactured in volume. With the STS tied-in in the manner shown, it can monitor MM functions "on-line" at close to a real-time basis, and can execute many diagnostics

automatically without human intervention. Furthermore, it provides a means of altering the management programs in the MM intelligent overhead unit.

Through the foregoing presentations, the complete structure of a mass memory having both full compatibility with use in a STARAN environment and adaptive compatibility with most presently foreseen configureable multiprocessor facilities has been established. This, however, is not intended to imply that all elements in the global MM system will be needed in every installation (nor, for that matter, in any of them). The hierarchy evolved incorporates all features foreseen as being needed under a diverse set of application requirements. Any specific installation may be configured to contain as little as the "core elements" identified in figure 15 to as much as the full capability system depicted in figure 19. The procuring agency must decide what level of complexity and performance features constitute the most cost effective tradeoffs in light of the minimum requirements of a particular installation.

3.2 STORAGE ORGANIZATION

Just as the heart of STARAN is its associative arrays, the heart of the mass memory is its nonvolatile main memory section. It is the characteristics of the NMM which determine the fundamental performance capabilities of the basic MM module and hence of any higher-ordered, application-oriented memory configuration. A fundamental observation to be made concerning the NMM is that it constitutes a basic, no frills, "barebones" mass memory. A supplemental requirement imposed on the development of this memory module, though, is the stipulation that it be compatible in every respect with future expansion and changes in any of its application customized support modules.

Secondary requirements included in the operational capabilities of the aggregate system which directly impact on the evolution of storage section designs include:

- Hamming (SECDED) encoding for error detection, correction, and counting (with the counter monitorable under software control)
- Translation of 256 bit width word channels, during I/O transactions involving less than four parallel processor words

- Allowance for unscheduled interruptions (temporary discontinuations or pauses) in I/O word sequencing of indeterminate duration.

Although subtle in context, the latter stipulation here is the key that unlocks the myriad facets of asynchronous operation. Its immediate points of interest lie with the introduction of on-line and off-line holds (idling states), dynamically variable I/O transaction rates, and concordant ability to interface with virtually any type peripheral, even including a totally serial device.

3.2.1 Access Modularity

The first point of major concern addressed in the study of mass memory storage section organizations is the systemological configurations that can be employed to adapt the physically fixed block organization of the MNOS memory chips to the floating or "sliding-window" type block access needed to take maximum advantage of the capabilities of a STARAN type associative multiprocessor. The concept of access to a variable width, word maskable block of multiple word data lines which can be chosen with a totally random starting line position is particularly useful in radar air traffic control and a multiplicity of file search applications. Figure 20 shows how this concept can be visualized in terms of a viewing slit laid over a page of paper containing lines of print. Note that for illustrative purposes each "line" contains 4 words made up of 256 bits.

With complete freedom to shift the viewing window about on the page, adjust its width, and open or close word viewing ports, one can obviously take a snapshot view of any 1-of-16 combinations of words on any contiguously adjacent set of lines. It is also possible to envision splitting the single, word maskable window into four separate word channel windows which can be positioned independently, so long as they all remain in different channels. An additional set of 4 word channel windows can also be conceived, with the restriction that the two windows in each channel do not overlap one another. If this process is carried to the extreme, the whole storage plane becomes covered by single word sized windows. Between the extremes of a single,

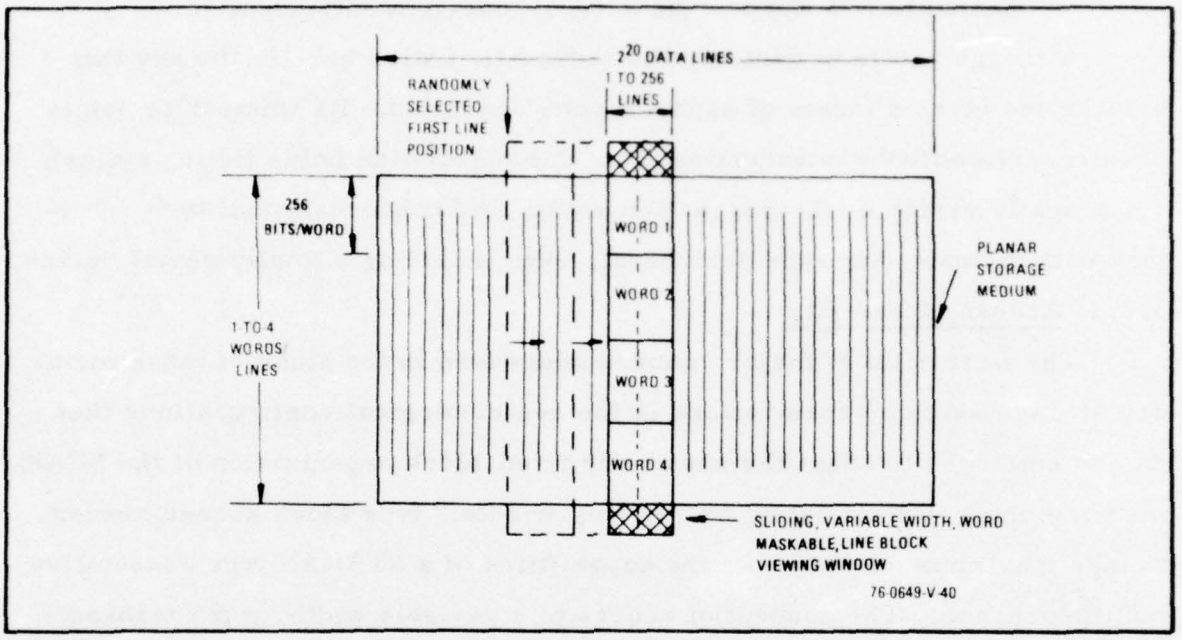

Figure 20. Sliding-Window Access Concept Diagram

full line width, nonmaskable, totally randomly positioned access window and 2^{20} separate, word unique, fixed position windows lies the modularity boundaries which will be optimum for the central storage sections of a mass memory intended for use with multiprocessor machines like STARAN.

Establishing one set of modularity boundaries along 256-bit-wide word channels as indicated provides the ideal path width for maximizing the transfer rate between the mass memory and STARAN's four associative arrays.[1] Review of the size of these arrays reveals that they are square matrices of 256 x 256 bits. Based on this, it appears that the greatest degree of assignment flexibility is secured by structuring the memory access function along

[1] P. H. Enslow, editor; Comtre Corporation; Multiprocessors and Parallel Processing; John Wiley, New York; 1974; p. 53.

the lines of word channel blocks, in the manner depicted in figure 21, such that each block contains 256 sequential channel words of 256 parallel bits each. From both a physical and electrical modularization standpoint, boundaries drawn along these lines have the twofold advantage of:

- Exactly matching the bit capacities of each STARAN array
- Permitting the stacking of modules for capacity expansion
  - back to back to increase the total channel depth.
  - side by side to increase the attainable line width.

RADC has confirmed that this breakpoint in the MM modularization appears to be readily adaptable to a wide variety of potential AAP system applications.

For purposes of addressing within the mass memory, it is necessary to correlate the storage-access organization with the physical-storage modularization boundaries possible with MNOS memory devices. Present schedules for chip developments call for evolution from the present 16 Kbit* device size up to 64 Kbits by early 1979 and to 128 Kbits by mid to late 1980. Based on this schedule, the evolutionary organizational parameters of the 64 Kbit chip are used here to establish the physical hardware boundaries of the MM word channel storage modules.

The storage section of MNOS IC's are, in general, organized in an X-Y or row-column matrix of cells. On the 16 Kbit chip, the storage matrix consists of 128 rows x 128 columns. Expansion to a 64 Kbit capacity is facilitated by doubling the size in both dimensions. Hence, the storage array will contain 256 rows x 256 columns. Addresses applied along the row axis result in the transfer of an entire row of cell data as a unit (or "block") to or from a set of IO shift registers lying along either side of the chip's orthogonal

---

*Note: 1 Kbit is defined to be 1024 bits.
Similarly, 1 Mbit $\triangleq$ 1 k^2bits = 1,048,576 bits
and, 1 Gbit $\triangleq$ 1 k^3 bits = 1,073,741,824 bits.

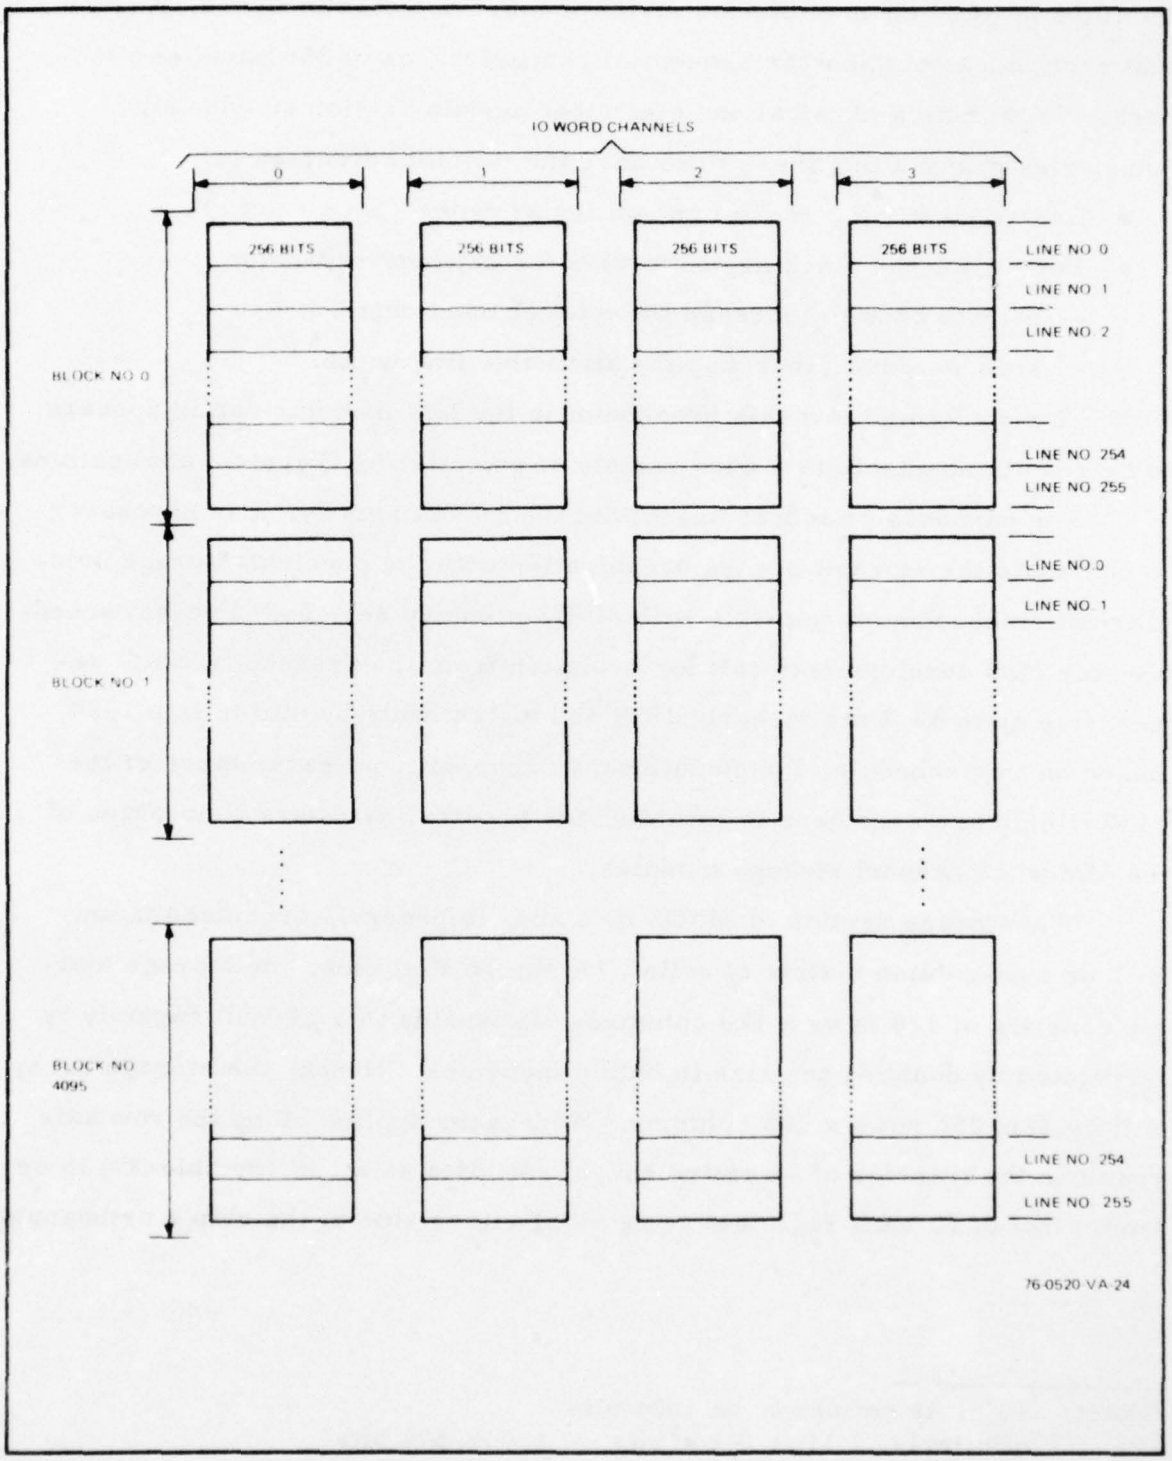

Figure 21. MM Storage Block-Access Organization

column axis. Four registers are presently planned, so each chip will possess 4 parallel IO bit channels with 64 sequently ordered bits present in each register.

To provide for 256 parallel bits/word channel, 64 MNOS memory chips must be stacked side-by-side. Similarly, to fill out a block depth of 256 words will require the end-to-end stacking of 4 chips. Each MM storage module will thus contain 4 x 64 = 256 MNOS memory IC's. Its storage capacity will accordingly be 256 bits/word x 256 words/block x 256 blocks/module = 16 Mbits/module (derivable alternately as 256 chips/module x 64 Kbits/chip). From this, it is seen that each module supplies 1/64th of the total 1 Gbit MM capacity (i.e., 1 Gbit/memory ÷ 16 Mbits/module = 64 modules/memory).

Since provision must be made to service 4 AAP word channels in parallel, it is necessary that the equivalent storage capacity of 16 modules (i.e., 64 modules/memory ÷ 4 word channels/memory) be allocatable to each array channel *during full IO port width transactions.* Viewed in this fashion, the average word channel storage capacity available to each of the 4 STARAN arrays is the equivalent to a backing store consisting of 4096 arrays. This is seen by recalling that 1 MM physical address block possesses the same capacity as 1 AA (associative array) in STARAN. From this, it is observed that 256 blocks/module x 16 modules/channel = 4,096 blocks/channel.

The relationship between the MM storage modules and the manner in which they are collectively accessed can be envisioned in a number of ways. One is to depict the stacking of the individual word channel blocks of figure 21 into four distinct piles having 4,096 layers (blocks) each. Another is to consider full line width groups of 256 blocks to be stacked in 16 distinct piles (called "planes"). From this latter visualization it is possible, through a simple change in terminology, to give a readily recognized physical

interpretation to the modularity interrelationships. In particular, upon making the following terminology substitutions,

- replace "block" with "page"
- replace "plane" with "chapter"
- replace "memory" with "text".

It is seen that the organizational modularity of the mass memory is strikingly analogous to that of an ordinary textbook (see figure 22).

Although this somewhat simpleminded interpretation of the MM accessing organization obviously does not directly allow for all possible manipulations of the data base which software can introduce, it does immediately bring to light the hardware boundaries along which capacity expansion or fractionalization can be accomplished. Starting from the full 1 Gbit baseline size shown, expansion or contraction of the capacity in the depth dimension involves addition or deletion of chapters (each of which represents a plane of memory comprised of 4 word storage modules). Alternately, to vary the memory capacity in the width dimensions involves modifying the page size (viz., by changing the number of words/line) in all 16 chapters.

Here it is to be noted that addition of chapters can easily be accommodated since this is the natural bussing axis within the memory, and no change is necessary in the porting facilities of the system. Conversely, retrofit expansion of the page size in a memory after it has been installed will be accomplished only with difficulty and considerable expense due to the change affecting virtually every physical modularity parameter built into the system. For emphasis of this point, contrast the relative ease with which chapters can be addended to a text with the near absurdity of trying to directly (without resorting to photolithographic processes) increase the size of all its pages after they have been printed!

For purposes of envisioning the manner in which the mass memory can be used in diverse application environments, it is also possible to consider each 256 bit word to be comprised of 64, 32, 16, or 8 bit "characters." If this

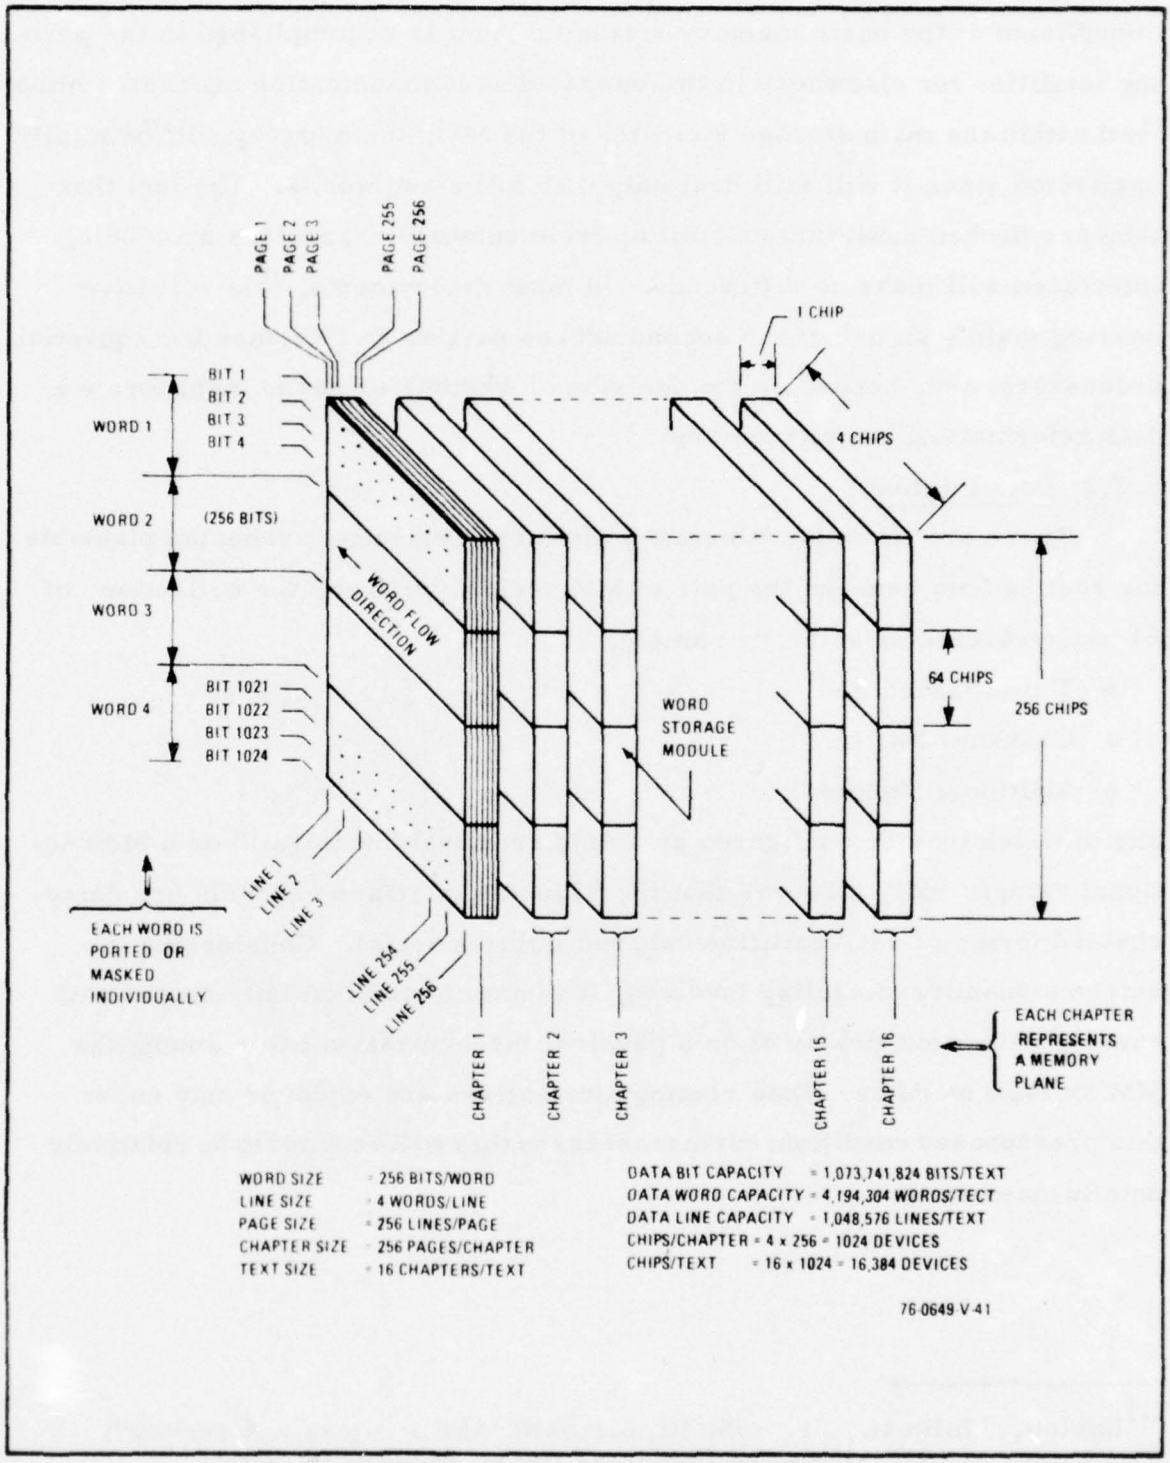

Figure 22. Chaptered Text Modularity Visualization of 1 Gbit MNOS Mass Memory

subdivision of the basic memory storage format is accomplished in the porting facilities (or elsewhere in the interdevice communication channel) rather than within the main storage facilities of the MM, the memory will be totally unaffected since it will still deal only with full sized words. The fact that they are broken down into or built up from subword characters upon being interfaced will make no difference. In most deployments, this will have bearing mainly on use of the second access port as an interface to sequential processors; and, hence, on the decision of whether or not to incorporate a data reformatting buffer memory.

3.2.2 Data Routing

There are three fundamentally distinct interconnect schemes plausible for routing data between the pair of MM access ports and the collection of 64 basic storage modules;[1] namely,

- Time-Shared Bus
- Crossbar Switch
- Multiport Module

any of which may be configured as a unidirectional "dual-rail" or a bidirectional "single-rail" (observe that the STARAN interface channels are daisy-chained forms of dual-rail,time-shared interconnects). Considering the extreme quantity of cabling involved, it is anticipated that only single-rail bussing will prove practical on a physical mechanization basis among the MM storage modules. Data routing alternatives are explored only under this presupposed condition; but in most cases they will be found to be relatively uninfluenced by it.

---

(1) Enslow, Philip H., Jr.; "Multiprocessor Architecture - A survey"; Procedings of 1975 Sagamore Computer Conference On Parallel Processing; August 1975; pp. 63-70.

Although many other factors must be considered in the global MM context, there are three key points upon which decisions regarding data routing networks within the mass memory storage section will be based. These are

- functional isolation of two concurrently operable ports
- facility for intermodule data transfer manipulation
- quantity of circuitry required for mechanization.

The first two evaluation base-points here are derived from the type of performance requirements imposed on an MM which is compatible with AAP type computers. The latter consideration serves as a final screening aid to determine the relative cost-effectiveness and reliability implications of alternatives capable of satisfying the basic performance criteria.

Of the three available approaches to data routing, time-shared bussing is the simpliest and cheapest. It is common in minicomputers because of the ease of expanding the system facilities by merely hanging additional peripheral devices on the bus. It is totally unsuited to use in raw form in the mass memory because it completely violates the first performance criterion and, without help in the interface unit, cannot satisfy the second. Accordingly, only cross-bar switch networks and multiport memory modules are explored here.

Cross-bar switch interconnect schemes constitute the most adaptable - and most complex - approach to configuring the MM data routing network. A completely generalized cross-bar switch interconnect network is shown in figure 23. In this type network, the following three core elements are employed:

- PIU - port interface unit
- MIU - module interface unit
- BIU - bus interface unit.

The PIU provides the hardware interface to the outside word. In a STARAN installation, it would contain the set of line repeaters that have been called

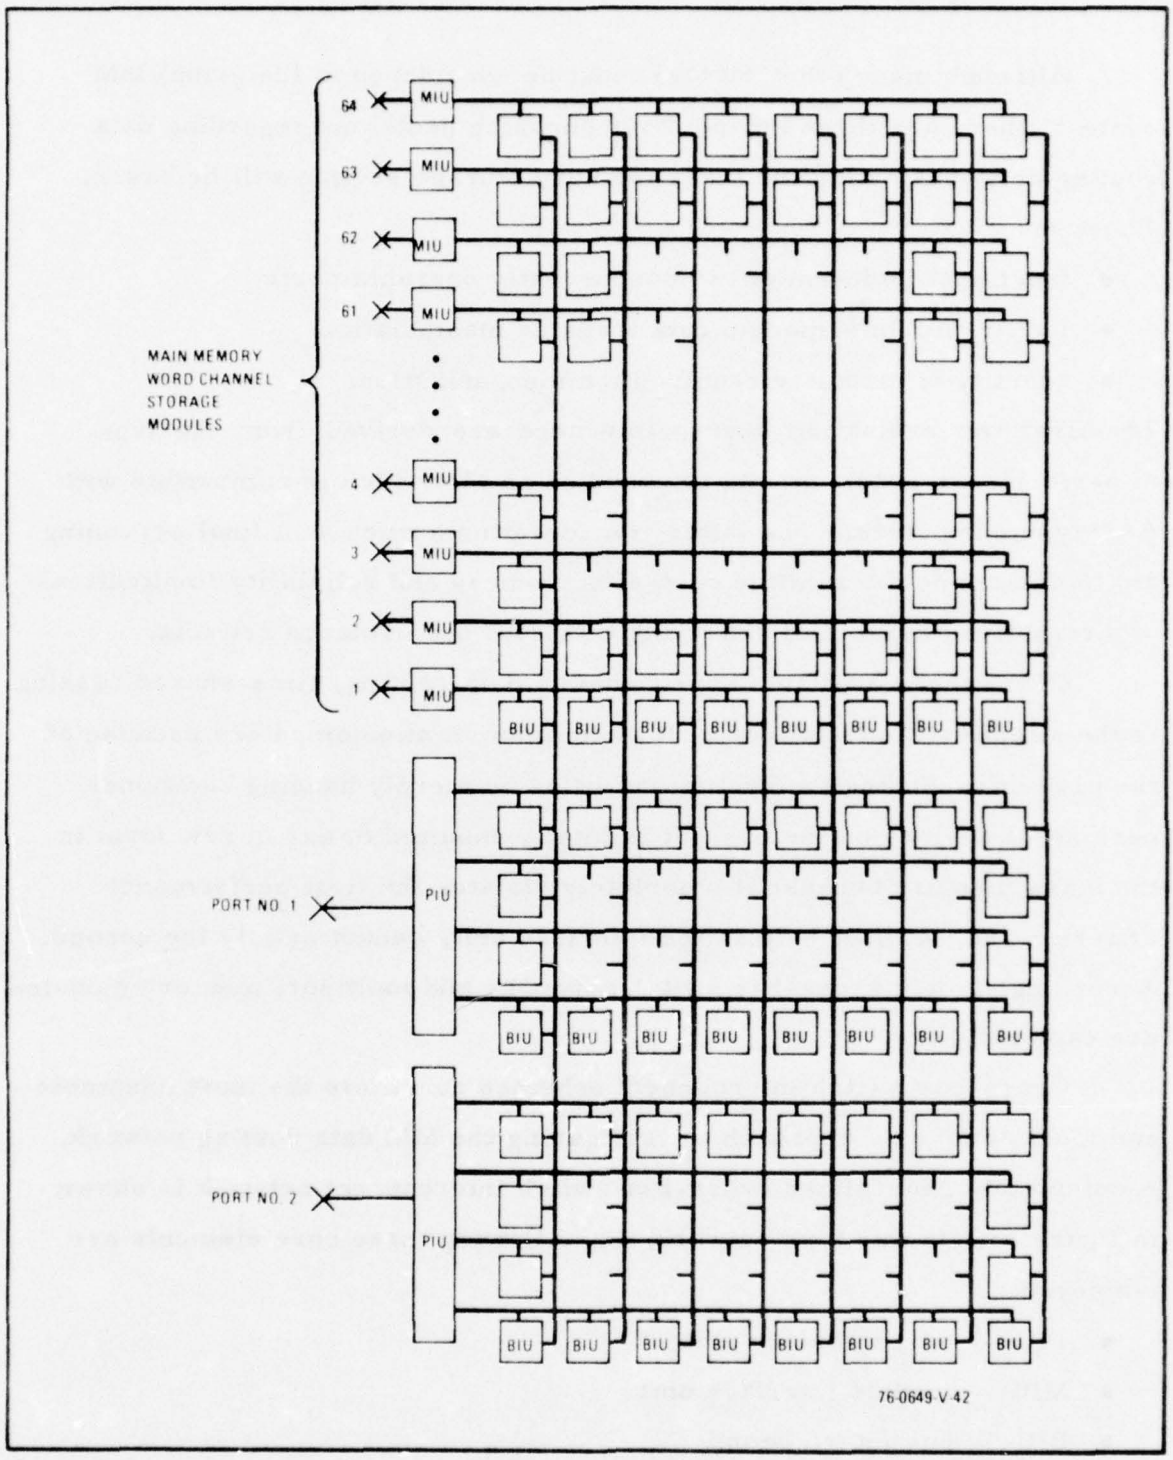

Figure 23. Port Generalized Cross-Bar-Switch Data Routing Network

for in future developments. The MIU supplies the equivalent interface to a word storage module. With MNOS memories it will consist simply of standard-logic line drivers and receivers. The BIU contains the cross-bar switch. It forms the logic connection between the various orthogonal bus lines.

This data routing network is capable of satisfying both of the performance criteria. Furthermore, it is readily expanded since all BIU's are identical, and additional buses can be added without affecting the BIU design. A key feature which makes this type of network attractive in general traffic switching applications, like telephone systems, is its high degree of redundancy and path routing freedom, which collectively provide excellent fail-soft characteristics. It is, however, very costly in terms of total quantity of hardware. Moreover, to capitalize fully on its adaptability and fail-soft options requires a complex network manager.[2]

Allowing for 256 lines in each bus, the simplest possible mechanization of each BIU will require 86 hex IC gate packs (128 would be needed if only quads are available). Since there are a total of 576 BIU's present in the network, a grand total of 49,536 IC's would be needed just for the BIU's. This includes no allowance for either the PIU's or the MIU's, nor for the overhead control structure. It is, therefore, not an attractive alternative for the mass memory.

First-order simplification of the cross-bar switch network of figure 23 can be realized in two different ways. The first is to eliminate the BIU blocks assigned uniquely to each PIU. This results in a network which appears as shown in figure 24. Although less flexible than the original

---

[2] J. H. Rowan, et al.; "Toward the Design of a Network Manager for a Distributed Computer Network"; Proceedings of the 1974 Sagamore Computer Conference on Parallel Processing; Springer-Verlag; pp. 148-166.

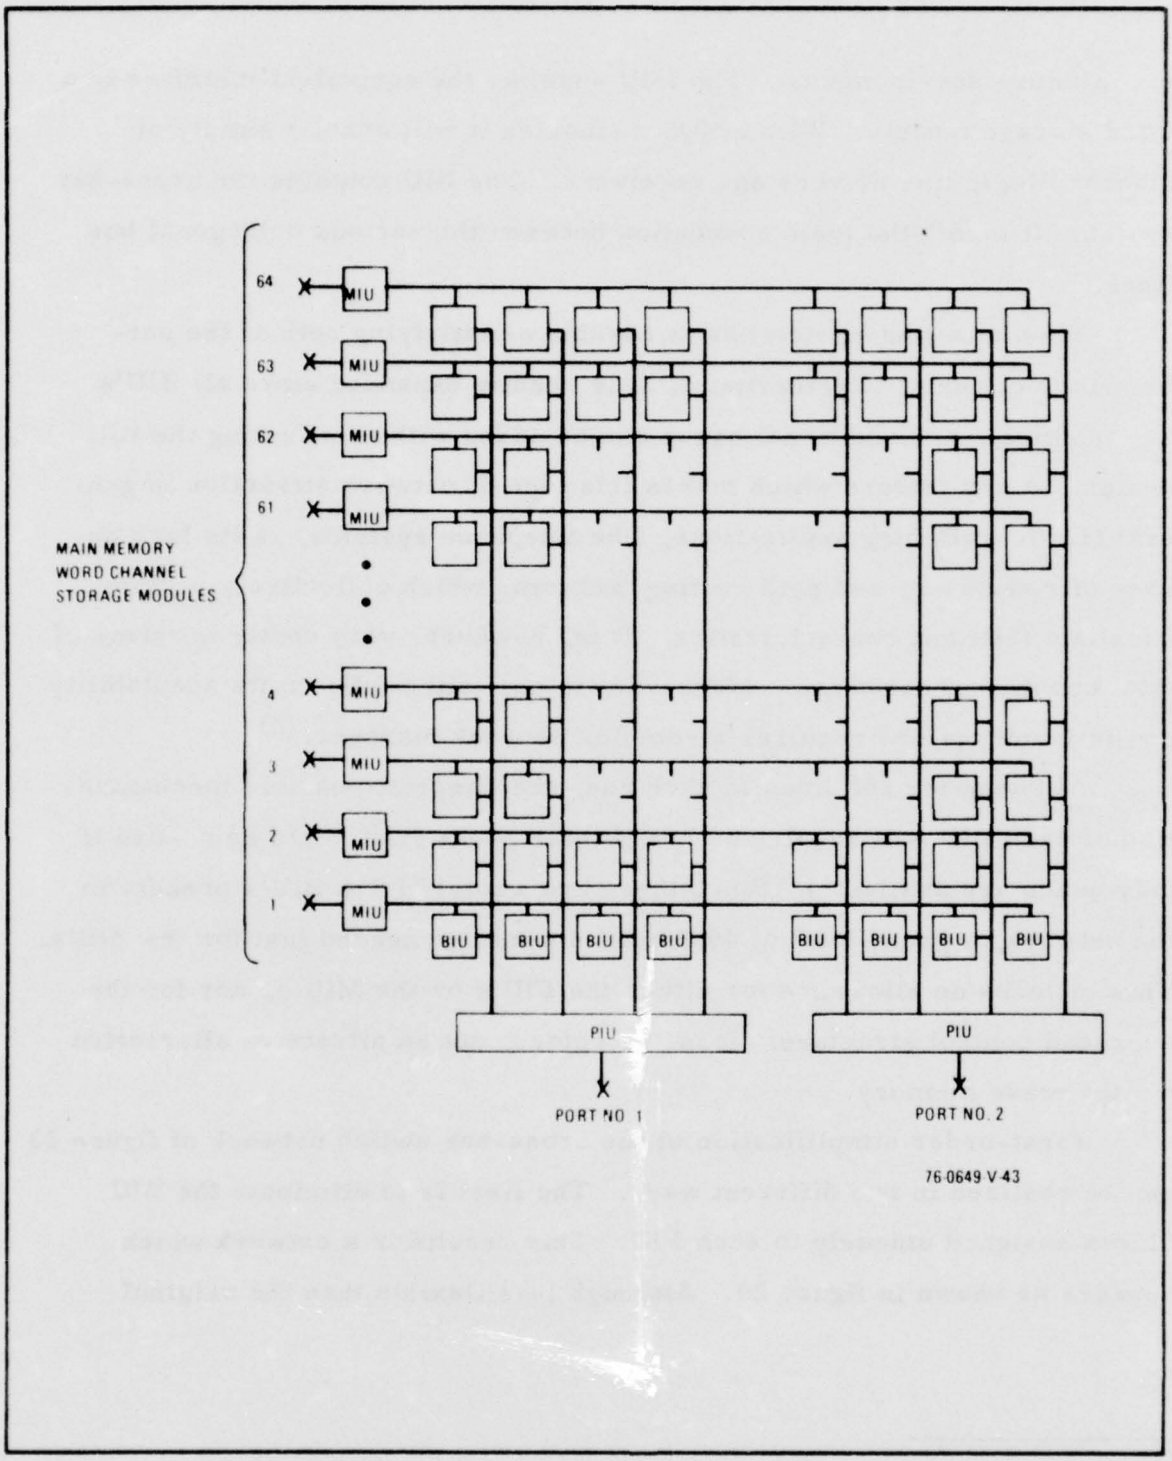

Figure 24. Port Bounded Cross-Bar-Switch Data Routing Network

version , this network is still fully capable of satisfying both performance criteria. However, a check reveals that it still requires an exhorbitant amount of circuitry. Specifically, 512 BIU's are present, so more than 44,032 hex IC's would be needed. It therefore is less attractive than the generalized network due to its loss in fail-soft reconfigurability.

The alternate approach to simplifying the full blown cross-bar switch network is to group the word storage memory modules into a 4 x 16 array and eliminate the large block of BIU's assigned to the MIU's. This process produces the network configuration depicted in figure 25. Review of this design shows that only 32 BIU's are now present. So, the cross-bar switch portion of the network would require only 2,752 hex IC gates. From this standpoint, the design appears quite attractive. Closer examination of the configuration principles applied in forming this network, however, reveals that it is a crossbreed between cross-bar switching and time-shared bussing. It does not provide complete, noninteractive functional independence between the two ports if both are active concurrently. It, therefore, fails the first performance criterion.

Despite the inability of this network to satisfy the port isolation requirement, its low mechanization device count will make it an excellent choice in situations where a limited degree of access conflict can be tolerated or worked around. An example of such a condition would be where one of the MM ports interfaces to a high-speed parallel processor communication channel while the second ties into a low-speed sequential processor channel. In such cases, the degree of cycle stealing from the high-speed channel necessary to service the low-speed channel, or alternately the hold-off delay in servicing the slow channel, would have negligible impact on throughput over either. For purposes of this study, though, the network must be ruled out due to its failure to fully meet the isolation criterion.

The remaining approach to structuring of the mass memory storage section is to employ dual-porting at the module level in the manner shown

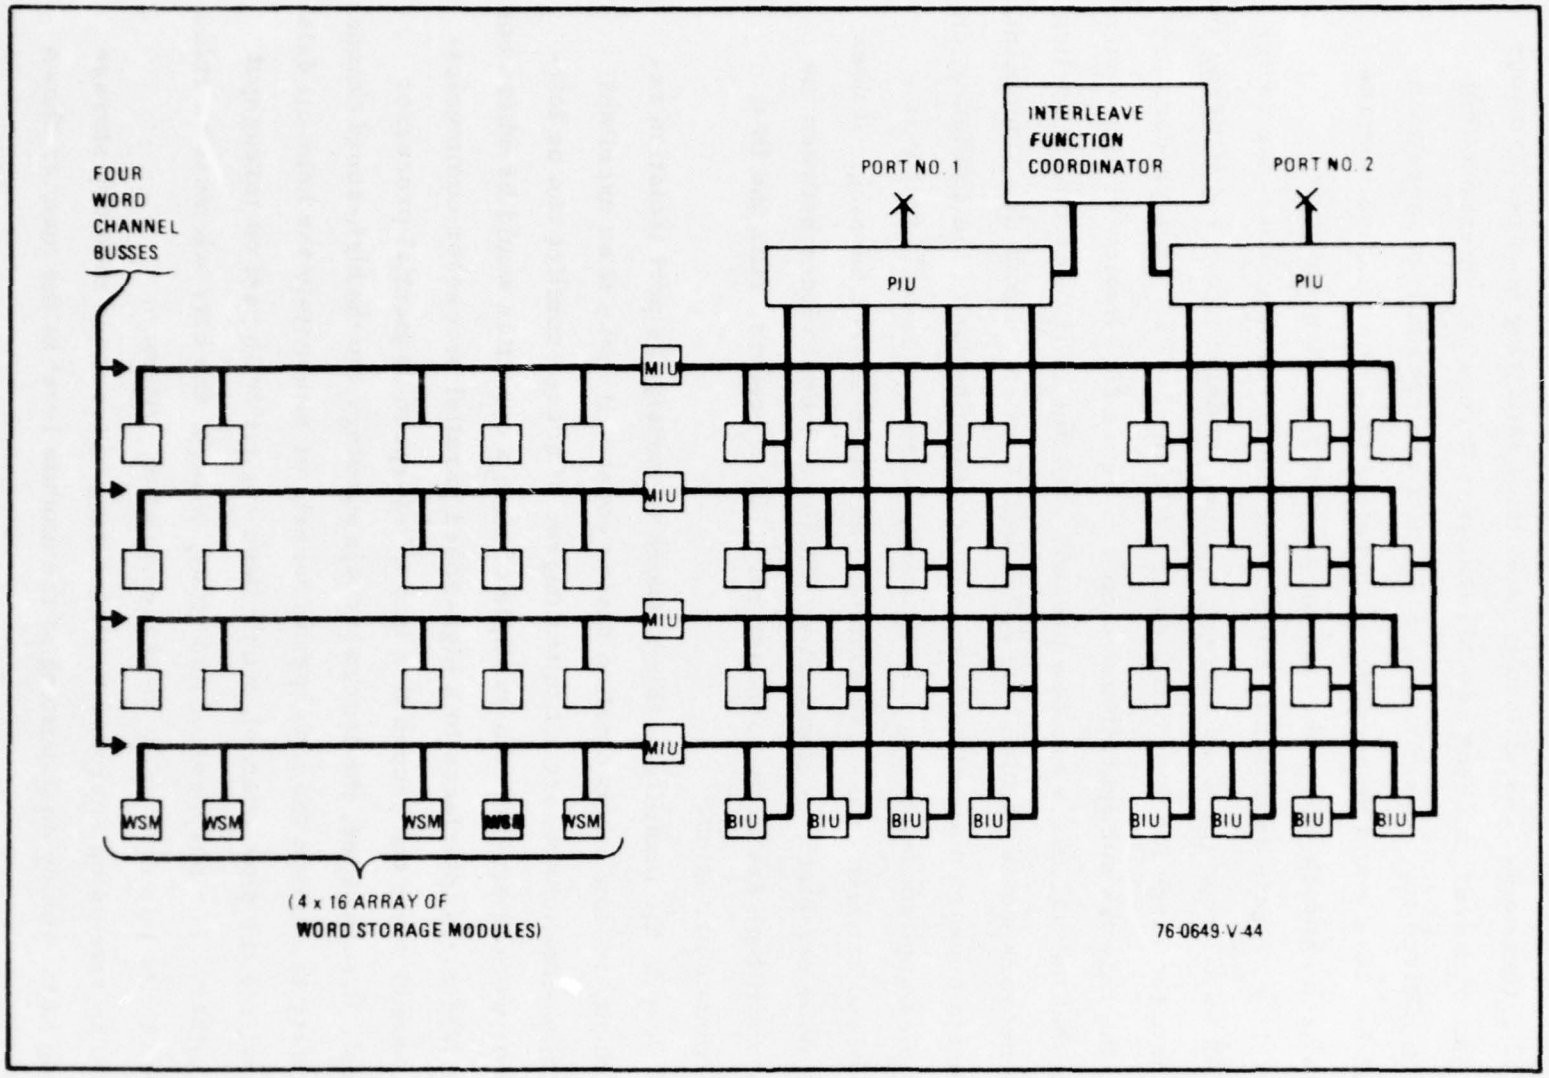

Figure 25. Port Coordinated Data Routing Network

in figure 26. This scheme intrinsically satisfies the first performance criterion but requires the presence of port unique channel permutation networks to meet the second criterion. Selection of the proper design option for these permutation networks constitutes a major tradeoff concern in the multiport module approach to configuring the MM storage section. It is these networks that control the composite system's ability to manipulate data among the various storage modules.

One obvious approach to designing a permutation network is, of course, use of a cross-bar switch arrangement like that in figure 25. On a per-port basis, some 1,376 IC's (minimum) would be required. Alternate networks capable of accomplishing the full N! permutations possible among the 4 word channels also exist. One of the best is the TDRN (triangular data routing network) as suggested by Kautz[3] and expounded on by Chen and Frank.[4] This type of network appears as shown in figure 27. The logical connection blocks indicated here are a modified form of cross-bar switch which provides programmable control of whether a crossing or bending mode of IO throughput occurs. It is possible to implement the basic, unidirectional network shown with 768 quad 2:1 multiplexer gate packs. Since permutations in both directions must be provided, a complete mechanization for either MM port would require 1,536 IC's. The point of note is that this is only about 10% more packs than were necessary with the basic cross-bar switch. Since the TDRN may require less control gating support, no clear-cut advantage exists between it and the CBS network on the basis of part count.

---

[3] W. H. Kautz, et al.; "A Cellular Interconnection Array", IEEE Transactions On Computers, Vol. C-17, No. 5; May 1968.

[4] C. J. Chen (Norand Corp.) & A. A. Frank (U. of Wisconsin); "On Programmable Data Routing Networks via Cross-Bar Switches for Multiple Element Computer Architectures"; Proceedings of the 1974 Sagamore Computer Conference on Parallel Processing; Springer-Verlag, New York; 1975; pp. 338-369.

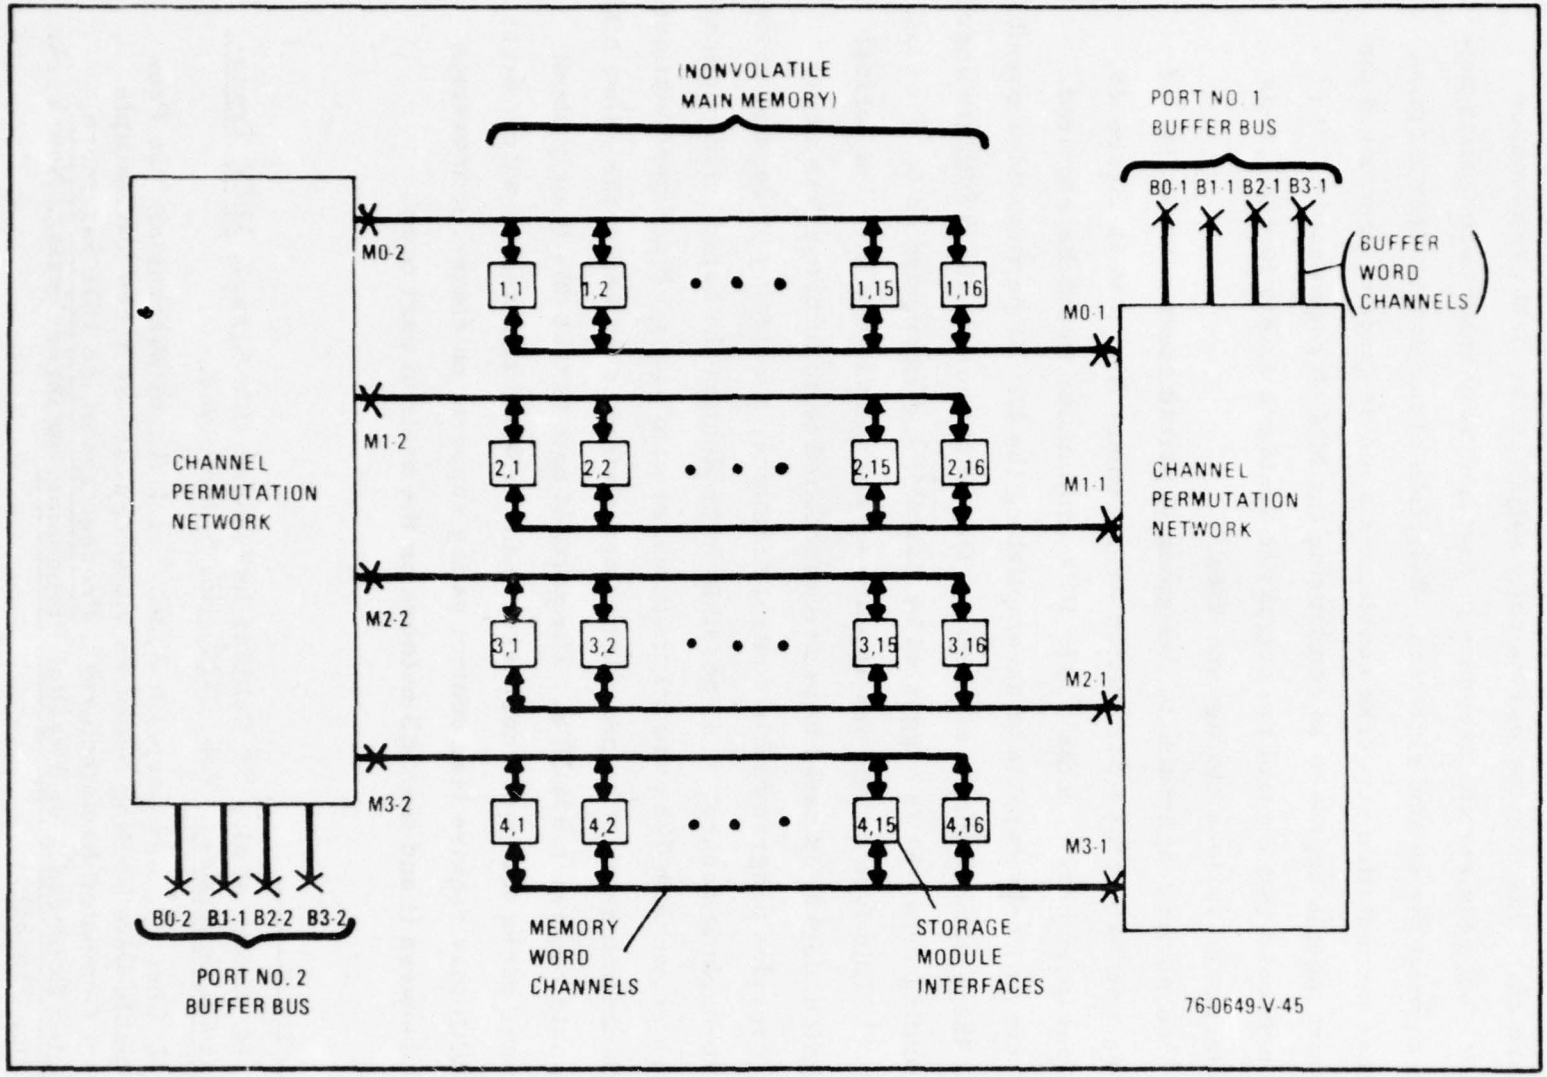

Figure 26. Port Dualed Main Memory Bussing Structure

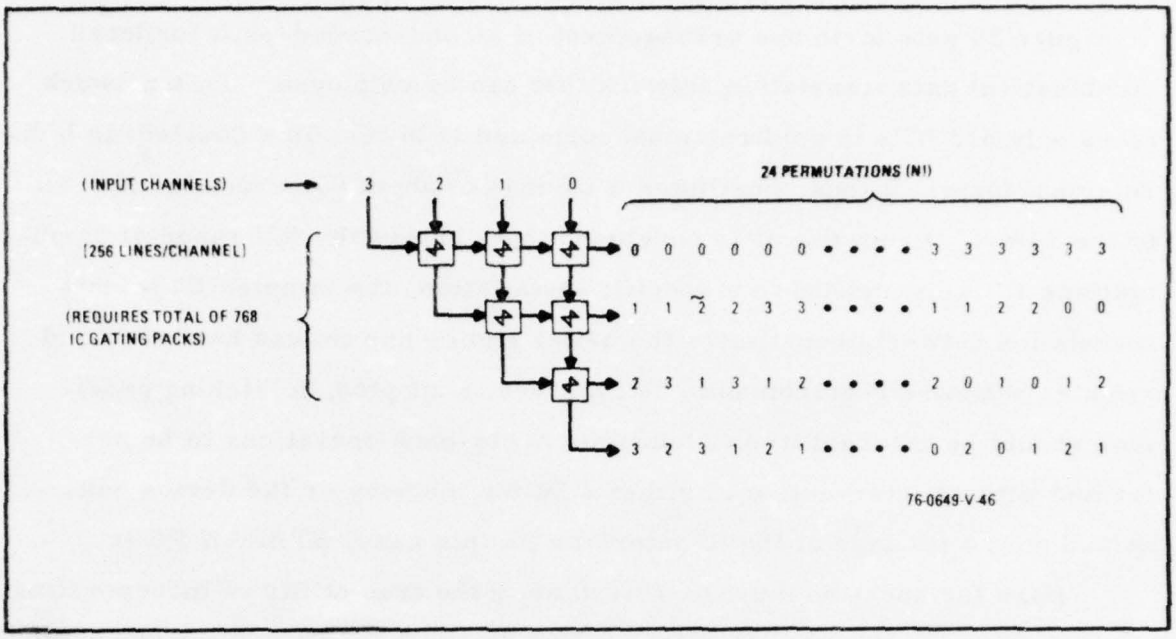

Figure 27. Full Permutation Capability Triangular Data Routing Network

In reviewing the control algorithms that would be necessary during transfers between STARAN and the mass memory, RADC personnel observed that system software could be simplified with little loss in general utility if the range of permutations was restricted to the four "end-around-shift" combinations. With these four available, transactions involving less than all four STARAN arrays would always be interfaced via the lower-most PIO ports. Specifically,

4 array transfers - via PIO ports 0, 1, 2, and 3

3 array transfers - via PIO ports 0, 1, and 2

2 array transfers - via PIO ports 0 and 1

1 array transfer - via PIO port 0

With this restriction imposed, it is possible to simplify the channel manipulation network still further.

Figure 28 sets forth one arrangement of an end-around-shift (ordered combination) data translation network that can be employed. This network takes only 512 IC's in unidirectional form and 1024 IC's in a doubled-up bidirectional form. It thus constitutes a savings of about 50 percent compared to the TDRN. From this it is concluded that, unless the full range of permutations is truly needed in a specific installation, the simpler DTN (data translation network) constitutes the better choice due to less hardware and simpler software requirements. Whichever is adopted, a latching provision should be incorporated to facilitate cycle-back operations to be performed without intervention of either a buffer memory or the device connected on the far side of the IO interface (in this case, STARAN PIO).

Here the question must be raised as to the true utility of incorporating a buffer memory into the MM data routing hierarchy. This question can be answered by addressing the following:

"What is the purpose of the buffer memory?"

"What functions must it perform?"

"What is its contribution to the aggregate system?"

The only characteristic of the MNOS main storage section devices which might be aided by the presence of a BM during normal IO read or write transactions is the minimum data transfer rate. Maximum flow rates of under 200 nsec/bit can be maintained continuously without any kind of help. Minimum rates with current MNOS device designs (which incorporate dynamic 2 phase IO data block interface shift registers) is limited over temperature to greater than 2 milliseconds/bit.

By way of the cycle-back capability specified for the MM channel permutation network, an accessed block of data can be recycle refreshed indefinitely. This allows a data block to be held in an IO ready state for whatever delay time might be necessitated by a system interrupt, without intervention by (or even existence of) a buffer memory. Contrarily, by making use of the BM, the MM storage devices could be powered down during any lull

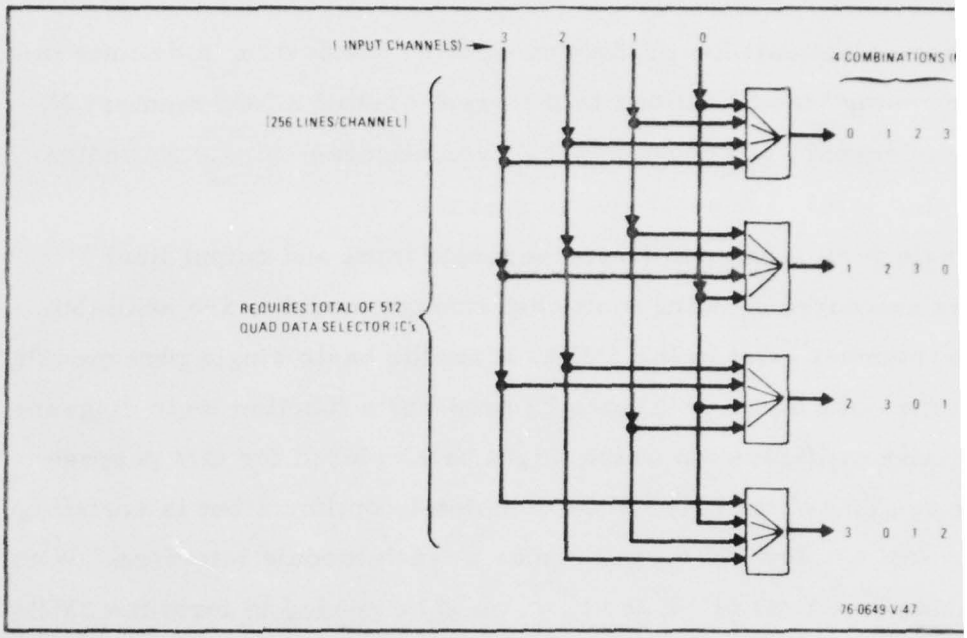

Figure 28. Simple (End-Around) Ordered-Combinati
Data Translation Network

periods in IO transactions. This has significant impact on relia
dead times occur frequently. Accordingly, the BM might be jus
these instances for use as a low-end rate buffer. Detailed stud
projections will have to be made on usage duty factors of applic
grams to make any final decision. This information was not av
the time of this study.

The remaining point of consideration in evaluating the mu
module MM interconnect scheme is the nature of the storage mo
face. Two extremes exist. One, the obviously most desirable
from the standpoint of mechanization hardware, is where the M
chips have been "groomed" to already possess dual-port, bidire
ing capabilities. In this case, no overhead support circuitry at
in the SMI (storage module interface) blocks of figure 26. Since

development plans do not call for production of this type device, a definite investment in "grooming" modifications to the basic 64-kbit MNOS memory IC design would be necessary to secure especially configured chips. No indication exists that this level of support can be guaranteed.

If only single port, dual-rail (i.e., separate input and output line) memory devices and corresponding multichip storage modules are available, sufficient circuitry must exist in the SMI to adapt the basic single port module to the dual port memory format. Figure 29 presents a function logic diagram of the type circuitry configuration which might be employed for this purpose. This particular design, which may not be completely optimum but is certainly representative, will employ 192 IC gate packs in each module interface. With 64 modules/memory, a total of 12,288 IC's would be needed to form the SMI's. Adding the device counts of the permutation (or translation) networks into this yields a final figure of either 15,360 IC's (using the TDRN) or 14,336 IC's (using the DTN) for the support circuitry of the port-dualed approach to MM storage section interconnect bussing.

Although the quantities of support circuits are still rather high ($\gtrsim$90 percent overhead in support devices compared to memory chips), especially when compared to the port coordinated CBS technique, anyone of the three versions of the multiport-module data routing system is a viable candidate for adoption into subsequent MM prototype developments. If full permutation is anticipated, the TDR (or CBS) approach should be taken. If this can be definitely eliminated in favor of a reduced end-around-shift, ordered combination subset of channel translations, then the best tradeoff occurs by using the DTN arrangement. Table 6 presents a summary of the relative merits of each approach in light of the three prime performance-evaluation criteria.

The third choice option is the port coordinated version of the cross-bar switch technique. The reason for this is that in situations where cycle stealing causes no problems, this option provides the best choice. Any final decision among the three prime contenders will have to be based on further

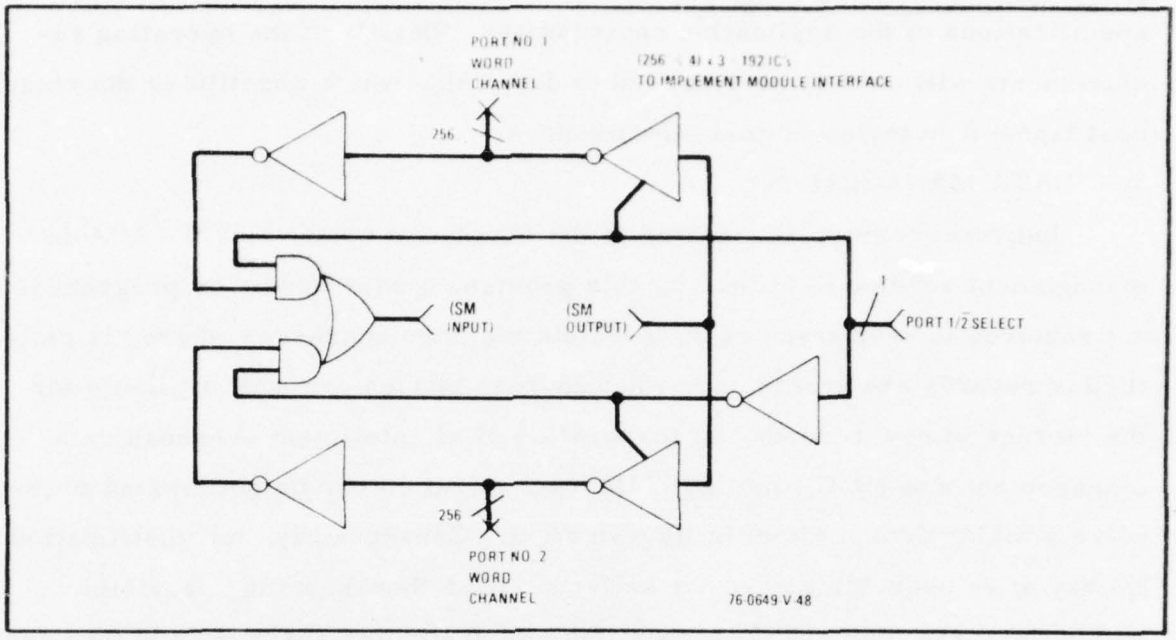

Figure 29. Dual-Porting Storage Module Interface

TABLE 6

MM INTERCONNECT TECHNIQUE TRADEOFF SUMMARY

INTERCONNECT TECHNIQUE	EVALUATION CRITERIA (Isolation)	(Permutation)	(Mechanization)	
Time-Shared Bus	No	Not Directly	–	
Cross-Bar Switch				
a. Port Generalized	Yes	Yes	49,536	
b. Port Bounded	Yes	Yes	44,032	
c. Port Coordinated	Yes Time Interleave	Yes	5,504	*(3)
Multiport Module				
a. CBS Interface	Yes	Yes	15,040	
b. TDRN Interface	Yes	Yes	15,360	*(1)
c. DTN Interface	Yes	Reduced Set	14,336	*(2)

* DENOTES PRIME CANDIDATES
(1) Full capability first choice
(2) Reduced capability second choice
(3) Qualified capability third choice

76-0649-T-49

specifications of the application environment. Details of the operating requirements will need to be assessed to determine which constitutes the singular best tradeoff in terms of cost effectiveness.

3.3 DATA MANAGEMENT

Indirect memory addressing is the foundation upon which the MM data management scheme is built. By this proviso, a user device or program is not required to keep track of the absolute memory addresses where his particular records are stored, nor what address spaces are still available for the storage of new records. Incorporation of an intelligent overhead data management unit (DMU) into any MM mechanization can be anticipated to involve a major design effort in its own right. Consequently, full justification by way of an overriding need for self-contained "bookkeeping" provision within the mass memory must exist for any particular application before such a development is undertaken.

Due to the complexity of the design task, it is not the intent of the mass memory organization study, and is not the intent of the present discussion, to establish final data management design requirements: only the general framework on which a coherent development can be based. Bookkeeping, or data management, constitutes a significant design concern having many as yet unresolved facets.

Considering all aspects of the data and file management task necessary in the mass memory, there are four global philosophies that can be adopted: historical, reflecting the original record entry order; frequential, reflecting the most recent record access order; spatial, reflecting an available space listing order; and preferential, reflecting a pointered page listing order. Of these, spatial filing is totally static in nature while frequential is fully dynamic. Historical and preferential lie between test two extremes and exhibit both static and dynamic aspects.

Of the four basic approaches to record filing, historical and frequential currently appear to possess the most attractive parameters for use in the fast search access ledgers of the access control unit (as was noted earlier). Each

uses comparatively simple and straightforward file entry and updating algorithms. *Functionally, the historical technique takes on the appearance of a* FIFO register stack. In contrast, the frequential approach is somewhat analogous to a LIFO stack. (Note: FIFO = First-in/first-out and constitutes a push-down-fall-through operation; LIFO = last-in/first-out and constitutes a push-down, pop-up operation.) Either provides the speed necessary to a practical access ledger, but device count at the final implementation level counter indicates their use in the DMU.

A spatial filing system is conceptually simple but too inflexible to constitute a truly viable alternative in either a parallel or sequential processor MM system. It does not, therefore, warrant further serious consideration. Preferential is the common technique used in conventional computer programming. Since presently it appears that the most judicious manner of mechanizing the DMU is by a standard minicomputer such as a PDP11, the preferential, pointer list approach is proposed for the DMU.

The general format of directory file entries foreseen for the DMU record directory is as shown in figure 30. The contents of each entry consist of

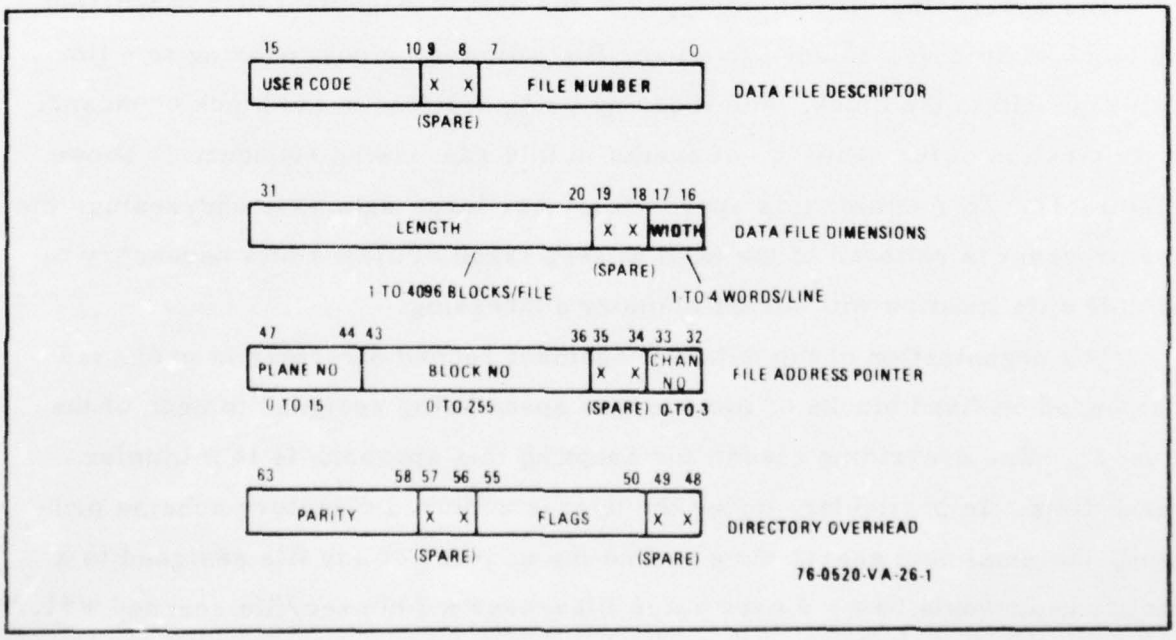

Figure 30. Indirect Memory Address Directory Record Fields

indirect addressing file discriptor information, the absolute address of the first (lowest ordered word channel) data block address, the dimensions of the storage space allocated to the specific user file, and miscellaneous management control flags and partity bits. The size of the entries in the data file directory is therefore anticipated to encompass on the order of 64 bits.

The minimum storage capacity requirement for the management directories is found from consideration of the maximum number of users anticipated and the average number of files each will be allowed. By direction of RADC, to establish the total capacity of the management memories, allowance is to be made for at least 64 users with each having 256 files. This implies that $64 \times 256 = 16K$ record storage spaces must be present in the DMU data file directory. Hence, the size of the DMU file-record memory is projected to be at least 16K entries x 64 bits/entry = 1 Mbit. In the case of the resources allocation memory, the available capacity conceivable may have to be >4 Mbit if it is necessary to allow for totally arbitrary data file boundaries (i.e., 1 Gbit $\div$ 256 bits/word).

The addressing format employed in the MM management unit is indirect to a block of storage assigned to a user file followed by postindexing to a line of storage within the block, with indexing being relative to the block boundary. An illustration of the sequence of events in this addressing sequence is shown in figure 31. By adopting this approach to very large data base addressing, the user program is relieved of the need to keep track of all the bits necessary to absolute data location with direct memory addressing.

The organization of the data management record directory is projected to be based on fixed blocks of file address space being assigned to each of the 64 users. The overriding reason for adopting this approach is to minimize search time. In particular, under the user-structured directory scheme proposed, the maximum search time to find the address of any file assigned to a specific user would be τ = 1 user x 256 files/user x 200 nsec/file scanned $\simeq 51.2$

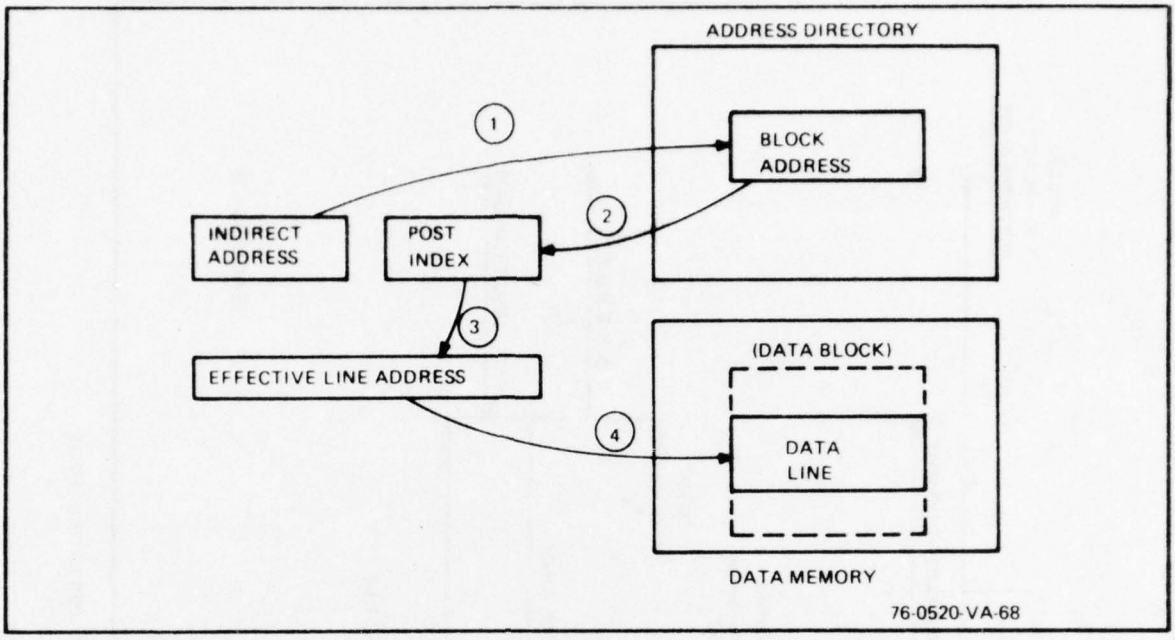

Figure 31. MM Indirect, Block-Relative, Post-Indexed Addressing

microseconds. Alternately, if a completely random ordering of any given user's file addresses is instituted, the total search time could extend to τ' = 64 users x 256 files/user x 200 nsec/file scanned $\simeq$ 3.28 milliseconds.

From figure 32, note that the total access time for each storage mode operation entails not only the direct access time of the memory chips, but also considerable overhead time. In all, five distinctly identifiable units of time are involved in getting the first line of a file set up at the IO port ready for transfer. The management search time cited above constitutes the first increment (τ'). Obviously, if the goal is to achieve an average access time of 50 μsec, a maximum of several milliseconds cannot be tolerated in the DMU.

Following the overhead DMU time, approximately 5 μsec is allotted to local storage supervisory functions in retrieving succeeding block pointers for the accessed file from their local listing. A maximum of 2 μsec is then allowed for activating the plane of memory chips containing the block that has been addressed. Finally, 3 μsec is set as the limit on delay in getting the first line of data off the block and to the output of the plane. At this point

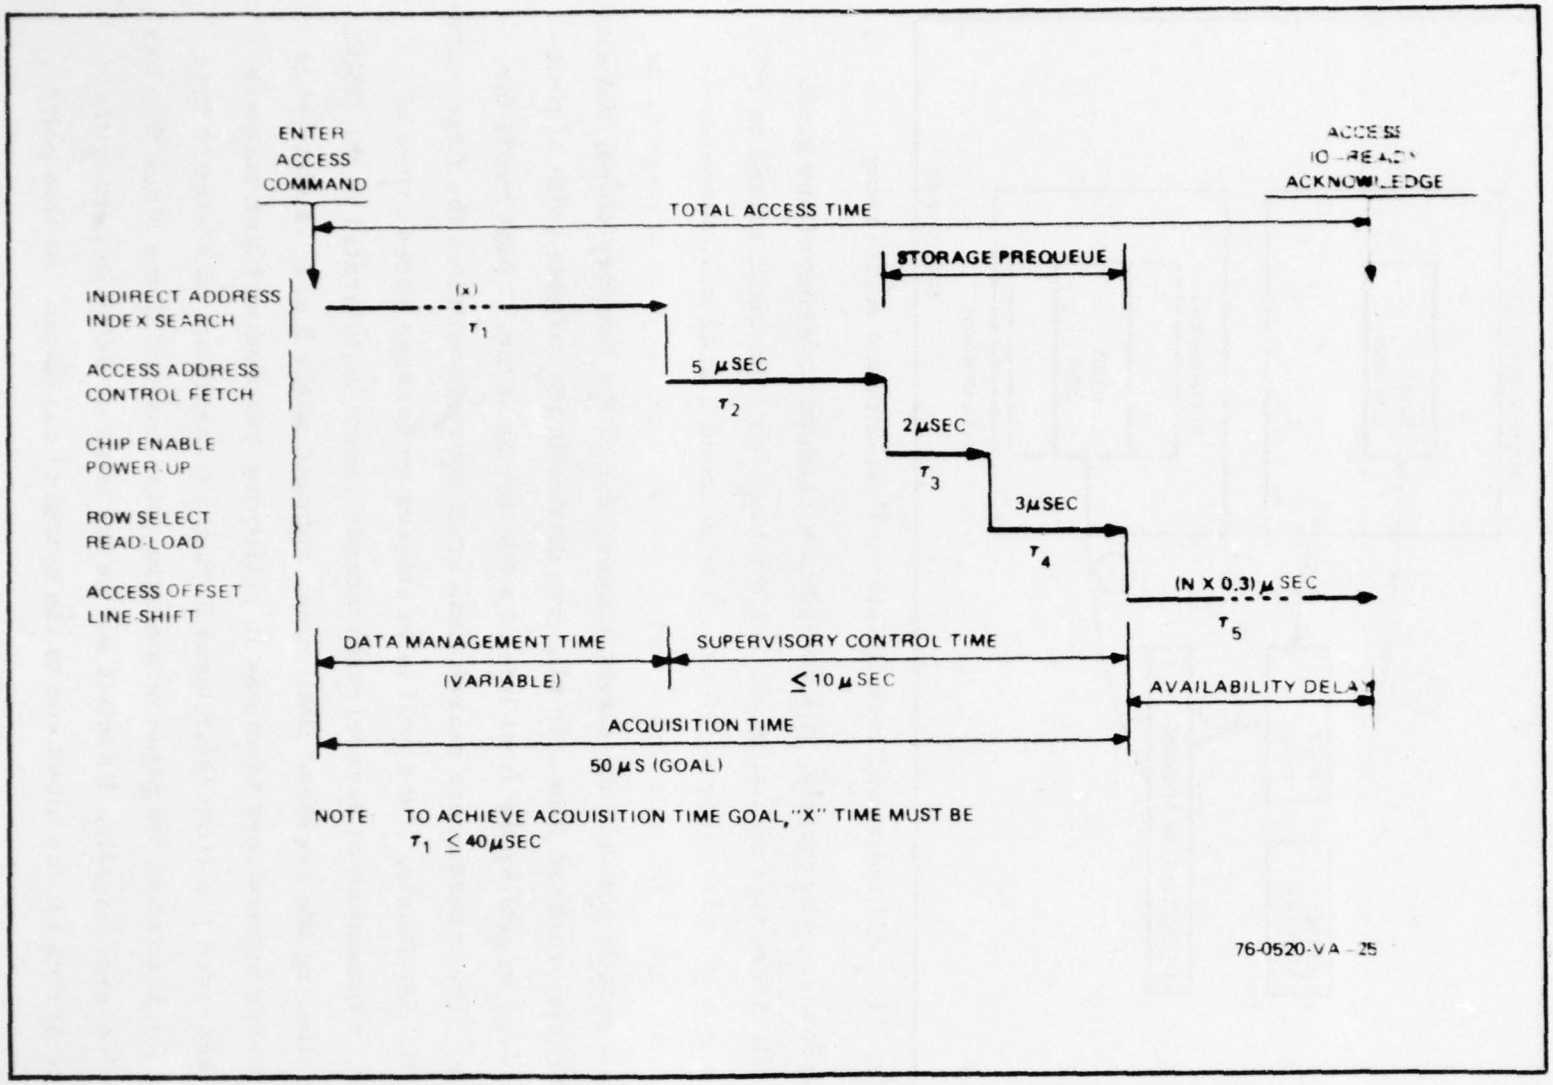

Figure 32. Access Time Considerations

another variable time period occurs: namely, that associated with indexing to a starting line position other than the first in a storage block. Each line of offset adds an increment of 300 nsec in the form of an availability delay. Neglecting this period, the "average" access time should be less than 25 + 3 + 2 + 5 = 35 μsec.

The allocation and control algorithms in the data management facilities must specifically address the problem of preventing unavailable space from accumulating. However, when the available space remaining becomes small, it may not be properly positioned in the various word channels to permit a specific data file to be entered. In those instances, human intervention or the application of external software programs via either the system test stand console (if this facility is included in the installation) or an optional tie-in to the DMU from the STARAN UNIBUS. This will facilitate user tailorable "shuffle" and "compress" operations to be performed on the storage space allocations to reorder the available space so as to fit in the desired data file.

3.3.1 Access Functions

There is virtually an infinite variety of ways in which instructions could be sent to the MM. As observed before, it is therefore not the intent of the present discussion to describe a set of final system specifications. Rather, it is to set forth a consistent approach to the handling of access commands via the lines present at the STARAN DMA port. The points of concern here are the schemes whereby command, control, and selection information is transferred between STARAN (or other user device) and the mass memory management sections. To maximize speed of operation, the condition is imposed that all transfers over the DMA channel should (if possible) be completed in one machine cycle. Although in the final analysis this stipulation may not be mandatory, it does for the present supply a working hypothesis on which organizational developments can be based.

Four distinct bits of information must be supplied to the memory in each access request: device select code, op code, address mode, and address. Since there is no separate "device select" bus provided from STARAN, the device selection function is accomplished by combining the device selection task with the op code and conveying this information over the least significant hexadecimal address of the DMA address channel (which has been allocated for the mass memory). Use of the DMA address channel in this manner is what makes the MM control memory effectively appear to STARAN as an extension of AP control memory. Utilization of the address lines in this manner to transmit the requisite storage mode op code leaves the full 32 DMA data lines for conveying the specific file location addressing information.

The single hex address (four binary digits) allows up to 16 storage mode op codes to be specified. At present, consideration is made for use of 14 of these: 12 denoting direct memory reference instructions and two calling for nonmemory reference instructions. One of these serves as a common call for a series of nonmemory reference operations that are brought into play by extending the effective number of op codes through the use of the lowest ordered two hex digits (eight least significant binary digits) of the DMA data bus. Table 7 lists the storage mode access instructions that are presently foreseen. The adjunct extension op codes for nonmemory reference operations are itemized in table 8.

A point of note here is that, as stated previously, storage mode access requests are signified by the state of the write/$\overline{\text{read}}$ line out of the DMA port: a low denotes Read, which is interpreted as a "memory reference" storage mode operation; and a high denotes Write, which is interpreted as "nonmemory reference" status readout operation. Accordingly, an equivalent set of 16 control memory readout instructions are available. Consideration of the possible form of this set is not made at this time since it necessarily depends on the detailed design structure of the MM controller. The instructions will, however, include those necessary to monitor all normal system activity,

TABLE 7

DMA ADDRESS CHANNEL DATA STORAGE MODE OP-CODE TABLE

Hex Code	No	Function Name	Execution/Termination Characteristics
0 0 0 F	(15)	CLEAR MEMORY	Executed Automatically Without PIO Interaction
0 0 0 E	(14)	CLEAR ALL FILES	Executed Automatically Without PIO Interaction
0 0 0 D	(13)	CLEAR FILE	Executed Automatically Without PIO Interaction
0 0 0 C	(12)	DUMP FILE	Taken to Completion on In-Process File
0 0 0 B	(11)	DUMP ALL FILES	Taken to Completion on In-Process File
0 0 0 A	(10)	DUMP MEMORY	Taken to Completion on In-Process File
0 0 0 9	(9)	READ MEMORY	Execution Stops at Point of Termination
0 0 0 8	(8)	READ ALL FILES	Execution Stops at Point of Termination
0 0 0 7	(7)	READ FILE	Execution Stops at Point of Termination
0 0 0 6	(6)	SWAP FILE	Execution Stops but Housekeeping Completed
0 0 0 5	(5)	MODIFY FILE	Execution Stops but Housekeeping Completed
0 0 0 4	(4)	WRITE FILE	Execution Stops but Housekeeping Completed
0 0 0 3	(3)	–(Spare)–	–(No Op)–
0 0 0 2	(2)	–(Spare)–	–(No Op)–
0 0 0 1	(1)	ACCESS TERMINATE	Institutes Instruction Close-Out & Relinquishes PIO Control
0 0 0 0	(0)	STATUS MODIFY	–(See Extended Op-Code Table)–

76-0520-TA-12

TABLE 8

DMA DATA CHANNEL STATUS MODIFY EXTENDED STORAGE MODE OP CODE TABLE

Binary Code	No	Function Name	Execution Effect
0 ■ ■ ■ ■ 0	(0)	OFF–LINE STATUS ABORT	Resets all Access & Control Functions to Null State
■ ■ ■ ■ 0 1	(1)	FILE INCREMENT REPEAT	Advances User File Number by One & Re-Executes Latest Instruction
■ ■ ■ 0 1 X	(3)	BLOCK INCREMENT CONTINUE	Skips Ahead One Block in Current File Without PIO Interaction
■ ■ 0 1 X X	(7)	INSTRUCTION INITIALIZE REPEAT	Re-Executes Latest Access Instruction in its Entirety
■ 0 1 X X X	(15)	–(Spare)–	–(None)–
0 1 X X X X	(31)	OFF–LINE INSTRUCTION SUSPEND	Relinquishes PIO Control but Maintains Access Status
1 0 ■ ■ ■ ■	(32)	ON–LINE INSTRUCTION RESUME	Resumes Control of PIO With Prior Status Re-Established
X 1 0 ■ ■ ■	(48)	–(Spare)–	–(None)–
X X 1 0 ■ ■	(56)	INSTRUCTION DECREMENT REPEAT	Re-Executes Instruction Before Last in its Entirety
X X X 1 0 ■	(60)	BLOCK DECREMENT CONTINUE	Steps Back One Block in Current File Without PIO Interaction
X X X X 1 0	(62)	FILE DECREMENT REPEAT	Retards User File Number by One & Re-Executes Latest Instruction
1 X X X X 1	(63)	ON–LINE STATUS TERMINATE	Terminates Latest Access Without Relinquishing PIO Control

76-0520-TA-13

readout storage parity error counters, and monitor progress of BIT (built-in-test) functional sequences. Also included must be commands to search and output all of the various data management memory entries (such as available storage spaces, complete listing of all files held, and listing of all files of a specific user). These instructions can be executed concurrently with PIO sequences with no interference.

Figures 33 and 34 illustrate the suggested DMA data channel encoding format for specifying the file addressing information related to the various storage mode op codes. As seen here, the 32 bit (eight hex digit) storage access address is broken into two 16 bit fields. The upper 16 bits specify the indirect address of the file or block of files being referenced. The lower 16 bits specify file handling mode directions and/or the postindexing offset to be applied to get the effective memory starting-line address.

3.3.2 Instruction Repertoire

In most cases, the stipulated op codes and related instructions are straightforward and largely self-explanatory. Points of note with regard to terminology are the following:

- Clear means total purging of a data unit from both the data storage and management directory memories without outputting the stored data.
- Dump means purging of a data unit from both the data storage and management directory memories after outputting the stored data.
- File operations involve a single file of a specific user.
- Swap specifies a combined read-write operation in which prior file data is readout, cleared, and replaced by new input data on a line-by-line basis.
- Modify specifies a type of overwrite operation where prior file data is cleared on a line-by-line basis and replaced by new input data.
- All File operations involve all files of only the specific user.
- Memory operations involve all files of all users.
- Appendix flagging of a write request signals that the user wants to have all remaining space in memory committed to his next entry.

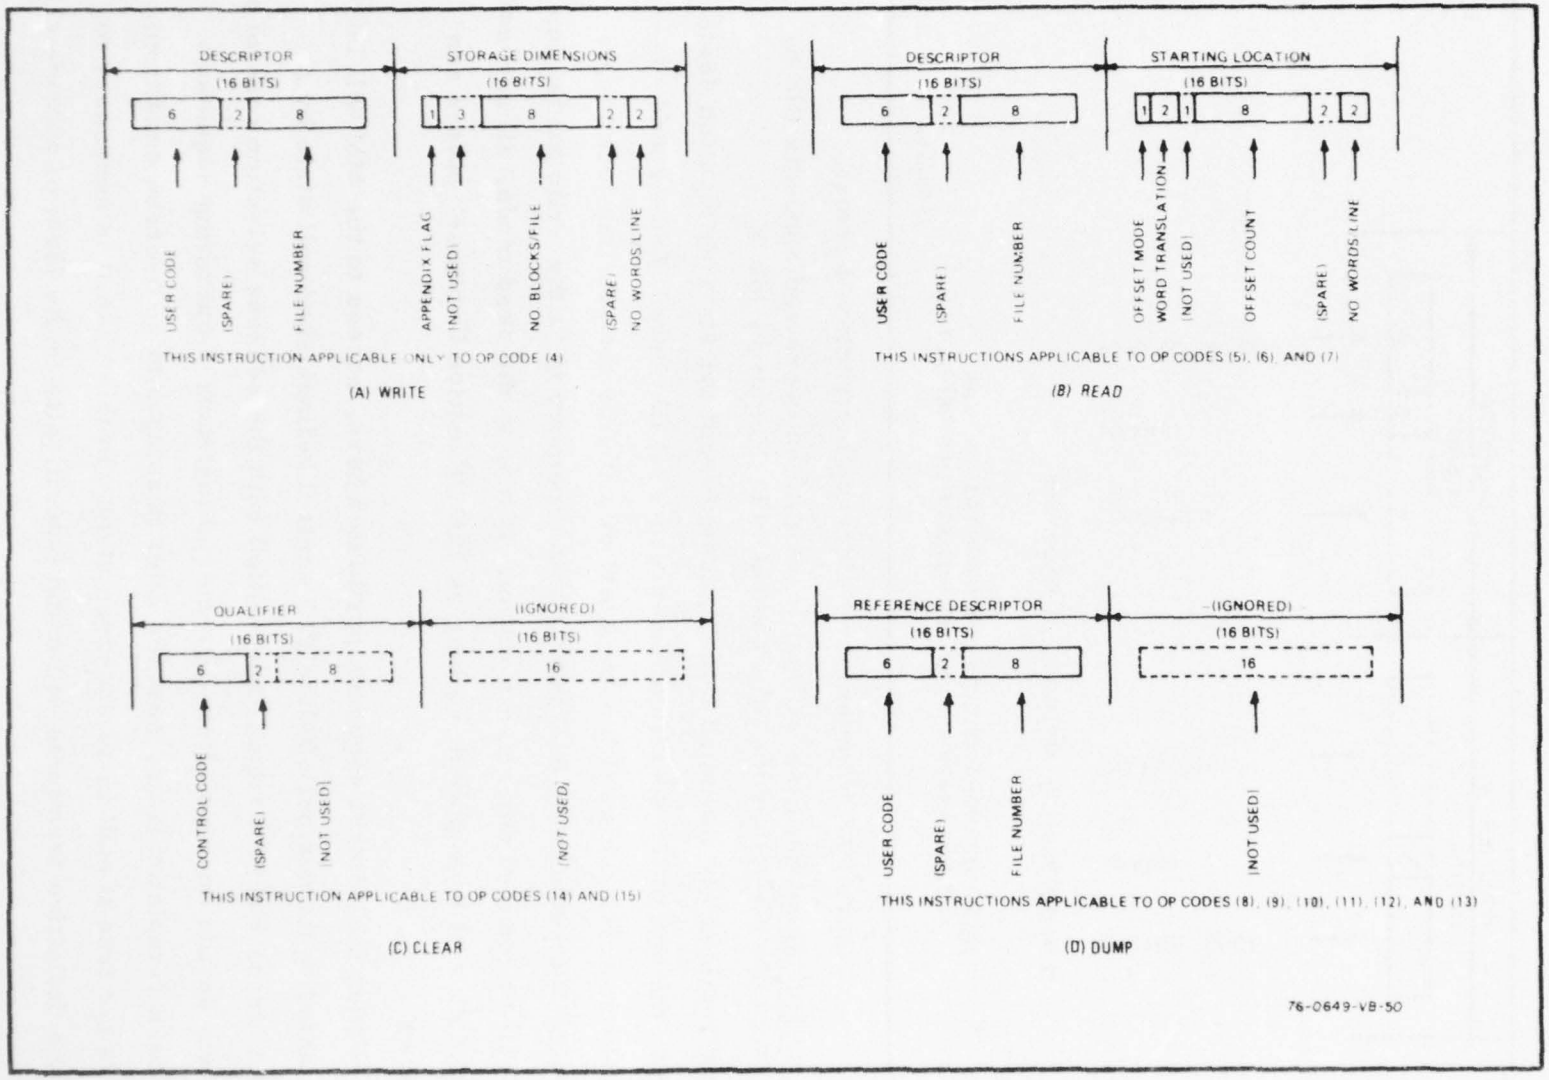

Figure 33. DMA Data Channel Instructions.

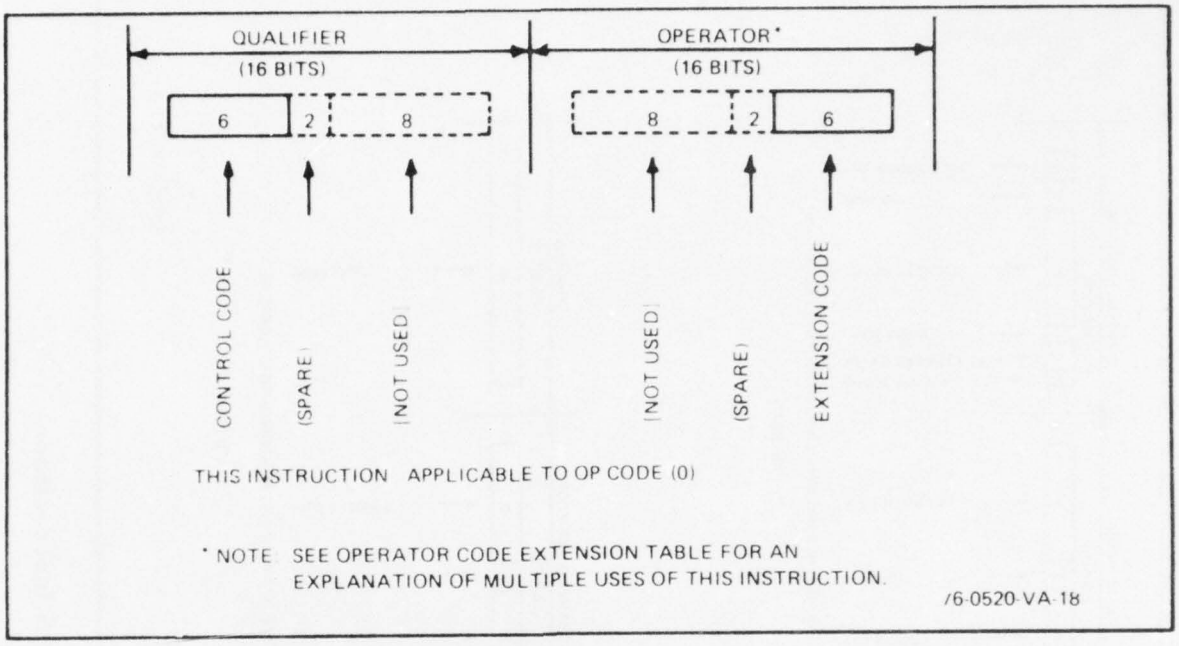

Figure 34. Nonmemory Reference Instruction Format

- Offset Mode facilitates offsetting within an oversized appendix file by signalling whether the offset count is by line or by block.

All instructions are assumed to be executed from the starting location designated to the end of the specified (or in-process) file only. End-around execution-continues, back at either the start of a file or the starting location, do not occur for sequencing overruns. If the memory is in the write mode, input data will be ignored during an overrun. If it is in the read mode, all zeroes output data will be supplied; so long as PIO "Function Continue" pulses are received.

Under the control sequence postulated here, access to the MM will be terminated to its normal off-line PIO control relinquished null state by a DMA Address Present signal transmitted with the address selection code of a different device anytime after all required instruction operating sequences have been completed (i.e., anytime after passing into an overrun condition). This is not true if still in an underrun state. In this case it is necessary to supply a definitive terminate command (specifically in the form of op code (1)).

AD-A033 715 WESTINGHOUSE DEFENSE AND ELECTRONIC SYSTEMS CENTER B--ETC F/G 9/2
MASS MEMORY ORGANIZATION STUDY.(U)
SEP 76 D L FARNSWORTH, C P HOFFMAN, J J SHUTT F30602-75-C-0317
UNCLASSIFIED NYD-56237-A002 RADC-TR-76-254 NL

2 OF 3
AD
A033715

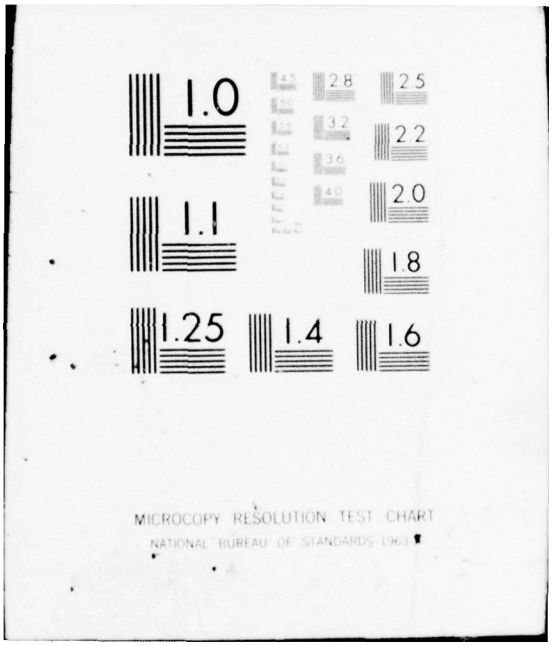
1.0
1.1
1.25
1.4
1.6
2.8
3.2
3.6
4.0
2.5
2.2
2.0
1.8
MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963

The only difference between the Swap and Modify instructions is that for a file modification sequence the PIO output line (inputs to STARAN) are not active. Instead, input data is merely used to replace former file data. Observe that this arrangement also facilitates erasing (or clearing) selected words, lines, or blocks in a file by supplying all zeroes as the replacement data.

Control of PIO functions is relinquished upon receiving a "Suspend" command, but all internal instructions and in-process management functions are maintained in the current status. Subsequent receipts of the "Resume" command causes control of PIO functions to be re-assumed with all conditions re-established as they were before being interrupted. If, however, any command other than Resume is received while in the Suspend state, off-line holding operations related to the Suspend function will all be terminated. On-Line Terminate accomplishes basically the same thing as Access Terminate, except that PIO control is not relinquished.

Although not included in the proposed op code listing, an additional operation which might be inserted in one of the spare code slots is a read prior /write new command. This operating mode conceivably could be used to duplicate an important data file with or without modification. The duplicate file would, of course, have to be assigned an unused name (i.e., code number) to avoid rejection of the command by the data management unit (which checks to ensure that a new file is never assigned the same "name" as a currently held file).

Note that the Clear Memory comn and will be overriden and interpreted instead as a Clear All Files command if the user code applied in the instruction field is not authorized to clear the memory. Similarly, the Instruction Decrement Command will be overriden and interpreted as an Instruction Initialize command if the prior instruction was not executed on a file assigned to the present user. This provides a limited degree of protection against unauthorized (or at least unintentional) file openings. The same observation applies to the File Increment and Decrement commands. If the

file selected at the time these commands are entered is either #00 (for Decrement) or #FF (for Increment), the command will simply be interpreted again as the Instruction Initialize command.

Checks performed by the data management facilities to validate authorization of a user to issue a Clear Memory command may be provided for either in firmware or software. If only a predetermined group of users will ever be so authorized, it could prove advantageous to make the check a hardware function. On the other hand, if the authorized user list is likely to be changed at frequent intervals, it should be made a software alterable function. This will ultimately have to be decided by mutual agreement between RADC and the contractor following further analysis.

4. RELIABILITY EVALUATION

A great many factors bear on the aggregate reliability of a fully deployed MNOS Mass Memory system. Table 9 provides a summary listing of the main items which ultimately must be addressed. Error rates (item VI) in an MNOS memory system, as with any high density storage technology, are affected primarily by the S/N ratio of the cell level detection mechanism.

Quantitative measurements of error rate levels generally presents a problem in that, without extensive field experience and feedback, this type of error is difficult to separate from and is frequently masked by equipment failure. With only one MNOS memory system presently in field service - namely, the BORAM module delivered under joint Army/Navy sponsorship - extensive error rate test data is not available.

Applicability comparisons of error rate experience with qualitative projections derived from intrinsic device parameters have, however, been made by J. E. Brewer of Westinghouse. The results of his findings are summarized in the graph of figure 35. This information reflects extensive survey data and, to the highest degree possible, depicts the present state of the art in the various digital device areas cited. Here it is seen that the anticipated MNOS memory error rates for a 16 Mbit capacity system lie in the range of 10^{-13} to 10^{-14} errors-per-bit processed, or roughly the same as the best available experience data indicates is secured wholly within a computer CPU.

The distinction between "soft" and "hard" error rates in the case of rotating memories pertains to the difference between recoverable (possible through rereading) and unrecoverable errors. Note that a "hard" (or unrecoverable) error may or may not be detectable and may, although not necessarily, signal a permanent storage medium failure. A "fuzzy set" type boundary thus exists between failures and errors at this point. For our purposes, all such conditions will be considered failures.

TABLE 9

MM RELIABILITY/MAINTAINABILITY ASSESSMENT FACTORS

I. Device Level Considerations
   - A. Failure Rate Schedules
   - B. Operating Environment
   - C. Qualification Screening

II. Module Level Considerations
   - A. Simplex vs Redundant Structures
   - B. Error Correction Techniques
   - C. Dormancy Characteristics

III. System Level Considerations
   - A. Modular Redundancy Effects
   - B. Support Circuitry Overhead
   - C. Interface & Power Requirements

IV. Reconfigurabiliy & Fail-Soft Provisions
   - A. Switchable Spares Techniques
   - B. Shrinking Memory Techniques
   - C. Implementation Hardware

V. Diagnostic & Testing Facilities
   - A. Hardware vs Software
   - B. Built-In vs Externally Exercised
   - C. Types of Routines Necessary

VI. Error Rate Projections
   - A. Predictability
   - B. Detectability
   - C. Recoverability

VII. Availability Time Forecasts
   - A. Mission Specifications
   - B. Operational Requirements
   - C. Location Accessibility

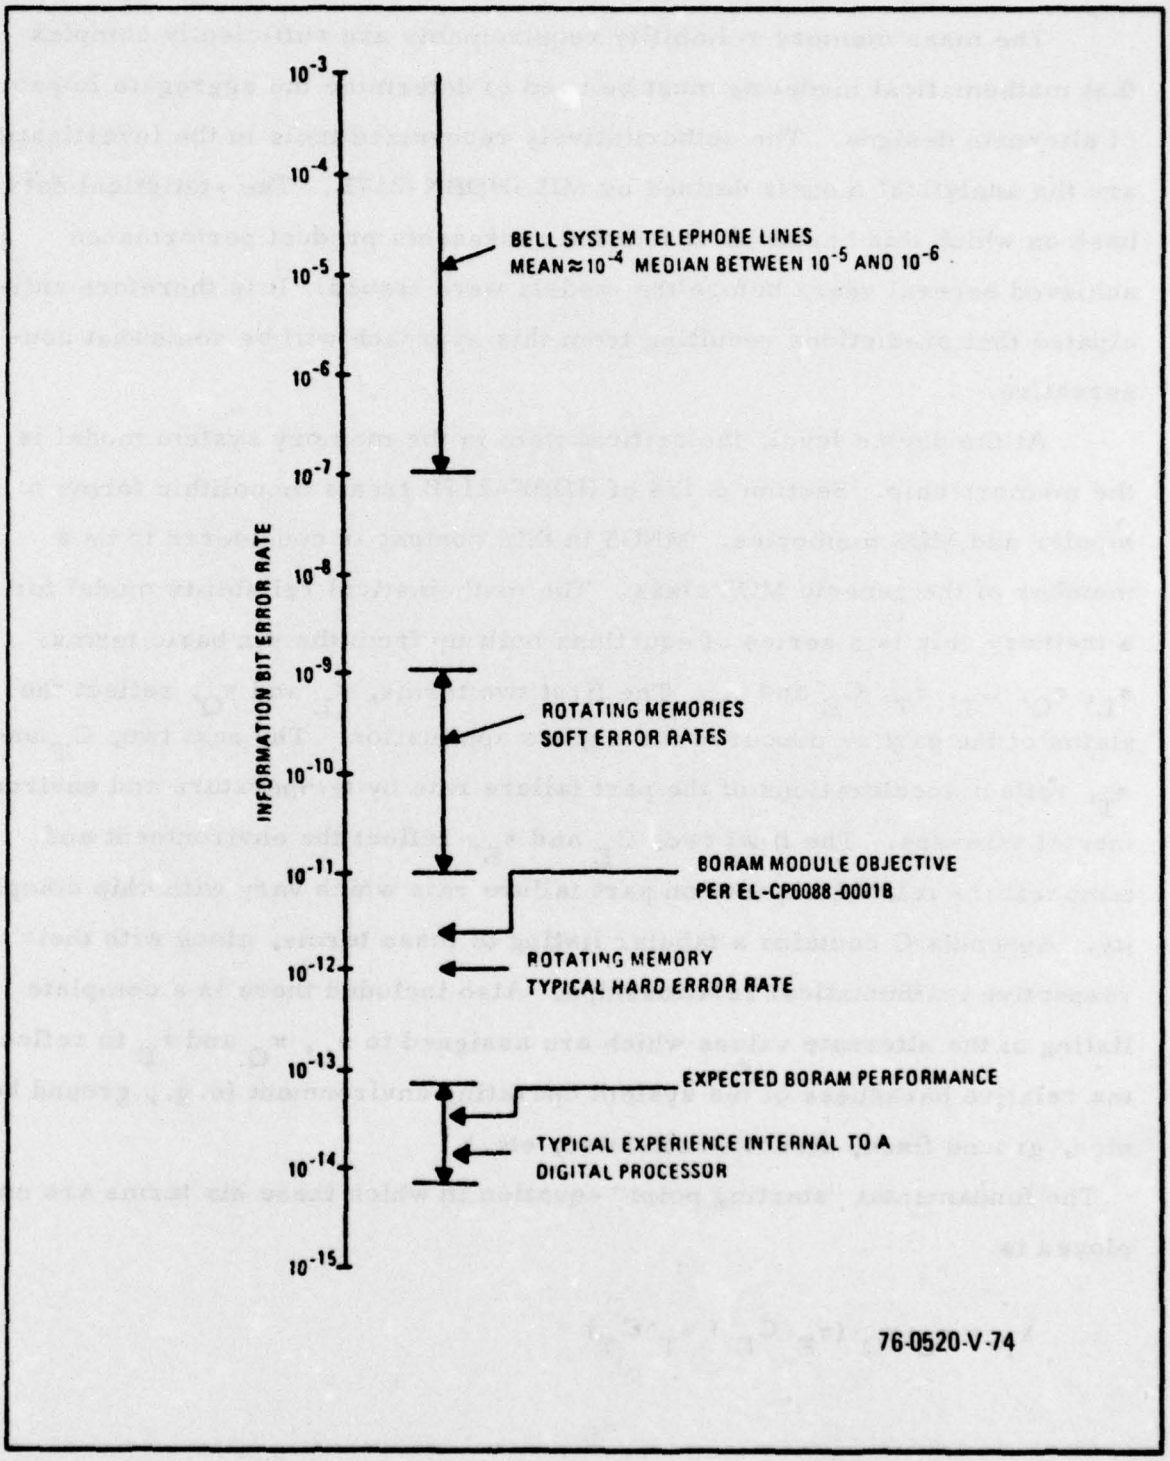

Figure 35. Comparison of Typical Error Rate Experience

4.1 MODELING APPROACH

The mass memory reliability requirements are sufficiently complex that mathematical modeling must be used to determine the aggregate impact of alternate designs. The authoritatively recognized tools in the investigation are the analytical models defined by MIL-HDBK-217B. The statistical data base on which this handbook is founded represents product performance achieved several years before the models were issued. It is therefore anticipated that predictions resulting from this approach will be somewhat conservative.

At the device level, the critical item in the memory system model is the memory chip. Section 2.1.4 of HDBK-217B treats monolithic forms of bipolar and MOS memories. MNOS in this context is considered to be a member of the generic MOS class. The mathematical reliability model for a memory chip is a series of equations built up from the six basic terms: π_L, π_Q, C_T, π_T, C_E and π_E. The first two terms, π_L and π_Q, reflect the status of the part as procured for a given application. The next two, C_T and π_T, reflect accelerations of the part failure rate by temperature and environmental stresses. The final two, C_E and π_E, reflect the environment and temperature related impacts on part failure rate which vary with chip complexity. Appendix C contains a tabular listing of these terms, along with their respective mathematical relationships. Also included there is a complete listing of the alternate values which are assigned to π_L, π_Q and π_E to reflect the relative harshness of the system operating environment (e.g., ground benign, ground fixed, airborne inhabited, etc.).

The fundamental "starting point" equation in which these six terms are employed is

$$\lambda_P = \pi_L \cdot \pi_Q \cdot (\pi_E \cdot C_E + \pi_T \cdot C_T)$$

which denotes the effective "failure rate" that can be expected from the given part under the constraints of a particular set of production and application conditions. Although calculation of chip failure rates using this relationship is a relatively straightforward "cookbook" procedure, preparation of reference failure rate tables for alternate chip sizes when operating under the range of conditions covered by HDBK-217B can be quite tedious.

To relieve some of the tedium and facilitate the generation of a sample set of tables, an equation solution program (see Appendix C) was written and used on an HP-25 calculator. No claim is made with regard to the efficiency or optimization of this program. All that can be said for it is that it was set up quickly, and it works. Using it, a group of tables (included for reference in Appendix C) was generated to reveal the effects of temperature on failure rate for memory chips in the 1 Kbit to 512 Kbit size range under a representative sample set of application conditions. The conditions specifically addressed include the following:

- Ground Benign Environment
  a. mature device in production >6 months with procurement to MIL-M-38510, class B.
  b. new device in production <6 months with procurement to MIL-STD-883, class B.
- Ground Fixed Environment
  (same as a)
- Airborne Inhabited Environment
  (same as a)

Availability of this tabular data will expedite the algebraic processes involved in the mass memory reliability analyses.

A point of note here is that failure rates derived from the calculator program as written are expressed in terms of failures per 10^9 hours rather than failures per 10^6 hours as in conventional in MIL-HDBK-217B. This fact is reflected in the values listed in the referenced tables, but conversion of specific failure rates to be in terms of 10^6 hours merely involves dividing the given values by 10^3.

The MOS and bipolar models of HDBK-217B differ in the form of the π_T factor. Detailed attention is here directed exclusively toward MOS classified memory devices. Part failure rates for all system logic devices are approximated on the basis of generic type, qualification level, environment, and functional complexity according to the estimation tables in Section 3 of HDBK-217B. It is assumed that standard or low power Schottky TTL devices are used exclusively within the memory module, except in the various memory areas which will be noted. With the circuit complexity estimated to be between 20 and 50 gates, on the average, with all devices in mature production and screened to MIL-M-38510, level B, the TTL device failures will be on the order of 0.02 failures/10^6 hours in a ground benign environment and 0.06 failures/10^6 hours if it is ground fixed. To match STARAN interface busing requirements, logic devices in the MM interface unit are assumed to be simple ECL types with a circuit complexity of less than 20 gates per IC. Their failure rates are accordingly taken as approximately 0.01 and 0.05 failures/10^6 hrs, respectively, in ground benign and fixed environments.

To use the mathematical model, it is necessary to make certain assumptions regarding the thermal environment. Since the intended usage for the MM is as a backing store for STARAN, that type of installation is assumed. The facilities in which the STARAN based AAP computer complex is housed at RADC fit into the classification of an engineering laboratory or ground benign environment. The external conditions to which the system is subjected are highly controlled with temperatures typically 20 ±5°C and relative humidity maintained at about 50 ± 10%. Furthermore, the equipment is serviced only by skilled technicians.

The class of construction considered for the mass memory is based on the use of standard 19-inch card rack cabinets, with forced-air ventilation used throughout. All devices are mounted in ceramic or kovar DIL (dual-in-line) packages of 24 pins or less. Groups of device DIL's are plugged

into PC boards, which in turn become part of a roll-out drawer assembly. An example of a drawer of this type which is commercially available from Augat is presented in figure 36. Three purposes are served by adopting this packaging technique. First, it minimizes initial memory construction costs. Second, it enhances the maintainability of the system (all parts and circuit boards are readily accessible). Third, it facilitates good thermal management of the internal cabinet environment. With proper ducting, the worst case hot spot ambient to which any device internal to a cabinet is subjected should be less than 15°C above the external temperature.

At the device level, use of ceramic or kovar DIL's means that the thermal resistance from device junctions to the surrounding environment (not just to the case) is on the order of 0.05°C/mW for 22 or 24 pin packages and about 0.06°C/mW for 14, 16, or 18 pin packs. In the nonvolatile main storage section of the MM system, the power dissipation of the MNOS chips under consideration is typically less than 200 mW, so a rise of less than 200 x 0.05 = 10°C is anticipated from junction to cabinet ambient. Supporting

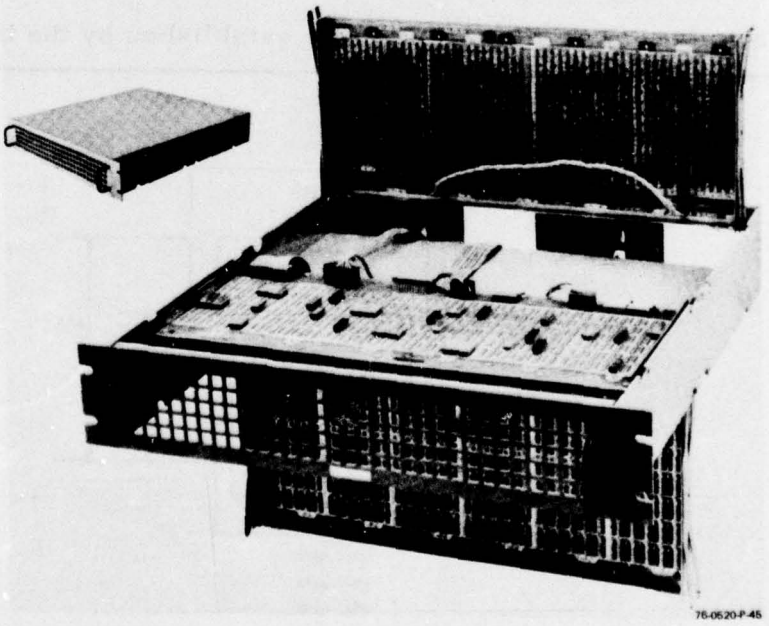

Figure 36. Augat Drawer Assembly

TTL logic devices can be expected to have about 300 mW or less dissipation on the average, so their junction-to-ambient rise will be about 300 x 0.06 = 18°C. Figure 37 describes the interrelated implications of the thermal management constraints which should be readily achievable with the type of packaging suggested.

For purposes of the reliability projections in this report, the external environment will be taken as being at its upper limit of 25°C. Furthermore, all IC's (both logic and memory types) will be assumed to operate at a junction temperature of 50°C when active. This assumption is admittedly conservative, but it provides a working baseline from which to make comparisons. Moreover, it is consistent with both the presently incomplete knowledge of the final system physical configuration and prescribed practices set forth in MIL-HDBK-217B for such situations. Clearly, control of the thermal environment is an important parameter. It forms a design option which can be massaged to improve overall system reliability.

An implicit assumption of *MIL-HDBK-217B is that* component failure rates are constant (i.e., that failures occur according to an exponential time distribution law). The failure rate established by the handbook model

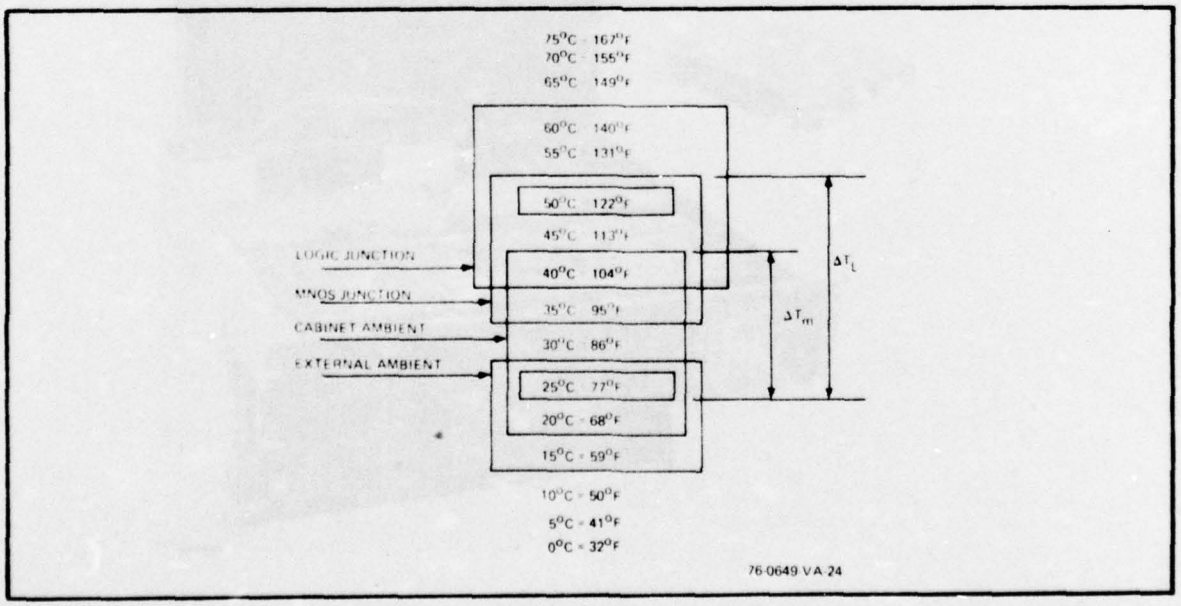

Figure 37. Thermal Management Constraints

represents what is known as the "active" failure rate. Active implies that the device has power applied. Contrastingly, when power is removed from the device it is considered to be "dormant" and exhibits a much lower failure rate. In particular, recent findings[1] indicate that real life dormant state failure rates, while not zero, are typically less than 1/5th of the active rate (for MIL-M-38510, level B qualified devices) with the junction having no delta relative to the prevailing environmental temperature during powered down periods. The less well qualified the device is, the greater the difference will be. Since the active junction temperature is being taken as 25°C higher than when it is powered down, the effective dormant failure rate is found to be typically less than 1/20th of the active rate.

Dormancy failure rates are applicable in two ways to the analysis of memory devices. First, for application in the RADC STARAN facilities, it can be anticipated that, out of any single year of operation, the system will actually be actively in use no more than 40 percent of the time out of any given year of operation. Note that this corresponds to the system being used an average of better than 9.6 hr/day, 7 days a week. Thus, for at least 60 percent of the mission the entire memory could be powered down and enjoy a dormant state failure rate - regardless of whether volatile logic or nonvolatile memory devices are being considered. Consultation with RADC reliability personnel has indicated, though, that to maintain a consistently conservative baseline for reliability comparisons, the system as a whole should be considered to have power applied 100 percent of the time; with the activity factor applied only to those devices which are intrinsically nonvolatile and their immediately surrounding interface circuits (i.e., only to those devices wholly within the storage section). Second, for the nonvolatile memory chips it is only necessary to apply power to those chips which, during burst periods, are actively engaged in I/O operations. Hence, most of the

---

[1] E.S. Bauer, et. al., (Martin-Marietta Aerospace); "Dormancy and Power-Off Cycling Effects on Electronic Equipment"; AD768-619; August 1973.

storage section remains powered down even during the 40 percent of the mission that the system is operating. The dormant failure rate, of course, continues to apply to all devices still having electrical stresses removed.

The fractionalized reference equation which illustrates the mathematical means by which dormant periods are introduced into the system life-cycle is given in MIL-HDBK-217B (Appendix A, section 30, item C, paragraph 2). For ease of algebraic manipulation, it is convenient to combine terms in this relationship. This process yields the basic reliability equation which has been applied in the analyses conducted during this study. The form of this equation is

$$R = \epsilon^{-N\lambda t/D} \qquad \text{(Eq. 4-1)}$$

where N = total device count

λ = active failure rate

t = total mission time

and "D" is defined as a dormancy improvement factor given by

$$D = \frac{K_\lambda}{1 + \frac{K_\lambda - 1}{K_N K_t}} \qquad \text{(Eq. 4-2)}$$

The terms employed in this relationship are in turn defined as follows:

$K_\lambda = \lambda_{active}/\lambda_{dormant}$

$K_N = N_{total}/N_{active}$ (during t_{active})

$K_t = t_{mission}/t_{active}$ (reciprocal of duty factor)

A summary listing of the terms in both equations is presented in table 10. Points to note with regard to the dormancy factor (D) are:

- all "K_j" terms must satisfy the inequality $1 \leq K_j \leq \infty$
- as the $K_N \cdot K_t$ product becomes $>> K_\lambda$, the value of D approaches its upper limit of $D \leq K_\lambda$.

Confirmation of the derivation accuracy of the above relationships can be secured by comparison of the starting point equation cited in MIL-HDBK-217B.

TABLE 10

RELIABILITY MODELING TERMS

Term	Definition	Relationship
t	total mission time	1 year = 8766 hours
λ	chip failure rate	*
N	total chip count	*
K_N	total-to-active chip count ratio	$K_N = \frac{N_{total}}{N_{active}}$
K_λ	active-to-dormant failure rate ratio	$K = \frac{\lambda_{active}}{\lambda_{dormant}}$
K_t	total-to-active mission time ratio	$K_t = \frac{t_{mission}}{t_{active}}$

*see text discussions

For reliability analysis purposes, it will be assumed that the Mass Memory is functionally configured in the manner shown in figure 38. As seen here, the MM system architecture described previously has been restructured slightly to delineate those portions of the overall circuitry which form uniquely characterizable subsystem "units" within the memory. Six different types of units are seen to exist, with three of the six duplicated as a result of the dual-port system configuration. Within each functional unit are shown the series of failure unit circuitry blocks which appear in the reliability system models. These blocks do not correspond specifically to electrical functions performed or physical boundaries existing within a subsystem unit, but rather identify those segments of circuitry which fail collectively.

In some subsystems there is only a single failure unit block. In others there are several. All aspects of failures occurring within these individual blocks is considered independent and noninteractive with that of any other block (this necessarily demands independent supplies). A single failure in any of

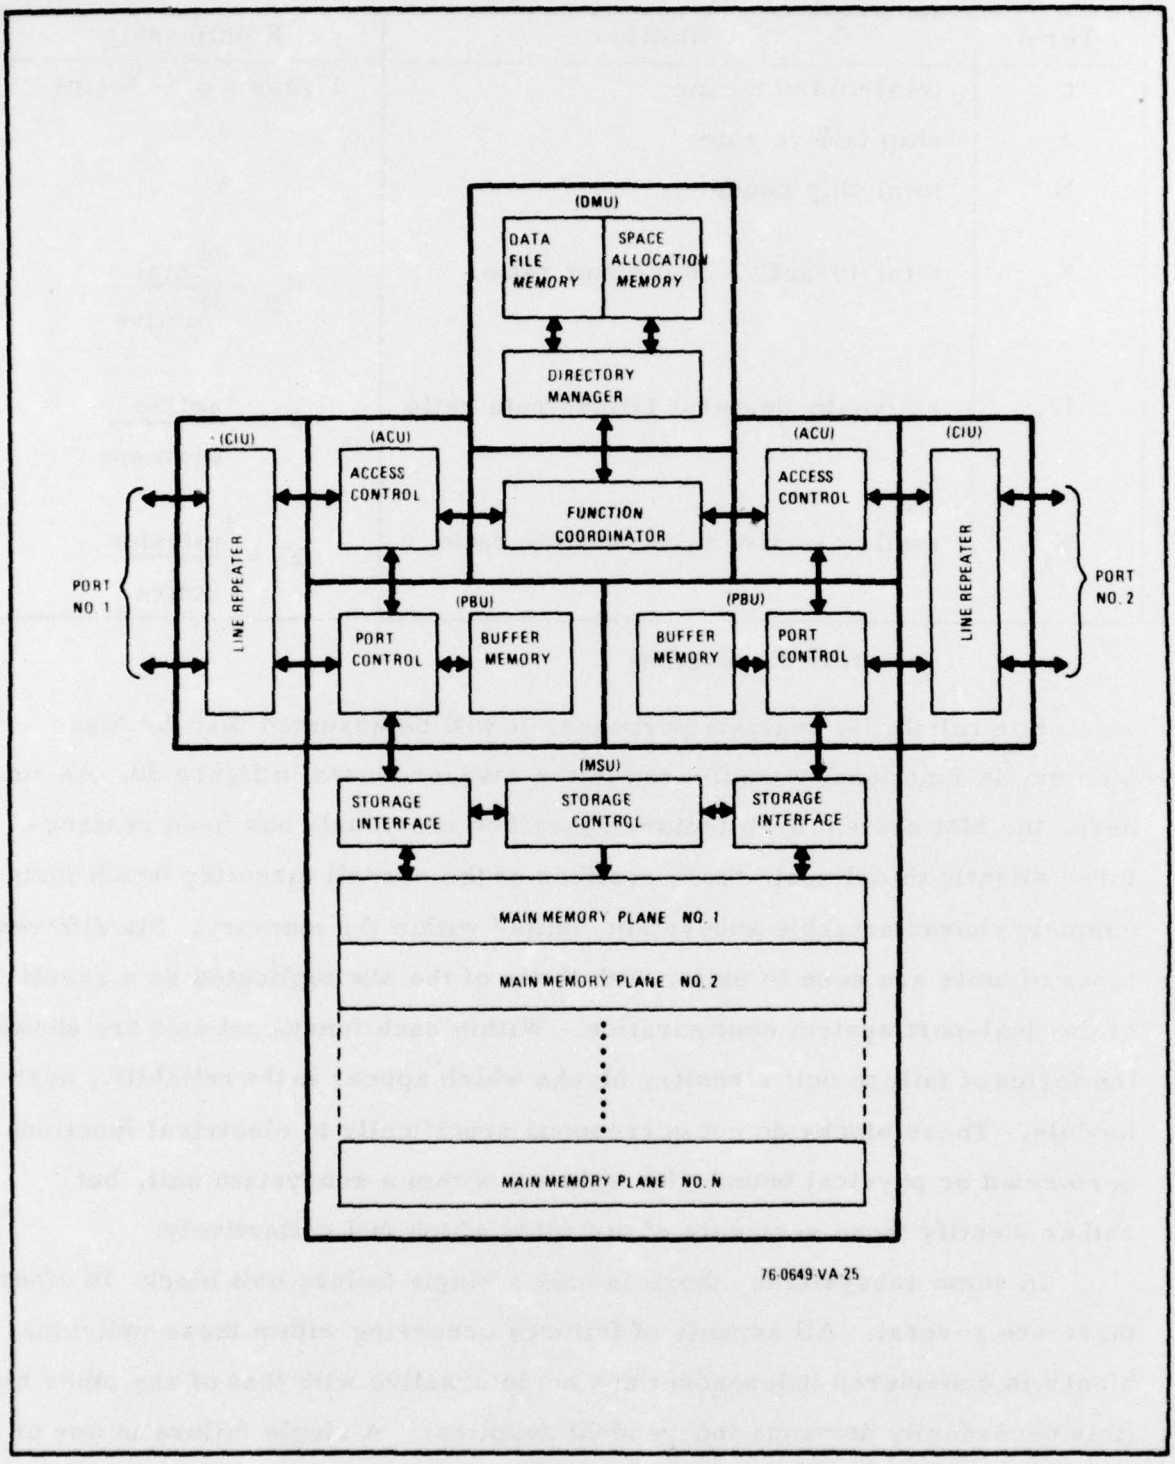

Figure 38. Reliability Analysis "Unitized" MM Block Diagram

the nonmemory circuitry blocks is considered to cause a mortal failure in that section of the memory. Whether or not loss of a specific failure unit amortizes the subsystem in which it is located, and whether this then causes failure of the entire memory, depends on its position in the memory hierarchy. In the PBU, for instance, the buffer memory can fail completely without amortizing either the PBU or the entire MM. The reason for this is simply that the buffer memory in an MNOS mass memory is an add-on which is included to provide extra system capability in accomplishing data reformatting. It can be completely deleted at any time with no loss of ability to enter or retrieve data from the main memory.

Those portions of the Mass Memory in which single, nonmemory device failures will result in loss of a subsystem unit whose failure is catastrophic to the entire memory are the FCU, the directory manager in the DMU, and the storage supervisor in the MSU. The reason for this is that no redundancy exists within these circuitry blocks, and they are not duplicated in the memory organization. It is therefore necessary that these portions of the system be very reliable. Use of simplest possible design principles requiring minimum component counts are indicated, with all devices qualified to MIL-M-38510, level B, as a minimum.

Portions of the system in which single failures are not necessarily catastrophic are the CIU, ACU, PBU, and the MSU storage interfaces. Each of the units existing in these areas is duplicated as a result of the dual port memory configuration. Loss of any one or more of these units which interact with a specific port will neither prevent nor hinder in any way the ability to process data via the opposite port. However, whether or not the loss of an access port is considered to be amortizing to the memory as a whole depends on the failure definition adopted for the system (i.e., the number and type of unit failures that can be tolerated before the memory is considered functionally useless).

Redundancy in the form of SECDED (single error correction/double error detection) type Hamming parity encoding is incorporated into each memory unit in the Mass Memory. This includes the data file and resource allocation memories in the DMU and the buffer memory in the PBU, as well as all of the memory planes in the MSU. Encoding is taken to be on a quarter-word basis, as outlined in table 11. With this proviso, single failures in any of the storage units will be automatically corrected. Depending on the statistical distribution, many failures can in fact occur before the memory fails in its ability to handle user data without introducing errors into it. Beyond transparent error (or failure) correction techniques of this type, no allowance is made for any form of degraded mode of operation in any of the storage sections. This restriction is imposed so that the analyses to be made will reflect hardware conditions exclusively, with no assumptions regarding the ability of software routines to recover from specific failure conditions.

Both volatile and nonvolatile memory device types are employed in the MM system. Volatile MOS memories are specified for use in the buffer memory. Nonvolatile MNOS devices are specified for both of the DMU memories and all planes in the MSU. Strictly on the basis of data integrity in the face of power loss, these storage structures must be nonvolatile in form. Obviously, loss of data from the main memory would constitute a failure. Similarly, loss of file location or storage allocation information from the data management memories would result in inability to retrieve data. It might be possible to incorporate a "restart sequence" into the directory manager's algorithms by which it would search the main storage data contents and retrieve "header" information. This would, however, require that sufficient storage space be set aside in the MSU for these headers. Although such a possibility is not ruled out, it will not be assumed to exist here.

TABLE 11

LIST OF PRINCIPAL MM QUARTER-WORD SECDED PARAMETERS

- Use of standard Hamming codes assumed
- Formation of 256-bit data words from four groups of 64 data bits and 8 parity bits
- Use of 72 storage chips in each encoded group with each chip having 4 I/O's, each of which is used in a different quarter-word segment
- Any single bit channel bad out of 72 is corrected
- Any double bit channel bad out of 72 is detected
- Any single chip failure in a 72 chip group -- causing a single bit failure in all four quarter-word segments -- is transparently corrected
- Any two complete chip failures in a 72 chip group are detected.

In the case of the MM main storage section, the need for nonvolatility is further strongly suggested by practical restrictions on the minimum system reliability level that can be tolerated. This fact is outlined in Appendix D, which reviews alternate memory device technologies. If dormancy capabilities are not included in the main storage section, its mean time between failure (MTBF) becomes vanishingly small. Preliminary analyses of the MSU conducted with dormancy effects ignored -- which is equivalent to using volatile memory devices -- revealed that the combined failure rates of the memory planes under this condition were so high that one would break down within every 20 hours of operation. All consideration of MM system reliability is therefore based on the assumption that the main storage section is implemented with nonvolatile MNOS devices, and that the dormancy related improvements realizable by their presence are indeed secured through the use of power switching during all periods of memory inactivity.

Device sizes employed in the nonvolatile MNOS storage sections are based on the projected establishment of military qualification standards for device sizes of 16 Kbits through 64 Kbits. At present, these standards exist only for memory chips of up to 4 Kbits. It is assumed that the device types deployed in the mass memory will have been in production a sufficient length of time to rate a production maturity factor (π_L) value of 1. Attainment of this condition appears essential but obviously will require considerable attention to formulate the conditions under which these parts must be fabricated, screened, and tested to provide MIL-M-38510 qualification levels.

Selection of the type device used in the buffer memory is free of any consideration of data integrity. It does not have to be nonvolatile, since any in-transit data which it may contain at the time of a power failure will either already be present in the nonvolatile main memory or can be rapidly dumped into it within the several microseconds decay time of the system main supply filter capacitors. As pointed out, the buffer memory is included for special feature purposes, rather than being a critical system item. Its capacity ($\lesssim$ 4 Mbits) and functional requirements indicate that it is most expeditiously and cost effectively implemented using volatile, commercially available MOS RAM's (such as present 4 Kbit units or industry projected 16 Kbit devices). Use of this type of memory IC has accordingly been assumed.

4.2 DETAILED ANALYSES

Two fundamental approaches exist to the calculation of the aggregate MM system reliability. One is a simple serial component count analysis which ignores all considerations of redundancy within the system. This approach yields a result which is normally expressed in terms of either average system failure rate (number of failures per unit time) or its reciprocal, mean time between failures (MTBF). The basic relationship facilitating this type of first order prediction of system failure rate (λ_s) is represented in closed summation form by the expression:

$$\lambda_s = \sum_{i=1}^{n} (N_p \lambda_p \pi_Q)_i \qquad \text{(Eq. 4-3)}$$

where the respective terms are defined as:

λ_s = total system (or subsystem) failure rate (normally expressed in terms of failures (10^6 hours)

λ_p = generic failure rate of ith generic part type in the stipulated operating environment (failures/10^6 hrs)

π_Q = procurement quality factor of ith generic part type (equals unity for MIL-M-38510, level B screening)

N_p = total number of the ith generic part type present in the system (or subsystem)

The generic failure rates and quality factors to be used in this relationship are found in section 3 of MIL-HDBK-217B and are reproduced for reference in Appendix C of this report.

The alternate approach to calculation of aggregate system reliability takes into account the organization of the system. It allows for the grouping of circuitry units and the manner in which they may be interconnected to achieve redundancy related improvements in reliability. This analysis approach yields a result which specifies the probability of the system performance not degrading below a certain level (defined as the amortization point) within a given length of time (mission time, t). The result is thus a probability of success prediction and is normally expressed as a percentage. The fundamental relationship by which system or subsystem reliability (R_s) is determined is a multielement product described in closed form by the expression

$$R_s = \prod_{i=1}^{n} R_i \qquad \text{(Eq. 4-4)}$$

where R_i is the reliability of the ith subunit in the system, or the ith part in a subsystem, as found from equation 4-1.

Once the system reliability, R_s, is determined, the effective mean-time-before-failure of the system ($MTBF_s$) can be derived by noting that

$$R_s = \epsilon^{-\lambda_s t} = \epsilon^{-t/(MTBF)_s} \qquad \text{(Eq. 4-5)}$$

which, when solved in terms of $(MTBF)_s$, yields

$$(MTBF)_s = 1/\lambda_s = t/\ell n(R_s) \qquad \text{(Eq. 4-6)}$$

Here it is to be noted that these last two relationships, in general, facilitate conversions between the two approaches to reliability prediction. Caution is advised in doing so, however, so as to ensure that the results obtained do not overlook redundancy factors, if it is intended that they be included, nor introduce them extraneously.

The approach which is suggested, and has been followed here, for assessing the aggregate reliability of the mass memory is to employ an expeditious mixture of the two analysis approaches. As an absolute baseline reference point (reflecting the most conservative perspective under a given set of conditions), exclusive use is made of the component count approach. This is then supplemented by the more detailed, integrated redundancy form of analysis to determine whether inclusion of selected types of redundancy provide sufficient improvement in critical areas to make the system practical.

To facilitate considerations of the mass memory on both a baseline, nonredundant basis and with all forms of redundancy included, it is first necessary to prepare a set of system and subsystem models. These models, supplementing the general system model of figure 37, have been assimilated and are included for reference in Appendix C. All but three of the MM system units to which these models apply possess self-contained (independent of port duplication) redundancy characteristics. These are the FCU, ACU, and the CIU. Here it is to be noted that a series of initial run-through analyses conducted during the study revealed weaknesses in the original PBU design. In particular, error correction provisions weren't at first included as part

of the buffer memory design. Furthermore, the PBU bussing was setup such that a buffer memory failure caused a failure of the complete PBU. As a result, its predicted reliability was unacceptably low and seriously degraded the aggregate reliability level realized in the PBU. Subsequently, the design of this unit was modified to correct for the noted deficiencies.

On the basis of the above models, the detailed analytical relationships necessary to performing a complete reliability analysis were derived. Tables 12 through 16 present, collectively, a complete summary description of the mathematical processes involved in projecting the reliability levels of the individual modularized units comprising the MM system and of the system as a whole. The relationships given in these tables are completely general in nature, though specifically applicable to the unit organizations from which they were derived, being based exclusively on the established practices set forth in MIL-HDBK-217B. The representative parameters included are, in part, strictly confined to use in considering what happens to the system in either a ground benign or a ground fixed environment. Moreover, the other parameters cited are, at present, engineering estimates of such items as system duty factor and parts counts in the various units.

Since the arithmetic processes are lengthy and tedious rather than difficult, they will not be duplicated at this time. Instead, two itemized summaries have been prepared (one for each environment) and are presented in tables 17 and 18. The points of note from these analyses are that, with the type of redundancies that have been provided in the aggregate system organization, the mass memory frequency of repair record, even allowing for mechanical failures not included here, should typically be on the order of a 1000 hours or more in a ground benign environment and better than 200 hours in a ground fixed environment. If, however, non-independence of the access ports occurs and redundancies are eliminated, the system would fail on the average every 100 or so hours even in a benign environment.

TABLE 12

MM SYSTEM RELIABILITY MODEL RELATIONSHIPS

General: $R_s = R_D \cdot R_F \cdot R_M \cdot R_x$

$R_y = R_A \cdot R_P \cdot R_C$

$R_x = f(R_y)$

a. Redundant Analysis: $R_x = 1 - (1 - R_y)^2$ { Equation allows for presence of two identical access ports.

b. Nonredundant Analysis: $R_x = R_y^2$

TABLE 13

NONREDUNDANT UNIT RELIABILITY MODEL RELATIONSHIPS

General: t = 1 year = 8766 hrs. (Mission time)

(Ground benign) $\lambda_{TTL} = 0.02$ failures/10^6 hrs.

(Ground fixed) $\lambda_{TTL} = 0.06$ failures/10^6 hrs.

(Ground benign) $\lambda_{ECL} = 0.01$ failures/10^6 hrs.

(Ground fixed) $\lambda_{ECL} = 0.05$ failures/10^6 hrs.

Powered duty factor all devices = 100%

a. Function Control Unit: N = 100 TTL devices

b. Access Control Unit: N = 200 TTL devices

c. Custom Interface Unit: N = 1900 ECL devices

TABLE 14

PBU RELIABILITY MODEL RELATIONSHIPS

General: $R_V = R_d \cdot R_b$

$R_b = (R_E)^m$; where $m = \left(\frac{N_b}{n}\right)$

Parameters: $t_{mission}$ = 1 year (8766 hr)

Powered duty factor all devices = 100%

N_L = 1200 TTL devices

N_d = 1400 TTL devices

*N_b = 2304 MOS devices

(Ground Benign) { λ_{logic} = 0.02 failures/10^6 hrs.; λ_{memory} = 0.54 failures/10^6 hrs.

(Ground Fixed) { λ_{logic} = 0.06 failures/10^6 hrs.; λ_{memory} = 0.73 failures/10^6 hrs.

n = 72 MOS devices/memory ELSEG

m = 32 ELSEG's/buffer memory

a. Redundant Analysis:

$$R_E = \sum_{k=(n-1)}^{n} \left(\frac{n!}{k!\,(n-k)!}\right) (R_{chip})^k (1-R_{chip})^{n-k}$$

Equation denotes that (n-1) of n must be good as facilitated by SECDED.

$R_P = R_L$ { Equation denotes that the buffer memory can be deleted without loss of entire PBU.

b. Nonredundant Analysis:

$R_E = (R_{chip})^n$

$R_P = R_L \cdot R_V$

*Note: Memory device count includes SECDED overhead redundancies. Memory device failure rates cited are for 4 Kbit capacity chips.

TABLE 15

DMU RELIABILITY MODEL RELATIONSHIPS

General:

$$R_D = R_\ell \cdot R_{DFI} \cdot R_{RAD}$$

$$R_{DFI} = (R_{EI})^2$$

$$R_{RAD} = (R_{ED})^2$$

Parameters:

$t_{mission}$ = 1 year (8766 hr)

System powered 100% of time and active 40% of time.

Control logic powered 100% of time.

Memories on-line 25% of time.

N_{logic} = 200 TTL devices

*N_{memory} = 304 MNOS devices

(Ground Benign)

- λ_{logic} = 30 fits
- λ_{memory} = 1.26 failures/10^6 hrs (active @ 50°C)
- λ_{memory} = 0.05 failures/10^6 hrs (dormant @ 25°C)

(Ground Fixed)

- λ_{logic} = 0.06 failures/10^6 hrs
- λ_{memory} = 1.72 failures/10^6 hrs (active @ 50°C)
- λ_{memory} = 0.14 failures/10^6 hrs (dormant @ 25°C)

n_{memory} = 72 chips/ELSEG (in DFI memory); 80 chips/ELSEG (in RAD memory)

a. Integrated System:

$$R_{\binom{EI}{ED}} = \sum_{k=(n-1)}^{n} \left(\frac{n!}{k!(n-k)!}\right) (R_{chip})^k (1-R_{chip})^{n-k}$$

Equation denotes that (n-1) of n must be good for successful mission.

b. Nonintegrated System:

$$R_{\binom{EI}{ED}} = (R_{chip})^n$$

*Note: Memory device counts include SECDED overhead redundancies. Memory device failure rates cited are for 16K bit capacity chips.

TABLE 16

MSU RELIABILITY MODEL RELATIONSHIPS

General:

$$R_M = R_K \cdot R_Q \cdot R_a$$

$$R_\beta = R_\delta \cdot R_\gamma$$

$$R_a = (R_\beta)^M$$

(Ground Fixed)

λ_{logic} = 0.06 f/10^6 hrs (active); 0.007 f/10^6 hrs (dormant)

λ_{memory} = 4.05 f/10^6 hrs (active); 0.35 f/10^6 hrs (dormant)

$$R_\gamma = (R_E)^m; \text{ where } m = \left(\frac{N_\gamma}{n}\right)$$

Parameters:

$t_{mission}$ = 1 year (8766 hr)

System powered 100% of time and actively powered 40% of time.

Off-plane devices powered 100% of time.

On-plane devices powered 80% of system active time (2 at a time).

N_K = 200 TTL devices
N_U = 1200 TTL devices (total = 1 x active)

N_δ = 2100 TTL devices
*N_γ = 1152 MNOS devices (total = M x active)

M = # memory planes/MSU
= 16 planes/fully system (using 64 K bit MNOS devices)

n = 72 MNOS devices/storage module ELSEG

m = 16 ELSEG's/module

(Ground Benign)

λ_{logic} = 0.02 failures/10^6 hrs (active); 0.002 failures/10^6 hrs (dormant)

λ_{memory} = 2.92 failures/10^6 hrs (active); 0.12 failures/10^6 hrs (dormant)

a. Integrated System:

$$R_E = \sum_{k=(n-1)}^{n} \left(\frac{n!}{k!(n-k)!}\right) (R_{chip})^k (1-R_{chip})^{n-k}$$

(SECDED Factor)
Equation denotes that (n-1) of n must be good for successful mission.

$$R_Q = 1-(1-R_U)^2$$

(Dual Port Factor)
Equation denotes that 1 out of 2 must be good for successful mission.

b. Nonintegrated System:

$$R_E = (R_{chip})^n$$

$$R_Q = (R_U)^2$$

*Includes overhead SECDED redundancies.

TABLE 17

MM SYSTEM RELIABILITY ANALYSIS SUMMARY

(Ground Benign Environment)

MM Unit	Nonredundant Analysis		Redundant Analysis	
	R(%)	MTBF (hr)	R(%)	MTBF (hr)
FCU	98.26	500,000	(No Redundancy in these units)	
ACU	96.55	250,000		
CIU	84.66	52,630		
DMU	3.388	2,590	95.56	193,100
PBU	1.164×10^{-3}	771.6	81.03	41,670
MSU	1.450×10^{-15}	226.1	3.522	2,620
Total	4.373×10^{-27}	134.2	2.930	2,483

TABLE 18

MM SYSTEM RELIABILITY ANALYSIS SUMMARY

(Ground Fixed Environment)

MM Unit	Nonredundant Analysis		Redundant Analysis	
	R(%)	MTBF (hr)	R(%)	MTBF (hr)
FCU	94.88	166,700	(No Redundancy in these units)	
ACU	90.02	83,330		
CIU	43.48	10,530		
DMU	0.9174	1,869	86.52	60,536
PBU	1.007×10^{-5}	544.1	53.20	13,890
MSU	3.177×10^{-35}	104.3	1.695×10^{-4}	659.7
Total	4.295×10^{-54}	71.34	5.191×10^{-5}	605.8

5. EXPLORATORY DEVELOPMENTS

5.1 DEVICE TEST STAND

The portable device screening and evaluation test stand shown in figure 39 was constructed during this study and is presently in use at the Westinghouse ATL facilities. It has been used to performance screen DIL packaged devices, unpackaged wafer form die, and multiple-device memory boards (via the adapter box shown at the left of the main control unit). Also seen here are two of the breadboard test system memory boards and a selection of 24 pin DIL IC packages.

This exerciser was designed to provide all of the waveforms needed to clear, write, and read 2 Kbit MNOS/BORAM memory chips. The unit consists fundamentally of a cyclic chain of monostables which generate pulses of various lengths. At the start of a cycle, a flag flip-flop is checked to set the timing to either a clear/write mode or a read mode. During IO operations that follow, the data registers are clocked at a front panel controlled rate of from 500 KHz to 5 MHz to affect transfer of data into or out of the memory chip. At the end of a clear/write cycle after the last row is written, the WRITE/READ flip-flop is cleared, and a new read cycle is started. In normal operation the exerciser reads the chip continuously but writes it only once upon command then returns to the continuous read sequence.

Either the data from the memory chip or an exclusive OR correlation of this data with the input to the device under test is displayed to reveal visually the placement of any invalid (erroneous) data storage sites in a chip operation. To facilitate rapid execution of these tests and to ensure thorough screening of all gross failure modes, the exerciser has been equipped to write and read 16 selectable data storage patterns into any chip under test. These patterns are depicted in figure 40. Although many other patterns are possible (and are in fact necessary to test all possible types of defects), this set was selected

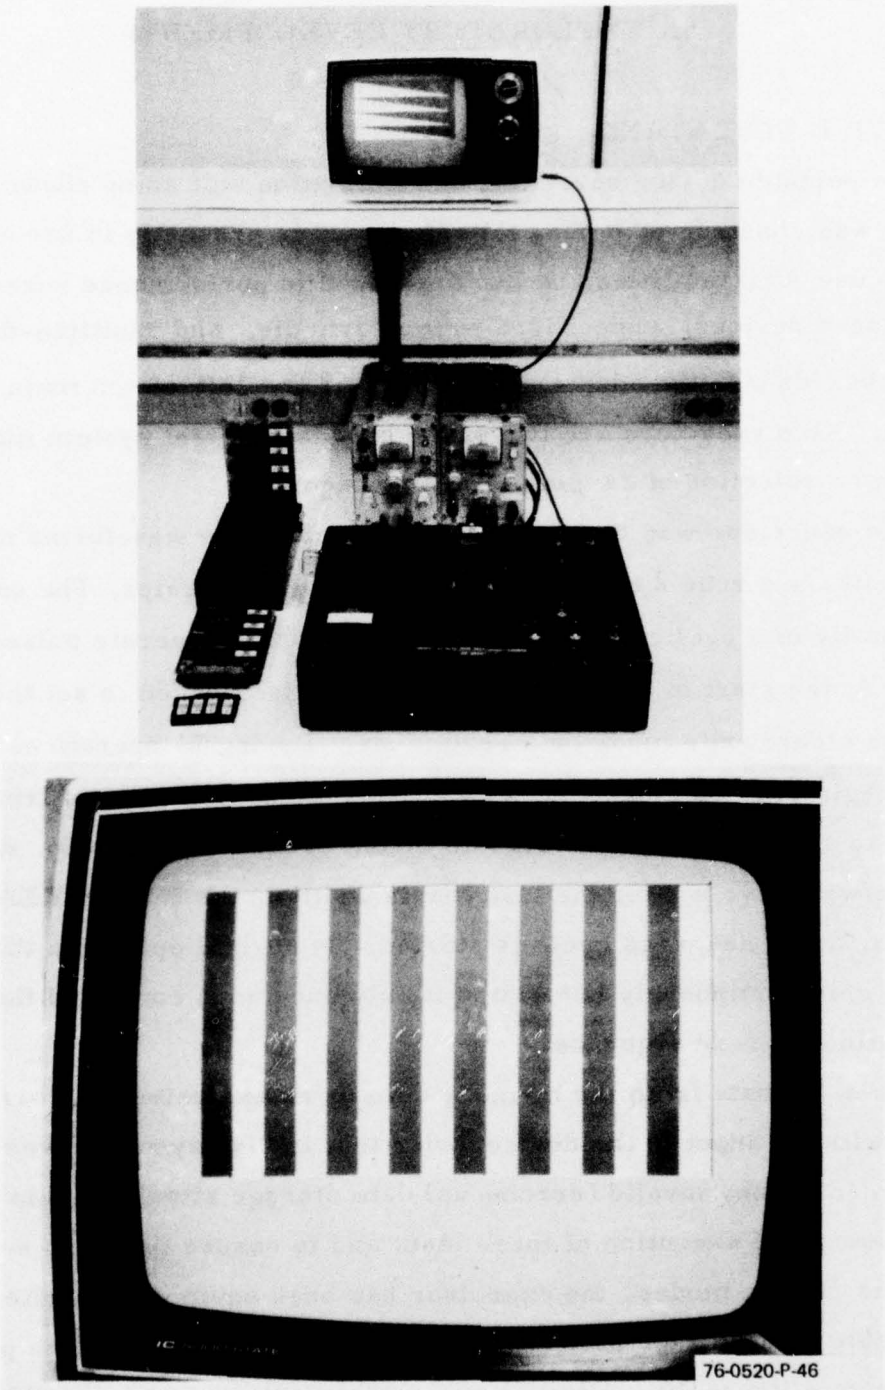

Figure 39. Developmental MNOS Device/Board Exerciser (Top) and Sample TV Monitor Display Pattern (Bottom)

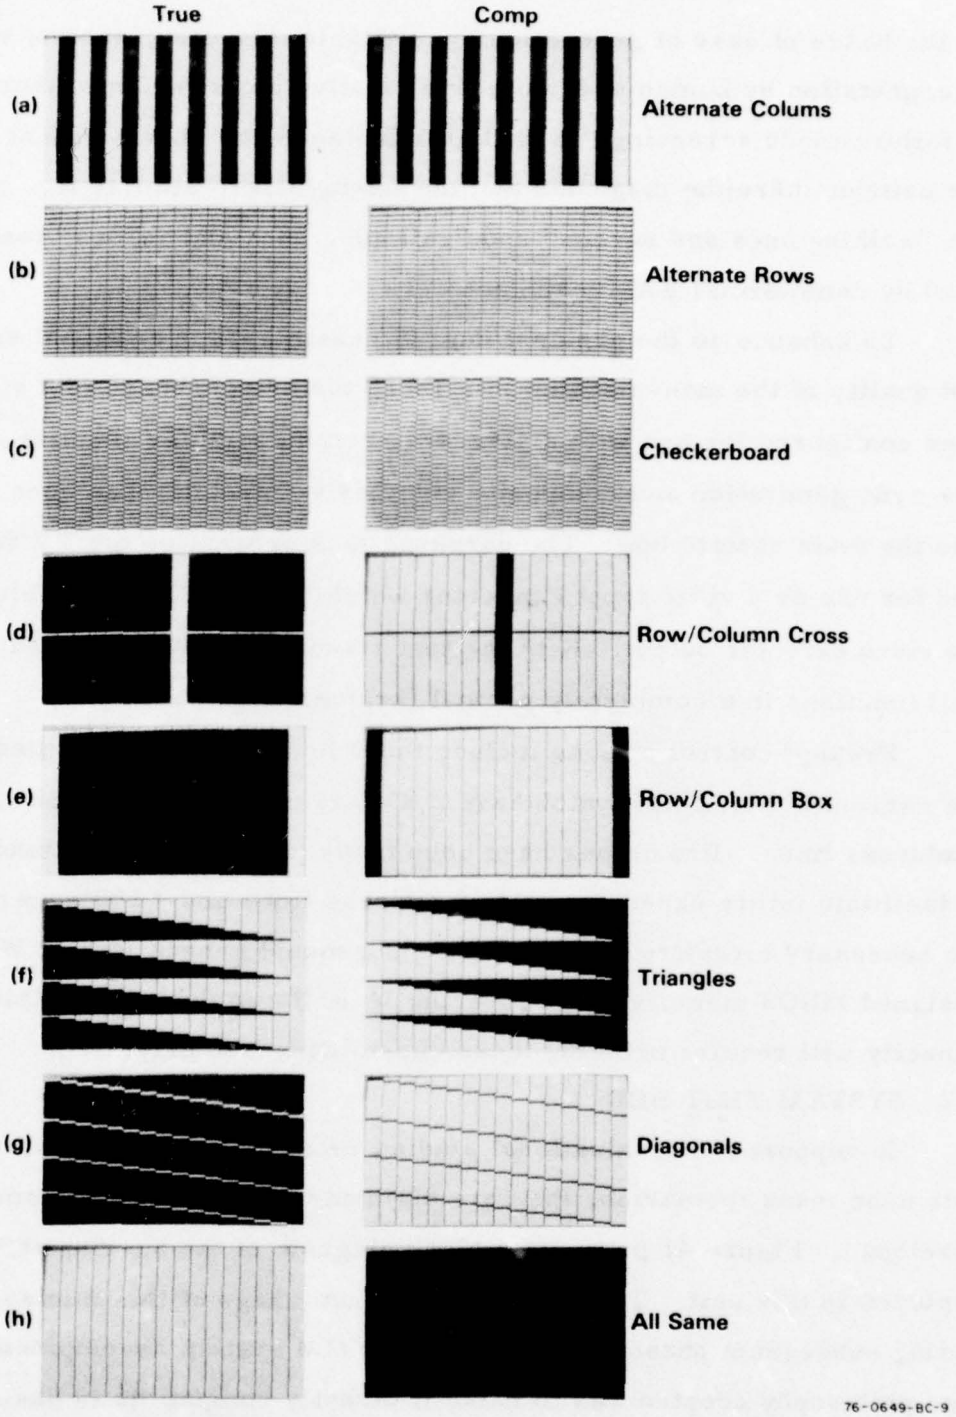

Figure 40. MNOS Device Screening Test Patterns

on the bases of ease of generation by available circuitry, ease of visual interpretation by human operator, and relative contribution to thoroughness of failure mode screening. It will, for instance, be observed that two of the patterns here(the diagonals and the triangles)are similar to, respectively, the "walking ones and zeroes" and "galloping ones and zeroes" test sequences used by commercial RAM manufacturers.

To enhance to the greatest degree possible its operational ease of use and quality of the man-machine interface, the device/board test stand has been configured for use with a standard, commercial TV monitor. All requisite sync generation and composite video drive circuitry has been incorporated into the main control box. The earphone jack present on the TV set was modified for use as a video input connector which, through a BNC cable, taps into the video detector output. With the test stand cable plug removed, the set still functions in a completely normal fashion.

Present control options include those necessary to fully exercise any of the current 2 Kbit chips which have 2 IO data channels per chip and require 6 address lines. Provisions have been made within the test stand, however, to facilitate future expansion up to 8 address lines and 4 I/O's per chip (i.e., the necessary circuitry is included). Subsequent generations of Westinghouse designed MNOS memory chips covering 16 Kbits up through 32 Kbits and 64 Kbits capacity will require bringing these capabilities into play.

5.2 SYSTEM TEST BED

In support of organizational studies relevant to the production of large data base mass memories, the core elements of a test bed system were developed. Figure 41 presents a block diagram of the bussing structure deployed in this unit. To facilitate maximum usage of this minisystem during subsequent phases of follow-on full MM system development, the design philosophy adopted was to make it directly compatible in basic organizational format with the interface seen at the DMA port between STARAN and the MM. The control registers and ledger stacks are full width to those which exist in the ACU's of the final system. Furthermore, the file "naming"

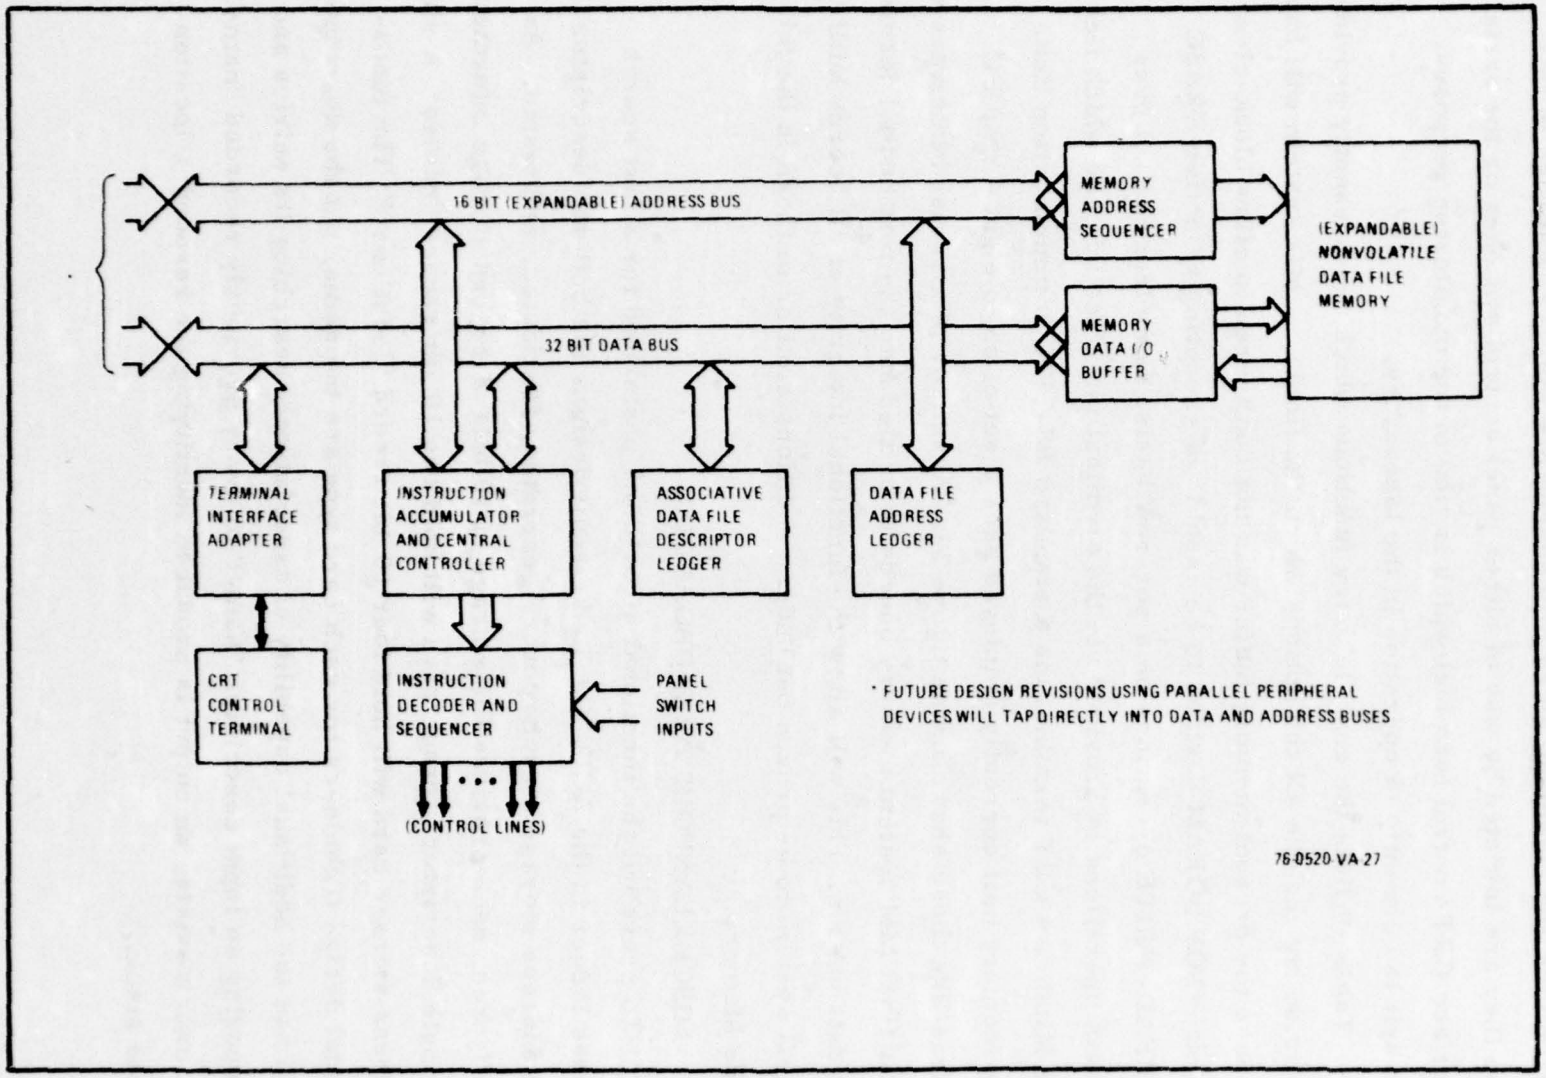

Figure 41. Test Bed MM System Organization

conventions and command entry procedures are similar, although in this case they are initiated by way of either panel control switches or the serial interface CRT control terminal which is tied in for monitoring purposes. The unit is currently in operation in the laboratory.

Table 19 lists the console entry functions which are presently provided. These do not include all the options which the final MM system will have because the present memory buffer was not configured to allow block-cleared 2 Kbit MNOS BORAM devices to be capable of executing selective ERASE or READ-WRITE operations on a per-row basis. Complete control over system operations is provided via the terminal interface adapter which ties to a Minibee 4 CRT terminal via a standard RS-232C communication link. The memory unit currently employed uses a set of up to eight 4-chip PC boards like those shown in the figure 39. Assembly of a revised interface will allow final MM system memory boards to be tied into the test bed's address and data busses. This will allow the functional integrity of all boards built for a final system to be proven out under conditions similar to those in the full-up Mass Memory.

5.3 MISCELLANEOUS BREADBOARDING

To prove out the functional algorithms postulated for a fast search access ledger in the ACU, detailed circuit designs for both the descriptor and address storage circuitry were generated, fabricated, and tested. As configured, each provides for storage of 10 bits of record storage information. A single descriptor storage card will hold two 10 bit record "names". A single address storage card will hold four 10 bit record "locations." The fundamental design techniques for each card type are the same, but the descriptor card has the additional capability of associatively searching its entries and comparing an input descriptor "name" to every previously recorded "name." If a match exists, an output is produced identifying the responder location in the stack.

TABLE 19

PANEL SWITCH ENTRY SYSTEM INSTRUCTIONS

	Panel Switch Entry System Instructions
	ON/OFF – System Power On-Off
	UP/DWN – Memory Power Up-Down
	SD – System Deactivate
	SA – System Activate
	WF XXXX – Write File (Named XXXX)
	FF XXXX – Find File (Named XXXX)
	RF XXXX – Read File (Named XXXX)
	FA – Read All (Data Only of Active Files Only)
	TC – Terminate Command
Storage Destructive Commands	AC – Abort Command
	CF XXXX – Clear File (Named XXXX)
	DF XXXX – Dump File (Named XXXX)
	DA – Dump All (Header and Data of All Active Files)
	CA – Clear All (All File Storage Space Fully Cleared)

76-0520-VA-29

Inputs and outputs to and from the register stacks on these boards are made via stack carry and feeder bus lines. The feeder bus also provides the channel whereby cycle-to-the-top updating is accomplished and accordingly allows securing access to a record located in the middle of the stack. In configuring these boards, a modularly commonalized design principle was applied so that the resultant record file stack could be expanded in either dimension by simply adding more identical boards. Horizontal stacking of cards (or card pairs) increases the number of address (or descriptor) bits available to each record. Vertical stacking, on the other hand, increases the number of record slots present in the directory stack. The results of this work confirmed the viability of the fast search access ledger algorithm.

Testing of the fast file descriptor and address storage boards was accomplished using a specially designed exerciser circuit. This circuit board provides all signals (except the master clocking strobe) necessary to the

functional checkout of either type file storage card. Availability of this exerciser board should expedite production of fast file cards during any subsequent prototype system developments.

To minimize development costs and maximize flexibility of fieldable MM systems, commonality of circuitry block designs appears to be a fundamental keyword. Following this principle, a detailed design was also generated for a multipurpose line repeater board which could be mass produced for universal application in either rail of all subsystem CIU's which must interface to an AAP like STARAN. An evaluation copy of this repeater design was breadboarded.

Dynamic testing of this board confirmed its universal interface functional capabilities. Analysis of the alternate functions which must be performed by the broadcast rail throughput repeaters and the response rail return repeaters produces the following findings:

Tapped Throughput Repeater Functions

1) Throughput
2) Line Tap
3) Input

Gated Return Repeater Functions

1) Throughput
2) Output
3) Cycle Back

In each case, the multipurpose 12 channel repeater card can supply the three requisite functional conditions.

Using 20 IC's, each multipurpose repeater card handles 12 channels. Further considerations of the most efficient way of implementing the line repeaters has shown that, despite the attractiveness of the universal board from a spare parts aspect, higher reliability and lower cost will be achieved by using two separate line repeater based designs: one for the tapped broadcast rail and one for the gated response rail.

A total of 5 memory boards were assembled for the MM test bed system discussed previously. All boards were fully populated with 4 each, no-defect, 2 Kbit MNOS memories. Four of these boards are now installed in the system test bed. The first of the five cards built used discrete device interfaces while all subsequent cards employed commercial IC's. Adoption of the latter construction concept serves to reduce inventory and construction costs. Moreover, it lends itself directly to assembly line production techniques (Such as wave soldering), which will be encountered in future system assemblies.

6. CONCLUSIONS AND RECOMMENDATIONS

This report has presented a number of wide-ranging discussions, some of which have been broad in nature and others of which have dealt with narrowly defined specific factors, related to the subject of high-speed, wide-bandwidth, multiport mass memories which may find use in future associative, parallel, and configurable processor facilities. To tie the diverse considerations together, three main points need to be covered:

- Identification of memory technology availability and grooming prospects for devices best suited to production of associative array processor compatible mass memories.
- Elucidation of suggested design implementations which are applicable to production of baseline through full-capability mass memory systems.
- Delineation of estimated scheduling for and major difficulties to be resolved during follow-on procurement of a developmental mass memory prototype.

It is the purpose of this section to treat these three topics in sufficient detail that RADC can proceed in an enlightened manner with the planning of near term (3 to 6 years hence) computer facilities which require the availability of reasonable cost, high-performance mass storage devices.

The choices made between memory technology alternatives and the suggested system design implementation techniques are derived from or are direct extensions of the material developed at length in prior sections. Consequently, certain points are called out or cited without further elucidation. Clarification of these points, when encountered, can be secured by referring back to the appropriate section in which they were first developed. The prototype follow-on scheduling estimates given are based on best engineering judgement and reflect present projections on both production device availabilities and required levels of effort to accomplish first-of-a-kind design mechanizations.

6.1 TECHNOLOGY PROSPECTS

6.1.1 Support

Westinghouse has full confidence in and is totally committed to the development and production of MNOS memory devices tailored to meet future mass memory system application requirements. In the past few years there has been increasing interest by the services in this technology as evidenced by the substantial support provided on ARMY, AIRFORCE, NAVY, and NASA contracts to help bring the technology base to maturity. Although initially slow in evolving, the MNOS technology is now justifying the attention that has been given to it.

In addition to Westinghouse supported efforts, there are at Westinghouse two separate ongoing funded contract programs aimed at the design and development of a family of higher capacity memory devices for future mass memory system applications. One of the contracts is for an 8 Kbit design, and the second is for a 16 to 32 Kbit design transition. The 16 Kbit design is presently near completion. It forms the basis of designs that will be configured for use in high-speed/high-capacity Mass Memory applications. All design modifications introduced into the 16 Kbit chip will subsequently be incorporated in the 32 Kbit device to provide for a logical growth sequence for securing lower overall memory costs. With the level of maturity that has been achieved in the basic MNOS technology, there is full confidence that the 16 Kbit devices will reach production status in the required program time frame.

Westinghouse is continuing a high level of support on the development of memory devices aimed at higher capacities and lower costs for future mass memory applications. These have a planned progression to 64 Kbits and then 128 Kbits, with a long range future goal of 1 Mbit per device. This overall planning and support includes the purchase and installation of new facilities to provide for higher resolution and therefore larger capacity devices with no increase in physical chip size. With no change in the level of support,

it is projected that the 64 Kbit device will reach production status by early 1979 with the 128 Kbit device following in 1980. Should demand prove sufficiently great that additional contract support to develop these higher capacity devices is obtained, the period to production maturity will be decreased sharply. These higher capacity units promise to make all future mass memory applications very highly cost-effective, both on an initial procurement cost basis and on an installed cost-of-ownership basis.

6.1.2 Status

At Westinghouse, the recent run of the 6000C BORAM device, developed under joint Army/Navy auspices, have produced an excellent, reproducible memory device with exceptionally stable performance characteristics and good production yields. This latest 2 Kbit chip achieves these characteristics through the use of a two transistor (2T) memory cell structure and a differential detection mechanism in conjunction with a simplified process sequence.

Experimental evidence has shown that the minimum geometry for an MNOS memory cell is well below the present design rules. The memory structure begins to malfunction due to fringing fields when the memory window width approaches 1 μm. Devices have been fabricated and successfully operated with a 1.6 μm memory window and a 1.6 μm diffusion. This indicates that improvement in production photolithography will allow the minimum device geometry to be reduced from 4 to 2 μm, with a subsequent increase in density.

The density can also be increased by adding a second layer of polysilicon. Development work has begun on a double polysilicon process which would produce, with no other design rule changes, a 32 Kbit memory on a die 225 x 200 mils. Double polysilicon provides new layout options that permit a significant decrease in row spacing in the memory array. By combining double poly with 2 μm geometries, memory elements can be built within 0.15 mils^2 of area.

In addition to these approaches, other MNOS development efforts are exploring possibilities that may have major impact. Foremost is development of both N and P channel MNOS devices using a self-aligned polysilicon process. Self-aligning processes reduce geometries as well as increase speeds and improve producibility by eliminating alignment tolerances. In total, the growth outlook for MNOS memories is seen to be excellent, with a high probability of realizing 64 Kbit LSI devices within 3 to 4 years.

6.1.3 Grooming

For the MM application it is proposed that the existing 16 Kbit MNOS memory chip design be modified to incorporate certain unique features which will enhance the performance of the entire memory system. A block diagram of the 16 Kbit chip as it would appear if specifically groomed for use in the dual-ported Mass Memory is presented in figure 42.

The storage array still comprises 16,384 cells arranged in a 128 x 128 matrix. Two additional 32 bit shift registers have been added on the groomed device to provide four 32 bit registers. This provides 4 I/Os per chip and reduces the number of chips that must be paralleled. Furthermore, a dual-port organization is incorporated directly on-chip, with 3 port-control lines determining the data port functions. Each shift register input or output can be linked to either port A or port B by the control logic. Each port interface line is made bidirectional, thereby allowing it to be used as either an input or an output as determined by the control lines. Input and output buffers swing CMOS voltage levels, and the output drivers can be switched to a high impedance state by the control logic. Inclusion of these features at the chip level greatly reduces the number of support logic circuits required to operate the memory in a system that must have two isolated primary access ports.

A unique characteristic of this memory is sequential addressing of the 128 rows. Since the system operates in a sequential mode, following location of the first row, the normal random access row decoder can be

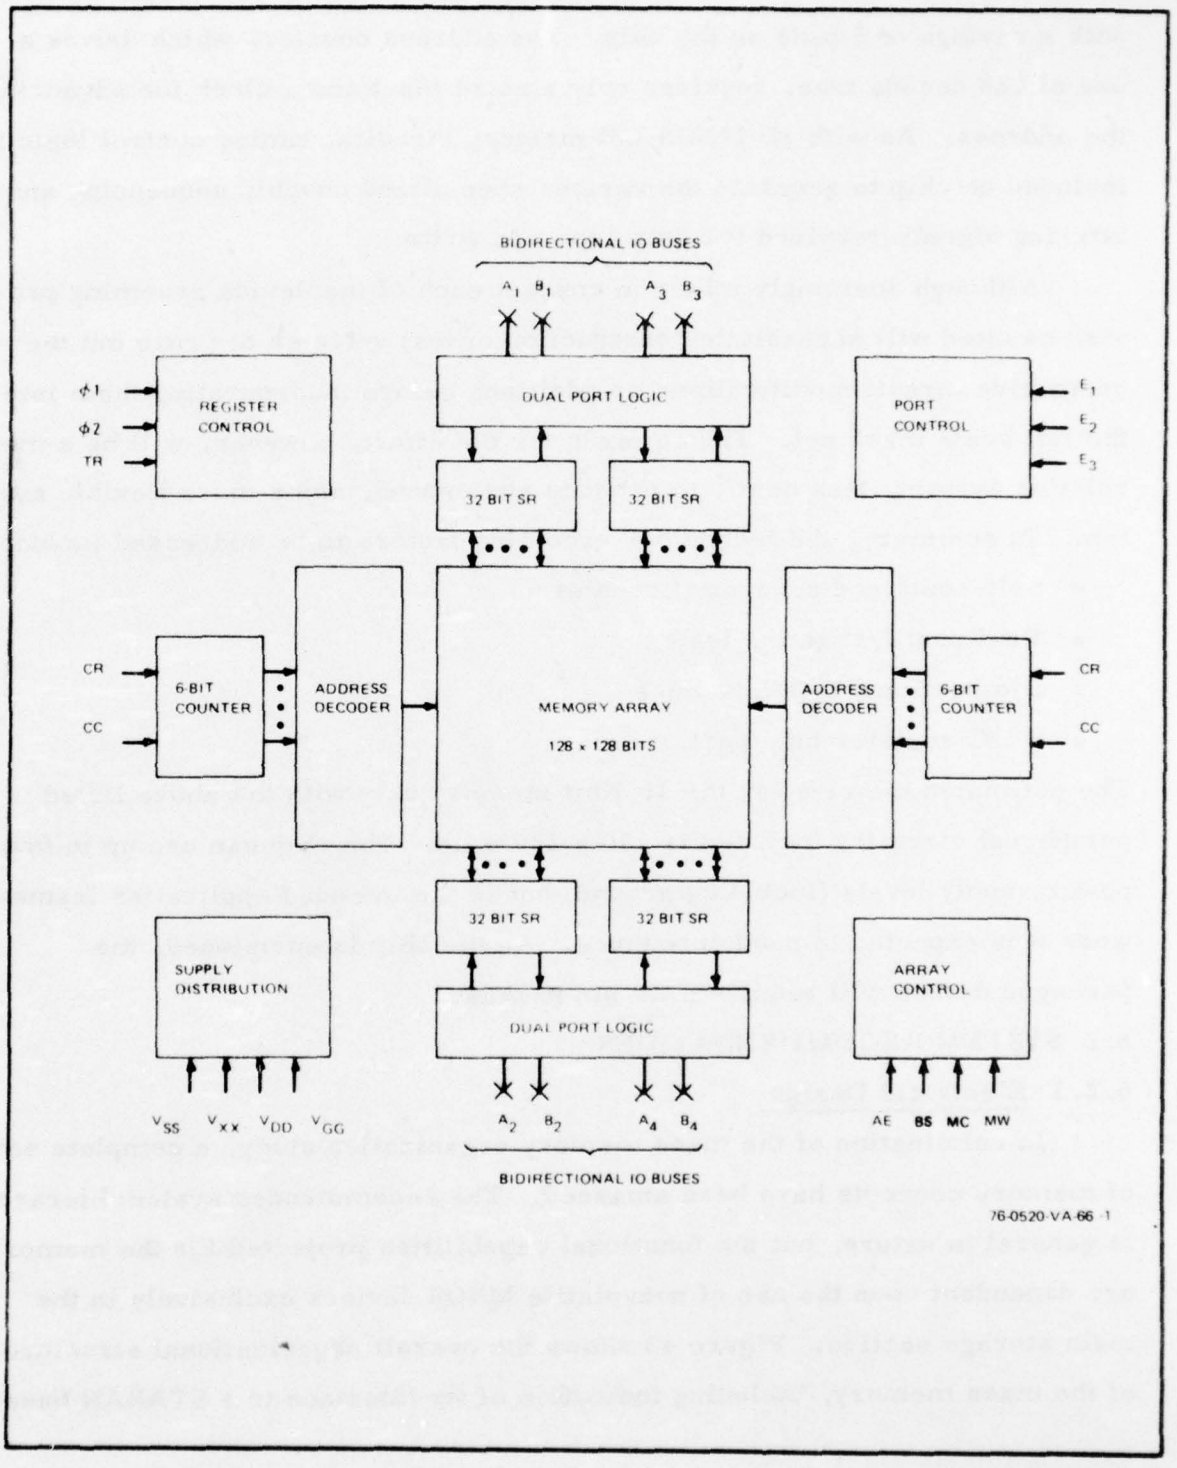

Figure 42. Dual-Port Application Groomed 16 Kbit MNOS Memory Chip

eliminated. A simple binary counter of 7 stages can replace 7 input buffers, with a savings of 5 pads on the chip. The address counter, which drives a one of 128 decode tree, requires only a reset input and a clock for advancing the address. As with all MNOS LSI memory circuits, timing control logic is included on chip to generate the various specialized on-chip sequencing and latching signals required for data read and write.

Although seemingly minor in concept each of the device grooming provisions cited will necessitate construction of test vehicles to prove out the respective circuit modifications or additions before incorporating them into the full scale mask set. The rewards for the effort, however, will be a more reliable system, less costly to produce and expand, and a more flexible system. In summary, the technology grooming factors to be addressed include:

- Self-contained sequential counter
- Dual port I/O gating logic
- Bidirectional I/O data lines
- 4 I/O register bits/port.

The estimated die size for the 16 Kbit memory chip with the above listed peripheral circuitry included is 190 x 220 mils. The chip can use up to four power supply levels (including ground) but in the intended application framework it is expected to need only three. As the chip is envisioned, the packaged device will require a 24 pin package.

6.2 SYSTEM RECOMMENDATIONS

6.2.1 Electrical Design

In culmination of the mass memory organization study, a complete set of memory concepts have been amassed. The recommended system hierarchy is general in nature, but the functional capabilities projected for the memory are dependent upon the use of nonvolatile MNOS devices exclusively in the main storage section. Figure 43 shows the overall organizational structure of the mass memory, including indication of its interface to a STARAN based

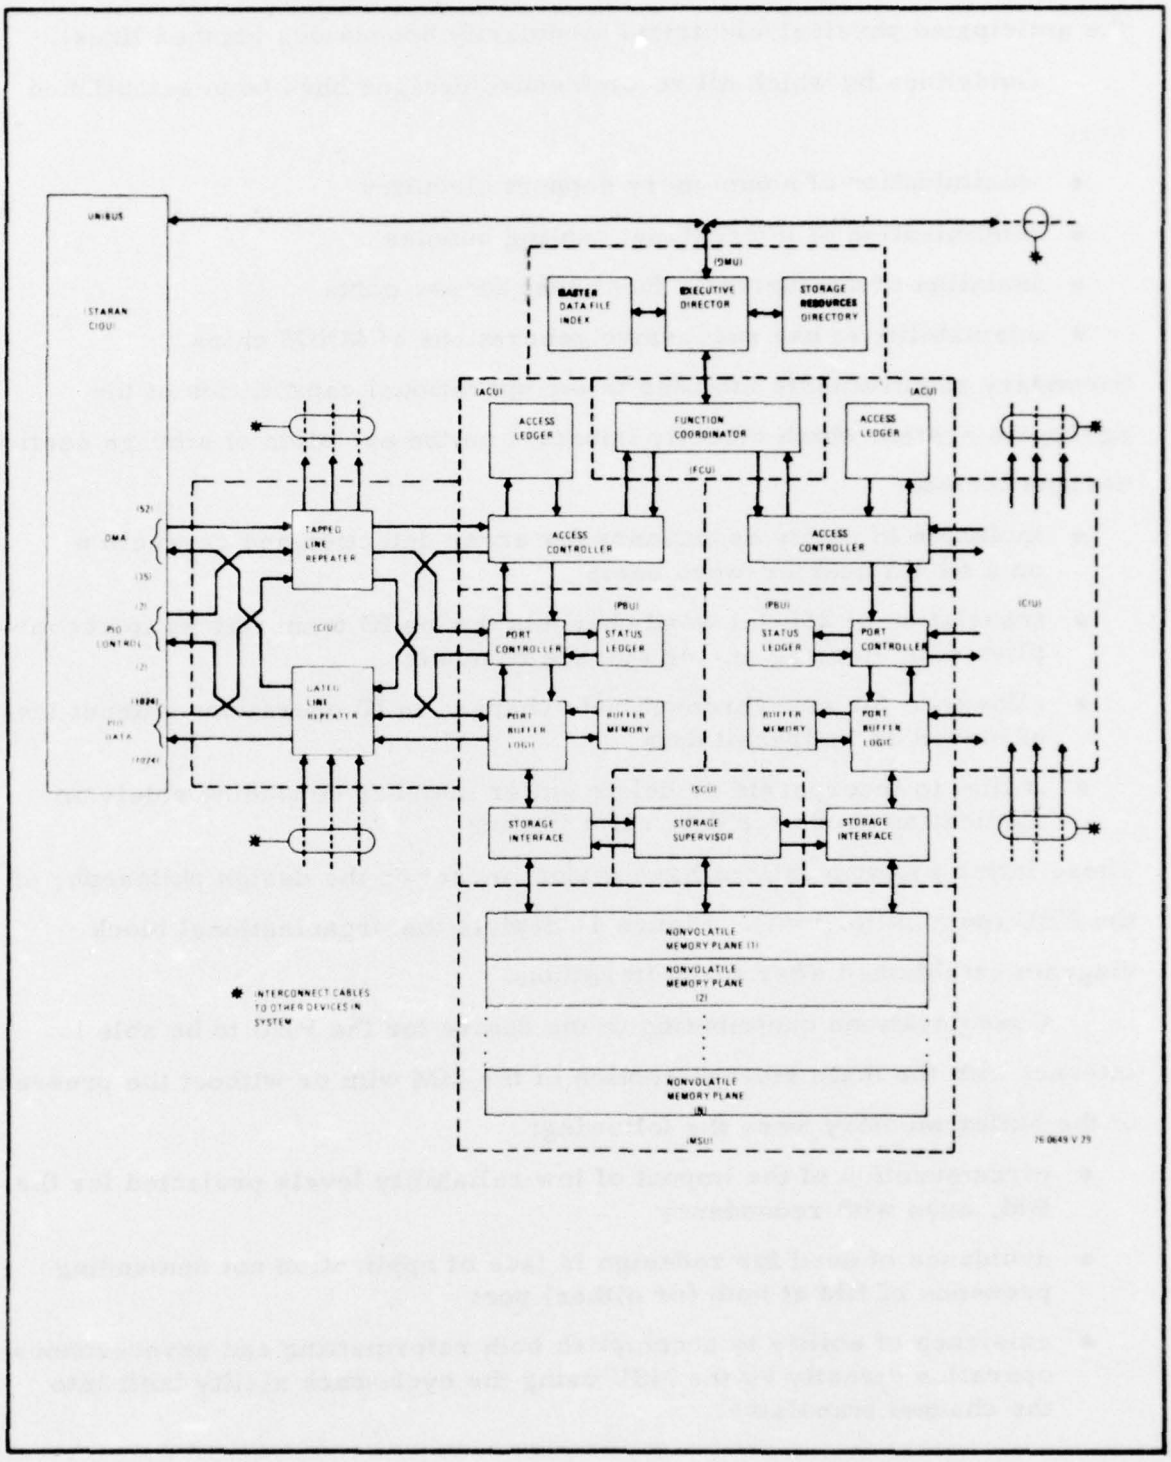

Figure 43. Mass Memory Organizational Block Diagram

associative array processor (APP) computer system. Also indicated are the anticipated physical/electrical modularity boundaries (dashed lines).

Guidelines by which all recommended designs have been established are:

- minimization of nonmemory support circuitry
- minimization of intercabinet cabling bundles
- isolation of concurrently functional access ports
- adaptability to use successive generations of MNOS chips.

Secondary requirements included in the operational capabilities of the aggregate system which directly impacted on the evolution of storage section designs include:

- inclusion of parity redundancy for error detection and correction on a 64-bit quarter-word basis
- translation of 256-bit word channels during IO transactions to accomplish data steering among storage channels
- allowance for asynchronous rate changes in IO operations without loss of stored or in-transit data
- ability to incorporate or delete buffer memory depending solely on application needs for data reformatting

These latter two stipulations have major impact on the design philosophy of the PBU (port buffer unit). Figure 44 depicts the organizational block diagram established after many iterations.

Considerations contributing to the desire for the PBU to be able to interact with the main storage section of the MM with or without the presence of the buffer memory were the following:

- circumvention of the impact of low reliability levels projected for the BM, even with redundancy
- avoidance of need for redesign in face of application not demanding presence of BM at both (or either) port
- existence of ability to accomplish both reformatting and asynchronous operation directly by the MSU using the cycle-back ability built into the channel translator.

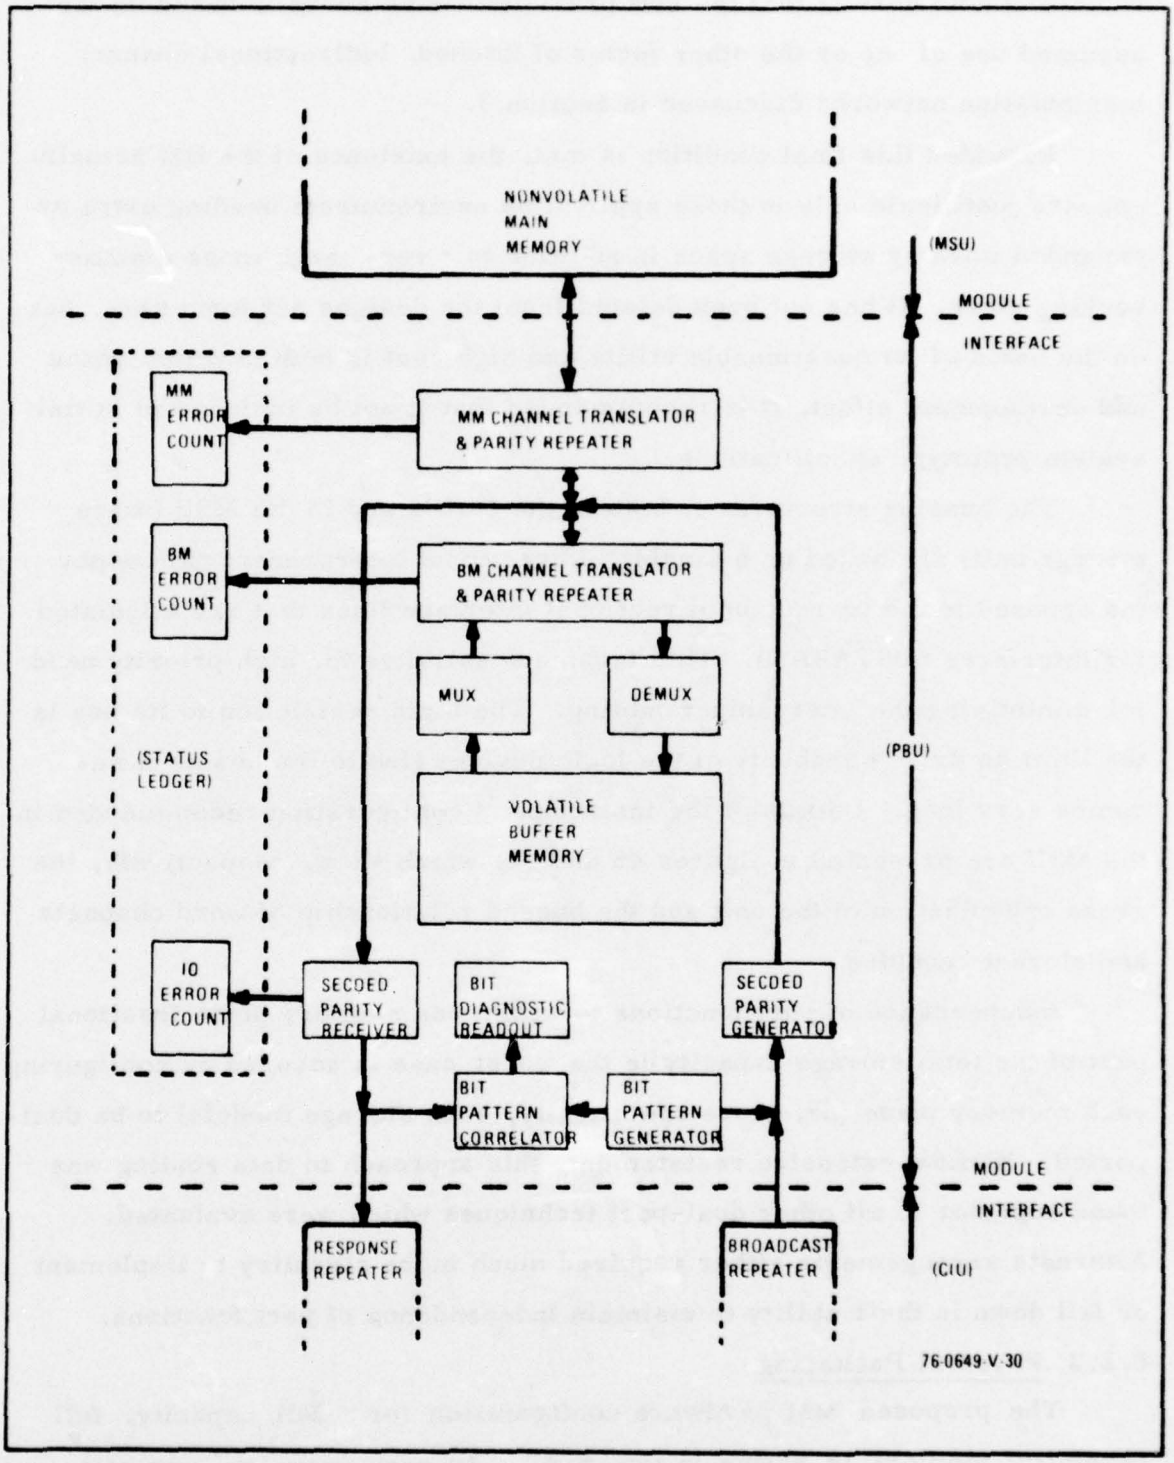

Figure 44. Port Buffer Unit Organizational Block Diagram

A point of note here is that the last provision listed above is based on the assumed use of one or the other forms of latched, bidirectional channel manipulation networks discussed in Section 3.

Provided this final condition is met, the existence of the BM actually appears justifiable only in those application environments needing extra or expanded working storage space in addition to a very large mass memory backing store. It has not been deleted from the designs set forth here, but on the basis of its questionable utility and high cost in both physical space and development effort, it is recommended that it not be included in initial system prototype specifications.

The bussing structures in both of the PBU's and in the MSU (mass storage unit) are based on a single, bidirectional interconnect philosophy (as opposed to the paired, unidirectional interface lines that are stipulated for interfaces to STARAN). This technique satisfies the high priority need for minimizing the intercabinet cabling. The main restriction to its use is the limit on drive capability of the logic devices tied to the bus as it becomes very long. Details of the interconnect configuration recommended in the MSU are presented in figures 45 and 46, which show, respectively, the gross organization of the unit and the bussed relationship of word channels and storage modules.

Independence of port functions to within one memory plane fractional part of the total storage capacity in the worst case is secured by configuring each memory plane (or, more specifically, each storage module) to be dual-ported. Without extensive restatement, this approach to data routing was found superior to all other dual-port techniques which were evaluated. Alternate arrangements either required much more circuitry to implement or fell down in their ability to maintain independence of port functions.

6.2.2 Physical Packaging

The proposed MM hardware configuration for a full capacity, full capability memory is shown in figure 47. As seen here, the complete system consists of 8 standard 7-foot-high 19-inch card racks plus one low-boy

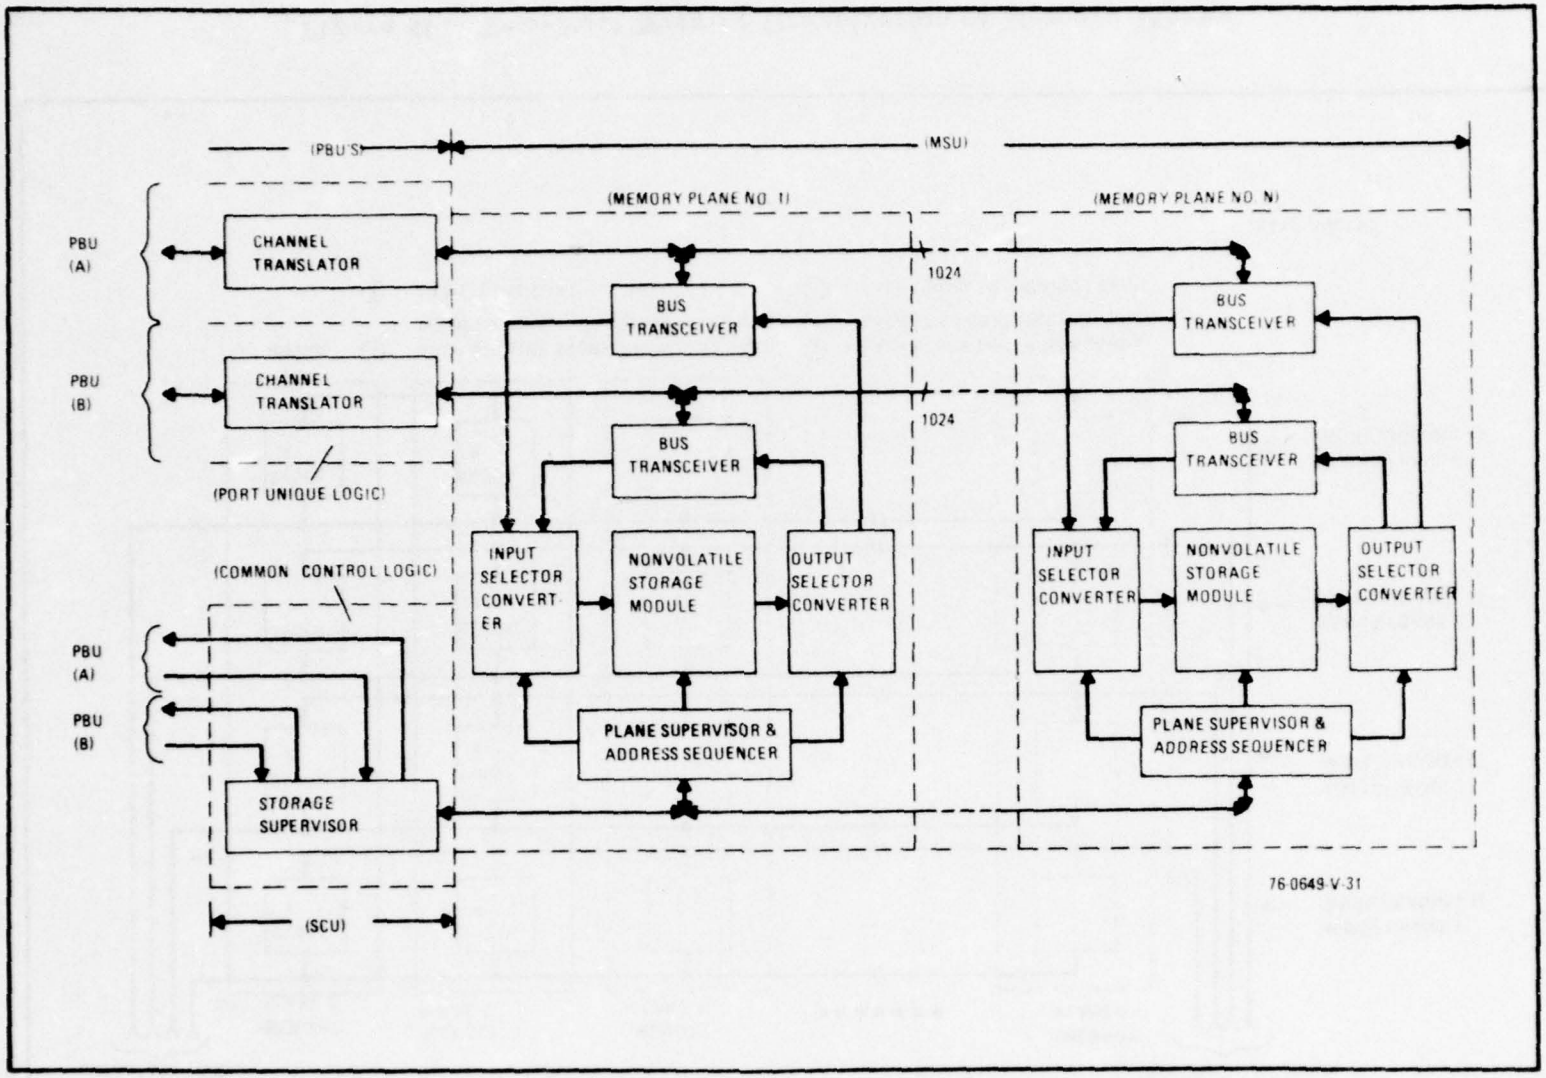

Figure 45. Mass Storage Unit Organizational Block Diagram

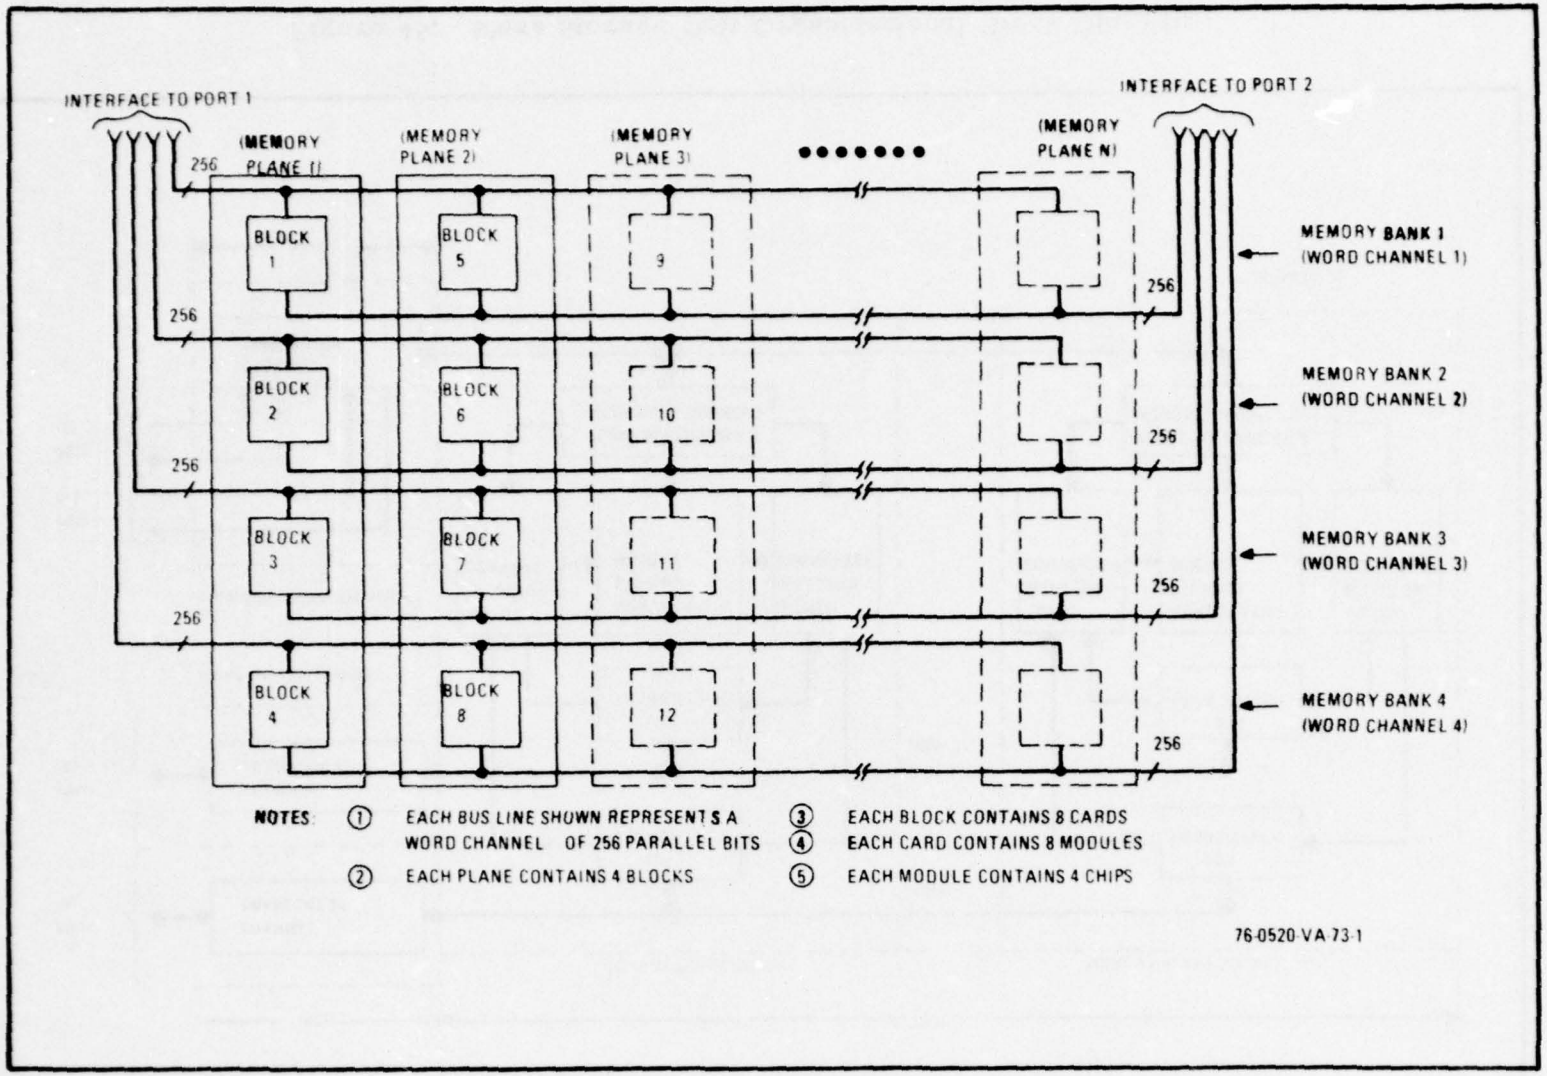

Figure 46. Two-Port Bussing Organization of Memory Planes

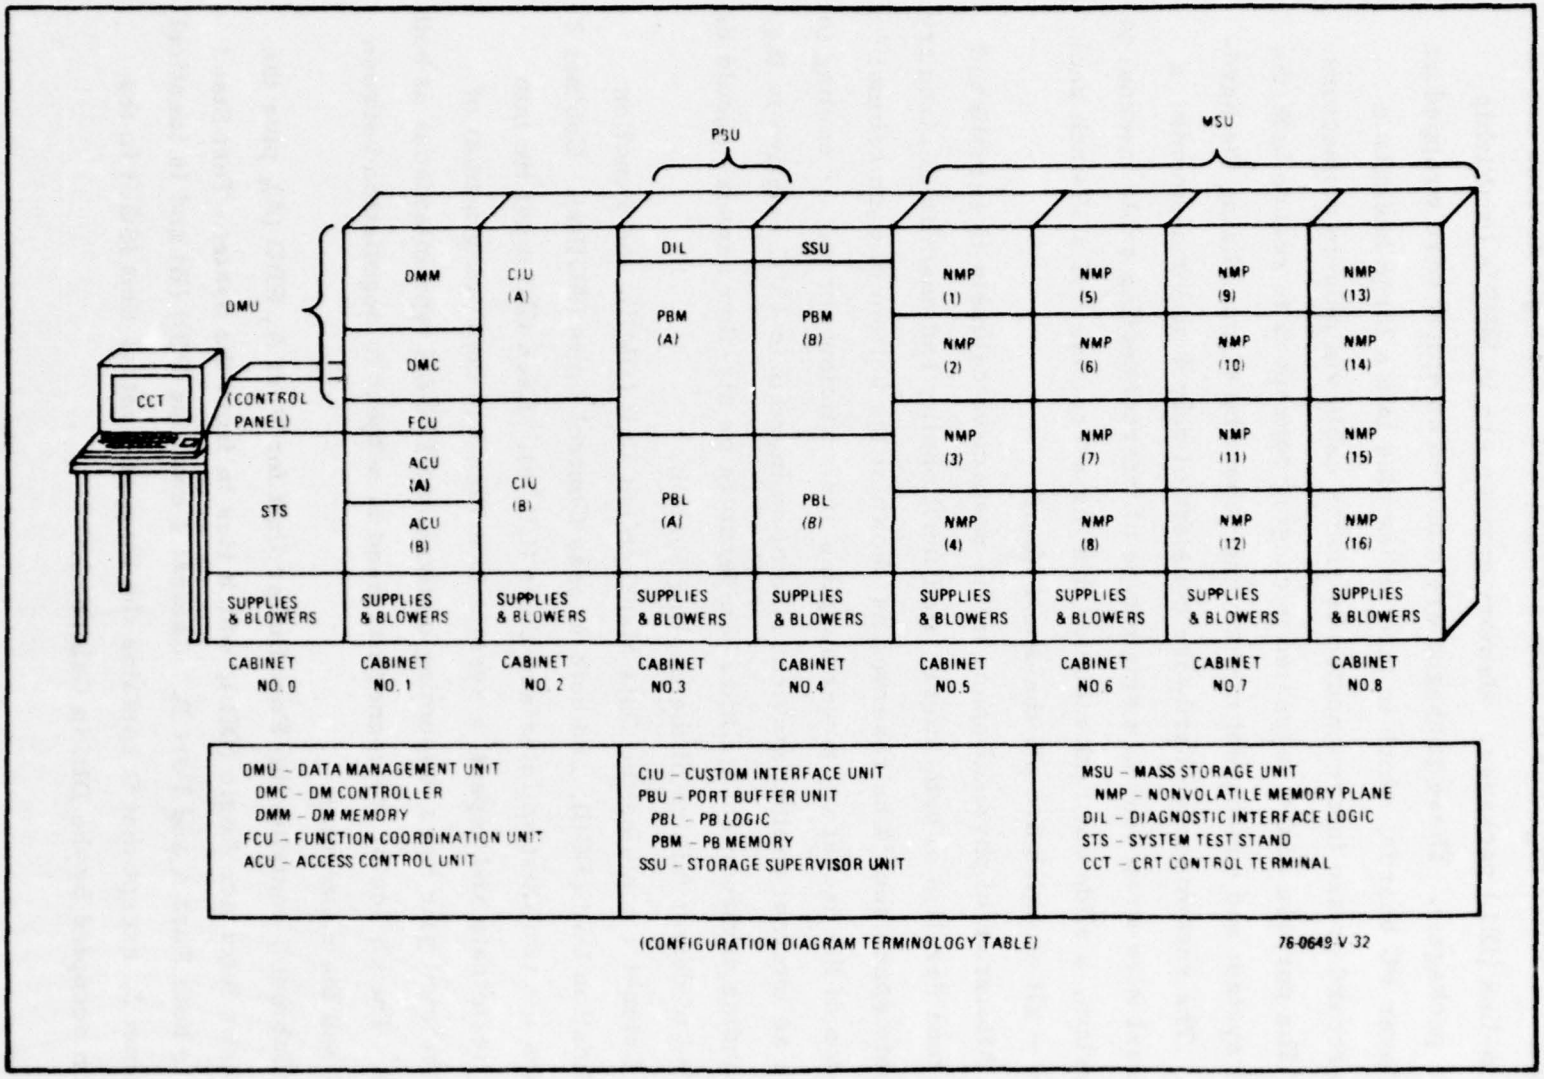

Figure 47. Full-Capacity, Dual-Ported, Commercially Packaged Mass Memory System Configuration Diagram

cabinet, all bolted together to form a single unit. All logic devices are in dual-in-line (DIL) packages. Memory devices are in MHP's (multichip hybrid packages). These packages are mounted on either wire wrapped or multilayer PC boards, which in turn either plug into a large backplane motherboard or are interconnected by ribbon cable via plug-in connectors.

The purpose of the suggested packaging concept is to reduce both the initial system and subsequent maintenance costs by standardizing the hardware. The recommended hardware consists of only 4 major elements: a universal wire wrap board, a single type of interconnecting cable internal to the system, a slide out rack mounted drawer assembly, and a 19-inch rack frame -- all of which are off-the-shelf items.

Allocation of physical space in the respective cabinets is established from consideration of both interrelated functionality and quantity of circuitry. Adequate space should be reserved as shown at the bottom of each cabinet for inclusion of the necessary power supplies and sufficient forced air cooling to ensure an internal cabinet ambient rise of no more than 15°C relative to the surrounding environment. Either temperature or air-flow sensors should be included to detect failure of the cooling system.

Cabinet 1 contains the Data Management Unit (DMU), the Function Coordination Unit (FCU), and both Access Control Units (ACU's). Cabinet 2 contains the two Custom Interface Units (CIU's). Each CIU must tie into (i.e., terminate and repeat) a system interface bus containing a total of 2,140 twisted pair lines (assuming equivalent STARAN type interfaces at both ports). The choice of the connectors used is subject to negotiation between RADC and the contractor.

Cabinet 3 contains the Port Buffer Unit for Port A, PBU (A), plus the Diagnostic Interface Logic (DIL), which ties in from the System Test Stand (STS) to both Port A and Port B. Cabinet 4 contains PBU (B) and is identical to Cabinet 3, except that it contains the Storage Control Unit (SCU) in the position occupied by the DIL in Cabinet 3.

Cabinets 5 through 8 contain the Mass Storage Unit (MSU), which consists of 16 memory planes with 4 planes housed in each cabinet. Additional memory cabinets can be chain-connected alongside of cabinet 8 to expand the system to whatever degree desired. Alternately, cabinets can be deleted for lower capacity applications.

Each functional unit (FCU, ACU, etc) is assumed to have a control panel associated with it. The respective control panels contain power supply switches, error and failure indicator lights, and all other controls and indicators uniquely required by the unit. The individual control panels are located in the same cabinets as the units they service.

6.2.3 Prototype Development

For purposes of prototype development, the full-up Mass Memory configurations that have been set forth can be cut back to appropriate levels to secure a functional evaluation vehicle on which detailed tests and feasibility analyses can be run. To accomplish this goal in the most cost-effective manner, it is recommended that the implemented capacity of the prototype be cut back to the equivalent of either two or four planes of storage. Further, that all circuitry and functional units associated with the 2nd port, which is presently still undefined as to interface requirements, be deleted. This means that only a single CIU, ACU, and PBU would be built.

The cabineting of a prototype system would be cut back accordingly. Cabinets 3 and 6 through 8 would be deleted, and cabinets 1 and 2 would not be fully populated. An option in this development exists in the case of the buffer memory. Its omission is suggested from considerations of cost to implement and questionable utility, particularly in an initial development.

Exclusion of the BM has the potential additional advantage of offering further cost savings later if expansion of the prototype is anticipated to include the second port. This comes about by virtue of the empty space existing in the PBU cabinet if the BM is omitted.

Another consideration of prototype development is capacity of the memory chips employed. Although the probability of production quantities of the proposed 16 Kbit MNOS is quite high based on present schedules and work progress, the possible need to back off to a smaller chip size cannot be overlooked. Table 20 illustrates the impact which alternate capacity memory components have on the total number of planes in the system at various capacity levels.

Accordingly, the advanced development model (ADM) prototype mass memory recommended by Westinghouse is a 1/32nd capacity system populated with fully groomed 16 Kbit MNOS memory chips. With this particular choice of memory device it will be possible to demonstrate full compliance with the MM prototype guidelines of table 21. In addition, much of the system support logic circuit requirements will be reduced as they will be

TABLE 20

VARIATION IN NUMBER OF MEMORY PLANES WITH CAPACITY & DEVICE SIZE

MM Fractional Capacity	Memory Chip Size					
	64Kbit	32Kbit	16Kbit	8Kbit	4Kbit	2Kbit
1	16	32	64	128	256	512
1/2	8	16	32	64	128	256
1/4	8	8	16	32	64	128
1/8	2	4	8	16	32	64
1/16	1	2	4	8	16	32
1/32	-	1	2	4	8	16
1/64	-	-	1	2	4	8

TABLE 21

MASS MEMORY PROTOTYPE GUIDELINE

Capacity	1 Gigabit (1,073,741,824 Bits)
Modularity	Fractional Capacities of 1/16 or 1/32
Technology	Nonvolatile MNOS
Integrity	Protected Against Power Transients and Loss
Compatibility	STARAN AAP System CIOU
Control Medium	STARAN DMA Channel
Data Medium	STARAN PIO Channel
Interface Busses	Differential Twisted Pair ECL Lines
Bus Configuration	Terminated and Repeated, Matching Original Line Characteristics
Bus Operations	DMA Channel Bus Lines Tapped and Repeated Continuously
	PIO Channel Bus Lines Gated to Prevent Transfer Interference
Selection and Control	Exercised via Block of DMA Channel Addresses
Control Provisions	Create, Modify Content and Delete User Files
	Declare All Available Storage Space to be One File
	Suspend Access and Relinquish PIO Channel
	Resume Access at Point of Suspension
	Terminate Access (With Orderly Close-Out)
	Read and/or Clear Individual Status Registers
	Select File Access Functional Mode
	Master Reset of All Control Functions
Functional Modes	Read, Write, Read-Write, Clear, and Read-Clear of Individual Files
	Plus Dump and Clear-All Contents of Memory
File Management	Provision for 64 Users with 256 Files/User
Resources Management	Preventive of Nonavailable Space Accumulation
I/O Facilities	Two Ports — Independent and Concurrently Operable
Port Size	Four Parallel Word Channels/Port
Word Size	256 Parallel Data Bits/Word
Channel Size	4096 Randomly Accessable Storage Blocks/Channel
Block Size	256 Sequentially Ordered Word Spaces/Block
File Structure	256 Lines by M Blocks Per File (M is User Assignable)
Line Size	1-To-4 Parallel Words/Line (User Designated)
Data Entity	1 Word (Smallest Retrievable File Element)
Address Entity	1 Block (Smallest Addressable File Element)
IO Entity	1 Line (Smallest Transferrable File Element)
Address Format	Indirect, Block Relative, Incrementally Post-Indexed (Via File Name, Lines of Offset, and No. Lines Transferred)
Acquisition Time	$< 50\ \mu s$ Average Latency to 1st Line of File
Availability Delay	< 400 ns Latency Add-On per Line of Offset from 1st
Transfer Control	Asynchronous Master/Save Handshake (With User Device, STARAN, Originating all Inter-Device Activity)
Transfer Format	1 Word/Line File Accesses via Word Channel 0
	2 Word/Line File Accesses via Word Channels 0 and 1
	3 Word/Line File Accesses via Word Channels 0,1, and 2
	4 Word/Line File Accesses via Word Channels 0,1,2, and 3
Data Flow Rate	∞ To 450 ns/Line (Asynchronously Variable)
Error Rate	$< 10^{-10}$ Errors/Bit Transferred
Error Detection	Automatic, With Running Tabulation Kept of Total No.
Failure Detection	Instituted via BIT Facilities, with All Failure Flagged
BIT Facilities	Off-Line Status Validation and Fault Isolation to Board Level
Environment	Ground Benign (As per MIL-HDBK-217B Definition)
Ambient Temperature	$21^{\circ}C \pm 5^{\circ}C$ ($70^{\circ}F \pm 10^{\circ}F$)
Relative Humidity	$50\% \pm 10\%$
Component Qualification	JAN MIL-M-38510, Class B
Reliability (MTBF)	$>$ 200 Hours
Maintainability (MTTR)	$<$ 60 Minutes (Average)
Power Consumption	$< 10 \mu W$/bit

76-0520-TA-33-1

added to the basic MNOS device. In the proposed configuration, the memory devices will be in a separate cabinet. The main memory will consist of two planes of memory, with each plane providing up to 1,024 parallel outputs. Additional memory planes can be provided with no system modifications required. Provision also is made allowing for subsequent addition of a second port buffer unit to facilitate an additional I/O port.

6.3 PROTOTYPE FOLLOW-ON PLAN

The technology prospects for an appropriate prototype mass memory system populated with MNOS memory devices are excellent for the design and fabrication of a follow-on unit to interface with STARAN. Design of the required basic prototype system will be a straightforward engineering task after all the interfaces with STARAN have been finalized and all data management functions agreed upon. An ADM meeting all of the basic performance requirements of the overall system can be readily designed and produced using the present 2 Kbit MNOS memory devices.

The MNOS memory technology has emerged from the development status and is now being matured for production reality military system applications. ECOM has recently awarded Westinghouse a Manufacturing Methods Techniques (MMT) program to bring all of the associated processing, assembly, and test techniques on the MNOS technology to a production maturity status. Although the MMT program is directed at the 2 Kbit BORAM MNOS memory device processing and subsequent MHP assembly and test, the results obtained on the program will have direct application to higher capacity memory devices, thereby expediting advancement of these units to production status. Pilot line production capabilities for the MNOS devices can readily reach or exceed around 30,000 chips per month and satisfy all production requirements now forseen.

Higher capacity MNOS memory devices are still slightly down stream. Westinghouse is currently working under ECOM contract on a 16 Kbit memory device which has the potential of being extended to 32Kbits, for future BORAM

type applications. Although this effort is progressing satisfactorily, the schedule being followed is marginal for reaching unqualified production status within the time frame of interest to development of a prototype mass memory. With minimal additional support, this design could be modified and groomed for the mass memory application and the time to production status accelerated.

Problems associated with the design and processing requirements of a family of higher capacity memory devices, with a long-term objective of 1 Mbit per chip, are being studied on Westinghouse funds. Included in this in-house effort are long-term plans to improve and upgrade the process facilities. This will provide the capability of handling larger wafers with more automated equipment and a much higher resolution in photoengraving. The end result will be the ability to obtain denser arrays on a smaller allotment of real estate with less personnel handling to provide the higher capacity memory devices at a lower cost per bit.

Since the technologies for the mass memory architecture and the 2 Kbit MNOS memory devices are both low risk in terms of providing a suitable prototype mass memory for the planned application, the overall follow-on program should consider how to best reach a full-up 10^9 bit system capability with a logical sequence of programmed events. Such a program could best be conducted on two parallel paths: one to design and deliver a partially populated memory, and the other to provide support to the higher capacity MNOS memory device effort.

Initially, a partially populated prototype memory would be designed, delivered, and interfaced with STARAN to provide a usable mass memory with which extensive investigations of parallel processing with STARAN could be carried out. This unit would be designed and have all support circuits for the full 10^9 bit capacity memory but would only be populated with memory devices to a 1/16th or a 1/32nd capacity level. This should be adequate to ensure

validity of the results of the initial concept and problem investigations. Design of this memory system would be configured to permit additional increments of memory to be added with any of the larger capacity MNOS memory chips with no changes in the basic memory system nor in its design being required. Full performance capability, with the exception of the smaller memory capacity, would thereby be provided by this unit with minimum cost and risk.

As there is no major risk forseen with this approach and it is fully anticipated that the investigations of parallel processing using STARAN will indicate that this concept is valid, it is practically assured that the full 10^9 bit capability will be desired. It would be advantageous at this point to have higher capacity memory devices to decrease the cost of procurement and the volume requirements of the supplemental memory units. In anticipation of this requirement for the full-up 10^9 bit memory, the parallel effort to expedite bringing higher capacity memory devices tailored for the application to production status is highly advantageous. This would be a cost effective approach to obtaining additional memory capacity to round-out population of the initial advanced development model and to mechanize the next generation mass memory system.

BIBLIOGRAPHIC ENTRIES

Batcher, Kenneth E.; "The Multi-Dimensional Access Memory in STARAN"; 1975 Sagamore Computer Conference on Parallel Processing; p. 167.

Bentchkowsky, D. F.; "A Fully Decoded 2048-Bit Electrically Programmable FAMOS Read-Only Memory"; IEEE JSSC, Special Issue On Semiconductor Memories and Digital Circuits; vol. SC-6, No. 5; October 1971; pp. 301-306.

Boyd, H.N.; "An Associative Processor Architective For Air Traffic Control"; Proceedings of 1974 Sagamore Computer Conference on Parallel Processing; pp. 400-416.

Brewer, J.E.; "MNOS Secondary Storage"; 1974 NAECON Record.

Brewer, J.E. and Hadden, D.R., Jr.; "Block-Oriented Random Access MNOS Memory"; National Computer Conference; May 1974.

Brewer, J.E.; "Army/Navy MNOS BORAM"; submitted for presentative at 1976 Government Microcircuit Applications Conf. (GOMAC), Lake Buena Vista, Florida, November, 1976.

Chang, Hsu, Editor; Magnetic Bubble Technology; IEEE Press; New York; 1975.

Chen, C.J. and A.A. Frank; "On Programmable Parallel Data Routing Networks Via Cross-Bar Switches for Multiple Element Computer Architectures"; Proceedings of 1974 Sagamore Computer Conference on Parallel Processing; pp. 338-369.

Cricchi, J.R.; "Direct Tunneling from the Selected State in MNOS Memory Devices," 5th IEEE Semiconductor Interface Specialists Conference.

Cricchi, J.R. and M.H. White; "Complementary MOS Arrays," IEEE International Electron Device Meeting, October 1966, Washington, D.C.

Cricchi, J.R., et al.; IEEE Trans. Aerospace, Electronic Systems, AES-3, 677 (1967).

Cricchi, J.R., Blaha, F.C., and Fitzpatrick, M.D.; "Drain-Source-Protected MNOS Memory Devices and Memory Endrance"; Technical Digest of 1973 IEDM; December 1973; pp. 126-129.

Cricchi, J.R., et. al.; "Nonvolatile Block-Oriented RAM"; Technical Digest of 1974 ISSC; February 1974; pp. 204-215.

Cricchi, J.R.; "MNOS Nonvolatile Semiconductor Storage"; 1974 IEEE International Magnetic Conference; Toronto, Canada.

Cricchi, J.R., D.W. Williams, F.C. Blaha, M.D. Fitzpatrick, and H.G. Oehler; "Nonvolatile MNOS Memory Arrays for BORAM"; GOMAC 1974; Boulder, Colorado.

Cricchi, J.R., F.C. Blaha, M. Fitzpatrick, F. Scuilli; "Space Charge Effects in MNOS Memory Devices and Endurance Measurements"; IEEE International Electron Devices Mtg.; Washington, D.C.; December 1975.

Danielsson, Per-Erik and Bjorn Gudmundsson; "Time-Shared Memory-Processor Interface"; 1975 Sagamore Computer Conference on Parallel Processing; pp. 90-98.

Enslow, Philip H. Jr., editor; Multiprocessors and Parallel Processing; Wiley-Interscience; 1974.

Enslow, Philip H. Jr.; "Multiprocessor Archetecture - A Survey"; 1975 Sagamore Computer Conference on Parallel Processing; pp. 63-70.

Fagan, J.L., M.H. White, and D.R. Lampe; "Digital Nonvolatile Memories from CCD/MNOS Technology," GOMAC-1974.

Fagan, J.L., M.H. White, and D.R. Lampe; "A High Density, Read/Write, Nonvolatile Charge Addressed Memory (NOVCAM)"; IEEE Solid-State Circuits Conf.; Philadelphia, Pa.; February 1976.

Foster, Claxton C.; Content Addressable Parallel Processors; Van Nostrand Reinhold; 1976.

Goldberg, J., Levitt, K.N., and Wensley, J.H.; "An Organization For A Highly Survivable Memory"; IEEE Transactions On Computers, vol. C-23, No. 7; July 1974; pp. 693-705.

Gordon, N. and Johnson, W.C.; "Switching Mechanisim in Thin-Oxide MNOS Devices"; IEEE Transactions On Electron Devices, vol. ED-20, No. 3; March 1973; pp. 253-256.

Goser, K. and Knauer, K.; "Nonvolatile CCD Memory With MNOS Storage Capacitors"; IEEE JSSC, Correspondence, vol. SC-9, No. 3; June 1974; pp. 148-150.

Heimerdinger, W.L., C.F. Muething, S.J. Nusph and L.B. Wing; "Architectural Considerations in Interfacing A Parallel Processor to the Air Traffic Control System"; Proceedings of 1974 Sagamore Computer Conference on Parallel Processing; pp. 372-382.

Hu, S.M.; "Properties of Amorphous Silicon Nitride Films"; Journal of Electrochemical Society, vol. 113; July 1966; pp. 673-688.

Hu, S.M., Kerr, D.R., and Gregory, L.V.; "Exidence of Hole Injection and Trapping in Silicon Nitride Films Prepared by Reactive Sputtering"; Applied Physics, Letters, vol. 10; February 1967; pp. 97-99.

Joynson, R.E., et. al.; "Eliminating Threshold Losses in MOS Circuits by Bootstrapping Using Varactor Coupling"; IEEE JSSC, vol. SC-7, No. 3; June 1972; pp. 217-224.

Kaplinsky, Cecil; "Parallel Processing By Virtual Instruction"; Proceedings of 1974 Sagamore Computer Conference on Parallel Processing; pp. 308-318.

Klaschka, T.F.; "A Method For Redundancy Scheme Performance Assessment"; IEEE Trans. On Comp.; November 1971; pp. 1371-1376.

Kuck, David J.; "Parallel Processor Architecture-A Survey"; 1975 Sagamore Computer Conference on Parallel Processing; pp. 15-39.

Lampe, D.R., M.H. White, J.L. Fagan, and F.J. Kub; "Nonvolatile Charge Addressed Memories," National Electronics Conference; Chicago, Ill.; October 1975.

Lin, H.C., J.L. Halsor, and M.N. Giuliano; "Effect of Silicon Gate Resistance on Frequency Response of MOS Transistors"; IEEE International Electron Devices Meeting; Washington, D.C.; 1974.

Lonky, M.L. and A.P. Turley; "Electrically Variable Channel - Length Changes in MNOS Transistors,"; Applied Physics, Letters; February 1976.

Lundstrom, K.I., and Svensson, C.M.; "Properties of MNOS Structures"; IEEE Transactions On Electron Devices, vol. ED-19, No. 6; June 1972; pp. 826-836.

Mathur, F.P.; "On Reliability Modeling and Analysis of Ultra-Reliable Fault-Tolerant Digital Systems"; IEEE Trans. On Comp.; Nov. 1971; pp. 1376-1382.

NASA, Goddard SFC; "Camera Memory Study For Large Space Telescope"; NAS 5-20512, Final Report; C.P. Hoffman, et. al.; February 1975.

Peck, D.S., and Zierdt, C.H. Jr.; "The Reliability of Semiconductor Devices In The Bell System"; Proceedings Of The IEEE, Special Issue On Reliability Of Semiconductor Devices, Vol. 62, No. 2; February 1974; pp. 185-211.

Reddi, S.S., and E.A. Feustel (Rice University); "An Approach to Restructorable Computer Systems"; Proceedings Of The 1974 Sagamore Computer Conference On Parallel Processing; Springer-Verlag; New York, 1975; pp. 319-337.

Turley, A. P., R. M. McLouski, P. R. Read, and D. H. McCann; "Fabrication of P-Channel Silicon Gate LSI Devices with $Si_3/N_4/SiO_2$ Gate Dielectric"; Journal of the Electrochemical Society; January, 1976.

Weller, Walter J.; Assembly Level Programming For Small Computers; Heath; 1975.

Westinghouse, Staff of Science and Technology; Integrated Electronic Systems; Prentice Hall (1970); New York.

White, M. H. and J. R. Cricchi; "Complementary MOS Transistors"; Solid State Electronics; 9, 991-1008 (1966).

White, M. H., F. C. Blaha, and D. S. Herman "A Study Relating MOS Processes to a Model of the AllSiO_2-Si System," Silicon Device Processing Conference, NBS Symposium; Gaithersburg, Maryland; June 1970 (NBS Special Bulletin 337).

White, M. H.; Semiconductors and Semimetals, Vol. 7a, Semiconductor Applications, Chapter MOS Transistors; Academic Press (1971); New York.

White, M. H. and Cricchi, J. R.; "Characterization of Thin-Oxide MNOS Memory Transistors"; IEEE Transactions On Electron Devices, vol. ED-19, no. 12; December 1972; pp. 1280-1288.

White, M. H., D. R. Lampe, J. L. Fagan, and D. A. Barth; "A Nonvolatile Charge Addressed Memory (NOVCAM) Cell"; IEEE International Electron Devices Meeting; Washington, D. C.; 1974.

White, M. H.; "Charge Transport and Storage in Charge Coupled Devices"; Seminar Department of Electrical Engineering, University of Chincinnati, Cincinnati, Ohio; November 1974.

White, M. H., D. R. Lampe, J. L. Fagan, and D. A. Barth; "A Nonvolatile Charge-Addressed Memory (NOVCAM) Cell"; 1975 Lehigh Valley Semiconductor Symposium, Lehigh University; May 9, 1975.

White, M. H., et. al.; "A Nonvolatile Charge-Addressed Memory (NOVCAM) Cell"; IEEE JSSC, Special Issue On Semiconductor Memory and Logic, vol. SC-10, no. 5; October 1975; pp. 281-287.

Yau, S. S. and H. S. Fung; "Associative Processor Architecture - A Survey"; 1975 Sagamore Computer Conference on Parallel Processing; pp. 1-14.

APPENDIX A

PROGRAM STATEMENT OF WORK

The SOW (statement of work) contained herein is included for the sake of completeness and as a point of reference. It sets forth the original directives prescribed by RADC for the Mass Memory Organization Study program which led to this report.

MASS MEMORY ORGANIZATION STUDY
PR NO. B-5-3107

RESEARCH AND TECHNOLOGY WORK STATEMENT

1.0 Objective: The objective of this effort is a study and investigation for a mass memory storage device that is directly compatible with the speed/bandwidth capabilities of an associative type computer processor. This device when interfaced with an associative processor (AP) can provide the necessary storage of large amounts of data and instructions, and facilitate the inherent higher speed operation of the AP.

2.0 Scope:

2.1 Numerous studies conducted by RADC have clearly demonstrated that effective use of the inherent power of an associative processor is highly dependent on the availability of a mass storage device capable of providing data at a rate matching the I/O rate of the associative array. However, before a mass memory can be specified for operational use in conjunction with an AP, experience and confidence must be gained in a technology suitable for implementing the mass memory. The intent of this effort is not to develop memory technology in general, but to "groom" a specific storage technology for direct application to the AP. The storage technology which best meets the requirements and can be readily available within the next three years will be selected within this effort. A mass memory organization will be designed based on this type memory, data capacity, data and memory speeds, reliability and cost to provide a practical AP storage device. The data acquired under this effort will lead to a mass memory storage device design having optimum data rates and minimal costs for future AP applications.

2.2 Typically, the AP has data channels of at least 1000 processing elements and can readily accept data rates of six billion bits per second. The design will incorporate the necessary computer architecture to maximize this transfer rate.

2.3 This study shall not be constrained by existing or planned AP designs or capabilities.

3.0 Background: Considerable basic research has been performed by commercial organizations in recent years oriented at solid state storage techniques to replace the mechanical disks, drums, and tapes for sequential computers. Since the requirements of a mass memory supporting a parallel processing system are unique (differs from one supporting a conventional sequential computer), a memory organization needs to be developed which encompasses memory type, capacity, cost, and accessing speeds. Work is required to organize the most promising of these storage technologies to provide a reasonably priced, ruggedized, random access, 10^9 bit mass memory capability to associative processing.

4.0 Tasks/Technical Requirements

4.1 The contractor shall accomplish the following:

4.1.1 Conduct a study and investigation for the eventual development of a mass memory storage device which shall be compatible with an associative type processor. As a minimum the contractor shall investigate the following areas:

4.1.1.1 Memory size of 10^9 bits or greater

4.1.1.2 Access of any 1024 bits data slice

4.1.1.3 Access time of less than 150 nanoseconds per bit slice

4.1.1.4 Read/write capability (stored data can either be read-out or erased and replaced by new data).

4.1.1.5 Non-volatile - even in the event of a power failure (non-volatility is meant to mean stored data that can be read-out after a lapse of a long period

of time; i.e., days, after having been written without the need of having been refreshed).

4.1.1.6 Memory type, availability and pertinent technical characteristics

4.1.1.7 Memory organization and structure

4.1.1.8 Low unit cost

4.1.1.9 High reliability

4.1.1.10 Size, weight and power

4.1.1.11 Environmental constraints

4.1.2 Based on the analysis stated above the contractor shall fabricate an Exploratory Development Model (See Note 1). This model shall demonstrate the feasibility of the design and shall be used to acquire pertinent data.

Note 1 - Exploratory Development. An item (preliminary parts or circuits) used for experimentation or tests to investigate or evaluate the feasibility and practicality of a concept, device, circuits, or system in breadboard or rough experimental form, without regard to the eventual overall fit or final form.

5.0 Data Items:

5.1 Research and Development Status Reports shall be prepared and submitted monthly to RADC/ISCA (2 copies) and RADC/PMRD (1 copy). Nine (9) reports shall be provided. The final (ninth) contract status report shall contain a duly executed ASPR Certificate of Current Cost or Pricing Data covering the cumulative total of hours expended in support of the term effort.

5.2 A Final Technical Report describing the work done, results obtained, and recommendations made shall be prepared and submitted to RADC (TIR) within 30 days after completion of all tasks in the term effort. This report shall be prepared in accordance with MIL-STD-847A, dated 31 Jan 1973. One set of reproducibles typed on Government furnished, "Format Sheet for Preparation of Technical Reports", and three (3) bond or high quality zerox copies shall be provided. All work shall be in final form with reproducibles suitable for direct recording on microfiche.

5.2.1 Technical Report Summary: As a separate document, the Technical Report will include a report summary specifying the purpose of the project and giving a description of any important equipment purchased or developed and the conclusions reached by the contractor. This summary, prominently identified, should normally not exceed a few pages.

APPENDIX B

STARAN DESCRIPTION

The material presented here contains the general information necessary to understand the overall associative array processor STARAN* system, the design tradeoffs plausible in an interrelated mass memory, and the manner in which the system interface is configured. All information presented here was selectively excerpted from Air Force supplied document number GER-15640, titled "STARAN Mainframe Maintenance Manual," September 1973; originally delivered under contract F30602-73-C-0134.

*TM, Goodyear Aerospace Corporation, Akron, Ohio.

STARAN SYSTEM CONFIGURATION

Many options and expansion capabilities are available with the STARAN system. Because of this, the configuration of systems may vary. Discussed here are the elements which comprise a minimal or basic STARAN system and some of the options available.

Basic STARAN System

A basic STARAN system contains, as a minimum, the following major assemblies:

1) Controller Cabinet. This cabinet contains the Sequential Processor (SP), the SP Memory, and the Bulk Core Memory Unit.
2) Control Logic Cabinet. This cabinet contains all Associative Processor (AP) control logic, the page memories (page 0, page 1, and page 2), and the High Speed Data Buffer (HSDB) memory.
3) Memory Cabinet. This cabinet contains at least one single-port associative array.
4) Custom Input/Output Unit (CIOU) Cabinet. This cabinet contains the input/output logic required to interface the STARAN system with its external operational environment.
5) Peripheral Device: These devices facilitate operational control over the STARAN system.

- Associative Processor STARAN Options Control Memory Expansion

In the basic STARAN system configuration, AP control memory consists of the following memory blocks:

1) Bulk core memory (16, 384 words)
2) Page memories (1536 words)
3) High Speed Data Buffer (HSDB) memory (512 words)

The AP control memory in a basic STARAN system contains a total of 18,432 words, each word containing 32 bits. However, provision is made in the basic STARAN system to double the size of any or all of the above memory blocks. Therefore, a fully expanded AP control memory may contain as many as 36,864 words if required.

In addition to the above AP control memory blocks, there is a block of AP control memory addresses reserved for the Direct Memory Access (DMA) channel to an external memory. In the basic STARAN system configuration this block contains 30,720 words. If the page memory block or HSDB memory block are expanded, however, the available DMA memory block is reduced accordingly.

- Sequential Processor Memory Expansion

In the basic STARAN system configuration, Sequential Processor (SP) memory contains 16K words of 16 bits each.

- Array Expansion

The basic STARAN system must contain at least one associative array. However, provision is made in the AP control logic to accommodate as many as 32 associative arrays. Arrays may be added as required, with no increase or change to AP control logic. The standard STARAN memory cabinet will accommodate up to three associative arrays, so that a fully expanded system would contain 11 memory cabinets.

The associative array(s) in the basic STARAN system are single-port. However, on an optional basis,

dual-port associative array(s) may be utilized. The dual-port associative array facilitates the transfer of data, 256 bits per array at a time, over the Parallel Input/Output (PIO) channel.

- Input/Output Options

The basic STARAN system has a variety of Input/Output (I/O) options available. The Custom Input/Output Unit (CIOU) cabinet containing buffered and/or unbuffered I/O channels to data gathering, data receiving, and data-storing devices is part of the basic STARAN system. STARAN can also be integrated with a variety of other computer systems. The Direct Memory Access (DMA) channel to a host-computer memory enables a rapid interchange of data between the systems in a common memory bank. A Buffered Input/Output (BIO) channel provides an alternate means of exchanging data, while an External Function (EXF) channel permits interrupts and/or other control information to be passed between the two systems.

The optional Parallel Input/Output (PIO) channel, with a width of up to 256 bits per array, can also be implemented in STARAN. The extreme width of this channel (up to 8,192 bits), plus its submicrosecond cycle time, gives STARAN an I/O bandwidth many times wider than that of a conventional computer. The PIO channel can easily accommodate the high data rates that arise in many real-time applications. Also, the PIO channel makes it possible for STARAN to interface with special high-bandwidth mass-storage devices,

permitting rapid retrieval, restructuring, and processing of data in a large data base.

- Peripheral Options

Typically, a basic STARAN system would include a paper tape reader/punch mounted in the controller cabinet and a stand-alone DECwriter. The sequential processor UNIBUS and the I/O options available in the basic STARAN system permit a wide variety of peripheral devices to be utilized. Depending on the user's application, any or all of the following peripherals may be integrated into a STARAN system:

1) DECwriter*
2) Paper Tape Reader/Punch
3) Card Reader
4) Line Printer
5) Alphanumeric CRT Display
6) Removable Cartridge Disk Unit
7) Magnetic Tape Unit

- Disk Operating System

Another peripheral system which may be integrated into the basic STARAN is the Disk Operating System (DOS). The DOS is extremely useful in that is provides a complete random access mass storage media for the STARAN system. Large volumes of data (approximately 1.2 million 16-bit words) may be transferred from a disk to SP memory and/or AP control memory and vice-versa. The disk itself provides a convenient and compact means of program storage. The DOS is supported by a complete software package

*TM, Digital Equipment Corporation, Maynard, Mass.

whereby operation of the disk drive can be controlled from a DECwriter or alphanumeric CRT display terminal.

STARAN SYSTEM FUNCTIONAL DESCRIPTION

The STARAN system consists of the following basic elements (see figure B-1):

1) Sequential control
2) External function logic
3) Program pager
4) AP control memory
5) Associative processor control
6) Associative arrays

Sequential Control

Sequential control in the STARAN system consists of a sequential processor (SP) with associated memory, interface logic to connect the SP to other STARAN elements, and peripheral units which interface via the SP UNIBUS. Sequential control provides the following functions:

1) A means to initially load the AP memory
2) A communication link between the operator and STARAN for on-line control and monitoring
3) Capabilities for assembling and debugging STARAN programs
4) Control for STARAN maintenance and diagnostic program routines

The sequential processor is a 16-bit, general purpose, parallel logic minicomputer using two's complement arithmetic. In a basic STARAN system, the 16K 16-bit words (32K 8-bit bytes) have SP octal byte

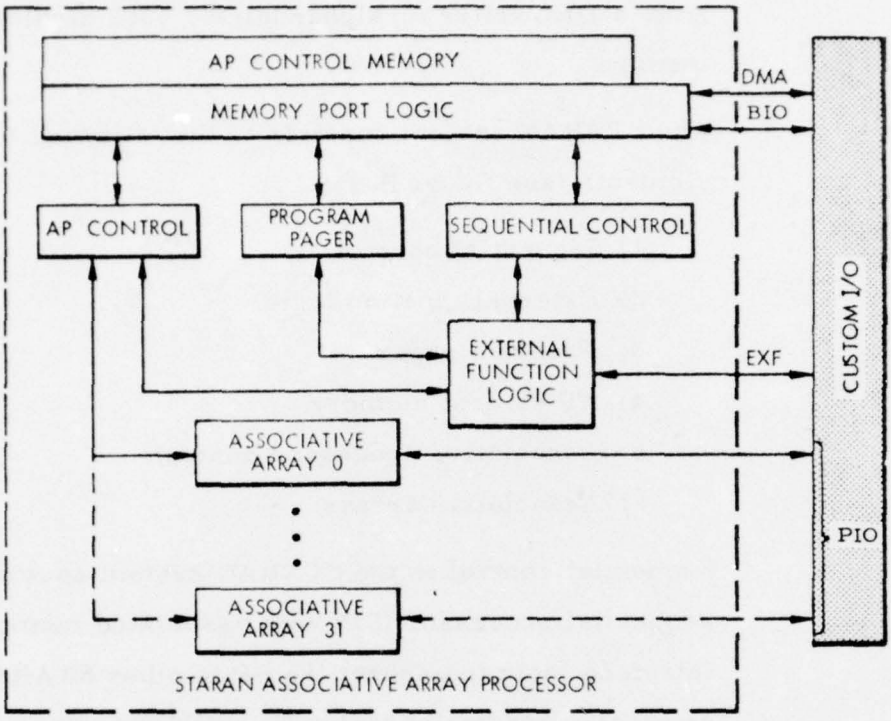

Figure B-1. STARAN Basic Block Diagram

Addresses 000000 through 077777. All communication between system components is performed on a single high-speed bus called the UNIBUS. Also included is a multilevel automatic priority interrupt system.

- Sequential Processor Interface

Communication between the sequential control and the associative processor takes place over various buses in AP control and the Sequential Processor UNIBUS. All AP buses which can transfer data to or from the UNIBUS are connected to the "sequential processor data switch". This logic module controls the flow of data between the AP buses and the UNIBUS. Four forms of communication between sequential control and the associative processor are provided:

1) Direct Access To AP Control Memory: Words in the AP control memory are given sequential processor addresses to facilitate transfer of data and instructions between AP control and sequential control.

2) Register Readout: Certain registers in STARAN can be read by sequential control. This facilitates program debugging and hardware maintenance and test.

3) External Functions: External function codes can be transmitted to the external function logic and sense bits received. This allows sequential control to activate and deactivate AP control and issue interrupts.

4) Interrupt Acceptance: Some elements of STARAN can issue interrupts to sequential control by issuing certain function codes to the external function logic. Also, when errors such as parity are detected, a sequential control interrupt is generated.

External Function Logic

Numerous hardware functions are under the control of the External Function (EXF) logic. These include page memory port switches, AP and sequential control interrupts, AP control and program pager activity control, resets and clears. Control and status sensing of these functions are accomplished by issuing 19-bit external function commands to the EXF logic and receiving one-bit sense signals in return. Three elements of the mainframe can issue EXF commands: AP control, program pager, and sequential control.

The EXF logic is expandable to allow receipt of EXF commands from the Custom Input/Output Unit (CIOU) and control of other hardware functions in the CIOU. A resolver in the EXF logic allows only one EXF command to be treated at a time.

The resolver in the EXF logic resolves conflicts among the four elements issuing function codes. One function code at a time is accepted by the EXF logic. The interrogation and/or control specified by a function code is performed and then another function code is accepted if one is present. A function code can interrogate and control an element in one operation without interference from another function code.

The classes of function codes are as follows: page memory port switches, interlocks, program pager, error control, AP control interrupts, sequential control interrupt, miscellaneous, and spare functions that may be used by the custom input/output unit (CIOU).

Program Pager

The function of the program pager is to transfer program segments from the bulk core memory to the page memories.

Under normal programming practice, the pager is activated by AP control when a new program segment is to be transferred to a page memory. The program pager transfers the segment one word at a time from source memory to a page memory while AP control executes instructions from previously loaded segments out of another page memory. When the pager completes

the transfer, it restores the page memory port switch to the AP instruction bus and halts.

AP Control Memory

The main function of the AP control memory is to store the assembled AP application programs. AP control memory can also be used for data storage and to act as a buffer between AP control and other elements of STARAN. Since the AP control memory is not an integral part of the associative arrays, AP control can overlap the AP control memory cycle time with the associative array cycle time.

AP control memory is divided into several memory blocks. One small, fast-memory, called the page memory, contains the current (active) AP program segments. Another high speed memory, the High Speed Data Buffer (HSDB), may contain frequently used data. The larger, slower bulk core memory contains the remainder of the AP program. A program pager is included in STARAN to facilitate transfer from the bulk core to the page memory.

Each word of AP control memory contains 32 bits of either data or instructions. Bit 0 is the left (most-significant) bit, and bit 31 is the right (least-significant) bit of each word. Each word is given a 16-bit address expressed in hexadecimal notation.

Associative Processor Control

The major function of AP control is to control the STARAN associative arrays. AP control fetches instructions from the AP control memory. A 16-bit program counter contains the address of the instruction, while a 32-bit instruction register contains the

instruction itself. Some instructions perform array operations, while others perform AP control functions. Internal control registers affect operations and other elements of the STARAN system.

Associative Arrays

The associative arrays are the heart of the STARAN system. There may be a maximum of 32 arrays in one system. Each associative array contains a multi-dimensional-access (MDA) memory with 65,536 (2^{16}) bits of storage. The 256 response-store elements in each associative array allow array data to be searched, restructured, and processed at a fast rate. The array memory is arranged in 256 words by 256 bits, with loading and storing permitted in either the vertical or horizontal direction and with a maximum of 256 bits per array transferred in a single operation. The multidimensional access feature allows data to be accessed in either dimension by simply changing a bit in the associative instruction format.

Sequential Control/ AP Control Memory

Sequential control can read and write into AP control memory. AP control memory is divided into groups. Each group contains 4096 32-bit words (16,384 bytes).

When data is accessed in the HSDB memory block, bulk core, or through DMA interface, data flow is from AP control memory, through the port priority switch logic to the 32-bit sequential processor data bus; through the sequential processor data switch to the sequential processor interface, and finally to the UNIBUS (or in reverse order if data is being written in AP control memory).

When data is transferred from the page memory block to the sequential processor, it is transferred from either page 0, page 1, or page 2 through the associated page memory port switch to the sequential processor data bus; through the sequential processor data switch to the sequential processor interface, and finally to the UNIBUS.

AP Control Memory to Page Memory

Data may be transferred from the HSDB, bulk core, or DMA interface portions of AP control memory to any one of the three page memories via the 32-bit program pager bus. This transfer is carried out under control of the program pager independent from AP control operation. Data flow for this type of transfer is from the HSDB, bulk core, or DMA interface portions of AP control memory through the port priority logic to the pager get bus; from the pager get bus through the pager put bus to the associated page memory.

AP Control/ AP Control Memory

Various data transfers may take place between AP control memory (HSDB, bulk core; and DMA interface) and various control registers in AP control. During a "store" operation (data transfer from a control register to AP control memory), data flow is from the affected control register through the bus shift logic to the 32-bit AP control data bus; from the AP control data bus through the port priority switch to AP control memory. During a "load" operation (data transfer from AP control memory to a control register), the above data flow is reversed.

In addition to the data transfer described above, the 32-bit instruction register in AP control may be loaded from any part of AP control memory including the page memory block. Data flow for this operation is from AP control memory through the port priority switch logic (or page memory port switch if data is being transferred from the page memory block) to the 32-bit AP control instruction bus, and from the AP control instruction bus to the instruction register.

AP Control/ EXF Logic

External function commands may be issued by AP control. Some of the 32-bit instruction words which are loaded into the instruction register via the AP control instruction bus and AP control memory contain 19-bit external function command fields. This external function command is transmitted to the external function logic, which then decodes the command and effects the proper AP action.

AP Control/ Associative Array

Various 32-bit data transfers may take place between AP control and the associative arrays via the 32-bit common register. The common register is one of the AP control registers and may be loaded from an associative array or from AP control memory or another AP control register through the bus shift logic. During a "load command register" instruction (data transfer from associative array to the common register) data is transferred from one of the following sources as specified by the associative instruction:

1) A bit slice containing one bit from each of the 32 words in an array memory field

2) A word slice containing 32 bits in one array word of an array memory field.
3) The contents of a 32-bit field of the M Register (MASK)
4) The contents of a 32-bit field of the 256-bit X Register
5) The contents of a 32-bit field of the 256-bit Y Register

A 256-bit word is transferred from its source, through the array flip network (which may flip, mirror, or shift it in various ways), through the array group fan-in/fan-out logic, through the main group fan-in/fan-out logic to the common register. The fan-in/fan-out logic modules transmit only the most significant 32-bit positions (after permutation by the flip network) to the common register. The flip network is controlled by the associative instruction effecting the transfer. Through utilization of the flip network, any 32-bit field in the 256-bit data word may be transferred to the common register.

Data flow is reversed during a "store common register" instruction (data transfer from the common register to an associative array). In this case, the fan-in/fan-out logic modules apply the contents of the 32-bit common register to bit positions 0 through 31 of the flip network input, and pad bit positions 32 through 255 with zeroes. The destination of the resulting 256-bit word may be one of the following depending on instruction format:

1) A bit slice containing one bit from each of the 256 words in array memory
2) A word slice containing all 256 bits in one array word
3) The M Register (MASK)
4) The Y Register, which may combine it logically with the current contents of Y register
5) The X Register, which may combine it logically with the current contents of the X register or with the contents of the Y register to modify the contents of the X register

STARAN SYSTEM I/O INTERFACE

Interface to the STARAN system is through one or more of the following I/O channels:

1) Direct Memory Access (DMA). This I/O channel can provide the STARAN system access to an external memory.
2) Buffered Input/Output (BIO). This I/O channel can provide an external processor with access to AP control memory.
3) Parallel Input/Output (PIO). This I/O channel permits direct access to the associative array(s) in the STARAN system.
4) External Function (EXF). The EXF interface permits the STARAN system to exercise control over an external processor, or vice-versa.

Figure B-2 shows the data and control functions which comprise each of the above I/O channels. The origination of data and control functions for each I/O channel is also indicated in figure B-2.

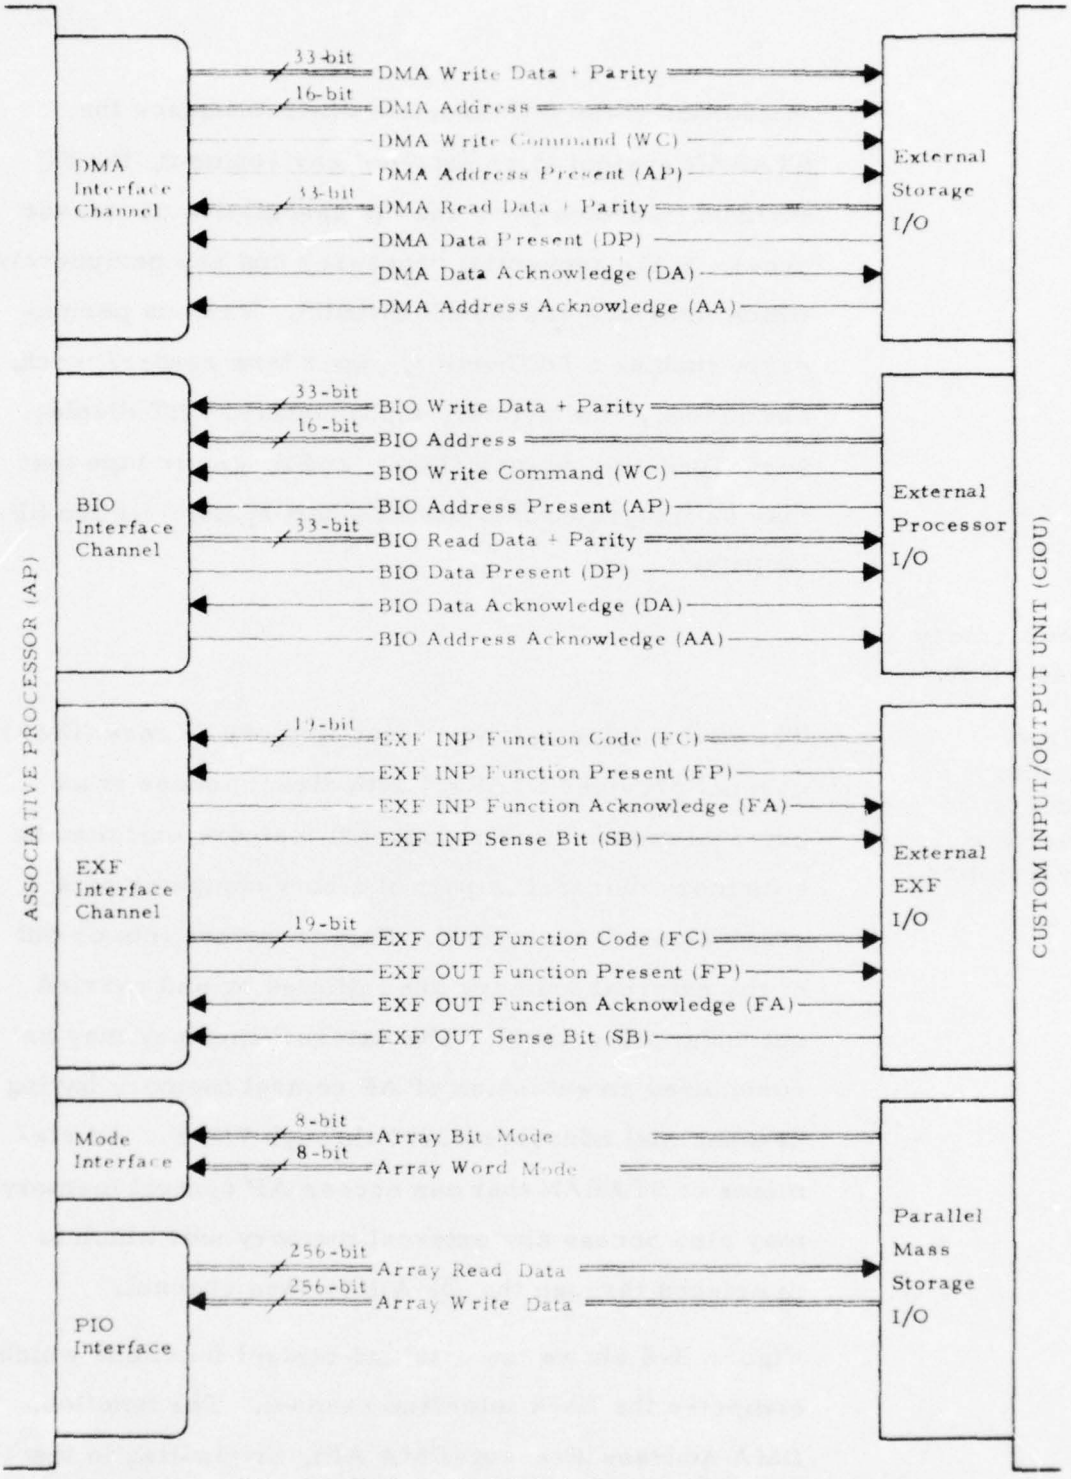

Figure B-2. STARAN System Interface

In addition these I/O channels which interface the STARAN system to an external environment, the SP UNIBUS interface provides the associative processor access to the sequential processor and any peripherals which interface via the SP UNIBUS. Various peripherals such as a DECwriter, paper tape reader/punch, and reader, line printer, alphanumeric CRT display, Disk Operating System (DOS), and magnetic tape unit may be integrated into the STARAN system via the SP UNIBUS.

Direct Memory Access Channel

- Direct Memory Access Interface Description

The 32-bit (plus parity) Direct Memory Access (DMA) channel provides STARAN with direct access to an external memory. The external memory unit may be a memory unit that is part of a host computer or a stand-alone memory unit. Data transfers into or out of the external memory are initiated by and carried out under AP control. The external memory may be considered an extension of AP control memory having hexadecimal addresses 0800 through 7FFF. All elements of STARAN that can access AP control memory may also access any external memory unit which is interfaced through the DMA interface channel.

Figure B-2 shows the data and control functions which comprise the DMA interface channel. The function, DMA Address Present (DMA AP), originating in the associative processor initiates all data transfers; the function, DMA Address Acknowledge (DMA AA),

originating in the Custom Input/Output Unit (CIOU) terminates all data transfers.

- DMA Write Timing

The DMA "write" operation refers to the transfer of data from the associative processor over the DMA interface to some external storage device (e.g. memory). Figure B-3 shows the timing relationship between data and control functions during a DMA "write" operation. The sequence of events for the DMA "write" operation is as follows:

1) WRITE DATA, ADDRESS, and Write Command (WC) raised by associative processor.
2) Address Present (AP) set by associative processor 25 ns (minimum) after (1).
3) Address Acknowledge (AA) set by external memory I/O when data has been stored.
   a) AP reset by associative processor upon detection of AA; WRITE DATA, ADDRESS, and WC lowered.
   b) AA reset by external memory I/O when AP reset detected.

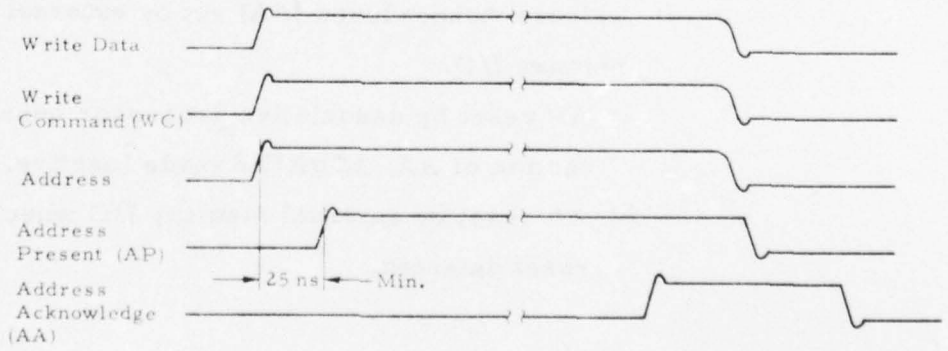

Figure B-3. Timing Relationships, DMA or BIO Write Operation

- DMA Read Timing

The DMA "read" operation refers to the transfer of data from an external storage device (e.g. memory) over the DMA interface to the associative processor. Figure B-4 shows the timing relationship between data and control functions during a DMA "read" operation. The sequence of events for the DMA "read" operation is as follows:

1) ADDRESS raised by associative processor.
2) Address Present (AP) set by associative processor 25 ns (minimum) after (1).
3) READ DATA raised by external memory I/O when read request honored.
4) Data Present (DP) set by external memory I/O when READ DATA has stabilized.
5) Data Acknowledge (DA) set by associative processor when data has been read.
   a) DP reset by external memory I/O upon detection of DA; READ DATA made inactive.
   b) DA reset by associative processor when DP reset detected.
6) Address Acknowledge (AA) set by external memory I/O.
   a) AP reset by associative processor upon detection of AA; ADDRESS made inactive.
   b) AA reset by external memory I/O when AP reset detected.

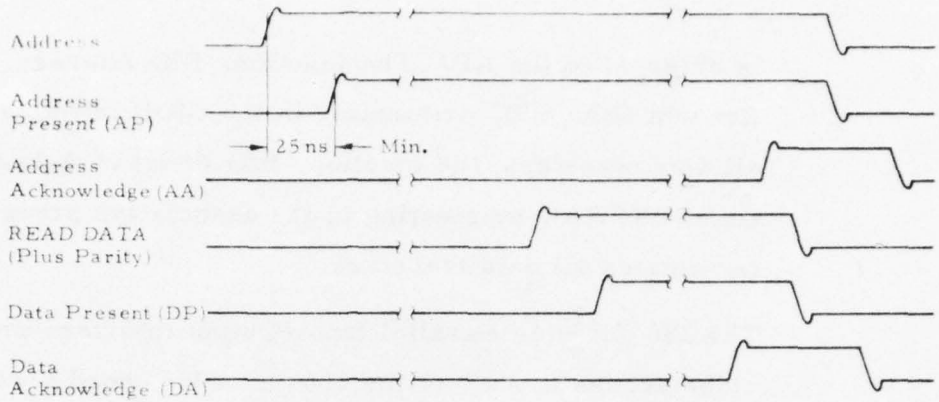

Figure B-4. Timing Relationships, DMA or BIO Read Operation

Buffered Input/ Output Channel

- Buffered Input/Output Interface Description

The 32-bit (plus parity) Buffered Input/Output (BIO) interface channel provides an external processor access to certain portions of AP control memory (viz. the HSDB and bulk-core) and any external memory interfaced through the DMA interface channel. Data transfers into or out of AP control memory are initiated and carried out under control of the external processor. Through use of the BIO interface, an external processor and the AP can share portions of AP control memory and any external memory interfaced through the DMA channel. Figure B-2 shows the data and control functions which comprise the BIO interface channel. These functions are identical to those which comprise the DMA interface channel; however, the origin of each BIO function is opposite to that shown for the DMA interface since control of data transfer

is external to the AP. The function, BIO Address Present (BIO AP), originating in the CIOU initiates all data transfers; the function, BIO Address Acknowledge (BIO AA), originating in the associative processor terminates all data transfers.

PIO Interface Channel

The 256-bit wide Parallel Input/Output interface provides access to each of the 256-bit x 265-word associative arrays used in a particular STARAN system. Each PIO channel connects the CIOU to one associative array through a 256-bit port. Addressing of the arrays may be carried out in the word mode or in the bit-slice mode. The array(s) which are to participate are enabled, and the array address is set up prior to initiation of each data transfer.

The data functions which comprise the PIO interface for an associative array are shown in figure B-2. Addressing and control of the array during a PIO data transfer is accomplished via an array control port by the AP mainframe or the CIOU.

- UNIBUS Interface

The Sequential Processor (SP) UNIBUS provides the interface required to the associative processor and to the peripheral units which utilize it. The UNIBUS, its function, and the timing relationships between UNIBUS signals are described herein. Interface between the SP UNIBUS and the associative processor is accomplished by the sequential processor interface and the sequential processor data switch logic modules located in the control logic cabinet.

The sequential processor interface functions as a "device controller" for the associative processor. This interface generates control signals required for communication on the UNIBUS.

The sequential processor data switch controls the transfer of data from various elements of the associative processor (viz, AP control memory, certain AP control registers, and the AP instruction register) to the sequential processor UNIBUS in the form of 16-bit data words. It also controls the transfer of 16-bit data items from elements of sequential control to either the 32-bit AP control memory (HSDB, bulk core, or DMA interface) or to EXF logic in the form of 19-bit EXF function commands.

- UNIBUS Function

All communication between the sequential processor, associative processor, and certain peripherals is accomplished on a single high speed bus called the UNIBUS. The UNIBUS is a single common path that connects sequential processor, SP memory, associative processor and any peripheral device which utilizes it. The form of communication is the same for every device on the UNIBUS.

Communication between two devices on the UNIBUS is in the form of a "master-slave" relationship. During any bus operation, one device has control of the UNIBUS. This device, the bus master, controls communication between itself and another device on the UNIBUS called the slave.

A priority structure determines which device obtains control of the UNIBUS. Every device on the UNIBUS capable of becoming bus master has an assigned priority. Each control signal issued by the master device must be acknowledged by a response from the slave to complete a transfer.

The request and granting of bus mastership is performed in parallel with data transfers on a completely independent set of bus lines. Thus, while one device is using the bus, the next request is being checked for priority and the next user is being assigned.

When a device other than the sequential processor gains control of the bus, it uses the bus to perform either a data transfer or an interrupt request as described in the following paragraphs.

- Data Transfer

Direct memory or device access data transfers can be accomplished between any two devices without sequential processor supervision. These are called NPR level data transfers. Normally, NPR transfers are made between SP memory and the associative processor.

During NPR transfers, it is not necessary for the sequential processor to transfer the informaton between the SP memory and the associative processor. The bus structure enables device-to-device transfers, thereby allowing controllers to directly access other devices on the bus.

An NPR device provides extremely fast access to the bus and can transfer data at high rates once it gains control. The sequential processor state is not affected by this type of transfer; therefore, the sequential processor can relinquish bus control while an SP instruction is in progress. This release of the bus can normally occur at the end of any bus cycle. An NPR device in control of the bus can transfer 16-bit words or 8-bit bytes to memory at the same speed as the memory cycle time.

- **Interrupt Requests**

Devices that gain bus control with one of the bus request lines can request an interrupt. The entire SP instruction set is then available for manipulating data and status registers. When a device is to be run, the task being performed by the sequential processor is interrupted, and the sequential device service routine is initiated. After the device request has been satisfied, the sequential processor returns to its former task. Note that interrupt requests can be made only if bus control has been gained through a BR priority level. An NPR level request cannot be used for an interrupt request.

- **UNIBUS Description**

The SP UNIBUS consists of 56 signal lines. All devices on the SP UNIBUS including the associative processor are connected to these lines in parallel. All signals on the SP UNIBUS are asserted when at the logical "0" level except for the bus grant lines.

- Data Transfer Lines

Forty bidirectional bus lines are used for data transfer. In a data transfer, one device on the SP UNIBUS is "bus master" and controls the transfer of data to or from another device on the SP UNIBUS called the 'slave'. The forty data transfer lines consist of the following:

1) Data (D00 through D15). The 16 data lines are used to transfer 16-bit SP data words between master and slave.
2) Address (A00 through D17). The 18 address lines are used by the master to select a unique SP memory location or the device register of a slave with which it will communicate.
3) Control Signals. UNIBUS control signals are divided into the following three groups:
   a) Control Lines (C1 and C0). These two bus signals are coded by the master device to control the slave in one of four possible data transfer operations.
   b) Master and Slave Synchronization. Master synchronization (MSYN) is a control signal used by the master to indicate to the slave that address and control information is present. Slave synchronization (SSYN) is the slave's response to the master (usually a response to MSYN).
   c) Parity Bit Low (PA) and Parity Bit High (BP). These signals are for devices on the UNIBUS that use parity checks. PB is the parity of the high-order byte (D08 through D-15), and

PA is the parity of the low-order byte (D00 through D07).

- **Priority Transfer Lines**

The UNIBUS contains 13 priority transfer lines. These signals are used to request and grant control of the UNIBUS and include the following:

1) Bus Request Lines (BR4 through BR7). These four bus signals are used by peripheral devices to request control of the UNIBUS.
2) Bus Grant Lines (BG4 through BG7). These signals are the sequential processor's response to a bus request. They are asserted only at the end of instruction execution, and in accordance with the priority determination.
3) Non-Processor Request (NPR). This signal is a bus request from a peripheral device to the sequential processor.
4) Non-Processor Grant (NPG). This signal is the sequential processor's response to an NPR. It occurs at the end of a bus cycle.
5) Selection Acknowledge (SACK). SACK is asserted by a bus-requesting device that has received a bus grant. Bus control passes to this device when the current bus master completes its operation. (If SACK is not received by the sequential processor within 10 μs of issuing a bus grant, timeout occurs and the bus grant is cleared automatically by the processor.)

6) Interrupt (INTR). This signal is asserted by the bus master to start a program interrupt in the sequential processor.

7) Bus Busy (BBSY). This signal is asserted by the master devices to indicate bus is being used.

• • Miscellaneous Control Lines

There are three additional lines on the UNIBUS which may be used by all devices. These control lines and their mnemonics are:

1) Initialization (INIT).
2) AC Line Low (AC LO).
3) DC Line Low (DC LO).

STARAN SYSTEM PERIPHERALS

Various optional peripheral equipment may be integrated into the STARAN system. A detailed functional description of each peripheral utilized in a particular STARAN system will be contained in the applicable commercial manuals supplied with that system. Optional peripherals available for use in the STARAN system include the following:

1) DECwriter
2) Paper Tape Reader/Punch
3) Card Reader
4) Line Printer
5) Alphanumeric CRT Display
6) Disk-Operating System (DOS)
7) Parallel Head Disk (PHD)
8) Magnetic Tape System

Only the last three are reviewed here.

Disk Operating System

The Disk Operating System (DOS) typically consists of one or more disk drive units which interface to the sequential processor UNIBUS through a disk drive controller. One disk drive controller could typically provide the necessary control and interface for eight (maximum) disk drive units. Four disk drive units, the disk drive controller, and necessary power supplies can be mounted in one stand-alone cabinet. An expanded system containing five to eight disk drive units would be contained in two stand-alone cablinets. The DOS option provides a complete random access mass storage system. Large data blocks such as system diagnostics and assembler programs can be stored on disk, accessed, and transferred to AP control memory.

A typical disk drive unit would use a disk similar to the IBM 2315 disk cartridge as its storage medium. The disk cartridge may contain approximately 200 accessible cylinders (tracks) on each side of the disk. Each of the 400 accessible tracks is divided into twelve 30 degree sectors, and each sector can store 256 words of 16 bits each.

A single disk drive unit may store as many as 1.2 million 16-bit words. Access time to the DOS is approximately 11.2 microseconds once a data block transfer has been initiated. Two movable heads, which "fly" above and below the rotating disk surface, are used by the disk drive unit to read or record on one of the 400 accessible data tracks.

The disk drive controller functions as the interface between the disk drive unit(s) and the sequential processor UNIBUS. The controller must convert disk drive serial data to parallel data for the UNIBUS, and vice-versa. The controller can typically transfer 2^{16} (maximum) consecutive words without reinitiation or sequential processor intervention.

Parallel Head Disk

The Parallel-Head-Disk (PHD) option provides direct access to the associative arrays in a STARAN system. Because data is transferred to or from the PHD 256 bits at a time, extremely high (512 million bits/sec) data transfer rates may be realized. This peripheral option may be used in high speed data processing systems or in digital video systems where high speed storage and/or retrieval of a large data base is required. The PHD system would consist of the PHD drive and its associated servo unit, controller, and associated power supplies. The entire system can be contained in one stand-alone cabinet. The PHD option must be interfaced through a 256-bit read/write port in the CIOU for a particular STARAN system. The 256-bit port would consist of separate "read" (data transfer from PHD to associative arrays) and "write" (data transfer from associative arrays to PHD) channels. The PHD interface also provides a "masked write" capability wherein only unmasked bits of a 256-bit word are written on the disk.

The storage unit for the PHD system would typically be a 12 inch 256-track disk similar to the 5250 disk

memory series made by Data Disk Inc. Word transfer rate for this system would be approximately 2 million words (256 bits/word) per second. Storage capacity of the disk would be 64,000 words. This disk itself contains 256 tracks with each track containing approximately 64,000 bits. Data is transferred one "sector" at a time. The PHD system would allow the number of sectors per disk to be selectable from the control panel. The sectors/disk ratio may be approximately 250, 500, 1000, 2,000 or 4,000 resulting in words/sector of 256, 128, 64, 32 or 16 respectively.

Magnetic Tape System

A magnetic tape system provides the means to store, access, and retrieve large data blocks for use in the STARAN system. An expanded system may include one or more magnetic tape transports and one control unit which interfaces the tape transports to the sequential processor UNIBUS. One control unit may provide control and interface for up to eight transports. The magnetic tape system is ideally suited for writing, reading and storing large volumes of data (6 million 16-bit words per transport). Typically, the tape transport utilizes two 10-1/2-inch reels to move 1/2-inch mylarbase tape, which is coated on one side with an iron oxide composition, past a read/write head. The magnetic tape may contain up to 9 tracks, eight of which are used for data transfer, and one for parity check. A maximum of 800 8-bit characters may be stored on one inch of the magnetic tape. The 2400 feet (maximum) of magnetic tape are organized into

approximately 1500 "record blocks" with each record block containing from 2 words to 4, 096 words (16 bits/ word).

Rewind speed for the tape transport is approximately 12-1/2 feet per second, maximum. The data transfer rate for the system may be as high as 18, 000 words/ second once a transfer has been initiated. This transfer rate is realized by utilizing the Non-Processor Request (NPR) method of data transfer to sequential control memory locations and vice-versa.

APPENDIX C

RELIABILITY DOCUMENTATION

The documentation presented here consists of various tables, programs, modeling block diagrams, and other related background or reference material on which the reliability analysis of the mass memory has been based, or which aid in the various mathematical calculations.

AD-A033 715 WESTINGHOUSE DEFENSE AND ELECTRONIC SYSTEMS CENTER B--ETC F/8 9/2
MASS MEMORY ORGANIZATION STUDY.(U)
SEP 76 D L FARNSWORTH, C P HOFFMAN, J J SHUTT F30602-75-C-0317
UNCLASSIFIED NYD-56237-A002 RADC-TR-76-254 NL

3 OF 3
AD
A033715

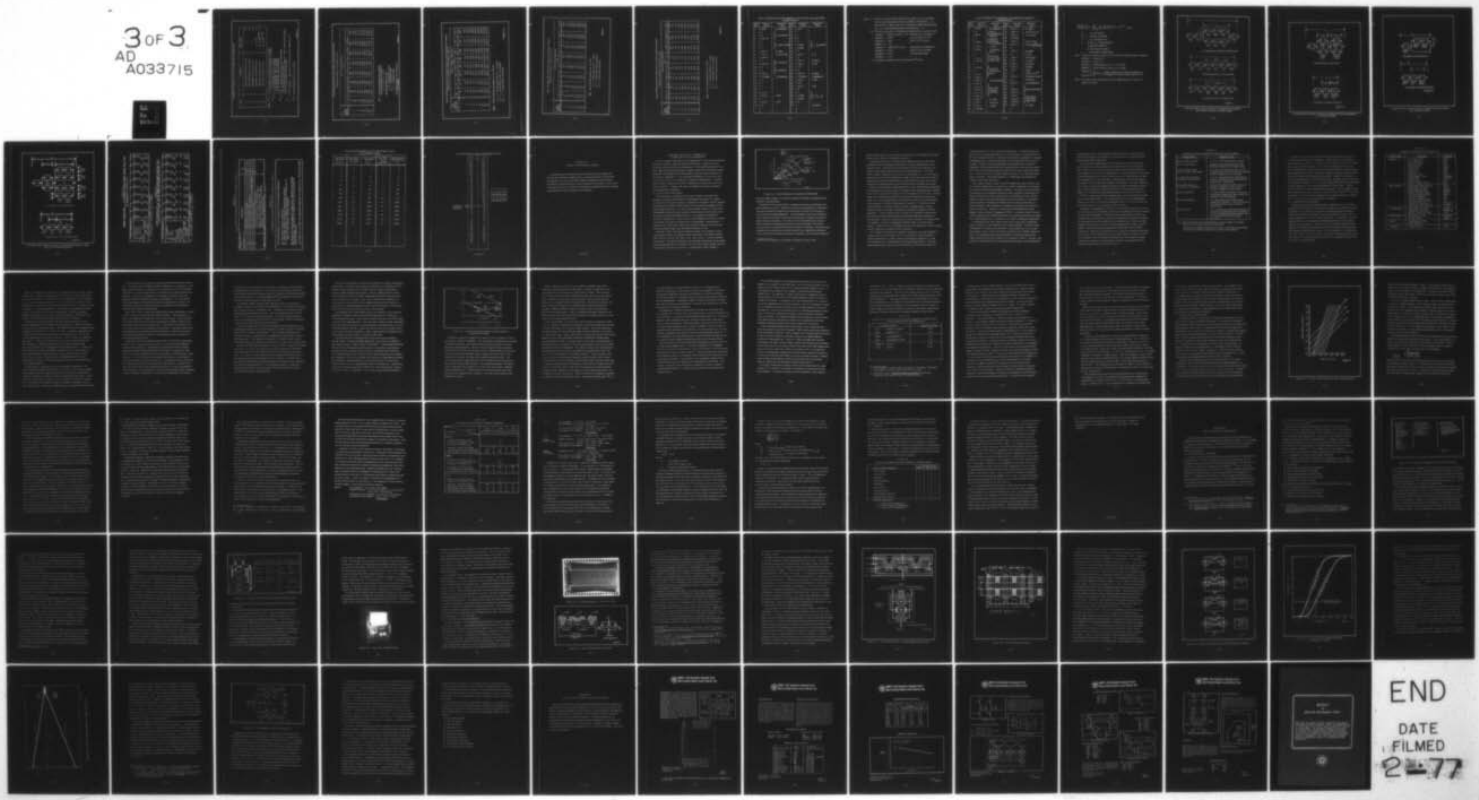

END
DATE
FILMED
2-77

MIL-HDBK-217B MOS Memory Chip Active Failure Rate Modeling Terms

Term	Definition	Relationship
π_L	device production maturity factor	* (1)
π_Q	device qualification level factor	* (2)
π_E	environment acceleration factor	* (3)
π_T	temperature acceleration factor	$\pi_T = 0.1e^x$ (4)
C_T	temperature complexity factor	$C_T = 1.99B^{(0.603)}$ (5)
C_E	environment complexity factor	$C_E = 0.56B^{(0.644)}$

*notes:

(1) $\pi_L = \begin{cases} 10; & \text{for device in continuous production} \leq 6 \text{ months} \\ 1; & \text{for device in continuous production} > 6 \text{ months} \end{cases}$

(2) $\pi_Q = \begin{cases} 1; & \text{for Mil-M-38510, class A screening} \\ 2; & \text{for Mil-M-38510, class B screening} \\ 5; & \text{for Mil-Std-883, class B screening} \\ 10; & \text{for Com-Eqv-883, class B screening} \end{cases}$

(3) $\pi_E = \begin{cases} 0.2; & \text{for ground benign \& space flight environments} \\ 1.0; & \text{for ground fixed environment} \\ 4.0; & \text{for ground mobile, airborne inhabited, \& naval sheltered environments} \\ 5.0; & \text{for naval unsheltered environment} \end{cases}$

(4) $x = -8121 \left(\frac{1}{T_j + 273} - \frac{1}{298}\right)$; where T_j = device junction temperature in ^oC

(5) B = device storage capacity in bits/chip

76-0649-V-12

New MOS Memory Chip Active Failure Rates in Ground Benign Environment

(production screening to Mil-Std-883, class B)

Chip Capacity (Kb/chip)	Device Junction Temperature (°C)							
	25	40	50	60	70	80	100	125
1	1136	2886	5845	11888	23701	45886	156361	612436
2	1747	4405	8899	18078	36020	69716	237514	930233
4	2687	6724	13550	27491	54743	105923	360786	1412937
8	4134	10265	20633	41808	83200	160936	548040	2146123
16	6360	15672	31420	63582	126451	244522	832484	3259766
32	9786	23931	47850	96699	192189	371524	1264564	4951294
64	15062	36547	72876	147072	292109	564496	1920908	7520578
128	23188	55820	110999	223694	443987	857708	2917923	11423103
256	35704	85268	169078	340247	674843	1303232	4432435	17350713
512	54986	130268	257564	517548	1025757	1980199	6733058	26354263

NOTES: (1) All failure rates expressed in terms of Failures/10^9 hours

(2) Basic modeling relationship from MIL-HDBK-217B15

$$\lambda_P = \pi_L \cdot \pi_Q \cdot (C_T \cdot \pi_T + C_E \cdot \pi_E)$$

(3) Device modeling parameters are as follows

$\pi_L = 10$ ···· chip in production for ≤ 6 months

$\pi_Q = 5$ ···· Mil-Std-883, class B screening

$\pi_E = .2$ ···· ground benign environment

(4) Chip capacitors are expressed in terms of Kbits/chip, where
1 Kb $\triangleq$ 1024 bits

76-0649-VA-13

Mature MOS Memory Chip Active Failure Rates In Ground Benign Environment
(production screening to Mil-M-38510, class B)

Chip Capacity (Kb/chip)	Device Junction Temperature (^{o}C)							
	25	40	50	60	70	80	100	125
1	45	115	234	476	948	1835	6254	24497
2	70	176	356	723	1441	2789	9501	37209
4	107	269	542	1100	2190	4237	14431	56517
8	165	411	825	1672	3328	6437	21922	85845
16	254	627	1257	2543	5058	9781	33299	130391
32	391	957	1914	3868	7688	14861	50583	198052
64	602	1462	2915	5883	11684	22580	76836	300823
128	928	2233	4440	8948	17759	34308	116717	456924
256	1428	3411	6763	13610	26994	52129	177297	694029
512	2199	5211	10303	20702	41030	79208	269322	1054171

Note: device modeling parameters are as follows

π_L = 1 --- mature production for > 6 months

π_Q = 2 --- Mil-M-38510, class B screening

π_E = .2 --- ground benign environment

76-0649-VA-14

Mature MOS Memory Chip Active Failure Rates in Ground Fixed Environment

(production screening to Mil-M-38510, class B)

Chip Capacity (kb/chip)	Device Junction Temperature (^{o}C)							
	25	40	50	60	70	80	100	125
1	123	193	312	553	1026	1913	6332	24575
2	191	298	478	845	1562	2910	9622	37331
4	297	459	732	1290	2380	4427	14621	56707
8	462	707	1122	1969	3625	6734	22218	86142
16	718	1091	1721	3007	5522	10245	33763	130855
32	1116	1682	2639	4593	8412	15586	51307	198777
64	1735	2595	4048	7016	12817	23713	77969	301956
128	2698	4003	6210	10718	19530	36078	118487	458694
256	4194	6177	9529	16376	29760	54895	180063	696794
512	6522	9533	14625	25024	45353	83530	273645	1058493

Note: device modeling parameters are as follows

π_L = 1 --- mature production for > 6 months

π_Q = 2 --- Mil-M-38510, class B screening

π_E = 1 --- ground fixed environment

76-0649-VA-15

Mature MOS Memory Chip Active Failure Rates in Airborne Inhabited Environment

(production screening to Mil-M-38510, class B)

Chip Capacity (Kb/chip)	Device Junction Temperature (^{o}C)							
	25	40	50	60	70	80	100	125
1	415	485	603	845	1318	2205	6624	24867
2	647	754	933	1301	2018	3366	10078	37787
4	1010	1171	1444	2002	3092	5139	15334	57420
8	1575	1821	2235	3082	4738	7847	23332	87255
16	2458	2830	3460	4747	7261	11984	35503	132594
32	3835	4400	5357	7311	11131	18304	54026	201495
64	5983	6842	8295	11263	17065	27960	82217	306204
128	9335	10641	12848	17355	26167	42716	125125	465332
256	14567	16549	19901	26748	40132	65268	190436	707167
512	22730	25741	30833	41233	61561	99739	289853	1074701

Note: device modeling parameters are as follows

π_L = 1 --- mature production for > 6 months

π_Q = 2 --- Mil-M-38510, class B screening

π_E = 4 --- airborne inhabited environment

76-0649-VA-16

HP-25 PROGRAM FOR MOS MEMORY CHIP ACTIVE FAILURE RATE COMPUTATION

Step No.	Function Name	Operation Note	Step No.	Function Name	Operation Note
0	STOP	λ_p displayed	25	g e^x	
1	1		26	.	
2	ϕ	T_j increment	27	1	
3	STO +ϕ		28	X	
4	1		29	2	
5	2	T_j upper bound	30	g NOP	$\pi_L \cdot \pi_Q$ product
6	6		31	g NOP	
7	RCL ϕ	T_j value	32	X	
8	f x < y	T_j bound test	33	RCL 3	1.99
9	GTO 15		34	X	
10	2	B increment	35	RCL 4	B value
11	STO X4		36	RCL 5	0.603
12	2		37	f y^x	
13	5	T_j initialize	38	X	
14	STO ϕ		39	RCL 4	B value
15	f PAUSE	T_j displayed	40	f PAUSE	B displayed
16	2		41	RCL 6	0.644
17	7		42	f y^x	
18	3		43	RCL 7	0.56
19	+		44	X	
20	g 1/x		45	2	
21	RCL 1	1/298	46	g NOP	$\pi_L \cdot \pi_Q \cdot \pi_E$
22	-		47	g NOP	
23	RCL 2	-8121	48	X	
24	X		49	+	λ_p result

Notes: (1) Specific program values given for $\pi \cdot \pi_Q$ and $\pi_L \cdot \pi_Q \cdot \pi_E$ products are based on assumptions of mature device production ($\pi_L = 1$), Mil-M-38510, class B device procurement qualification screening ($\pi_Q = 2$), and ground fixed type environment ($\pi_E = 1$).

(2) To execute program, the following parameters are stored in stack

Register 0 - initial T_j value - 10°C - [successive increments of T_j are steps of +10°C up to +125°C bound]

Register 1 - 1/298

Register 2 - -8121

Register 3 - 1.99

Register 4 - initial B run value - [successive increments of B are in 2X multiples of preceding value]

Register 5 - 0.603

Register 6 - 0.644

Register 7 - 0.56

(3) λ_p values are given in terms of failures/10^9 hours.

HP-25 PROGRAM FOR BINOMIAL EXPANSION RELIABILITY COMPUTATION

Step No.	Function Name	Operation Note	Step No.	Function Name	Operation Note
0	STOP	R_u displayed	25	RCL 2	$\rho!$ value
1	ϕ	Clear Σ Accumulation Register	26	STO ÷3	Tab of $(\eta!/\rho!)$
2	STO 1		27	RCL 4	$(\eta-\rho)$ value
3	2	Initialize Factorial Tab Register	28	GTO ϕ6	
4	STO ϕ		29	RCL 2	$(\eta-\rho)!$ value
5	RCL 6	η value	30	STO ÷3	Tab of $\left[\frac{\eta!}{\rho!(\eta-\rho)!}\right]$
6	STO 2	Loop Entry Value	31	1	
7	1	Factorial Loop Entry Value Test	32	RCL 7	R_e value
8	f x ≥ y		33	-	$(1-R_e)$
9	GTO 30		34	RCL 4	$(\eta-\rho)$ value
10	-	Factorial Calculation Loop	35	f y^x	$(1-R_e)^{\eta-\rho}$
11	STO X2		36	RCL 7	R_e value
12	1		37	RCL 5	ρ value
13	f x ≠ y		38	f y^x	$(R_e)^{\rho}$
14	GTO 1ϕ		39	X	$(R_e)^{\rho}(1-R_e)^{\eta-\rho}$
15	STO -ϕ	Tab Increment	40	RCL 3	Factorial Term
16	RCL ϕ	Factorial Storage Tab Test	41	X	Partial Σ
17	g x < ϕ		42	STO +1	Accumulated Σ
18	GTO 29		43	1	Index Increment
19	g x = ϕ		44	STO +5	
20	GTO 25		45	STO -4	
21	RCL 2	$\eta!$ value	46	RCL 4	Calculation Terminate Index Test
22	STO 3	Tab on $\eta!$	47	g x ≥ ϕ	
23	RCL 5	ρ value	48	GTO ϕ3	
24	GTO ϕ6		49	RCL 1	R_u value

Equation Computed $\left\{ R_u = \sum_{\rho=(\eta-k)}^{\eta} \left[\frac{\eta!}{\rho!\ (\eta-\rho)!} \right] (R_e)^{\rho}\ (1-R_e)^{\eta-\rho} \right.$; where

R_u = Unit Reliability

R_e = Element Reliability

η = No. Elements/Unit

k = [Max No. failed elements permitted before a unit becomes useless

ρ = summation running index

Note: To execute program, register stack must be initially loaded as follows:

Register 7 - value of R_e

Register 6 - value of η

Register 5 - initial value of ρ; i.e., $(\eta-k)$ value

Register 4 - initial value of $(\eta-\rho)$; i.e., k value

Register 3 - ↓ Register 0 - } Clear (These registers are used by program, so anything left in them is destroyed by run)

Note: Factorial portion of program will not handle value of $\eta > 69$ due to register overflow.

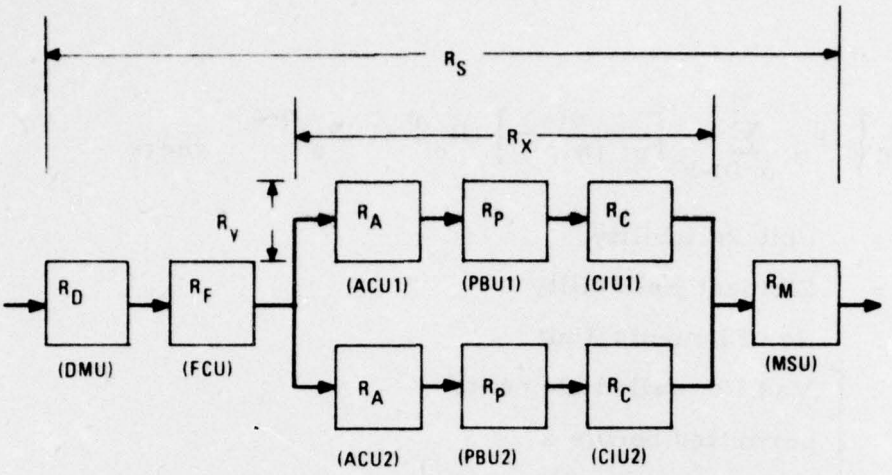

A. SYSTEM VIEWED AS A FUNCTIONALLY INTEGRATED ENTITY

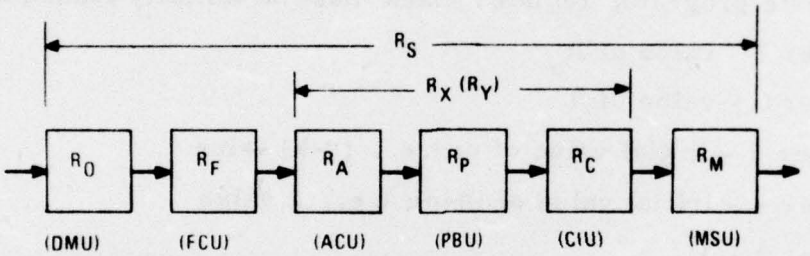

B. SYSTEM AS SEEN INDIVIDUALLY FROM EITHER PORT

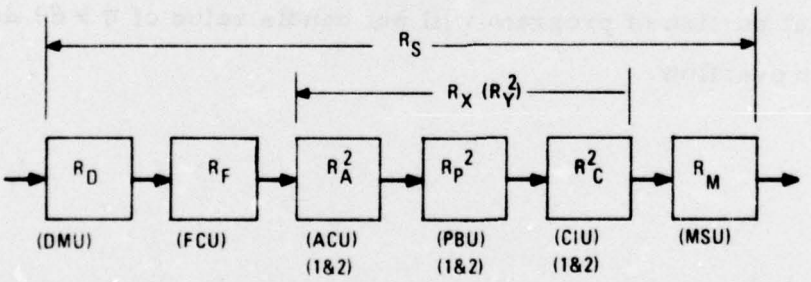

C. SYSTEM AS SEEN COLLECTIVELY FROM BOTH PORTS

76-0649-VA-17

Alternate Block Diagram Analysis Perspective of Composite Mass Memory System Reliability Model

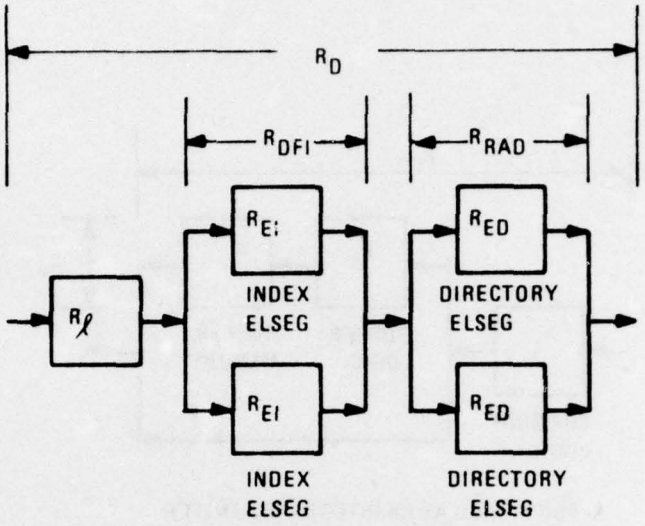

A. DMU VIEWED AS AN INTEGRATED ENTITY

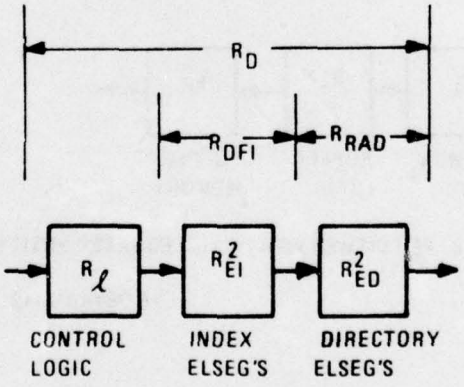

B. DMU VIEWED AS A NONINTEGRATED ENTITY

76-0649-VA-18

Alternate Block Diagram Analysis Perspective of Data Management Unit Reliability Model

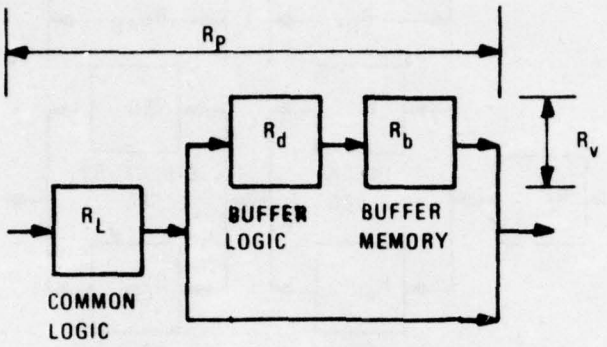

A. PBU VIEWED AS AN INTEGRATED ENTITY

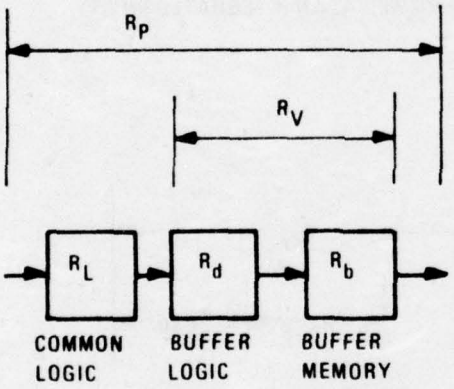

B. PBU VIEWED AS A NONINTEGRATED ENTITY

76-0649-VA-19

Alternate Block Diagram Analysis Perspectives of Port Buffer Unit Reliability Model

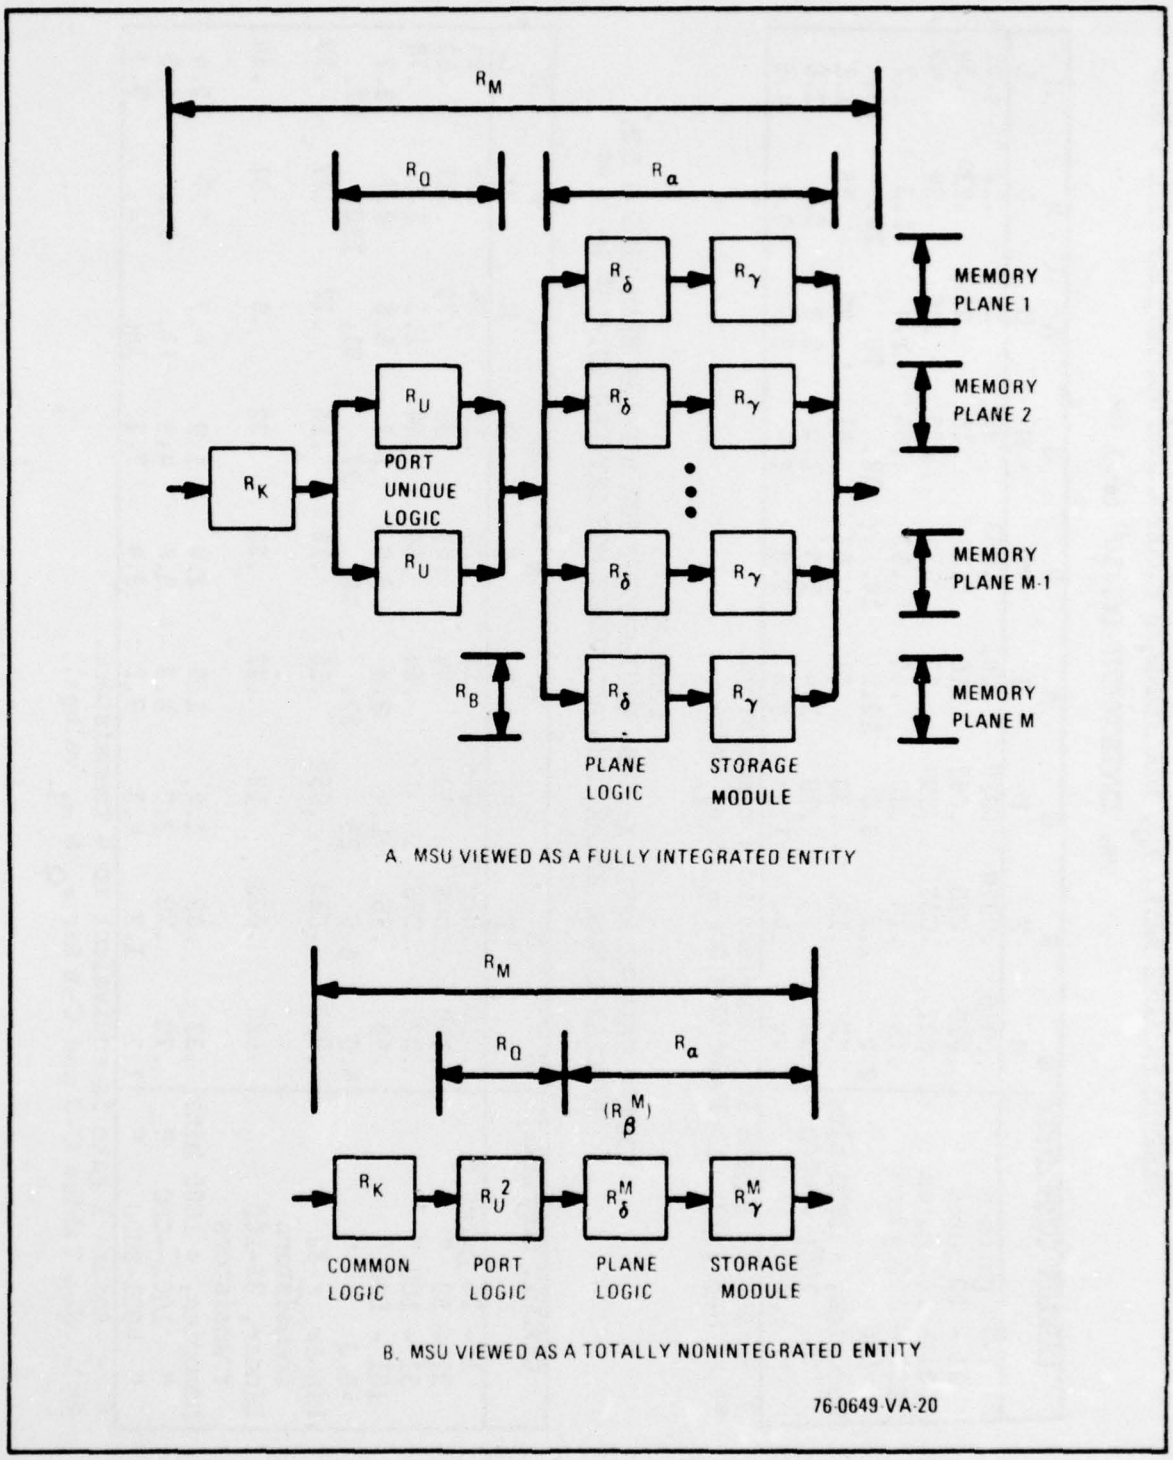

Alternate Block Diagram Analysis Perspective Extremes for Mass Storage Unit Reliability Model

TABLE C-1

GENERIC FAILURE RATE, λ_G, FOR STANDARD BIPOLAR DIGITAL DEVICES (TTL & DTL) vs. ENVIRONMENT (f./10^6 hr.) **

CIRCUIT COMPLEXITY	G B	S F	G F	A I	N S	G M	A U	N U	M L
1-20 gates *	.0070	.0070	.029	.091	.093	.091	.14	.12	.21
21 - 50 gates	.020	.020	.062	.16	.17	.16	.25	.23	.34
51 - 100 gates	.032	.032	.094	.23	.24	.23	.36	.34	.47
101 - 500 gates	.14	.14	.38	.86	.92	.86	1.4	1.3	1.7
>500	2.2	2.2	6.0	12.0	14.	12.	20.	20.	23.
Memories, < 1000 bits	.12	.12	.30	.61	.67	.61	.99	.98	1.1
" 1001 - 4000 "	.26	.26	.70	1.4	1.5	1.4	2.3	2.3	2.6
" 4001 - 8000 "	.44	.44	1.2	2.4	2.6	2.4	3.9	3.8	4.5

* - Assume 1 gate is equivalent to 4 transistors.

** - See Tables C-3 and C-4 for π_Q & π_L values.

TABLE C-2

GENERIC FAILURE RATE, λ_G, VS. ENVIRONMENT FOR BIPOLAR BEAM LEAD & ECL, BIPOLAR & MOS LINEAR, AND ALL OTHER MOS DEVICES (f./10^6 hr.) **

CIRCUIT COMPLEXITY	G B	S F	G F	A I	N S	G M	A U	N U	M L
1 - 20 gates *	.010	.010	.048	.12	.14	.12	.24	.25	.24
21 - 50 "	.048	.048	.19	.34	.43	.34	.79	.92	.52
51 - 100 "	.076	.076	.31	.54	.68	.54	1.3	1.5	.78
101 - 500 "	.36	.36	1.4	2.3	3.0	2.3	5.6	6.7	3.2
>500 "	6.0	6.0	23.	37.	48.	37.	91.	110.	47.
Linear, ≤ 32 transistors	.012	.012	.052	.15	.16	.15	.27	.27	.33
Linear, 33-100 transistors	.026	.026	.12	.32	.35	.32	.6	.61	.68
Memories, < 1000 bits	.32	.32	1.2	1.9	2.4	1.9	4.7	5.7	2.4
" 1001-4000 "	.70	.70	2.7	4.3	5.6	4.3	11.	13.	5.5
" 4001-8000 "	1.2	1.2	4.5	7.2	9.4	7.2	18.	22.	9.3

* - Assume 1 gate is equivalent to 4 transistors.

** - See Tables C-3 and C-4 for π_Q & π_L values.

TABLE C-3

π_Q, QUALITY FACTORS FOR USE WITH TABLES C-1 AND C-2

Quality Level or Screen Class	Description	π_Q
A	Mil-M-38510, Class A (JAN)	0.5
B	Mil-M-38510, Class B (JAN)	1
B-1	Mil-Std-883, Method 5004, Class B	2.5
B-2	Vendor Equivalent of Mil-Std-883, Method 5004, Class B	5
C	Mil-M-38510, Class C (JAN)	8
D	Commercial (or non-mil standard) part, with no screening beyond the manufacturer's regular quality assurance practices. The indicated π_Q value represents an average for all grades of commercial parts.	75

* These values are different from the values in the Table on page C-3.

TABLE C-4

π_L, LEARNING FACTORS FOR USE WITH TABLES C-1 AND C-2

The learning factor π_L is 10 under any of the following conditions:

(1) New device in initial production.
(2) Where major changes in design or process have occurred
(3) Where there has been an extended interruption in production or a change in line personnel (radical expansion).

The factor of 10 can be expected to apply until conditions and controls have stabilized. This period can extend for as much as six months of continuous production.

π_L is equal to 1.0 under all production conditions not stated in (1), (2) and (3) above.

INTERRELATIONSHIP OF DATA AND PARITY BITS
FOR HAMMING CODES

Max Number Data Bits $2^N-(N+1)$	Min Number SEC Parity Bits (N)	Total Encoded SEC Bits 2^N-1	Min Number SECDED Parity Bits (N+1)	Total Encoded SECDED Bits 2^N
1	2	3	3	4
4	3	7	4	8
11	4	15	5	16
26	5	31	6	32
57	6	63	7	64
120	7	127	8	128
247	8	255	9	256
502	9	511	10	512
1013	10	1023	11	1024
2036	11	2047	12	2048
4083	12	4095	13	4096
8178	13	8191	14	8192
16369	14	16383	15	16384
32752	15	32767	16	32768

TEMPERATURE SCALES CONVERSION TABLE

°C	°F
130	266
125	257
120	248
115	239
110	230
105	221
100	212
95	203
90	194
85	185
80	176
75	168
70	158
65	149
60	140
55	131
50	122
45	113
40	104
35	95
30	86
25	77
20	68
15	59
10	50
5	41
0	32
-5	23
-10	14
-15	5
-20	-4
-25	-13
-30	-22
-35	-31
-40	-40
-45	-49
-50	-58
-55	-67
-60	-76

NORMAL AMBIENT →

$$C = \frac{5}{9}(F-32)$$

$$F = \frac{9}{5}C+32)$$

APPENDIX D

MEMORY TECHNOLOGY REVIEW

The dissertations contained in this enclosure provide a guided tour through the frequently confusing maze of alternate memory technologies. The purpose is to select the most suitable technology for use in large secondary or mass memories, and to demonstrate the basis for that choice. To achieve this end, a process of elimination is employed based on tradeoff comparisons to a set of guideline requirements.

An Evaluation Of Memory Technologies For Secondary Mass Memory Applications

Critical computer secondary storage applications, such as those typically existing in military, DOD, and spacecraft systems, constitute a major technical challenge to memory development. The combined parameters of large capacity and high performance impose restrictions on possible implementation technologies and on the designs of both the memory system and the storage components. In the following discussions, an overview treatment is given first to technology and component requirements. Available and emerging technologies are then examined and screened on the basis of gross suitability. Finally, the major attributes of the prime contenders are evaluated and compared.

D. 1 SELECTION GUIDELINES

The first point of clarification that must be made is the size of the storage mechanism inferred by the terminology "secondary" or "mass." At times there has been a degree of confusion on this point, and the capacity boundaries have tended to both deviate from one memory proponent to the next and shift with time. To correct this situation, figure D-1 sets forth a nomograph of unified memory storage classifications covering both structural and organizational nomenclatures. The commonly encountered structural nomenclature deals exclusively with capacity variations in terms of total bits of storage (horizontal axis of nomograph). Organizational nomenclature, conversely, takes into account not only the bit capacity but also the grouping of data into commonly addressed blocks.

Note that "block" has no fixed definition insofar as the number of data bits it must contain. In different system situations, it has been variously referred to as a word, line, byte, slice, field, sector, or some combination thereof. Its absolute minimum and maximum boundaries are set solely by the obvious physical restriction that a block can contain neither less than one bit nor more than all of the bits. Beyond about 1 Kbit, though, a

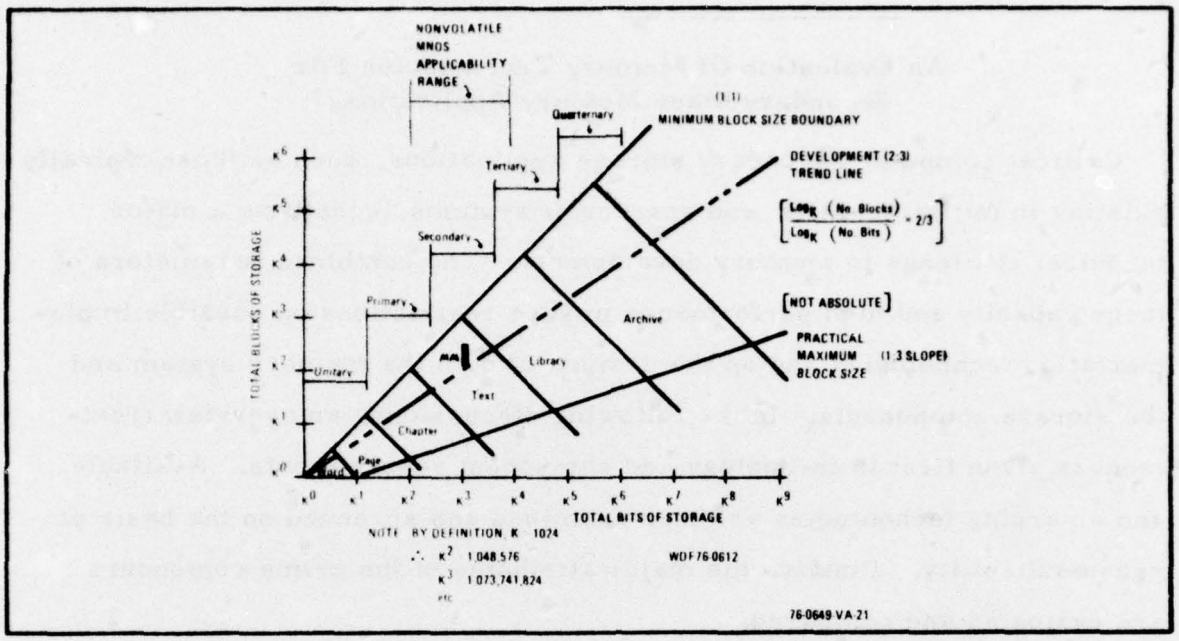

Figure D-1. Unified Memory Classification Nomograph

practical upper boundary is generally imposed as shown by implementation hardware considerations.

Assignment of the organizational classifications follow established computer industry practices. Respective regions are established in conformity with best engineering judgement of both present and future structuring of information processing equipment. The development trend line is based on the increased use of parallelism and multiprocessor computer facilities at higher storage capacities. For reference, the operational mass memory (MM) requirement established by RADC for their STARAN* associative array processor (AAP) based computer complex is indicated. This particular multiprocessor installation presently employs four (expandable ultimately to 32) concurrently operable associative processors, each of which transacts

*STARAN is a trademark of Goodyear Aerospace, Akron, Ohio.

logical operations on blocks of data containing up to 256 sequentially ordered words (or bit slices) of 256 parallel bits each.

Also shown in figure D-1 is the storage capacity range over which the nonvolatile MNOS technology (which will be reviewed subsequently) finds greatest utility. By human nature and obvious self-serving interests, the author is admittedly biased in favor of this technology base. No furtive effort is made to disguise such an easily recognizable fact. For the sake of engineering as well as self-integrity, though, every effort has been made to maintain the greatest possible degree of objectivity and equity in all statements made regarding memory alternatives and the tradeoffs among them.

Based on the definitions established by way of the classification nomagraph, it is seen that the capacity range of secondary mass storage systems extends from 10 K^2 (or about 10^7) bits up to 100 K^3 (slightly over 10^{11}) bits. Furthermore, from consideration of the organizations of memories falling in this range (i.e., text configurations), the format for data access is determined to be a mixture of random and sequential processes. Structuring of the storage medium is such that data is grouped and subgrouped at several different levels, in a manner analogous to subroutine nesting in software programs. Totally random access to single bits is physically impractical and extremely inefficient timewise when transferring large quantities of data between a main CPU or primary store and a secondary or backing store. Accordingly, access is accomplished randomly only to multiword data blocks. Within a block, access operations are word sequential with successive words consisting of a number of bits in parallel.

At a rapidly increasing rate, computer complexes are being expanded to contain a variety of processor units. In some cases these processors remain independent. In others they may be pipeline dependent, or else interconnected via a reconfigureable communications network. Each normally possesses its own local or working store, but collectively they are

commonly tied in to a central bulk or backing store. To prevent such complexes from becoming bogged down in access functions, the latency delay of the mass memory medium — excluding overhead data management or system bookkeeping time — generally should be no more than a few tens of microseconds. Further, once the setup delay to the first word in a block is completed, the word serial flow rate (reciprocal data rate) of the accessed data block must be asynchronously variable from a low-end full stop condition — such as may occur during a system interrupt — to a high-end maximum of typically < 200 nanoseconds/bit (> 5 MHz) over each parallel bit channel comprising a data word.

When a great number of user records and data files, some of which may be virtually irreplaceable, are placed in a central repository, it is mandatory that this memory possess high reliability. Extensive and varied forms of redundancy are commonly inserted in these memories to both detect the onset of failure and correct certain forms of transistory errors, thereby helping to ensure the continued integrity of all entered data. If the basic storage medium technology contributes sufficient innate reliability as a starting baseline, then redundancy techniques prove highly successful and quite cost effective. However, if the memory technology base cannot supply at least a 100-hour unsupported MTBF, a crossover point is quickly reached beyond which breakdowns become so frequent and intricate to repair that a diminishing return is realized on all further redundancy bulwarking.

An additional factor that must be taken into account when considering memory integrity is the volatility of the storage medium. Since it is next to impossible to unequivocally guarantee uninterrupted service from any prime power source, either the memory module or the computer system in which it is deployed must include provisions for protecting against data losses induced by any form of power line transient or failure. Further, this protection should extend over a minimum safety margin in the eventuality of

a lengthy power loss. Such situations typically arise as a result of intentional power-downs for maintenance or the occurrence of failures during unattended periods of operation.

Limits on the acceptable power drain of any given mass memory are heavily dependent on the application environment in which it is deployed. There exists, at one end of the scale, a special class of applications for which power is an overriding consideration. In those instances, permissible average dissipation levels may be on the order of 10 nanowatts/bit or less. At the opposite end of the spectrum, certain types of commercial applications place very little premium on power consumption. The only consideration is one of cost-of-ownership in buying the required quantities of power. Up to 100 microwatts/bit could be tolerated if it was necessary in some cases. For a mid-range mass memory having a capcity of K^3 bits (i.e., $\sim 10^9$ bits), this means that the allowable power may be anywhere between 10 watts and 100 kilowatts. As a working base from which valid assessments can be made, a mid-scale dissipation limit of 1.0 microwatt/bit can be taken as generally representative from the family.

Similarly, the purchase price which a user is willing to pay for a secondary storage unit is dictated by the nature of his application and the economic resources at his disposal. Presently, and into the immediate future (up through 1980), development of very large memories will initially be funded by some agency of the government or the armed forces, owing to the high engineering overhead involved. Once these nonrecurring costs are paid, though, commercial end-users generally demand that, to be a viable commercial commodity, initial procurement costs for production quantities of any basic memory module must be no greater than 0.1 cent/bit. If this is not possible, then the memory technology will not be widely accepted. This factor, along with the other points outlined, is summarized in the guideline requirements listing of table D-1.

TABLE D-1

SECONDARY MASS MEMORY GUIDELINE REQUIREMENTS

Characteristic	Guideline Value
Storage Capacity Range	$> 10^7$ to $< 10^{12}$ bits/memory module
Data Access Format	Random to multiword data block and parallel-bit word serial within a block
Access Latency Time (direct memory addressing)	< 10 μsec to first parallel-bit word of a multiword data block, exclusive of system variable address control overhead
Accessed Data Flow Rate (reciprocal of data rate)	< ∞ to > 200 nsec/bit per individual bit channel, unaided by time interleaving or multichannel multiplexing
Stored Data Integrity (volatility of stored data)	Intrinsically protected against any form of power transient or loss, even for periods extending to several days
System Reliability	> 100-hr MTBF, exclusive of all forms of overhead redundancy, unless they are inherent to storage medium
Power Consumption	< 1.0 μwatt/bit on full module basis (application dependent — up to two orders of magnitude less may be allowed in some cases)
Cost Potential	< 0.1 cent/bit for production quantities of commercial grade basic storage modules (excluding application unique system add-on circuitry which affects pricing structure)

Note: Not included are such considerations as mass (weight) and volume, which are critical only in certain situations.

Furthermore, specific attention is not given to operating environment factors such as ambient temperature, shock, and radiation.

A number of factors have been intentionally excluded from the list of secondary storage selection guidelines. Among these are such items as mass (weight) and volume. Although obviously extremely critical in certain areas (such as flight systems), the variance in these parameters is too heavily influenced by external factors to assign meaningful screening limits. For purposes of first-order technology evaluation, specific attention is also not given to the ultimate operating environment. Under this category are end-use considerations of ambient temperature, humidity, shock, vibration, acceleration, radiation, etc. It is tacitly assumed that specialized packaging could be employed to overcome any problems in these areas. Finally, intrinsic error rate capabilities have been ignored. Current trend in computer memories calls for a soft error rate of no more than 10^{-10} errors/bit processed. Error correction techniques are commonly applied to depress the error rate below this level whenever necessary.

D. 2 PRELIMINARY SCREENING

Almost every physical phenomenon which exhibits two or more stable states has been considered, at one time or another, as a possible memory technology. An exhaustive survey of all spurious attempts at memory realization is obviously beyond the intent of the present review. Still, it seems desirable to begin by "casting a wide net." Accordingly, table D-2 presents the full list of plausible technologies identified during a broad ranging literature search. Included here is every major implementation approach to suitably sized memories found to have been mentioned in the open literature which is either under active development or in production. In certain cases, the level of effort in a particular technology area is very low keyed, or it is being pursued principally for an application other than its use as a storage medium.

TABLE D-2

MEMORY TECHNOLOGY SHOPPING LIST

Generic Class	Technology	Acronyms
Magnetics	Fixed Head Drum	DRUM
	Fixed Head Disc	FHD
	Moving Head Disc	DISC
	Tape Recorder	TAPE
	Ferrite Core	CORE
	Plated Wire	WIRE
	Closed Flux	CFM
	Oligatomic Film	OFM
	Ferromagnetic Film	CROSSTIE
	Domain Wall	DYNABIT
	Domain Tip	DOT
	Magnetic Bubble Domain	MBD
Semiconductor	Metal Oxide Semiconductor	MOS
	Complementary MOS	CMOS
	Beam Addressed MOS	BEAMOS
	Nitride Interfaced MOS	MNOS
	Alumina Interfaced MOS	MAOS
	Silicon On Sapphire	SOS
	Charged Coupled Device	CCD
	Integrated Injection Logic	I^2L
	Bipolar Semiconductor	BIPOLAR
	Amorphous Semiconductor	OVONIC
Ferroacoustics	Magnetostrictive Wire	FAME
	Magnetostrictive Film	SONISCAN
	Surface Wave Delay Line	SAW
Electrostatics	Ferroelectric Film	MENTOR
	Storage Tubes	EBAM
Optical	Holographic Storage	--
	Beam Scanned	--
Cryogenic	Josephson Effect	JEM

Memory technologies in table D-2 are grouped initially according to their generic class. Specific technologies in each class are further subgrouped according to similarities in either their storage phenomena, mode of operation, or device fabrication features. The one noticeable exception is beam addressed MOS (BEAMOS), which has been here classified as a semiconductor technology. It actually represents a cross-breed, encompassing both high vacuum, hot cathode construction and monolithic IC wafer processing. Equally valid arguments can thus be made for grouping it under semiconductors or with storage tubes (like EBAM) under electrostatics.

The generic technology classes that seem most promising for mass memory applications are magnetics and semiconductor. Optical memories have yet to reach the product stage. When they do they are not expected to be competitive much below a K^3 ($\sim 10^{12}$) bit level. Cryogenic approaches are not mature and do not seem well suited to secondary stores. The ferroacoustic approaches have diverse problems. SAW is volatile and requires high power. FAME offers high data rates, but it has low bit density and poor reliability. SONISCAN fabrication is a complex materials and processing problem, and development has been discontinued. Ferroelectric film memories are being investigated by several firms, but evidence of progress significant enough to warrant serious consideration for large memories is lacking. Certain storage tube approaches have recently experienced a rebirth, but they do not appear to offer any advantages over the BEAMOS approach.

Of the magnetic technologies a few are obviously not suitable. The electromechanical approaches (DRUM, FHD, DISC, TAPE) can be dropped from consideration immediately because of the access time requirement. Similarly, DYNABIT and DOT can achieve neither the access time, nor the flow rate requirements. CROSSTIE is too immature to warrant consideration. This leaves five plausible magnetic technology candidates: ferrite core, plated wire, closed flux, oligatomic film, and magnetic bubble domain.

Core and plated wire are mature technologies with limited growth potential. They have not been applied to gigabit systems for cost reasons. Core has recently achieved 0.5 cents per bit in commercial 6×10^5 bit systems. Plated wire (2.5 mil) is projected at 2 cents per bit if volume production is achieved. A ferrite core approach to secondary memories would have difficulty with the reliability and power requirements. A 2.5-mil plated wire mass memory could probably achieve the performance, power, and reliability objectives, but the production cost potential of such systems would be nearly $20,000,000 per unit.

From time to time, various developments have been undertaken to achieve a planar format for the plated wire type of memory. The motivation for planar film approaches is potential cost and size savings. These magnetic film technologies store and retrieve information by domain rotation in an anisotropic media similar to plated wire. They differ from plated wire in fabrication processes and equipment requirements. Generally, planar film developments have encountered serious problems in reduction to practice. Many different firms have, however, embraced particular schemes and performed exploratory R&D.

Closed flux (CFM) and oligatomic film (OFM) are the planar equivalents of plated wire. The literature is unclear as to the current development status of the technologies. Manufacturing feasibility of CFM was being explored at Ampex Corporation under Navy funding. Arrays with 5,000 bits/in.^{2} have been demonstrated, with a potential of 10,000 bits/in.^{2} being claimed. In a CFM version of a secondary mass memory, performance and reliability goals could probably be met. Developmental units described by Ampex required about 70 microwatts/bit when active. Although this is much too high to achieve the stipulated goal directly, it may be possible to take advantage of the nonvolatility of their storage and apply power switching to lower the power requirement. No cost data is available, but certainly

an improvement over plated wire is expected. Information on the related OFM approach is very sparse. UNIVAC has been working on OFM but to date is known only to have demonstrated the feasibility of individual elements for a memory system. One would expect the OFM approach to have much the same characteristics and capabilities as CFM.

The magnetic bubble domain (MBD) approach is the only nonelectromechanical magnetic technology with the potential of realizing mass memory size storage systems at reasonable cost. The primary advantage of MBD is potential high bit density per chip. Devices containing on the order of 100 K^2 (> 10^5) bits are under development. The major disadvantages of MBD are slow speed and the fact that the nonvolatility of MBD is dependent on the rotating fields being shut down properly.

From the vantage point of a general secondary storage application, the performance question is of central importance. MBD is limited to various forms of shift register organizations. This necessarily imposes some rotational latency. A major-minor loop approach requires some delay in the major loop. It will defnitely, therefore, be a difficult task to meet a 10-microsecond access delay at the chip level, let alone at the module level.

Present day MBD systems operate at about 100 KHz. Research personnel claim that 1 to 10 MHz rates will be achieved, but personnel responsible for system design note that 300 KHz is closer to the practical maximum rate. To achieve the mass memory data flow specification, some form of on-chip and/or off-chip multiplexing will be required. No doubt, there are choices of bubble size, detection scheme, and number of I/O terminals that can meet the requirement. The chip design, however, must also consider the system factors of chips per field coil and number of field coils that must operate simultaneously.

The coil driving problem of MBD becomes difficult at higher frequencies. In particular, to prevent data storage perturbations it is necessary that all dynamic fields associated with an MBD system be deactivated within a narrow ±5 degree phase angle window. This is a difficult task at a field rate of 100 kHz, and one obviously compounded as the field rate and drive current levels are increased. Also, during critical field shutdown, an MBD system requires some form of supplemental nonvolatile storage to remember the rotational position of data in the shift register so that subsequent access requests retrieve the correct block of data.

If one assumes that these difficulties can be overcome, it is of interest to predict the physical features of an MBD realization of a 1 Gbit (i.e., K^3 bits) memory module. Such a unit, if based on the use of 64 Kbit chips, could be built in approximately 7 cubic feet with a weight of about 300 pounds. Given that MBD devices prove to be producible at the same order of cost as LSI semiconductor devices, a production system price on the order of 0.07 to 0.10 cents/bit may be achieved. Power requirements will depend on how the multiplexing is performed to achieve the data rate.

The semiconductor technologies offer both volatile and nonvolatile storage media. Of the nonvolatile approaches, only three are capable of fast read and write: MNOS, MAOS, and BEAMOS. Volatile approaches like MOS and CCD, while possessing the potential of realizing mass memory systems, exhibit marginal at best reliability and power characteristics. The combined capacity and 100-hour MTBF imply constraints on bit density per component. Figure D-2 presents the results of an approximate analysis of the memory chip contribution to an intermediate sized (1 Gbit) mass memory failure rate for a nonredundant system configuration. Chip failure rates were estimated using MIL-HDBK-217B, MOS RAM models, based on T_j (active) = 40°C, T_j (dormant) = 25°C, and a ground fixed environment.

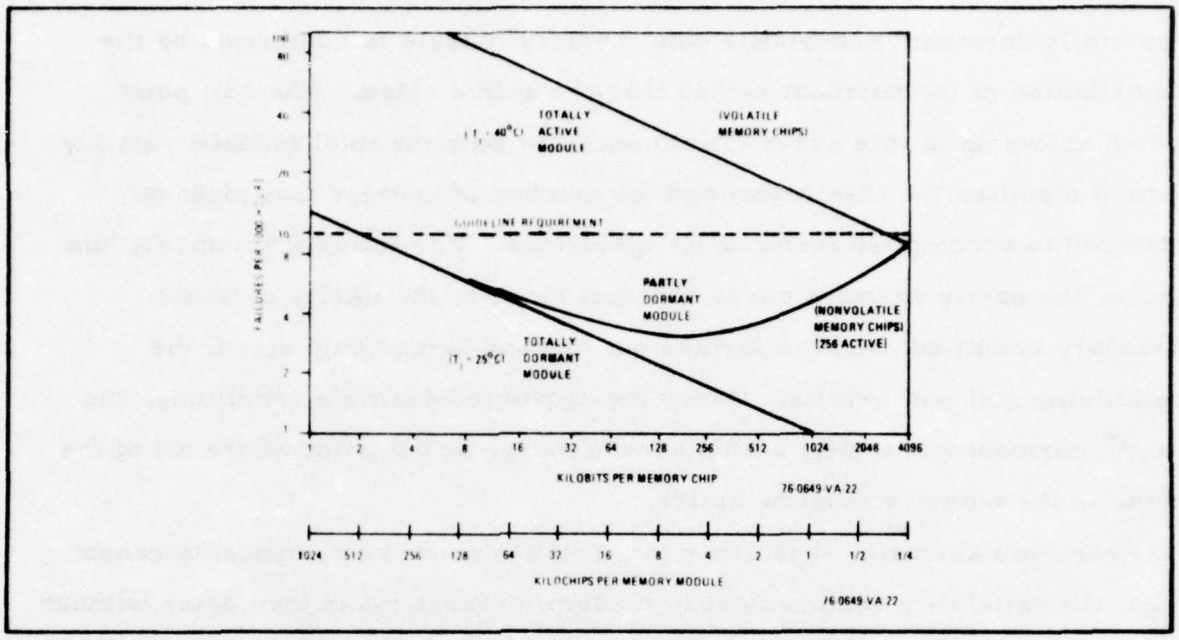

Figure D-2. Memory Chip Contribution to Gigabit Memory Module Failure Rate

Nonvolatile memory components are shown to have a significant advantage from a reliability standpoint. The nonvolatile components are assumed to be power switched. Only those devices which are to emit or receive data have voltages applied. Devices which are powered down (dormant) enjoy a reduced failure rate because no electric field stresses are present. The curve for volatile memory chips is simply the failure rate for one chip times the number of chips in the complete memory module. The curve for nonvolatile chips assumes that 256 chips must be energized but that all other chips remain dormant. For modeling purposes, the failure rate of a dormant chip is taken as 1/10 that of an active chip at reduced junction temperature. A lower bound for the module failure rate exists when all of its components are dormant. This state is the antilog of the totally active state.

Up to a chip size on the order of 128 Kbits, the failure rate of the "partially dormant" nonvolatile chip memory module is dominated by the contribution of the dormant rather than the active chips. The null point which shows up in this curve is influenced by both the total module capacity (which modifies the chip count) and the number of devices that must be powered to accomplish memory IO operations. The fewer the number, the longer the partly dormant curve remains close to the totally dormant boundary condition, and the further out (toward larger chip sizes) the optimizing null point shifts. Under the stipulated example conditions, the partly dormant and totally active curves merge at the point where all of the chips in the memory must be active.

From this analysis, it is clear that volatile memory components cannot meet the reliability requirement in medium to large mass memories without the aid of redundancy in the basic storage section. Even with redundancy, a volatile device will not provide the same level of reliability as an equivalent nonvolatile chip. These results tend to indicate that volatile technologies should be rejected for this size of memory on reliability grounds alone. Doing so without at least a cursory review of their other characteristics is somewhat arbitrary, however, since in actual practice some form of redundancy will normally be employed. The effective failure rate contribution of the storage section devices will thus typically be significantly lower than either of the nonredundant curves shown indicates. In fact, physical constraints arising from memory component count and module labor content frequently prove to be the limiting considerations rather than the MTBF goal.

Among the semiconductor technologies, BIPOLAR is the oldest and most mature. It has been used for some time to manufacture small memories. It possesses the capability of high speed, but also exhibits high power, low bit density, and high cost per bit. BIPOLAR memory technology is therefore not suited to use in large backing stores. Integrated injection logic (I^2L) is

a recent addition to the semiconductor family. It is fundamentally a BIPOLAR technology in which processing has been simplified through the elimination of epitaxial layers and the exclusion of isolation diffusions. I^2L promises higher bit packing densities and lower speed-power products than its predecessor. Although a number of major semiconductor houses are heavily involved in investigating I^2L devices, the technology is still too immature to be a serious contender.

Amorphous (or OVONIC) semiconductor devices are inherently slow write memories. Because of the read-mostly nature and high dissipation of the technology, it is rejected for high-speed mass memory use. Silicon-on-sapphire (SOS) technology does not constitute a stand-alone approach to memory device production. SOS processing is combined with a major device fabrication technique, such as CMOS or MNOS, to secure dielectric isolation of the individual active elements. Its major assets are very high radiation resistance and rapid recovery from high radiation exposure. It is unacceptable for large memories due to low bit density and high cost per bit (resulting from complex manufacturing and low processing yields).

Characteristics of the hybrid BEAMOS technology are mixed in their fit to mass storage applications. To date the feasibility of a 32 Megabit storage tube has been demonstrated, with access times of 30 microseconds and serial bit rates of 10 MHz claimed. A tube and associated electronics are estimated to have an 8,000-hour MTBF (with cathode replacement on a maintenance schedule). Data readout is not completely nondestructive. Disturbed data must be refreshed after approximately 100 reads. Non-volatile retentivity of the stored data is adequate for about a month of powered down dormancy. Volume costs as low as 0.02 cents/bit have been projected.

The BEAMOS approach (as developed principally by General Electric) appears to be capable of realizing large, low-cost storage systems; but the single serial bit channel nature of its output seriously hampers its use in secondary stores calling for the transfer of data words having many parallel bits. For example, 32 tubes could provide a Gigabit (K^3 bits) of storage and an I/O flow rate of 100 nanoseconds/bit in each of 32 simultaneous bit streams. A minimum of 128 tubes (allowing for a 2:1 external circuitry multiplexing factor) would, however, have to be deployed to establish a 256 parallel stream interface with all streams flowing at 200 nanoseconds/bit. Since all tubes would have to be continuously powered in this case, no dormancy related improvement in their composite reliability could be realized. By component capacity and IO structure, these devices more aptly belong in the tertiary memory size classification; but their qualified nonvolatility makes them totally unsuited for such uses.

Of the remaining semiconductor technologies, all are based on field effect structures of one form or another. MOS (metal-oxide-semiconductor) is the grandaddy of them all. Both p-channel and n-channel memory devices in static as well as dynamic form and in sizes up to 4 Kbits/chip are presently in widespread use. Units of 16 Kbits are expected to become commonplace by 1980, with still higher capacities following soon thereafter. These devices have access times ranging from about 100 to 800 nanoseconds and cycle times on the order of 200 nanoseconds to over 1 microsecond. They are conventionally configured to provide totally random access by bit or word (2 to 4 parallel bits) rather than block oriented accesses. Nothing in the storage media design or fabrication prevents such structuring, however, should there be sufficient economic demand for it.

Bit density capabilities of dynamic MOS memory chips based on single transistor (1T) memory cell designs are adequate for production of very large systems. The only technology typically exhibiting a better density

capability is CCD. Table D-3 shows how the most widely recognized semiconductor and the one plausible nonsemiconductor (MBD) memory technologies stack up in this category. The data in this table is based on the correlated findings of a broad ranging NASA study and an IEEE Press publication on the MBD technology.[1, 2] Principal drawbacks to the use of either the CCD or MOS technologies are the volatility of stored data, the resulting marginality in large memory module reliability, and the relatively greater amount of system support electronics (compared to static or nonvolatile approaches) necessary to ensure continuous refreshing of dynamic memory cells.

TABLE D-3

STORAGE CELL SIZES OF ESTABLISHED MEMORY TECHNOLOGIES GIVEN COMMON PHOTOLITHOGRAPHY

Technology Base		Relative Cell Size
CCD	(Dynamic 3 phase)	0.60
MOS	(Dynamic 1T)	0.80
MNOS	(Nonvolatile 2T)	1.00
MBD	(Nonvolatile T-Bar)	1.45
BIPOLAR	(Static)	9.50
CMOS	(Static)	16.00

1. NASA Goddard; "Camera Memory Study for Large Space Telescope, Final Report;" February 1975; table 3-10, p. 3-48.
2. Hsu Chang, editor; Magnetic Bubble Technology; IEEE Press; New York; 1975; chap. 6, sect. 6, table 3, p. 214.

Complementary MOS (CMOS) memory structures are based on the use of counter-balanced p-channel and n-channel transistors on a common monolithic die. Due to this, it is a more complex technology than either PMOS or NMOS alone. CMOS memory cells typically include 8 transistors. Consequently, at any given state of development, larger capacity components will be available with standard MOS than with CMOS. Low dc (quiescent) power is the sole unique advantage offered by CMOS. In all other respects it constitutes a poor choice for secondary storage applications.

Nitride interfaced MOS (MNOS) technology also involves additional processing steps compared to standard MOS (8 to 10 versus 5 to 7, typically). It does not, however, suffer any major loss in memory cell density. MNOS density capabilities based on single transistor (1T) memory cells are in fact nearly a factor of two better than those of conventional dynamic MOS due to simpler layout geometries (planar rather than cross-sectional). It is unfair to base density comparisons on a 1T cell structure, though, since MNOS device designs founded on 1T cells have proven impractical for large die -- sensitivity to processing variations leads to unacceptably low yields. The data of table D-3 therefore reflects the use of a differentially matched two-transistor (2T) MNOS memory cell, as originally developed by Westinghouse. This cell configuration has demonstrated high yields in production runs, exhibits extremely wide operating supply margins, and withstands greater than 10^{12} clear-write operations without degradation.

Structurally, MNOS devices are formed as a dual dielectric sandwich of aluminum, silicon nitride (Si_3N_4), silicon dioxide (SiO_2), and silicon. Memory in an MNOS transistor is produced by electric field induced tunneling of charges from the silicon through the oxide layer to trap sites at the oxide-nitride interface. Stored data retention capabilities extend anywhere from several months to many years depending on how the product parameters are tailored during manufacture. Access times are typically

on the order of 1 to 5 microseconds. In block-oriented (BORAM) configurations, MNOS data flow rates of multichanneled memory chips are better than 200 nanoseconds/bit/channel. Write times commonly used with 2T MNOS memory devices range from 50 to 200 microseconds. Clearing or erasue of blocks of stored data is readily accomplished in under 30 microseconds.

MNOS devices easily meet the mass memory performance requirements. Power drain of an MNOS memory module in most installations can be expected to be an order of magnitude lower than that allowed. Using redundancy, an effective MTBF of 10,000 hours should be realizable in a mid-sized (K^3 bits) module. If an optimally packaged memory of this capacity were fabricated using 64 Kbits MNOS chips, the module would occupy approximately 3 ft^3 and would weigh roughly 200 pounds. Final system prices of 0.08 to 0.10 cents/bit for production quantities should be readily achievable.

A variety of other approaches to nonvolatile semiconductor memories have been proposed and fabricated in limited quantities. One, referred to as FAMOS (floating-gate avalanche-injected MOS), makes use of an electrically floating polysilicon gate to hold data related quantitites of charge. With these devices, charge can readily be injected electronically but can only be removed with great difficulty. X-ray or UV radiation is, in fact, commonly employed to remove the trapped charge. They are intended for use as read-only memories and are totally ill suited to use in high-speed secondary stores. Another, MAOS (alumina interfaced MOS), stores data charges within an alumina film. This technology is quite similar to MNOS. It is projected to have similar or better characteristics, but presently is still too immature for serious consideration.

Upon considering the mass memory prerequisite factors collectively, there appears to be only two candidate technologies whose viability for implementing secondary stores (10 K^2 to 100 K^3 bits) is truely plausible: namely, MNOS and MBD. Both technologies are capable of providing high

bit density per chip and offer similar cost potentials. Each supplies non-volatile data storage, although a degree of control conditionality is attached to that of MBD. Of the two, MNOS can achieve higher performance. A related point (which, although not part of the general selection guidelines, should at least be mentioned) is that MBD operating margins are relatively sensitive to ambient temperature and require compensation if the environmental conditions change. MNOS storage, on the other hand, being based on a temperature insensitive tunneling phenomenon, is unaffected by wide variations in the operating environment.

D.3 TECHNOLOGY COMPARISONS

To determine whether grounds exist for making a final definitive choice between the two prime technology candidates, it is necessary to examine their respective characteristics in greater detail and evaluate any differences in performance that accrue in a representative mass memory application. Comparisons will be based on the postulated existence of comparable capacity devices from both MBD and MNOS in a 1978 to 1980 time frame. Furthermore, it will be assumed that these devices have been configured and IO interface "groomed" in select ways to optimize their use in wide ported (i.e., many parallel bit streamed) secondary storage environments. In making any type of cross-comparison, it is desirable to have a point of reference from which to validate conclusions. Toward this end, similar treatments of a representative high density volatile technology (viz., CCD) are presented where appropriate for orientation or contrast.

Presently, the maximum practical data transfer shift rate for MBD memories is under 250 KHz using bubble domains of approximately 4 μm in diameter. Experimental devices in the laboratory have operated at slightly over 1 MHz but employed bubbles on the order of 50 μm in size. A fundamental materials problem exists (as shown in figure D-3) which hampers the simultaneous attainment of small (sub-micrometer) sized bubbles and

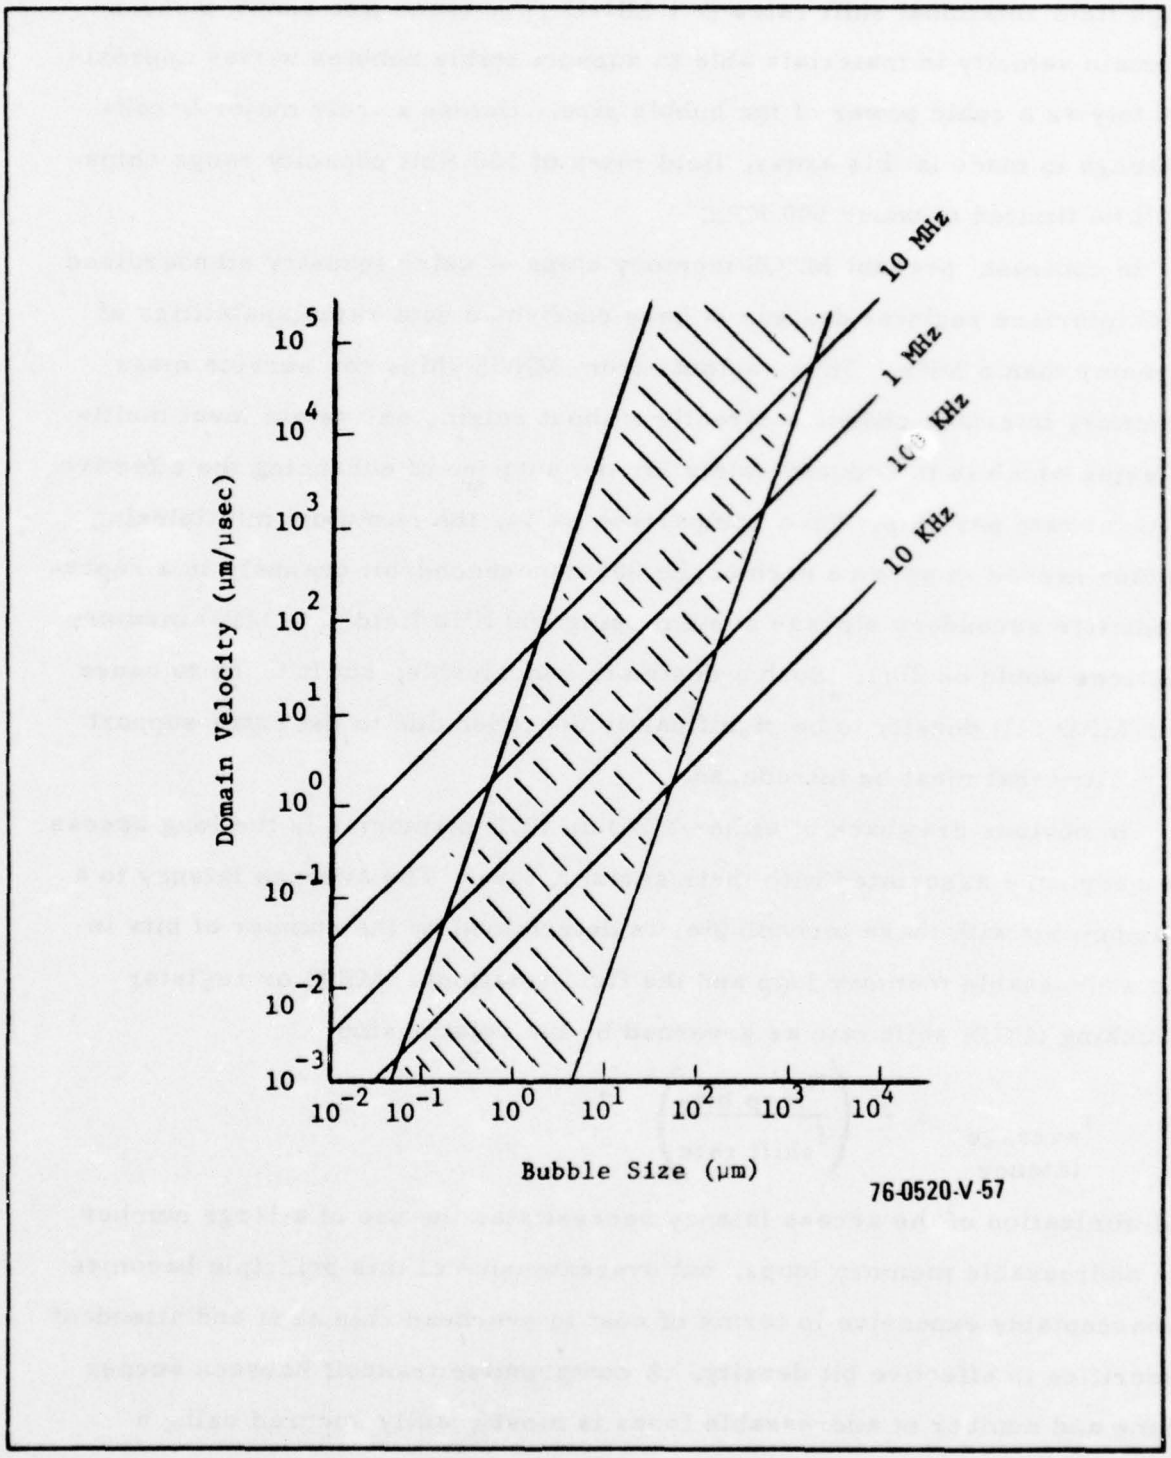

Figure D-3. Magnetic Bubble Materials "Size-Velocity" Band

high field rotational shift rates (> 1 MHz). The trend line shows that the domain velocity in materials able to support stable bubbles varies approximately as a cubic power of the bubble size. Unless a truly major breakthrough is made in this areas, field rates of 100 Kbit capacity range chips will be limited to under 500 KHz.

In contrast, present MNOS memory chips — using industry standardized I/O interface register designs — have confirmed data rate capabilities of greater than 6 MHz. Thus, outputs from MNOS chips can service mass memory interface channels directly without relying on system level multiplexing which is introduced solely for the purpose of enhancing the effective I/O bit rate per chip. On a comparison basis, the minimum multiplexing factor needed to service each of the 200 nanosecond/bit channels in a representative secondary storage system using 250 KHz field rate MBD memory devices would be 20:1. Such a design is conceivable, but it tends to cause the MBD cell density to be significantly degraded due to the extra support circuitry that must be introduced.

An obvious drawback of either MBD or CCD memories is the long access latency time associated with their serial nature. The average latency to a random bit with these technologies is determined by the number of bits in an addressable memory loop and the field rotational (MBD) or register clocking (CCD) shift rate as governed by the relationship:

$$t_{\substack{\text{average} \\ \text{latency}}} = \frac{1}{2}\left(\frac{N_{\text{loop bits}}}{f_{\text{shift rate}}}\right)$$

Minimization of the access latency necessitates the use of a large number of addressable memory loops, but overextension of this principle becomes unacceptably expensive in terms of cost in overhead chip area and attendant sacrifice in effective bit density. A compromise tradeoff between access time and number of addressable loops is most readily secured using a

major/minor loop structure. For chips having a capacity on the order of 100 Kbits, though, the latency time would still lie somewhere in the millisecond range. MNOS memories, on the other hand, are natually random and suffer no inherent latency delay to the first bit of a block other than address decode and gate propagation delays.

As has been noted, both MNOS and MBD technologies exhibit excellent bit density capabilities. Cell size within the storage area of a memory device is, however, only part of the story when considering total system package dimensions. In an MBD memory it is impractical to place more than about 16 chips in a common bias and rotating field coil assembly. At greater than this level, matching of the sets of chips becomes critical and significantly degrades the operating bias margins. For a K^3 bit capacity system built using 16,384 MBD chips of 64 Kbits each, this implies that a minimum of 2,048 coils and 1,024 permanent magnets would be needed just to build the field coil assemblies.

As an interesting contrast in ultimate packaging densities achievable, consider the following. Bell Laboratories recently announced a dual field coil assembly "bare bones" 512 Kbit MBD memory module which uses 32 chips of 16 Kbits capacity each and is housed in a package measuring 3.75 by 1.375 by 0.8125 inches. This is a volume of 4.19 cubic inches. To make a comparison to an equivalent MNOS memory, consideration is made of developmental 16 Kbit Westinghouse chips of current design which are less than 250 by 250 mils. Allowing for a liberal hybridization factor of 3:1, the substrate assembly surface area needed would be $(250 \times 250) \times 32 \times 3 \times 10^{-6} = 6.00$ in^2. Using a 40-mil-thick ceramic substrate with 55-mil clearances above and below, 50-mil clearances on either side, 50-mil chassis wall thickness all around, and spacing of 0.5 inch for connectors off one edge of the substrate, the package size of a "bare bones" 512 Kbit MNOS memory module would be no greater than 2.5 by 3.0 by

0.25 inches. This is a volume of only 1.88 cubic inches, or less than half that of the equivalent capacity MBD module.

Consideration of the type of support electronics available to implement a K^3 bit MBD memory with the minimum necessary process control functions, I/O multiplexers, transfer drivers, generator drivers, replicate/annihilate drivers, field coil drivers, and sense amplifiers (but exclusive of intelligent access control and data management overhead) leads to an estimated need for some 20,000 IC's, 22,000 resistors, and 15,000 capacitors. An MNOS memory would need less than one-fourth of these components: viz, approximately 5,000 IC's, 2,000 resistors, and 3,000 capacitors. Accordingly, in the area of system support circuits, MNOS exhibits a distinct advantage over MBD.

To facilitate analysis of the power dissipation of MNOS and MBD (and, for comparison purposes, CCD) memory modules, it is necessary to first be cognizant of certain rudimentary aspects of the memory architectures that are needed with each technology to provide the necessary macroscopic interface conditions. Studies have shown and experience has confirmed that, in general, the power of a volatile type store (such as CCD) will be minimum if all devices are operated in parallel at their slowest possible rate consistent with having their outputs multiplexed together to achieve the necessary system data rate. Nonvolatile stores, contrarily, should be configured in such a manner as to have the minimum number of devices powered up at any given time so as to secure maximum advantage of their nonvolatility in the form of improved reliability and reduced support circuitry.

Strictly for equivalency in the analyses to be made, it will be assumed that 64 Kbit[1] devices of each technology are available. Accordingly, each memory module will employ 16, 384 chips to provide identical capacities of 1 Gbit[2]. Only power drawn by the memory devices and their immediate interface driver circuits located within a minimum configured storage module will be considered.

Assumptions regarding the MBD chip structure are that a factor of 8 multiplexing has been incorporated directly on the chip and that 8 parallel outputs are provided from each chip. It is recognized that the assumed ability of MBD manufacturers to readily accomplish this may be somewhat optimistic. MNOS chips are assumed to have 8 parallel I/O bit channels per device with 64 stages in each IO register. For the CCD chips, two cases will be considered: (1) a single IO channel per chip and (2) eight parallel IO channels per chip. For the system level interfaces, it will be assumed that 4 concurrently operable 256-bit-wide word ports must be serviced at 200 nanoseconds/word. No direct consideration will be given to implications of differences in system overhead support electronics. Only storage section drivers will be included.

Based on 8 channels available from alternate memory devices, the memory planes of the respective modules must contain 128 chips to provide the necessary 1, 024 paralleled data interface channels. Grouping of the MBD devices is assumed to be in the form of 16 chip matched sets. This means that there are 8 field coil assemblies per MBD memory plane. Furthermore, it is to be noted that there will be 128 planes in each memory (given 16, 384 chips/memory and 128 chips/plane).

---

(1) 1 Kbit is defined as 1, 024 bits; so, 64 Kbits = 64 x 1 Kbit = 65, 536 bits.

(2) 1 Gbit is defined as 1 K^3 bits; so, 1 Gbit = $(1, 024)^3$ bits = 1, 073, 741, 824 bits.

The final stage-setting factor to be considered is the number of memory planes which are active. For a volatile store, all planes are active. With the MBD memory, four planes will be activated to ensure achieving the specified IO data flow rate. In the case of an MNOS store, only one plane has to be active to provide the IO rate, but two will be assumed to be activated and operated in an alternating flip-flop like manner where first one and then the other handles the transfer of successive data blocks. This process ensures that block addressing functions are completely masked, even at successive block boundaries.

A summary of device characteristics relevant to the power calculations to be made here is presented in table D-4. Confirmed 1976 parameters are itemized in the first column, then sets of both realistic and optimistic projections for a 1978 technology base are listed in the latter two columns. The realistic projections will be employed since they represent the most likely extent of advancement, based on engineering judgement, that can be expected in each technology field. Sample calculations for the condition where continuous data I/O exchanges are taking place at a rate of 200 nanoseconds/bit over each of 1,024 parallel data lines are itemized in the following listing. In each, allowance is made for the dissipation of 50 percent efficient field or clock drivers. For simplicity, no other form of overhead support circuitry power has been added in.

MNOS:

On chip power = $0.15 \times 256 = 38.4$ watts

Capacitive power = $0.4 \times 10^{-12} \times 64 \times 8 \times 2 \times 128 \times (12)^2 / 0.20 \times 10^{-6} = 37.75$ watts

50% efficient driver power = $0.5 \times 37.75 = 18.87$ watts

Total power (less support) = $38.4 + 37.75 + 18.87$

$= \underline{95.67 \text{ watts}}$

TABLE D-4

ALTERNATE TECHNOLOGY CHARACTERISTICS

Technology \ Timeframe	1976 Confirmed Parameters	1978 Realistic Projections	1978 Optimistic Projections
MNOS			
Number Clock Phases (ϕ)		2	
Clock Drive Amplitude (V/ϕ)		12	
Minimum Clock Rate (Hz)		0	
Maximum Clock Rate (MHz)	6	8	10
Clock Capacitance (pF/SR bit/ϕ)	0.8	0.4	0.2
On-Chip Static Power (mW/chip)	200	100	50
MBD			
Number Field Phases (ϕ)		2	
Field Drive Amplitude (mA/ϕ)		120	
Minimum Field Rate (Hz)		0	
Maximum Field Rate (kHz)	150	250	500
Field Inductance (ηH/bit/ϕ)	4	3	2
On-Chip Static Power (mW/chip)	200	100	50
CCD			
Number Clock Phases (ϕ)		2	
Clock Drive Amplitude (V/ϕ)		12	
Minimum Clock Rate (kHz)	100	64	32
Maximum Clock Rate (MHz)	2	4	6
Clock Capacitance (fF/bit/ϕ)	40	30	20
On-Chip Static Power (μW/bit)	4	2	1

MBD:

On chip power = $0.15 \times 512 = 76.8$ watts

Inductive power = $3 \times 10^{-9} \times 65,536 \times 512 \times 2 \times (0.12)^2 / 6.4 \times 10^{-6} = 453$ watts

50% efficient driver power = $0.5 \times 453 = 226$ watts

Total power (less support) = $76.8 + 453 + 226$

= 775.8 watts

CCD:

On chip power = $2 \times 10^{-6} \times 16,384 \times 65.536 = 2,147$ watts

Case 1 (8 IO's/chip)

Capacitive power = $2 \times 10^{-6} \times 65,536 \times 2 \times 16,384 \times (12)^2 / 15.625 \times 10^{-6} = 593.7$ watts

50% efficient driver power = $0.5 \times 593.7 = 296.9$ watts

Total power (less support) = $2147 + 594 + 297$

= 3038 watts

Case 2 (1 IO/chip)

Capacitive power = $30 \times 10^{-15} \times 65,536 \times 2 \times 16,384 \times (12)^2 / 6.4 \times 10^{-6} = 1450$ watts

50% efficient driver power = $0.5 \times 1450 = 724.8$

Total power (less support) = $2147 + 1450 + 725$

= 4322 watts

A check of the results listed here reveals that a significant power advantage accrues to the MNOS technology. At the specified 5.12 GHz effective bit I/O rate (i.e., 1,024 divided by 200 nsec), MNOS needs no more than 1/8th the power of MBD and over 40 times less than CCD's. For other data rates above and below the target specification, the power drain of the two nonvolatile technologies vary proportionately. Dissipation of the CCD based memory module, on the other hand, buttoms out at, and never drops below, the 3 KW level denoted for the 8 I/O chip configuration. The reason for this is simply due to the fact that CCD's are both volatile and dynamic. Consequently, all devices in the memory must remain powered at all times and must receive refresh energy at some minimum rate (i.e., 64 KHz, in this instance).

The power requirements considered here pertain to worst-case conditions existing during active data input-output operations. Obviously, if there is dead-time in the memory activity, the average power for both nonvolatile type storage units will drop as a result of their being powered

down during such intervals. Volatile stores such as CCD's enjoy no such advantage. As noted, this factor impacts heavily on a memory's reliability since devices capable of being placed onto a powered-down state with all electrical stresses removed typically experience a failure rate improvement of at least one order of magnitude.

By following the established practice of MIL-HDBK-217B, it is possible to conduct detailed, in-depth analyses of any memory system. Only a simplified treatment will be undertaken here, solely to investigate innate differences. For this effort, the basic MTBF (mean-time-before-failure) of the respective memory devices will be evaluated using the fundamental relationship:

$$MTBF = D/N\lambda$$

where:

N = total number devices

λ = active device failure rate

D = dormancy improvement factor

Since the cross-comparison form of treatment given to the alternate memory technology modules has been based on the postulated existence of equivalent capacity chips from each technology, the device count (N) will be the same for each. As to the active failure rate (λ) in the MTBF equation, MIL-HDBK-217B covers only the MNOS technology; and even it not up to the size chips being considered here. Until such time as a firm data base is established for both higher capacity memory components and the CCD and MBD technologies, values of confirmed validity cannot be assigned to λ. For the present first-order review, they will therefore simply be taken as being identical; even though some degree of error may be inherent in doing so.

Without delving into overhead support electronics and redundancy considerations, it appears that the dormancy factor (D) is the only term that may differ on a gross scale among the technologies. To see whether and how this occurs, it is necessary to evaluate D according to

$$D = \frac{K_\lambda}{1 + \left(\frac{K_\lambda - 1}{K_N K_t} \right)}$$

where:

$K_\lambda = \lambda_{active} / \lambda_{dormant}$ - (failure rate ratio)

$K_N = N_{total} / N_{active}$ - (device count ratio) [during t_{active}]

$K_t = t_{mission} / t_{active}$ - (reciprocal duty factor)

Note that, by virtue of the manner in which they are defined, each of these "K" terms must satisfy the inequality:

$$1 \leq K_i < \infty$$

Under the complementary stipulations that stored data cannot be sacrificed and that the memory will always contain some data, the K_N and K_t factors for volatile memory structures will both be equal to unity. The corresponding value of D is therefore 1.0, and no dormancy related improvement in reliability is realized. Conversely, for nonvolatile memory devices, neither K_N nor K_t will be equal to unity.

During quiescent periods when the memory is neither inputting nor outputting data, both nonvolatile MNOS and MBD storage modules would be completely powered down. Both technologies would presumably be subjected to the same quiescent intervals, so the K_t factor of each would increase by a like amount and no difference would exist in the degree of reliability improvement thus secured. Comparison of the K_N factors, on the other hand, reveals that, due to the MBD module having twice as many chips active

during IO operations, the MNOS based module will enjoy a slightly better dormancy related improvement in reliability (on the order of 12 percent).

D. 4 SUMMARY

Upon drawing together the diverse technology evaluation factors which have been considered, evaluating them in light of the guideline requirements set forth for secondary mass memories, and comparing the performance, operation, and mechanization tradeoffs which are available, it is found that the single most suitable candidate is MNOS. In all categories reviewed, MNOS meets or exceeds by a wide margin the requirements of the postulated memory systems. Furthermore, it ranks an unqualified first in most of the categories. Table D-4 provides a summary rundown of the relative ranking of MNOS, MBD, and CCD based memory modules.

TABLE D-4

RELATIVE RANKING OF MASS MEMORY TECHNOLOGIES

Memory Characteristic	Technology Ranking		
	MNOS	MBD	CCD
Cell Density	2	3	1
Access Time	1	3	2
Data Rate	1	3	2
Storage Integrity	1	2	3
Reliability	1	2	3
Power Drain	1	2	3
Production Cost	1	2	3
Environmental Stability	1	3	2
Physical Size/Weight	1	2	3

Ranking Explanation:

1 - Best suited to application
2 - 2nd best suited to application
3 - Least suited to application

An MNOS memory cell can be made extremely small. No storage capacitor is required as in dynamic memories. MNOS devices currently being developed have cell sizes of under 0.75 mil^2. As photolithographic capabilities improve, MNOS density will improve. Block-oriented devices are capable of efficiently transferring blocks of data at high rates because of on-chip parallel transfer of data to a shift register. The shift registers may be operated at 6 MHz over the temperature range. In contrast, the data rate of most standard MOS RAM's is about 2 MHz. CCD minimum shift rates are limited by transfer efficiency. To overcome the problem, refresh stages are installed, but this results in significant loss in maximum rate. As a result, the upper limit on CCD operating rates is typically less than 3 MHz.

MNOS provides a system advantage in that it avoids some of the off-chip overhead of other technologies. For example, CCD and MOS systems must have refresh logic. MBD devices are dependent on off-chip electronics for operation and have a large ratio of off-chip components to memory chips. To take maximum advantage of nonvolatility, MNOS circuits will generally use power switching circuitry. These circuits are low power and usually consist of a few transistors per printed circuit card. In contrast with MBD which requires special interface buffer circuits, MNOS interfaces readily and directly with standard LSI logic families.

The final and perhaps most important consideration is cost. Mature production quantities of memory products from all the major technologies are determined by die size and manufacturing yield. Cost advantages on a per-bit basis at the component level hinge on the number of bits per chip. Little advantage exists for any of the main contenders at this point. At the system level, costs are affected by the amount and type of support circuitry needed and the ease of module assembly. MNOS provides benefits in both areas compared to either MBD or CCD. Procurement costs on the order of

0.08 cents/bit should be realized. Ultimately, the cost of ownership for an MNOS secondary mass memory will be lower than that of competing technologies due to its higher performance, lower power, and higher reliability.

APPENDIX E

MNOS TECHNOLOGY DESCRIPTION

The information presented herein gives the Metal-Nitride-Oxide Semiconductor (MNOS) memory technology a broadbrush treatment. It is intended for purposes of familiarization. All comments and discussions relating to MNOS structures are derived from current practices at Westinghouse Electric Corporation.

E.1 EVOLUTIONARY OVERVIEW

Beginning with the earliest reported work with compound, dual dielectric MOS structures in which a nitride (Si_3N_4) layer was introduced into the normal Al-S_iO_2-S_i sandwich of modern integrated circuits, the potential of an embryonic memory technology was recognized.[1][2] Over a period of years immediately following its discovery, the new technolgy went through a laboratory budding stage in which its physical theory was explored in detail. It subsequently passed from a curious materials phenomenon to functional test vehicle memory devices. Lack of endurance and good retention characteristics plagued early developments, however, and it was not until the invention of the drain-source-protected FET structure by Cricchi

(1) S.M. Hu; "Proerties of Amorphous Silicon Nitride Films;" Journal of Electrochemical Society, vol. 113; July 1966; pp. 673-688.

(2) S.M. Hu, D.R. Kerr, and L.V. Gregor; "Evidence of Hole Injection and Trapping in Silicon Nitride Films Prepared by Reactive Sputtering;" Applied Physics, Letters, vol. 10; February 1967; pp. 97-99.

in 1972 that the great promise of MNOS passed from an engineering gleam in the eye to a practical reality.[3]

MNOS memory technology has undergone continuous development by a number of different organizations in several countries over the past decade. Figure E-1 provides partial listings of both currently active firms and some of the diversified MNOS applications. Collectively, the companies cited are engaged in producing a variety of MNOS products. Although millions of LSI memory devices have, as a result, been delivered and are now in end-use field serivce, the full potential of the MNOS technology has only recently become generally understood. This is largely due to its development emphasis being oriented toward military applications rather than commercial uses. Consequently, it has not been broadly publicized.

As a technology base, MNOS encompasses a tremendous realm of applications. It actually consists of a number of subtechnologies aimed at specific functions; namely:

- Electrically Alterable ROM's (EAROM)
- Block Organized RAM's (BORAM)
- Random Access Memories (NOVRAM)
- Sequential Access Memories (NOVSAM)
- Nonvolatile LSI Logic (NOVLOG)

Various hybrids representing technology cross-breeds also exist. Of these, the ones of greatest import are:

- Charge Addressed Memories (NOVCAM)
- Dielectric Substrate RAM's (DISRAM)
- Volatile Intermittently RAM's (VINRAM)

---

(3) J.R. Cricchi, F.C. Blaha, and M.D. Fitzpatrick; "Drain-Source-Protected MNOS Memory Device and Memory Endurance;" Technical Digest of 1973 International Electronic Devices Meeting; December 1973; pp. 126-129.

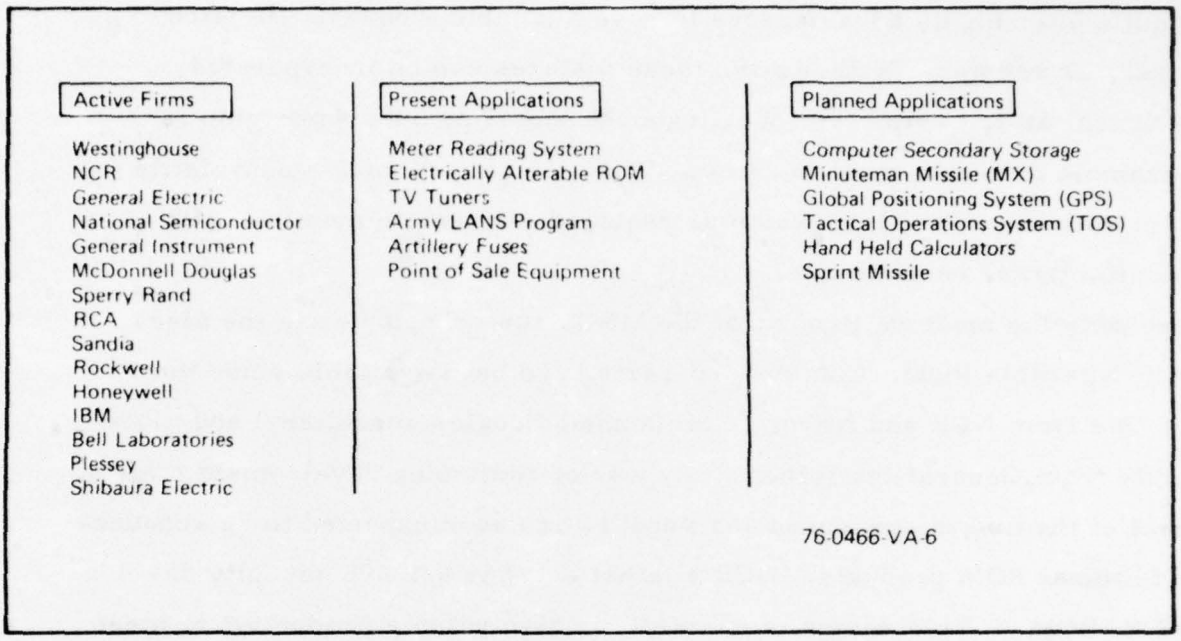

Figure E-1. Developers and Applications of MNOS Technology

In NOVCAM structures a marriage of CTD (charge transfer device) principles with MNOS processing is made. A similar technology marriage is realized with DISRAM devices, in which all MNOS elements are dielectrically isolated (rather than junction isolated). The most recent addition to the list, VINRAM, combines the properties of a nonvilatile backing store with those of a volatile working storage medium on a common monolithic chip.

Each of the MNOS subtechnologies institutes nonvolatile storage via some form of field effect transistor (FET), but the respective transistor designs and functional characteristics in certain cases are quite different. In addition, the memory element structure and manufacturing sequence often varies widely from company to company. A wide variance in on-chip peripheral circuitry also exists among some of the categories. Universally, though, the ability to form memory and nonmemory transistors on the same chip is an important asset. This is manifested by the fact that in logic systems

it is quite often highly advantageous to have available a nonvolatile latch, register, or counter. With MNOS, these features can be incorporated directly on an LSI chip. The Westinghouse Metering Subsystem chip is one example of such a mechanization. This device combines a nonvolatile counter with the other logic functions required for remote monitoring of residential power consumption.

Probably the most well known of the MNOS subtechnologies is the electrically alterable ROM. Commercial parts have been available since the early 70's from NCR and Nitron (a McDonnell Douglas subsidiary) and more recently from General Instrument. By way of continuing development efforts, several of the major semiconductor vendors are soon expected to be announcing additional ROM products. NCR's latest entry is a 4,096 bit fully decoded EAROM.* This device is currently in high volume production to meet system requirements within NCR.

EAROM type devices are expressly designed to emphasize the length of data retention but, as a result, require comparatively long write times — typically 10 to 100 miliseconds. Their read time is a function of on-chip interface circuit design. Although access and cycle times can be made quite small, the major applications for such devices do not require it and it is generally not done. The NCR part access time, for instance, is under 2.0 microseconds. In contrast, a current Nitron 1,024 bit part has a typical access time of nearly 20 microseconds. Neither part was specifically optimized for this parameter. Nitron is, however, planning to come out with a series of much faster devices.

At the opposite end of the spectrum, NOVRAM constitutes a special class of MNOS device expressly designed for very high speed operation. For such devices, write as well as read speeds of below 1.0 microsecond are commonly required. Short retention times on the order of minutes are, however, normally considered adequate. The reason for this is simply that the

*EAROM is a trademark of NCR.

prime economic driving force behind their development has been military needs for radiation hardened nonvolatile RAM's in missile systems. Because of the high speed and radiation considerations, special on-chip circuit designs and device fabrication techniques (including dielectric isolation) are employed. Power dissipation and bit density are secondary considerations. In the next generation Minuteman missile, MNOS NOVRAM memories are expected to replace the heavy and expensive plated wire memories used in present systems.

Intermediate on the speed/retention-time scale, MNOS BORAM and NOVSAM are very close in terms of their memory transistor and data IO interface structures. They differ mainly in their address and control features. Depending on the particular device and application constraints, clearing and writing operations are accomplished in the range of 10 to 200 microseconds. Retention time varies with writing time over this range but is normally prescribed to be on the order of a few days to several years. At the I/O ports of these devices, much higher data rates are maintained than at the individual memory cell level. Information is stored in an X-Y array with multiple bits written or read simulataneously. Data bit blocks are bidirectionally loaded between the storage matrix and a set of on-chip shift registers which transfer data into or out of the chip serially at up to 6 MHz. With access to the first bit requiring no more than 1 microsecond, readout of an entire block of 64 or less sequentially ordered, parallel-bit data words can be accomplished in under 12 microseconds.

BORAM devices are intended for secondary and high-end primary storage applications where data is typically handled in blocks of about 1, 024 parallel-bit bytes (1 byte in this instance is taken to mean 8 bits). A BORAM module will thus range in size from 10^6 to 10^9 bits. NOVSAM devices are intended for secondary and low-end tertiary storage applications involving larger block sizes. The capacity of this type module will fall in the range of 10^8 to 10^{12} bits. Figure E-2 identifies the manner in which these MNOS

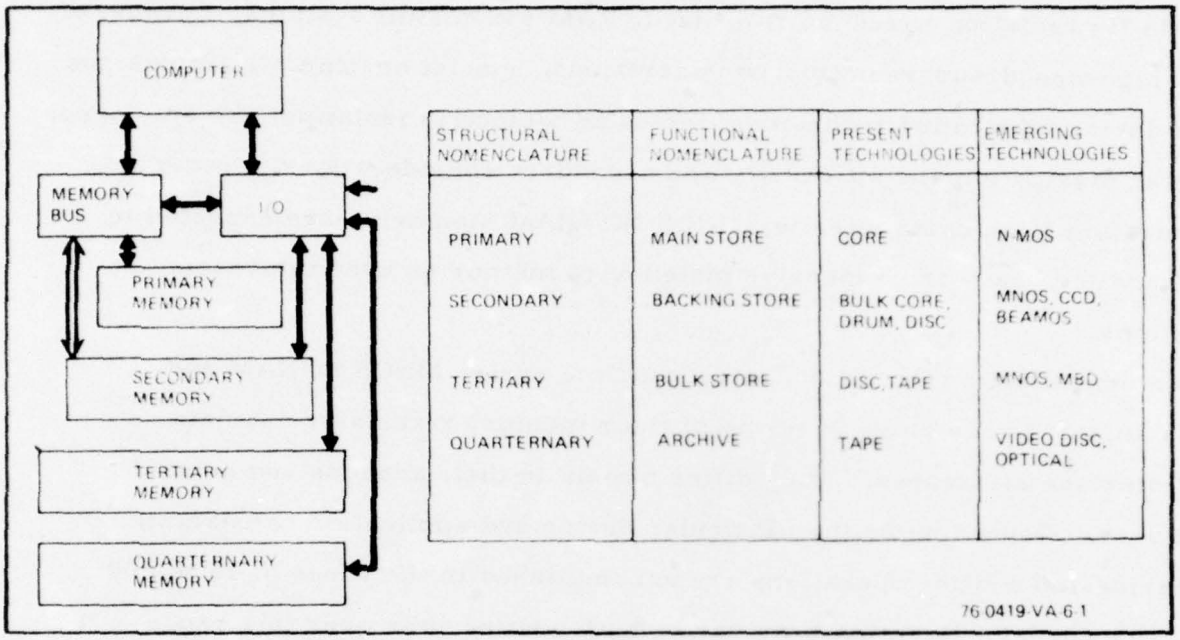

STRUCTURAL NOMENCLATURE	FUNCTIONAL NOMENCLATURE	PRESENT TECHNOLOGIES	EMERGING TECHNOLOGIES
PRIMARY	MAIN STORE	CORE	N-MOS
SECONDARY	BACKING STORE	BULK CORE, DRUM, DISC	MNOS, CCD, BEAMOS
TERTIARY	BULK STORE	DISC, TAPE	MNOS, MBD
QUARTERNARY	ARCHIVE	TAPE	VIDEO DISC, OPTICAL

Figure E-2. Computer System Memory Hierarchy Classification

technology memories fit into the general hierarchial computer storage classifications.

Of all the MNOS subtechnologies, BORAM has received the greatest attention due to the pressing military need to replace unreliable electromechanical secondary storage devices in critical field applications. At least two BORAM programs are presently underway. The US Navy is sponsoring a module development with Univac (a Sperry Rand subsidiary) in St. Paul, Minnesota. The US Army Electronics Command (ECOM) is sponsoring a similar module development with Westinghouse in Baltimore, Maryland. Recently the Naval Air Systems Command (NASC) joined with ECOM in sponsorship of further BORAM applications programs.

At Westinghouse, BORAM system development has proceeded along the lines of establishing a joint services, multimission module design. In October 1975, the first prototype module was delivered to ECOM with a

partial memory population of 152 first-generation 2 Kbit MNOS BORAM chips. This module, shown in figure E-3, has since been retrofit with a complement of MHP's (multichip hybrid packages) containing second-generation (6000C) 2 Kbit chips. Both the Army and the Navy have undertaken BORAM module evaluations. For information on Army tests, inquires should be routed through the ECOM technical monitor ---Mr. E.J. Gallagher, DRSEL-NL-BP-1, (201-544-2213), Computer Techniques and Development Team, COMM/ADP Lab, Fort Monmouth, N.J. 07703. Information on test results at the Naval Air Test Center (NATC) can be obtained from Mr. F.A. Phillips, CODE SY43, (301-863-4787), Computer Technology Group, Systems Engineering Test Directorate, NATC, Patuxent River, Md 20670.

The 2 Kbit MNOS BORAM 6000C nonvolatile integrated circuit memory chip deployed in the Westinghouse developed joint services module possesses exceptionally stable and highly reproducible characteristics. It employs a single level of metal interconnect pattern and requires only seven active masks.

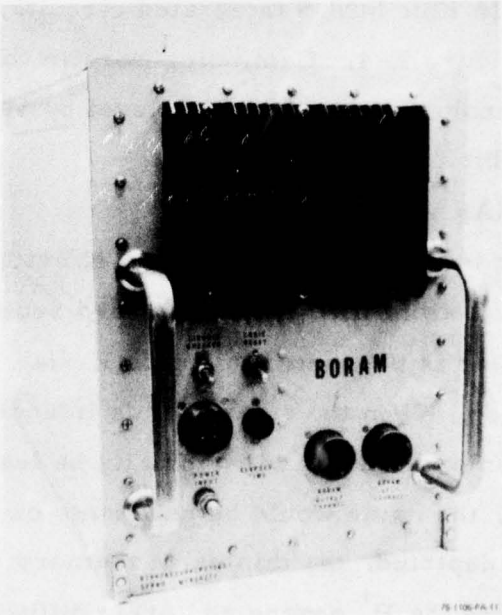

Figure E-3. Army-Navy BORAM Module

Its fabrication sequence is believed to be the simplest presently in existance for the manufacture of fully decoded MNOS arrays. The two-transistor per cell, differentially balanced detector approach taken in the 6000C circuit design tolerates wide variations in both initial process parameters and in final operating voltages. Yield experience for initial quantity production lots has proven consistently high. The advent of this highly producible 6000C chip provides the final element confirming the commercial viability of the MNOS technology.

As a corollary effort, development of a family of similar, higher density MNOS devices - extending in capacity from 8 to 64 Kbits - has been instituted at Westinghouse. This work has led to a major simplification in the memory cell and processing mask configurations. Orthogonal straight-line geometries, with no offsets, pads, or elbows of any kind are used exclusively to form the memory array. This facilitates realization of the much higher cell packing densities necessary with larger capacity devices. With this arrangement, photolithographic tolerances are not critically determined by contact windows or vias in the storage area. By applying these along with other advanced circuit design techniques, this family progression program has recently succeeded in producing a 16 Kbit MNOS integrated circuit. A microphotograph of the die is shown in figure E-4. Continuing work in this area is scheduled to bring the MNOS technology to the 128 Kbit level by 1980, with an ultimate long-range goal of 1 Mbit per chip.

E.2 FUNCTIONAL CHARACTERISTICS

An MNOS transistor is an insulated gate field effect transistor (FET) with a two layer insulator. A simplified unscaled cross section of a p-channel MNOS memory transistor is presented in figure E-5(a). The nitride layer is typically 300 to 500 Å. When the transistor is intended to be a memory cell component, the thin oxide layer will typically be less than 50 Å. For a nonmemory transistor, the oxide would be uniformly on the order of 1000 Å.

In the configuration depicted, the thin oxide memory region is located some distance away from the P^+ source and drain diffusions. Thick nonmemory oxides are employed in the interspersed regions. This results in

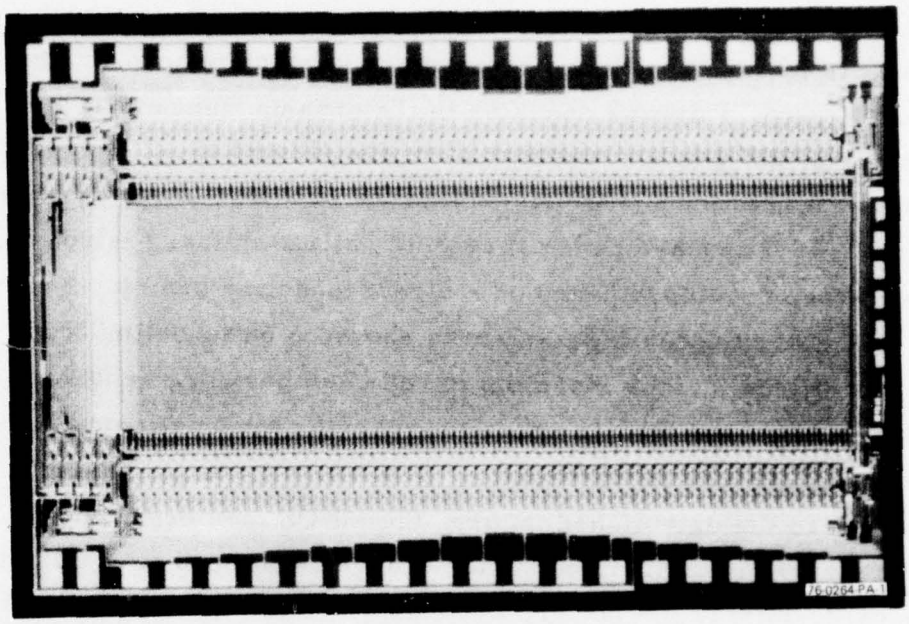

Figure E-4. Microphotograph of 16 Kbit MNOS Chip

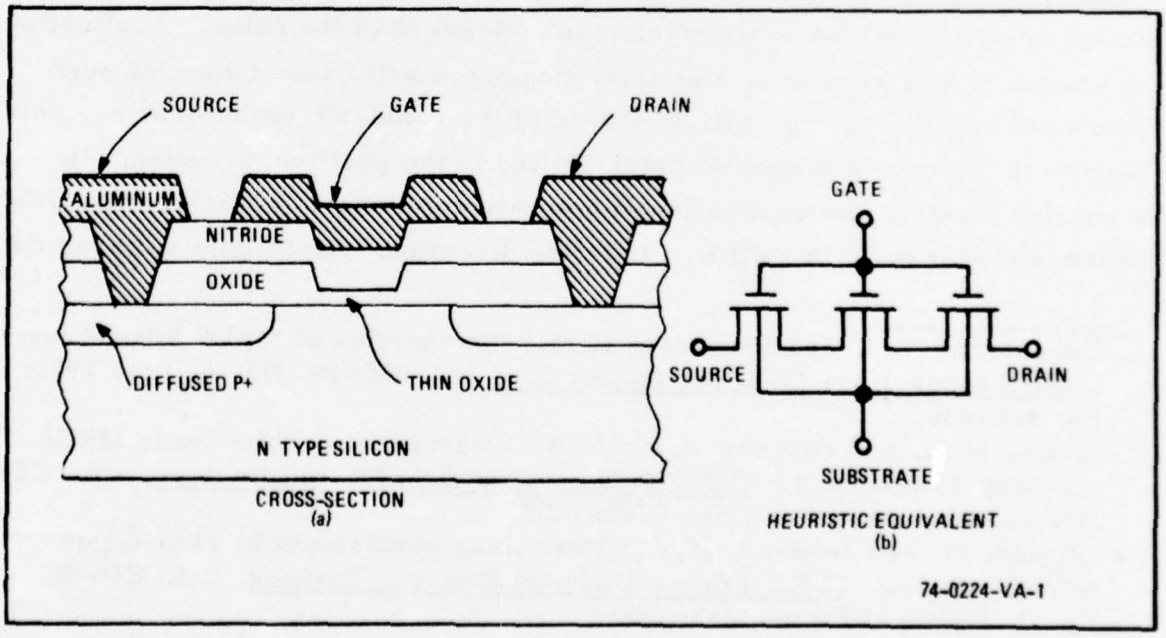

Figure E-5. Basic MNOS Memory Transistor

the structure exhibiting functional equivalence to the three-transistor model of figure E-5(b). The two outer transistors are nonmemory devices which can tolerate large gate-to-drain or gate-to-source voltages. The central memory device can be pulsed into a depletion mode positive threshold state, but since it is in series with two nonmemory enhancement mode devices, the overall structure always operates in the enhancement mode.

A number of independent workers have contributed to the formulation of the mathematical models which explain the behavior of MNOS memory transistors.[1, 2, 3] Expressions for threshold voltage shift and retention characteristics have been developed based on a direct tunneling theory. From these models, the charge decay mechanism is shown to be tunneling from deep traps to the silicon. Since tunneling is not a temperature sensitive process, both retention time and device operating margins are confirmed as being relatively temperature independent.

According to the established theory, the MNOS memory transistor provides storage by injecting or removing charge from deep trap sites located near the oxide-nitride interface. Because a relatively large amount of energy is required to lift the charge out of these traps, long retention times can be achieved. Charges are transported to the oxide-nitride interface by tunneling through the thin oxide. Current flow in the nitride during this action is negligible due to its being much thicker than the oxide. Application of a positive gate voltage of sufficient magnitude will cause tunneling such that a net negative charge will be stored at the oxide-nitride interface. This results in the device threshold being shifted in the positive direction. In a similar fashion, the application of a negative voltage of sufficient magnitude causes the threshold to shift in a negative direction. Small gate voltages do

---

1. Lundström, K. I. and Svensson, C. M.; "Properties of MNOS Structures;" IEEE Transactions On Electron Devices, vol. ED-19, No. 6; June 1972; pp. 826-836.
2. White, M. H. and Cricchi, J. R.; "Characteristics of Thin-Oxide MNOS Memory Transistors;" IEEE Transactions On Electron Devices, vol. ED-19, No. 12; December 1972; pp. 1280-1288.
3. Gordon, N. and Johnson, W. C.; "Switching Mechanism in Thin-Oxide MNOS Devices;" IEEE Transactions On Electron Devices, vol. ED-20, No. 3; March 1973; pp. 253-256.

not disturb the threshold, so a nondestructive readout of the memory device is readily realized.

The basic drain-source-protected (DSP) transistor structure of figure E-5 is now standard in MNOS integrated circuit memories. It provides continuous enhancement mode operation, minimizes voltage restrictions on the source and drain, and exhibits virtually no threshold window closure with clear-write cycling. Life tests have in fact revealed no significant degradation after 10^{12} clear-write cycles. First generation MNOS memory chips were built using the basic DSP structure directly in a single transistor per memory cell array. Although fully functional, full capability devices were produced, the fabrication yield achieved proved to be unacceptably low due to sensitivity to processing nonuniformities. The circumvent this condition, a refined two-transistor (2T) DSP memory cell was developed and applied in second-generation MNOS memories. Figure E-6 reveals the structural details of the cell. As a result of the exceptional stability, processing insensitivity, and high fabrication yields secured with the 2T cell approach, it has been adopted as the fundamental building block in all on-going MNOS product developments at Westinghouse.

Layout details of the 2T cell structure and the manner in which it is integrated into an X-Y matrix to produce a memory array are set forth in figure E-7. This figure represents a planar view of a two-by-two section taken from the composite processing masks of a current high-density MNOS memory chip. In evidence here are the major geometrical factors contributing to the high cell density and fabrication ease of modern MNOS integrated circuits. In particular, observe that the entire storage area design rules are based totally on orthogonal, straight-line layout principles. No offsets, elbows, or stepped pads of any kind are employed. It is this geometric simplicity which provides a two-transistor per cell memory array with a per cell area of less than 0.75 square mil without approaching photolithographic limits on readily resolvable lines and spaces.

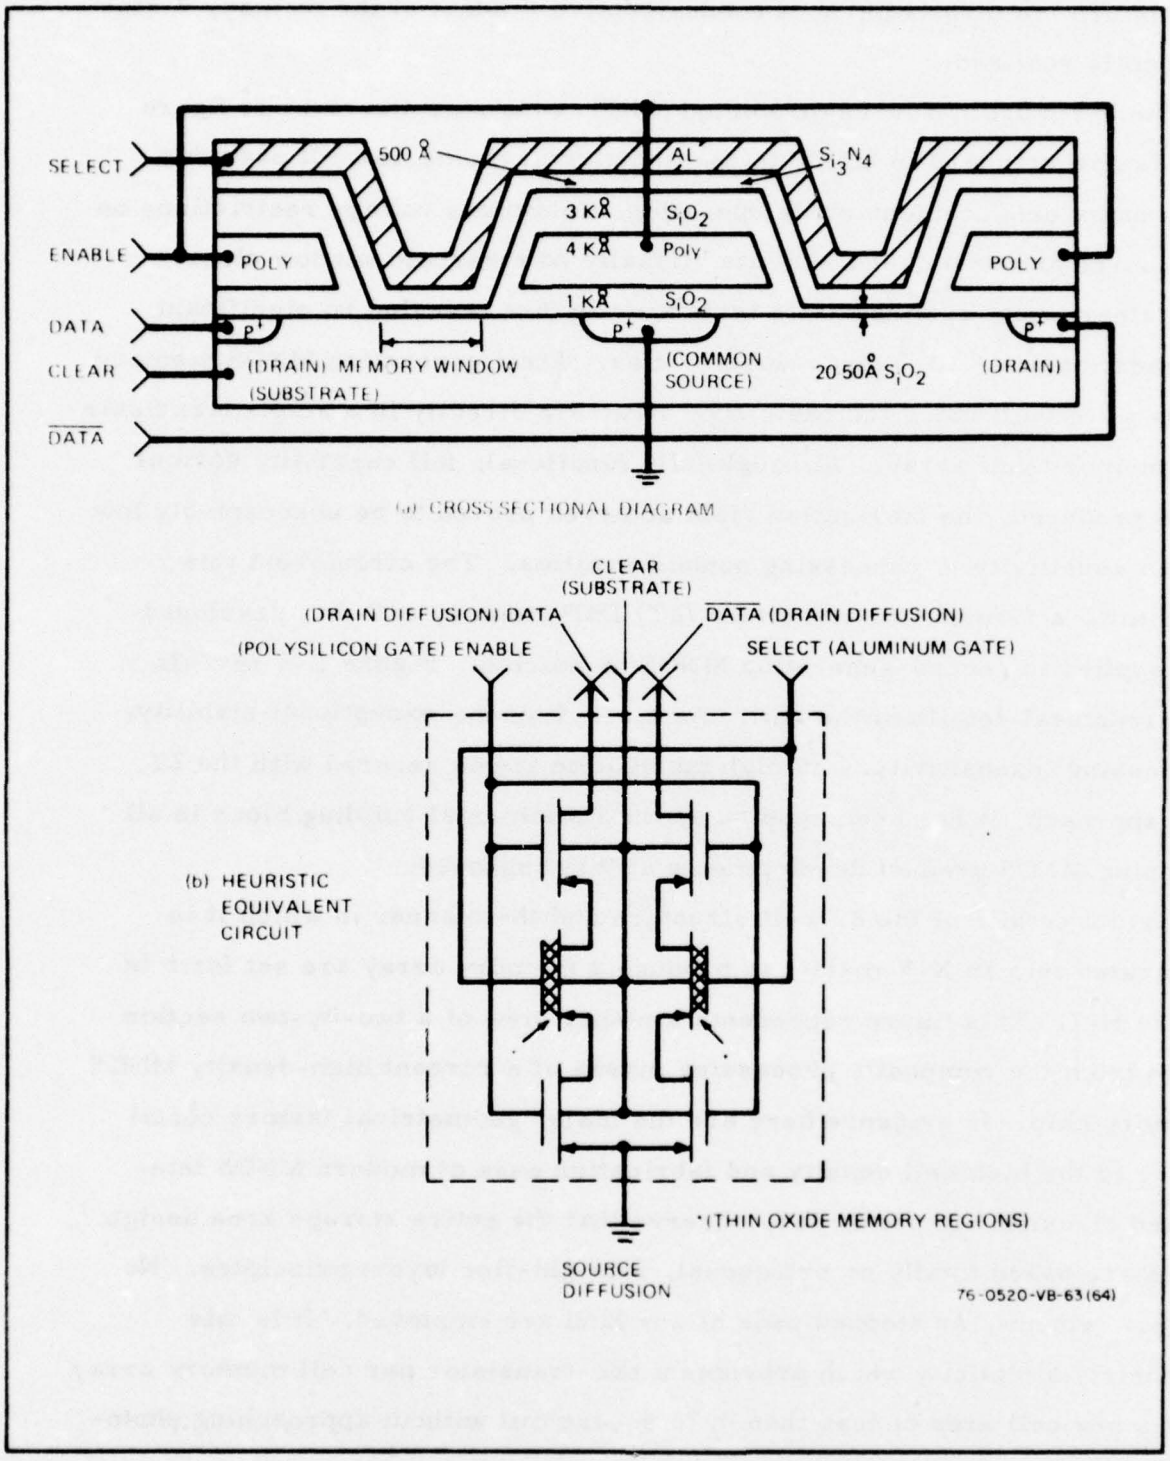

Figure E-6. Two-Transistor (2T) P-Channel MNOS Memory Cell

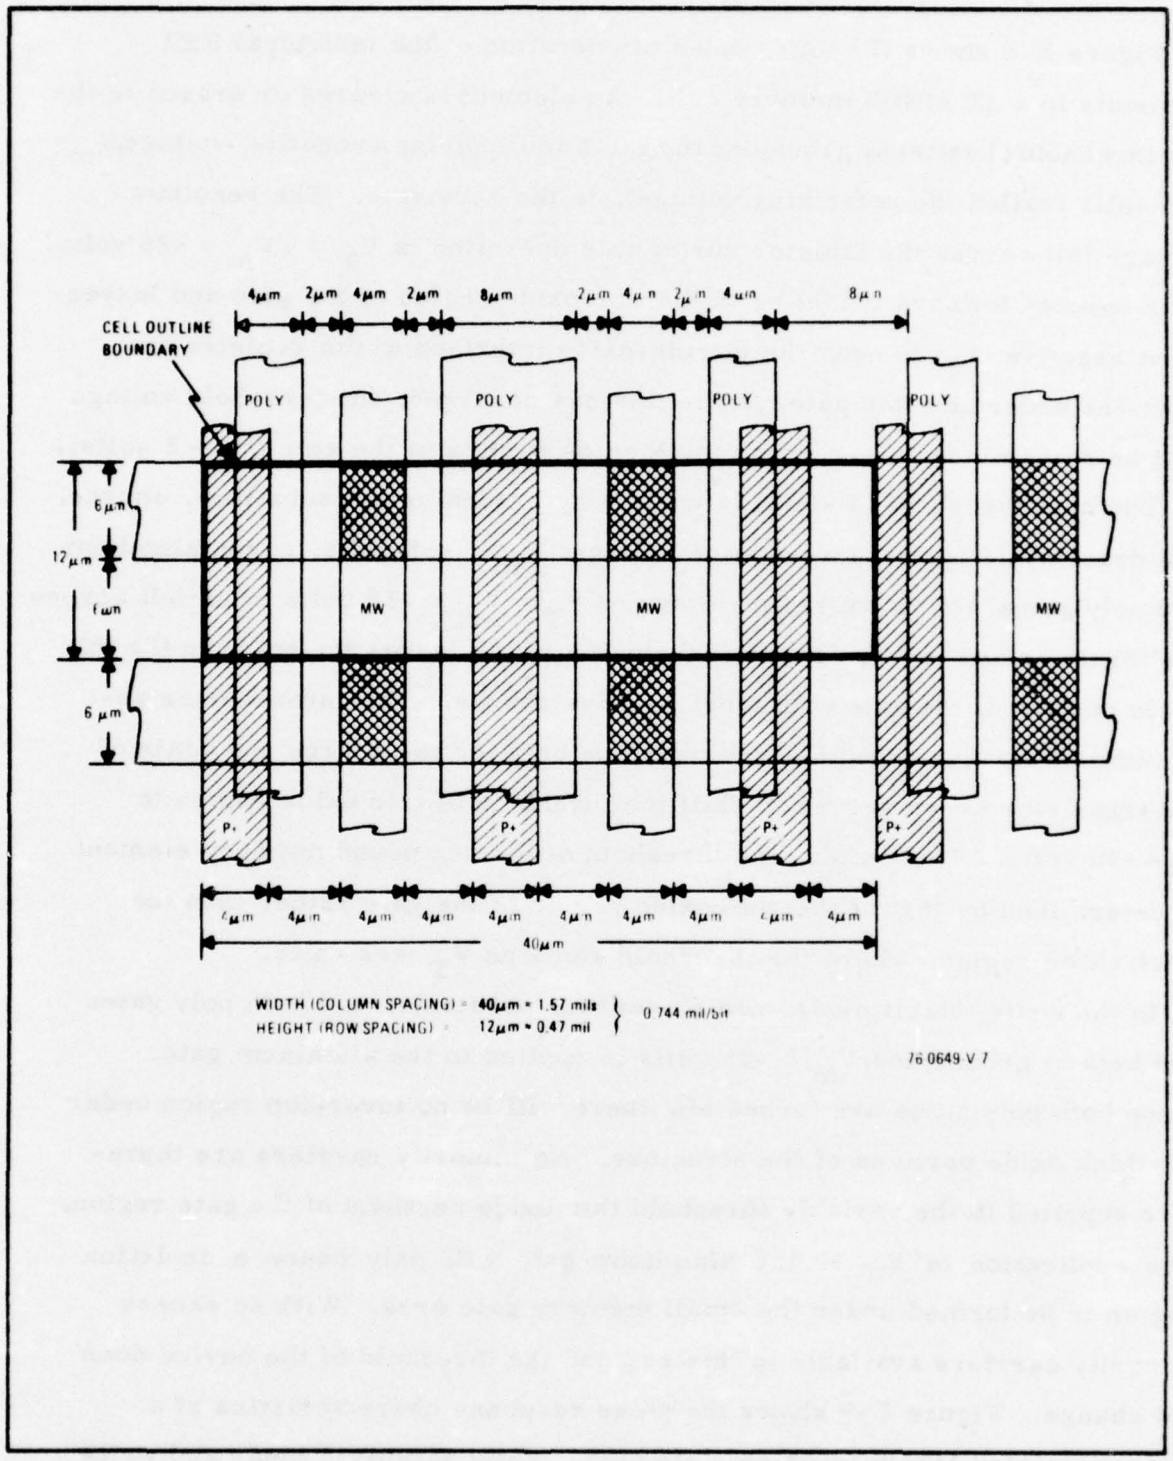

Figure E-7. 2T MNOS Memory Matrix

Figure E-8 shows the four modes of operation of the individual FET elements in a 2T MNOS memory cell. An element is cleared or erased to the low threshold (1 state) by grounding the gates and applying a negative voltage, V_m = -25 volts (called the polarizing voltage), to the substrate. The resultant voltage felt across the isulator during this operation is $V_T = +V_m = +25$ volts. This induces forward tunneling in the thin oxide region of the gate and leaves a net negative charge near the nitride-oxide interface which depletes the carriers under the thin gate. Once this has occurred, the threshold voltage will be determined solely by the thick oxide portion of the gate ($V_T \sim 2$ volts).

The high threshold (0 state) is written by grounding the substrate, source, and drain while applying a negative voltage, $V_m = -25$ volts, to the aluminum and poly gates. This causes a voltage of $V_T = V_m = -25$ volts to be felt across the insulator and induces reverse tunneling which leaves the traps in the thin oxide portion of the gate with a net positive charge. This produces an accumulation layer along the silicon surface between the source and drain of the transistor which serves to shift the threshold of thin oxide region to $V_T \sim -10$ volts. Thereafter, the threshold of the compound memory element is determined by that of the thin oxide region of the gate rather than the thick oxide region, where the threshold remains $V_T \sim -2$ volts.

In the write inhibit mode, drain, source, substrate, and both poly gates are held at ground and $V_m = -25$ volts is applied to the aluminum gate. Since both poly gates are turned off, there will be no inversion region under the thick oxide portions of the structure. No minority carriers are therefore supplied to the variable threshold thin oxide sections of the gate region. The application of V_m to the aluminum gate will only cause a depletion region to be formed under the small memory gate area. With no excess minority carriers available in this region, the threshold of the device does not change. Figure E-9 shows the pulse response characteristics of a typical isolated MNOS transistor element. Fully saturated clear and write times take on the order of 1 to 10 milliseconds. When deployed in a balanced

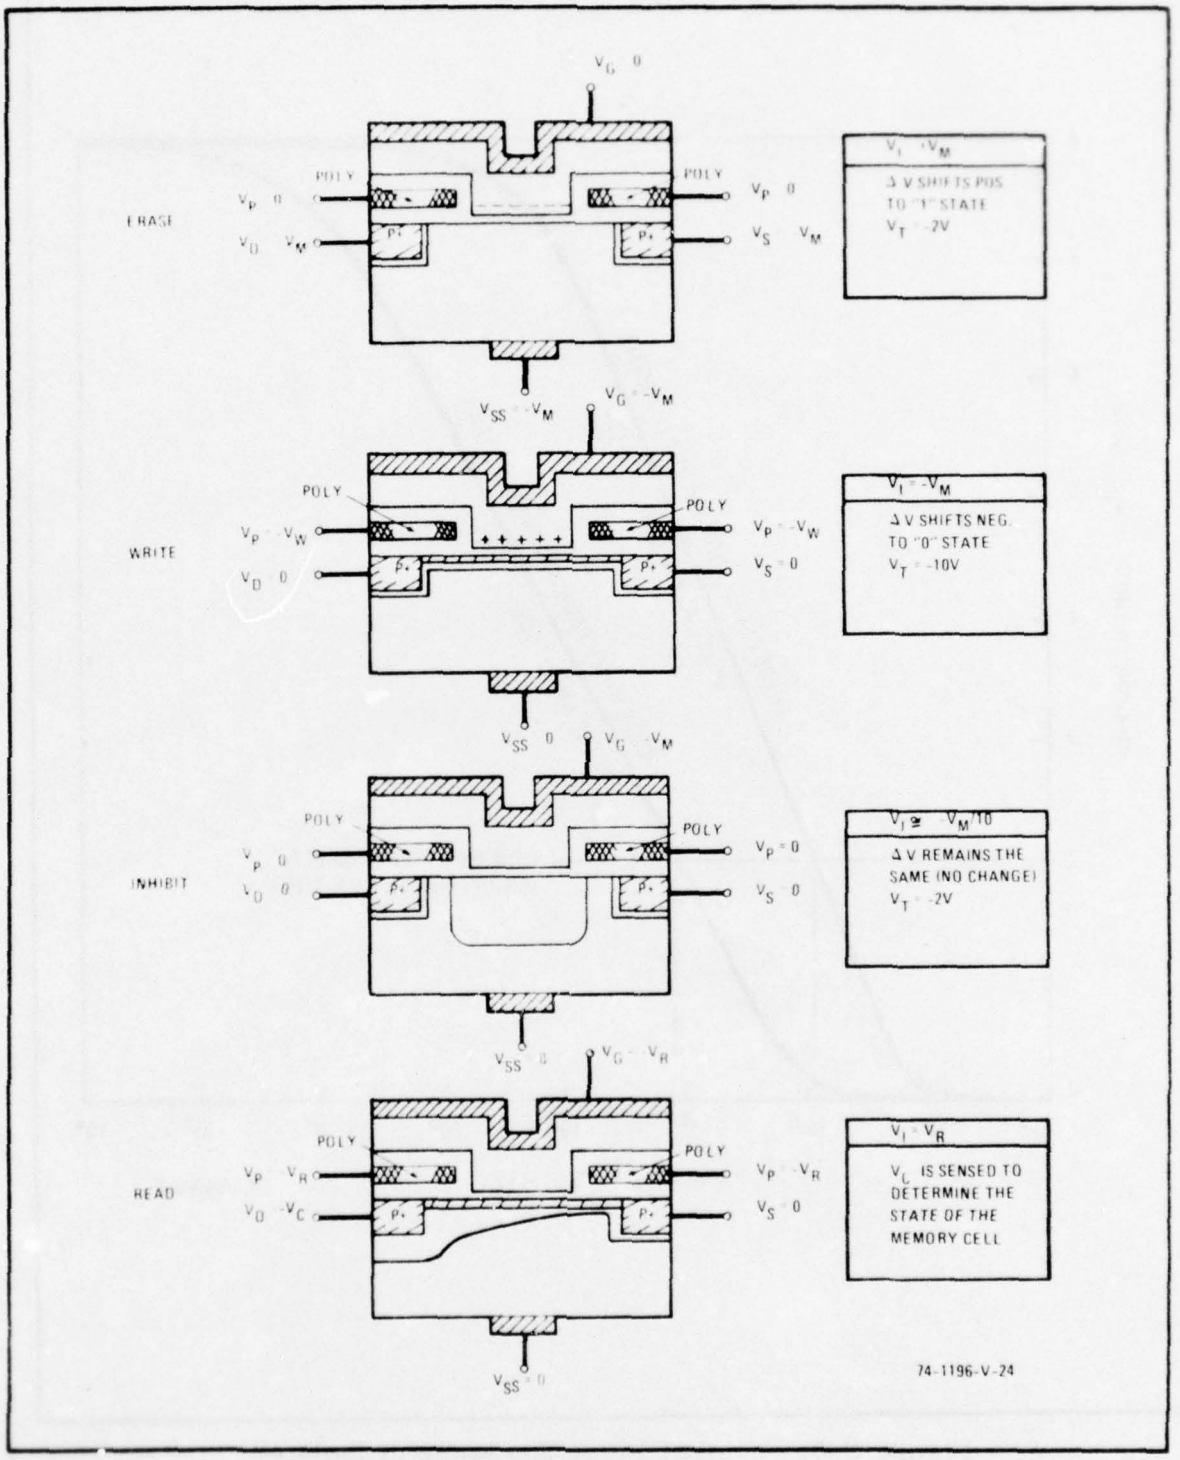

Figure E-8. Operating Modes of Individual MNOS Memory Element

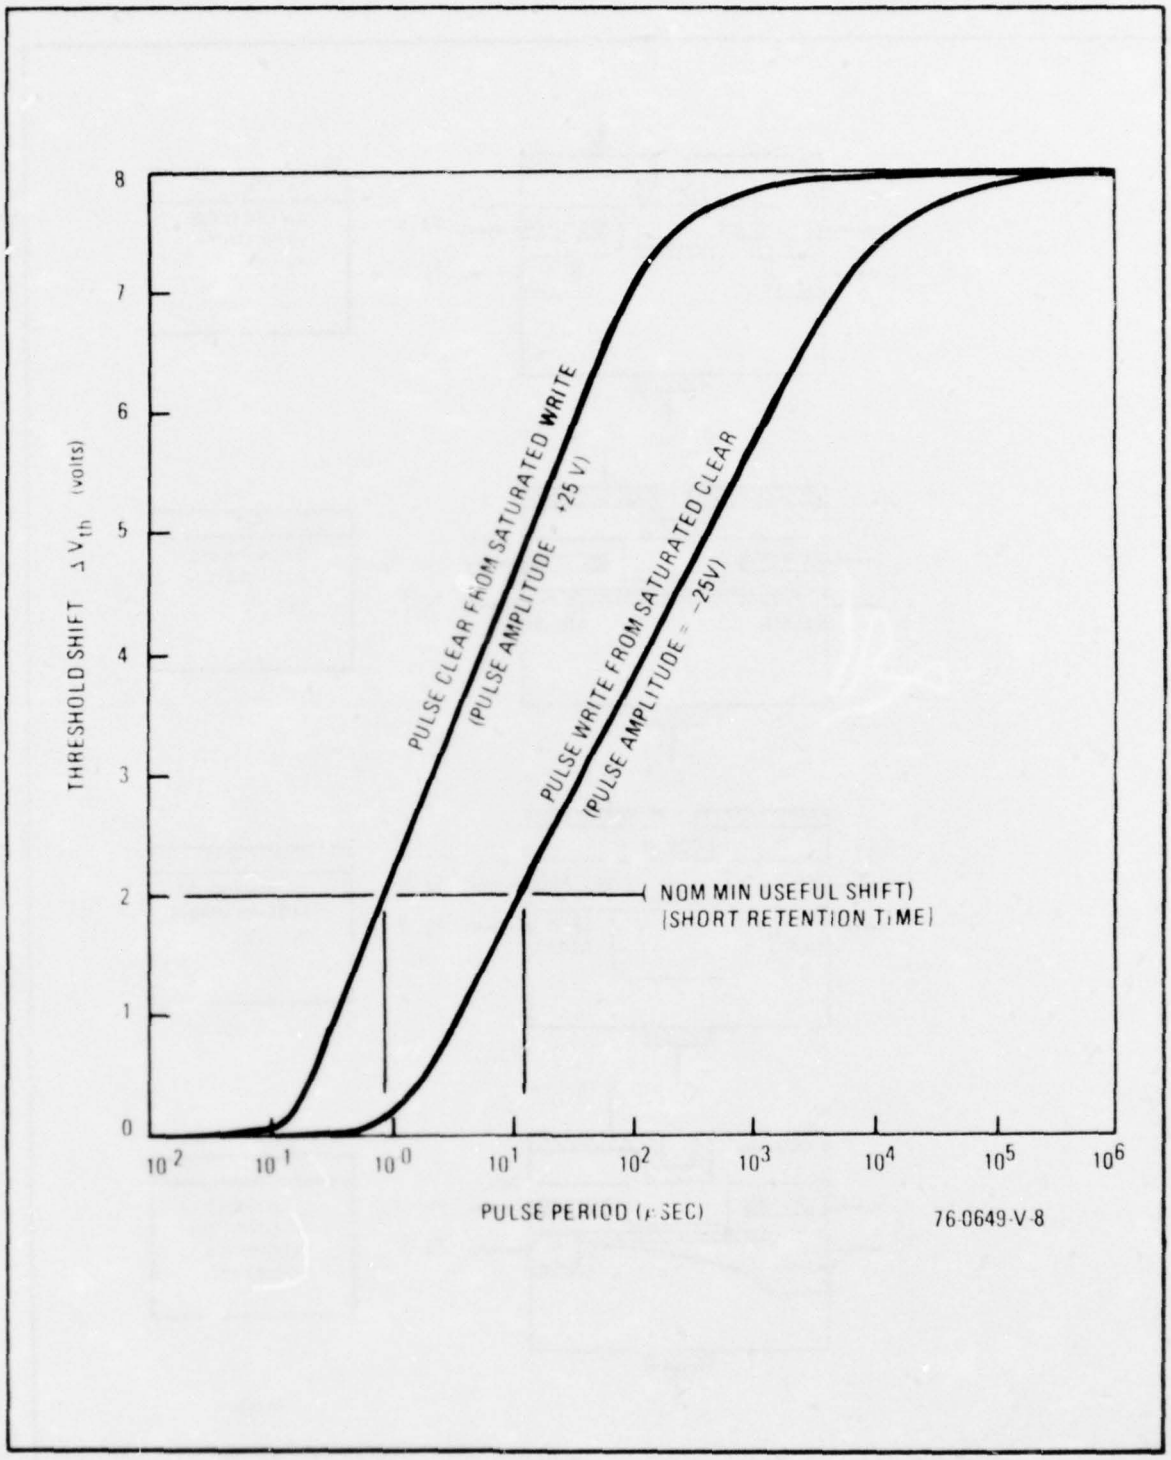

Figure E-9. Typical Pulse Response of Isolated MNOS Memory Element

2T memory cell, useful operation with retention times of a few days can be secured with 1 microsecond clear times and write times as short as 10 microseconds.

To read the MNOS memory element, the drain is first precharged to $V_C = -15$ volts, and the source is grounded. The drain is then allowed to float and $V_R = -6$ volts is applied to both the poly and the aluminum gates. Since this exceeds the -2 volt threshold of the thick oxide poly controlled regions of the gate, an inversion layer is formed in these areas. If the memory element is in the low threshold state, an inversion region will also be formed under the thin oxide portion of the gate and cause the drain to be discharged toward ground. If the memory element is in the high threshold state, no inversion region will be formed under the thin gate and the drain voltage will remain charged at $-V_C$.

By virtue of using a differential detection circuit in conjunction with the close LSI matching of devices in the 2T memory cell, current MNOS memories approach closely the theoretical limits on data retention. This can be see from figure E-10. Here the decay slopes for a matched set of MNOS elements in a 2T cell are plotted. From the high state, decay proceeds at about a quarter of a volt per decade of time. From the low state, the decay is approximately twice as fast, or about a half of a volt per time decade. Projecting these slopes produces an intersection at around 10^8 minutes, which is roughtly 190 years. Allowing for even a very sloppy 400 to 500-millivolt minimum differential detection window, the realizable retentivity extends to at least 3×10^7 minutes, or approximately 50 years. This is for a fully saturated write time. Shorter write times produce correspondingly shorter retention times.

At all capacity levels, high-speed MNOS memory chips are implemented in a manner functionally analogous to that originally evolved for and refined

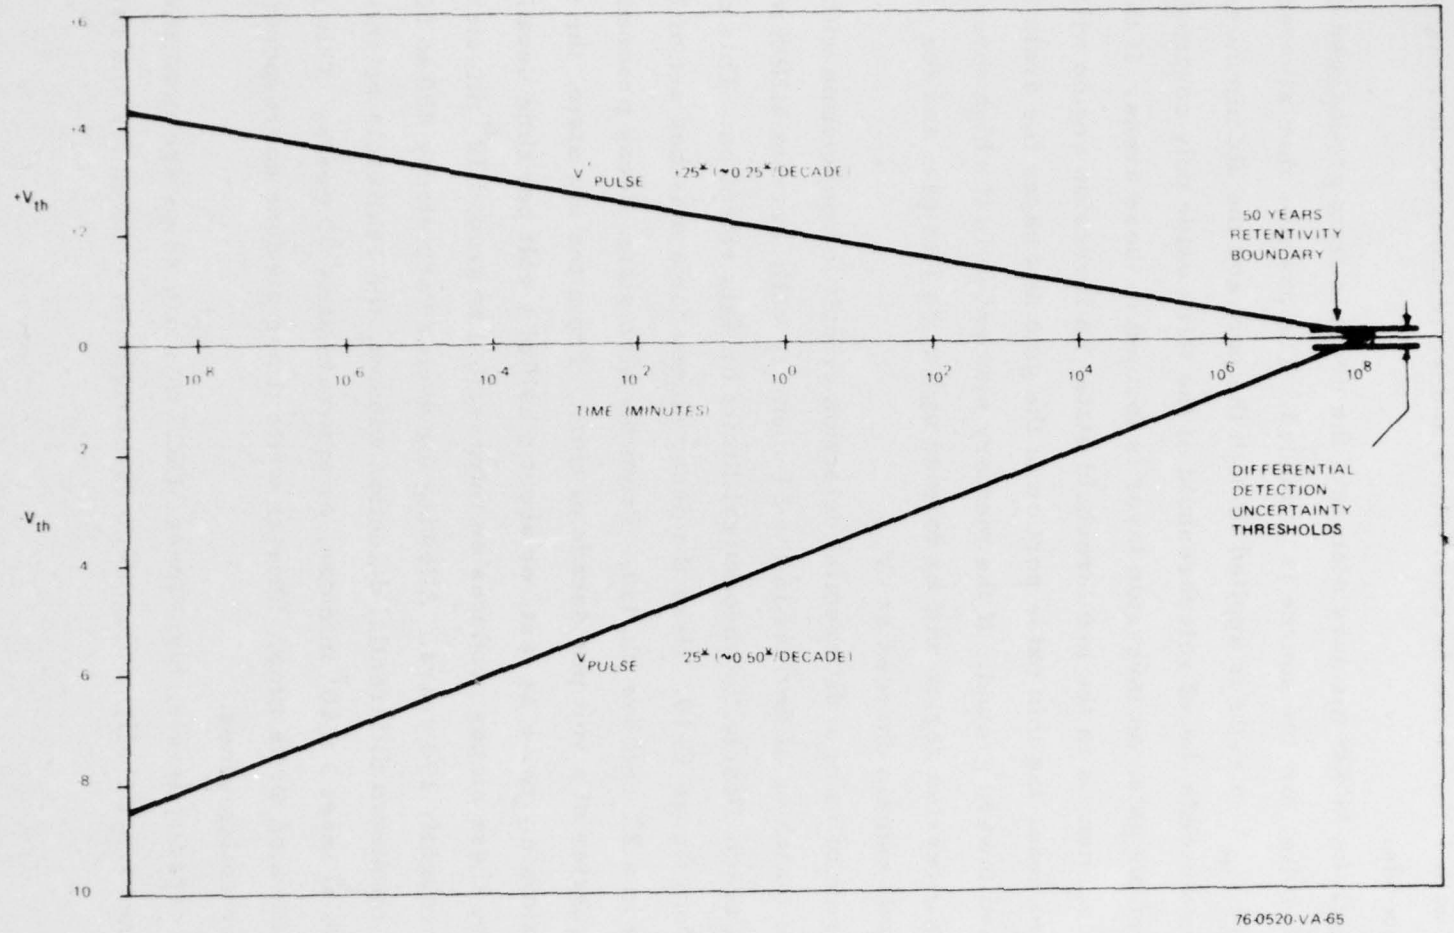

Figure E-10. Retention Characteristics of 2T MNOS Memory Cell

on the BORAM program.[4] In each, careful attention is given to the important aspect of on-chip peripheral circuit design to ensure providing a practical memory product. Care is taken in the array interface area to ensure that the gate to substrate potential of elements not being addressed is held solidly at zero volts. This prevents nitride leakage effects from affecting retention time. Especially configured read circuits are employed to guard against the possibility of read disturb effects. Writing and clearing circuit designs are based on considerations of both static and transient conditions so that safe operating margins are maintained at the individual MNOS transistor gates. Since MNOS writing and clearing functions require voltages on the order of 20 to 30 volts, on-chip interface buffers are incorporated to simplify the off-chip driver requirements.

Figure E-11 presents a block description of the most recent Westinghouse 16 Kbit MNOS integrated circuit memory configuration. The differences between it and earlier 2 Kbit MNOS BORAM chips - besides higher capacity - lie in the array interface and address structures and in the ability to clear on a per-row basis as well as on a full-chip bases. Just as with BORAM, though, clearing is executable any time before writing and is asynchronously command initiated for the most recently enabled row address. Data access to the memory array is provided by the two dynamic shift registers and the set of transfer switches. Each shift register stage controls alternate column bit lines via the respective transfer switches. The designs for these array interfaces are based on industry proven varactor coupled MOS scanning circuitry.[5]

---

(4) J.E. Brewer and D.R. Hadden, Jr.; "Block-Oriented Random Access MNOS Memory;" National Computer Conference; May 1974.

(5) R.E. Joynson, et.al.; "Eliminating Threshold Losses in MOS Circuits by Bootstrapping Using Varacter Coupling;" IEEE ISSC, vol. SC-7, No. 3; June 1972; pp. 217-224.

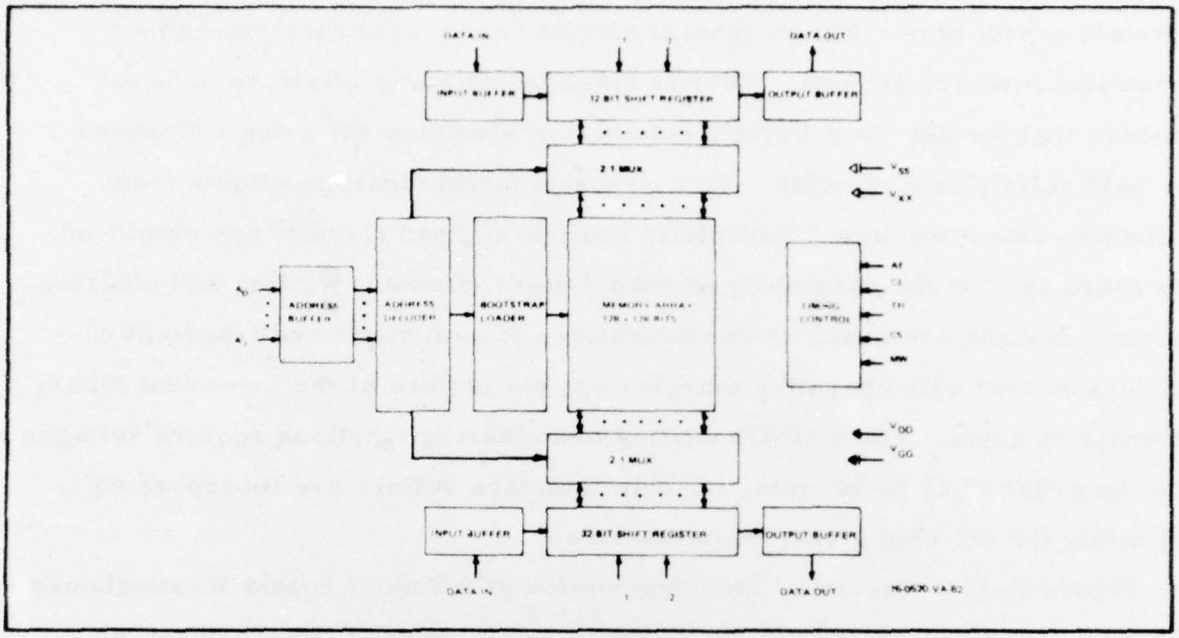

Figure E-11. Westinghouse 16 Kbit MNOS Memory

In operation, address signals (A0-A7) are applied to the chip and pass through the address buffer and decoder to select the row containing the desired block of data. Once the row address is enabled (by AE), it is decoded and applied to the selected row of the storage matrix. The columnized data contents of the cells along the selected row are then either loaded out into the I/O registers during a read mode access or loaded in from these SR's during the write mode. During this data exchange between the registers and the storage matrix, the register clocks are held quiescent so that they dissipate no CV power. Accessed blocks of data in this particular memory chip contain a total of 64 bits. Through the serializing action of the I/O shift registers, each block is interfaced to the outside world in the form of 32 serial words which are each 2 parallel bits wide.

As a generic class, MNOS memory devices are fundamentally MOS structures, and as such share the same failure modes and exhibit the same level of failure rate. Test data available to data reveals no pronounced differences in active mode failure rates between similarly sized MNOS and MOS parts. However, it is recognized that when electrical stress is removed from an MOS part its failure rate (dormant mode) drops by at least an order of magnitude. This has significant impact on memory systems in which only a few chips are involved in data transactions at any given time or which have extended periods of inactivity during which their stored data must be maintained. If the chips are volatile, they must all be in the active mode in order to retain data. If the chips are nonvolatile, only the chips engaged in the data transaction need be active. This results in a reliability advantage for a nonvolatile system which becomes more dramatic for larger systems and/or longer mission times.

Obviously, a power switched nonvolatile system will have a power advantage over any competitive volatile memory. Total system power will be determined by how many chips must be turned on simultaneously to meet data rate requirements. If minimum power were the single overriding design objective, it is feasible to build a 10^9 bit MNOS memory module with less than 1 watt of dissipation. In less esoteric applications, the benefits of low power are reduced power supply cost, weight, and volume; and lower temperature operation. This latter feature enhances reliability and reduces the cost, weight, and volume problems associated with heat removal.

As was mentioned earlier, MNOS charge storage is not a temperature sensitive phenomenon. Consequently, the operating temperature range is established by the design margins allowed in the on-chip peripheral circuitry rather than by leakage in the memory cell. No critical temperature problems therefore exist, and operation at 125°C is readily achievable. The radiation resistance of MNOS memory cells has also proven to be quite good. Dielectrically isolated MNOS parts are in fact presently being developed for use in

radiation hardened systems. The limiting feature of these devices for radiation tolerance has been shown to be the on-chip peripheral circuitry. Even without special hardening though, it is generally found that the peripherals will withstand greater than 10^5 rads total dose. This is, of course, a function of the circuit design and must be confirmed by tests on individual part types.

Upon reviewing all facets of the MNOS technology, it is seen that its major merits for use in building memory systems stem principally from its comparatively simple LSI structure and the intrinsic nonvolatility of its basic storage elements. In summary, the advantages citeable for MNOS include:

- Nonvolatile Storage
- High Bit Density
- High Reliability
- High Data Rate
- Low Power Drain
- Radiation Resistance
- Temperature Stability
- Direct Logic Interfaces
- Minimum Off-Chip Overhead
- Low-Cost Batch Fabrication.

APPENDIX F

MNOS 2 KBIT MEMORY CHIP SPECIFICATIONS

The set of specification sheets enclosed herewith describe the operational characteristics and functional capabilities of a 2,048 bit two-transistor cell MNOS nonvolatile integrated circuit memory, chip developed by Westinghouse Electric Corporation. This device is now offered in volume for fully deployed military and DOD memory systems. The specifications given represent the initial production guaranteeable capabilities of the device and are somewhat conservative in certain areas. A family of similar, higher capacity MNOS memory IC's, extending in size up through 64 Kbits, is presently under development.

MNOS 2 Kbit Nonvolatile Integrated Circuit
Block Oriented Random Access Memory Chip

The chip contains a fully decoded 64-word by 32-bit random access memory and two dynamic 2-phase 16-bit shift registers. All data I/O takes place serially through the shift registers. The shift registers and RAM may operate independently. Data is transferred in parallel between the RAM and registers via an external 32-bit latch. Data output drivers are three-state devices capable of sinking a low power TTL gate load. This 2048-bit memory is intended for application in computer secondary storage systems. It is normally packaged in multi-chip hybrid form. Use of MNOS (metal-nitride-oxide-semiconductor) technology allows nonvolatile information storage and low power operation. Bonding pads are $> 5\ mils^2$, and are positioned on opposite sides of the die. The chip measures 163 by 169 mils. A glass overcoat guards against scratches due to handling. Protective devices on all inputs avoid damage by static charge.

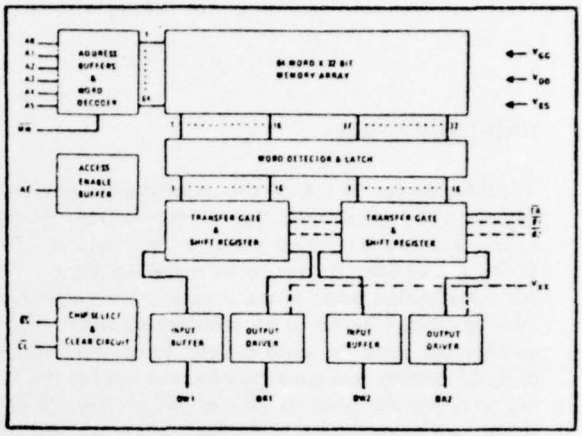

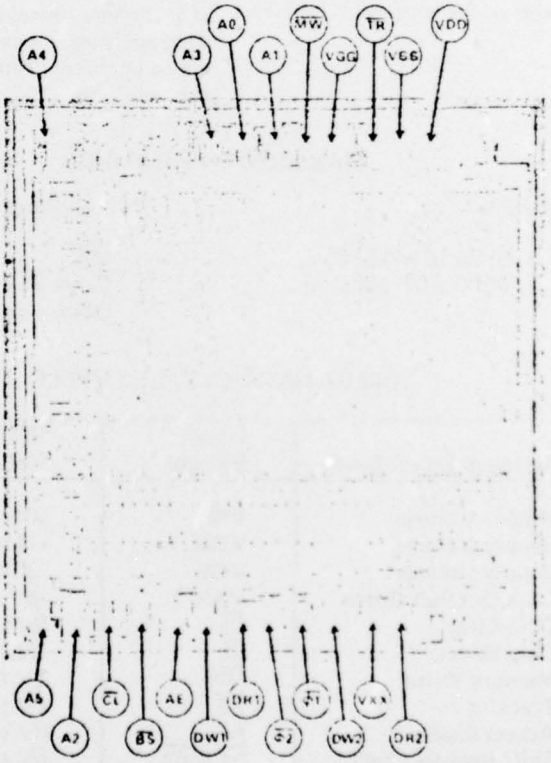

Westinghouse Electric Corporation
Defense and Electronic Systems Center
Baltimore, Maryland

Revised
76-0649-V-1

This page, although not completely legible, is included for information purposes only.

MNOS 2 Kbit Nonvolatile Integrated Circuit
Block Oriented Random Access Memory Chip

RAM OPERATION

The RAM portion of the 2K chip is particularly easy to use because it is fully decoded, has no external read reference voltage and has no critical power supply voltages. RAM timing is controlled by the access enable AE signal. While AE is low critical nodes in the array and word detector are clamped. After the A0 to A5 addresses are stable, AE may go high and allow the word detector to operate. If $\overline{TR}$ is high, the detector will sense the addressed word in the memory array and will latch. If $\overline{TR}$ and $\overline{\Phi 2}$ are low, the detector will sense the contents of the shift register and will latch. If $\overline{MW}$ is held low, the contents of the latch will be written into the memory array.

CLEARING OR INITIALIZING

In large block oriented memory systems it is an advantage to be able to erase an entire data block at one time. The $\overline{CL}$ pulse clears all 2048 cells in the chip of previously stored data. PN junction isolation separates the memory array from the shift registers and control circuitry. To avoid forward biasing the isolation junctions, the high state of $\overline{CL}$ must be < a silicon diode drop, VD, more positive than VSS. VSS should be more positive than all other chip terminals. Chip clear is gated internally by the chip select $\overline{BS}$, and unselected chips will not be affected by $\overline{CL}$. After clearing, the threshold levels of all memory transistors will be shifted into the most positive state; and, from a logical point of view, the memory contents will be undefined. After writing, each bit will assume a one or zero state.

MAXIMUM CHIP RATINGS

Temperature Range

Operating	– 55°C to +125°C
Storage	– 65°C to +150°C

Voltage Range Referenced to VSS

$\overline{MW}$	+0.5V to -36V
$\overline{CL}$ & $\overline{BS}$	+0.5V to -34V
Other Inputs	+0.5V to -30V

OPERATING VOLTAGE LEVELS

Terminal Description	Signal Symbol	Voltage Level Referenced to VSS
Supply Voltage	VSS	0V
Supply Voltage	VDD	-15V ± 5%
Supply Voltage	VGG	-30V ± 5%
Data Out Pull Down	VXX	-5V (TTL), -15V (CMOS)
Chip Clear	$\overline{CL}$	0V to -30V
Chip Select	$\overline{BS}$	0V to -30V
Memory Write	$\overline{MW}$	0V to -30V
Transfer	$\overline{TR}$	0V to -15V
Access Enable	AE	0V to -15V
Shift Register Clocks	$\overline{\Phi 1}$ & $\overline{\Phi 2}$	0V to -15V
Data Write Inputs	DW1 & DW2	0V to -15V
Data Read Outputs	DR1 & DR2	0V to VXX
Address Inputs	A0 to A5	0V to -15V

Westinghouse Electric Corporation
Defense and Electronic Systems Center
Baltimore, Maryland

Revised
76-0649-V-2

MNOS 2 Kbit Nonvolatile Integrated Circuit
Block Oriented Random Access Memory Chip

MAXIMUM POWER REQUIREMENTS

Conditions: VSS = 0V, VDD = −15 V, VGG = −30 V, $T_A = 25^oC$

Electrical Parameter	Test Units	Operating Mode		
		$\overline{BS}$ = 0V	$\overline{BS}$ = −30V	
		Deselected	Read/Shift	Write/Shift
ISS	mA	+0.01	+7	+10
ICL	mA	+0.01	+2	+4
IDD	mA	-0.01	-6	-6
IGG	mA	-0.01	-3	-8
PDD	mW	0.2	90	90
PGG	mW	0.3	90	240
P_{CHIP}	mW	0.5	180	330

POWER VS TEMPERATURE

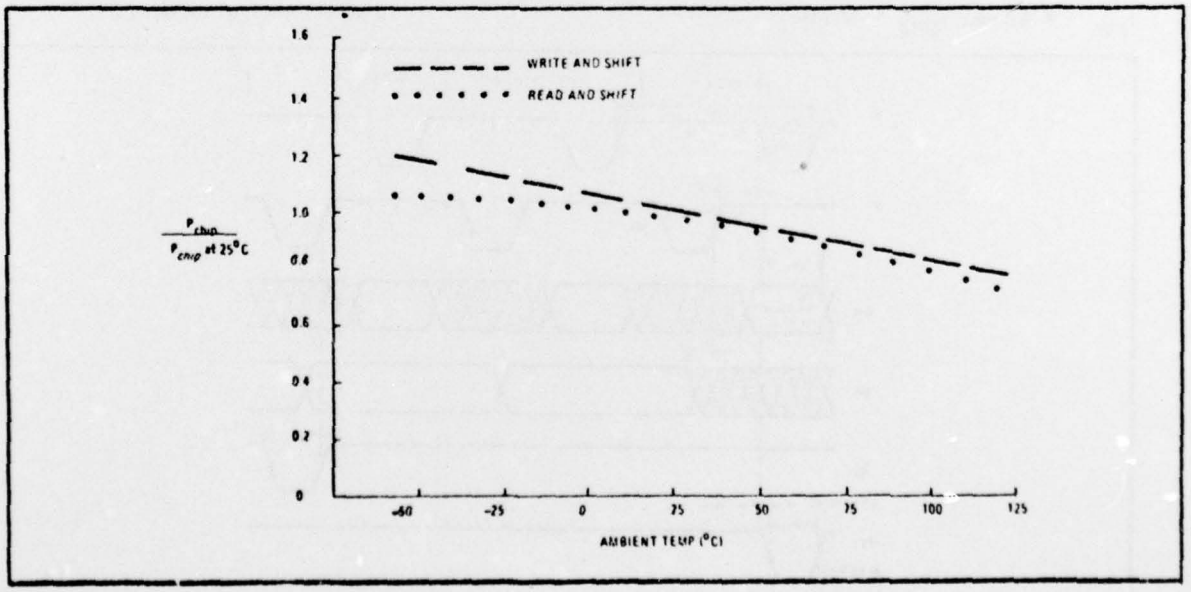

Westinghouse Electric Corporation
Defense and Electronic Systems Center
Baltimore, Maryland

Revised
76-0649-V-3

MNOS 2 Kbit Nonvolatile Integrated Circuit
Block Oriented Random Access Memory Chip

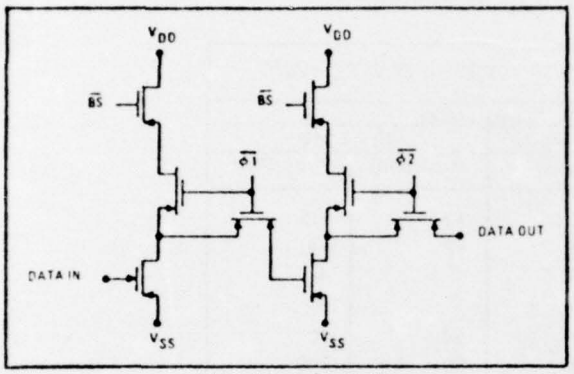

Shift Register Cell

SHIFT REGISTER OPERATION

The 2K-bit chip contains two 16-bit ratio-type dynamic 2-phase shift registers. The output stages feature a convenient pull-down terminal, VXX, to allow TTL compatibility. When $\overline{BS}$ is high, the DR terminals assume a high impedance state, which permits wired OR connections of output data lines. Input data lines will accept either CMOS or TTL level signals. Shift rates > 5 MHz can be achieved.

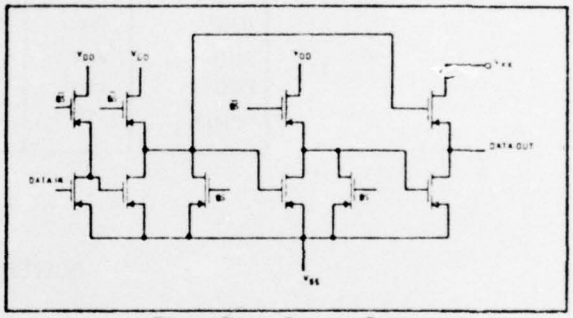

Three-State Output Stage

SYSTEM REGISTER OPERATING CONDITIONS

$t\phi \geq 200$ nsec, $td\phi \geq 20$ nsec

$PW_{\phi 1} = PW_{\phi 2} \geq 40$ nsec, $t_{rise} = t_{fall} \geq 20$ nsec

$t_{DWs} \geq 50$ nsec, $t_{DRd} \leq 50$ nsec

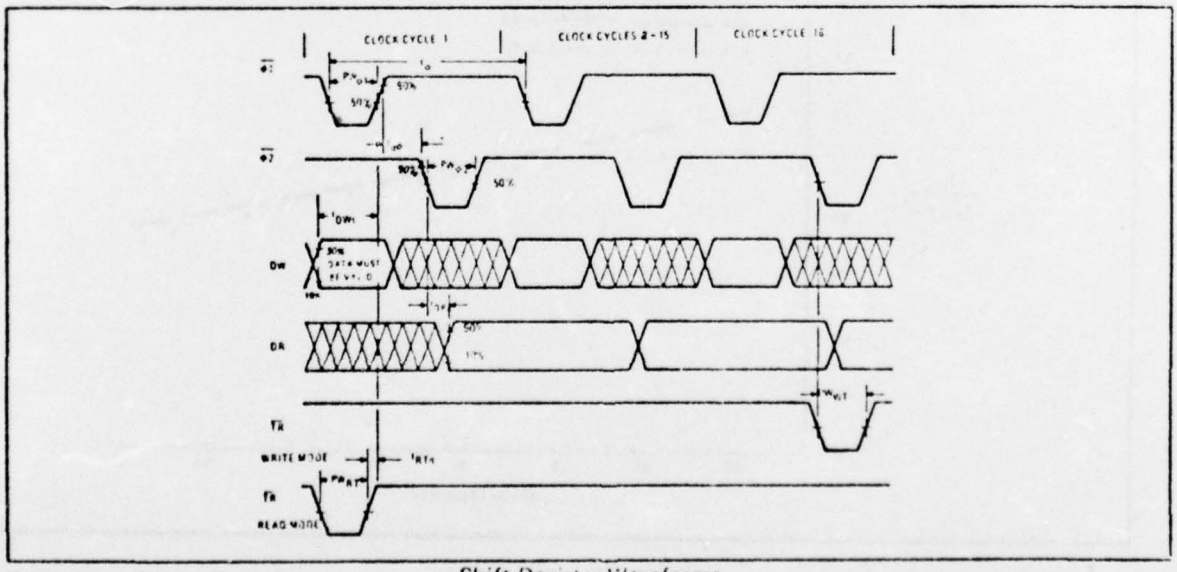

Shift Register Waveforms

Westinghouse Electric Corporation
Defense and Electronic Systems Center
Baltimore, Maryland

Revised
76-0649-V-4

MNOS 2 Kbit Nonvolatile Integrated Circuit
Block Oriented Random Access Memory Chip

SYSTEM CLEARING CONDITIONS

$t_{CLd} \geq 100$ nsec
$t_{fCL} \geq 1\ \mu$sec
$PW_{CL} \geq 100\ \mu$sec
$t_{rCL} \geq 1\ \mu$sec
$t_{Ad} \geq 0.0$ nsec

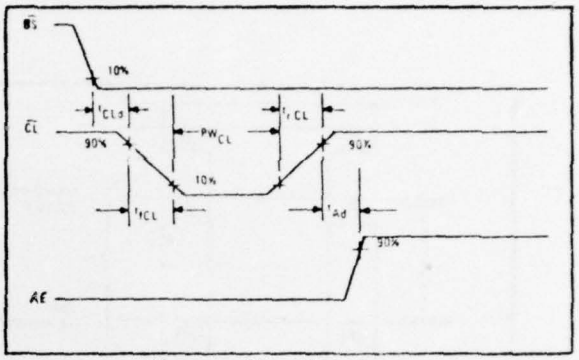

Memory Clear Waveforms

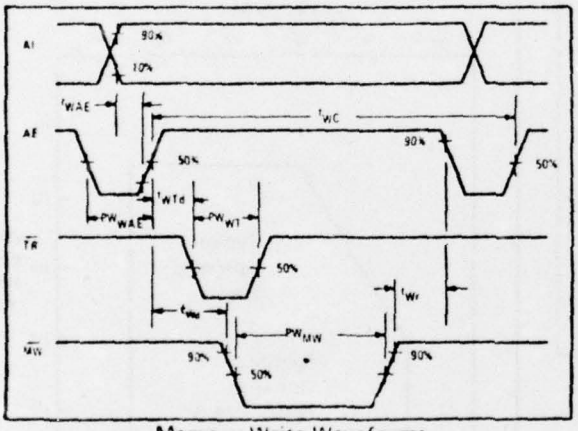

Memory Write Waveforms

SYSTEM WRITE CONDITIONS

$PW_{WAE} \geq 200$ nsec
$t_{WAE} \geq 100$ nsec
$t_{WTd} \leq 0.0$ nsec
$PW_{WT} \geq t_{wd} + 100$ nsec
$t_{Wd} \geq 200$ nsec
$PW_{MW} \geq 100\ \mu$sec
$t_{Wr} \geq 0.0\ \mu$sec
$t_{WC} \geq 101\ \mu$sec

SYSTEM READ CONDITIONS

$t_{BS} \geq 100$ nsec
$t_{RAE} \geq 100$ nsec
$t_{RTd} \geq 200$ nsec
$PW_{RT} \geq 200$ nsec
$t_{AEd} \geq 100$ nsec
$t_{AC} \geq 200$ nsec

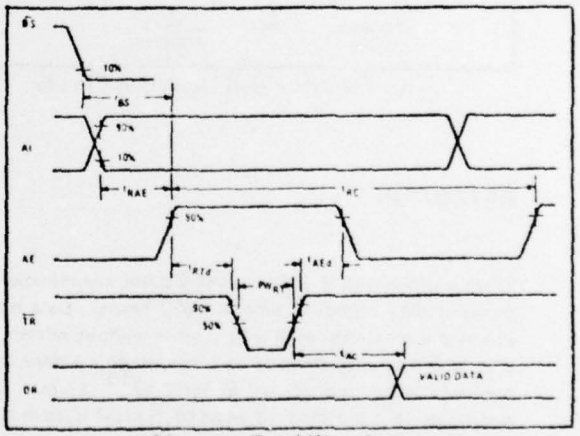

Memory Read Waveforms

Read conditions in a typical system do not approach the performance capability of the device. First bit access time (total) from stable address to stable data < 2 μsec can be achieved. Data can be shifted at 5 MHz, providing a 3.6 μ sec read cycle and an 8.8 MHz data rate.

TYPICAL READ CAPABILITY

$t_{RAE} = 200$ nsec
$t_{RTd} = 600$ nsec
$PW_{RT} = 400$ nsec
$t_{RC} = 3.6\ \mu$sec

Westinghouse Electric Corporation
Defense and Electronic Systems Center
Baltimore, Maryland

Revised
76-0649-V-5

MNOS 2 Kbit Nonvolatile Integrated Circuit
Block Oriented Random Access Memory Chip

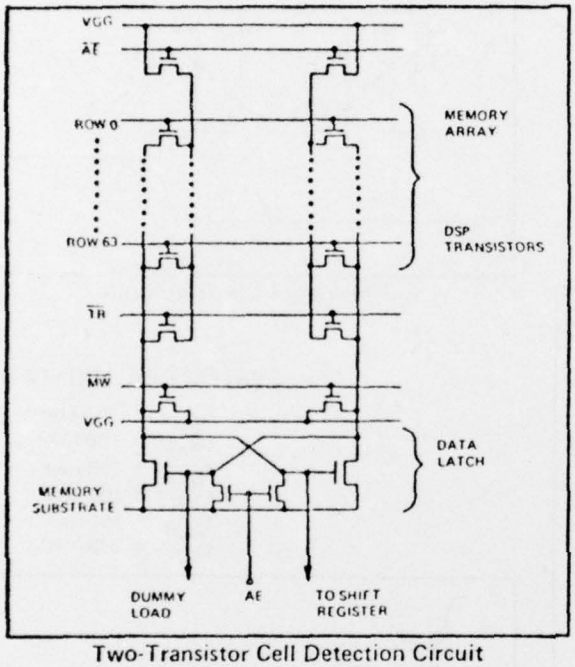

Two-Transistor Cell Detection Circuit

TWO-TRANSISTOR CELL

The MNOS 2K chip stores information in cells that consist of two drain-source protected (DSP) memory transistors. A balanced detection circuit provides high sensitivity and immunity from any power supply voltage sensitivity problems. No external read reference voltages are required, and chips do not have to be matched or sorted in any way for use in systems.

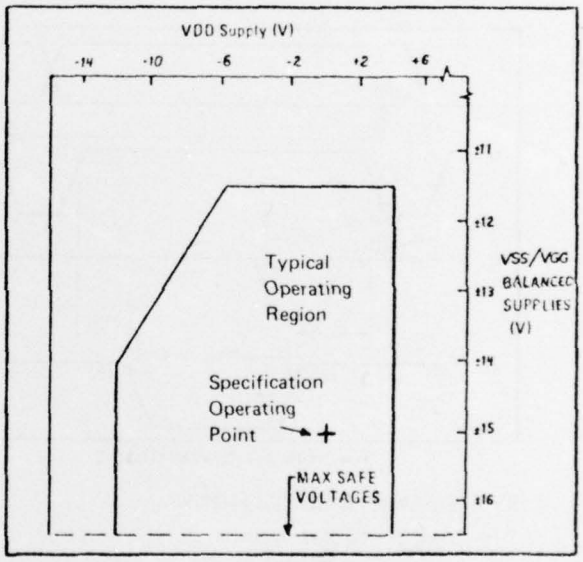

RETENTION

When written with a 100 μsec $\overline{MW}$ pulse, the minimum unpowered data retention time is 4000 hours. Data may be accessed a minimum of 2×10^{11} times without refresh. The chip is intended to be used as a read/write memory, and it may be cleared and written at least 10^{10} times. This is equivalent to more than 10 years of typical system continuous operation.

POWER SUPPLY RANGE

System supplies may vary over a wide range with disturbing the normal operation of the memory.

INPUT CAPACITANCE

$\overline{Ai}$, DWi, DRi, AE	< 5 pF
$\overline{CL}$	< 8 pF
$\overline{MW}$	< 10 pF
$\overline{TR}$, $\overline{\phi i}$	< 15 pF
$\overline{BS}$	< 38 pF

Westinghouse Electric Corporation
Defense and Electronic Systems Center
Baltimore, Maryland

Revised

76-0649-V-6

MISSION
of
Rome Air Development Center

RADC plans and conducts research, exploratory and advanced development programs in command, control, and communications (C^3) activities, and in the C^3 areas of information sciences and intelligence. The principal technical mission areas are communications, electromagnetic guidance and control, surveillance of ground and aerospace objects, intelligence data collection and handling, information system technology, ionospheric propagation, solid state sciences, microwave physics and electronic reliability, maintainability and compatibility.