NAVAL POSTGRADUATE SCHOOL
Monterey, California

THESIS

A 32-BIT MICROPROCESSOR BASED SOLID
STATE DATA RECORDER
FOR SPACE BASED APPLICATIONS

by

Thomas J. Frey, Jr.

March 1986

Thesis Advisor: R. Panholzer

Approved for public release; distribution is unlimited.

T226309

REPORT DOCUMENTATION PAGE

a REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
a. SECURITY CLASSIFICATION AUTHORITY	3 DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution is unlimited
b. DECLASSIFICATION / DOWNGRADING SCHEDULE	

PERFORMING ORGANIZATION REPORT NUMBER(S)	S. MONITORING ORGANIZATION REPORT NUMBER(S)

a. NAME OF PERFORMING ORGANIZATION aval Postgraduate School	6b. OFFICE SYMBOL (If applicable) 62	7a. NAME OF MONITORING ORGANIZATION Naval Postgraduate School
c. ADDRESS (City, State, and ZIP Code) onterey, California 93943-5000		7b. ADDRESS (City, State, and ZIP Code) Monterey, California 93943-5000

a. NAME OF FUNDING / SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
:. ADDRESS (City, State, and ZIP Code)		10 SOURCE OF FUNDING NUMBERS

PROGRAM ELEMENT NO.	PROJECT NO	TASK NO	WORK UNIT ACCESSION NO

TITLE (Include Security Classification)
A 32-BIT MICROPROCESSOR BASED SOLID STATE DATA RECORDER FOR SPACE BASED APPLICATIONS

PERSONAL AUTHOR(S)
Thomas J. Frey, Jr.

a TYPE OF REPORT Master's Thesis	13b TIME COVERED FROM _____ TO _____	14 DATE OF REPORT (Year, Month, Day) 86 March	1S PAGE COUNT 76

SUPPLEMENTARY NOTATION

COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Magnetic Bubble Memory; Space Data Recorder; Tape Recorder Emulator; Non-Volatile Memory

ABSTRACT (Continue on reverse if necessary and identify by block number)
A need exists for non-volatile, reliable and radiation hardened mass ata storage devices for space applications. Tape recorders modified for se in the space environment have proven both expensive and unreliable. he use of magnetic bubble memory as a recording medium in recent space rojects offers a promising alternative to tape recorders. This thesis evelops the design of a 32-bit microprocessor based advanced solid state ata recorder for space based applications utilizing two-megabyte agnetic bubble memory boards as the storage medium.

DISTRIBUTION / AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS	21 ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED
a NAME OF RESPONSIBLE INDIVIDUAL rofessor Rudolf Panholzer	22b TELEPHONE (Include Area Code) (408)646-2154	22c. OFFICE SYMBOL 62Pz

D FORM 1473, 84 MAR   83 APR edition may be used until exhausted   SECURITY CLASSIFICATION OF THIS PAGE
All other editions are obsolete

A 32-bit Microprocessor Based Solid State Data Recorder
for Space Based Applications

by

Thomas J. Frey, Jr.
Lieutenant, United States Navy
B.S., United States Naval Academy, 1977

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN ELECTRICAL ENGINEERING

from the

NAVAL POSTGRADUATE SCHOOL
March, 1986

ABSTRACT

A need exists for non-volatile, reliable and radiation hardened mass data storage devices for space applications. Tape recorders modified for use in the space environment have proven both expensive and unreliable. The use of magnetic bubble memory as a recording medium in recent space projects offers a promising alternative to tape recorders. This thesis develops the design of a 32-bit microprocessor based advanced solid state data recorder for space based applications utilizing two-megabyte magnetic bubble memory boards as the storage medium.

3

TABLE OF CONTENTS

I. INTRODUCTION

Space consists of vacuum, fields, and radiation. Information is the principal commodity of space based systems. The problem: Design a mass data storage system that functions reliably in a vacuum, in the presence or absence of electrical, magnetic, and gravitational fields, and in spite of continuous exposure to high levels of radiation.

In navigation, telecommunications, meteorology, the physical sciences, and engineering, space systems produce or manage data. If the environment of space has expanded the horizons of our knowledge, it has also complicated our manipulation of that knowledge. Unless one is able to provide (and can afford!) continuous downlinks to earth, the need arises for a data storage system which is rugged, reliable, small, light, and low power. This thesis proposes the design of a 32-bit microprocessor based solid state data recorder utilizing multiple magnetic bubble memories as the recording media.

II. NATURE OF THE PROBLEM

The ideal data system would be able to process DC to light for an infinite length of time, run on a watch battery, fit in a cigarette pack, have a zero error rate, and possess a mean time between failure of infinity. Obviously, such a recorder does not exist.

For mundane uses, tape recorders and disk drives are used to store large amounts of data. Adapting them for space use is difficult at best. They must be packaged in an atmosphere, because the gas acts as a lubricant between the recording medium and the heads. Fragile drive systems contain moving parts, susceptible to mechanical shock and wear.   The resulting waste heat compounds the complex thermal control problem in a space system. Finally, the recording media themselves are fragile.

For data storage systems in a space environment, magnetic bubble memories offer an attractive alternative to tape recorders and disk drives. Although lacking the storage capacity of a tape system, bubble memories compare quite favorably to disk drives. Figures 1 and 2 are adapted from a
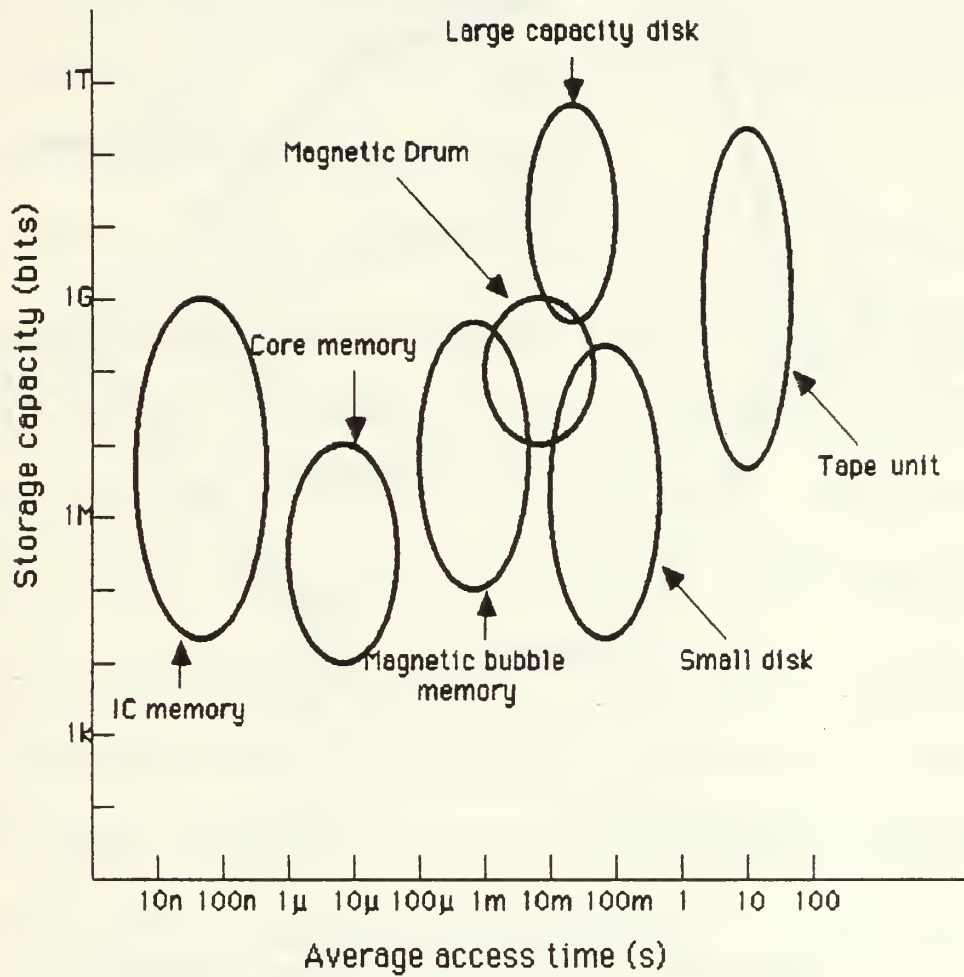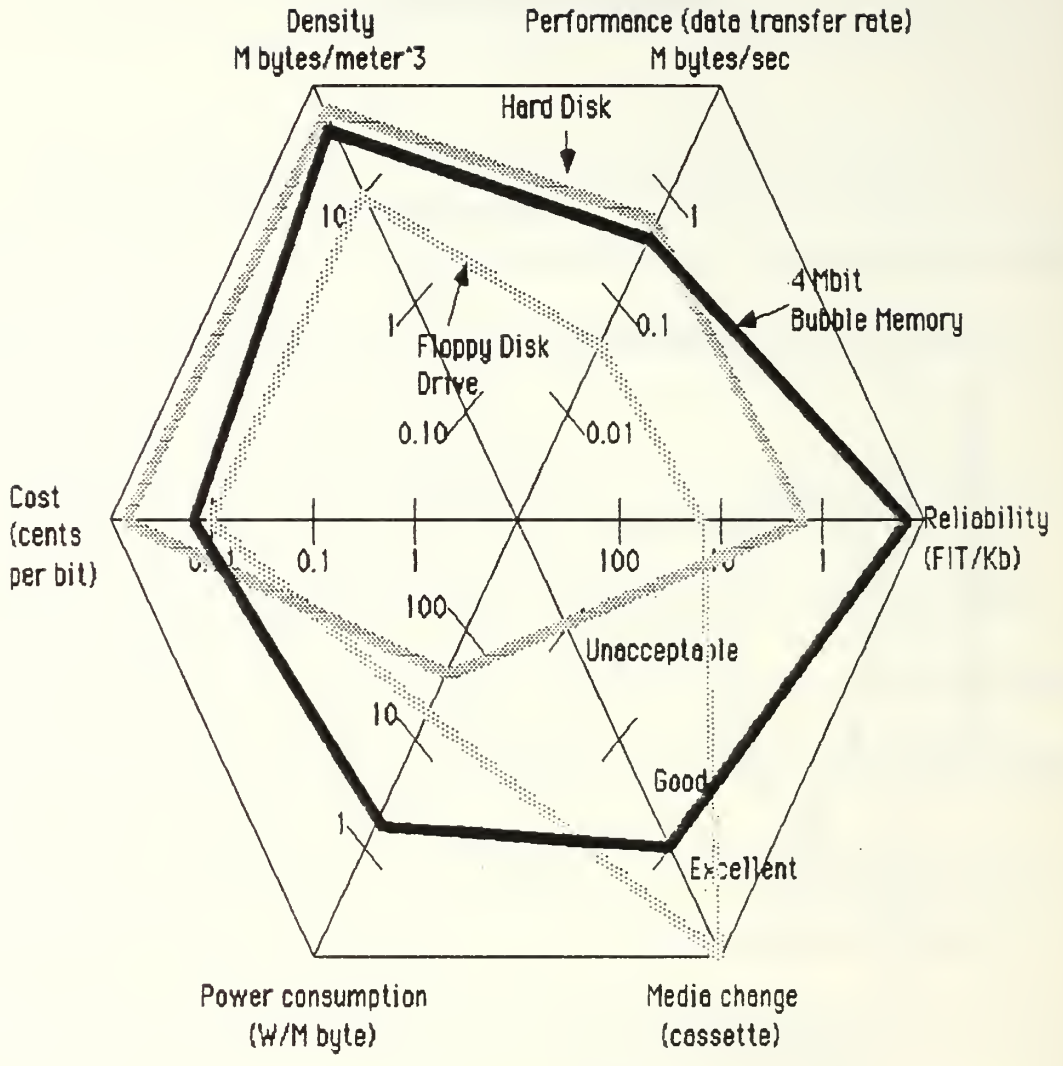
Figure 1. Storage Capacity versus Access Time

7

Figure 2 Bubble Memory Performance
(Magnetic bubble memory compared with floppy and hard disk)

8

Hitachi technical note [Ref. 1:p. 3]. Figure 1 compares the average access times and storage capacities of integrated circuit random access memory (RAM), magnetic core memory, magnetic bubble memory, magnetic drum systems, small disk drives, large capacity hard disk drives, and tape units. Figure 2 depicts the performance of magnetic bubble memory compared with floppy and hard disk drives.

In terms of non-volatility, density, reliability, and survivability, bubble memories are superior to other types of solid state, multiple read and write memories. The tradeoffs involve bandwidth, access time, complexity, and price. As an example of the suitability of magnetic bubble memories , consider the impracticality of building a twelve megabyte solid state data recorder for a space application from 256 kilobit integrated circuit memories (DRAM's). However, this recorder has been built using magnetic bubble memories. A more complete comparison of magnetic bubble memories and other solid state memories, in addition to a description of the twelve megabyte device is found in [Ref. 2].

To appreciate the utility of bubble memories in a space environment, a basic understanding of magnetic bubble memory operation is required. The

basic material for any magnetic bubble memory is garnet. The general

formula is either $M_3^{II}M_2^{III}(SiO_4)_3$ or $M_3^{III}M_5^{III}O_{12}$ where M^{II} is a valance

2 metal ion, and M^{III} is a valence 3 metal ion [Ref. 3: p. 508]. If the metals

used in making a synthetic garnet are properly chosen, the resulting

material will be ferromagnetic. This means that the garnet will consist of

many small magnetic domains. A magnetic domain is a group of molecules,

in a magnetic material, with parallel magnetic orientations. Current

magnetic bubble memory technology uses samarium lutetium garnets,

although extensive work has been done by Bell Labs and others with bismuth

yttrium garnets in an effort to widen the temperature range for magnetic

bubble memory performance [Ref. 4]. In unmagnitized bulk material, the

domains are oriented at random in three dimensions. Forming the material

into a thin film, with proper crystalographic orientation, limits the random

orientation to two directions, perpendicular to the surface of the film. This

is usually accomplished by liquid phase epitaxy on a gallium gadolinium

garnet substrate. For example, consider a thin, ferromagnetic garnet wafer

(Figure 3.). In the absence of an external magnetic field, called a bias field,

approximately half of the domains are oriented in each of the two possible

directions, represented in the figure as "+" and "-".

Garnet Wafer with No External Magnetic Field

Figure 3.

The application of a weak external magnetic field, in this case oriented

in the "-" direction, causes the domains that are oriented with the field to

grow by the motion of their walls. This causes the domains of opposite

orientation to substantially narrow, as shown in Figure 4.

Garnet Wafer with Weak External Magnetic Field

Figure 4.

If the bias field is increased, the oppositely oriented domains continue

to decrease until they become cylindrical in shape, surrounded by a much

larger domain oriented with the applied field (Figure 5). If the bias field

were to continue to increase, all of the oppositely oriented domains would

disappear. The effect can be reversed by decreasing the bias field. First, the

oppositely oriented cylindrical domains would reappear; as the field

continued to decrease, they would expand to long, thin domains; and finally,

in the absence of the external field, half of the material would again be

oriented in each direction.

Garnet Wafer with Strong External Field

Figure 5.

It is these cylindrical domains oppositely oriented to the applied bias

field (called, appropriately enough, "bubbles") that are used for non-volatile

data storage in magnetic bubble memories. A "bubble" represents a binary

"1", absence of a "bubble" represents a binary "0".

A permalloy film pattern is overlaid on the garnet to control the

location of the bubbles. It is this pattern which organizes the bubbles into

storage loops and input and output tracks. A seed bubble, which uses pulses

of current to create additional bubbles by cutting the seed bubble into two

bubbles, each of which expands to the size of the original seed bubble, is

used for data input. The seed bubble remains and the newly created bubble

is swapped onto the input track. This bubble replication mechanism is used

at the end of each storage loop to duplicate the data onto the output track.

A magnetoresistive bridge then acts as a bubble detector to output the data.

A rotating magnetic field, generated by current passing through two

orthogonal coils wrapped around the bubble memory is used to manipulate

and move the data. The magnetic bubble memory architecture is shown in

Figure 6.   A detailed treatment of the mechanics of magnetic bubble

memories may be found in [Ref. 5: pp. 6-1 to 6-10].

The performance of a particular bubble memory is determined by the

coil frequency, the number of storage loops, and the number of bits per

storage loop. The usual performance specifications are: data rate, access

time, page size, and power consumption.

Data rate is a measure of how many bits (or bytes) per second a bubble

memory device can read or write. It is usually presented in two formats,

peak data rate, and continuous data rate. Peak, or burst, data rate measures

the ability of the bubble memory to handle data assuming the proper storage
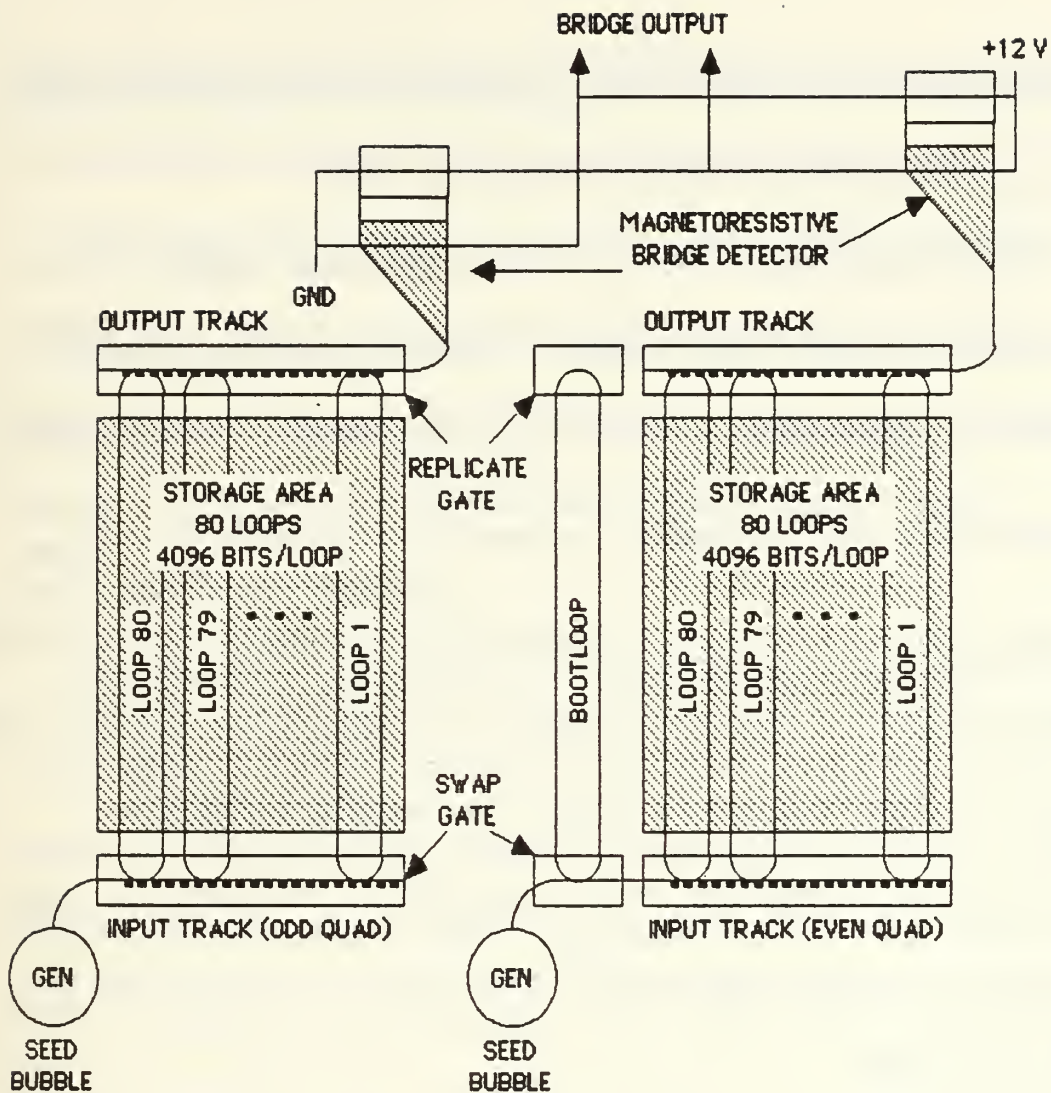
Figure 6. Magnetic Bubble Memory Architecture (One Half Shown)

location is accessed already. No provison is made for access time by the

peak data rate. The peak rate can only be maintained for milliseconds.

15

Continuous data rate accounts for access time and other system overhead and represents system performance in the "real world."

Access time measures how long it takes the bubble memory to obtain data from a particular storage location. It is usually quoted as average access time, and maximum, or worst case, access time. In timing diagrams of bubble read and write operations, access time is often shown as t_{seek}.

Page size refers to the smallest quantum of data with which the bubble system can perform a read or write operation. Page size is a function of the number of storage loops, since each page uses an integral number of locations on each storage loop. In most designs it is one location per loop for each page, although some architectures provide for two locations per loop for each page.

Power consumption is ususally reported as peak power, and either standby or typical power. Peak power is the maximum power consumed by a bubble memory device running at its peak data rate. Standby power is the power consumed by a bubble memory that is on, but not reading or writing data. Typical power is the time averaged power consumption of a bubble

device during all phases of operation. Typical power is highly dependent on the "benchmark" operations used when making this measurement.

Data rate is proportional to the coil frequency, and the number of storage loops. Access time is proportional to the number of storage locations per storage loop, and inversely proportional to the coil frequency. Page size is proportional to the number of storage loops. Power is proportional to coil frequency.

Since the total data capacity of a bubble memory is simply the product of the number of storage loops and the number of bits in each loop, the tradeoffs in a bubble memory recorder system design now become apparent. Increasing the data rate will adversely affect power and size, since either the coil frequency must be increased or additional storage loops provided. Decreasing the access time will increase both power and page size, as the coil frequency must be increased or fewer long storage loops must be replaced by many short storage loops. As examples of these design choices, consider the Intel 7110 one megabit and 7114 four megabit bubble memory chips, and the Hitachi BDL 0134 one megabit and BDN 0151 four megabit bubble memory chips.

Installed in an Intel BPK 72 magnetic bubble memory kit the 7110 has 256 data storage loops each with 4,096 storage locations on each loop. Each page uses two storage locations in each loop. Therefore, the page size is 64 bytes (256 x 2 = 512 bits = 64 bytes). The coil frequency is 50 kHz. The result is a peak data rate of 12.5 kbytes/second, a maximum continuous data rate of 8.5 kbytes/second, an average access time of 48 milliseconds, and a maximum access time of 82 milliseconds. Peak power consumption is 6.72 watts and Intel quotes a typical power consumption of 3.9 watts [Ref. 5,6].

There are two primary differences between the Intel one megabit chip and the BDL 0134 from Hitachi. Hitachi uses a coil frequency of 100 kHz. The second is in the loop configuration. The BDL 0134 is organized as two blocks of 256 loops each with 2048 storage locations. The page size is 32 bytes which means only one of the two blocks is accessed per page. The primary benefit of the Hitachi design is a reduction in access time. Hitachi lists an average access time of 13 milliseconds, but does not list a maximum access time. Even though doubling the coil frequency enables the Hitachi BDL 0134 to handle twice as many pages per second as the Intel 7110, the pages are only half as large, and the peak data rate is still 12.5

kilobytes/second. Although no maximum continuous data rate is given in the specifications, it must be better than 8.5 kilobytes/second due to the reduced access time [Ref.1, p. 10]. The obvious penalty is in power consumption. Hitachi lists a typical power consumption of 10.7 watts. There is no peak power figure listed [Ref. 7: p.5].

By reducing the feature size, Intel developed a four megabit magnetic bubble memory with a volume only 32 per cent larger than the 7110. There are 512 storage loops with 8,192 bits per loop. The coil frequency remains 50 kHz; however, a new thin-film bubble detector allows the data rate to double. The access time is increased to 88 milliseconds on the average and a maximum of 168 milliseconds. Peak power is 13.4 watts and a typical power of 6.5 watts is specified [Ref. 5,6].

Hitachi chose another route. Its four megabit bubble is essentially four separate one megabit devices on a single wafer. The result is a very low access time, only 18.4 milliseconds. There is no improvement in data rate unless a controller accesses all four blocks in parallel. The coil frequency

is 100 kHz, and the listed typical power is 6.5 watts. Although not specified, peak power is in excess of 16.2 watts due to the current draw of the coils [Ref. 1: pp. 8-10].

All bubble memories are serial, and require an interface to connect to a parallel data bus. The interface is a bubble memory controller, and it is essentially a specialized microprocessor. A single bubble memory controller may control more than one bubble memory, usually up to eight. However, Hitachi has a bubble memory controller capable of handling 256 bubble memories [Ref. 8].

An important element of the bubble memory controller is the data buffer or first-in, first-out buffer (FIFO). This buffer is where the parallel data bus connects to the bubble memory. Early bubble memory controllers had 40 byte buffers, although two pages, or 128 bytes, are now standard.

To achieve the maximum data rate in a bubble memory system, several pages of data are accumulated in a system memory cache and then transferred to the bubble memory controller FIFO buffer via direct memory access (DMA) on demand of the bubble memory controller. Since, to date, most of this programming has been done at the assembly language level, a

128 byte FIFO simplifies the coding considerably and is a major

improvement over the 40 byte FIFO.

The single bubble memory per controller is the simplest design to

implement, and was the configuration of the first complete bubble memory

kits, the BPK 72 and BPK 5V75A from Intel. These kits had a single bubble

memory and all the required support circuits mounted on a single printed

circuit board, with the appropriate assembly language drivers for interface

to an 8-bit microprocessor system. The solid state data recorder developed

by the Naval Postgraduate School for the NASA Get Away Special program

aboard the space shuttle used 24 of the four megabit BPK5V75A bubble

memory kits as the recording medium. By taking advantage of the

non-volatility of the bubble memories and switching off the power to

temporarily inactive bubble memories, a recorder with a maximum

continuous data rate of 15 kilobytes per second, and a peak power of 15

watts was achieved. The 12 megabytes of storage allows almost 14

minutes of data to be recorded at the maximum continuous data rate [Ref. 2].

In a proposal for the mass memory unit of the MIlstar satellite,

Motorola, proposed using its military version of the 7110 one megabit

bubble memory as the recording medium. In a radiation hardened design, in which each component was designed to withstand a total dose of 100,000 rads, a redunancy scheme was developed with two bubble memory controllers and 24 bubble memories. It should be noted that the radiation hardness is not limited by the bubble memories, which can withstand a total dose of 5,000,000 rads, but by the radiation hardness of the various support chips. The goal was to have at least one controller and 16 bubble memories functioning after 10 years in orbit. Any single controller is able to address any eight bubble memories at any one time. Thus, total storage is three megabytes, degrading gracefully to two megabytes after 10 years. The maximum continuous data rate is 84 kilobytes per second, at a peak power of 93 watts. Under the expected operating conditions of a ten percent duty cycle, with eight bubble memories operating, the data rate is to be about seven kilobytes per second at eight watts. The primary limitation to this system is the use of one megabit instead of four megabit bubble memories [Ref. 9].

In sum, the objective is to increase the data rate and storage capacity while minimizing increases in power and physical size. Other desired

22

improvements include software written in a high level language,

programmable operation by the user, and increased radiation hardness.

III. AN IMPROVED SOLID STATE DATA RECORDER

Consider Figure 7, the data bus schematic of the 12 megabyte magnetic bubble memory recorder developed at the Naval Postgraduate School for the Get Away Special Program.
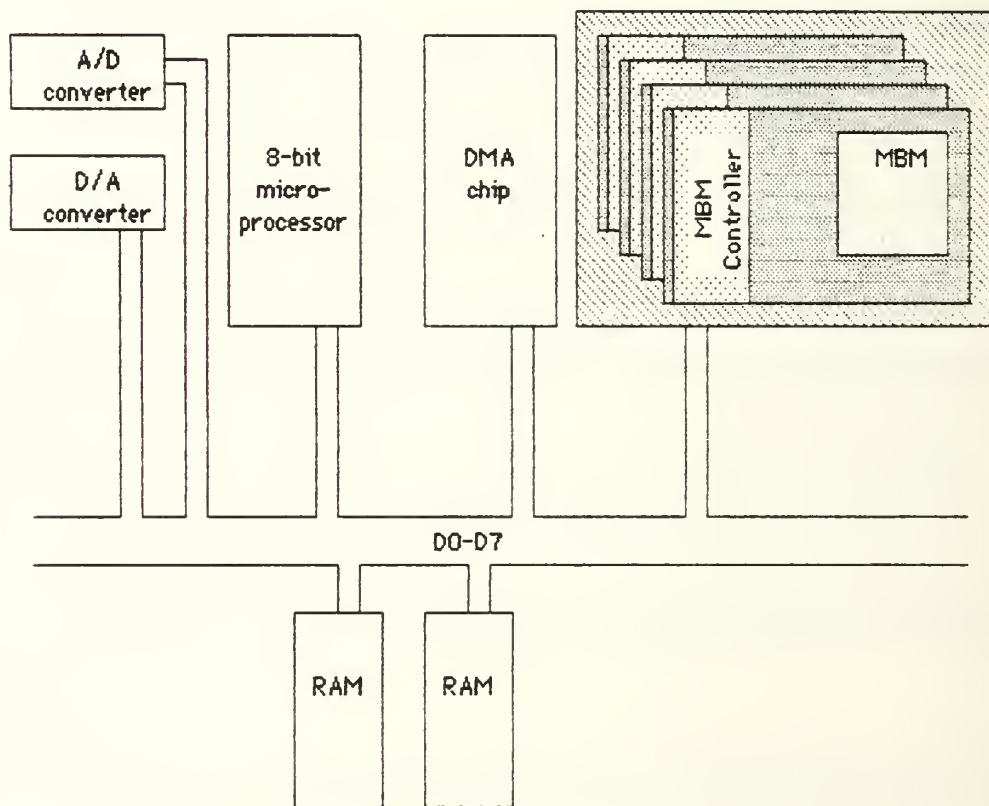

Figure 7. Data Bus Schematic for NPS Solid State Data Recorder

24

Sequential selection of individual BPK 5V75A cards by the microprocessor in order to conserve power, results in a solid state data recorder with a maximum continuous data rate of 15 kilobytes per second.

There are two methods for improving a solid state data recorder; improve the recording medium, the magnetic bubble memories, or improve the microprocessor. Desired improvements in the medium include increased data density, throughput, and reduced power consumption. For the microprocessor the goals are a faster clock rate, expanded instruction set, larger memory address space, and 32 bit bus size to enable multiple bubble memory controllers to operate in parallel.

By reducing the feature size, 16 megabit magnetic bubble memories are possible. However, initial production runs of 16 megabit devices are at least nine months to a year away. Indications are that while offering a significant improvement in data density, the throughput will be similar to that of the four megabit bubble memories [Ref. 10].

Multiple bubble memories with a single bubble memory controller offer the advantage of a higher data rate, and, because of a reduction in the chip count of support circuits, a saving in space and power consumption

compared to a single bubble memory with a bubble memory controller. By using a larger page size, the access time remains identical to that of the single-bubble memory-per-controller configuration. Therefore, an obvious improvement to the NPS 12 megabyte design would be to replace the BPK 5V75A cards with Intel iSBC 264-4 magnetic bubble memory boards. The 264-4 boards contain four of the Intel four megabit magnetic bubble memories, a single bubble memory controller, an 8257 direct memory access chip, and the required support circuits for 8- or 16-bit I/O addressing and an 8-bit data bus via the Multibus I interface. A functional block diagram of the 264-4 board is found in Appendix A [Ref. 11].

Because of the access time, the efficiency of any bubble memory system will be affected by the size of the block of data to be transferred. To read or write data, a single page at a time will result in a much lower data rate than if a thousand pages of data are transferred at once. The 264-4 card supports the transfer of up to 2048 (256 byte) pages of data at one time. The total time required, worst case, is:

$$164 \text{ ms} + 2630 \text{ } \mu s + (4100 \text{ } \mu s * N) = T \text{ seconds}$$

where N is the number of pages to be transferred, in this case 2048, 164

milliseconds is the maximum access time, 2630 microseconds is the read operation overhead (150 microseconds for a write), and 4100 microseconds is the time to read each page of data (4040 microseconds to write each page), resulting in a time of 8.56 seconds for the complete transfer. Therefore, the maximum continuous data rate for a single 264-4 card is 60 kilobytes/second.

Figure 8 is a data bus schematic of the modified data recorder. The maximum continuious data rate would be 60 kilobytes per second at a peak power of 37 watts. The volume occupied by the 12 megabytes of storage would be 0.2 cubic feet instead of 0.3 cubic feet.

There is, of course, nothing magical about 12 megabytes of total storage; more 264-4 boards could be added. Each additional board would add two megabytes of storage and occupy less than 0.03 cubic feet. Power and thermal considerations would be unchanged as would the data rate.

Since each bubble memory controller can control up to eight bubble memories, a case can be made for the Intel iSBC 268-4 board with 4 megabytes of total storage on each board. The primary reason for not using the 268 board is that current production versions allow for only two active
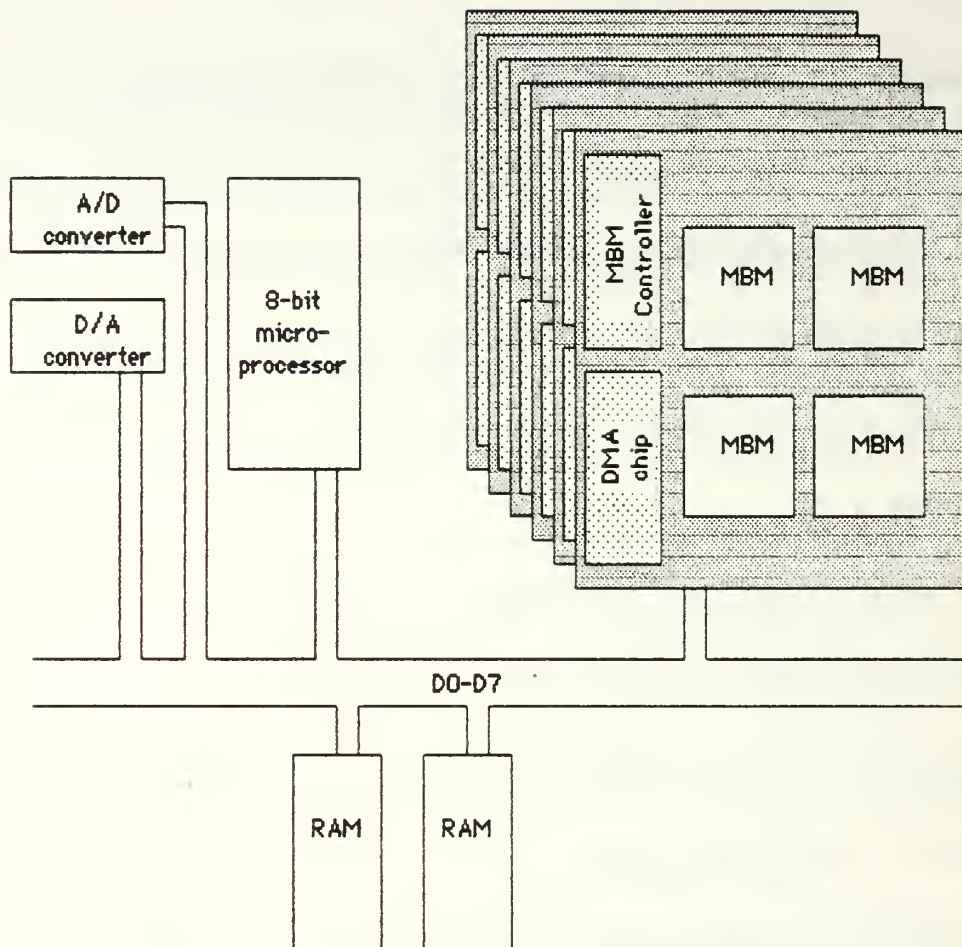
Figure 8. Modified Solid State Data Recorder

bubble memories at a time on an individual board. Thus, the continuous data

rate is less than that for a 264 board with four active bubble memories.

Further, the 268 board does not support direct memory access transfers of

data; only polled and interrupt driven modes are provided. The result would

be an enormous increase in microprocessor overhead, and increased complexity in the software.

The microprocessor performs two functions. The first is to transfer data between the system inputs or outputs and RAM. The second is to program the bubble memory controllers for DMA transfers from RAM to the magnetic bubble memories. At a data rate of 16 kilobytes per second, an 8-bit microprocessor is operating at close to peak speed. To obtain a quantum jump in bubble memory data recorders, a faster, more powerful microprocessor is required. A 16-bit data bus would allow two bubble memory controllers to be active at the same time. Use of a 32-bit microprocessor would increase the number of active bubble memory controllers to four. Therefore, a 32-bit microprocessor based solid state data recorder which used iSBC264-4 bubble memory boards as the recording medium would have a maximum continuous data rate of 240 kilobytes per second at a peak power of less than 145 watts. For the purpose of this design, the Intel 80386 32-bit microprocessor is used. A detailed description of the 80386 is contained in Appendix B. This microprocessor provides for four gigabytes of physical memory and 64 kilobytes of

addressable input/output (I/O) [Ref.12]. The iSBC264-4 cards would be I/O

mapped. However, data would be transferred using direct memory access

via memory mapped RAM buffers. The data bus schematic is shown in Figure

9.

The crucial engineering task would be the design of the random access

memory buffers and the buffers for the control, address, and data buses.

Appendix C discusses the buffers in detail.

The size of the random access memory buffers will have the primary

impact on system data rate since increasing the number of pages of data

transferred with each read or write command reduces the percentage of

time required for access and overhead by the bubble meory controller. To

obtain the maximum possible performance from the iSBC 264-4 cards, a

total of one megabyte of random access memory will be required for each of

the four buffers. This megabyte would be divided into two equal 512

kilobyte buffers, buffer 0 and buffer 1, each capable of holding 2048 pages

of data. Thus, the microprocessor would access one of the two buffers, for

example buffer 0, while the 264-4 card accessed the other. By keeping the

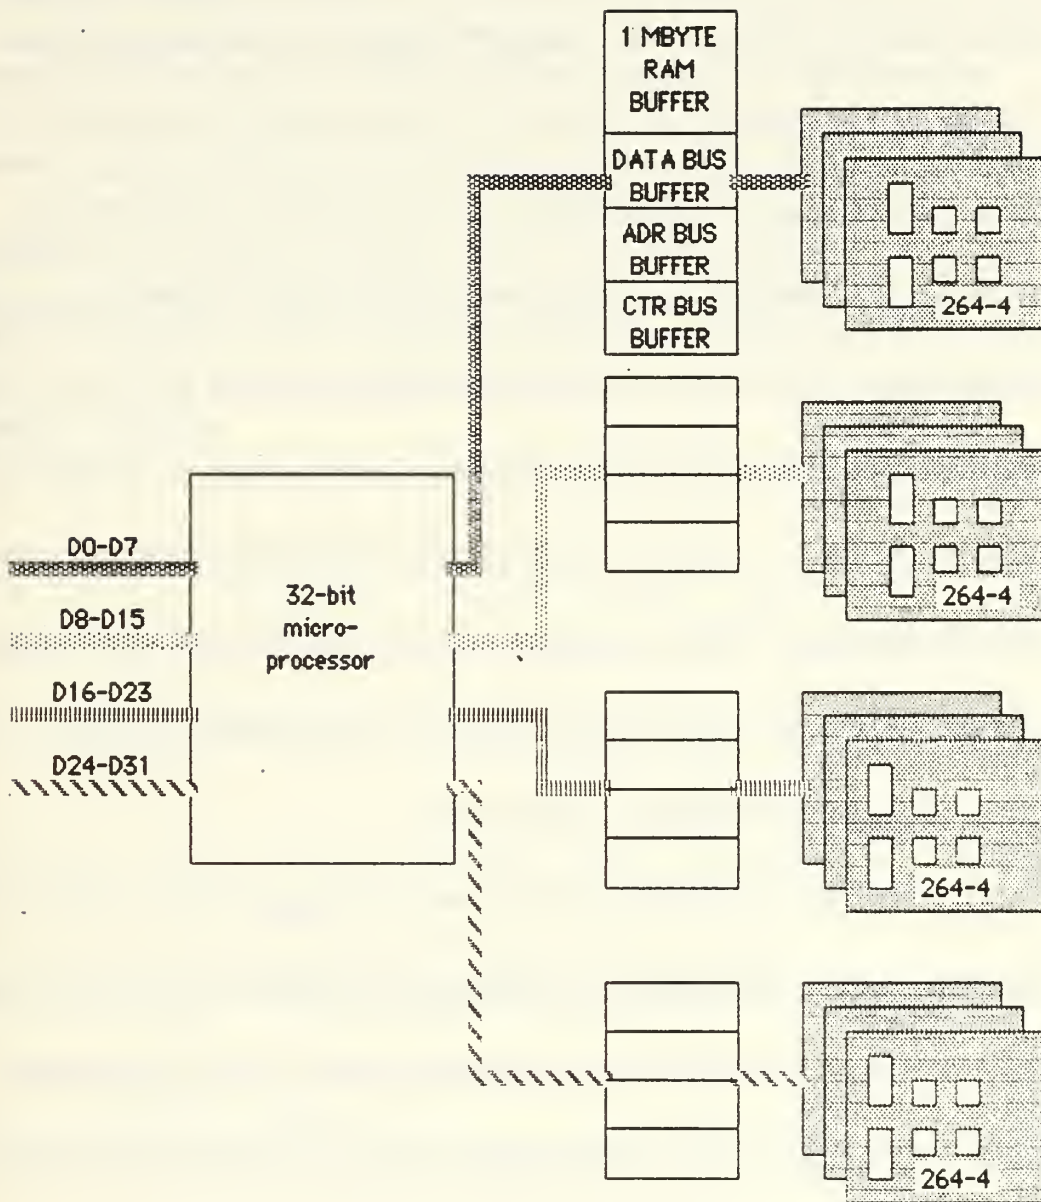system data rate below the maximum continuous data rate for the 264-4

Figure 9.

Data Bus Schematic for a 32-bit Microprocessor based
Solid State Data Recorder

cards, the microprocessor would always take longer to fill or empty its buffer, than the 264-4 cards. Therefore, when the microprocessor had finished filling or emptying its buffer, the microprocessor and 264-4 card would synchronously switch buffers, and the process would continue.

Power consumption would depend on the number of bubble memories active on each 264-4 card and the system data rate. For four active bubble memories on a single 264-4 card, the peak current draw is 7.0 amps at 5 volts, or 35 watts. The entire recorder will have a peak power of under 145 watts if low-power CMOS components are used for the microprocessor and buffers. The actual percentage of time that the recorder will draw 145 watts would be a function of the data rate.

Physical size is primarily a function of storage capacity. Each 264-4 card provides two megabytes of memory and occupies a volume less than 0.03 cubic feet. Allowing for the required buffers, the microprocessor, power switching circuits, and sufficient spacing betwwen circuit boards, a 48 megabyte recorder would have a volume of 3.1 cubic feet.

Finally, there is the required software. The Naval Postgraduate School design for the 8-bit, 12 megabyte solid state data recorder has 1800 lines

of assembly language software. The task of writing such code for a 32-bit microprocessor based solid state data recorder is dauntingly complex. However, using a microprocessor which supports UNIX and C keeps the task manageable. For system software such as this, the ability to move from high-level functions to bit and byte manipulation, makes C the programming language of choice.

IV. <u>CONCLUSIONS</u>

Figure 10 is a graph of maximum continuous data throughput rate versus peak power for a one megabit magnetic bubble memory, the NPS 12 megabyte solid state data recorder, the NPS design utilizing the iSBC 264-4 card as the recording medium, the Motorola design for the MILSTAR mass memory unit, and the subject of this thesis, a 32-bit microprocessor based design, which uses four groups of iSBC 264-4 cards as the recording medium.

Obviously, a higher data rate results in increased power consumption. However, the 32-bit based, 48 megabyte design has a data rate 15 times higher than the 8-bit based, 12 megabyte solid state data recorder with a less than tenfold increase in power. In terms of physical size, the improvements are equally promising: four times the total data storage in less than three times the volume. For those applications requiring the increased data rate and larger storage capacity, this advanced solid state data recorder design is superior to multiple 8-bit data recorders.
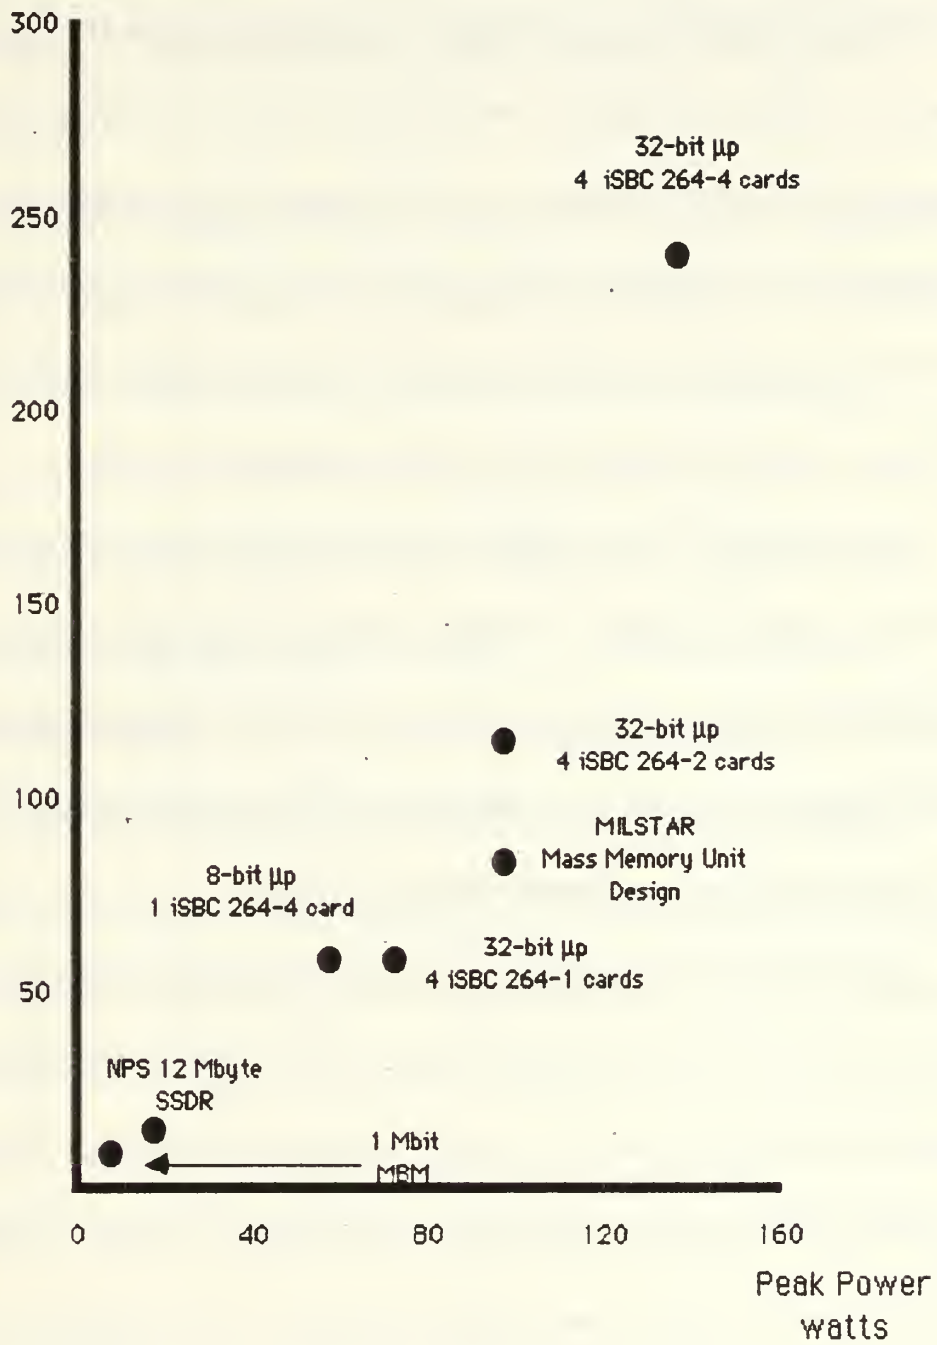
Figure 10. Solid State Data Recorder Performance

A brief comment about radiation hardness is in order here. The hardness of a recorder using magnetic bubble memory as the recording medium is dependent upon the radiation hardness of the components supporting the bubble memories, as the magnetic bubble memories themselves are capable of withstanding a total dose of 5,000,000 rads. For the CHMOS technology used in the 80386 microprocessor, Intel claims the ability to withstand a total dose of 500,000 rads [Ref. 13].

As a result of initial interest by the United States Air Force Space Division, several research proposals were initiated by the Naval Postgraduate School. These proposals examined the scheduling and management of producing this design in various configurations. The results of this planning may be summarized as follows.

Experience in the development of the 12 megabyte data recorder indicates an optimum development team of four to five individuals. A breadboard prototype could be up and running in nine months. The cost, not including plant overhead and salaries, would be on the order of $500,000. A first production, flight ready model would take 15 to 18 months, and cost around $2,000,000.

APPENDIX A: THE iSBC 264-4 MAGNETIC BUBBLE MEMORY BOARD

The iSBC 264-4 Magnetic Bubble Memory Board provides 2 megabytes of

non-volatile storage in a single, fully assembled printed circuit board with

four Intel 7114 4-megabit magnetic bubble memories, a 7225 bubble

memory controller, and an 8257 direct memory access controller (DMAC).

The board is 12.0 inches long, 6.75 inches high and has a depth of 0.59

inches. The block diagram is shown in Figure 11. The 264-4 supports

polled, interrupt-driven, and direct memory access data transfers in page

sizes of 256 bytes. The peak, or burst, data rate is 100 kilobytes per

second, and the maximum continuous data rate is 60 kilobytes per second.

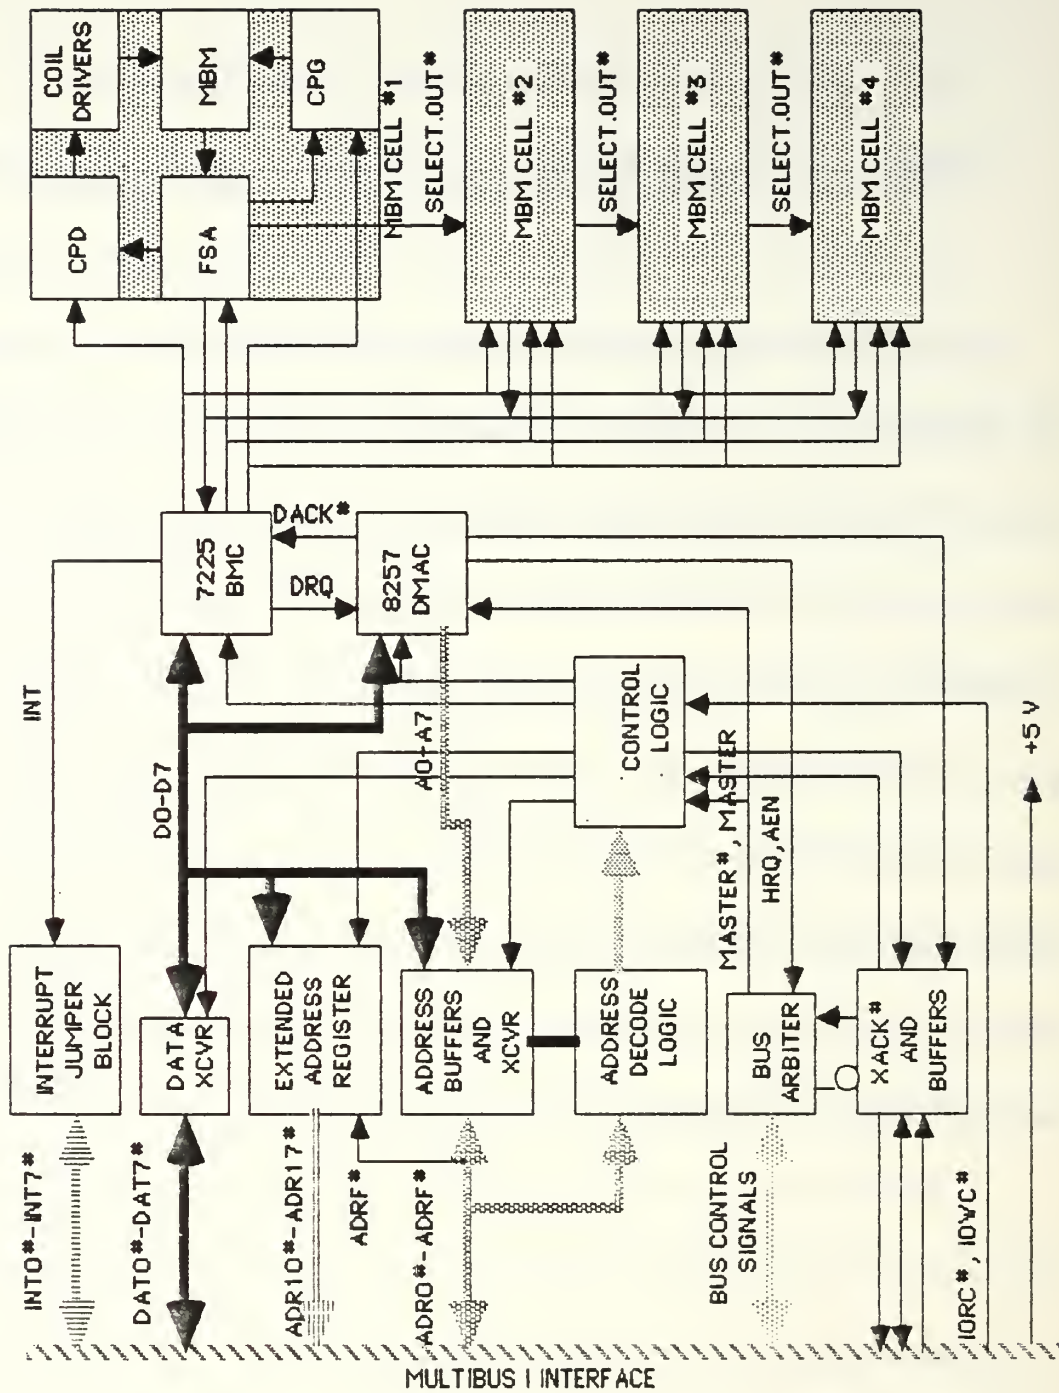The board requires +5 volts, and has a maximum current draw of 7.0 amps.

Figure 11. iSBC 264 Board Block Diagram

38

APPENDIX B: THE INTEL 80386 32-BIT MICROPROCESSOR

The 80386 is an advanced 32-bit microprocessor capable of addressing four gigabytes of physical memory and 64 kilobytes of I/O. The microprocessor can handle 8, 16, and 32-bit data types. There are two available clock speeds, 12.5 megahertz and 16.0 megahertz.

The 80386 supports Unix, C, and PL/M in addition to being object code compatible with all iAPX 86 (8086, 8088, 80186, 80188, 80286) microprocessors. Figure 12 depicts the functional signal groups of the 80386. Figure 13 shows the memory and I/O maps of the microprocessor for this application. Since there are six iSBC 264-4 cards connected to each of the four bytes of the data bus, each of the I/O addresses listed corresponds to the base address of one of the 24 iSBC 264-4 cards. The cards are numbered according to which byte of the data bus (byte 0 through 3) and which of the six iSBC cards (card 0 through 5) connected to that particular byte is being referenced.
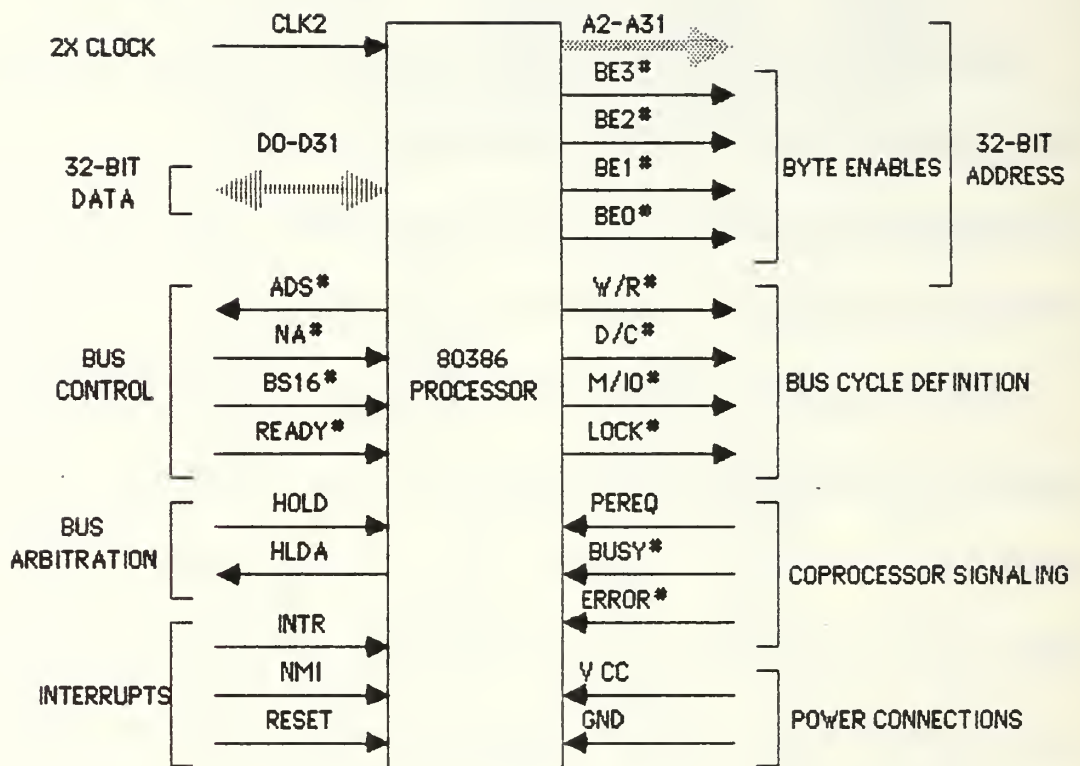
Figure 12. Functional Signal Groups of the 80386

MEMORY MAP (4GBYTES)   I/O MAP (64KBYTES)

FFFFFFFF H	RESET ROUTINES	
FFFFFFF0 H		
	NOT USED	
802FFFFF H	4 MBYTE RAM BUFFER	
80000000 H		
	AVAILABLE FOR SYSTEM MEMORY	
000003FF H	INTERRUPTS	
00000000 H		

I/O	Address
AVAILABLE FOR OTHER I/O	FFFF H – 029C H
BYTE 0 CARD 5	0283 H
BYTE 0 CARD 5	0282 H
BYTE 0 CARD 5	0281 H
BYTE 0 CARD 5	0280 H
BYTE 3 CARD 4	0203 H
BYTE 2 CARD 4	0202 H
BYTE 1 CARD 4	0201 H
BYTE 0 CARD 4	0200 H
BYTE 3 CARD 3	0183 H
BYTE 2 CARD 3	0182 H
BYTE 1 CARD 3	0181 H
BYTE 0 CARD 3	0180 H
BYTE 3 CARD 2	0103 H
BYTE 2 CARD 2	0102 H
BYTE 1 CARD 2	0101 H
BYTE 0 CARD 2	0100 H
BYTE 3 CARD 1	0083 H
BYTE 2 CARD 1	0082 H
BYTE 1 CARD 1	0081 H
BYTE 0 CARD 1	0080 H
BYTE 3 CARD 0	0003 H
BYTE 2 CARD 0	0002 H
BYTE 1 CARD 0	0001 H
BYTE 0 CARD 0	0000 H

Figure 13. System Memory and I/O Maps

41

APPENDIX C <u>80386 MICROPROCESSOR TO iSBC 264-4 MAGNETIC BUBBLE MEMORY BOARD BUFFERS</u>

The 24 iSBC 264-4 magnetic bubble memory boards are organized as four groups of six boards each. Each group provides data storage for one of the four bytes of the 32 bit data bus. For the 80386, the last two bits of the address determine which byte of the data bus is in use. Thus, any memory or I/O address with A0 and A1 set to 0 is connected to D0-D7. If A0 is a 1 and A1 is a 0, D8-D15 is the active byte. D16-D23 is active for A1 set to 1 and A0 set to 0. Finally,if both A1 and A0 are high, then D24-D31 is active. The microprocessor generates the signals BE0*, BE1*, BE2*, and BE3*, (for Byte Enable, the * means the active state is a logic low) in lieu of A1 and A0. Therefore, the least significant two bits of any address indicate which group of magnetic bubble memory boards is being addressed.

Each 264-4 board has its own unique set of I/O addresses for its parametric, control, and status registers. Further, each board has its own 8257 Direct Memory Access Controller (DMAC), which is capable of

addressing 512 kilobytes of RAM for memory-to-I/O DMA transfers. The

DMAC is also capable of acting as a bus master. The iSBC 264-4 card

automatically increments its Register Address Counter (RAC) and its DMAC

Address Register after many of its operations. To take advanatge of this

hardware feature to simplify software, the microprocessor address lines

A2-A21 must be "mapped" into the address lines A0-A19 on each iSBC 264

card.

   If two seperate 512 kilobyte RAM buffers are provided for each byte of

the 32 bit data bus, then the microprocessor can engage in two different

types of external data transfer. It may communicate with any single 264-4

card via an I/O address, or it may perform a memory access of one of the

eight RAM buffers. On the other hand, each 264-4 card also engages in two

types of external data transfer. It may communicate with the

microprocessor via I/O, or it may access one of the two memory buffers

connected to its 8-bit data bus via DMA. The goal of the buffer system is to

prevent conflict between the microprocessor and the 264-4 cards. Figure

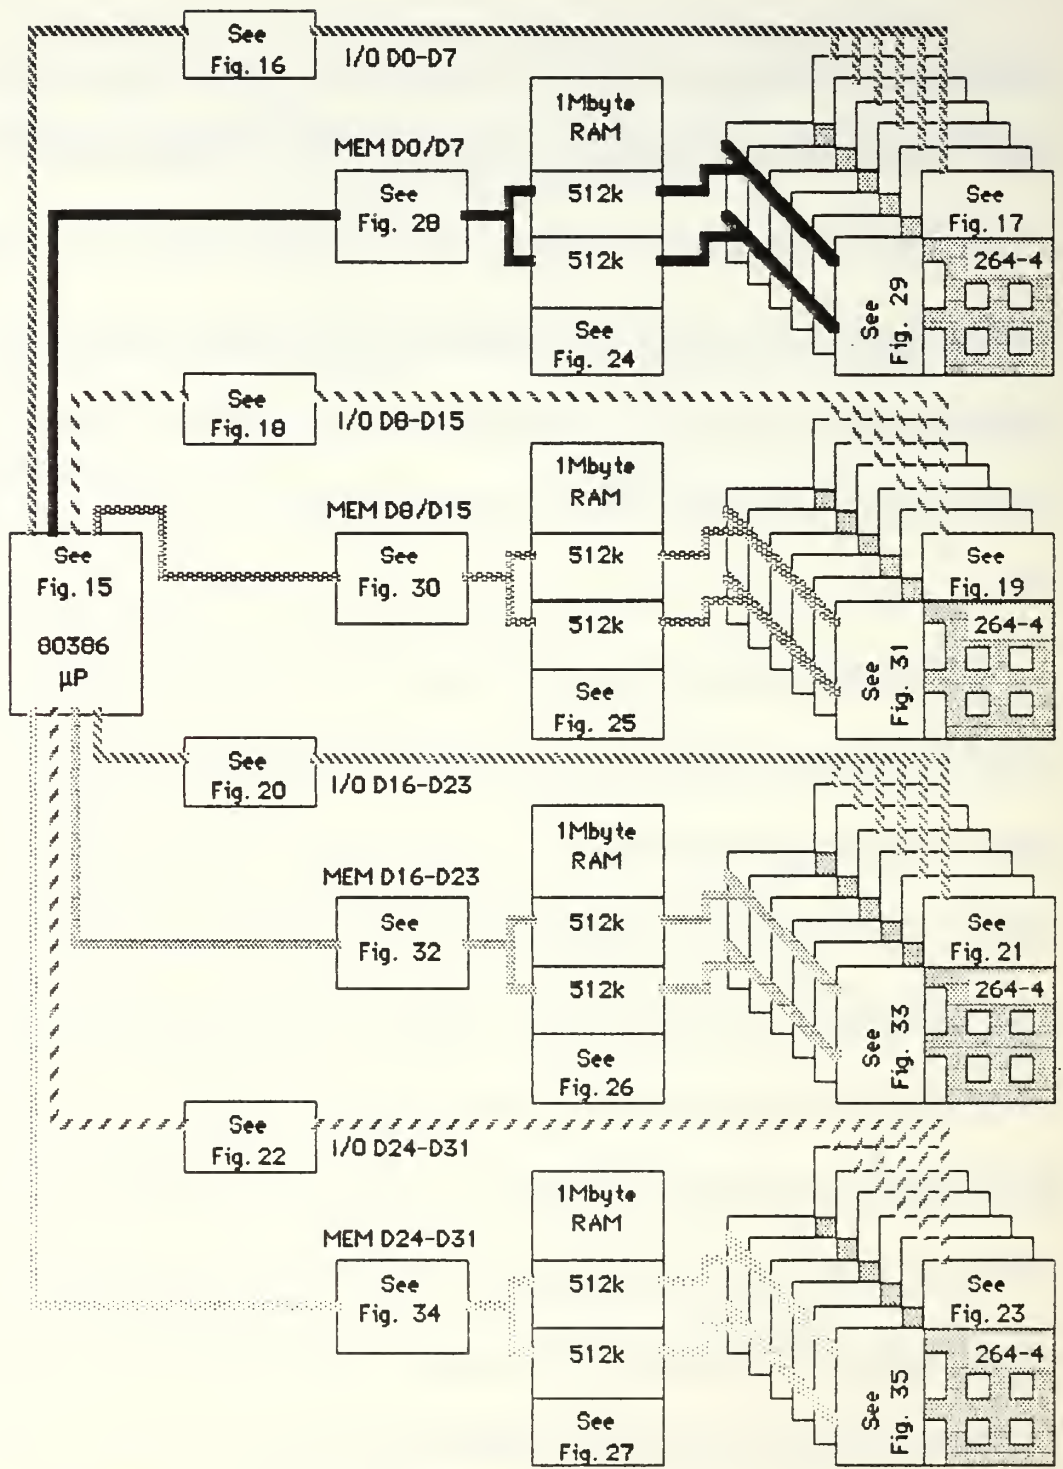14 presents a general overview of the buffer structure.

Figure 14. Advanced Solid State Data Recorder Buffer System

Figure 15 is a diagram of the logic for generating the control signals

IOR, IOW, MEMW, MEMR, and the address lines A0 and A1 from the 80386.

These signals control the activation and direction of other components in

the buffer system.

* NOTE: The * symbol indicates the active state occurs
when the signal is at a low voltage

Figure 15. Generation of Control Signals

46

Figure 16 depicts the microprocessor side of the I/O buffer for 264-4

cards connected to D0-D7. This buffer allows its local address and data bus

to be active if the microprocessor is attempting to read or write to the

control or status registers of one of the six 264-4 cards attached to D0-D7.
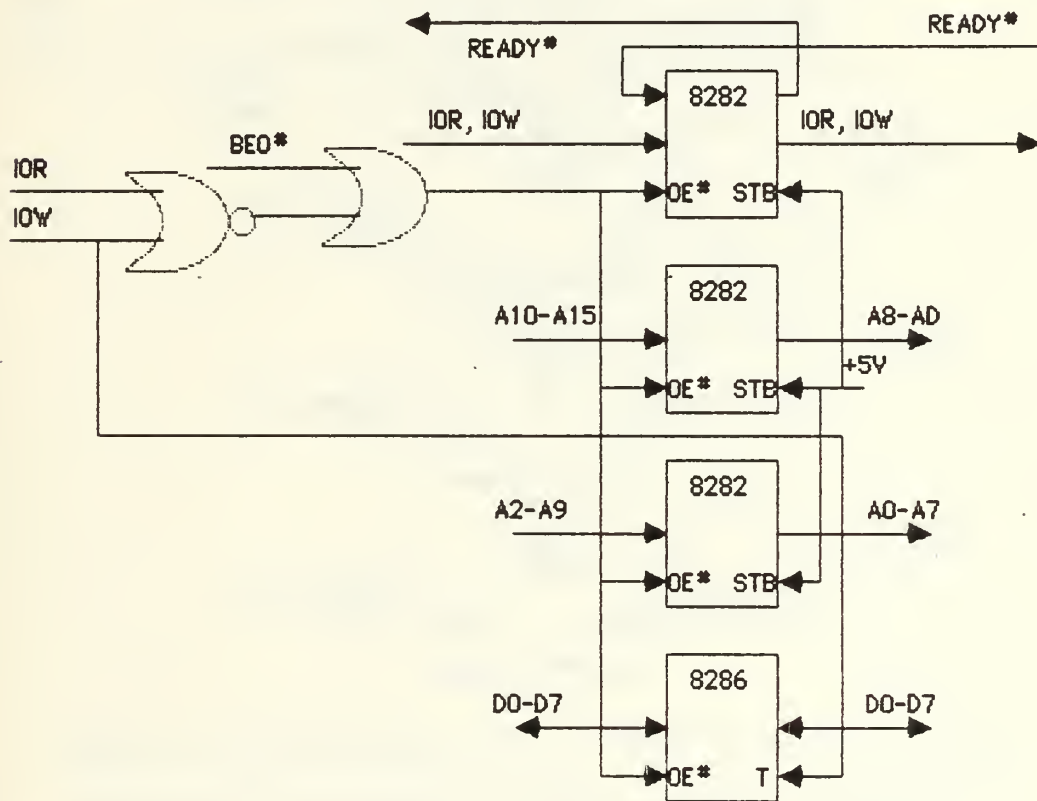

Figure 16. Microprocessor Side of I/O Buffer for D0-D7 Bus

Each of the six 264-4 cards has a corresponding I/O buffer which uses inverting 8283 and 8287 chips, because of the inverted Multibus signals, allowing it to respond to the microprocessor only if it is not conducting DMA transfers of data to its RAM buffer. This is illustrated in Figure 17.


Figure 17. iSBC 264-4 Side of I/O Buffer for DO-D7 Bus

Figures 18 through 23 depict the I/O buffers for the other three bytes of the data bus. The byte enable signal and the byte of the data bus that is buffered are the only differences in these figures.


Figure 18. Microprocessor Side of I/O Buffer for D8-D15 Bus
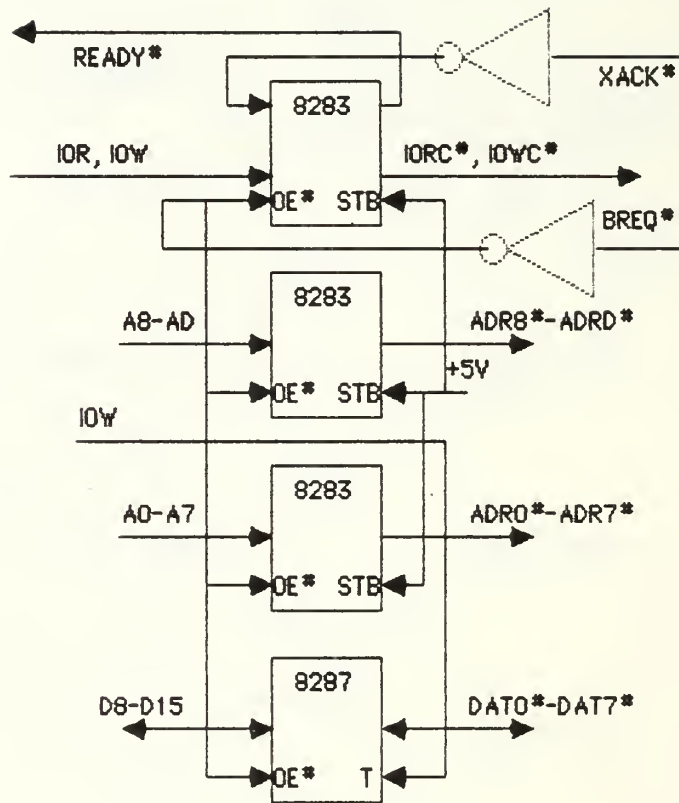
Figure 19. iSBC 264-4 Side of I/O Buffer for D8-D15 Bus

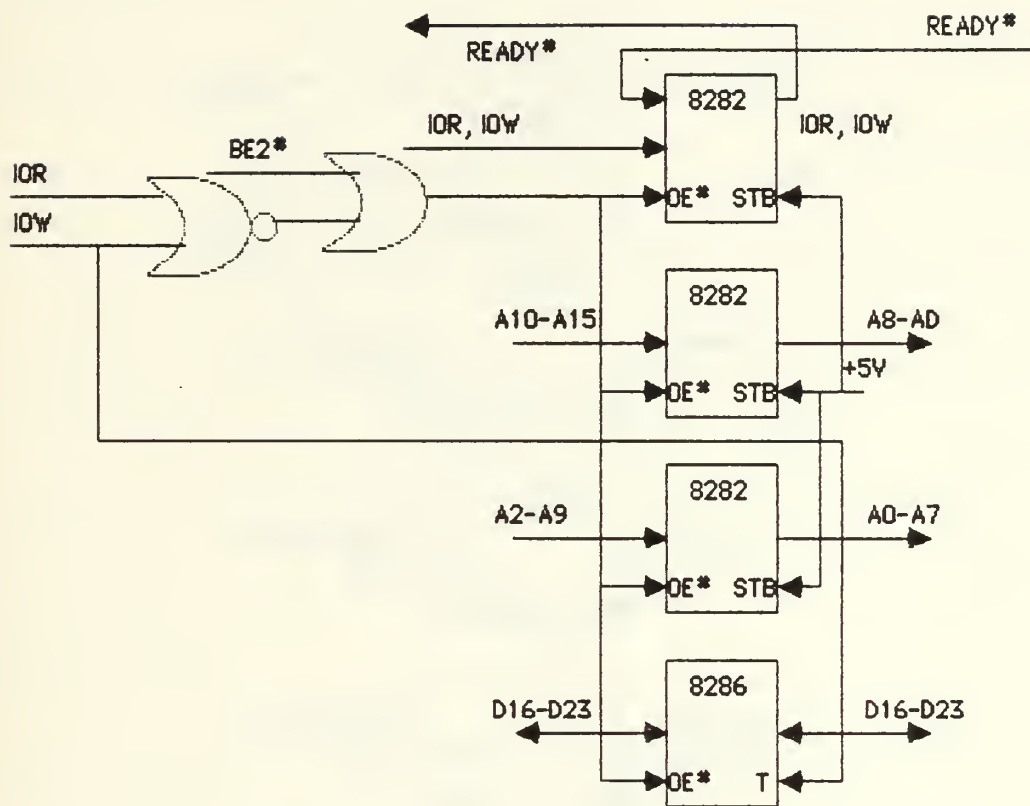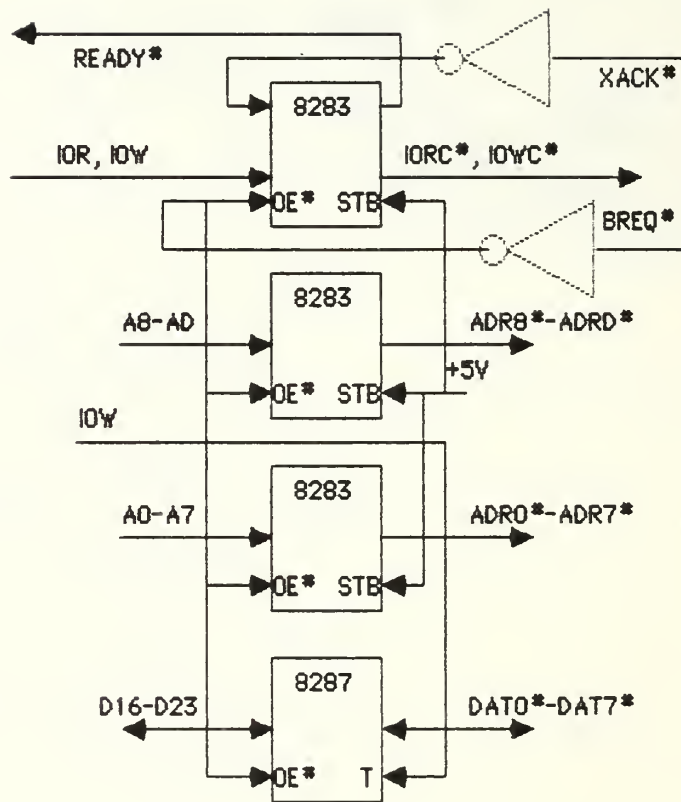Figure 20. Microprocessor Side of I/O Buffer for D16-D23 Bus

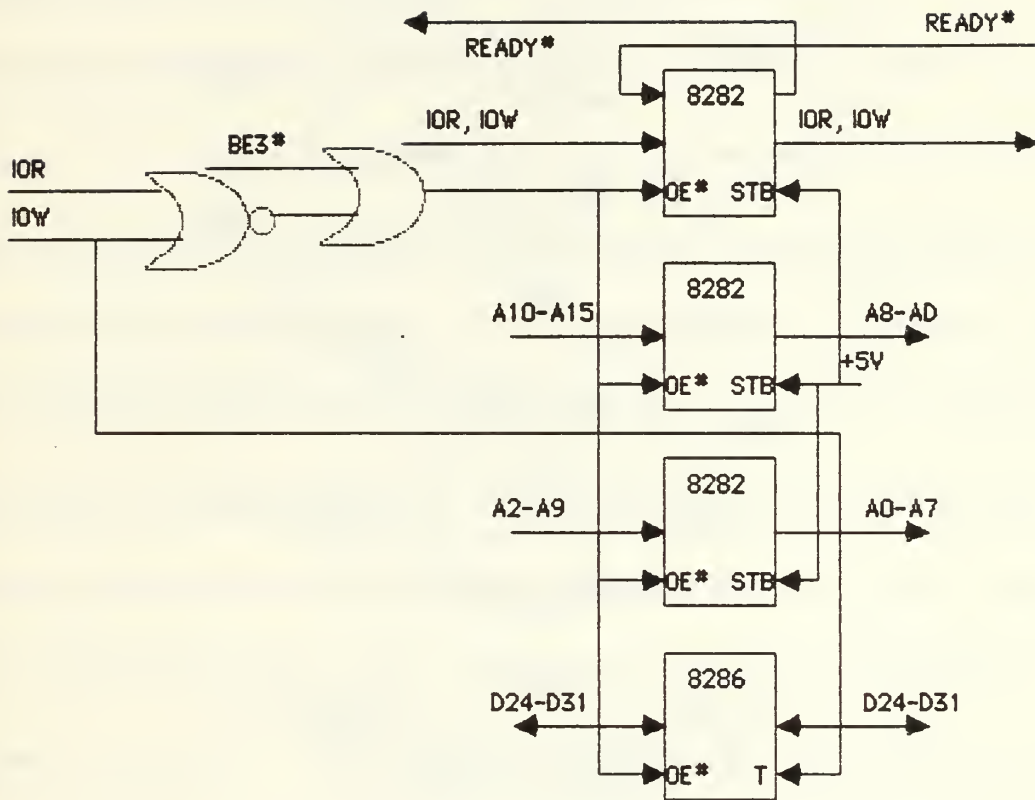51

Figure 21. iSBC 264-4 Side of I/O Buffer for D16-D23 Bus

52

Figure 22. Microprocessor Side of I/O Buffer for D24-D31 Bus
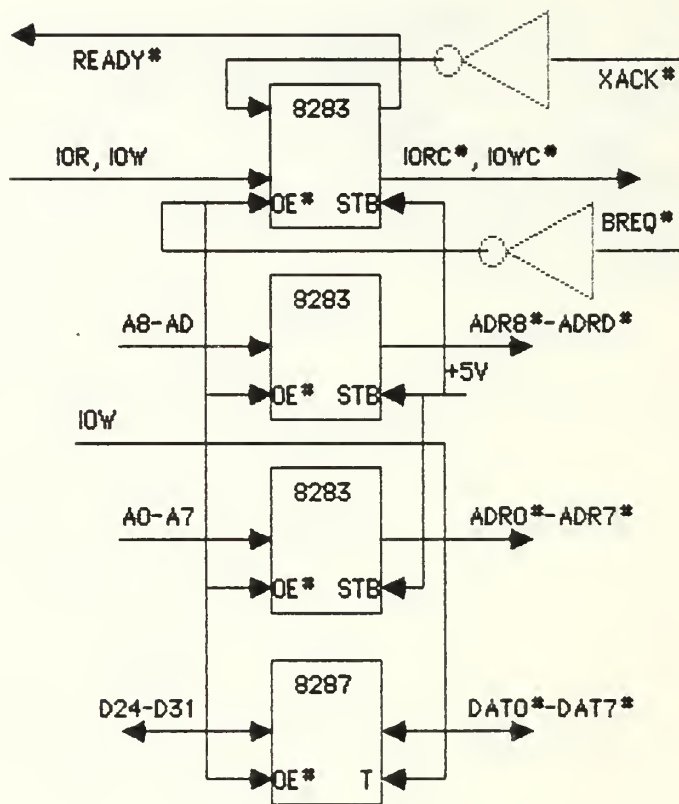
Figure 23. iSBC 264-4 Side of I/O Buffer for D24-D31 Bus

54

To facillitate data transfer between the microprocessor and the iSBC 264 cards, a one megabyte RAM buffer is provided for each byte of the data bus. Each of these megabyte buffers is divided into two separate 512 kilobyte buffers, buffer 0 and buffer 1. At any given time, the microprocessor may access one or the other, but not both, buffers. The status of A21 determines whether buffer 0 or buffer 1 is being accessed by the microprocessor. When the microprocessor has completed reading data from (or writing data to) one of the buffers it switches to the second buffer and programs the iSBC 264 card, via the I/O interface, to begin writing to (or reading from) the first buffer. This mutually exclusive alternation of buffers continues until the data transfer is complete. As the RAM is single port, the microprocessor address lines A2-A21 become the iSBC 264-4 address lines A0-A19. Similary, each byte of the microprocessor data bus is recognized as DAT0-DAT7 by the iSBC 264-4 cards connected to that particular bute of the data bus. Figures 24-27 depict the RAM buffers. The dual set of addresses denote the address as seen by the microprocessor and the address as seen by the iSBC 264-4 card.
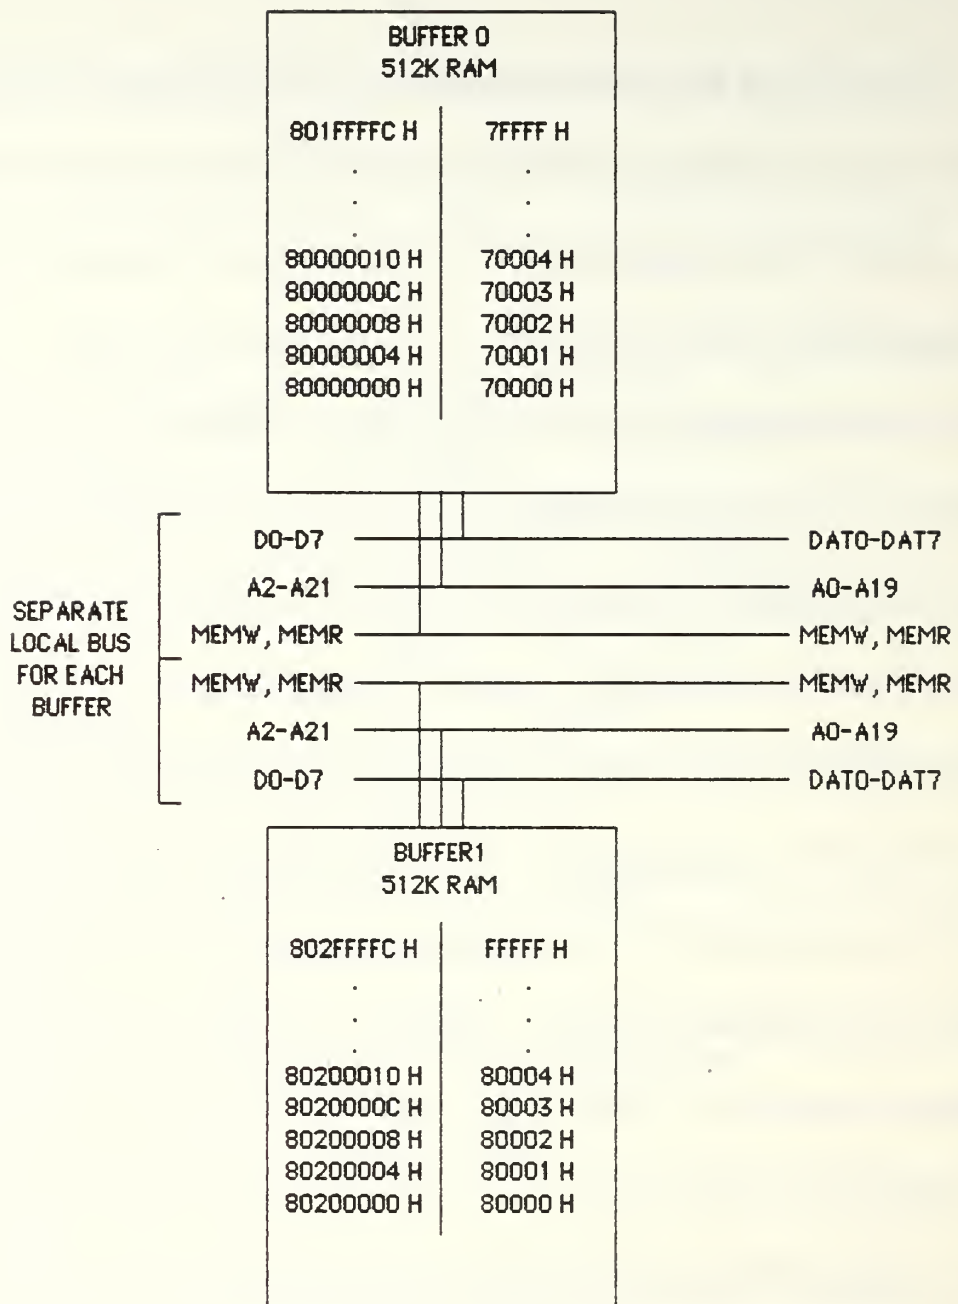
Figure 24. Ram Buffer for D0-D7 Bus

Figure 25. Ram Buffer for D8-D15 Bus

BUFFER 0
512K RAM

801FFFFE H	7FFFF H
.	.
.	.
.	.
80000012 H	70004 H
8000000E H	70003 H
8000000A H	70002 H
80000006 H	70001 H
80000002 H	70000 H

SEPARATE
LOCAL BUS
FOR EACH
BUFFER

D16-D23 ——————————————— DAT0-DAT7
A2-A21 ——————————————— A0-A19
MEMW, MEMR ———————————————
MEMW, MEMR ——————————————— MEMW, MEMR
A2-A21 ——————————————— MEMW, MEMR
D16-D23 ——————————————— A0-A19
 ——————————————— DAT0-DAT7

BUFFER 1
512K RAM

802FFFFE H	FFFFF H
.	.
.	.
.	.
80200012 H	80004 H
8020000E H	80003 H
8020000A H	80002 H
80200006 H	80001 H
80200002 H	80000 H

Figure 26. Ram Buffer for D16-D23 Bus

```
                    ┌─────────────────────────────┐
                    │         BUFFER 0            │
                    │         512K RAM            │
                    │                             │
                    │   801FFFFF H      7FFFF H   │
                    │       .             .       │
                    │       .             .       │
                    │       .             .       │
                    │   80000013 H      70004 H   │
                    │   8000000F H      70003 H   │
                    │   8000000B H      70002 H   │
                    │   80000007 H      70001 H   │
                    │   80000003 H      70000 H   │
                    │                             │
                    └─────────────────────────────┘
```

            D24-D31 ──────────────────────── DAT0-DAT7

            A2-A21 ───────────────────────── A0-A19

SEPARATE   MEMW, MEMR ───────────────────── MEMW, MEMR
LOCAL BUS
FOR EACH   MEMW, MEMR ───────────────────── MEMW, MEMR
BUFFER
            A2-A21 ───────────────────────── A0-A19

            D24-D31 ──────────────────────── DAT0-DAT7

```
                    ┌─────────────────────────────┐
                    │         BUFFER 1            │
                    │         512K RAM            │
                    │                             │
                    │   802FFFFF H      FFFFF H   │
                    │       .             .       │
                    │       .             .       │
                    │       .             .       │
                    │   80200013 H      80004 H   │
                    │   8020000F H      80003 H   │
                    │   8020000B H      80002 H   │
                    │   80200007 H      80001 H   │
                    │   80200003 H      80000 H   │
                    │                             │
                    └─────────────────────────────┘
```

Figure 27. Ram Buffer for D24-D31 Bus

The interface between the microprocessor and the one megabyte RAM buffer for D0-D7 is shown in Figure 28. For memory transfers to addresses 80000000 H through 801FFFFC H (with A1 and A0 both 0) the data and address buses to buffer 0 are active. For memory transfers to addresses 80200000 H through 802FFFFC H (again, both A1 and A0 low) the data and address buses to buffer 1 are active.

The bus connections from the RAM buffer to the iSBC 264-4 card are shown in Figure 29. DMA transfers using locations 70000 H through 7FFFFF H (as seen by the iSBC 264 card) cause the data and address buses to buffer 0 to be active. DMA transfers involving locations 80000 H through FFFFFF H enable the data and address buses to buffer 1.

There are two requirements for the system software to prevent conflicts. The first is that the microprocessor not enable (via I/O) a iSBC 264 card to access a buffer (for example, buffer 0) until the microprocessor has shifted its memory transfers to the other buffer (in this case, buffer1). The second is that the system data rate be chosen such that it always takes the microprocessor longer to fill or empty a buffer than the iSBC 264 card.
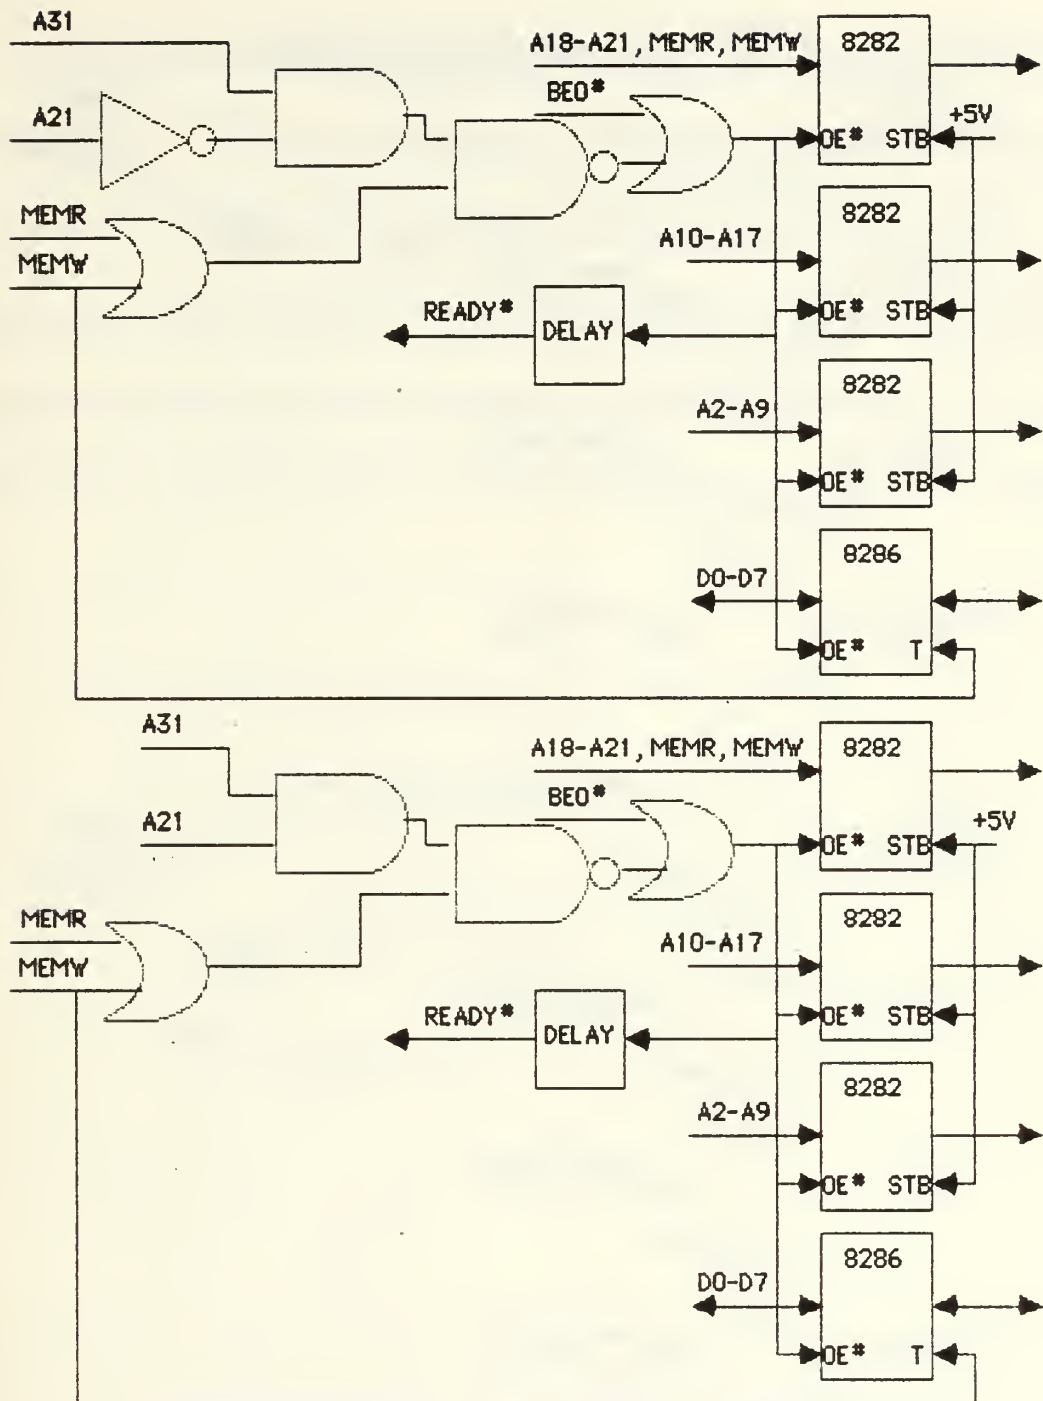
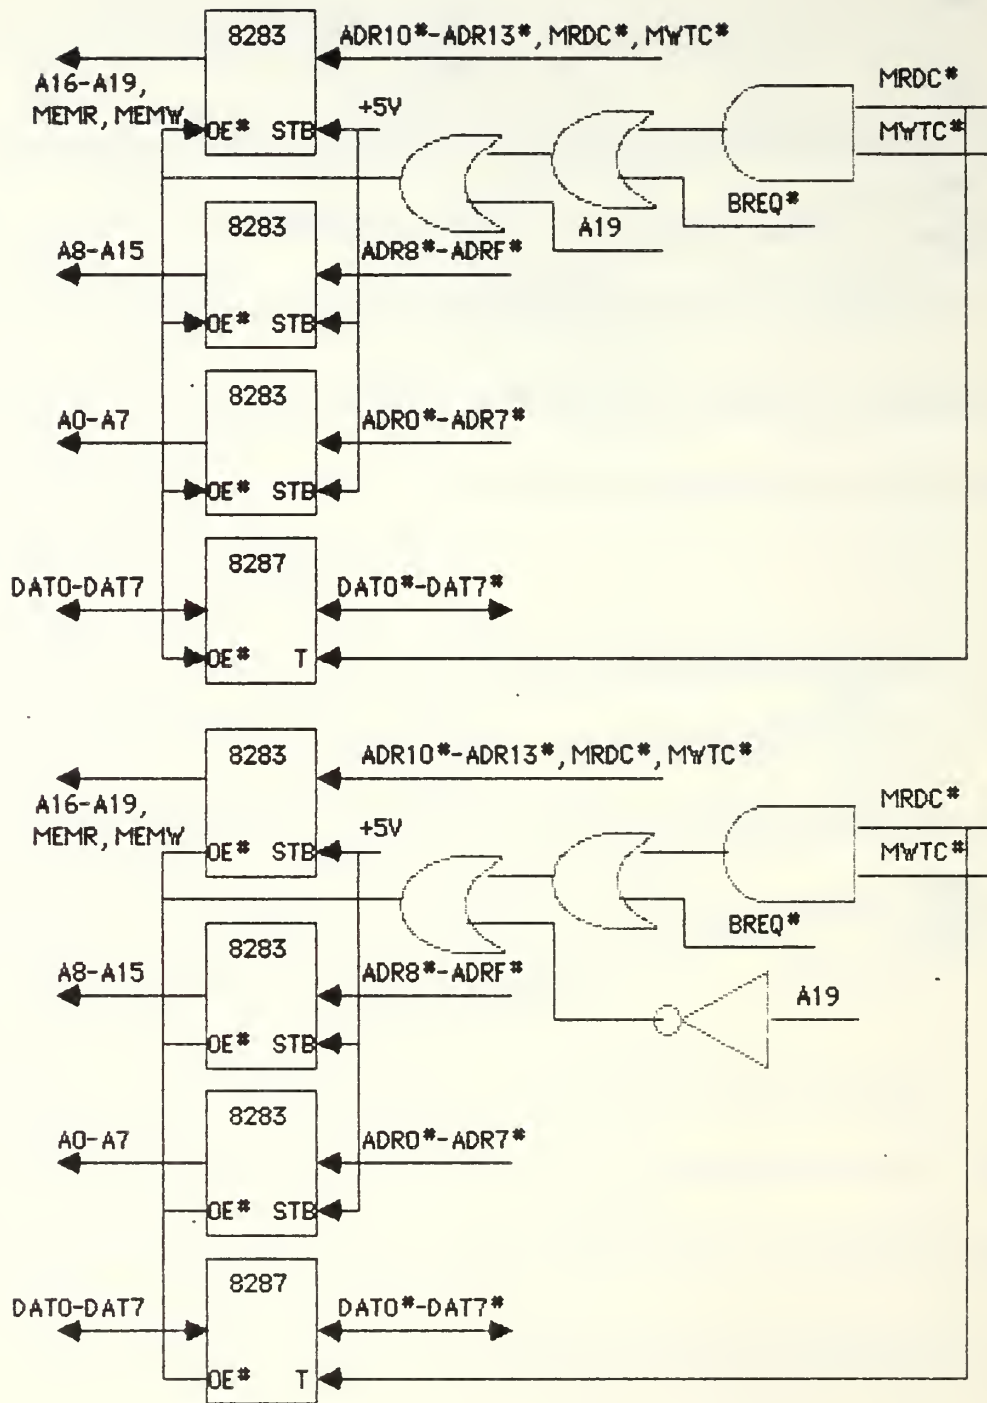Figure 28. Microprocessor Side of MEMR/MEMW Buffer for D0-D7

61

Figure 29. iSBC 264-4 Side of MEMR/MEMW Buffer for D0-D7 Bus

62

Figures 30 through 35 provide the schematics for the other three bytes of the data bus. On the microprocessor side, the difference is in which byte of the data bus is buffered, and the use of the appropriate byte enable signal, BE1*, BE2*, or BE3*. There are no differences between the various bytes of the data bus as they are connected from the RAM buffers to the iSBC 264-4 cards. Therefore, Figures 31, 33, and 35 are identical to Figure 29.

Figure 30. Microprocessor Side of MEMR/MEMW Buffer for D8-D15

64

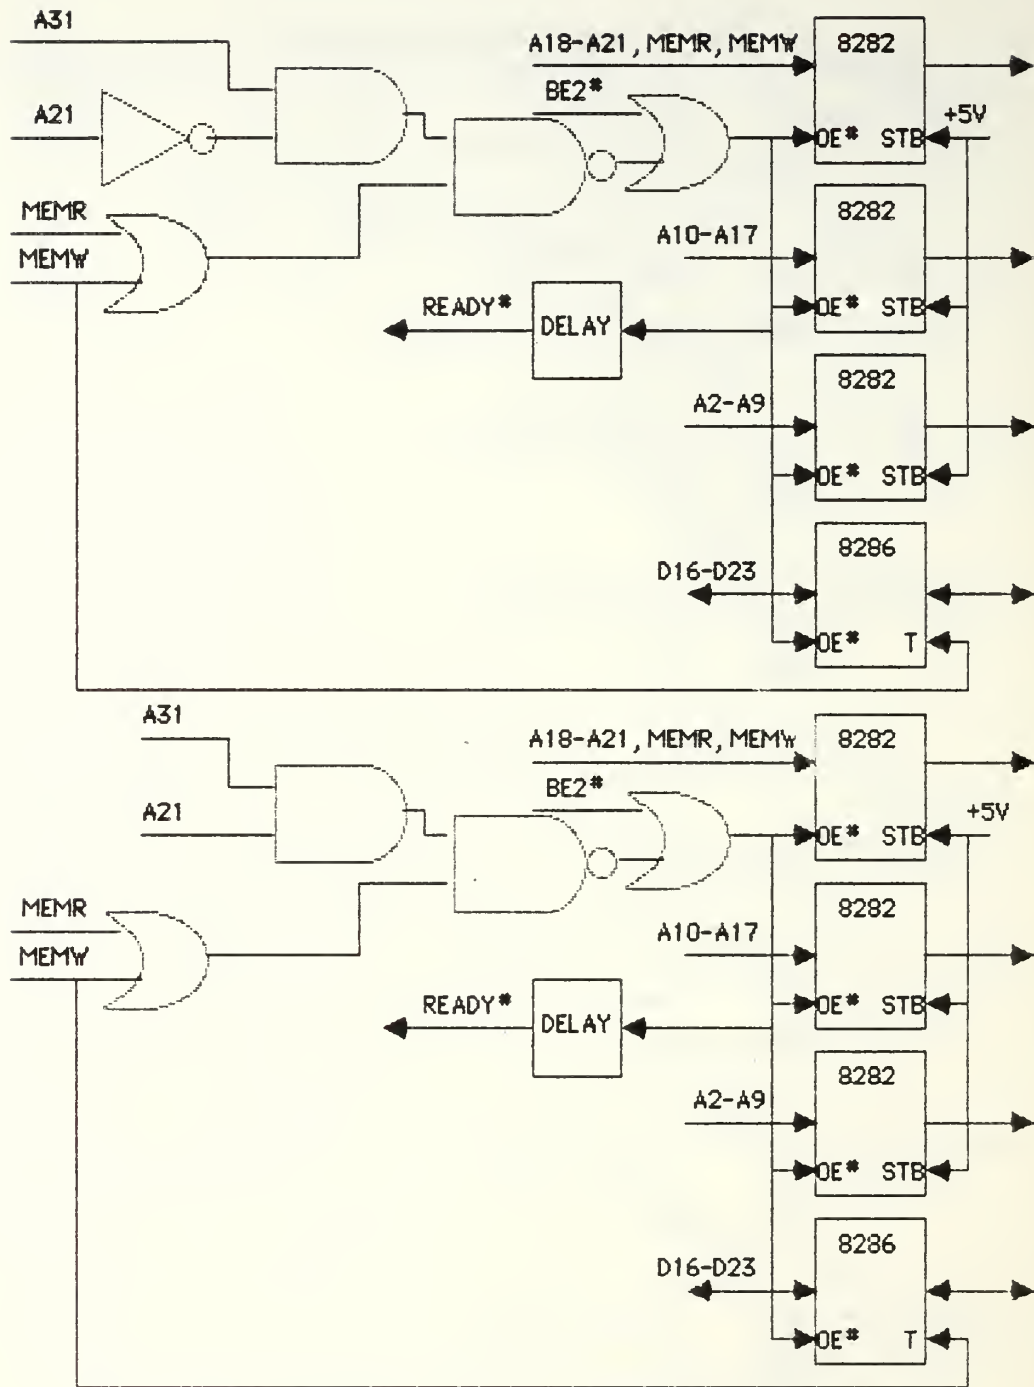Figure 31. iSBC 264-4 Side of MEMR/MEMW Buffer for D8-D15 Bus

Figure 32. Microprocessor Side of MEMR/MEMW Buffer for D16-D23
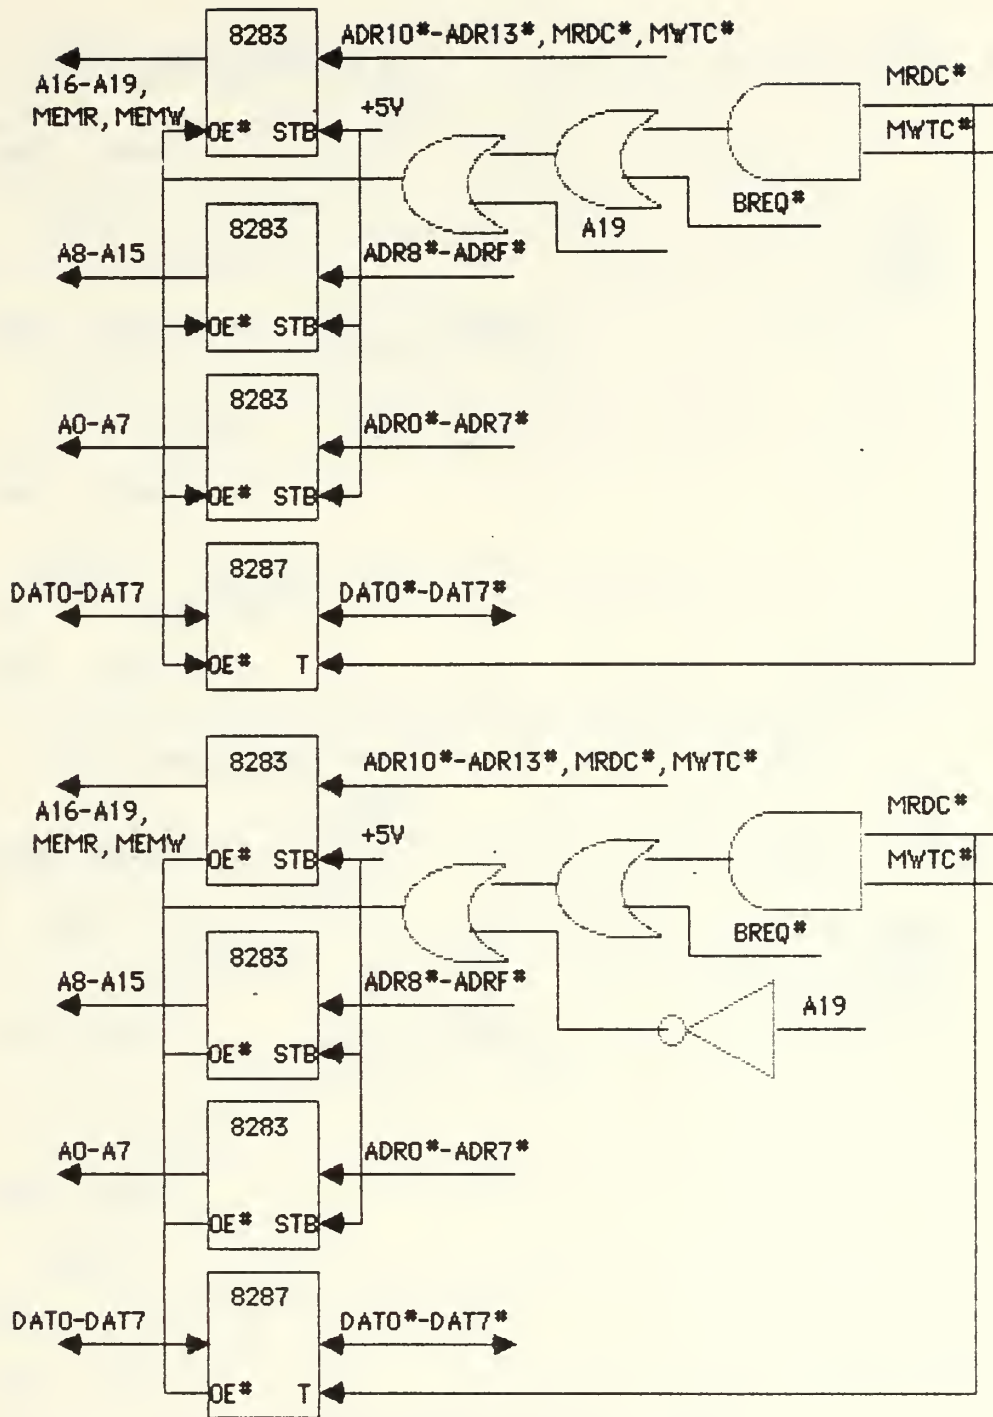
66

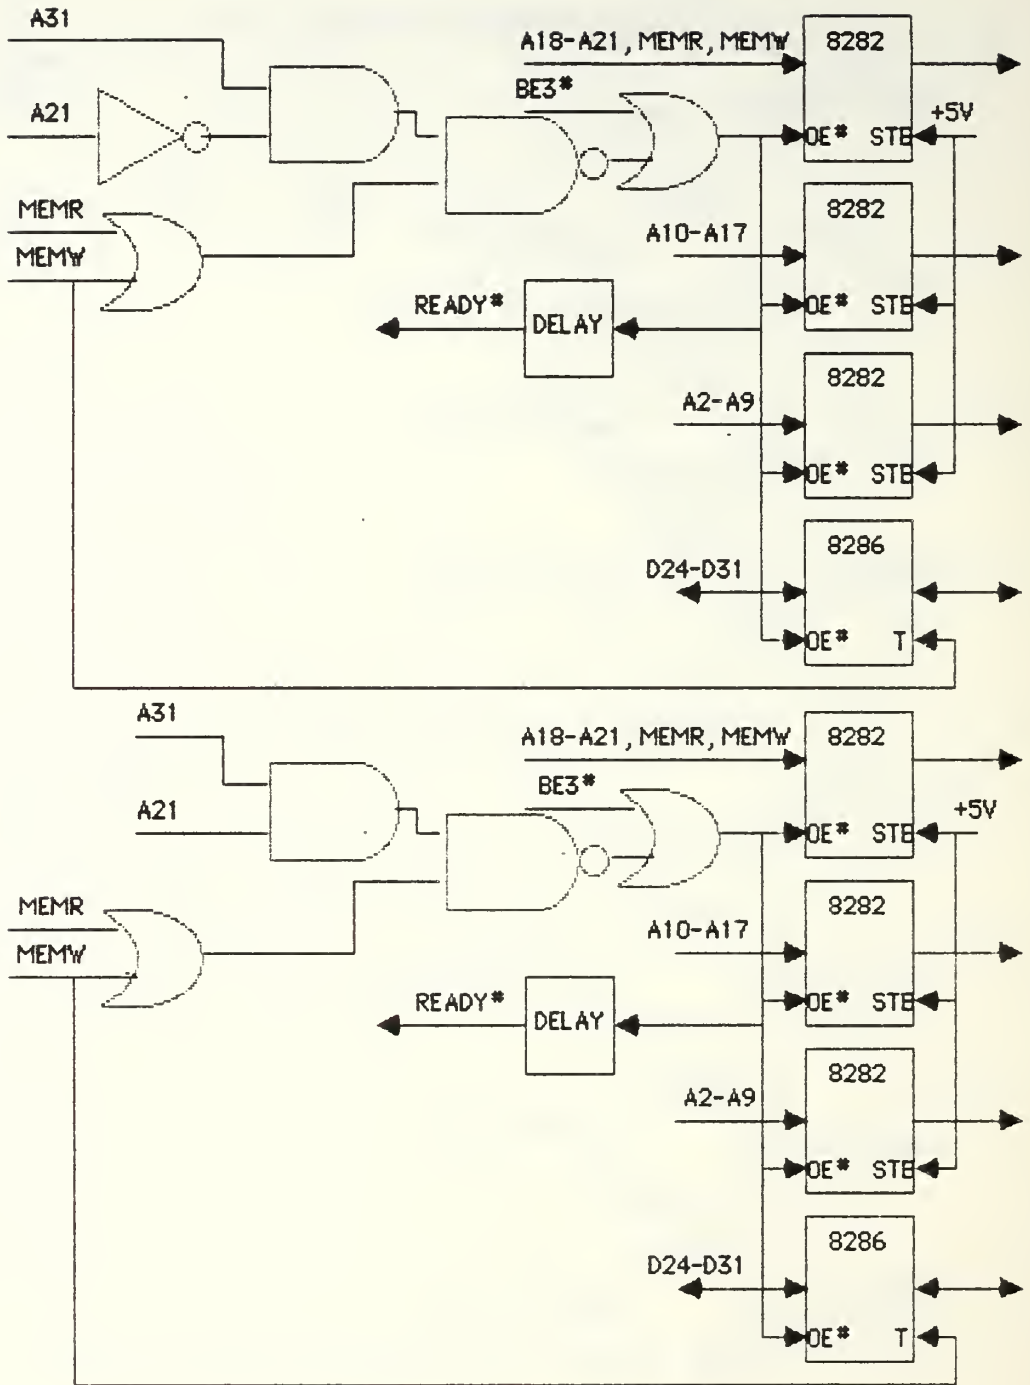Figure 33. iSBC 264-4 Side of MEMR/MEMW Buffer for D16-D23 Bus

Figure 34. Microprocessor Side of MEMR/MEMW Buffer for D24-D31
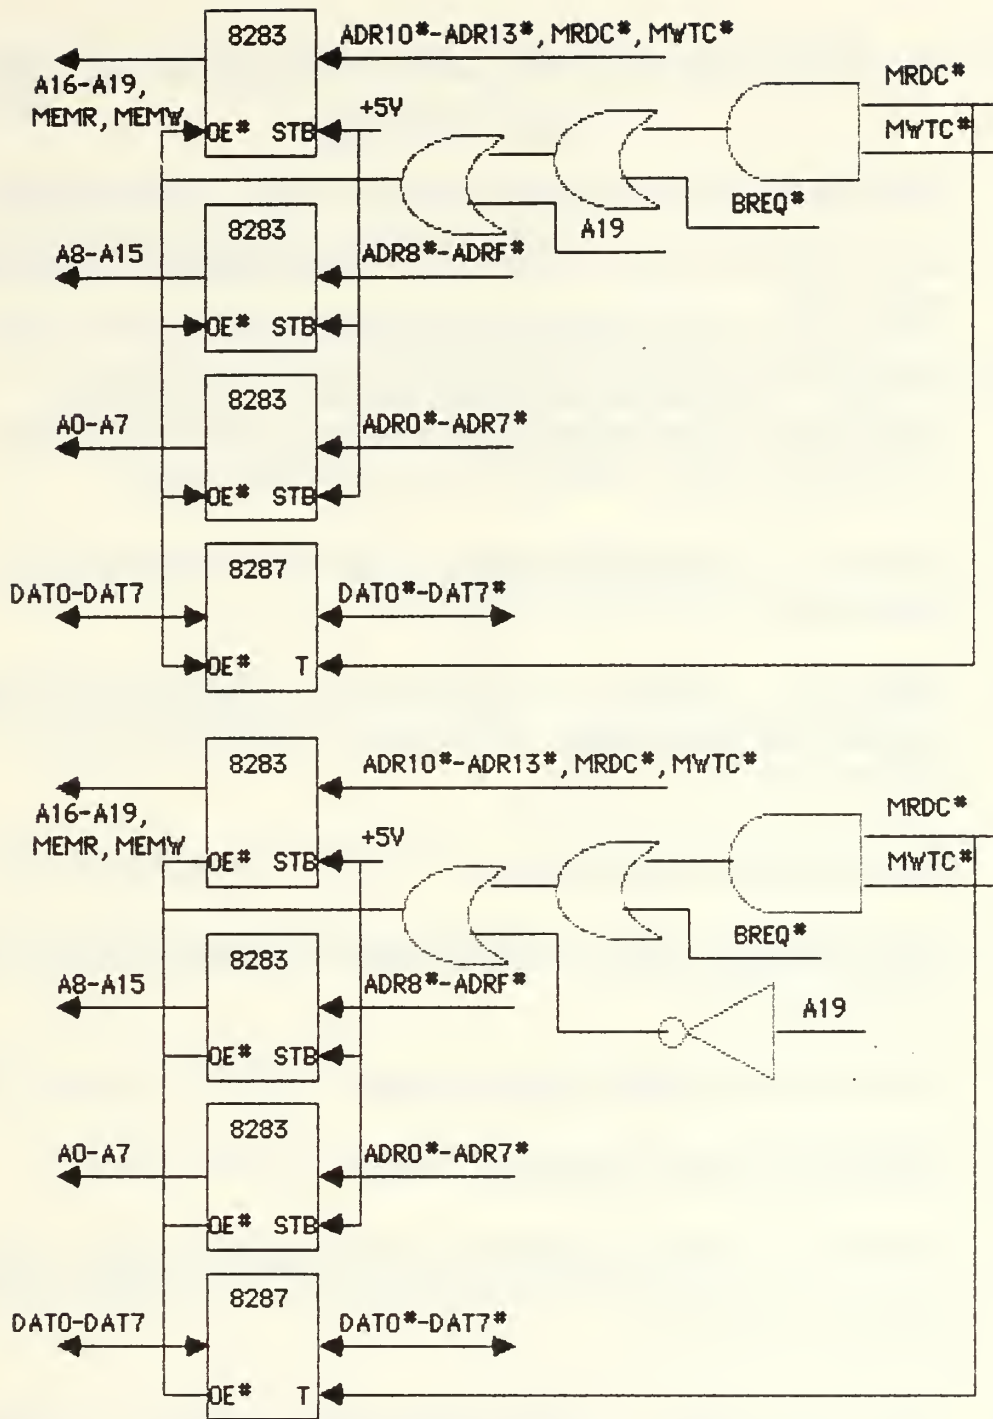
68

Figure 35. iSBC 264-4 Side of MEMR/MEMW Buffer for D24-D31 Bus

69

LIST OF REFERENCES

1. Electron Tube Division, Hitachi, Ltd., <u>Hitachi Magnetic Bubble Memory</u>, April, 1985

2. D'Ercole, T. M., <u>A Solid State Data Recorder for Space Based Applications Using Magnetic Bubble Memory</u>, MSEE Thesis, Naval Postgraduate School, Monterey, California, March 1986.

3. Huheey, J. E., <u>Inorganic Chemistry: Principles of Structure and Reactivity</u>, p. 508, Harper & Row, 1972

4. Mayberry, M., "Magnetic Film Development," <u>Magnetic Bubble Memory Research and Development</u>, conference at Intel Corp., Fulsom, California, March 11, 1986.

5. <u>Memory Componets Handbook</u>, pp. 6-1 to 6-10, Intel Corp., 1984

6. <u>BPK 5V75A Four-Megabit Bubble Memory Prototyping Kit User's Manual</u>, Intel Corp., 1984.

7. <u>Hitachi Magnetic Bubble Memory Boards</u>, Hitachi, Ltd., Japan, 1985.

8. <u>CMOS Bubble Memory Controller</u>, Hitachi, Ltd., Japan, 1985.

9. Motorola, Inc., <u>Radiation Hardened Bubble Meory for Space Applications</u>, technical presentation at Wright Patterson Air Force Base, Dayton, Ohio, April 26, 1985.

10. Phone conversation with Lance Day, Program Manager, Intel Magnetics Corp., February 11, 1986.

11. iSBC® 264 Magnetic Bubble Memory Board Reference Manual, Intel Corp.,
     1985.

12. 80386 High Performance Microprocessor with Integrated Memory
     Management, Intel Corp., October, 1985.

13. Howard, Steve, Intel's 80386 32-bit Microprocessor, Monterey,
     California, 1985. (Lecture presented at the Naval Postgraduate School.)

INITIAL DISTRIBUTION LIST

		No.	Copies
1.	Defense Technical Information Center Cameron Station Alexandria, Virgina 22304-6145		2
2.	Library, Code 0142 Naval Postgraduate School Monterey, California 93943		2
3.	Department Chairman, Code 62 Department of Electrical and Computer Engineering Naval Postgraduate School Monterey, California 93943		2
4.	Commander Naval Space Command Dahlgren, Virginia 22448		2
5.	Professor R. Panholzer, Code 62Pz Department of Electrical and Computer Engineering Naval Postgraduate School Monterey, California 93943		3
6.	Professor F. Terman, Code 62Tz Department of Electrical and Computer Engineering Naval Postgraduate School Monterey, California 93943		1

7. LT Thomas J. Frey, Jr., USN   2
    425 South Elm Street
    Palatine, Illinois. 60067

8. David Rigmaiden   1
    1135 Third Street
    Monterey, California 93940

9. Professor S. L. Garrett, Code 61Gx   1
    Department of Physics
    Naval Postgraduate School
    Monterey, California 93943

10. Professor O. B. Wilson, Jr., Code 61W1   1
    Department of Physics
    Naval Postgraduate School
    Monterey, California 93943

11. Space Projects Group, Code 72   1
    Naval Postgraduate School
    Monterey, California 93943

12. Professor John Simpson   1
    Enrico Fermi Institute
    University of Chicago
    Chicago, Illinois 60010

13. Jim Ferguson   1
    Intel Corporation
    Mail Stop LF 2-56
    1998 Praire City Road
    Folsom, California 95630-4760

14. Jim Tobin   1
    United States Naval Research Laboratory
    4555 Overlook Avenue SW
    Washington, DC 20375

15. Stewart Cummins   1
    Air Force Wright Aeronautical Lab/AADE
    Wright Patterson Air Force Base, Ohio 45433

16. Steve Knapp   1
    Intel Corporation
    Mail Stop LF 2-63
    1998 Praire City Road
    Folsom, California 95630-4760

17. Lt. John Wisman USAF   1
    AFSTC/TPA
    Kirtland Air Force Base, New Mexico 87117

18. Aerospace Corporation   1
    Attention: Carol Tanner
    P.O. Box 92957
    Los Angles, CA 90009

19. Jack Belove   1
    General Manager
    Intel Magnetics Operation
    Intel Corporation
    Mail Stop LF2-53
    151 Blue Ravine Road
    Folsom, CA 95630

20.   CAPT Roger P. Onorati USN   1
       Commanding Officer
       Naval Ocean Research and Development Activity
       NSTL MS 39529

21.   Roman Fedorak   1
       Code 5023
       Naval Air Development Center
       Warminster, PA 18974

22.   William T. Lee   1
       Rocketdyne Division
       Rockwell International Corporation
       6633 Canoga Avenue
       Canoga Park, California 91303

23.   Col. Ross USAF   1
       Code 9110
       United States Naval Research Laboratory
       4555 Overlook Avenue SW
       Washington, DC 20375