DISK INTERFACE AND CONTROLLER

REFERENCE MANUAL

Mark L. Greenberg

July 12, 1972

Center for Research in Management Science
Systems Group

Technical Document

The Disk System

The disk system is comprised of up to four Century Data Model 215 type disk drives (two drives per cabinet), up to four disk interface boards (one per drive) which plug into the IO cage, and the disk controller microcode. The disk system is driven by commands which are stored in core memory. The disk controller (microcode) interprets these commands, causes correct sequencing of the appropriate operations to the disk hardware, and detects error conditions resulting from operation of the disk drive and interface. The disk interfaces perform serial/parallel conversion of data between the disk drive and the disk controller, generate cyclic redundancy characters and detect read and write errors, and provide a command and status interface between the controller and the drives.

A complete documentation of the disk interface and controller includes this document, the disk controller microcode listing, and the disk interface hardware drawings.

The concept of microcode controlling of disk drive data transfer was developed by the author and Charles A. Grant. The disk controller was designed and implemented by the author and the disk interface by A. Ross Harrower.

Disk Formatting

A disk pack is divided into 406 cylinders or head positions. Each cylinder is divided into 20 tracks, one track for each head. Each track is divided into 3 sectors around the disk. Thus, a disk is formatted into 24,360 sectors (406 x 20 x 3). Each sector can store 512 36-bit words of data (plus other information). The total disk capacity is then 12,472,320 words or about 449 million bits of data.

Each sector is divided into 3 records:

(1) header record

(2) unique name record

(3) data record

Each record is separated from other records by gaps in which no information is recorded. The purpose of the gaps is to allow enough time for safe switching between read and write operations.

Each record begins with a 4 word preamble which is 136 zero bits followed by an 8 bit sync character. The zeros allow sufficient time for the disk interface decoder phase lock-loop circuit to lock in on the clock pulses from the disk. The sync character indicates when detected by the disk interface that the first word of information immediately follows.

Each sector ends with a 1 word cyclic redundancy character (CRC). The CRC is 18 bits long with 18 unused bits following. The CRC is generated from the data written when a record is written. When a word is read, the CRC's is generated and compared with the written CRC. If the CRC's do not match exactly, then a data read or write error has occurred. The CRC check is designed to detect most burst errors and bitwise systematic errors.

(1) Header Record.

Contains a sixteen bit disk address for the sector composed of a cylinder address (9 bits), a head address (5 bits) and a sector address (2 bits). The remaining 56 bits of data in the header record are unused.

(2) Unique Name Record.

Contains a 48 bit unique name supplied by the disk manager. The unique name is used to check for software consistency and to aid in system crash recovery. The remaining 24 bits of data in the unique name record are unused.

(3) Data Record.

Contains 2 words for the time the data record is written and 512 words for the data block. The time is only 48 bits long so that 24 bits are unused.

Sector Size Analysis.

1. The bit transfer rate is 2.5 megacylces or one bit every 400 nanoseconds. Therefore a 36-bit word is transferred in 14.4 microseconds (.4 x 36).

2. Each sector consists of:

3 eight word gaps	24 words
3 four word preambles	12
3 CRC's	3
disk address	2
unique name	2
time written	2
data block	512
	557 words

3. The total time to write a sector is 8.021 msec. (554 x 14.4).

4. The disk rotates in 25 msec. ± 2%. Therefore, in the worst (fastest) case the disk rotates in 24.5 msec. (25 - .02 x 25) and a minimum sector time is 8.166 msec. (24.5 ÷ 3).

5. Therefore, the minimum idle time at the end of a sector will be 145 msec. (8.166 - 8.021). The outside diameter of the disk is about 16" and the circumference is about 50". Therefore, the outside edge travels at a speed of about 500 msec. per inch. Thus, the 145 msec. of idle time represents about one-third inch of unused space at the end of every sector. Clearly, a sector of information fits on the disk with room to spare.

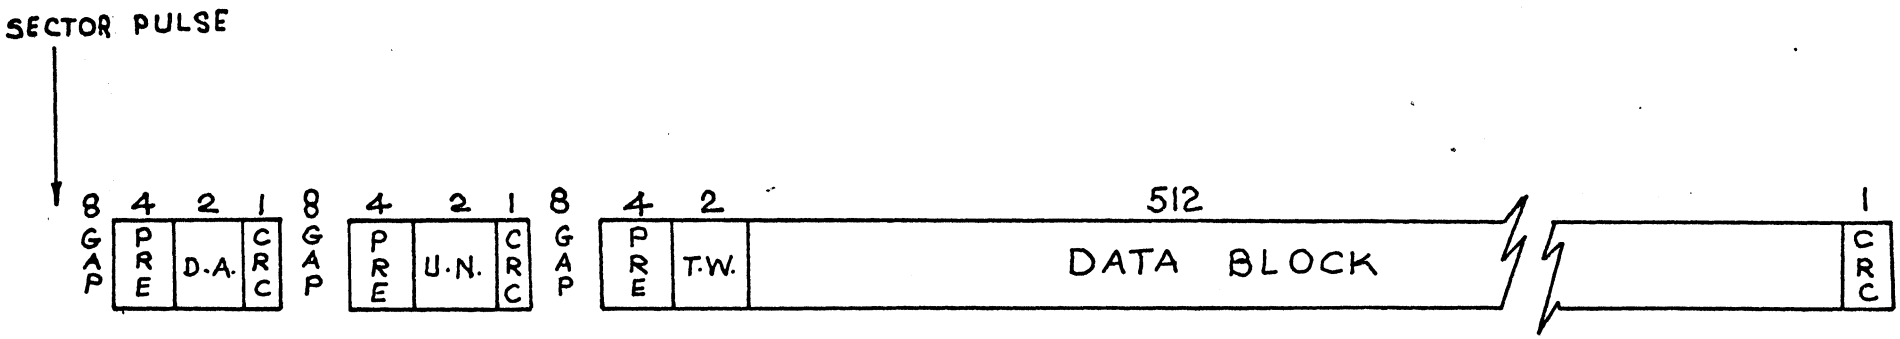

TW = time written
UN = unique name
DA = disk address

PRE = Preamble, 4 words or 144 bits long, 136 zero bits followed by an 8 bit sync character of all ones.

GAP = 8 words, long, no information written

CRC = 1 word, first 18 bits contain cyclic redundancy check character

Disk System Commands

Commands for the disk system are in the form of five-word elements which are stored in the first 253 words of core memroy. All commands for a particular disk drive are held on a simple linked list. The list is strung through the pointer fields (CHAIN) in the disk commands which contain the core address of the next command in the list. The last command in the list has a zero pointer field. The pointer to the current command, i.e., the first command in a list, is maintained by the disk controller, generally in a high speed register. A disk command has the following format:

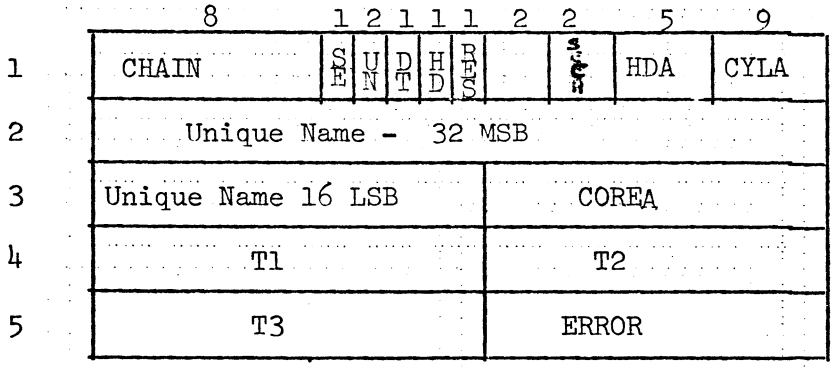

Option Bits.(HD, UN, DT, SE, RES)

HD = 1	write headers
ϕ	read and compare header
UN = $\phi\phi$	read unique name and store in command
ϕ1	read unique name and compare with command
10 or 11	write unique name from command
DT = 1	write time and data
ϕ	read time and data
RES = 1	disk restore operation
ϕ	no disk restore, normal command

Disk Address (SECA, HDA, CYLA)

CYLA = cylinder address

HDA = head address

SECA = sector address, if HD = 1, then SECA must be zero.

Unique Name.

This 48-bit unique name field is written in the unique name record on the disk of UN = 10 or 11, and is compared for equality against the unique name read from the disk if UN = ϕ1, or is loaded from the unique name read if UN = $\phi\phi$.

Disk Command List.

CHAIN = core address of next disk command or zero if the last disk command in the list.

Core Address.

COREA = the core address at which the data block transfer operation should begin.

Time.

If the disk command successfully completes and DT = ϕ, then the 48-bit real time is stored in T1, T2, T3, the 32 MSB in T1, T2, and the 16 LSB in T3.

Errors.

If an error occurs during a disk command operation, ERROR will be non-zero.

ERROR = 1 disk status error

2 index scan error

3 unique name mismatch error

4 header scan error

5 header write protect error

If ERROR = 1, then T2 receives the disk status word.

If ERROR = 4, then T2 receives the last disk address read.

Otherwise T2 is undefined.

For all errors, the 6 MSB of T1 receive the last microprocessor state routine number (with drive unsafe bit merged in if ERROR = 1), and T3 receives the SIMPLE machine Program Counter at the time the error is detected. T1 is undefined if ERROR = 3.

Command Descriptions.

There are three kinds of disk commands.

(1) Restore. (RES = 1).

Causes the disk to perform a restore operation. The heads are positioned to cylinder 0 and certain error conditions such as emergency head retract in the disk drive are removed. Restore should be attempted as part of the recovery procedure after error conditions occur.

(2) Header Write. (HD = 1)

This command is used to initialize the header records on a disk pack. Once the headers are successfully written, they should not be rewritten since all data stored on the disk would then be lost. The write header protect-switch, if on, will prohibit accidental writing of headers. The header-write command will cause three headers to be written for the three sectors at cylinder address CYLA and head address HDA. Unique name and data records will be read or written as with a normal command. However, a zero unique name and a zero data block should normally be written to initialize the disk pack. SECA must be zero initially. On successful completion, SECA = 3. The disk controller will begin writing headers at the sector after the sector with the index pulse. Therefore, all sectors with sector address (SECA) equal zero will be written at the same physical sector position and likewise for SECA = 1 and SECA = 2.

(3) Normal Command. (RES = 0 and HD = 0)

This command causes reading or writing of the information at the

sector specified by the command disk address (CYLA, HDA, SECA). The disk controller will cause the disk to select head HDA and position the heads to cylinder CYLA and then will read header records until a match between the header disk address and the command disk address is found. Then the unique name record is written from the command unique name (UN = 10 or 11) or read into the disk command (UN = $\phi\phi$) or compared for equality with the disk command unique name (UN = ϕ1). Then the data record is written (DT = 1) or read (DT = 0). If written, the real time of the write operation is written with the data. If read, the real time is read into the disk command. A data block transfer begins at COREA and continues to the last word of the 512-word core page in which the transfer began. Thus, the number of words transferred is given by 512 - (COREA MOD 512). Normally, a full 512-word page is transferred in which case the 7 MSB should be the page number and the 9 LSB should be zero.

Error Conditions.

There are five error conditions which can occur during disk command processing.

Status Error (ERROR = 1)

At the completion of every disk drive operation, i.e., read, write, seek, or restore, the error status bits from the disk interface are tested for error conditions. A status error completion occurs if any of these conditions are true.

Index Scan Error. (ERROR = 2)

During a header write command, if the disk controller fails to

find an index pulse after scanning three sectors, then an index scan error command completion occurs. A hardware malfunction is indicated.

Unique Name Error. (ERROR = 3)

If a unique name compare is specified (UN = 01) and the disk command unique name is not identical to the unique name read from the unique name record, then a unique name error command completion occurs. The unique name field of the command will be indeterminate.

Header Scan Error. (ERROR = 4)

During a normal command, if the disk controller read three header records without matching the disk address from the disk command with a disk address read from a header record, then a header scan error command completion occurs. A badly written header, a faulty seek operation, or a faulty head selection may be indicated.

Header Write Protect. (ERROR = 5)

If the write protect switch is on during a header write command, then a header write protect error command completion will occur.

DISK CONTROLLER FLOW CHART

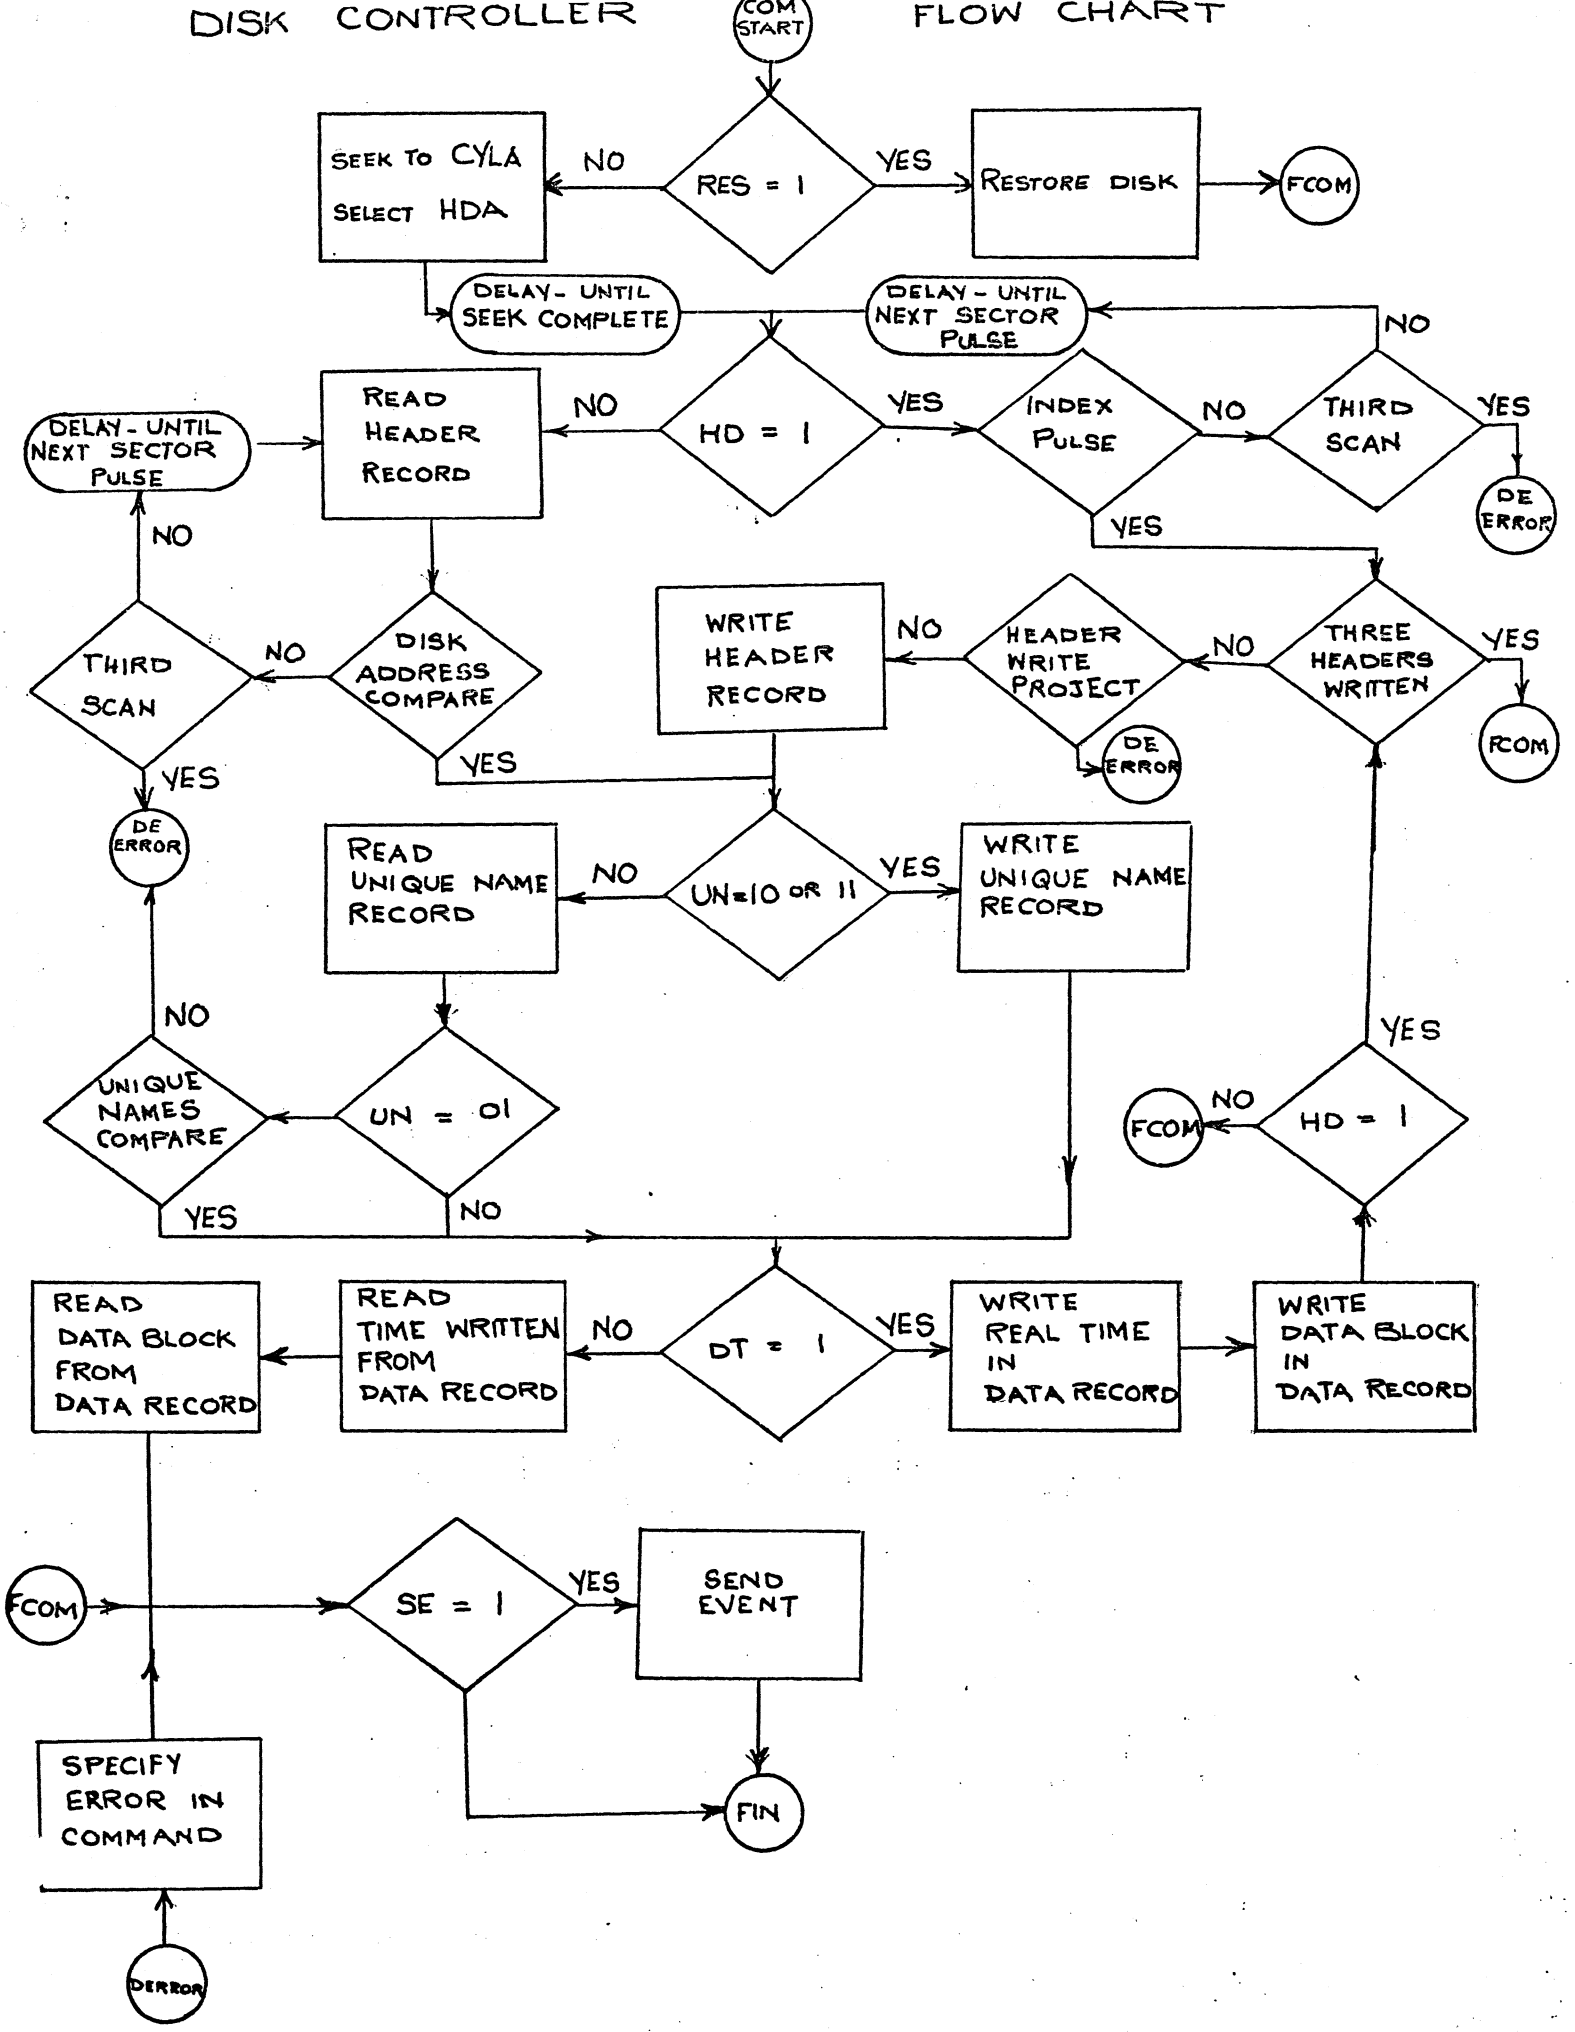

Disk Interface Description.

A disk interface can be logically divided into five sections:

(1) disk control

(2) disk status

(3) data write

(4) data read

(5) CRC generation and checking.

Disk Control.

This section interprets one of seven different commands from the IODATA output bus and generates the levels or pulses on the disk drive lines which cause the specified operation. The detailed function of each command is given later.

Disk Status.

This section generates 13 status signals on the IODATA input bus which indicate the condition of the disk drive and disk interface. Most of these conditions are flip-flops which are reset after their outputs are referenced on the IODATA input bus and are set when a particular error condition occurs. A detailed description of these is given later.

Data Write.

This section logically consists of an eight 18-bit word buffer which is serviced first-in first-out, an 18-bit shift register, and a data encoder. Data words are loaded into the buffer (holding registers) from the IODATA output bus on signal from the disk controller, or from the CRC section on a stop-write command. Words are moved

from the buffer to the shift register and shifted one bit at a time into the data encoder which puts the data into the code for the serial data line to the disk drive.

Data Read.

This section logically consists of an eight 18-bit word buffer, an 18-bit shift register, and a data decoder. Actually, the same physical buffer is used for both reading and writing. Signals on the serial data line from the disk drive are decoded and the resulting successive bits of data are shifted into the shift register. At the start of a read operation, these bits are compared against the sync character bit pattern. A match indicates the beginning of valid data. Thereafter, when the shift register becomes full, its contents are loaded into the buffer. The buffer is serviced first-in first-out. Words are gated from the buffer to the IODATA input bus and to the CRC section for CRC generation and error checking.

CRC Generation and Checking.

This section logically consists of an 18-bit register used to generate a cyclic redundancy character. At the beginning of a read or write operation, the register is cleared. Successive words read or written on the IODATA bus are bit-wise exclusive ored with the contents of the register and then the contents of the register are right cycled one bit position. At the end of a write operation, the contents of the register are loaded into the write buffer to be written at the end of the word. At the end of a read operation, the register is tested for zero. A non-zero register indicates a read or write error has occurred.

Disk Controller - Disk Interface Communication.

All disk interfaces (one per disk drive) send status and data to the disk controllers on a common 18-bit bus called the IODATA input bus. Eight disk device codes (see appendix) determine if data or status from each of the four interfaces are presented on the bus.

All disk interfaces receive commands and data on an 18-bit bus called the IODATA output bus. The eight disk device codes determine whether commands or data to each of the four disk interfaces are on the bus.

A six bit bus called IODEVICE is used to present the device code to the disk interfaces.

Finally, each disk interface has two unique attention lines on a bus called IOSTATUS. These lines are used to notify the disk controller of control, read, and write attention conditions.

When a command is placed on the IODATA output bus by the disk controller with a microinstruction which references IODATA on the D-bus, then both the PZ and IO bits should be set. The PZ bit will insure completion of the previous command by the disk interface before a new command is initiated and the IO bit will generate a GO signal to the interface which initiates execution of the command specified on the bus. When data are placed on the IODATA output bus, the IO bit should be set in the micro-instruction. The generated GO signal tells the disk interface that the next 18 bits of data are now stable on the output bus.

When status or data is received on the IODATA input bus by the disk controller by referencing IODATA in a microinstruction on the B-bus,

the IO bit should be set. The IO bit generates an ACK signal to the disk interface which tells the disk interface that status or data has been received and status conditions may be cleared or the next word of data may be gated onto the bus.

Except where here specified, IO and PZ bits should never be used while a disk device is selected on the IODEVICE bus.

Disk Interface Commands.

The disk interface accepts commands on the 12 low order bits of the IODATA output bus in the following format.

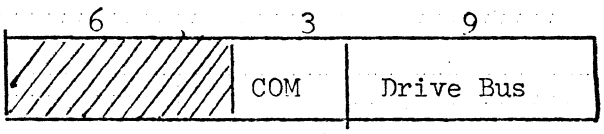

The COM field is used to encode seven different commands to the disk interface. Zero means no command.

START READ (COM = 1)

The disk interface raises the appropriate disk drive signal lines to select heads and activate the read electronics (i.e., read gate and head select). Word data transfer begins when the sync character is detected by the disk interface.

START WRITE (COM = 2)

The disk interface raises the appropriate disk drive signal lines to select heads and activate the write and erase electronics (i.e., head select, write gate and erase gate). Three zero words are written by the disk interface. The first word given to the disk interface by the disk controller should be the sync character word.

STOP READ OR WRITE (COM = 3)

If a write is in progress, the CRC that has been generated by the disk interface is loaded into the holding registers and the raised signal lines are lowered, thus stopping the write as soon as all holding registers

in the disk interface are empty. If a read is in progress, one more word is read from the disk if the holding registers are empty, and then the read signals to the drive are lowered. The word in the holding register is the CRC which is compared against the generated CRC for errors.

CONTROL PULSE (COM = 4)

The command sends the Drive Bus field of the command to the disk and pulses the Control Tag line, thus issuing control tag commands such as seek, start, restore, and head reset to the disk drive.

SET HEAD ADDRESS (COM = 5)

The drive bus field is sent to the disk and the set head tag line is pulsed, thus loading the disk drive head address register.

SET CYLINDER ADDRESS (COM = 6)

The drive bus field is sent to the disk and the set cylinder tage line is pulsed, thus loading the disk drive sylinder address register.

SET CONTROL ATTENTION (COM = 7)

The control attention condition to the disk controller is set. This command is used to notify the disk controller that the disk command list is no longer empty. Control attention is reset when the disk control device code is selected.

Disk Interface Status.

A disk interface provides a 13-bit status word to the disk controller on the IODATA input bus when the disk control device is selected in the following format:

5							13						
ϕ	DUS	CRC	STL	OVF	SWH	RO	WCS	EOC	SIN	DNR	NOL	WHI	IDX

WRITE HEADER INHIBIT (WHI)

True if write header protect switch is on.

INDEX PULSE (IDX)

True if index pulse occurred during previous sector.

CYCLIC REDUNDANCY ERROR (CRC)

True if previous operation was a read with a cyclic redundancy error detected.

START READ OR WRITE LATE (STL)

True if previous operation was a read or write that was started by disk controller after falling edge of read or write disable. Indicates a microcode timing error.

BUFFER OVERFLOW OR UNDERFLOW (OVF)

True if the previous operation was a read or write that lost or picked up data as a result of disk interface holding register overflow or underflow. Indicates a microcode timing error.

SECTOR PULSE WHILE HEAD SELECTED (SWH)

True if the previous operation was a read or write that was in

progress when a sector pulse occurred, Indicates the formatted data did not fit in a sector.

READ ONLY (RO)

True if the previous operation was a read some time during which the read-only switch on the disk drive was on.

WRITE CURRENT SENSE (WCS)

True if the previous operation was a write some time during which the write current went low. Indicates data incorrectly written.

END OF CYLINDER (EOC)

True if the previous operation was a read or write to a non-existent head. Head addresses must be in the range 0-19.

SEEK INCOMPLETE (SIN)

True if the previous operation was a seek to a non-existent cylinder or if the seek did not complete within one second.

DRIVE NOT READY (DNR)

True if the previous operation was a read or write some time during which the drive was not ready. In particular, a drive is not ready if the heads are not detented.

NOT ON LINE (NOL)

True if the previous operation was a read or write some time during which the drive was not on line. A drive is on line if the heads are not retracted and a retract operation is not in progress.

DRIVE UNSAFE (DUS)

True if a drive unsafe condition exists in the disk drive. A drive unsafe condition can be removed only by powering down the disk drive.

Disk Controller Implementation.

The disk controller consists of eleven ROM pattern boards holding about 350 micro-instructions. Four pattern boards contain disk transfer units, one per disk drive and seven boards contain the disk control state routines which are common to all four disk drives.

The disk controller is organized somewhat like a finite state machine with ten states numbered ϕ through 9. The current state of the disk controller is maintained in a sixteen word table (called DSTATE), starting at core location 256. The state of each drive is kept in a four-word entry in DSTATE. Only the first word of the entry is used and has the following format.

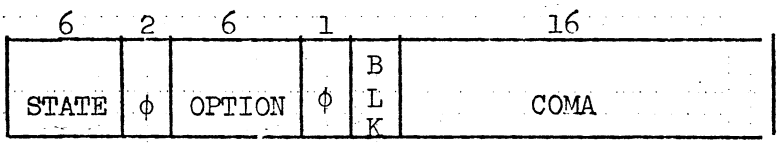

STATE = current state number

OPTION = option bits from current command

COMA = core address of current command

BLK = block transfer bit

The disk controller state routines are entered for a particular disk when the disk interface generates a control attention for the disk. A control attention is generated:

(1) on completion of a write operation.

(2) on completion of a read operation.

(3) on any sector pulse occurring after success or error completion of a seek operation.

of a seek or restore operation and before initiation of the next seek or restore operation.

(4) on evoking a set control attention command to the disk interface from the microcode.

Thus, in general, a control attention indicates that a disk operation is completed and it is time to initiate another operation.

The state number (STATE) indicates which state routine should be entered when a control attention occurs. Upon a control attention, the microcode functional unit scheduler will transfer control to locations D1COMMAND, D2COMMAND, D3COMMAND, and D4COMMAND for the four disk drives, respectively. These locations will load LINK with four times the disk number (numbered 0 through 3) and branch to the command state routine entry logic at DCOMMAND. Four times the disk number is kept to facilitate generation of disk device select codes which are left shifted 2 bits in the IODEVICE register.

The DCOMMAND entry logic:

(1) tests for a status error (no status errors in state ϕ).

(2) gets DSTATE from core and sets up the working registers.

(3) increments the state number in DSTATE.

(4) dispatches to the state routine address.

Control Sequences.

There are three control sequences that the disk controller may take depending on whether the disk command is (1) a restore (RES = 1); (2) a header write (HD = 1) or (3) a normal command.

RESTORE

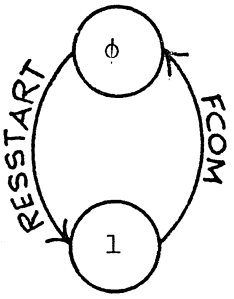

command start

restore complete

The COMSTART routine will initialize a restore operation to the disk interface. State 1 is entered when the restore is complete and the command finish sequence (FCOM) is entered.

HEADER WRITE

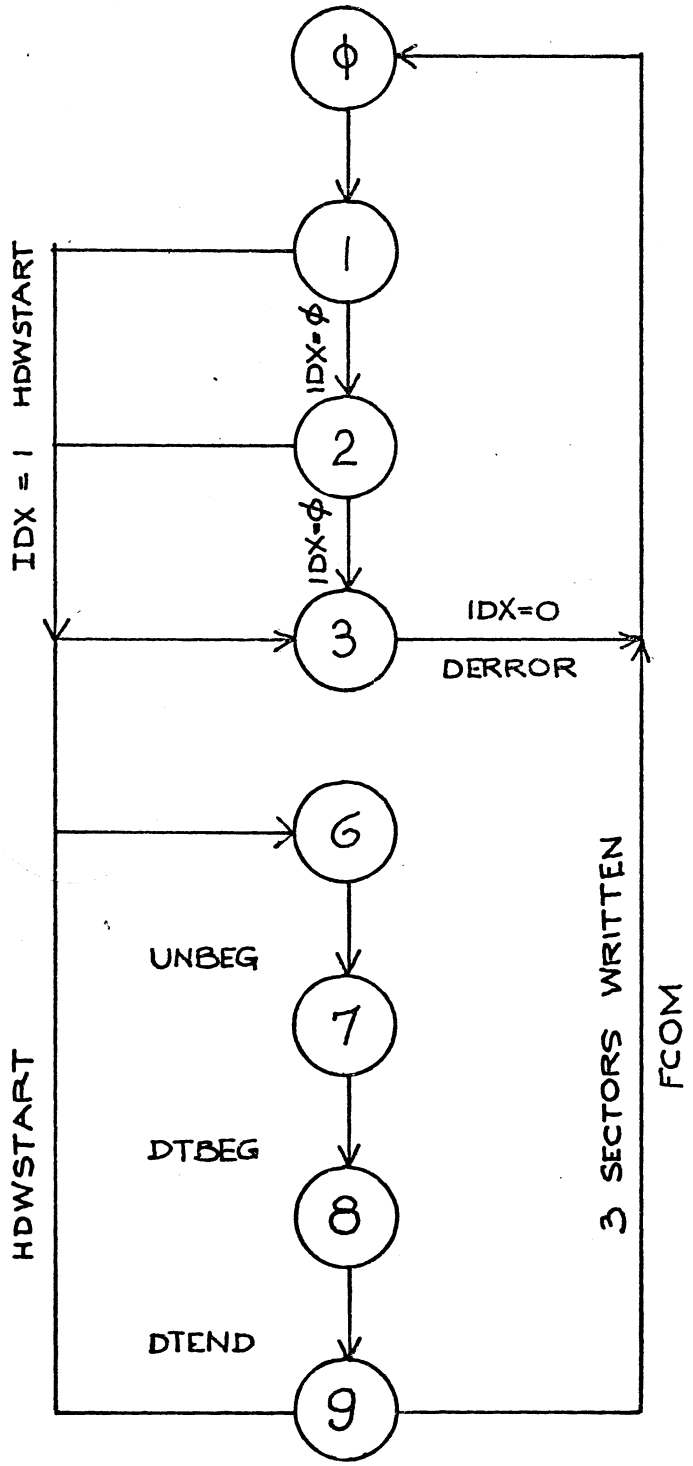

The COMSTART routine initializes a seek start operation to the disk interface. The seek operation is complete at state 1. States 1, 2, and 3 scan successive sectors for an index pulse. An index scan error occurs if not found by State 3. When an index pulse is found a header write is initiated and the state number is set by SETSTATE to be in state 6 when the header write completes. At state 6 a unique name record read or write is begun (UNBEG). At state 7, the unique name record is complete and a data record read or write is begun (DTBEG). At state 8 the data record operation is done (DTEND). At state 9 the next sector pulse has been reached and a new header is written returning to state 6 unless all three headers have been written already in which case a command finish sequence is entered (FCOM).

NORMAL COMMAND

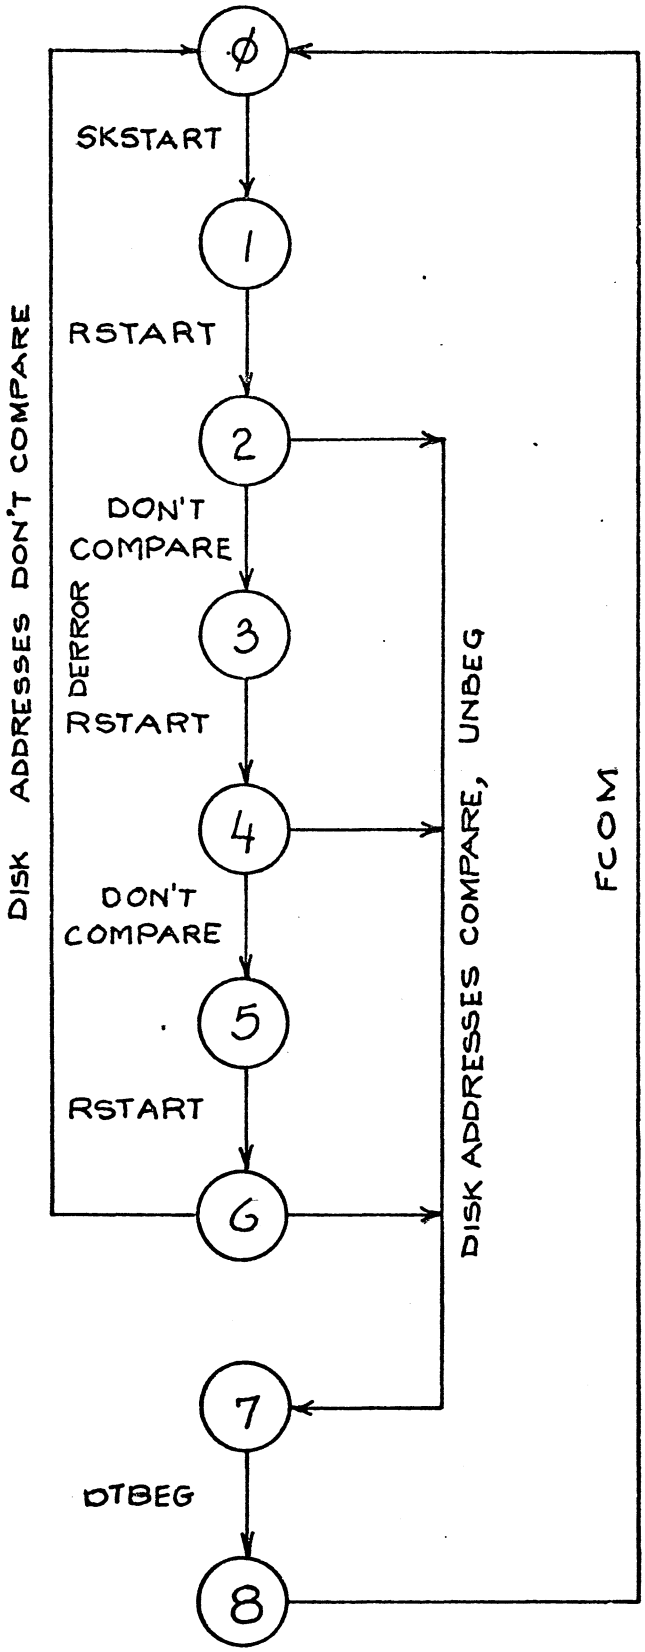

The COMSTART routine initializes a seek start operation to the disk interface. The seek operation is complete at state 1. State 1, 3, and 5 initialize header record read operations. States 2, 4, and 6 compare the header disk address against the command disk address. If they compare identically, the unique name record operation is begun and the state number is set to state 7. If they don't compare, the disk drive advances to the next sector at the next state, except for state 6 where a scan error is indicated and the DERROR logic is entered. State 7 causes initiation of the data record operation. It is done in state 8 and the command finish sequence is entered.

Data Transfer Units.

The data transfer units (one per disk drive) respond to read and write attention conditions when entered from the scheduler. A read attention occurs during a read operation whenever there is at least one full 36-bit word heading register in the disk interface. A write attention occurs during a write operation whenever there is at least one empty 36-bit word holding register in the disk interface. One call of a transfer unit causes a single word to be transferred between core memory and the disk interface. Then the address register (D1TEMP, D2TEMP, D3TEMP, D4TEMP for the four disk units, respectively) which points to the word transferred, is incremented and tested to determine if the block transfer is complete. All block transfers complete at 512 word boundaries except the first 1024 words of core are divided into four 256 word blocks. The last three words of each of these four blocks are used by the disk controller as buffers for the four disk units respectively. The buffers are used as a place into

which to read the first two data words of a record or a place out of which to write the fourth preamble word and the first two data words of a record. If a data transfer unit determines that a block transfer is complete, then the DFINISH code sequence is entered. The DFINISH sequence will stop a read or write operation (RWSTOP) unless the data block of the data record is still to be transferred (state 8 and block transfer bit set) in which case the block transfer bit is cleared and the core address for the data block transfer is loaded into the address register, e.g., DITEMP) from the disk command (RWBLOCK).

Disk Controller Routines.

In general, routines are named by the label on the first instruction of the routine:

HDASTART - starts a header write. First enters FCOM routine if 3 sectors have been written. Otherwise, loads disk address from command into core buffer, increments sector number in command disk address, goes to DERROR if write header protect switch is set, starts write operation (WSTART) and causes state to be set to state 6 and causes block transfer bit to be set (SETSTATE).

WSTART - start a write operation after loading last word of preamble in buffer and setting address register to first word of buffer.

RSTART - starts read operation after setting address register to second word of three-word buffer.

SETSTATE - is entered after RSTART or WSTART if state is to be changed or block transfer bit is to be set (DS $\neq \phi$)

and such action is then taken by merging DS into DSTATE entry.

UNBEG - starts a unique name record read operation or loads buffer with unique name from disk command and starts a write operation.

UNCOM - compares unique name in buffer against unique name in disk command and generates an error (DERROR) or mismatch or reads unique name in buffer into disk command.

DTBEG - starts data record read operation or reads real time clock and stores time in buffer and then starts a write operation (DTWRITE).

DTEND - if a data record read end, the time written is copied from the buffer to the disk command. Then, if a header write, HDWRITE is entered; otherwise the command finish sequence is entered (FCOM).

SCAN - compares disk address in buffer against disk address in disk command and if they compare, the state is set to state 7 and the block transfer bit is set as UNBEG is entered. Otherwise, a scan error is detected (DERROR) if in state 6.

DERROR - loads disk command words 4 and 5 with error indicators and enters command finish sequence (FCOM).

FCOM - an event is sent to the disk manager process if specified and then the COMSTART sequence is reentered.

DISK CONTROLLER CORE MAP

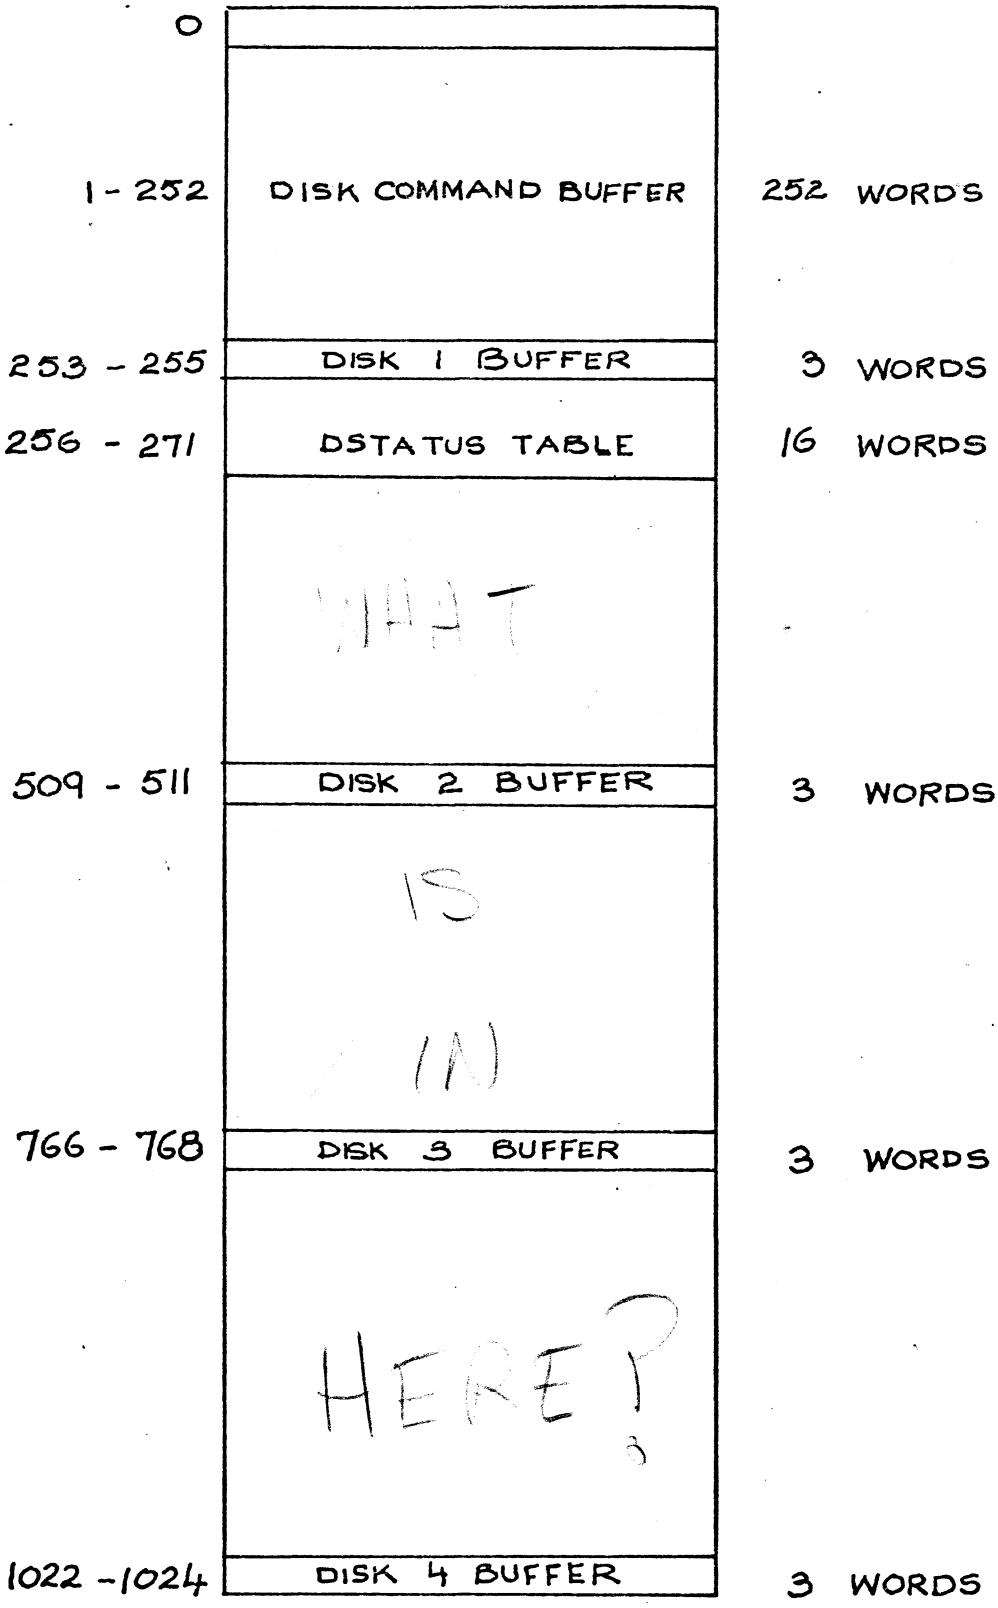

Disk Controller Scheduling.

A disk controller routine is called in response to disk attention conditions. For each of the four potential disk units, there are three different attention conditions.

(1) read attention

(2) write attention

(3) control attention.

The read and write attentions together are referred to as the transfer attentions. Disk controller routines are divided into those that respond to transfer attentions and those that respond to control attentions.

At any time, the microprocessor is executing in some routine or functional unit or is executing in the scheduler. In addition to the disk control routines, there are function units for the IPU, teletype multiplexor, etc. All other functional units are of lower priority than the disk controller functional units. Furthermore, the disk control routines are lower priority than the disk transfer routines.

At frequent intervals in the various functional units, a test is made to see if there is any higher priority functional unit that requires service. If there is, then the scheduler is entered. It is the function of the scheduler to determine which functional unit should be executed next and to transfer control to that unit. The highest priority functional unit with its attention condition set is the one scheduled.

If the scheduler is entered from any functional unit except a disk transfer routine, then the scheduler will assign the highest transfer routine priority to DISK No. 1 and lowest to DISK No. 4. If the scheduler is entered

from a disk transfer routine, then highest transfer routine priority is given to the next transfer routine and lowest priority to the transfer routine that just finished. Since one call of a transfer routine causes one word to be transferred between core and the disk interface, this scheduling rule insures that no disk may transfer more than one word without first giving all other transfer routines a chance to transfer a word.

The typical execution time of a disk transfer routine is less than 2 microseconds. A worst case of 5 microseconds occurs only on the last word of a block transfer. A disk read transfer routine must complete a word transfer within 3.5 word times (3.5 x 14.4 msec. = 50.4 msec.) of that word entering the disk interface word buffers; otherwise, a buffer overflow will occur and data will be lost. An equal time constraint exists for write transfer operations.

As a worst case for analysis for data transfers, assume the highly improbable event of all four transfer units raising their attention conditions for the last word of a block transfer immediately after some lower priority functional unit is scheduled. All four transfer units must complete one full worst case execution of 5 msec. within 3.5 word times (50.4 msec).

This total execution, including scheduling would take 20 msec. leaving 30.4 msec. of execution for the lower priority functional unit.

For disk control routine scheduling, a worst case control routine executes for about 10 msec., and due to the width of inter-record.

gaps need not complete execution for five-word times (72 msec.) after their attention condition is raised. In the worst case, if all four control attentions raise simultaneously immediately after a lower priority functional unit is scheduled, then the 40 msec. of execution required for four worst case control routine calls leaves 32 msec. of execution for the lower priority functional unit.

The following rule is demanded for programming lower priority functional units:

NO FUNCTIONAL UNIT MAY EXECUTE FOR MORE THAN 30 MICROSECONDS WITHOUT OFFERING CONTROL TO THE SCHEDULER.

Disk System Device Codes.

10B	D1CONTROL
11B	D2CONTROL
12B	D3CONTROL
13B	D4CONTROL
14B	D1DATA
15B	D2DATA
16B	D3DATA
17B	D4DATA

Disk IOSTATUS Bits

	Bit	Bit
Disk 1	6	7
Disk 2	4	5
Disk 3	2	3
Disk 4	0	1

nothing	0	0
control attention	0	1
read attention	1	0
write attention	1	1

Comments on Gap Structure.

(1) gap delay is triggered by leading edge of sector pulse or falling edge of write gate.

(2) read disable and write disable are triggered by the falling edge of gap delay or the falling edge of read gate.

(3) write gate is raised 3 us. after falling edge of write disable.

(4) read gate is raised on falling edge of read disable.

(5) control attention conditions occur on falling edge of read gate or write gate or rising edge of sector pulse.

(6) the eight word gap size allows 5-word times from beginning of CRC to fall of read gate and 5-word times from fall of read gate to rise of write gate.

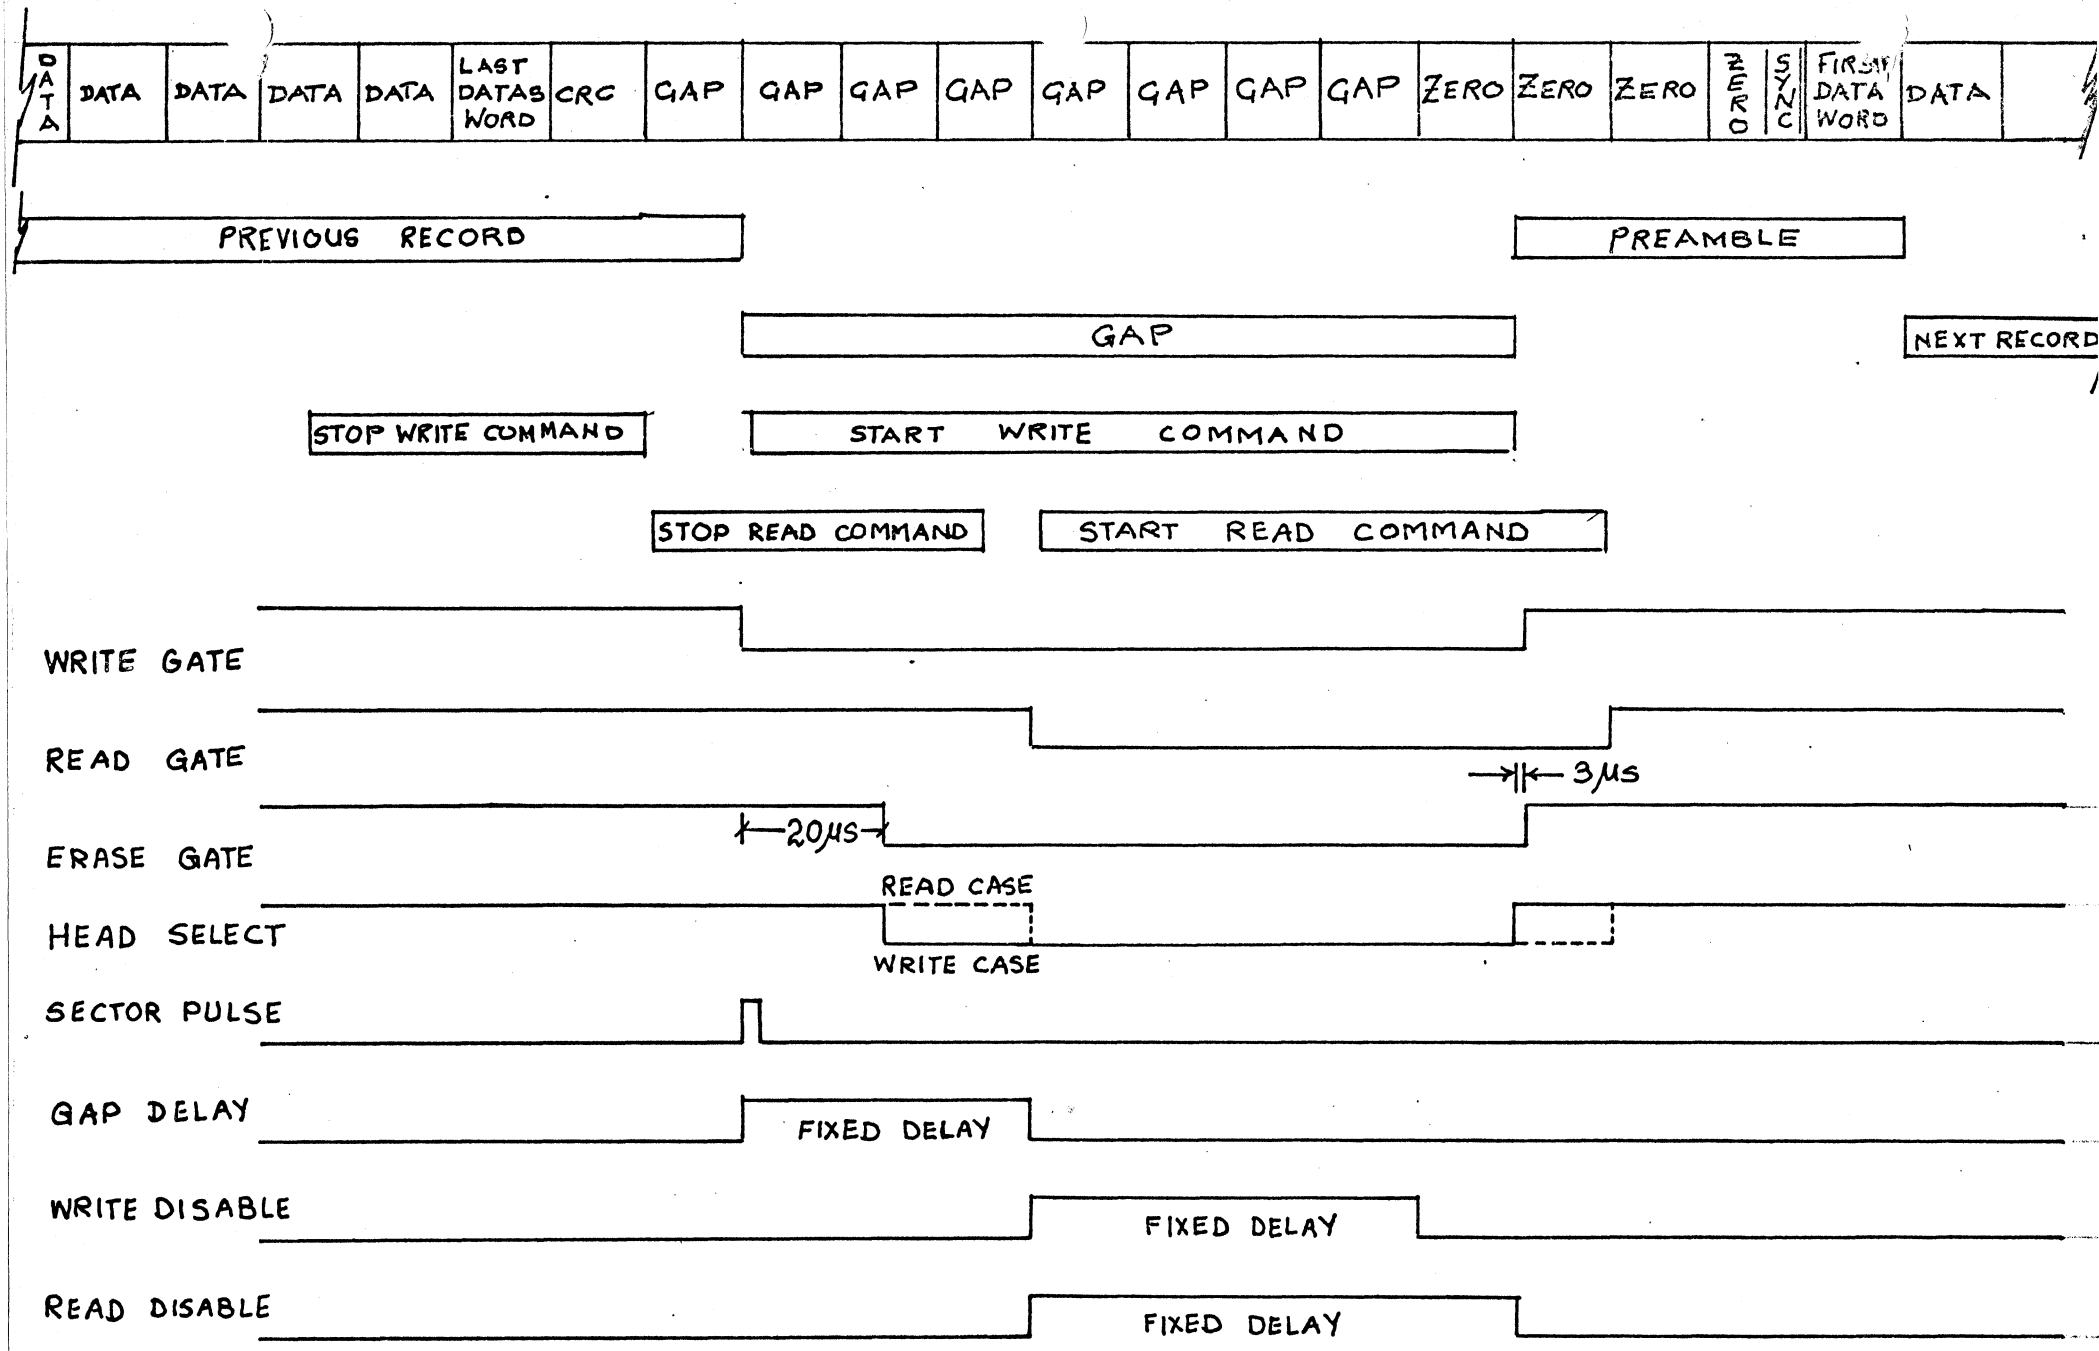

GAP STRUCTURE